'Photogrammetric Week 05' Dieter Fritsch, Ed. Wichmann Verlag, Heidelberg 2005.

Distributed Processing – Added Value to Geospatial Data Production

ORRIN LONG and MOSTAFA MADANI, Huntsville

ABSTRACT

Distributed processing involves taking a linear process which runs on a single computer, and breaking it into a series of independent processes which can be run independently in parallel on several computers. This is desired to increase overall throughput in processing to decrease production cost and delivery time of geospatial data production. Intergraph Corporation has recently implemented distributed processing in the TerraShare and ImageStation OrthoPro commercial software products for image management and orthophoto production, and is implementing distributed processing in other compute-intensive products and modules as well. This paper discusses the design and experiences to date with distributed processing for orthophoto production.

1. INTRODUCTION

The demand for geospatial data, especially imagery and orthophotos, has been steadily increasing. Goverments, private companies, and individuals use geospatial data to make better and faster operational decisions. Photogrammetry production shops are increasingly under pressure to reduce the cost of their jobs. This means they need software products that have higher performance and many automation tools. There are many processes in the photogrammetric workflow, such as orthorectification, that are repetitive and well suited to automation as they require little operator intervention. Distributed processing, a function of TerraShare Advanced Server, addresses this by allowing users to automate tasks by using a group of connected nodes (processors) to carry out these tasks simultaneously in parallel. This allows a process that would take many hours running on one computer to be completed much faster on multiple connected processing nodes.

2. TERRASHARE FRAMEWORK

TerraShare, introduced by Z/I Imaging in 2001, is an enterprise system for geospatial data management and earth imaging production and integrates storage infrastructure with end-user production and exploitation tools to address users' geospatial data management, access, and distribution needs (Wuescher, 2005).

The central data management structure is a virtual file system composed of files and folders. TerraShare files link one or more physical files together, along with metadata stored in the database to form a single logical entity. TerraShare folders contain one or more TerraShare folders and/or one or more TerraShare files. Together they form a virtual hierarchical file system similar to the Windows file system.

A standard TerraShare system is composed of a server running the TerraShare Server application and one or more desktop clients running the TerraShare Client application. TerraShare Server provides the "back office" data management functionality and is designed to run with either a SQL Server or Oracle relational database system. TerraShare Client provides the standard user interface to the system and allows users to add, delete, or browse TerraShare files and folders. TerraShare Client exposes the TerraShare file system to the user with a Windows Explorer namespace plug-in.

As such the TerraShare file system appears integrated directly into the Explorer window and appears as a second file system parallel to the “My Computer” icon.

In June 2004 TerraShare Distributed Processing was introduced which augmented what was primarily a geo-data management system with production monitoring and management tools. The first application to make use of the distributed processing engine was ImageStation OrthoPro for orthophoto production. A standard TerraShare distributed processing configuration is shown in figure 1. A TerraShare Server should always be a designated computer in the customers’ network. The submitter can be each application such as OrthoPro. The TerraShare database either resides on the TerraShare server or on a separate database server. The processing nodes can consist of any computer in the customers’ offices as long as they fit the requirement of the applications’ distributed software. A File Server is optional but recommended as long as the network is capable to handle read/write of the processing nodes without delay or in an appropriate time. If that’s not the case a decentralized data storage system can be used and managed by TerraShare.

