Linking Continental Databases in the Oracle RDBMS with Project Data in the Arc/Info GIS

by

Prame N Chopra & Roderick J Ryburn

RECORD 1993/12

AGSO
AUSTRALIAN GEOLOGICAL
SURVEY ORGANISATION

RECORD 1993/12

Linking Continental Databases in the Oracle RDBMS with Project Data in the Arc/Info GIS

Prame N. Chopra and Roderick J. Ryburn
Information Systems Branch

DEPARTMENT OF PRIMARY INDUSTRIES AND ENERGY

Minister for Resources: Hon. Michael Lee
Secretary: Greg Taylor

AUSTRALIAN GEOLOGICAL SURVEY ORGANISATION

Executive Director: Harvey Jacka

ISSN: 1039-0073

ISBN: 0 642 19223 5

Table of Contents

Abstract

Experience with a method of linking continental scale geoscience data in AGSO's Oracle relational database management system to the project data holdings in AGSO's Arc/Info geographic information system is outlined in this Record. The details of the hardware and software links between Oracle and Arc/Info are outlined and illustrated. The differences between the data models used by AGSO in Oracle and Arc/Info are also discussed.

Particular reference is made to the Database Integrator module in Arc/Info Rev. 6.1 and the issues that must be addressed in order to link the data in Oracle to Arc/Info. The role and implementation of the Arc/Info *relate environment* as it applies to AGSO's databases is crucial to a successful link. Practical examples of how to establish an appropriate *relate environment* are given for the NGMA SITES and ROCKCHEM databases.

The usefulness of the Oracle-Arc/Info link is illustrated with reference to the current AGSO data holdings for the Lake Carey, WA and Blayney, NSW, 1:100,000 sheets.

Introduction

AGSO uses an Oracle Relational DataBase Management System (RDBMS) to store and query much of its very extensive holdings of geoscience information for the continent of Australia, its offshore territories and continental shelf. The majority of the holdings currently stored in Oracle are point data and most of these are georeferenced (e.g. Chopra, 1989; Ryburn, 1990).

The Arc/Info geographic information system (GIS) is used, in the main, to store and manipulate very different kinds of data to those held in Oracle. Arc/Info is the primary information technology (IT) system for the assembling of geoscience mapping information under the National Geoscience Mapping Accord (NGMA). As such, Arc/Info is principally concerned with vector data at scales of 1:100,000 to 1:250,000.

Much is to be gained by AGSO in terms of productivity and synergy if the Oracle and Arc/Info systems can be effectively linked so that data in one system can be accessed by the other. At present it is not possible to access Arc/Info data directly from an Oracle session so as to incorporate the GIS data into queries framed against the RDBMS's datasets. However, the reverse is possible. The Arc/Info Rev 6.1 system includes a module called the Database Integrator which permits GIS users to query the Oracle system and, provided the appropriate access privileges are granted, also allows users to update the Oracle data holdings.

Examples of procedures that can be followed in order to establish working links between the Oracle and Arc/Info systems and between the data models that they use to store and access AGSO's data are the subjects of this Record. We illustrate the power of these links and how they can be used through examples involving field location and geochemical data holdings in the SITES and ROCKCHEM databases in Oracle for the Lake Carey, WA and Blayney, NSW 1:100,000 sheets.

Connecting Arc/Info to Oracle Using SQL*Net

The Oracle SQL*Net product has been installed on the AViiON 6240 corporate database server and on the SUN 670 corporate Arc/Info server. This software provides a link between the RDBMS and the GIS and makes the Oracle system accessible from Arc/Info.

In order to be able to access Oracle on the AViiON from the SUN computers in AGSO, users must update the .cshrc file in their home directory (which is usually /mnt/<username>). The following line must be added to .cshrc using a plain text editor such as **vi** or **textedit**:

source /mnt/protos/cshrc-oracle

Once this line has been added, the user must log off and log on again to bring the change into effect. Alternatively, issuing the command:

source ~/.cshrc

at the UNIX prompt will put the change into effect immediately.

Arc/Info users can then access the Oracle RDBMS and its data in a number of ways from within the GIS. The simplest way, but probably the least useful, is to initiate a subshell from within Arc/Info as follows:

Arc: ***sqlplus USERNAME/PASSWORD@T:av:oraprod*** [use your own username/password here]
Submitting command to Operating System ...
ORACLE_HOME = [/usr/local/oracle] ?

SQL*Plus: Version 3.0.11.1.2 - Production on Mon Mar 1 16:06:19 1993

Connected to:
ORACLE RDBMS V6.0.34.3.1, transaction processing option - Production
PL/SQL V1.0.34.2.0 - Production

SQL> ***select * from tab;***

TNAME	TABTYPE
MAG3_DATA	TABLE
MAGSAT_DATA	TABLE
MAGSAT_KEY	TABLE
MAGSAT_OBS	TABLE
TEMP_MAG	TABLE

6 rows selected.

SQL> *quit*
Disconnected from ORACLE RDBMS V6.0.34.3.1, transaction processing option - Production
PL/SQL V1.0.34.2.0 - Production
Arc:

This procedure suspends Arc/Info while the user accesses the Oracle RDBMS. There is no opportunity for the user to incorporate into Arc/Info any of the data selected by a SQL query. However, all the normal SQL*Plus functionality in Oracle is available to the user.

Normally a better procedure is to use the **dbmsexecute** command in Arc/Info. This command which is part of the Database Integrator, issues SQL commands to Oracle. It does not however support the Oracle SQL*Plus extensions so commands such as **describe** are not available (the user must use the Arc/Info **columns** command instead).

The **dbmsexecute** command can be used in two modes. In the first mode, the SQL command can be included on the **dbmsexecute** command-line. In this case the SQL statement is limited to a maximum of 128 characters. The other mode is called interactive mode, and in this case, the command can be as long as necessary. An example of the latter usage of the **dbmsexecute** command follows:

Arc: ***connect av USERNAME/PASSWORD*** [use your own username/password here]
Connection to AV successful.
Arc: ***show connects***
av
Arc: ***dbmsexecute av***

Enter DBMS specific command. (Enter "END" or a blank line when finished.)
>: ***select passno from magsat_data***
>: ***where latitude > -10***
>: ***and longitude between 150 and 152***
>: ***end***
PASSNO

1323
2217
2271
1060
1323
2217

Arc: ***disconnect av***
Arc:

The **dbmscursor** command provides another way to link Arc/Info and Oracle. This command provides the ability for a user to interactively step through a selected set of records. Alternatively, the **dbmscursor** command can be used from within AML (Arc Macro Language) routines to step through the selected set of records in a procedural way. Either way, the **dbmscursor** command provides a powerful means of accessing related records in the RDBMS. This can be particularly useful in handling one-to-many relationships because the user is able to access all the records, not just the first one (see below).

We do not go into the details of the **dbmscursor** command in this Record because its functionality is beyond the needs of most AGSO GIS users at present.

Linking the Arc/Info and Oracle Data Models

Once the Arc/Info and Oracle systems can communicate with one another through hardware and software, the next task is to establish links between the data holdings on the two systems. This is not as straightforward a task as it might be.

Establishing a *relate*

A *relate* in Arc/Info is the equivalent of a join in a relational database. The *relate* is a named relationship between an item in an INFO table and another instance of that item in some other INFO file or an external RDBMS attribute database. Note that a *relate* is not restricted to any particular table. A *relate* operates on any INFO table item which has the name specified by the user when the *relate* was created. Thus the *relate* **sites** created with the command:

Arc: ***relate add***
Relate Name: ***sites*** [NB. maximum length of name is 8 characters]
Table Identifier: ***ngma.sites*** [NB. maximum length of name is 128 characters]
Database: ***av*** [NB. maximum length of name is 8 characters]
INFO Item: ***siteid*** [NB. maximum length of name is 16 characters]
Relate Column: ***siteid*** [NB. maximum length of name is 32 characters]
Relate Type: ***first***
Relate Access: ***RO***
Arc: ***relate save sites***

would attempt to join any INFO item called siteid to the Oracle siteid field. Hence any active coverages with feature attribute tables containing the item siteid would be linked to Oracle.

Clearly care must be taken in establishing *relates* so as to avoid unwanted joins. The item names in the INFO tables should be unambiguous and should be unique when they contain different types of data. It is also a good idea to avoid creating *relates* on the recno, cover-ID and cover# items in INFO feature attribute tables. The first two items can change following a **clean** or during an Arcedit session. The latter two items will be renamed if the coverage is renamed or copied.

