

# Autonomous Orchard Vehicles for Specialty Crops Production

**Bradley Hamner**

Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, bhamner@cmu.edu

**Marcel Bergerman**

Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, marcel@cmu.edu

**Sanjiv Singh**

Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, ssingh@cmu.edu

**Written for presentation at the  
2011 ASABE Annual International Meeting  
Sponsored by ASABE  
Gault House  
Louisville, Kentucky  
August 7 – 10, 2011**

**Abstract.** *Comprehensive Automation for Specialty Crops (CASC) is a four-year project to develop crop intelligence and agricultural automation technologies for the apple and nursery tree industries. Supported by the USDA Specialty Crop Research Initiative, CASC aims at impacting industry where it matters most: increasing production efficiency and reducing labor costs. In this paper we present our most recent results with the Autonomous Prime Movers, a family of reconfigurable vehicles designed to augment or automate a variety of orchard operations. We discuss the modifications made on the first APM to make it robust to uneven terrain and variance in canopy types, along with field tests that totaled 159 km in autonomous orchard driving. We also present the design and initial deployment of two new autonomous orchard platforms.*

**Keywords.** Specialty crops, Efficiency improvement, Labor cost reduction, Crop intelligence, Agricultural automation, Autonomous orchard vehicle, Autonomous orchard platform.

## Introduction

Comprehensive Automation for Specialty Crops (CASC) is a multi-institutional, multi-disciplinary initiative to address the needs of the specialty crops sector, with a focus on apples and nursery trees. In CASC we develop methods to improve production efficiency, identify threats from pests and diseases, and detect, monitor and respond to food safety hazards. At the end of this four-year project, we will achieve significant advances from the integration of robotics technology and plant science; understand the socio-economic barriers that prevent technology adoption and devise ways to overcome them; and make our results available to growers and stakeholders through a nationwide outreach program.

The project is organized into three themes that together address growers' needs for increased efficiency and labor cost reduction: Crop Intelligence, Automation, and Technology Adoption. Each thematic area is subdivided into several themes (Figure 1).

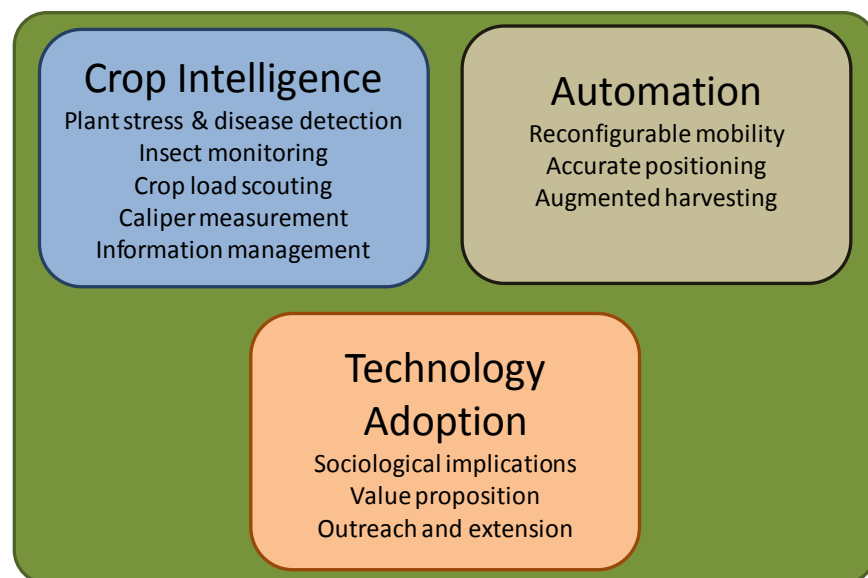


Figure 1. CASC's three main themes and their underlying thematic areas. CASC impacts growers by providing crop intelligence that increases farm efficiency and automation solutions that reduce production costs. All technologies developed are validated via thorough socio-economic analyses aimed at overcoming technology adoption barriers, and demonstrated in actual field conditions via outreach and extension activities.

