

An interoperable XML encoding for the exchange of Spatial OLAP data cubes in SOA environments

Etienne Dubé^{1,2}, Thierry Badard¹ and Yvan Bédard^{1,2}

1. Centre for Research in Geomatics (CRG)

2. NSERC Industrial Research Chair on Geospatial databases for Decision Support
Université Laval
Québec, Canada

Phone: +1 418 656 7116 Email: etienne.dube.2@ulaval.ca, thierry.badard@scg.ulaval.ca, yvan.bedard@scg.ulaval.ca

Abstract - XML and Web Services technologies have revolutionized the way data are exchanged on the Internet. Meanwhile, Spatial OLAP (SOLAP) tools have emerged to bridge the gap between the Business Intelligence and Geographic Information Systems domains. While Web Services specifications such as XML for Analysis enable the use of OLAP tools in Service Oriented Architecture (SOA) environments, no solution addresses the exchange of complete SOLAP data cubes (comprising both spatial and descriptive data and metadata) in an interoperable fashion.

This paper proposes a new XML grammar for the exchange of SOLAP data cubes, containing both spatial and descriptive data and metadata. It enables the delivery of the cube schema, dimension members (including the geometry of spatial members) and fact data. This XML format can be used in the development of Web Services to be deployed in various situations, not limited to traditional client-server platforms but also ubiquitous mobile computing environments.

I. INTRODUCTION

The World Wide Web Consortium's (W3C) XML (eXtensible Markup Language) specification [1] has gained large popularity as a data exchange format on the Internet. Along with the growing acceptance of Service Oriented Architecture (SOA) and Web Services technologies [2], it provides the means to exchange information between heterogeneous systems, independently of programming language, operating system, hardware and software platforms. XML grammars have been developed for various applications. This is also the case for geographic information systems (GIS), with the Geography Markup Language (GML) specification [3] from the Open Geospatial Consortium (OGC) and International Organization for Standardization (ISO), along with Web Services proposals such as Web Feature Service (WFS) [4], enabling the exchange of geographic features in distributed GIS applications.

In the meantime, SOLAP (Spatial OLAP) tools have been introduced to support the exploration and analysis of geospatial decision-support data. It is described as "a type of software that allows rapid and easy navigation within spatial databases and that offers many levels of information granularity, many themes, many epochs and many display modes that are synchronized or not: maps, tables and diagrams" [5]. Based on OLAP (On-Line Analytical Processing) [6] technologies merged with GIS components, these tools provides interactive tabular, graphical and cartographical views of spatial data cubes, which are multidimensional datasets comprising spatial data such as the geometry of countries and regions. This enables ana-

lysts to fully exploit the spatial component characterizing about 80% of all enterprise data [7].

Emerging SOLAP applications, especially in mobile computing environments, have prompted the need for an interoperable and reliable way to deliver SOLAP data cubes to heterogeneous clients. Unconventional usage patterns, such as disconnected operation when roaming between mobile networks, and inherent limits of mobile computing such as limited bandwidth and processing capabilities, renders traditional client-server OLAP poorly suited for these cases. SOA and Web Services have been considered as a suitable solution for these new applications [8].

GML and WFS specifications well cover the interchange of transactional geospatial data, traditionally found in operational systems. However, analytical systems such as SOLAP are often based on the multidimensional paradigm, in which data is represented in terms of hypercubes, composed of objects such as dimensions, members and facts. This contrasts with transactional data, typically represented by tables containing columns and rows. Several proposals have been made for the transmission of OLAP data through XML documents and Web Services. XML for Analysis (XMLA) [9], introduced by Microsoft and Hyperion, is targeted towards client-server applications using the MDX ("MultiDimensional eXpressions") query language [10], thus encapsulating MDX queries and result sets in XML messages. It is therefore not suited for the exchange of complete data cubes (data and metadata) for disconnected operation in mobile clients, and it lacks support for spatial data. XCube, a set of XML document templates for OLAP cubes [11], is capable of representing complete cubes, but has no provision for spatial data either. Another research paper proposes a Web Service, named "GML for Analysis" (GMLA) [12], which leverages the capabilities of XMLA and WFS, providing a way to combine multidimensional and spatial queries, and embed spatial data in the result sets. However, like XMLA, this technology is suited for client-server applications using a query-response paradigm, with result sets intended for presentation (i.e. in spreadsheets or other analytical clients), and not for the transmission of complete SOLAP cubes.

