

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 December 2002 (12.12.2002)

PCT

(10) International Publication Number
WO 02/099641 A1

- (51) International Patent Classification⁷: **G06F 11/00**
- (21) International Application Number: PCT/US02/15321
- (22) International Filing Date: 14 May 2002 (14.05.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/873,039 2 June 2001 (02.06.2001) US
- (71) Applicant: **REDBACK NETWORKS INC.** [US/US];
250 Holger Way, San Jose, CA 95134 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- (72) Inventor: **LAL, Sanjay**; 20380 Stevens Creek Blvd., #320, Cupertino, CA 95014 (US).
- Published:**
— with international search report
- (74) Agents: **MALLIE, Michael, J.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 12400 Wilshire Boulevard, 7th floor, Los Angeles, CA 90025 (US).
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: METHOD AND APPARATUS FOR MAINTAINING CONSISTENT DATA

(57) Abstract: A method and apparatus for maintaining consistent data is described. A computer implemented method comprises generating a first command for a set of network data to be executed on a local memory, executing a second command for the set of network data on a remote memory in response to generation of the first command, determining whether the second command has been executed successfully on the remote memory, (413) executing the first command on the local memory upon determining the second command is executed successfully, and generating an error upon determining the second command is not executed successfully (417).



WO 02/099641 A1

METHOD AND APPARATUS FOR MAINTAINING CONSISTENT DATA

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The invention relates to the field of communications. More specifically, the invention relates to communication networks.

Background of the Invention

[0002] Control card redundancy improves reliability of network elements. Having a secondary control card with the same information as the primary control cards enables a network element to withstand hardware and/or software failure on the primary card. In addition, the secondary control card enables maintenance of network elements. The primary control card can be pulled for repairs or testing because the secondary control card will take over operations.

[0003] To avoid interruptions in service, consistency must be maintained between control cards. Information received by the primary control card is passed to the secondary control card. After the information is stored on the primary control card it is copied to the secondary control card.

THIS PRIOR ART TECHNIQUE HAS THE DISADVANTAGE OF EXPOSING THE NETWORK ELEMENT TO DISRUPTION. A FAILURE MAY OCCUR BEFORE INFORMATION IS COPIED TO THE SECONDARY CONTROL CARD. WHEN THE SECONDARY CONTROL CARD TAKES OVER OPERATION OF THE NETWORK ELEMENT, THE SECONDARY CONTROL CARD WILL BE LACKING INFORMATION. THE NETWORK ADMINISTRATOR WILL HAVE TO DISCOVER THE INCONSISTENCIES AND CORRECT THESE INCONSISTENCIES. IN ADDITION, THE INCONSISTENCIES MAY CAUSE SERVICE TO BE DISRUPTED TO A CUSTOMER. BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the drawings:

[0005] Figure 1 is a diagram of a network element according to one embodiment of the invention.

[0006] Figure 2 is a diagram of the memory 109 and the memory 107 of Figure 1 according to one embodiment of the invention.

[0007] Figure 3 is a diagram of a configuration manager interacting with the memory 107 and 109 of Figure 1 according to one embodiment of the invention.

[0008] Figure 4 is a flowchart for storing information to backup memory according to one embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

[0009] In the following description, numerous specific details are set forth to provide a thorough understanding of the invention. However, it is understood that the invention may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the invention.

[0010] Figure 1 is a diagram of a network element according to one embodiment of the invention. In Figure 1, a network element 102 includes two control cards 103 and 105 and an interface 101. Examples of a network element include routers, switches, bridges, edge devices, etc. The control card 103 is connected to the control card 105. The control card 103 includes a processor 113 and a memory 109. The control card 105 also includes a processor 111 and a memory 107. During operation of the network element 102, one of the control cards 103, 105 is designated as the primary control card while the other is the secondary control card. If the primary control card fails then the secondary control card operates in place of the primary control card. Both of the control cards 103, 105 are connected to an interface 101.

[0011] The interface receives input from a user. Although the interface is shown as part of the network element, the interface can be a remote terminal in another embodiment of the invention. This input is sent from the interface to the primary control card. If the control card 103 is the primary control card, then the interface 101 transmits configurations, data, etc. to the control card 103. To maintain consistency between the

control cards, information stored on the primary control card is duplicated on the secondary control card. Continuing the example of the control card 103 as the primary control card, the control card 103 transmits information received from the interface 101 to the control card 105 before storing the information in its memory 109. The control card 105 receives information (configurations, network data, etc.) from the control card 103 and writes the information to its memory 107. Before information is written to the memory 107, the information is written to the memory 109.

