GEODATABASE CENTRIC ORTHOIMAGE PRODUCTION USING ARCGIS IMAGE SERVER

P. Becker[a]

[a]ESRI, 380 New York St., Redlands, CA, 92373, USA - pbecker@esri.com

Commission IV, WG IV/2

KEY WORDS: Imagery, Management, Production, Orthoimage, Parameters, Metadata, Distributed, Online

ABSTRACT:

In traditional image production workflows, imagery is processed in multiple steps, with the imagery being sampled multiple times, and being read and written to disk multiple times. This reduces image quality and production efficiency. The linear workflows also are susceptible to bottlenecks when a single parameter in the workflow is not available or needs to be changed. ArcGIS Image server enables a geodatabase centric workflow by which the parameters and models for processing imagery are stored within a database and the imagery product is generated on demand as required directly from the base imagery. This methodology provide a number of advantages. For example, it enables serving of dynamic image services that can be updated as revised parameters become available. It enables the creation of graded products that change over time with the revised parameters and enabling improvements in quality assurance processes. The same image services can also be used to generate imagery products in the form of caches or the traditional tiled images that are required by most orthoimage mapping projects. The quality of the imagery is superior due to reduced sampling of the images. By optimizing the processing and reducing disk access, the production system is very efficient and scalable, enabling high production rates.

1. TRADITIONAL ORTHOIMAGE PRODUCTION

Traditional orthoimage production workflows, result in imagery being processed in multiple steps. Typically, an image will go through the following steps: radiometric correction, pan-sharpening, aerial triangulation, orthorectification, mosaicking, reprojection, then product generation. Depending on the software used, the order of these steps may change, but typically each step is performed separately with imagery being read, sampled, or enhanced and then written to another location.

With each radiometric or geometric enhancement some information is lost, resulting in the final image quality being non-optimum. Even processes, such as mosaicking multiple images together, result in unnecessary sampling if pixels of all the input and output images are not aligned. Additionally, if the sampling of the input and output imagery is a very similar resolution (which is often the case when re-projecting) then aliasing artefacts can also become apparent especially in imagery covering areas that are near featureless, but have good textures such as over water or gravel desert. In some areas the imagery appears to have slightly higher contrast than in other areas. These artefacts are caused by pixels in some regions being sampled to be close to the average of four neighbouring pixels (which reduces the local contrast), while in other areas the pixel is the nearly solely derived from a single input pixel (maintaining the local contrast). Although different sampling methods may be defined to reduce these effects, they are generally at the cost of accuracy. To reduce the data volumes some workflows apply lossy compression methods to the intermediate products, which further increases the creation of artifacts and degrades quality.
Such orthoimage generation workflows where the imagery is read and written multiple times are also not truly scalable. Traditionally, such workflows are scaled using technologies such as CORBA, enabling distributed processing over multiple machines. The processing gains are quickly mitigated by network and harddisk bottlenecks, which are caused by the multiple reads and writes of large image data volumes to disks saturating the available bandwidth as well as fragmenting the disks.

1.1 Why do we need a geodatabase centric approach?

Linear process workflows can also cause substantial bottlenecks in production. A delay in one step will stop subsequent steps. For example, if the accurate orthoimage product is not created due to the non-availability of a terrain model, the color balancing cannot be performed, since color balancing in such workflows are dependant on the orthoimages. As projects become larger, the chances of one part of the process being delayed increases which in turn increase the delay and risks for the complete project.

Such production workflows can be considered a set of independent tasks and not all tasks need to be performed necessarily in a fixed order. Each of the production steps can be considered to consist of two components: The determination of process parameters, and/or the application of these parameters on the imagery. The process of aerial triangulation is a typical example of parameter determination with no pixel processing being applied. Orthorectification is a process that utilizes the parameters of orientation and a terrain model to apply a pixel process. During the complete image production workflow, there are many parameters that affect the resulting product. These include parameters of orientation, radiometric enhancement, pan-sharpening, terrain models, and mosaic seamlines. These production stages are actually quite loosely coupled. For example, the parameters of pan-sharpening have no effect on aerial triangulation. The determination of color balancing parameters for imagery is often dependant on the orthoimages, but does not require accurate orthoimages nor does it require the orthoimagery at high resolution. Even parameters such as mosaic seamlines are not dependant on the requirement for orthoimage with final accuracy, Most of the dataflow and production bottlenecks can be alleviated by using geodatabase and server centric workflows.

2. INTRODUCTION TO ARCGIS IMAGE SERVER

ArcGIS Image Server is a product from ESRI that provides management, processing, and dissemination of very large volumes of imagery with on-the-fly server based processing. This product can serve traditional pre-processed image tiles, and serve imagery that has not be pre-processed. The required processing is applied in real time as the imagery is being served. The processing to be applied is defined in a database. This enables the same imagery to be served in different forms, as required by different users, without additional data storage. Also, the imagery can remain in its native or original form removing the requirement to convert the data. Additionally, the image products can be served as virtual mosaics. When an image is requested, the server not only extracts the required pixels, but applies the required processes, returning the pixels as well as the metadata to the user's application.

