Spatially Enabling the Public Petroleum Data Model: Migrating from a Standard E&P Industry Database to a Geodatabase

Prepared for PPDM Association by
Ian Batty, Maguffyn Consultants Ltd.

Spatially Enabling the Public Petroleum Data Model: Migrating from a Standard E&P Industry Database to a Geodatabase

Summary

Executive summary

Spatial data is an inherent part of Exploration and Production (E&P) petroleum enterprises. Spatially enabling business information provides the ability to display and query spatial and non-spatial data in a Geographic Information System (GIS) using map based user interfaces. The Public Petroleum Data Model (PPDM) Association has developed a methodology to enhance an existing PPDM relational database with the spatial functionality of a Geodatabase. A PPDM version 3.7 database can store data in any or all of 25 different business and support modules. These modules can be spatially enabled as required to provide a spatial extension to the database. This allows Geoscientists access to the data in a manner that ultimately provides a better and more efficient way to conduct subsurface modelling, surface mapping, searching or analysis by integrating the existing relational business database with the spatial extensions provided by a Geodatabase.

Introduction

Volume of data managed in E&P

The search for oil and gas is a highly visual, data rich and spatially oriented endeavour. After approximately 130 years the exploration and production of oil and gas has resulted in approximately 1.2 million wells in continental U.S., 380,000 wells in Canada and 360,000 wells in the rest of the world (Batty 2002a). Exploration has created more than 1 million 2D seismic lines and 10,000 3D seismic volumes.

Existing data management structure

These data are saved and maintained by many organisations and many different business units within those organisations. Currently very few organisations are able to make all this data available via a GIS to all users; however, almost all the data is stored in relational databases and a large amount is stored in a standard data model (PPDM).

Challenges facing data management

In today's world of mergers, acquisitions, joint ventures, farm-ins and collaborative projects, the demand to exchange information has grown exponentially. The issues of spatial correctness and data standards have also grown correspondingly. Reliance on referential integrity amongst attributes in a database is the hallmark of an effective E&P information management system. The problem is that many types of spatial relationships among business objects cannot be effectively managed through referential integrity.

Achievements of PPDM Spatial Enabling projects

Since 1999, the PPDM Association has worked collaboratively with industry to develop standardized ways of integrating the detailed business information in a PPDM database with the powerful search, display, analysis and research capability of spatial engines. Later the project developed a methodology to link a PPDM relational database with the spatial functionality of ESRI Geodatabase.

Background to Spatially Enabling PPDM

Users require the ability to navigate into any geographical area such as a state, province, or sedimentary basin using a map to display any combination of features from a database. Many themes, including surface geology, geophysical data, engineering information, hydrography, and land ownership are required to support data analysis and interpretation. Although GIS has been used for a long time, historically the spatial information used to create these map displays was extracted from simple numeric columns in a relational database. Performance from these systems was generally poor, and true spatial analysis was difficult.

Functionality of a spatial database

The solution to this shortcoming is to use a spatial engine that supplies geometry information for business objects from a database as requested by a client GIS application. The business objects can then be displayed for the area specified. GIS technology allows a user to select the information needed to achieve their goal. A GIS can link data based on geographic locations from multiple databases to create a virtual integrated database that can be shared by all geoscientists. As other data stores become available, they can be integrated into the shared database as required. Effective query and retrieval tools display business objects on a map-based browser that allows data to be queried with "point and click" functions.

History of Spatially Enabling PPDM

Link geometry to business data

The first phase of Spatially Enabling PPDM was completed in September 2000. It provided for the simple use of spatial storage formats and a relational association between these formats and the data model. It delivered a methodology and associated code to create geometries for point, linear or polygonal business objects. The results were incorporated into PPDM version 3.5. This phase used spatial storage formats designed by ESRI and Oracle. However, there was no explicit involvement from either organisation.

