Comprehensive Automation for Specialty Crops

Funded by the USDA NIFA Specialty Crop Research Initiative

Accomplishment Report

October 1st, 2009 through September 30th, 2010

Project Director: Sanjiv Singh

Carnegie Mellon University

ssingh@cmu.edu

Table of Contents

1. Introduction

CASC is a multi-institutional, multi-disciplinary initiative to comprehensively 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 expect to demonstrate significant advances from the integration of robotics technology and plant science; to understand and overcome the socio-economic barriers that prevent technology adoption; and to make our results available to growers and stakeholders through a nationwide outreach program.

CASC's partners include academic institutions (Carnegie Mellon University, Pennsylvania State University, Washington State University, Oregon State University, and Purdue University), agricultural machinery companies (Vision Robotics, IONco, Toro, and Trimble), and a federal research laboratory (USDA Agricultural Research Service). CASC activities are overseen by an advisory panel of growers and stakeholders, including the Washington Tree Fruit Research Commission, US Apple Association, National Wine and Grape Initiative, California Citrus Quality Council, California Canning Peach Association, International Fruit Tree Association, and several US universities, among others.

In this document the CASC team presents the accomplishments it attained in the first year of the project in each of the three main themes that comprise the project (Figure 1).

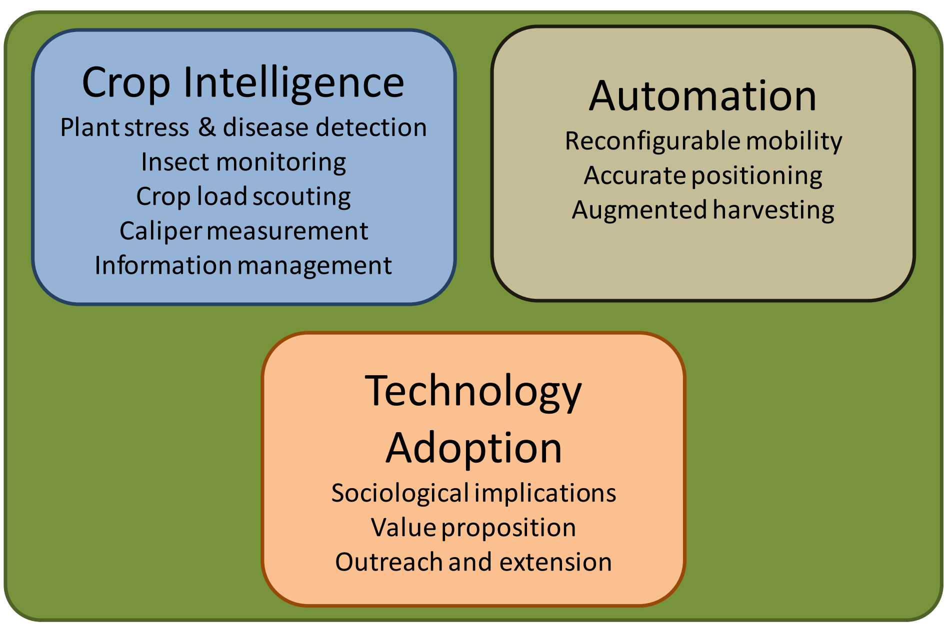

Figure 1. CASC's three main themes and their underlying eleven thematic areas. CASC impacts growers by providing crop intelligence that increase farm efficiency and automation solutions that reduce production costs. All technologies developed are validated via thorough socio-economic analyses aimed at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

Nomenclature

In this document we use the following acronyms and abbreviations:

ARS	USDA Appalachian Fruit Research Station in Kearneysville, WV
CMU	Carnegie Mellon University
FB	Fire blight
FREC	PSU Fruit Research and Extension Center in Biglerville, PA
GIS	Geographical information system
IFM	Internal fruit moth
MMDAQ	Multi-modal data acquisition system
NDVI	Normalized difference vegetation index
OSU	Oregon State University
Purdue	Purdue University
PSU	Pennsylvania State University
ROI	Region of interest
USDA	US Department of Agriculture
VRC	Vision Robotics Corp.
WSU	Washington State University
WTFRC	Washington Tree Fruit Research Commission

Acknowledgements

CASC is primarily funded by the USDA under award no. 2008-51180-04876 and matching funds from growers, industry associations, agricultural equipment manufacturers, the states of Pennsylvania, Oregon, and Washington, and the institutions members of the project.

The following people and institutions provided valuable information and access to facilities for the execution of the activities described in this document: Reed Soergel, Soergel's Orchard, Wexford, PA; Bruce Hollabaugh, Hollabaugh Bros., Biglerville, PA; Corey McCleaf, McCleaf Orchard, Biglerville, PA; Terry Salada, Penn State Fruit Research and Extension Center, Biglerville, PA; John Baugher, Adams County Nursery, Aspers, PA; Brian Knouse, Knouse Fruitlands, Arendtsville, PA; Joy Cline, Bear Mountain Orchard, Aspers, PA; Neil Starner, Fruit Haven, Aspers, PA; Ken Guise, C & G Orchard, Gardners, PA; John Verbrugge, Valley Fruit, Yakima, WA; Jay Brunner, WSU Tree Fruit Research and Extension Center, Wenatchee, WA; Qin Zhang, WSU Center for Precision Agriculture, Prosser, WA; Patrick Scharf, WSU Center for Precision Agriculture, Prosser, WA; Roger Adams, Willow Drive Nursery, Ephrata, WA.; Paul Tvergyak, Cameron Nursery, El Topia, WA; McDougall Farms, Royal City, WA; Washington Fruit and Produce, Royal City, WA.

2. Outputs and Outcomes

The USDA National Food and Agriculture Institute (NIFA) developed a logic model that describes the interrelationship between inputs, activities, outputs, and outcomes of its projects. These concepts are defined as:

- Inputs: What is invested, such as resources, contributions, and investments that are provided for the program.
- Activities: What the program does with its inputs to services it provides to fulfill its mission.
- Outputs: Products, services and events intended to lead to the program's outcomes.
- Outcomes: Planned results or changes for individuals, groups, communities, organizations or systems.

NIFA's complete logic model for the Specialty Crop Research Initiative (SCRI) defines each of these four concepts in detail. We refer the reader to their web site for more information (http://www.csrees.usda.gov/funding/integrated/integrated_logic_model.html). Here we describe how CASC researchers contributed to the success of the SCRI in terms of outputs and outcomes.

2.1 CASC Outputs

Since the beginning of the project, CASC researchers produced the following outputs (classified according to the SCRI logic model):

Basic and Applied Research

<u>Stress and disease detection</u>

- "On-the-go" multimodal sensing system for plant stress detection based on canopy reflectance and temperature, including a graphical user interface for sensor fusion and data visualization.
- Algorithms for detecting fire blight damage in digital images taken in the greenhouse and the field.

<u>Insect monitoring</u>

- Two models of electronic trap prototypes based on IR sensors (IR-trap, six units) and on bio-impedance sensors (Z-trap, twenty units).
- Image processing-based internal feeding worm damage detection algorithms.
- 5,800 images of internal feeding worm-damaged apples and healthy apples across four varieties, to validate damage detection algorithms.

- Web-accessible database of images of internal feeding worm-damaged and healthy apples collected in 2009.

Crop load scouting

- Upgraded crop load estimation scout ("Newton") system including new mast, additional cameras, industrial flash, GPS receiver.
- Improved crop load (fruit count and size) estimation software.

Caliper measurement

- "On-the-go" automatic caliper for measuring tree diameter in the field.
- "On-the-go" automatic counter for counting nursery trees in the field.
- Four weeks of field tests in Ephrata, WA, Royal City, WA, and Hickman, CA.
- Large field experiment with Crop Tech LLC to count 1 million nursery trees to validate an insurance claim using CASC automatic counter.

Information management

- Web-based geographic information system tool for collecting and managing crop information from the field.

Reconfigurable mobility

- Upgraded Autonomous Prime Mover based on Toro electric utility vehicle, including new computing, sensors, and environmental protection.
- Hydraulic orchard platform converted to autonomous orchard vehicle, including computing and sensing.
- Improved row following and turning algorithms for the Autonomous Prime Mover family of orchard vehicles.
- New user-friendly Autonomous Prime Mover command interface.
- Four weeks of field tests in Pittsburgh, PA, Wexford, PA, Biglerville, PA, and Doubs, MD.
- Two weeks of field tests in Wenatchee, WA and Royal City, WA.

Accurate positioning

- Positioning algorithm that eliminates the need for artificial infrastructure in the orchard.

Augment harvesting

- Passive bin-filling prototype using a variety of energy absorbing materials.

- Harvest assist transport and bin filling system for apple orchards.

Outreach to Producers, Processors, and Consumers

- 6,218 direct contacts through field days and presentations.
- 45 students involved in CASC activities.
- Socioeconomic surveys of Eastern and Pacific Northwest growers on barriers to technology adoption.
- Economic models for four CASC thematic areas.
- Business plan for digital traps.
- AgTools website.
- AgFinance software (alpha version).
- Licensing agreement with Spensa Corp. to commercialize digital traps.
- Partnership with Crop Tech LLC to commercialize counter.
- Outreach and collaboration tools:
    - CASC web site: http://www.cascrop.com
    - CASC on YouTube: http://www.youtube.com/user/TheCASCrop
    - CASC on Facebook: http://on.fb.me/awfoFB
    - CASC on Slideshare: http://www.slideshare.net/CASCrop
- Publications
    - 39 research papers and written reports
    - 71 invited presentations
    - 61 trade magazine and newspaper feature stories
    - 14 bulletins, fact sheets, computer software and videos
    - 47 workshops, tours, field days, and reports to grower organizations

2.2 CASC Outcomes

Likewise, CASC researchers contributed the following outcomes to the SCRI program:

Short-Term

Generate new knowledge for specialty crop systems

- Determined that spectral signature (NDVI) and canopy temperature are highly correlated to plant water stress

- Detected up to 36% photosynthesis reduction due to plant water stress

<u>Adapt existing knowledge to specialty crop systems</u>

- Image processing-based internal feeding worm detection algorithm achieves >85% accuracy with <4% false alarm rate

<u>Engage broadest possible scientific community in challenges faced by specialty crop industries</u>

- Engaged team of 50+ horticulturists, entomologists, engineers, computer scientists, economists, social scientists

<u>Web based and other digital information</u>

- CASC web site, YouTube channel, Facebook page, and Slideshare account.

Medium-Term

<u>New processes and products for specialty crop producers</u>

- Built two autonomous prime movers (APMs)
- Deployed APMs in orchards in WA and PA
  - Toro electric utility vehicle autonomously traversed 300+ km of orchard rows
  - Hydraulic orchard platform autonomously traversed 10 km of orchard rows
- Designed and tested grower-friendly APM user interface
- Demonstrated sub-meter positioning accuracy without GPS in over 100 km of field trials
- Integrated GPS-free positioning system with APM software architecture
- Designed, built, and field-tested harvest assist system
  - Initial trials showed 10% improvement in harvesting speed with 5% reduction in bruising
  - Determined that singulation of fruit during transport from tree to bin is essential to reduce bruising
- Designed, built, and field-tested digital traps with a capture rate comparable to standard delta traps
  - In laboratory, digital traps achieve >95% detection accuracy
  - In the field, preliminary data indicate that digital traps achieve >80% detection accuracy with <5% false alarm rate
- Designed, built, and field-tested crop load scout system

  - For red fruit, the crop load estimate for six acres was 75% of the load extrapolated from hand count; median apple diameter estimate within 1% of the median of 2.58”
  - For green fruit, the crop load estimate for eleven acres was 94% of the load extrapolated from hand count; median apple diameter estimate was within 3% of the median of 2.63”
- Designed, built, and field-tested on-the-go automatic caliper
  - Determined caliper within 2 mm of harvested bareroot tree in the warehouse
  - Determined caliper within 3 mm of nursery trees while moving at 2 mph
  - Caliper data used to effectively determine bareroot tree grade
- Designed, built, and field-tested on-the-go automatic tree counter
  - Counted in-ground nursery trees of >= 1/4” in caliper with 95% accuracy while moving at 2.5 mph
  - Counted in-ground nursery trees of >= 1/2” in caliper with 97% accuracy while moving at 3 mph

<u>New professionals engaged in specialty crop systems</u>

- Directly involved 45 students in CASC activities

<u>Producers and processors adopt newly developed technologies and innovations</u>

- A quarter of mid-Atlantic producers who attended CASC presentations are adopting new technologies for precision agriculture
- PNW growers who attended a CASC field day rated interest in adoption 3.8 to 4.4 on a 1-5 scale
- Thirty percent of PA producers who attended field days are adopting trellised planting systems and 65% plan to make this change

<u>Networks that improve the flow of information among all components of specialty crop systems</u>

- Created open source information management system to aid growers in data capture and analysis and decision-making

Long-Term

<u>Profitable systems for specialty crop production/processing</u>

- Socio-economic survey findings include
- Crop projection sensor more popular in PNW than in the East

- Justifiable price point of technology that increases fruit packout is higher in the Northwest
- Eastern growers are less concerned with water availability effects on crop production.

Increased competitiveness of U.S. specialty crop producers and processors

- Preliminary findings suggest that CASC technologies could have an impact ranging from $20 to $1,047/acre/year
- Management efficiency trials in pilot orchards demonstrated increases in efficiency as high as 78%

3. Crop Intelligence

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.

This section presents the goals and accomplishments in the following five thematic areas:

- Plant stress and disease detection;
- Insect monitoring;
- Crop load scouting;
- Caliper measurement;
- Information management.

3.1 Plant Stress and Disease Detection

Thematic area leaders

Name	Institution	Email
Johnny Park	Purdue University	jpark@purdue.edu
Henry Ngugi	The Pennsylvania State University	hkn3@psu.edu
D. Michael Glenn	USDA Agricultural Research Service	michael.glenn@ars.usda.gov

Year 2 goals

Activities	Deliverables	Success Criteria
1. Improve sensor fusion-based stress detection algorithm. 2. Develop stereo vision algorithms for locating shoots where fire blight infection starts. 3. Integrate sensors into APM. 4. Carry out field and greenhouse stress and fire blight detection experiments.	1. Report on the performance of stress detection and shoot localization algorithms. 2. Report on the results from greenhouse and field experiments. 3. Publicly available dataset of images with accompanying ground truth.	1. Identify 25% photosynthesis reduction due to stress in field conditions. 2. Localize > 70% of shoots in the field-of-view with < 30% false alarm rate.

Notable results:

- Developed an on-the-go multimodal sensing system for plant stress detection based on canopy reflectance and temperature and a graphical user interface for sensor fusion and data visualization.
- Determined that spectral signature (NDVI) and canopy temperature are highly correlated to plant water stress.
- Measured photosynthesis reduction in growth chambers and detected up to 36% photosynthesis reduction due to plant water stress (75% water) by multimodal sensors.
- Developed several algorithms for detecting fire blight damage in digital images taken in the greenhouse and the field, but all of the methods still have a high rate of false detections.

Data Analysis from Year 1

Multi-modal sensor data was collected in the PSU-FREC greenhouse in 2009 for the study of water stress detection on five groups of young apple trees (Buckeye Gala) in different water treatment (100%, 90%, 75%, 60%, and 45% at field capacity) with four replications per group. Data collection was continued on a weekly basis from late April to June. Water treatment levels were defined as the percentage replacement of water content to bring the tree to field capacity and were measured by weighing the containers before irrigation in order to calculate the needed replacement amount.

Hyperspectral Image Analysis

The hyperspectral camera was calibrated with dark and white reference images for every group of ten trees prior to image acquisitions (Kim et al., 2010a). The calibrated images were processed to calculate plant stress index. Thirteen different spectral indices that have been published were selected for plant stress detection, as shown in Table 1. Seventeen individual spectral band images for the spectral index calculations were extracted. The vegetation area in each image was segmented out using the corresponding NDVI value and processed to calculate canopy reflectance.

Table 1. Spectral indices evaluated for plant stress detection. The three digit numbers following the letter *R* in each formula represent the response of the hyperspectral camera at that particular wavelength.

Broadband Greenness	NDVI (Normalized Difference Vegetation Index)	$= \frac{R800 - R680}{R800 + R680}$
	SRI (Simple Ratio Index)	$= \frac{R900}{R680}$
	EVI (Enhanced Vegetation Index)	$= 2.5 \times \frac{R800 - R680}{R800 + 6 \times R680 - 7.5 \times R450 + 1}$
	ARVI (Atmospherically Resistant Vegetation Index)	$= \frac{R800 - (2 \times R680 - R450)}{R800 + (2 \times R680 - R450)}$
Narrowband Greenness	Red Edge NDVI	$= \frac{R750 - R705}{R750 + R705}$
	Modified Red Edge NDVI	$= \frac{R750 - R705}{R750 + R705 - 2 \times R445}$
	Modified Red Edge SRI	$= \frac{R750 - R445}{R705 - R445}$
	VOG Red Edge Index (REI) 1	$= \frac{R740}{R720}$
	VOG REI 2	$= \frac{R734 - R747}{R715 + R726}$
	VOG REI 3	$= \frac{R734 - R747}{R715 + R720}$
Light Use Efficiency	PRI (Photochemical Reflectance Index)	$= \frac{R531 - R570}{R531 + R570}$
Dry or Senescent Carbon	PSRI (Plant Senescence Reflectance Index)	$= \frac{R680 - R500}{R750}$
Canopy Water Content	WI (Water Index)	$= \frac{R900}{R970}$

The broadband indices compare a reflectance peak in near infrared (NIR) to another peak in red range where chlorophyll absorbs photons for photosynthesis. Since these features are spectrally broad, they can work effectively with broadband multispectral sensors. The narrowband indices use red edge that is a steeply sloped region of the vegetation reflectance curve between 690 nm and 740 nm, and caused by the transition from chlorophyll absorption to NIR leaf scattering.

The results of the four best indices from the images acquired on May 4, 7, 11, and June 4, 2009 are shown in Figure 2. NDVI had the highest average correlation among broadband greenness indices with $R^2 = 0.68$ (Figure 2a). The highest average correlation among narrowband greenness indices was found at Red Edge NDVI with $R^2 = 0.89$ (Figure 2c).

Table 2 presents the correlation values of selected indices to water stress over the tree images. Since narrowband measurements in the red edge are more sensitive to smaller changes in vegetation health than broadband indices, narrowband indices are suitable for hyperspectral sensors. The result of our experiment with the hyperspectral camera supports this result with better performance in narrowband indices.

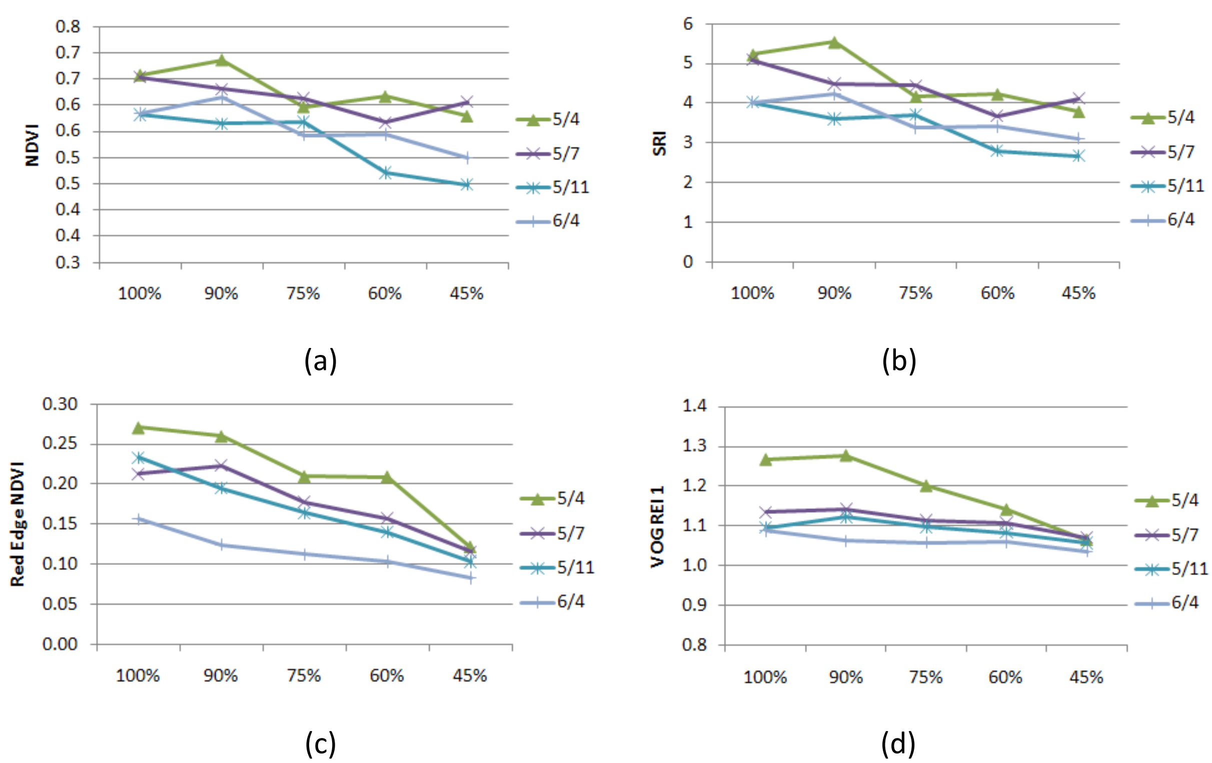

Figure 2. Hyperspectral image responses to water treatments (a) NDVI, (b) SRI, (c) Red Edge NDVI, and (d) VOG REI 1.

Table 2. Spectral indices that were evaluated for plant stress detection.

		4/27	5/4	5/7	5/11	6/4	R	R^2
Broadband Greenness	NDVI	0.70	0.83	0.76	0.93	0.89	0.82	0.68
	SRI	0.55	0.89	0.81	0.94	0.90	0.82	0.67
Narrowband Greenness	Red Edge NDVI	0.86	0.95	0.96	0.99	0.96	0.94	0.89
	VOG REI1	0.95	0.98	0.94	0.82	0.89	0.91	0.84

GreenSeeker NDVI Analysis

NDVI values of both the vegetation area of tree canopy and the non-vegetation area between trees were obtained with the NDVI sensor. Figure 3a shows continuous readings of NDVI in which higher values are from the vegetation area of canopy and lower values are from the non-vegetation background. NDVI values corresponding to each apple tree were manually selected by comparing a digital image in the same time stamp (marked with circles in Figure 3a). Average NDVI responses over 13 measurement dates generally match water treatment patterns with a correlation of R = 0.6 (Figure 3b). Low correlation may be due to measurement errors in the data manually selected from unsynchronized NDVI readings and digital image acquisition.

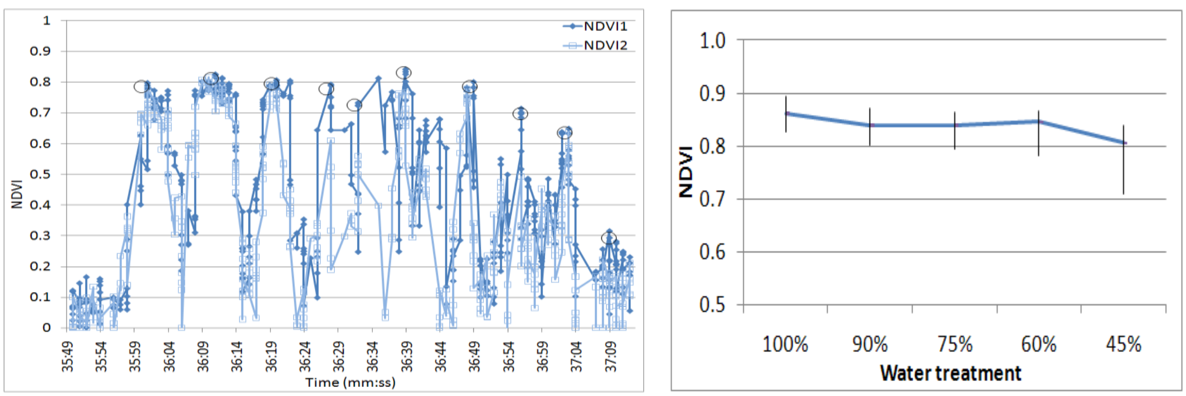

(a) (b)

Figure 3. NDVI sensor responses to water treatments: (a) continuous readings of vegetation of tree canopy and non-vegetation area between trees, (b) average NDVI over 13 measurement dates indicating correlation (R = 0.6) between NDVI measurements and plant water treatment.

Canopy Coverage Analysis

Leaf area index (LAI) provides an important parameter to detect stress, as leaf area influences photosynthesis and plant growth. A value of LAI = 3 is obtained when almost 100% of the incoming light is intercepted and thus corresponds to a canopy coverage of 100% (Potato Explorer, 2010). Digital images were acquired from nadir view and processed for canopy coverage. Figure 4 displays an image sequence of apple trees with five different water treatments. Images are arranged in chronological order from April 27 to June 18 with an interval of 3–5 days. Canopy coverage was calculated by image segmentation of canopy area with NDVI image masking and presented as a percentage of canopy coverage. As plants grow over the time, the growth pattern increases in the beginning of the time period, followed by a relatively static period of growth in the middle (May 10 to June 10) (Figure 5a). After June 10, there are clear decreases of canopy coverage in 60%- and 45%-watered groups, while in the other three groups the canopy coverage remains high and static. The average percentage canopy coverage over the period is 50% in 100%-watered trees and 27% in 45%-watered trees (Figure 5b), resulting in an overall correlation R = 0.7. The result indicates that canopy coverage measurement can be used as a supplemental index of chronic plant water stress.

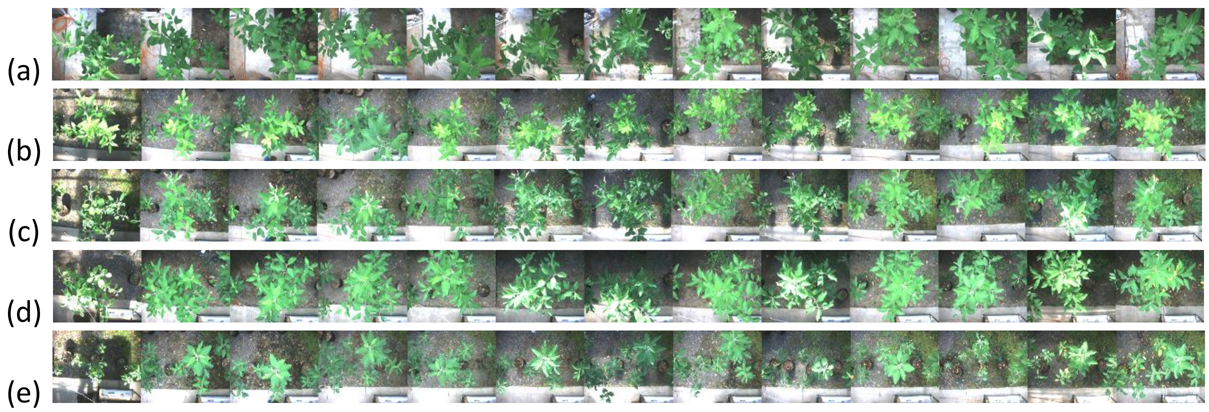

Figure 4. Canopy growth (top-view) in 13 measurements from April 27 to June 18 on young apple trees with five different water treatments: (a) 100%, (b) 90%, (c) 75%, (d) 60%, and (e) 45%.

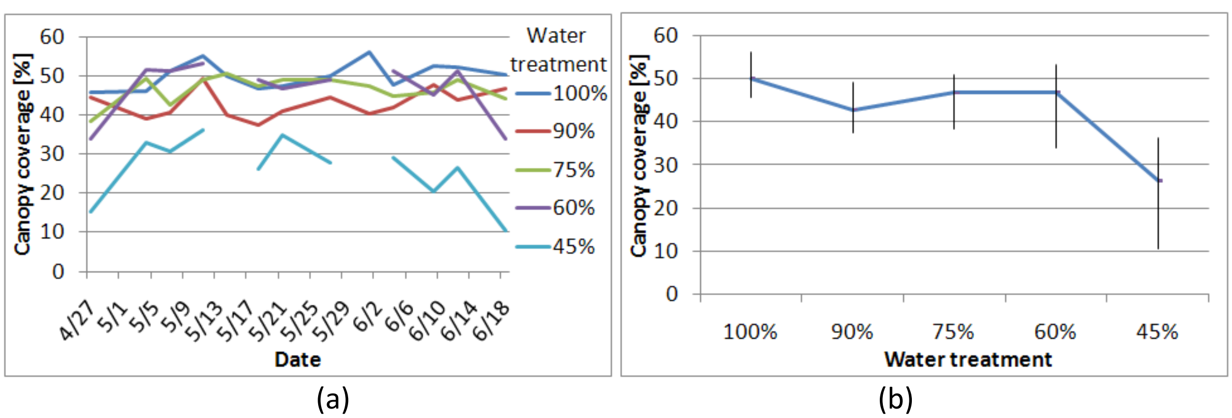

Figure 5. Percentage canopy coverage to different water treatments: (a) measurements along the plant growth from April 27 to June 18 (partly missing data in 60%- and 45%-watered trees due to an image formation problem), (b) average canopy coverage to water treatments with correlation of R = 0.7.

Field Data Analysis for Disease Detection

Field data obtained during the Year 1 growing season were processed for disease detection. The data had been collected at the Pennsylvania State University Fruit Research and Extension Center (FREC) on a 5-cultivar plot. Each treatment block consists of five cultivar apple trees:

Rome Beauty (RB), Golden Delicious (GD), Stayman (Stay), Cortland (Cort), and Red Delicious (RD). The plot was arranged in a randomized complete block design consisting of 20 fungicide treatments each replicated four times (Figure 6), thus generating data for 400 trees in total. GreenSeeker NDVI data was collected in continuous mode while driving down the row in the same direction in each path.

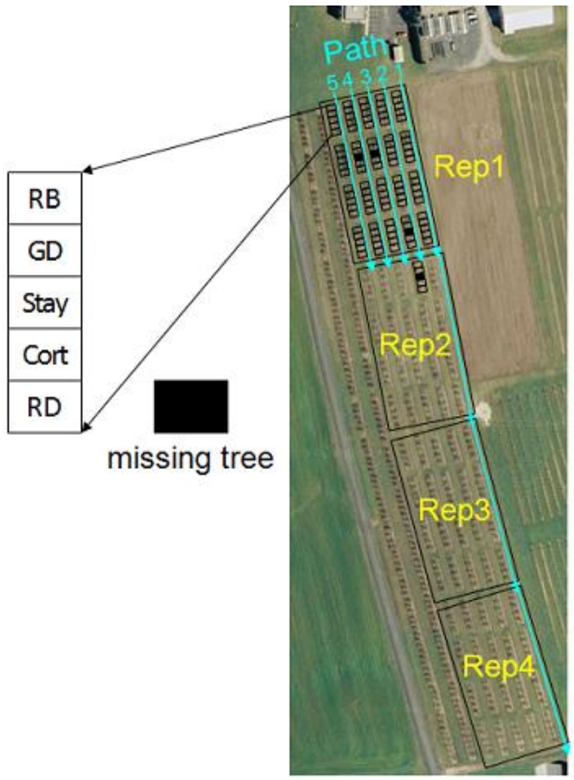

Rep 1	14	3	15	16	8
	19	2	4	9	13
	1	20	11	12	7
	17	6	18	5	10
Rep 2	4	18	19	5	3
	9	6	20	14	15
	11	10	2	1	7
	13	17	16	12	8
Rep 3	3	8	11	9	6
	20	10	1	19	15
	2	18	14	17	5
	16	12	4	13	7
Rep 4	19	7	17	2	16
	5	4	15	9	12
	3	8	6	1	13
	11	10	20	18	14

Figure 6. Field experiment layout for disease detection on five cultivar apple trees: Rome Beauty (RB), Golden Delicious (GD), Stayman (Stay), Cortland (Cort), and Red Delicious (RD), with 20 treatments and four replications.

Three types of diseases, powdery mildew (pm), cedar apple rust, and scab were measured by hand-sampling to obtain a disease severity value on leaves. Each disease was compared with 20 treatments and 5-cultivar apple trees. NDVI response was also compared among treatments and apple variety. The diseases varied significantly between treatments and varieties. Overall correlation to NDVI values are shown in Table 3. The highest correlation to NDVI responses found were: Golden Delicious (R = -0.62 with scab), Stayman (R = -0.58 with rust), and Cortland (R = -0.38 with powdery mildew). Table 3 also shows the comparison of NDVI sensor performance with (R_f) and without (R) filtering by canopy coverage conditions. Better corrections were found on average when the filter was applied.

Table 3. Correlation of NDVI sensor responses to three diseases (powdery mildew, rust, and scab) in 5-cultivar apple trees. R_f indicates correlation to NDVI values that were filtered by canopy coverage conditions. In each row the entry with highest value is marked in bold.

Disease	Correlation	RB	GD	Stav	Cort	RD	Average
Powdery Mildew	R =	-0.09	-0.04	-0.22	**-0.41**	-0.19	-0.19
	R_f =	-0.15	-0.02	-0.25	**-0.38**	-0.19	-0.20
Rust	R =	-0.37	-0.53	**-0.58**	-0.46	-0.27	-0.44
	R_f =	-0.43	-0.51	**-0.58**	-0.43	-0.33	-0.46
Scab	R =	-0.58	**-0.60**	-0.09	-0.16	-0.44	-0.37
	R_f =	-0.56	**-0.62**	-0.13	-0.19	-0.48	-0.39

System Update in Year 2

In this year we made several updates to the multimodal sensor system. Hardware updates include:

- a multi-spectral camera (ADC Lite from TetraCam, Chatsworth, CA) that captures light from 520 nm in the visible spectrum to 920 nm in the near infrared, delivering the red, green, and NIR bands needed for calculation of NDVI;
- two rangefinder ultrasonic sensors (MB7060 from MaxBotix, Baxter, MN) to validate the NDVI sensor's target distance (Kim et al., 2010b);
- a low-cost WAAS-enabled GPS (18xPC from Garmin, Olathe, KS) for geo-referenced field mapping; and
- six IR temperature sensors (five MLX90615_10FOV and one MLX90615_90FOV from Parallax, Rocklin, CA) and a microcontroller (BS2-IC from Parallax, Rocklin, CA) to integrate them to a sensor array.

Software updates include:

- integration of the ultrasonic rangefinders to the multimodal data acquisition (MMDAQ) software to filter NDVI readings based on target distance;
- integration of the GPS to the MMDAQ to read and append latitude/longitude positions to sensor readings; and
- an interface with IR thermometer array and display of 2-D thermal map in real-time or retrieval mode.

Greenhouse Experiment

The greenhouse experiment was prepared at the PSU-FREC with young apple trees (cultivar 'Gale Gala') grafted on M9 rootstocks and potted in Berger BM1 potting medium—twenty trees for water stress and five trees for fire blight (FB) studies. For the water stress study, each set of five trees were treated at different water amounts of 100% (control), 75%, 65%, 55%, and 45%

(most stressed) at field capacity with four replications per set. Water treatment levels were defined as the percentage replacement of water content to bring the tree to field capacity and were measured by weighing the containers before irrigation in order to calculate the needed replacement amount.

The multimodal sensor system was installed on a cart with extended support pipes and positioned to get nadir view over the top of the canopy (Figure 7, left). The sensor system includes a multispectral camera, an NDVI sensor, a digital camera, an ultrasonic rangefinder, and a thermal imager (Figure 7, right). Solar radiation, air temperature, and relative humidity in the greenhouse were separately recorded by a datalogger (WatchDog 450 from Spectrum Technologies, Inc., Plainfield, IL). Soil water status at two trees of each water group was also monitored using soil moisture sensors (Watermark from Irrometer, Riverside, CA).

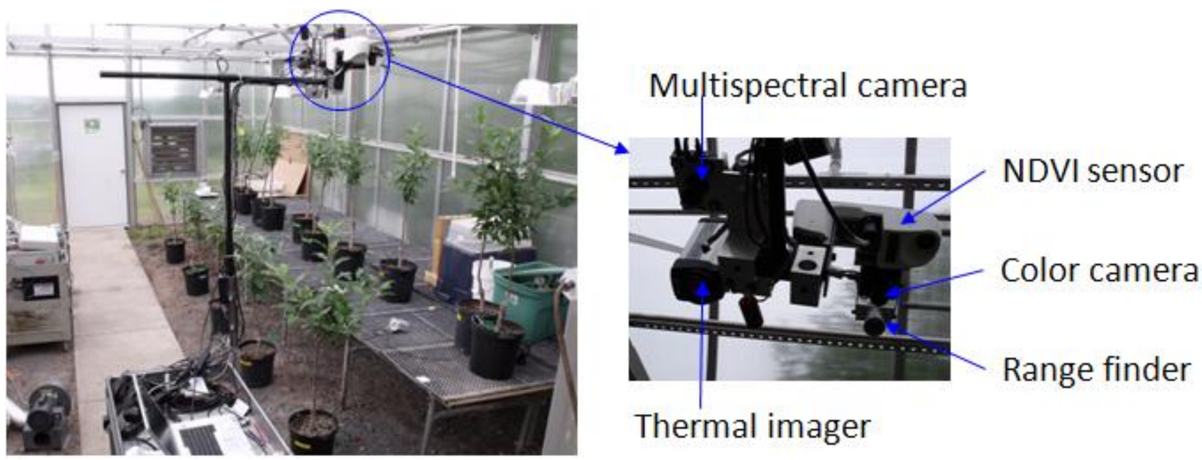

Figure 7. Greenhouse experiment apparatus for plant stress study.

