

KEEPING GROUND ROBOTS ON THE MOVE THROUGH BATTERY & MISSION MANAGEMENT

BY TULGA ERSAL, YOUNGKI KIM, JOHN BRODERICK, TIANYOU GUO, AMIR SADRPOUR, ANNA STEFANOPOULOU,
JASON SIEGEL, DAWN TILBURY, ELLA ATKINS, HUEI PENG, JIONGHUA (JUDY) JIN, AND A. GALIP ULGÖY

From cleaning floors, handling warehouse materials, and inspecting sewer pipes, to locating and disarming explosives, ground robots are increasingly helping humans in new ways with dangerous, tedious, or inconvenient tasks. Their utility and impact, currently limited by their on-board batteries, would be enhanced if their energy and power capabilities can be extended. Understanding the dynamic limitations of the batteries with stringent design criteria (e.g., no cooling due to volume and weight considerations), and better management of the battery and the mission, can lead to longer and safer operations.

This article summarizes a recent collaboration at the University of Michigan's Automotive Research Center that considered best use of batteries in ground robots from several perspectives, such as planning the mission and tracking energy during its execution. **Figure 1** illustrates the four main subproblems addressed in this collaboration. Specifically, an area coverage problem is considered using a tracked robot, and the development of an energy efficient coverage plan is first addressed. Track-terrain interaction is then modeled to better predict the power consumption due to locomotion on different types of terrains. An electro-thermal

model of the battery is developed and used in a model predictive control framework to ensure that the battery is always operated within its electrical and thermal limits. A current-limiting approach is implemented to prevent the battery from overheating. Finally, a framework is developed that can combine the prior information from simulations or experiments with the online measurements to provide an adaptive and probabilistic estimate of the mission energy requirements. This framework allows the robot to predict the likelihood for a given mission to be completed with what energy remains in the battery. If mission failure is expected, coverage for the remaining area can be replanned with the available energy.

This article discusses these subproblems and the ongoing efforts to address them. A case study is presented to highlight the importance of the interactions among these subproblems.

The key message is that battery and mission management play a multi-faceted role in ground robotics. The most effective use of batteries in ground robots requires an integrated framework that considers all these factors. This is an important and exciting area to which the dynamic systems and control community can contribute.

ENERGY EFFICIENT COVERAGE PLANNING

Area coverage is a common task for a ground robot; applications include floor cleaning, lawn mowing, or sweeping a sensor through a region, searching for explosives or chemicals, or performing surveillance. Unlike traditional motion planning tasks, which involve moving from a start to a target point, the coverage task determines a path along which the robot passes near every point in the region. It is a time and energy intensive task. Previous research in creating algorithms to generate coverage paths typically did not seek to achieve an optimal path with time and energy constraints. Our work extends the previous coverage planners by planning an optimal velocity trajectory along the route using the techniques of optimal control to balance the mission time, the efficiencies of the electrical system, and the area remaining to be covered.

The proposed solution to this problem can be summarized in the following three steps. First, a coverage path is planned using any existing algorithm from the literature, treating the path as a series of waypoints. In our study, we used the boustrophedon decomposition method. The area is broken down into polygonal

regions and each region is covered by simple back and forth movements. Other methods for path planning are summarized in reference 2 and can also be used for the purposes of this step. In addition to knowing the region, the key parameter for these methods is the search distance from the robot. This value is based on the current sensor in use and defines how far apart the adjacent path segments must be for coverage. Second, a cost function is defined as a linear combination of the track forces, the ratio of remaining area to be covered to the total area, and motor efficiency. Hence, this cost function penalizes (1) the energy expended by the robot while completing the mission through the track forces, (2) the uncovered areas, and (3) the operation of the motor in inefficient regions. Finally, in the third step, an optimal control problem is solved for this cost function to find an optimal velocity profile on the path generated in the first step. In addition to the cost function, the optimal control formulation includes a dynamic model for the robot that uses position, heading, forward velocity, and yaw rate as the states. At each waypoint, the robot must turn towards the next waypoint. A moving turn is used instead of having the robot come to a complete stop at each waypoint and turn in place. This promotes efficiency due to the terramechanics as discussed in the next section, albeit at the cost of missing small patches of area to be covered.

