

Edge Computing for Situational Awareness

Mahadev Satyanarayanan
School of Computer Science
Carnegie Mellon University

Abstract—Situational awareness involves the timely acquisition of knowledge about real-world events, distillation of those events into higher-level conceptual constructs, and their synthesis into a coherent context-sensitive view. We explore how convergent trends in video sensing, crowd sourcing and edge computing can be harnessed to create a shared real-time information system for situational awareness in vehicular systems that span driverless and drivered vehicles.

I. FOCUSED ATTENTION IN A HECTIC WORLD

Situational awareness is defined as “up-to-the-minute cognizance or awareness required to move about, operate equipment, or maintain a system” [1]. Events in the world around us can be perceived and comprehended in many different ways, and from many different viewpoints. Some of these events may be directly relevant to one’s current mission; other events may be irrelevant. We use the term “mission” here broadly, as “a specific task with which a person or a group is charged” [2]. To someone driving home from work in a snowstorm, knowledge of which roads have recently been plowed is extremely valuable. During the same snowstorm, knowledge of road conditions is irrelevant to someone who is taking the metro home. However, if the metro rider uses a wheelchair, it would help to know in advance that the elevator at his regular metro stop is not operational. **This paper focuses on the creation and maintenance of a system-wide real-time knowledge base from which many context-sensitive and mission-specific worldviews can be extracted.**

Situational awareness involves the timely acquisition of knowledge about real-world events, distillation of those events into higher-level conceptual constructs, and their synthesis into a coherent holistic view that is specific to one’s mission. In this paper, we focus on situational awareness as it applies to the residents and administrators of metropolitan areas. These are areas that may range in size from small compact cities to sprawling county-sized entities. The convergence of three technology trends offers the potential to greatly enhance situational awareness in such areas.

The first trend is the growing deployment and acceptance of *always-on video cameras in public spaces*. A 2013 survey in the U.K. estimated one surveillance camera in a public space for every 11 people [3]. By 2012, virtually every automobile in Russia had a video camera on its dashboard to record incidents for insurance purposes [4], [5]. Extrapolating from these trends, the report of the 2013 *NSF Workshop on Future Directions in Wireless Networking* [6] predicts that “It will soon be possible to find a camera on every human body, in every room, on every street, and in every vehicle.”

The second trend is the *growing social acceptance of crowd sourcing* for acquiring real-world information. Yelp, TripAdvisor and Waze are three examples of widely accepted forums for crowd-sourced information on restaurants, travel services, and traffic conditions respectively. While work continues on optimizing incentive structures for attracting high-quality crowd-sourced information, it is clear that successful business models can be built on crowd sourcing.

The third trend is *real-time video analytics using edge computing* [7]–[9]. There is growing recognition that live video offers several advantages relative to other sensing modalities. Most important is its flexibility and open-endedness: new image and video processing algorithms can be developed to enhance the information extracted from an already-deployed video camera. Additionally, video offers high resolution, wide coverage, and low cost relative to embedded sensors. **A critical requirement for scalability is that analytics be performed close to the point of capture.** Shipping video to the cloud from myriad cameras places excessive bandwidth stress on the ingress networks of a metropolitan area. **The solution is to perform the video analytics on dispersed elements called cloudlets** [10] that have wired or wireless LAN connectivity to associated cameras. We assume that each cloudlet has sufficient compute power and hardware accelerators to perform real-time video analytics on all its associated cameras. It also possesses ample storage to preserve video at full fidelity for a significant retention period before being overwritten.

We explore how these convergent trends can be harnessed to create a shared real-time information system upon which mission-specific software can be built to provide enhanced situational awareness. Since situational awareness is primarily about support for human decision making, “real-time” is on the order of seconds to tens of seconds. In the extreme case of immersive decision-making based on virtual or augmented reality, end-to-end latencies as low as a few tens of milliseconds may be involved. We focus on vehicular use cases because they are likely to offer the highest payoff in the shortest time for the ideas expressed here. However, these concepts have much broader relevance. With modifications, they also apply to pedestrians, bicyclists and others who are in wireless contact with cloudlets. They can also be used in the context of drones, delivery robots, and other autonomous vehicles.