[0012] In one embodiment of the invention, global mapping is implemented in the network element 102. Global mapping provides a unique address for hardware throughout the network element. With global mapping, the control card 103 can modify the memory 109 directly. Likewise, the control card 105 can modify the memory 107 directly.

[0013] Figure 2 is a diagram of the memory 109 and the memory 107 of Figure 1 according to one embodiment of the invention. In Figure 2, a dashed line 205 indicates a region 201 of the memory 107 being mapped to a region 203 of the memory 109. The memory management process will associate addresses of the region 201 to addresses of the region 203. In another embodiment of the invention, all of the memory 107 is mapped to the memory 109.

[0014] Figure 3 is a diagram of a configuration manager interacting with the memory 107 and 109 of Figure 1 according to one embodiment of the invention. In Figure 3, a configuration manager 301 receives configurations and processes the configurations. The configuration manager breaks down a configuration into a component or components and determines if a write, update or delete operation or operations is required. The configuration manager 301 then requests a write for write and update operations and a delete for delete operations. The request from the configuration manager 301 corresponds to an address in the memory 109, assuming the memory 109 is being used by the primary control card. Before the request is carried out or executed, an exception causes the operation to be performed on a corresponding address in the memory 107 as indicated by the dashed line 303.

[0015] Figure 4 is a flowchart for maintaining consistency of data between local and remote memory according to one embodiment of the invention. In Figure 4 at block 401, remote memory is mapped. It is assumed that local memory is on the primary control card and remote memory is on the secondary control card. At block 402, local memory is marked with a flag that will cause an exception. In another embodiment of the invention, local memory has read-only permission until a command is successfully performed on remote memory. In such an embodiment, once local memory is modified, read-only permission is reinstated. At block 403, configuration(s) are received from the interface 101 of Figure 1. At block 405, the configuration manager 301 of Figure 3 processes the configurations. At block 407, the configuration manager 301 submits modification requests (i.e., write or delete operations) corresponding to the processed configurations. At block 409, the submitted modification request(s) for local memory trigger an exception(s). In an embodiment of the invention, this exception is spawned by the operating system when attempts are made to modify given areas of local memory. In one such embodiment, given locations of local memory are associated with given exceptions. For example, referring to Figure 2, a write command is generated for the area 203 of the local memory 109. The exception spawned by the request to write to the area 203 is associated to the area 201 of the remote memory 107. In another embodiment of the invention, a write or delete is only allowed on local memory once remote memory is modified. Otherwise, read-only permission is enforced on local memory.

[0016] At block 411, the exception(s) modifies the remote memory in accordance with the requests (i.e., data is written or removed from the remote memory). At block 413, it is determined if the exception request was executed successfully. If the exception request was executed successfully, then at block 415 the request is executed on the local memory. If the exception request was not executed successfully on the remote memory, then at block 417 an error is generated. The error can be transmitted to an error parser or directly to the user via the interface.

[0017] In one embodiment of the invention, the exception are software interrupts. In another embodiment of the invention, these exceptions are hardware interrupts.

[0018] Insuring data is written or deleted from a secondary memory of a secondary control card before being written or deleted from a primary memory of a primary control

card avoids inconsistency of data between control cards. If an error occurs with the control card before the data is written to the primary memory, then the secondary control card has the most recent version of data or configurations. If the data is not written to the secondary memory, then the primary and secondary control cards still have the same information.

[0019] Furthermore, implementing the embodiments herein of the invention at the lowest level, such as with kernel exceptions, provides a simple backup mechanism with high performance. The complications of higher level backup applications are avoided. In addition, the secondary control card can take control immediately.

[0020] The network element 102 includes line cards in addition to the control cards 103 and 105. Line cards in the network element 102 and the control cards 103 and 105 include memories, processors, and/or Application Specific Integrated Circuit ("ASICs"). Such memories and the memories 107 and 109 include a machine-readable medium on which is stored a set of instructions (i.e., software) embodying any one, or all, of the methodologies described herein. Software can reside, completely or at least partially, within this memory and/or within the processor and/or ASICs. For the purpose of this specification, the term "machine-readable medium" shall be taken to include any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"), random access memory ("RAM"), magnetic disk storage media, optical storage media, flash memory devices, electrical, optical, acoustical, or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.), etc.

[0021] While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described. The method and apparatus of the invention can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting on the invention.