This paper introduces a new XML format for the exchange of multidimensional geospatial data, addressing the shortcomings of existing proposals. This format leverages the GML grammar for the representation of geographic features, and is inspired in some aspects by existing XML and Web Service standards used in the OLAP industry. Possible uses of such a format include the dissemination of SOLAP cubes (or subsets of these cubes) from a central data warehouse to distributed data marts, the technical integration of heterogeneous systems (e.g. migrating SOLAP

data between products of different vendors) and the implementation of Web Services for the delivery of SOLAP data to clients on the Internet or to mobile clients over wireless networks. The paper is divided as follows: first, requirements for the XML exchange of multidimensional spatial data are described. This is followed by the technical specification of the XML format itself. An application of this format, in the context of a geospatial decision-support Web Service for the delivery of spatial cubes to mobile devices, is also proposed. Finally, future outlooks and research challenges are discussed.

II. REQUIREMENTS FOR A SOLAP DATA CUBE XML FORMAT

Existing specifications and proposals for the XML exchange of either spatial data (GML, WFS) and multidimensional data (XMLA, GMLA and XCube) do not fully address the issue of delivering complete spatial data cubes, comprising both the metadata (the definition of the cube schema including dimensions, levels and measures) and the data (the dimension members and the fact cells). This prompts the design of a new XML format to overcome limits identified in existing ones. Several requirements have been identified for such a format, and are presented in this section.

A. Representation of both data (facts and members) and metadata (schema)

The exchange of complete data cubes requires encoding of both cube data and metadata (schema). Query services such as XML for Analysis specify only the transmission of data and a subset of the schema (query axes) describing the resulting cell set. While such a response format is suitable for presentation purposes (e.g. for displaying data in a spreadsheet), it is not sufficient for transmitting the full cube schema, i.e. dimensions, hierarchies and levels.

By including both the full cube schema along with fact and member data, a client application retrieving an XML-encoded data cube will have all the information necessary to use the cube in an autonomous manner. This is a first reason why the proposed XML encoding must provide for a representation of both data and metadata. In addition, metadata are necessary to assess the quality of measures.

B. Based on existing standards

In order to address interoperability with current and future systems, the XML encoding for SOLAP data cubes should be based on existing standards and specifications.

Geography Markup Language (GML) is a standardized way to encode geographic features in XML form, which can be used for spatial members in SOLAP data cubes. GML elements can be embedded inside the XML cube schema, seamlessly including geographic features. Nevertheless, other XML specifications, such as Scalable Vector Graphics (SVG), KML (as used in Google Earth) or GeoRSS could be substituted for this purpose, depending on the specific application context.

With regard to the XML representation of multidimensional data structures in data cubes, the only existing de-

facto specification is XMLA. However, as explained in the introduction, it does not fulfill all the requirements for the exchange of complete SOLAP data cubes.

C. Be independent of particular implementations in OLAP and GIS products

One of the reasons of using open standards is to avoid lock-in to proprietary software platforms. Since the proposed XML format for SOLAP data cubes aims to be open and interoperable, it must not be tied to particular OLAP software. This can be achieved by choosing a common subset of features supported by most OLAP servers. Most of them support some form of a multidimensional data model, composed of dimensions, hierarchies, levels, members, measures and facts; these objects have to be expressed in a homogeneous form in the XML encoding. As previously explained, using an open standard like GML also ensures interoperability with regards to geographic data. Independence from proprietary query languages and APIs is also a requirement. In this way, systems implementing the proposed XML encoding could be based on almost any OLAP and GIS products from the market. The possibility of exchanging SOLAP data between systems from different vendors is also a benefit of this approach.