II. HISTORICAL ROOTS

The inspiration for our system architecture comes from work that is over 75 years old. During the Battle of Britain in 1940, the superior situational awareness of the Royal




Fig. 1. RAF Uxbridge Control Room During the Battle of Britain (1940) (Source: Imperial War Museums, UK)

Air Force (RAF) proved crucial to victory. The ability to allocate just the right quantity of very scarce resources (pilots and aircraft) just in time to precisely the right places was priceless. In far less dire circumstances, the administrators of every city, county, or other metropolitan area face the challenge of optimally directing scarce resources (salt trucks, fire trucks, police officers, etc.). The circumstances are even more challenging during recovery after a major disaster (e.g., earthquake, flood, terrorist attack, etc.)

The RAF's situational awareness system was based on a new type of sensor, *radar*, that had just been invented. While primitive by today's standards, the radar of 1940 was capable enough to make the crucial difference to survival. However, by itself, a single radar station could only offer a tiny sliver of knowledge about the pattern of action evolving in real time. This high-level information (e.g. number of aircraft, heading, altitude) had to be combined with similar inputs from many other radar stations to form a composite picture of the unfolding threat. This was further combined with static map information as well dynamic information pertaining to the state of readiness and engagement of RAF resources. The resulting synthesis provided decision makers with the situational awareness necessary for command and control.

Figure 1 illustrates how static map information was combined with dynamic annotations obtained through real-time input streams from edge analytics. In the technology of 1940, "edge analytics" consisted of human operators interpreting radar signals at a multitude of stations. Notice that the analysis of sensor inputs (i.e., radar signals) happened at the edges. Only the distilled, and hence much lower bandwidth, information from the edge analytics were transmitted (via human telephony) to the control room shown in Figure 1. Human operators performed updates to the crucial map-based data structure shown in Figure 1, which represented the synthesis and visualization of current mission-specific knowledge. A subset of this data structure, scoped to the appropriate coverage zone, was also available to lower echelons of command. This accurate and highly agile shared information system enabled optimal decision making at every level of command.

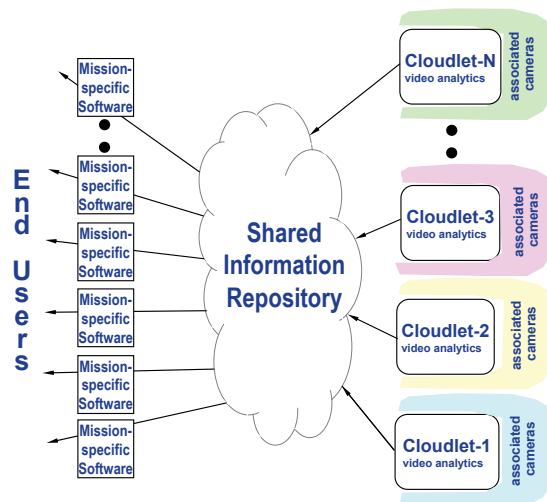


Fig. 2. System Architecture for Situational Awareness

Fast forward to 2017. Figure 2 illustrates a system architecture inspired by Figure 1. The entire room and people shown in Figure 1 are replaced by a smartphone screen (for a mobile user) or a wall-sized screen (for an administrator in an office). The software driving each screen is mission-specific. It obtains its real-time inputs from a common shared information system, but analyzes, filters and customizes its output to be mission-specific. The real-time inputs to the shared information system are provided by video analytics running on cloudlets and performing real-time processing of input streams from video cameras on vehicles, people, buildings, and so on.

The architecture shown in Figure 2 is extensible because mission-specific video analytics code can be dynamically inserted into cloudlets. For example, if a parent wishes to follow her child as he makes his way about a city, code to recognize that child's face can be dynamically installed on relevant cloudlets. As the child comes into view of different cameras, and is recognized, the cloudlet transmits details of the sighting (possibly with a short video clip) to the shared repository. Access controls on this data can restrict it to the parents and a small set of trusted friends and relatives. This architecture could also be generalized to support a wide range of multi-modal inputs, of which video is just one modality.

III. MAPS: CURATED SPATIAL KNOWLEDGE

For a mobile user, a substantial part of situational awareness is captured in spatial knowledge that is relatively static. For example, knowing that the road narrows sharply after the next bend and that there is a steep uphill climb beyond that, is valuable to someone who is approaching that bend in a car or bicycle. This information only changes rarely (e.g., the road is widened), typically over timescales of months or years. Other components of spatial information are much more dynamic. For example, the presence of dense fog around the next bend is valuable information, but its lifetime is likely to be measured in minutes or hours. Fallen rocks or tree branches, temporary lane closures, traffic congestion, icy road conditions, and many other phenomena relevant to situational awareness have short lifetimes, on the order of a few minutes to a few hours.

