

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
31 December 2008 (31.12.2008)

PCT

(10) International Publication Number
WO 2009/002722 A2

(51) International Patent Classification:
G06F 9/38 (2006.01)

(21) International Application Number:
PCT/US2008/066775

(22) International Filing Date: 12 June 2008 (12.06.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/821,647 25 June 2007 (25.06.2007) US

(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **DUFFY, John Joseph**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **CALLAHAN, David**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **DETLEFS, David**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **MORRISON, Vance**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **GRUNKEMEYER, Brian**; One Microsoft Way, Redmond, Washington 98052-6399 (US).

TRIBBLE, Eric Dean; One Microsoft Way, Redmond, Washington 98052-6399 (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, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, 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, MT, NL, NO, PL, PT, RO, SE, SI, SK, 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 applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

(54) Title: CONCURRENT EXCEPTION HANDLING

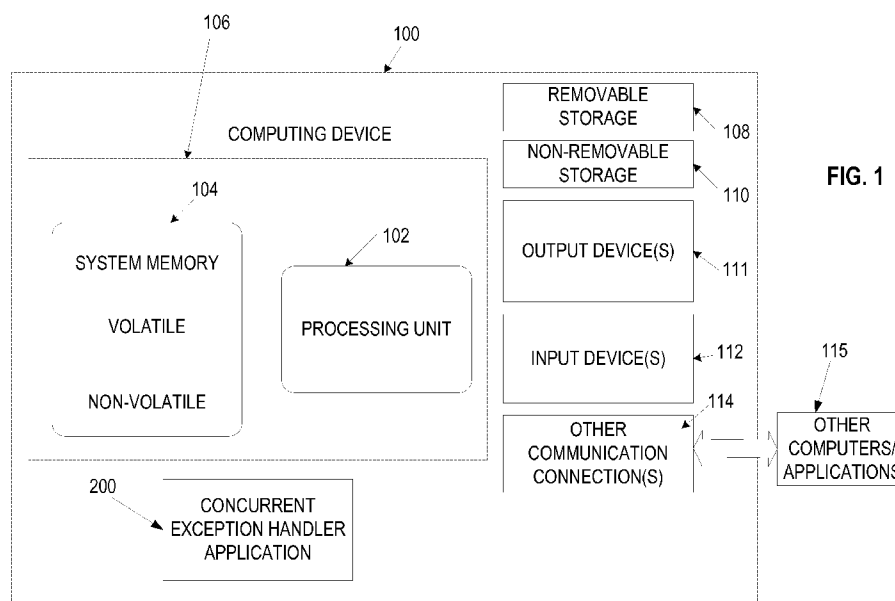


FIG. 1

(57) Abstract: Various technologies and techniques are disclosed for providing concurrent exception handling. Exceptions that occur in concurrent workers are caught. The caught exceptions are then forwarded from the concurrent workers to a coordination worker. The caught exceptions are finally aggregated into an aggregation structure, such as an aggregate exception object. This aggregation structure is rethrown and the individual caught exceptions may then be handled at a proper time.



-
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *without international search report and to be republished upon receipt of that report*

CONCURRENT EXCEPTION HANDLING

BACKGROUND

[001] Software programs have been written to run sequentially since the beginning days of software development. In a sequential program, tasks execute one after another, i.e. one task finishes completely first before the next is begun, and if a previous task does not finish subsequent tasks do not execute. Steadily over time computers have become much more powerful, with more processing power and memory to handle advanced operations. This trend has recently shifted away from ever-increasing single-processor clock rates and towards an increase in the number of processors available in a single computer, i.e. away from sequential execution and toward parallel execution. Software developers want to take advantage of improvements in computer processing power, enabling their software programs to run faster as new hardware is adopted. With parallel hardware, however, this requires a different approach: developers must arrange for one or more tasks of a particular software program to be executed “concurrently” (sometimes called “in parallel”), so that the same logical operation can utilize many processors at one time, and deliver better performance as more processors are added to the computers on which such software runs.

[002] Concurrent programs are fundamentally different than sequential ones, because many tasks happen at once instead of one after the other. This raises one particularly difficult problem: whereas in a sequential program, execution of a logical operation stops at the first occurrence of an exception, in concurrent programs exceptions may occur in many tasks at once. Current exception handling mechanisms assume linear, sequential execution, and so communicating such failures to the program presents problems. To make matters worse, the software developer is often working with an application programming interface that gives the appearance of a traditional sequential interface, even though its implementation runs work concurrently, which makes it increasingly difficult to encapsulate the implementation details of an interface when concurrency is used internally.

SUMMARY

[003] Various technologies and techniques are disclosed for providing concurrent exception handling. Exceptions that occur in concurrent tasks are caught and forwarded to a coordination task. These exceptions are then aggregated into an aggregation structure, such as an aggregate exception object, and rethrown by the coordination task. The aggregation structure may then be caught and the individual exceptions inside handled appropriately by the program.

[004] In one implementation, the exceptions are handled in the program once all concurrent tasks have completed. In such an implementation, a coordination task initializes a shared flag and a shared list of exceptions to be aggregated.

Concurrent worker tasks are started, and should one throw an exception, the shared flag is set and the exception is added to the shared list. All worker tasks poll the flag to determine whether it has been set; when they see that it has been set, they voluntarily terminate. If any additional exceptions occur on concurrent tasks before they are able to voluntarily terminate, those further exceptions are also added to the shared list. Once all of the concurrent worker tasks terminate, an aggregate exception object is created by the coordination task containing references to each of the exceptions in the shared list and then the aggregate exception is thrown by the coordination task.

[005] In another implementation, when the concurrent work is spawned, one or more exception handler descriptors are supplied: each such descriptor is comprised of an exception type and a handler function. For each respective exception that occurs on a given concurrent worker task, the task determines if the respective exception type is of a kind of handled by the one or more exception handler functions provided. If the respective exception is of such a type, then the particular handler function is run and the exception is considered handled. Any exceptions that remain unhandled are then processed in the manner described above.

[006] In one implementation, a sequential process is provided that uses concurrent worker tasks, with a collection of input data being processed to produce a corresponding collection of output results. The exceptions are interleaved alongside non-exceptional output results. The collection of output results is made

available to the sequential process. The worker tasks continue to produce the output results completely, as exceptions are encountered. This technique defers the evaluation of exceptions, permitting them to be handled during the consumption of the results rather than as part of the production of said output results.

5 [007] In another implementation, a single particular exception is selected that represents the multiple exceptions, leading to only one exception being thrown out of the coordination task. This is used when many homogeneous exceptions are caught and/or when a clear preference can be determined by the system.

[008] This Summary was provided to introduce a selection of concepts in a
10 simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

15 [009] Figure 1 is a diagrammatic view of a computer system of one implementation.

[010] Figure 2 is a diagrammatic view of a concurrent exception handling application of one implementation operating on the computer system of Figure 1.

[011] Figure 3 is a high-level process flow diagram for one implementation of the
20 system of Figure 1.

[012] Figure 4 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in catching and forwarding from the concurrent tasks to a coordination task.

[013] Figure 5 is a process flow diagram for one implementation of the system of
25 Figure 1 illustrating the stages involved in aggregating many exceptions into a single aggregate exception object.

[014] Figure 6 is a process flow diagram for one implementation of the system of Figure 1 illustrating the stages involved in terminating concurrent tasks after the first concurrent exception has been thrown.

[015] Figure 7 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in collapsing exceptions that occurred concurrently on separate tasks into fewer exceptions.

5 [016] Figure 8 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in handling multiple exceptions once gathered in the aggregate and rethrown from the coordination task.

[017] Figure 9 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in providing a deferred concurrent exception process.

10 [018] Figure 10 is a process flow diagram for one implementation of the system of Figure 1 that illustrates the stages involved in selecting an exception to raise out of many exceptions.

DETAILED DESCRIPTION

15 [019] For the purposes of promoting an understanding of the principles of the invention, reference will now be made to the embodiments illustrated in the drawings and specific language will be used to describe the same. It will nevertheless be understood that no limitation of the scope is thereby intended. Any alterations and further modifications in the described embodiments, and any further applications of the principles as described herein are contemplated as would
20 normally occur to one skilled in the art.

[020] The system may be described in the general context as an application that provides concurrent exception handling, but the system also serves other purposes in addition to these. In one implementation, one or more of the techniques described herein can be implemented as features within a framework program such
25 as MICROSOFT® .NET Framework, or from any other type of program or service that handles concurrent exceptions occurring in programs.

[021] In one implementation, a concurrent exception handling system is provided that can manage exceptions that occur concurrently on many worker tasks. The exceptions can be handled in various ways and at various times, such as by
30 aggregating exceptions and throwing from a coordination task after all concurrent worker tasks finish, such as placing exceptions interleaved among non-exceptional

data in the output stream for deferred handling, and so on. The term “task” as used herein means a logical unit of work which may or may not result in an independent thread. The term “coordination task” as used herein means a task that is responsible for coordinating multiple worker tasks. The term “worker task” as used
5 herein means a task that is responsible for carrying out one or more particular logical units of work associated with a higher level program operation.

[022] As shown in Figure 1, an exemplary computer system to use for implementing one or more parts of the system includes a computing device, such as computing device 100. In its most basic configuration, computing device 100
10 typically includes at least one processing unit 102 and memory 104. Depending on the exact configuration and type of computing device, memory 104 may be volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. This most basic configuration is illustrated in Figure 1 by dashed line 106.