Note that the Relate Column in the RDBMS must be the same format as the INFO item if the relate is to work. Thus if the foreign key (the INFO item) is of type integer, then the primary key in the RDBMS must also be an integer field. It is generally not a good idea to use real number fields for the foreign and primary keys because the internal representations of these in INFO and the RDBMS may differ (e.g. differences in precision) and they may therefore not match properly. Similarly, character fields containing strings of blanks will not map correctly to character fields in the RDBMS. Lastly, it is important to note that INFO does not support the NULL data type (see *Managing Tabular Data, Arc/Info 6.1 Manuals,* p.5-20) so *relates* cannot be made to RDBMS columns containing NULL values (i.e. the NOT NULL constraint must be set for any columns in the RDBMS which are to be used in Arc/Info).

Relates created between INFO files and external RDBMSs such as Oracle can only be of the relate type *first* wheras relates between INFO tables can also use the *linear, ordered* and *link* relate types (see *Managing Tabular Data,* Arc/Info 6.0, p.2-9).

When the **relate add** command is used, a *relate* is added to what is known as the *relate environment* for that session of Arc/Info. Up to 100 *relates* can be defined as part of the *relate environment* at any one time but only 5 can be used simultaneously in an Arc/Info operation. The relate environment for a particular Arc/Info session can be saved to disk and subsequently recalled for use by using the **relate save <filename>** and **relate restore <filename>** commands.

Arc/Info supports nested *relates* up to a maximum of two levels deep (see *Arc Command References, Commands J-Z, Arc/Info Rev. 6.1 Manuals*, p. Relate-3). Thus it is possible for two tables to be joined via an intermediate table with a *relate* established between each of the main tables and the intermediary. This arrangement is called a *stacked relate* in Arc/Info.

Creating a new foreign key

The **additem** command in Arc and Arcedit allows users to add items in any position in an INFO file. This includes feature attribute tables which are just another type of INFO file. However, it is very important not to add or drop an item before the cover-ID item in feature attribute tables (see *Managing Tabular Data, Arc/Info 6.0 Manuals,* p.1-25). The cover-ID is the item in the feature attribute table with the form COVERNAME-ID (e.g. for a cover called rockchem, the cover-ID would be called ROCKCHEM-ID).

Indexes

Arc/Info uses two types of indexes. The first type are called spatial indexes and these are used to increase the speed of graphical queries on spatial features in Arc/Info coverages. The second type of indexes, called item indexes, are concerned with the attribute data in the Arc/Info system. Thus attributes such as the concentration of MgO at stream sediment sampling points in a point coverage can be indexed using an item index. These indexes are used to speed up the selection of attribute data based on logical expressions. For example:

Arcedit: ***editcoverage geochem label***
Arcedit: ***select MgO > 150***

This query would be much faster if the item MgO was indexed in INFO. Such an index could be created with the **indexitem** command:

Arc: ***indexitem***
Usage: INDEXITEM <info_file> <index_item>
Arc: ***indexitem geochem.pat MgO***

On the negative side however, item indexes slow down update operations in Arcedit very significantly because the index must be regenerated with each update. For this reason, ESRI generally recommend that if updates are likely to be made to a coverage, and if the database management is all to be done by INFO, then no more than three separate item indexes should be maintained for that coverage.

For coverages which use a *related* external RDBMS however, the situation is quite different. In this case, the foreign key in INFO must be indexed. Moreover, ESRI recommend that users should index all the appropriate columns in the external RDBMS. The preferred procedure from ESRI's perspective, is to index both the *relate* column (i.e. the primary key in the RDBMS) and the attribute columns being queried (e.g. MgO). This indexing must be done using the RDBMS and can only be carried out by the "system owner" of the database. The index can be created either by logging into Oracle directly or by accessing it through the Arc/Info Database Integrator functionality.

The overhead involved with updating these indexes is of course negligible to Arc/Info since it is all carried out by the RDBMS. Similarly, the storage requirements for the indexes, which can be considerable, must be met not by Arc/Info, but by the RDBMS. Clearly however, the RDBMS is also a finite resource and hence some selectivity must be exercised when choosing columns to index.

Sorting files in INFO

Records selected from INFO files are always returned in the order that they were inserted unless the INFO file has been sorted using the INFO or Arcedit **sort** command. INFO files can be sorted in ascending or descending order on one or more columns using these commands. Sorting can be very useful when doing *relates* because it can greatly speed up data retrievals. It can also provide some degree of logic in data retrievals involving one-to-many relationships. The latter point is important because INFO will only return the first matching record in such a relationship and the actual identity of that record can be haphazard. Moreover, Arc/Info does not test for one-to-many relationships and if it encounters one, it will not warn the user that it has done so.

Look up tables and other INFO files are often best sorted (and re-sorted following any inserts and updates). Feature attribute tables must not be sorted however because these must be ordered by the cover# attribute alone. Sorting these files will corrupt the coverage.

The type of sorting that INFO employs is not normally used in external RDBMS such as Oracle to sort attributes. The non-procedural nature of data storage and retrieval in an RDBMS generally makes it impossible to predict the actual order of data retrieval from query to query unless an explicit **order by** clause is used. Thus great care should be taken when handling one-to-many relationships between the data in Arc/Info coverages and Oracle databases. If such relationships do occur then only the first matching row in the RDBMS will be returned to Arc/Info unless the **cursor** or **dbmscursor** command is being used (see above).

Some Practical Examples

In this section we present two practical examples illustrating how appropriate *relate* environments can be set up in Arc/Info to allow access to data in the AGSO Oracle RDBMS. The procedures which we present specifically concern the SITES and the ROCKCHEM Oracle databases, though the general principles outlined can be applied to any of AGSO's Oracle databases.

Using an existing Arc/Info coverage with the SITES database (Lake Carey, WA)

The bedding_1 polygon coverage for the Lake Carey 1:100,000 sheet in Western Australia is typical of the data structure used in Arc/Info by AGSO for its NGMA mapping projects. This structure is summarised in Table 1:

DATAFILE NAME: BEDDING_1.PAT 03/02/1993
6 ITEMS: STARTING IN POSITION 1

COL	ITEM NAME	WDTH	OPUT	TYP	N.DEC
1	AREA	4	12	F	3
5	PERIMETER	4	12	F	3
9	BEDDING_1#	4	5	B	-
13	BEDDING_1-ID	4	5	B	-
17	ORIGNO	5	5	C	-
22	LOCNO	16	16	C	-
** REDEFINED ITEMS **					
17	UNIQUE_ID	21	21	C	-

Table 1 A typical Arc/Info data structure used with NGMA data

The UNIQUE_ID redefined term is a concatenation of the ORIGNO and LOCNO character items in the feature attribute table. UNIQUE_ID is type CHARACTER and it can be used as a foreign key to the Oracle RDBMS's NGMA databases.

The format of the NGMA SITES database which is central to the NGMA system of databases in AGSO is summarised in Table 2 and is more fully described by Ryburn et al (1993).

Name	Null?	Type
ORIGNO	NOT NULL	NUMBER(5)
SITEID	NOT NULL	CHAR(16)
FIELDID		CHAR(16)
OBSDATE		DATE
OBSTIME		NUMBER(4,2)
COUNTRYID	NOT NULL	CHAR(3)
STATE		CHAR(3)
GEOPROVNO		NUMBER(5)
SUBPROVNO		NUMBER(5)
DOMAINNO		NUMBER(5)
LOCDESC		CHAR(64)
HMAPNO		NUMBER(4)
QMAPID		CHAR(6)
EASTING		NUMBER(6)
NORTHING		NUMBER(7)
ACCURACY	NOT NULL	NUMBER(4)
HEIGHT		NUMBER(5)
HEIGHTACC		NUMBER(3)
DLAT		NUMBER(8,6)
NS		CHAR(1)
DLONG		NUMBER(9,6)
EW		CHAR(1)
METHOD	NOT NULL	NUMBER(3)
MAPID		CHAR(10)
AIRPHOTO		CHAR(36)
ST		CHAR(1)
OC		CHAR(1)
RO		CHAR(1)
TS		CHAR(1)
RC		CHAR(1)
OZ		CHAR(1)
MD		CHAR(1)
SC		CHAR(1)
RT		CHAR(1)
RP		CHAR(1)
ENTEREDBY		CHAR(16)
ENTRYDATE		DATE
GRIDREF		CHAR(10)
GEOGAREA		CHAR(64)

TABLE 2 The data structure of the NGMA SITES Oracle table

Note that the ORIGNO and SITEID fields, when concatenated (i.e. ORIGNO||SITEID), constitute the unique site identifier for the SITES database and all the other NGMA databases which depend on it.