D. Support for spatial dimensions and spatial members

SOLAP cubes are characterized by the inclusion of spatial members in geographic dimensions, and sometimes spatial measures [13]. The proposed XML encoding must provide the mechanisms for including the spatial component in these members and measures. This requirement is fulfilled by using an XML-based standard such as GML for encoding geospatial objects, as previously explained.

E. Support of shared dimensions

More than one cube can refer to the same dimension: for example, a “sales” cube and another “payroll” cube could refer to a same “employee” dimension. In the literature, such shared dimensions are often referred to as conformed dimensions [14]. In addition, this technique allows one to perform analysis across different cubes, by joining them by their shared dimensions.

The encoding should support this feature, by not restricting a dimension to the context of a single cube. Multiple cubes can refer to a shared dimension document, to avoid replicating the dimension schema (hierarchies and levels) and data (members) redundantly across cubes.

F. Describing calculated measures and members

Sometimes, calculated measures (or members) are defined to display expressions computed at run-time. Languages such as MDX allow users to define calculated members within a query. Most OLAP products also offer the capability to declare calculated members in the cube schema; in this way calculated measures and members are accessible to all users of the cube. These predefined calculated members should also be expressed as part of the cube metadata encoded in XML form, so clients would have access to these members. However, the question of calcu-

lated spatial measures and aggregate operators on spatial members is still a research issue, currently being investigated by our research group. By now, only numerical calculated measures are supported.

G. Separation of contents and presentation

Separation of contents and presentation is a design practice that aims to distinguish between the expression of semantics, as represented in data, and the presentation of the information to end users. For example, the same data in a SOLAP cube can be presented in different forms: cross-tabs, charts, reports and maps. Moreover, the output device can vary: contents can be displayed in an interactive SOLAP desktop client or in a Web page, it can be printed or it can be displayed on the small screen of a mobile device such as a personal digital assistant (PDA) or a cell phone. Finally, the graphic semiology used to represent data depends on the context (purposes, users, standards).

Accordingly, the XML cubes should contain only the data and its structure, i.e. cube schema. Any information specifying how to display the data to users, i.e. text format, chart colors and style, classes (ranges) for measure values, map symbols and so on, would be external to the XML data cube documents. The rendering and display format of SOLAP data can vary greatly and is dependant of the capabilities of the client application. Moreover, presentation is subject to customization by the user in some applications. For these reasons, the proposed XML format does not specify a way to encode presentation information.

III. SPECIFICATION OF THE XML ENCODING

Considering the requirements presented in the last section, a specification of the XML encoding for SOLAP data cubes has been designed. The proposed encoding is inspired from the XCube paper [11] but uses grammar and semantics closer to the data models of most OLAP software (specifically, terminology used in the MDX query language). A SOLAP data cube encoded with this specification would consist of three documents, respectively for the cube schema, dimension members and cell data. Separating the cube contents in three parts allows for progressive discovery of cube metadata and data. There are three root elements in the model, one for each type of document, i.e. *CubeSchema*, *CubeMembers* and *CubeCells*. A W3C XML Schema (WXS) definition [15] was developed for formalizing the semantics of the encoding and for validation of data cube documents¹.

A. CubeSchema document

The *CubeSchema* document encodes information pertaining to the structure of the cube. It identifies the cube name and all its dimensions with their respective hierarchies, levels and properties. An example *CubeSchema* document is presented in Fig. 1.

At the root of the document lies the `<cubeSchema>` element. Optional `<include>` elements can be used to refer

to external XML documents, to be included in the current cube schema. The `<cube>` element is used for defining the schema of a cube; there can be more than one in a single document. The optional `<sharedDimensions>` element provides for the definition of dimensions shared across cubes. Shared dimensions can come from external files, as specified in `<include>` elements; this allow sharing dimension definitions across many CubeSchema documents.

The `<cube>` element has one or more `<dimension>` elements as children. A dimension is identified by its *name* attribute. An *isMeasure* optional attribute specifies if this is the measure dimension, i.e. whose members identify the measures. A dimension can be defined in two manners: either with a `<dimensionDef>` element directly embedded in `<dimension>`, defining a dimension which is local to the cube, or with a `<dimensionUsage>` element, which refers to the name of a shared `<dimensionDef>`, defined in the `<sharedDimensions>` element (child of `<cubeSchema>`).