15 [023] Additionally, device 100 may also have additional features/functionality. For example, device 100 may also include additional storage (removable and/or non-removable) including, but not limited to, magnetic or optical disks or tape. Such additional storage is illustrated in Figure 1 by removable storage 108 and non-removable storage 110. Computer storage media includes volatile and nonvolatile,
20 removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Memory 104, removable storage 108 and non-removable storage 110 are all examples of computer storage media. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash
25 memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can accessed by device 100. Any such computer storage media may be part of device 100.

30 [024] Computing device 100 includes one or more communication connections 114 that allow computing device 100 to communicate with other

computers/applications 115. Device 100 may also have input device(s) 112 such as keyboard, mouse, pen, voice input device, touch input device, etc. Output device(s) 111 such as a display, speakers, printer, etc. may also be included. These devices are well known in the art and need not be discussed at length here. In one

5 implementation, computing device 100 includes concurrent exception handling application 200. Concurrent exception handling application 200 will be described in further detail in Figure 2.

[025] Turning now to Figure 2 with continued reference to Figure 1, a concurrent exception handling application 200 operating on computing device 100 is

10 illustrated. Concurrent exception handling application 200 is one of the application programs that reside on computing device 100. However, it will be understood that concurrent exception handling application 200 can alternatively or additionally be embodied as computer-executable instructions on one or more computers and/or in different variations than shown on Figure 1. Alternatively or additionally, one or
15 more parts of concurrent exception handling application 200 can be part of system memory 104, on other computers and/or applications 115, or other such variations as would occur to one in the computer software art.

[026] Concurrent exception handling application 200 includes program logic 204, which is responsible for carrying out some or all of the techniques described herein.

20 Program logic 204 includes logic for providing a system that performs operations for an application concurrently on multiple concurrent worker tasks 206; logic for catching exceptions that occur in the concurrent tasks 208; logic for forwarding exceptions to the shared exception list and setting the shared flag 210; logic for producing an aggregation of the exceptions from the coordination task 212; logic
25 for optionally collapsing exceptions into fewer exceptions 214; logic for optionally selecting a single exception to raise out of multiple exceptions 216; and other logic for operating application 220. In one implementation, program logic 204 is operable to be called programmatically from another program, such as using a single call to a procedure in program logic 204.

30 [027] Turning now to Figures 3-10 with continued reference to Figures 1-2, the stages for implementing one or more implementations of concurrent exception

handling application 200 are described in further detail. Figure 3 is a high level process flow diagram for concurrent exception handling application 200. In one form, the process of Figure 3 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 240 with providing a system that performs operations for an application concurrently (i.e. executing multiple concurrent worker tasks on multiple OS threads) (stage 242). The system recognizes exceptions that occur in the concurrent tasks (stage 244), forwards them to a coordination task, and aggregates them (e.g. in an aggregate exception object or other structure) (stage 246). The system propagates the aggregate from the coordination task once all concurrent worker tasks have completed, enabling the program to handle the concurrent exceptions that happened during the operation at a proper time (stage 248). The logic performed to handle any given exception within the aggregate can be equivalent to the logic it would perform on a sequential program, provided it is capable of recognizing the aggregate exception type and retrieving individual exceptions from within. The process ends at end point 250.

[028] Figure 4 illustrates one implementation of the stages involved in catching and forwarding exceptions from the concurrent worker tasks to a coordination task. In one form, the process of Figure 4 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 270 with providing a system that performs operations for an application concurrently on multiple worker tasks (e.g. running in separate OS threads) (stage 272). The system catches exceptions that occur in the concurrent worker tasks (stage 274) and forwards the recognized exceptions from the concurrent tasks to a coordination task (stage 276), e.g. by placing them into a shared location. The coordination task is then responsible for tracking the received exceptions from other concurrent worker tasks, if applicable (stage 278). The process ends at end point 280.

[029] Figure 5 illustrates one implementation of the stages involved in aggregating many exceptions into a single aggregate exception object. In one form, the process of Figure 5 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 290 with providing an aggregate exception object that can store multiple exceptions and related details (stage 292).

In one implementation the aggregate exception object can store the stack trace, thread id, and/or API specific information about the concurrent operations which threw the forwarded exceptions to help debuggability, etc. (stage 292). In other implementations, additional and/or fewer details can be stored. The system
5 populates the aggregate exception object with exceptions that occurred in concurrent worker tasks (optionally in the order in which they were thrown by assigning a global timestamp ordering) (stage 294). The system then throws the aggregate exception, enabling an exception handler in the program to handle it at a proper time (stage 296). The process ends at end point 298.

10 [030] Figure 6 illustrates one implementation of the stages involved in terminating concurrent worker tasks after the first concurrent exception has been thrown. In one form, the process of Figure 6 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 310 with providing a coordination task that initializes a shared flag and a shared location
15 (e.g. a list of exceptions) (stage 312). The coordination task initiates one or more concurrent worker tasks (e.g. by assigning logical units of work called tasks to OS threads, processes, etc.) (stage 314). As soon as an exception is thrown by such a task, the shared flag is set and the exception is added to the shared location (stage 316). When a concurrent worker task subsequently polls the flag to see if it was
20 set, it voluntarily terminates (stage 318). If any further exceptions occur before a respective task terminates, those exceptions are also logged to the shared location (stage 320). The coordination task then waits for all concurrent tasks to terminate (stage 322). An aggregate exception object is created out of the logged exceptions and is then thrown by the coordination task for proper handling (stage 324). In one
25 implementation, an optional mechanism is provided for telling the user which operations completed, which operations threw an exception, and/or which ones were not processed. The process ends at end point 326.

[031] Figure 7 illustrates one implementation of the stages involved in collapsing exceptions that occurred concurrently on separate worker tasks into fewer
30 exceptions. In one form, the process of Figure 7 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 340

with aggregating exceptions that occurred on separate tasks (stage 342). The aggregated exceptions are analyzed to see if any can be eliminated as redundant, such as because they are homogeneous (i.e. of the same exception type and the same root cause) (stage 344). The system eliminates the redundant exceptions so that only the exceptions necessary to convey the issues remain in the aggregated exceptions structure (stage 346). If all exceptions are homogeneous, an implementation may decide to propagate just one of the exceptions instead of aggregating into an aggregate exception object. The process ends at end point 348.

[032] Figure 8 illustrates one implementation of the stages involved in handling

multiple exceptions once gathered in the aggregate and rethrown from the coordination task. In one form, the process of Figure 8 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 370 with receiving exception results (i.e. catching an aggregate exception object that was thrown) (stage 372). The catch handler invokes an API and supplies one or more exception handler functions (stage 374), each of which is comprised of an exception type and a handler function. For each of the exceptions (decision point 376), the system determines if the exception is one of the kinds handled by the handler function(s) (decision point 378). If so, the particular handler is run and the exception is marked as handled (stage 380). Once all of the exceptions are processed to see if there are handlers, the system determines if any unhandled exceptions are still present (decision point 382). If none remain, then the process ends at end point 386. If one exception remains, then the system rethrows the specific exception; if more than one exception remains, a new aggregate exception object is constructed with just the unhandled exceptions and the new aggregate exception object is rethrown (stage 384). The process ends at end point 386.

[033] Figure 9 illustrates one implementation of the stages involved in providing a deferred concurrent exception process. In one form, the process of Figure 9 is at least partially implemented in the operating logic of computing device 100. The process begins at start point 400 with a sequential process that uses concurrent worker tasks. A collection of input data is processed to produce a corresponding

collection of output results (stage 402). The collection contains either an actual output value or an exception value for each result, with the actual output values being interleaved with the exception values (stage 402). In one implementation, the output results can contain the completed results, aggregated exceptions, and/or an indicator that the work was not processed.

[034] In one implementation, as exceptions are interleaved among ordinary output, the system provides the option of preserving the original input ordering in the resulting output. Suppose there is an input of {a,b,c,x,y,z}, and the task is to map elements via some projection function f . The system can provide the capability of ensuring output ordering corresponds to input ordering, such that the results will always be: { $f(a),f(b),f(c),f(x),f(y),f(z)$ }. Similarly, when exceptions are deferred, the system can ensure the resultant exception ends up in the corresponding output element position. If exceptions $e1$ and $e2$ happened while executing $f(b)$ and $f(y)$, but all others succeeded, for instance, then the results will be:

{ $f(a),e1,f(c),f(x),e2,f(z)$ }.

[035] The collection of output results is made available to the sequential process, and the worker tasks continue to produce the output results completely, as exceptions are encountered (stage 404). In one implementation, the output results can then be processed as usual, which causes any exceptions to be rethrown lazily during processing as exceptional output elements are seen. This allows the exceptions to be handled as part of the consumption rather than the production of output. In another implementation, exceptions may be written to the output results first, and then if an exception of a certain type occurs, the operation stopped. For example, if a serious exception occurs that would affect the result of the computation, then it could be better to stop the computation than to continue writing exceptions to the output results.

[036] A gather-exceptions operation is optionally provided to accumulate all exceptions in the output results into a list which can be further analyzed, logged, etc. (stage 406). An eliminate-exceptions operation is optionally provided to remove exception values in the output results and optionally store them in a list which can be further analyzed (stage 408). The process ends at end point 410.

[037] Figure 10 illustrates one implementation of the stages involved in selecting an exception to raise out of many exceptions. In one form, the process of Figure 10 is at least partially implemented in the operating logic of computing device 100.

