Reducing Run Times of Complex Regional Models

Using Cube Cluster on the St. Louis regional model

Background

Major cities around the world are implementing the next generation of travel demand models—disaggregated models that use activity-based processes. While such models offer new analytical possibilities, the additional calculations and data requirements often result in prohibitive run times.

East-West Gateway Coordinating Council, the metropolitan planning organization for the St. Louis, Missouri region, experienced such results. East-West Gateway recently developed a next-generation model using Cube Voyager. The model includes an eight-county network disaggregated into 2,527 traffic analysis zones, including all regional roads classified as collectors or higher. The model considers 17 trip types, and uses a single-constraint, combined mode- and destination-choice algorithm.

After developing the model, East-West Gateway had difficulty validating the model due to long run times. Even after upgrading their model workstation to a state-of-the-art machine, they found that the model required at least 48 hours to run. Wanting to maintain the model's complexity and geographic coverage, they requested Citilabs' assistance. Citilabs upgraded their model to use Cube Cluster.

Method

Cube Cluster, an extension to Cube Base, allows analysts to run models across multiple computer processors and substantially reduce model run times.

How Cube Cluster works

Cube Cluster distributes run processes across multiple computer processors. Cube Cluster uses two forms of distributed processing:

- Intrastep distributed processing — Splits the zone-based processing from a single step into zone groups, and distributes the groups among available processors for separate processing.
- Multistep distributed processing — Distributes blocks of one or more independent modeling steps across available computer processors.

With appropriate configuration, Cube Cluster enables you to run model processes simultaneously, optimizing use of your computer resources and reducing your overall run times.

To test the Cube Cluster implementation on East-West Gateway's preliminary model, Citilabs ran the entire model and ran key components—mode choice and assignment— individually. The tests ran the model simultaneously on up to ten computers with dual-core processors.

Citilabs, Inc.
312 Clay Street, Suite 180
Oakland, California 94607, USA

World Wide Web
www.citilabs.com

Document Revision 001-0

December 5, 2006

citilabs

Analysis

In East-West Gateway's model, the combined mode- and destination-choice step consisted of a single call to Cube Voyager's Matrix program. On one processor, this step took more than 17 hours, the longest single step in the run. Using intrastep distributed processing to run this step across multiple processors, however, reduced the run time nearly proportionally with the number of processors.

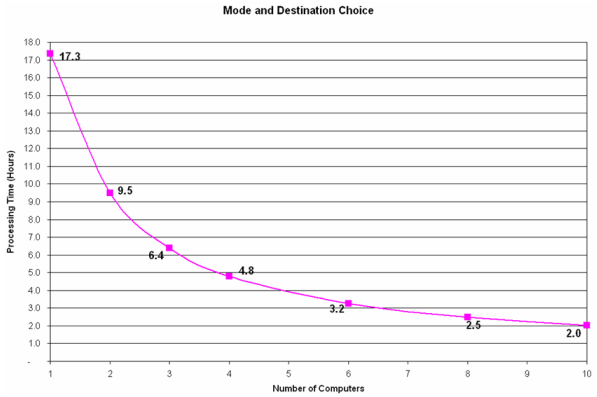

Similarly, for highway assignment, Cube Voyager's Highway program uses an iterative process that finds the best paths between each origin and destination. Distributing the path-finding process across multiple computers can reduce run time. However, after each iteration, each processor must recompile the data. This recompilation time is "lost time," which cannot be reduced by adding additional processors. For East-West Gateway, the lost time was about 14 minutes, or 12 percent of the step's run time. In most cases, lost time amounts to 10 to 20 percent of a highway-assignment step's run time.

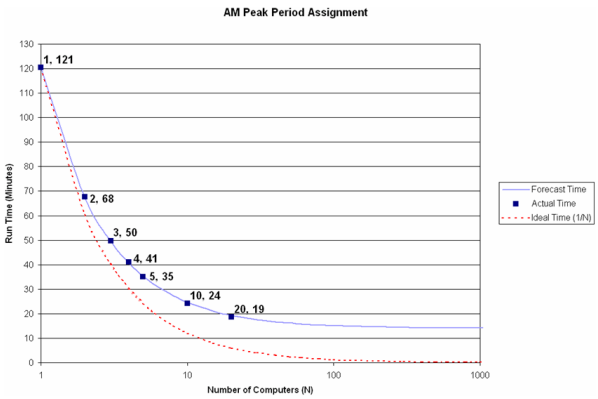

For some steps, Citilabs deployed both intrastep and multistep distributed processing. The highway assignment step, for example, completes a series of independent sub-steps to assign traffic for both peak and off-peak periods. With multistep distributed processing, one computer processor can assign traffic for the peak period and a different computer processor can assign traffic for the off-peak period. With more processors, intrastep distributed processing can share the computing tasks from a single sub-step across available processors.

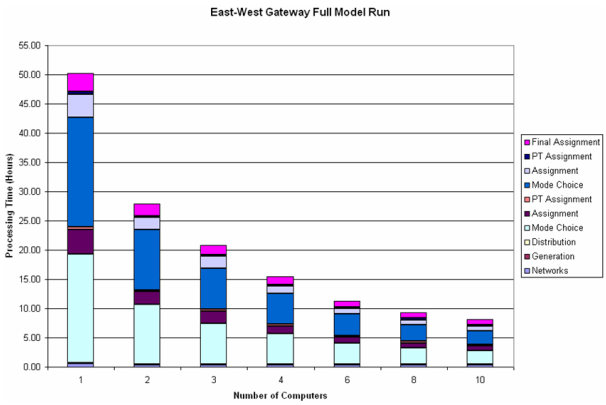

Results

The tests showed that distributing the run of East-West Gateway's preliminary model across two processors reduced the run time by 44 percent, and distributing the run across ten processors reduced the run time from more than 50 hours to just over eight.

Therefore, by implementing Cube Cluster in a ten-processor environment, East-West Gateway was able to reduce run times of their preliminary model by nearly 85 percent. As a result, their model offers a usable solution, one they can execute daily as they validate and refine the model.

More information

To learn more about East-West Gateway, please visit http://www.ewgateway.org/.

To learn more about Citilabs and Cube Cluster software, please visit http://www.citilabs.com.

Cube Cluster is one of the many innovative and market-leading solutions developed by Citilabs—the leader in transportation planning software solutions.