A `<dimensionDef>` element has one or more `<level>` children, defining which levels compose the dimension. Each `<level>` element can contain zero or more `<rollupTo>` children, identifying the parent level. More than one parent level can be specified when using multiple hi-

```
<?xml version="1.0" encoding="UTF-8"?>
<cubeSchema xmlns="http://geosoa.scg.ulaval.ca/GeoCubeML"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://geosoa.scg.ulaval.ca/GeoCubeML
  CubeSchema.xsd">

  <include fileLocation="WholesalersDimensions.xml" />

  <cube name="Sales">
    <dimension name="Measures" isMeasure="true">
      <dimensionDef>
        <level name="MeasuresLevel">
          <property name="aggregators">
            <contentType ref="aggregator"
              minOccurs="1" maxOccurs="unbounded" />
          </property>
          <property name="measureType">
            <contentType ref="contentType" />
          </property>
        </level>
      </dimensionDef>
    </dimension>
    <dimension name="Store">
      <dimensionUsage defName="WallyWorld Stores" />
    </dimension>
    <!-- ... other dimensions can be defined here ... -->
  </cube>

  <sharedDimensions>
    <dimensionDef defName="WallyWorld Stores">
      <hierarchy name="Administrative limits" default="true"/>
      <hierarchy name="Population areas" />
      <level name="(All)" isAll="true" />
      <!-- ... (Country, Province, Region levels) ... -->
      <level name="City">
        <inHierarchy hierarchy="Administrative limits" />
        <rollupTo level="Province" />
        <property name="City geometry">
          <contentType ref="gml:AbstractSurface" />
        </property>
      </level>
      <level name="Agglomeration">
        <inHierarchy hierarchy="Population areas" />
        <rollupTo level="Province" />
        <property name="Agglomeration geometry">
          <contentType ref="gml:AbstractSurface" />
        </property>
      </level>
      <level name="Retail outlet">
        <rollupTo level="City" />
        <rollupTo level="Agglomeration" />
        <property name="Retail outlet geometry">
          <contentType ref="gml:Point" />
        </property>
        <property name="Number of floors">
          <contentType type="xs:integer" />
        </property>
      </level>
    </dimensionDef>
  </sharedDimensions>
</cubeSchema>
```

Fig. 1. CubeSchema document example

¹ The WXS definition (XSD file) and UML data model are available at: <http://geosoa.scg.ulaval.ca/GeoCubeML/>

erarchies. In this case, `<hierarchy>` elements must figure under `<cube>`, to declare the hierarchy names and which to use by default; `<inHierarchy>` elements are used under `<level>` to specify which hierarchies (one or more) the level is part of. Omitting `<inHierarchy>` means that a level is shared by all the hierarchies of the current dimension. Levels can also have optional `<property>` elements, declaring member properties (i.e. attributes). Properties have a name, and contain a `<contentType>` element, specifying either which WXS simple data type (with the `type` attribute) or which complex XML elements (with the `ref` attribute, referring to an element declared in a WXS definition, with optional `minOccurs` and `maxOccurs` attributes defining cardinality) can be used as property values. This design choice, leveraging W3C XML Schema (WXS), allows the inclusion of any XML element to be used as member attributes. In the example, properties are declared with “gml:AbstractSurface” and “gml:Point” element references, respectively allowing the use of GML surface (region) or point geometries. Any other XML encoding of geometry could be substituted instead (for example SVG).