CLAIMS

What is claimed is:

1. A computer implemented method comprising:
generating a first command for a set of network data to be executed on a local memory;
executing a second command for the set of network data on a remote memory in response to generation of the first command;
determining whether the second command has been executed successfully on the remote memory;
executing the first command on the local memory upon determining the second command is executed successfully; and
generating an error upon determining the second command is not executed successfully.
2. The computer implemented method of claim 1 wherein the set of network data is a set of configurations.
3. The computer implemented method of claim 1 wherein the first command is a write command to an address of the local memory and the second command is a write command to an address of the remote memory.
4. The computer implemented method of claim 1 wherein the first command is a delete command for an address of the local memory and the second command is a delete command for an address of the remote memory.
5. A computer implemented method comprising:

mapping a first memory to a second memory;
receiving a set of configurations;
processing the set of configurations;
generating a set of commands for the first memory, the set of commands
corresponding to the processed set of configurations;
triggering an exception when beginning to execute the set of commands on the
first memory, wherein the exception performs the following:
executing the set of commands on the second memory in response to the
exception;
upon determining the set of commands are executed successfully on the
second memory, completing execution of the set of commands on
the first memory; and
upon determining the set of commands are not executed on the second
memory successfully, generating an error.

6. The computer implemented method of claim 5 wherein the set of configurations are for a set of network processes.
7. The computer implemented method of claim 5 wherein the error is displayed as a text message.
8. The computer implemented method of claim 5 wherein the error is passed to an error parser.
9. The computer implemented method of claim 5 wherein the mapping the first memory to the second memory comprises associating a set of addresses of the first memory to a set of addresses of the second memory.

10. The computer implemented method of claim 5 wherein the set of commands are a set of write commands.
11. The computer implemented method of claim 5 wherein the set of commands are a set of delete commands.
12. A computer implemented method comprising:
 - receiving a request to modify configuration data located at a local address in local memory in an active control card;
 - generating an exception when the configuration data located at the local address in the local memory is attempted to be modified, wherein the exception performs, within a processor in the active control card, the following:
 - modifying configuration data located at a remote address in remote memory in an inactive control card, wherein the configuration data located at the local address corresponds to the configuration data located at the remote address;
 - modifying the configuration data located at the local address in local memory in the active control card upon determining that the modification of the configuration data located at the remote address in the remote memory in the inactive control card was successful;
 - and
 - generating an error without modifying the configuration data located at the local address in local memory in the active control card upon determining that the modification of the configuration data located at the remote address in the remote memory in the inactive control card was not successful.

13. The computer implemented method of claim 12 wherein the request is to write the configuration data.
14. The computer implemented method of claim 12 wherein the request is to delete the configuration data.
15. The computer implemented method of claim 12 wherein the error is displayed on an interface.
16. The computer implemented method of claim 12 wherein the error is passed to an error parser.
17. An apparatus comprising:
 - an interface to receive a set of configurations;
 - a first control card coupled to the interface, the first control card having a first memory;
 - a second control card coupled to the interface and the first control card, the second control card to receive the set of configurations from the interface, to generate a set of commands for the set of configurations, to execute the set of commands on the first memory in response to an exception, and to execute the set of commands on a second memory of the second control card.
18. The apparatus of claim 17 wherein the first memory and the second memory are a first main memory and a second main memory.

19. The apparatus of claim 17 wherein the first memory and the second memory are a first storage and a second storage.
20. The apparatus of claim 17 wherein the first memory is mapped into the second memory.
21. The apparatus of claim 17 further comprising the second control card to generate an error if the set of commands are not executed successfully on the first memory.
22. The apparatus of claim 17 wherein the set of commands are a set of write commands.
23. The apparatus of claim 17 wherein the set of commands are a set of delete commands.
24. An apparatus comprising:
 - a processor to execute a configuration process, the configuration process to receive a set of configurations, process the set of configurations, and to submit a set of requests to the processor;
 - a first memory coupled to the processor, the processor to perform the set of requests on the first memory in response to an exception triggered by the set of requests; and
 - a second memory coupled to the first memory and the processor, the processor to perform the set of requests on the second memory if the set of requests are performed successfully on the first memory.