It is standard practice to separate the large static component of spatial information into the well-known form of encoding known as a *map*. This process of separation necessarily involves considerable abstraction, and is best viewed as an act of *curation* by the map-maker. For example, two high-resolution satellite images of the same area taken just a few minutes apart may differ significantly at the pixel level due to the motion of individuals and vehicles, or due to the change in angle of reflection of sunlight from surfaces. Yet, a map of the same area does not change at that timescale. The curation of maps from raw pixel data can be significantly automated, but still involves human mediation at least in the final stages for sanity checks and certification. From the viewpoint of distributed systems technology, a map may be viewed as cloud-sourced read-mostly data that is rarely updated. Standard eventual-consistency techniques can be used for caching and prefetching, with little runtime cost or implementation complexity required for maintaining consistency.

The dynamic component of spatial knowledge is typically structured as an *overlay* on top of a map. Overlays can be compact, because they annotate the much larger volume of static information in the underlying map. For example, to show that a segment of a road is congested, one only needs to convey the road ID, the road coordinates between which the congestion information applies, and an indication of the degree of congestion. This could be encoded in 100 bytes or less. None of the complexity of the twists and turns of the road over its length, its changing gradient, details such as number of lanes, or topology of the surrounding areas need to be conveyed in the overlay. That voluminous information, whose encoding may be many megabytes in size, remains unchanged.

Maps can be combined with mission-specific overlays to provide situational awareness in a way that is highly customized, yet efficient in storage, transmission and processing. In the rest of our discussion, we therefore use maps and overlays as the organizing data structures for situational awareness.

IV. DRIVERLESS AND DRIVERED VEHICLES

A. Proactive versus Reactive Actions

Research on driverless vehicles is nearing commercialization by companies such as Google and Uber. An important lesson that has been learned is that extremely accurate maps of high resolution are crucial to the success of driverless vehicles. Real-time sensing (for example, based on computer vision or radar) complements map-based knowledge. The sensing is essential to providing important details (such as new obstacles on the road) that are missing from the map. However, a map that has already been updated to reflect those obstacles would be even better. It would enable proactive actions, rather than reactive just-in-time actions. For example, moving to a different lane well in advance of an obstacle is safer than last-minute detection and avoidance. In other words, superior situational awareness can produce improved responses to real-world events and thus lead to better outcomes. Knowing that there is poor visibility around the next curve of a windy

mountain road, or that deer have recently been seen crossing the road can prepare you better for those hazards.

Today, driverless vehicles treat knowledge from sensing as a completely separate channel of knowledge from map information. There is an opportunity to close the loop: i.e., to automatically update the shared information repository in Figure 2 as a side effect of real-time sensing for reactive control. By doing this in a crowd-sourced manner, every driverless vehicle becomes both a real-time producer and a real-time consumer of information in that repository.

B. Human Inputs Plus Edge Analytics

Human-in-the-loop crowd-sourced systems such as Waze already exist. Our proposed approach would reduce the cognitive load on humans in such a system. They would still be welcome to offer manual inputs, but we expect that most updates would be generated automatically through edge analytics. Each vehicle would emit a symbolic content stream that compactly encodes sensed information. In some cases, images or short video snips may annotate the symbolic encoding. Once a new sensed observation is incorporated into a map and distributed to vehicles, it can be omitted from further reporting by vehicles. This can be viewed as a form of deduplication that reduces the volume of total transmitted data. Periodic reconfirmation of short-lived observations may be needed: e.g., “Is the dead animal still in the left lane?” may result in the response “No, the left lane is now clear.” There may also be value in getting human confirmation of automatic updates, especially to resolve conflicting updates from multiple sources. This can be implemented as a hands-free, speech-based system in which a driver is asked to verbally confirm an observation that should be currently visible to him.

