From Video to Knowledge

Unmanned aerial vehicles, such as the Predator shown here, use multiple cameras for persistent surveillance. (Courtesy of Department of Defense.)

Livermore's Persistics data-processing pipeline combines graphics-based computer hardware and clever software to extract meaning from wide-area overhead surveillance video.

UNMANNED aerial vehicles, such as the drones Predator and Reaper, are rewriting the rules of battle in Iraq and Afghanistan and are being considered for monitoring U.S. borders. Operated by remote control, sometimes thousands of kilometers from the battlefield, the aircraft use multiple cameras to track and destroy the enemy.

These drones can fly several kilometers high, hover over targets for 24 hours and longer, and are often armed with bombs and air-to-ground missiles. By finding, tracking, and monitoring people, vehicles, and events of interest on a continuous basis with wide-area video cameras, the aircraft provide persistent intelligence, surveillance, and reconnaissance (ISR). In theory, persistent ISR coverage can prevent enemies from evading overhead surveillance systems, thereby enabling fast decision making for tactical missions, with reduced risk to U.S. forces and noncombatants.

A new Livermore computational system is designed to help the Department of Defense and other agencies monitor tens of square kilometers of terrain from the skies, with sufficiently high resolution for tracking people and vehicles for many hours at a time. The system, called Persistics, promises to overcome a severe and growing problem: the overwhelming volume of video data generated by modern overhead imagery. Military wide-area video surveillance cameras such as Constant Hawk and Angel Fire are collecting ever-increasing amounts of valuable video imagery over large geographic areas. However, the technology to promptly process, store, and extract meaning from data, and then transmit this information over long distances for further analysis, has lagged far behind.

Data Overwhelm Capabilities

"Sensing capabilities for wide-area video reconnaissance have evolved to far exceed our wildest expectations," says physicist Sheila Vaidya, principal investigator for Persistics and deputy program director of Defense in Livermore's Global Security Principal Directorate. "There exists an enormous—and growing—gap between the amount of information continuously collected by existing sensors and our ability to quickly take advantage of that information.

Advanced and accurate analysis techniques are urgently needed by the people who are tasked with categorizing, indexing, annotating, and drawing conclusions from the petabytes of data collected daily in theatre operations."

Computer scientist Mark Duchaineau, one of Persistics' key architects, says, "The data-processing infrastructure for national security is not designed for the amounts and types of data being generated by unmanned aerial drones relative to the scale of human resources available to analyze them through conventional 'VCR playback' style of viewing." In addition, the communication bandwidth supporting data transmission from air to ground and the archive storage capability are much too slow or too small to support fast-turnaround human analyses.

Duchaineau says, "Several years ago, Department of Defense managers presented us with this interesting processing challenge." The Livermore scientists knew they needed to work on the software side to get the most information out of the streaming video imagery. The greatest opportunity for helping tackle this challenge lay in the development of advanced algorithms (short computer programs) implemented on new computer architectures.

Persistics, the product of several years' effort, is an innovative data-processing "pipeline" that takes a radically different approach to addressing the video-data overload challenge. (See the box at right.) The technique retains the level of detail necessary for detecting anomalies while at the same time compressing the unchanging "background" and everything in motion by about 1,000 times without losing pertinent information. As such, the approach ameliorates the dearth of communication bandwidth for transporting video without losing image fidelity. Indeed, Persistics technology can produce subpixel resolution for the background and any "movers" (people and vehicles), thereby allowing for additional analyses of suspicious activities. A single pixel can correspond to anywhere from several square meters to less than 1 square meter of real estate.

The Persistics architecture can support near real-time monitoring for tactical combat missions as well as forensic analysis of past events. Its analysis algorithms permit surveillance systems to "stare" at key people, vehicles, locations, and events for hours and even days at a time while automatically searching with unsurpassed detail for anomalies or preselected targets. The Livermore breakthrough combines optimized hardware featuring the newest generation of graphics chips (typically used for computer gaming) with innovative algorithms. Some algorithms focus on compressing data while others analyze the streaming video content to automatically extract items of interest.

The system stabilizes incoming video imagery through extremely accurate calibration of onboard cameras and correction of pixel-to-pixel intensity variations. In this way, a mosaic made from hundreds of individual sensors

Inside Persistics

The Persistics data-processing "pipeline" uses a technique called pixel-level dense image correspondence to stabilize video, eliminate parallax (the slight differences in the apparent positions of objects viewed from different cameras), compress background, and provide unprecedented subpixel resolution of moving objects of interest. Dense image correspondence compares the scene from one pixel to another and corrects for individual differences.

