Solution Brief

Cities & Critical Infrastructure
Traffic Management

Derq, Intel Collaborate on Video AI for Traffic Safety

The Derq platform uses AI-based video inference to analyze traffic patterns to inform cities about safety challenges and alert drivers, pedestrians, and others in real-time. In a test, the Derq platform supported over 200 camera streams on one Intel server.

Cities are embracing video-based artificial intelligence (AI) to improve traffic safety. Thanks to cameras that are less expensive yet have improved video resolution and features, many cities have deployed video cameras in traffic intersections for safety and to deter speeding and reckless driving. What's missing is analysis of the camera streams to add context and insight and to infer the likelihood of future traffic incidents.

Using technology initially created at MIT, Derq has developed its Real-Time AI Perception Platform that collects data from video cameras and any available sensor to generate cutting-edge analytics. The collected data includes accurate and reliable detection of traffic violations, pedestrian compliance issues, and road-user (near-miss) conflicts to improve safety for drivers, pedestrians, and cyclists.

In this paper, Derq and Intel have come together to test the capacity of the Derq platform in terms of number of simultaneous video streams that can be processed in real-time per server. These tests demonstrate the computational efficiency of the Derq platform and its capability to run high-performance deep learning models with minimal processing and low latency leading to highly accurate data outputs on both edge and cloud-based architectures.

Derq Makes Transportation Safer

The Derq platform aggregates data from different sources such as traffic sensors, signal controllers and connected infrastructure. It then runs real-time AI-powered processes to extract intelligence from the data to enable traffic control applications, connected and autonomous vehicles (CAV) alerts as well as actionable safety and traffic insights.

There are two components to the Derq platform:

- **Derq Sense:** This edge-based solution provides real-time infrastructure-based analytics and is designed to run in a multi-access edge compute (MEC) node where it can communicate with cameras, sensors, roadside units (RSUs) and can pass on alerts to drivers about safety events or traffic congestion (see left-hand side of Figure 1). Derq Sense is a containerized application that can run on MEC services that feature the appropriate compute performance and transport latency.
- **Derq Insight:** this solution can run on either an on-premises server or MEC node to generate real-time safety and traffic insights, providing up-to-date, accurate and granular data. Derq Insight enables in-depth analysis of safety issues, countermeasure development and traffic performance assessment. Derq Insight delivers the dashboard (right side of Figure 1) that traffic engineers use to have access to various events, heatmaps of safety hotspots, and flexible and comprehensive reporting.

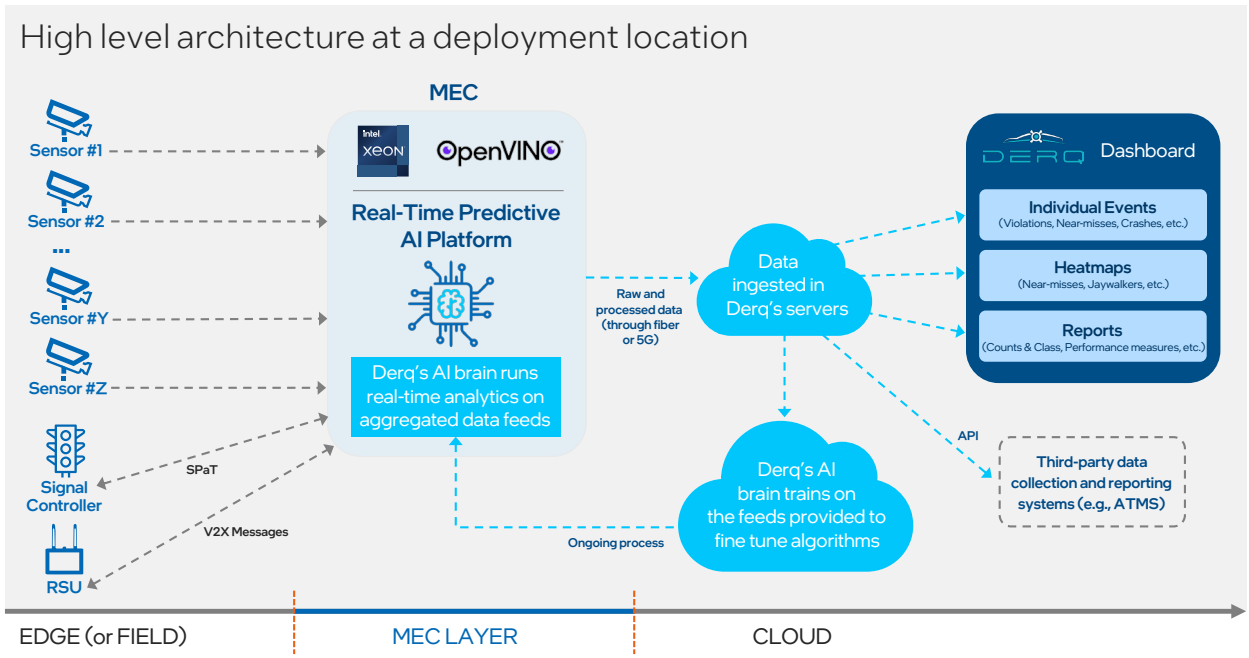

Figure 1. Derq platform solution architecture from edge to cloud.