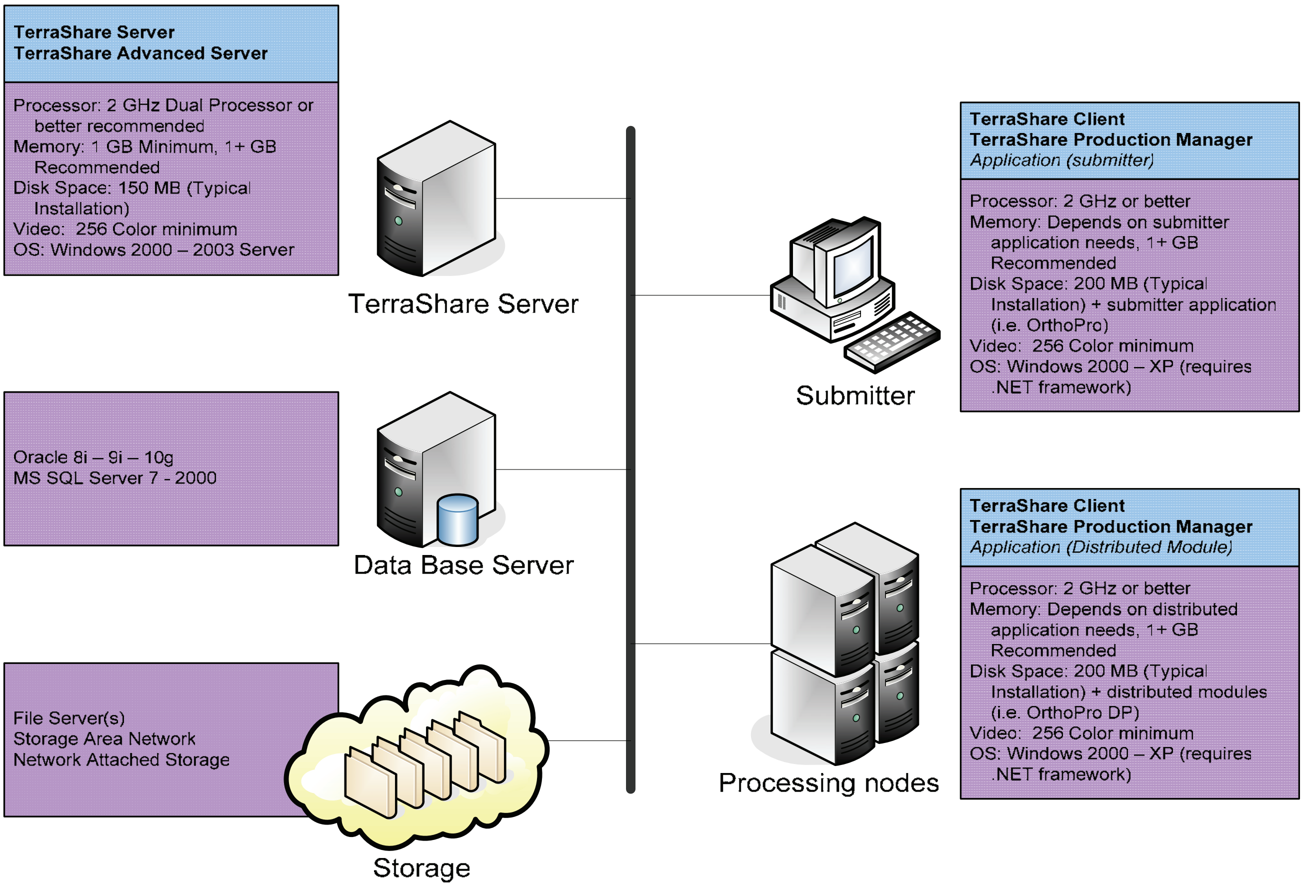

Figure 1. TerraShare Distributed Processing Configuration

3. IMAGESTATION ORTHOPRO

ImageStation OrthoPro is an integrated product for orthophoto mosaic production from aerial frame and line scanner, and satellite imagery (Madani, 1999). OrthoPro includes functionality for project planning, orthorectification, dodging, true ortho, seam line definition (manual or automatic), tone balancing, mosaicking, and quality control. OrthoPro takes as input orientation data and images from a photogrammetry project, and a set of one or more DTM files covering the project area, and outputs orthorectified image mosaics. Input orientation data in a variety of common formats can be imported, and input images and DTMs in a variety of standard formats can be used directly without translation. A key feature of OrthoPro is the ability to process data in different coordinate systems on-the-fly without translation. That is, the input orientation data and DTMs can each be in different coordinate systems, and the output orthophoto mosaics can be in yet a different coordinate system. OrthoPro also allows for manual or automatic processing. With manual processing, the user runs each step in sequence. With automatic processing, the user can define all the job parameters and hit "Start", and have all the steps completed without operator intervention. Output products can be either a single mosaic, or a tiled set of mosaics based on a user-defined database. For output format, the user can specify GeoTIFF, Intergraph, USGS DOQ, or JPEG 2000. Figure 2 shows the basic OrthoPro environment and control panel for processing.

