Image Exploitation for the Enterprise

BRAD SKELTON, MLADEN STOJIC, RYAN STRYNATKA, Atlanta

ABSTRACT

Growing standards, increased bandwidth, service oriented architecture and affordable multiprocessing are changing the way that we work. Isolated desktop systems will be augmented with federated clusters of systems which can be located in the same building or distributed across the world. Collaborative capabilities will enable users to consume and share geospatial data and services with the producers of GIS data improving the decision making process.

1. INTRODUCTION

The typical photogrammetric production workflow (Fig. 1) has changed little since the development of digital photogrammetric workstations. Imagery is ingested as image files, projects are setup to manage the collections of data, interior orientation is established, exterior orientation is established through point matching and bundle adjustment, stereo models are created and used for interactive extraction or automated extraction. The final products (digital terrain models, orthos, mosaics, vectors, 3D models etc) are then delivered to the customer. A great deal of work has gone into accelerating the various production steps because of their highly compute intensive nature (faster CPUs, multiple processors, distributed computing) but each of these has been a specialized or proprietary system. The overall workflow at the project level and at the business level has been disconnected from the photogrammetric workflow. Photogrammetric systems remain contained within enterprise departments as standalone silos. Enterprise and information technologies have now matured to such a level that they can now be brought to bear on the photogrammetric workflow, or perhaps it is better to say that the photogrammetric workflow can now be part of the larger enterprise business system.

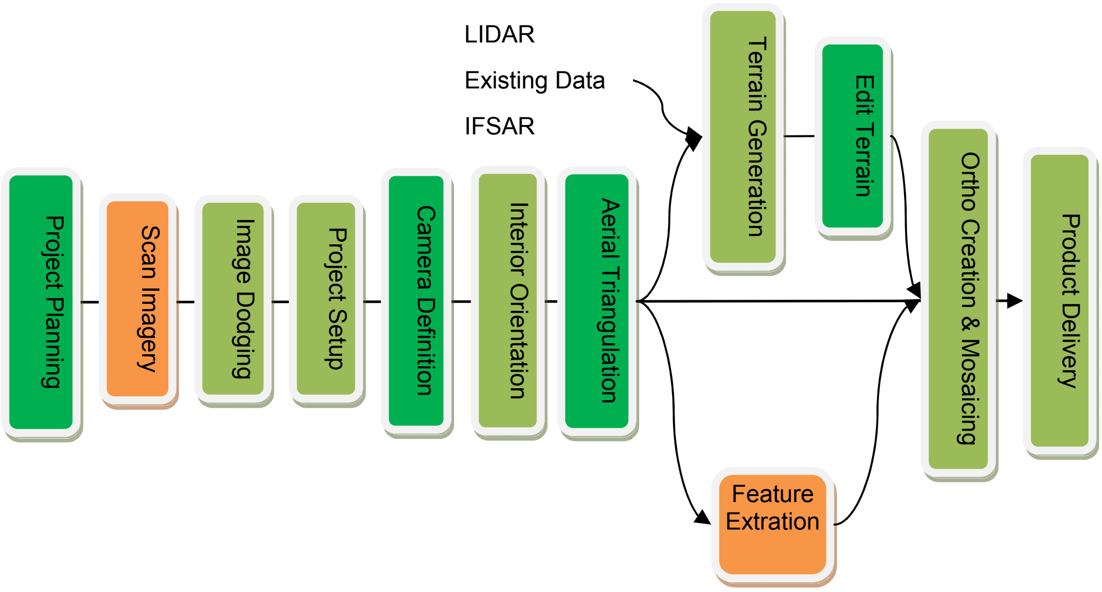

Fig. 1: Traditional Photogrammetric Workflow

It is useful to begin with a definition of the term enterprise. A system which is comprised of the following can be thought of as an enterprise system:

- True multi-user, simultaneous access to the same production project from any workstation in the production network
- Transaction processing against a central database
- Real-time automatic status updates of client workstations in the network as the project progresses
- Project access and security on par with the system domain
- Rational schemes for managing high volume, highly transient data types
- Capabilities provided as interoperable services
- Scalable to meet the growing production and throughput demands of an organization
- Ability to persist all variables and parameters associated with the photogrammetric workflow
- Extensible platform for customizing the photogrammetric workflows and integrating them with other business workflows