- **Crop Intelligence** encompasses the work aimed at increasing the amount of information and corresponding level of detail growers have on their crops, especially crop load (for tree fruit) and caliper (for nursery trees); at providing early detection of stress due to water or nutrient deficiency and of disease such as fire blight; at monitoring insect populations and infestation; and at organizing the information in georeferenced databases for faster and improved decision making.
- **Automation** encompasses the work aimed at increasing farm efficiency and reducing production costs via the deployment of self-guided, low-cost agricultural machines to automate sensor data collection and farm operations such as spraying and mowing; and at increasing worker efficiency and reducing worker load via the deployment of pruning, thinning, and harvesting assist technologies.

- **Technology Adoption** encompasses the work aimed at understanding and overcoming the socio-economic barriers, perceived or real, that inhibit growers' incorporation of new technologies and methods; and the nationwide outreach and extension activities that we conduct to demonstrate, in actual field conditions, the technologies developed in the project and their applicability to growers in different industries and states.

In this paper we report on recent results obtained within the Reconfigurable Mobility thematic area. Here we are developing autonomous orchard vehicles and platforms that to date drove over 300 km in experimental and commercial orchards. We conclude with a discussion on the steps we are taking toward realizing this technology's full potential in the field.

## Autonomous Prime Movers

### *APM Deployment Roadmap*

At the core of our agricultural automation work is a family of vehicles we call the Autonomous Prime Movers, or APMs. The APM embodies the concept of a vehicle that can be rapidly tasked to support a variety of operations, including tree pruning and training, blossom and fruit thinning, fruit harvesting, mowing, spraying, sensing, etc. Vehicles in the family share a common sensing and computing architecture, and may differ in size, propulsion, carrying capacity, and applications that they enable (Figure 2).



Figure 2. Examples of vehicles in the APM family. On the left is Laurel<sup>1</sup>, based on a Toro MDE eWorkman chassis. The two laser rangefinders on the front are used to find the tree rows closest to the vehicle. The information is passed on to the onboard laptop in front of the passenger seat; it finds the center of the row and guides the vehicle along it. Steering and wheel encoders complement the sensor suite and provide data for the computer to localize the vehicle along the row. On the right is Nittany, based on an Italian platform from N. Blosi. Because their sensing and computing architectures are the same, we are able to transfer row following and other autonomy software from one to the other with minimal or no rework.

Successful commercial deployment of APM vehicles will be the result of a progression of steps aimed at proving the system's feasibility, incorporating robustness to the target environment and users, and making them easy to use and maintain—in addition to ensuring they are cost

---

<sup>1</sup> Currently the vehicles in the APM family are named after American mountains and mountain ranges. Laurel is named after the Laurel Highlands region in southwestern PA, and Nittany is named after Nittany Mountain in Centre County, PA.

effective. In our work, we associate quantitative success criteria with each step, and only proceed to a subsequent step when the previous one has been completed—i.e., when its success criteria have been reached. The criteria for Years 1 to 3 of our work in Reconfigurable Mobility are shown in Table 1. Note that the goals progress to the point where the vehicles are supervised by extension educators, growers, and farm workers, with the Carnegie Mellon engineering team hundreds or thousands of miles away.

Table 1. Success criteria progression for the work in the Reconfigurable Mobility area. Each project year runs from October 1 to September 30. Reaching these challenging goals is a necessary condition for successful technology commercialization.

Year	Goal	Success Criterion	Actual Result
1	Demonstrate vehicles feasibility (proof-of-concept)	100 km of autonomous orchard traversals with the vehicle supervised by the Carnegie Mellon engineering team	130 km
2	Make vehicles robust to canopy and terrain variation	100 km of autonomous orchard traversals with the vehicle supervised by the Carnegie Mellon engineering team, consisting only of runs over 10 km without operator intervention	159 km
3	Make vehicles usable by growers	100 km of autonomous orchard traversals with the vehicle supervised by extension educators, growers, and farm workers	15 km (as of 6/30/2011)

### ***Proof-of-Concept***

In Year 1 of the project (Oct 2008—Sep 2009) we focused on proof-of-concept. We equipped Laurel with sensors, onboard computers, and a software suite to enable the vehicle to drive between rows of trees and turn around at the end of rows (Hamner et al., 2010; Singh et al., 2010). We also showed how such vehicle could be used to tow a mower or carry a sprayer and thus execute typical orchard maintenance operations. By the end of Year 1 the vehicle had driven 130 km in both experimental and commercial orchards in Wexford and Biglerville, PA, and Rock Island and Royal City, WA. Typical runs varied between 4 and 6 km, with the longest run topping 11.8 km at Penn State’s Fruit Research and Extension Center (FREC).