The ORIGNO field in the Oracle SITES database is type NUMBER but in the Arc/Info coverage structure in AGSO the equivalent item, again called ORIGNO, is type CHARACTER (see Table 1). While Oracle treats the ORIGNO||SITEID concatenation as type CHARACTER and Arc/Info allows a *relate* to be constructed between UNIQUE_ID and ORIGNO||SITEID, in practice the *relate* fails. The error returned in this case is a "Column-Type mismatch".

Moreover, the UNIQUE_ID redefined term in INFO contains blanks whereas the ORIGNO||SITEID concatenation in Oracle does not. The blanks in UNIQUE_ID can be seen in the following listing:

```
---------------------------------------------------------------------------------------------
Arcplot: list
Usage: LIST <cover> <feature_class> {range} {item...item}
Usage: LIST <info_file> INFO {range} {item...item}
Arcplot: list test.pat INFO unique_id
Record  unique_id
   1  127  92967592
   2  127  92967096
   3  127  91967473
   4  127  91967477
   5  127  92967415
   6  127  92967153
   7  127  92967124
   8  127  92967126
   9  127  92967122
---------------------------------------------------------------------------------------------
```

The number of blanks occurring in the UNIQUE_ID redefined term depends upon the value of ORIGNO. The entries for ORIGNO in the NGMA SITES database in Oracle currently include one, two and three digit numbers only, though there is provision for numbers of up to five digits. These numbers are also stored in INFO in the ORIGNO item which, in this case, is a character field of width 5 (see Table 1). However, because character fields in INFO are always left justified (i.e. filled from the left hand side), any string of digits written to the ORIGNO item which is less than five characters long will have trailing blanks added to it. There can be two, three or four blanks inserted by INFO in this way depending upon whether the ORIGNO entry is three, two or one digit long.

To get around these problems we have constructed a view in Oracle of the SITES database called NGMA.ARCSITES. This view contains a new field called ORIGSITE which is a character field concatenated from the ORIGNO and SITEID fields in SITES but which contains the appropriate number of blanks within the character string as expected by INFO. The procedure for creating this view in Oracle is illustrated in Table 3:

```
SQLplus> create view arcsites as
   select
         rpad(to_char(ORIGNO),5,' ') || SITEID origsite,
         FIELDID,
         OBSDATE,
         OBSTIME,
         COUNTRYID,
         STATE,
         GEOPROVNO,
         SUBPROVNO,
         DOMAINNO,
         GEOGAREA,
         LOCDESC,
         HMAPNO,
         QMAPID,
         EASTING,
         NORTHING,
         ACCURACY,
         HEIGHT,
         HEIGHTACC,
         DLAT,
         NS,
         DLONG,
         EW,
         METHOD,
         MAPID,
         AIRPHOTO,
         OC,
         ST,
         RO,
         TS,
         RC,
         OZ,
         MD,
         SC,
         RT,
         RP,
         ENTEREDBY,
         ENTRYDATE
   from ngma.sites;

grant select on arcsites to public;
```

Table 3 The SQL statement used to create the ARCSITES view

A *relate* can now be made between UNIQUE_ID in the standard Arc/Info data structure and the ORIGSITE field in the arcsites view. The procedure for doing this is as follows:

Arc: ***CONNECT av USERNAME/PASSWORD*** [use your own username/password here]
Connection to AV successful.
Arc: ***relate add***
Relation Name: ***sites***
Table Identifier: ***ngma.arcsites***
Database Name: ***av***
INFO Item: ***UNIQUE_ID***
Relate Column: ***origsite***
Relate Type: ***FIRST***
Relate Access: ***RO***
Relation Name:
Arc: ***relate list***
Relate Name: ***SITES***
Table: ***ngma.arcsites***
Database: ***av***
Item: ***UNIQUE_ID***
Column: ***ORIGSITE***
Relate Type: ***FIRST***
Relate Access: ***RO***

Arc: ***relate save /isb/pramec/capeyork/arcinfo/sites***
1 Relates saved to file /isb/pramec/capeyork/arcinfo/sites

With the *relate* created, it is now possible to query the SITES database from Arc/Info. For example, the INFO **.pat** file for the polygon coverage test contains the following data:

Arcplot: ***list test points***

Record	AREA	PERIMETER	TEST#	TEST-ID	ORIGNO	LOCNO
2	0.000	0.000	2	2	127	92967096
36	0.000	0.000	36	36	127	92967220
38	0.000	0.000	38	38	127	92967217
115	0.000	0.000	115	115	127	92967521
131	0.000	0.000	131	131	127	92967517
152	0.000	0.000	152	152	127	92967520
158	0.000	0.000	158	158	127	91967121
159	0.000	0.000	159	159	127	91967119
160	0.000	0.000	160	160	127	91967120
161	0.000	0.000	161	161	127	91967117
163	0.000	0.000	163	163	127	91967116
166	0.000	0.000	166	166	127	91967126
169	0.000	0.000	169	169	127	91967101

Associated data held in the Oracle SITES database can be retrieved in a number of ways. First, the **list** command can be used in Arcplot to display the data in tabular form. For example we can use the sites *relate* that we set up with **list** to view the ORIGSITE and OBSDATE fields held in Oracle:

Arcplot: ***list test points sites//origsite sites//obsdate***

Record	sites//origsite	sites//obsdate
2	127 92967096	1992-05-08 00:00:00
36	127 92967220	1992-05-19 00:00:00
38	127 92967217	1992-05-19 00:00:00
115	127 92967521	1992-06-14 00:00:00
131	127 92967517	1992-06-14 00:00:00
152	127 92967520	1992-06-14 00:00:00
158	127 91967121	1991-07-26 00:00:00
159	127 91967119	1991-07-26 00:00:00
160	127 91967120	1991-07-26 00:00:00
161	127 91967117	1991-07-26 00:00:00
163	127 91967116	1991-07-26 00:00:00
166	127 91967126	1991-07-26 00:00:00
169	127 91967101	1991-07-25 00:00:00

Arcplot:

Data in an Arc/Info coverage can also be differentiated graphically on the basis of attributes held in the Oracle SITES database as can be seen in the following sequence of commands and in Figures 1 and 2.

Arc: ***&stat xterm***
Arc: ***arcplot***
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
All Rights Reserved Worldwide.
ARCPLOT Version 6.1 (June 30, 1992)

Arcplot: ***connect av USERNAME/PSSWORD*** [use your own username/password here]
Connection to AV successful.
Arcplot: ***relate restore sites***
Arcplot: ***mapex test***
Arcplot: ***markersize 0.4***
Arcplot: ***markersymbol 17***
Arcplot: ***points test*** [see Figure 1]
Arcplot: ***markersymbol 5***
Arcplot: ***reselect test points ^sites where qmapid = 'SH5106' and airphoto ~***
Arcplot: ***like '%/4/%'***
TEST points : 17 of 176 selected.
Arcplot: ***points test***
Arcplot: ***aselect test points***

TEST points : 176 of 176 selected.
Arcplot: ***markersymbol 1***
Arcplot: ***reselect test points ^sites where qmapid = 'SH5106' and airphoto ~***
Arcplot: ***like '%/6/%'***
TEST points : 9 of 176 selected.
Arcplot: ***points test***
Arcplot: ***markersymbol 49***
Arcplot: ***aselect test points***
TEST points : 176 of 176 selected.
Arcplot: ***reselect test points ^sites where qmapid = 'SH5106' and airphoto ~***
Arcplot: ***like '%/8/%'***
TEST points : 18 of 176 selected.
Arcplot: ***points test***
Arcplot: ***aselect test points***
TEST points : 176 of 176 selected.
Arcplot: ***markersize 0.5***
Arcplot: ***markersymbol 9***
Arcplot: ***reselect test points ^sites where qmapid = 'SH5106' and airphoto ~***
Arcplot: ***like '%/8/%'***
TEST points : 85 of 176 selected.
Arcplot: ***points test*** [see Figure 2]