These are common characteristics of traditional business systems like banking, but they are now key components of Geographic Information Systems (GIS). The success of these traditional systems has been based on the development of standards which allow the various components to interoperate, as it is rare that a single vendor can provide a single enterprise system and certainly no single vendor dominates the industry to provide the interoperability which is needed for large distributed systems to interact. Standards like the hypertext transfer protocol (HTTP) allow for the transmission of information across disparate systems, while standards like the extensible markup language (XML) allow this information to be understood by these same disparate systems. These standards are now a part of the GIS world with specific applications of XML for feature data known as the Geographic Markup Language (GML). These same standards bodies are working to produce standards in the photogrammetric world such as SensorML and the Transducer Markup Language as a means of persisting sensor information and the Community Sensor Model (CSM) as a definition of a reusable sensor model API.

2. STANDARDS

The majority of standards driving the geospatial industry are driven by one of two organizations. The International Organization for Standardization (ISO) located in Geneva, Switzerland and the Open Geospatial Consortium (OGC) located in Wayland, Maryland in the US. While the set of ISO standards work is very broad (ranging from agriculture, to mechanical engineering, to information technology) the OGC focuses specifically on standards that affect Geospatial Information Technology.

A short list of OGC standards includes:

- Coordinate Transformation Service - provides interfaces for general positioning, coordinate systems, and coordinate transformations
- Web Map Service (WMS) - provides three operations in support of the creation and display of registered and superimposed map-like views of information that come simultaneously from multiple remote and heterogeneous sources

- Geography Markup Language (GML) – provides an XML encoding for the modeling, transport and storage of geographic information including the spatial and non-spatial properties of geographic features
- Web Feature Service (WFS) – provides the means for a client to retrieve and update geospatial data encoded in Geography Markup Language (GML) from multiple Web Feature Services
- Catalog Service (CS) – defines common interfaces to discover, browse, and query metadata about data, services, and other potential resources
- Transducer Markup Langue (TML) – provides an efficient method for transporting sensor data and preparing it for fusion through spatial and temporal associations

A short list of relevant ISO standards:

- ISO 19130 – defines sensor and data models for imagery and gridded data
- ISO 19115 – defines the schema required for describing geographic information and services.

In addition to the efforts ongoing at ISO and OGC to define various aspects of sensor models and their representations there has been work going on for the United States National Geospatial Administration (NGA) to create a common model for the implementation of sensor models called the Community Sensor Model (CSM). A key part of any photogrammetric system is the definition of the mathematics which govern the mapping of the 3D world coordinates onto the 2D image plane called the sensor model. While common sensors such as metric aerial cameras have well known models, more elaborate airborne and space borne sensor systems have sensor models which are either very complex, confidential or both. It is often very difficult (if not impossible) to get sufficient detail about a given sensor to write an effective sensor model for it. As NGA deals with a large number of vendors it has found that it is very problematic for each of its different system providers to have to independently implement sensors models (often with varying degrees of accuracy). So NGA began a program to define a common API for sensor modeling which would enable a single implementation of a sensor model to be used with different vendors' software, assuring common capability and accuracy across a broad range of applications. By writing to a common API it is possible for a sensor model to be written by a single organization and then supplied in the form of a dynamically loaded library (DLL) to different organizations with the expectation that it will work with various software applications. Software from different vendors such as BAE, Sensor Systems and Leica Geosystems now use this standard and can be expected to work with the new models. This API, called CSM, is now gaining adoption among vendors of software to the US Defense, but more importantly it is also being looked at by international organizations and organizations such as the OGC.

With a standard sensor model API the end user wins because he can begin using new sources of data more rapidly. Instead of waiting for all of the vendors of his different software packages to incorporate a new model he only has to wait until one of them supports the model (ideally the sensor maker would actually provide the model as a working DLL). Once the model is on the system all of the different software packages can use it. This is not unlike having drivers for hardware which free all of the application programs from worrying about the specific details of each piece of hardware.

While CSM defines an interface for dealing with sensor models it does not address persistence of the metadata which is used by the sensor model. This is being addressed by the work of the ISO 19130 working group. As a result it is necessary to design a system (Fig. 2) which integrates these two mechanisms. The domain layer will use CSM as the computational engine and the data layer will look to ISO 19130 to provide system data interoperability.