The process begins at start point 430 with aggregating multiple exceptions that occurred on concurrent worker tasks (stage 432). The system selects one exception that represents the multiple exceptions (e.g. one that is heuristically determined to be most important, one determined programmatically by a policy, one that was thrown first, etc.) (stage 434). In one implementation, the selected exception is an exception that was not one of the multiple exceptions. In another implementation, the selected exception is selected out of the multiple exceptions as one exception to represent the others because the multiple exceptions all have a same root problem. The system then throws the selected exception (stage 436), discarding the rest. The process ends at end point 438.

[038] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims. All equivalents, changes, and modifications that come within the spirit of the implementations as described herein and/or by the following claims are desired to be protected.

[039] For example, a person of ordinary skill in the computer software art will recognize that the client and/or server arrangements, user interface screen content, and/or data layouts as described in the examples discussed herein could be organized differently on one or more computers to include fewer or additional options or features than as portrayed in the examples.

What is claimed is:

1. A computer-readable medium having computer-executable instructions for causing a computer to perform steps comprising:
 - catch exceptions that occur in concurrent worker tasks;
 - 5 forward the caught exceptions from the concurrent worker tasks to a coordination task;
 - aggregate the caught exceptions into an aggregation structure; and
 - rethrow the aggregation structure from the coordination task so that the caught exceptions can be handled by the program at a proper time.
- 10 2. The computer-readable medium of claim 1, wherein the aggregation structure is an aggregation exception object.
3. The computer-readable medium of claim 2, wherein the aggregate exception object is provided to an exception handler so the exceptions can be handled by the program at the proper time.
- 15 4. The computer-readable medium of claim 1, wherein the coordination task is responsible for tracking the recognized exceptions from the concurrent worker tasks.
5. The computer-readable medium of claim 1, wherein one or more of the caught exceptions are removed from the aggregation structure as being redundant
- 20 exceptions.
6. The computer-readable medium of claim 1, wherein the aggregation structure is an output stream that contains concurrent exceptions interleaved with normal output.
7. The computer-readable medium of claim 1, wherein the exceptions are only
- 25 aggregated into an aggregation structure if an unhandled exception occurs on one of the concurrent worker tasks.
8. The computer-readable medium of claim 1, wherein the proper time for aggregating and rethrowing the caught exceptions in the aggregation structure is when all of the concurrent worker tasks finish executing.
- 30 9. A method for aggregating many exceptions into a single aggregate exception object comprising the steps of:

providing an aggregate exception object that can store a plurality of exceptions;

populating the aggregate exception object with the plurality of exceptions as the plurality of exceptions occur in concurrent worker tasks; and

5 providing the aggregate exception object to an exception handler at a proper time.

10. The method of claim 9, wherein the aggregate exception object can also store related information about each of the plurality of exceptions.

11. The method of claim 10, wherein the related details include an identifier of
10 the thread which originally threw the exception.

12. The method of claim 10, wherein the related details include an original stack trace.

13. The method of claim 10, wherein the related details include application programming interface specific information about a concurrent operation that
15 resulted in a particular exception.

14. The method of claim 9, wherein the one or more of the plurality of exceptions are removed from the aggregate exception object as being redundant exceptions.

15. The method of claim 9, wherein the proper time at which the aggregate
20 exception object is provided to the exception handler is after all of the concurrent worker tasks have finished executing.

16. The method of claim 9, wherein the proper time at which the aggregate exception object is provided to the exception handler is after a first exception occurs on one of the concurrent worker tasks.

25 17. A computer-readable medium having computer-executable instructions for causing a computer to perform the steps recited in claim 9.

18. A method for terminating concurrent worker tasks after a first concurrent exception occurs comprising the steps of:

30 providing a coordination task that initializes a shared flag and a shared location;

starting concurrent worker tasks;

as soon as a first concurrent exception is thrown, setting the shared flag and logging the exception to the shared location;

when each respective concurrent worker task subsequently polls the flag to see if the flag was set, the respective task voluntarily terminates;

5 if any further exceptions occur before each respective task terminates, logging those further exceptions to the shared location;

waiting for all of the concurrent worker tasks to terminate; and

creating an aggregate exception object out of any exceptions logged to the shared location and then throwing the aggregate exception object for proper

10 handling.

19. The method of claim 18, wherein the shared location contains a collection of exceptions.

20. A computer-readable medium having computer-executable instructions for causing a computer to perform the steps recited in claim 18.