Some of the key features of the Derq technology include:

Open System: The Derq platform is compatible with a wide range of sensors (video detection, CCTV, radar, etc.), traffic controllers and roadside units (RSUs). The technology can be integrated with both existing and newly deployed infrastructure.

Flexible Deployment: The Derq platform can be deployed at the roadside edge in a decentralized architecture for low latency safety-critical applications such as vehicle-to-everything (V2X) collision avoidance, or signal and rectangular rapid flashing beacons (RRFB) actuation. The platform can also be deployed in a centralized server – either on-premises or in the cloud – for advanced safety and traffic insights as well as real-time data for active traffic management systems (ATMS), integrated corridor management (ICM) and dynamic messaging signs (DMS) operations.

Multi-Application Platform: Applications are run on top of a robust, versatile, and modular platform, allowing the operator the freedom to subscribe to new applications with minimal added installation and configuration. The data can be made available through a dedicated web-based dashboard for real-time insights reporting as well as robust APIs to the customers' existing dashboards and systems.

Real-Time Advanced AI Analytics: The analytics are aided by proprietary and patented computer vision and machine learning techniques enabling accurate road user detection, classification and multi-sensor tracking with predictive capabilities.

User Persona	Use Case
Traffic Engineer	Analyze traffic patterns and safety hotspots to help engineers make informed decisions on potential corrective actions at targeted intersections to enhance traffic operations and safety.
	Identify traffic patterns and safety hotspots in an intersection to help engineers understand root causes and make better design decisions.
	Immediately notify if the infrastructure is offline so that engineers can take necessary actions to get the system up and online.
	Send information on real-time safety and traffic patterns at an intersection and send smart recommendations to adapt traffic light systems in an automated manner.

Vehicle Driver	Notification of dangerous events happening on a driver's route to help them stay safe from potential accidents.
	Assist driver to avoid traffic congestion based on analysis of historical and real-time data analysis. Allows driver to choose an alternate path that saves travel time.
	Notify driver of need to be extra careful based on analysis of their route and the historical near-miss or collision patterns identified for that turn or movement.
	Notify driver if a particular area of the road they are driving on is usually congested with pedestrians at that time so they can be extra cautious.
	Notify driver of real-time signal phase and timing to optimize his/her travel time and speed.
Public Transport Vehicle Driver	Send driver safety insight alerts and trends to alert to the need for extra caution on parts of their route.
	Guide driver on how to reduce total travel time and lower fuel usage based on historical traffic patterns.
	Notify driver of real-time signal phase and timing to request signal priority as the vehicle approaches an intersection.
Public Transport User	Send real-time alerts and information on when the bus is expected to arrive at a particular stop.
Vulnerable road users (bicyclists / pedestrians)	Inform all vehicles on the road to be aware of presence of a vulnerable road user and vice versa so that the vehicle drivers can take preemptive action to avoid a conflict/collision and vulnerable road users can safely navigate the road and quickly reach their destination.

Table 1. Prominent use cases enabled by the Derq system.

Derq Works with Intel Technology to Power AI

The Derq platform uses a CPU-based object detection pipeline that leverages Intel® OpenVINO™ on Intel® Xeon® Scalable processors and Intel's edge solutions.

- Intel® Xeon® Scalable processors deliver optimized performance, scale and efficiency across a broad range of data center, edge and workstation workloads. The test results in this paper are based on 4th Gen Intel® Xeon® Scalable processors, which are built with a unique CPU architecture that integrates up to eight accelerators into the processor providing a power and performance efficient way to achieve higher performance.
- Intel® OpenVINO™ Toolkit accelerates the inferencing and development of machine learning solutions. It seamlessly works with popular AI frameworks and offers a variety of model compression techniques to optimize, inference performance along with a common API to enable deployment across Intel.
- The solution includes edge centric software building blocks for hosting Derq remote applications. The stack is an edge-native distributed computing platform that provides cloud-like ease, resiliency and security for deploying and managing demanding container-based workloads like AI at the edge.

Testing Set Up

Derq developed a series of tests for this paper to demonstrate the ability of the Derq platform to run a large number of sensor streams with minimal processing and hardware constraints providing the data needed for high performance traffic management, CAV and analytics solutions.

The tests utilized two systems under test (SUT) based on Intel CPUs. The number of streams delivered by each server provides a guideline for how many intersections can be covered by each server. The Platinum SUT utilized a 52-core 4th Gen Intel® Xeon® Scalable Platinum 8470N processor (2.0 GHz), and the Gold SUT was based on dual 32-core 4th Gen Intel® Xeon® Scalable Gold 6438N processors (2.2 GHz).