B. CubeMembers document

Once a cube schema is specified with the *CubeSchema* document, the members in its dimensions have to be described. This is done using the *CubeMembers* document (Fig. 2). Members are represented as `<member>` elements, and are contained in successively embedded `<cubeRef>`, `<dimensionRef>` and `<levelRef>` elements, referring to the cube, dimension, and level defined in *CubeSchema*. A member has a `name` attribute, uniquely identifying it, and can contain `<rollupTo>` elements, describing the roll-up relationships to other members. These can be to members of a higher level, or to members of the same level as it is the case in parent-child hierarchies. More than one `<rollupTo>` element can be specified for a member, to account for multiple hierarchies (i.e. a member in a level rolling-up to members in two alternate levels, in different hierarchies). In this case the parent level must be specified (with the `level` attribute). Values for member properties are represented with `<propertyValue>` elements. This element contains data in any form authorized by XML, and constrained by the WXS type specified for the property in *CubeSchema*. The client application is responsible for ensuring conformity to the property’s data type (i.e. by validating `<propertyValue>` contents against the WXS definition for the namespace of the data type.) Finally, optional `<include>` elements can also be used under `<cubeMembers>`, to specify members stored in external files (e.g. for shared dimensions).

In the example, we see a special case of member in the “Measures” dimension. These are measures, defined using `<member>` elements, and having the following special properties which are to be interpreted by the client application. The “aggregator” property contains `<aggregator>` elements, defining the type of operator used for aggregation (sum, average, etc.). More than one `<aggregator>` element can be specified, to account for non-additive or semi-additive measures (for example stocks or population, which cannot be summed across time) [14]. Additional aggregators must specify which dimension they apply to (using the `dimension` attribute), indicating the client application which aggregative formula to use, instead of the default one, when summarizing along this specific dimen-

sion. Finally, the “measureType” property is used to specify the WXS data type or element for cell values corresponding to this measure, using the same `<contentType>` element introduced for the *CubeSchema* document. This allows for a great flexibility in the representation of non-numerical measures, for example geometric spatial measures [13].

C. CubeCells document

CubeCells, the last type of document, is shown in Fig. 3. It contains the individual cells representing the facts of the cube. A `<cubeRef>` element identifies the cube (described in a *CubeSchema* document) to which the cells belong. Then, each cell is represented as a `<cell>` element, composed of a `<tuple>` with several `<memberRef>` elements (each one referring to a member in each dimension) and a `<cellValue>` for the value of the fact. A value can be of either a simple data type (number, string) or a complex one (e.g. a GML geometry representing a spatial measure). Usually, a *CubeCells* document will contain only the cells corresponding to leaf members of each dimension, or in

```
<?xml version="1.0" encoding="UTF-8"?>
<cubeMembers xmlns="http://geosoa.scg.ulaval.ca/GeoCubeML"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://geosoa.scg.ulaval.ca/GeoCubeML
  CubeSchema.xsd">

  <cubeRef name="Sales">
    <dimensionRef name="Measures">
      <levelRef name="MeasuresLevel">
        <member name="Sales amount">
          <propertyValue property="aggregators">
            <aggregator><sum /></aggregator>
          </propertyValue>
          <propertyValue property="measureType">
            <contentType type="xs:decimal" />
          </propertyValue>
        </member>
        <!-- other measures defined here -->
        <member name="Quantity in stock">
          <propertyValue property="aggregators">
            <aggregator><sum /></aggregator>
            <aggregator dimension="Time">
              <avg />
            </aggregator>
          </propertyValue>
          <propertyValue property="measureType">
            <contentType type="xs:integer" />
          </propertyValue>
        </member>
      </levelRef>
    </dimensionRef>
    <dimensionRef name="Store">
      <levelRef name="(All)">
        <member name="All Stores" isAll="true" />
      </levelRef>
      <!-- ... members in other levels (city, region, ...) -->
      <levelRef name="Province">
        <member name="Prince Edward Island">
          <rollupTo member="Canada" />
          <propertyValue property="Province geometry">
            <gml:MultiPolygon>
              <gml:polygonMember>
                <gml:Polygon>
                  <gml:outerBoundaryIs>
                    <gml:LinearRing>
                      <gml:coordinates>
                        -63.17146699969915,46.126305000168074
                        <!-- ... -->
                        -63.17146699969915,46.126305000168074
                      </gml:coordinates>
                    </gml:LinearRing>
                  </gml:outerBoundaryIs>
                </gml:Polygon>
              </gml:polygonMember>
              <!-- ... -->
            </gml:MultiPolygon>
          </propertyValue>
        </member>
        <!-- ... other members ... -->
      </levelRef>
    </dimensionRef>
  </cubeRef>
</cubeMembers>
```