C. Human Interactions

Improved situational awareness is clearly valuable to human drivers in conventional drivered vehicles. However, there are many open questions regarding how best to present the information to the driver. Should new information be presented through a synthesized voice, much like a GPS navigation system today? Should it just be a visual annotation on a display built into the vehicle or separately carried by the user? Should it use the new visualization capabilities made possible by “smart windshields” that allow augmentation of the scene in front of the vehicle [11]? Should a user wear a smart helmet or head-up display to benefit [12]? With human drivers, these HCI questions about *how* to present the information, become as important as the deeper challenge of *what* information to present. The latter component is common to both driverless and drivered vehicles.

Note that these HCI challenges do not affect the ability of a drivered vehicle to *contribute* data. If a drivered vehicle is equipped with sensors (e.g. video cameras, accelerometers, GPS location sensing, etc.) and the computational capability to perform edge analytics (i.e., an on-board cloudlet), it can contribute to map updates without involvement of the human driver. The incentive structure to encourage such involvement

needs to be worked out, but there are no technical impediments beyond those faced by driverless vehicles.

D. Beyond Automobiles

These concepts generalize to use cases beyond terrestrial vehicles. For example, it applies equally well to the aerial context where drones co-exist with piloted aircraft. On-board edge analytics of sensor data can provide map updates from both types of sources; both can benefit from the improved situational awareness obtained through real-time map updates. In some use cases, such as a small drone, the mobile entity may be too small to host a cloudlet powerful enough to completely perform all necessary real-time processing. This will require development of strategies that partition the processing between the mobile entity and a static cloudlet. There will be difficult bandwidth and scalability challenges that result from such partitioning. Intelligent onboard preprocessing and data sampling will be necessary requirements to address these challenges. Our vision, as expressed in Figure 2, is a single software architecture with rich interfaces that provide ample scope for context-specific customization through parameterization and plug-in software modules for customized edge analytics. The deployment of such an architecture could greatly improve situational awareness in a wide range of use cases.

V. SYSTEM ARCHITECTURE

A. Scaling, Coverage and Resolution

One of the first questions to be answered is how large a coverage zone is desired. It is within this coverage zone that real-time inputs are collected, interpreted, fused, and distributed. The granularity and resolution of detail has to be fine enough to influence the actions of driverless vehicles. End-to-end network latency and scalability are important considerations in sizing coverage zones. It is hard to see how to create a single coverage zone that spans the entire continental United States. What appears more feasible is a federation of many coverage zones that are each much smaller. Across that federation, the timeliness and granularity of knowledge propagation may be significantly poorer than within a single coverage zone. Observers outside a coverage zone can “zoom in” at fine granularity to details within it, but there will be inevitable lag in seeing updates. Our intuition, which needs to be validated in real implementations, is that a county-sized coverage area is likely to be what we can handle today. This is roughly 500 square miles, using the smaller average county sizes of the eastern US. It is a substantial coverage area for situational awareness at the fine granularity of small potholes, small rocks, icy spots on roads, stalled vehicles, dead animals, etc. As vehicles move, they will eventually cross from one zone to another. That transition will trigger a handoff. There are many open questions surrounding the design and implementation of such a cross-zone handoff mechanism.

Independent of latency and scalability considerations, there are sound non-technical reasons why a global situational awareness architecture is always likely to remain a federation of quasi-independent zones. In particular, there are compelling

national security reasons to disallow very fine grain real-time knowledge at street level to be visible outside a country. Only coarser-grain or stale knowledge may be allowed. These reasons are even more compelling in military use cases — control and restriction of knowledge from within a zone to entities outside the zone is an obvious requirement for operational security. Even within a zone, there may be restrictions. For example, all vehicles may be allowed to report observations, but only low-fidelity information may be released to vehicles whose occupants lack appropriate security privileges. In other words, situational awareness may be degraded for security reasons. Using a unit of local government such as a county for a situational awareness zone is an approach that aligns naturally with boundaries of administrative trust and responsibility.

B. Aggregation and Access Control

For each coverage zone, we envision a single logical entity that is responsible for collecting inputs from all the driverless and drivered entities in that zone. Such a single point of synthesis provides a clear point of control for curation of data and for enforcement of security and privacy policies. The alternative approach of decentralized collection is technically feasible, but is harder to implement and control. For these reasons, it is likely that each coverage zone will have a single point of collection and synthesis.

We refer to this single point as the *zone cloudlet* and envision it being a modest data center with ample compute and storage resources. Although a single point of control raises concerns about failure resiliency, there are well-understood replication techniques that can be leveraged to alleviate this concern. Which replication technique to use, and how best to implement failure resiliency are open issues at this point. The zone cloudlet has an important role to play in controlling access to situational awareness information. It is the guardian of knowledge about its zone, and hosts the mechanisms for enforcing security and administrative policies.