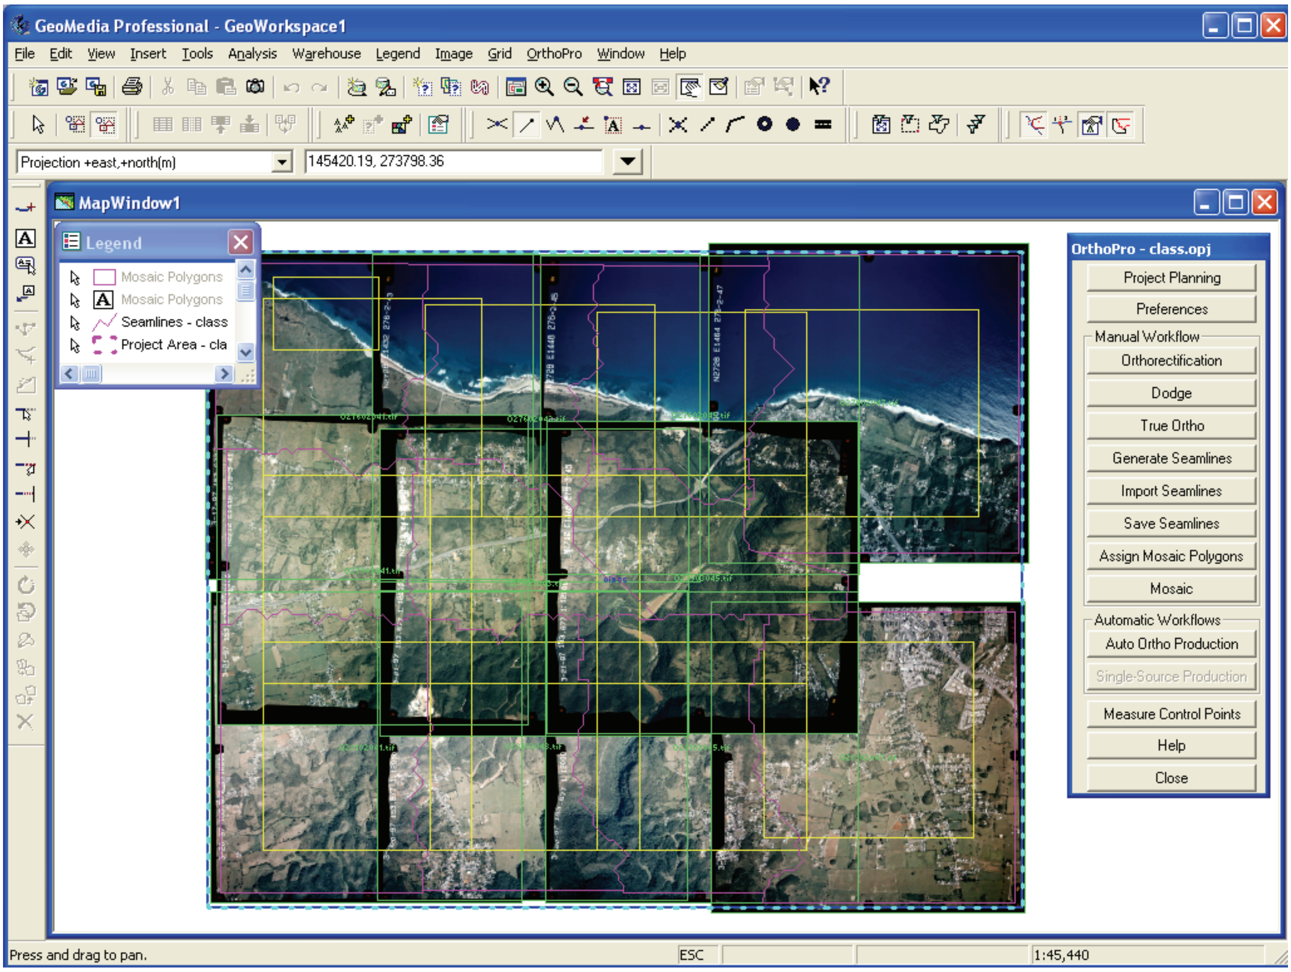

Figure 2. OrthoPro User Environment and Control panel

4. DISTRIBUTED PROCESSING IN ORTHOPRO

OrthoPro previously included the ability to process more than one image at a time on a single machine by threaded processing. The number of images which could be processed in parallel was a user setting, and typically would be set to the number of processors on the machine, which was limited to 1 or 2 by current technology. Thus OrthoPro was already well-structured to have distributed processing implemented, with the promise of breaking the 2 processor limit on a single machine to an essentially unlimited of machines on a network. In practice, however, the maximum number of machines which can be used to speed up a linear set of processes will be limited by the user’s hardware configuration, including network and disk speed, and also data distribution. Initial testing and comments from customers indicate that disk rather than network speed is the primary limiting factor. This provides an opportunity for TerraShare to provide additional value in the future, by automatically distributing the imagery across a decentralized storage system that it is managing. This will effectively distribute the disk activity in addition to distributing the processing, thus assuaging the primary limiting factor of disk activity.

Figure 3 shows the simple user interface change to OrthoPro that was needed to initiate and control distributed processing. Each of the major processing steps in OrthoPro – orthorectification, dodging, true ortho, and tone balance/mosaicking has a dialog which lets the user set processing parameters, initiate processing, and monitor status. To this dialog was added a Job Processing group with Local and Distributed Processing choices, and a sub-dialog to set distributed processing parameters including Type (First Available, Input File, Output File, or Specific), Priority, Start Date & Time, and Notification Method. Type=First Available is most commonly used, and defaults are typically taken for the other settings.