Canopy Reflectance

Canopy reflectance in near-infrared (NIR) and red spectral bands of young apple trees was measured by a NDVI sensor and a multispectral camera. The NDVI sensor generates a point value from a line scan at 10 Hz and was used on a continuous mobile mode. The sensor obtained NDVI values of both the vegetation area of tree canopy and the non-vegetation area between trees along a driving path. An ultrasonic rangefinder was deployed to determine the distance from the sensor to the target canopy and filter out NDVI values if synchronously measured target distance is above a set point. Figure 8a shows continuous readings of NDVI and target distances of five trees under different water treatments and four replications. NDVI values at target distance above 170 cm were considered as non-canopy area; this allowed us to filter out the majority of noisy NDVI measurements (Figure 8b). Still, there were low NDVI values that were suspected as non-vegetation but not filtered due to sparse presence of leaves and unmatched field-of-view (FOV) shapes of the NDVI sensor (line) and the rangefinder (oval). Improved results are expected in outdoor application with mature dense canopy trees. The highest NDVI values for each tree were manually selected and confirmed by comparing a digital image in the same time stamp and marked in circles in Figure 8a.

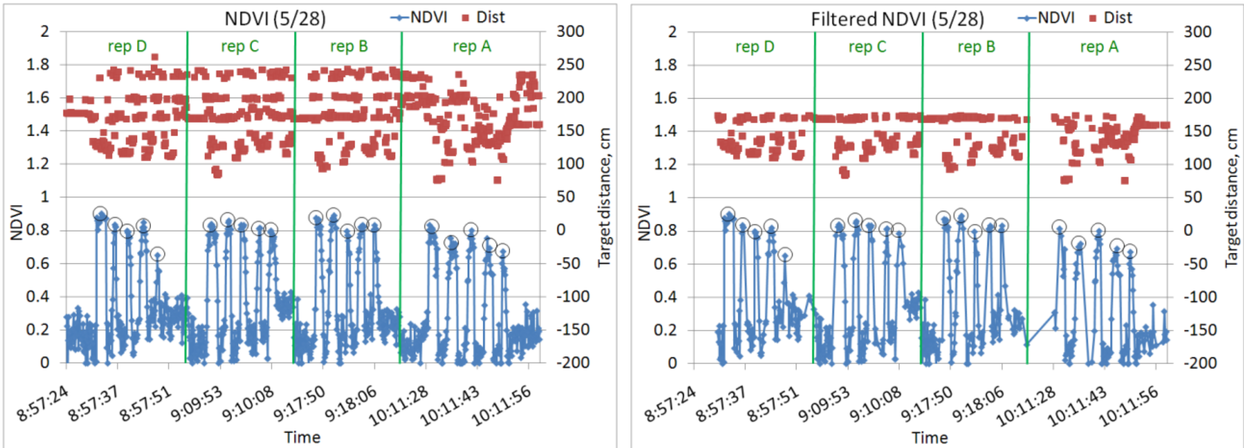

(a) (b)

Figure 8. Continuous readings of NDVI and target distances of five trees under different water treatments and four replications: (a) before filtering, (b) after filtering non-vegetation (measurements above 170 cm).

The multispectral camera captures an image (2048 x 1536 pixels) in the NIR, red, and green spectral bands. The image is displayed as a false color image by mapping the original bands to red, green, and blue, respectively (Figure 9a). The image was processed for canopy segmentation (Figure 9b) using five threshold or sieve values that separates canopy from the soil spectral space within the red—NIR region, and then converted to a gray-scaled NDVI image (Figure 9c) by spectral band math (NIR-red)/(NIR+red). A calibration coefficient was determined by reading pixels on a white reference tag contained in the image as shown at the bottom of Figure 9b and applied for reflectance calibration for all rest of pixels in generating the NDVI image. Canopy NDVI was calculated by manually selecting a region of interest (ROI) on the canopy area.

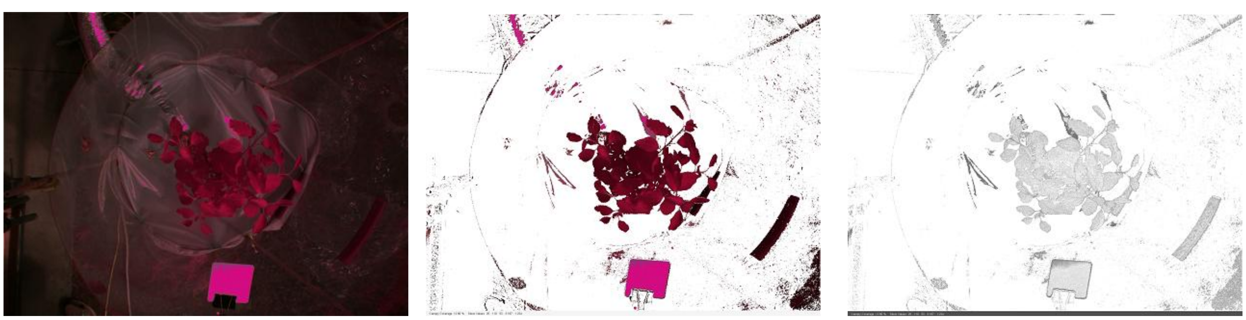

(a) (b) (c)

Figure 9. Multispectral image of a tree in a photosynthesis chamber under 45% water treatment acquired on June 3, 2010. The white reference tag at the bottom is used for reflectance calibration. (a) A false color image of NIR, red, and green mapped to RGB. (b) Canopy segmented image using five threshold values. (c) Gray-scaled NDVI image in which canopy ROI was selected for canopy NDVI.

Comparison of NDVI responses between the NDVI sensor and the multispectral camera is plotted in Figure 10a for six measurement dates from June 3 to 23. Because the two sensors perform in different ways, they resulted in different NDVI ranges. The NDVI sensor had a higher resolution range of 0.61-0.86 in comparison with 0.73-0.8 for the multispectral camera, but suffered from wider noisy data due to small canopy coverage of the young apple trees within the FOV (Figure 10b/c). The average correlation between the two sensors was R = 0.85, with correlation of both sensors to water treatments resulting in R = 0.75.

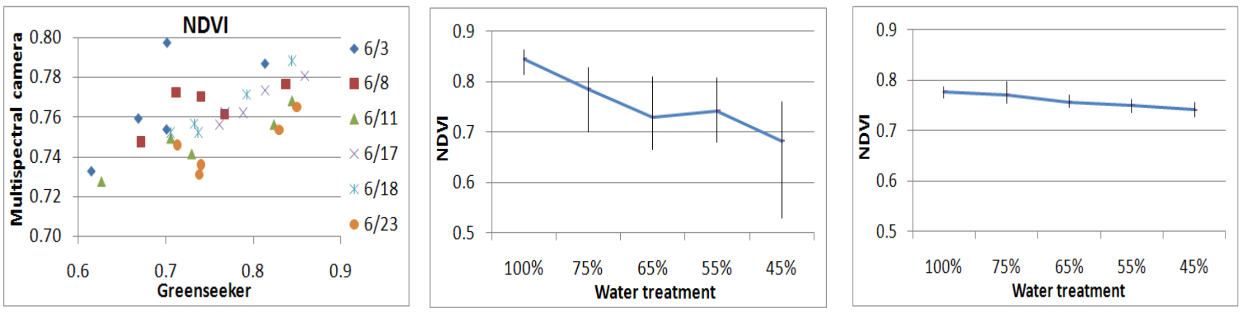

(a) (b) (c)

Figure 10. Canopy spectral response: (a) comparison of NDVI responses between the NDVI sensor and the multispectral camera, (b) NDVI sensor with a range of 0.61-0.86, (c) multispectral camera with a range of 0.73-0.8.

Canopy Temperature

The thermal camera was used to measure plant canopy temperature and determine if plant water stress affects plant canopy temperature. Figure 11 illustrates thermal images of five trees acquired on May 14, 2010 with the temperature range 22.3-40.7 °C. There were clear changes of canopy temperature across the water treatments from cool (bluish) at the 100%-watered trees to hot (yellowish) at the 45%-watered ones. Since the thermal image contains both canopy and background areas and their temperature distributions are different, the canopy area only was selected in rectangular regions of interest (ROI) for the calculation of canopy temperature only. The vertical line next to temperature scale indicates the temperature range of pixels within the ROI.

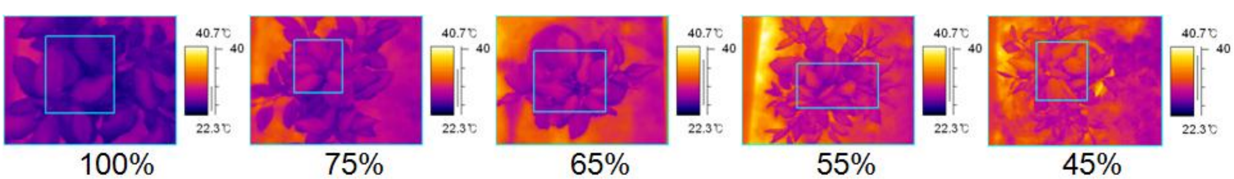

Figure 11. Thermal images of apple trees in five different water treatments: 100%, 75%, 65%, 55%, and 45%. The rectangles in each image indicate the regions of interest (ROI) for the calculation of canopy temperature.

The canopy temperature variations of trees under different water treatments were measured teen times from May 14 to June 22 (Figure 12). All measurements showed the same trend of temperature increase as water availability decreases. Air temperature in the greenhouse was affected by weather and caused different ranges of canopy temperature from 0.75 to 7.13°C in the five water groups. An average 3.9 °C temperature difference was found from most stressed trees (45% water) to the control (100% water). The canopy temperature indicated a strong correlation (R = -0.96) to plant water stress, and thus further research was performed to develop and deploy an array of low-cost infrared temperature sensors for outdoor and real-time applications.

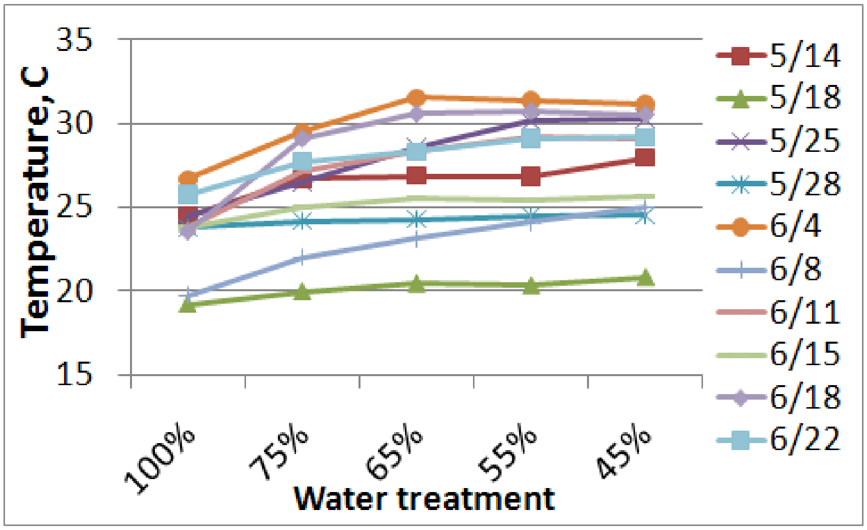

Figure 12. Canopy temperature variation of trees under different water treatments acquired by a thermal camera from May 14 to June 22.

Soil Water Status

Soil water condition was monitored every hour continuously from May through June, 2010 (Figure 13). Higher soil matric potentials indicate more water stress. Average daily soil matric potential remained low, mostly under 20 kPa, in control trees (100% water) indicating well-watered soil conditions. It gradually increased as water treatment decreases. The repeated peaks in the graph were generated by daily watering based on the water treatment plan. Average daily soil water status of five groups is shown in Figure 13 and indicates a well-designed experiment with a clear separation of water treatments in soil (R = -0.89).

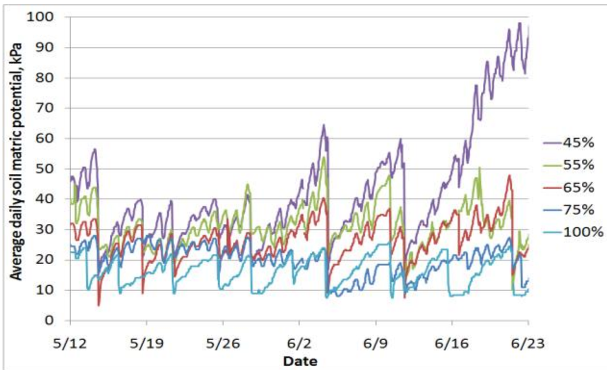

(a)

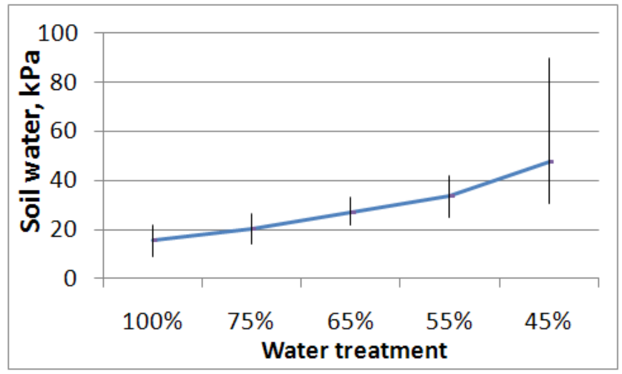

(b)

Figure 13. Soil water status monitored with a Watermark logger. (a) Every hour continuously in May–June, 2009. (b) Average daily soil water status. Higher soil potentials at 45% and 60% indicate water-stressed plants.

Photosynthesis Measurements

Photosynthesis measurements at a tree of each water group were prepared along with the plant water stress experiments to determine photosynthesis reduction by water stress and compare with multimodal sensor data. Each tree was enclosed in a chamber by a polycarbonate board on the side and polyester films on the top and bottom (Figure 14). Air was supplied to the chamber by a motor fan through a 4"-diameter aluminum air duct and ventilated via a 4"-diameter opening in the middle of the top cover. Input reference air at the air duct and output air at 6" below the opening in the top of the enclosure were monitored by a CIRAS I infrared gas analyzer (PP Systems, Amesbury, MA) through a 0.5"-diameter Tygon tube and vacuum pump. Plant carbon dioxide (CO_2) assimilation and water vapor pressure deficit were measured by comparing input and output air by the gas analyzer and recorded by a datalogger (CR7 from Campbell Sci., Logan, UT), as well as temperatures inside and outside the chamber. Air velocity was manually measured in each chamber by taking an average of six readings at 0.5" incremental depths in the input air duct. Air ducts were cut and positioned in order to provide the same air velocity of about 1200 ft/min to all five chambers.

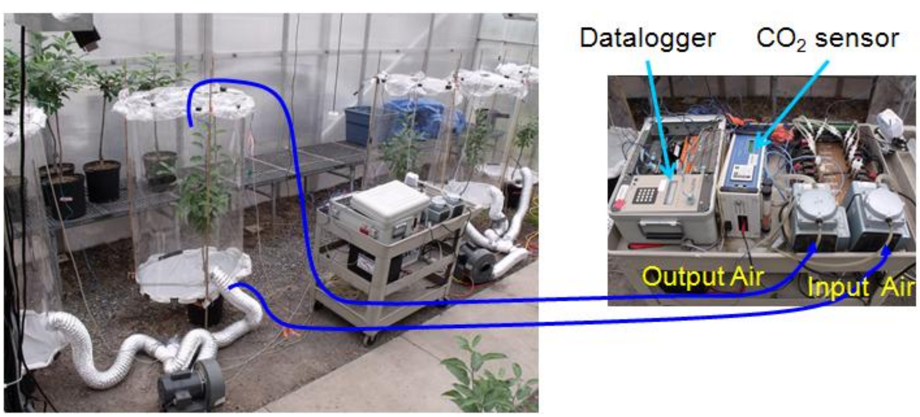

Figure 14. Photosynthesis measurement setup on five trees in different water treatments.

Photosynthesis measurements during the entire period of greenhouse experiments (May 12 to June 21, 2010) are shown in Figure 15a. Reference CO_2 stayed low at around 370 ppm during the day and increased to more than 500 ppm at night due to leaf respiration; CO_2 consumption (delta CO_2) stayed above zero during the day and decreased below zero at night. Daily photosynthesis (g CO_2/h) of five trees under different water treatments and solar radiation on May 15, 2010 is shown in Figure 15b, which shows a treatment separation during the day with high photosynthesis levels at the 100%-watered tree and lower levels as water availability decreases to 45%.

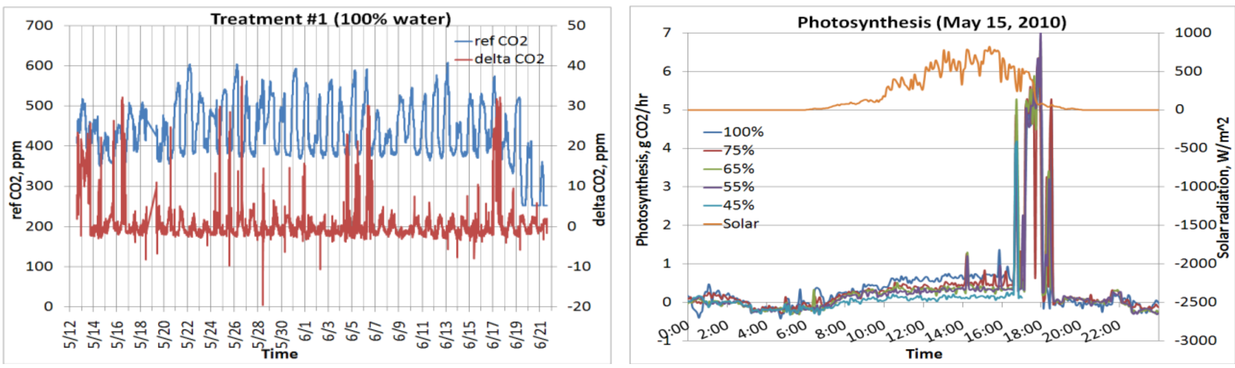

(a) (b)

Figure 15. Photosynthesis measurements. (a) CO_2 consumption in a chamber of treatment #1 tree with 100% water supply during the entire period of greenhouse experiments (May 12 to June 21). (b) Daily photosynthesis of five trees under different water treatments on May 15.

Photosynthesis (Ps) data typically stays below 1 g CO_2/h, but also presents high peaks as shown in Figure 15; we believe these peaks are likely due to sensor noise. Therefore, data was further processed to predict photosynthesis during the noisy period based on filtered Ps, vapor pressure deficit (*VPD*), and photosynthetically active radiation (*PAR*). All photosynthesis data above 1 g CO2/hr and outliers were filtered. The *VPD* was calculated from saturated vapor pressure (*SVP*) as:

$$SVP = 0.0548\,T^2 - 0.8396\,T + 18.353 \tag{1}$$

$$VPD = SVP \times \left(1 - \frac{H}{100}\right) \tag{2}$$

where *T* is air temperature in degrees Celsius and *H* is relative humidity. The greenhouse *PAR* was estimated from solar radiation of a weather station nearby the greenhouse as:

$$PAR = SR \times 2.424 \times 0.75 \tag{3}$$

where *SR* is solar radiation in W/m^2, 2.424 is a unit conversion factor from W/m^2 to µmol/m^2/s, and 0.75 is a conversion factor from ambient *PAR* to greenhouse *PAR*.

These three values, filtered Ps, *VPD*, and greenhouse *PAR*, were used for multiple regression analysis to generate parameter estimates for predicted photosynthesis. The predicted photosynthesis value closely matched the measured ones (Figure 16a) and used to substitute noisy data measurements. Figure 16b supports the successful estimation with clear separation of water treatment levels by the predicted Ps during the entire period of experiments.

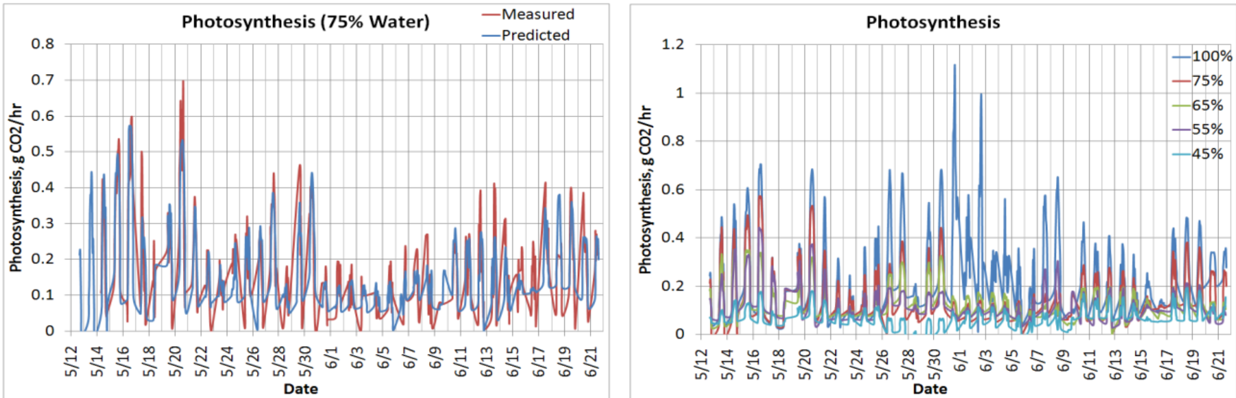

(a) (b)

Figure 16. Predicted photosynthesis. (a) Comparison with measured photosynthesis at the tree under 75% water treatment. (b) Predicted photosynthesis of five trees under different water treatments.

The predicted photosynthesis data was further processed to calculate daily photosynthesis reductions at each water treatment as a ratio to the average daily Ps of control trees (see Table 4 and Figure 17). Photosynthesis reduction varied from 15% to 69% with an average of 36% in 75%-watered trees and increased up to 73% in the least watered trees (45% water). Daily photosynthesis reductions of the trees created by water stress were compared with multimodal sensor data: NDVI response of GreenSeeker (Figure 17a), NDVI provided by the multispectral camera (Figure 17b), and canopy temperature of the thermal imager (Figure 17c). Correlation of multimodal sensors to daily photosynthesis reduction is listed in Table 5. The highest correlation to photosynthesis reduction was found in canopy temperature (R = -0.91), followed by GreenSeeker NDVI (R = 0.87) and multispectral camera (R = 0.8).

Table 4. Photosynthesis (Ps) reduction by water stress as a ratio to average Ps of control trees (100% water).

Water treatment	5/13	5/14	5/25	5/28	6/4	6/8	6/11	6/15	6/18	average
100%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
75%	15%	24%	26%	43%	65%	69%	28%	33%	24%	36%
65%	28%	25%	54%	25%	50%	75%	38%	37%	52%	43%
55%	40%	27%	50%	36%	62%	55%	53%	43%	53%	47%
45%	72%	68%	65%	104%	72%	85%	57%	69%	63%	73%

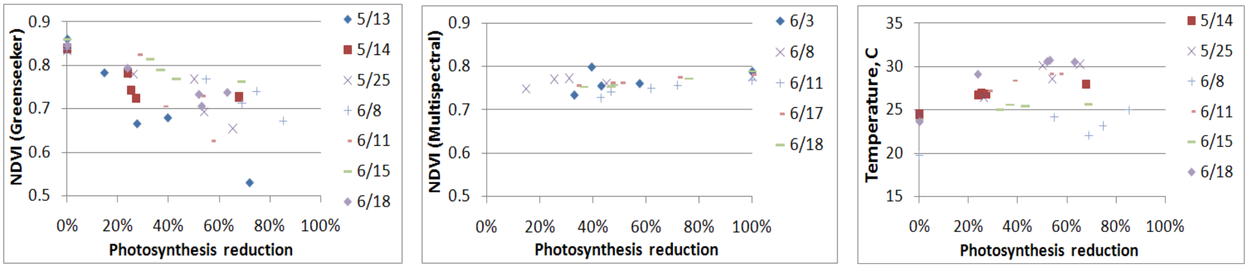

(a) (b) (c)

Figure 17. Photosynthesis reduction by water stress compared with: (a) NDVI response of GreenSeeker, (b) NDVI response of the multispectral camera, (c) canopy temperature of the thermal imager.

Table 5. Correlation of multimodal sensors to photosynthesis (Ps) reduction by water stress.

	GS	ADC	Ps	Thermal	SoilW
GreenSeeker (GS)	1.00				
ADC multispectral camera	0.85				
Photosynthesis (Ps)	0.87	0.80			
Thermal imager	-0.86	-0.89	-0.91		
Soil water status (SoilW)	-0.75	-0.88	-0.74	0.77	1.00

Field Experiment

The outdoor experiment was conducted at the PSU-FREC with 10 irrigated and 10 non-irrigated apple trees during the period August 11 through September 24, 2010. The irrigated trees were regularly irrigated every 2-3 days based on conventional soil water management practice. The multimodal sensor system was installed on a Toro eWorkman utility vehicle with extended support pipes and positioned to aim the side of the tree canopy (Figure 18). The sensor system includes a multispectral camera, two sets of NDVI sensors with a digital camera and an ultrasonic range finder, an IR thermometer array, and a thermal imager. Solar radiation, air temperature, and relative humidity were separately recorded by a WatchDog datalogger. Soil water status was also monitored using soil moisture sensors that were installed in 6”-deep soil for both irrigated and non-irrigated trees.

The multispectral camera and thermal imager acquired images on a static mode and were used as references for the NDVI sensor and IR thermometer array, respectively, running on continuous mode. Preliminary data from the multimodal sensor system from the non-irrigated apple trees acquired on August 17, 2010 is shown in Figure 19. A complete data analysis is under process and will be reported in the first period of Year 3.

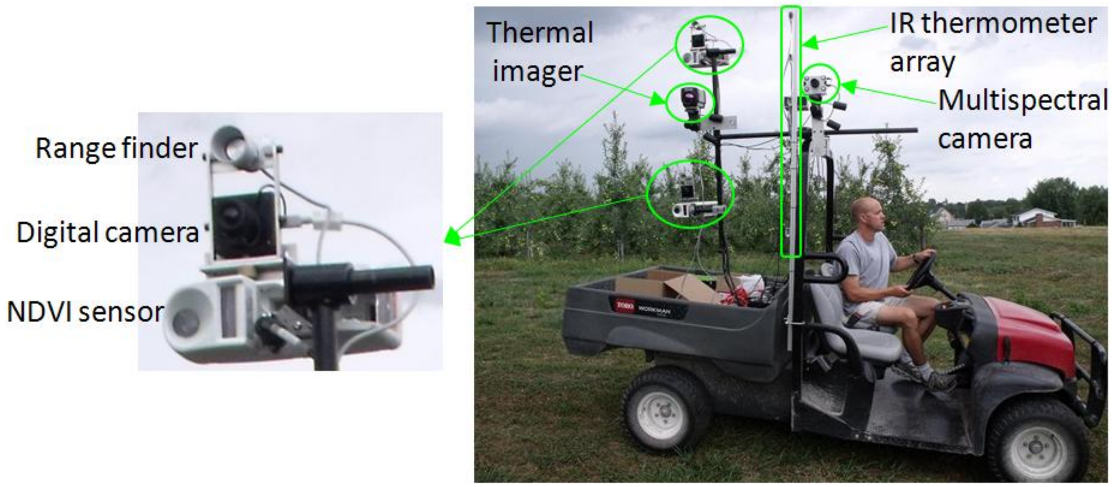

Figure 18. Multimodal sensor system mounted on a Toro utility vehicle for plant stress detection by measuring canopy temperature and reflectance.

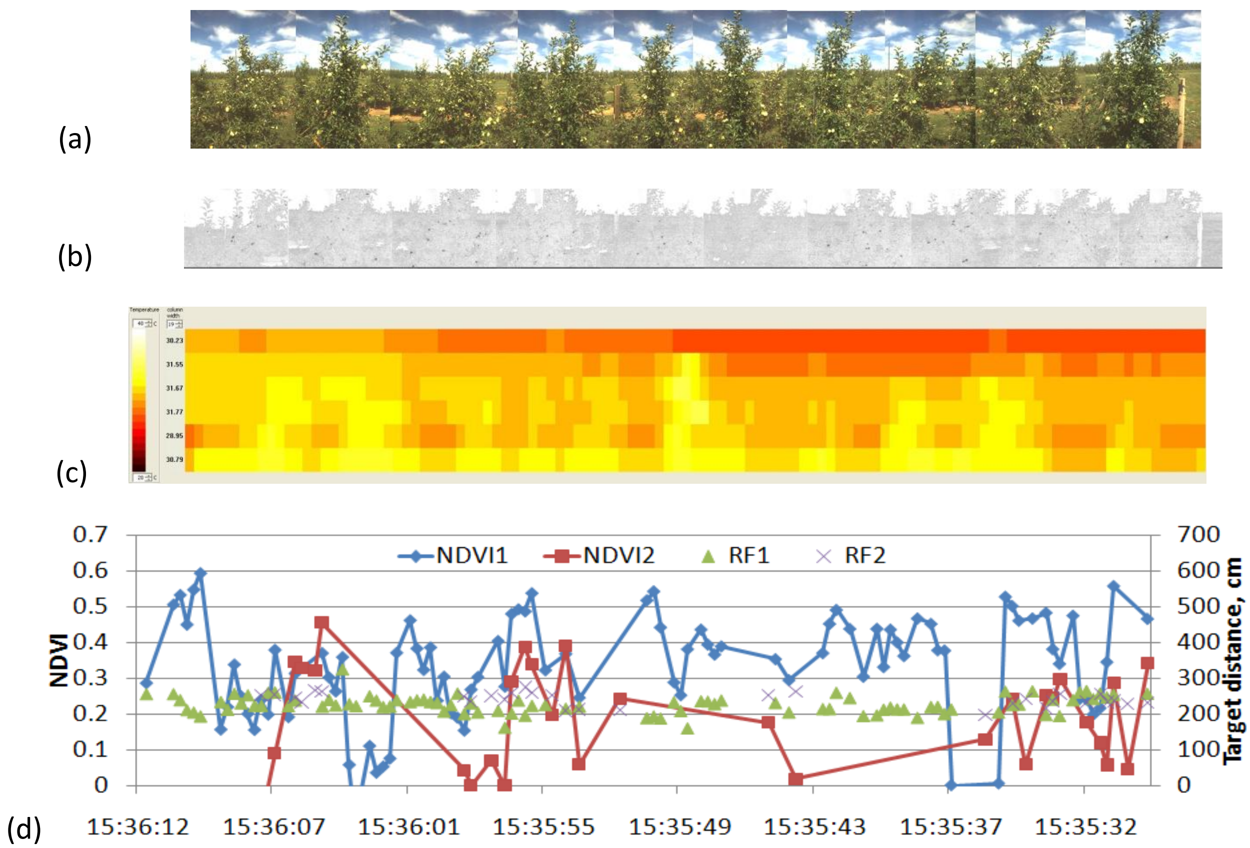

Figure 19. Multimodal sensor data acquired at non-irrigation trees on August 17, 2010. (a) Digital color image, (b) gray-scaled NDVI image of the multispectral camera, (c) 2-D canopy temperature map of the IR thermometer array, (d) NDVI responses of two GreenSeeker sensors filtered by a rangefinder (RF).

Fire Blight Detection

We acquired a large number of images of fire blight (FB), both in early stage (i.e., when only leaf veins are brown) and in advanced stage (i.e., when the entire leaf becomes brown). Table 6 summarizes the number of fire blight images taken in the greenhouse and in the field. In addition to the fire blight images, we also acquired approximately 600 images of healthy trees in the field. Figure 20 shows some examples of the images in the database.

Table 6. Number of images in database for fire blight detection.

	Early stage damage	Late stage damage	Total
Greenhouse	137	144	281
Field	189	335	524
Total	326	479	805

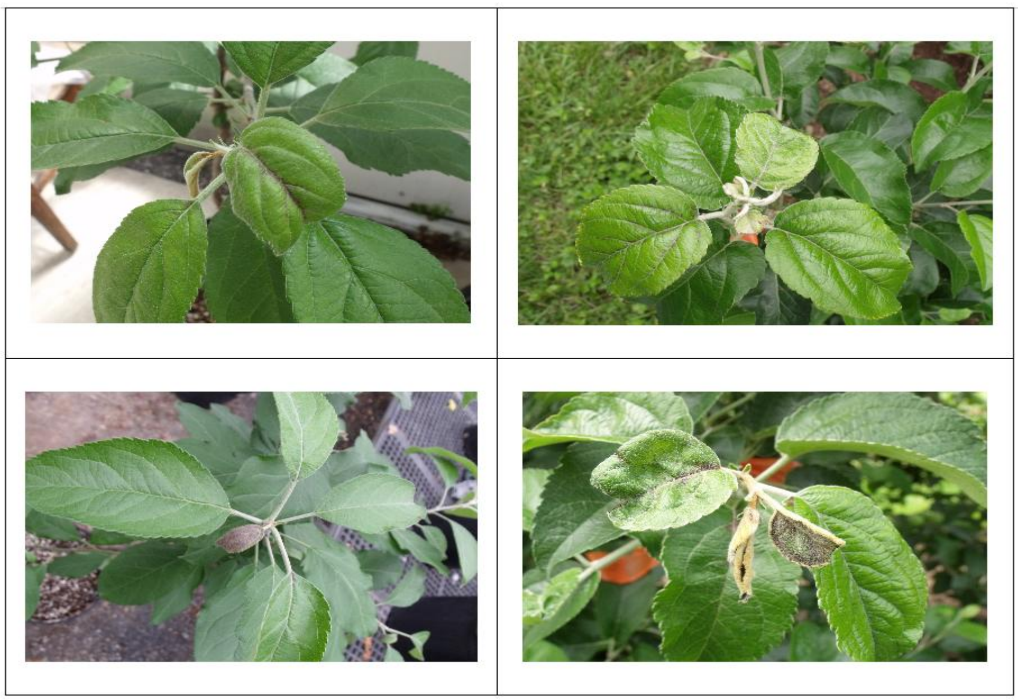

Figure 20. Sample images of fire blight-infected trees.

In Year 2 we tried several different image processing methods to automatically detect fire blight symptoms. Unfortunately, none of the methods performed well. In this section, we describe two of the methods we have tried and their results.

- Method 1: Leaf detection followed by fire blight detection

Recognizing that early-stage fire blight symptoms appear on leaf veins, this method first extracts leaf regions in the image, then attempts to detect fire blight. A simple color-based support vector machine (SVM) (Joachims, 1999) is applied to extract leaves in the image. A second SVM-based classifier attempts to detect fire blight around the extracted leaf regions in the image. Some qualitative results of this method are shown in Figure 21 and Figure 22. In these figures, the top row shows the original images, the pink regions in the center row represent the detected leaf areas, and the white regions in the bottom row represent the pixels detected as non-fire blight regions. All other (i.e., non-pink and non-white) pixels are the regions detected as fire blight-infected. Although this method does not give rise to false detections in the background, several false detections occur around the edges of the leaves.

Figure 21. Fire blight detection results on greenhouse images using Method 1.

Figure 22. FB detection results on field images using Method 1.

- Method 2: SVM using a histogram of hue values

In this method, a 36-bin histogram of hue values in an image window is used as a feature vector. Similar to other SVM-based methods, a large number of positive and negative examples are used to train the SVM classifier. Some qualitative results of this method are shown in Figure 23. Similarly to other methods, this one presents a high rate of false detections. Without considering the background, this method yields a detection rate of 76.6% with a false detection rate of 58.9%.

In Year 3 we expect to improve the results obtained by deploying other sensors, in particular forward-looking infrared cameras.