C. On-board Processing and Storage

Each participating vehicle is equipped with a *vehicular cloudlet* that has substantial processing capability and local storage. This cloudlet performs edge analytics on external sensor readings (e.g. video cameras, possibly multiple per vehicle) and internal sensor readings (such as speed, engine performance parameters, occupant alertness, etc.). These edge analytics transform the high data rate of raw sensor data into a data stream of much lower bandwidth to the zone cloudlet.

Figure 4 shows the interactions between a vehicular cloudlet and the zone cloudlet with which it is currently associated. In typical usage, most of the interactions will consist of event reports (①) from the vehicular cloudlet to the zone cloudlet. From time to time, the zone cloudlet may explicitly request more information or ask for confirmation of an observation from another vehicle (②). Each vehicular cloudlet caches data from its zone cloudlet. Within a single zone, it is reasonable to expect strict cache consistency across all connected vehicles and the zone cloudlet. The details of the scalable cache

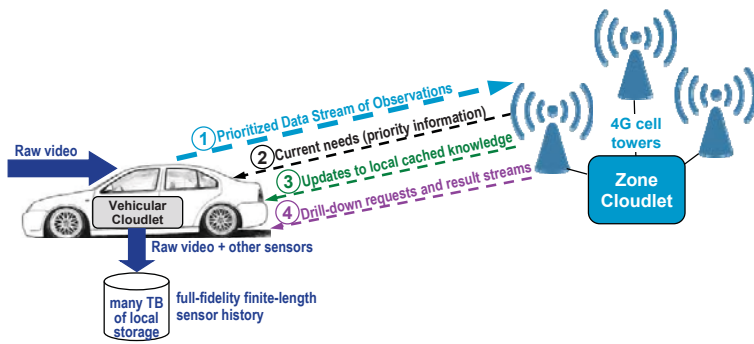


Fig. 3. Interactions Between Vehicle Cloudlet and Zone Cloudlet


consistency protocol need to be worked out, but it is likely to include notifications from the zone cloudlet (③). From time to time, *ad hoc* queries may be presented by the zone cloudlet to the vehicular cloudlet (④). These are human-originated queries that are routed through zone cloudlets to an appropriately scoped subset of vehicles: e.g., a help request such as “Does any vehicle that was in zone X at time Y have a video frame in which this lost dog appears?”. An authorization mechanism and policy to determine who can present such requests and how they are routed to specific vehicles will be needed in a real-world implementation.

The ability to respond to ad hoc requests relies on medium-term storage of full-fidelity sensor data on vehicular cloudlets. Netflix’s storage estimate of 3 GB per hour of HD video suggests that a 3 TB disk (less than \$100 today) can retain nearly 6 weeks of data at full fidelity from a single video sensor. Even with multiple cameras and other sensors, full-fidelity retention periods of a day or more are quite feasible. Such full-fidelity data retention is valuable if a need arises later to re-process the data with fresh analytics, or to drill down for more details. These circumstances can often arise in use cases such as debugging/troubleshooting, forensics, law enforcement and public service.

D. Reducing Transmission Volume

For the foreseeable future, 4G LTE offers the most plausible wide-area Internet connectivity from a moving vehicle. The demand for this resource is intense, and its spectrum-limited supply is scarce [13]. Hence, it is important to be frugal in terms of wireless transmission. Both peak bandwidth demand and total volume of data transmitted should be minimized.

The challenge is to offer real-time awareness of events sensed by vehicles, while transmitting as few bits as possible. Having every vehicle report every observation would achieve the best responsiveness. However, it would also transmit an enormous amount of redundant information. In crowded areas, many vehicles may transmit a slight variant of the same update. A small amount of redundancy is desirable, serving as cross-validation across vehicles, but the payoff drops rapidly as redundancy increases. In contrast, there may be few recent reports from lightly-traveled roads. Since sensing in our context is purely opportunistic, it is essential to make the most of rare observations. Exploring alternative transmission control



Brighter colors indicate more complete/recent knowledge. Each vehicle caches this data structure, including its underlying data, using a strong consistency protocol. Darker shades of gray represent areas from which no reports have been received in a long time. As more closely-spaced reports are received from an area, the zone cloudlet changes its color to lighter shades of gray, through yellow and orange to red hot. Colors decay over time in the reverse direction if no further reports are received. The exact encoding of the heatmap data structure and the management of color transitions are implementation-specific.

