Development and Implementation of a Team of Robotic Tractors for Autonomous Peat Moss Harvesting

David A. Johnson
Commercial and Consumer Equipment Division
John Deere
Charlotte, North Carolina 28273

David J. Naffin
NavCom Technology
John Deere
Torrance, California 90503

Jeffrey S. Puhalla
Phoenix International
John Deere
Fargo, North Dakota 58102

Julian Sanchez*
MITRE Corporation
McLean, Virginia 22102

Carl K. Wellington
National Robotics Engineering Center
Carnegie Mellon University
Pittsburgh, Pennsylvania 15201

Received 28 July 2008; accepted 19 February 2009

This paper describes the key system components of a team of three tractors in a peat moss harvesting operation. The behavior and actions of the tractors were designed to mimic manual harvest operations while maintaining a safe operating environment. To accomplish this objective, each of the three tractors was equipped with a bolt-on automation package, and a human operator (team leader) was given a remote user interface to

*Previously with John Deere.

Published online in Wiley InterScience (www.interscience.wiley.com). • DOI: 10.1002/rob.20297

command and monitor the mission. The automation package included positioning, planning, and control, as well as coordination and perception systems to preserve field harvesting order, detect obstacles, and report physical changes in the operating environment. The system performed more than 100 test field harvesting missions during one season in a working peat bog, including three complete system tests with the end users.

1. INTRODUCTION

A team of robots performing tasks in an optimal, coordinated manner has been a goal of robotic research for many years. Only recently have robotic technologies matured enough to explore the possibility of developing multirobot systems that can reliably perform tasks in outdoor, real-world applications over an extended timeframe. Multirobot systems are complex examples of systems engineering, which require the integration and implementation of factors such as coordination, perception, path planning, obstacle avoidance, and a number of additional subsystems. This paper describes the design, implementation, and initial deployment of three autonomous tractors (Figure 1) for the purpose of harvesting peat moss.

Peat moss is accumulated, partially decayed plant material found in bogs and is the main ingredient in many potting mixes and professional growing media. The ability of peat moss to retain water and still allow oxygen to plants makes it very useful for horticulture applications. An active peat bog is divided into smaller, rectangular fields that are surrounded by drainage ditches on three sides to facilitate drying (Figure 2). When the top layer of peat is dry, the fields can be harvested. Weather permitting, the harvesting process is repeated daily during the summer months. Originally harvested by hand (with shovels), this labor-intensive process took a leap forward in the late 1960s when tractors pulling power take-off (PTO)–driven vacuum harvesters were brought into the peat bogs (Figure 1). A similar leap forward in efficiency is now possible with automated vehicles.

Figure 1. Autonomous tractor with sensor pod and peat moss vacuum harvester, showing a storage pile in the distant background.

2. KEY PROJECT DRIVERS

There were several key drivers for the development of this system. First, the peat moss industry is experiencing a labor shortage. Because of the remote locations, long work hours on harvest days, and unpredictable weather conditions, it is very difficult to ensure that enough qualified operators are available and that those same operators have enough consistent work throughout the season. Therefore, an autonomous solution would help alleviate some of the labor shortage strains that are currently experienced in the industry. Second, peat bogs offer a fairly well-structured environment in remote locations. When drainage ditches are added to a bog, they create a known fixed layout, which changes only slightly as new areas of the bog are added to the set of active fields. Furthermore, the bog is generally kept clear of obstacles and does not contain vegetation, which makes the perception challenge more tractable. From a technical point of view, all of these factors yield a desirable location and application to develop a fully automated off-road system.

A final motivation for this effort was that the existing manual harvesting operation paradigm lent itself to a smooth transition to automation. In the manual operation, a single very skilled team leader supervises a team of three or more tractors by communicating with the operators of each tractor via radio and/or hand signals. The team leader is the main and key decision maker in the field and is responsible for determining the ground speed of

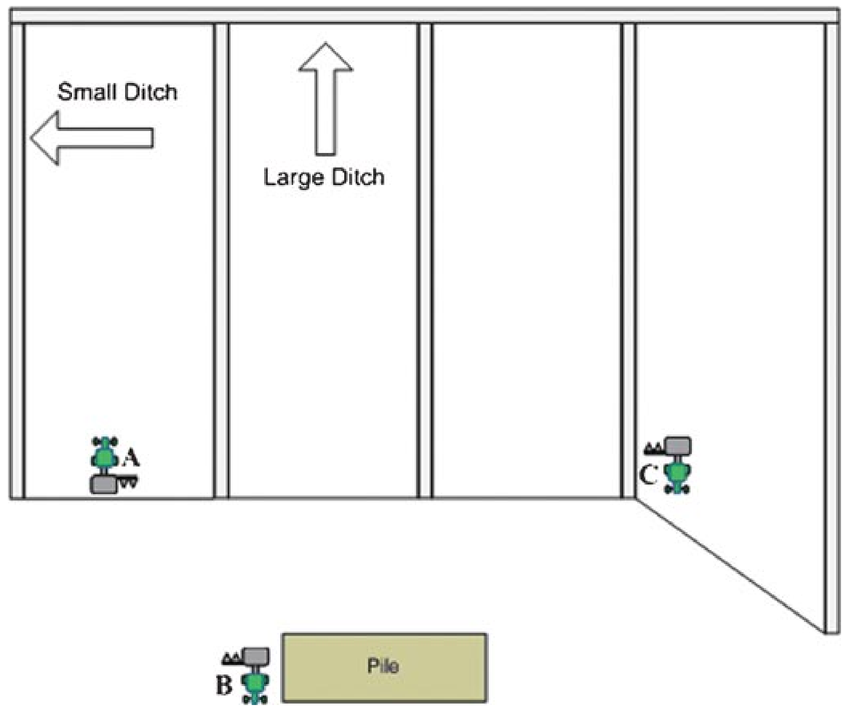

Figure 2. A typical bog layout with small ditches between fields, a pile of material outside a set of fields, and a larger ditch at the perimeter of the set of fields. In this scene a tractor was paused at the end of the first field while another tractor unloaded peat at a pile. Note that the figure is not to scale.

the tractors, the engine speed, and where the tractors should unload the peat, as well as numerous other settings on the harvesting equipment. By maintaining the same paradigm, the team leader remained in the loop and instead of communicating with human operators, he issued commands directly to the individual robots and monitored their performance via a remote user interface.

To achieve success in the transition to automation, it was necessary to ensure that the automated system (team of robots) harvested fields at a rate comparable to that of a human operator. Each tractor's movements needed to be coordinated to minimize idle time and ensure full coverage of each field. Furthermore, the individual tractors were required to perform most of the key tasks carried out by human operators, which included safeguarding themselves against obstacles and ditches, as well as detecting unload points at the edge of temporary storage piles. It is worth noting that this automated harvest system was only one component of the much broader peat moss harvesting operation. For example, the stockpiles of peat created by the harvesters (manually driven or robotic) were loaded onto wagons by a manually operated front-end loader and hauled to an on-site processing plant. Also, after each field was harvested, a manually driven tractor would harrow the area to prepare for the pass of harvesting, which could occur as soon as 4 h later. Again, the emphasis of this project was on harvesting the fields, not all of the other parts of the process.

3. HISTORY AND PRIOR ART

This peat moss harvesting system built on and benefited from progress in robotics and precision agriculture from academics and industry alike. Vision- and global positioning system (GPS)–based steering systems for tractors have been successfully demonstrated in experimental environments for a number of years. Early work included a vision algorithm that could segment cotton rows and determine heading and offset errors (Reid & Searcy, 1987), and similar visual guidance approaches have been used to track straight row crops at speeds up to 25 km/h (Billingsley & Schoenfisch, 1995; Gerrish, Fehr, Van Ee, & Welch, 1997). Instead of using local visual features, researchers have also developed automatic tractor steering systems using GPS (Bell, 1999; O'Conner, Bell, Elkaim, & Parkinson, 1996; Rekow, 2000), and additional work has been done on fusing vision-based crop row detection with GPS to try to provide more robust vehicle guidance (Zhang, Reid, & Noguchi, 1999). A summary of additional research in the automatic guidance of farm equipment is provided in Reid, Zhang, Noguchi, and Dickson (2000).

GPS-based guidance has matured to the point that a number of agricultural equipment and service providers (e.g., Autofarm, John Deere, and Trimble) have successfully delivered assisted-steering solutions for agricultural vehicles with "light bar" and "autosteer" products. Similar to steering control, implement control has recently seen high levels of automation-assisted operations in commercially available precision agriculture products (e.g., John Deere's iTec Pro). These operator assist products, combined with the availability of commercial tractors equipped with computer-controlled, infinitely variable transmissions, provide researchers with a rugged and reliable platform for pursuing full-automation solutions.

Researchers have demonstrated several fully automated agricultural systems that included obstacle detection in which a human operator was not needed. For example, an automated harvester used an adaptive vision-based classifier to track the cut/uncut line in an alfalfa field (Ollis & Stentz, 1997). This system was able to autonomously harvest hundreds of acres of crop in various fields and lighting conditions and included vision-based techniques for end-of-row detection and simple color-based obstacle detection (Pilarski et al., 2002). Another system used GPS-based guidance of a tractor in an orange grove combined with image-based obstacle detection (Stentz, Dima, Wellington, Herman, & Stager, 2002). Using a pretaught path, it drove autonomously for 7 km at speeds ranging from 5 to 8 km/h while pulling a sprayer. In addition to fully autonomous vehicle applications, there are also early examples of teams of robots, which include Balch and Arkin (1995), Chen and Luh (1994), and Parker (1994). More recently, Zlot and Stentz (2006) demonstrated a market-based framework for coordinating teams of robots while performing an outdoor mapping task. Vail and Veloso (2003) demonstrated their own market-based approach that included potential functions in the Robo-Soccer domain. Bochtis, Vougioukas, Tsatsarelis, and Ampatzidis (2007) provided a hierarchical decomposition approach to the problem of complex agricultural tasks and a method of task allocation that results in an optimal resource utilization. Finally, Gerkey and Matarić (2004) provide a formal analysis and taxonomy for many competing multirobot task allocation techniques.