"If we know precisely how any two video images relate to one another, we can remove parallax and determine what is background and what is moving, increase signal-to-noise ratio, and increase resolution to see details on vehicles," says computer scientist Mark Duchaineau. The biggest benefit, however, is the ability to compress data without throwing away important information.

The technique allows for image stability, which is essential for accurate

(left) In the manual stitch image, the road does not line up between the collected images. (right) Dense image correspondence corrects for intensity variations between cameras to provide a seamless image.

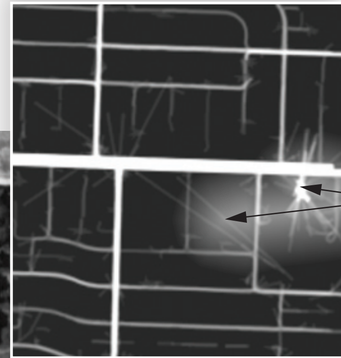

(left) Persistics allows analysts to automatically follow individual vehicles. In this scenario, a vehicle is tracked after a meeting at a presumed safe house (red box) as it travels (green line) to another building. (above) Persistics can extract a road network from flowing traffic and show vehicle motion in a “heat map” representation of traffic density to automatically discern if the traffic pattern changes. Spurious tracks can be caused by the glint of sunshine and other factors.

and high-resolution object identification and tracking. It also permits “seamless stitching,” a process that optically combines images from multiple cameras to create a virtual large-format camera. To do this, individual cameras must be finely calibrated to account for relative stitch position and intensity differences, lens distortion, and pixel-to-pixel variations. The resultant mosaic behaves like a single sensor for Persistics to perform the necessary huge data reduction.

A requirement for Persistics’ thousandfold data reduction is the continual revisit of a surveillance platform over a specified area. The increasing coherence of a scene built up over many passes produces high-fidelity backgrounds from which it is easy to spot movers. By having an aircraft circle the same area, Persistics can calculate the effect of parallax, which causes the perspective to constantly rotate. The corrected image provides a straight-down view of highly detailed terrain. “Removing parallax makes it much easier to determine what is in reality moving in the background,” says Persistics project leader Holger Jones. Current processing techniques use global registration, which generates many false-positive movers (people and vehicles).

Persistics is efficient in allotting processing resources to where they are most needed. For example, it can automatically direct additional computing power to track a particular convoy of vehicles in real time or keep watch over a group of Marines who are asleep in a poorly defended area.

Global registration

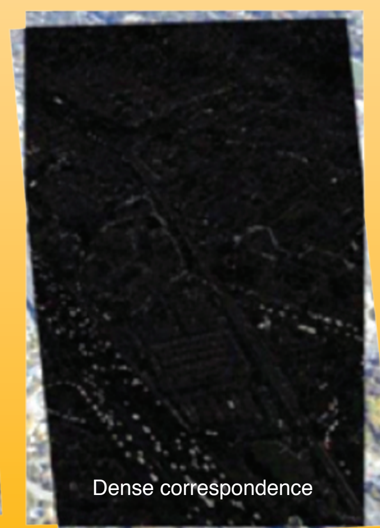

Dense correspondence

Persistics uses parallax-removing pixel-level dense correspondence to register two images and detect what has changed in the time period between capturing each image. The left image shows the original aerial video view. In the center, two images of the aerial view have been corrected with the traditional method of global registration. Because of parallax motion, some buildings as well as moving vehicles are detected. The right image corrected with dense correspondence shows only vehicles in motion (the bright areas); Persistics has removed the static background.

Persistics software is currently housed on a high-density 12-node minicluster, providing an aggregate 40 trillion floating-point operations per second of computing capability. Each node contains a combination of central processing units and seven graphics processing units. Pictured are (left) Sheila Vaidya and Holger Jones.

can be represented as an image from a single camera. This procedure allows the system to exploit the spatial and temporal redundancy in the data and tease out movers from the static background.

In Use at Ground Stations

In collaboration with the National Geospatial-Intelligence Agency, Defense Advanced Research Projects Agency (DARPA), Air Force, Army, and several Department of Defense laboratories, Livermore is incorporating the Persistics pipeline into data-processing ground stations fed with video data from Constant Hawk video cameras onboard both unmanned and manned aircraft. Recent tests of Constant Hawk imagery have validated Persistics' approach, demonstrating 1,000-times compression of raw video collections while maintaining high fidelity. In comparison, compression of still images can reduce data content tenfold, while standard video compression can achieve at best a 30-times data reduction.

