Airphoto processing simplified

The growing volume and diversity of digital images have created the need for an automated image processing system that can handle imagery efficiently and economically. PCI Geomatics takes the call with its GeoImaging Accelerator that offers high-performance image processing

David Piekny

THE VOLUME, AVAILABILITY AND diversity of digital images are increasing steadily due to new high-resolution satellites and commercial sales of digital airphoto cameras worldwide. Tens of thousands of images are becoming increasingly frequent and the trend for precision moves towards smaller pixels and more finely tuned bandwidths; this creates a wide and varied selection of products which has become a daunting processing task. Whether for company growth or competitiveness, the need of the hour is an automated processing system to handle imagery efficiently and economically.

Responding to the needs of the digital airphoto production market, PCI Geomatics has introduced a high-performance image processing system called the GeoImaging Accelerator (GXL), based on multi-core CPUs and Graphical Processing Units (GPUs) with distributed processing. The GXL system uses six areas of development including system automation, modular components, distributed processing, processor optimisation, system sustainability, and risk mitigation, to achieve high-performance and quality throughput for today's large airphoto projects.

This integrated hardware and software solution enables multifold speed improvements for orthorectification and significant gains for mosaic creation – up to 100 per cent increase in speed.

Overall, system automation in GXL provides several important benefits including consistency, increased production and revenues, higher quality assurance, and reduction in costs for process management, operations and training. Another benefit of automation is a more effective use of skilled operators and technological capital, since automating a manual process allows transfer of skills from basic tasks to value creation by managing quality assurance and product creation.

Overall, system automation in GXL provides several benefits including consistency, increased production and revenues, higher quality and reduction in costs.

Automation also ensures effective hardware usage and the GeoImaging Accelerator incorporates a Job Processing System (JPS) to manage multiple users, job classes, definitions and scheduling. Parallel processing is used for large image sets, i.e. orthorectification, as well as for tiled processing of

singular tasks such as mosaic creation. Searching, prioritisation and job sorting are all features of the JPS, resulting in efficient job management (Figure 1).

Scalability is another element of the GXL and is applicable to the addition of new processing nodes in the JPS, which will be discussed along with distributed processing, as well as to the migration from single, monolithic, stovepipe processes to easily-arranged, modular workflows built from Off-The-Shelf (OTS) components. An OTS approach reduces implementation time and increases the flexibility of any image processing system , for instance in an ortho-mosaic workflow used with UltraCam imagery.

Workflows can be fully automated and allow breakpoints for QA/QC and process control. This means image processing jobs can be chained together, run with multiple parameter sets, and components can be re-used to reduce migration and update costs. The modules and workflows are managed by the JPS (Job Processing System), which directly handles both data and metadata ingest and output formatting modules. Individual modules are described below:

IO IMPORT: Import interior orientation

EO IMPORT: Import exterior orientation, PAT-B or Bingo model.

AP MODEL: Calculate airphoto model

AUTOTIE: Optional automatic tie-point collection

AUTOGCP: Optional automatic GCP collection

ORTHO: Orthorectification (at 1:1 sampling)

AUTOMOSAIC: Automatic mosaic, preview generation, seam line selection and colour-balancing

FORMATTING: Output format specifications.

Just as OTS modules provide flexibility and reusability, so does a distributed processing environment where multiple computers are managed and controlled from a simple interface – critical to smooth job management. The capabilities of the Job Processing System (JPS) include load balancing and node availability for new tasks with classes of jobs (for instance 'orthorectification') given a priority and a resource flag indicating relative ease of processing. The JPS overhead is minimal and this component is installed on any node, including laptop, desktop and rack

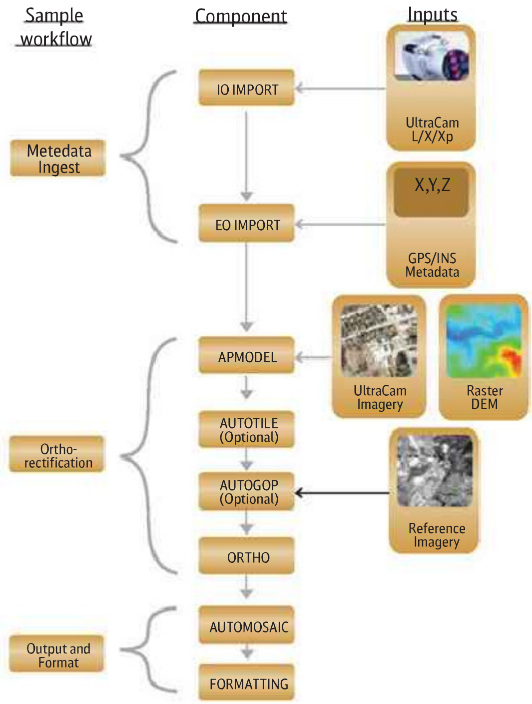

***Figure 1**: Modular workflow showing extension of UltraMap pre-processing into ortho-mosaic creation. Manual QA/QC can be done before or after any module*

