

(12) United States Patent
Buskens et al.

(10) Patent No.: US 6,298,039 B1
(45) Date of Patent: Oct. 2, 2001

- (54) **HIGH AVAILABILITY DISTRIBUTED CALL PROCESSING METHOD AND APPARATUS**
- (75) Inventors: **Richard Wayne Buskens**, Middletown, NJ (US); **Thomas F. La Porta**, Thornwood, NY (US); **Yow-Jian Lin**, Edison, NJ (US); **Kazutaka Murakami**, Freehold, NJ (US); **Ramachandran Ramjee**, Matawan, NJ (US)
- (73) Assignee: **Lucent Technologies Inc.**, Murray Hill, NJ (US)
- (*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

- (21) Appl. No.: **09/017,105**
- (22) Filed: **Jan. 31, 1998**
- (51) Int. Cl.⁷ **H04L 12/50**
- (52) U.S. Cl. **370/216; 714/4; 714/16**
- (58) Field of Search **370/216, 242, 370/244, 310, 328, 260; 709/100, 226, 303; 714/1, 2, 5, 4, 15**

- (56) **References Cited**
- U.S. PATENT DOCUMENTS
- | | | | |
|-----------|---|------------------------|---------|
| 4,905,181 | * | 2/1990 Gregory | 709/100 |
| 5,105,420 | * | 4/1992 Ardon et al. | 370/216 |
| 5,883,939 | * | 3/1999 Friedman et al. | 379/9 |
| 6,085,086 | * | 7/2000 La Porta et al. | 455/432 |

OTHER PUBLICATIONS

“Structuring Call Control Software Using Distributed Objects”; H. Blair, S. J. Caughey, H. Green and S. K. Shrivastava; International Workshop on Trends in Distributed Computing, Aachen, Germany, 1996; pp. 95–107.

“Distributed Call Processing For Personal Communications Services” T. F. LaPorta, M. Veeraraghavan, P. A. Treventi and R. Ramjee; IEEE Communications Magazine, vol. 33, No. 6, Jun. 1995; pp. 66–75.

“The Role Of New Technologies In Wireless Access Network Evolution”; T. F. LaPorta, A. Sawkar, W. Strom; Proceedings of International Switching Symposium (ISS 1997), IS-03.18, 1997; pp. 533–539.

“Signaling System No. 7: A Tutorial”; A. R. Modarressi and R. A. Skoog; IEEE Communications Magazine, vol. 28, No. 7, Jul. 1990; pp. 19–35.

“A Survey of Rollback–Recovery Protocols in Message-Passing Systems”; E. N. Elnozahy, D. B. Johnson and Y. M. Wang; Tech. Report CMU–CS–96–181, School of Computer Science, Carnegie Mellon University, Oct. 1996; pp. 1–46.

“Optimistic Recovery in Distributed Systems”; R. E. Strom and S. Yemini; ACM transactions on Computer Systems, vol. 3, No. 3, Aug. 1985; pp. 204–226.

“Why Optimistic Message Logging Has Not Been Used In Telecommunications System”; Y. Huang and Y–M. Wang; Proceedings of the 25th International Symposium on Fault-Tolerant Computing 1995; pp. 459–463.

“Software Fault Tolerance in the Application Layer”; Y. Huang and C. Kintala; Software Fault Tolerance, John Wiley & Sons Ltd.; pp. 231–248.

“Implementation of On–Line Distributed System–Level Diagnosis Theory”; R. P. Bianchini, Jr. and R. W. Buskens; IEEE Transactions On Computers, vol. 41, No. 5, May 1992; pp. 616–626.

* cited by examiner

Primary Examiner—Hassan Kizou
Assistant Examiner—Inder Mehra
 (74) *Attorney, Agent, or Firm*—Jeffery J. Brosemer

(57) **ABSTRACT**

A method of delivering highly-reliable, fault-tolerant communications services in a telecommunications network of distributed call processing systems. The method advantageously identifies a set of objects within the telecommunications network requiring checkpointing; checkpoints the objects; and subsequently restores the checkpointed objects in the event of a failure. Various aspects of the method are disclosed, including restoration strategies.