Fig. 4. Heatmap Data Structure

strategies in this space will be important. Both centralized strategies (in which the zone cloudlet controls the reporting) and decentralized strategies need to be explored.

One possible decentralized strategy would be for each vehicle to base its probability of transmission on how well-informed the zone cloudlet already is about the vehicle’s current surroundings. This state of knowledge can be captured in a *heat map* data structure, as shown in Figure 4. The zone cloudlet maintains the master copy of this heatmap, and each vehicle caches those parts relevant to its current location. Although Figure 4 shows a rectangular grid for simplicity, its actual shape may be highly irregular to match local map topography. The cached copy of the heatmap at each vehicle can be used to modulate reporting. A vehicle entering a gray area knows that any reports from it will be of high value. In a red area, it can remain silent more often.

E. Known Unknowns and Unknown Unknowns

Knowledge at the zone cloudlet closely tracks reality, but it can never be perfect. Important events may have happened recently, but there may be no vehicle nearby to report them or their lingering consequences. Delays of many seconds, possibly stretching to many minutes or tens of minutes, may occur in obtaining reports from isolated areas. This constitutes a *known unknown* because the zone cloudlet can be aware of its lack of knowledge. This is easiest to see in the context of a heatmap. Awareness of known unknowns can be used in modulating reporting by vehicles. If routing is centralized (e.g., for military or industrial vehicles), a zone cloudlet can direct vehicles to preferentially traverse gray zones to dispel darkness. The ideal would be red-hot grid squares everywhere, yet maintained with lowest transmitted bytes and few deviations from optimal routes.

A deeper source of ignorance is end-to-end system latency. The state of the shared information repository in Figure 2 will always lag reality by at least an amount equal to this quantity. It includes network transmission delays, as well as queueing and processing delays at vehicle and zone cloudlets. These can be significant under conditions of high load, when each of many vehicles reports a burst of serious events. The peak processing demand from video analytics can be especially challenging. Unfortunately, events in the real world tend to be correlated (e.g., a chain reaction leading to many accidents).

Better algorithms, improved networks, and more powerful hardware at both ends can help, but there will always be room for improvement. It will always be the case that the most recent sensor observations will be available at a vehicle, but not yet known at the zone cloudlet.

From the viewpoint of the zone cloudlet, this gap in knowledge is an *unknown unknown*. Fortunately, a reporting vehicle can detect this situation and transform it into a “known unknown” by examining its cached data. A real-world observation that does not appear in the cached map data is likely to be an unknown unknown to the zone cloudlet (likely, but not certain, because some other vehicle may have reported the observation recently). The vehicle can prioritize the reporting of this event, possibly separating the occurrence of a serious event (transmitted and processed at highest priority) from full details such as a video segment of the event (arriving later, and processed at normal priority).

VI. FROM HIGH-LEVEL VISION TO REALITY

Situational awareness is a context-sensitive and personalized view of the world. A shared real-time information repository for situational awareness can be the unifying force that holds together diverse sensor-based real-time information producers and consumers. For the vehicular theme explored in this paper, there is already a convergence of technology availability, user demand, societal commitment, and commercial opportunity. For example, the Traffic21 initiative at Carnegie Mellon University [14] aims to “design, test, deploy and evaluate information and communications technology based solutions to address the problems facing the transportation system of the Pittsburgh region and the nation.” In 2016, the US Department of Transportation announced that Columbus, OH was the winner of its \$40 million Smart City Challenge. The goal of this effort is “to implement a holistic vision for how technology can help all residents to move more easily and to access opportunity” [15]. All over the world, city-scale entities are looking for ways in which their quality of life can be improved through sensing and wireless technologies.

