Autonomous Soil Investigator

Matt Bird, *birdmw@gmail.com*
Josh Gillham, *jgillham@msudenver.edu*
Scott McGimpsey, *scott.mcgimpsey@gmail.com*
Max Lichtenstein, *mnlichtenstein@gmail.com*
Matt Wicke, *mwicke@msudenver.edu*

Metropolitan State University
Dr. Megan Paciaroni
mpaciaro@msudenver.edu
April 8th, 2012

Abstract

The development of strategies for managing forests and forest fires is hampered by the difficulties in gathering data from a wide geographic region in rough terrain with many obstacles. An inexpensive autonomous robotic soil collector could be deployed in large numbers to reduce the cost of large soil sampling operations. In this paper, a proof-of-concept design is presented for a robot that can autonomously gather topsoil samples from a mapped region of moderately dense forest. This design was developed as part of the 2013 IEEE Region 5 student robotics competition, which challenged undergraduate teams to develop robots capable of retrieving soil samples in a simulated forest environment.

1. Introduction

Forest fires present a particular challenge in land management. The economic impact of these fires can be severe, but suppression of fires can cause long-term negative effects on forest ecosystems, as well as create the potential for more damaging fires. Efforts to manage these effects are hampered by a lack of data about forest ecosystems and by the need to improve understanding of large-scale forest disturbances [1][3]

One promising avenue for the development of forest land management strategies is the use of geostatistical soil surveys. The aim of this technique is to use a small number of soil samples from discrete points spread out across a geographical area to interpolate information about areas between these points. This form of analysis is useful in implementing precision agriculture and in conducting environmental assessments [2][4].

Several barriers prevent the use of geostatistical forest surveys. Firstly, the cost of taking samples distributed across a wide geographic area is often prohibitive. The sampling locations must necessarily include remote regions. Additionally, measurements at multiple points in time are often crucial [4]. Secondly, effective surveys should be conducted before and after a fire [1], but because the locations of future fires are unknown, effective surveys need to cover very large areas. Finally, knowledge about a geographic system is fundamentally constrained by the density of sampling locations [2]. These factors together limit the cost effectiveness of a given survey, reducing the utility of an otherwise plausible methodology for investigation.

These costs can all be reduced by the introduction of an automated system that can collect samples and perform measurements [5]. Automated systems thus far have been limited to those which require direct human implementation at the location of the soil sample. This paper presents the Metropolitan State University of Denver Autonomous Soil Investigator (MSUD ASI), a proof-of-concept robotic system designed to navigate a forest landscape and collect of soil samples. This robotic system can function autonomously at a total cost less than $2000 in raw materials. The development and implementation of this robotic system has the potential to dramatically reduce the cost of sampling efforts while also improving the geographic density and frequency of data collection, providing forest researchers with a powerful avenue for understanding and managing forest fires.

2. Overview of the ASI

The ASI is a four-wheeled rover that can configure itself to fit within a 1’x1’x1’ cube. The chassis is constructed from lightweight plywood, and both the motors and wheels are mounted to the direction control servos with steel brackets. The onboard computer is an

ARM-based embedded-linux platform, which communicates with an Arduino Mega 2560 microcontroller, allowing for processing-intensive applications while interfacing with a wide range of sensing and control devices. The high-level code is written in Python, with individual functions written in C++. The sensors include a 60Hz camera, infrared triangulators, and sonar rangefinders. The control systems include DC motors and servo motors. The power is supplied by a 11.1v hobbyist battery supply with two power transistor banks to handle a peak power of 96 watts each to the servos and the computer system.