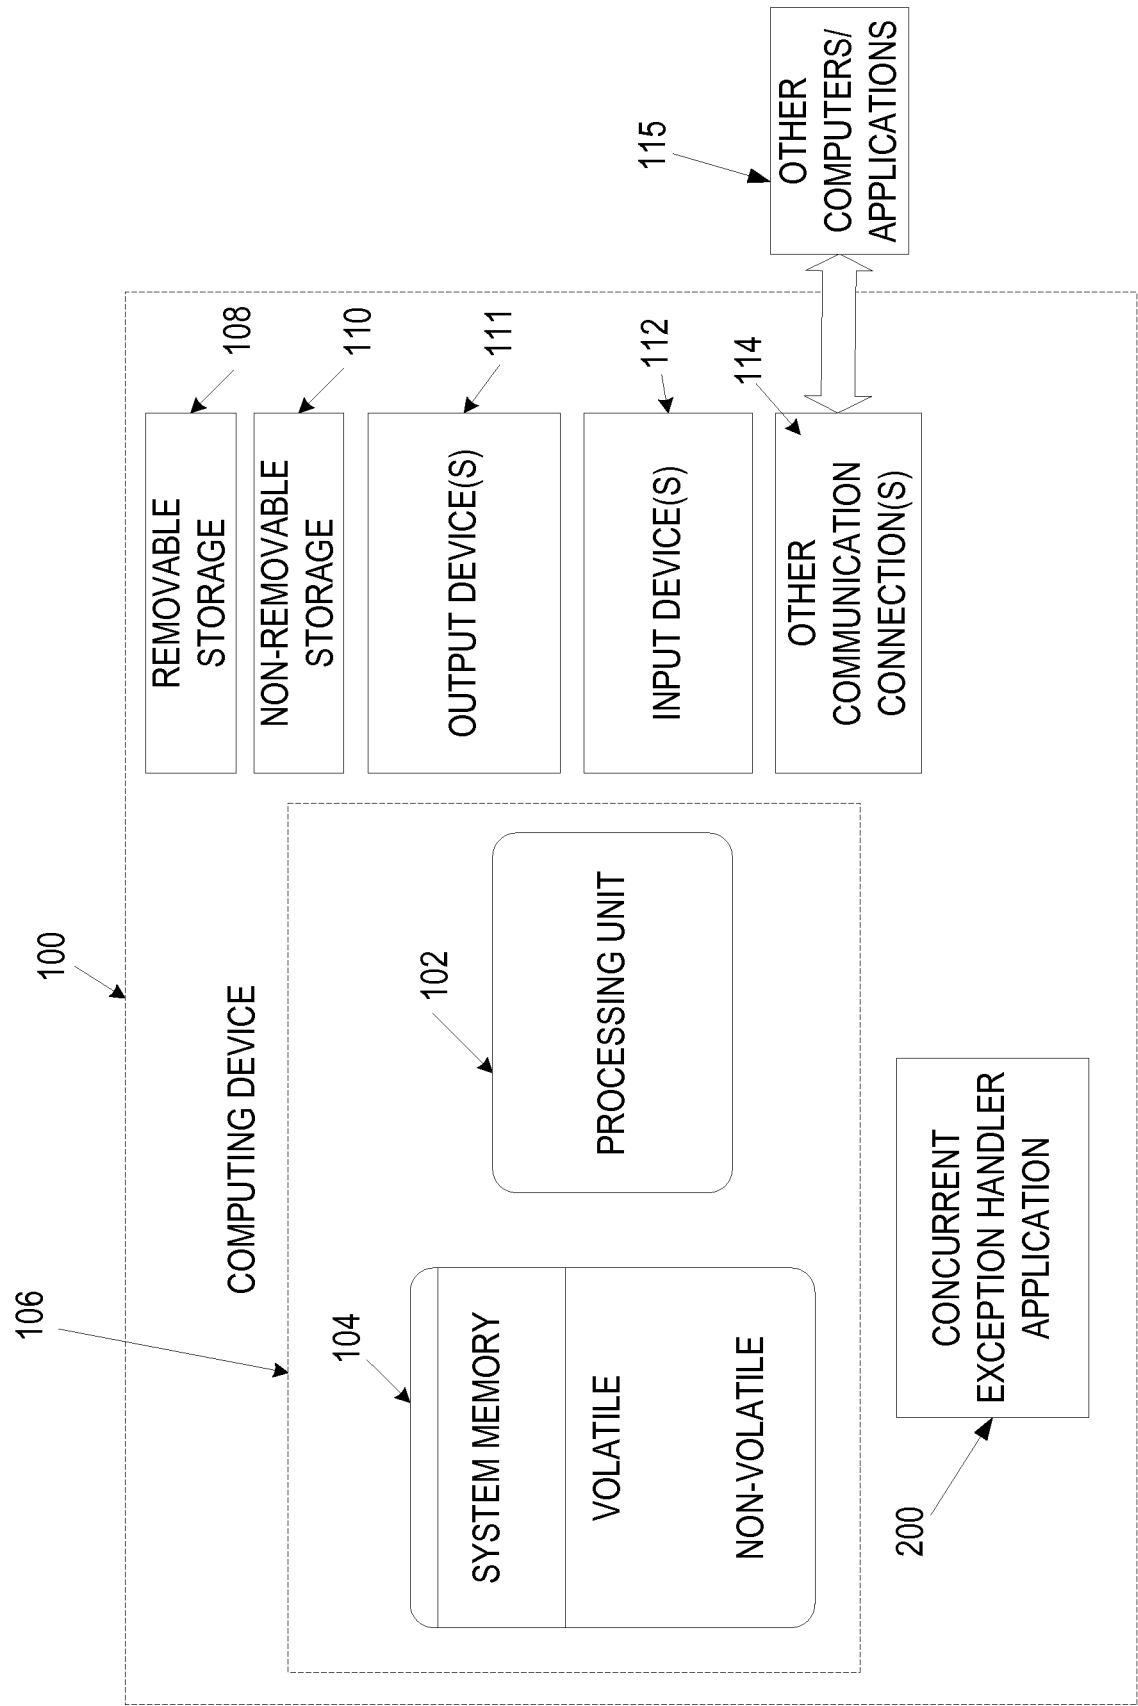


FIG. 1



FIG. 2

3 / 10

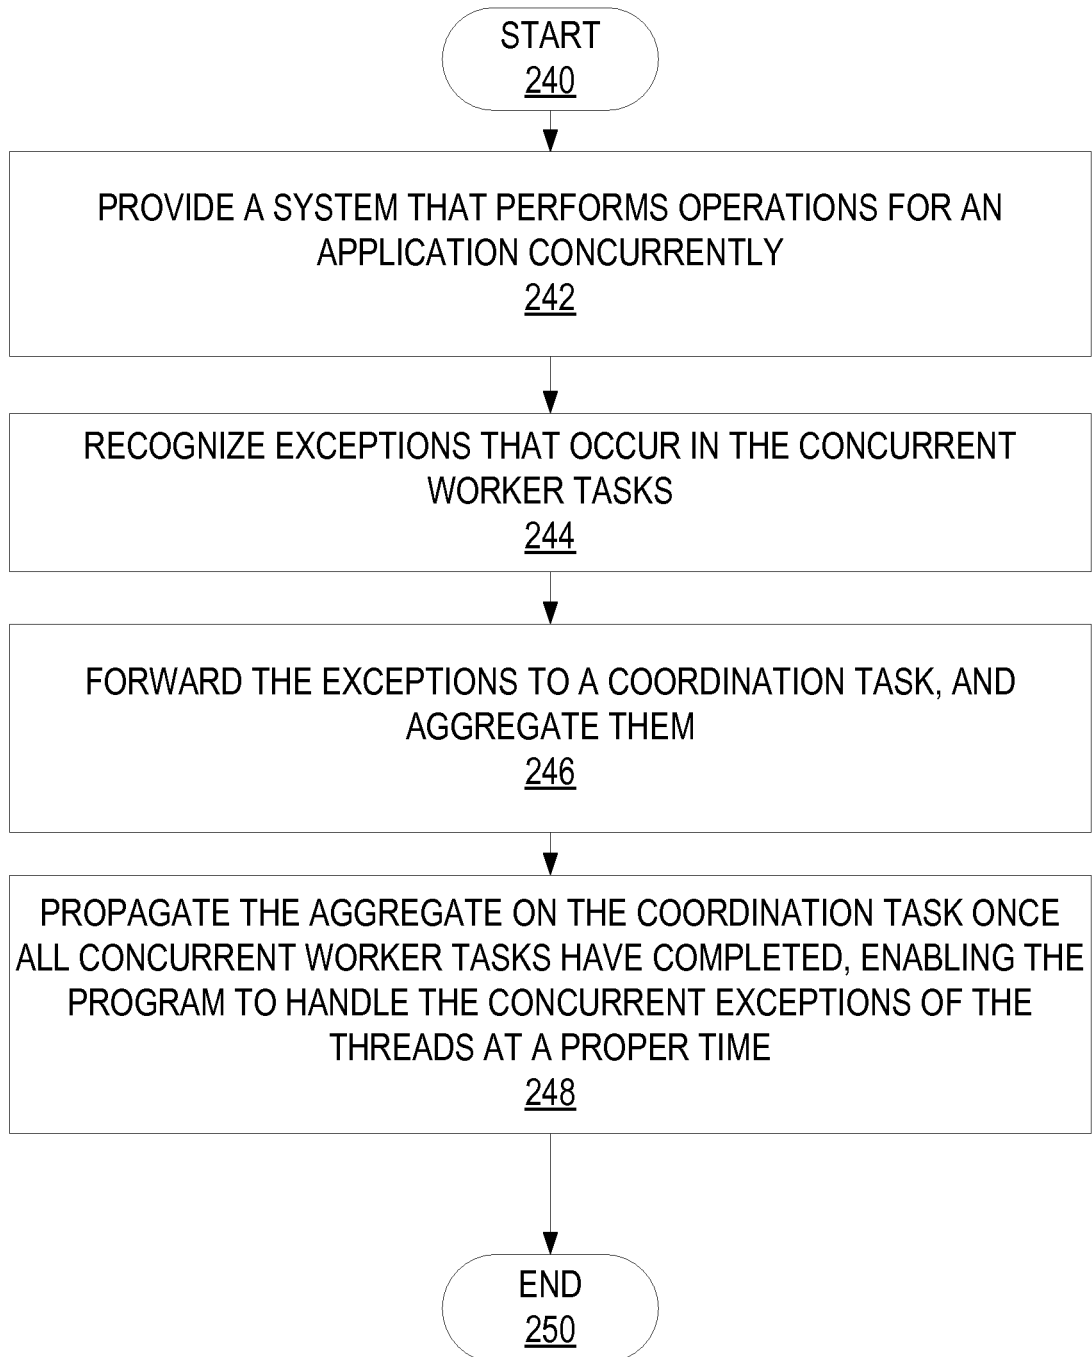


