

Customizable Parallel Scientific Stream Processing

Tore Risch, Milena Ivanova, Erik Zeitler
Uppsala Database Laboratory
<http://user.it.uu.se/~udbl/>

The talk presents work done at Uppsala University on processing high volume scientific events from the application area of space physics. The events are produced by distributed digital receivers of signals from space. The event streams are collected in a central processing facility where computations are made based on combining events from one or several receivers. The often expensive computations are specified in continuous queries. Clusters must be used for achieving acceptable throughput. For performance it is particularly important that users can specify parallelizations of computations and such parallelization often depends on what kind of computations are made in the continuous queries. To enable customized parallelization, our systems provide user extensions for parallel processing of high volume continuous queries. The talk presents the user definable parallel event stream processing facilities shown to improve stream throughput significantly on a Linux cluster. Furthermore, the architecture is presented of our new parallel stream processor, SuperComputer Stream Query Processor, SCCQ, for heterogeneous computing environments and applied on our application environment consisting of Linux clusters, a BlueGene supercomputer, and other computers. The talk is based on publications in the VLDB conference and an ICDE workshop.