



US007185003B2

(12) **United States Patent**
Bayliss et al.

(10) **Patent No.:** **US 7,185,003 B2**
(45) **Date of Patent:** **Feb. 27, 2007**

(54) **QUERY SCHEDULING IN A
PARALLEL-PROCESSING DATABASE
SYSTEM**

5,471,622 A 11/1995 Eadline
5,495,606 A 2/1996 Borden et al.
5,551,027 A 8/1996 Choy et al.
5,555,404 A 9/1996 Torbjørnsen et al.
5,655,080 A 8/1997 Dias et al.
5,732,400 A 3/1998 Mandler et al.

(75) Inventors: **David Bayliss**, Delray Beach, FL (US);
Richard Chapman, Boca Raton, FL
(US); **Jake Smith**, London (GB); **Ole
Poulsen**, Bend, OR (US); **Gavin
Halliday**, Royston (GB); **Nigel Hicks**,
London (GB)

(Continued)

(73) Assignee: **Seisint, Inc.**, Boca Raton, FL (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 774 days.

Eike Schallehn et al., "Advanced Grouping and Aggregation for
Data Integration," Department of Computer Science, Paper ID: 222,
pp. 1-16.

(Continued)

(21) Appl. No.: **10/293,489**

Primary Examiner—Joon Hwan Hwang
(74) *Attorney, Agent, or Firm*—Hunton & Williams LLP

(22) Filed: **Nov. 14, 2002**

(65) **Prior Publication Data**

(57) **ABSTRACT**

US 2004/0098374 A1 May 20, 2004

(51) **Int. Cl.**
G06F 17/30 (2006.01)
G06F 15/16 (2006.01)
G06F 9/46 (2006.01)

(52) **U.S. Cl.** **707/3; 707/10; 718/102;**
709/201

(58) **Field of Classification Search** **707/3,**
707/10; 718/102; 709/201–203
See application file for complete search history.


(56) **References Cited**

U.S. PATENT DOCUMENTS

4,543,630 A 9/1985 Neches
4,860,201 A 8/1989 Stolfo et al.
4,870,568 A 9/1989 Kahle et al.
4,925,311 A 5/1990 Neches et al.
5,006,978 A 4/1991 Neches
5,276,899 A 1/1994 Neches
5,303,383 A 4/1994 Neches et al.
5,423,037 A 6/1995 Hvasshovd

A system and method for scheduling database operations to one or more databases in a parallel-processing database system are described herein. After a query server generates a dynamic-link library (DLL) or other executable representative of one or more database operations to a database, the query server notifies a scheduling services module of the generation of the DLL and submits the DLL to a query agent. The query agent notifies the scheduling services module of its receipt of the DLL. Based on any of a variety of considerations, the scheduling services module schedules a time of execution for the DLL by one or more processing matrices that store the database. At the scheduled time, the scheduling services module directs the query agent to submit the DLL to the indicated processing matrices. The scheduling services module also can be adapted to monitor the execution of previously submitted DLLs by one or more processing matrices and adjust the scheduled times of execution for subsequent DLLs accordingly.

25 Claims, 21 Drawing Sheets



U.S. PATENT DOCUMENTS

5,745,746	A	4/1998	Jhingran et al.	
5,857,180	A *	1/1999	Hallmark et al.	707/2
5,878,408	A	3/1999	Van Huben et al.	
5,884,299	A	3/1999	Ramesh et al.	
5,897,638	A	4/1999	Lasser et al.	
5,983,228	A	11/1999	Kobayashi et al.	
6,006,249	A	12/1999	Leong	
6,026,394	A	2/2000	Tsuchida et al.	
6,081,801	A	6/2000	Cochrane et al.	
6,266,804	B1	7/2001	Isman	
6,311,169	B2	10/2001	Duhon	
6,427,148	B1	7/2002	Cossock	
6,990,503	B1 *	1/2006	Luo et al.	707/200
2003/0037048	A1 *	2/2003	Kabra et al.	707/4

OTHER PUBLICATIONS

Vincent Coppola, "Killer APP," Men's Journal, vol. 12, No. 3, Apr. 2003, pp. 86-90.

Eike Schallehn et al., "Extensible and Similarity-based Grouping for Data Integration," Department of Computer Science, pp. 1-17.
Rohit Ananthakrishna et al., "Eliminating Fuzzy Duplicates in Data Warehouses," 12 pages.