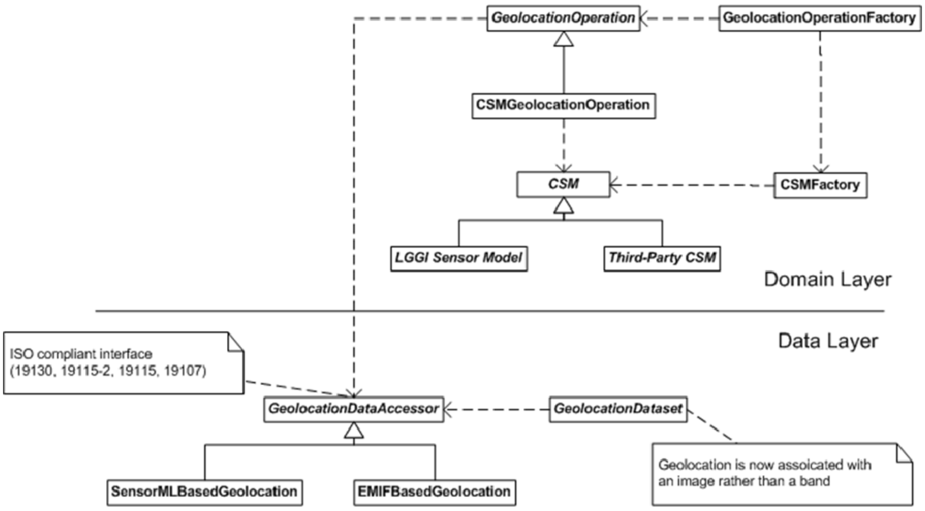

Fig. 2: Integrating CSM and ISO19130

3. HIGH PERFORMANCE COMPUTING

The digital photogrammetric workflow has always put high demands on the processing power of the workstation and the current set of high resolution digital aerial cameras are pushing the demands even higher. Leica Geosystem's Aerial Digital Sensor (ADS40) is capable of collecting 100GB of raw data for every hour of flight. On a typical single CPU workstation this type of data can require about 40 hours of ground processing for each hour of flight. The timely delivery of finished products is key to the business success of these systems so many means of bringing greater computer power to bear have been explored and exploited. There are many different technologies available in this domain:

- MPI: Message Passing Interface
- PVM: Parallel Virtual Machine
- Condor: High Throughput Job Scheduler
- DCOM: Distributed Component Object Model
- OpenMP: Fine Grained Shared Memory Parallelization

These can be categorized as coarse or fine grained parallelization technologies based on the overhead of the interprocess communication. Photogrammetric processing lends itself well to parallelization. Since the typical scenario involves processing hundreds (if not thousands) of images in the same way, it is clear that there can be a great deal of opportunity for distributing the work. For example, the production of orthopotos from a large block of photography requires the repetitive application of an orthocorrection process to each image. The result of one does not depend on the result of any of the others and in fact the results within an image are likewise independent. There can be a number of strategies for attacking the problem and subdividing it into smaller units which can be distributed. On a multiple CPU machine (or even a vector machine) it makes sense to break an image into smaller units which are handed to each processer, depending upon the architecture this could be subdivided at the pixel level (OpenMP). More typically this would be done by break the image into sub tiles and having each processor handle a tile. This is referred to as fine grained parallelization. Because the data is divided into separate files due to the means of collection and storage it is convenient to split the processing at the level of the file. Then there are no issues associated with subdividing the images and recombining the results. This is referred to as coarse grained parallelization.

The first two methods (MPI and PVM) are models which provide a means of transparent message passing between cooperating processes which may be on the same computer or distributed across the network. MPI is a standardized interface with many implementations across a broad range of machines while PVM is an integrated set of software tools with a large user base, though it seems to have been superseded by MPI. While each of these could be used to attempt a fine grained approach the overhead could eliminate the gains of the distributed processing.

Condor is a sophisticated distributed job scheduler developed by the Condor research project at the University of Wisconsin-Madison Department of Computer Science. It provides facilities for managing collections of computing nodes which can participate in the distributed computing. This ability to manage and schedule the computing nodes makes it unique and well suited to photogrammetric processing applications. It provides parallelization at the coarse grained level.

Working with partners, Leica has explored these options and ultimately selected Condor to increase the throughput of its ground processing software (GPRO) for the ADS40 (Fig. 3). The GPRO Rectifier already used multithreading as a means of fine grained parallelization to improve performance on a single image. Integration required minimal changes to existing software and provided dramatic improvement in processing times for collections of images.

