ORACLE®

Conference Proceedings

Oracle Corporation
International Users' Conference
August 18-21, 1985

David R. Crowe
NovAtel Communications Ltd.
700 Ninth Avenue S.W.
Calgary, Alberta
Canada, T2P 3V4
(403) 298-0523

ABSTRACT

NovAtel Communications Ltd. is the developer of the AURORA family of Cellular Communications System. AURORA has a unique distributed control and switching architecture that minimizes network facility costs and allows for cost-effective application in both low density rural and high density urban environments. The Oracle DBMS is used to provide the man-machine software interface to the system and maintain the distributed database.

Intelligent radio subsystems in cell sites are connected by voice and data links to Satellite Mobile Centres (SMCs) which perform time critical call processing and switching functions. SMCs are the interface to the public switched telephone network (PSTN) and are connected by data links to the Master Mobile Centre (MMC).

Unlike the real time SMC environment, user-friendliness and fast development time are more important in the MMC. Oracle has been used to rapidly develop forms that provide convenient and controlled access to the database and to trigger the distributed network server to transmit each database update to one or more SMC's.

Guaranteed distribution of database updates is implemented by keeping updates on disk at the originating node (SMC or MMC) until the destination node acknowledges receipt. Data on disk is protected at each SMC, by having dual disks and, at the MMC, by Oracle's After-Image Journaling. As a result no single failure in the system can cause loss of a database update.

INTRODUCTION

NovAtel Communications Ltd. is the developer of the AURORA family of Cellular Communications System. AURORA has a unique distributed control and switching architecture that minimizes network facility costs and allows for cost-effective application in both low density rural and high density urban environments. The Oracle DBMS is used to provide the man-machine software interface to the system and maintain the distributed database.

WHAT IS CELLULAR COMMUNICATIONS?

Cellular Communications is a means of extending telephone communications to vehicles or to areas where it is too expensive to construct telephone lines. Unlike conventional radio-telephone systems, operator intervention is not required, conversations are private and many more simultaneous calls can be supported in a given serving area.

The most important feature of cellular communications is 'frequency re-use'. The area to be served is divided into 'cells' (hence the name 'Cellular Communications') each allocated a subset of the available radio frequencies. Cells are arranged so that those with the same frequency allocation are far enough apart that they don't interfere with each other. Various frequency reuse plans are possible; Figure 1 illustrates a 4 cell reuse plan. At the centre of each cell is a 'cell site' which contains the radio equipment capable of transmitting and receiving on the allocated frequencies. As more cell sites are added to the system (to increase its capacity) cell boundaries are contracted by reducing the transmit power of cell sites.

Another important feature of cellular communications sytems is 'handoff'. This is a mechanism for maintaining a call when the cellular terminal (e.g. a cellular mobile phone in a car) moves outside the limits of a cell. The cell site that is receiving the strongest signal from the terminal assigns it a frequency and the terminal automatically switches over to it. The people involved in the call will not notice any break in their conversation.

THE AURORA 800 ARCHITECTURE

The unique feature of the AURORA 800 system is its distributed architecture. As illustrated in Figure 2 the system is a hierarchy of cell sites (at the bottom), Satellite Mobile Centres (in the middle) and a Master Mobile Centre (at the top).