Initiation of upstream Geodatabase project

During the ESRI PUG meeting in the spring of 2001 an initiative arose among a number of oil companies, data vendors and software developers within the E&P industry to develop an upstream Geodatabase for the petroleum sector. The result was a project, sponsored by oil companies and involving data vendors, software developers, information management consultants and ESRI to evaluate the feasibility of creating an industry standard Geodatabase. To validate the feasibility of creating a Geodatabase a subset of the PPDM data model was spatially enabled and incorporated into a Geodatabase.

Release of deliverables

The results from this project were released to industry in October 2002. They generated sufficient interest for another project to create a Geodatabase for the complete data model. The results of this project are available on the PPDM Association web site (www.ppdm.org).

Outline of Spatially Enabling PPDM

Overall objective

The objective of Spatially Enabling PPDM was to create a methodology that would allow member companies to enhance their existing PPDM databases and incorporate all relevant business data into a GIS by creating an upstream E&P Geodatabase. A key requirement was that existing processes and functions continue unimpeded.

Traditional Geodatabase design approach

A typical approach to designing a Geodatabase takes six steps (ESRI 2002)

1. Identify concepts, attributes and associations
2. Create classes, identifying subclasses, superclasses and potential subtypes
3. Create attributes for the classes and identify the range domains and coded value domains
4. Create the relationships between the classes, identifying if they are composite, simple or attributed
5. Design the spatial abstractions required to model the spatial classes
6. Create networks, identifying connectivity and junction rules

Steps completed using relational model

The fact that the PPDM Association has already created a data model means that the first four steps have been completed. However, as PPDM is entirely relational there are no spatial abstractions and no networks. The spatial abstractions can be developed for a spatially enabled database. The networks, connectivity and junction rules require the development of a Geodatabase.

Migrating from a Relational to Spatially Enabled Model

PPDM modelling methodology

The PPDM Association has a well-defined approach for deriving business requirements and then modelling the concepts into physical tables using a series of development principles. For full details see www.ppdm.org/whyppdm/ppdmway.html. These development principles cover defining modules, tables, columns and constraints. The resulting data model (PPDM version 3.6) consists of approximately 900 tables, 17,000 columns and over 5,000 referential constraints grouped in 15 business modules and 10 support modules. Developing a relational model is similar to the object oriented approach as can be seen in the following table

Relational	Object oriented
Module	Concept
Table	Class
Column	Attribute
Referential constraint	Relationship

Table 1: Relational/ Object cross reference

Comparison of PPDM modelling vs. object oriented modelling for Geodatabase

The different approaches do result in slight differences in the outcome: Some of the modules naturally map to a single spatial concept; some contain many spatial concepts. The tables have been designed to provide the most efficient storage of data from a relational point of view; however, the design of the table may not be suitable for access via the data management tools provided with a GIS. In an object oriented modelling process a class will either be a spatial object or will store information that a spatial object refers to. In the relational modelling process this is not the situation. Finally, and obviously, the columns and the referential constraints do not support spatial considerations.

However, despite all these differences, there are sufficient similarities that the relational model can be used as a basis to create a spatially enabled database.

Relational module to spatial concept

Identifying spatial concepts

In any data modelling process, either object oriented or relational, the intent is to store data associated with a number of real world objects or entities. Some of the objects are intrinsically spatial, for example a country. Some objects are spatial but are entirely dependent on the spatial location of another object, for example a test that occurs in a well. Some are potentially spatial, such a company or a person. Finally, some objects are not obviously spatial, for example a legal contract to provide services to an organisation.

Spatial modules in PPDM

The first phase in moving from a relational, non-spatial model to an object oriented, spatial model is to work out which relational modules map to spatial concepts.