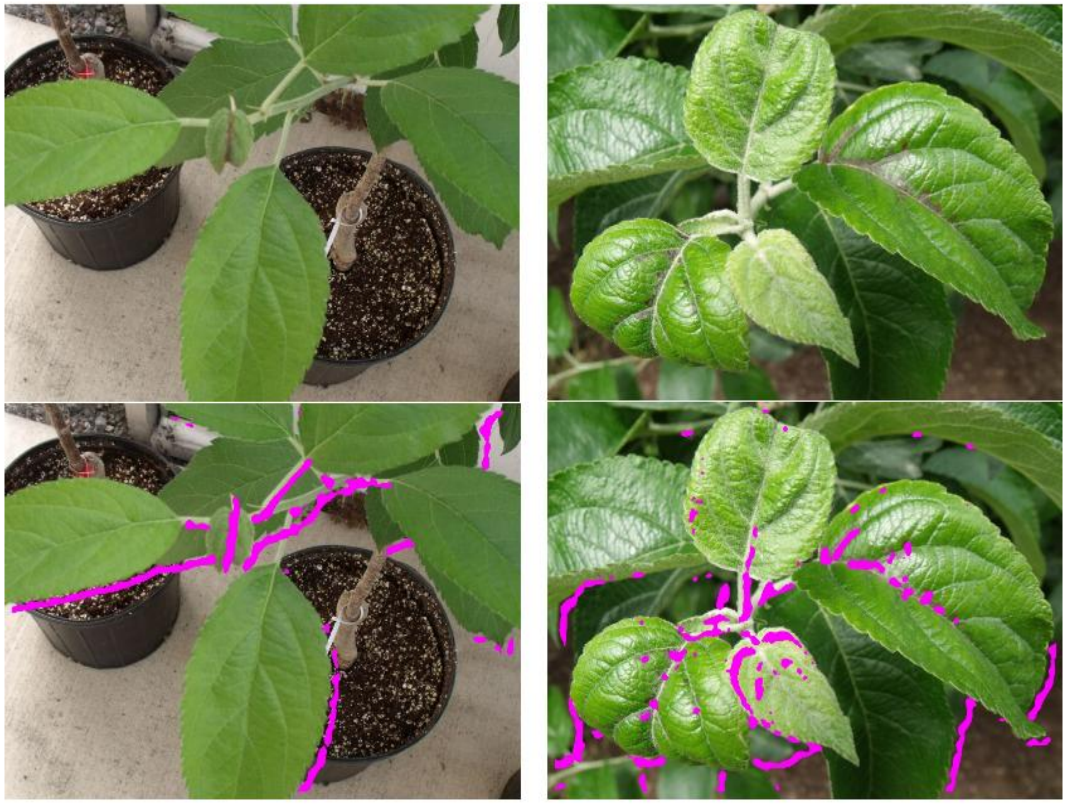

Figure 23. Fire blight detection results using Method 2

References

- Kim, Y., D. M. Glenn, J. Park, H. K. Ngugi, and B. L. Lehman. 2010a. Hyperspectral image analysis for plant stress detection. ASABE Paper No. 10-09114. St. Joseph, MI: ASABE.
- Potato Explorer. 2010. Leaf area index. Available at: www.potato.nl/explorer/pagina/leafarea.htm. Accessed 20 April 2010.
- Kim, Y., D. M. Glenn, J. Park, H. K. Ngugi, and B. L. Lehman. 2010ba. Active spectral sensor evaluation under varying conditions. ASABE Paper No. 10-09111. St. Joseph, MI: ASABE.
- Joachims T. 1999. Chapter 11: Making large-Scale SVM Learning Practical. Advances in Kernel Methods - Support Vector Learning, B. Schölkopf and C. Burges and A. Smola (ed.), MIT-Press.

3.2 Insect Monitoring

Thematic area leaders

Name	Institution	Email
Larry Hull	The Pennsylvania State University	lah4@psu.edu
Vincent Jones	Washington State University	vpjones@wsu.edu
Johnny Park	Purdue University	jpark@purdue.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Improve electronic pheromone monitoring traps and algorithms for detecting moths. 2. Test trap under various orchard and pest conditions. 3. Compare electronic trap to current standard trap. 4. Improve computer vision algorithms for fruit injury detection. 5. Explore the use of stereo vision for fruit injury detection.	1. Reports on the orchard studies of automatic moth detect technology. 2. Reports on the results from IFM fruit injury detection studies. 3. Publicly available dataset of images with accompanying ground truth.	1. 10 traps installed in WA and PA. 2. Ability to automatically count moths entering the trap with > 90% accuracy. 3. Ability to ID moth and limit non-target species with > 70% accuracy while withstanding various orchard conditions. 4. Ability to automatically detect IFM-damaged apples with > 80% detection rate and < 20% false alarm rate.

Notable results:

- Developed two models of electronic trap prototypes: one based on IR sensors (IR-trap) and the other based on bio-impedance sensors (Z-trap).
- Deployed six IR-traps and twenty Z-traps in orchards in WA and PA, and carried out field experiments.
- Determined that the IR-traps have a low-capture rate compared to standard delta traps.
- Determined that the original Z-traps also have a low-capture rate. After modifying the trap shape and the size of the high-voltage coil, the new Z-traps achieved a similar or better capture rate compared to standard delta traps.
- In the laboratory wind tunnel, the Z-traps achieved >95% detection accuracy. Preliminary field data results seem to indicate that the new Z-traps achieved >80% detection accuracy with <5% false alarm rate.

- Acquired an additional 5,800 images of internal feeding worm (IFW)-damaged apples and healthy apples across four varieties.
- The latest IFW damage detection algorithm achieves >85% accuracy with <4% false alarm rate.
- Created a database that contains the images of IFW-damaged and healthy apples collected in 2009 and made it available on the web.

Bio-Impedance Sensor-Based Traps (Z-trap)

In collaboration with Spensa Technologies and partial funding from the Washington Tree Fruit Research Commission, we have devised a new electronic trap based on a bio-impedance sensor, the "Z-trap." A schematic representation of the Z-trap is shown in Figure 24. A lure is located inside a cylinder-shaped high-voltage electric grid, which is the bio-impedance sensor. When an insect, attracted by the lure, touches the electric grid, it gets electrocuted and falls down to the bottom collector. The bottom collector is designed in such a way that it is easy for temporarily-stunned pests to fall through the top opening of the collector but difficult for them to get out of it. Whenever an electric discharge occurs (i.e., an insect gets electrocuted), the microcontroller records the time of the event. Furthermore, the microcontroller analyzes the properties of each electric discharge (i.e., the amplitude and duration of the electric current pulse) in order to distinguish whether the event was caused by a target insect or by a non-target insect.

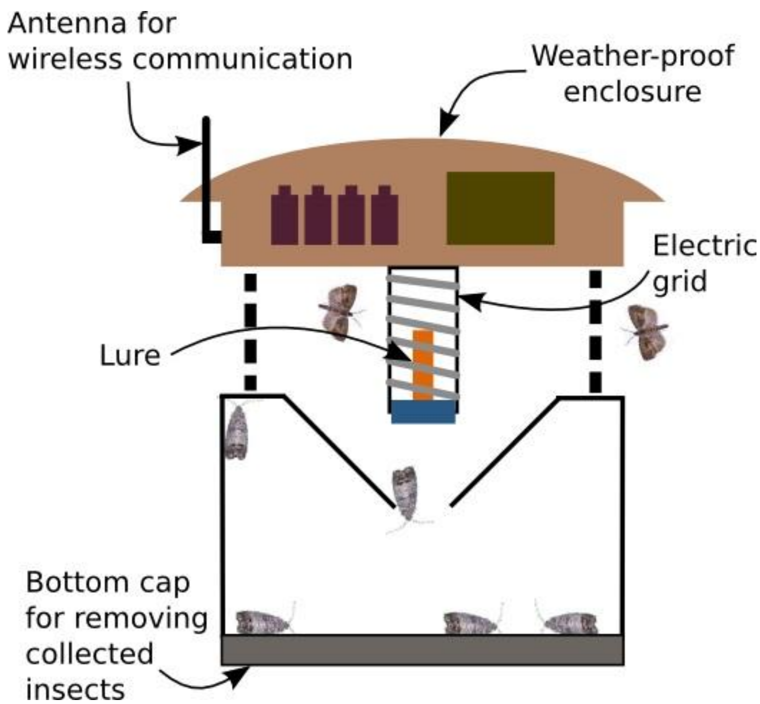

Figure 24. Schematic design of the bio-impedance based trap (Z-trap).

The high-voltage electric grid consists of a pair of metallic coils spaced approximately 1/5" apart from each other (although the inter-coil spacing may vary according to the target insect species). An electric voltage is applied to each coil, but since they form an open circuit, there is no current flow. As an insect approaches/touches the grid, the circuit is closed and an electric current flow is initiated, which electrocutes the target insect. Figure 25 shows a picture of the grid. By eliminating a supporting frame for the coils on this grid design, we were able to avoid accumulation of chemicals used in the field that could potentially short-circuit both coils. In addition, there is no non-conductive landing surface in the grid; therefore, target insects cannot touch it without being electrocuted and hence detected.

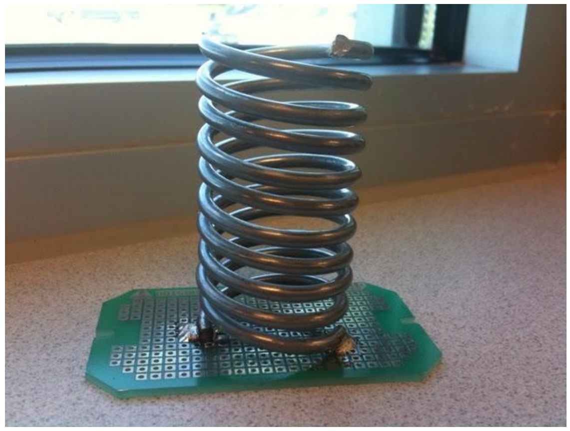

Figure 25. Zapper coils prototype.

Based on several laboratory experiments, we found voltage levels that allow the target insects to be temporarily stunned whenever they contact the electric grid without sticking to its surface. For both codling moth (CM) and obliquebanded leafroller (OBLR), the average voltage level was around 750 V, whereas for the oriental fruit moth (OFM), it was around 450 V. After finding the optimal voltage levels for the detection of the different insect species, a prototype was built and tested in the wind tunnel at the Washington State University Tree Fruit Research and Extension Center (WSU-TFREC), in Wenatchee, WA with adult CM and OFM. The experiments showed that the zapper trap is capable of accurately counting the number of insects captured, obtaining >95% detection accuracy in this controlled environment.

In June 2010, ten Z-traps were deployed in an experimental orchard at the WSU-TFREC and ten were deployed at the PSU-FREC. Figure 26 shows one of the traps deployed in WA. For a period of three months, the traps were used to monitor populations of CM and OFM in the orchards. During this period, the performance of the traps was carefully monitored and the data generated by each trap was periodically collected. For comparative purposes, the same number of standard delta traps was deployed in nearby locations within the same orchards.

Although the Z-traps proved functional and operated uninterruptedly without major problems for the entire test period, they achieved significantly lower capture rates than standard delta traps. This is in contrast with the results of the wind tunnel experiments carried out at WSU, which achieved a comparable, if not better, capture rate than the conventional traps. We believe that the exterior shape of the current trap is somehow disrupting the dispersion of the pheromone plume. We evaluated the initial trap design in the wind tunnel and in the field. In the field, during the month of July, the original trap design caught between 8 and 11% of the moth catch of standard delta traps. In the wind tunnel, we used smoke (i.e., stannic oxychloride) to evaluate the pheromone plume emitted from the trap. We found that the shape of the trap caused a vacuum downwind from the large top, which caused the pheromone to curl back to the top, so that moths spent more time around the top of the trap and did not approach the zapper coil. The lower part of the trap was also problematic, because moths that went below the bottom part of the trap also lost the pheromone signal and were unable to locate the plume again, which resulted in the moths staying below the lower disk.

Figure 26. Z-trap prototype (left); A Z-trap deployed at WSU-TFREC orchard for field testing (right).

In order to identify and understand the reasons behind the low catch rate of the new trap prototype, several modifications to the design of the external structure of the trap were evaluated, some of which are shown in Figure 27. We tested eight different modifications of the trap either by modifying the bottom, or by incorporating portions of a bucket trap, which appeared to improve airflow around the trap and moth capture in the wind tunnel. We also evaluated two other modifications, one a bucket trap with the zapper coil attached at the top, and the other a standard delta trap attached over the coil, as illustrated in Figure 28. Because we only had a limited number of traps to test, these parts of the test were replicated only over time. From early August to September 8th, we were able to release large numbers of sterile adult CM from Canada roughly once a week in an orchard adjacent to the lab at the WSU-TFREC.

Figure 27. Experimental trap exterior designs.

In the field, we found that the bucket trap modified with a coil (Figure 28, left) captured ≈56% (85 moths) of the total delta trap capture (152 moths), and on most days was within 1-2 moths of the delta trap. On one date where the delta trap was placed high in the tree (instead of at the same level as the electronic trap), trap capture was roughly 2.5 times higher (62) than the bucket trap modification (24). In the WSU-TFREC experiments, the delta trap modification was less successful, at least in part because the coil was of smaller gauge wire and was easily distorted by the trap, potentially resulting in the zapper short-circuiting. The other modifications were consistently less efficient than either the delta or bucket modifications, despite not having the same issues with the coil. The bucket trap modification shows the potential of the zapper, and would easily allow for the combination of the zapper with the IR traps. Future designs of the zapper should simply place the electronics into a convenient, water-proof enclosure with the zapper coil on a cable that can easily be attached to any trap design for testing purposes to determine what works best. We anticipate that this design could be easily tested in the lab wind tunnel during the winter and modified relatively easily without having to re-engineer the trap design.

Figure 28. Modified Z-traps employing existing bucket trap (left) and delta trap (right) exterior designs.

Figure 29 (oriental fruit moth) and Figure 30 (codling moth) show some quantitative results obtained in the orchard at PSU. The figures show the cumulative capture of two modified Z-traps with delta bottoms as well as the corresponding nearby standard delta traps used to evaluate the efficiency of the Z-traps. As one can see, the number of OFM and CM captured in the modified Z-trap with the delta bottom (Figure 28, right) easily surpassed the number captured by the manual delta trap at PSU.

During the field experiments—conducted at WSU and PSU to improve the Z-trap exterior—we also collected additional electronic and environmental data. The data collected is accompanied by ground truth information based on the daily observations made in the field. That is, along with the digital signals obtained by the sensor, the data included the number of target and non-target insects captured by each trap as well as the corresponding weather information. Unlike the digital signals obtained in the wind tunnel experiments, which contained only target insect detections, the signals obtained in the field experiments also contained a large number of other unintended/unwanted events caused by things such as non-target insects, meteorological conditions or spraying events. These signals must be pre-processed to remove the unwanted events before they can be used to count the actual number of detected pests.

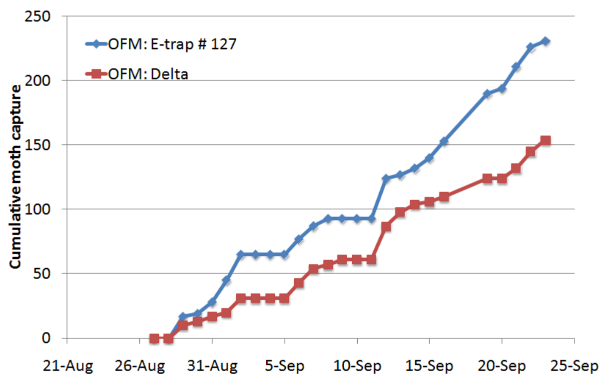

Figure 29. Cumulative OFM capture of one of the zapper traps with a delta bottom compared with the captures of a nearby large delta trap.

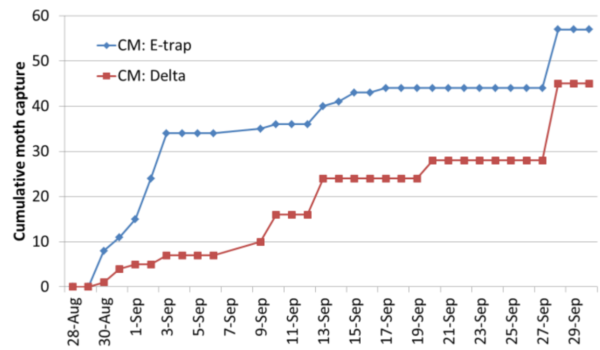

Figure 30. Cumulative CM capture of one of the zapper traps with a delta bottom compared with the captures of a nearby large delta trap.

Processing and analyzing the massive amounts of data produced during the field experiments is a formidable task. The information obtained by the traps consists of hundreds of millions of data points and several thousand events of interest. Just to put in perspective, using a regular laptop computer, it takes approximately 10 minutes to load and plot the data corresponding to two weeks of monitoring by a single trap. Manual inspection of the data allowed us to initially identify some of the signal features such as pulse width and slope that could be employed to distinguish between target insects, non-target insects, and irrelevant events, but developing algorithms to carry out these tasks and robustly filter out undesired data will require further investigation.

Infrared Sensor-Based Traps (IR-Trap)

We have also developed a new generation of infrared sensor based traps ("IR-traps"). Figure 31 shows the new IR-trap deployed at the PSU-FREC. It has a small weather-resistant enclosure on the bottom that contains the electronics and batteries. Figure 32 shows how the circuit board was mounted on the back of a 3-cell D battery holder and placed inside the enclosure. The sensor funnel was mounted inside the bucket trap, as shown in Figure 33, and the sensors were connected to the circuit board internally with a pair of ribbon data cables.

Figure 31. IR-trap deployed at PSU-FREC.

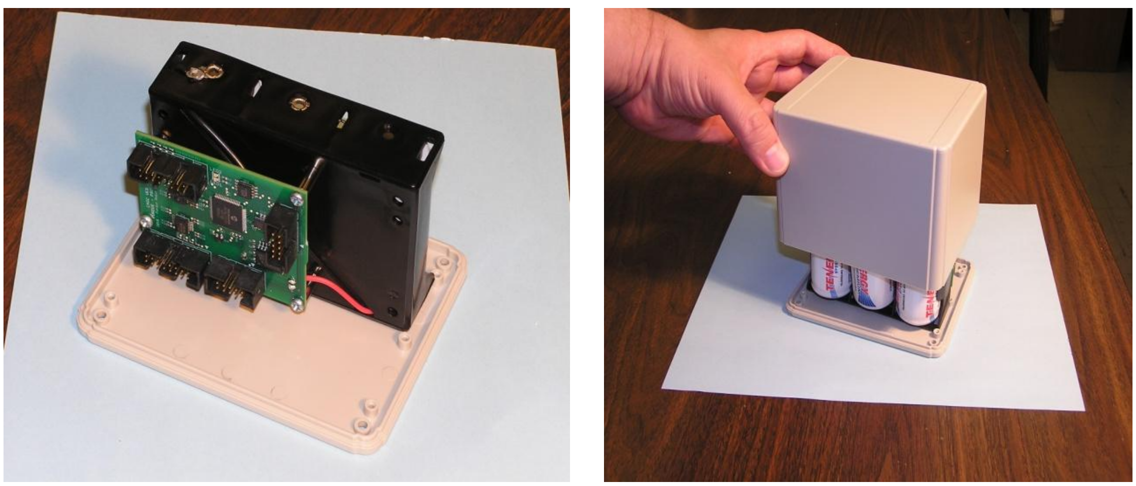

Figure 32. Circuit board and batteries inside the enclosure.

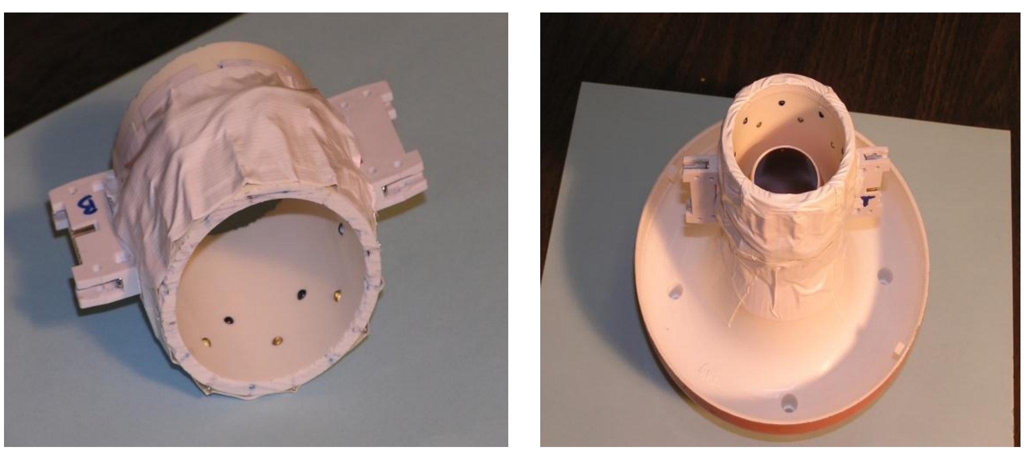

Figure 33. Sensor funnel.

The new IR-trap has eight IR emitter/receiver pairs; last year's model had four pairs. The IR sensors can be operated in two ways: as a single ring of eight emitter-receiver pairs, or as two rings of four emitter-receiver pairs. Using a single ring allows for detecting smaller objects, and using two rings allows for determining whether the object is moving upwards or downwards.

A 32 Mbyte flash memory integrated circuit board was added to the IR-trap in order to collect raw data. Every event is recorded with a time stamp and can be transferred later to a PC for further analysis.

Power efficiency was significantly improved in the new IR-trap. Some of the new features for increasing the power efficiency include individual power control strategies over every single peripheral and a new signal conditioning circuit that works without the need of inverters or dual sets of batteries. Table 7 compares the power consumption between the Year 1 and Year 2 prototypes. In order to make the comparison fair, the Year 2 prototype is not using any algorithm for smart duty-cycling and all sensors were kept active at all times to compare only the efficiency improvement in the electronics.

Table 7. Power consumption for Year 1 and Year 2 prototypes for one day of operation.

	Operation Time Power Consumption	Idle Time Power Consumption	Average Operation Time
Year 1 prototype	400 mA	30 mA	2 days
Year 2 prototype	35 mA	0.2 mA	60 days

In order to maximize the trap operation time, we designed a solar battery charger and controller. There is a wide range of solar energy products on the market. Integrating a solar panel solution to a customized electronic application, however, is still a difficult task. This year, we have successfully designed, implemented, and tested a solar charger and controller prototype. Integration of the solar charger and the electronic trap is part of our future work.

We deployed six IR-traps at PSU-FREC in early September. Unfortunately, low capture rate was again the main problem. Figure 34 and Figure 35 show the cumulative capture for OFM and CM, respectively. The number of moths captured by larger delta traps used as controls is shown in red and that by the IR-traps is shown in green. We believe the bucket traps in general are not effective in capturing OFM and CM, unless there is a mechanism to force the insects to fall through the funnel (e.g., by stunning them with a high-voltage electrical grid similar to the one in the Z-trap). We also believe that the Vapona™ "kill" strip placed inside the bucket trap (i.e., as was used in the 2009 field studies) likely repels some moths from entering the trap.

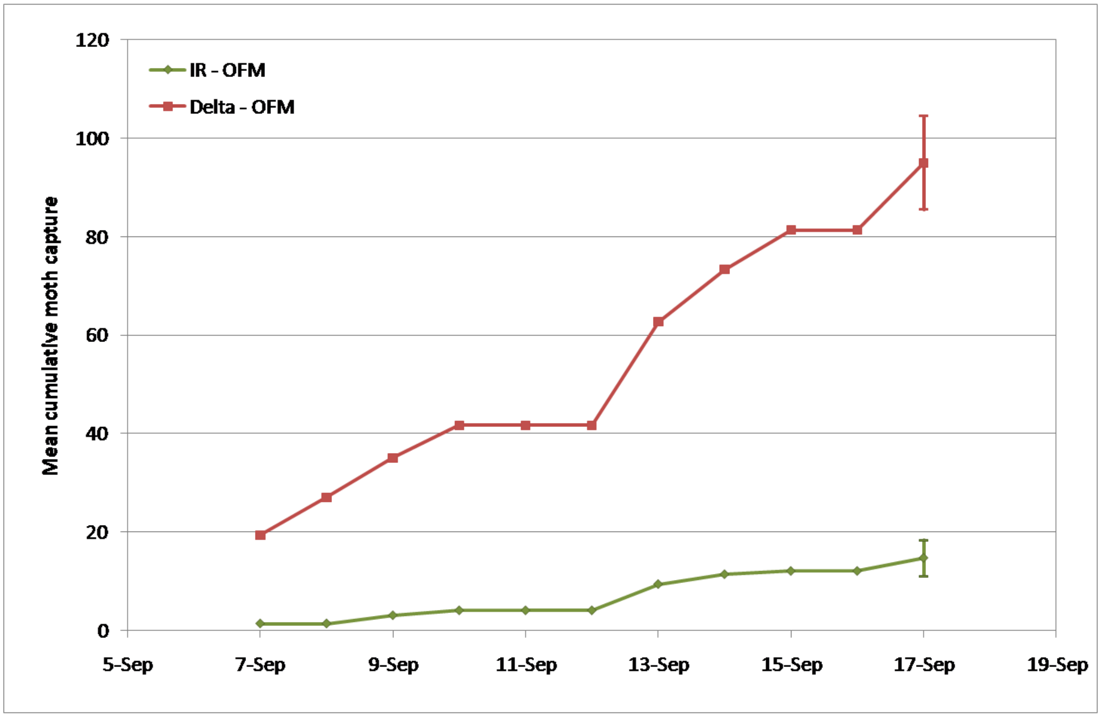

Figure 34. Mean cumulative capture of OFM for three IR-traps.

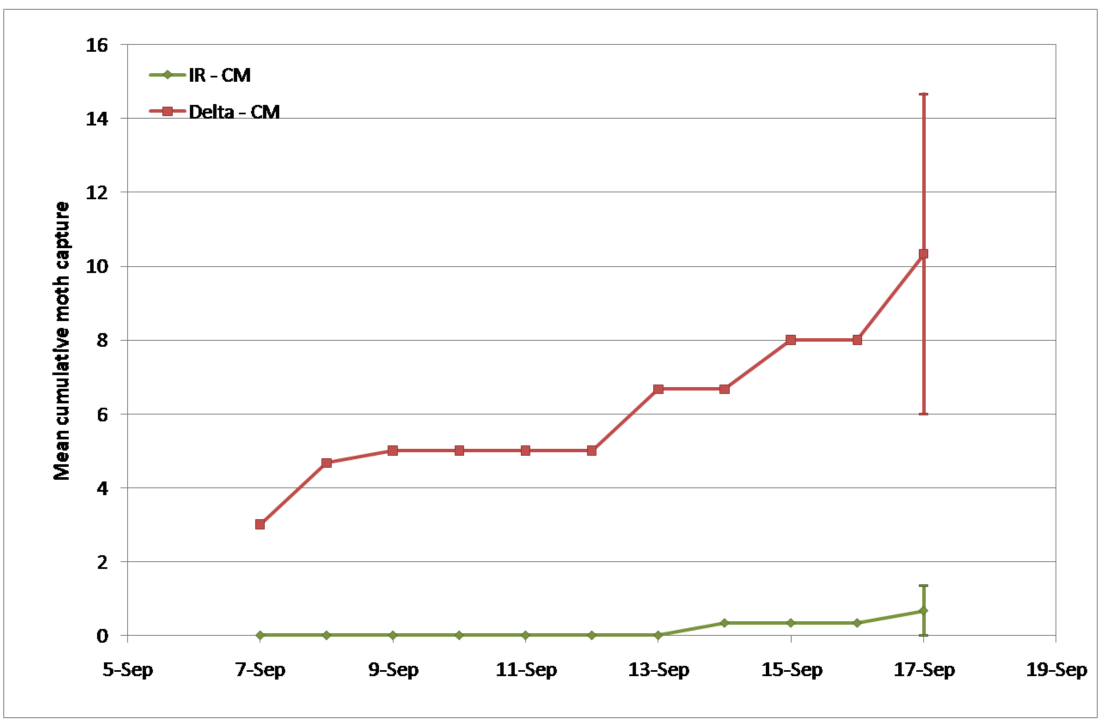

Figure 35. Mean cumulative capture of CM for three IR-traps.

Automated Detection of Internal Feeding Worm (IFW) Damage

We have developed a two-step classification algorithm for detecting IFW damage on apples. A fast classifier is used in the first step to quickly remove most of the non-damaged regions. In the second step, a multiple kernel learning (MKL) approach is used on the remaining regions to obtain the final IFW damage detection.

Figure 36 shows some examples of IFW-damaged apples across three varieties in different maturation stages. These images illustrate the wide variability in color of the apple and in the size and shape of IFW damage regions. These images also show the complexity of the image background that includes leaves, branches, ground, sky, etc. All of these factors make IFW damage detection a challenging task.

Figure 36. Examples of IFW-damaged apples.

Figure 37 shows the four fundamental operations of the current IFW damage detection algorithm: (1) feature extraction, (2) feature dimensionality reduction, (3) first-step classification based on support vector machine (SVM); and (4) second-step classification based on multiple kernel learning (MKL). Given an image, color and gradient features are extracted first. The dimension of the feature vector is reduced to 20 using principal component analysis (PCA). A two-step classification approach is applied after feature extraction and dimensionality reduction. In the first step, a standard SVM classifier with color and gradient features is utilized. After this first step classification, normally over 80% of the non-damaged regions are removed, thus only a small portion of the image regions (candidate regions) need to be evaluated in the second step. A more complicated MKL classifier (Figure 38) is used to detect the damaged regions from the candidate regions in the second step. A histogram of gradients (HoG) and color-shape features are calculated on a 21 x 21 patch surrounding the candidate region. The color-shape feature includes the average RGB intensities inside the candidate region, the average RGB intensities surrounding the candidate region, the difference of the RGB intensities between the candidate region and its surroundings, the size of the candidate region, and the aspect ratio of the candidate region.

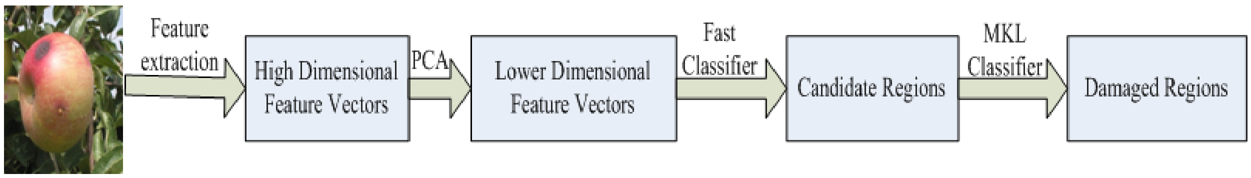

Figure 37. IFW damage detection algorithm.

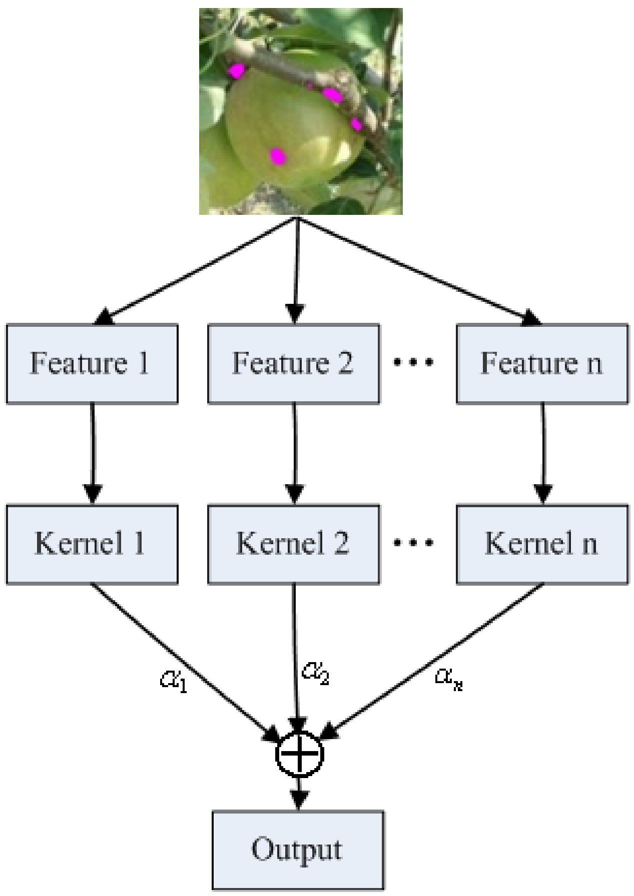

Figure 38. Multiple kernel learning classifier procedure.

In 2009, we collected a total of 2,774 apple images across three varieties to test our IFW detection algorithms (Table 8). These images are available for other research groups pursuing similar image processing problems at http://cobweb.ecn.purdue.edu/RVL/Database/IFW/. Half of the images in the 2009 database were used to train the algorithms and the other half to test them. Figure 39 through Figure 41 show some examples of detection results for Fuji, Golden Delicious and York. In these figures, the left image is the original image and the right image shows the detection result with pink regions detected as IFW damage.

Table 8. IFW damage image database collected in 2009.

	Damaged	Healthy	Total
Fuji	493	499	992
Golden Delicious	490	401	891
York	461	430	891
Total	1444	1330	2774

Figure 39. IFW damage detection results for Fuji apples.

Figure 40. IFW damage detection results for Golden Delicious apples.

Figure 41. IFW damage detection results for York apples .

Figure 42 shows the receiver operating characteristic (ROC) curves of the previous and the new detection algorithms. (An ROC curve is a plot of the sensitivity of a binary classifier when its discrimination criterion threshold is varied. The curve can be obtained by plotting the true positives vs. the false positives of the classifier under test. ROC analysis is usually regarded as a cost-benefit analysis because it allows one to select an optimal model and discard those that are suboptimal.) ROC curves generated from the test results to evaluate the first-step classifier, the single kernel classifier (SKL), and the multi-kernel classifier (MKL). The red line represents the ROC curve for the first-step classification, the blue line corresponds to the second-step classification using a SKL classifier, and the pink line the second-step classification using a MKL classifier. As one can see, the MKL-based classifier achieves about 85% detection accuracy with less than 4% false alarm rate.

During the 2010 season, the PSU team acquired additional images of IFW damaged apples and healthy apples across four varieties (Table 9). We plan to use these images to retrain the detection algorithms in Year 3.

Table 9. IFW damage image database collected in 2010.

	Fuji	**York**	**Golden Delicious**	**Red Delicious**	**Total**
IFW Damaged	742	1061	1022	1005	3830
Healthy	394	647	422	496	1959
Total	1136	1708	1444	1501	5789

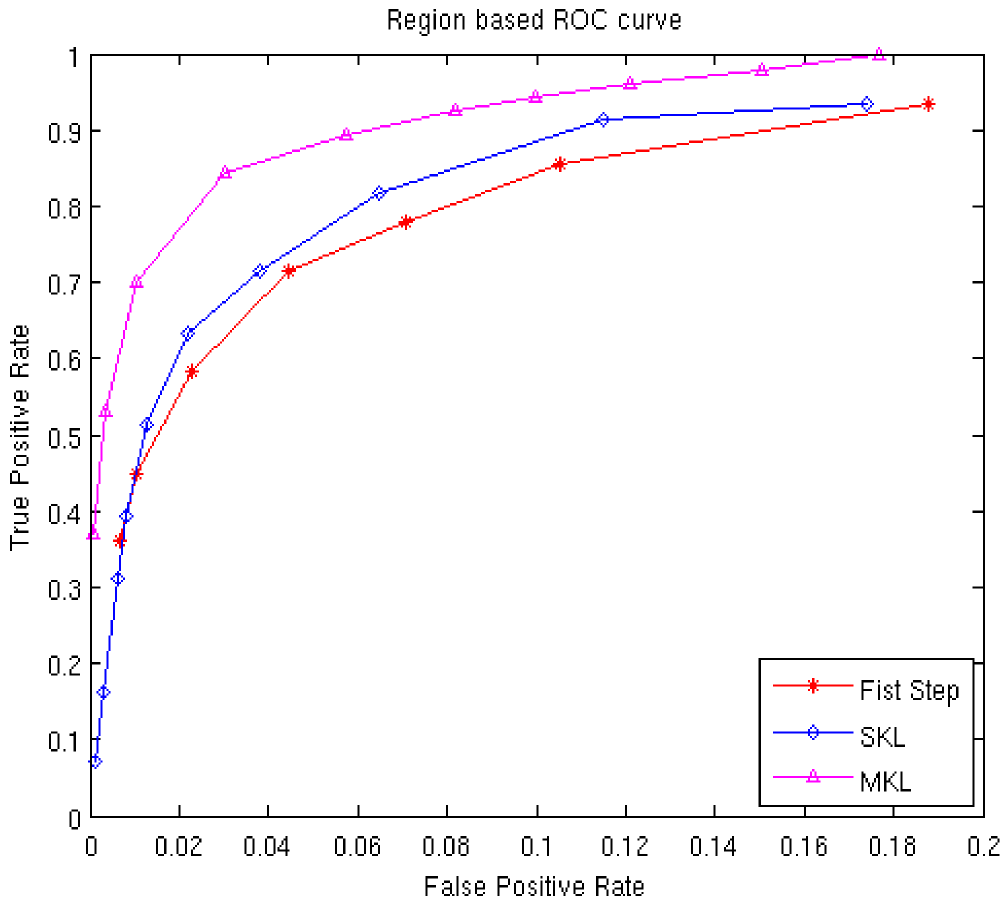

Figure 42. ROC curves for the IFW damage detection algorithms.