Fig. 2. CubeMembers document example

```

<?xml version="1.0" encoding="UTF-8"?>
<cubeCells xmlns="http://geosoa.scg.ulaval.ca/GeoCubeML"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://geosoa.scg.ulaval.ca/GeoCubeML
  CubeSchema.xsd">
  <cubeRef name="Sales">
    <cell>
      <tuple>
        <memberRef dimension="Measures" level="MeasuresLevel"
          member="Sales amount" />
        <memberRef dimension="Product" level="Item"
          member="Fuzzy dice" />
        <memberRef dimension="Time" level="Day"
          member="2007-08-23" />
        <memberRef dimension="Store" level="Retail outlet"
          member="Downtown Quebec City" />
      </tuple>
      <cellValue>120.00</cellValue>
    </cell>
    <cell>
      <tuple>
        <memberRef dimension="Measures"
          level="MeasuresLevel" member="Sales amount" />
        <memberRef dimension="Product" level="Item"
          member="Fuzzy dice" />
        <memberRef dimension="Time" level="Day"
          member="2007-08-23" />
        <memberRef dimension="Store" level="Retail outlet"
          member="Sainte-Foy" />
      </tuple>
      <cellValue>210.00</cellValue>
    </cell>
  <!-- ... -->
</cubeRef>
</cubeCells>

```

Fig. 4. CubeCells document example

other words the most granular data. The client application can dynamically rebuild all the cube cells at less detailed levels, by applying the aggregation operators defined in *CubeMembers* for each measure. Cells corresponding to non-leaf members can also be included if needed (to include pre-computed aggregates for performance reasons or for cells corresponding to non-summarizable measures which cannot be obtained from an aggregative formula).

D. Relevance of the encoding for spatial data and SOA

The proposed encoding has two characteristics making it suitable for spatial data and SOA: first, it takes into account the specificity of spatial data in a generic and extensible manner. By considering that any complex XML data can be part of members and facts, this encoding can leverage external XML Schemas such as GML, SVG and Google KML, allowing for the representation of rich spatial data and extensibility for future versions of these standards. Second, its structure is particularly suited to coarse-grained, document-oriented Web services, which have a central role in Service Oriented Architectures (SOA) [16]. By having a loosely-coupled way to integrate SOLAP systems, promoting component reuse and interoperability, productivity gains in development can be achieved.

IV. EXAMPLE OF AN IMPLEMENTATION

The described XML format was first designed for a Web Service for the delivery of SOLAP mini-cubes to mobile clients [17]. A mini-cube is a data cube of reduced size, extracted from existing cubes, and adapted with respect to the constraints of mobile computing (i.e. reduced CPU speed and storage size of PDAs or cell phones, limited bandwidth and coverage of wireless networks).

The Web Service is invoked using the SOAP protocol [18]. The service contract, i.e. the operations accessible to clients, is described in WSDL (Web Services Description Language) [19]. The following operations are supported:

1. *GetCapabilities*: this operation informs the client about the service version and gives information on the identification (name and address of the organization, etc.).

2. *DescribeSchema*: this operation is used to get the metadata describing cubes, dimensions, hierarchies, levels and measures, returned in a *<cubeSchema>* element.

3. *GetMembers*: this operation allows querying of cubes' dimensions to retrieve their members. It must support the navigation into a hierarchy of members, i.e. get the parents or children of a specific member, get the members of a level, etc. This operation also offers mechanisms to select members based on their spatial attributes, with metric or topological operators. The selected members are returned in a *<cubeMember>* element.

4. *GetCells*: this operation is intended to query the facts of a data cube. The argument of this operation forms a multidimensional "select query" that defines the set of cells to be obtained and returned as part of a *<cubeCells>* element.