The current work combines a broad range of capabilities found in fully autonomous vehicle platforms and multiple-robot teams in a team of three autonomous tractors. These capabilities include mature vehicle guidance and control, multiple-vehicle coordination, run-time path planning, and perception for obstacle detection and pile estimation. Furthermore, considerable effort was also allocated to developing a successful end-user experience by understanding and clearly defining the role and tasks of the human operator (or team leader).

4. SYSTEM OVERVIEW

During manual operation of an individual harvester, an operator completes one lap in a field, vacuuming a very thin layer of peat off the surface. When the lap is completed, the operator drives to a nearby pile, unloads the material, and drives to the next field to repeat the process. This sequence of events was preserved in the automated system and provided the process framework for this project. The overall system architecture is shown in Figure 3. This section introduces the major components of the system.

4.1. User Interface

The team leader, in manual operation, was responsible for assigning tasks to teams of harvesters, including the range of fields that needed harvesting as well as the order in which those fields had to be harvested. In the automated system, the team leader utilized the team leader user interface (TLUI) to create a scenario, or work plan, for the day. In addition, the TLUI displayed the individual status (telematics) of each robotic harvester, as well as the overall team progress as the scenario was executed. The user interface and its development are discussed in more detail in Section 5.

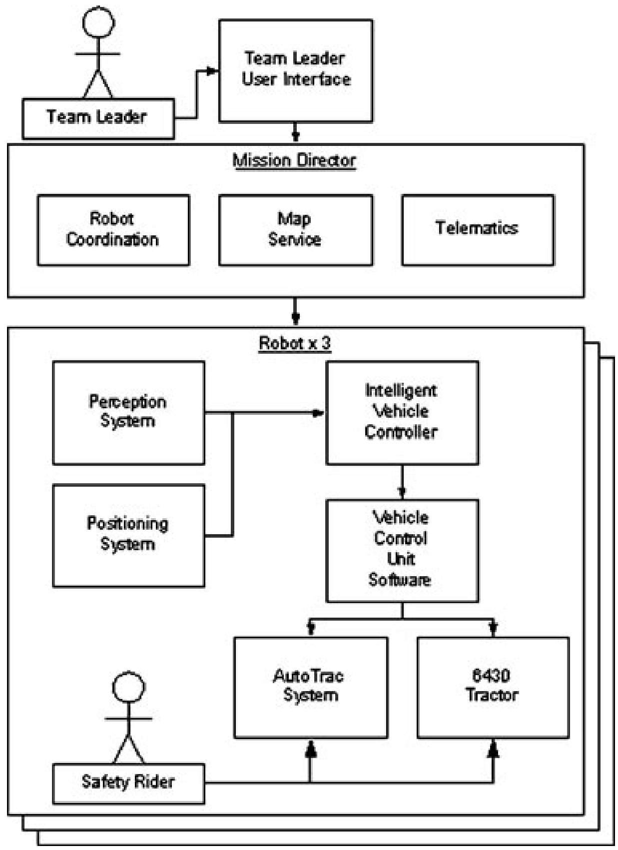

Figure 3. System block diagram.

4.2. Mission Director

The scenario generated by the team leader through the TLUI was first communicated to the mission director module. It was the mission director's responsibility to allocate individual tasks to each robot and monitor these tasks to ensure their completion. During the execution of each scenario, the mission director consulted a list of "rules" to ensure that coordination constraints were being met. As tasks were completed, the mission director communicated this progress to the TLUI.

In addition to its other responsibilities, the mission director provided a centralized repository for a set of maps. The world map was a map of static features, including field and ditch boundaries, created at the start of the season with data collected from a GPS survey of the bog. A field coverage map (or the field map) contained waypoints for use during field harvest planning. The field map's waypoints were calculated and verified a priori based on these field boundaries. Nonstatic objects such as piles were stored in the pile map and could be updated by the team leader through the TLUI or by any of the robots in the system. The pile map included information on the size and location of each pile of material used during the unload phase of a mission. Section 6 provides further details on the design of the mission director.

4.3. Robots

Each robot in this application included perception and positioning systems, an intelligent vehicle controller (IVC), vehicle control unit software (VCU), an AutoTrac^{TM} GPS-based steering system, and a John Deere 6430 tractor. The perception system fused data from a variety of sensors to provide the IVC with vehicle-centric maps of ditches, piles, and obstacles. Positioning was accomplished by merging real-time kinematic (RTK) GPS, vehicle odometry, and a model of the vehicle's dynamics. The IVC provided a set of robot behaviors that integrated the information from the mission director, the perception maps, and the vehicle pose to perform the task assigned to it by the mission director. The set of behaviors included path planning, perception driving, and safe-speed obstacle avoidance.

4.3.1. Path Planning

All path planning was performed at run time. A point-to-point planner was used in combination with the maps from the mission director to find safe and efficient paths to each objective. Paths included information about where to drive, the speed of the vehicle along a path, and any implement actions that needed to be performed. Section 7 describes path planning in more detail.

4.3.2. Perception-Based Driving

During unloading of the harvested peat, it was necessary to perceive the changing location of the continually growing pile and then control the robotic vehicle to align with the edge of the pile in order to dump in the proper location. The perception system is described in Section 9.

4.3.3. Safe Speed and Obstacle Avoidance

All paths that the vehicle drove were continuously checked for obstacles using the perception maps. The obstacle avoidance behavior for this application reduced the vehicle speed to avoid hitting obstacles (path replanning was not used). Additional information on the safe-speed obstacle avoidance behavior is described in Section 8.

The IVC determined the low-level actions to be executed by the robot and communicated those to the VCU. The VCU translated these low-level commands to signals that were interpreted by the AutoTrac path tracking system and the tractor for steering, as well as transmission and implement commands.

During this test phase of the project, each robot had a safety rider onboard at all times. The safety rider had the ability to pause and resume that specific robot as well as take manual control of all tractor functions at any time.

5. HUMAN FACTOR CONSIDERATIONS

The form and design characteristics for the user interface of an autonomous or semiautonomous vehicle largely depend on its level of automation. At its most basic state, the semiautonomous vehicle is fully controlled by a human operator who is in a remote location but has visual contact with it. With such a system, information about the vehicle's intent is not necessary. However, in systems with high levels of automation, it is critical to provide all information that will help ensure that the human operator (supervisor) maintains a high level of situational awareness (Endsley, 1995). This is particularly important in systems in which there is more than one vehicle.

Over the past 15 years, the body of research that has investigated potential human–computer interaction issues with highly unmanned vehicle systems has grown steadily. Although a large portion of this work has been focused on unmanned air vehicles, several researchers have begun using unmanned ground vehicles as their main platform. Furthermore, human factors researchers have begun to examine the challenges associated with one human operator controlling and/or supervising a team of vehicles (Adams, 2007; Humphrey, Gordon, & Adams, 2006; Humphrey, Henk, Sewell, Williams, & Adams, 2007; Vig & Adams, 2006). Regardless of the application, the aim of these efforts has been directed at providing human operators with an intuitive interface to monitor and/or control a team of vehicles. Much attention has been allocated to developing visualization solutions to represent the vehicles' location in the environment and vehicle intention, as well as the presentation of diagnostic information. In the current effort, those challenges were met successfully, as the human team leader was able to easily monitor vehicle position information and information about the state of each individual robot.

5.1. Developing the TLUI

As automation is introduced into any system, it is important to consider the performance impact on any humans who remain in the loop. Specifically, it is critical to ensure that the operation of the system remains safe and the human interaction with the automation is driven by intuitive and error-free design. To accomplish this objective, a detailed task analysis of the manual peat moss harvesting operation was initially conducted. A task analysis is a method to identify the individual elements that make up a task, their relationship to one another, and the logical or time-sequence structure of these elements (Luczak, 1997).

The task analysis results yielded a list of items that highlighted the sensing, information processing, decision making, and/or control action requirements for the automation; in other words, the tasks that were performed by humans in the manual system but now would have to be performed by the automation. Furthermore, the task analysis results provided a list of information items that the team leader required for the successful supervision of the autonomous vehicles and the natural sequence of tasks in his typical interaction with the manual system.

The data from the task analysis were a key driver of the TLUI design. A screenshot of the main TLUI page is shown in Figure 4. The TLUI was housed in a 10″-in. (25.4 cm), rugged, tablet personal computer (PC) with a stylus pen as the primary input device. The left-hand side of the screen contains diagnostic information for each vehicle (ground speed, rpm, and vacuum height). This diagnostic information was deemed to be the most important and time critical for the team leader. The main portion of the page illustrates a map of the peat fields and, within that map, the real-time location of the vehicles. The fields were also color coded to represent different field conditions. For example, red fields are fields that the team leader has designated as "keep-out" zones. To navigate within this portion of the screen, a set of

Figure 4. Screenshot of main TLUI.

tools shown on the top-left-hand corner of this section were available. At the top of the page, a set of navigation tabs were available to access other screens and functionality.

5.2. Emergent Display Features of the TLUI

One of the challenges in designing a user interface to supervise and control multiple vehicles was incorporating all of the necessary information about the state of the vehicles and the operational environment. Ensuring that the team leader had easy access to all of this information was one of the keys to guaranteeing high situational awareness. However, to achieve this goal, information needed to be included in a limited amount of screen real estate and its layout and visualization had to minimize the team leader's workload and cognitive demands.