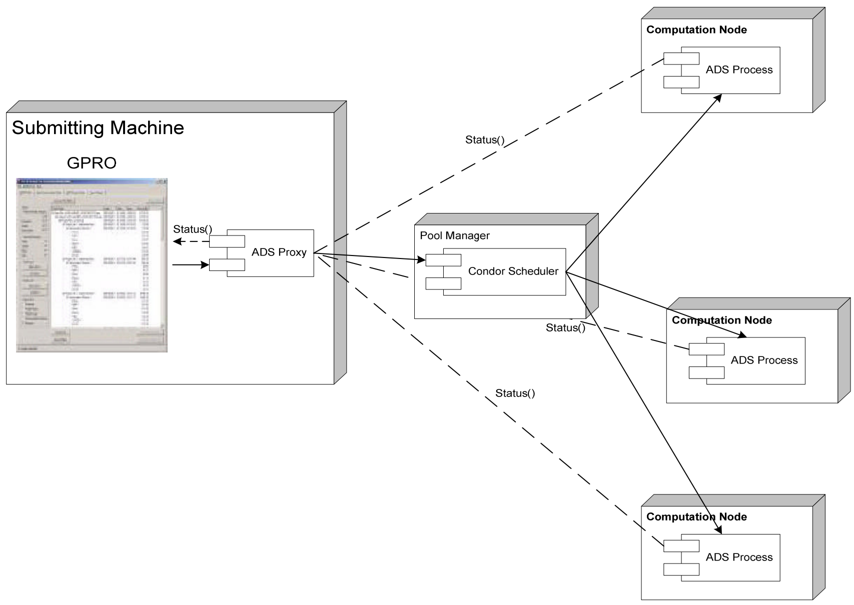

Fig. 3: GPRO Distribution Model

In a similar fashion Leica has worked with GeoCue to create the Leica Ortho Accelerator which is a generalized tool for production of orthophotos (see *Workflow Management* below). In this example the project contained 858 8-bit TIFF images of about 962MB each. The DTM was physically loaded onto each workstation for the orthorectification process. Each 10M cell size DTM is 27MB, which means that in addition to the source images, 27 MB of DTM is also copied from the server to each node's caching folder.

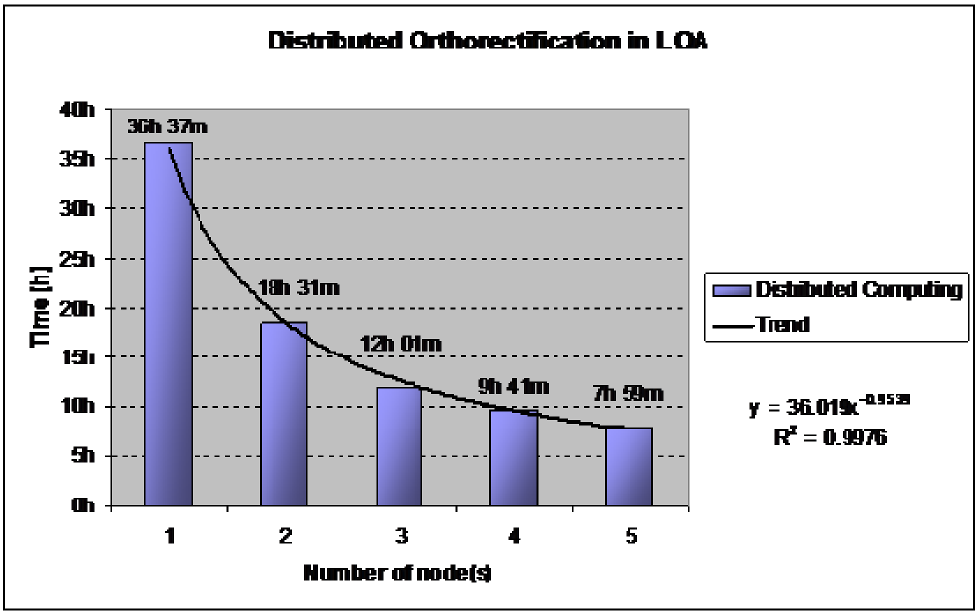

Fig. 4: Orthorectification time in hours and minutes for 858 images

4. PHOTOGRAMMETRIC DATA MODEL

Any complex process such as the photogrammetric production workflow must be modeled and persisted or saved. The photogrammetric workflow has a very specific set of data entities that can be described within a photogrammetric data model. These data entities may include imagery (raw, oriented, and orthorectified), ground control points, automatically or manually measured tie points, terrain data, vector data, camera information and so forth. Photogrammetric projects also use metadata to describe these data assets. Important metadata include image identifiers, interior orientation parameters, exterior orientation parameters, coordinate system, descriptive data, and