Module	Spatial	Geometry stored in PPDM	Table Group
Contracts	Dependent on contracted object		
Facilities	Yes	Yes	Land Legal Description
Geodetic	No		
Interest sets	No		
Information and Records Management	Dependent on object in product		
Land Mineral Rights	Yes		Land Legal

			Description
Lithology	Yes	Yes	Land Legal Description
Obligations	Dependent on business object		
Production Reporting	Dependent on producing entities		
Production Entities	Yes	Yes	Land Legal Description, Well Geometry, Area Geometry
Projects	Yes	Yes	Land Legal Description
Restrictions	Yes	Yes	Land Legal Description
Seismic	Yes	Yes	Seismic Geometry
Stratigraphy	Dependent on area		
Support Facilities	Yes	Yes	Land Legal Description
Wells	Yes	Yes	Well Geometry

Table 2: PPDM Modules and geometry tables

As can be seen in Table 2, many of the modules are spatial; all modules that are spatial have geometry information stored in relational tables, but the geometry information for all the modules is being stored in just four locations

The process of spatial enabling a PPDM database is reduced to deriving the spatial classes and then creating the relevant geometry representation for each class.

Relational table to spatial class

Dividing tables into business and spatial classes

There is a large amount of data within each module; for example, the well module contains over 90 tables. Many of the tables contain business data that have no inherent spatial value. These tables are required to support business requirements and may be accessed using queries or reports from the GIS. Some tables contain business data that do have inherent spatial value; these tables are mapped into spatial classes. Spatially enabling a database creates this mapping. The full richness of the underlying database is exposed to the end user, through the use of the GIS, without modification to the data structure.

Dividing spatial classes into base and dependent classes

Within the spatial classes, some classes are intrinsically spatial and some are dependent on another spatial feature. The spatial classes are divided up into base spatial class i.e. classes that are intrinsically spatial, and dependent spatial classes. An example of a base class is a well or a seismic line. Examples of a dependent spatial class are well tests (a well test can only occur somewhere along a wellbore) or a seismic horizon (the seismic pick is made at a certain time associated with a common midpoint in the seismic line or volume).

The actual process to make these two divisions requires a thorough understanding of the original data model. However, certain features of the data model process can be used. If the relational tables are modelled using super-types and sub-types the super-type will probably not be spatially enabled although the sub-type may be.

Deciding on geometry representation for classes

Developing standard geometry representations

Generally, each spatial class can have multiple geometry representations. Each geometry representation will have different benefits to the application. Historically, application developers have created geometry representations to solve their immediate business need. The objective of the PPDM Spatial Enabling Project is to develop a standard geometry representation for each spatial class that allows the full richness of the data model to be accessed using existing GIS technology and maximizes the benefits while avoiding needless risk (Curtis, T. et al, 2002). The representation needs to allow all the dependent spatial classes to use the base class geometry and, at the same time, provide access to all the business tables in the data model.

Using GIS technology

The only way to meet all the requirements is to utilise existing GIS technology. Depending on the type of data being spatially enabled, the technology required, such as linear referencing, geocoding or topology would be different each time.

Relational storage of geometry

The information required to create the geometry for each class may be stored in a number of different tables within the module or it may reference another part of the data model. A GIS system typically uses the "best" data available to it. One of the advantages of moving from a CAD system to a GIS is that the data need not be optimised for a particular scale. Of course sometimes performance issues mean that certain assumptions need to be made, but there is no implicit requirement from the GIS not to use the most accurate and high-resolution data available.

Wells

The first spatial class to be considered was a well. Deriving the base and dependent spatial classes resulted in a single base class, `WELL`, and several dependent classes.

Base class relational geometry storage

The well geometry can be stored in three ways: as a directional survey, as a "node" which may have many versions, or as a denormalised value in the well header table. The tables used to store the location information depends on the type of data available for the geometry- a well may have a complete directional (deviation) survey, in which case the geometry is stored in the `WELL_DIR_SURVEY` table and its associated child tables. If the location of the top and base of a well is all that is known, the data is stored in `WELL_NODE` and its child tables. Finally, some applications require the surface and/ or bottom hole locations to be entered into the denormalised `WELL` table.

Geometry representation options

The geometry storage of the base class and the relationship between the dependent classes and the base class mean that the natural geometry of a well is a line. The dependent classes can then be spatially enabled by storing a measure value in the linear geometry and using a linear referencing system to locate the dependent feature.