Peter Christen et al., "Parallel Computing Techniques for High-Performance Probabilistic Record Linkage," Data Mining Group, Australian National University, Epidemiology and Surveillance Branch, Project web page: <http://datamining.anu.edu.au/linkage.html>, 2002, pp. 1-11.

Peter Christen et al., "Parallel Techniques for High-Performance Record Linkage (Data Matching)," Data Mining Group, Australian National University, Epidemiology and Surveillance Branch, Project web page: <http://datamining.anu.edu.au/linkage.html>, 2002, pp. 1-27.

Peter Christen et al., "High-Performance Computing Techniques for Record Linkage," Data Mining Group, Australian National University, Epidemiology and Surveillance Branch, Project web page: <http://datamining.anu.edu.au/linkage.html>, 2002, pp. 1-14.

William E. Winkler, "Matching And Record Linkage," U.S. Bureau of the Census, pp. 1-38.

Peter Christen et al., "High-Performance Computing Techniques for Record Linkage," ANU Data Mining Group, Australian National University, Epidemiology and Surveillance Branch, Project web page: <http://datamining.anu.edu.au/linkage.html>, pp. 1-11.

William E. Winkler, "The State of Record Linkage and Current Research Problems," U.S. Bureau of the Census, 15 pages.

William E. Winkler, "Advanced Methods For Record Linkage," Bureau of the Census, pp. 1-21.

William E. Winkler, Frequency-Based Matching in Fellegi-Sunter Model of Record Linkage, Bureau Of The Census Statistical Research Division, Oct. 4, 2000, 14 pages.

William E. Winkler, "State of Statistical Data Editing And Current Research Problems," Bureau Of The Census Statistical Research Division, 10 pages.

The First Open ETL/EAI Software For The Real-Time Enterprise, Sunopsis, A New Generation ETL Tool, "Sunopsis™ v3 expedites integration between heterogeneous systems for Data WAREHOUSE, Data Mining, Business Intelligence, and OLAP projects," <www.suopsis.com>, 6 pages.

Alan Dumas, "The ETL Market and Sunopsis™ v3 Business Intelligence, Data Warehouse & Datamart Projects," 2002, Sunopsis, pp. 1-7.

Teradata Warehouse Solutions, "Teradata Database Technical Overview," 2002, pp. 1-7.

WhiteCross White Paper, May 25, 2000, "wx/des-Technical Information," pp. 1-36.

Teradata Alliance Solutions, "Teradata and Ab Initio," pp. 1-2.

Peter Christen et al., The Australian National University, "Febrl—Freely extensible biomedical record linkage," Oct. 2002, pp. 1-67.

William E. Winkler, "Using the EM Algorithm for Weight Computation in the Fellegi-Sunter Model of Record Linkage," Bureau Of The Census Statistical Research Division, Oct. 4, 2000, 12 pages.

William E. Winkler et al., "An Application Of The Fellegi-Sunter Model Of Record Linkage To The 1990 U.S. Decennial Census," U.S. Bureau of the Census, pp. 1-22.

William E. Winkler, "Improved Decision Rules In The Fellegi-Sunter Model Of Record Linkage," Bureau of the Census, pp. 1-13.

Fritz Scheuren et al., "Recursive Merging and Analysis of Administrative Lists and Data," U.S. Bureau of the Census, 9 pages.

William E. Winkler, "Record Linkage Software and Methods for Merging Administrative Lists," U.S. Bureau of the Census, Jul. 7, 2001, 11 pages.

Enterprises, Publishing and Broadcasting Limited, Acxiom-Abilitec, pp. 44-45.

TransUnion, Credit Reporting System, Oct. 9, 2002, 4 pages, <<http://www.transunion.com/content/page.jsp?id=/transunion/general/data/business/BusCre...>>.

TransUnion, ID Verification & Fraud Detection, Account Acquisition, Account Management, Collection & Location Services, Employment Screening, Risk Management, Automotive, Banking-Savings & Loan, Credit Card Providers, Credit Unions, Energy & Utilities, Healthcare, Insurance, Investment, Real Estate, Telecommunications, Oct. 9, 2002, 46 pages, <<http://www.transunion.com>>.

White Paper An Introduction to OLAP Multidimensional Terminology and Technology, 20 pages.