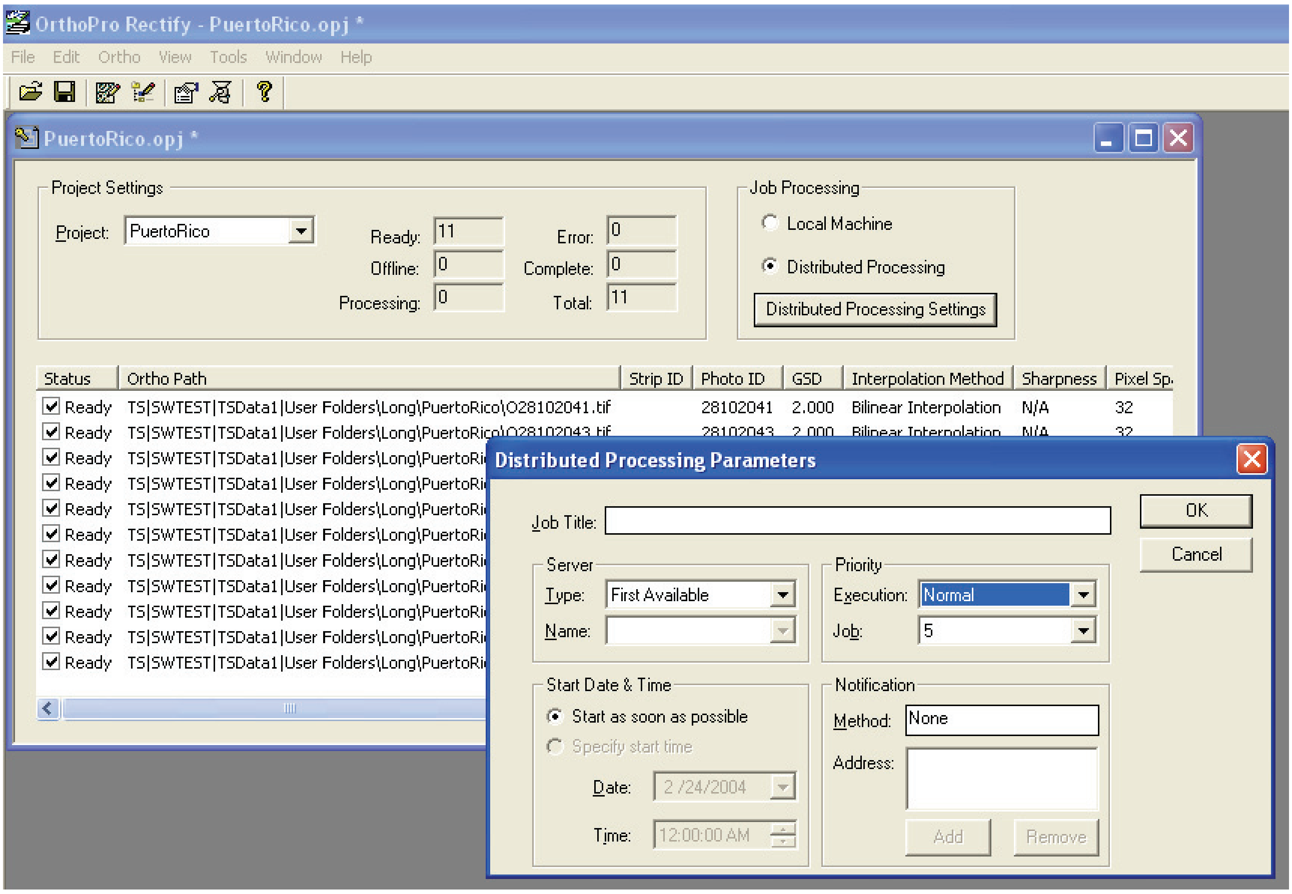

Figure 3. Distributed Processing User Interface

The TerraShare Distributed Processing (TSDP) architecture requires that a client application such as OrthoPro deliver Command Components (CC) to be used to submit a distributed job from the client machine and to execute a distributed job on a processing machine. A Command Component is a COM object that exposes the ICommandComponent interface. A Rectify Command Component (RectifyCC) and others were created in order to add the ability to submit and execute orthorectification jobs using TSDP. Hence, RectifyCC will be installed on the submitting machine and on each processing machine.

When OrthoPro submits a set of defined jobs to TSDP, it invokes the RectifyCC. For each defined job, it passes the Command Component the input and output filenames, tells the Command Component to submit the job, and starts monitoring the returned job Queue Id. In this manner, all of the defined jobs are submitted to TSDP for processing and OrthoPro monitors the status of each job until all are complete. When OrthoPro detects that an orthorectification job is complete, it will retrieve the completion status and log data and add it to the OrthoPro project database, which can be viewed using the OrthoPro Rectify dialog shown above. Figure 4 shows the architecture of the distributed processing implemented in OrthoPro using the TerraShare framework.