At that time Laurel’s autonomous system was functional, but not reliable enough to undertake long missions. In particular it was not robust to uneven terrain and variance in canopy types. While we met our 100 km distance goal for Year 1, we were frequently restricted to blocks or rows which were neatly maintained or on relatively level terrain. Our new focus was then to improve the reliability of the system and its robustness to unexpected variations. We aimed for longer missions, more variety in the canopy and terrain of our test sites, and fewer failures. Our goal for Year 2 was again to drive 100 km autonomously, except that now with a mean distance between system failures of greater than 10 km (in other words, we did not accrue kilometers if the distance traveled on a certain test was less than 10 km). Not written explicitly into the goals was that we would achieve this by driving entire blocks, not picking and choosing rows that were “nicer” for the autonomous system.

### ***Improving Row Detection***

Laurel’s original design had two laser rangefinders located on the corners of the vehicle about one foot off the ground. The lasers scan 180° in a horizontal plane and provide distance to objects in front of them. We oriented the lasers to give a 270° field-of-view in front of and to the

sides of the vehicle. This is a simple, low-cost setup which allows the autonomous system to see the world around it, though only the objects that are one foot high relative to the vehicle. This includes not only tree trunks and canopy, but weeds and sometimes the ground itself when the terrain slopes upwards (Figure 3, left). To the autonomous system such obstacles are indistinguishable from trees and make row detection difficult.

To overcome this limitation, we modified Laurel to use only one laser rangefinder for row following. This is a new model of sensor which provides a  $270^\circ$  field-of-view. We mounted the sensor centered left-right and at a height just above the vehicle's hood (Figure 3, right). This configuration has two advantages. First, because the laser is above the hood we can take advantage of its increased field-of-view. Second, the sensor at this height is less susceptible to the spurious obstacles mentioned above.




Figure 3. (Left) Typical orchard map produced by the laser rangefinders. The vehicle is represented by the gray rectangle, with the red part indicating the front. The white dots are the rows of trees or other objects within the lasers' field-of-view (weeds, terrain, etc.) The yellow and blue lines represent the best reconstruction of the tree rows given the laser data; note how slopped terrain to the right of the vehicle introduces error in the tree row estimation, which in turn leads to inconsistent row following. (Right) Laurel's new configuration. The laser rangefinder mounted just above the hood (yellow circle) is the only one used for row following.

Our first row detection system used a Hough transform (Duda and Hart, 1972), which looks for the most likely pair of parallel lines in the laser range data. In normal circumstances, when the robot only sees the tree trunks and canopy, the two lines of tree canopy are the most likely. The autonomous system then finds the midway line between the two, and that is selected as the desired path to drive. When, however, spurious obstacles are present, the system may perceive them as being part of a good pair of parallel lines, which results in the vehicle drifting off center, sometimes into the canopy (Figure 4, left). Although the new laser height mitigated this problem by presenting fewer such obstacles, we required a row detection system that was more robust.

We focused then on developing a row detector that uses a particle filter (Dellaert et al., 1999; Gustafsson et al., 2002), a common estimation tool used in robotics. The particle filter makes multiple guesses of where the tree rows could be, and scores each guess by how much it agrees with the laser data. Furthermore, high-scoring row lines are kept from one iteration of the detection to the next, so that detections get better over time. When spurious obstacles appear, the filter remembers the previous row lines that had been detected, and can select the correct row (Figure 4, right).




Figure 4. Comparison of the performance of two row detection methods among tall weeds or uneven terrain. (Left) The Hough transform method used initially is confused by the unexpected data, resulting in bad row detections. (Right) The new particle filter method correctly ignores the additional data and finds the tree rows.