A design feature of the TLUI that enhanced its overall usability was the use of emergent display information. An emergent display is one in which individual data elements are grouped such that their emerging features convey information inherent to their relationship. A unique characteristic of emergent displays is the reduction of information processing time, because the human does not have to sense and interpret each individual element but rather can use the overall emerging patterns to draw information about the system (Treisman, 1986). In the TLUI, the concept of emergent displays was used to represent the diagnostic information for each vehicle, as well as the real-time location of the vehicles in the map. For the diagnostic information, the speed, rpm, and vacuum height indicators for each tractor were connected by a line. As the value of any of these indicators changed, the connecting line would adjust accordingly. Therefore, the user was able to quickly assess the performance of each vehicle by simply looking at the emerging pattern of the lines, rather than gazing at each parameter of the vehicle. For example, Figure 5 illustrates a pattern representing the desired settings of speed, rpm, and vacuum height for the morning (when material is moist) and the afternoon. After time, users learn to recognize the

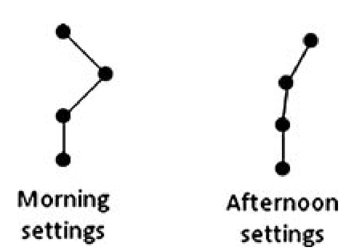

Figure 5. Example pattern representing machine parameters in the morning and afternoon.

specific patterns and can make judgments about the performance of the vehicles more efficiently than if attending to each piece independently. Similarly, the user was able to rapidly establish the real-time location of the vehicles in the map, by looking at the emergent triangular shape created by the vehicle indicators.

5.3. TLUI Navigation

A common outcome when introducing automation into an existing system is that the nature of the tasks required from the human changes, mainly because the cognitive and physical demands on the human are impacted. Because this transition in task demands can sometimes lead to confusion, as the user is not clear on his/her new role, it is important to develop a system that is consistent with the mental model of the current operation and system to ensure a smooth transition. Figure 6 illustrates the natural sequence of tasks that the team leader was accustomed to performing with the manual system configuration. The navigation scheme of the TLUI for setup and startup of the autonomous system was designed to mimic this sequence.

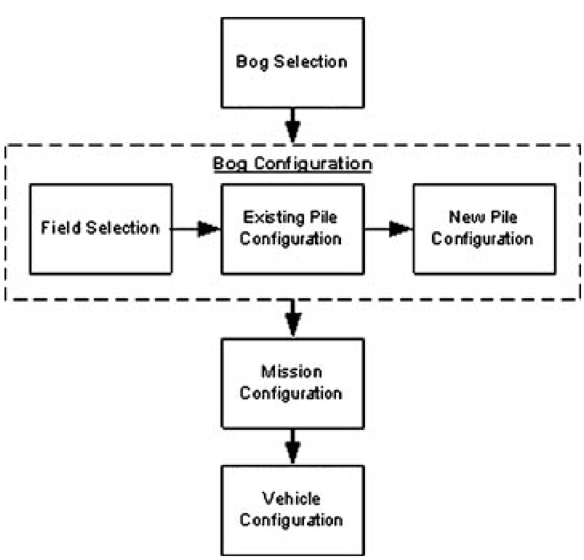

Figure 6. Natural task sequence for the team leader with manual system.

5.4. Safety Features

During the initial deployment of this system, one requirement was to have safety riders in the autonomous tractors/harvesters. Safety riders were responsible for "taking over" control from the automation in situations that they deemed unsafe. Therefore, a visual display was installed in the tractor cabs to convey the intended path of the tractor in autonomous mode and provide control functionality that allowed the operator to stop the vehicle and restart movement. In addition to the many safety considerations and design features, one of the fundamental safety elements of the autonomous system was the allocation of control for the stop/start of the vehicles during operation. The underlying design philosophy was that only the agent who commanded a stop had the authority to restart the vehicle. However, both agents had authority to stop the vehicle at any point. Furthermore, when any vehicle was commanded to restart movement by the team leader, the safety rider was notified and given 10 s to override the command. Finally, information about the current intentions of the vehicles was displayed in both the TLUI and the in-cab displays. This information helped set the expectations of the team leader and safety riders about the behavior of the vehicles.

6. MULTIVEHICLE CONTROL

One of the initial decisions that was made during the early design stages was to allocate a great deal of onboard autonomy to the tractors. Each tractor had a high-level controller that was capable of performing its own behaviors such as path planning, mission execution, and obstacle avoidance. However, the coordination of these tractors was handled by a centralized component known as the mission director.

6.1. Mission Director

The mission director's primary purpose was to provide the system with a mechanism for multirobot coordination by acting as a gateway between all of the pertinent vehicle telemetry and the TLUI. Additionally, it provided a central repository for shared data and configuration information, as well as a mission generator (task allocator). It was the responsibility of the mission director to allocate tasks to the individual robots as well as maintain overall coordination between robots. To accomplish this goal, the mission director maintained a database that represented the overall state of the system (i.e., all pertinent individual robot states, as well as the overall team's harvesting progress). To avoid collisions, deadlocks, and other basic coordination problems, the mission director implemented a set of rules that queried from this state database and placed constraints on the individual robots.

6.2. Data Repository

The world, field, and pile maps were all housed and maintained in the mission director. At startup, the mission director sent each IVC the current version of these maps. In addition, any updates to these maps were replicated to all IVCs.

6.3. Scenarios

The TLUI communicated the team leader's intentions via a scenario data structure. A scenario specified the set of fields and the order that the fields should be harvested. It also specified the set of vehicles selected to perform the harvesting and a set of piles to unload the peat. The scenarios referenced fields and piles simply by named keys. Later these keys were used to look up the exact coordinates from the field or pile maps.

6.4. Mission Generation

Once a scenario was defined by the TLUI and was sent to the mission director for execution, the mission director generated a set of "missions" or tasks for all IVCs referenced in the scenario. These missions were a sequence of parameterized commands for performing one of four simple tasks: drive to a given field, harvest a given field, drive to a given pile, or dump at a given pile.

6.4.1. Drive to a Field

Each IVC instantiated a point-to-point planner that was used to generate a path from the robot's current position to the entrance point of a given field. The entrance points for all fields were known a priori and stored in the field map. The point-to-point planner used the world map to locate any obstacles between the robot's current position and its goal. Once a path was generated, the IVC instantiated a path tracking behavior that utilized the robot's StarFire GPS receiver to track the vehicle to its goal position. During this period, the vehicle vacuum was disengaged.

6.4.2. Harvest a Field

Once in a field, the IVC utilized a list of stored waypoints in the field map to plan a lap of harvesting. Once the path had been generated, the IVC instantiated its path tracking behavior to traverse the path. When the vehicle reached the top of a field, it raised its vacuum brooms before making a U-turn and then lowered them after it completed this maneuver.

6.4.3. Drive to a Pile

After making a full lap of a field, a vehicle planned a route to a pile to unload its newly harvested peat. Again the IVC instantiated its point-to-point planner to generate a path from the end of a field to a point near the expected dump point of a given pile (stored in the pile map) as well as a path tracking behavior to move it to its goal. During this period, the vehicle vacuum was disengaged.

6.4.4. Dump at a Pile

Once near the pile, the vehicle utilized a perception-based behavior to locate the true end of the pile and generated a path to position the vehicle's vacuum at the proper orientation. Perception information was also used to update the pile's dump point in the shared pile map.

6.5. Rules

Before the mission director commanded an IVC to perform a task, it checked the task against a list of rules, which can be classified into two types, passive and active. Passive rules were checked only once prior to a new task being assigned to an IVC, whereas active rules were continuously checked. The identifier of the IVC and the task it had been assigned were both passed on with each rule. The rule queried the mission director's state database for any other details needed to determine whether that task could be performed safely. If the task passed all the rule checks, it was sent to the IVC for execution. If even one rule failed, the task was placed in a queue and rechecked at a later time. Meanwhile the IVC waited idle until its task had passed all of the mission directors' rules. If an active rule failed, the mission director sent a

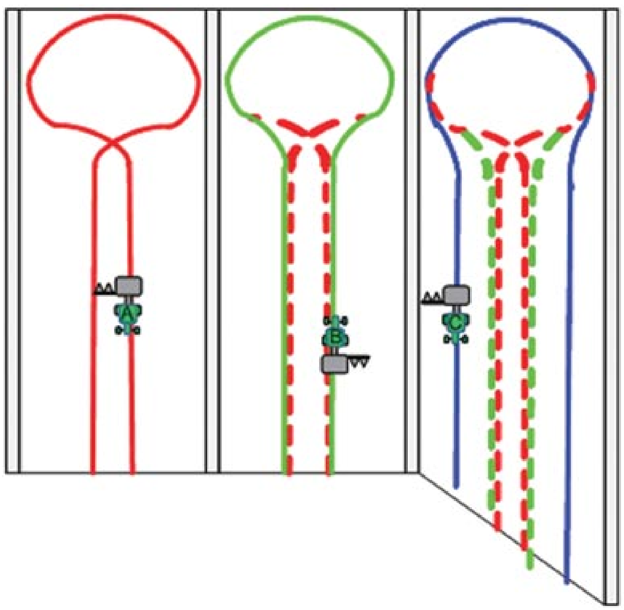

Figure 7. Harvesting order.

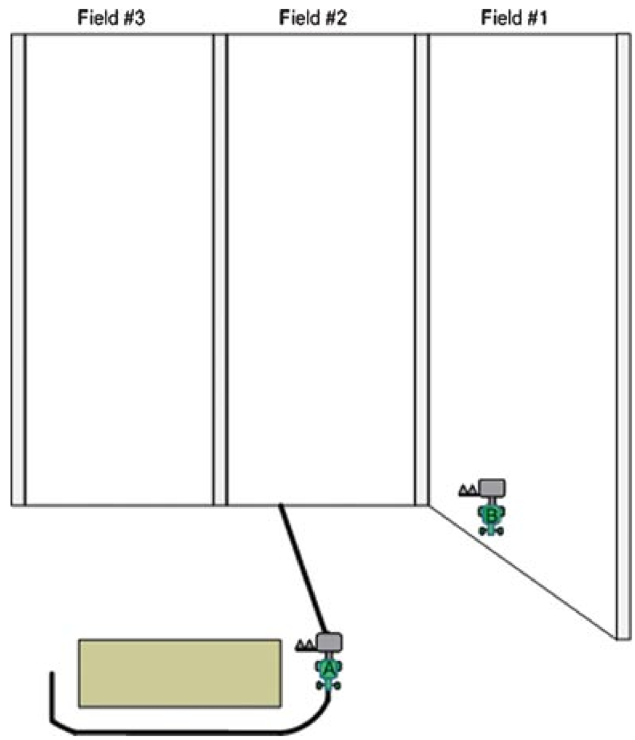

Figure 8. Example of "one robot out of a field" rule.

pause signal to the "offending" vehicle(s) until the condition that caused the rule to fail no longer existed. The mission director implemented the following set of rules.

