2012 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
ROBOTIC SYSTEMS (RS) MINI-SYMPOSIUM
AUGUST 14-16, MICHIGAN

CONVOY ACTIVE SAFETY TECHNOLOGY – ENVIRONMENTAL UNDERSTANDING AND NAVIGATION WITH USE OF LOW COST SENSORS

David Simon
Lockheed Martin MFC, Grand Prairie, TX
Bernard Theisen
U.S. Army TARDEC, Warren, MI

ABSTRACT

This paper will document the development of the Convoy Active Safety Technology (CAST) program, which was created to design a low cost, optionally manned vehicle (OMV) solution for tactical wheeled vehicle (TWV) fleet. This paper will describe the approach taken to integrate low cost sensors for understanding the environment sufficiently to accomplish convoy missions. This paper will also discuss the approach taken to develop the low cost guidance and navigation solution used in the CAST program.

INTRODUCTION

The Convoy Active Safety Technology (CAST) development program sought to develop a low cost, optionally manned vehicle (OMV) solution. An objective of the CAST program was to overcome some of the barriers to transitioning autonomous capabilities out of the lab. First, the system would need to be relatively low cost, a fraction of the target platform cost. While there were systems with higher levels of autonomy, their cost increased the risk that they would not be transitioned. Second, transitioning technology would need to follow a model of crawl, walk, run. The services would need to see a system working well while under close human supervision in order to develop trust in automation. Lastly, the system could not substantially increase the training, maintenance and logistics burden of the vehicle system.

With these requirements, a low cost appliqué kit was developed that could be installed on any tactical wheeled vehicle (TWV). This paper will describe the approach taken to integrate low cost sensors for understanding the environment sufficiently to accomplish convoy missions. This paper will also discuss the approach taken to develop the low cost guidance and navigation solution used in the CAST program.

BACKGROUND

TARDEC began an effort in 2005 to develop an appliqué kit to provide autonomous control of tactical wheeled vehicles for eventual retrofit of existing vehicles in the fleet. This autonomy would provide a spectrum of selectable capabilities, from driver assist functions up to optionally manned vehicles (e.g., unmanned vehicles). The goal of the autonomy kit is to provide robotic assistance of the driving task, which will reduce the occurrence and severity of vehicular accidents, increase the tempo and effectiveness of missions, relieve the vehicle operator of the continuous driving task, and increase the occupants' situational awareness through scanning of the surrounding area not simply the road ahead.

Figure 1: CAST convoy during Warfighter Experiment.

A key realization in the development of this effort was that in order to economically justify equipping the Warfighter with this technology, the appliqué kit must be affordable as a fraction of the cost of the platform. While prior robotics efforts such as the Robotic Follower (RF) Advanced Technology Demonstration (ATD) achieved some measure of success on combat vehicles, the cost of the retrofit equipment (projected at $650k/system) was economically prohibitive for tactical vehicles.

TARDEC began the CAST program to develop a solution for convoy automation, with the hypothesis that by heavily utilizing Commercially available Off-The-Shelf (COTS) components, emphasizing software centric solutions, and utilizing an open development paradigm (i.e., eliminating custom and proprietary development), a low kit cost could be achieved. A production price point of $30k/system was set, which is a fraction of $150-$500k vehicle platform cost. TARDEC assembled a Government/Industry team to develop CAST and validate the cost and performance of the system through rigorous production cost studies and through rigorous test, demonstrations, and subsequent evaluation.

The resulting system has been demonstrated to be an effective autonomy solution for logistics missions. The system has been tested in five engineering evaluation and test (EET) events and three Warfighter Experiments, each conducted in relevant environments and conditions similar to current theaters of operation. During the Warfighter Experiments, Soldiers operated the CAST enabled systems following a brief training period. Training of CAST functions for the Military Occupational Specialties (MOS) 88M truck drivers requires less than an hour. A simple user interface with intuitive controls and failsafe “return to manual mode” interaction makes the CAST system user friendly, similar to cruise control functions on an automobile. The results of these objective tests show that the CAST system maintains near-human level cross-track error (a key parameter to ensure safe operations), and convoy interval spacing better than a human driver. The performance of the vehicles operating in a convoy is increased which allows the vehicle driver to focus on other tasks such as counter improvised explosive device (IED) detection.