An OTS approach reduces implementation time and increases the flexibility of any image processing system.

configurations, with or without Graphical Processing Units (GPUs). The result is ease of scaling to enterprise processing levels by taking advantage of multi-core and GPU-based processing, where each node reports the availability of each processing core contained within. This two-fold approach allows management of a cloud of physical nodes with diverse specifications, as it handles parallel processing on many machines with multi-threading on each.

Beyond a distributed environment and node management, further speed gains are found in processing-core optimisation, specifically GPU-based processing. Using nVidia Tesla servers (dedicated to performing graphical operations), the speed of orthorectification and other tasks have been increased five-

***Figure 2**: GXL Hardware*

fold or more: Producing an ortho from a 330-MB UltraCam scene takes less than ten seconds using a combination of GPU and multi-core processing. The advantages of this acceleration are multiplied as the hardware scales up.

During development of the ORTHO module (Figure 1), numerous techniques were used to increase the processing speed, and the achieved performance figures reflect this result. In general, the sequence of algorithmic improvements proceeded as follows:

1. Grid-spacing dependencies were removed from math models
2. Raster I/O library was improved
3. Separate input and output drives were used
4. The code was compiled in multithreaded mode
5. Processing moved to the GPU

The tests yielded some interesting observations:

1. Removing grid-spacing constraints allowed greater modularity
2. Improving the raster I/O library reduced the runtime
3. Multithreaded processing is helped by drive I/O separation
4. For an average scene, processing time was reduced by over 90 per cent
5. Moving all of the operations possible onto GPUs further reduced processing time by greater than 10 per cent

From interviews and vendor-experience, massive throughput (50 TB weekly) is not only attractive, but necessary to the airphoto market. Hardware infrastructure must be flexible and easily upgradeable, and new sensors and processes need to keep pace with the Earth Observation / Remote Sensing industry – this is a GXL design goal , therefore IT and hardware standards like OpenMP for multithreading and CUDA for GPU optimisation have been adopted. Like the modules previously discussed, multi-core and GPU-based hardware is OTS, resulting in lower total cost of ownership and fast implementation. Both the computing architecture and interface to the GPUs, called CUDA (Compute Unified Device Architecture) have been adopted as general industry standards.

Development environment and tools

openSuSE Linux 10.3 (64-bit)

Intel C++ compiler (version 10.1)

OpenMP 2.5 libraries for multithreading

NVIDIA CUDA SDK release 2.0 for programming the GPUs

Hardware infrastructure must be flexible and easily upgradeable, and new sensors and processes need to keep pace with the earth observation / remote sensing industry – this is a GXL design goal.

Development system hardware

HP Blackbird 002 LCi

1x Intel® Core2™ Extreme Quad-Core 3.0GHz QX9650 CPU

2x 4GB RAM

1x 7200 RPM SATA HDD

2x 10,000 RPM SATA HDD

2x NVIDIA GeForce GTX-280 GPU with 1GB of GDDR3 SDRAM

System Cost: US$~7,600

Mitigation of risk is a benefit for both technical and business reasons, and increasing processing demands requires a stable foundation for predictable growth. Because the GXL is built on a foundation of automated processing (increasing consistency and quality while reducing operational costs), and speed (meaning faster product availability), modular components must be easily configured without the need for complex development or training. The JPS distributes processing for large batch processes, which provides power that is flexible and cost-effective, both of which are absolute requirements. Multi-threading and GPU processing yield speed increases and allow for faster project completion in varied environments. Providing expertise to the geomatics community for over 28 years, PCI Geomatics recognised the need for automation. Profitability, global diversity and active industry participation at all levels provide the stability required to support long-term growth and development of the high-performance GXL system.