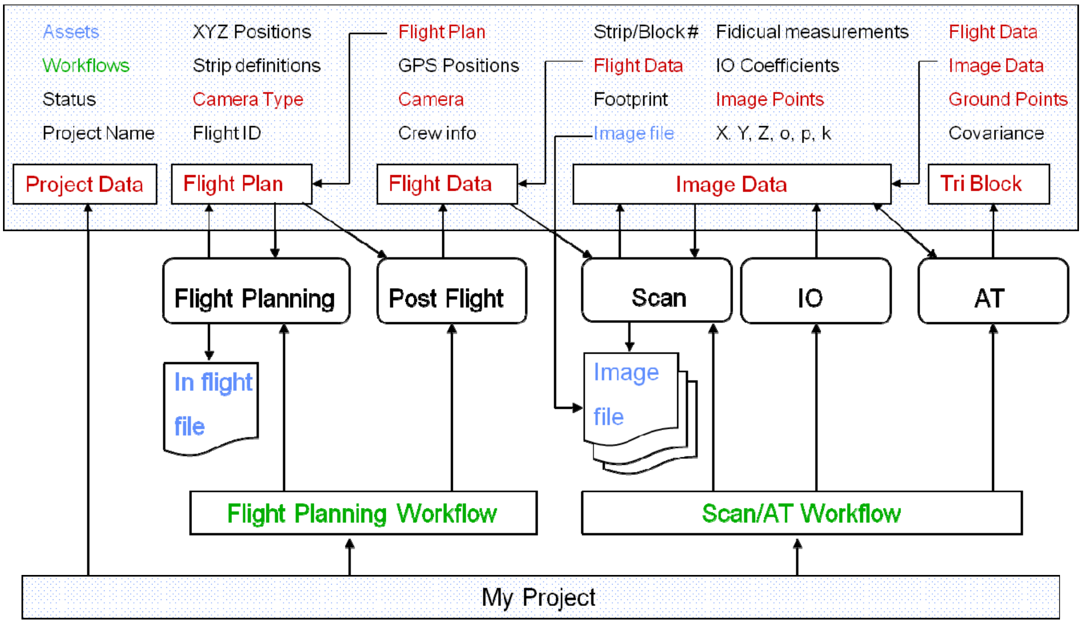

Fig. 5: Photogrammetric Data Model

more. Together the photogrammetric data entities and the metadata describing them comprise the photogrammetric data model. In most systems this model will take the form of some type of project. Fig. 5 shows such a model.

We can think of the various entities managed in the system as assets which are described by their various metadata. In this case the model supports an asset from its inception through its processing and should continue to support the asset through its final dissemination. The more closely the data model agrees with the actual process the more automation can be brought to bear. In this example, the photogrammetric model begins with flight plans which are created at the time the project is planned. By including the flight plan in the model and carrying it through the system it is possible to increase the automation at all points. In this case the flight plan is loaded onto the control system and the camera can be controlled to define exposures. The exposure metadata are also carried in the model as flight data which can then be used to automatically establish the project on the desktop workstation. This type of coupling can be found in the Leica Geosystems flight planning (FPES), flight control (FCMS) and ground processing (GPRO) software.

5. ENTERPRISE MODEL

Enterprise geospatial processing has the potential to offer greater throughput than the desktop approach to the photogrammetric workflow. For example, current workflows are very project specific. While data is often (but not always) saved and archived, it is often kept offline and is difficult to manage. The enterprise approach embraces the notion of “collect once, use many times”. Available data should be used to continuously refine photogrammetric products. In a hypothetical enterprise system, a bundle adjustment should be performed and orientation data automatically updated when a control point is added on a client machine – possibly even from the field. Likewise automatic terrain extraction should automatically identify and use available input data to seed the correlation process. The system should ideally harvest seed data from a variety of sources: global terrain sources (e.g. GTOPO30), available raster and vector data (e.g. breaklines), previously processed terrain files, and online data stores.