SYSTEM DESCRIPTION

The CAST Kit is a complete set of hardware and software to retrofit an existing vehicle for convoy automation and when engaged, actuates steering, throttle, and braking. When installed, any vehicle can serve as either a leader or follower vehicle. In a convoy, the system acts on the follower vehicles to drive in the same path as the leader, while maintaining convoy mission speeds and longitudinal spacing between vehicles. The system also senses and avoids obstacles and reverts back to manual mode with any of several override actions by the operator. Safety has been paramount during the development of the CAST Kit solution.

The sensor suite of the Kit is comprised of a number of low cost COTS complementary sensors. Any single sensing modality exhibits limited functionality under certain conditions (e.g., Electro-Optical/Infra-Red (EO/IR) sensors in obscured atmospheric conditions), and is possibly limited in utility in a particular mission profile (e.g. sensors with active emissions cannot be used during stealthy missions). Through the use of an innovative software approach to dynamically and continuously fuse inputs from a variety of sensors, limitations of individual modalities are mitigated, and a diverse operating capability is achieved, all while basing the suite on low cost COTS items. Figure 2 highlights the CAST Kit as equipped on an FMTV, including all sensing, computing, and actuation components.

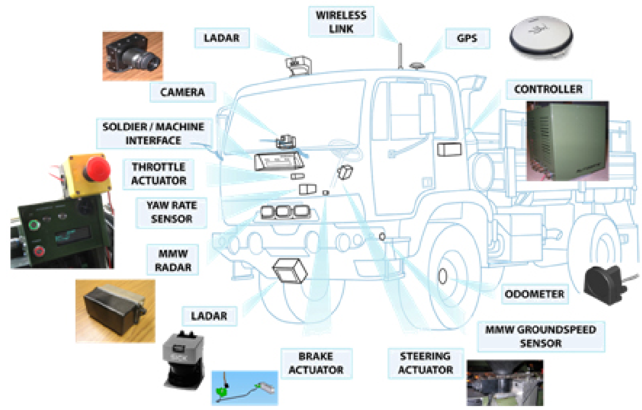

Figure 2: Low cost sensors and innovative algorithms provide an affordable optionally manned convoy solution.

LOW COST SENSOR FUSION

With the overriding requirement of a low cost system for CAST, it is necessary to select sensors that will cost, at volume, less than $1,000. The challenge with using these sensors is that they do not provide a sufficient solution on their own. A case in point is the vehicle following algorithm, Perceptive Vehicle Follower (PVF). The PVF module is responsible for detecting the lead vehicle and outputting a trajectory that the lead vehicle was observed to traverse. PVF can be thought of as three independent, sequential tasks executed in an endless loop:

1. Vehicle Detection - process data from each sensor to produce multiple estimates of immediate leader position.
2. Vehicle Tracking - combine all available information (e.g., sensor measurements from step 1, data from previous cycles, etc.) to produce a best estimate of current leader position.
3. Trajectory Generation - output a trajectory relative to our current position that will lead us to drive the same path that was driven by the immediate leader.

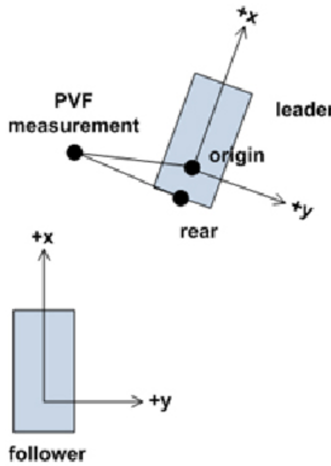

Figure 3: Coordinate frames and points of reference for PVF performance statistics.

VEHICLE DETECTION

The problems of using low cost sensors present themselves in each step. For Vehicle Detection, each sensor is plagued with false positives or environmental occlusions. In a recent performance upgrade of the system, a ground truth characterization of the PVF sensor was gathered. Figure 1 shows the relative position and measurement of the systems. The data was collected at Fort Carson and was post processed using the NovAtel RTK packages, which give sub-centimeter accuracy. Table 1 shows the sensors used in PVF and their individual sensor performance. Obviously no single sensor is sufficient for use with the idea that a follower will not exceed 0.5 meters of error in following the wheel tracks of the leader.