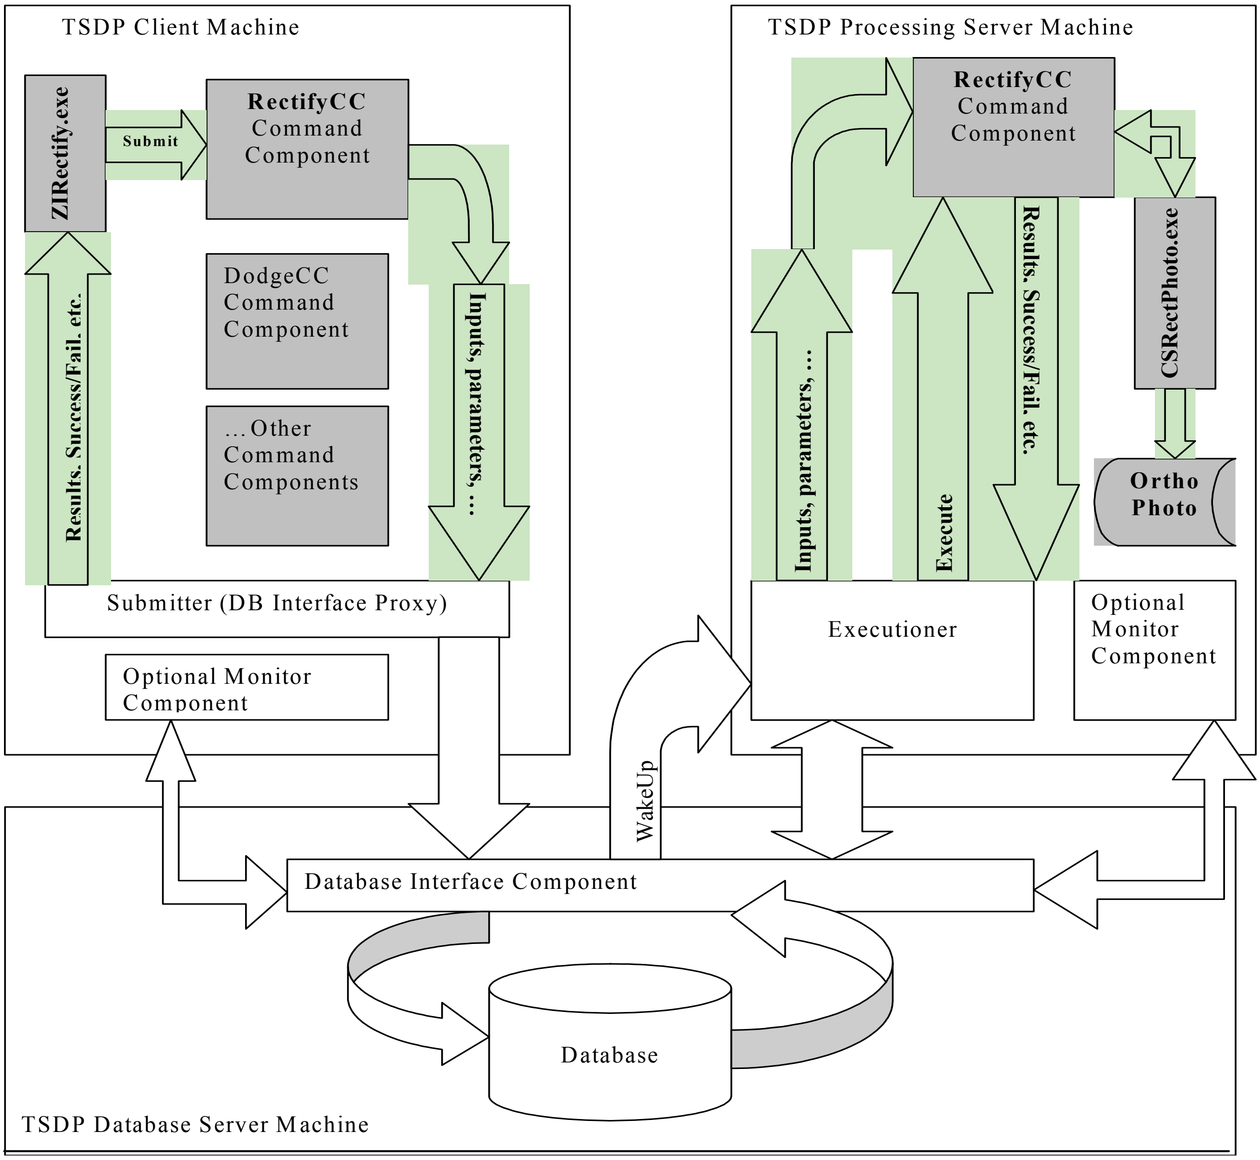

Figure 4. Distributed Processing Architecture

5. PERFORMANCE COMPARISON BETWEEN CONVENTIONAL AND DISTRIBUTED PROCESSING

There is some debate about the theoretical increase in performance which is possible with parallel processing. Amdahl's Law (Amdahl, 1967) states that maximum expected speedup is given by:

$$Speedup = 1 / (s + p/N)$$

Where s is the serial portion of the process, p is the parallel portion of the process, and N is the number of processing nodes. This means that even when the serial portion of a process is small, the maximum amount of speedup is limited by the serial portion, giving a very steep curve of diminishing returns with additional processing nodes:

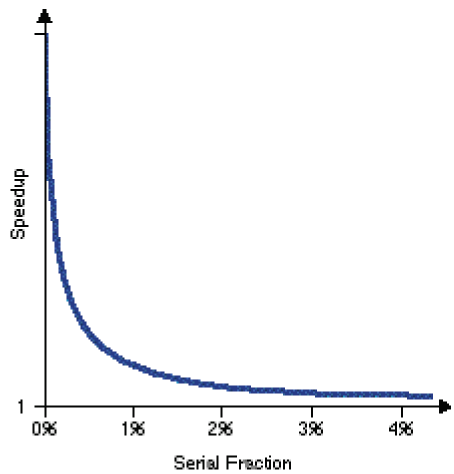

For example, if 10% of a process was serial, and 90% was done in parallel, then using 50 processing machines would only give an 8.5x speedup:

$$\text{Speedup} = 1 / (0.1 + 0.9/50) = 1 / (0.1 + 0.018) = 1 / 0.118 = 8.47$$

In initial tests of speedup of the DMC post-processing by distributed processing, Dörstel (2005) observed a speedup of 2.2x with 3 processors, in line with the prediction of Amdahl's Law. In this case, 11.4% of the process was considered as serial, and 88.6% was considered as parallel. That test concluded that 6 processors would be a practical optimum.