For the first set of tests, the SUT servers were configured to process 640 x 480 video streams from multiple cameras at 10 frames per second (fps) per container. For each test, the inference time and compute resources needed to process the streams were analyzed.

Test Results

The tests confirmed that at a specific configuration of video streams per container and total number of containers, the Platinum SUT processed 208 streams across all cores, with the Gold SUT processing 128 streams. In a real-world use case, these servers can provide complete analytics on up to 208 traffic camera streams per server (or 128 with the Gold SUT). The performance also gives traffic engineers a sense of how hosted on-premises or MEC servers could be used to deliver Derq services.

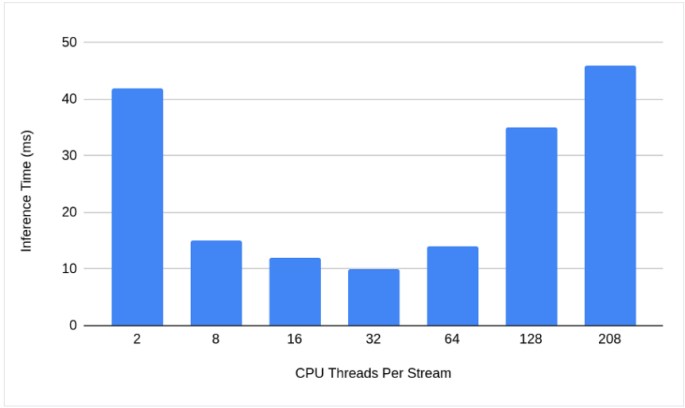

Figure 2. Chart of inference times for a single data stream at super-high 100 fps. This test shows that optimal latency is 32 streams. More broadly, the test shows the configurability of the system for a variety of applications with differing framerates.

In the subsequent tests, the frame rate was increased up to 100 fps – ten times normal – to show how the inference time can vary based on the number of cores used to process a stream. Figure 2 shows the threads ranging from two to 208 with inference time ranging from zero to 50 milliseconds. A 100-fps stream is processed by each of the core configurations with the lowest inference time coming from the 32-thread configuration. While most applications will not use 100 fps, this test shows the configurability of the system for a variety of applications that do feature differing framerates.

Although not tested, Derq expects that the stream density could improve significantly in a server with an Intel GPU to accelerate stream processing. The company is expecting to test this stream density in a future paper.

Conclusion

Video AI has the potential to improve traffic, reduce traffic-related fatalities and facilitate traffic management, CAV and analytics applications. Using AI, the Derq platform is able to analyze historical events at intersections all across a city and use that analysis as actionable information to warn drivers, bikers and pedestrians of high risk for accidents while also presenting this data to traffic engineers to help ongoing traffic planning.

Working with Intel, Derq demonstrated an ability support over 200 traffic camera streams with accuracy and low latency using Intel® Xeon® Scalable processors along with OpenVINO and edge focused software stacks. The solution can be deployed in a distributed architecture that places servers at the edge for very low latency or centralized in the cloud or in the customer's data center.

Intel and Derq continue to partner on optimization of workloads on the Derq platform for future live deployments within cities and critical infrastructure.

Learn More

Derq Homepage

Intel® Network Builders

OpenVINO™ Toolkit

Intel® Xeon® Scalable processors

About Derq

Derq is an award-winning MIT-spinoff powering the future of roads for safer and more efficient movement of road users and autonomous vehicles. Through its proprietary and patented technology, Derq provides cities and fleets with an artificial intelligence (AI) platform that powers advanced analytics and connected & autonomous vehicle (CAV) applications to help them improve road safety and better manage traffic.

About Intel

Intel (NASDAQ: INTC) is an industry leader that focuses on developing technologies that change the world, drive global progress, and enrich human lives. Driven by Moore's Law, Intel is dedicated to the continuous innovation of semiconductor design and manufacturing to provide customers with solutions to major challenges. By integrating intelligence with the cloud, network, edge, and various computing devices, Intel unlocks the potential of data and helps improve both business and society.

Platinum Server: 1-node, 2x Intel® Xeon® Platinum 8470N processors with 256 GB (16 slots/ 32GB/ 4800 total DDR5 memory, HT on, Turbo off, OS: Ubuntu 22.04, 5.15.0-83-generic kernel, 1x Intel SSDSC2KG960G701 S4600 Series 960GB 2.5" SSD, DERQ 1.2.1, cmake, OpenVino (version 2022.3), gstreamer, test by Derq on <08/03/2023>.

Gold Server: 1-node, 2x Intel® Xeon® Gold 6438N processors with 256 GB (16 slots/ 32GB/ 4800 total DDR5 memory, HT on, Turbo off, OS: Ubuntu 22.04, 5.15.0-72-generic kernel, 1x Intel SSDSC2BX016T4R 1.6TB MLC SATA 6gbps 2.5" SSD, DERQ 1.2.1, cmake, OpenVino (version 2022.3), gstreamer, test by Derq on <08/03/2023>.

1123/CCI/DJA/HO9/PDF ♻Please Recycle