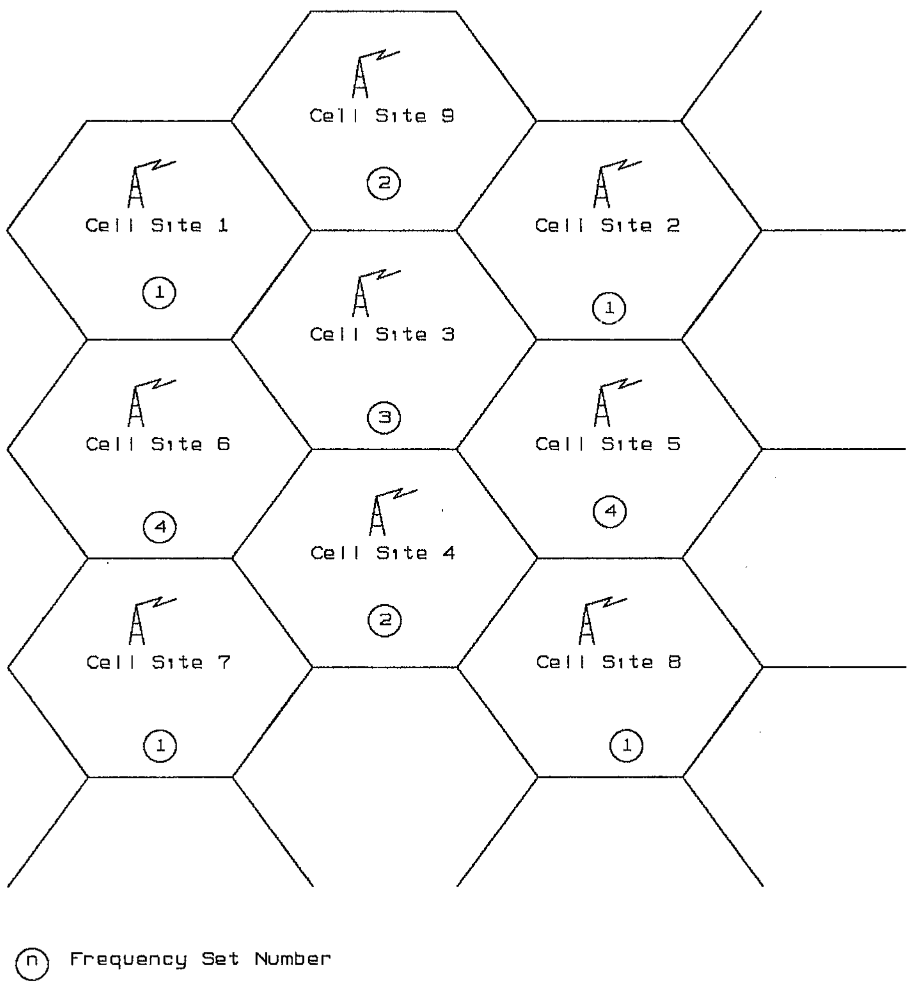

Figure 1: A Cellular Communications System

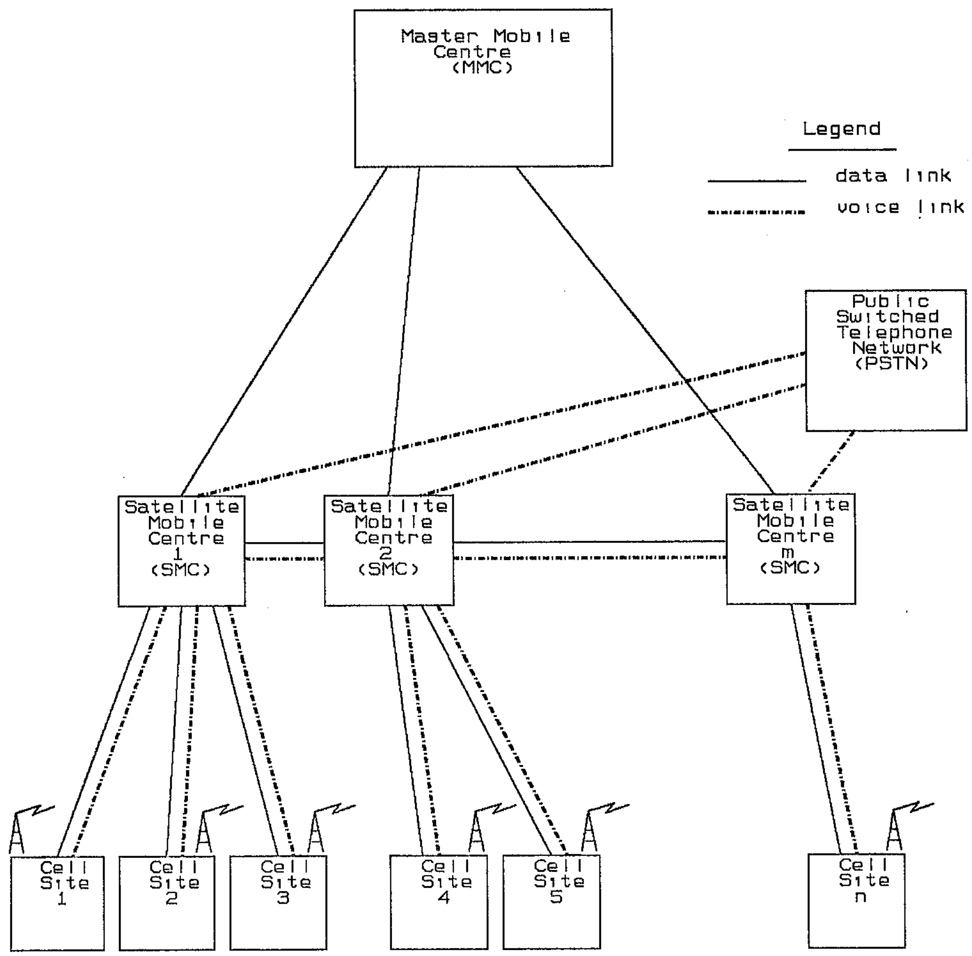

Figure 2: The AURORA-800 Architecture

Cell sites are the interface between the wireless cellular terminals and the SMCs. Each cell site is connected by voice and data links to one SMC.

SMCs manage calls originating or terminating in the cells they control. They also interface the AURORA system to the Public Switched Telephone Network (PSTN). Each SMC is connected to one or more cell sites by voice and data links, to the PSTN by voice links, to the MMC by data links and to other SMCs by voice and data links. Each SMC is a multi-processor computer with up to 22 micro-PDP 11s, split into two identical sides. At any time one side is active and the other is standby (ready to become active) or out of service (not ready to become active). SMCs can operate even when isolated from the MMC and other SMCs.

The MMC does not participate in call processing. It is the man-machine interface of the system and provides a single control point for it. At the MMC they system database can be queried or modified, equipment can be put in or out of service and billing and system status information from all SMCs is accumulated. The MMC is a VAX computer running the VMS operating system. The model of VAX chosen for a system, and the amount of disk storage, depends on the requirements of the customer. A minimal system requires a VAX 11/730 with two 70MB disks, larger systems need a more powerful CPU and more, or larger, disks.

In summary, the network of SMCs allows calls to be connected without having to run (expensive) voice and data links from each cell site back to a single point. Yet, the MMC maintains the advantage of a centralized system by providing a single point for system monitoring, control and data gathering. This contrasts with other cellular communications systems that have cell sites connected to a sub-system that performs both call processing and maintenance functions. In these systems facility costs can be reduced by having several small, isolated systems without the benefit of a centralized maintenance facility. Or, a centralized maintenance facility can be retained at the expense of providing one voice link from each cell site to a central location for every allocated frequency.

THE DISTRIBUTED DATABASE