Analysts working at ground stations will interact with the transmitted airborne video data. For example, Persistics has been integrated into the Air Force Research Laboratory–developed Pursuer viewer to allow analysts to pan, zoom, rewind, query, and overlay maps and other

(a)

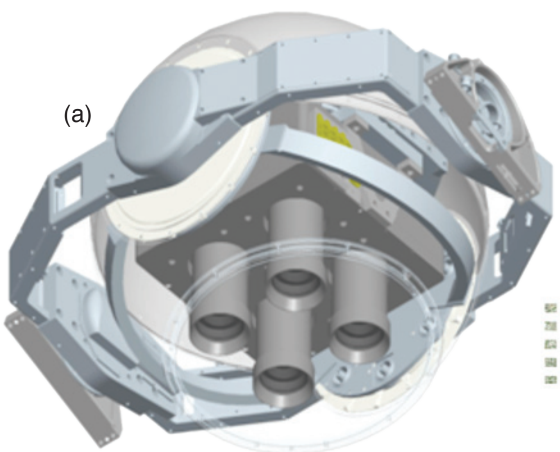

(b)

(a) The Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System (ARGUS-IS) has four stabilized lenses (telescopes), which together use 368 5-million-pixel cell phone cameras. (b) A mosaic of the 368 individual images is stitched together to form one wide-angle seamless image from which details can be extracted.

metadata. With this viewer, they can use Persistics to make requests such as, “Give me the frames that recorded this vehicle from one to two o’clock this afternoon,” or “Show me all the vehicles that stop at this location today.” Says Persistics project leader Holger Jones, “With Persistics, analysts can determine the relationships between vehicles, people, buildings, and events.”

Persistics data-processing modules rely on commercial video-processing hardware designed for computer graphics, video editing, and games. “GPUs [graphics processing units] continue to grow in power, while shrinking in size and energy requirements,” says Jones. In fact, their processing power is outpacing that of central processing units (CPUs), which run most computer operations. “We’re riding the GPU technology wave,” says Jones.

Vaidya led a former project that researched how GPUs might be programmed and used in knowledge-discovery applications relevant to national security. “We realized these processors—traditionally designed for fast rendering of visual simulations, virtual reality, and computer gaming—could provide efficient solutions to some of the most challenging computing needs facing the intelligence and military communities,” says Vaidya. (See *S&TR,* November 2005, pp. 19–20.)

The Persistics software is currently housed on a 12-node minicluster, and each node contains a combination of CPUs and GPUs. To help meet ISR requirements in a cost-effective manner, the Livermore Persistics team has optimized a combination of microprocessors and high-end graphics cards found in both gaming boxes and many personal computers. When combined with a high-speed network and software tools written in open-source (not vendor-proprietary) code, the clusters outperform larger and more expensive proprietary engines in extracting