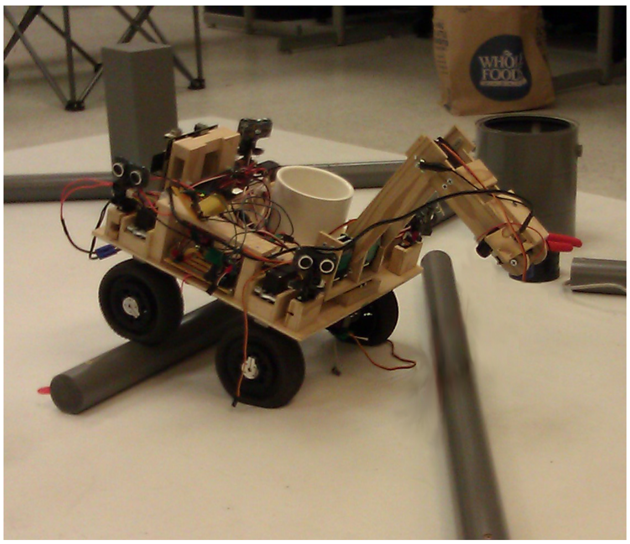

Figure 1. The ASI Rover

The design of the ASI is modular, implementing the functions of localization, pathfinding, traversal, and sample collection independently. Each system occupies a separate physical space in the robot's body, and the robot operates by conducting the actions of each system in polymorphic sequence.

3. Traversal

The traversing and pathfinding systems are designed in concert, since the two are interdependent. In code, the robot must strategically plan its movements based on goals and abilities. Mechanically, the system must be able to traverse the path given by the pathfinder. Choosing whether to solve a traversion problem through mechanical design or with careful pathfinding requires thorough empirical testing. Several traversion systems were tested in the laboratory, and a compromise was reached that utilized an inexpensive set of RC hobby wheels but demanded that the pathfinder guide that ASI across log obstacles at approximately right angle.

The major hurdles to traversal are space requirements, and climbing low-friction obstacles. Limited space prohibits a rocker-bogie suspension system or the use of large wheels. Obstacle avoidance was facilitated by allowing each wheel to be steered independently, allowing zero point turning, and the use of ackerman geometry. For obstacle climbing, the tires were deflated to gain contact surface area and improve traction.

4. Pathfinding

Intelligently navigating through an environment is one central requirement of an autonomous robot. As the complexity of the environment grows, so do the number of possible solutions. For this reason, the ASI references user generated map data. This map data is built using a custom GUI which allows users to store location specific information about the environment.

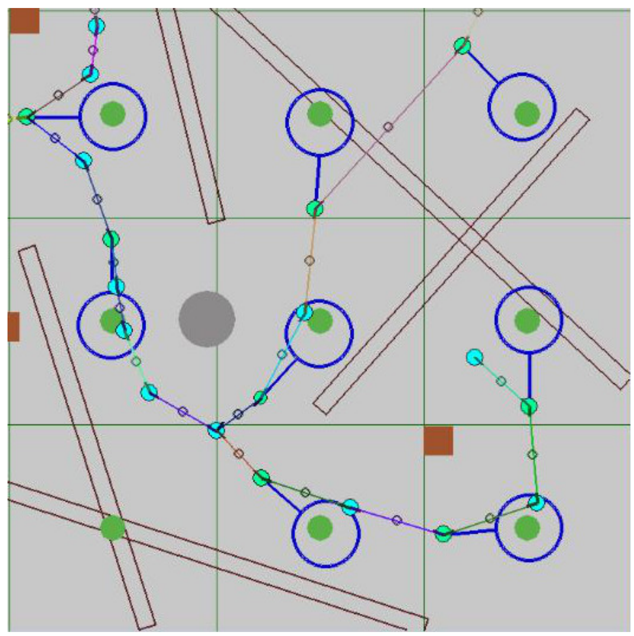

Figure 2. Pathfinder Node and Link Map

The pathfinding function utilizes a map of the region to be investigated, as well as a human-designed graph that represents the traversable paths available. Using a predetermined graph allows the robot to handle the challenges of different map regions in specialized ways. For example, map data can be used to cause the robot to localize more frequently while navigating narrow channels, or to power its wheels for a longer duration while climbing a slippery obstacle.

5. Localization

Localization is conducted with data from four "eye" modules, located on the four corners of the robot. The modules consist of an infrared triangulator, an ultrasound sonar unit, and a precision servo. These servos control the sensors as they sweep out 145 degree fan-shaped regions about the robot, returning ranges to nearby objects.