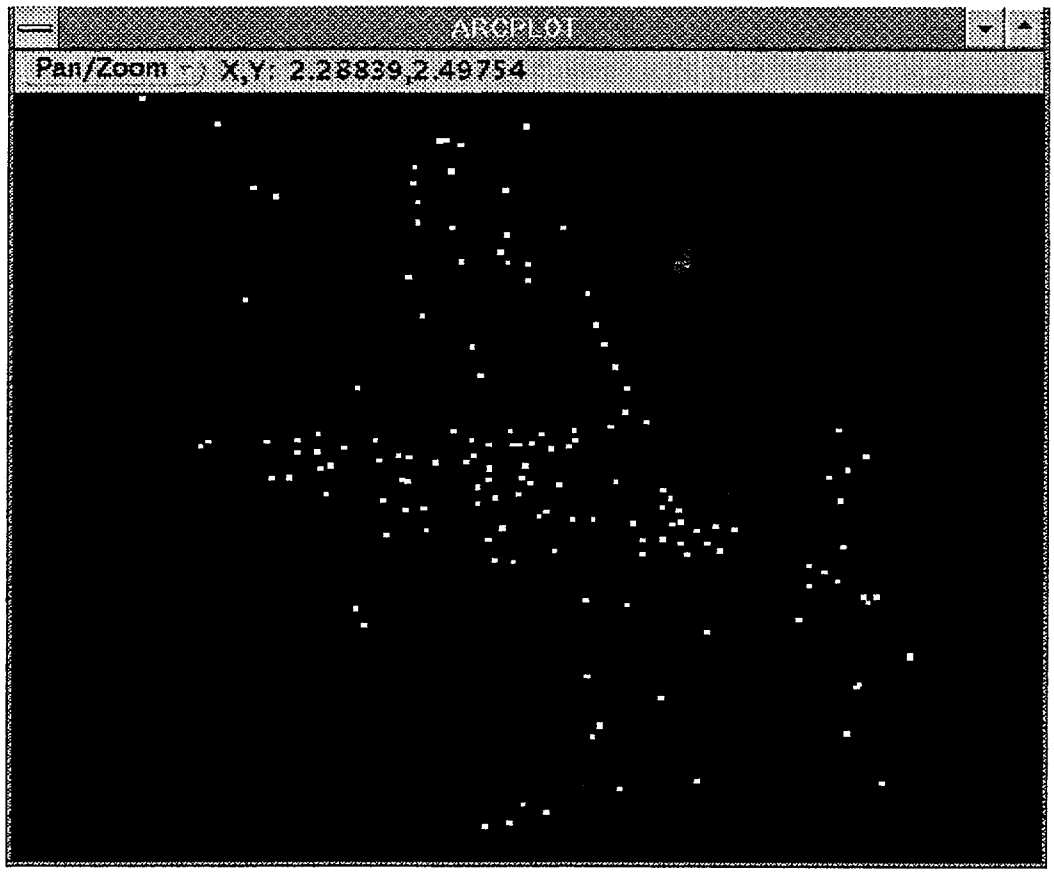

Figure 1 Point data in a demonstration Arc/Info coverage plotted by Arcplot. The points cannot be differentiated on the basis of any useful attributes.

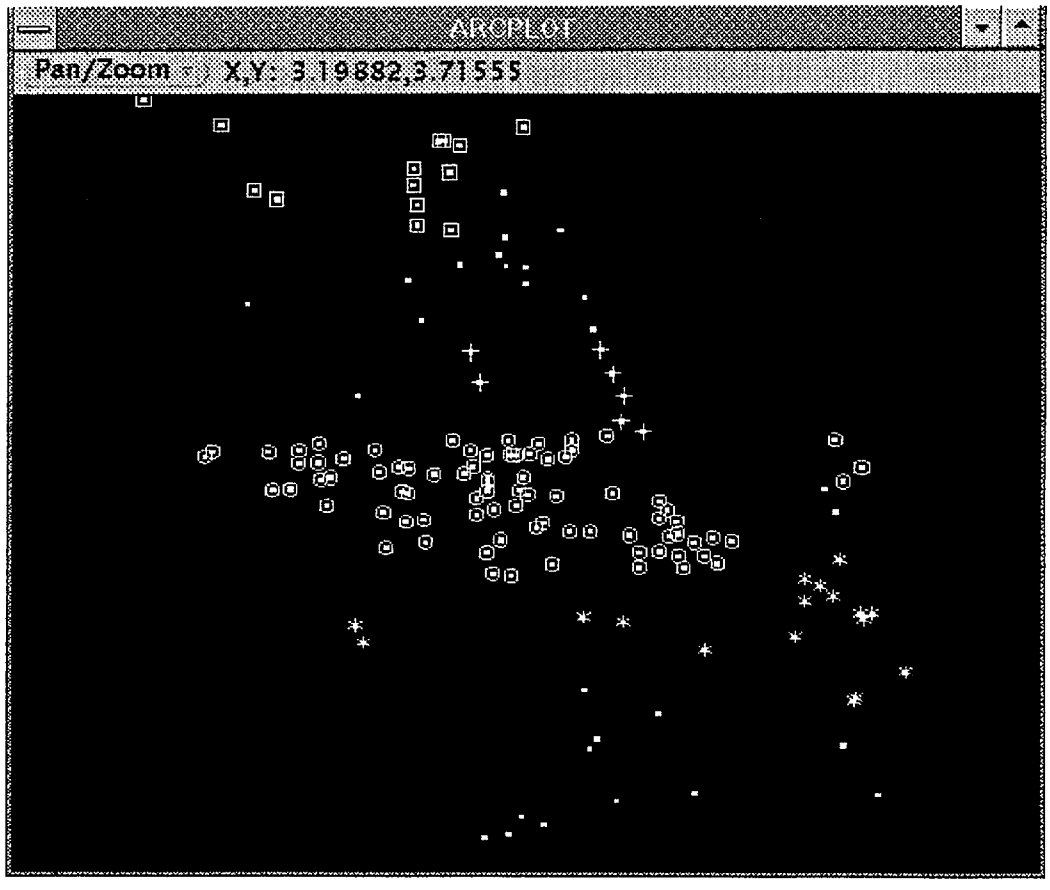

Figure 2 The same points plotted by Arcplot as in Figure 1 but differentiated on the basis of attribute data held in the AGSO SITES Oracle database.

Constructing a new Arc/Info coverage from the ROCKCHEM database (Blayney, NSW)

Attribute data in AGSO's ROCKCHEM database can be linked to Arc/Info project data in an analogous fashion to that described above for SITES. In this case we have created a simplified view of the data in the MAJORS table in ROCKCHEM combined with attributes from the NGMA SITES database which can be presented to Arc/Info as if it were a single table. This view is called ARCMAJORS and it is listed in Table 4. The SQL commands used to create the ARCMAJORS view are listed in Table 5.

To illustrate the use of the ARCMAJORS view we have created an Arc/Info points coverage of the data for the Blayney 1:100,000 sheet which is part of the Lachlan Fold Belt NGMA area. As with many things in Arc/Info, there are several ways to achieve the same result, so the example we present here should not be seen as the only way to build such a coverage, nor indeed may it necessarily be the best way to do it.

Name	Null?	Type
ORIGSITE		CHAR(21)
SUSER		CHAR(16)
SENTERED		DATE
FIELDID		CHAR(16)
OBSDATE		DATE
OBSTIME		NUMBER(4,2)
COUNTRYID	NOT NULL	CHAR(3)
STATE		CHAR(3)
GEOPROVNO		NUMBER(5)
SUBPROVNO		NUMBER(5)
DOMAINNO		NUMBER(5)
GEOGAREA		CHAR(64)
LOCDESC		CHAR(64)
HMAPNO		NUMBER(4)
QMAPID		CHAR(6)
EASTING		NUMBER(6)
NORTHING		NUMBER(7)
ACCURACY	NOT NULL	NUMBER(4)
HEIGHT		NUMBER(3)
DLAT		NUMBER(8,6)
NS		CHAR(1)
DLONG		NUMBER(9,6)
EW		CHAR(1)
METHOD	NOT NULL	NUMBER(3)
MAPID		CHAR(10)
AIRPHOTO		CHAR(36)
OC		CHAR(1)
ST		CHAR(1)
RO		CHAR(1)
TS		CHAR(1)
RC		CHAR(1)
OZ		CHAR(1)
MD		CHAR(1)
SC		CHAR(1)
RT		CHAR(1)
RP		CHAR(1)
SAMPNO	NOT NULL	CHAR(16)
ANALNO	NOT NULL	NUMBER(5)
BATCHNO		NUMBER(6)
RESTRICTED		CHAR(1)
RELEASED		DATE
SOURCENO		NUMBER(5)
METHODNO		NUMBER(5)
SIO2		NUMBER(4,2)
TIO2		NUMBER(4,2)
AL2O3		NUMBER(4,2)
FE2O3TOT		NUMBER(4,2)
FE2O3		NUMBER(4,2)
FEO		NUMBER(4,2)
MNO		NUMBER(4,2)
MGO		NUMBER(4,2)
NA2O		NUMBER(4,2)
P2O5		NUMBER(4,2)
H2OPLUS		NUMBER(4,2)
H2OMIN		NUMBER(4,2)
CO2		NUMBER(4,2)
LOI		NUMBER(4,2)
REST		NUMBER(4,2)
TOTAL		NUMBER(5,2)