### ***Improving Turning and Row Entry***

When the robot reaches the end of one row it attempts to turn into the next row. Due to inconsistencies in tree plantings and in the robot's location estimate, the system does not rely solely on the prescribed row width for its target position. Instead, as it is turning, it runs the row detection to find the lane to go into. The method of turning used in Year 1 was to make a sharp turn towards the next row, then run row detection when the robot was perpendicular to the row. This was suboptimal in a couple of ways. First, the sharp turns were hard on the motor controlling the steering wheel, which may have led to a motor failure sustained in June 2009. Second, row detection was difficult with the vehicle pointed perpendicular to the row. From this point, the canopy of the nearest trees can block the far row, giving the autonomous system little data to work with (Figure 5). Detections from this point were unreliable, with a success rate less than 50% in dense canopy.

We designed new turn methods to make row entry smoother and more reliable. We implemented a path planner that takes the nominal starting point of the vehicle in the next row and generates a smooth path that respects the vehicle's steering constraints (Ferguson, 2010). The planner optimizes the path to avoid hard steering angles that stress the motor, while minimizing the total path distance (Figure 6, left). This planner aligns the vehicle with the row as it enters, so the vehicle is pointed along the row early. Before the vehicle completes the turn the autonomous system runs a row detection to refine its target entry point (Figure 6, center).

Frequently, in order for the vehicle to be aligned with the row before entering, the planner creates a large-radius, bulb-shaped turn. This turn is not always possible due to space constraints. Therefore we have also developed a three-point, "K" turn. The vehicle first turns past the row, then backs up and points towards the row, and finally enters the row (Figure 6, right). During the third leg of the turn the system runs row detection to refine the final goal point.

We have successfully tested both of the new turn methods. The limitation at this point is that we need to manually program which type of turn to do, smooth or three-point. We also tell the system how much space it will have in which to turn the robot (for planning collision-free paths). In future work we will program the system to autonomously assess the drivable space and decide which turn maneuver to execute.






Figure 5. Laurel attempts to enter the row lane indicated by the yellow arrow. The canopy of the trees at the near row, however, blocks the rest of the near edge and most of the far edge. The row detector has little data to work with in this situation.




Figure 6. (Left) The new planner generates a smooth trajectory for Laurel to follow, staying away from hard steering angles that stress the motor, while aligning the vehicle to be able to see into the row before entry. (Center) Once the vehicle has turned towards the row the system performs a row detection to refine the desired path of the vehicle. (Right) This row is too narrow for Laurel to smoothly turn into. The system plans a “K”-turn, where the vehicle turns past the row, backs up, and then enters. In the third leg of the turn the system performs a row detection to refine the final desired path.

## Field Tests

We conducted tests of the improved autonomous navigation system in Pennsylvania and Washington during the summer of 2010. Tests were conducted in large blocks because our goals required missions longer than 10 km. In places where blocks were not this long, we set the robot to repeat rows. Frequently upon arrival at a new block we would discover that minor tweaks to the system were necessary to deal with a new type of canopy previously unseen. In that case we had many short runs before we were ready to test long missions. For this reason, we only list the test missions longer than 10 km in the results below.

We logged a total of 159 km of autonomous travel in the period July-October 2010 (Table 2). The first 20 km were achieved in commercial orchards in Washington State. The next 80 km came at our Robot City test planting in Pittsburgh. At Robot City our primary concern was to find and fix problems such as the software memory errors that occurred on August 24<sup>th</sup> and 31<sup>st</sup>. These errors occurred when software elements—e.g., a desired path for the robot—were handled incorrectly, causing the computer to slowly run out of memory and the software controller to terminate unexpectedly. They are typically hard to find and require such endurance runs to uncover. After each of the tests on the 24<sup>th</sup> and 31<sup>st</sup> we ran memory check tools and corrected the software to fix the problem. Following the corrections for the test on the 31<sup>st</sup> we ran the system in a simulator to confirm that there were no more memory errors. The simulated system ran for over 30 km in eight hours of operation, so we were ready to proceed with more on-vehicle tests. Tests at Robot City became progressively longer, culminating in a 25 km run over five hours.