3.3 Crop Load Scouting

Thematic area leader

Name	Institution	Email
Tony Koselka	Vision Robotics Corp.	tkoselka@visionrobotics.com

Year 2 goals

Activities	Deliverables	Success Criteria
1. Improve hardware for more reliable field operation. 2. Enhance detection and sizing performance. 3. Increase processing speed. 4. Geo-reference data.	1. Field tests using enhanced system will scout orchards with fruiting walls up to 3 ft., and red, green and mixed color apples sizes 2” diameter and larger. 2. Analysis report validating performance.	1. Continuous scouting in excess of one hour in typical orchards with temperatures exceeding 95^{o}F. 2. Continuous operation in excess of eight hours (with software resets). 3. One mile of row data analyzed by system and incorporated into GIS database. 4. Data from 50’ foot scan of red apple trees analyzed and output within 60 minutes. 5. Average relative error in fruit count over reasonably-sized 3D regions less than 10%. 6. Sizing performance quantified.

Notable results:

- Six acres of green apples and 5.5 acres of red apples were scanned and estimates of fruit count and size were produced.
- Bias correction is required for accurate fruit count estimation.
- Raw fruit sizing estimates had median size within 3% of hand-measured median size.
- Ground traversal speed was increased to 1 mph.
- Successful operation in temperatures in excess of 100^{o}F.
- Analysis time decreased to 30 to 70 minutes for 100’ of data (30 minutes for the 2009 data used as the baseline).
- Data visualization through yield maps shows crop variability.

Introduction

The goals for Year 3 were to improve the apple detection and sizing accuracy for both red and green fruit, and to upgrade the robustness and operating scale of the prototype Scout. The year was to culminate with field tests to demonstrate crop load estimation on a significantly larger scale than in previous years.

Summary of Field Test Results

McDougall Farms, Ambrosia Block (Green Fruit, Figure 43)

- The crop load estimate for six acres was 560,358, or 75% of the 742,388 load extrapolated from hand count of 240'. Raw count was 299,997.
- Based on manual estimates of three bins per acre remaining in the block after harvest, the Scout's bias-corrected estimates are 376, 335, 300 and 262 bins if the harvested apples had average sizes 100, 88, 80 and 72, respectively.
- Counts from four trials each of two 60' hand-counted sections had consistencies (ratio of range to average) less than 7.1%.
- Ratios of Scout count to hand count from the four hand-counted sections had consistency (ratio of range to average) of 13.7%.
- Median apple diameter estimate for six acres was 2.60" which was within 1% of the median of 2.58" of the 240' of hand-measured apples.
- Raw median apple diameter for six acres was 2.62".

Figure 43. McDougall Farms, ambrosia block (green fruit).

Washington Fruit and Produce, Gala Block (Red Fruit, Figure 44)

- The crop load estimate for eleven acres (with every other row scanned and the results doubled) was 1,445,143. This is 94% of the 1,532,232 load extrapolated from hand count of 240'. Raw count was 284,477 for half of the block (5.5 acres).
- Based on manual estimates of one bin per acre remaining in the block after harvest, the Scout's bias-corrected estimates are 1,004, 899, 810 and 712 bins if the harvested apples had average sizes 100, 88, 80 and 72, respectively.

- Counts from four trials each of four 60' hand-counted sections had consistencies (ratio of range to average) less than 12.1%.
- Ratios of Scout count to hand count from the four hand-counted sections had consistency (ratio of range to average) of 19.2%.
- Median apple diameter estimate for eleven acres was 2.56”, which was within 3% of the median of 2.63” of the 240' of hand-measured apples.
- Raw median apple diameter for eleven acres was 2.49”.

Figure 44. Washington Fruit and Produce, gala block (red fruit).

Crop Distribution

- The Scout scanned 11.5 acres over two blocks, which is a significant amount of data collected about the orchards. While two blocks do not represent a large statistical data set, it does demonstrate the potential for the Scout to enable precision farming.
- The six acres of ambrosia apples consisted of 30 rows and approximately 5 miles of trees. When broken into approximately 5 m sections, the Scout count/section values had a median of 340 and a standard deviation of 158.7.
- The 5.5 acres of gala apples consisted of 26 rows and approximately 4.5 miles of trees. When broken into approximately 5 m sections, the count/section values had a median of 513 and a standard deviation of 128.4.
- The four ambrosia 60’ hand-count sections had median 1574 and standard deviation 162.6, while the four gala 60’ hand-count sections had median 1988 and standard deviation 104.4.
- Figure 45 and Figure 46 show the hand counts and bias-corrected Scout estimates for the four 60’ sections, where each bar represents the count for a section with width equal to twice the tree spacing (for the ambrosia sections) or equal to the tree spacing (for the gala sections). Two key observations can be made. First, the Scout counts closely follow the hand counts for the vast majority of the trees. Second, the hand counts on such a small basis vary greatly, namely by a factor of three between the lowest and highest density trees. Hence, sampling one or two trees and extrapolating is not necessarily an accurate approach to crop load estimation. It should be noted that the exact transition between the equal width sections may not be the same for the hand count and the Scout count, i.e., apples on the border may be in different sections for the different counting methods.

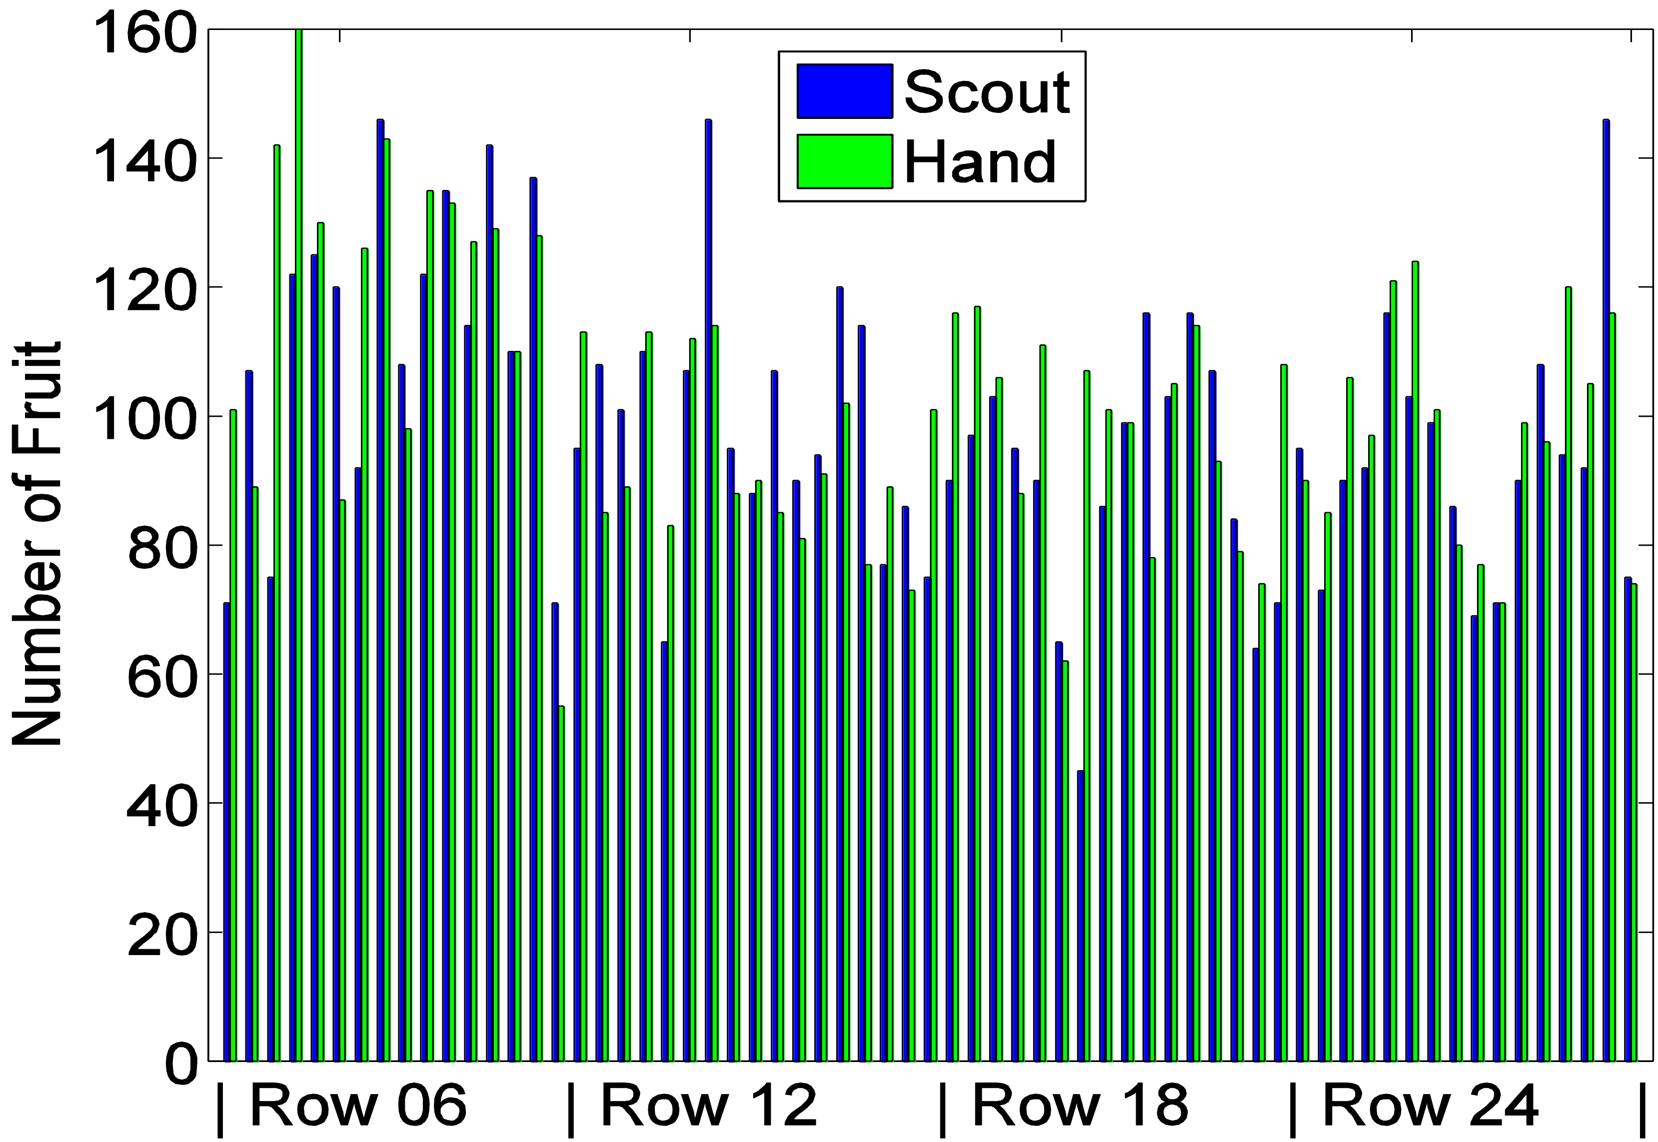

Figure 45. Hand count and bias-corrected Scout estimate on a tree by tree basis for ambrosia trees.

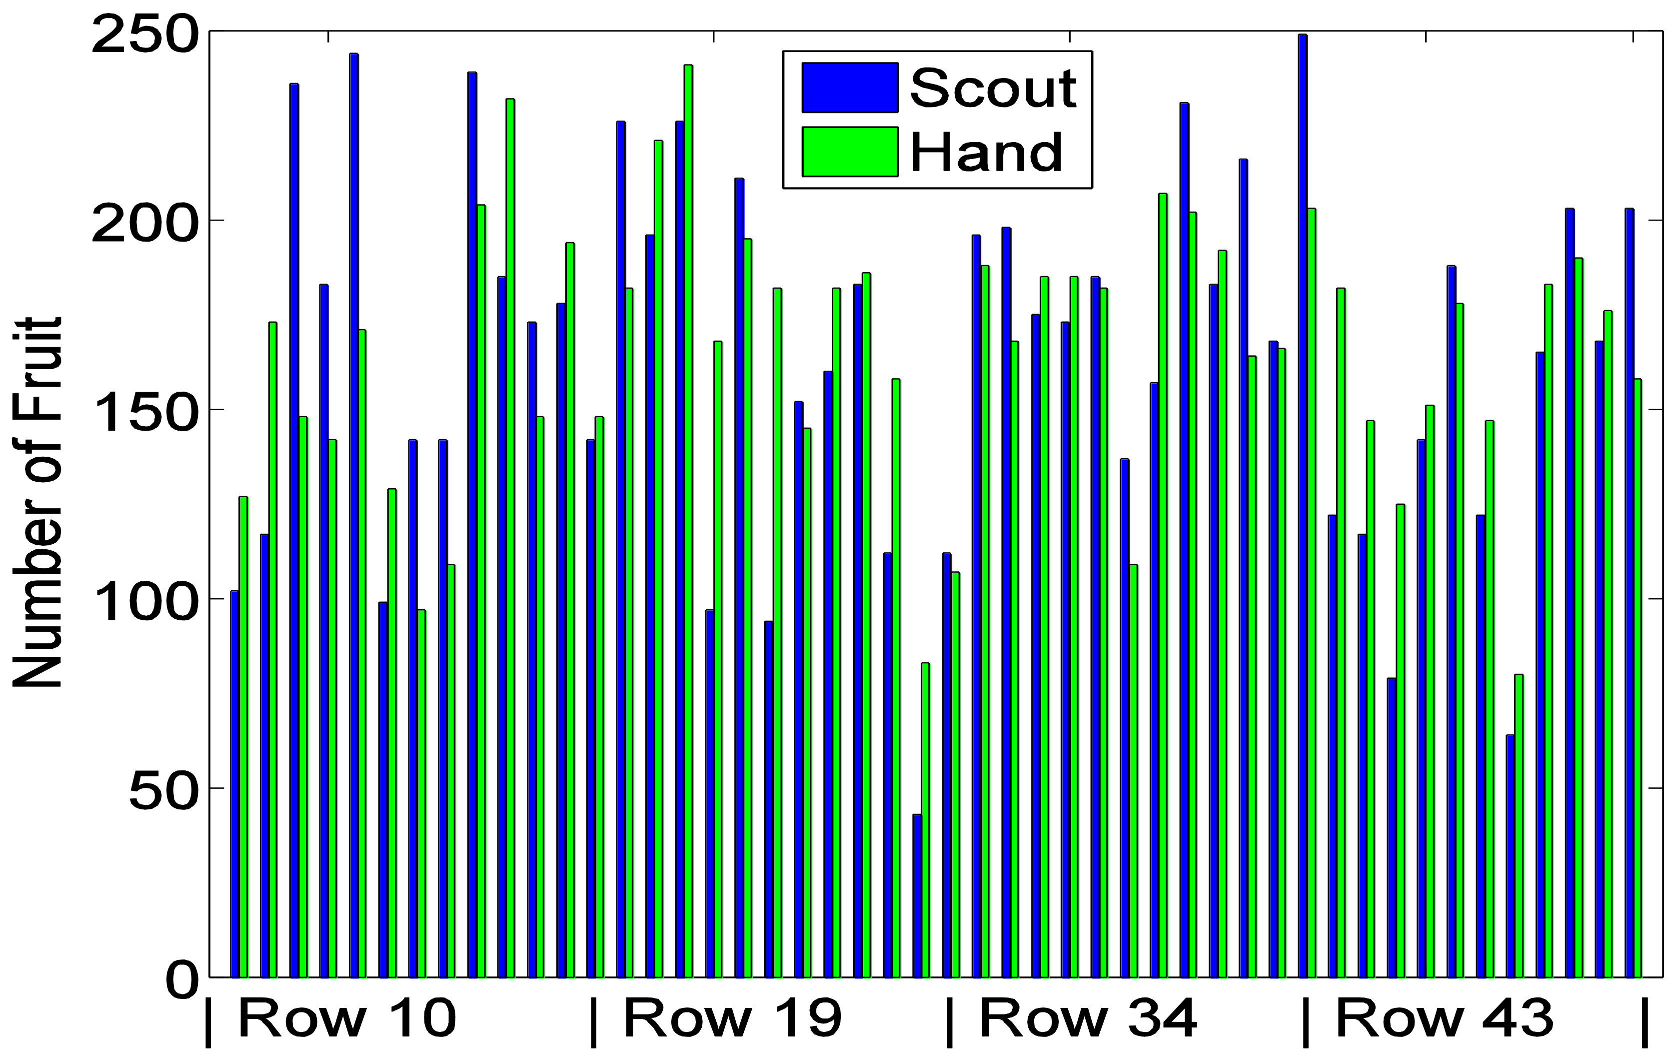

Figure 46. Hand count and bias-corrected Scout estimate on a tree by tree basis for gala trees.

2009 Sizing Benchmark

- As noted above, the 2010 sizing performance was extremely accurate, within 3% for the median apple size in the ground truth sections for both blocks scanned. This strong performance was predicated on benchmarking the 2009 results, analyzing actual sizing performance, improving sizing algorithms, and realizing benefits resulting from enhancements elsewhere in the Scout system.
- Vision Robotics benchmarked the 2009 software's sizing performance by scanning apples of known sizes in several different configurations.
- The raw average apple diameter was 2.2” while the true average diameter was 2.9”.
- The data have been used to create a statistical model that can be applied to correct for such system biases.

Speed and Robustness

- Using standard off-the-shelf computers, industrial flashes, and forced air cooling, the Scout worked continuously for over one hour in temperatures in excess of 100°F during the field tests.
- Ground speeds in excess of 1 mph were achieved using a camera frame rate of 20 images/s and the new flashes.
- The industrial flashes provide sufficient light to enable the Scout to operate in all Sun conditions, a flash rate high enough for production, and robustness that approaches that required for production.
- The system has been optimized such that the median time to analyze 100' of data was 30 to 70 minutes, depending on the block from which the data were collected.

Results and Discussion

Moving the Scout system towards production requires strong performance in counting and sizing the apples, high processing speed, high system robustness and an effective means of displaying the data. In 2010, all these components were significantly improved.

Estimation Performance

At this time, it is difficult to measure the accuracy of the Scout’s crop load estimates because accurately determining the actual crop load even after the harvest is largely impractical. Additionally, unknowns, such as an accurate value for how much of the crop was left in the orchard after it was scouted, are present. Furthermore, comparisons of Scout count estimates and after-harvest estimates are difficult because using bin counts as a unit can introduce inaccuracies since accurately converting a bin count to a count of apples depends on precise knowledge of the size of the fruit within the bins. Finally, at the time of this writing, post-harvest estimates are not yet available for comparison. Thus, the main ground truth data available for comparison are hand counts of 60' sections; however, as described below, these data were used for bias correction. Nevertheless, as discussed below, study of the consistency of the raw Scouting estimates suggests that it is reasonable to believe the Scout's bias-corrected crop load estimates approach the 10% maximum relative error goal; however, sufficiently accurate ground truth data is not currently available to make this determination with certainty.

Bias Correction

In a crop load estimation system such as the Scout, it is likely that biases will inherently be present. Causes of such biases could include software factors such as a systematic tendency to undersize fruit during analysis, as well as physical factors such as high tree thickness causing some fruit to not be visible in captured images (and thus not counted). While the variability of such biases across different orchards and apple varieties is not yet known, they are likely to remain consistent through a block. Consequently, a small set of hand collected data can be used to develop a statistical model to correct the bias, transforming raw data into more accurate, statistically-adjusted data. For example, by applying bias correction, an estimate is equally accurate regardless of whether the Scout consistently identifies 99% of the apples correctly with zero double counts or false positives, or consistently estimates 80% of the apple count regardless of the number of correct versus incorrect detections. Undoubtedly, the better the system is at correctly distinguishing apples, the more likely it is to have a consistent count.

In 2010, hand counts from 60' sections were used to perform bias correction. The ratios of the median hand-measured diameter to the median Scout estimated diameter were 0.98 and 1.06 for the green and red apples, respectively. These near-unity diameter scaling factors demonstrate that only a small bias was present in the median size estimates, and that potentially either an extremely small or no bias correction will be required for sizing.

To correct for bias in count estimates, a scaling factor was determined for each of the green and red apple scans by averaging the ratio of the hand count to the Scout count over the four 60' sections. The resulting count scaling factors were 1.87 and 2.54 for the green and red apples, respectively. The relatively large separation of these values from unity shows that, despite being rather self-consistent, a fairly large bias was present in count estimates. A brief review of images in which detected fruit were outlined with circles suggested that a large percentage of more distant fruit were not visible to the Scout's cameras because they either were obscured by closer foliage and fruit, or were insufficiently illuminated by the flashes. In general, as illustrated in the images below, the majority of the fruit detected by the Scout were on the close side of the trunk, which roughly corroborates a scaling factor of two. One potential approach to combating such a source of bias would be to detect tree trunks, filter detected fruit to include only those on the near side of the trees, and then double the estimated count. Another is for the Scout to determine the thickness of the canopy and determine the depth into which is sees a large and uniform number of apples. It could then filter out all apples beyond that depth and use a scaling factor between that depth and the total canopy thickness. Both these logical refinements would lead to a more stable bias factor across different orchards.

The McDougall block scanned in 2010 was representative of the configurations initially anticipated for the Scout. The Washington Fruit and Produce block reached beyond the Scout's initial objectives; however, the 2010 field tests demonstrated that bias correction can extend the capabilities of the Scout to operate in such orchard configurations.

Figure 47. Two flash images with detected apples indicated with circles.

Apple Count

To improve raw estimation performance, the first step in 2010 was to further analyze the Scout's 2009 performance. The estimates of the number of apples in the 100' sections of jazz rows in the Allan Brothers orchard scanned in 2009 were typically accurate (based on 4-6 runs) to within 25% using 2009 software, with some of the runs achieving 98% accuracy. The Scout count includes correctly identified apples, doubly counted apples, missed apples and false detections. The results of these analyses, coupled with experience gained improving algorithm performance in 2009, provided target areas for enhancement in 2010.

VRC has implemented an improved detection algorithm with better identification of individual apples (particularly those of mixed color) and stronger performance in identifying individual apples within clusters. The new algorithm uses the same software for detecting both red and green apples, but requires different input filters. The visual odometry software module is used to determine how much the cameras have moved between pictures and to determine the relative positions between cameras. The module has been enhanced to better correlate the portions of the images that overlap between cameras, which helps to eliminate double counts due to multiple cameras seeing the same apple (the largest number of errors in 2009) and enables improved location and size determination by incorporating more views of the same apple from different perspectives.

Significant improvements to the collection software and the prototype design also improved detection performance. Using data collected from the 2009 field tests, the team optimized the camera/mast configuration, adding one additional camera pair (nine as opposed to eight in 2009) and changing the relative locations and orientations to better image entire trees. Additionally, camera settings were adjusted and the auto-exposure algorithms were updated. Taken together, these modifications have improved the capability of the Scout to see the apples on the trees, which leads directly to improved estimation accuracy.

Figure 48 and Figure 49 show the raw and statistically-adjusted counts for the 30 rows of green apples, and the 26 rows of red apples. The median and standard deviation of the statistically-adjusted row counts were 17,820 and 2,799 for the green apples and 28,709 and 4,548 for the red apples. The plots illustrate that while a significant bias correction factor has been applied, the raw counts are overall fairly consistent from row to row. Note that row 31 of the red apple block was of substantially shorter length than the other red apple rows, giving rise to a correspondingly low count.

To view the data on a finer scale, each row was divided into approximately 16' sections along the row. Histograms giving the statistically-adjusted counts per section are shown on the top in Figure 50 for green fruit and on the bottom for red fruit. These plots demonstrate the Scout's ability to detect load variability within the block.

In order to develop a bias correction model and to study the consistency of the Scout's estimates, the hand-counted 60' sections were scanned repeatedly; the results are shown in Table 10. In all cases, the latter three trials were performed in succession, while the first trial was performed at a different time of day. The counts clearly demonstrate the consistency of the Scout estimates over these trials.

Table 10. Estimation consistency over multiple scans of the same sections.

Row	Scout Counts	Consistency (max-min)/average
Green 18	829, 780, 772 and 818	7.1%
Green 24	807, 812, 780, and 768	5.6%
Red 10	769, 817, 833 and 819	7.4%
Red 19	781, 737, 744 and 726	7.6%
Red 34	777, 851, 844 and 878	12.1%
Red 43	764, 733, 724 and 704	8.2%

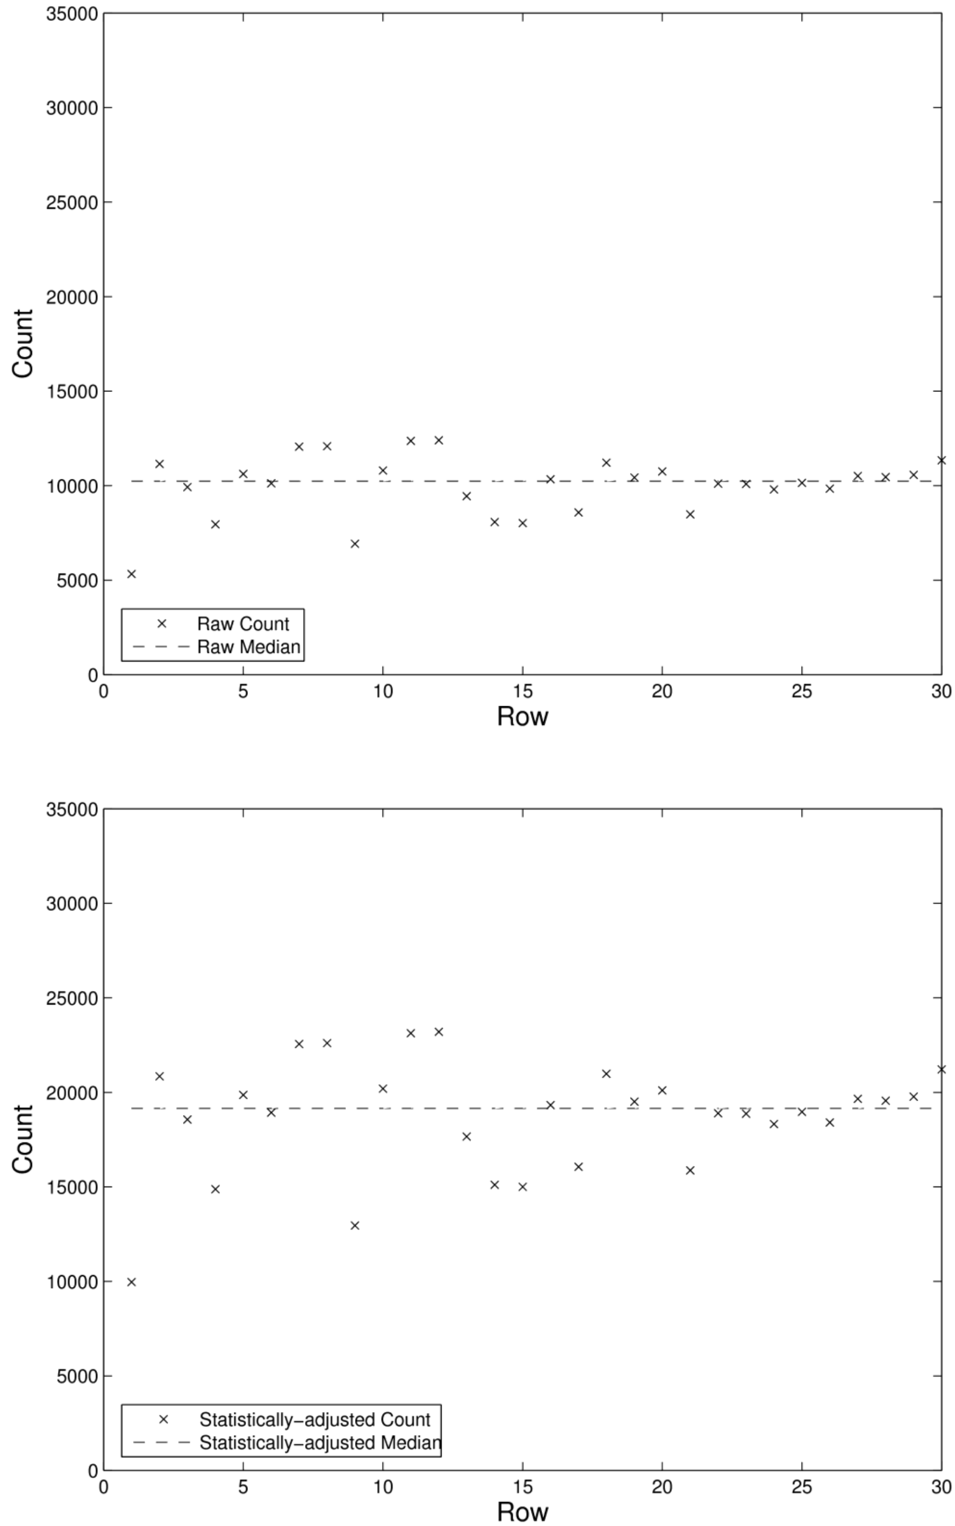

Figure 48. Raw counts (top) and statistically-adjusted counts (bottom) for 30 scanned rows of green fruit.

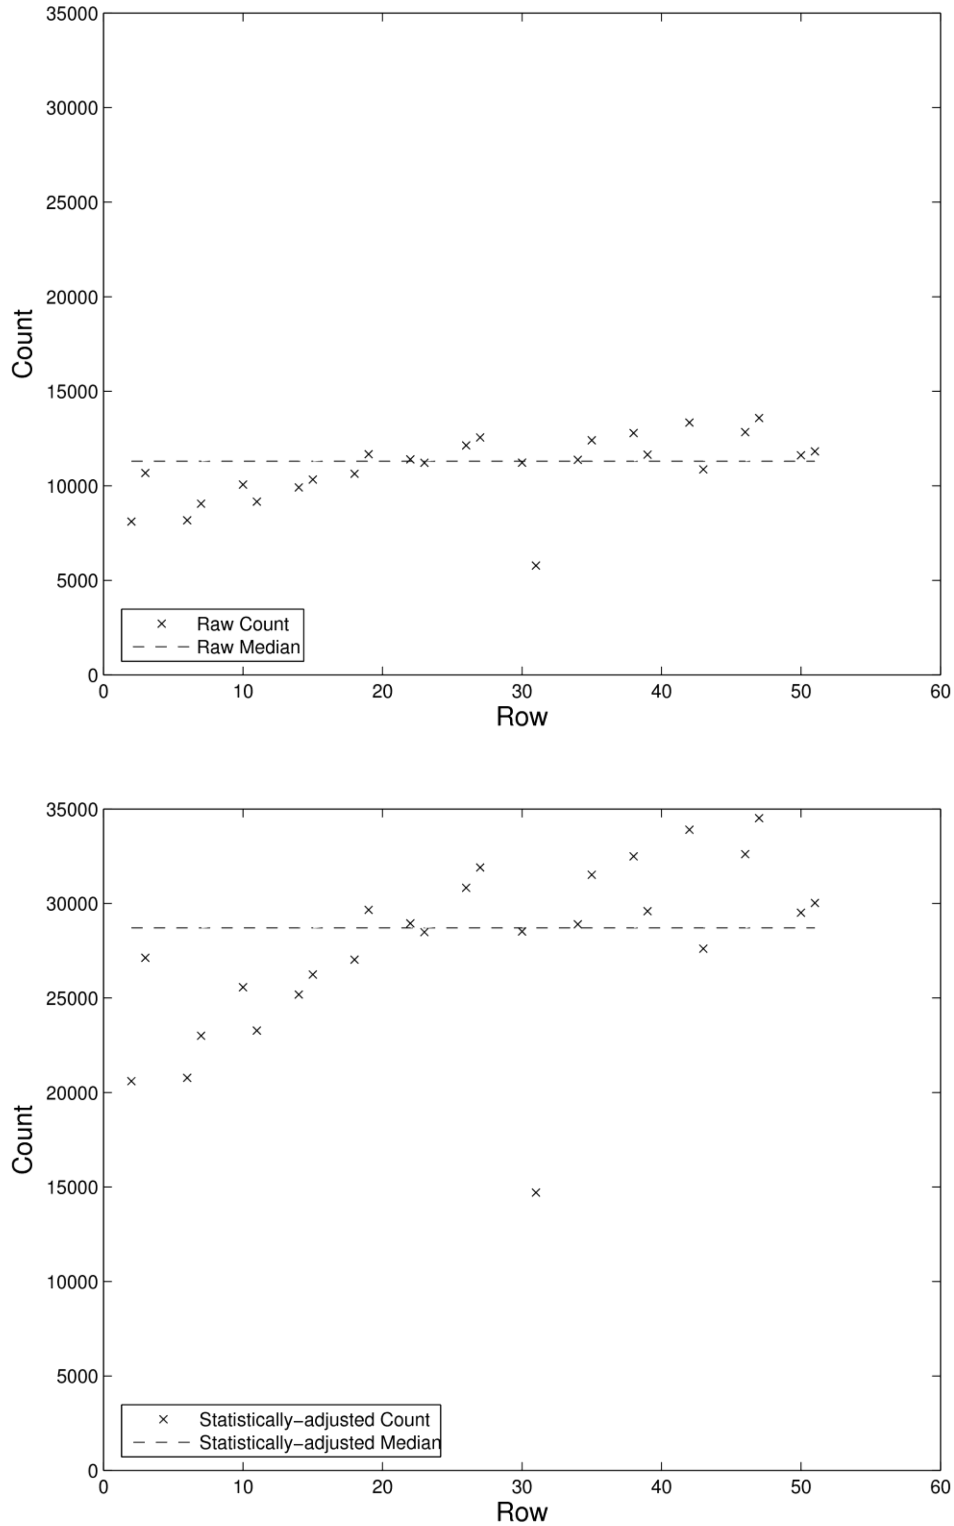

Figure 49. Raw counts (top) and statistically-adjusted counts (bottom) for 26 scanned rows of red fruit.

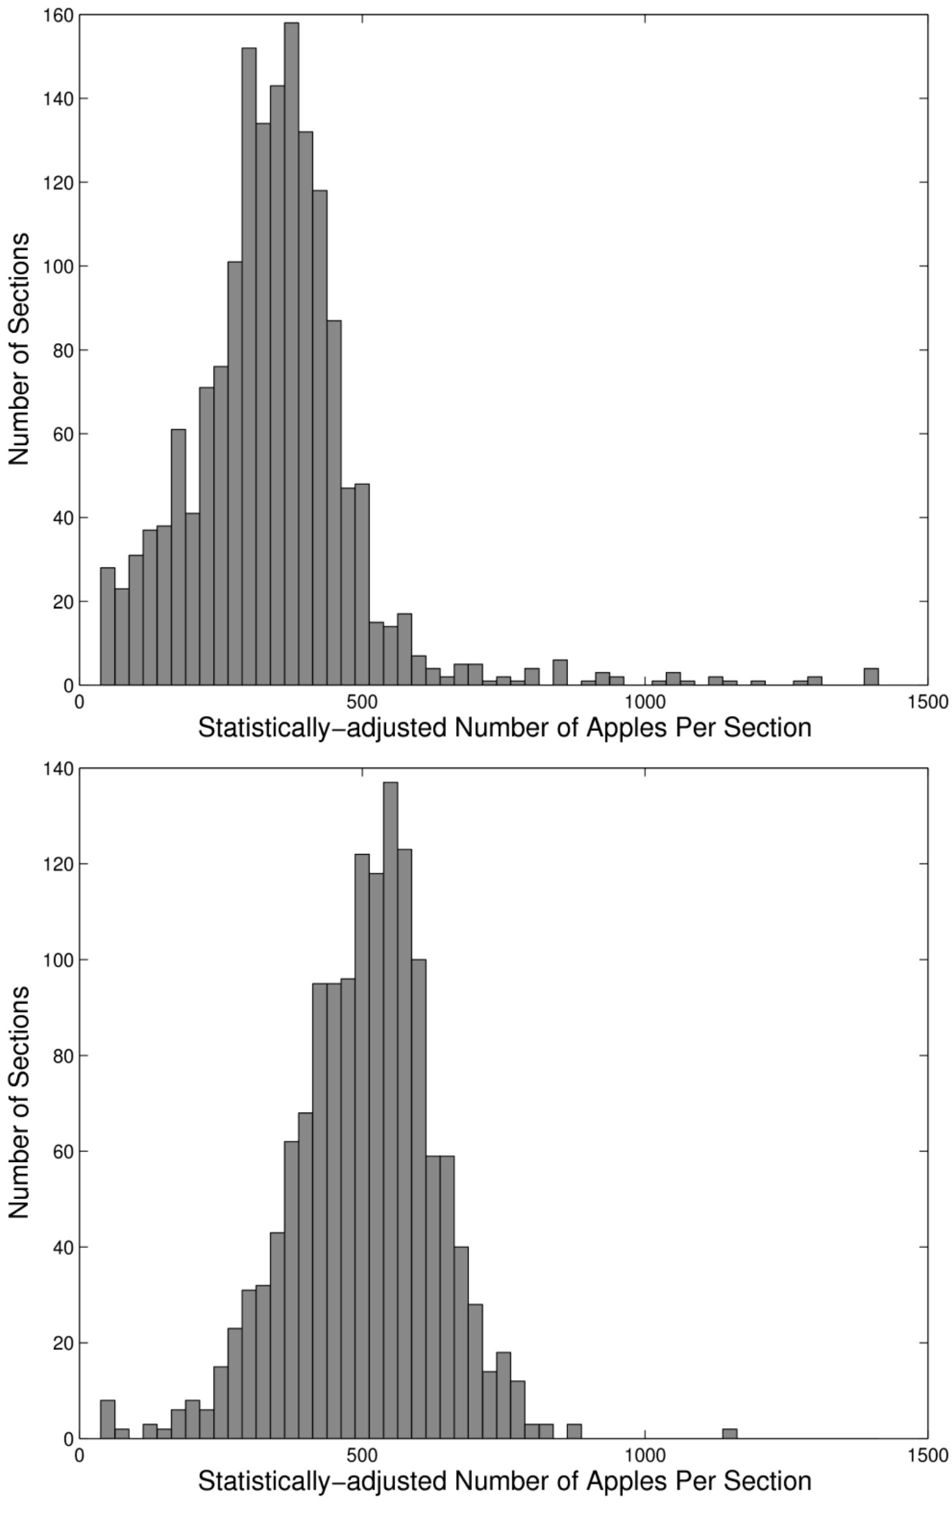

Figure 50. Statistically-adjusted histograms of fruit size for 30 scanned rows of green fruit (top) and 26 scanned rows of red fruit (bottom).