In order to represent the linear geometry of the well bore as accurately as possible the GIS needs the "best" data available. This means that the geometry generation looks for a directional survey first, as this should be the best, most complete data. The next most complete data is in the `WELL_NODE` tables and finally, if there no data exists in either the directional survey or the nodes tables, the well header table will probably contain a well surface or bottom hole location

Dependent class relational geometry storage

The dependent spatial classes (such as completions, well tests or logs) do not have explicit geometry information stored in relational tables. Instead the dependent features are located at a known distance (the measured depth) along the base class (the well bore). Dynamic segmentation allows multiple features associated with a single deviated well to be located correctly in `(x,y)` space.

Traditionally, a well has been modelled as a point, using the data from `WELL` or a single value from `WELL_NODE`. This makes it difficult to access the dependent classes, as there typically is a one-to-many (`1:m`) relationship between the base class and the dependent classes. The data can be accessed using `relate` functionality in a GIS, but the features cannot be displayed directly on the map. Using a linear geometry and linear referencing data can be displayed on the map without using relate or join technology.

Link business data to geometry

Drilling technology allows a well to be deviated, which means that the top and base of the well are not coincident in `(x,y)` space. Therefore, point geometry cannot accurately represent the location of the wellbore, which may be an extremely complex line as it deviates to avoid subsurface hazards or targets multiple geological horizons.

As a result, location dependent classes i.e. events within a wellbore, may have different `(x,y)` coordinates from either the top or the base of the well and may not be anywhere near the straight line that is drawn between the two points. Examples of this include the point along a wellbore where it penetrates a particular geological formation or where a type of test is performed. Generating the geometry from the WELL_DIR_SURVEY tables creates the geometry closest to the real world possible with the data stored in PPDM. It allows multiple features associated with a single deviated well to be located correctly in `(x,y)` space using dynamic segmentation. Each feature will also have the correct depth and measured depth to be displayed on a map.

Spatial classes used to display business data

As a result of this approach the only class that is truly spatially enabled in the well module is the WELL_SDE table. This table inherits from the ESRI Feature class (i.e. a collection of spatial data in Geodatabase with the same shape type). All other classes inherit only from the ESRI Object class (i.e., the storage format for non spatial object in a Geodatabase).

The WELL_SDE class holds polyline geometry with Z (depth) and M (length) domains. Each line (wellbore) within this class contains information about the depth and length dimension along this line. The length is measured from zero at the starting point (reference point) to the total depth (TD) at bottom of the hole and acts a linear reference system for all events that occur along this path.

Seismic

Base class relational geometry storage

Seismic data is typically recorded in lines (2D) or over areas (3D). Each line or area is constructed from a series of points in the `SEIS_POINT` table. To create the relevant line or survey the data is grouped according to attributes in other columns (typically the `seis_set` column)

Geometry representation options

The two different types of seismic data seem to lend themselves naturally a line and area geometry representation respectively. Again, though, the challenge is to ensure that the full richness of the data model can be accessed using a GIS.

2D seismic is modelled as a line, again using measures, typically the shot point, to link to the rest of the data in the data model.

A 3D survey is actually a collection of in-lines and cross-lines that create a grid, consisting of a number of bins that represent a point in the subsurface. Therefore, a 3D seismic survey can be represented by a polygon (the overall area of the survey), a line (the individual shot and receiver lines) or a point (the bin centres). The actual geometry chosen will depend on the business requirement.

Link business data to geometry

The final issue was, as with wells, how to represent information from the subsurface that is associated with the seismic data. The PPDM Association method creates new subsurface linear geometries. Instead of being horizontal, a linear object is projected from each point in the subsurface, where data is available, to the surface and to an arbitrarily large subsurface depth. The linear object is created with a measure value that can be either time or depth as required.

Land data

Relational geometry storage

Land data is stored as polygons or as a list of parcels that can be accessed to group the relevant land parcels into a legally defined entity.