* cited by examiner

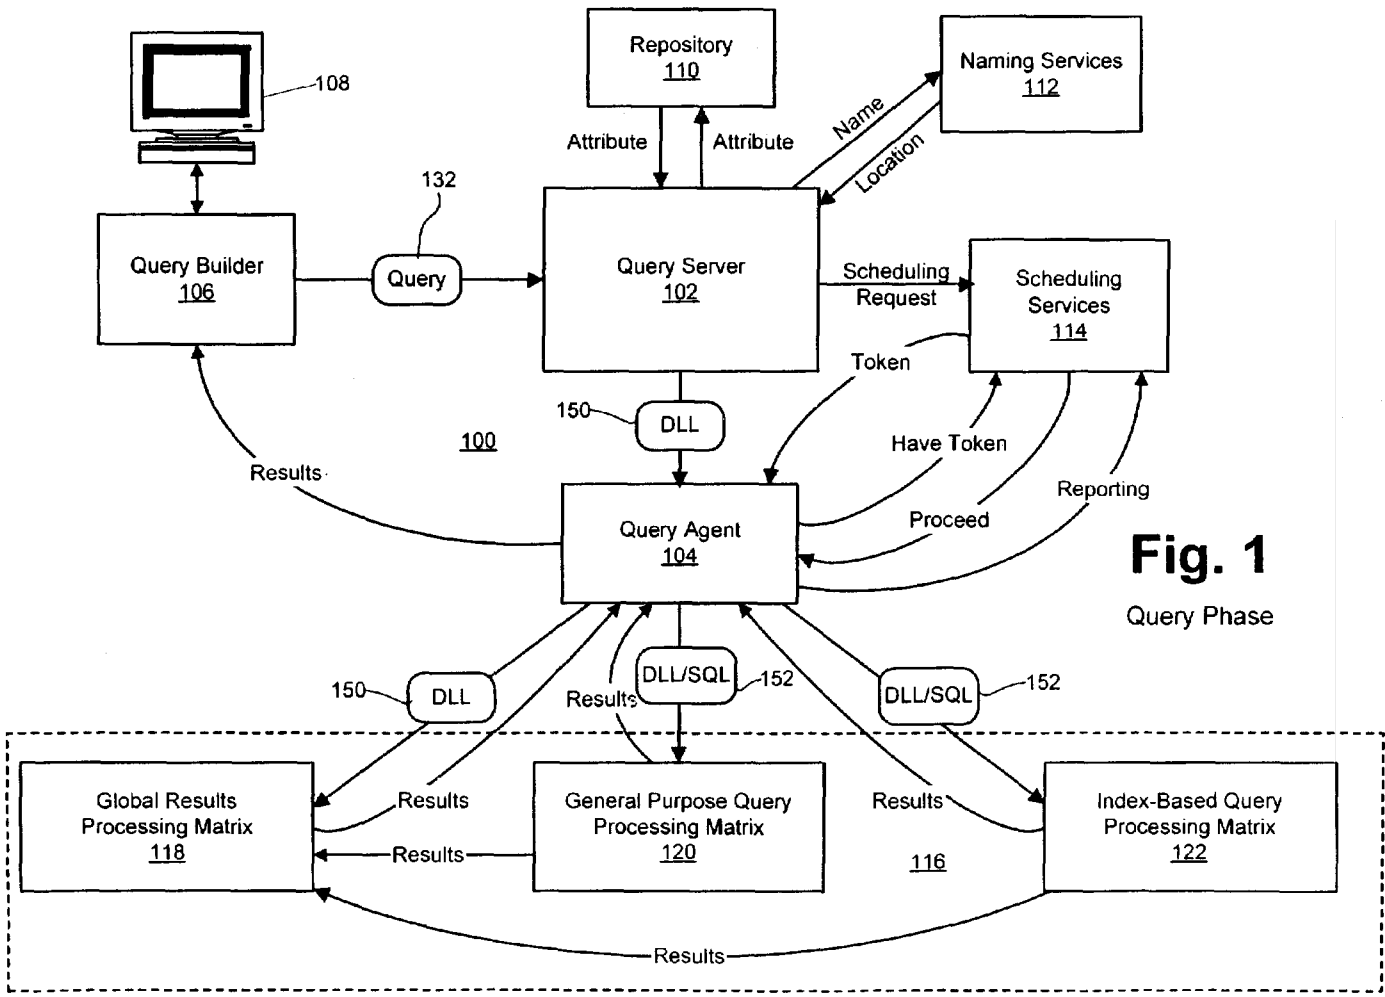


Fig. 1
Query Phase

Fig. 2