5. *GetCube*: this operation is used to retrieve a complete cube. It is based on the *GetCells* operation to get the data, and adds all the members that correspond to the mini-cube selection and the schema description of this mini-cube (metadata, i.e. dimensions, hierarchies, levels and measures structure). The XML document that is returned is a complete representation of the requested mini-cube (data and metadata), i.e. regrouping *<cubeSchema>*, *<cubeMembers>* and *<cubeCells>* elements.

A prototype implementation of this Web Service has been developed (Fig. 4). It is based on the open source Mondrian OLAP server and a PostgreSQL DBMS (with PostGIS spatial extensions) for the data warehouse backend. It uses Mondrian to query data cubes and extract a subset of these (i.e. a sub-cube) corresponding to a user-defined selection, done with operations from the service contract. This sub-cube is then serialized to the XML en-

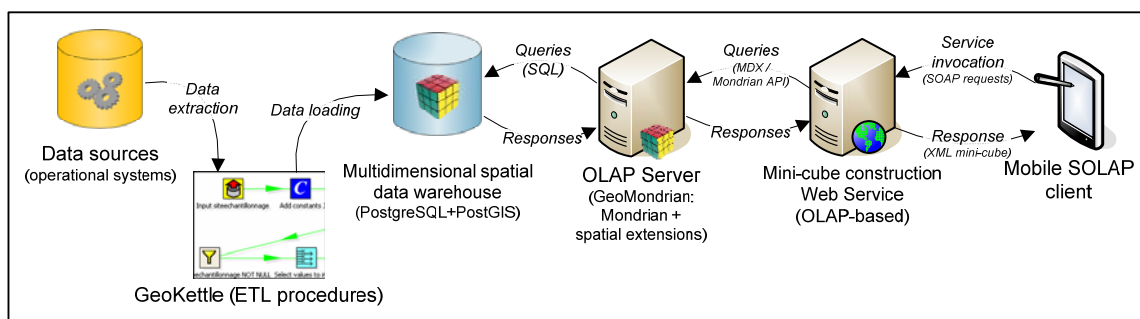


Fig. 3. Architecture of the Web service for SOLAP mini-cube delivery [17]

coding, in order to deliver a mini-cube for use with a mobile client. The mobile SOLAP client then de-serializes the mini-cube to an embedded database prior to use, since XML is useful as an exchange format, but not suited for direct query (due to parsing overhead and lack of indexes).

V. CONCLUSION AND OUTLOOKS

The requirements and the specification of an XML encoding for SOLAP data cubes have been presented. It allows the exchange of geospatial decision-support data in an interoperable manner. It enables the representation of cube metadata and data, and supports spatial members and measures. Its use was also illustrated in a Web Service for delivering SOLAP cubes to mobile clients. Other uses include any situation where cubes have to be exchanged between different computing environments, as often required in Enterprise Application Integration (EAI) systems.

It should be noted that XML data size and verbosity could pose problems with regards to the *CubeCells* documents. Since a data cube can contain millions of facts, encoding all the individual cells with references to members would yield a very large document. This can be mitigated with data compression. However, this solution does not address parsing overhead. A better approach could consist of a more compact encoding: for example, a multidimensional cell set can be stored by first defining axes with ordered positions referring to members, and then sequentially writing all cell values. In addition, binary XML formats such as EXI [20] or Fast Infoset [21] could be used to further improve parsing speed and reduce data size, especially for mobile computing environments.

The proposed encoding addresses the requirements listed in section 2, except for F. Further work is required to enable the representation of calculated members: a new language, with a corresponding XML encoding, will have to be designed to describe calculated members in an interoperable manner. In addition, while the XML format provides for technical interoperability between SOLAP systems, it does not solve the problem of semantic interoperability, i.e. the integration of heterogeneous data cubes with conflicts in meaning, interpretation and intended uses. Ongoing research work aims to address these issues [22]. Another research novelty which could eventually extend the XML cube model is the concept of raster-based spatial data cubes, based on a continuous grid division of space instead of discrete (vector) geometric features [13].

ACKNOWLEDGMENTS

This work has been made possible with grants from the NSERC Industrial Research Chair in Geospatial Databases for Decision Support (<http://mdspatialdb.chair.scg.ulaval.ca>). We would also like to thank the Canadian Institute of Geomatics Scholarship Program.