Although we reached the 100 km mark at Robot City, we continued to go out to Soergel Orchards in Wexford, PA and the FREC to demonstrate Laurel's successful operation in fully-developed orchards. We logged an additional 50 km at these test sites. Also at the FREC, but not listed in the results table, we wished to show the ability of the autonomous navigation system to handle different blocks and canopy types. We went to multiple blocks, programmed in the number of rows as well as their length and width, and told the system to start driving, without adjusting any other parameters. Laurel successfully drove six different blocks, despite differences in row width, canopy training, and growing systems (Figure 7).

Table 2. Autonomous driving runs obtained in 2010. Only those runs 10 km or longer were logged for the purposes of reaching the 100 km goal.

Date	Location	Distance (km)	Reason for Stoppage
07/27	Skyline East Orchards, Royal City, WA	10.0	Bad row detection caused by person walking in front of vehicle
07/27	Skyline East Orchards, Royal City, WA	10.0	Vehicle battery critically low
08/24	Robot City, Pittsburgh, PA	11.0	Software memory error
08/31	Robot City, Pittsburgh, PA	11.3	Software memory error
09/02	Robot City, Pittsburgh, PA	13.8	Missed turn caused by communication delay with vehicle controller
09/03	Robot City, Pittsburgh, PA	10.2	Manually stopped (no error)
09/08	Robot City, Pittsburgh, PA	15.6	Bad row detection due to dust cloud
09/09	Robot City, Pittsburgh, PA	25.8	Manually stopped (no error)
09/23	Soergel Orchards, Wexford, PA	10.0	Manually stopped (no error)
09/29	Penn State FREC, Biglerville, PA	15.1	Manually stopped (no error)
10/01	Penn State FREC, Biglerville, PA	12.6	Manually stopped (no error)
10/02	Penn State FREC, Biglerville, PA	13.3	Manually stopped (no error)





Figure 7. A sampling of the blocks driven by Laurel in 2010. Growing systems varied from recent fruiting wall plantings (bottom middle) to completely standalone, untrained trees (bottom right). All pictures are from blocks at the Penn State FREC.

### ***Making APMs Usable by Growers***

Having successfully demonstrated over 150 km of autonomous driving in experimental and commercial orchards, we started looking toward commercial deployment of the APM vehicle family. The first step was to create vehicles that can not only survive the harsh orchard environment year round, but also be operable by growers and farm workers. To that end we commissioned three new vehicles. One was delivered to Carnegie Mellon for use in pace mode, where an off-board operator can command the vehicle to work on an entire orchard block. The other two are two-person platform vehicles belonging to The Pennsylvania State and Washington State Universities. Their purpose is to operate in scaffold mode, where the vehicle very slowly follows a row of trees while on-board workers perform tree maintenance, including pruning, thinning, training, and harvesting.

The computing demands of our software system could be satisfied by an average personal computer. Laptops and other common home computers, however, are not robust to the high temperatures found in all-day orchard operation during the summer, nor are they sealed to prevent damage from dust or rain. To operate a vehicle outdoors continuously we require a specialized computing solution. We chose an embedded computer from SmallPC which is completely sealed and rated to a temperature of 105 degrees Fahrenheit. One of these computers was installed in each of the three new vehicles, to keep the vehicles as similar to each other as possible.

The embedded computer runs the autonomy software, receiving data from the laser rangefinder and sending commands to the vehicle's drive-by-wire controller. The user controls the system from a device that communicates with the embedded computer via WiFi through a ruggedized wireless access point (Figure 8, left).

Finally, we needed to consider remote maintenance of the vehicles. As we continue to improve the software throughout the year we will need a way to provide updates to the users. Also, in

case of abnormal vehicle behavior we need a way to access data logs on the embedded computer. To solve these problems we purchased two cellular modems, one each for Washington State and Penn State. The embedded computer runs a script such that, when the modem is turned on, the script automatically creates a connection to a server at Carnegie Mellon. This way, we can provide updates and perform remote debugging with no work required from the user.