Geometry representation options

A legal agreement for mineral land rights typically describes the areal extents of the land, the subsurface stratigraphic zones over which it operates and the specific mineral rights that are granted in each stratigraphic zone. The stratigraphic zones are derived from well logs identified in the land agreement.

In order to accurately represent land data in a GIS a true three-dimensional storage model is required. This would allow the top of the land parcel (i.e. the surface), the areal extent (described in terms of the land parcels) and the base (the depth of the geological horizon) to be stored. Unfortunately, current storage formats, such as SDE Binary, do not support this functionality. Therefore, land data can only be represented as a surface polygon.

Representing mineral rights as a surface polygon creates a risk when the surface polygon does not accurately represent granted subsurface rights, such as areas where a mineral right may only include specific stratigraphic units, or even only specific substances within a stratigraphic unit.

Multiple polygons can exist for a single surface location, due to the subsurface grants. Therefore, spatial integrity and topological relationships must be based on additional information contained in relational tables.

Spatial queries between mineral lands and other business objects (such as wells) may identify a false intersection between a land right and a well. Real relationships must be identified on the basis of other business information, such as a contract or a three dimensional intersection. In each case, relational tables must be used to support the actual business relationships between objects.

Link business data to geometry

The data associated with the spatial class can still be accessed because PPDM has been designed to allow complete access to the land data from a purely relational point of view.

Implementing a Spatially Enabled PPDM database

Spatial data integration

The final step to creating a spatially enabled database is to actually implement it. There are a number of different technological solutions available. All need to create the geometry from the relational tables and then maintain consistency between the binary geometry and the relational tables. Integration between the two formats can be implemented, with either the binary geometry or the relational tables acting as the "master" store. The replication is a managed process and can be tailored to the specific business requirements of an organisation.

Minimising data duplication

However, the replication does create redundant data. It is possible to eliminate the duplication, by creating a database view that replaces the relational tables. The view is populated with values stored in the binary geometry. However, current technology generally renders this unfeasible due to performance considerations

Migrating from a Spatially Enabled Model to Geodatabase

The process of migrating from a relational database to a spatially enabled database is different to the process of designing a Geodatabase. A spatially enabled database assumes that 4 of the 6 steps to build a Geodatabase were completed when building the relational database. However, those steps were completed without consideration of the requirements of creating a Geodatabase. To design a Geodatabase the complete six-step process needs to be reconsidered i.e.

1. Identify concepts, attributes and associations

2. Create classes, identifying subclasses, superclasses and potential subtypes
3. Create attributes for the classes and identify the range domains and coded value domains
4. Create the relationships between the classes, identifying if they are composite, simple or attributed
5. Design the spatial abstractions required to model the spatial classes
6. Create networks, identifying connectivity and junction rules

Identify Concepts

Deciding on the concepts in the Geodatabase is no different from grouping the tables in a relational database into modules. Both approaches are attempting to solve business needs. The approaches that were used to create the spatially enabled database are directly applicable to the Geodatabase.

Create Classes

A Geodatabase is concerned with both the storage of data and the presentation of data to users. PPDM is designed to store all the data required to meet the business requirements of member companies in the most efficient and accurate manner.

Subsetting PPDM

The basic technique used to create the Geodatabase is to subset the existing PPDM database. However, there are a number of factors that need to be considered when subsetting the model:

1. Ensure that the subset is a recoverable subset. This means that it must be possible to edit the data in the subset and migrate it back to the complete data model. Obviously, the subset must be capable of reflecting any changes in data in the complete model.
2. The subset process must not result in any behavioural changes to the subset data model.

As long as these conditions are met it is possible to create classes in the Geodatabase.

UML diagrams

UML tools help design Geodatabase

Using UML to design the Geodatabases classes is useful to capture how and what data should be exposed to GIS users. Using the UML tools associated with Geodatabase design accommodates the limitations inherent in designing a data model that is spatially oriented. For example, it is not possible to create a join between a spatial class and business data if the relationship is a one-to-many (1:m) relationship. The only way to display this data is to use `relate` functionality or to denormalise the data.