6.5.1. Harvesting Order

Three full laps were required to completely harvest each field (Figure 7). Each lap required a different vacuum broom height. The current implementation required that each vehicle use a fixed vacuum broom height. Therefore, three different vehicles types (A, B, and C) were associated with each vehicle, depending on the height of the vehicle's brooms. Complete harvesting of a given field required not only that all three types of vehicles complete a lap but that the order of the laps (A followed by B followed by C) also is enforced. The "harvesting order" rule guaranteed that the vehicle ordering was preserved. It also ensured that only one vehicle was in a given field at a time. This rule was checked prior to an IVC performing a "drive to field" task.

6.5.2. Only One Robot out of a Field

To avoid collisions and deadlock situations, only a single robot from the team was allowed to leave a field at any given time. In Figure 8, tractor B is paused waiting in field 1 while tractor A dumps at the pile. This rule was checked prior to an IVC performing a "drive to pile" task.

6.5.3. Next Field Ready

To avoid a possible deadlock situation around the piles, the "next field ready" rule required that before a vehicle left a field, the next field it planned to enter was not occupied by another vehicle. This rule prevented vehicles from sitting idle at the pile's unload points while waiting to enter their next field. Instead, the vehicles waited at the exit point of their last harvested field. In Figure 9, tractor B waits for tractor A to finish field 2 and dump at the pile before leaving field 1.

6.5.4. No Robot Left Behind

To facilitate the team leader's monitoring task, it was important that all vehicles in a team harvest the same general region of the bog. Therefore it was important

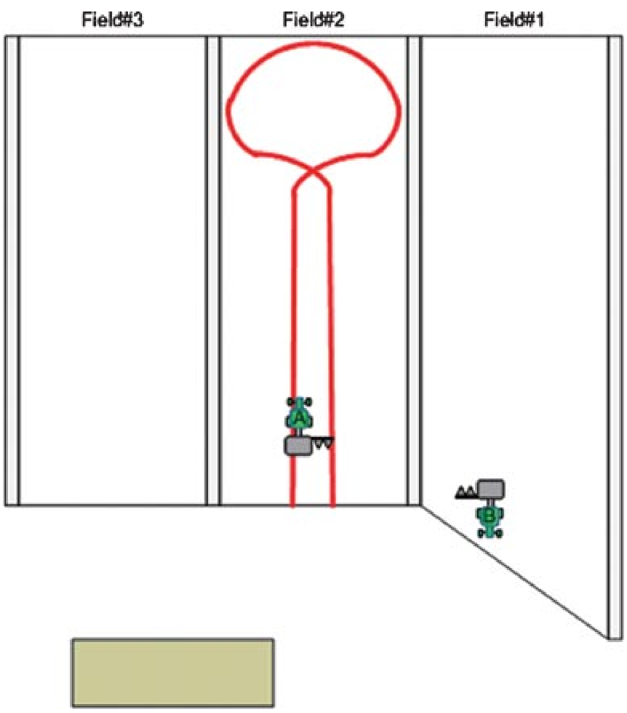

Figure 9. Example of "next field ready" rule.

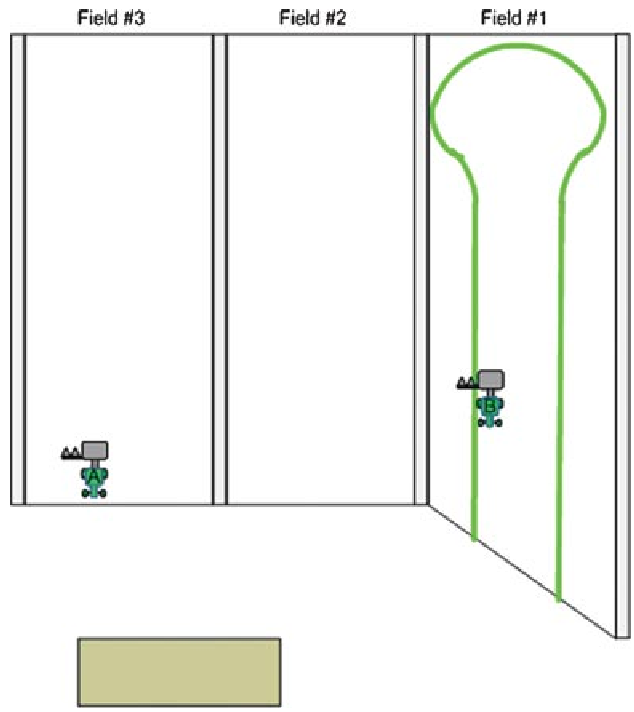

Figure 10. Example of "no robot left behind" rule.

to prevent any one vehicle from getting too far ahead of the other two. This rule required that an A-type vehicle not leave its current field unless a B-type vehicle was within two fields behind it (see Figure 10). In addition a B-type vehicle could not leave its field unless a C-type vehicle was within two fields behind it.

6.5.5. Minimum Safe Distance

This is the only active rule. This rule required that any two vehicles maintain a minimum safe distance between them. If this rule was broken, the lower priority vehicle (e.g., an A-type vehicle was higher priority than a B, which was higher than a C-type) was paused. If the distance between two vehicles dropped below a critical distance, then both vehicles were paused and the safety drivers and team leader were alerted to the problem.

7. PATH PLANNING

Each autonomous tractor performed real-time, point-to-point path planning for each of the four tasks: drive to field, harvest a field, drive to a pile, and dump at a pile. Two methods of point-to-point path planning were utilized to provide planning for these four tasks. The two methods were a visibility graph planner and a grid-based planner.

7.1. Visibility Graph Planner

The visibility graph planner was based on a visibility graph, which was generated from polygonal features of the world map and the pile map (Figure 11). Before visibility graph construction, the planner expanded the keep-out areas of the map by at least half the width of the vehicle. In addition, terrain boundaries were shrunk on all sides by at least half the width of the vehicle. Because the vehicle had a minimum turning radius, the corners of the boundary and keep-out areas were rounded to ensure that the minimum turn radius requirement was not violated. For narrow objects, the sides were expanded so that a minimum turn radius could be employed between adjacent corners. All of these steps facilitated the inclusion of drivable path segments around obstacles into the final solution. Figure 11 shows rounded shapes representing the expanded obstacles in the map.

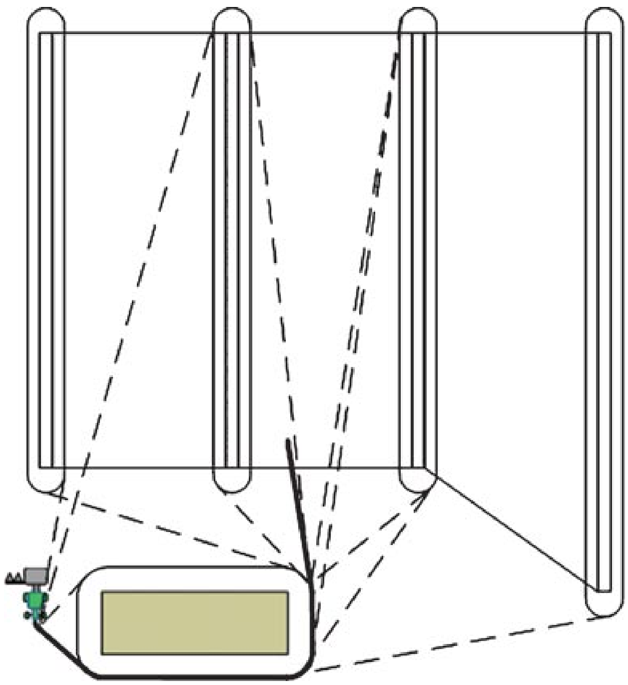

Figure 11. Visibility graph from peat bog area for point-to-point planning. Dashed lines represent the edges of the graph. The chosen path from the dump location to the field is shown by a boldface line.

After the expansion/contraction of the map objects, the visibility graph was constructed. The graph consisted of nodes and edges, where the edges represent drivable paths through free space and the nodes represent points where paths merge or split. For planning in free space (off-road, no roads in the map), nodes represent the starting point, ending point, and tangent connections to the obstacles from the starting and ending points. Tangent connections between obstacle pairs would also be considered node points. Figure 11 shows the edges of the graph (dashed lines) connecting node points.

The edges generated from the visibility graph are searched using the A* algorithm to find the optimal path from the start location to the stop location. The computed path either connects the start location with the goal location directly or connects the start location to the nodes surrounding each of the obstacles until arriving at the goal location. Figure 11 shows the edges connecting the vertices of each polygon. The resultant optimal path calculated by the A* algorithm is shown in boldface.