FIG. 3

4 / 10

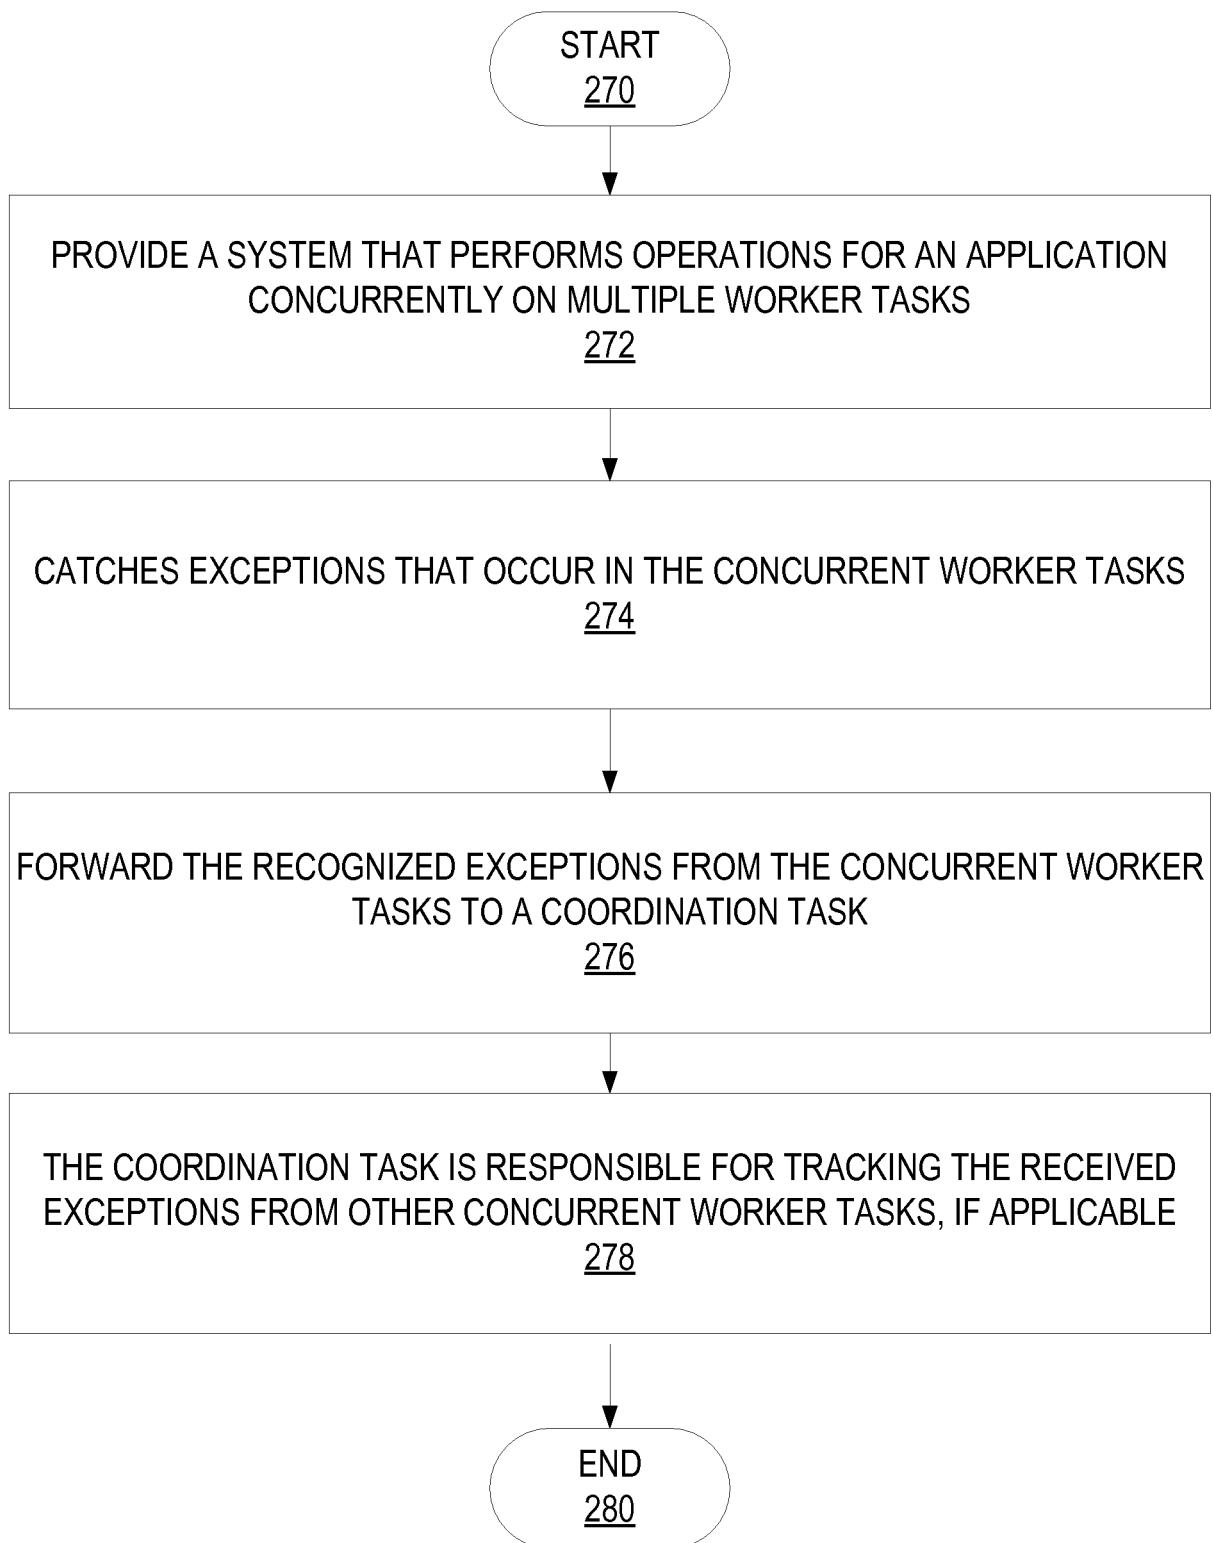


FIG. 4

5 / 10

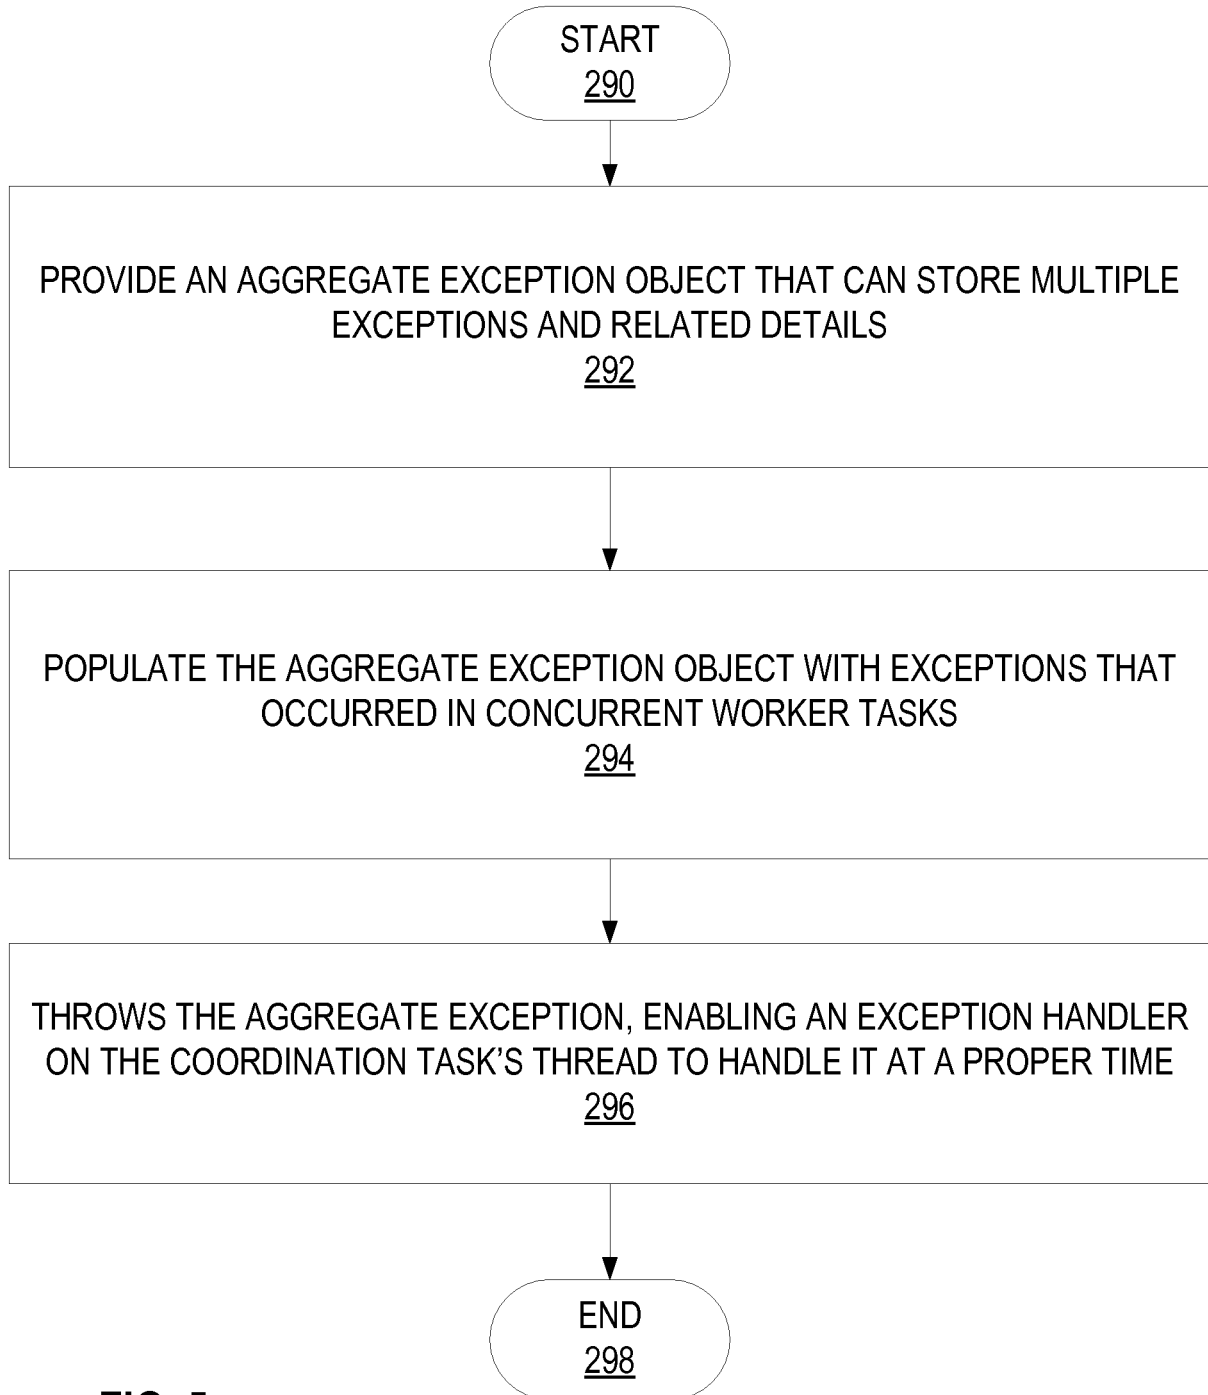