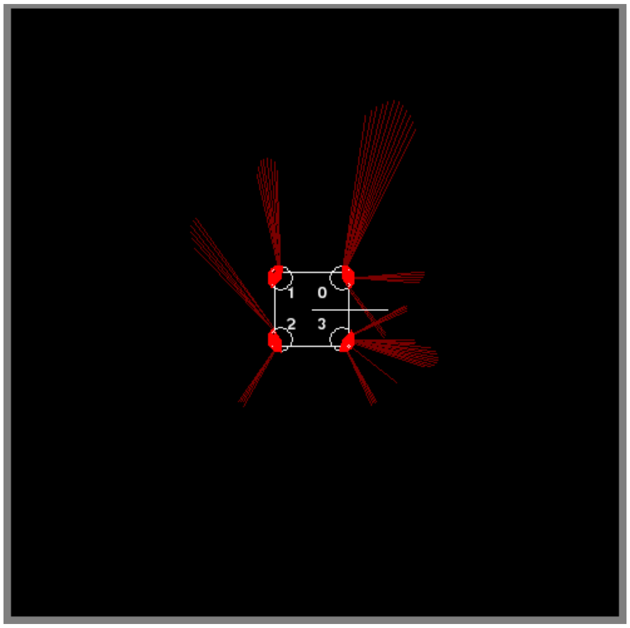

Figure 3. Proximity Data "Fan"

ASI implements a Monte-Carlo particle filter-like approach to localization. The robot's pose is considered as a probability distribution in {x,y, theta} space, where x and y are the cartesian coordinates of the robot, and theta is the angle between the robot's heading and the x axis of the map. This probability distribution is represented as a cloud of discrete hypotheses about the orientation and position of the robot, each with an associated weight. When the weights are normalized across all hypotheses, then each hypothesis' weight is approximately the probability that its pose is the most correct of any in the cloud.

At the beginning of the round, the hypothetical poses (or "hypobots") are arranged in a dense flat distribution across the starting region, each with a weight of one. As the robot moves, the hypobots mimic the expected motion of the robot through its pose space, with stochastic adjustments in x, y and theta at fixed time intervals. The cloud of hypobots then resembles a collection of particles diffusing across the map.

When the standard deviation in the cloud of hypobots becomes too large, or if the pathfinding algorithm deems it necessary, the robot performs a scan. As the eye modules rotate, the robot simulates ideal measurements from the eye modules for each hypothesis of its pose. Each hypobot then stores the results, populating lists with a simulated data point for each real data point measured by the physical sensors.

Each hypobot compares each real data point with its corresponding simulated data point, and multiplies its weight by the approximate bayseian probabilty of its own correctness given that data point. This probability is taken as a gaussian function of the difference between the real and ideal data point using empirically determined, non-linear noise from a sensor. This process is equivalent to applying a one-step extended Kalman filter to each hypobot.

Once all the hypobots are weighted, the weights are normalized and a weighted average pose is taken as the best guess pose. This result is passed to the pathfinder, but the hypobot cloud is retained. Hypobots with very low weights are discarded and new hypobots are cloned from the remaining cloud. This avoids wasting computational time on very unlikely hypobots while still maintaining a dense cloud.