There are strengths and weaknesses with the visibility graph approach. The strengths include its ability to maintain start- and end-point location and heading accuracy. Also, the minimum turn radius of the vehicle is considered at all times. The weaknesses of the approach include its algorithm complexity when faced with multiple and/or complex obstacles, which results in long planning times. Also, in planning paths to/from each field, the start/end points of these paths are often required to be near an obstacle that is not much farther than the half-width of the vehicle. This presents a problem with the visibility graph approach because the first step of its algorithm is to grow the keep-out areas according to its vehicle dimensions as well as its turning radius. The minimum turning radius of this vehicle/implement pair was approximately 6 m; therefore, obstacles that were long and narrow, such as ditches and sometimes piles, tended to grow by an additional amount because of the turning radius of the vehicle. This led the visibility planner to fail when planning a path in which the start/end point was along the side of a pile or a ditch. Another shortcoming with the visibility graph approach is that it grows the obstacles by an amount based on a static vehicle model. The vehicle/implement pair in this application is asymmetric about the center of the vehicle (the left-hand vacuum nozzles stuck out much farther than the right-hand side of the implement). One could give the planner the half-width corresponding to the left-hand side, giving the right-hand side ample room as well, but then paths arriving at obstacles on the vehicle's right-hand side could not get sufficiently close to the obstacle. Likewise, if the obstacles are grown according to the right-hand-side dimensions, paths would be planned where the left-hand side of the implement would get hit by the obstacle.

7.2. Grid-Based Planner

The grid-based planner subdivided the entire local region where the robotic vehicle was positioned into a two-dimensional (2D) grid of cells. The size chosen for each cell in this application was 1×1 m. A graph was constructed in a fashion similar to the visibility graph approach. In this case, the nodes of the graph represent the center point of each of the grid cells. The edges of the graph represent straight-line trajectories to each nearest neighbor of a grid cell. In this application, eight nearest neighbors to each cell were chosen so that the edges represented horizontal, vertical, and diagonal connections to the next grid cell neighbor. The cost of travel to a neighboring cell was calculated by adding the distance required to travel to this cell and the cell costs of all the cells that the vehicle model would occupy during travel. Free-space cells were given a cost of 0, and cells with obstacles were given a max cost value. The A* algorithm was used to search this graph for the optimal path from the start location to the final location. The result of the A* algorithm would provide an optimal path that would not allow the vehicle model to intersect any obstacles. This path was then smoothed to get drivable path segments.

Two of the key strengths of this planning approach were its efficiency in the presence of multiple obstacles and its ability to plan paths very close to its left- or right-hand side. The latter is possible by calculating the cost to travel to an adjacent cell by taking into account the vehicle's shape and dimensions across all grid cells as it makes the transition to the next cell. However, there are three shortcomings with this approach. One is that the final goal position accuracy is limited by the resolution of the grid. Another is that to guarantee end-point heading accuracy, virtual obstacles are placed in the map to force the planner to calculate the start/stop heading appropriately. This adds clutter in the map that should not be needed. Finally, for minimum turning radii that significantly exceed the half-width of the vehicle, the path plans become undrivable. This final problem did not affect the peat moss application significantly because of the large size of the implement.

7.3. Planning Tasks

By taking into account the strengths and weaknesses of both path planners, an acceptable solution was developed for the path planning of this application. For "planning in a field," shown as path 1 in Figure 12, end-point accuracy was needed. Also, a map with no obstacles was used, because safe field waypoints required to produce the desired pattern were preplanned and present in the field map. Therefore, the visibility graph approach was used for this task. Because obstacles were not present in the planner's map, fast execution times were possible with this approach.

For the task "planning to a pile," shown as path 2 in Figure 12, multiple obstacles (6+) needed to

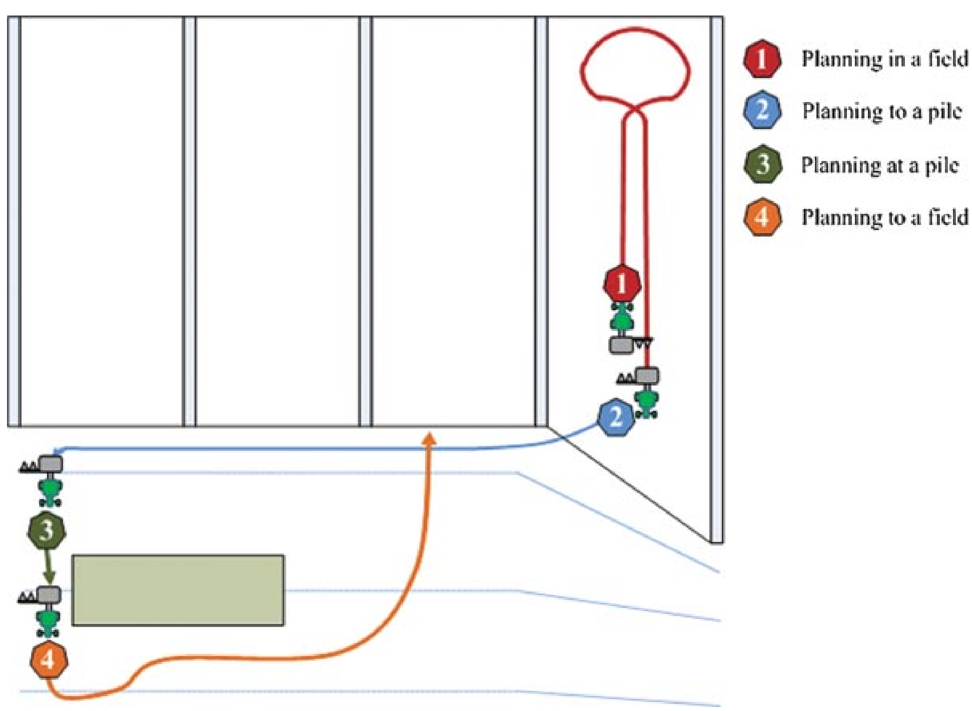

Figure 12. Four path planning operations during typical field operations.

be considered. Also, the planner needed to plan to positions that included an obstacle near the right- or left-hand side of the vehicle. Therefore, the grid-based planner was used for this task to take advantage of its fast execution with multiple obstacles and its ability to plan near keep-out areas.

For the task "planning at a pile," shown as path 3 in Figure 12, only the pile obstacle that needed consideration in the planner's map and end-pose accuracy was critical. Therefore, the visibility-based planner was used for its superior end-point accuracy. With only one obstacle in the planner's map, the execution speed was sufficient.

For the task "planning to a field," shown as path 4 in Figure 12, multiple obstacles needed to be considered. Again, the planner needed to plan to positions that included an obstacle near the right- or left-hand side of the vehicle. Therefore, the grid-based planner was used for its fast execution with multiple obstacles and ability to plan near keep-out areas.

8. SAFE-SPEED CONTROL

During autonomous operation, the system monitored the commanded turning angle of the vehicle. Using the perception map information along with the vehicle model, it calculated the distance the vehicle needed to travel before hitting an obstacle (or reaching the edge of the map). The calculated distance to the obstacle was reduced by the length of the nose of the vehicle, a speed-dependent distance that accounted for system lag and a user-defined safety buffer. This new calculation provided the distance that the vehicle had to come to a complete stop. Based on a conservative deceleration factor, the safe speed to travel to ensure that the vehicle could decelerate to a stop was calculated. If the current speed of travel exceeded this calculated safe speed, then the vehicle speed was reduced until it was equal to or less than the computed safe speed. This procedure allowed the tractor to gracefully slow to a stop when approaching an obstacle in its path, and it caused only a brief slowdown if it was a false obstacle.

The various objects detected by the perception system were reported in separate maps, and each was treated differently by the speed control algorithm. Obstacles such as people or other vehicles were a hazard for both the tractor and the implement, and an extra safety buffer was required so that the tractor and implement did not get close to these obstacles. Piles were also a hazard for both the tractor and the implement, but during a dump the tractor and implement travel extremely close to the pile so there was no extra safety buffer for the pile. Ditches were a hazard for the tractor, but the vacuum heads of the implement had to travel over the ditch during turns, so the side section of the implement was allowed to intersect the ditch map.

9. PERCEPTION

Although GPS was used for mission planning and vehicle control, a fully autonomous system required onboard perception to detect unmapped obstacles and also to estimate the changing location of the pile used for determining where to dump a load of harvested peat moss. As described above, three different maps were generated by the perception system:

- Obstacle map: people, tractors, loader, ATV
- Pile map: estimation of pile location for dumping
- Ditch map: drainage ditches between fields

These output maps were produced using a National Robotics Engineering Center (NREC)–designed sensor pod and set of perception algorithms. Figure 13 shows the sensor pod and perception modules, described below, that were used to produce the required output maps.

9.1. Sensor Pod

As shown in Figure 13, the NREC sensor pod used for perception included a number of different sensing modalities. The primary sensor was a SICK laser range finder (LADAR) that produced a 180-deg horizontal scan of range measurements of the area in front of the tractor. This LADAR was mounted on an actively controlled nodding mechanism to produce three-dimensional (3D) range data by rotating the LADAR from 5 to 45 deg below horizontal (one full sweep up and down every 2 s). The sensor pod also contained a thermal infrared camera and a pair of color cameras with a neutral density filter on one of the color cameras to produce a higher effective dynamic range.

The sensor pod also included an inertial measurement unit that was combined with the tractor speed sensor to produce a smooth local pose estimate. The tagged range points from the nodding LADAR were

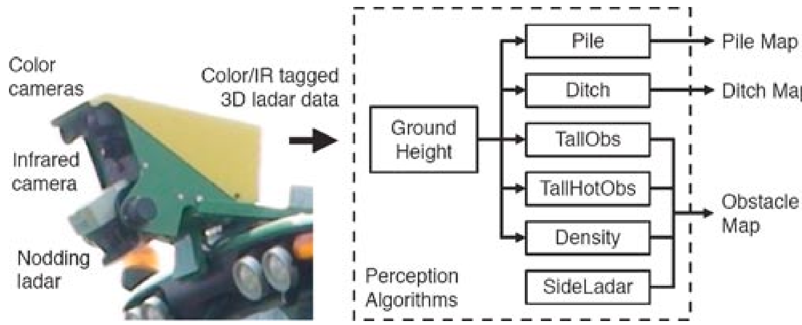

Figure 13. Sensor pod mounted on tractor cab and perception system modules that process the sensor pod data to produce pile, ditch, and obstacle maps. (The fixed side LADAR sensor that feeds the SideLadar module is not shown here but is visible above the right rear tire in Figure 1.)

tightly synchronized with this local vehicle pose estimate, allowing the system to accumulate data over time into a 3D point cloud. The nodding LADAR was also calibrated with respect to the cameras to allow the 3D range points to be projected into the camera images and tagged with color and infrared information, as shown in Figures 14(a) and 14(b). Note that color information was not used in the final fielded system described below.