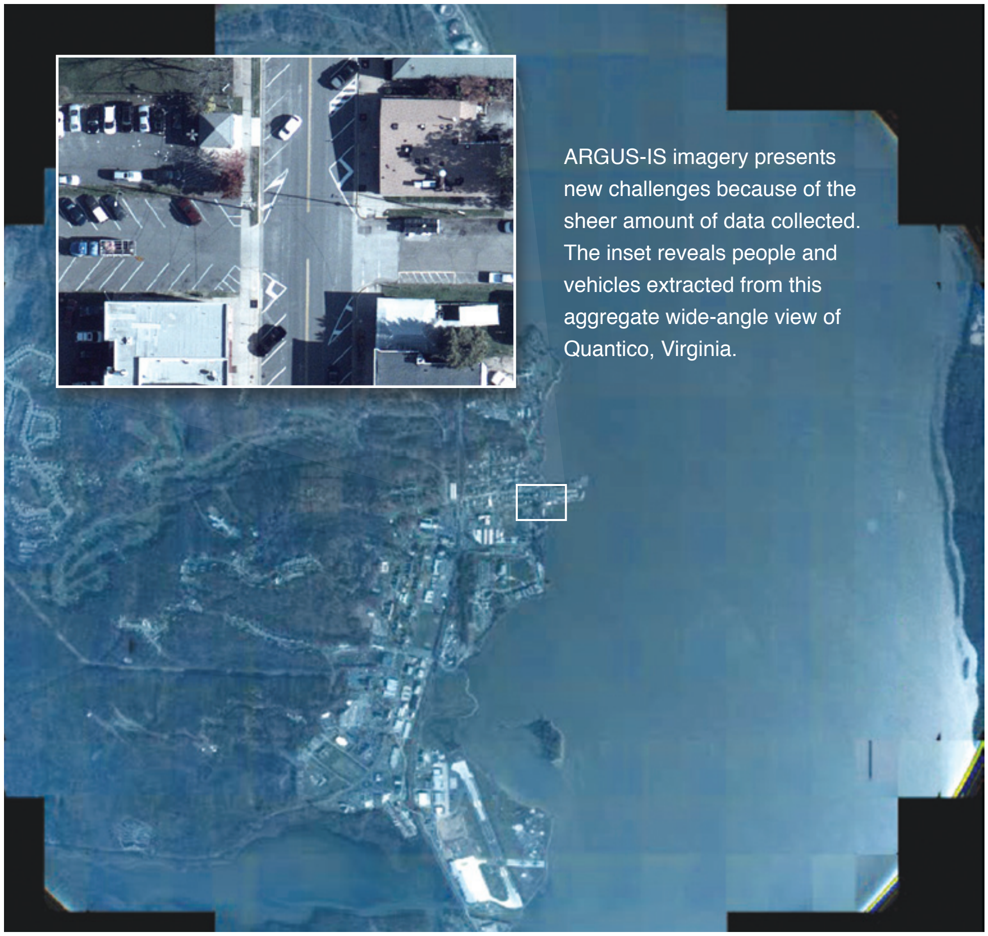

ARGUS-IS imagery presents new challenges because of the sheer amount of data collected. The inset reveals people and vehicles extracted from this aggregate wide-angle view of Quantico, Virginia.

information from visual data. (See *S&TR,* November 2004, pp. 12–19.)

Persistics Coming On Board

Although the current Persistics configuration is designed for ground-station processing, Livermore researchers are working to shrink the size, weight, and power requirements to permit mounting of the system onboard surveillance craft within two years. Developers are exploring the trade-off between data compression, image fidelity, and computing requirements to successfully evolve the next-generation architecture that will support onboard flight processing of massive data sets.

Persistics is being further enhanced to work with DARPA’s newest generation of real-time persistent surveillance capability called the Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System (ARGUS-IS). This system to detect and track events on battlefields and in urban areas can cover 100 square kilometers, a significantly greater aerial footprint than current systems. The ARGUS-IS video-collection rate of 12 hertz (frames per second) is considerably greater than previous systems, which operate at 2 hertz. ARGUS-IS comprises 368 cameras of about 5 million pixels each, identical to those used in cell phones. The cameras operate together behind four high-quality telescope lenses. In all, ARGUS-IS has 1.8 billion pixels compared with 4 million pixels in the first Sonoma sensor (the predecessor to ARGUS-IS) developed at Livermore in 2003, 176 million pixels in the most advanced Sonoma sensor developed in 2007, and 800 million pixels in the sensor developed in 2009 by the Massachusetts Institute of Technology's Lincoln Laboratory.

Persistics can simultaneously and continuously detect and track the motion of thousands of targets over the ARGUS-IS coverage area of 100 square kilometers. ARGUS-IS can generate several terabytes of data per minute, hundreds of times greater than previous-generation sensors. "Until now, we had no practical way to store that much data," says Jones. "With Persistics, we have an innovative method to compress the equivalent of thousands of hard drives to just a few drives."

Decades of Visualization Research

About 30 people are involved in the Persistics effort, including computer scientists, electrical engineers, optics experts, statisticians, and machine-learning specialists who draw meaning out of visual data, a core Livermore strength. The Laboratory has been conducting research in ISR for two decades. Previous efforts include designing advanced sensors and large, lightweight optics; analyzing space imagery; and inventing instruments to detect adversarial clandestine activities associated with weapons of mass destruction. In addition, Livermore computer scientists are developing methods to manage large volumes of visual data and automated techniques to extract meaning and make discoveries from that information for national security applications.

Among other accomplishments, researchers have pioneered methods to locate areas of interest in three-dimensional simulations of nuclear weapons physics generated by supercomputers performing trillions of operations per second. Laboratory scientists have also developed computer-automated techniques to search for microscopic flaws in glass optics used in the National Ignition Facility, the world's most energetic laser.