25. The apparatus of claim 24 wherein the first memory and the second memory are a first main memory and a second main memory.
26. The apparatus of claim 24 wherein the first memory and the second memory are a first storage and a second storage.
27. The apparatus of claim 24 wherein the set of requests are a set of write commands.
28. The apparatus of claim 24 wherein the set of requests are a set of delete commands.
29. The apparatus of claim 24 further comprising the processor to generate an error if the set of requests are not performed on the first memory successfully.
30. The apparatus of claim 24 further comprising a second processor coupled to the first processor, the first memory, and the second memory.
31. A machine-readable medium that provides instructions, which when executed by a set of processors of one or more processors, cause said set of processors to perform operations comprising:
- generating a first command for a set of network data to be executed on a local memory;
 - executing a second command for the set of network data on a remote memory in response to generation of the first command;
 - determining whether the second command has been executed successfully on the remote memory;

executing the first command on the local memory upon determining the second command is executed successfully; and
generating an error upon determining the second command is not executed successfully.

32. The machine-readable medium of claim 31 wherein the set of network data is a set of configurations.

33. The machine-readable medium of claim 31 wherein the first command is a write command to an address of the local memory and the second command is a write command to an address of the remote memory.

34. The machine-readable medium of claim 31 wherein the first command is a delete command for an address of the local memory and the second command is a delete command for an address of the remote memory.

35. A machine-readable medium that provides instructions, which when executed by a set of processors of one or more processors, cause said set of processors to perform operations comprising:

- associating a first memory to a second memory;
- receiving a set of configurations;
- processing the set of configurations;
- generating a set of commands for the first memory, the set of commands corresponding to the processed set of configurations;
- triggering an exception when processing the set of commands for the first memory, wherein the exception performs the following:
 - executing the set of commands on the second memory;

upon determining the set of commands are executed successfully on the second memory, executing the set of commands on the first memory; and
upon determining the set of commands are not executed on the second memory successfully, generating an error.

36. The machine-readable medium of claim 35 wherein the set of configurations are for a set of network processes.

37. The machine-readable medium of claim 35 wherein the error is displayed as a text message.

38. The machine-readable medium of claim 35 wherein the error is passed to an error parser.

39. The machine-readable medium of claim 35 wherein the associating the first memory to the second memory comprises associating a set of addresses of the first memory to a set of addresses of the second memory.

40. The machine-readable medium of claim 35 wherein the set of commands are a set of write commands.

41. The machine-readable medium of claim 35 wherein the set of commands are a set of delete commands.

42. A machine-readable medium that provides instructions, which when executed by a set of processors of one or more processors, cause said set of processors to perform operations comprising:

receiving a request to modify configuration data located at a local address in local memory in an active control card;

generating an exception when the configuration data located at the local address in the local memory is attempted to be modified, wherein the exception performs, within a processor in the active control card, the following:

modifying configuration data located at a remote address in remote memory in an inactive control card, wherein the configuration data located at the local address corresponds to the configuration data located at the remote address;

modifying the configuration data located at the local address in local memory in the active control card upon determining that the modification of the configuration data located at the remote address in the remote memory in the inactive control card was successful; and

generating an error without modifying the configuration data located at the local address in local memory in the active control card upon determining that the modification of the configuration data located at the remote address in the remote memory in the inactive control card was not successful.