Creating subset from PPDM

If the subsetting conditions are correctly met, the tables required to create the classes already exist in PPDM. The Geodatabase classes can either be created as database views (either spatial or traditional views) on top of the existing relational tables to represent the data or as managed data replication between the underlying PPDM database and the Geodatabase. Either approach can be used; the business requirements of an individual organisation will dictate which one is chosen.

Create Attributes

Attributes will be subsetted in the same manner as classes. For most attributes this process is driven by the business requirements. However, some attributes possess special qualities and must be treated differently- the most common example of this is the use of lookup values or reference values.

Reference value lookups

Reference values improve quality

Many classes use reference values to improve data quality by controlling data input. In a relational database this is enforced through external tables and referential integrity constraints. However, in order to fully take advantage of a Geodatabase reference values are converted into coded domains. Three approaches can be used to manage integration of relational reference values with Geodatabase coded domains (after Grise, 2003)

Implementation options

First, the Geodatabase client can run a process on a regular basis that looks for changes in the relational tables and updates the relevant Geodatabase domains. This will run as a scheduled process and should only run when there are no connections to the Geodatabase, as the system will require a schema lock on the tables affected by the update.

Second, you can register the domain tables with the Geodatabase and join to the feature tables to get the descriptive values for the codes that will be stored in the feature class. This will use relationship classes between the domain tables and the feature (or object) classes. This allows end users to control when the code values are displayed and when the reference values are displayed.

Finally, you can join the domain tables to the feature class to get the descriptive values for the codes that will be stored in the feature class. This join is performed on the database server, giving the Geodatabase designer complete control over the display of the data, but removing control from the end user. This approach will typically be used when the values in the reference tables will not be modified through a Geodatabase client application.

Create relationships

Types of relationships

Categorising relationships

PPDM has already modelled the objects into database tables and defined business relationships between those tables. However, these relationships are not explicitly spatial (though they might be). Relationships generally fall into four categories: spatial, spatial/relational, relational and process

Spatial relationships

Spatial relationship implementation

A spatial relationship is one that can **only** be expressed in a spatial manner. For example, a well must be assigned to the field that the well completion is located in. These relationships can be coded in a number of ways: as a class extension, using the SDE API to enforce validation as part of a database trigger or as part of topology validation.

A class extension will typically be used when data manipulation will occur through the Geodatabase client and a database trigger when the data manipulation occurs on the server. The topology approach will be used when data could have been manipulated in either fashion and the data control and validation needs to occur after the edits have occurred.

Full code to demonstrate how to implement a class extension or topology is beyond the scope of this article. However, the database validation code using the SDE API to test if the well is within the field, written in Java, might look like:

```
se_shape_well.isWithin(se_shape_fieldArea)
```

Spatial/ Relational relationship

A spatial/relationship is a relationship that is spatial, but can be expressed in terms of traditional relational constructs. A trivial example would be to ensure that the bottom of a well core is deeper than the top of the well core. This example might look like:

```
alter table well_core add constraint depth_check check (base_depth
- top_depth >= 0)
```

The actual implementation of this kind of relationship will depend on the skill set of the implementer in an organisation; it may be more efficient to code these relationships relationally or spatially.

Referential relationship

A referential relationship is one that is just that: a normal, referential foreign key constraint. Examples include simply checking that a well test is performed on a well that actually exists

```
alter table well_test add constraint wt_well_fk foreign key(uwi)
references well (uwi)
```

Process relationship

Finally, a process relationship is one that cannot be coded using existing technology. This class of relationship is fairly rare for those organisations that have an unlimited budget, but is often encountered in the real world where the cost of developing the code far outweighs the cost of documenting the process, publishing it and then ensuring that all the people involved follow the procedure.

Obviously, process relationships will not be stored in the database; therefore, they will not be considered further.

Relationship Implementation

Architectural considerations and technological limitations both mean that all relationships within any one type of relationship should be implemented using a single technology. This prevents conflicts between the relational referential constraints and the relationship classes created within the Geodatabase.