Sensor	Reliable Measurements	Reliable %	Unreliable Measurements	Unreliable %	Error Avg	Error Dev	Error Max
Camera	2766	16.10%	14416	83.90%	0.9	0.33	2.63
RADAR	7018	97.69%	166	2.31%	0.89	0.42	5.65
GPS	12880	74.49%	4412	25.51%	0.23	0.26	1.96
LVD	8476	93.38%	139	1.61%	0.96	0.32	2.41

Table 1: Measurement statistics show that no single sensor is sufficient for vehicle following.

VEHICLE TRACKING

After all sensor inputs have been collected, the most troublesome problem for PVF has always been the task of deciding which sensor inputs to trust, and which ones represented false positives. On path to take would be to develop heuristics to determine which sensors to believe. This logic makes for a system that is brittle and difficult to maintain or extend and result in unnecessary takeovers.

For example, the Figure 4 shows a case where the PVF estimate of leader position (the yellow line) is oscillating between estimates from two different sensors: the red trail of dots representing estimates based on camera image processing, and the magenta trail representing GPS position estimates. This was caused by arbitration between camera and GPS measurements. The net result was an output filtered trajectory (the white line) that was between the trajectories from the two individual sensors, with roughly a 0.5 meter bias relative to the true leader position.

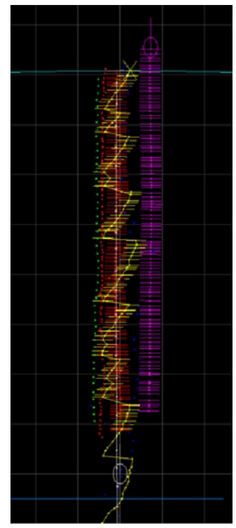

Figure 4: PVF output oscillating between two sensors.

A more robust approach is to implement a systematic method for combining an arbitrary number of generic sensor measurements and output a best estimate of the lead vehicle position. This approach is outlined below:

1. For each sensor measurement, initialize a kalman filter at each measurement's location. These kalman filters represent multiple hypotheses of the lead vehicle's position, speed, heading, and steering radius. On subsequent input cycles, update these hypotheses by executing the following steps.
2. Drop kalman filters that haven't been recently updated with new sensor data.
3. Compare each new sensor measurement with each existing kalman filters' prediction. This is done by calculating the mean between the measurement and the prediction, and then calculating the Mahalanobis distances between this mean and the measurement and prediction. If the distances are both within a configurable threshold, then the measurement is considered to "agree" with the prediction (ie, it's likely that the measurement and the kalman are both tracking the same physical object).
4. Update each kalman with each sensor measurement that "agrees" with it (as determined in step 2).

5. For each sensor measurement that did not "agree" with any kalman filters, initialize a new kalman at the measurements' locations.
6. Compare all possible pairs of kalmans (using the Mahalanobis distance metric described earlier) and consolidate kalmans that agree with each other (ie, they are probably tracking the same physical object).
7. Pick the kalman that most likely represents the leader position based on whether its history agrees qualitatively with the trajectory reported by the lead vehicle (see discussion of the TranRel paths in the Vehicle Detection section above), and based on the number of times it has been updated with new sensor measurements. This decision also includes some hysteresis so that once a kalman is selected; PVF will not switch to a different kalman on future cycles unless the original one becomes significantly inferior to another one.
8. Output the current "best" kalman state as the best estimate of leader position.