9 Claims, 3 Drawing Sheets

FUNCTIONAL OBJECTS FOR DISTRIBUTED CALL PROCESSING




FIG. 1

FUNCTIONAL OBJECTS FOR DISTRIBUTED CALL PROCESSING




FIG. 2

TYPICAL STATE MACHINE IN CALL PROCESSING SYSTEMS


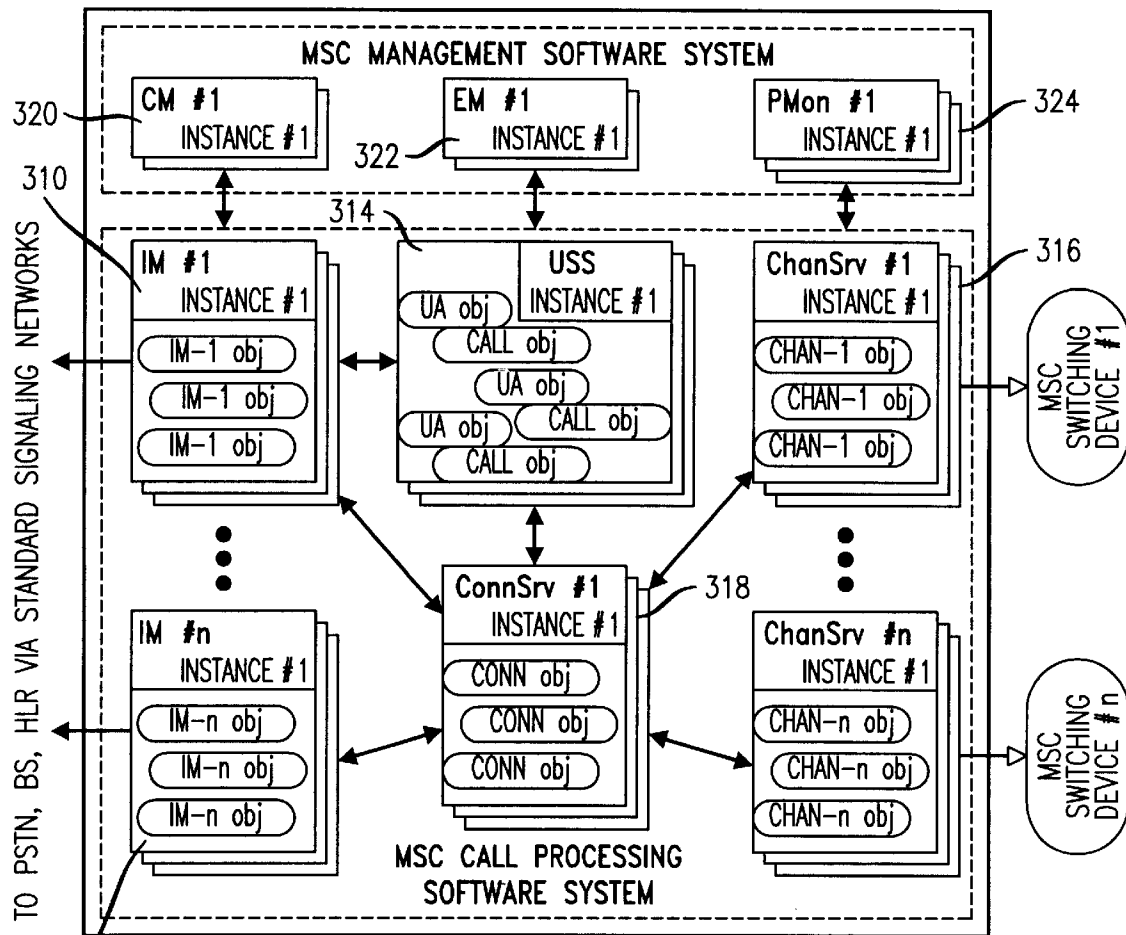


FIG. 3

MSC CALL PROCESSING SOFTWARE STRUCTURE

300




310 USS: USER SIGNAL SERVER
 ConnSrv: CONNECTION SERVER
 CM: CONFIGURATION MANAGER
 PMon: PROCESS MONITOR

ChanSrv: CHANNEL SERVER
 IM: INTERWORKING MANAGER
 EM: EVENT MANAGER

FIG. 4

FAILURE-FREE OVERHEAD: LATENCY VERSUS LOAD WITH AND WITHOUT CHECKPOINTING



HIGH AVAILABILITY DISTRIBUTED CALL PROCESSING METHOD AND APPARATUS

TECHNICAL FIELD

This invention relates generally to the field of telecommunications and in particular to a method for imparting high availability and fault tolerance to distributed call processing systems.

BACKGROUND OF THE INVENTION

The development of telecommunications call processing or switching systems constructed from a distributed set of general purpose computing systems is emerging as an area of particular interest in the art. See, for example, H. Blair, S. J. Caughey, H. Green and S. K. Shrivastava, "Structuring Call Control Software Using Distributed Objects," International Workshop on Trends in Distributed Computing, Aachen, Germany, 1996; T. F. LaPorta, M. Veeraraghavan, P. A. Treventi and R. Ramjee, "Distributed Call Processing for Personal Communication Services," *IEEE Communications Magazine*, vol.33, no.6, pp. 66-75, June 1995; and TINA-C, Service Architecture Version 2.0, March 1995.

