

Geographical Information System Application of Multiprocessor Multidisk Image Servers

Benoit Gennart and Roger D. Hersch

Département d'Informatique, Ecole Polytechnique Fédérale (DI-EPFL), CH-1015 Lausanne, Switzerland (DI-EPFL). {gennart,hersch}@di.epfl.ch

This contribution analyses the behavior of two kinds of multiprocessor multidisk storage server architectures for a data-intensive application, namely for spatial queries in geographical information systems (GIS). The two kinds of servers are : (1) a workstation cluster architecture with multiple processors, multiple disks, and a shared-bus shared-memory architecture ; (2) a distributed-memory architecture, similar to the T9000 transputer-based architectures, where processing nodes and disk nodes are connected by a high-throughput crossbar switch. The GIS application under investigation is the map overlay with map at different scales and resolutions.

1. Multiprocessor multidisk (MPMD) architectures

The *WorkStation* cluster architecture (Figure 1) consists of a single high-speed backplane bus connected to processors, SCSI-channels and main memory. The SCSI-channels connect secondary storage devices (typically, magnetic disks) to the backplane-bus. We assume that it is possible to transfer data directly from secondary storage to main memory by DMA. The processor local memories are 64KB caches. The raw throughput of a typical shared-bus is 400MB/s (50Mhz,8Bytes).

The *Gigaview* distributed-memory architecture consists of a server interface processor connected through communication links and a crossbar switch to an array of intelligent disk nodes (Figure 2). The server interface processor provides the network interface. Each disk node consists of one or more standard disk(s) connected through a SCSI-II bus to a local disk node processor. The disk nodes support disk access, extent caching, image part extraction and image (de)compression. The crossbar switch supports 32 links, each rated at 12.5MB/s, for a total raw throughput of 400MB/s, equal to the WorkStation backplane throughput.




FIGURE 1. Shared-memory architecture




FIGURE 2. Distributed-memory architecture