The database in an AURORA-800 system is distributed because the system is physically distributed. SMCs reference data frequently throughout the processing of each of the many thousands of calls processed each hour -- therefore database queries must be processed quickly. The MMC, being the central point of the system must provide a friendly interface for querying and modifying the database as well as collecting call records (for billing customers), traffic statistics etc. The network, not being fast enough to support queries from the SMCs to the MMC tables in real time, must provide a means for updating SMC tables from the MMC. Also, to allow data gathering at the MMC, the network must support updating of MMC tables by SMCs.

Data tables at the SMC are stored in memory or on disk (depending on the size and frequency of access) and are usually organized as arrays of records, accessed by array index. Because each SMC is duplicated in hardware, both sides must contain the same information. Therefore, every update to an SMC table, whether generated from the MMC or within the SMC must be applied to both sides of the SMC. Oracle was not chosen to implement this portion of the database because of the memory and speed constraints of the SMCs.

Data tables at the MMC are managed entirely by Oracle. It was chosen to help provide a man-machine interface that would be friendly, maintain the consistency of the database and could be implemented quickly. The following features of Oracle have proved most important.

- o SQL, Oracle's query language, provides sophisticated data manipulation capabilities with a shorter development time that 3rd generation languages. Also, even where 3rd generation lanaguage programs are required, optimizations of the database access algorithm can often be made without modifications to the program (e.g. creation of indexes and clusters).

- o IAF, Oracle's form processor provides an easy to use, 'fill in the blanks' tool for AURORA-800 users to query and modify the MMC tables with no knowledge of SQL. It allows screen applications to be implemented without writing any 3rd generation language code.

- o IAF SQL Triggers are SQL statements that are executed before or after any operation on a form. They can be specified to validate data entered on the form or to perform complex operations which appear to the user as a single operation (i.e. a transaction). If any of the triggers fail, the entire transaction can be rolled back. This can be used to prevent invalid records from being entered and partial transactions from being performed.