Test OS for the test-system described is RED HAT Enterprise Linux 5 and uses NVIDIA Tesla GPU S1070-500 Server Systems. These 1U servers include dual

Table 1: *Performance Metrics (h:mm:ss)*

Process	Average time per scene - baseline	Total processing time – baseline	Average time per scene - GPU	Total processing time – GPU
Data ingest	0:00:11	0:50:19	0:01:01	4:30:00
Orthorectification	0:02:38	11:46:12	0:00:09	0:40:02
Mosaic preparation	0:49:51	0:49:51	0:49:51	0:49:51
Mosaic generation	0:08:02	19:10:15	0:08:02	19:10:15
Total elapsed time	1:00:42	32:36:37	0:59:03	24:10:08

and quad-core Xeon processors and 10,000-rpm Velociraptor HDDs. Additional 4-GB, 2.5-GHz servers can be added, along with a Storage Area Network. The test dataset was 330-MB UltraCam Xp Imagery at 11,500 pixels by 7,500 lines. 267 images were used, each with three channels (RGB) at 10-cm GSD (resolution). The tasks used for processing included airphoto data ingestion, orthorectification, mosaic preparation with seam line and colour-balancing, and finally mosaic creation and geocoding of the output.

The baseline test results were gained by running the original code on the test hardware – no optimisation, multi-threading or GPU processing was used. The GPU timings reflect the results of the optimisation described above and at the time of writing, the mosaic tasks had also been optimised, thought not reflected in Table 1.

This means, overall processing time on the GPU shows a 25 per cent speed increase with no ingest or mosaic optimisation. The ortho-processing time on GPU, which has been optimised as previously described, shows a full 18x speed increase. To put this into perspective, the test system is capable of processing over 400 GB of orthos every day, which for this data-set equates to over 1,000 km² or over 9,500 scenes. The key points to note are the orthorectification production increase, and the importance of data-management for workflow optimisation. With the subsequent completion of GXL optimised mosaic components and a dedicated QA/QC interface for seam line editing, the GXL is poised to meet the demanding performance needs of digital airphoto production around the world.

David Piekny
Product Manager, PCI Geomatics
piekny@pcigeomatics.com

BOOKMARK

INTERACTING WITH GEOSPATIAL TECHNOLOGIES

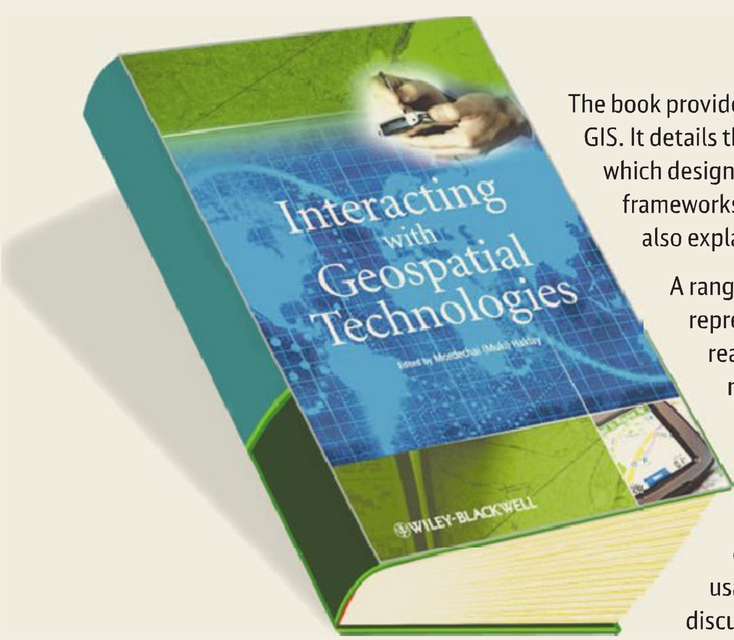

The book provides insights into Human-Computer Interaction (HCI) and usability of GIS. It details the principles of HCI and discusses special usability aspects of GIS which designers and developers need to take into account. Tried and tested frameworks, matrices and techniques that can be used within GIS projects are also explained.

A range of topics from the cognitive models of geographical representation, to interface design have been discussed to provide readers with frameworks and techniques that can be used for computer mapping. Case studies depicting real time application of these techniques for computer mapping application are also discussed. The book is divided into three parts: theory, framework and practicalities and techniques. It discusses the 'HCI for GIS' aspect, providing readers an overview of the history and evolution of these fields. Principles and scientific theories fundamental to usability are discussed in the theory section. The framework section discusses the principles of user-centered design.