However, Gustafson (1988) achieved much better results than predicted by Amdahl in tests on a 1024-processor system, and concluded that speedup was instead linear. These results indicated the s component of the problem did not scale with problem size, rather the problem size scaled with the number of processors. E. Barsis at Sandia National Laboratories suggested an alternative formula for speedup:

$$\text{Scaled speedup} = s' + p' \times N$$

Where s' and p' are the serial and parallel time spent on the parallel system, and N is the number of processing nodes.

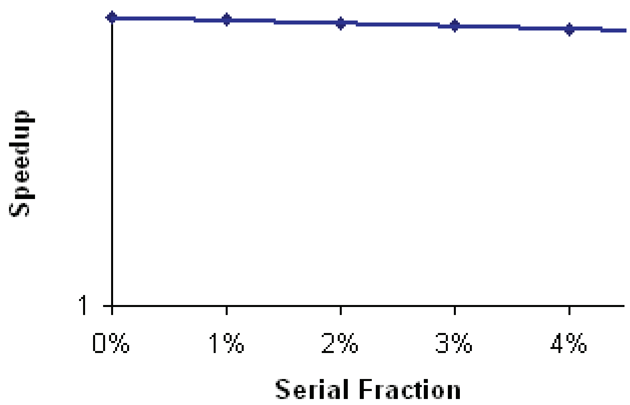

Considering the same example given above of 10% serial, 90% parallel, 50 processing machines would predict a scaled speedup of 45.1x rather than the 8.47x predicted by Amdahl's Law.

To test the speedup of distributed processing applied to the case of orthophoto production, a test project was used to compare processing times when using one workstation and a number of workstations using distributed processing of the orthorectification process. The test project consisted of 819 Images (40.1 GB compressed) and 294 DTMs in ttn format (9.42 GB). The computers used in the tests were:

Workstation	Processor	Memory	Network
IS2003	2x3 GHz	2 GB	94 Mbps
IS2002	2x1.5 GHz	2 GB	94 Mbps
IS2001	2x1 GHz	1 GB	94 Mbps
TerraShare server			
TSServer	1*1.5 GHz	1.2 GB	94 Mbps
Input image server			
Denlmz2k2	2*2 GHz	1 GB	94 Mbps

The input DTM files were located on IS2001, and the output images were written to the TerraShare server. IS2003 and IS2002 were set up to run 2 jobs at a time, and IS2001 was set up to run 1 job at a time due to limited memory. Using only workstation IS2003 (the fastest) the orthorectification of all images was completed in 32.12 hours. Using all 5 processors on 3 workstations, the rectification completed in 18.78 hours. Normalizing for processor speed, this gave a 4.65x increase in throughput of orthorectification, or 0.93x per additional processor – slightly less than linear. This practical results appears to be more in line with Barsis' prediction than Amdahl's. More testing is needed to determine increase of total system throughput including mosaicking, and upper limit of increase based on network and disk characteristics.

6. CONCLUSION

Distributed processing has been implemented in the ImageStation OrthoPro commercial software product from Intergraph Corporation based on the TerraShare framework, and has been shown to increase geospatial data production throughput by distributing the processing of individual images onto multiple computers on the network. The actual improvement in time will vary based on the user's hardware configuration, including network and disk speed, and also data distribution.

7. REFERENCES

Amdahl, G. (1967): Validity of the Single Processor Approach to Achieving Large-Scale Computing Capabilities, In: AFIPS Conference Proceedings vol. 30, AFIPS Press, Reston, Virginia, pp. 483-485.

Dörstel, C., S. Traub, D. Wuescher (2005): Towards Fully Automated Processing of DMC Images. In: Proceedings High-Resolution Earth Imaging for Geospatial Information, ISPRS Workshop 2005, Hannover.

Gustafson, J.: Reevaluating Amdahl's Law. In: http://www.scl.ameslab.gov/Publications/Gus/AmdahlsLaw/Amdahls.html

Madani, M. (1999): OrthoPro: The Z/I Imaging Approach for the Orthophoto Production. In: Photogrammetric Week ’99, Eds. D. Fritsch/R. Spiller, Wichmann, Heidelberg, pp. 247-258.