9.2. Ground Height Estimation

An accurate ground estimate was required for finding piles and ditches and differentiating between obstacles on top of the ground from rises in the ground (such as the pile). However, ground height estimation was difficult for several reasons. Range points did not align perfectly due to sensor noise and drift in the position estimate. There were gaps in the data from the limited density of measurements and because of occlusion. Finally, the observed lowest range measurements were only an upper bound on the true ground height because the range sensor was not always able to measure the ground surface directly.

To handle these challenges, the system used a Markov random field (MRF) probabilistic model for smooth ground based on previous work (Wellington, Courville, & Stentz, 2006). This approach uses a 2D grid to represent the ground surface as shown in Figure 14(c). The MRF models the noise in the sensor data and also includes spatial correlations between neighboring patches of ground to enforce the assumption that the ground surface is generally smooth. The approach can be visualized as trying to conform a rubber sheet to the lowest range data in each ground patch.

Although the MRF smoothing naturally handles noisy sensor data, data expiration is used to further limit the effect of data misalignment. In addition to the range points, the ray from the LADAR sensor

(a) View from tractor (b) Colorized 3D point cloud data (c) Perception output

Figure 14. Example autonomous run where the system had acquired the pile for dumping but then had to stop because the Gator ATV was in its desired path. The perception output shows a visualization of the computed ground height, the pile map, and the obstacle map.

to the range point is included as an inferred upper bound on the ground height to help fill in gaps in the range measurements. However, these inferred measurements are often significantly above the true ground surface in occluded areas behind obstacles so the MRF smoothes areas with inferred measurements much more than areas with actual range measurements.

Figure 14(c) shows an example of the ground height estimation output. The ground surface is smooth, which results in accurate pile detection, and the algorithm found the correct ground height under the obstacle to allow it to be differentiated from the pile. A simple approach that uses the lowest range data in each ground patch or uses a spatial median filter on the lowest range data results in a rise in the ground estimate near the all-terrain vehicle (ATV) because the sensor does not directly sense the ground there. This false rise in the ground height causes a false pile to be detected and therefore could result in incorrect autonomous behavior.

9.3. Obstacle Detection

As shown in Figure 13, a number of different modules were combined to find obstacles. Because of the complementary nature of the different algorithms, an obstacle was reported if any of the following detectors was above a threshold.

9.3.1. Tall Obstacle (TallObs)

This detector looked for regions of hits above the ground. This was a simple detector that found obstacles larger than 1 m tall, even if they were moving. Because it was looking for hits above the ground, it was important that the ground height estimate was correct underneath obstacles.

9.3.2. Tall Hot Obstacle (TallHotObs)

Using the infrared-tagged range points, this detector looked for regions above the ground that had a high temperature. The ground surface often gets very hot, but any object above the ground that has a high temperature is likely to be an obstacle. This detector was useful for differentiating between warm objects such as people or vehicles and objects that stay cool such as weeds or other vegetation.

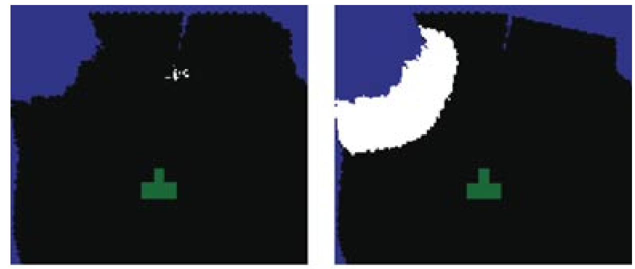

(a) Obstacle (b) Pile

Figure 15. Obstacle and pile and obstacle maps for example from Figure 14.

9.3.3. Density

A 3D grid of voxels (small volumes of space) was maintained around the vehicle, and then each LADAR measurement was ray traced through this grid to maintain the number of hits and passes for each voxel of space. The ratio hits/(hits + passes) was then used as the density score for that voxel. Density measurements have been used in the past to differentiate between vegetation and solid obstacles (Lacaze, Murphy, & DelGiorno, 2002; Wellington et al., 2006), but in this application they were used to detect stationary obstacles shorter than the 1-m threshold used by TallObs without suffering from false positives due to poor data registration or dust near the ground. In either of the latter cases the data near the ground tend to be a mixture of hits and passes and result in a low density score that can be discriminated from a short solid obstacle. However, a moving obstacle also will result in a mixture of hits and passes, so the TallObs and TallHotObs detectors are required to detect moving obstacles. This density-based approach was able to filter out many false LADAR returns caused by light dust, but heavy dust appeared solid to the LADAR and often resulted in a false positive.

The obstacle maps in Figures 14(c) and 15(a) show that the combination of these modules detects both the tall part of the ATV in the center and the dense shorter rear section.

9.4. Pile Detection

The piles were generally maintained by the loader to have a characteristic shape as shown in Figures 1 and 14(a), but there was often a secondary mound at the edge of the pile right after another tractor/harvester had unloaded, as shown in Figure 16(a).

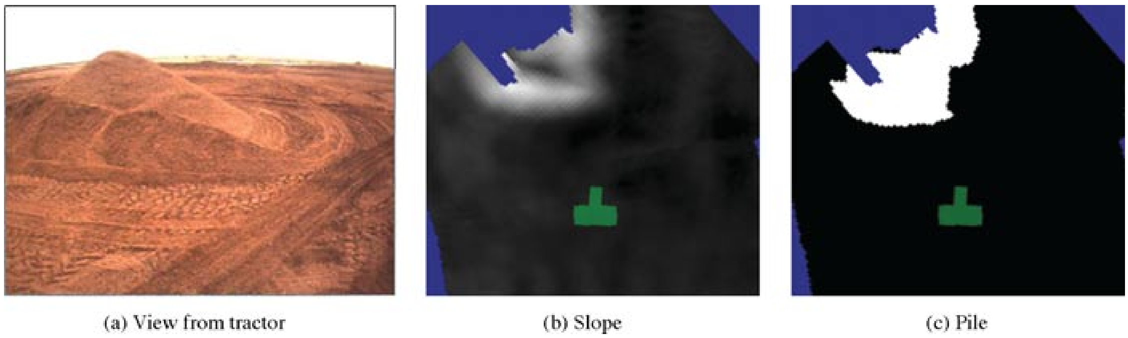

Figure 16. Pile detection using slope and blob filter.

Given an accurate ground height map, the pile detector used the slope operator in Figure 17 to find locations where the maximum slope θ over all orientations was within a range that is characteristic for peat moss piles. The edges of large drainage ditches often had a similar slope, so the lowest point of any potential pile location was required to not be significantly below the ground under the tractor.

Because of the multiple ridges in a pile after a dump, there were areas of the pile that did not match the characteristic slope range, as shown by the low-slope region in the center of the pile in Figure 16(b). Figure 16(c) shows the results of a blob filter that was used to connect regions of high slope into a single pile estimate. To reduce false piles due to sensor noise or small terrain features, the resulting pile estimate had to cover a minimum required area. Figure 15(b) shows the output pile map for the example in Figure 14.

9.5. Ditch Detection

As shown in the background of Figure 14(a), there were drainage ditches between the fields in the peat bog. These ditches were known hazards, and their GPS coordinates were included in the world map used for path planning, but as an added safety measure the ditches were detected using the onboard perception system.

The ditch operator shown in Figure 17 was convolved over the ground height map over all orientations to find ditches with some required depth d. Figure 18 shows the ground height map and the result of the ditch operator.

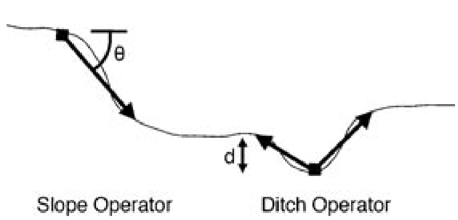

Figure 17. Side profile of an example ground surface showing slope and ditch operators used to find maximum slope θ and ditch depth d.

9.6. Implement Protection

As shown in Figure 1, there was an area in front of the vacuum heads of the implement that was behind the field of view of the sensor pod on the cab (during a left turn, this area is even larger). Static obstacles could be detected by the sensor pod, but a person or vehicle could enter this area without being sensed. To protect against this case, a fixed side-facing LADAR was mounted above the right rear wheels

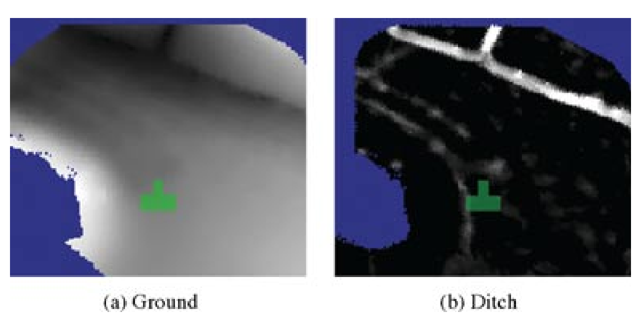

Figure 18. Ditch detection for the area behind the pile in Figure 14.

Figure 19. Automated tractors performing turns at the end of fields 1 and 2 during harvest.

and pointed slightly down from horizontal to act as a safety curtain (see Figure 1). The SideLadar module in Figure 13 reported obstacles to the side of the tractor just like the other algorithms detect obstacles in front of the tractor.