Table 4 The ARCMAJORS view of the ROCKCHEM MAJORS table and data from SITES

```
REM  VIEW ARCMAJORS  PROVIDES A VIEW OF MAJORS AND SITES FOR ARCINFO

drop view arcmajors;
create view arcmajors as
  select
    rpad(to_char(SITES.ORIGNO),5,' ') || SITES.SITEID origsite,
    SITES.ENTEREDBY suser,
    SITES.ENTRYDATE sentered,
    FIELDID,
    OBSDATE,
    OBSTIME,
          COUNTRYID,
          STATE,
          GEOPROVNO,
          SUBPROVNO,
          DOMAINNO,
          GEOGAREA,
          LOCDESC,
          HMAPNO,
          QMAPID,
          EASTING,
          NORTHING,
          ACCURACY,
          HEIGHT,
          HEIGHTACC,
          DLAT,
          NS,
          DLONG,
          EW,
          METHOD,
          MAPID,
          AIRPHOTO,
          OC,
          ST,
          RO,
          TS,
          RC,
          OZ,
          MD,
          SC,
          RT,
          RP,
    SAMPNO,
    ANALNO,
          BATCHNO,
          RESTRICTED,
          RELEASED,
          SOURCENO,
          METHODNO,
          SIO2,
          TIO2,
          AL2O3,
          FE2O3TOT,
          FE2O3,
          FEO,
          MNO,
          MGO,
          CAO,
          NA2O,
          K2O,
          P2O5,
          H2OPLUS,
          H2OMIN,
          CO2,
          LOI,
          REST,
          TOTAL
  from ngma.sites, rockchem.majors
  where sites.origno = majors.origno
  and   sites.siteid = majors.siteid;
grant select on arcmajors to public;
```

Table 5 The SQL statement used to create the ARCMAJORS view

The first step is to extract the geo-referencing data and the unique foreign key for the entries in the SITES and MAJORS databases in Oracle. This can be done as follows:

```
SQL-Prod> set pagesize 0
SQL-Prod> set feedback off
SQL-Prod> spool blayney_majors.dat
SQL-Prod> select rownum||','||easting||','||northing||','||sites.orignno||','||sites.siteid
  2 from ngma.sites, rockchem.majors
  3 where hmapno = 8730                                    [the Blayney 1:100,000 sheet]
  4 and sites.orignno = majors.orignno
  5 and sites.siteid = majors.siteid
SQL-Prod> spool off
SQL-Prod> quit
Disconnected from ORACLE RDBMS V6.0.34.3.1, transaction processing option -
Production
PL/SQL V1.0.34.2.0 - Production
```

In our example, a file called blayney_majors.dat has now been created which contains the data we need for the coverage. This file should be edited with **vi** to remove all but the lines of data, and a line containing the word "end" (without inverted commas) should be added at the end of the file.

The file should then be transferred from the AViiON database server to the GIS server, zircon. This can be done with the TCP/IP file transfer protocol, **ftp** as follows:

```
/home1/pchopra% ftp zircon
FTP user (DG/UX TCP/IP Release 5.4.1 Jan 92) ready.
Connected to zircon.
220 zircon FTP server (SunOS 4.1) ready.
Name (zircon:pchopra):
331 Password required for pchopra.
Password (zircon: pchopra):
230 User pchopra logged in.
ftp> put blayney_majors.dat
200 PORT command successful.
150 ASCII data connection for blayney_majors.dat (192.104.43.110,2783).
226 ASCII Transfer complete.
16868 bytes sent in 0.060 seconds (2.7e+02 Kbytes/s)
ftp> bye
221 Goodbye.
/home1/pchopra%
```

Creating the coverage

In our case we use an existing **.tic** file in Arc/Info so that the geographic extent of our coverage is precisely that of the Blayney 1:100,000 sheet. If one is not available for the sheet area you are working with then it is a good idea to build a **.tic** file first. The procedure for doing this is described in the Arc/Info manuals.

We copy the existing tic points in the coverage chem into a new empty coverage called 8730_majors and then use the **generate** command to read in the data file we created with Oracle. We then **build** the topology for the new point cover and in so doing, create a **.pat** file:

Arc: ***create 8730_majors chem***
Creating coverage 8730_majors
Arc: ***generate 8730_majors***
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
All Rights Reserved Worldwide.
GENERATE Version 6.1.1 (December 23, 1992)

Generate: ***input blayney_majors.dat***
Generate: ***points***
Creating points with coordinates loaded from blayney_majors.dat
Generate: ***quit***
Arc: ***build 8730_majors point***
Building points...
Arc:

Part of this **.pat** file is listed in Table 6 and a plot of the data at this point in the process is shown in Figure 3.

Record	AREA	PERIMETER	8730_MAJORS#	8730_MAJORS-ID
1	0.000	0.000	1	1
2	0.000	0.000	2	2
3	0.000	0.000	3	3
4	0.000	0.000	4	4
5	0.000	0.000	5	5
6	0.000	0.000	6	6
7	0.000	0.000	7	7
8	0.000	0.000	8	8
9	0.000	0.000	9	9
10	0.000	0.000	10	10

Table 6 The **.pat** file created with **generate** and **build**

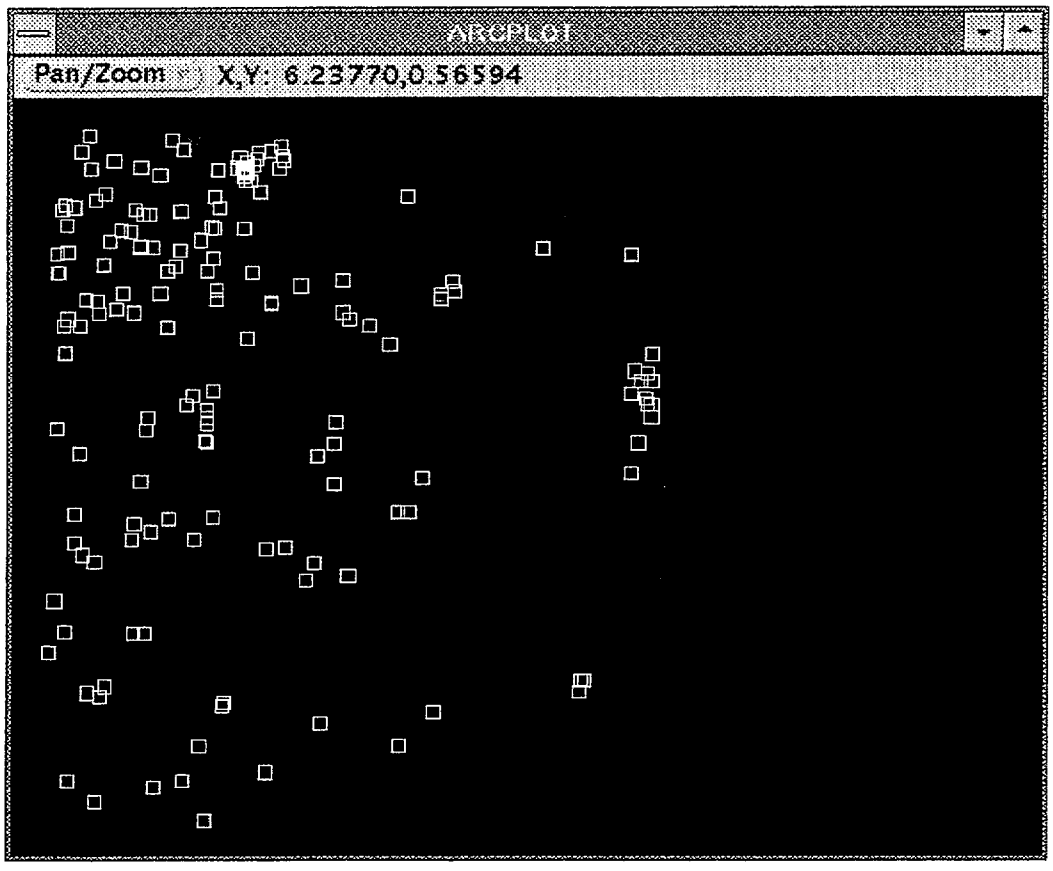

Figure 3 The point data from MAJORS as an Arc/Info coverage plotted by Arcplot. The points can't as yet be differentiated on the basis of their geochemistry.