As with the other ArcGIS Server family of products, the workflow is defined in three stages: Author, Serve, and Use.

2.1 Author

In authoring an image service definition database is created, which defines a catalog of all the imagery. This image service definition contains information on each of the individual images including all properties, metadata, footprint, and location. The large volume of pixel data are not loaded into the database, but remain in their original form and are referenced. Also, associated with each image are the properties of how the images should be processed. These include the processes to be applied and the parameters of the processes. Such processing can include how to radiometrically enhance the images as well as how to process the image geometrically. A geometric process for georeferencing can be a simple affine transform or more complex transforms such as orthorectification that is dependant on camera calibration information, orientation data, and a reference to an elevation model. Associated with each image is a footprint shape that defines the extent of the pixels to be used, enabling the exclusion of pixels that represent the data frame of a scanned image or exclusion of areas such as clouds. The same set of images can be included in different image service definitions enabling different imagery products to be defined with no additional imagery being stored. Such products could include, different band combinations or enhancements that utilize the multispectral capabilities of many digital sensors.

Central to ArcGIS Image Server is the image service definition that defines within a database structure the processing to be performed on the imagery as well as all the parameters that effect the processing. As processes and parameters become available in the production flow they can be added or modified within the image service. The service definition consists of both processes defined using XML as well as the parameters of these processes defined either within the XML or as attributes of the service tables. This provides great flexibility and enables the database table to define only those parameters that may change. For example, parameters such as interior orientation that do not change can be defined as constants in the XML while parameters such as the six external orientation parameters can be defined in the tables. These orientation parameters may be defined initially from the GPS and flight navigation data and be updated when the GPS/IMU parameters are available, then updated after an aerial triangulation is performed.

The open and documented data structures used enable simple integration of these parameters from a wide range of existing applications that perform such parameter determination. This enables ArcGIS Image Server to be quickly integrated into existing workflows.

2.2 Serve

The image service definition can be published on the server and then is simultaneously accessible to many different client applications. These applications can either connect directly to ArcGIS Image Server using RPC (Remote Procedure Calls) or through standards such as OGC WMS, WCS as well as SOAP and REST by serving through ArcGIS Server. These enable a large range of client applications to access the imagery as services, either over local area networks or over the web.

2.3 Use

Clients connecting to image services can view the imagery as a large, single image; although they often consist of thousands of individual images. Importantly, metadata about the image service as well as the individual images being viewed can be accessed, providing the client with information about the imagery that is critical to many decision making tasks. The client applications can also interact with the image services to dynamically refine the processing. For example, a client may change the order of overlapping imagery to obtain imagery that more closely matches their requirement, such a being closer to a specified date or having a required sun elevation. When working over low bandwidth connections, a client can change the compression used for the data transmission. The compression used for the data transmission is therefore independent of the compression used to store the data, enabling an analysis, for example, to use highly compressed imagery for faster navigation, while still being able to reduce or remove the compression for detailed analysis. Therefore, the required image product can be created on demand, based on parameters defined in the database or specified by the client application.

3. PROCESSES

For all requests, the multiple geometric processes such as orthorectification, reprojection, pan-sharpening, and mosaicking are concatenated internally by ArcGIS Image Server to a single sampling of the original imagery. As a result, the quality of the resulting pixel data is improved. Prior to sampling the imagery the geometry is analysed such the only the pixels required for processing are read, therefore, the volume of data to be read and processed is substantially reduced. Typically, with aerial imagery being flown with high overlaps only a small percentage of the original pixels may be needed in the final product and require processing. Performing pan-sharpening on-the-fly without the need to store the large pan-sharpened file also substantially reduces data volumes.

4. GRADED PRODUCTS

The image service can be dynamically updated. Imagery can be added, and various processes and parameters can be changed without stopping the published image services. This is important for the use of the system in time-critical applications, and also enables the creation of a graded product. In this method, a preliminary image service is created based on the best available parameters and when the parameters and models are refined, they are updated within the image service definition. This provides a preliminary product as soon as the imagery is available on a server. For example, an image service can be created based on imagery from a digital camera as soon as the data is available. The initial orientation for orthorectification may come from preliminary orientation parameters available from the GPS/IMU. The initial radiometric processing could be default values. Such image services can be created as soon as an aircraft has landed and made available to users—when the value of the imagery is highest for applications, such as emergency response. As different steps in the orthoimage production flow are performed, the processes and the parameters in the image service definition are updated. For example, after performing an aerial triangulation the orientation parameters may be updated increasing accuracy. After updating the DTM the accuracy further increases. Similarly, the radiometric processing or pan-sharpening parameters can be updated as improved values become available. The initial product may not use seamlines and after the system has generated the seamlines this can be used to create a more seamless product. Thereby the graded product improves over time.