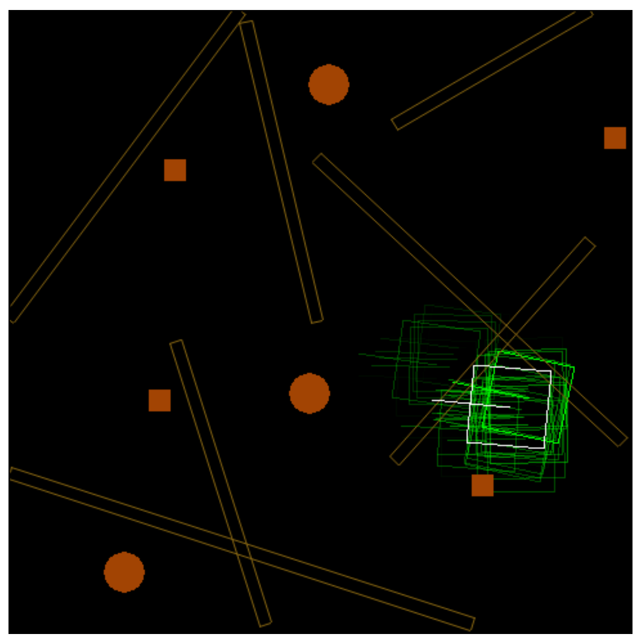

Figure 4. Hypobots With Associated Weights

This process has the advantages of being robust and adaptable, maintaining its effectiveness in a variety of maps. This adaptability comes at the cost of computational cycles, but this can be mitigated by performing expensive simulation calculations during the scan phase of operation.

6. Collection

The Autonomous Soil Investigator utilizes an arm with three rotational degrees of freedom mounted on one side of the robot's body for collecting soil samples. The links are 6 inches in length, and can nest compactly for transport. The joints of the arm are 360 degree servo motors, allowing for maximum versatility. The shoulder joint of the arm is driven by two servos, allowing the arm to produce torque sufficient to separate and lift a sample. The outermost link of the arm carries a 60hz 640x480 camera, a free hanging end effector mounted on a pivot, and an array of ultrabright LEDs for additional illumination.

To collect a sample, the robot first navigates to the approximate location of the desired sample, then extends its arm above the target region. A series of images taken by the camera whose view fits fully within the workspace of the arm. These images are parsed using tools from the OpenCV suite to identify a suitable point for sampling. The coordinates of this point are then parsed by an inverse kinematics function, yielding an efficient arm-movement solution. At this point, the arm extends to the configuration using a smoothing filter and the end effector is activated. If no suitable point is found by OpenCV, the robot relocates to another location within the target region, and the process is repeated. Once the end effector collects soil, the arm brings the sample to a hopper near the base of the arm and releases the sample.

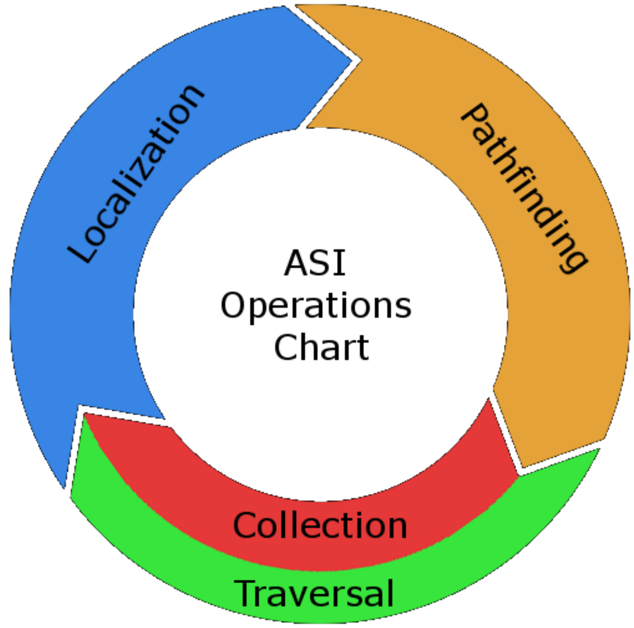

Figure 5. ASI Operations Chart

Currently, the end effector is implemented as an electromagnet for simple testing, which is not suitable for the collection of soil. This end effector can be replaced with a scoop or clamshell. The arm could also implement a derrick-style hoist to support heavier end effectors.

7. Limitations and Scalability

The Autonomous Soil Investigator has several important limitations, requiring further research and development prior to real-world deployment. A notable limitation is due to its purpose driven nature. Simply put, the ASI was not designed to move through extremely dense forests, or to be a rock climber. Other limitations, such as those stemming from water resistance, have plausible solutions which may be implemented in future versions of the ASI.

The ASI can effectively traverse small obstacles such as stones and branches, allowing use in more environments than other current solutions. However, this ability is still quite limited. The ASI does not perform well on severe slopes, and cannot traverse objects significantly larger than its wheel radius. Mountainous terrains, or terrains with high concentrations of obstacles, cannot be investigated by our system. To add perspective, these environments are likely to be inaccessible to any wheeled robot, since larger wheels would detract from the ability to navigate between trees and rocks.