FIG. 5

6 / 10

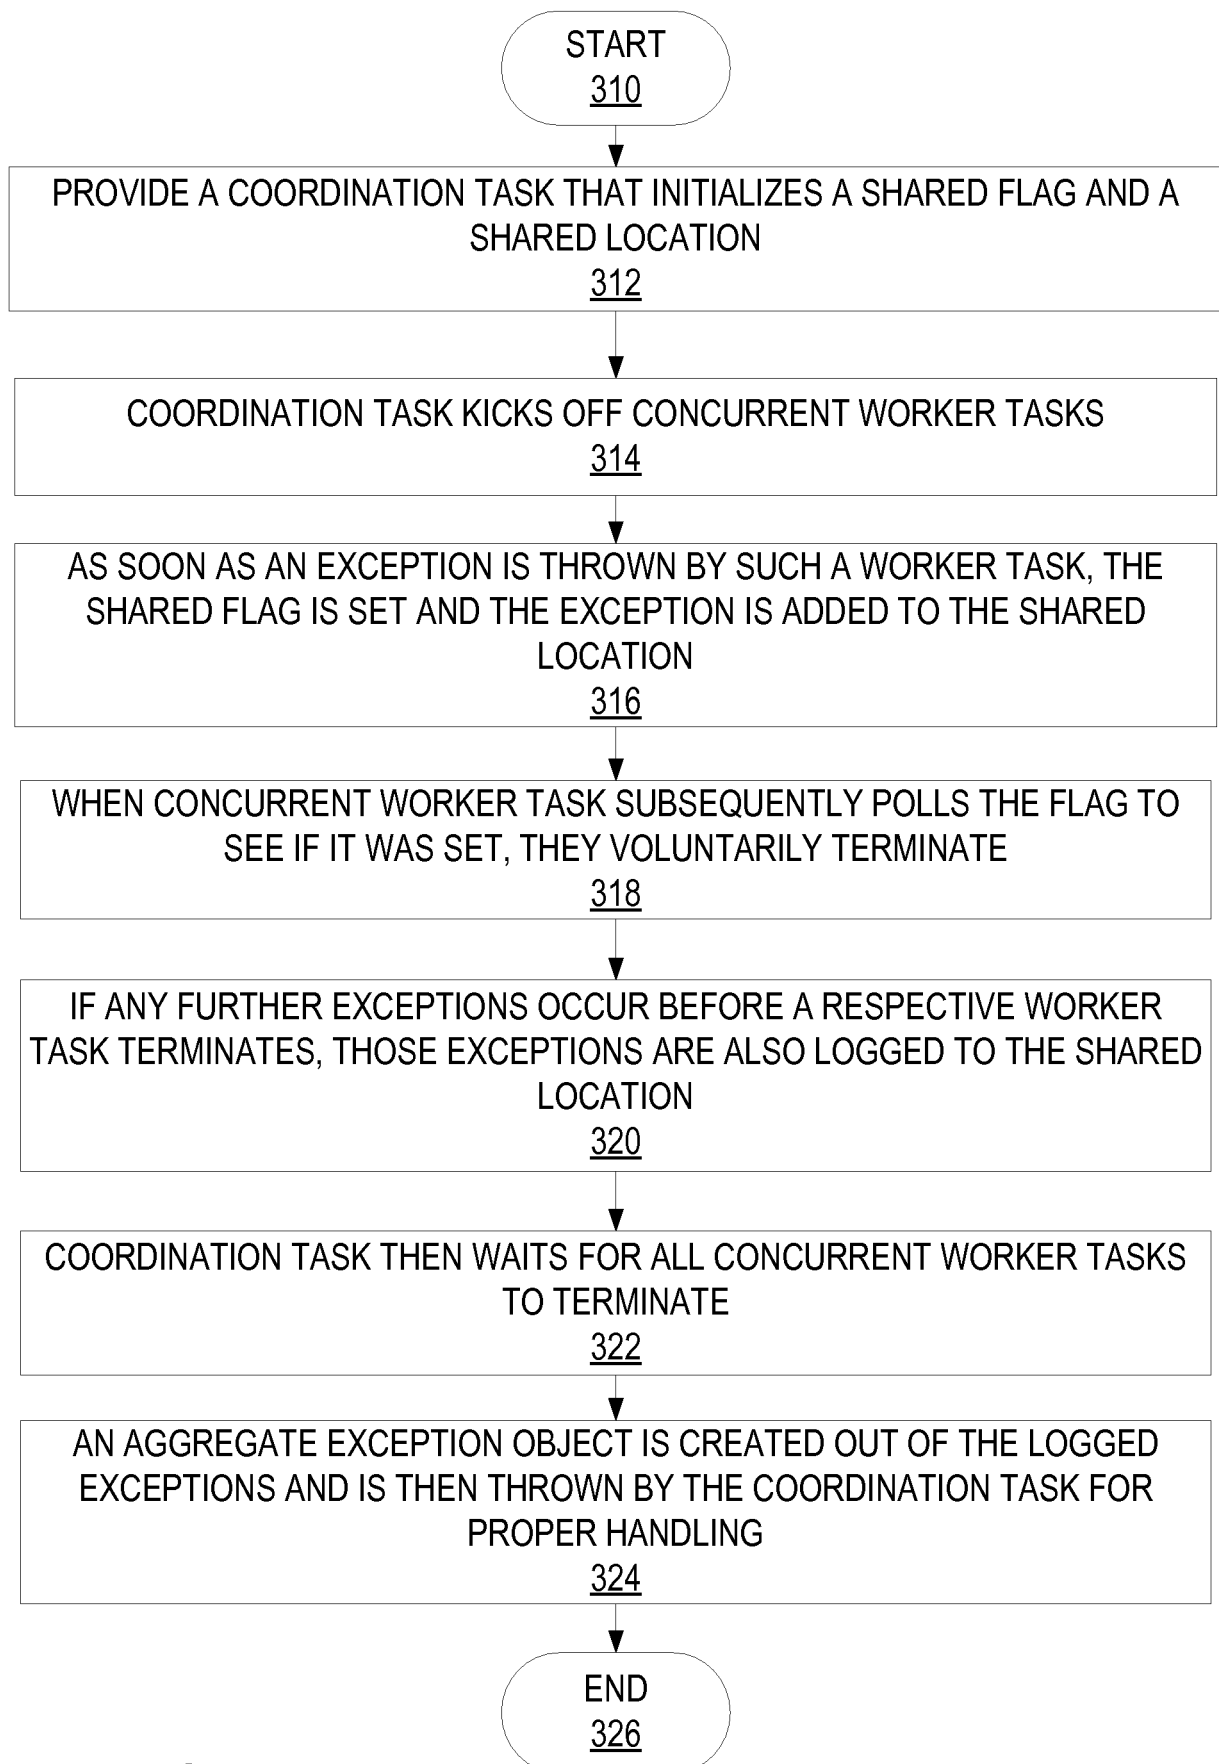


FIG. 6

7 / 10