5. QUALITY CONTROL

The methodology of ArcGIS Image Server not only enables use of the imagery sooner which increases the value of the imagery, but it is also used for efficient quality assurance processes and is an integral part of Quality Control (QC).

In most production workflows, QC becomes a major component of the total workflow. Some of the QC can be performed in an automated manner, for example, in analyzing the histograms of the output images, while other steps must be viewed by operators to check the accuracy of control points or the junctions of seamlines. If these QC processes require the final product to be produced, then the processing must be repeated if any errors are found. Since the processing with ArcGIS Image Sever is performed on the fly, the QC steps can be performed prior to the extensive processing. Corrections can be updated in the database the corrections quickly reviewed. The integration of ArcGIS Image Server into multiple GIS client applications further enables the implementation of geospatially enabled QC processes. With the ability to serve both the imagery and QC tasks on the web, QC can be checked by the end user of the product prior to production and delivery.

6. CACHING

Such dynamic image services with on-the-fly processing have significant advantages over image servers that only serve pre-generated and mosaicked static imagery. Dynamic imagery servers that are driven from such databases and generate products directly from the base imagery can be expected to quickly augment the traditional static image servers that require the pre-generation of large images or lots of individual tiles. ArcGIS Image Server can easily be scaled to serve hundreds of simultaneous users. To scale to very large numbers of simultaneous users map caches are used. Web applications that provide interactive display capabilities make requests to servers using predefined tiles that are assumed to be static. Such static tiles can be stored on the web servers as small static image files. Such static tiles can be easily distributed on the Web and are utilized by edge servers to reduce server requests, data transfers and speed so up applications. Many web applications also cache tiles locally to further reduce requests to the server and provide faster interaction in areas that the user has already visited. Such static tiles can be generated in advance and stored on the Web server, but for large project would require extensive processing that can take a long time. With ArcGIS Server 9.3, such tiles need not be pre-generated. Instead they can be generated on demand when a user first accesses an area at a specified resolution. As the majority of high resolution requests to servers repeatedly cover only a small percentage of the full area, the high performance is maintained without the need to pre-generate high resolution caches for the complete area. Such caching can be driven directly from ArcGIS Image Server, where the cache is generated directly from the original image data. Caches therefore have significant advantages over solutions that utilize large static files that can not utilize web caching infrastructures and need to be pre-generated.

7. ORTHOIMAGE PRODUCTION

Although dynamic serving of imagery directly from the source and using on-demand caches have many advantages, there is currently a demand to create orthoimages as tiles that conform to a specific grid or map sheet layout. Many large orthoimage mapping projects currently specify the creation of image products as tiles in a range of formats including TIFF, NITF, and JPEG 2000. It can be expected that the requirement to deliver such pre-generated tiled products will continue as this delivery mechanism is required by many organizations and provides a clearly defined product that conforms to the pre-defined standards. Therefore, the need will remain for companies to produce large sets of tiled products. With the size and number of these projects increasing, production companies are being challenged to create very large numbers of such tiled images.

With ArcGIS Image Server, the actual processing of the pixels is performed on demand as required for a user's specified area of interest, resolution, and projection. This processing can be performed in near-real-time for screen sized requests as a user zooms into a specific area. Large image requests, that define the extent of a map sheet or large format plot at high resolution, can also be made. This feature can be used to process and extract large tiles of imagery from the server. The output can be written to a number of formats such as GeoTIFF or JPEG 2000. As with the smaller requests, the important image level metadata can also be derived as output. This metadata is extensive and includes the mosaic polygons that define which part of the output comes from which input raster, as well as the processing performed and metadata from the original imagery.

8. PRODUCTION SCALING

An important scalability feature of ArcGIS Image Server is its ability to process multiple requests simultaneously and spread the processing load over multiple service providers. The server acts as a broker and performs load balancing of multiple requests over one or more service provider machines. Service providers are the work horses of the system that perform the processing on request. These are multi-threaded and on a multi-CPU or multi-core computer environment enable the full processing power to be utilized. Combined with the optimized data access and very efficient algorithms, ArcGIS Image Server can rapidly generate orthoimage products.

Another feature added to ArcGIS Image Server is the Derive Images tool. This acts as a special client that takes a list of required images as input including their spatial reference system, extents, resolution, and output format. The Derive Images tool is capable of making multiple simultaneous requests. Combined with the distribution of these requests by the server of multiple cores and machines, enabling very high production rates.

9. CONCLUSION

ArcGIS Image Server, through its database-centric management approach and ability to generate image products on demand, enables workflows to be optimized and removes bottlenecks inherent in linear production workflows. The server can be used to dynamically server imagery products to large numbers of users as well as form the base for the generation of cached map services. The fast and optimized processing combined with the removal of intermediate products being written to disk results in very high production rates with higher image quality through reduced resampling.