


GEO INTELLIGENCE CASE STUDY



SRIS DELIVERS REAL-TIME GEOSPATIAL PROCESSING RESULTS

Innovative GPU-Based Architecture 72x Faster, 12x Cheaper than CPUs

CUSTOMER PROFILE

- SRIS is a leading technology provider and system integrator serving the defense and intelligence community, focusing primarily on GPU technology

PROJECT CHALLENGE

- Deliver real-time processing of geospatial data enabling users to make informed decisions based on timely, actionable information
- Create a system to manage GPU clusters to achieve maximum output

NVIDIA SOLUTION

- NVIDIA® GPU Accelerator with CUDA met all project challenges, outperforming CPUs by 72x

“GPUs allow us to process and filter complex events in near real time. We can now receive vast amounts of data from multiple sources, merge it with an additional reference dataset, process the enriched data utilizing 18 complex algorithms, extract the relevant information, and present it back to the users in a matter of seconds.”

*Srinivas Reddy
Chief Technology Officer, SRIS*

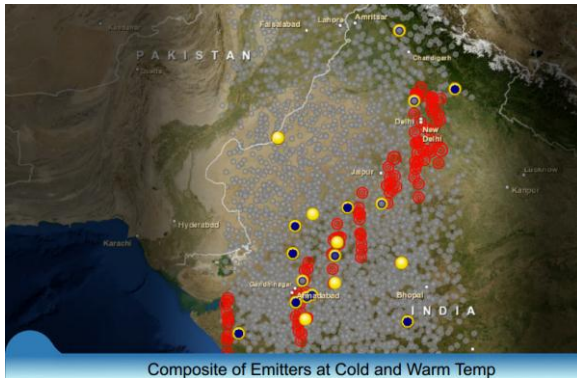
The Real-Time Challenge for Geospatial Processing

Geographic information systems (GIS) are essential to understanding our world. Until recently, the idea of having real-time access to geospatial information was hard to imagine. Our ability to create vast amounts of data has outpaced our ability to leverage the wealth of information this data could provide. For example, processing emitter and sensor data to track wildlife in a remote geography presents a monumental challenge, requiring billions of calculations. Data is flowing continuously, sensors and emitters are in constant motion, not to mention variations in terrain and weather, both of which can impact data collection.

SRIS has created an innovative architecture accelerated by NVIDIA Tesla® GPUs for processing massive amounts of data in real time. “We found that the increased processing capacity in GPUs makes desired results attainable,” said Srinivas Reddy, Chief Technology Officer at SRIS. “We also found that as data scales in size exponentially a new problem begins to arise, which is managing immense data in a distributed GPU clustered environment.” For that, SRIS designed a custom system to manage and process data flows to the GPU clusters.

Cost and Time Improvements with GPU Acceleration vs. CPUs

SRIS put together a cluster for processing geospatial data in real time. They were receiving data from multiple sources, performing required analysis, and sending data to the ESRI server for visualization. “Our first version was a CPU configuration,” said Reddy. “The cost of the CPU system was \$1.2M, with a performance time of 3.6 minutes, so we were unable to achieve true real-time results. After implementing GPU acceleration, we reduced cost by 12x to \$100K and improved performance time by 72x to 9 seconds, near real time.”



“The cost of the CPU system was \$1.2M, with a performance time of 3.6 minutes, so we were unable to achieve true real-time results. After implementing GPU acceleration, we reduced cost by 12x to \$100K and improved performance time by 72x to 9 seconds, near real time.”

Srinivas Reddy, Chief Technology Officer, SRIS

SRIS continues to push the envelope. Very recently, SRIS improved their GIS CUDA algorithm and their data loading technique. They are now able to process the same geospatial data in 0.0017 seconds utilizing the NVIDIA K20 GPU. This is as close to real-time as one can get.

In order to facilitate a real-time data flow, SRIS chose the Storm environment and ran geospatial processing and computation on a GPU cluster of Tesla K10 GPUs and Tesla K20 GPUs. “The algorithms are in CUDA to efficiently access the massively parallel processing provided by the GPU clusters,” continued Reddy. “We used a HyperDex cluster for the storage and retrieval of reference data because it has extremely fast searching capability.”

SRIS developed a platform called MonsterWave to manage and process data flows to the GPU clusters. “MonsterWave allows us to maximize the speed and power of the GPU clusters,” continued Reddy. “It schedules and manages jobs based on GPU awareness, queuing jobs based on GPU status for the most efficient output.” SRIS believes this is the first efficient approach to accelerate real-time geospatial computations with dynamic entities and events. “With this system, we can now receive vast amounts of data from multiple sources, merge it with an additional reference dataset, process the enriched data utilizing 18 complex algorithms, extract the relevant information, and present it back to the users in a matter of seconds. Having the ability to process data this quickly allows users to formulate decisions in a time span where it truly makes a difference.”

Bright Future for Real-Time GPU Enabled Applications

The SRIS team sees lots of potential for these capabilities. “We are currently working on geospatial algorithms for the GPU and providing a capability to tightly couple databases to GPUs,” said Reddy. “This will allow data warehouse appliances and database systems to harness the power of GPUs in the geospatial domain. We are confident we can achieve millisecond responses when analyzing specific geospatial information within the GPUs.”