- o Before-Image journaling allows transactions that are incomplete when the MMC crashes to be rolled back when it is restarted. This, again, helps prevent partial transactions from being performed.

- After-Image journaling allows the MMC tables to be recovered after a catastrophic disk failure. All updates to the database are also written to an after-image file (on a separate disk from the database). If the database disk fails the database can be recovered from the after-image files and a backup copy of the database.

- IAF User Escapes are portions of 3rd generation language code that can be executed before or after any operation on a form. They allow operations to be performed that are beyond the capabilities of SQL. The AURORA-800 system uses one escape, written in C, that activates the process responsible for distributing updates to a particular table (via a message written to a VMS mailbox; not to be confused with the VMS MAIL utility). This escape is also used to initiate queries on tables that are not stored at the MMC because the information in them is required at the MMC much less frequently than it is updated at the SMCs (e.g. equipment status tables).

The network connecting the MMC and SMCs uses a standard protocol (X.25) so that it can be implemented by connecting to a data network or with dedicated data links. Cost considerations dictate that the speed of the network may be as low a 1200 bits per second so that it can be implemented with regular telephone lines.

Distributed Database Updates

There are two types of distributed database updates, those initiated at the MMC and those initiated at the SMC. In both cases it must be guaranteed that every database update is distributed to the node(s) it is destined for, even with a failure in the system. In most cases of multiple failure, delivery must also be guaranteed.

Updates initiated at the MMC may have to be distributed to one SMC (e.g. cell site configuration data) or all SMCs (e.g. AURORA-800 subscriber records). Updates from an SMC to the MMC are initiated after important system events (e.g. detection of an alarm or completion of a call). In both cases, a similar scheme is used to guarantee delivery.

The originating node (SMC or MMC) stores the database update on stable storage (i.e. disk) before sending an update message to the destination node. The destination node, if it receives the update message, stores the update on disk before sending an acknowledgement back to the originator. If the originator receives the acknowledgement it deletes that database update from its disk. Otherwise, the database update will be sent again at a later time. Updates are guaranteed to be delivered in sequence because each update must be acknowledged as having been completely processed at the destination before the next update is sent. The following classes of failures have been considered.

- Disk Failure. Tables on disk at both the originating and destination nodes have to be protected. At the SMCs this is accomplished by having two disks containing the same information. Every update to a table on one disk will also be made to an identical table on the other. Normally the failure of an SMC disk will cause the standby side to become active immediately. However, an out of service SMC must bring its disk up to date with the active side before becoming standby.

  The MMC tables are protected by Oracle's After-Image journaling. A disk failure would require the accumulated after-image files to be merged with a backup copy of the database file before the MMC could be restarted. While the MMC is processing these files it will not communicate with the SMCs (i.e. it will appear as a failed node to the rest of the network).

- Link Failure. Link failure can result in the loss of an update message or acknowledgement. The originating node cannot tell the difference because in neither case will an acknowledgement be received -- assuming the worst it will retransmit the update. Updates will not be lost because they are stored on protected disk storage at the originating node. The AURORA-800 database has been designed so that all tables can be updated with the same data many times as long as the updates are received in the sequence they were generated in. It would be possible, if a table could not be updated with the same data more than once, to make special provision to recognize receipt of a duplicate update.

- Node Failure. Node failure can be caused by a power failure, CPU failure, software failure or, at the MMC, by a disk failure. The symptom is that the failed node does not respond to any attempt to communicate with it over the network. Because this type of failure is indistinguishable from a link failure, the same procedure is used to guarantee delivery of the update once the node recovers.

CONCLUSION

The AURORA-800 cellular communications system illustrates some of the tradeoffs of a distributed system. The system attempts to attain the advantages of lower facility costs due to decentralized call processing and lower maintenance costs due to a centralized man-machine interface. This is accomplished by having the distributed call processing nodes (SMCs) connected to a central control point (MMC) by data links only.

It was decided to maintain the database with copies of data tables at whichever network nodes needed them. This requires distribution of database updates but allows queries to be processed within the local node.

Oracle is used to efficiently implement the man-machine interface and to protect data tables stored at the Master Mobile Centre (MMC). It triggers distribution of database updates whenever data is entered through the man-machine interface and provides reliable storage of database updates at the MMC until they are known to have been stored at the destination node.