An enterprise model for geospatial processing must certainly involve all of the elements of the photogrammetric workflow previously discussed but it must also support working with the data in a fashion which is distributed amongst different users acting in different roles. The traditional three tiered model is used, providing a clear separation between processing (business logic), data storage and presentation (interaction with the user).

Most of the traditional relational database providers either have or will soon have versions of their databases which are spatial enabled, allowing for the direct storage and manipulation of spatial data including features and raster. Oracle was the first RDBMS with such integration, however the open source Postgres system now has a spatial version called PostGIS and IBM and Microsoft have both announced spatial versions of their systems. Managing huge amounts of photogrammetric data as assets in a relational database allows management and scaling which is not available with simple file system based approaches. Including spatial awareness at the database level allows spatial queries to be as fast any other query made against the database.

As a result photogrammetric processes are made available as services through an application server. Perhaps the most widely used application server is the Apache Web Server plus Active Server Pages (ASP). However, this is not a highly scalable platform suitable for such applications. The J2EE (Java 2 Platform, Enterprise Edition) specification provides a scalable distributed multi-tiered application server environment. Sun, Oracle, IBM, BEA, JBOSS all have implementations of this platform. The J2EE platform serves as a shell in which various services run. The platform takes care of scaling to fit the load with the major difference in various implementations being on how well they scale.

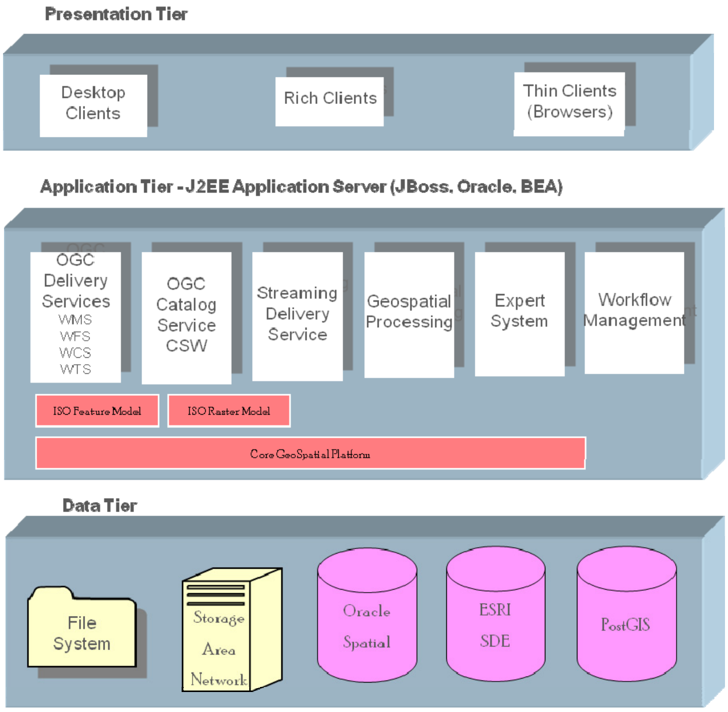

Fig. 6: Enterprise Architecture

Key services to be provided in the application tier are:

- OGC Catalog Service
- OGC Delivery Services
- Streaming Data Delivery Service
- Geospatial Processing
- Expert Management Services
- Workflow Management

The first two sets of services are currently widely available as either open source or commercial products and are at the core of a number of online geospatial services. GeoServer is an open source project which provides the OGC WMS, WFS, WFS-T services. A richer set of these services plus an OGC compliant cataloging service is available as a set of commercial tools called RedSpider Enterprise from IONIC Software in Belgium. This is the most mature implementation of the OGC and ISO standards in the geospatial marketplace.

Spatial asset catalogs have been around for a long time. In the early 90's the University of California at Santa Barbara started a librarian project which is still around today as the Alexandria Digital Library. Similarly an early pioneer in this area was CORE Software, who has the TerraSoar Suite for cataloging, searching and distributing Geospatial data. Each of these was originally built upon proprietary schemes. The Catalog Service from OGC defines an open standard for cataloging and querying metadata. It defines a common query language (the OGC Common Catalog Query Language) which allows disparate systems to share a common catalog. Perhaps, more significantly, is the provision for the catalog query protocol to be propagated to other compliant catalogs. This enables the creation of federations of catalogs which allows for the creation of networks of distributed catalogs around the globe.