Attaching the foreign key to the new coverage

The coverage has been created and its **.pat** file has been built. The next step is to extract the information necessary to build the foreign key in INFO which will be used to establish the *relate* to the geochemical and other attribute data which are stored in Oracle.

First, a temporary INFO table in the correct format (see Table 7) into which we can load the data in the file we generated in Oracle (i.e. in this example the file blayney_majors.dat) has to be created. The procedure is as follows:

Arc: ***tables***
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
All Rights Reserved Worldwide.
TABLES Version 6.1.1 (December 23, 1992)

Enter User Name: arc

Enter Command: ***define 8730_temp***
1
Item Name: ***row_number***
Item Width: ***3***
Item Output Width: ***3***
Item Type: ***i***
4
Item Name: ***easting***
Item Width: ***6***
Item Output Width: ***6***
Item Type: ***i***
10
Item Name: ***northing***
Item Width: ***7***
Item Output Width: ***7***
Item Type: ***i***
17
Item Name: ***origno***
Item Width: ***5***
Item Output Width: ***5***
Item Type: ***c***
22
Item Name: ***siteid***
Item Width: ***16***
Item Output Width: ***16***
Item Type: ***c***
38
Item Name: [press the RETURN key on an empty line]
Rounding record length up to produce even record length.
Enter Command: ***q stop***

COL	ITEM NAME	WDTH	OPUT	TYP	N.DEC
1	ROW_NUMBER	3	3	I	-
4	EASTING	6	6	I	-
10	NORTHING	7	7	I	-
17	ORIGNO	5	5	C	-
22	SITEID	16	16	C	-

Table 7 The format of the temporary Info table 8730_TEMP

The next task is to fill this temporary table with the data in the file generated by Oracle. To do this we again use the **tables** command in Arc:

```
---------------------------------------------------------------------------------------------
Arc: tables
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
        All Rights Reserved Worldwide.
TABLES Version 6.1.1 (December 23, 1992)

Enter User Name: arc
Enter Command: add from blayney_majors.dat
Enter Command: q stop
```

The structure of the resulting Info table is illustrated by the first 10 records:

Record	ROW_NUMBER	EASTING	NORTHING	ORIGNO	SITEID
1	1	685700	6281600	49	87840009
2	2	685800	6280100	49	87840010
3	3	705000	6255800	76	79620325
4	4	729800	6281400	76	79620326
5	5	702300	6277700	49	87840094
6	6	698100	6278700	49	87840096
7	7	698100	6277900	49	87840099
8	8	690300	6277200	49	87840101
9	9	690800	6278500	49	87840102
10	10	700100	6288200	49	90844001

The only data in the temporary table which we now need are the values in the last two fields because these constitute the unique foreign key for the data in the Oracle ARCMAJORS view. We extract these data with the **pullitems** command in Arc and write them back into the temporary table as the sole data, with the following command sequence:

Arc: ***pullitems 8730_temp 8730_temp***

Enter item names (type END or a blank line when done):
==========
Enter the 1st item: ***origno***
Enter the 2nd item: ***siteid***
Enter the 3rd item: ***end***

The temporary table now has the following structure as illustrated by the first 10 records:

Record	ORIGNO	SITEID
1	49	87840009
2	49	87840010
3	76	79620325
4	76	79620326
5	49	87840094
6	49	87840096
7	49	87840099
8	49	87840101
9	49	87840102
10	49	90844001

The data in the temporary table can now be added to the **.pat** file of the coverage by using the **joinitem** command in Arc:

Arc: ***joinitem 8730_majors.pat 8730_temp 8730_majors.pat 8730_majors-id 8730_majors-id LINK***
Joining 8730_majors.pat and 8730_temp to create 8730_majors.pat
Arc:

The final tasks are to create the UNIQUE_ID redefined term which will represent the foreign key for the *relate* to the Oracle ARCMAJORS view and to create the mandatory index for this foreign key. The steps to follow are:

Arc: ***tables***
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
All Rights Reserved Worldwide.
TABLES Version 6.1.1 (December 23, 1992)

Enter User Name: arc

Enter Command: ***select 8730_majors.pat***
200 Records Selected.

Enter Command: ***items***

COLUMN	ITEM NAME	WIDTH	OUTPUT	TYPE	N.DEC	ALTERNATE NAME	INDEXED?
1	AREA	4	12	F	3	-	
5	PERIMETER	4	12	F	3	-	
9	8730_MAJORS#	4	5	B	-	-	
13	8730_MAJORS-ID	4	5	B	-	-	
17	ORIGNO	5	5	C	-	-	
22	SITEID	16	16	C	-	-	

Enter Command: ***redefine***
Enter starting column: ***17***
Item name: ***UNIQUE_ID***
Item width: ***21***
Item output width: ***21***
Item type: ***c***

Enter starting column: [press the RETURN key on an empty line]

Enter Command: ***items***

COLUMN	ITEM NAME	WIDTH	OUTPUT	TYPE	N.DEC	ALTERNATE NAME	INDEXED?
1	AREA	4	12	F	3	-	
5	PERIMETER	4	12	F	3	-	
9	8730_MAJORS#	4	5	B	-	-	
13	8730_MAJORS-ID	4	5	B	-	-	
17	ORIGNO	5	5	C	-	-	
22	SITEID	16	16	C	-	-	
** REDEFINED ITEMS **							
17	UNIQUE_ID	21	21	C	-	-	

Enter Command: ***q stop***

Arc: ***indexitem***
Usage: INDEXITEM <info_file> <index_item>
Arc: ***indexitem 8730_majors.pat UNIQUE_ID***
Arc:

Using the link to ARCMAJORS with the Blayney Sheet data

Data in the Oracle ARCMAJORS view are now available to Arc/Info. The following four sequences of Arcplot commands and Figures 4 and 5 illustrate the usefulness of the link.

1) Listing the MgO data

Arc: ***connect av USERNAME/PSSWORD*** [use your own username/password here]
Connection to AV successful.
Arc: ***relate restore majors*** [use the same procedure as described on Page 14 to create the relate file "majors" first]
Arc: ***relate list***
Relate Name: SITES
Table: ngma.arcsites
Database: av
Item: UNIQUE_ID
Column: origsite
Relate Type: FIRST
Relate Access: RO

Relate Name: MAJORS
Table: rockchem.arcmajors
Database: av
Item: UNIQUE_ID
Column: origsite
Relate Type: FIRST
Relate Access: RO

Arc: ***arcplot***
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
All Rights Reserved Worldwide.
ARCPLOT Version 6.1.1 (December 23, 1992)

Arcplot: ***list 8730_majors points majors//origsite majors//mgo***

Record	majors//origsite	majors//mgo
1	49 87840009	2.07
2	49 87840010	3.55
3	76 79620325	0.50
4	76 79620326	0.97
5	49 87840094	6.48
6	49 87840096	9.14
7	49 87840099	8.95
8	49 87840101	6.65
9	49 87840102	5.52
10	49 90844001	1.32
11	49 90844005	5.84
12	120 LFB485	1.65
13	120 LFB486	2.30
14	120 LFB487	1.34
15	120 LFB488	0.69

[only the first 15 of the 200 points in the coverage are shown here]