As an example, **Figure 2** shows a simple region with the path generated in Step 1 shown on the x-y plane, where the shaded region represents an obstacle. The optimized velocity profile over the path is shown as the third dimension of the plot. While many sections of the profile look similar, there are differences in the velocity as the robot covers more area. Full analysis of the tradeoffs between time and energy are presented in reference 3. Once the full trajectory has thus been planned, the robot can be driven along this trajectory using a trajectory-tracking controller.

LOCOMOTION POWER ESTIMATION

Locomotion power (straight-line travel and skid steering) can consume a large percentage of the total power in a tracked ground robot. For example, order-of-magnitude power requirements might be 100W for locomotion, 10W for computation, and 1W for communication. Therefore, a terramechanics model is critical for accurate power and energy analysis.

Compared to straight-line travel, modeling skid steering of tracks is more difficult on soft soils because of the track-soil interaction and the distributed nature of shear stress along the large contact area. Thus, researchers developed several methods to approximate skid steering in steady state operation. Our power modeling of skid steering is based on Wong's theory, in which the turning resistance coefficients vary with both turning radius and forward vehicle velocity. However, our approach includes some simplifications to achieve computational efficiency assuming uniform pressure under the tracks and a small slip ratio. The basic idea of this fast computation method is that the skid steering equations can be separated

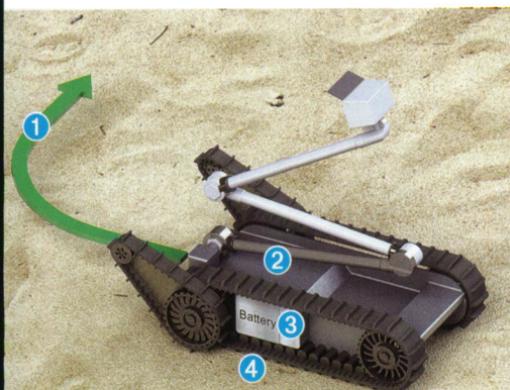


FIGURE 1

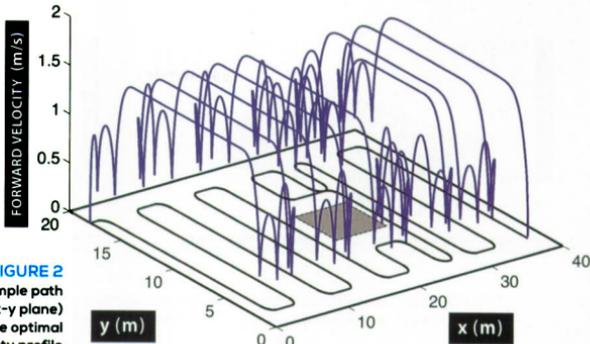
Overview of the four major subproblems considered in the collaborative effort

- 1 ENERGY EFFICIENT COVERAGE PLANNING** The velocity profile along the path is optimized to minimize energy consumption, avoid inefficient motor operation, and maximize coverage.
- 2 ONLINE ENERGY TRACKING** The probability of completing the mission with the remaining energy is tracked combining prior knowledge about the mission with real time data.
- 3 BATTERY POWER MANAGEMENT** Using a thermo-electric model, the maximum battery power is controlled to avoid violation of voltage, temperature, and SOC limits.
- 4 LOCOMOTION POWER ESTIMATION** Terramechanics models are used to predict the power needed for locomotion on a given terrain type depending on velocity and turning radius.

into the computationally expensive part of accounting for the shear displacement - shear stress distribution on the track and the computationally cheaper part of solving the force and moment balance equations for the vehicle. Our formulation solves the first part separately and stores the solution in look-up tables to be used in the second part. This approach gives results very close to solving the two parts in a coupled manner. Details can be found in reference 6 along with experimental validations.