WMS provides for the delivery of geospatial data in the form of a spatially referenced map. While the data is delivered as a web graphic (PNG or JPEG) this protocol is not well suited to the delivery of raster data in a streaming fashion. Instead, WMS is highly focused on delivering a well formed map (hence the M in WMS). The performance is fine as for casual usage, but it tends to put large demands on the server and is simply not responsive enough for highly interactive work. This warrants the development of dedicated high performance image serving protocols to provide imagery to applications at a rate that allows comfortable interaction. ER Mapper offers Image Web Server (IWS) based on the ECWP protocol. There is also a proposed standard for the transmission of imagery using the JPEG2000 standard called the JPEG 2000 Internet Protocol (JPIP). Currently the ECWP appears to give the best performance among commercially available products. However, it should be noted that Microsoft has demonstrated very impressive capabilities with the acquisition of the SeaDragon technology, providing completely smooth interaction which is only limited by the bandwidth of the connection and the size of the screen.

There are currently no widely adopted standards driving the area of Geospatial Processing, though there is work underway to develop such standards. For example, the OGC has a draft standard called the “OpenGIS Web Processing Service”. Similar to parallelization this can be looked at as either a fine grained or coarse grained problem. Some of the proposed standards strive to keep the specification as light as possible and thus put as few constraints as possible on the processing; this would be the coarse grained approach. At the other end there is an effort to create a standard for processing services which are supported by a rich vocabulary of operations and expressions providing, in effect, a spatial modeling language, this can be thought of as the fine grained approach. With such a fine grained approach it is possible to have authoring systems which are used by experts to create models to solve specific problems. These models are then “published” making them widely available for across the enterprise for consumers of geospatial information products.

For example, Leica’s ERDAS IMAGINE has a rich Spatial Modeling Language (SML) with an associated spatial modeling engine. This language has well over one hundred operators for basic image processing, thematic mapping, resampling, etc. Complex operations can be created using concise expressions in a compact, “C-like” notation. Providing such an engine as a service can make the analytic capability of such models available to more casual users through a web based interface.

A key to making such the capability available to casual user is the combination of the modeling service and catalog service with an expert system. Every model needs input, but what serves as usable input to a model? This can be answered by querying the catalog and then using the expert system to select the “best” inputs.

6. WORKFLOW MANAGEMENT

Workflow management in the geospatial market has largely been left to project managers. It is not uncommon for a large photogrammetric project to be managed by a series of paper maps hung on the wall. Footprints of the various components to be processed are drawn on the map with various pieces of status information being recorded (often hand written) on the map. Integration with organizational business systems consists of the manual transcription of this information to the appropriate accounting systems. While project management systems have been a part of accounting systems for years, the integration with geospatial processing software has been slow to develop.

A few years ago this need was recognized and addressed by NIIRS10 (now called GeoCue Corporation). The first implementation was focused on enterprise production management with an emphasis on LIDAR data trying to solve:

- The management of large projects operating with different teams of people
- The management of large data volumes
- The streamlined workflow of a common set of LIDAR processing tools.

The resulting product called GeoCue is an enterprise geospatial production workflow management tool which maintains a central, synoptic view of the project, launching tools directly from the data. In effect the operator is focused on performing the task at hand and is not worried about the underlying files or tools (or how to share them for that matter). The information regarding the state of the project is maintained in a relational database allowing for the creation of a true multiuser project management system.

Leica has worked with GeoCue to provide tools for enterprise orthophoto production operating within the GeoCue framework. The resulting Leica Ortho Accelerator product provides a highly productive and automated system for the production of orthophotos combined with a system for the production of DEM from LIDAR data.

7. COLLABORATION

While large centralized data shares are a key part of geospatial enterprise systems there is a place for a more decentralized form of collaboration and sharing. Internet messaging (IM) and peer to peer media sharing (Napster) are a common place in the current internet space. Leica has combined