Apple Sizing

Analysis of the 2009 sizing performance revealed that partially occluded apples tended to lead to size underestimation with a larger variance than true distributions; the estimates of the range of apple sizes tended to be wider, flatter and shifted to smaller sizes than the actual crop. Sources such as this introduce system biases which can be reduced through the use of statistical modeling. A statistical model was created to adjust the size distribution to address these expected inaccuracies. More and improved data from this year's field tests will enable refinement of the model in 2011.

As reported last year, the 2009 raw average size estimates were approximately 20% too small. For the run shown in Figure 51, the raw average apple size was 2.2" diameter, which is 24.1% less than the hand-measured average of 2.9" diameter. The histogram of the data after the statistical model was applied shows a mean size of 3.0" diameter, or 3.4% larger than ground truth. Thus, the statistical modeling is effective in adjusting the mean of the size distribution. The variance of the distribution, however, remains larger than the true variance. This result is, in part, due to the need for a larger sample size when developing the statistical model. Such a larger set will be available when statistical models are developed based upon field data as opposed to laboratory data.

When compared to 2009, the 2010 Scout produces a better distinction between an individual fruit and the surrounding fruit, leaves and branches in the images. This delineation directly leads to the software more accurately detecting the perimeter of the apples in the images, thereby significantly improving the raw average sizing performance. The 2010 raw and statistically-adjusted (using the simple scaling model) size distributions for the six acres of green fruit are shown in Figure 52, and for the 5.5 acres of red apples in Figure 53. The aggregate hand-measured size distributions of fruit in the four 60' sections for green and red apples are shown in Figure 54. Visually comparing the histograms for the raw and hand-measured fruit diameters immediately illustrates that very little bias is present in the median size; however, the Scout estimates display a larger variance, as is expected. After bias correction with a near-unity scaling factor, the median green apple diameter was 2.60”, which was 0.78% larger than the hand-measured median diameter of 2.58”. Similarly, the bias-corrected median red apple diameter was 2.56”, which was 2.66% smaller than the hand-measured median diameter of 2.63”.

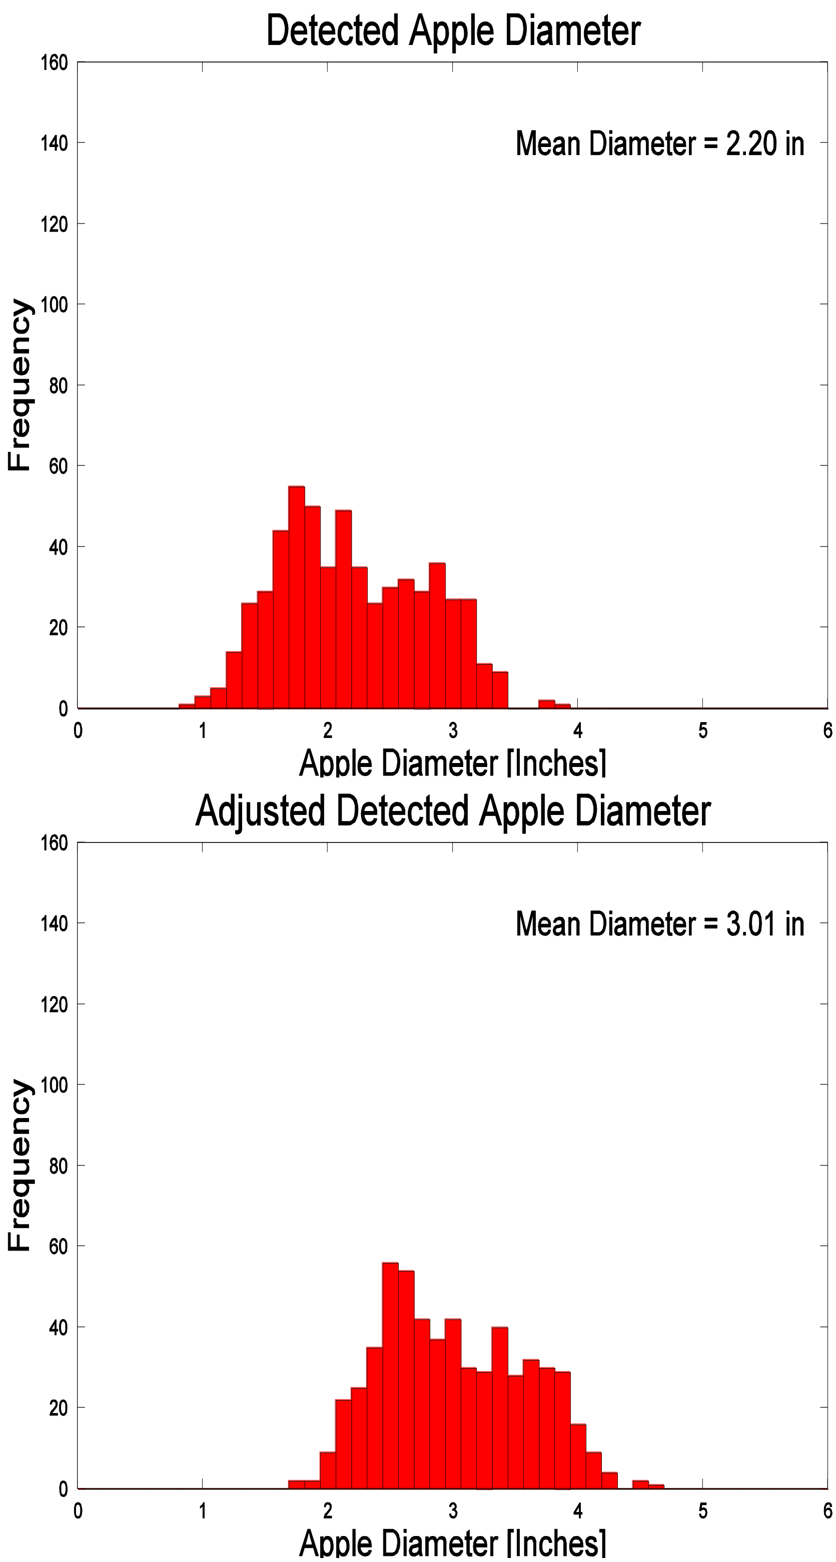

Figure 51. Raw (top) and statistically-adjusted (bottom) size distribution for a 2009 scan.

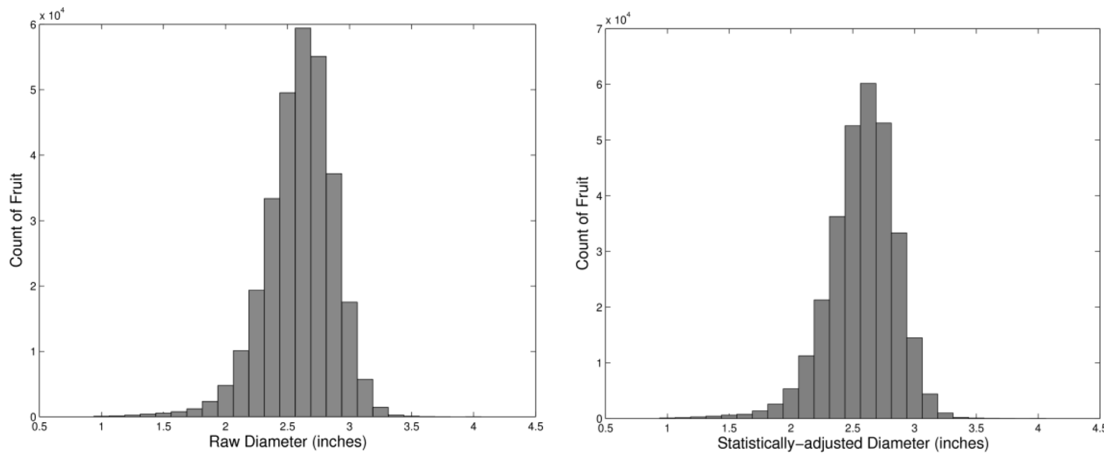

Figure 52. Raw size estimates (left) and statistically-adjusted size estimates (right) for 30 scanned rows of green fruit.

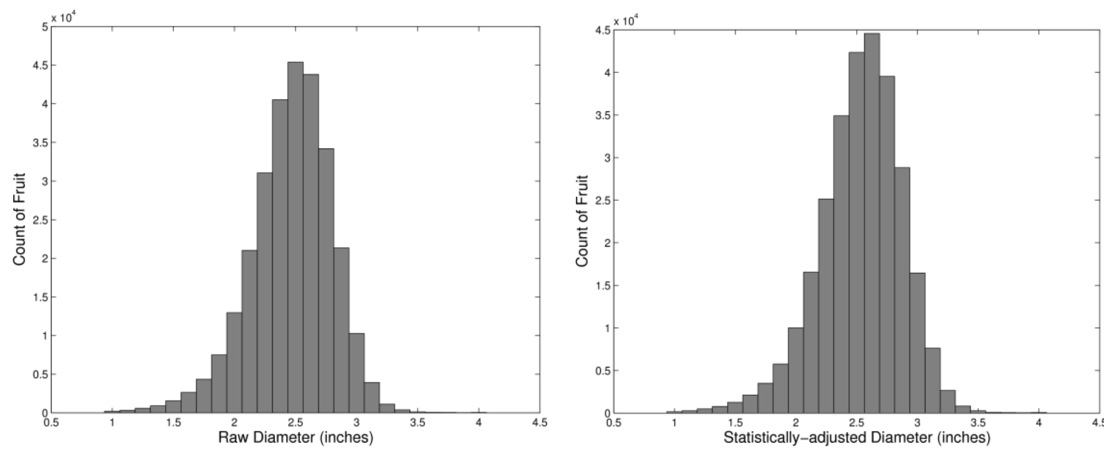

Figure 53. Raw size estimates (left) and statistically-adjusted size estimates (right) for 26 scanned rows of red fruit.

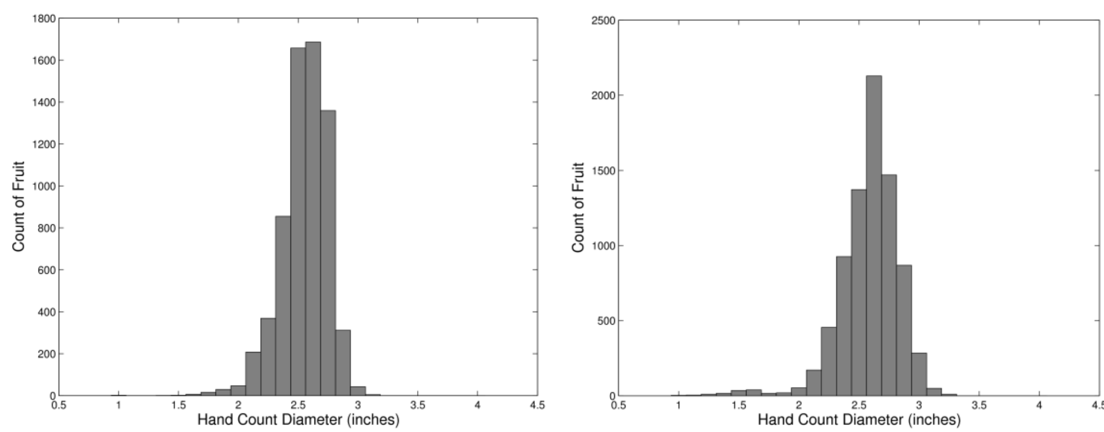

Figure 54. Hand-measured apple sizes for 240' of green (left) and red (right) apple.

System Robustness

Improving the system robustness was a key goal for this year and all related objectives were met. The 2010 Scout prototype and camera module (Figure 55) represent significant upgrades towards a production design, but the resources necessary to fully weatherize them were not expended. For example, each camera pair is now in a closed module that is straightforward to fully seal, but the time and expense were not taken to use IP65 connectors and gaskets. Similarly, the computers are more robust than those used in 2009, but they are still standard desktop models. Fully weatherized and robust computers are available, but only represent a marginal robustness improvement that, as expected, was not required this year. Active cooling of the electronics cabinet through a fan system was included in the 2010 prototype. VRC conducted several local field tests to debug the prototype and the final unit operated virtually flawlessly during the week of field tests in Washington where the temperatures were in excess of 95°F every day and above 100°F a couple of days.

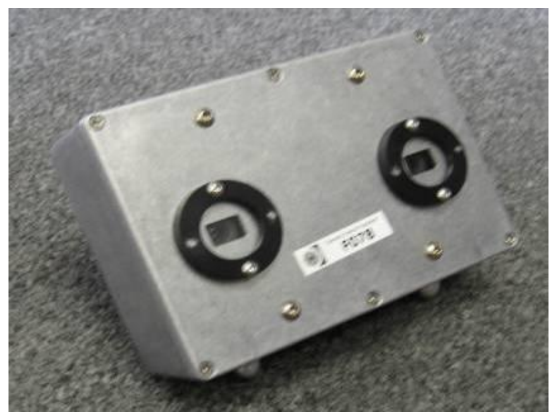

Figure 55. Scout prototype (left) and camera module (right).

As noted, the improved robustness was a part of the requirement to ensure that the Scout can operate at a production scale. Additional improvements introduced to achieve this goal include:

- increased Scout scanning speed to 1 mph;
- increased camera frame rate;
- decreased image density (pictures per inch);
- incorporated GPS system to geo-reference data;
- decreased number and increased robustness of electrical connections;
- debugged software to eliminate crashes;
- incorporated industrial flashes.

Analysis Speed

The current Scout is approximately 20 times faster than in 2009 when analyzing data from 2009, meeting the goal of analyzing 50' of data within 60 minutes. In fact, 50’ of 2009 data was process in under 30 minutes, or approximately twice the initial success criteria. The speed gains were achieved primarily through parallelization and decreased analysis time because of the new detection algorithms. Analysis times for 2010 data are somewhat longer, with median times of 70 and 59 minutes for 100' of green and red apples, respectively. This increase can be attributed primarily to the blocks, which have significantly more fruit than those scanned in 2009, and to the fact that the scans were significantly longer (the longer runs require the software to track more fruit during each run).

Integration with the APM and CASC GIS System

In 2010, the APM towed the Scout throughout the field tests at the Washington Fruit and Produce orchard. VRC and the CMU team spent almost a full day integrating the two systems, primarily updating the APM software to correctly turn between the rows when towing the Scout. The two robots completed the red apple scans over the course of the next day and a half. VRC has provided the crop load estimate data to CMU for integration into the GIS database.

Data Visualization

VRC has created a framework for viewing data output by the Scout to provide detail and a debugging environment at VRC. The crop load data can be overlaid onto a Google Earth map of the block to show the crop load and sizes for various resolutions. The yield and median size maps for the red apple data broken down into approximately 5 m sections along each row are shown in Figure 56 for red fruit and Figure 57 for green fruit. In each case, red indicates lower counts (or smaller sizes), yellow indicates medium counts (or sizes) and green indicates higher counts (or larger sizes). Note that some variability is present due to inaccuracies in raw received GPS data. Such inaccuracies likely account for instances where data which should truly appear in rows which are quite red being shifted to appear in neighboring rows (making them very green). Count and size data can be shown at any resolution from the entire block, to a row, to any meter along the row.

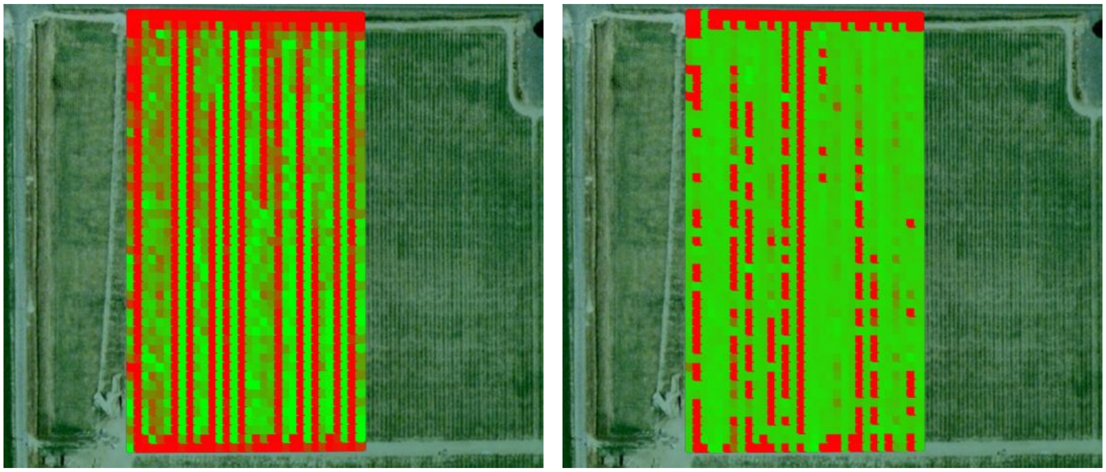

Figure 56. Yield map (left) and median size map (right) for the 26 scanned rows of red fruit.

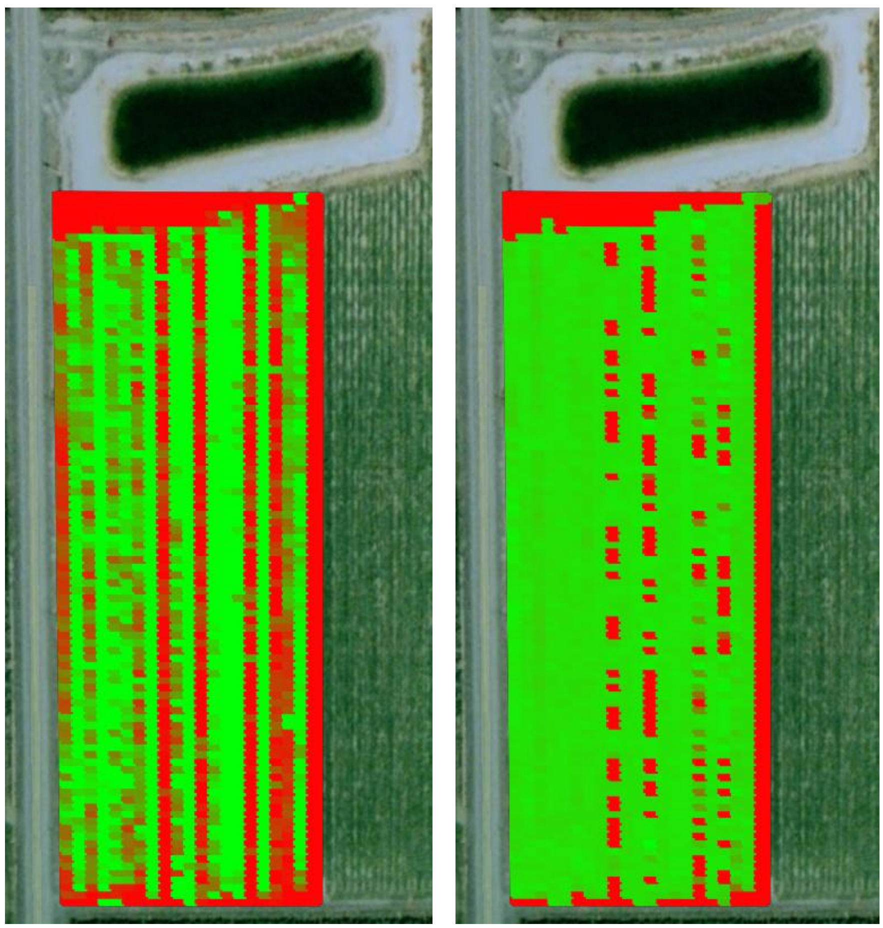

Figure 57. Yield map (left) and median size map (right) for the 30 scanned rows of green fruit.

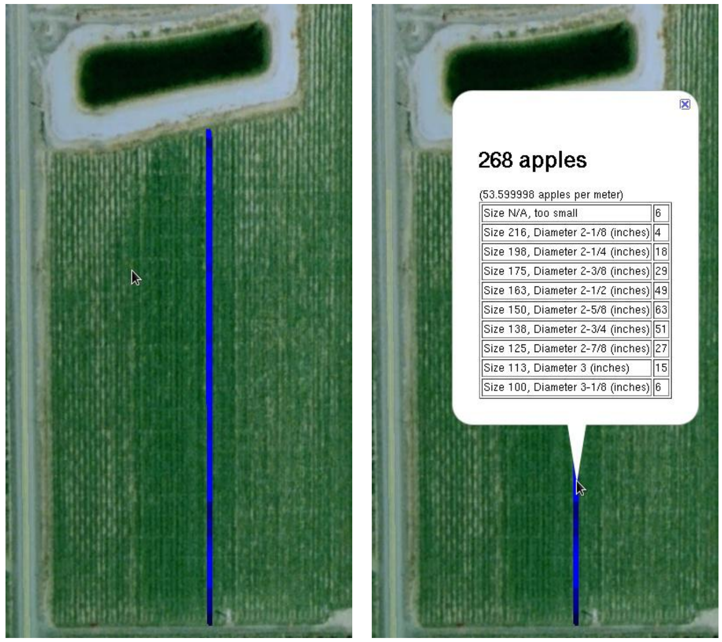

Figure 58. Yield map for a single row (left) and size information for a 16' section (right) for a green apple scan.

Field Tests

The VRC team wishes to acknowledge and thank everyone that helped make our field tests a success. This includes the McDougall's and Washington Fruit and Produce who let us into their orchards as well as helped keep us dry during those hot days. Similarly the CMU and WTFRC teams went way above and beyond reasonable effort. Collectively, we worked from before dawn to late into the night, and even through thunderstorms. Finally, we appreciate and thank the Commission for collecting the ground truth data, both the estimates for the full blocks and the hand counts of small sections within the blocks. During the tests, the Scout collected approximately nine terabytes of data for analysis. Despite the temperature, rain and sprinkling, it operated all three days and collected data without a failure except for a couple of hard disk-related crashes with one of the twelve disks used.

The Future

VRC is pleased with the 2010 progress; the detection software performance (particularly with respect to sizing) and speed was improved, and the Scout demonstrated its ability to scan large blocks with high consistency within a block. The goals for the future include a plan for continued refinement of the apple detection and sizing performance, and further increasing the processing speed with the ultimate goal of achieving real time. One specific goal for 2011 is to analyze the data collected this year to determine a statistical scanning plan to create accurate crop load estimates while scanning only portions of the orchards. An additional key goal is to collect data from a larger and more diverse set of blocks (in terms of varieties and tree configurations) to analyze the variability in the statistical models used for bias correction.

3.4 Caliper Measurement

Thematic area leader

Name	Institution	Email
James Owen	Oregon State University	jim.owen@oregonstate.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Develop tree counter. 2. Develop 2nd version of caliper. 3. Quantify caliper performance in field conditions in OR (shade trees) and WA & PA (fruit trees).	1. Report on performance of counter and caliper devices. 2. At least one each of the counter and caliper prototypes operational in WA, OR, and CA. 3. Database of images with accompanying ground truth.	1. Caliper accuracy to within 1 grade. 2. Count accuracy to within 2%.

Notable results:

For the caliper:

- Determined caliper within ±2 mm of harvested bareroot tree passing in front of the stationary device in a warehouse.
- Determined caliper within ±3 mm of nursery trees at an unspecified height (≤10”) while moving at approximately 2 mph in the field. Caliper was only obtained when no visual obstructions or interference was present.
- Caliper data was used to effectively determine bareroot tree grade.

For the counter:

- Counted in-ground nursery trees of ≥ ¼” in caliper with 95% accuracy while moving at approximately 2.5 mph.
- Counted in-ground nursery trees of ≥ ½” in caliper with 97% accuracy while moving at approximately 3 mph.
- Counter is undergoing commercialization trials by Crop Tech LLC, to be used in WA fruit tree nurseries.

Introduction

In Year 2 we focused on redesigning and field-testing the on-the-fly caliper and counter devices (Figure 59). Extensive testing was conducted in various production systems over three states: Pennsylvania, Oregon, and Washington (Table 11).

Figure 59. Second generation caliper, constant height wheel-mount and infrared sensors for counting.

Table 11. Year 2 sites utilized to test the counter and caliper devices.

Production System/Application	Location
Fruit tree nursery (bareroot trees in warehouse	Adams County Nursery, Aspers, PA
Shade tree nursery (in field)	J. Franck Schmidt and Son, Canby, OR
Flowering tree nursery (in field)	J. Franck Schmidt and Son, Canby, OR
Ornamental tree nursery (in field)	J. Franck Schmidt and Son, Canby, OR
Pot-in-pot ornamental tree nursery (in field)	Oregon Turf and Tree, Hubbard, OR
Ornamental tree nursery (in field)	Oregon Turf and Tree, Hubbard, OR
Established, trellised orchard	Skyline East Orchards, Royal City, WA
Newly planted, trellised orchard	Skyline East Orchards, Royal City, WA
Fruit tree nursery (in field)	Willow Drive Nursery, Ephrata, WA
Fruit tree nursery (in field)	Willow Drive Nursery, Ephrata, WA
Fruit tree nursery (in field)	Dave Wilson Nursery, Hickman, CA

Caliper

Ornamental and fruit tree nursery producers identified tree caliper as an important indicator of growth and yield, respectively. For nurseries, measuring caliper of in-ground trees would increase accuracy of current inventory and improve predicting of future inventory. Fruit tree growers identified caliper as most meaningful for monitoring young orchards, one to three years old, to determine rate of growth and establishment. Knowledge of young tree establishment allows orchard owners to better predict fruit yield, and therefore overall return on investment.

Data analysis of the 2009 caliper device field evaluations in Washington, Oregon, Maryland, and Pittsburgh was completed in the winter of 2009-2010. The measurements from the prototype caliper device were within the suggested 3 mm specifications when imaging was performed at speeds under 2 mph. Further improvements were necessary to improve speed processing, widen the field of view and measure at a larger range of depths.

The second version of the caliper device provides an increased field-of-view, depth-of-view, and incorporates a visible laser to assist the user with maintaining height and distance from the tree being measured. The issues of uneven ground conditions, the need to measure a fixed distance above a variable budding point and crop and non-crop related vegetation interference remain as challenges to be addressed. Additionally, this new caliper uses Class III laser diodes that pose eye safety issues. Sustained exposure to the laser source such as in the case a person places their eye at the aperture of the caliper device can cause injury. Proper workplace precautions need to be observed to ensure OSHA compliance.

The second generation caliper device, when mounted to a wagon or equipment, is able to measure 97% of unstaked bareroot fruit or ornamental trees caliper within 3 mm accuracy moving at < 2.5 mph if tree stems are straight, planted in line, and perpendicular to the ground. At 3 mm accuracy, grade can accurately be determined with the exception of those trees following below or above a grading threshold such as 1” caliper. This is not a regular occurrence. In Year 2 it was observed that the speed of travel can increase with increasing tree caliper. This concept was successfully tested by measuring caliper travelling at 3 mph when determining caliper of in-ground wholesale ornamental cherries at Oregon Turf and Tree. When sorting or grading indoors the second generation caliper, as a stationary device, can measure caliper within 2 mm (Figure 60).

Obstacles we still face are measuring caliper at 6” above the soil line or bud graft, interference from the tree stake, and curved stems (Figure 61). All of these obstacles occur only under field conditions. The issue of measuring caliper at a given height and ground interference is worsened by the non-uniform graft height and irregular grade or mounds/hills that occur within the tree row. Caliper cannot be determined on trees in which the stake is in front or behind the tree due to interference. The crooked tree stem is primarily a concern in fruit tree nurseries or orchards in which curved stems are acceptable. These stems move towards or away from the device as much as 10” if parallel with the row. A third generation caliper device will be designed and constructed in Year 3 using new camera and lasers to address and attempt to overcome these barriers.

Figure 60. From top to bottom, left to right: caliper being used to determine tree diameter in a warehouse sorting operation, bareroot ornamental nursery, established orchard, bareroot fruit tree nursery and ornamental wholesale nursery in both field and pot-in-pot production systems.

Figure 61. Current challenges when obtaining caliper in a nursery or orchard include tree stakes, variations in grade, tree sleeves, trellis posts, strings, curved stems, pots, irrigation stakes, and varying heights of graft union (bottom right, blue lines).

Counter

In Year 1 the focus of our research was to define existing methodology utilized by industry, possible obstacles, field conditions and industry integration. This occurred simultaneously with the development and evaluation of the on-the-fly caliper device. During Year 1 fruit tree bareroot nurseries identified counting as or more important than caliper measurement. Ornamental nurseries also identified counting as important.

In Year 2, we determined that counting alone could be accomplished using a simpler approach with much less expensive hardware then that needed for measuring caliper. We designed and constructed a low-cost infrared device that works by detecting the presence or absence of trees. Initial field tests of the counter were conducted at Adams County Nursery in Aspers, PA on November 11th, 2009. Eighteen tests were conducted on 70 tree whips > ½" caliper spaced 24" at a speed < 2 mph. For the 70 trees measured the prototype counter device was able to remain within 4% accuracy (± 3 trees). Our device and a more expensive, commercially-available infrared sensor were sent to WA and OR for additional field evaluations.

In Year 2 both counters were evaluated at fruit and ornamental nurseries. Testing began at J. Frank Schmidt and Son Co. in Canby, OR on June 17th, 2010. Forty-five Sergeant crabapples were counted five times with an average accuracy of 97% when travelling at approximately 3 mph. Also at J. Frank Schmidt and Son, forty-five coffee tree seedlings were counted four times with an average accuracy of 97% while travelling at approximately 2.5 mph. In Ephrata, WA at Willow Drive Nursery on July 27th, 2010, twenty small caliper apple tree seedlings were counted ten times with an average accuracy of 95% while travelling at approximately 2 mph. Many passes did not work initially because of ground interference. The last field evaluation was conducted at Dave Wilson Nursery in Hickman, CA on July 29th, 2010. Counters were mounted on a tractor and used to count fifty cherry trees ten to twenty times. Trees were counted with an average accuracy of 97% while travelling at 2.5 to 3.0 mph (Figure 62).

Figure 62. Infrared devices used in Year 2 to count trees in bareroot ornamental nursery, fruit tree seedling nursery, bareroot fruit tree nursery and ornamental tree seedling nursery.

Commercialization

Crop Tech LLC, a new agriculture technology startup led by three WSU graduates, met Karen Lewis in Year 2 to discuss opportunities of creating a service-based business that counted and callipered trees for central and eastern WA tree fruit nurseries. Karen Lewis and entrepreneurs of Tech LLC developed a business plan and began discussions with CASC engineers. Shortly after, Crop Tech LLC traveled to Oregon and met with Jim Owen, observed on-farm trials of caliper and counter devices at Oregon Turf and Tree Farms, met with CASC industry participants, and attended a presentation by Wenfan Shi that explained the technical details of construction and application of the caliper and counter devices. Furthermore, Crop Tech LLC entrepreneurs met with CASC scientists and industry participants to discuss opportunities and pitfalls of the technology, business plan, and potential clientele.

Crop Tech LLC identified the counter to be the first device employed by the company. In the fall of 2010 an opportunity arose to count trees as part of an insurance claim for a WA tree fruit nursery than incurred damage from the winter of 2009-2010 (Figure 63). The company acquired two second generation counters and provided service and software at no cost. Counting is underway.

Figure 63. Crop Tech LLC utilizing CASC counter and software to count trees at a Washington tree fruit nursery.

3.5 Information Management

Thematic area leader

Name	Institution	Email
Ben Grocholsky	Carnegie Mellon University	grocholsky@ri.cmu.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Integrate plant science sensor suite into GIS. 2. Deploy limited APM field-testing.	1. Demo in test orchard.	1. Produce map of 10 km block with three complementary sensors.

Notable results:

- Developed a web-based geographic information system (GIS) tool for collecting and managing crop information from the field. Re-implemented Year 1 GIS capability to remove scripting and supervision of data entry and processing by engineers or scientists. Processing is now performed automatically on the system server. Data query, editing and display use free and open source GIS tools (PostGIS, OpenLayers and Geo-Django).
- Used new GIS tools to display Vision Robotics crop load estimates from WA field trials. The system provides display of geo-referenced crop data that includes total fruit count for each region and counts within specified size classes.
- Leveraged APM mobility to collect data on continuous runs over 10 km in length; the scale of whole orchard blocks. Three sensors (laser, camera, NDVI) were used to collect canopy data that complements crop load measurement. It is anticipated that this data will be valuable in designing efficient crop load assessment schemes that utilize sparse sampling rather than dense canopy measurement.

Information Management Task

The Year 1 geographic information systems (GIS) effort demonstrated the value of orchard data collection and display in terms of supporting awareness and high-level decision making by growers and managers. Reviews identified two drawbacks with the previous approach:

1. Unacceptable level of scientist involvement in supervising data entry and processing;
2. Poor scalability to large data sets without operations on small space and time selections.

Development effort for Year 2 focused on re-implementing Year 1 capability, using existing open database and network frameworks to deliver a GIS tool that is scalable and easy to use.

Approach

A streamlined systematic GIS workflow for orchard information collection and management has been implemented using free and open source software. The software tools utilized provide a flexible web-based user interface (Django), geospatial database functionality (PostGIS), data visualization overlays (OpenLayers) and automated numerical procedures such as data interpolation (Python). This combination facilitates the shift from manual data processing to automated processing on a GIS server triggered via connections to network clients. The GIS server data storage modeling and the network interfaces that provide user interaction are described in the following sections.

CASC Data Model

A data model is implemented to provide storage primitives that describe types of measurements, information and infrastructure typical to orchard and nursery operations. Time and location are common to all data primitives. Fields to store specific numeric values, text descriptions and geometry have been added to represent the following entities:

- **Orchard Structure:** Site boundaries, individual Tree, Row and Block geometry.
- **Sensing/Scouting:** Sensor type and characteristics, measurements at point, area or voxel locations, registered imagery and manual incident descriptions.

The data modeling methodology is extensible. Foreseeable additional features include representing orchard roads and paths, geometries such as irrigation lines, and maps of landmarks used for APM autonomy and localization.

Network and User Interfaces

The system provides network interfaces for data access and graphical web-based interfaces for user interaction. The following actions are supported by these interfaces:

- **Import:** Data is entered into the system via either on-line connection to the server or off-line from standard text or data base file formats logged on field devices.

- **Edit/Administration:** A web user interface is used to provide dynamic listing, inspection and editing of the data on the server.
- **Data Selection:** User input supports queries based on measurement type, time range and spatial bounds allowing focused recall and inspection of specific information of interest.
- **Display:** Selected data is visualized via map spatial overlays to provide inspection, comparison, analysis and decision aid to orchard managers.
- **Export:** Export of selected data to standard file formats supports further data use and analysis via import to specialized scientific and agricultural software tools.

Field Trials

Data collection trials were conducted in Pennsylvania and Washington orchards using the APM as a sensor platform and in Washington using the Vision Robotics Scout Newton integrated with the APM to provide mobility and data registration. Three additional sensors were mounted to the APM (laser, camera, NDVI) as shown in Figure 64 to collect canopy data complementing the crop load measurements made by Vision Robotics. Notable features of the data collection trails were:

- Low operator effort: Collection was conducted in the background during APM testing.
- Scale and diversity: Over fourteen 10 km datasets were collected over a variety of research and production orchards.