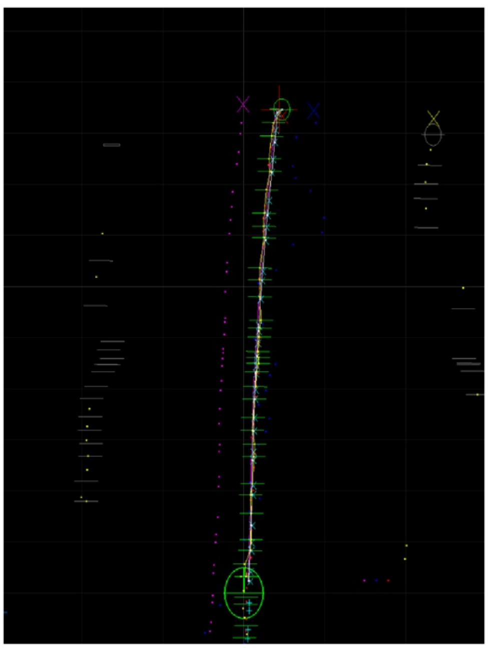

Figure 5: Ignoring false positives from camera.

Live testing has shown that this algorithm performs remarkably well. Instead of trying to decide up front which sensor inputs are reliable, PVF accepts all sensor measurements and assumes that the false positives will be implicitly ignored as they only cause short-lived Kalman filters to be added to the list of hypotheses of the leader position. The robustness of this approach is also due to utilizing the entire history of sensor inputs i.e. through the Kalman filter theory, and tracking multiple hypotheses

As applied to the previously mentioned scenario involving the zigzag trajectory, PVF will now correctly output a trajectory that follows the camera based estimate. If the GPS based estimates are sufficiently close (determined by the Mahalanobis Distance between their measurements), then they will be combined (through the Kalman filter framework) with the camera based estimates to increase the accuracy of the final position estimate. On the other hand, if they are far enough away (ie, sufficiently divergent based on the Mahalanobis Distance), then they will be fed into a separate Kalman filter, which will only come into play if other sensors drop out.

In Figure 5 the color camera (yellow trail of dots) was artificially induced to output sensor measurements that were roughly 2m to both the left and right of the actual leader position. Kalman filters (gray trail of hash marks) are seen to track both trails on the left and right since they represent somewhat consistent, and potentially valid, estimates of the leader's position. However, other sensors, GPS (magenta dots) and radar (blue dots) are judged to be inconsistent with the camera based hypotheses, so a third hypothesis (green trail of hash marks) is formed. It is this kalman filter that is judged to be most reliable estimate of the leader trajectory because it is founded on more sensor inputs, and the inputs are more consistent with the predicted motion of the lead vehicle. Therefore, this hypothesis is used to output the best estimate of leader position, and it is completely unaffected by the other inaccurate sensor inputs.

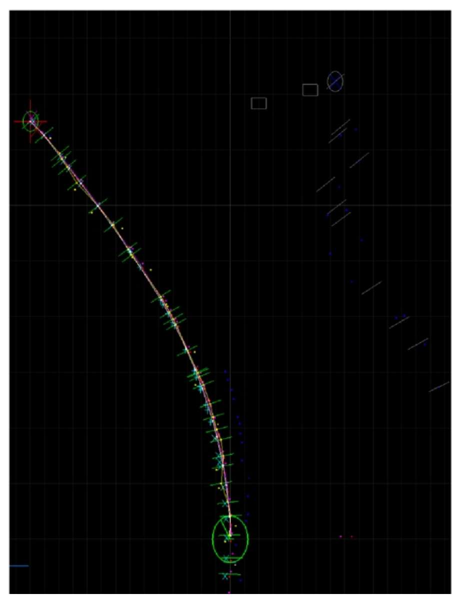

Figure 6: Ignoring false positives from radar.

Additionally, Figure 6 shows a similar problem that PVF often encounters. Since the radar has a narrow FOV, it will often lose sight of the lead vehicle on sharp bends and begin to report false positives. The figure above demonstrates that the radar (blue trail of dots at the bottom) is accurately contributing to PVF's estimate of the leader position until we approach a bend. As the leader leaves the radar's FOV, it suddenly begins reporting measurements far off to the right. The figure shows that several Kalmans (the gray circles and squares) are born to track the inconsistent output from the radar at this point. But PVF's best estimate of leader trajectory (lines leading to the green circle at the upper left) remains unaffected by this bad radar data as it follows other sensors that are more accurate and consistent with the predictive model.