Creating networks

The final stage of creating a Geodatabase is to create the networks and connectivity rules. However, true networks are rare in PPDM and those that do exist generally require full three-dimensional geometry storage. Therefore, no networks, topology or other Geodatabase specific functionality have been created at this time.

Conclusion

Migrating from a traditional relational database to a Geodatabase provides many benefits in terms of data quality, data integrity and data availability.

Spatial enabling the relational model provides benefits to end-users as this makes the full richness of the relational model available to GIS users. Migration from a spatially enabled database to a Geodatabase must ensure that the business rules in the relational database are correctly captured in the Geodatabase. This requires a thorough understanding of the relational data model, the data content and the Geodatabase technology. Using the methodology developed by PPDM to migrate from a relational PPDM database via a spatially enabled database to an upstream Geodatabase can ensure that those rules are maintained.

About PPDM

PPDM Mission

The Public Petroleum Data Model Association (PPDM Association) is an organization that builds a worldwide standard model (PPDM) for the petroleum industry. It was founded in 1988 and incorporated in 1991. The mandate of the PPDM Association is to develop products based on standardized, open and platform independent technology. Currently, the PPDM Association has about 100 members.

PPDM model history

The first data model released was PPDM version 1.0 in 1989 and the latest release (PPDM version 3.6) was released in late 2002. The data model is built on open SQL 92 technology and has adapted since the first version to support new technology that has allowed more complex functionality.

PPDM modules

The PPDM Association data model (PPDM) is comprised of business and support modules. Business modules are developed entirely through work group processes. The modules support data and information requirements that enable operational activities of PPDM Association members. At the current time (version 3.6), the business modules consist of: Contracts, Facilities, Geodetic and Spatial, Interest sets, Information and Records Management, Land Mineral Rights, Lithology, Obligations, Production Reporting, Production Entities, Projects, Restrictions, Seismic, Stratigraphy, Support Facilities and Wells. For the latest developments, check at the PPDM Association (www.ppdm.org) and ESRI Petroleum User Group (PUG) website (www.esri.com/industries/petroleum).

The spatial developments described in this paper integrate with all existing modules and is adaptable to new modules that will be developed as PPDM continues to evolve

References

Batty I., T. Curtis, and R. Taylor, 2003. PPDM Version 3.7: Spatially Enabling 2003 Deliverable Release 1, The PPDM Association.

Batty, I. 2002a. PPDM Spatial II. Presentation at the public meeting PPDM, 31 May 2002.

Batty, I. 2002b. Spatially enabling the Public Petroleum Data Model. Presentation at the ESRI User Conference, San Diego, 10 July 2002.

Batty, I. 2002c. PPDM Association Spatial II Project. Presentation at the PPDM Fall Conference 2002. PPDM Association.

Batty, I. 2002d. Spatially enabling the Public Petroleum Data Model. Presentation at the PPDM workgroup meeting, 3 October 2002.

Batty, I., T. Curtis, and M. Diakuw. 2002. PPDM version 3.6: spatially enabling 2002 reference guide. PPDM Association, November 2002.

Curtis, T., C. Smith, M. Loughlean, P. Boorman, J. Connick, A. Stoness, S. Waters, A. Zolnai, B. Dick, L. Fabbro, 2002. PPDM projects: spatial pilot project chapter. February 21, 2002

Curtis, T. and I. Batty, 2002. Spatially enabling E&P data; 2003 project charter. October 2002

ESRI, 2001. Linear Referencing and Dynamic Segmentation in ArcGIS 8.1. ESRI Technical Paper May 2001.

ESRI, 2002. Geodatabase Design Course Notes. January 2002

MacDonald, A., 2001. Building a Geodatabase. GIS by ESRI. ESRI Press, 480 pp.

PPDM, 2002. PPDM version 3.6: road map, revised in March 2002.

PPDM, 2000. PPDM version 3.5: road map, revised in October 2000.

Grise S., 2003. Personal Communication