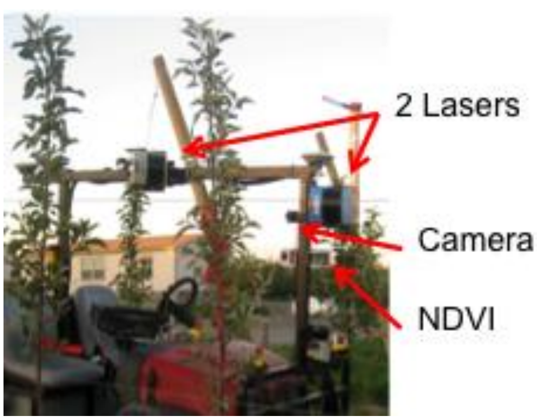

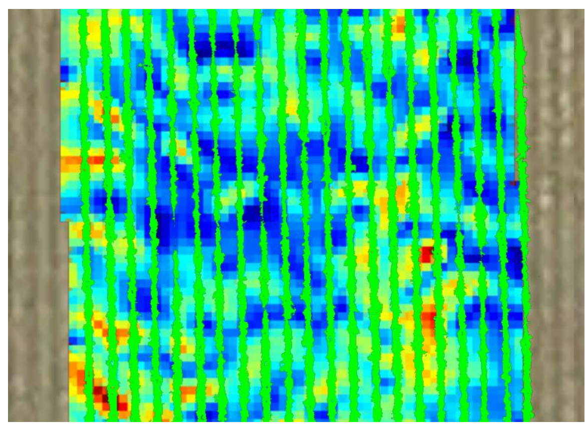

Figure 64. Canopy sensing instrumentation on the APM (left), and example GIS data layers (right) showing a visual comparison between yield data collected by traditional methods and canopy shape sensed by APM-mounted laser scanners.

Lessons Learned

- Automating data collection and management is challenging. Significant payoffs include elimination of data entry effort and large-scale/duration operation without conscious monitoring by equipment operators.
- Standardized open source web and database frameworks provide basis for building powerful free GIS tools.
- APM mobility and localization depend on maps of orchard features. Usability would benefit from having future GIS provide orchard maps to APM system
- Tracking and auditing APM activity and performance is a tedious task for operators but naturally suited to GIS. This task is highly valuable in demonstrating APM usability. Future GIS development will monitor and report APM activity.

4. Automation

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.

This section presents the goals and accomplishments in the following three thematic areas:

- Reconfigurable mobility;
- Accurate positioning;
- Augmented harvesting.

4.1 Reconfigurable Mobility

Thematic area leaders

Name	Institution	Email
Sanjiv Singh	Carnegie Mellon University	ssingh@cmu.edu
Brad Hamner	Carnegie Mellon University	bhamner@cmu.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Integrate payload for assessment and treatment tasks. 2. Integrate low-cost localization. 3. Perform field tests in WA and OR. 4. Extend APM automation to one more platform. 5. Create, simulate and test control policies for thinning.	1. APM integrated with GIS and crop load assessment. 2. APM integrated with precision spraying. 3. APM automation package installed and tested on hydraulic orchard platform. 4. LIDAR integrated with string thinner control in test orchard.	1. 100 km low-cost APM scout safe operation with a MDBF of 10 km. 2. 10 km of autonomous row following with hydraulic orchard platform. 3. Working closed-loop LIDAR-based control of the string thinner with quantitative results. 2. Quantitative comparison of the ultrasonics and LIDAR for control.

Notable results:

- 159 km of autonomous driving with the APM, each segment longer than 10 km.
- 10 km of autonomous driving with hydraulic orchard platform.
- APM controlled by orchard workers using a user-friendly graphical interface.

Improving Reliability

Year 1 in the area of reconfigurable mobility focused on proof-of-concept. We automated the first vehicle in the Autonomous Prime Mover (APM) family, a Toro eWorkman, and added sensors and a software suite to have the vehicle drive between rows of trees and turn around at the end of rows. We also showed how such a 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 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. We met our 100 km distance goal for that year, but were frequently restricted to blocks or rows which were neatly maintained or on relatively level terrain. Our focus for Year 2 was on improving 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 this year 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.

Modifications to the APM

The original design of the APM had two laser range finders located on the corners of the vehicle about one foot off the ground (Figure 65, left). The lasers scan 180^{o} in a horizontal plane and provide distance to each object in front of them. We oriented the lasers to give a 270^{o} 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 65, right). To the autonomous system such obstacles are indistinguishable from trees and make row detection difficult.

In Year 2 we modified the APM to use only one laser range finder for row following. This is a new model of sensor which provides a 270^{o} field-of-view. We mounted the sensor centered left-right and at a height just above the hood of the APM (Figure 66). This configuration has two advantages. The first is that since the laser is above the hood we can take advantage of its increased field-of-view. The second is that the sensor at this height is less susceptible to the spurious obstacles mentioned above.

Figure 65. (Left) The Autonomous Prime Mover (APM) in its original configuration, with laser range finders mounted low on the corners of the vehicle. (Right) With the lasers at this height it was easy for sloped terrain to come into view, which presented a challenge to the row detection algorithm.

Figure 66. The APM in its new configuration. The laser range finder mounted just above the hood is the only one used for row following.

Improvements in Row Following

One major concern in row following in Year 1 was performing row detection in the presence of spurious obstacles like weeds or sloped terrain. Our first row detection system used a Hough transform, 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 67, left). Although the new laser height mitigated this problem by presenting fewer such obstacles, we required a row detection system which was more robust.

In Year 2 we developed a row detector which uses a particle filter, 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 range 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 67, right).

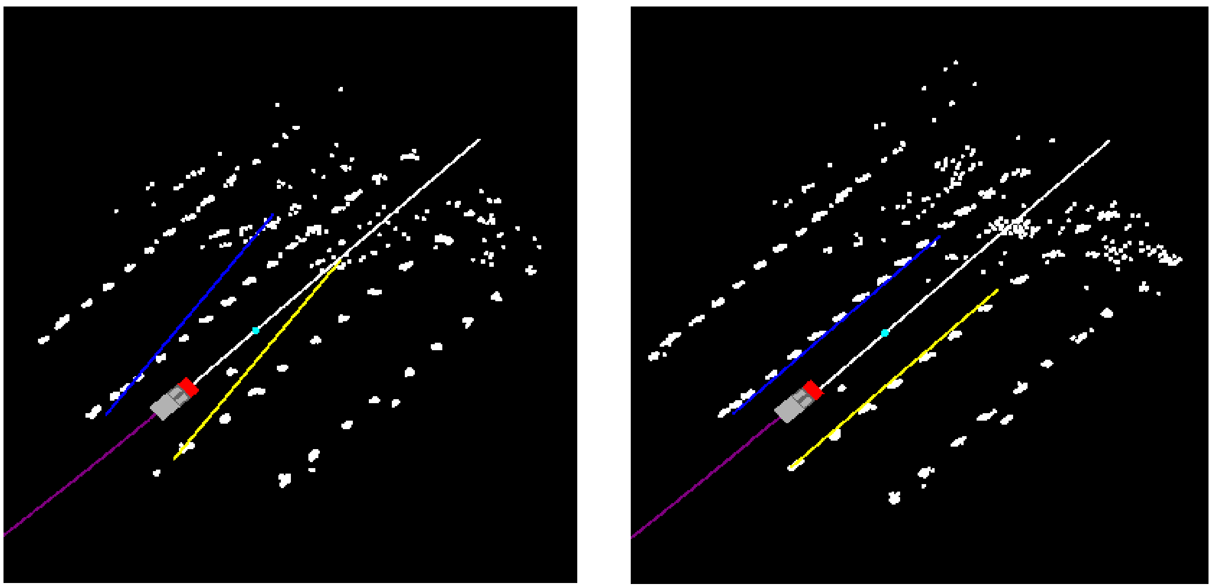

Figure 67. Comparison of row detection methods among tall weeds or uneven terrain. (Left) The Hough transform method used in Year 1 is confused by the unexpected data, resulting in bad row detections. (Right) The particle filter method used in Year 2 correctly ignores the additional data and finds the tree rows.

Improvements in Turning and Row Entry

When the robot has reached the end of one row it attempts to turn into the next row. Due to inconsistencies in the 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 we had in June. 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 68). Detections from this point were unreliable, with a success rate less than 50% in dense canopy.

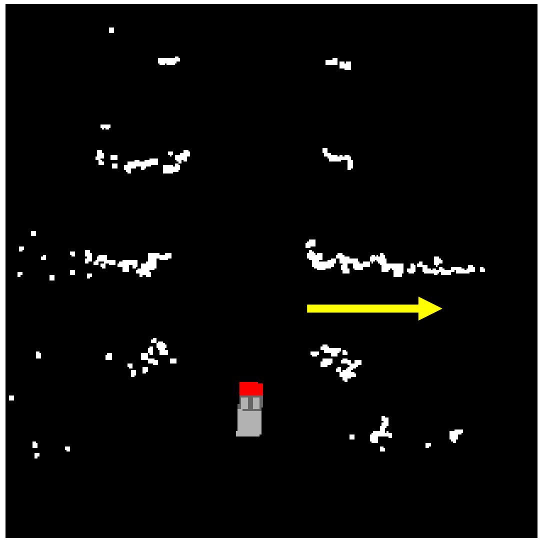

Figure 68. The APM attempts to enter the row indicated by the yellow arrow. The canopy of the trees at the near side of the row, however, blocks the rest of the near edge and most of the far edge. The row detector has little data to work with in this situation.

In Year 2 we designed new turn methods to make a smoother row entry. We implemented a path planner that takes the nominal starting point of the vehicle in the next driving lane and generates a smooth path that respects the vehicle's steering constraints. The planner optimizes the path to produce gentle steering, avoiding hard steering angles that stress the motor, while minimizing the total path distance (Figure 69, 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 69, right).

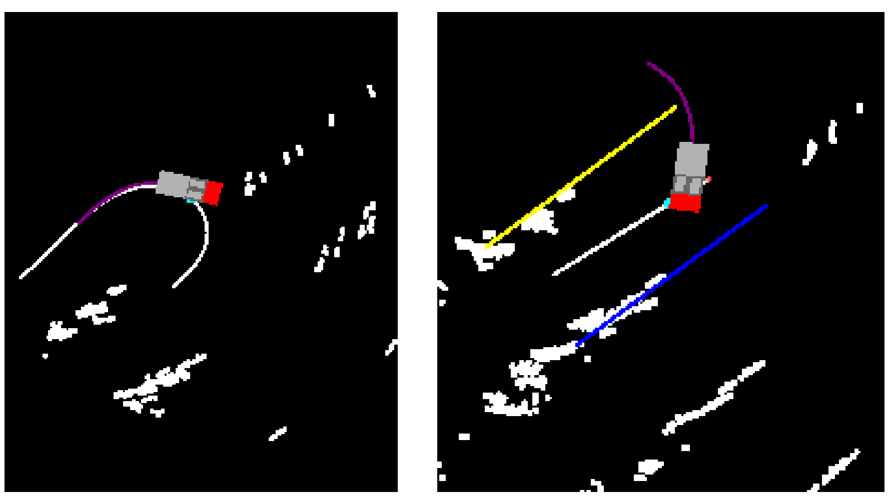

Figure 69. (Left) The new planner generates a smooth trajectory for the APM to follow, staying away from hard steering angles that stress the motor, while aligning the vehicle to be able to see into the row before entry. (Right) Once the vehicle has turned towards the row the system performs a row detection to refine the desired path of the vehicle.

Frequently, in order for the vehicle to be aligned with the row before entering, the planner creates a path with 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 70). During the third leg of the turn the system runs the row detector to refine the final goal point.

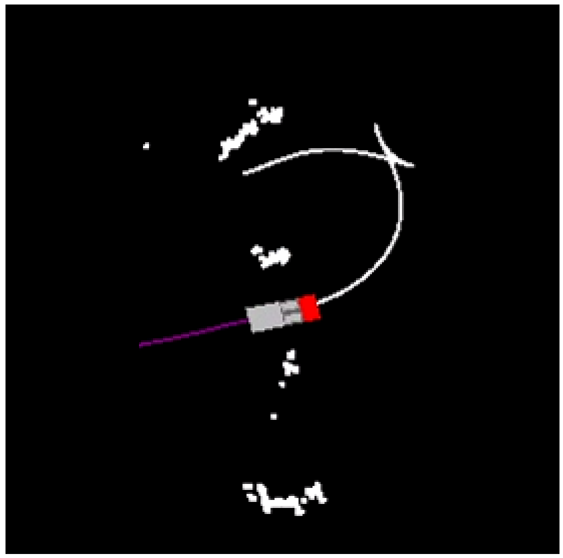

Figure 70. This row is too narrow for the APM to smoothly turn into. The system plans a "K"-turn, where the vehicle turns past the row, backs up, and then enters. In the third leg of the turn the system performs a row detection to refine the final desired path.

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.

Tests

We conducted tests of the improved autonomous navigation system in Pennsylvania and Washington during the summer months. Test missions 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 the block. 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 Year 2 (Table 12). 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 24th and 31st. 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 crash. They are typically hard to find and require such endurance runs to uncover. After each of the tests on the 24th and 31st we ran memory check tools and corrected the software to fix the problem. Following the corrections for the test on the 31st we ran the system in a simulator to confirm that there were no more memory errors. The simulated system ran for over 30 simulated km in 8 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 that lasted 5 hours, which was manually stopped by the test lead for time scheduling reasons.

Although we reached the 100 km mark at Robot City, we continued to go out to Soergel Orchards and Penn State's Fruit Research and Extension Center to demonstrate the APM'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. In this way the APM successfully drove six different blocks at the FREC, despite differences in row width, canopy training, and growing systems (Figure 71).

Figure 71. A sampling of the blocks driven by the APM this year. 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.

Table 12. Autonomous driving runs obtained in Year 2. 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	Person walked in front of robot causing bad row detection
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	Brief delay in communication with low-level vehicle controller causing missed turn
09/03	Robot City, Pittsburgh, PA	10.2	Manually stopped (no error)
09/08	Robot City, Pittsburgh, PA	15.6	Dust cloud obscured laser causing bad row detection
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)

Automation of the N. Blosi Hydraulic Orchard Platform

One of the main objectives of the Reconfigurable Mobility theme is to show how autonomy technology can be mapped between different types of vehicles, thus resulting in a family of autonomous vehicles all using the same sensors and software. In the Year 1 report we described initial work in adding our autonomy package to an N. Blosi platform owned by the Penn State Fruit Research and Extension Center. Our goal for Year 2 was to complete this work and demonstrate the platform driving autonomously in an orchard. We added valves to control speed and steering, connecting to a microcontroller similar to that used on the Toro eWorkman. The navigation software's interface to the microcontroller is identical, so no software had to be changed between the eWorkman and the N. Blosi platform. Two Sick LMS laser range finders were added to the front of the platform scanning a horizontal plane low to the ground, similar to the initial configuration of sensors on the eWorkman (Figure 72).

Figure 72. The automated N. Blosi platform. Two laser range finders provide a view of the world similar to the original configuration of the Toro eWorkman APM. All software, including the user interface, runs on the laptop mounted on the top rail.

Testing of the N. Blosi platform took place at the Penn State FREC in May 2010. The platform drove over 10 km autonomously, thus satisfying the Year 2 goals. There were, however, some setbacks which would prevent wider use. One is that the vehicle's speed controller performed poorly when traveling uphill, with the platform driving much slower than intended. In manual mode we were able to drive the platform as fast as desired, suggesting that the problem is a lack of throttle. This should be repairable at the microcontroller level, and we are in talks with the microcontroller's manufacturer to have that repair done.

Another problem is related to the laser height. The lasers for the N. Blosi platform were designed to mimic their height on the original eWorkman APM. At this height it was easy for weeds and unlevel terrain to cause poor row detections, leading to the possibility of poor control. In one situation (Figure 73) the terrain dipped such that the lasers on the platform could not see beyond a few meters. We had raised the lasers on the APM for this very reason; the APM traveled this same row later in the year with no problems. We will raise the lasers on the platform before further use with our navigation system.

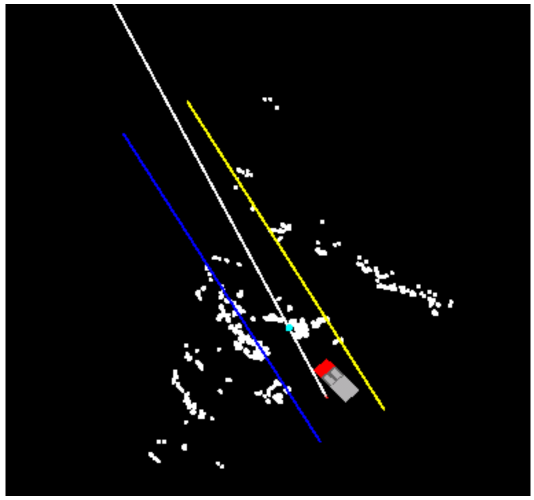

Figure 73. In unlevel terrain the N. Blosi cannot see more than a few meters ahead, and the travel lane is filled with weeds. It is impossible for the row detection to get a good result with this data.

User Interface Development

In Year 2 we began to take the steps necessary to deploy autonomous vehicles in orchards. Our first action was the development of a new user interface for the APM. We had developed an interface for our own purposes (Figure 74). This interface, however, is unsuitable for use by someone unfamiliar with the inner workings of the autonomous navigation system: it refers to the vehicle's position in GPS coordinates as opposed to its location in the orchard; commands are sent via custom-built text files and sequences of button presses. Details of the autonomous system are displayed, which is useful for engineers but too much information for end users. An interface for end users would need to be simpler, focusing on the user controlling the vehicle with a minimum of effort.

We enlisted a team of three design students from the Human-Computer Interaction department at Carnegie Mellon University to work from August to December, 2009. We asked them to design a graphical user interface intended for an end user of an autonomous vehicle. Since our goal is reconfigurable mobility we asked them to consider both an off-board mode where a vehicle is sent out to mow or spray an entire block as well as an on-board mode where the user would be on a vehicle (like a platform) and commanding the vehicle to go to the end of the row.

The designers employed a formal top-down design process, starting with interviews with growers where they asked what their intended use of the vehicle would be. Growers from four farms in Wexford, PA and Adams County, PA were interviewed. The designers learned about the management of an orchard, what vehicles are typically used for, who would be the likely users of an autonomous vehicle, and their level of comfort with computers and software. From there they developed prototypes of the interface on paper (Figure 75). They presented these to growers and other people involved in agriculture to get a preliminary sense of what control concepts made sense and what needed to be redesigned.

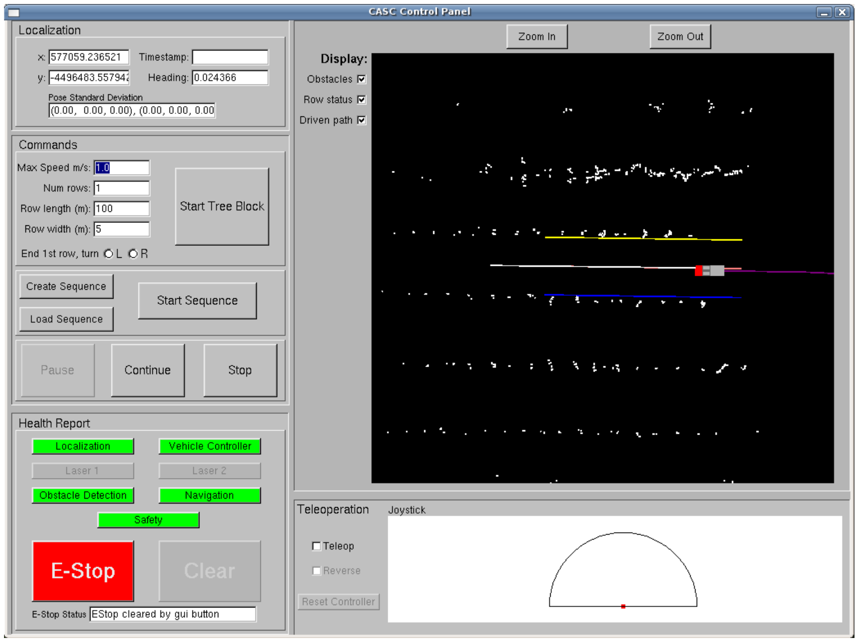

Figure 74. Our user interface tells the software engineers everything they need to know when testing the functionality of the APM, but the controls are not intuitive enough for a lay user.

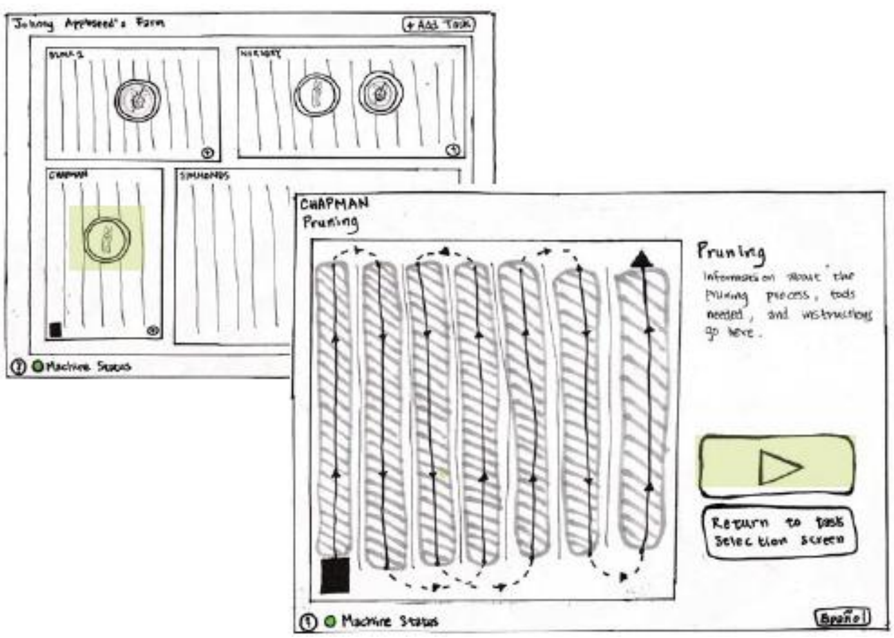

Figure 75. An initial concept of the user interface presented on paper to potential users. Top left, the farm view shows a map of the blocks and tasks that are available to perform on those blocks. Bottom right, the block view shows the vehicle's progress through the block and lets the user start and pause operation.

The next phase of the process was a working prototype (Figure 76) using lessons learned from the paper prototype. We worked with the designers to interface their program with our autonomous navigation software. The biggest realization from the design perspective was that orchard managers want to specify how tasks are accomplished, while workers just want to start the vehicle without having to set parameters. For example, the manager would decide at what speed the vehicle should move while towing a sprayer. It would be the worker's job to hook up the spraying machine and start the vehicle, but he should not change the parameters set by the manager. The interface accomplishes this with a set of buttons in the top-left corner of the screen that are associated with predefined tasks. The worker would take the vehicle to the desired block, click on the desired task, and start the vehicle. The interface keeps track of completed rows, which the manager could use at the end of the day to track the work being done. Of course, our intention of using the interface for on-board use required a method to tell the vehicle to speed up, nudge left or right, etc. The designers placed an array of sliders along the bottom of the interface screen for this purpose.

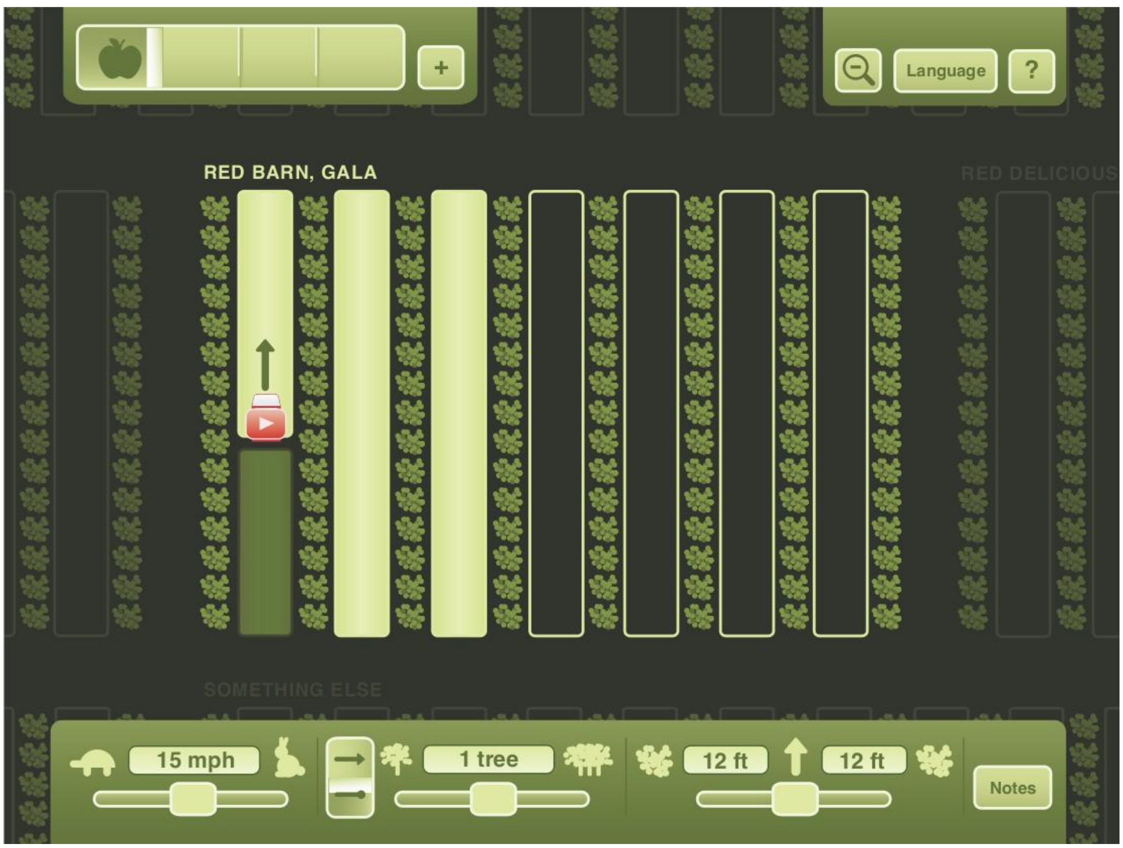

Figure 76. Screen shot of the new user interface (working prototype). Pre-defined tasks available for the block are shown in the top left (in this example, harvest). Sliders along the bottom allow the user to control the desired speed and lateral offset of the vehicle, as well as whether it should drive to the end of the row or stop after a certain number of trees.

Tests of the prototype were conducted in a simulator running on recorded data from commercial orchards. The designers presented this version of the interface to employees at

Soergel Orchards in Wexford, PA. These employees were given scenarios of orchard work and asked to start and stop the vehicle, adjust its speed, lateral offset in the row, etc., all without previous training. The workers who were comfortable with computers executed every command correctly the first time. The workers with little computer experience typically found the correct commands, but were unsure of their actions. They explained that the vehicle would be an expensive piece of equipment, and they were afraid of doing something incorrectly. The main conclusion from this round of interviews was that the user interface required a training step, and very little modification to the interface itself would be required.

For the final round of testing we brought the APM with the new user interface to Hollabaugh Orchards in Biglerville, PA. We had four potential users, the orchard manager and three full-time workers, operate the APM following a five-minute explanation of the vehicle and the interface. Each user was able to operate the vehicle easily. Each was allowed to experiment with the controls, starting and stopping the vehicle, increasing and decreasing its speed, and switching between going to the end of the row and a mode where the vehicle would only go the prescribed number of trees and then stop.

This interface has shown its promise for an orchard management situation. Using it, a worker should easily be able to complete any off-board task required of the vehicle. We have found, however, the interface to be ill-suited to on-board use. The sliders to control speed and lateral offset in the row are sufficient to control the vehicle, but are small on the screen and difficult to use while the vehicle is moving. The orchard management perspective of the interface may be seen as overkill for a worker on a platform trying to prune trees. He may not care about where in the orchard he is, and would rather have larger buttons to operate. We will move forward with this interface when considering the problem of organizing a fleet of autonomous vehicles, but we have determined that it would be best to have a separate interface for use of the vehicle within a row by an on-board worker.

In the fall of 2010 we are working with another student design team to explore the issue of on-board use of an autonomous vehicle. They are attempting to answer the questions of how best to control the vehicle, not limited to a screen interface. Possibilities include joysticks, levers, or even voice commands. We will present the conclusions at the of Period 1 in Year 3.

Mechanized Thinning

In Year 2 we collaborated with the SCRI project "Innovative Technologies for Thinning of Fruit," led by Prof. Paul Heinemann at Pennsylvania State University. We focused on a control and perception system to automate the operation of the Darwin 300 string thinner, made by Fruit Tec. The Darwin has been demonstrated to be a cost-effective method of thinning[1] but currently requires the operator to weave in and out of the trees using the steering on the tractor to maintain engagement with the canopy. This causes operator fatigue and wear on the tractor. Working with Prof. Heinemann's team we created a system that automatically actuates

[1] Schupp, J.R.; Baugher, TA; Miller, S.S; Harsh, R.M; Lesser, K.M. "Mechanical Thinning of Peach and Apple Trees Reduces Labor Input and Increases Fruit Size." HortTechnology. 2008 Oct-Dec, v. 18, no. 4, p. 660-670.

the thinner's tilt angle and offset based on the tree profile while the operator focuses on driving straight down the row (Figure 77). We intend to show that this method is as effective as or more effective than the current manual control of the thinner.

Our goals were to create, simulate and test algorithms for planning a trajectory for thinning, and compare ultrasonic and laser sensors for creating this trajectory. Penn State is developing the mechanism, ultrasonic sensors, and embedded control, while CMU is providing the laser sensing and planning. Our test platform is a Darwin string thinner which has been modified with hydraulics that control the piston offset and angle.

Offset and Tilt Control

In automating the Darwin thinner, the basic questions we need to answer are: (1) how do we obtain a representative profile of the trees, and (2) given a tree profile, what are the optimal offset and tilt of the thinner? At this moment we are investigating the use of laser rangefinders and ultrasound sensors to respond question #1, and experimenting with simple algorithms to advance solutions to question #2. For most of the types of trees we are studying (especially peaches and apples), horticulturists determined that it is desirable to have as much of the thinner spindle next to the outer edge of the tree as possible, while avoiding collisions of the spindle with the tree. This maximizes thinning of inner branches and approximates human behavior with the thinner.

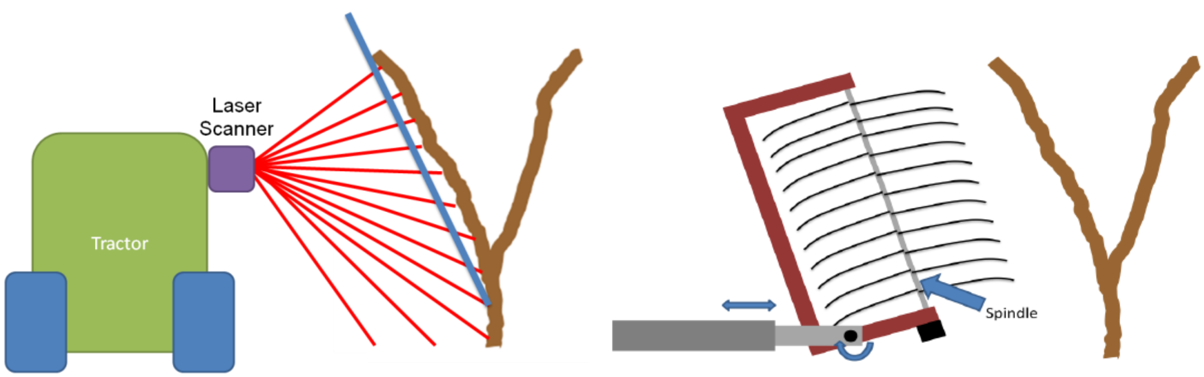

Figure 77. The proposed automation system for the Darwin string thinner. A laser or ultrasound sensor captures the profile of the tree (left) and sends this information to a computer (not shown here). Software running on the computer analyzes the tree profile and actuates the Darwin's tilt angle and offset (blue arrows) accordingly (right). The piston has 24" of travel in offset and 35^{o} in angular range.

To begin our investigation we built a simulation environment using data collected by mounting laser and ultrasound sensors on a tractor and driving it along rows of tree fruit. We created a kinematic model of the thinner, and we feed it commands as the data replays. From this, we can play back the expected positions of the thinner and therefore observe the performance of the algorithms before testing them on the field.

For sensing, we mounted a Sick scanning laser rangefinder to the side of the vehicle such that it is scanning vertically as the tractor moves forward. We also added a Trimble Ag GPS mounted near the roof of the tractor. By logging the laser data and interpolating the GPS we are able to build a map of the canopy as the vehicle drives down the row. The control algorithm uses this map to decide where the thinner should be at each time step and then sends commands to the embedded controller to move the thinner. The sensor is mounted sufficiently far ahead of the thinner to give the computer time to process the laser data and actuate the thinner (Figure 78).

Figure 78. The modified Darwin thinner mounted onto a tractor. Note the ultrasonic mast visible above the tractor roof.

To implement an algorithm that maximizes contact while avoiding collisions, we build a local voxel grid and store the number of laser returns in each 10 cm cube in space. This gives a discrete representation of the tree which can be used for testing whether or not the spindle is in collision with the tree. We then find the set of collision-free positions of the thinner based on a discrete representation of the machine, breaking up the different possible angle and offset commands into those reachable within 200 ms. From the available thinner positions which are not in collision, we choose the one which is the furthest extended into the trees and therefore is closest to the majority of the tree without colliding with it. This choice is made using a heuristic that gives preference to lateral extension over angle extension since we expect the trees to be naturally angled towards the tractor.

In addition to simulating this algorithm and visually observing that it follows the trees well, we successfully tested the method in an apple orchard in runs with an inactive spindle. For the physical test, we made the modification that the voxel grid is entirely in the frame of the tractor and neglects the GPS. Instead, a set of occupied cells and commands are built based on each

200 ms of laser scans and new commands are sent to the thinner with a delay of 2.4 s. This time interval was determined assuming that the tractor is travelling at about 1 m/s (2.2 mph), the thinner is about 3 m behind the sensor, and that the system needs a bit of lead time to compensate for the controller lag.

Comparison of Ultrasound and Laser Sensors

Figure 79 presents a preliminary comparison of the orchard representations that can be obtained with the laser rangefinders and the ultrasound sensors. Clearly the laser provides much denser data for the control algorithm to use, albeit at significantly higher computational and installation costs. These are initial results and do not take into account the fact that we observed hardware problems with the ultrasound sensors, which we are in the process of diagnosing.

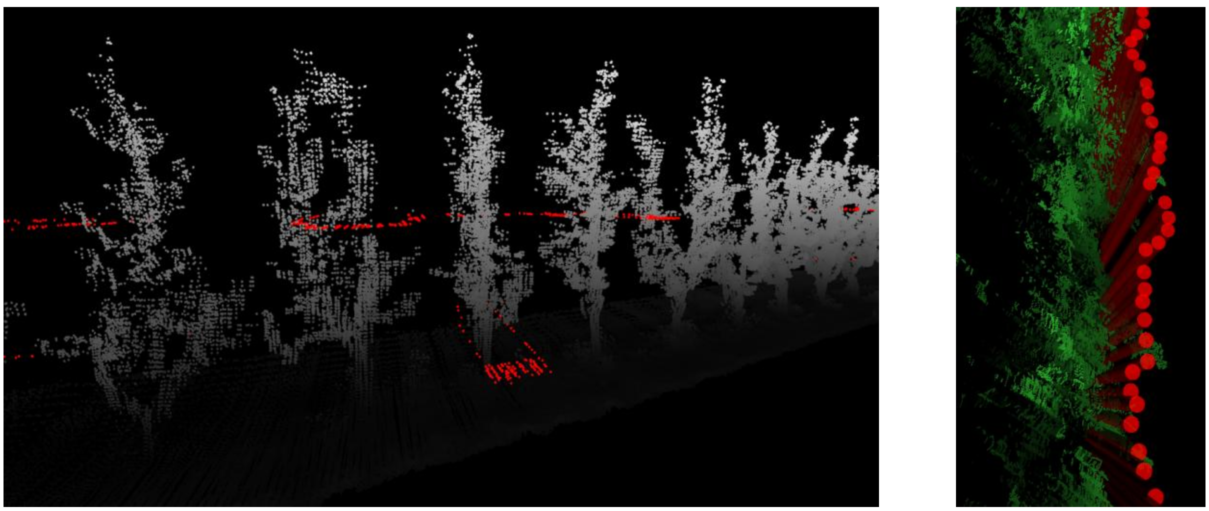

Figure 79. (Left) A 3D rendering of the trees obtained with the laser rangefinder (grayscale) and the ultrasound sensors (red). The ultrasound sensors detect less trunks and trees than the laser. This may be due to hardware difficulties we experienced, or a fundamental problem with sensing trees using ultrasound. (Right) A 3D rendering of the trees obtained with the laser rangefinder (green) and the thinner's path (as red cylinders) during field testing. Note that the thinner "hugs" the edge of the trees while mostly avoiding collisions.