During turns, the position of the implement relative to the tractor changed. Instead of adding a sensor to measure this angle, the implement kinematics from Bevly (2001) were used to track the location of the implement vacuum heads based on the known vehicle motion, and then obstacles were detected only within the portion of the side LADAR scan that was in front of the implement.

10. FIELD TESTING

The automated peat moss harvesting system was deployed and tested under the supervision and control of the authors, for one complete season of peat harvesting. Figures 19 and 20 show examples of the system operating autonomously. To minimize the impact on the day-to-day harvesting operations at the farm, a majority of the testing was conducted without an implement attached to the robots. This facilitated testing in wet conditions and allowed other manually driven tractors to use the implements while the autonomous system was being developed. Including tests without the implement, these robots autonomously "harvested" more than 100 fields of peat moss. Harvesting, in this case, included making a lap through a field, planning to an unload point that is updated by perception, unloading, and navigating to the next field.

Position estimate logs were collected for many of the harvesting sequences. Figure 21 shows the position estimates collected during a customer test, overlaid on a plot of the world map and pile map used during the mission. This figure shows the distinct patterns driven by each member of the team as introduced earlier in Figure 7. During this test, tractor 1 began in field 1 (the field farthest to the right), harvested the field, and unloaded at the left-most pile (the pile to the right had been closed by the team leader). The first tractor then continued to the second field (left of the first field) and began to harvest. At this point, tractor 2 began to harvest the first field, driving nearly on top of the first tractor's path but in the opposite direction (see implement coverage pattern in Figure 7). Once the first tractor finished the second field, it unloaded material at the pile and continued to the third field for harvest. As the first tractor moved into the third field, the second tractor made its way to the pile, unloaded, and transported to the second field. This progression was repeated as tractor 3 entered field 1, bringing the full system into operation. In the data log used to create Figure 21, each tractor was stopped at the end of its respective field; however, during normal operation this coverage pattern was continued to the left as the tractors harvested the fields planned for that day.

Figure 20. Tractor autonomously dumping at a storage pile that had been acquired using perception.

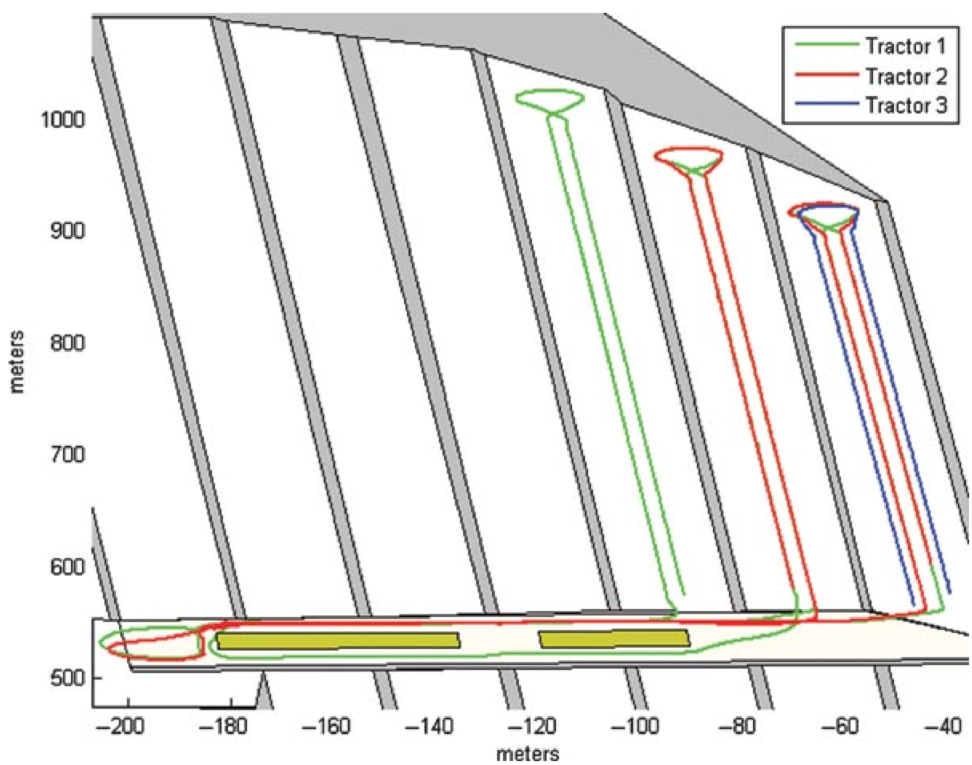

Figure 21. Plot of vehicle position logs shown on the world map. The horizontal boxes near the bottom represent the positions of two peat piles. The long vertical lines mark the locations of drainage ditches between fields. The large white spaces between gray areas are the harvestable fields.

Although the collection of quantitative performance data over extended test runs was not a priority for this phase of the project, limited data were available from the final three customer tests. Results from the full season testing and customer tests are presented below by subsystem.

10.1. Human Factors

One of the key elements of the successful implementation of this system was to provide a natural and simple transition for the team leader from the task of managing a manual peat harvesting operation to supervising a team of peat harvesting robots. To that end, it was critical for the TLUI to facilitate a rapid learning curve, as well as promote high situational awareness and an error-free user experience. The first exposure of the team leader to the TLUI occurred in the first customer test, during which a brief (5 min) training session was provided. The purpose of this training was to familiarize the team leader with the features of the tablet PC and the general functionality of the user interface. Overall, the training session was successful, as the team leader assimilated all the relevant information and appeared comfortable with the idea of "controlling" the vehicles via the TLUI. The only hesitation on his part was not being sure that the system was going to have a significant impact on the sequences of tasks he was used to performing. The main instruction provided to the team leader was "follow the same steps as if you were managing a manual operation, but instead of calling the harvester operator on the radio or signaling him to change the configuration of the vehicles, you can make the changes through here [TLUI]." After a couple of trial runs with all the robots, the team leader's interaction with the system appeared effortless. Specifically, the concept of being able to command the vehicles to stop and start, as well as set their engine speed, was very appealing to him.

At the start of the second customer test, which was approximately 1 week later, another short training session to review the functionality of the TLUI was provided to the team leader. During these trials the comfort level of the team leader with the TLUI was noticeable, as he was able to interact with the display without errors and while moving (in his ATV). A key issue that came up during the second test was his inquiry about the safeguarding capabilities of the robots. The team leader asked, "Will the harvester know to stop if I drive in front of it?" During manual operation, the harvester operators have full control of the vehicle, so the team leader has specific expectations about how close he can safely maneuver around those harvesters. The authors proceeded to explain that although the robots were equipped with perception and obstacle detection technology, safeguarding his vehicle (and himself) was his responsibility. Therefore, any time he needed to drive or park in front of the robot, he should acquire the habit of stopping the vehicle through the TLUI. An important safety feature of this system was that if a robot was stopped via the TLUI, it could be restarted only through the TLUI.

Future work should also document the "trust in automation" of the human by capturing behaviors such as his willingness to forfeit line of sight with the autonomous harvesters for extended periods of time. The development of *appropriate trust* (Lee & See, 2004) is likely the greatest human factors challenge for the implementation of this type of system in a commercial application. The critical issue associated with appropriate trust is the extent to which the automation is used (or relied on) by the human. To achieve appropriate trust, the human has not only to have an accurate perception of the overall reliability of the automation but also to understand the capabilities and limitations of the automation and use it accordingly. A well-documented impact of introducing automation into many systems is that when the automation operates without noticeable failures, for extended periods of time, humans become overcomplacent and begin to overrely on it (Parasuraman & Riley, 1997).

10.2. Multivehicle Control

The task allocation and constraint rules allowed multiple vehicles to perform the harvesting operation in a coordinated and safe manner. The rule set was tested in simulation prior to field testing and continued to work as expected throughout the full season and customer tests.

Each of the vehicles spent the majority of operating time executing the field coverage and harvesting task. The harvesting task is run at a constant ground speed for both automatic and manual operation and therefore was completed by the automated system in a time approximately equal to that of the manual operation.

Pile dumping data from the final customer test showed that the automated dump procedure typically took around 1 min, whereas a skilled human operator was able to complete a dump operation in approximately a quarter of this time. Unlike with human operators, the automated dump was a sequence of discrete actions. The system drove slowly near the pile, used the perception system to update the pile location information, planned a dump path, and drove forward while raising and then lowering the storage tank. Skilled human operators quickly drove to the pile and minimized the time spent dumping based on the actual amount of peat that had been collected. Although the automated dump took substantially longer than the manual dump, it represented a small portion of the overall field operation for any given vehicle. For both manual and automated operations, a nonstop complete harvesting cycle took between 15 and 21 min for a single vehicle to complete, depending on the field length.

The technical challenge remaining in the task allocation and constraint rules is to relax constraints that caused unnecessary delays not present in the manual operation. Specifically, the rules restricting the start-up sequence and access to the road need to be relaxed. Throughout the testing season, at system start-up, the harvesting order rule set forced the vehicle start times to be staggered by as much as 20 min. Similarly, the restriction on the number of vehicles allowed outside of the fields created scenarios in which a tractor would sit idle for up to 5 min while waiting for the other tractor(s) to clear the road.

10.3. Path Planning

The utilization of the two path planning methods to satisfy the four types of mission tasks worked well. Depending on the complexity of the map (number and complexity of keep-out areas) and the start and stop locations, execution times varying from less than 1 s to 5 s were observed.

Early customer tests revealed that 5-s delays between the completion of one task and the start of the next task were unacceptable from the safety rider's point of view. These planning delays were present in the "plan to pile" and "plan to field" tasks due to the relatively complex maps and vehicle model used by the grid-based planner that was employed for these tasks. The final system initiated these two planning tasks several seconds before the path solution was required by the vehicle to ensure no noticeable delay by the team leader or safety riders.

For the final customer tests, the path planning system experienced no problems. All generated paths were qualitatively reasonable, and the execution times required to generate those paths were acceptable when combined with the preplanning strategy described above.