2) Plotting the MgO data

Arc: ***connect av USERNAME/PASSWORD*** [use your own username/password here]
Connection to AV successful.
Arc: ***relate restore majors***
Arc: ***&stat 9999***
Arc: ***arcplot***
Copyright (C) 1989,1990,1991,1992 Environmental Systems Research Institute, Inc.
All Rights Reserved Worldwide.
ARCPLOT Version 6.1.1 (December 23, 1992)
Arcplot: ***mapex 8730_majors***
Arcplot: ***markerset mineral.mrk***
Arcplot: ***markersymbol 129***
Arcplot: ***markersize 0.1***
Arcplot: ***reselect 8730_majors points ^majors where mgo < 1***
8730_MAJORS points : 29 of 200 selected.
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***markersize 0.2***
Arcplot: ***reselect 8730_majors points ^majors where mgo > 1 and mgo < 3***
8730_MAJORS points : 62 of 200 selected.
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***markersize 0.3***
Arcplot: ***reselect 8730_majors points ^majors where mgo > 3 and mgo < 7***
8730_MAJORS points : 64 of 200 selected.
Arcplot: ***points 8730_majors***
Arcplot: ***markersize 0.35***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where mgo > 7 and mgo < 12***
8730_MAJORS points : 32 of 200 selected.
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***markersize 0.4***
Arcplot: ***reselect 8730_majors points ^majors where mgo > 12***
8730_MAJORS points : 12 of 200 selected.
Arcplot: ***points 8730_majors***

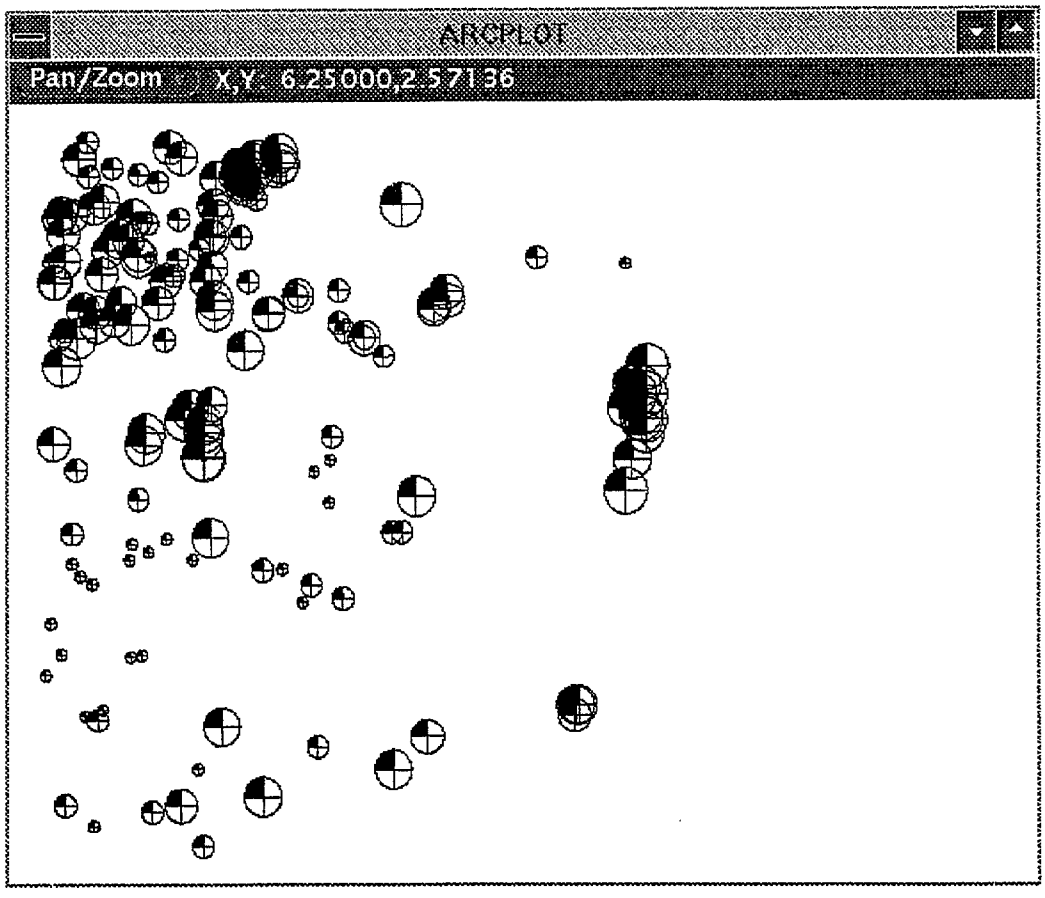

Figure 4
A plot of the concentration of MgO in rocks from the Blayney 1:100,000 sheet, NSW. Symbols represent:

Geochemical Parameter	Symbol Size	Symbol Colour
MgO < 1%	0.1	black
1% < MgO < 3%	0.2	black
3% < Mgo < 7%	0.3	black
7% < MgO< 12%	0.35	black
MgO > 12%	0.4	black

3) Plotting the data for Na_2O and K_2O for rocks with $SiO_2 < 60\%$

Arcplot: ***markersymbol 101***
Arcplot: ***reselect 8730_majors points ^majors where SiO2 < 60.0***
8730_MAJORS points : 131 of 200 selected.
Arcplot: ***markersize 0.15***
Arcplot: ***reselect 8730_majors points ^majors where (Na2O + K2O) < 2***
8730_MAJORS points : 28 of 200 selected.
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where SiO2 < 60.0***
8730_MAJORS points : 131 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where (Na2O + K2O) > 2 and ~***
Arcplot: ***(Na2O + K2O) < 4***
8730_MAJORS points : 20 of 200 selected.
Arcplot: ***markersize 0.25***
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where SiO2 < 60.0***
8730_MAJORS points : 131 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where (Na2O + K2O) > 4 and ~***
Arcplot: ***(Na2O + K2O) < 6***
8730_MAJORS points : 49 of 200 selected.
Arcplot: ***markersize 0.3***
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where SiO2 < 60.0***
8730_MAJORS points : 131 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where (Na2O + K2O) > 6 and ~***
Arcplot: ***(Na2O + K2O) < 8***
8730_MAJORS points : 20 of 200 selected.
Arcplot: ***markersize 0.35***
Arcplot: ***markercolor grey***
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***
8730_MAJORS points : 200 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where SiO2 < 60.0***
8730_MAJORS points : 131 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where (Na2O + K2O) > 8 and ~***
Arcplot: ***(Na2O + K2O) < 10***
8730_MAJORS points : 10 of 200 selected.
Arcplot: ***markercolor 'dark slate gray'***
Arcplot: ***markersize 0.4***
Arcplot: ***points 8730_majors***
Arcplot: ***aselect 8730_majors points***

8730_MAJORS points : 200 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where SiO2 < 60.0***
8730_MAJORS points : 131 of 200 selected.
Arcplot: ***reselect 8730_majors points ^majors where (Na2O + K2O) > 10***
8730_MAJORS points : 4 of 200 selected.
Arcplot: ***markercolor 'dim gray'***
Arcplot: ***markersize 0.45***
Arcplot: ***points 8730_majors***

