Sentinel: Real-Time In-Situ Intelligent Video Analytics for Mobile Platforms

Shawn Recker & Christiaan Gribble
Applied Technology Operation
SURVICE Engineering

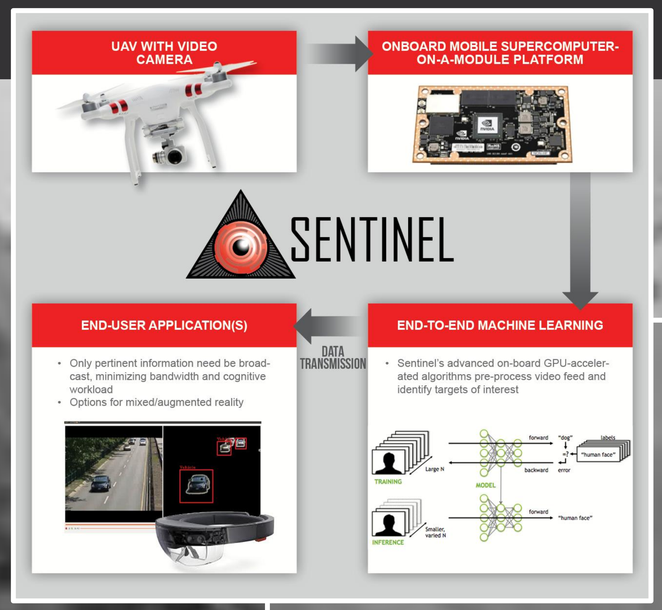

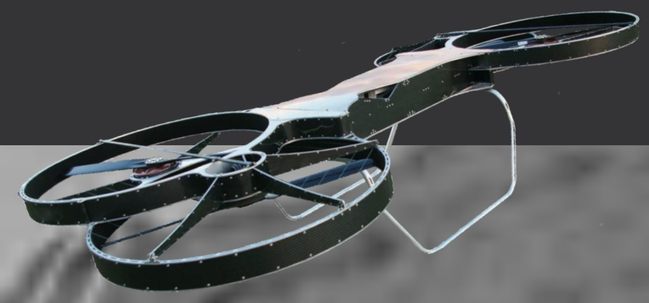

Overview

Sentinel is a system for real-time in-situ intelligent video analytics (IVA) on mobile platforms. Sentinel combines state-of-the-art techniques in high performance computing (HPC) with Dynamic Mode Decomposition (DMD), a proven method for data reduction and analysis. We are currently working to enhance and extend our initial prototype to create an affordable and scalable module to provide real-time in-situ IVA across a wide range of mobile platforms.

By leveraging current miniaturization and performance trends in modern computing hardware, our carefully crafted computer vision (CV) algorithms and high-performance implementations now enable the ability to execute traditionally advanced video analysis tasks—including object detection, identification, and tracking—on low-cost, low-power mobile chipsets and platforms.

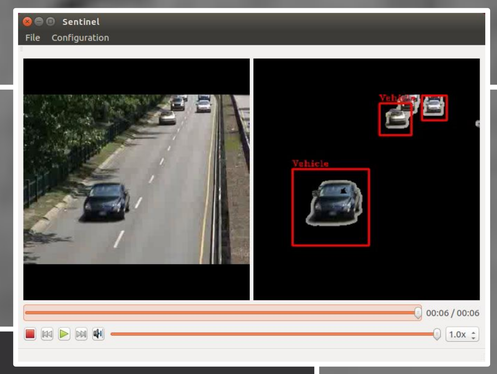

Data Reduction & Analysis

Our collaborators at the University of Washington recently introduced compressed DMD (cDMD) as a fundamental mathematical framework for robustly performing real-time analysis of data, including background/foreground separation and feature extraction for downstream object classification. When combined with compression, the resulting cDMD technique processes extremely high-dimensional data in real-time on modern computing architectures while still maintaining high-quality data reconstructions. We are utilizing and extending cDMD in a streaming and multi-resolution software architecture to scale our system to the challenges inherent to real-time object detection, identification, and tracking.

Massively Parallel Computing

Massively parallel computing architectures offer a compelling means to satisfy the demands of compute-intensive applications, including real-time data analysis. These architectures boast tens, hundreds, or even thousands of processing cores that provide a massively parallel computational environment at a fraction of the cost of traditional HPC systems. We are leveraging these architectures for accelerated CV, thereby scaling our system across the full range of computing architectures, from traditional workstations to mobile and embedded environments.

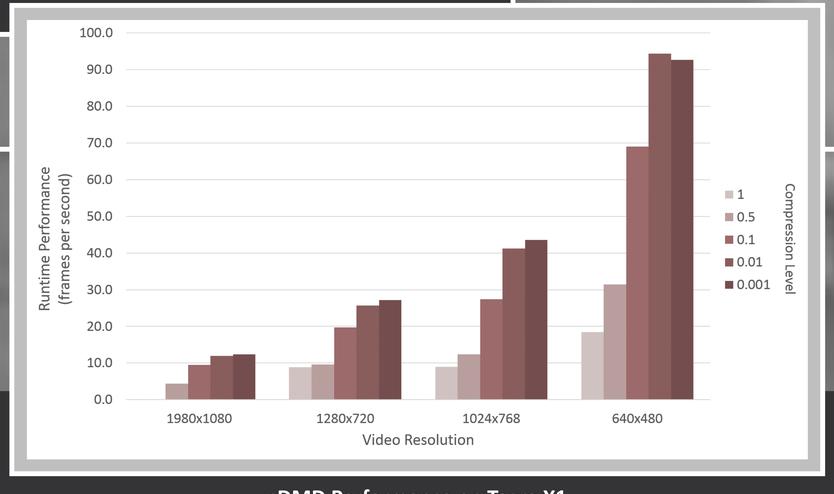

cDMD Performance on Tegra X1

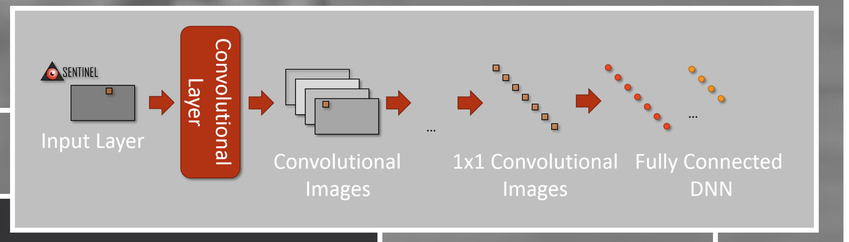

Machine Learning

Machine learning via deep neural networks (DNNs), or so-called *deep learning*, is at the forefront of modern data analysis R&D. Deep learning uses many-layered DNNs to learn levels of representation and abstraction that make sense of data such as images, sound, and text. We are exploiting deep learning on modern GPUs to quickly design, optimize, and refine effective neural networks for object detection, identification, and tracking.

SENTINEL

SURVICE