10.4. Perception

The perception subsystem was responsible for pile detection and safeguarding the vehicle. Reliable pile detection was critical in this operation because the location of the edge of the pile continually changed as more peat was unloaded. By probabilistically modeling the noise characteristics of the range sensor data and the smoothness of the ground, the perception system was able to reliably estimate the ground surface and find piles effectively. The system accurately detected the edge of the pile for dumping in all operating conditions including heavy dust and rain (although actual harvesting would not happen during rain). Additionally, the performance of the pile edge detection was consistent, without any failures noted during the months of field testing.

In general, the fields were clear of hazards, and the coordination rules prevented the automated harvesters from being near each other. However, the harvesters operated in the same area as the team leader on his ATV and the loader driver that maintained the piles. Therefore, obstacle detection was required for safe operation. During system testing, the authors performed a number of obstacle tests with standing and walking people, other tractors, trucks, and ATVs. For this class of relevant obstacles that the system was designed to detect, the robots always detected the hazard and stopped at a safe distance.

As described above in Section 10.1, the operating procedure for the team leader was to stop the vehicles before approaching them, so the obstacle detection system primarily operated as a backup. For this reason, there generally were no obstacles present during operation for customer tests (other than placing an ATV in front of the tractor as a demonstration). During one of the customer tests, the team leader drove quickly past the front of an active autonomous tractor to pass it. The tractor came to a stop and then resumed its mission a few seconds later once the range data from the passing ATV cleared away.

The majority of peat moss harvesting occurs during dry conditions with low wind. Under these conditions, the obstacle detection system had very few persistent false positives that stopped the vehicle and required human intervention. Occasionally, some dust or sensor data misalignment would result in a false obstacle in front of the vehicle and the safe-speed algorithm would gently slow the tractor down. Additional sensor updates would generally eliminate the false positive and allow the system to seamlessly continue on its mission. During the customer tests, these slowdowns were infrequent, so they did not bother the safety rider or meaningfully increase the time to complete the harvesting operation. The automated system did not have any persistent false positives that required human intervention during the customer tests.

During windy conditions, there was a lot of dust, which caused problems for the obstacle detection system. Dust was sometimes detected as an obstacle and would result in a tractor slowdown. On occasion, this dust caused an obstacle to be detected close to the vehicle at the limits of the sensor field of view so that new sensor data were not able to clear out the false obstacle. In this case the harvester required human intervention to continue the mission. The side LADAR in particular produced many false obstacles due to its proximity to the wheels and the dust being produced during operation, and it often needed to be disabled. There were a few times when the wind caused so much dust that the entire obstacle detection system needed to be disabled to allow continued testing of the other components of the system (the safety operator in the cab was still present in this case).

Although the system was not set up to record the total number of false positives during all of the field tests, we did collect a set of perception logs, including some complete autonomous runs. Over four harvesting missions logged during typical calm conditions that included approximately an hour of operation and over 4 km of autonomous driving, the system had one brief false positive that caused a momentary slowdown. Conversely, during one 15-min, 1.2-km harvesting mission in very windy conditions, the system had six false positives, with five of them near the pile where the dust was worst. These numerical results match our qualitative experience that the system performed well during calm conditions but suffered from false positives in windy, dusty conditions.

The main technical challenge remaining for the perception component is increased robustness to dusty conditions, while maintaining good detection of obstacles. This is particularly challenging with moving obstacles because the laser signatures of dust and moving obstacles are similar. As described in Section 9.3.3, maintaining density measurements was an effective way to detect stationary obstacles and ignore light dust, but moving obstacles have low density similar to that of dust, so additional detectors were needed to reliably find moving obstacles and these detectors were more sensitive to dust. Future work includes the use of cameras and radars in addition to the LADAR sensors to improve performance in dusty conditions.

11. CONCLUSION

A system of three tractors for coordinated autonomous peat moss harvesting was successfully developed and tested in a working peat bog. The system and user interface were built to mimic existing manual processes, which made it easy for the human operators to use and integrate them into the rest of their operations. A set of coordination rules was applied to a combination of surveyed map data and onboard perception output to safely perform field missions. Although the current implementation does not yet match the speed of the manual operations, the extensive field testing of the system showed the potential viability of this approach.

ACKNOWLEDGMENTS

The authors would like to thank all of the team members involved in this project, including those from NREC (Cris Dima, Carlos Vallespi, and Herman Herman), T8 Design (Alex Stevenson), John Deere (Cameron Mott, Boyoon Jung, Robert Mayfield, Andrew Rekow, Shufeng Han, Chad Ackerman, Daniel Mairet, Jason Meredith, Zachary Bonefas, and John Reid), and our hosts at the peat bog.

REFERENCES

Adams, J. A. (2007, March). Unmanned vehicle situation awareness: A path forward. In Proceedings of the 2007 Human Systems Integration Symposium, Annapolis, MD.

Balch, T., & Arkin, R. C. (1995). Motor schema-based formation control for multiagent robot teams. In Proceedings of the First International Conference on Multiagent Systems (pp. 10–16). Menlo Park, CA: AAAI Press.

Bell, T. (1999). Precision robotic control of agricultural vehicles on realistic farm trajectories. Doctoral dissertation, Stanford University, Stanford, CA.

Bevly, D. (2001). High speed, dead reckoning, and towed implement control for automatically steered farm tractors using GPS. Doctoral dissertation, Stanford University, Stanford, CA.

Billingsley, J., & Schoenfisch, M. (1995). Vision guidance of agricultural vehicles. Autonomous Robots, 2(1), 65–76.
Bochtis, D., Vougioukas, S., Tsatsarelis, C., & Ampatzidis, Y. (2007). Field operations planning for agricultural vehicles: A hierarchical modeling framework. Agricultural Engineering International: the CIGR Ejournal, IX, February, Manuscript PM 06 021.
Chen, Q., & Luh, J. Y. S. (1994, May). Coordination and control of a group of small mobile robots. In Proceedings of the IEEE International Conference on Robotics and Automation, San Diego, CA (pp. 2315–2320).
Endsley, M. R. (1995). Measurement of situation awareness in dynamic systems. Human Factors, 37(1), 65–84.
Gerkey, B. P., & Matarić, M. J. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. International Journal of Robotics Research, 23(9), 939–954.
Gerrish, J. B., Fehr, B. W., Van Ee, G. R., & Welch, D. P. (1997). Self-steering tractor guided by computer vision. Applied Engineering in Agriculture, 13(5), 559–563.
Humphrey, C. M., Gordon, S. M., & Adams, J. A. (2006, October). Visualization of multiple robots during team activities. In Human Factors and Ergonomics Society Annual Meeting Proceedings, San Francisco, CA (pp. 651–655).
Humphrey, C. M., Henk, C., Sewell, G., Williams, B., & Adams, J. A. (2007, April). Assessing the scalability of a multiple robot interface. In Proceedings of the 2nd ACM/IEEE International Conference on Human–Robotic Interaction, Arlington, VA (pp. 239–246).
Lacaze, A., Murphy, K., & DelGiorno, M. (2002, July). Autonomous mobility for the Demo III experimental unmanned vehicles. In Association for Unmanned Vehicle Systems International Conference on Unmanned Vehicles, Orlando, FL.
Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46, 50–80.
Luczak, H. (1997). Task analysis. In G. Salvendy (Ed.), Handbook of human factors and ergonomics (pp. 340–416). New York: Wiley.
O'Conner, M., Bell, T., Elkaim, G., & Parkinson, B. (1996, June). Automatic steering of farm vehicles using GPS. In Proceedings of the 3rd International Conference on Precision Agriculture, Minneapolis, MN (pp. 767–778).
Ollis, M., & Stentz, A. T. (1997, September). Vision-based perception for an autonomous harvester. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robotic Systems, Grenoble, France (vol. 3, pp. 1838–1844).
Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39, 230–253.
Parker, L. E. (1994). Heterogenous multi-robot cooperation. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
Pilarski, T., Happold, M., Pangels, H., Ollis, M., Fitzpatrick, K., & Stentz, A. (2002). The Demeter system for automated harvesting. Autonomous Robots, 13(1), pp. 9–20.
Reid, J. F., & Searcy, S. W. (1987). Vision-based guidance of an agricultural tractor. IEEE Control Systems Magazine, 7(2), 39–43.
Reid, J. F., Zhang, Q., Noguchi, N., & Dickson, M. (2000). Agricultural automatic guidance research in North America. Computers and Electronics in Agriculture, 25(1), 155–167.
Rekow, A. (2000). System identification, adaptive control and formation driving of farm tractors. Doctoral dissertation, Stanford University, Stanford, CA.
Stentz, A., Dima, C., Wellington, C., Herman, H., & Stager, D. (2002). A system for semi-autonomous tractor operations. Autonomous Robots, 13(1), 87–104.
Treisman, A. (1986). Features and objects in visual processing. Scientific American, 255(5), 114–125.
Vail, D., & Veloso, M. (2003). Dynamic multi-robot coordination. MultiRobot Systems: From Swarms to Intelligent Automata, II, 87–100.
Vig, L., & Adams, J. A. (2006). Market-based multirobot coalition formation. In Distributed autonomous robotic systems 7 (pp. 227–236). New York: Springer.
Wellington, C., Courville, A., & Stentz, A. (2006). A generative model of terrain for autonomous navigation in vegetation. International Journal of Robotics Research, 25, 1287–1304.
Zhang, Q., Reid, J. F., & Noguchi, N. (1999, August). Agricultural vehicle navigation using multiple guidance sensors. In Proceedings of the International Conference on Field and Service Robotics, Pittsburgh, PA (pp. 293–298).
Zlot, R., & Stentz, A. (2006). Market-based multirobot coordination for complex tasks. International Journal of Robotics Research, Special Issue on the Fourth International Conference on Field and Service Robotics, 25(1), 73–101.