TRAJECTORY GENERATION

The last step in the PVF cycle is to output a trajectory representing all past estimates of leader position. If PVF has a good estimate of the transform between the leader and follower TranRel frames (see discussion in the Vehicle Detection section), then it simply converts the leader's self-reported trajectory points into the follower's TranRel frame. Otherwise, PVF generates a trajectory by smoothing and filtering the path of previous leader position estimates i.e., "best" kalman states generated from sensor measurements.

Sensor	Reliable Measurements	Reliable %	Unreliable Measurements	Unreliable %	Error Average	Error Deviation	Error Max
Fiducials	2766	16.10%	14416	89.90%	0.455	0.183	4.1
Fiducials	16366	99.86%	23	0.14%	0.22	0.15	1.3
Radar	7018	97.69%	166	2.31%	1.152	0.287	3.262
Radar	15968	95.96%	672	4.04%	0.58	0.32	1.71
GPS	12880	74.49%	4412	25.51%	0.779	0.242	4.766
GPS	15718	99.98%	4	0.02%	0.47	0.23	2.39
LVD	8576	98.39%	139	1.61%	0.424	0.106	1.392
LVD	13331	96.70%	455	3.30%	0.25	0.11	1.27
Final	17125	100%	0	0%	0.23	0.15	1.44

Table 2: Measurement statistics show that no single sensor is sufficient for vehicle following.

PVF GENERAL RESULTS

At the beginning of this development effort, high-accuracy ground-truth data was collected to document and characterize PVF baseline performance. This data is presented in Table 2 as statistics on the average, standard deviation, and maximum errors of each individual sensor. These numbers are split into two categories: statistics on measurements that PVF judged to be unreliable (false positives) and measurements that PVF judged to be reliable (i.e., used to update the best estimate of the leader position). The statistics for each sensor are also split into the errors in the X and Y directions (longitudinal and lateral directions) with respect to the follower vehicle. This is of interest because errors in the longitudinal direction are less likely to cause problems during autonomous operation, whereas lateral errors in the estimates lead directly to cross-track error that might cause the follower to go off the road and result in a takeover.

After completing the PVF improvements outlined above, ground-truth data was again recorded on the same test course to capture the final performance of PVF. This data is represented in the highlighted rows in Table 2, and is interspersed with the baseline data for comparison.

The row labeled "Final" represents the total number of position estimates that PVF calculated from sensor inputs. These are marked 100% reliable because, by definition, only "reliable" measurements were used to update the Kalman filter that represented the best estimate of the leader position. Note that the data above should not be interpreted literally as an indication of sensor reliability. For example, the baseline data indicates that fiducials were only 16% reliable. However, this was a fluke mostly due to the issue described in the "Vehicle Tracking" section above where the camera fiducial based estimates were actually correct, but PVF was judging them to be unreliable after comparing them with inaccurate GPS measurements. So the data in this table is only useful in conjunction with the tables below to get an idea of how well PVF is distinguishing accurate measurements from false positives.

CONCLUSION

The Convoy Active Safety Technologies program has been a successful development program that has shown, through independent testing, the potential effectiveness of autonomous driving technologies in military convoys and provides a feasible solution to reduce task loading associated with military convoy driving. Today's vehicle operators are expected to drive, but also maintain situational awareness, attend to communication and navigation as well as respond to attacks. Addressing task load through the use of the CAST system provides a potential means of mitigating these demands on the driver and also improving overall convoy performance. The CAST system demonstrated its ability to augment convoy driving in a way that not only reduced task loading on the operator, but improved several aspects associated with convoy vehicle control for both day and night driving conditions. These improvements include convoy integrity by reducing the accordion effect inherent in convoys and providing faster responses to unanticipated stopping of the leading vehicle. The CAST system also demonstrated its utility by enhancing the operator's ability to perform critical local security tasks thereby enabling operators to improve upon survivability and attain mission accomplishment. Studies have also shown, an affordable autonomous system can be developed and produced. As Soldiers and commanders become more accustomed to driver assist or autonomous mode, the acceptance and utility of CAST like capabilities will become combat multipliers for the operational commander.