Many federal agencies are interested in Persistics applications, including the Department of Energy, Department of Defense, and Department of Homeland Security. Collaborations in ISR for nuclear safeguards and treaty monitoring support extend to international organizations such as the International Atomic Energy Agency and the Comprehensive Nuclear-Test-Ban Treaty Organization.

Over the past decade, Livermore researchers have studied the vexing issue of sensor output growth far outpacing human capabilities to ingest and react to the collected data in a timely manner. Early work, funded through the Laboratory Directed Research and Development Program, focused on developing software techniques that could first "quiet" the jittery video taken onboard a moving airborne vehicle buffeted by the atmosphere and then compress the data. The research effort grew into the Department of Energy project called Sonoma, which developed

Persistics researchers are developing the capability to produce three-dimensional images from overhead video data.

wide-area sensors for monitoring nuclear nonproliferation. The Sonoma Persistent Surveillance System featured wide-area views at high resolution, real-time onboard data processing, and high-performance visualization at the receiving end. In 2006, the Sonoma Project team received an R&D 100 Award for their innovation. (See *S&TR,* October 2006, pp. 4–5.)

"People in the Department of Defense got excited because they realized they could use the same techniques to look for terrorist activities," says Duchaineau. In 2005, the video-camera effort that began with Sonoma was passed to the Department of Defense, which developed Angel Fire for the Air Force and Marine Corps and Constant Hawk for the Army. ARGUS-IS, built by BAE Systems, Inc., is the newest imaging system.

During Sonoma's development, DARPA managers turned to Livermore to explore avenues for reducing the massive volume of data collected by these new sensors without sacrificing image quality. Vaidya recalls, "Our inability to transport the data to ground in an efficient manner had become a bottleneck." Persistics provides not only an efficient means of stabilizing and compressing video to transport it across limited bandwidth channels but also back-end anomaly detection and behavior analysis to differentiate between normal and abnormal patterns of behavior.

Fast Forward or Rewind

Persistics video processing is event-driven, meaning its algorithms are designed to detect in real-time potentially important events in the streaming video data. The pipeline can detect moving people and vehicles and track them over many hours. Automatic recognition of certain scenarios can provide early warning of a threat and greatly strengthen tactical combat support. For example, allied forces could intercept terrorists in the process of planting an improvised explosive device or setting up for an allied base attack. Livermore scientists are also developing automated methods for identifying patterns of behavior that could indicate deviations from normal social and cultural patterns as well as networks of subversive activity.

At the other end of the spectrum, Persistics supports forensic analyses. Should an event such as a terrorist attack occur, the archival imagery of the public space could be reviewed to determine important details such as the moment a bomb was placed or when a suspect cased the targeted area. With sufficiently high-resolution imagery, a law-enforcement or military user could one day zoom in on an individual face in a heavily populated urban environment, thus identifying the attacker.

Looking Ahead

Livermore researchers continue to optimize the Persistics platform so it can be carried aloft. That effort should be made easier this year because the size of GPUs is expected to shrink yet again, with power consumption reduced in half by late 2011. "As we ride the GPU technology wave, we should have the system up in the air by 2013," says Vaidya. "We plan to stream to ground only the information that is pertinent to the end user, which may require a synergy of multiple sensors in the air to provide robust estimates of what's important."

Persistics technology will soon be made available as open source to the government and subcontractors. The technology is modular to allow future plug and play, as Livermore scientists develop additional automated techniques. For example, they are researching ways to make possible the three-dimensional viewing of targets, which could further enhance data compressibility. They are also exploring methods to overlay multiple sensor inputs—including infrared, radar, and visual data—and then merge data to obtain a multilayered assessment.

Vaidya notes that unmanned aircraft have demonstrated their ISR value for years in Afghanistan and Iraq. As U.S. soldiers return home, the role of overhead video imagery aided by Persistics technology is expected to increase. Persistics could also support missions at home, such as monitoring security at U.S. borders or guarding ports and energy production facilities. Clearly, with Persistics, video means knowledge—and strengthened national security.

—Arnie Heller

Key Words: Angel Fire; Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System (ARGUS-IS); Constant Hawk; graphics processing unit (GPU); intelligence, surveillance, and reconnaissance (ISR); Persistics; pixel; Predator; Reaper; Sonoma.

For further information contact Sheila Vaidya (925) 423-5428 (vaidya1@llnl.gov).