Obviously, the path to real-world deployment will not be easy or short. There are major technical issues to be resolved and questions to be answered. Some of these issues have already been mentioned in the preceding sections of this paper. Others will come to light in the course of detailed design. A foundational assumption is that video analytics algorithms of sufficient speed, accuracy and descriptive power can be developed for reporting events in real-time. Today’s driverless vehicles demand the speed, but not necessarily the descriptive power required for event reporting. For example, it is only necessary to know that an obstacle is a few feet ahead in order to trigger accident avoidance. Reporting that obstacle, however, requires knowing whether it is a rock, a dead animal, a pothole or other type of obstacle. On the other hand, it may be acceptable to sample the video stream at low frame rate for purposes of reporting, thus relaxing the speed requirement. Significant research is needed to develop video analytics algorithms and necessary specialized hardware for

the appropriate combination of speed, accuracy and descriptive power. Companies such as RoadBotics [16] have already begun this task. Privacy is a major concern with continuous video capture in public spaces. Fortunately, recent work by Wang et al [9] has shown that video can be denatured in real time on cloudlets to preserve privacy. The denaturing is reversible under carefully controlled circumstances, to allow retrospective queries such as searching for a lost child.

In closing, cloudlet-based edge computing is the key to success. Only by processing video and other high data rate sensors close to the point of data capture can scalability be achieved. Otherwise, continuous uploading of video over 4G would pose daunting bandwidth challenges at large scale. Situational awareness thus emerges as a “killer application” for edge computing.

ACKNOWLEDGEMENTS

In writing this paper, I have benefited from discussions with Wenlu Hu, Ziqiang Feng and Padmanabhan Pillai. This work was supported by the National Science Foundation (NSF) under grant number CNS-1518865. Additional support was provided by Intel, Vodafone, Deutsche Telekom, Verizon, Google, Crown Castle and the Conklin Kistler family fund. Any opinions, findings, conclusions or recommendations expressed here are those of the author and should not be attributed to Carnegie Mellon University or the funding sources.

REFERENCES

- [1] R. Pew and A. M. (Editors), *Modeling Human and Organizational Behavior: Application to Military Simulations*. National Academy Press, 1998.
- [2] “Merriam-Webster Dictionary,” Accessed March 18, 2017, <https://www.merriam-webster.com/>.
- [3] D. Barrett, “One surveillance camera for every 11 people in Britain, says CCTV survey,” *Daily Telegraph*, July 10, 2013.
- [4] A. Davies, “Here’s Why So Many Crazy Russian Car Crashes Are Caught On Camera,” *Business Insider*, December 2012.
- [5] D. Lavrinc, “Why Almost Everyone in Russia Has a DashCam,” *Wired*, February 15, 2013.
- [6] S. Banerjee and D. O. Wu, “Final report from the NSF Workshop on Future Directions in Wireless Networking,” National Science Foundation, November 2013.
- [7] M. Satyanarayanan, P. Simoens, Y. Xiao, P. Pillai, Z. Chen, K. Ha, W. Hu, and B. Amos, “Edge Analytics in the Internet of Things,” *IEEE Pervasive Computing*, vol. 14, no. 2, April-June 2015.
- [8] Simoens, P., Xiao, Y., Pillai, P., Chen, Z., Ha, K., Satyanarayanan, M., “Scalable Crowd-Sourcing of Video from Mobile Devices,” in *Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services (MobiSys 2013)*, Taipei, Taiwan, June 2013.
- [9] J. Wang, B. Amos, A. Das, P. Pillai, N. Sadeh, and M. Satyanarayanan, “A Scalable and Privacy-Aware IoT Service for Live Video Analytics,” in *Proceedings of ACM Multimedia Systems*, Taipei, Taiwan, June 2017.
- [10] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, “The Case for VM-Based Cloudlets in Mobile Computing,” *IEEE Pervasive Computing*, vol. 8, no. 4, 2009.
- [11] A. Moscaritolo, “Toyota Patents New Augmented Reality Windshield Tech,” *PC Magazine*, March 2016.
- [12] S. Tsantilas, “BMW shows motorcycle helmet head-up display and laser light concepts,” *New Atlas*, January 2016.
- [13] R. N. Clarke, “Expanding mobile wireless capacity: The challenges presented by technology and economics,” *Telecommunications Policy*, vol. 38, pp. 693–708, 2014.
- [14] Carnegie Mellon University, “Traffic21: A Smart Transportation Research Institute,” <http://traffic21.heinz.cmu.edu/>, Accessed on March 27, 2017.
- [15] Columbus, OH, “SmartColumbus,” <https://www.columbus.gov/smartcolumbus/>, Accessed on March 27, 2017.
- [16] “RoadBotics,” <https://www.roadbotics.com/>, Accessed on March 27, 2017.