Two Years In: The Direction of Reconfigurable Mobility

To this point we have chosen to use a limited sensing model with one or two laser range finders mounted rigidly to the vehicle. The technology exists to pan or tilt the lasers, but it more than doubles the cost of sensing. It also introduces moving parts which would need maintenance and repair, which is not ideal for growers who are unfamiliar with this technology. Our model has the advantages of low cost and reliable hardware. The fixed, horizontal-plane range scans, however, provide only a one-slice view of the world. This is the largest limiting factor in our work to improve the reliability of the autonomous navigation system. That the robot can only see objects at one height has a profound effect on safety. Multiple growers have asked us at demonstrations and field days what would happen if the vehicle came across a worker or trespasser lying on the ground in their orchard. With this model the vehicle cannot see such people or crates or other tools left in the orchards. For now we do not have to worry about such a question, since all tests are monitored, but it is something to consider as the system moves towards being put into use by growers.

We have also described above (and in previous reports) the many limitations this sensor model puts on the reliability of the autonomous navigation system. Since the sensors cannot see all of the canopy, there will be the possibility of canopy growing into the travel lane at a height not seen by the sensors, which the vehicle could then collide with. For this reason, and for safety considerations, in the coming year we will experiment with a new laser range technology called flash ladar that provides a complete range image, much like a camera. We will test the capability of these new sensors to deal with sunlight, dust, and other environmental conditions. For now, though, our plan is to move forward with the existing sensor model while we focus on other matters.

In the first two years of the project we explored the breadth of possibilities for the use of autonomous vehicle technology in orchards. We asked what operations could be completed with an autonomous vehicle, and how to achieve them. We have shown the ability of an autonomous vehicle to mow grass and carry a selective sprayer, the WeedSeeker. We have worked to make autonomous navigation as reliable as possible within the limits of a low-cost sensor model. This work culminated in a series of experiments which showed the ability of the system to drive in many canopy types for long periods of time.

No matter how reliable we make the autonomy software, though, an autonomous vehicle will at some time fail. Sensors will become occluded; cable connectors will break; and network connections will drop. How the system deals with such failures, and how the owner or user of an autonomous vehicle will recover from it, is an important question. It will remain a barrier to adoption until it is solved. Our work in Year 3 will begin to answer this question.

At this point in the project we believe it is more important to focus on the issue of usability, to make a concerted effort towards deployment of autonomous vehicles in orchards, rather than to continue major effort towards improving the reliability of the software in those few cases where it fails now. This is not to say that we will not try to improve reliability, but that it should not be the focus of our work in the next two years. In Year 3 we will deploy two new autonomous vehicles, one at Penn State and one at Washington State. Each will be a small-scale platform vehicle intended to carry two workers engaging in tasks like pruning, training of branches, and thinning. Our major work this year will be in the deployment of these vehicles for on-board use. We must prepare for our software to be robust to the kinds of errors described above. We must sort out the details of how to start up computers and run the system. And most importantly, we must continue our work to produce a simple, easy to use interface. It is our hope that growers and extension workers will see these vehicles being put into use and understand how the vehicles would improve their own operation. This way we will move closer to a reality of autonomous vehicles in orchards.

4.2 Accurate Positioning

Thematic area leader

Name	Institution	Email
George Kantor	Carnegie Mellon University	kantor@ri.cmu.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Refine APM localization algorithm by incorporating laser odometry. 2. Implement new algorithm on APM and test extensively. 3. Develop localization approach for outside of rows.	1. In-row localization software implemented onboard APM. 2. Report with analysis of results of in-row localization software. 3. Out-of-row localization software implemented and integrated with APM control system.	1. Demonstrate on-line positioning with less than 25 cm error 90% of time in 20 hours of general operation (in rows and during turning). 2. Demonstrate positioning during end-of-row turning and show accuracy sufficient for automated row entry. 3. Integrate positioning with end-of-row turning and demonstrate successful row entry with < 1% failure in at least 100 trials.

Notable results:

- Developed a new positioning algorithm that eliminates the need for reflective features inside rows (reflective features required at ends of rows only).
- Demonstrated sub-meter positioning accuracy in over 100 km of field trials (on-line and off-line).
- Integrated positioning system with APM autonomy to enable position-based navigation during both in-row and out-of-row operations.

Accurate positioning efforts during Year 2 were primarily focused on the development of new algorithms that are designed to improve the reflective feature-based positioning algorithm by making it less dependent on the use of engineered features (e.g., cones with reflective tape manually placed along the rows). Specifically, two approaches were investigated: laser odometry and row detection. Both of these techniques can be easily fused into the extended Kalman filter (EKF) framework for positioning that was developed in Year 1 of the project. Both laser odometry and row detection work using naturally existing structure, so they should provide a means of improving positioning accuracy while simultaneously reducing required landmark density. In addition to these primary efforts, a new technique for laser calibration was developed and implemented. This provides a means of quickly recalibrating the laser while in the field, which is necessary from time to time due to the bumps and bruises that the robot sometimes sustains during field experimentation and transportation.

In order to understand how the laser odometry and row detection efforts fit into the EKF framework, it is important to understand how the EKF works. Figure 80 shows the structure of the EKF localization filter that was developed in Year 1. It has two basic components: a prediction step and an update step. In the prediction step, the vehicle position estimate is advanced using the wheel encoder measurements of distance traveled and steering angle. In the update step, the vehicle position estimate is improved using the laser scanner measurement (range and bearing) to a reflective feature in a known location. By iterating these two steps of prediction and update, it is possible to keep an accurate estimate of where the vehicle is as it moves through the orchard.

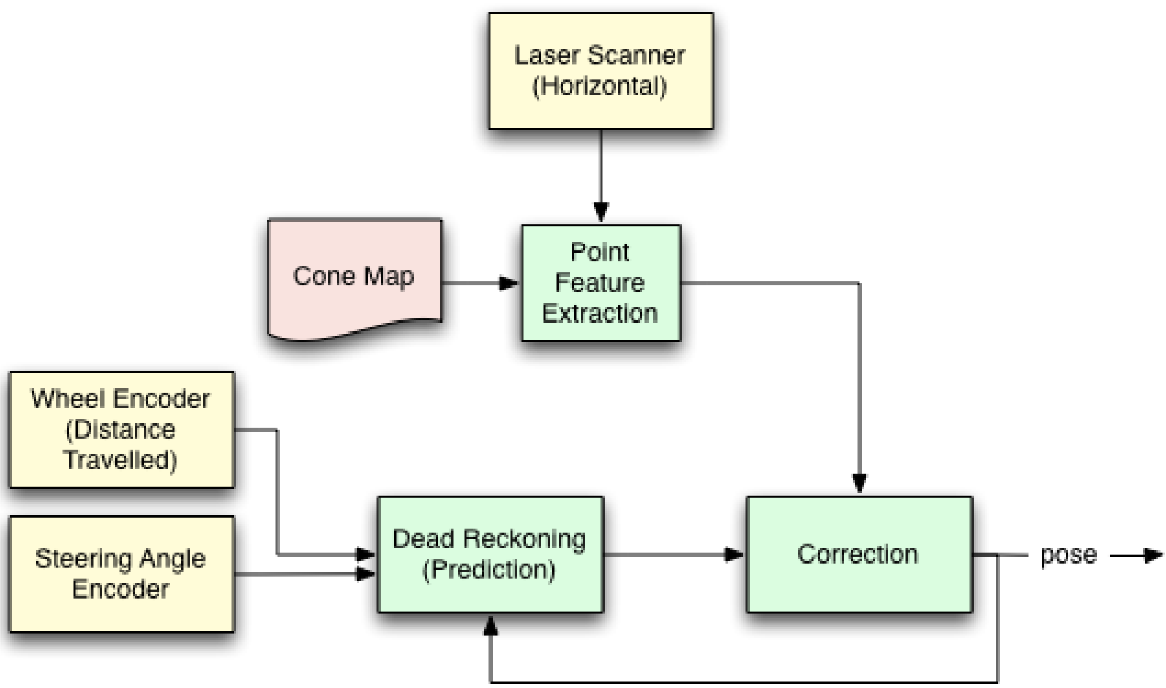

Figure 80. The extended Kalman filter structure of the localization filter developed in Year 1.

Laser odometry uses the laser range scanner to provide an estimate of vehicle motion that is more accurate than the motion estimate provided by wheel odometry, so the laser odometry and wheel odometry measurements can be fused together to provide a more accurate EKF prediction step. Row detection provides a means to correct crosstrack (side-to-side) errors in position by comparing the detected rows to a known map of rows in the orchard. This provides a means to correct the position estimate when the vehicle is not within range of a reflective feature. A block diagram showing how these two techniques can be combined with the Year 1 filter is shown in Figure 81.

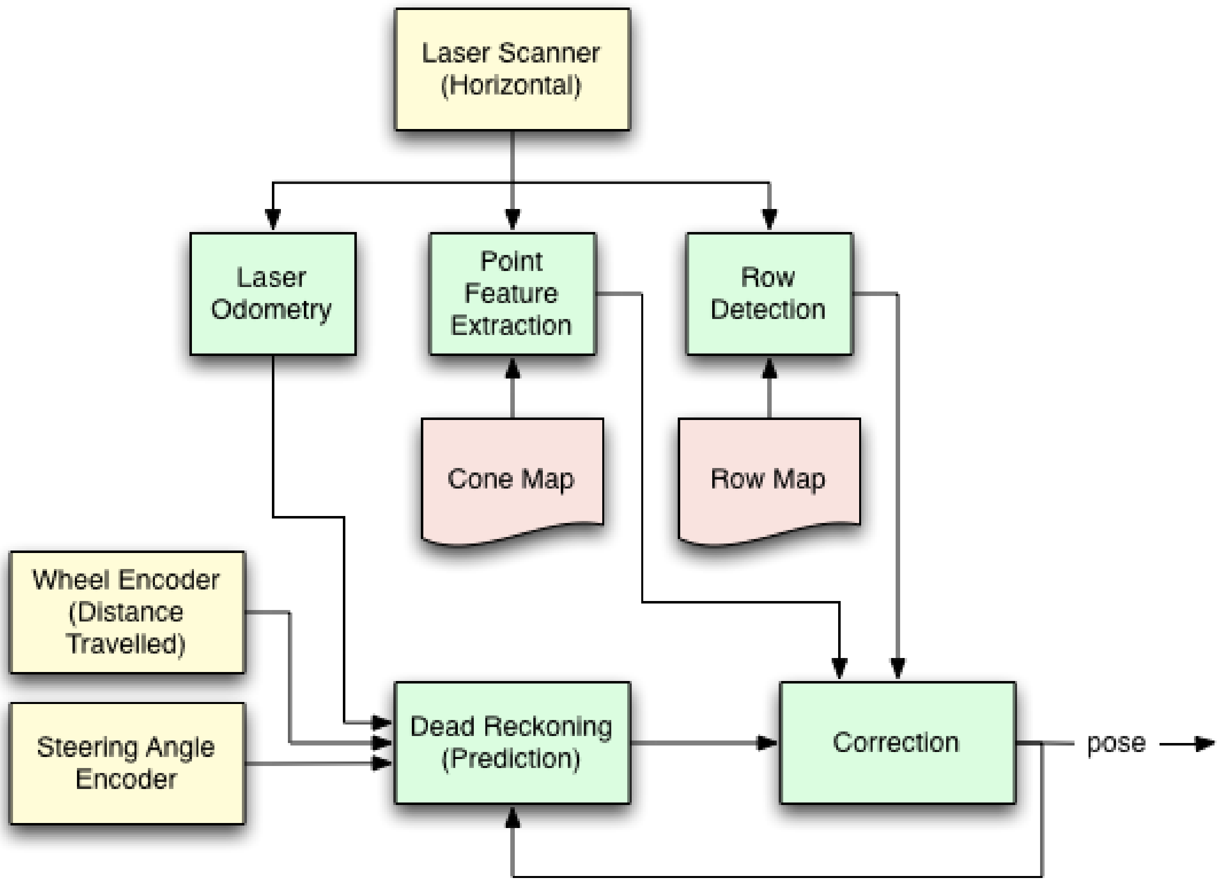

Figure 81. The new EKF with laser odometry and row detection incorporated.

Laser Odometry: In this technique, consecutive laser scans are matched to provide an accurate estimate of the vehicle motion between the two scans. This is achieved by searching for a translation and rotation so that, when the first scan is translated and rotated by those amounts, it lines up as closely as possible with the second scan. The resulting translation and rotation then provide an estimate of the rotation and translation by which the vehicle moved between the first and second scans. Since we know the time elapsed between the two scans, the translation and rotation found in scan matching can be used to get estimates of the linear and angular velocities of the vehicle, respectively.

Theoretically, this approach naturally exploits features in the environment without the need to do explicit feature extraction. We have found, however, that in practice the algorithm works better when features can be extracted and the rotation and translation are found by matching the extracted features rather than the raw scans. Some orchard environments contains many features that are useful for this purpose, most notably tree trunks and trellis poles.

We have implemented this feature-based scan-matching laser odometry and tested in post-processing on real data with varying results. In orchards where the tree trunks can clearly be seen by the laser, the laser odometry works reliably and provides a more accurate estimate of vehicle motion than can be had with wheel odometry alone. In orchards where leaves or weeds occlude tree trunks, the results were not as reliable. Further, we found that the majority of our test sites fall into this second category. For this reason, we decided not to incorporate laser odometry into the combined filter for Year 2. The approach, however, does show some promise, so we may attempt to improve it in future work.

Row Detection: Our row detection approach uses the line extraction algorithm developed for autonomous row following for the purpose of localization. This provides a means of estimating the lateral position of the vehicle within a row, allowing an EKF update step to be used to correct crosstrack error without the need for reflective features. The dominant error observed in the reflective landmark-based positioning algorithm developed in Year 1 was due to crosstrack drift in the sections where no landmark was visible, so row detection should provide a means to correct those errors, allowing the reflective features to be place further apart. We characterized the noise present in the row detections and used the resulting row model to develop an EKF update step for row measurements.

Combined Filter: The row detection work was merged with reflective-feature filter from Year 1 to create an integrated filter. This combined filter and its integration with the rest of the APM autonomy system was the principal result of the Year 2 localization efforts. This filter was tested extensively during Year 2 in over 100 km of on-line and off-line tests using the APM. This filter proved to be capable of localizing the APM to within about 70 cm crosstrack and about 1% of the distance between reflective features in downtrack (along the direction of travel). The error metric used here is 3σ (“three-sigma”) error, meaning that the distance between the actual vehicle position and the estimated vehicle position is smaller than the 3σ error metric 99.9% of the time. Table 13 shows some typical results from a subset of Year 2 trials. In all of these experiments, reflective landmarks were place at the ends of rows only, so the distance between landmarks is the same as the row length.

Integrating with APM Autonomy System: The combined filter was integrated with the APM autonomy system so that the position estimates it generates can now be used for autonomous vehicle navigation. Two of the 10 km autonomous runs that were conducted by the APM in Year 2 were executed with the vehicle navigating from the combined filter position estimate.

Table 13. Summary of combined filter localization results from Year 2.

Test Site	Row Length (m)	Total Distance (km)	Mean Crosstrack Error (m)	Mean Downtrack Error (m)	3σ Crosstrack Error (m)	3σ Downtrack Error (m)
FREC	125	2.05	0.18	0.28	0.78	1.13
FREC	125	2.00	0.14	0.34	0.60	1.82
Sunrise	53	1.38	0.15	0.17	0.54	0.65
Sunrise	53	1.78	0.18	0.16	0.62	0.61
Sunrise	53	1.59	0.17	0.25	0.60	1.01
Skyline	345	2.94	0.23	0.66	0.67	3.16
Skyline	345	10.21	0.22	0.73	0.64	3.48

Directions for Year 3: We learned lessons during Year 2 that will inform our efforts in the coming year. First, the accuracy of 25 cm that we had originally set out to achieve is not required for any known orchard operation and this level of accuracy is difficult to achieve without placing landmarks inside the orchard rows. Second, the localization system in its current form is difficult to use. For all practical purposes, the system can only be used with the direct involvement of its developer. With these two lessons in mind, our direction for Year 3 will be to relax the accuracy requirement and develop the reliability, ease-of-use, and interface that make it possible for non-technical personnel to use the localization system for APM operation and data collection. In fact, for most uses, our goal is to make the localization system “invisible” in the sense that a non-technical user would not make direct use of the localization system itself. Rather, they would operate the APM, and the associated localization functions would be automatically invoked and managed without direct attention from the user.

4.3 Augmented Harvesting

Thematic area leader

Name	Institution	Email
William Messner	Carnegie Mellon University	bmessner@andrew.cmu.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Develop platform-mounted of bin filler with integrated fruit transport.	1. Demo in test orchard. 2. Report with quantification of speed and bruising.	1. Speed up individual worker productivity by 25% (e.g., from 3,000 to 3,750 apples picked and loaded per hour) with no increase in bruising.

Notable results:

- Discovered that a variety of energy-absorbing materials can reduce bruising in passive bin fillers, including large sheets of inexpensive industrial bubble pack.
- Determined that apple to apple impacts are the most significant cause of bruising, and thus singulation of fruit during transport from tree to bin is essential.
- Early trials with DBR vacuum-assist transport and bin filling system showed at least a 10% improvement in harvesting speed with 5% reduction in bruising.

Introduction

In Year 2, after careful evaluation of the direction and productivity of IONco, we decide not to renew IONco's contract. We instead opted to focus on the passive bin filling designs while searching for another partner to work on transport of fruit from the tree to the bin. Thanks to the efforts of graduate student Brian Kliethermes, in January 2010 we identified DBR Conveyor Concepts of Conklin, MI as a new partner to develop a vacuum-based fruit transport system with a rotary "elephant-ear" bin filler. The company formally joined CASC in April 2010 and proved to be a successful partner.

Passive Bin Filling

Bins remain the prevailing method for transporting apples from the field to storage and processing facilities. The process of filling bins is a bottleneck for the efficient harvest of apples because workers must trade-off careful handling of apples to avoid damage with the desire to empty bags quickly into the bins.

Carnegie Mellon, Penn State, and USDA-ARS developed designs for two prototype bin fillers that showed promise in laboratory testing for reducing damage to fruit during bin filling:

- Energy absorbing grate bin filler: frames of energy absorbing materials strung on elastic bands.
- Pneumatically self-adjusting bin filler: parallel inflatable soft polymer cylinders attached to a frame and of an external air supply for pneumatic inflating of the cylinders.

Figure 82 shows one instantiation of the energy absorbing grate bin filler. Figure 83 shows the concept for the pneumatic self-adjusting bin filler, and Figure 84 shows the test rig for the full-scale prototype. Our objective was to determine if either of these two approaches could be successfully adapted to use in the field to both reduce fruit damage and increase harvest speed.

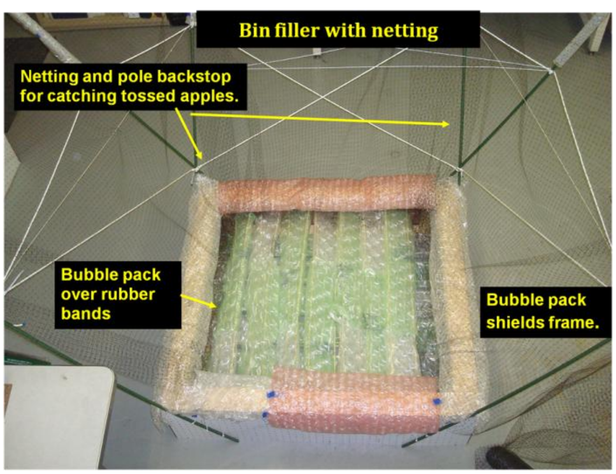

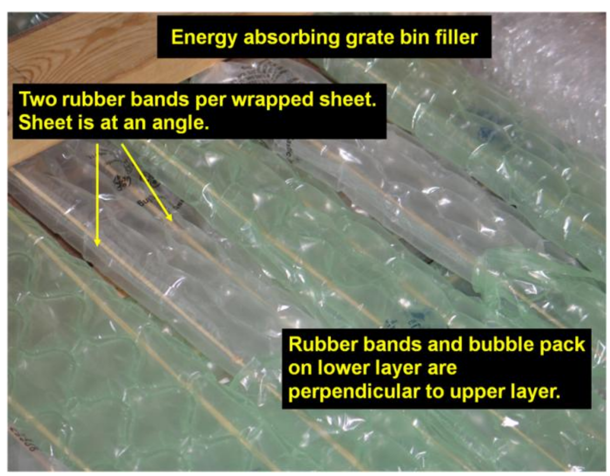

Figure 82. (Left) Energy absorbing grate bin filler. The netting was for experiments in guiding tossed apples into the bin filler. (Right) Energy-absorbing grate bin filler with bubble pack.

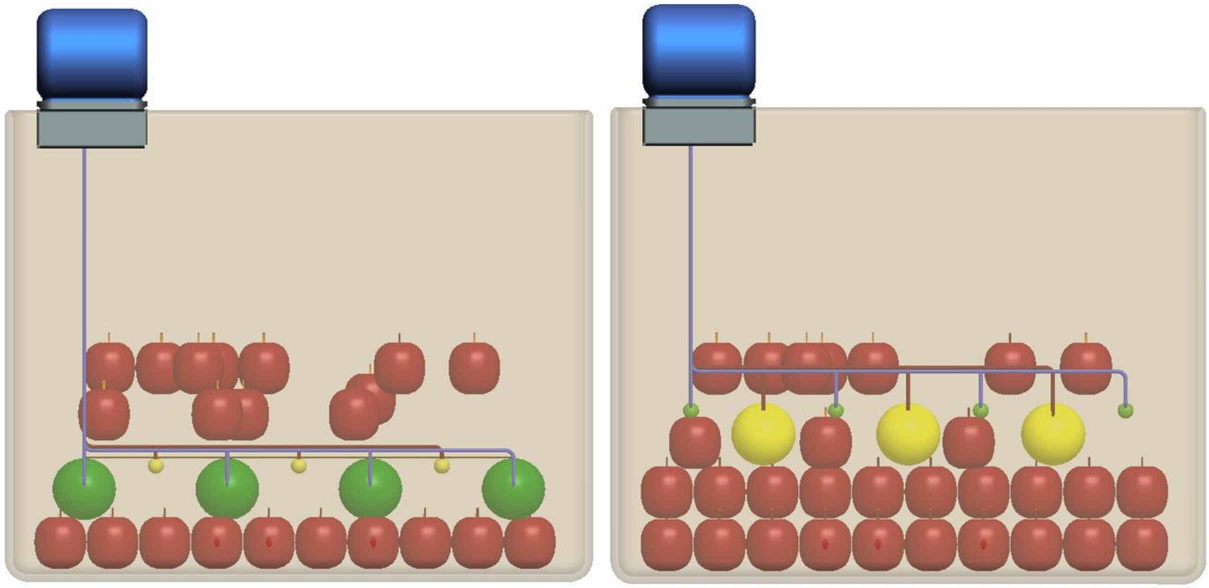

Figure 83. Pneumatic self-adjusting apple bin filler concept.

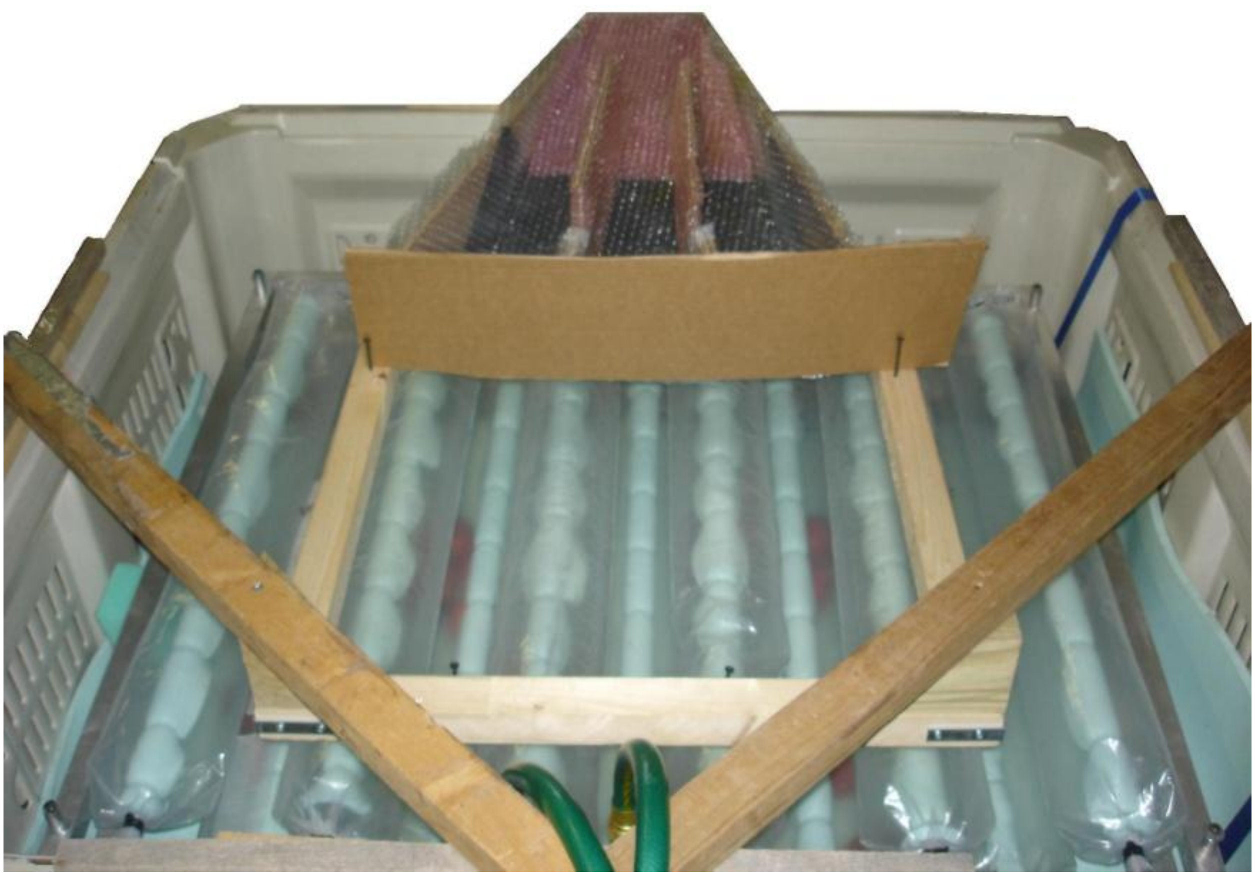

Figure 84. Pneumatic self-adjusting bin filler with test ramp.

Significant Findings

Our major findings regarding the passive bin fillers were the following:

- Both types of bin fillers can reduce bruising when apples are dropped into them one at a time, and they are positioned within 2-4 cm of the top layer of apples in a bin.
- A variety of different energy-absorbing materials successfully reduced bruising, including large sheets of inexpensive industrial bubble pack.
- The pneumatically self-adjusting bin filler could not lift itself because of the compliance of the polymer cylinders.
- When apple were dropped into a chute leading to the bin fillers, the impacts between the apples in the chute and in the bin filler caused significant bruising.
- Nets for guiding apples into the bin show promise for reducing the need for picking bags.
- Singulation of the apples during transport into and through the bin filler is essential to reduce bruising.

Results and Discussion

We determined that a variety of energy absorbing materials, such as foam balls strung on rubber bands, are suitable for the bin fillers themselves. Significantly, even inexpensive, easily replaceable bubble pack can work, provided the pressure in the bubbles is high enough and there is the right amount of space between the energy absorbers so that they slow down the apples without completely stopping them.

The pneumatic self-adjusting bin filler needs some modification to work as intended. The soft polymer cylinders are too compliant, and they bend under the weight of the rest of the mechanism when inflated. The middle of a cylinder remains in contact with the central inflation tube, and thus cannot lift the mechanism. The solution may be as simple as tying off the cylinder into discrete, separately inflatable sections.

Both types of bin fillers were effective at reducing bruising as intended when apples were dropped one at a time into the bin filler with enough time between successive drops to prevent apples from hitting each other. On the other hand, pouring a bag of apples into the bin filler from the side or into a chute leading to the bin filler resulted in excessive bruising.

To address the problem of singulation, we also tried a system of netting to catch and guide apples into the passive bin filler. This method still had excessive bruising, but the problem was with impacts between the apples and the sides of the bin filler. There was insufficient padding on the edges of the bin filler and the netting did not sufficiently guide the apples to the middle of the bin filler away from the edges. Considering that apples were tossed from heights above 1.8 m (6 feet), the fact that the apples survived at all was encouraging. We believe this concept hold considerable promise for moving apples from the tree to bin without bags.

Summary of Passive Bin Filling Effort

The primary result on the passive bin fillers effort was determining that the singulation of apples during transport from the tree through the bin filler is the key to reducing bruising for passive bin filling (Figure 85).

Overall, the effort on passive bin filling provided several valuable insights that will be useful for the industry even though we did not yet achieve the goal of having a field-deployable bin filler. The fact that we did not reach our goal in Year 2 was due in part to changes in personnel during the project, which delayed the implementation of modifications to our prototypes, rather than insurmountable technical challenges. We believe that what we learned will help the industry develop cost-effective active or passive bin filling and transport mechanisms.

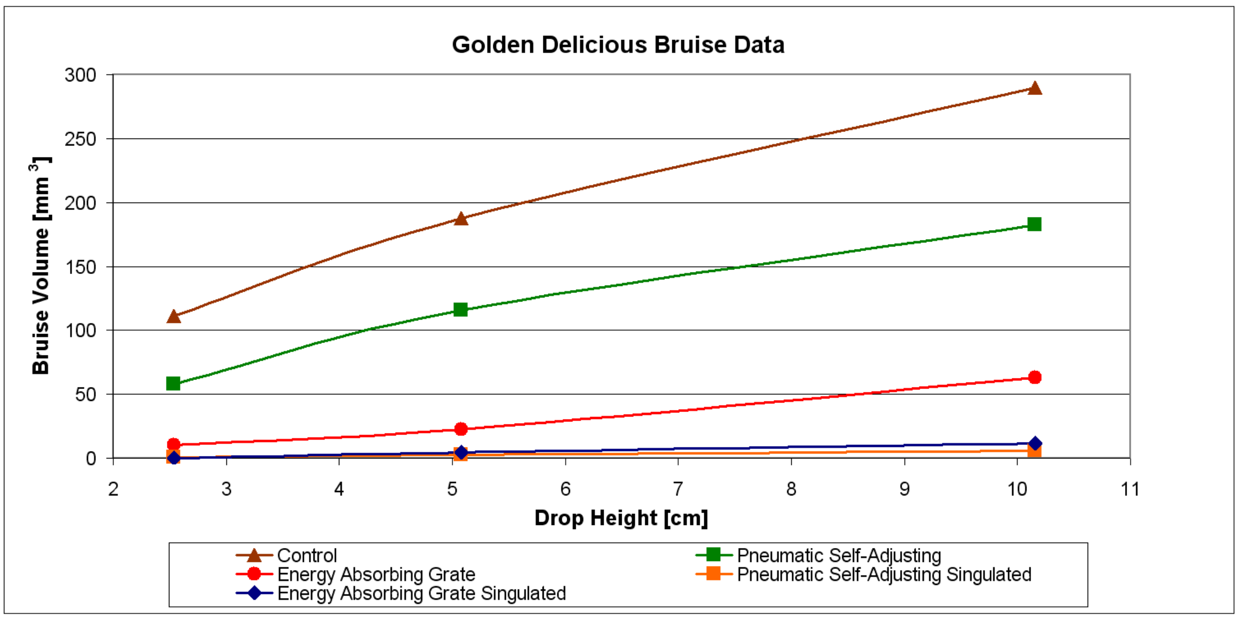

Figure 85. Trend lines comparing the performance of two full-scale bin filling prototypes across three drop heights with and without singulation. Singulation improves performance by a factor of 10 for the energy absorbing grate and by a factor of about 100 for the pneumatic self adjusting bin filler.

DBR Conveyor Concepts Vacuum Transport System

Eliminating ladders and picking bags has the potential to greatly increase harvest efficiency provided that bruising of the fruit is prevented. Phil Brown Welding of Conklin, MI unveiled a prototype vacuum transport system in fall 2009 in which pickers placed fruit in vacuum hoses that led to a “decelerator” mechanism that dropped the fruit into rotary “elephant ear” bin filler. They demonstrated low bruising, but the system was integrated with a very large vehicle making it impractical for use in orchards due to a high turning radius.

Based on the demonstration of the vacuum technology, CASC enlisted the three inventors of the technology, Chuck Dietrich, Phil Brown, and Mike Rasch, who together formed DBR Conveyor Concepts to develop a smaller modular system that could be mounted on a harvest platform or on a tractor with hydraulic forks. Only four months after contracts were signed in April 2010, DBR completed their first prototype mounted on an N. Blosi platform towing a bin trailer (Fig. 6). After preliminary testing in Michigan, DBR delivered it to the FREC in Biglerville, PA in early September for field trials.

The DBR system consists of two subsystems—the vacuum transport and the elephant ear bin filler. Each of these two subsystems can operate independently of the other. The vacuum transport consist of dual vacuum transport hoses, a new dual proprietary single-wheel decelerators (patent pending), dual vacuum return hoses, and two vacuum pumps with a single internal combustion engine providing power (Figure 86). Figure 87 shows a close-up of a worker picking apples and the vacuum hose inlet. The primary innovations of the vacuum transport design are (1) the lining of the vacuum hoses with neoprene to prevent bruising as apples move through the hose and (2) the single wheel decelerator, which “catches” the apples moving at high speed, and gently lowers them out of the vacuum into ambient pressure while maintaining a tight seal. The single wheel decelerator employs a hydraulic motor with significantly low size, weight, and power requirements.

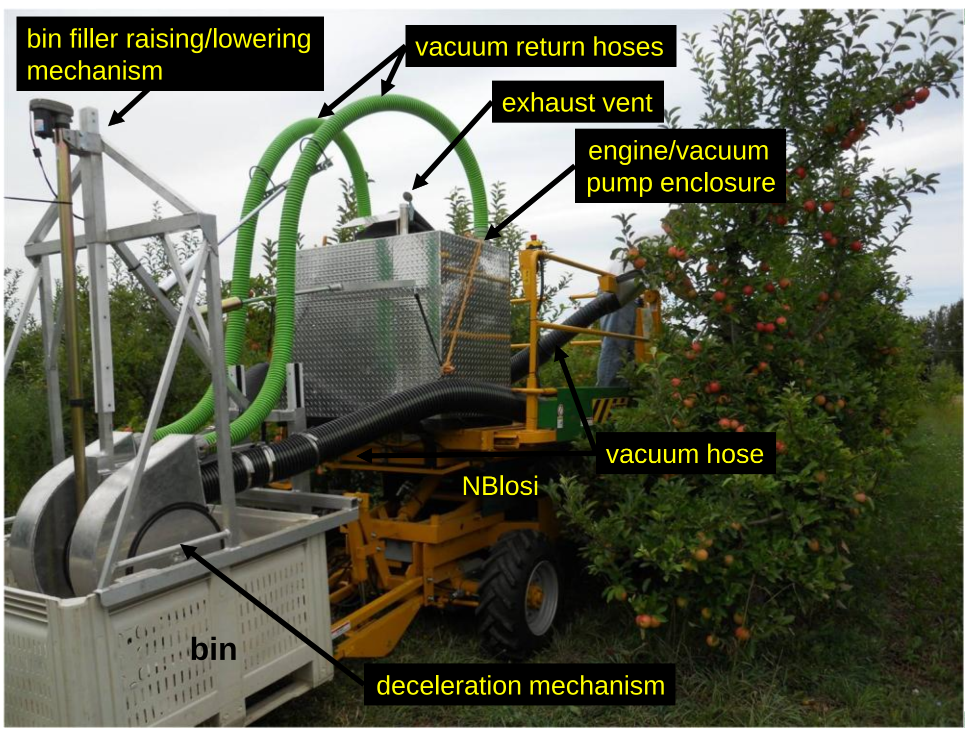

Figure 86. DBR Conveyor Concepts’ vacuum transport mechanism.