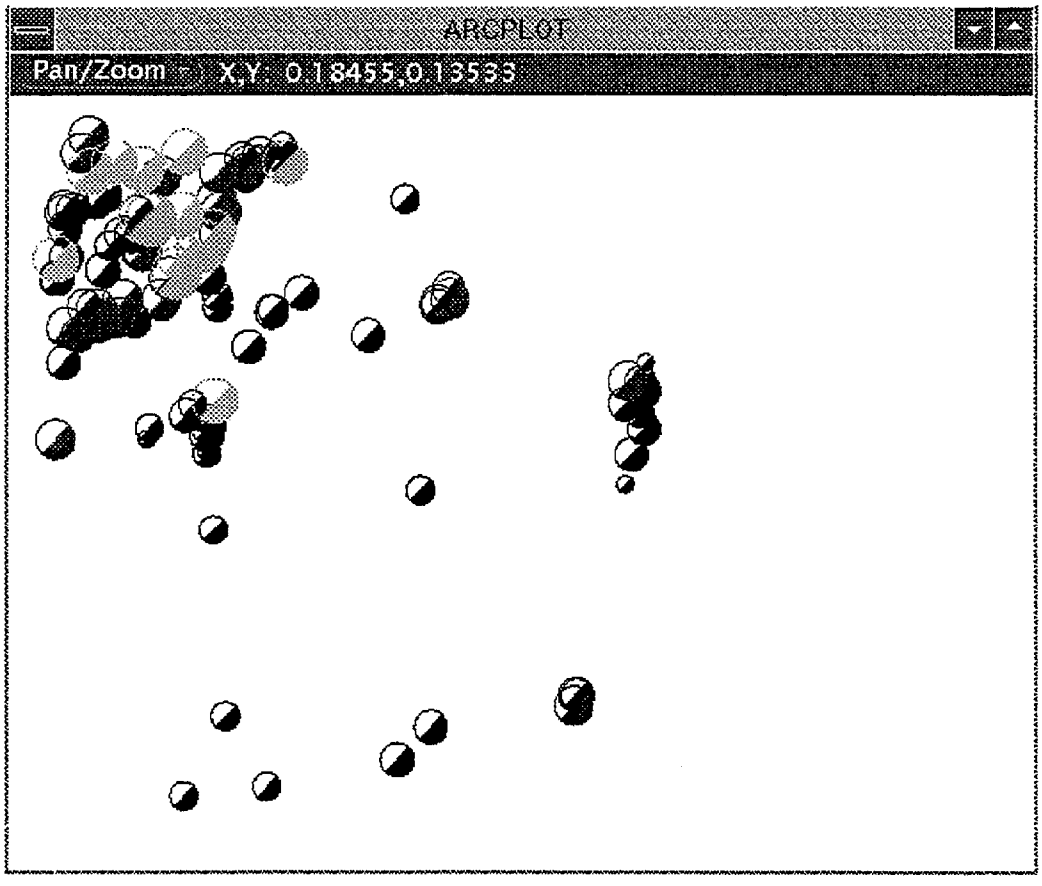

Figure 5
A plot of $Na_2O + K_2O$ for rocks with a total SiO_2 of < 60%. The Blayney 1:100,000 sheet, NSW.
Symbols represent:

Geochemical Parameter	Symbol Size	Symbol Colour
$Na_2O + K_2O < 2\%$	0.15	black
$2\% < Na_2O + K_2O < 4\%$	0.25	black
$4\% < Na_2O + K_2O < 6\%$	0.3	black
$6\% < Na_2O + K_2O < 8\%$	0.35	black
$8\% < Na_2O + K_2O < 10\%$	0.4	grey
$Na_2O + K_2O > 10\%$	0.45	grey

4) Plotting the relative proportions of Na_2O and K_2O.
In this case, the plot represents data from the NW corner of the Blayney 1:100,000 sheet, NSW.

Arcplot: ***mapex 8730_majors***
Arcplot: ***spotsize 85 0.4***
Arcplot: ***mapex ****
Define the box
Arcplot: ***pointspot circle 8730_majors majors//Na2O,1 majors//K2O,3***

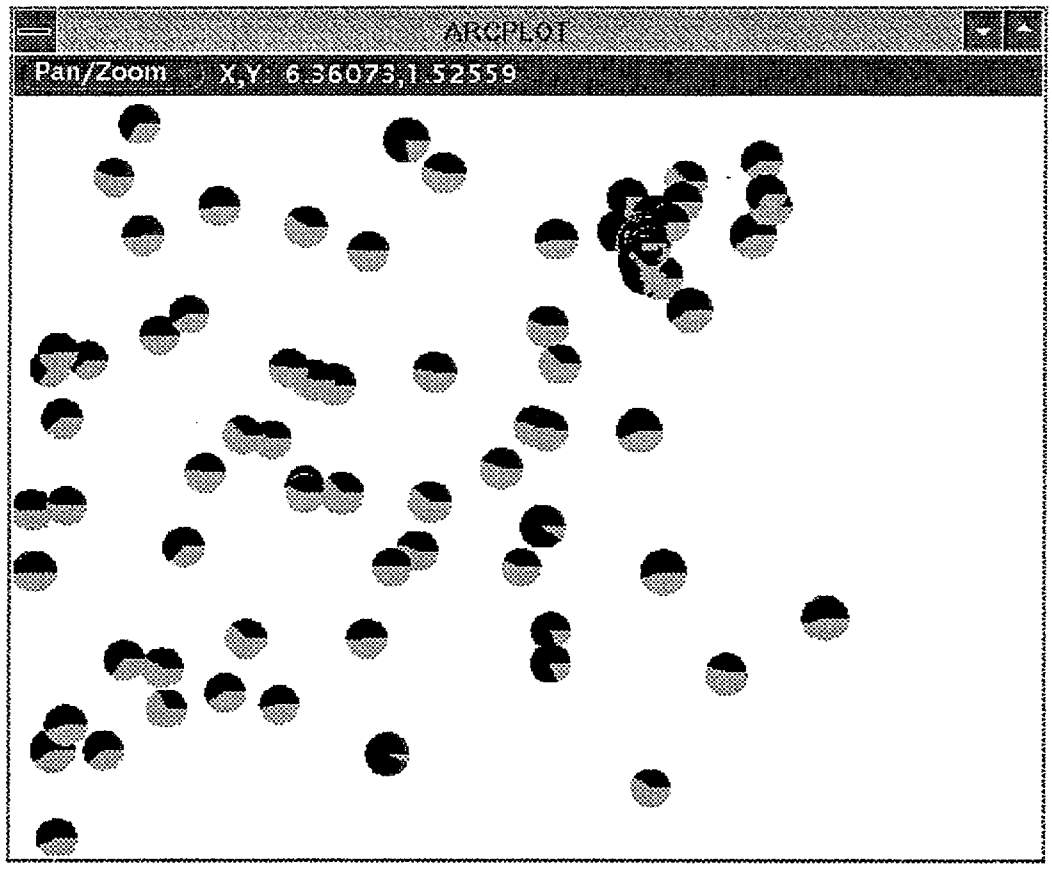

Figure 6
The relative proportions of Na_2O and K_2O for rocks in the NW corner of the Blayney 1:100,000 sheet, NSW.
The black portions of each symbol represent Na_2O and the grey K_2O.

Some Performance Aspects of Arc/Info Access to Oracle

The performance of Arc/Info operations which involve access to Oracle databases can vary very widely. For example, the queries that were used to generate Figures 1 - 5 were generally completed in a minute or two. The command used with Figure 6 on the other hand took several hours to complete and placed a very heavy load on Oracle.

It was apparent from the way that Figure 6 was drawn that Arc/Info queried Oracle for the Na and K values for each point on a one-by-one basis. Thus for each symbol in Figure 6, a separate SQL query was generated for Oracle and each of these took several minutes to complete at a time when overall Oracle usage was very low.

Robyn Gallagher has provided us with some unpublished written comments from ESRI on optimising Arc/Info access to an external RDBMS. In essence, these comments distinguish between two types of access: SQL access and INFO-style access. The **dbmsexecute** command in Arcplot is an instance of SQL access (see Page 6 for an example). INFO-style access is illustrated with the **reselect** command in Arcplot (Figures 4 and 5 and their associated command sequences are examples).

ESRI suggest that when there are few features selected from a large dataset, the INFO-style of access will give faster response than SQL access. Conversely, they say that when there are a large number of features selected, SQL style selects give faster results.

For routine query types in Arc/Info which rely on Oracle data, it will probably be worthwhile to determine which type of access to the RDBMS is more efficient and to then use this method by choice. Ultimately, the decision on access type must be made by the user and it is difficult to provide any hard and fast rules to guide in the choice. Practical experience is the only judge.

Conclusions

The software and data model links necessary to link AGSO's continental scale geoscience data holdings in the Oracle RDBMS to its NGMA project data in the Arc/Info GIS are described in this Record. Practical examples of how to establish links are given for the NGMA SITES and ROCKCHEM databases using project data from the Eastern Goldfields and Lachlan Fold Belt NGMA areas.

The procedure described for creating new Arc/Info coverages from Oracle data, while it is effective, is rather cumbersome. An Arc/Info Macro Language (AML) menu-driven routine could be developed which would take most of the tedium out of this process. This is an issue that could be addressed by ISB in the future if there is sufficient demand.

Acknowledgements

We wish to thank Mirek Kucka of ISB, and Mark Rattenbury and Robyn Gallagher of MLUP for their assistance and for useful discussions. This is not to say that they necessarily endorse all of our methods.

References

Chopra P.N., 1989. The BMR MAGSAT and 3rd-Order geomagnetic databases, *Bureau of Mineral Resources, Australia, Record,* 1989/44.

Ryburn, R.J., 1990. User's guide to the PetChem database, *Bureau of Mineral Resources, Australia, Record,* 1990/19.

Ryburn, R.J., Blewett, R.S. and Williams, P.R., 1993. User's guide to the NGMA field database, *Australian Geological Survey Organisation, Record,* 1993/49.