In-the-field failures present another barrier to implementation. Debilitating failures can cause the robot to become immobile, requiring human retrieval. If these failures occur often enough, the cost of retrieving ASI units could drive its use to be more costly than human-driven soil sampling. These failures can be caused by threats in the environment such as shifting terrain, precipitation, interaction with wildlife, or theft. Additionally, failures can occur as a result of internal malfunctions or battery depletion. The ASI has not been field-tested in forest environments, so it is unlikely to perform robustly without significant additional development. Further development demands this field testing in order to fully asses the range of threats and obstacles to the ASI.

The prototype of the ASI is built primarily using off-the-shelf components. Other components, such as its power supplies and structural elements, are custom built. These components could be mass-produced using existing production technology. Altogether, the unit cost of an ASI is less than $2000, presenting the possibility of using fleets of hundreds of ASIs. Such a fleet could be placed on a truck, and units could be placed by the side of a road, activated, and collected later. In this way, a forest survey could be conducted in a short time by a two-person team with limited skills.

Several barriers exist to such an application, in addition to the technical limitations described above. Currently, the ASI requires an aerial map, in which obstacles and landmarks must be identified by a human. While this task is fairly easy, the size of the geographic regions to be surveyed is immense. Software solutions that can parse aerial photographs would mitigate the expense of this process. Such a solution is likely a prerequisite to implementation.

Another barrier to scalability is the technical expertise required to manage ASIs. While they can be deployed by untrained personnel, more training is required to store ASIs, charge them for use, and maintain them. Design revisions are required in order for the robots to be user and maintenance friendly.

While the ASI itself has many limitations, the fundamental design of the system enables the possibility of a multitude of applications for the device. For example, the use of computer vision on the end effector allows for a large suite of object recognition based on color or general shape. Few modifications to the overall design of the ASI would be necessary to convert the device from a soil sample collector to a fruit picker to even a rock sample collector. In the future, autonomous robots implementing these techniques could conduct soil surveys on other celestial bodies. With exception to the limits noted, as long as pre-implementation mapping is a possibility, the use of the ASI becomes an option.

8. Conclusion

The ASI is early in its development, and requires ADJECTIVE testing and further development. However, early tests suggest that it is an effective system for localizing, pathfinding, and traversing in environments with high concentrations of obstacles and collecting samples within that environment, given sufficient prepared map data. The modular design of the ASI allows the integration of new technologies and specialized instruments. This approach contains the promise of a wide range of applications, including large-scale automated techniques that can provide forest ecologists with revolutionary abilities to collect data.

References

[1] J.G. Germán Flores, D.A. Moreno G., and F.R. Rincón,“Planning, and Policy: A Global View - Reconstruction of Forest Areas Using Geostatistics as an Aid in the Evaluation of Burned Areas”, *Proceedings of the Second International Symposium on Fire Economics*, U.S. Department of Agriculture; Forest Service, Albany, CA, 2006-4-01, pp. 291-302.

[2] R.A. Olea, “A Practical Primer on Geostatistics”, *Open File Report*, U.S. Geological Survey, 2009-11-03, p. 346.

[3] A.L. Maclean, and D.T. Cleland, “The Use of Geostatistics to Determine the Spatial Extent of Historical Fires as an Aid in Understanding Fire Regimes for Northern Lower Michigan", *USDA Forest Service Proceedings*, RMRS-P-29, Rocky Mountain Research Station Publications, Fort Collins, CO, 2012-4-16, pp. 289-300.

[4] Soil Survey Division Staff, *Soil Survey Manual*, Soil Conservation Service, U.S. Department of Agriculture, 03/28/2011.

[5] M.J. Sadowsky, W.C. Koskinen, J. Seebinger, B.L. Barber, and E. Kandeler, “Automated Robotic Assay of Phosphomonoesterase Activity in Soils”, *Soil Science Society of America Journal Volume 2 Issue 70*, Alliance of Crop, Soil, and Environmental Science Societies, 2006-03-01, pp. 378–381.