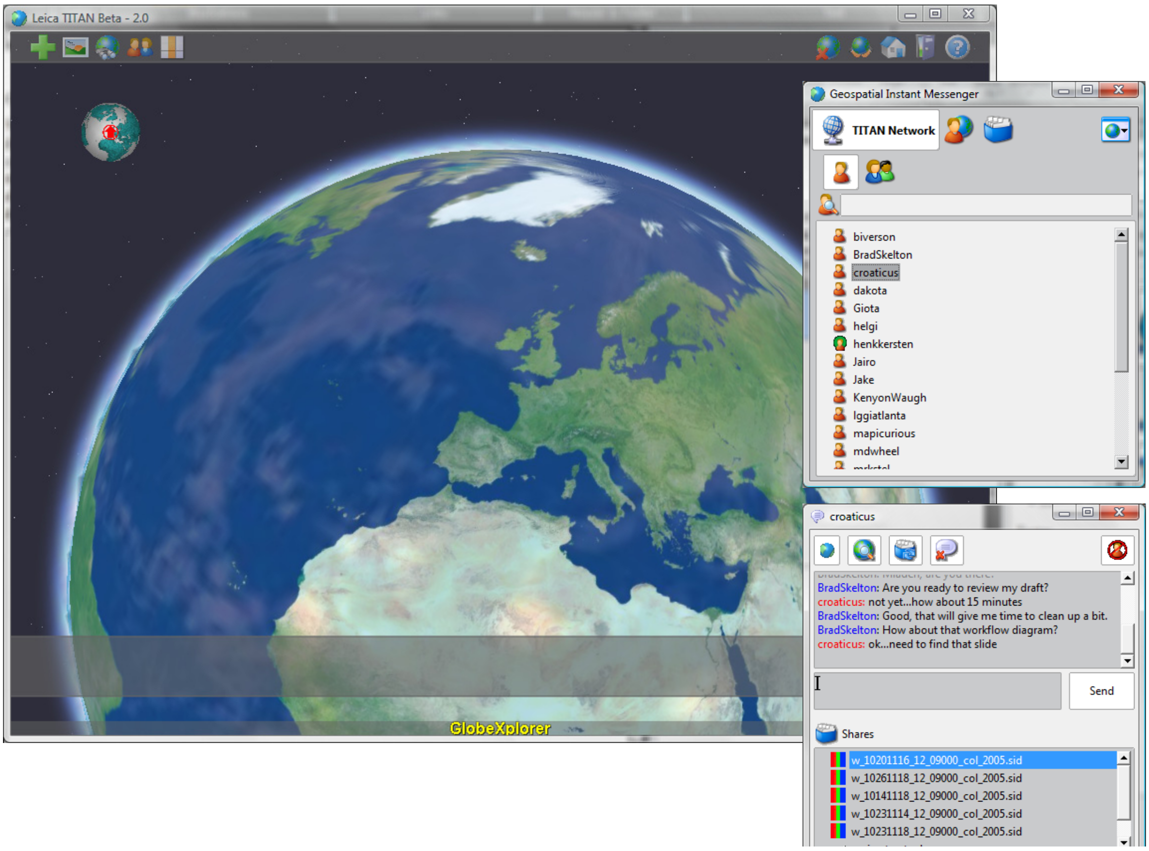

Fig. 7: Global Collaboration

these ideas with geospatial awareness to create a unique sharing environment called Leica TITAN (Fig. 7). This tool combines the now familiar spinning globe environment with ability to chat with any of the members of the service and to exchange geospatial data by dragging and dropping to the globe. This could be integrated with the previous workflow scenario to provide a means for remotely located people to work on the same large geospatial project. For example, operators in different locations working on adjacent areas could discuss how a particular feature might be captured and attributed via the collaborative capabilities, while sharing local data with each other.

8. CONCLUSION

The rapid pace of development in enterprise systems and the standardization of geospatial technologies is extending the application and utilization of desktop workstation applications to other parts of an organization. Geospatial processing will become a mainstream tool in the general enterprise toolbox which can be integrated at many levels throughout an organization. This integration will allow wider collaboration, reduce production timelines, and make the sharing of geospatial data with those that need easier and intuitive access. Ultimately this serves to push previously niche processing operations into the hands of a more mainstream audience.

9. REFERENCES

Books and Journals:

Beshah, B.T., Welter J. and Moran K. (2004): Distributed Computing in Ground Processing, *Proceedings of the 20th ISPRS Congress, 12-23 July, Istanbul Turkey*.

Graham, L. (2007): Enterprise Geospatial Production, *PE&RS May 2007.*

Y. Xue et al. (2005): High Throughput Computing for Spatial Information Processing (HIT-SIP) System on Grid Platform, *EGC 2005, LNCS 3470,* pp. 40-49.

www:

Open Geospatial Consortium: http://www.opengeospatial.org/

The International Organization for Standardization: http://www.iso.org/iso/en/ISOOnline.frontpage