PART OF THE PROCESS

MATT BROWN EXPLORES THE CHALLENGES ORGANISATIONS FACE WHEN DEALING WITH THE ONSLAUGHT OF EXTREME DATA AND HOW THEY CAN USE THE POWER OF GRAPHICS PROCESSING UNITS TO GAIN INSIGHTS FROM DATA

In the era of digital transformation, location intelligence plays a vital role in improving operational productivity and performance across telecom, logistics, and a variety of other industries.

Gartner predicts that by 2022, 30% of customer interactions will be influenced by real-time location analysis and the businesses that generate location data will be able to create new revenue streams by building data services for other companies and industries.

More importantly, businesses that leverage location intelligence to provide contextual, location-aware experiences and make location-driven decisions will have a competitive edge. According to Dresner Advisory Services, 58% of enterprises rank location intelligence as either critical or very important to ongoing revenue growth strategies. While there's clearly a lot of value to be gleaned from real-time location analysis, there are still a number of challenges that are holding back innovative GIS deployments.

Existing GIS, as well as traditional relational databases that support spatial capabilities, are really struggling to meet the requirements of businesses to build next-generation location apps. Specifically, they are struggling with the inability to perform complex processing and visualisation of spatial and temporal data at scale with legacy systems. Traditional database management systems for GIS have long experienced significant issues with data latency, where batch processing would take hours, if not days to complete. When matched with outdated or generalised maps, businesses struggled to implement a system that is suitable for real-time decision support.

As more location data becomes available from sensors, customers and transactions, there are increasing demands to analyse these datasets interactively and visualise the results on client-side mapping technology. However, traditional databases weren't designed to perform in a world where IoT systems require tracking millions of sensors generating frequent updates, while running complex analysis on that data simultaneously.

The root cause of most of these problems have long been structural. Most databases are simply not designed to perform large-scale geospatial analytics in a reasonable amount of time. The database architecture that underpins most traditional GIS have different components that are in charge of ingestion, storage, analysis and decisioning processes. Massive data latency occurs as a result because these systems are required to ship data internally to perform the processing. Moreover, traditional GIS struggles with large compute workloads because operations such as spatial joins, geometry functions and spatial aggregations are computationally expensive, and it doesn't take advantage of the full capabilities that modern servers and workstations provide.

Challenges at scale

We now live in a world where organisations are trying to analyse tremendous amounts of data. Imagine trying to analyse millions of static data points from customer purchases

and aggregating those based on ad-hoc proximity to retail stores. The types of intersection calculations needed to run these are computationally expensive.

Now imagine running historical analysis on billions of records – your query may leave you waiting all week.

Assuming you can process this volume of data, how do you visualise large geospatial datasets with any sort of interactivity? Web browsers running client-side applications struggle to handle tens of thousands of features, and it takes time to send large volumes of data over the wire. Users requesting more than several thousand points or a thousand complex polygons will likely see the browser slow to a crawl. Eventually, there is a threshold where it's not practical to send all the data across the wire for the client to sort out. Visuals need to be rendered on the server-side.

The other huge challenge for traditional systems is analysing and visualising streaming data. As IoT devices increase in number, advanced location analysis of streaming data is a key requirement of modern location intelligence applications. However, existing systems built on disk-based and CPU architectures are hitting both I/O and compute bottlenecks.

GPU-accelerated technologies for GIS

Graphics processing units (GPUs) present an opportunity to solve the computational challenge. GPUs were originally designed to speed up video games, but their parallelised vector processing capabilities can be used to accelerate geospatial analysis. When coupled with the high performance of in-memory computing, they can deliver a converged engine for advanced analytics, complex geospatial analysis and rendering.

Almost any database can be used to store geospatial data: coordinates can be stored as floats, and shapes can be converted into well-known text (WKT) format and stored in a text column. But while a database can store data this way, it isn't readily available for query. A separate geospatial system needs to retrieve these records, convert them into a geometry object and evaluate the match – one record at a time.

What is needed is a spatially-aware database that has a geometry engine built in, that can work natively with geospatial data and compute relationships between geometries in a single system.

Geospatial queries can be very compute-intensive. As the volume of data increases, performance becomes a critical issue. GPU-based technologies also enable fast query response, even across the largest datasets. Spreading spatial computations across thousands of GPU nodes, across multiple cards and multiple machines, is an exceptional solution for the types of brute-force calculations needed for advanced analysis of large and streaming geospatial-temporal datasets.

With the rise of IoT data such as social media feeds, moving vehicles and sensors, a modern geospatial database must also be able to handle high-velocity streaming data. Multi-head ingest design enables each node in a cluster to share the work of absorbing streaming data. Serving data from system memory, combined with GPU-accelerated queries, offers tremendous query performance improvements. Full table scans are accelerated and less indexing is required, which allows the data to be made available to query the moment it arrives.

New opportunities

This opens up a world of new opportunities for real-time and predictive analytics.

GPU-based technologies are ideal because spatial computations can often benefit from GPU-acceleration during spatial queries and visualisations are built fast locally during the query process. This compresses the stack by reducing communication overhead, and delivers the best performance for both computation and visualisation in a single platform.

Location data is increasingly playing a foundational role in organisations' path to digital transformation. Businesses need technologies that can deliver a new level of real-time data analytics, allowing them to generate new revenue streams and growth opportunities. GPU-accelerated technologies offer a real stepping-stone for companies to set themselves apart from the competition and ensure they are ready to thrive in the Fourth Industrial Revolution.

Matt Brown is a product manager at Kinetica (www.kinetica.com)

Location intelligence is transforming industries to enable better customer experiences and power smarter decisions