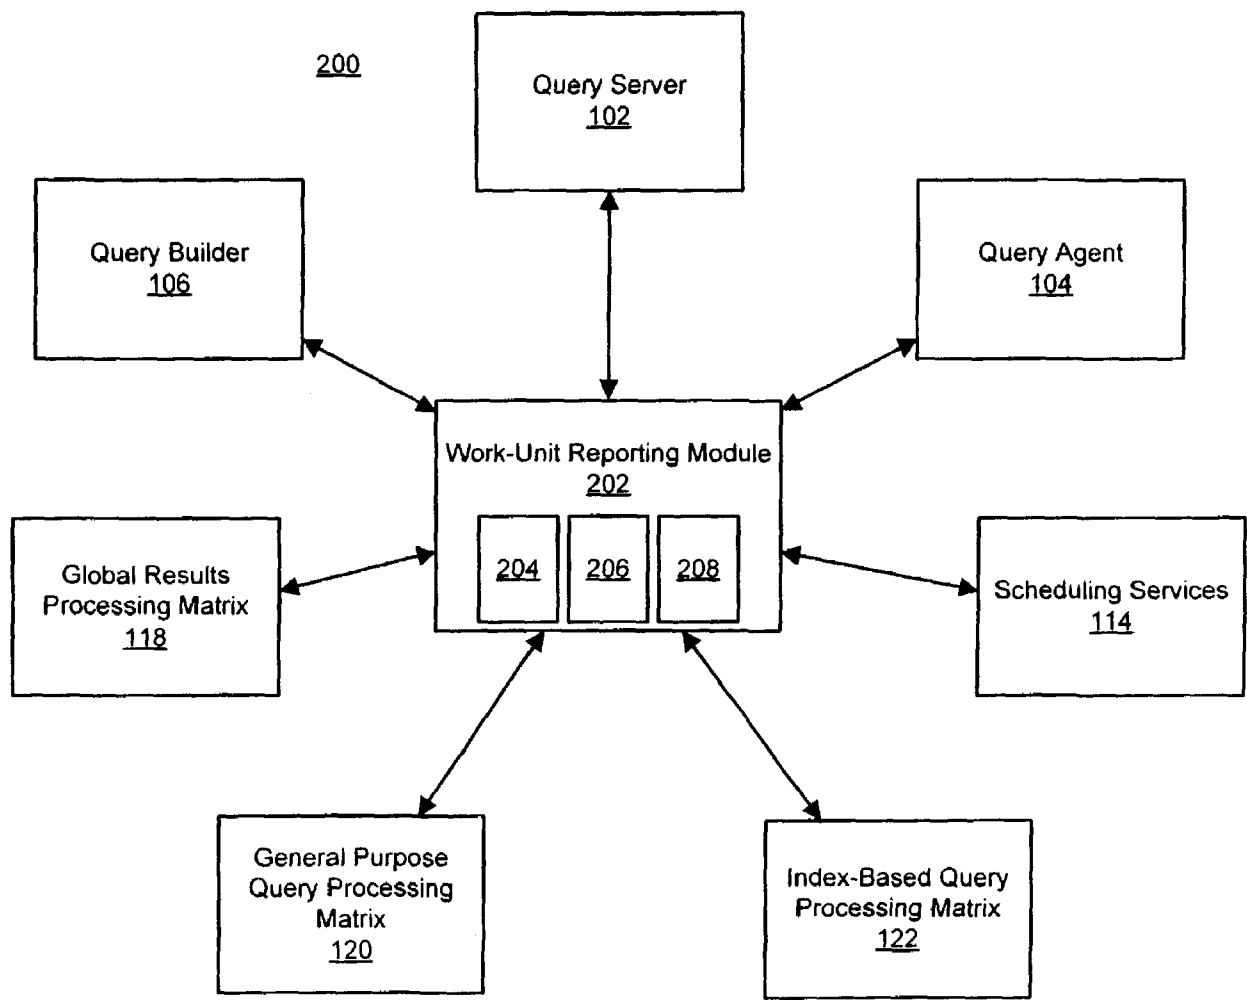



Fig. 3



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.