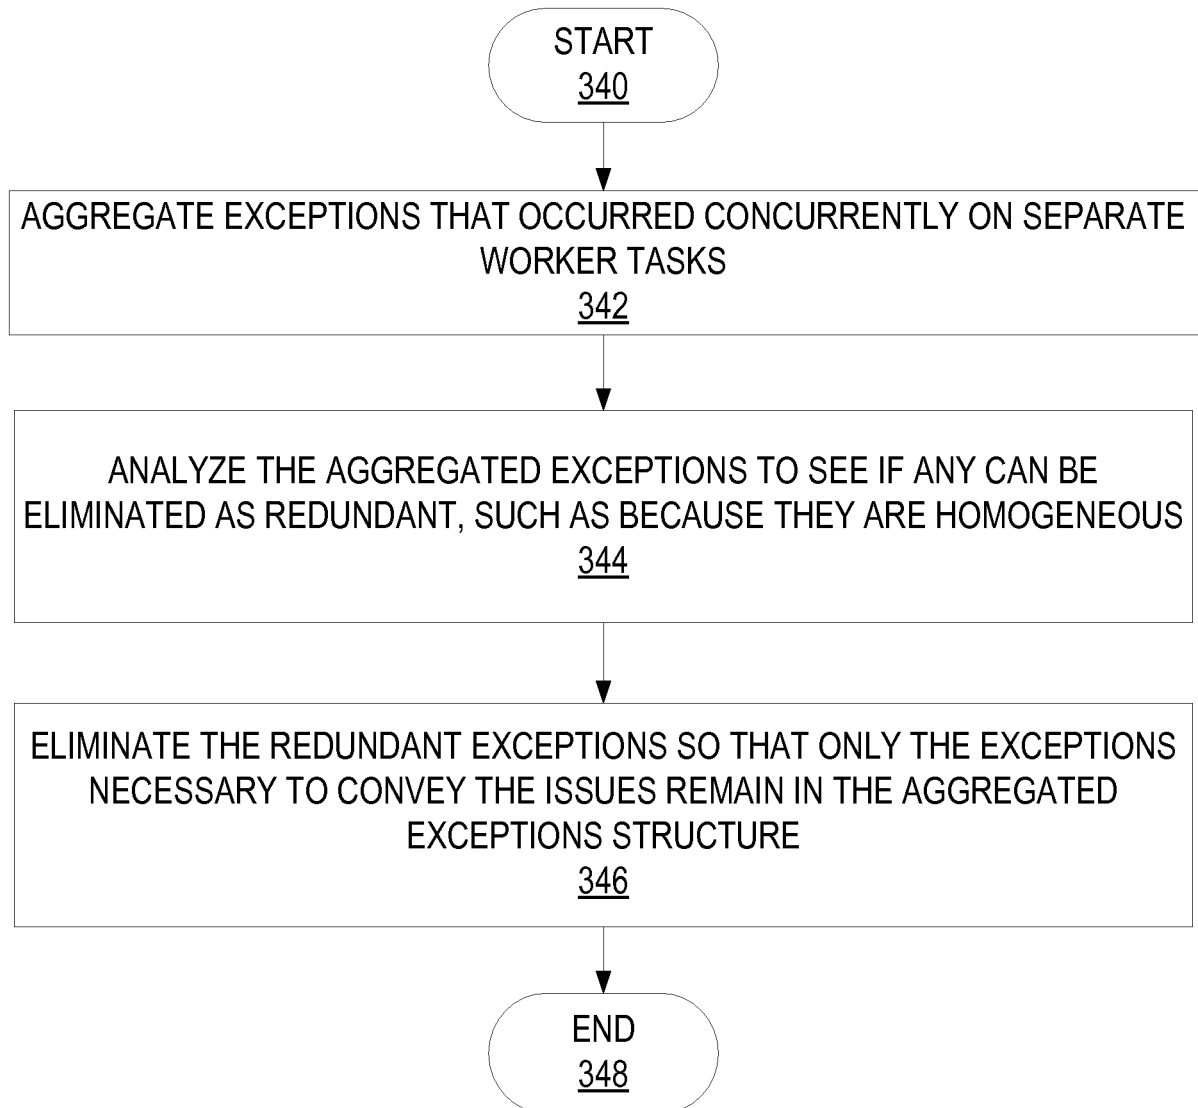


FIG. 7

8 / 10

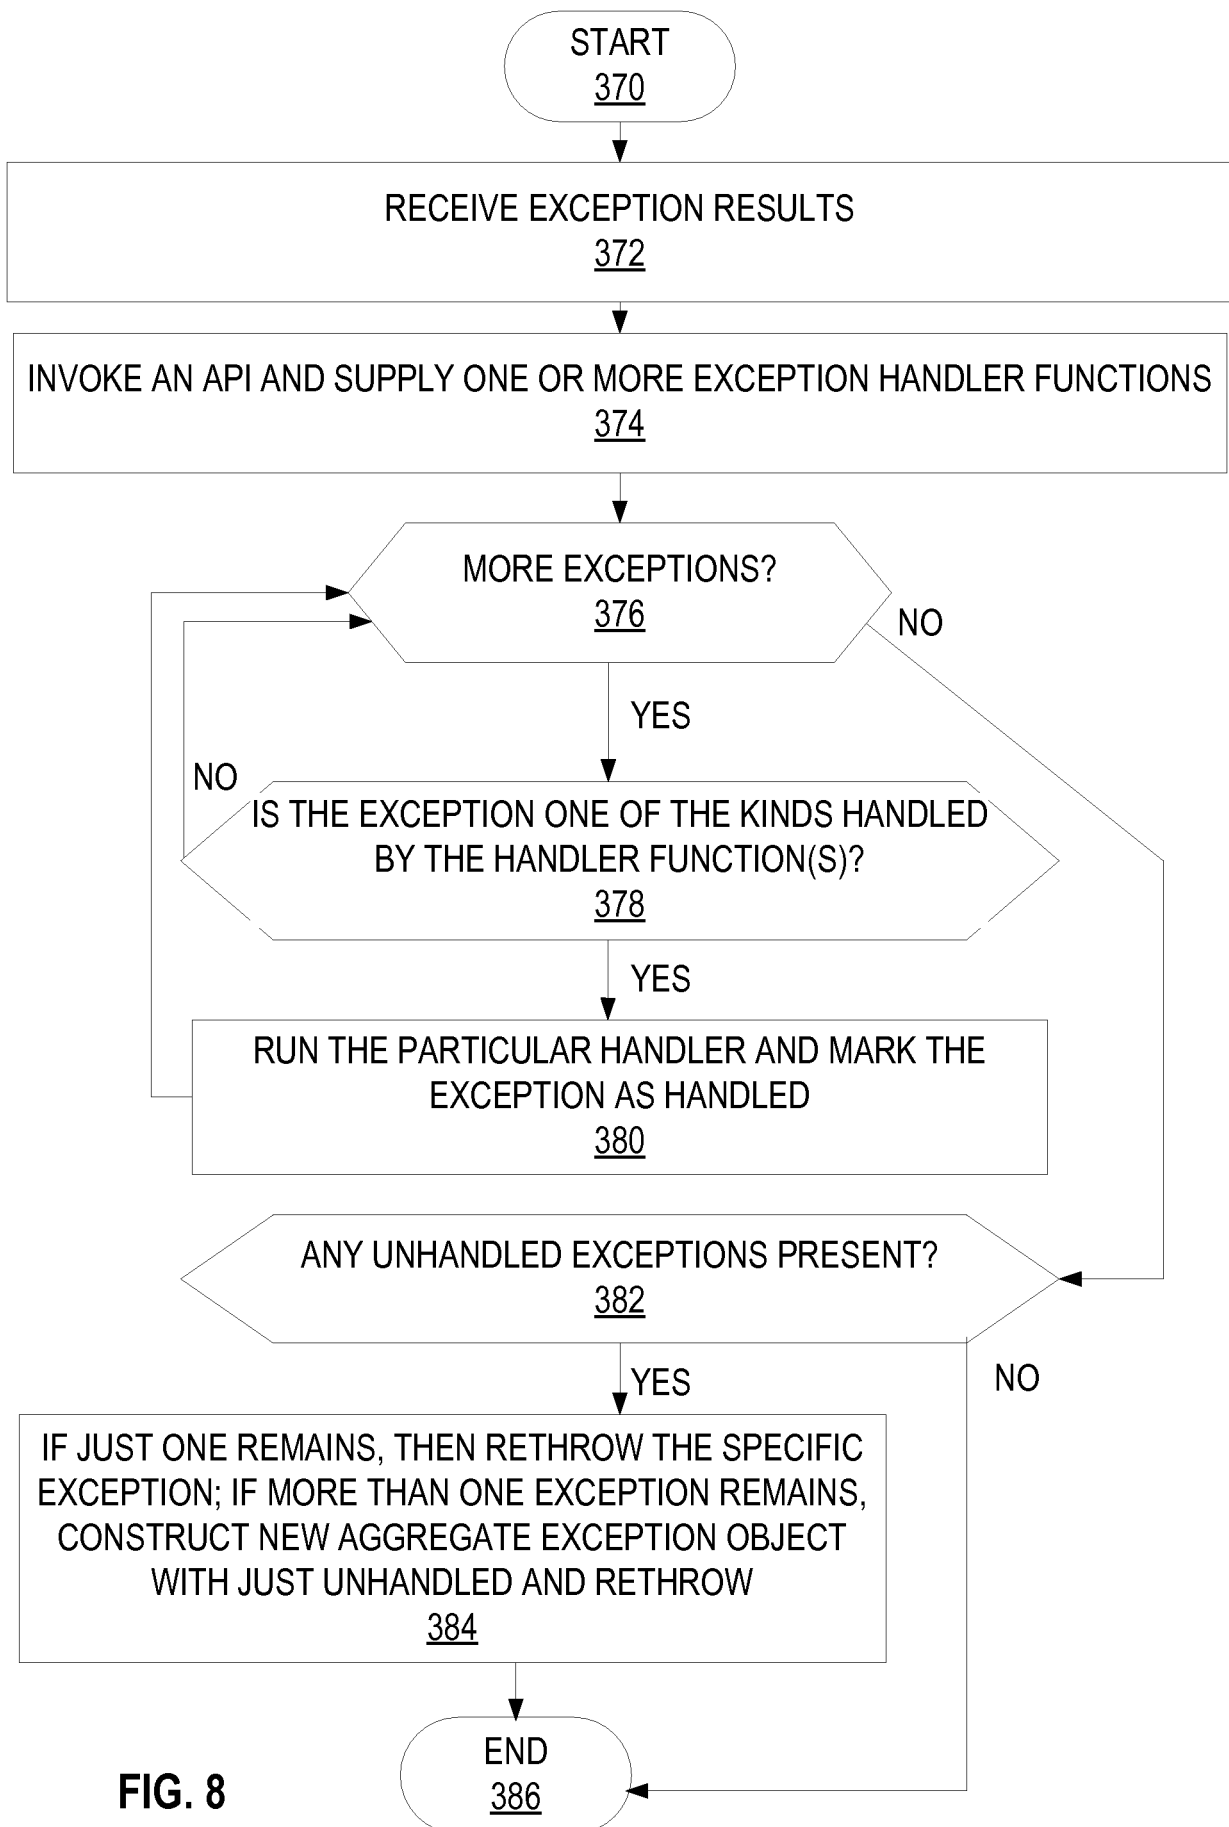


FIG. 8

9 / 10