As noted in a paper published by T. F. LaPorta, A. Sawkar and W. Strom, entitled "The Role of New Technologies in Wireless Access Network Evolution," that appeared in Proceedings of International Switching Symposium (ISS '97), IS-03.18, 1997, systems employing distributed call processing architectures exhibit increased system scalability, performance, and flexibility. Additionally, advances in open distributed processing, such as the Common Object Request Broker Architecture (CORBA), described in "The Common Object Request Broker: Architecture and Specification," by the Object Management Group (OMG) Rev. 2.0, July 1995, facilitate portable and interoperable implementations of distributed software architectures in a heterogeneous computing environment. As is known, systems employing such technologies advantageously leverage a rapidly increasing price/performance ratio of "off-the-shelf" computing components.

The stringent performance and availability requirements of public telecommunications systems pose particular challenges to developing highly available distributed call processing systems which incorporate these off-the-shelf computing components. Specifically, and as noted by A. R. Modarressi, R. A. Skoog, in an article entitled "Signaling System No. 7: A Tutorial", which appeared in *IEEE Communications Magazine*, Vol. 28, No. 7, pp. 19-35, in July 1990, call processing software must process each call request within a few hundred milliseconds, and a switching system may not be out of service for more than a few minutes per year. As such, present day switching systems employ custom-designed fault-tolerant processors and special-purpose operating systems to meet these stringent requirements. In order for next generation switching systems to be built using general purpose computing platforms, software-based fault-tolerant methods and systems are required to achieve the same or similar performance and availability goals.

Two software methods for enhancing the level fault tolerance in a distributed computing environment that have been described in the literature are checkpointing and message logging. See, for example, E. N. Elnozahy, D. B. Johnson and Y. M. Wang, "A Survey of Rollback-Recovery Protocols in Message-Passing Systems," Tech. Report CMU-CS-96-181, School of Computer Science, Carnegie Mellon University, October 1996, and R. E. Strom and S.

Yemini, "Optimistic Recovery in Distributed Systems," *ACM Transactions on Computer Systems*, vol.3, no.3, pp.204-226, August 1985. Briefly stated, checkpointing involves periodically taking a "snapshot" and saving an entire state of a software process while messages sent or received by the software process are logged (message logging) between subsequent checkpoints. Assuming a piecewise deterministic execution model, and as described by Y. Huang and Y. M. Wang, in an article entitled "Why Optimistic Message Logging has not been used in Telecommunications Systems," that appeared in the Proceedings of the 25th International Symposium on Fault-Tolerant Computing, pp. 459-463, 1995, the state of the process can be later reconstructed during a recovery process by replaying logged messages in their original order. As observed by Y. Huang and C. Kintala, in "Software Fault Tolerance in the Application Layer," which appeared in *Software Fault Tolerance* (M. R. Lyu, Ed.), John Wiley & Sons, Chichester, England, pp.231-248, 1995, checkpointing, message logging, and "rollback" recovery techniques can be embedded into the operating system while remaining virtually transparent to application software.

Unfortunately, however, there are numerous disadvantages to these approaches when applied to distributed call processing systems. First, taking a snapshot of the entire process state may create a long period of time during which the process is unable to service requests from its clients, thereby increasing end-to-end call setup latency. Second, a single call request may involve a significant number of message exchanges between functionally distributed servers. Consequently, logging every message becomes too time-consuming to meet stringent call setup latency requirements of only a few hundred milliseconds associated with call processing. Additionally, if checkpoint intervals are made sufficiently long in an attempt to minimize checkpoint overhead, a prohibitively large number of messages may need to be replayed after a failure, thereby making recovery time unacceptably long. Consequently, a continuing need exists in the art for software-based fault-tolerant computing systems suitable for demanding telecommunications applications.

SUMMARY OF THE INVENTION

An advance is made over the prior art in accordance with the principles of the present invention directed to a method of delivering highly-reliable, fault-tolerant communications services in a telecommunications network of distributed call processing systems. The method advantageously identifies a set of objects within the telecommunications network requiring checkpointing; checkpoints the objects; and subsequently restores the checkpointed objects in the event of a failure. Additionally, the method accommodates the selective determination of particular states requiring restoration, and reduces, where desired, duplicate restorations within the system.

Further features and advantages of the present invention, as well as the structure and operation of various embodiments of the present invention are described in detail below with reference to the accompanying drawing.

BRIEF DESCRIPTION OF THE DRAWING

The teachings of the present invention can be readily understood by considering the following detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 is a bubble diagram depicting functional objects associated with distributed call processing;

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.