REFERENCES

- [1] World Wide Web Consortium (W3C), *Extensible Markup Language (XML) 1.0 (Fourth Edition)*, 2006.
- [2] E. Cerami, *Web Services Essentials*, O'Reilly, 2002.
- [3] ISO TC/211, *Geographic information – Geography Markup Language (GML) (Committee Draft, version 3.1.0)*, Open Geospatial Consortium, 2004.
- [4] OGC (Open Geospatial Consortium), *Web Feature Service Implementation Specification, version 1.1.0*, 2005.
- [5] S. Rivest, Y. Bédard, M.-J. Proulx, M. Nadeau, F. Hubert and J. Pastor. 2005. SOLAP: Merging Business Intelligence with Geospatial Technology for Interactive Spatio-Temporal Exploration and Analysis of Data. *Journal of ISPRS “Advances in spatio-temporal analysis and representation”*, vol. 60, no. 1, pp. 17-33.
- [6] E. F. Codd, S. B. Codd and C. T. Sally. 1993. “Providing OLAP to User-Analysts: An IT Mandate.” Codd Ass.
- [7] C. Franklin, “An introduction to geographic information systems: linking maps to databases”, *Database*, vol. 15, no. 2, pp. 13-21, 1992.
- [8] T. Badard, Y. Bédard, F. Hubert, E. Bernier, E. Dubé, “Web Services Oriented Architectures for Mobile SOLAP Applications”, *International Journal of Web Engineering and Technology (IJWET)*, 2008 (in press).
- [9] Microsoft Corporation, Hyperion Solutions Corporation, *XML for Analysis Specification (version 1.1)*, 2002.
- [10] Microsoft Corporation, *Multidimensional Expressions (MDX) Reference*, 2007.
- [11] W. Hümmer, A. Bauer and G. Harde, “XCube – XML For Data Warehouses”, In: *DOLAP’03*, ACM, pp. 33-40, 2003.
- [12] J. Silva, V. C. Times, R. N. Fidalgo and R. S.M. Barros, “Providing Geographic-Multidimensional Decision Support over the Web”, In *Proc.: Web Technologies Research and Development – APWeb 2005, 7th Asia-Pacific Web Conference (Shanghai, China)*, pp. 477-488, 2005.
- [13] Y. Bédard, T. Merrett, J. Han, *Fundamentals of Spatial Data Warehousing for Geographic Knowledge Discovery*, In: *Geographic Data Mining and Knowledge Discovery*, 2nd edition, Chapter 3, Taylor & Francis, 2008 (in press).
- [14] R. Kimball, M. Ross, *The Data Warehouse Toolkit*, Second Edition, New York, John Wiley & Sons, 2002.
- [15] World Wide Web Consortium (W3C), *XML Schema*, 2004.
- [16] E. Ort. 2005. “Service-Oriented Architecture and Web Services: Concepts, Technologies, and Tools.” Sun Microsystems (online).
- [17] E. Dubé, T. Badard, Y. Bédard, “Web Service for the Real-Time Delivery of SOLAP Mini-Cubes to Mobile Clients – A Service Oriented Architecture for Mobile Use of Geospatial Decisional Data”. In *Proc.: AGEO/GQFD-Géo. Clermont-Ferrand, France*. 18 p., 2007 (in French)
- [18] World Wide Web Consortium (W3C), *Simple Object Access Protocol (SOAP) 1.1*, 2000.
- [19] World Wide Web Consortium (W3C), *Web Services Description Language (WSDL) 1.1*, 2001.
- [20] World Wide Web Consortium (W3C), *Efficient XML Interchange (EXI) Format 1.0 (W3C Working Draft)*, 2007.
- [21] ITU-T, *Information technology – Generic applications of ASN.1: Fast infoset (ITU-T Rec. X.891)*, 2005.
- [22] T. Sboui, Y. Bédard, J. Brodeur and T. Badard, “A Conceptual Framework to Support Semantic Interoperability of Geospatial Datacubes”, *ER Workshops 2007, LNCS 4802*, pp. 378-387, 2007.