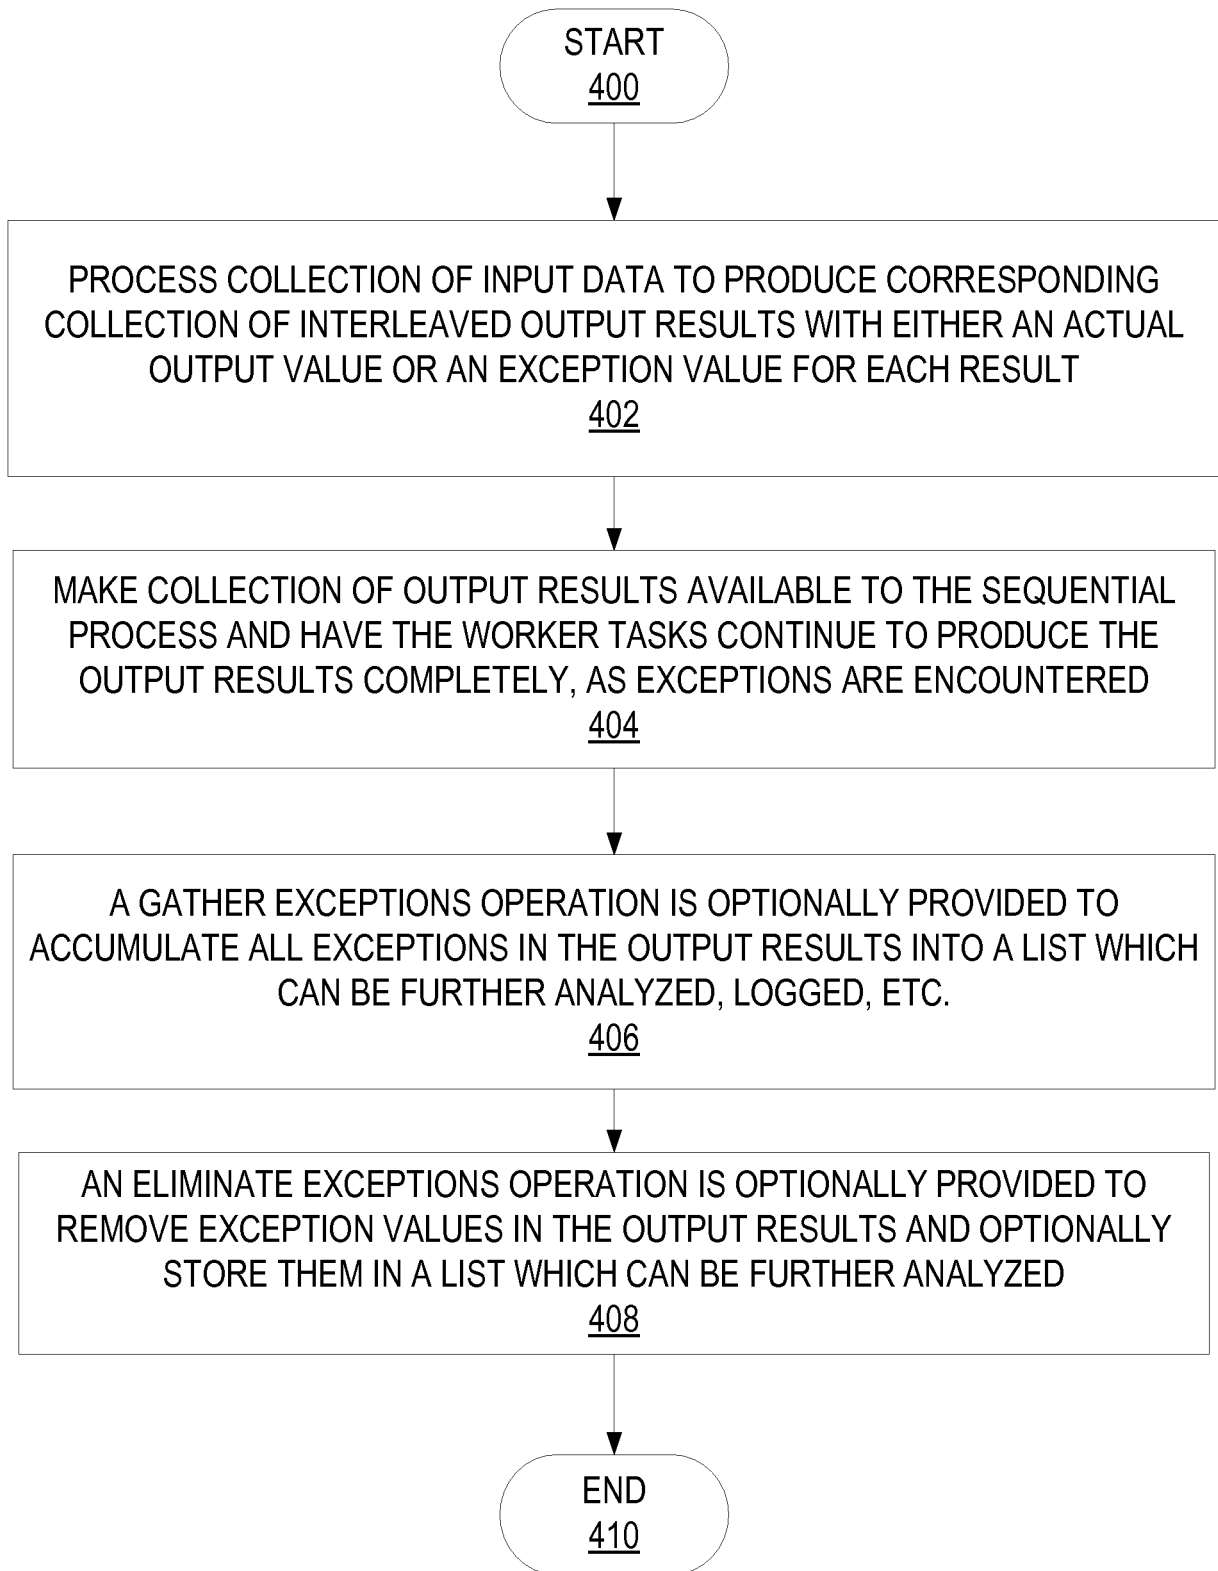


FIG. 9

10 / 10

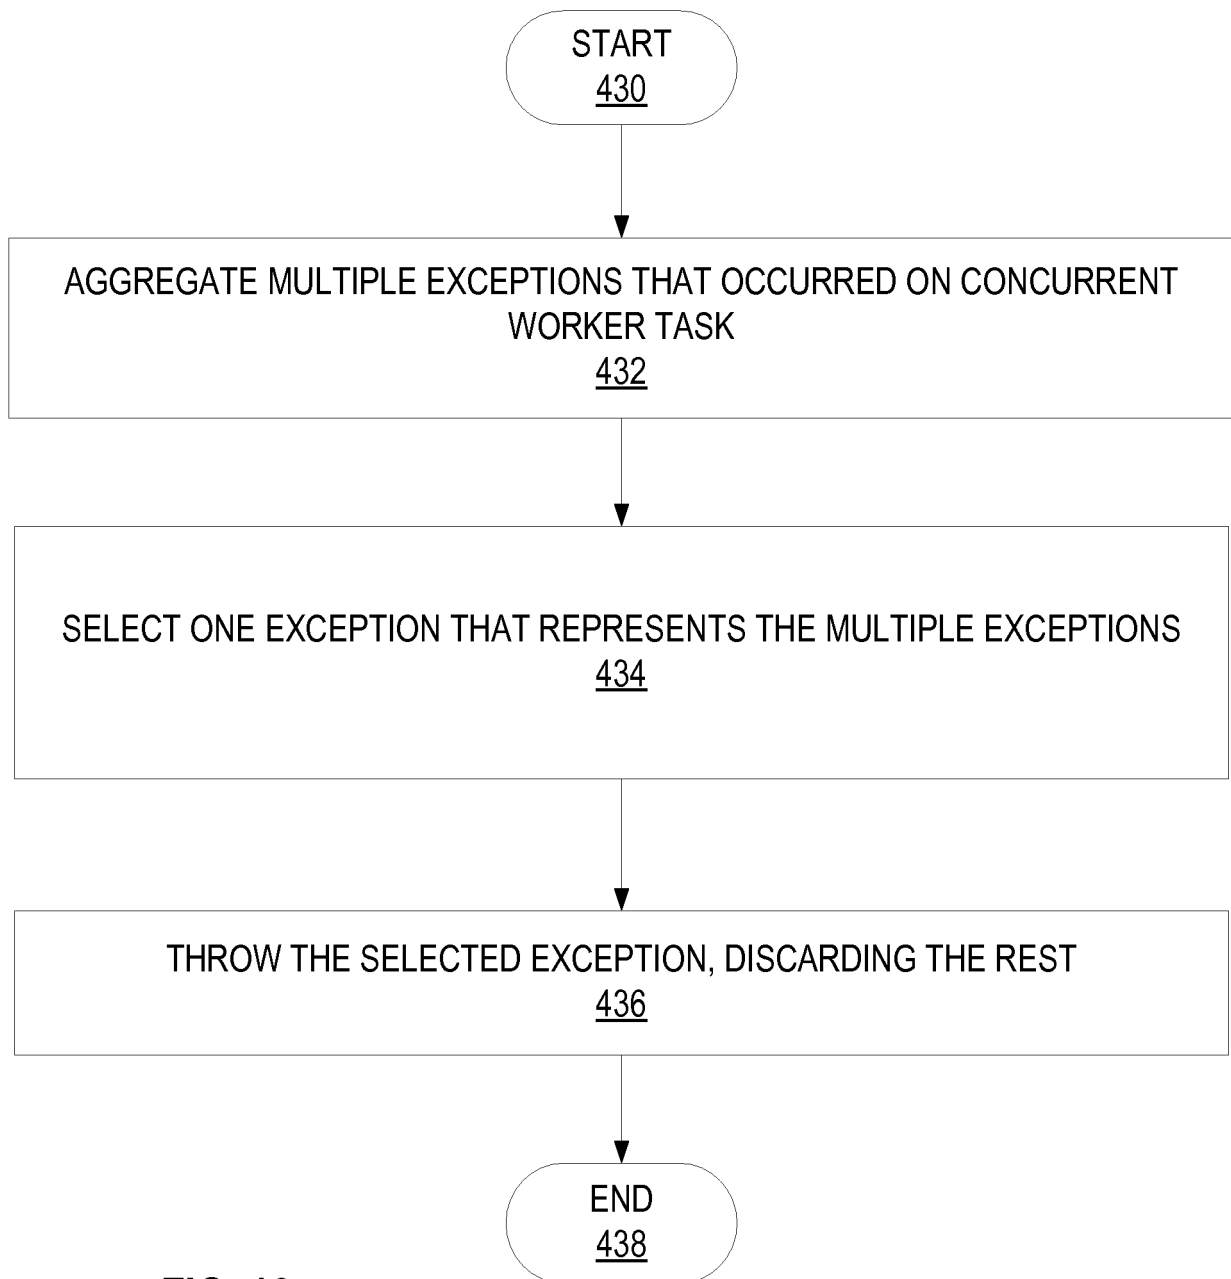


FIG. 10