43. The machine-readable medium of claim 42 wherein the request is to write the configuration data.

44. The machine-readable medium of claim 42 wherein the request is to delete the configuration data.

45. The machine-readable medium of claim 42 wherein the error is displayed on an interface.

46. The machine-readable medium of claim 42 wherein the error is passed to an error parser.

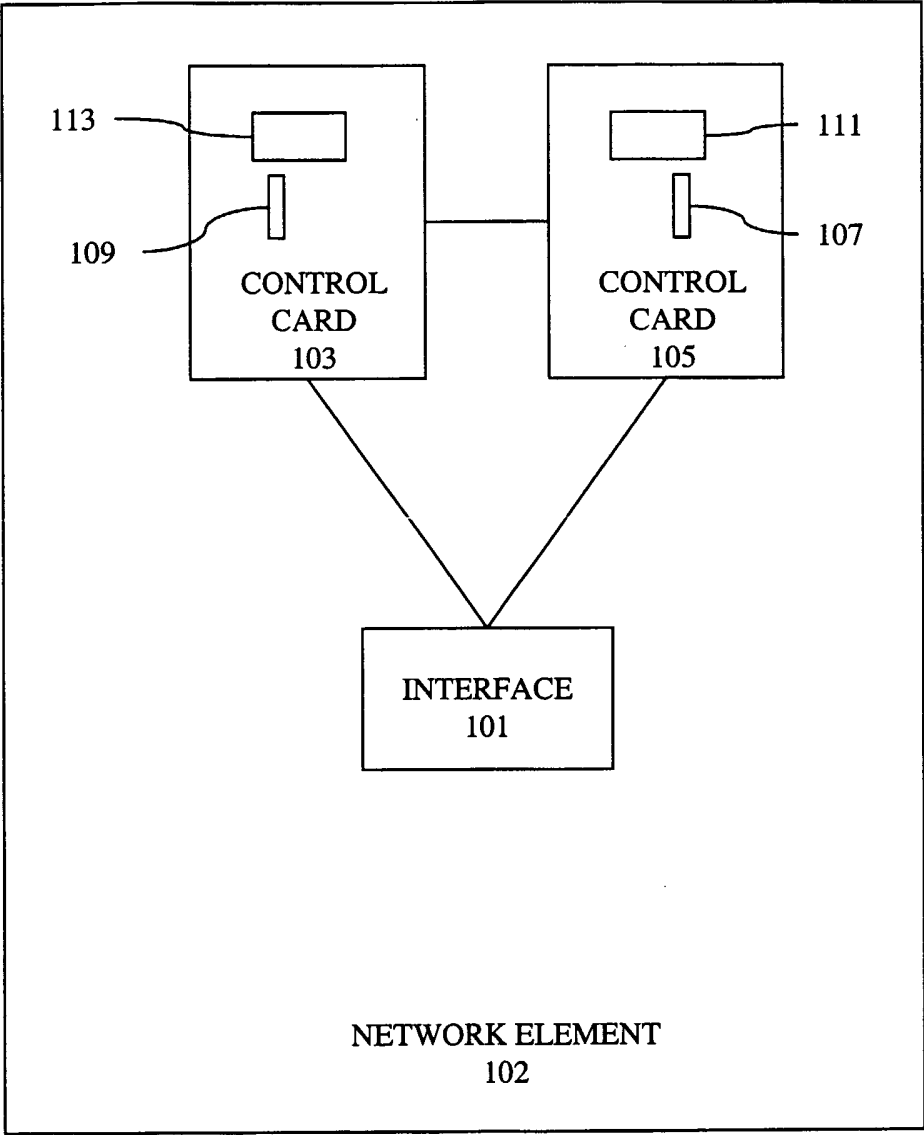


FIG.1




FIG.2

3/4




FIG.3

4/4

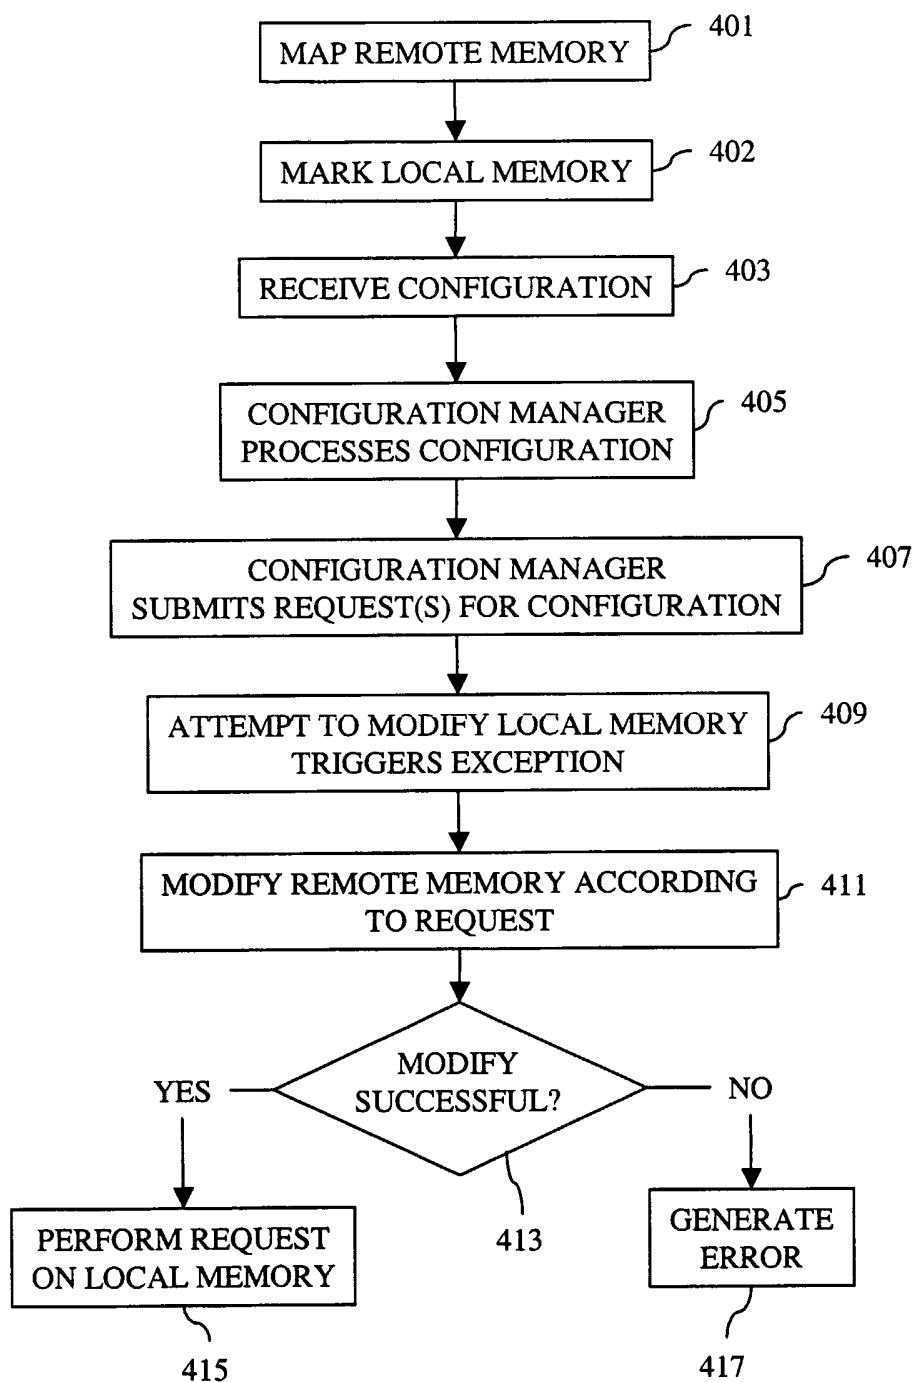


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/15321

A. CLASSIFICATION OF SUBJECT MATTER																						
IPC(7) : G06F 11/00																						
US CL : 714/5																						
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) U.S. : 714/5																						
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched USPAT FUL, EPO, JPO, IBM TDB																						
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) global and local memory , coherency																						
C. DOCUMENTS CONSIDERED TO BE RELEVANT																						
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																				
X	US 5,574,922 A (JAMES) 12 November 1996 (11.12.1996) column 1, line 1- column 36, line 22	1-46																				
X	US 5,801,938 A (KALANTERY) 01 September 1998 (09.01.1998) column 1, line 1-column 26, line 22	1-46																				
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.																						
* Special categories of cited documents: <table border="0"> <tr> <td>"A"</td> <td>document defining the general state of the art which is not considered to be of particular relevance</td> <td>"T"</td> <td>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</td> </tr> <tr> <td>"E"</td> <td>earlier application or patent published on or after the international filing date</td> <td>"X"</td> <td>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</td> </tr> <tr> <td>"L"</td> <td>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)</td> <td>"Y"</td> <td>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</td> </tr> <tr> <td>"O"</td> <td>document referring to an oral disclosure, use, exhibition or other means</td> <td>"&"</td> <td>document member of the same patent family</td> </tr> <tr> <td>"P"</td> <td>document published prior to the international filing date but later than the priority date claimed</td> <td></td> <td></td> </tr> </table>			"A"	document defining the general state of the art which is not considered to be of particular relevance	"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	"E"	earlier application or patent published on or after the international filing date	"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	"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)	"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	"O"	document referring to an oral disclosure, use, exhibition or other means	"&"	document member of the same patent family	"P"	document published prior to the international filing date but later than the priority date claimed		
"A"	document defining the general state of the art which is not considered to be of particular relevance	"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																			
"E"	earlier application or patent published on or after the international filing date	"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																			
"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)	"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																			
"O"	document referring to an oral disclosure, use, exhibition or other means	"&"	document member of the same patent family																			
"P"	document published prior to the international filing date but later than the priority date claimed																					
Date of the actual completion of the international search 12 July 2002 (12.07.2002)		Date of mailing of the international search report 07 AUG 2002																				
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230		Authorized officer Robert Beausoliel <i>James R. Matthews</i> Telephone No. (703) 305-3900																				