Spatial Data Streaming Or Streaming Spatial Data: Just Stream It the Way You Like.

Abstract

Have you ever counted the number of times the word "streaming" has occurred in a geospatial oriented conference proceedings over the past few years? Have you ever monitored the growth of the geospatial research and industrial community? Have you ever noticed that geospatial researchers are living the luxury of an era where real-time data is streamed at the convenience of their finger tips? Thanks to advances in the "geosensing" technologies, the geospatial community is "*almost*" ready for their next revolutionary jump. But not quite ready yet! Until we understand how we can process, mine, and analyze the massive amount of data being streamed from geo-sensors every second, the challenge is still there.

However, the horizon looks bright. There have been several success stories to continuously monitor and manage spatiotemporal stream data, each of which has its own position and promise. They ranged from leveraging a geographic information system with streaming capabilities (call it, *spatial-data streaming*) to leveraging a full-fledged data streaming system with spatial libraries (*streaming spatial data*, with the emphasis on the word *streaming* coming first). The spectrum in between these two extremes has been investigated as well.

This talk covers the "Today of Geospatial" and introduces to the audience several geospatial directions at Microsoft, e.g., SQL Server Spatial, Bing Maps, SQL Server BI, SQL Spatial Library, and then, goes into the "Future of Geospatial": *geostreaming* and, more specifically, *geostreaming* in the cloud. The talk is divided into two parts: The first part provides a 10,000 foot view of various *geospatial* efforts at Microsoft and, then, zooms-in through selected angles to highlight key milestones that have advanced the *geostreaming* state of the art. The second part of the talk introduces the Microsoft SQL Server StreamInsight approach to *geostreaming* and highlights its impact on the future of the geo-world. This talk provides the unique lessons that have been taken over the last few years, an industrial perspective of the problem, and definitely a vision of how the ".geo" term will be one of the hottest terms over the coming decades (if not over the coming years).

Short Biography

Balan Sethu Raman is a Distinguished Engineer at Microsoft. He began his career with the aspirations to become a hardware designer. While designing software to help chip designers he developed an interest in software which led him to accept the opportunity to develop software over other opportunities. At Microsoft he was a part of the Windows team developing file systems. During this time he noticed the need to offer more services over the increasing amounts of data and pioneered efforts with the SQL Server team to offer richer services over various forms of data. He subsequently joined the SQL server team and is now leading efforts to extend the Microsoft Data Platform to streaming data.

Mohamed Ali's main research interests focus on advancing the state of the art in the design and implementation of data stream systems to cope with the requirements of emerging applications. In summer 2006, he visited the database group at Microsoft Research (*MSR*) where he and his colleagues

ramped up the Complex Event Detection and Response (*CEDR*) project. Few months later, Mohamed joined the SQL Server group at Microsoft Corporation to incubate the *CEDR* project into Microsoft SQL Server. *CEDR* has become Microsoft SQL Server StreamInsight. Mohamed is also an active member of the ACM SIGSPATIAL group and is leading the *Geo-Streaming* efforts at Microsoft StreamInsight.

Recent Publications

- Mohamed H. Ali, Ciprian Gerea, Balan Sethu Raman, et al. Microsoft CEP Server and Online Behavioral Targeting**.** In Proceedings of the International Conference on Very Large Data Bases (VLDB), Lyon, France, 2009.
- Badrish Chandramouli, Jonathan Goldstein, David Maier: On-the-fly Progress Detection in Iterative Stream Queries. In Proceedings of the International Conference on Very Large Data Bases (VLDB), Lyon, France, 2009.
- Roger S. Barga, Jonathan Goldstein, M. H. Ali, and Mingsheng Hong. Consistent Streaming Through Time: A Vision for Event Stream Processing. In Proceedings of the International Conference on Innovative Data Systems Research (CIDR), Asilomar, CA, January 2007.

Interesting Links

The Microsoft SQL Server 2008 R2 StreamInsight documentation provides detailed information about the product. The documentation is updated periodically. To view the latest version online, use

http://msdn.microsoft.com/en-us/library/dd631799(SQL.10).aspx.

You can also download the latest documentation from the StreamInsight download site

http://go.microsoft.com/fwlink/?LinkID=161910

The Microsoft SQL Server 2008 R2 StreamInsight samples provide hands-on Visual Studio projects that illustrate the key concepts of Microsoft StreamInsight in running code. You can download the samples from CodePlex at

http://streaminsight.codeplex.com/

Refer to the StreamInsight blog for late-breaking information. The StreamInsight blog URL is

http://blogs.msdn.com/streaminsight/

The StreamInsight MSDN portal contains further pointers to the forum and other documentation:

http://msdn.microsoft.com/en-us/sqlserver/ee476990.aspx