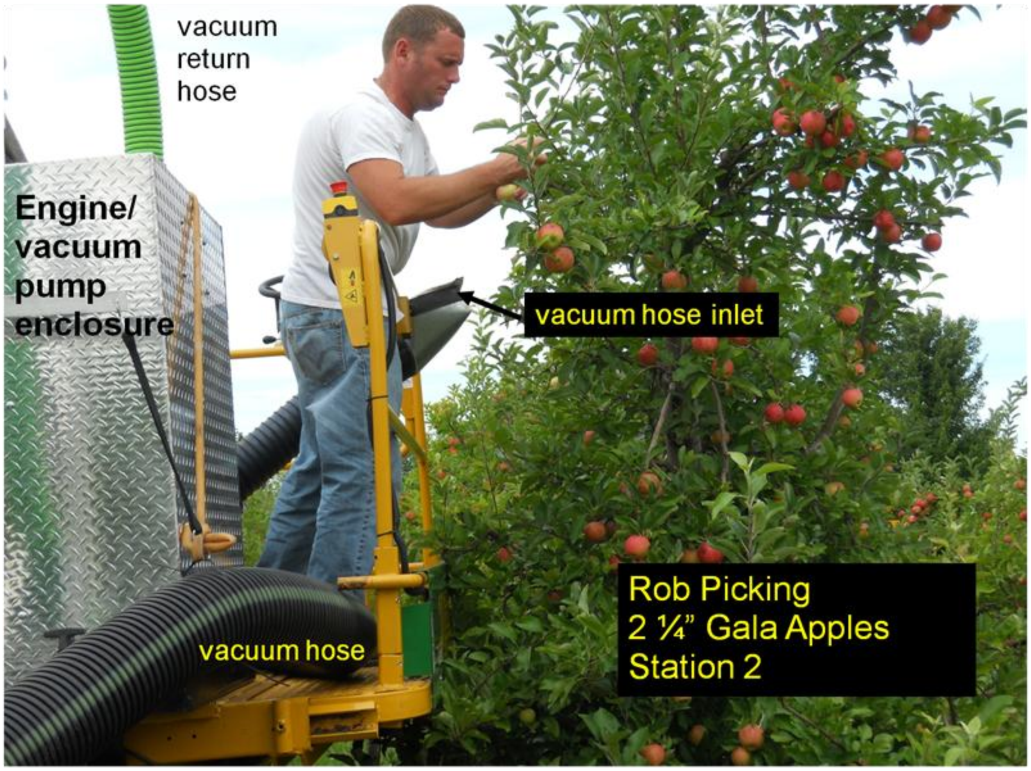

Figure 87. Close-up of DBR system showing worker and vacuum hose inlet.

The elephant ear bin filler is raised and lowered into the bin with a hydraulic motor. An electric eye measures the height of the bin filler above the top layer of apples in the bin. This sensor provides a signal to automatically raise the bin filler as the bin fills. When the bin is completely full, the workers manually control the lifting mechanism to raise the bin filler high enough to remove the full bin and replace it with an empty one. The workers then lower the bin filler into the empty bin.

Field Test Results

Bruise testing and efficiency trials were conducted with several cultivars. The complete results are shown in the appendices. Overall, the system increased efficiency by approximately 10% and increased the percentage of Extra Fancy apples by about 5% compared to hand-picking. On the other hand, the total percentage of culls increased from 2.3% to 3.0% while the percentage of apples with no detectable bruising decreased from 4.8% to 2.8%. The estimated increase in economic benefit due to the DBR system was $245/acre.

In summary, the DBR harvest assist system shows great promise for increasing efficiency while decreasing bruising. The initial bruise testing was on particularly soft (overripe) Golden Delicious, which would favor hand-picking. We believe that even greater efficiency gains of up to 25% are possible in orchards with architectures more suitable for the harvest platform. Development of a tractor-mounted system will reduce its cost and may be more suitable for orchard architectures in Washington state. This will be investigated in Year 3.

5. Technology Adoption

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.

This section presents the goals and accomplishments in the following three thematic areas:

- Sociological implications;
- Value proposition;
- Outreach and extension.

5.1 Sociological Implications

Thematic area leaders

Name	Institution	Email
Katie Ellis	The Pennsylvania State University	kag298@psu.edu
Karen Lewis	Washington State University	kmlewis@wsu.edu
Tara Baugher	The Pennsylvania State University	tab36@psu.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Execute socio-economic survey of WA growers. 2. Share WA and PA socio-economic survey results with advisory panel & industry stakeholders.	1. Surveys' results posted at Comprehensive Automation web site. 2. Surveys' results detailed in outreach, trade, & refereed journals.	1. 5% increase in growers' support in the form of participation, interest, and in-kind matching for Comprehensive Automation.

Notable results:

- Socio-economic survey conducted in the Pacific Northwest; results were compared with data obtained in the Mid-Atlantic region in Year 1. Some interesting findings include:
    - Crop projection sensor data more popular in Pacific Northwest than in the East;
    - Justifiable price point of technology that increases fruit packout is higher in the Northwest;
    - Eastern growers are less concerned with water availability effects on crop production.
- Survey outcomes described in the article "Results from Survey Instruments Used to Assess Technology Adoption for Tree Fruit Production," to appear in HortTechnology in Dec. 2010.
- Recommendations for CASC value proposition and outreach activities based on survey data.

Introduction

We drafted a socioeconomic survey in January 2009 to assess stakeholder input and potential obstacles to industry adoption of new technologies in tree fruit. The survey consisted of eleven pages of questions preceded by a brief explanation of the CASC project and the purpose of the survey. The questions were grouped into related sections: (1) Specific information about the participant's farm enterprise; (2) Needs/potentials for automation and sensor technologies in specialty crops; (3) Potential benefits of harvest assist (semi-automated harvest) technology; (4) Potential benefits of automated disease detection and pest monitoring technologies; (5) Potential benefits of automated technologies for monitoring plant stress; (6) Benefits of fully automated harvest technologies; and (7) Specific orchard system planting information.

Several questions were designed to address regional differences between the Eastern and Western U.S. fruit industries; specifically, information about irrigation systems and current tree training may help determine more precise regional barriers to adoption of orchard technologies. The survey was initially distributed to audience members at the Tree Fruit session held during horticultural conventions in Pennsylvania and New York in February 2009 (results summarized in the CASC Year 1 report). The same survey was conducted at the December 2009 Washington State Horticultural Association NW Hort Expo in Wenatchee, WA.

The NW Hort Expo yielded 38 respondents. Compared with Eastern participants, the Western growers grew less diverse crops but tended to manage more acreage. The distribution of acreages managed by each respondent (Figure 88) was similar to that of the Eastern growers, but most grew only one, two, or three different crops. The annual gross operational revenues were also more evenly distributed across categories, thanks to more operations in the >$5M category (Figure 89).

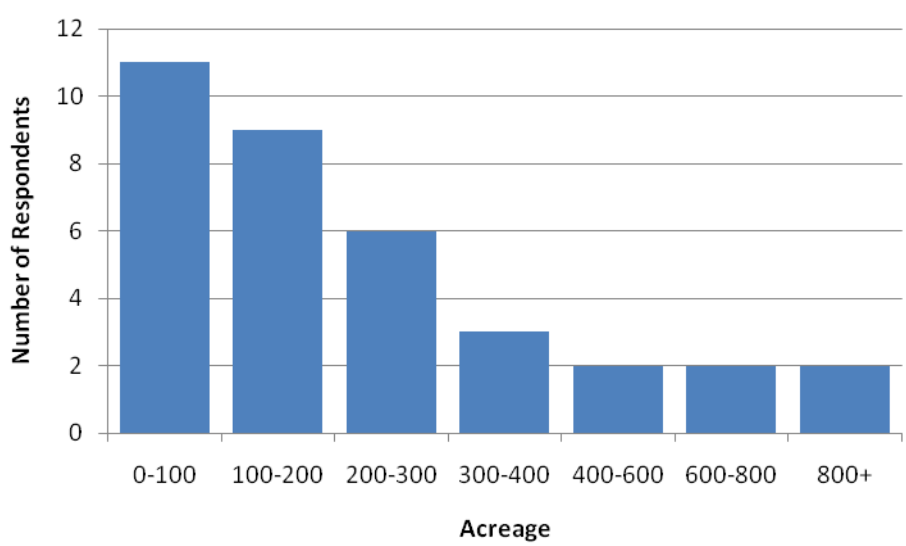

Figure 88. Total number of respondents that own or manage 0-100, 100-200, etc. acres of tree fruit or vine crops.

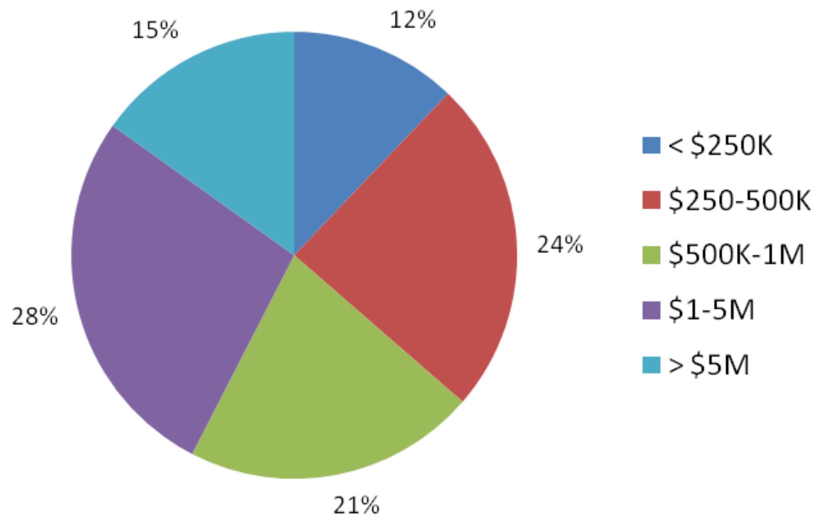

Figure 89. Percentage of growers in each annual gross revenue category.

Needs and Potentials for Automation and Sensor Technologies in Specialty Crops

To improve precision and efficiency in orchard enterprises, the participants rated fruit thinning, spraying and monitoring water/nutrients the areas of greatest need. The need scores for all other orchard activities were lower than those from the Eastern growers. Mowing was seen as least important. Harvesting, spraying, and monitoring water and nutrient status were identified as the areas of greatest need for improving environmental stewardship and sustainability (Figure 90). New spraying technologies received the highest ratings overall. These opinions were similar to those of the Eastern participants.

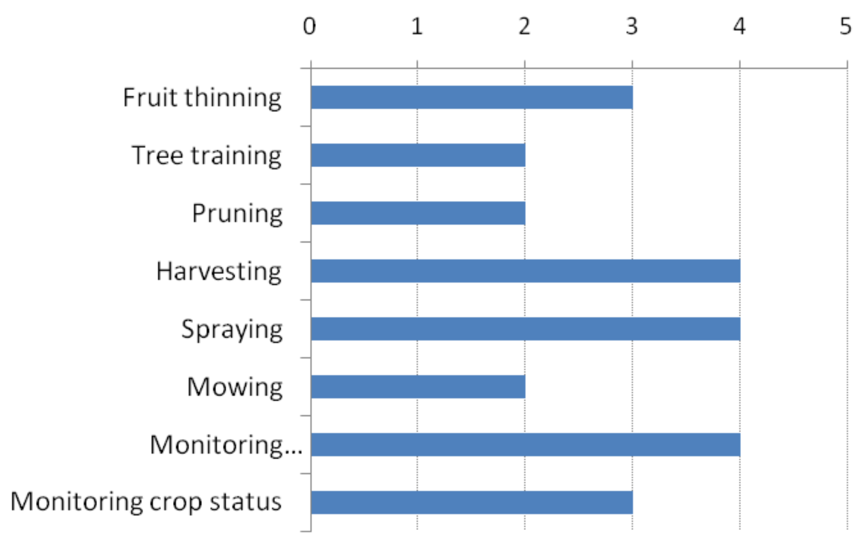

Figure 90. Median rating for each area of need in environmental stewardship and sustainability.

Potential Benefits of Harvest Assist (Semi-Automated Harvest) Technology

Many Western participants anticipated benefits with the suggested areas of improvement with harvest assist technologies; however, they were generally less enthusiastic than Eastern participants. As in the East, increased workforce productivity was selected most often (Figure 91). Eastern growers anticipated reduced costs with new technologies more than Western growers, who saw more benefits with the reduced need for a steady workforce. Western growers also did not anticipate increased fruit quality/packout as the Eastern growers did.

Figure 91. Percentage of participants anticipating benefits with harvest assist technologies.

Respondents were also asked to choose the three main obstacles to industry adoption of harvest assist technologies from the list shown in Figure 92. Cost and equipment reliability were most often selected as potential obstacles, while decreased employee retention, decreased safety, reduced control over management of harvest operations, and the need for specialized employee training were not seen as major obstacles.

Participants were asked to identify the maximum equipment price justified by an increased efficiency of harvest employees by 30-40%. The median price justified by Eastern and Western growers was $35,000. If harvest technologies increased fruit packout by 10-15%, Eastern respondents felt $25,000 was reasonable; however, Western growers justified a much higher price point of $55,000.

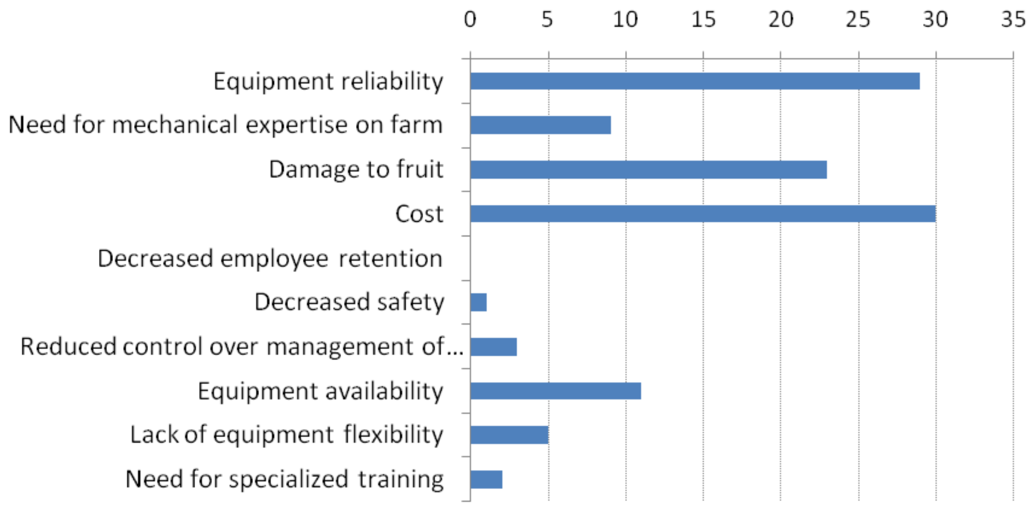

Figure 92. Number of respondents that selected each potential obstacle to industry adoption of harvest assist technology.

Potential Benefits of Automated Disease Detection and Pest Monitoring Technologies

Internal fruit feeding insects, such as codling moth and Oriental fruit moth, were seen as a problem for more of the Western participants (82%) than Eastern (52%). Apple pack-out losses, however, were distributed about the same in both areas (Figure 93). Many Western growers (as in the East) would consider increasing trap density with new technologies but not expand their program to larger areas.

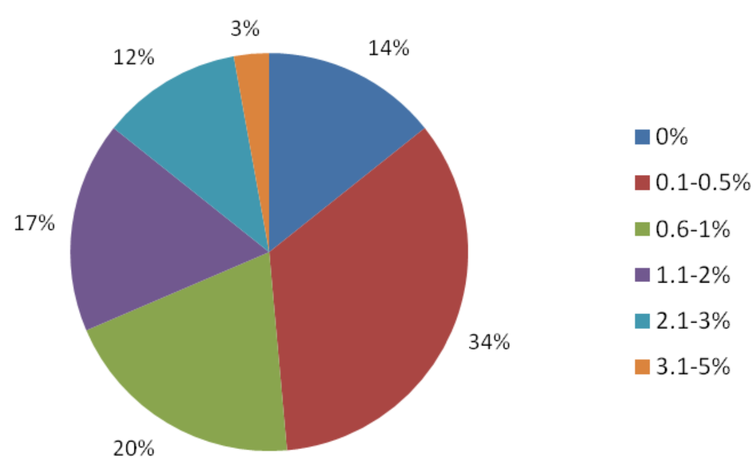

Figure 93. Percentage of Western participants reporting apple pack-out losses to internal fruit feeding insects in six categories.

Of the four potential obstacles posed by the survey to industry adoption of new sensor technology for insect monitoring, most felt that cost and equipment reliability under varying environmental conditions were most prohibitive. These responses were very similar to those from the East. Research validating efficiency in finding problems was regarded as potentially successful. Across both regions, early involvement of commercialization partners was seen as least helpful (Figure 94).

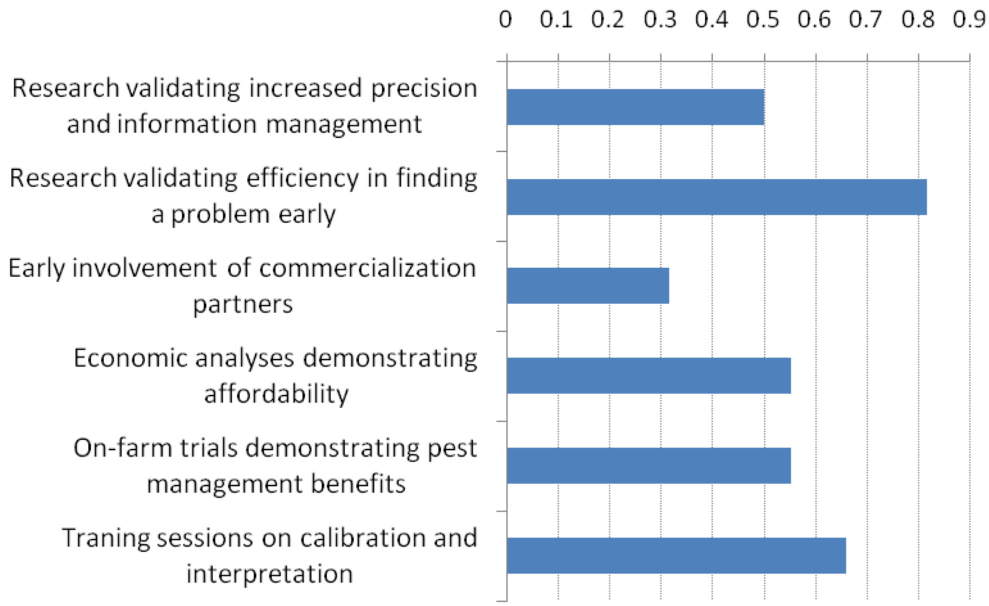

Figure 94. Percentage of respondents that were more likely to adopt precision systems for detecting disease and monitoring insects due to various research and outreach efforts.

Potential Benefits of Automated Technologies for Monitoring Plant Stress

Western participants were willing to incorporate sensor technologies into their irrigation systems with 35% energy and/or 50% water savings. In the East, willingness to redesign irrigation systems was fairly low, at a median response of 25% for each consideration.

Imaging Systems that can Scout and Map Crop Load

Improved preparation for harvest/fruit storage and accurate pre-harvest crop projections were seen as the most beneficial result from imaging systems that can scout and map crop load (Figure 95). Only 47% of Western growers saw a benefit from more efficient fruit thinning, vs. over 90% of Eastern growers.

As in the harvest assist responses, the major obstacles to overcome for adoption of fully automated harvest technologies are cost and equipment reliability.

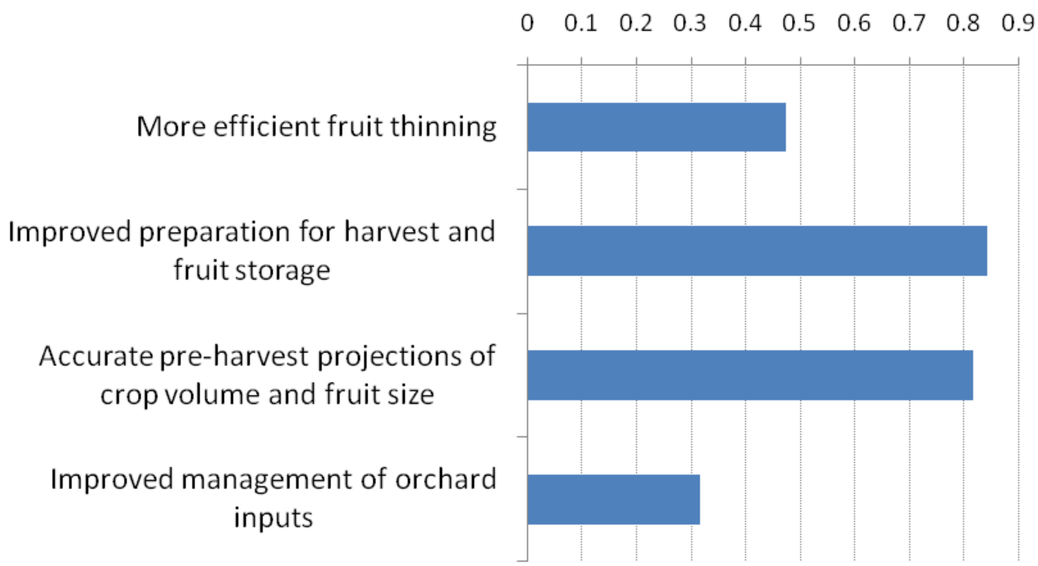

Figure 95. Percentage of respondents perceiving benefits in imaging/information management in various areas of crop management, harvest, and fruit storage.

Specific Orchard System Planting Information

Survey-takers were finally asked to specify the number of apple acres currently trained to a vertical fruiting wall, angled fruiting wall, and/or narrow fruiting wall; they were also asked to provide this information as a percentage of their total apple acreage, as well as planned tree training over the next ten years. Most growers that have trained their trees to fruiting walls have done so vertically, though angled fruiting walls were more common in the Northwest than the East. Those growers that intend on increasing their high density acreage over the next 10 years will do so by about 20%. In the Northwestern U.S., vertical, angled, and/or narrow fruiting wall systems will likely increase more in sweet cherries than in other tree fruit crops.

Regional Implications for Improvements in Technology Adoption

While most responses were consistent across the country, Western and Eastern growers indicated differences in irrigation concerns and justifiable price points for harvest-assist technology. This result suggests a benefit in using region-specific outreach topics to emphasize local needs. Western growers with larger pack-and-ship operations associate a benefit with packout improvement, while smaller Eastern retail-based businesses would relate better to emphasis on reduced labor costs and fruit quality improvement.

Western growers were also particularly interested in sensor data for crop projections, which may be due in part to recent disparities between projected and actual crops. Several months prior to the Western survey, Washington growers were faced with a large crop of unusually small apples. In addition, the 2008 actual fresh crop was nearly 11% more than the total projected that summer (assessment by Yakima Valley Growers-Shippers Association).

Therefore, recent reminders of crop variability may influence the participants' opinions and interest in specific sensor technologies.

These survey results are summarized in the HortTechnology article "Results from Survey Instruments Used to Assess Technology Adoption for Tree Fruit Production" to be published in December 2010.

Activities

Year 2 goals were accomplished by sharing the results with advisory panel members and industry stakeholders at several meetings:

- CASC showcase webinar (November 2009);
- Mid-Atlantic Fruit and Vegetable Convention advisory panel meeting (February 2010);
- American Society of Agricultural and Biological Engineers annual conference (June 2010);
- Engineering Solutions Technology Showcase and Advisory Panel Meeting (October 2010).

Survey results were also posted at the CASC web site (http://www.cascrop.com) where the ASABE slide show received 57 visits and the showcase received 36 visits.

5.2 Value Proposition

Thematic area leader

Name	Institution	Email
Clark Seavert	Oregon State University	clark.seavert@oregonstate.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Continue to collect field data from researchers. 2. Conduct a marketing assessment for the automated caliper and VRC's Scout. 3. Discuss with growers and others in PA and WA about the economic and financial information needed for project. 4. Schedule and conduct AgProfit™ workshops.	1. AgTool™ website. 2. AgProfit™ program, including training modules. 3. AgFinance™ program (beta version). 4. Report on the marketing potential of the caliper and Scout. 5. One AgProfit™ workshop each in PA and WA reaching a total of 30 growers.	1. Validation from growers and others that economic and financial information is representative of their industries in each region. 2. 20 growers will have an understanding of the economic and financial risks involved in adopting technologies proposed in this proposal.

Notable results:

- Technologies being developed by CASC researchers could derive a benefit at the grower level from $20 to $1,047 per acre per year. These benefits represent the upper limit of net returns an end user would receive from adopting a technology. This value does not include the costs to acquire, repair, service, or train employees to operate the technology or a return on investment to the end user for the risks associated with implementing the technology. If these costs were equal to the net present value annual equivalent, the end user would receive no benefit from adopting that technology.
- The market potential of the insect monitoring station for the primary market is estimated at $10 million in U.S. Plant and Protection Quarantine applications and an additional $50 to $100 million in U.S. tree fruit and nut crops. The potential in the secondary market is $10 million in international tree fruit and nut producers with a long range market of less than $100 million in U.S. cotton and other row crops.
- The development of the AgTools website allows growers, processors, packers, and technology providers access to risk management tools. These tools can assess the profitability and feasibility of CASC projects if implemented in their own businesses.

In Year 2 we conducted work in the areas described in the sequel.

(1) Research was conducted to develop economic models for four CASC thematic areas. Our preliminary findings suggest the technologies being developed by CASC researchers could have an economic impact ranging from $20 per acre per year to $1,047 per acre per year, if they achieve their stated objectives (Table 14). This value does not include the costs to acquire, repair, service, or train employees to operate the technology or a return on investment to the end user for the risks associated with implementing the technology. If these costs were equal to the net present value annual equivalent, the end user would receive no benefit from adopting that technology. Relating this information to the size of farm operation, a grower with 345.9 acres of apples in Washington State could achieve $95,126 of benefits per year by implementing the insect monitoring station (Table 15). A grower in Pennsylvania could achieve $34,580 and $27,758 of benefit per year from implementing the insect monitoring station on 235.2 acres for fresh or process apples, respectively.

Table 14. The net present value equivalent[1], per acre and by average farm size within size category, of selected CASC technologies.[2]

		Impact by average farm size [3]		
Selected CASC Projects	Per acre	Less than 50 acres	50 to 99 acres	Over 100 acres
1.4 Augmented Fruit Harvest, WA	$515	$5,105	$33,042	$178,146
2.1 Detecting Plant Stress, WA	$32	$317	$2,053	$11,069
2.1 Detecting Plant Stress, Fresh, PA	$20	$85	$1,258	$4,705
2.3 Monitoring Insect Pop., WA	$275	$2,726	$17,644	$95,126
2.3 Monitoring Insect Pop., Proc., PA	$118	$504	$7,423	$27,758
2.3 Monitoring Insect Pop., Fresh, PA	$147	$627	$9,247	$34,580
2.5 Auto. Crop Load Scout, WA	$1,047	$10,379	$67,175	$362,172

1 Assumes a ten year life for the investment with a 6% discount rate.
2 These values do not account for the costs to acquire and implement the technologies. These estimates provide an upper limit to the economic value a grower would receive from adopting one of these technologies. If the cost to acquire and implement a technology were equal to the net present value annual equivalent, economically, a grower would receive no benefit from that technology.
3 See Table 2 for acreage of average farms by region

Table 15. Number of farms by size of apple orchard, in acres, for the US and selected regions.*

		Less than 50 Acres	Percent of Total	50 to 99 Acres	Percent of Total	Over 100 Acres	Percent of Total	Total
US	Farms	20,122	93%	788	4%	806	4%	21,716
	Acres	93,448	26%	49,514	14%	217,234	60%	360,196
	Average Size	4.6		62.8		269.5		16.6
WA	Farms	2,324	80%	263	9%	325	11%	2,912
	Acres	23,039	15%	16,874	11%	112,422	74%	152,335
	Average Size	9.9		64.2		345.9		52.3
PA	Farms	1,482	94%	43	3%	50	3%	1,575
	Acres	6,324	30%	2,705	13%	11,762	57%	20,791
	Average Size	4.3		62.9		235.2		13.2

* Data: USDA, 2007 Census of Agriculture

(2) Commercialization potential of a CASC technology: Product Innovation and commercialization (OSU ENGR 599) class project

Oregon State University offers a graduate course that presents the process of taking an intellectual property developed by a university and creating a plan to commercialize the respective product. The course is divided into two sessions; in the first, the students assess the product's viability and during the second, the students develop a plan to take the product to market. In a recent course, a group of students choose to study the insect monitoring station under development by members of the CASC team.

The initial step was to assess the viability of the intellectual property as a product by working through the following list of topics (see the appendices for the completed report):

- Product Functionality
- Product Direct Costs
- Product Value
- Market Requirements
- Product Differentiation
- Manufacturing Explanation & Estimated Costs
- Delivery Explanation & Estimated Costs
- Support Explanation & Estimated Costs
- Profit Justification

After working through each of these topics for the technology, a summary assessment can be made as to whether the technology has the potential to be developed into a standalone enterprise, a product line extension for an existing enterprise, or simply isn't a viable product as currently envisioned due to costs, market size, or competition.

In the case of the insect monitoring station, the assessment resulted in a reasonable opportunity but the market size and estimated adoption rate suggested that the device may have better success as a product line extension rather than a product to build a company around. The student group decided that they would continue with this device into the second term of the course and develop a business model around the technology.

The second term project was to develop a presentation for potential investors describing an estimated market size and value, market channels and business model (see the appendices for the presentation). To do this the students contacted potential customers to get feedback as to how it might fit into their operation and any suggestions they may have to improve the device, as well as their general opinion of the technology.

During the first term there were five students involved in developing the product viability assessment. During the second term, only two of the original students returned and continued the project; the others did not register for the second term for various reasons. All the students involved with the CASC-related group projects earned the grade of "A" for their efforts. This project was presented as a showcase during the CASC meeting in September (see the appendices for the slides). The presentation described the economic assessment performed by members of the CASC team and the results of the student project.

(3) On March 4th, 2010, Clark Seavert and Jim Julian presented a "Got Economics?" workshop to 30 growers, processors and input suppliers in Biglerville, PA. They also met with 5 local lenders the next day to discuss improvements to the output of the AgFinance^{TM} software program. A similar meeting with lenders in Yakima, WA was held to seek advice from a western lender perspective.

Other accomplishments include: release of the AgTools^{TM} website, including download and installation videos (details in the appendices); development of AgFinance^{TM} software (alpha version) after advice from lenders (release date Winter 2010); brochures highlighting the AgTools^{TM} programs; presentation of the CASC project to Columbia Gorge Technology Alliance in Hood River, Oregon on March 16th.

5.3 Outreach

Thematic area leaders

Name	Institution	Email
Tara Baugher	The Pennsylvania State University	tab36@psu.edu
Karen Lewis	Washington State University	kmlewis@wsu.edu
Gwen-Alyn Hoheisel	Washington State University	ghoheisel@wsu.edu

Year 2 goals

Activities	Deliverables	Success Criteria
1. Conduct educational programs to prepare stakeholders for future adoption of new technologies. 2. Conduct management efficiency trials in whole-farm sustainable demonstration plantings. 3. Track rates of adoption of trellis-based planting architecture of growers in areas of influence of CASC outreach. 4. Increase student interest in developing engineering solutions for specialty crops. 5. Improve project relevance and impact through structured input and feedback.	1. In-depth risk-management workshops for producers and farm employees. 2. Comprehensive Automation web site, video-cast, and extension publications. 3. Management efficiency and economic impacts from on-farm trials reported in research and outreach publications. 4. Engineering interns trained to assist in developing engineering solutions for specialty crops. 5. Report on rates of adoption of trellis-based planting architecture of growers in areas of influence of CASC outreach.	1. Seventy percent of growers who attend CASC field days/workshops indicate planned changes in tree architecture as a result of CASC outreach programming.

Notable results:

- A total of 3,671 direct contacts were made through field days and presentations while an improved Internet presence accounted for hundreds more. A quarter of mid-Atlantic producers who attended presentations are already adopting new technologies for precision agriculture. PNW growers who attended a CASC field day rated interest in adoption 3.8 to 4.4 on a 1 to 5 scale.

- Management efficiency trials in pilot orchards demonstrated increases in efficiency as high as 78%.
- Thirty percent of producers who attended field days are already adopting trellised planting systems and 65% plan to make this change.
- During the first two years, 42 students were involved in CASC projects and now have an interest in engineering solutions for specialty crops.

CASC outreach was comprehensive in the development of outreach programs, conducting applied field research, and training other extension educators and students.

Outreach Programming: The outreach programs were focused on preparing stakeholders for the future adoption of new technologies. In Year 2 we conducted eight field days and thirty presentations at industry and professional meetings, and published forty trade journal/ extension and thirty research articles. This resulted in 3,671 direct contacts through field days and presentations. Additionally, our website (http://www.CasCrop.com) was redesigned to include social media and YouTube presence. Fourteen new videos, a fact sheet, and legislative portfolio were all created and customized for the intended audience.

Applied Field Research: In response to stakeholder suggestions in the CASC socioeconomic survey summary, an increased focus is being placed on field trials to present evidence of increased management efficiency, improved fruit quality, and increased net income. Trials were conducted to compare harvest and pruning from a semi-autonomous platform versus from ladders, and efficiency increased by 44% and 78%, respectively. In trials with a vacuum harvest system, fruit quality improved by 5%. Savings were seen in targeted spray application and harvest-assist trials. Specifically, there was a 32-67% savings in herbicides ($14-29/acre) with new precision technologies and a $245/acre savings with vacuum assisted harvest.

Tracking Rates of Adoption: We believe that through structured input and feedback, we can improve our project relevance and impact. The CASC team has an advisory panel that includes growers; engineering, horticultural, and environmental scientists; commodity association leaders; consumer representatives; community and workforce development educators; and commercial machinery developers. An advisory panel meeting was held in February 2010. CASC advisors also are invited to attend regularly scheduled webinars (six per year). During these webinars, project leaders showcased results that are at a point where industry recommendations are needed on technology transfer.

Many of our field days and workshop presentations include a feedback mechanism. In the Pacific Northwest field days, CASC presentations rated 3.8-4.4 (1-5 scale), with highest interest in automated traps. At Mid-Atlantic field days and meetings, nearly half of respondents (47%) were very likely to adopt new technologies for increasing precision and 26% were already making changes (Figure 96). Sixty-five percent of respondents were very likely to plant high density systems and 30% were already making changes. A third of respondents (32%) were very likely to adopt new labor-saving technologies and 16% were already making changes.

Training: To build capacity in agricultural engineering and competitive orchard systems, we devoted time to training students and other extension faculty. Throughout the entire project, 42 students have been involved in CASC-related activities. Five faculty members from UC Extension, University of Maryland, and Rutgers were trained on labor saving technologies.

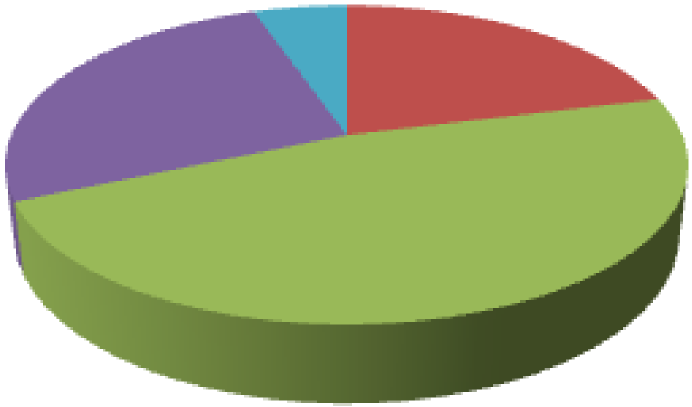

Figure 96. Producers in the Eastern US are embracing and preparing for novel technologies that improve efficiencies and reduce environmental impact.

Three CMU courses included projects on developing engineering solutions for specialty crops.

In the Human-Computer Interaction Institute's senior design project course, three students tackled the problem of redesigning the APM interface to make it more grower-friendly. The students undertook a formal design approach to the problem, starting with requirements collected in interviews with actual growers and ending with a functional prototype that was used by a farm owner and three farm employees to drive the APM in their commercial orchard.

In the Mechanical Engineering Department Prof. Bill Messner advised a group of seven students studying new methods to reduce fruit bruising and worker stress during harvest operations. The members of the Apple Picking Team took two approaches to the problem of bin filling. At first they considered modifying the bin, but changed their approach after the complexity of that task became apparent and after Prof. Messner advised them of the difficulty of getting such a product adopted. Their second approach was to modify the picking bag and combine that design with a removable padded insert for the bin. The insert created separate rectangular sections in the bin. The modified bag would reduce the jostling of the fruit as it was released into a section of the insert. Once an entire layer was filled, the insert would be lifted. The resulting prototype showed promise, but still had ergonomic issues. Providing a detachable mechanism on the side of the bin that allowed the workers to lower the picking bag over the middle of the bin might make the system practical.

In the Robotics Institute Prof. Martial Hebert used images of apples collected by CASC scientists to train undergraduate students in image processing-based automatic fruit detection and counting. The algorithms that yielded the best results may be incorporated in our work on crop load sizing and counting.