By April 2011 we had completed the embedded computing and wireless capability on the two platforms, Allegheny and Cascade (Figure 8, right). The vehicles were shipped to Washington State and Penn State in May. Preliminary results where workers performed tree pruning and training and blossom and fruit thinning indicate efficiency improvements up to 58% when compared to workers performing the same tasks with ladders. Additionally, the preliminary results indicate energy savings on the order of 97%, with the electric APMs coming in at \$0.16/acre versus diesel platforms that come in at \$5.95/acre. We also learned, however, that the stock eWorkman chassis is not appropriate for the very low speed, continuous motion that such tasks require. Achieving consistent low speed is difficult for the vehicle, and it overheats after extended operation. We are investigating solutions to this problem, including a possible vehicle redesign in cooperation with Toro.



Figure 8. (Left) Grower-friendly interfaces are an important element of the overall usability aspect of autonomous vehicles. Our current interface, based on formal human-computer interaction methodology, runs through a web browser on any mobile device. It allows workers to control the desired speed and lateral position of the vehicle in the row, and also provides workers with feedback on the status of the autonomy system. (Right) To begin investigating the commercial prospects of autonomous orchard vehicles, we commissioned three APMs: Tuscarora for pace mode and Allegheny and Cascade for scaffold mode, the latter two equipped with a scissors lift. Here a worker at Hollabaugh Orchards in Biglerville, PA, uses Allegheny APM to train apple trees.

We will report on the performance of these vehicles throughout the summer of 2011 in a future publication.

## Conclusion

As fruit production moves from large, tall trees to high-density fruit walls, the potential for automation increases significantly. While we may be many years away from robotic fruit picking that is economically viable, we already have the technology necessary to create cost-effective autonomous orchard vehicles that can augment or automate tree pruning and training, blossom and fruit thinning, fruit harvesting, mowing, spraying, sensing, etc. In fact, orchard platforms are a daily reality in many tree fruit production operations in Europe, and their presence in the US has been growing steadily. Here we described a family of autonomous orchard vehicles that can be reconfigured and retasked year-round for a variety of operations. To achieve the ultimate goal of turning these vehicles into commercial products, we follow a stepped approach in which feasibility is followed by robustness and usability. As of the end of June 2011, we had completed over 300 km of autonomous orchard traversals in experimental and commercial orchards in Pennsylvania and Washington. By the end of September 2011 we expect to demonstrate 100 km of driving without any supervision by the Carnegie Mellon engineering team—a challenging goal, but one we believe is necessary before this technology is accepted by industry.

## Acknowledgements

This work is supported by the US Dept. of Agriculture National Institute for Food and Agriculture under award no. 2008-51180-04876.

## References

- Dellaert, F., D. Fox, W. Burgard, and S. Thrun. 1999. Monte Carlo localization for mobile robots. Proceedings IEEE International Conference on Robotics and Automation, 1322-1328.
- Duda, R. O. and P. E. Hart. 1972. Use of the Hough Transformation to Detect Lines and Curves in Pictures. Comm. ACM, 15: 11–15.
- Ferguson, D. 2010. Path Generation and Control for End of Row Turning in an Orchard Environment. Master Thesis, Carnegie Mellon University Robotics Institute.
- Gustafsson, F., F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund. 2002. Particle filters for positioning, navigation, and tracking. IEEE Transactions on Signal Processing, 50(2):425-437.
- Hamner, B., S. Singh, M. Bergerman. 2010. Improving orchard efficiency with autonomous utility vehicles. American Society of Agricultural and Biological Engineering Annual Convention Proceedings Abstract and Paper.
- Singh, S., M. Bergerman, J. Cannons, B. Grocholsky, B. Hamner, G. Holguin, L. Hull, V. Jones, G. Kantor, H. Koselka, G. Li, J. Owen, J. Park, W. Shi, J. Teza. 2010. Comprehensive Automation for Specialty Crops: Year 1 Results and Lessons Learned. *Journal of Intelligent Service Robotics, Special Issue on Agricultural Robotics*, 3(4): 245-262.