As an example, **Figure 3** shows the power consumption predictions of the model during skid steering for two different soil types, dry sand and sandy loam. This figure serves two purposes. First, it shows that the typical trend is that the power consumption increases for smaller turning radii and higher velocities. This is the reason why it is beneficial to slow down while negotiating the turns as seen in Fig. 2 and why a moving turn was preferred in the coverage planning task instead of stopping and turning on the spot. Second, Fig. 3 illustrates that the locomotion power requirement can change by as much as 100% due to the soil type. This is why it is critical to have a terramechanics model as part of the robot battery and mission management framework.

FIGURE 2
Example path
(on the x-y plane)
with the optimal
velocity profile
shown as the
third dimension.



BATTERY POWER MANAGEMENT

The third major component in this collaborative effort is the development of a new battery power management strategy. In this context, battery power management refers to ensuring that the battery is operated within its voltage and temperature limits.

In this study, a lithium-ion battery is considered as a typical battery type. To avoid aging and capacity loss, and to ensure safe operation, limits on the operating temperature of the battery must be enforced. In general, temperature regulation of battery packs involves using either active thermal management systems or limiting the peak current drawn from the pack. These strategies increase the rate of heat rejection or limit the rate of internal heat generation, respectively. In a mobile robot application, the first strategy is not feasible, since mobile robots rarely have a cooling system due to volume and weight limits. It is critical to control the discharge current of batteries so that operating temperatures do not exceed maximum value. Thermostatic or proportional-integral-derivative (PID) controllers are traditionally used to limit current or power drawn from the battery when temperature exceeds the pre-defined limits. But calibrating thermostatic thresholds, dead-bands, and PID gains and integrating them with the overall power allocation strategies in battery management systems is a challenge. As an alternative, we have developed a model-based method to estimate the maximum power capability of

the battery that accounts for not only electrical constraints such as terminal voltage and battery state-of-charge (SOC), but also thermal constraints.

In estimating the maximum power capability, the following factors are considered:

- The thermal and electrical dynamics of a lithium-ion cell are intrinsically coupled. For a constant current, any arbitrary increase in cell temperatures will cause reduced internal losses, and subsequently generate less heat.
- The rate of change of internal resistance with respect to temperature decreases with increasing temperatures.
- Over a reasonably short horizon, the temperature increase can be assumed to be bounded. Similar arguments can be made for the change in the electrical quantity SOC.

The above statements are valid insofar as the temperature of the cell does not exceed the threshold temperature at which thermal runaway is initiated. Since thermal dynamics are much slower than electrical dynamics, considering electrical and thermal constraints independently over a short horizon yields conservative estimates of power capability. Consequently, the thermal and electrical constraint problems are addressed separately in our approach. This is done by developing models capturing the electrical and thermal dynamics, and using them to calculate the maximum constant current over a prediction horizon of 10 seconds that does not lead to any voltage, SOC, or temperature violations. If the demanded current from the battery exceeds this maximum allowable current, the actually delivered current is limited by the maximum amount. The maximum power capability can then be found as the product of the maximum allowable current and the terminal voltage.

Figure 4 illustrates the performance of the battery power management algorithm during repeated duty cycles. Note that all constraints are inactive initially. Hence, the battery can meet the power demand up to 4705 s until the voltage constraint is violated first. This is because as the power is drawn from the battery, the battery SOC is reduced. The corresponding decrease in open circuit voltage and voltage drop caused by internal resis-

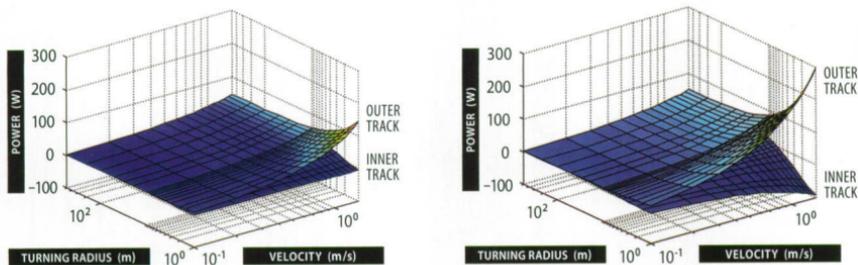


FIGURE 3 Simulation of power consumption of a tracked robot running on dry sand (left) and sandy loam (right).

tances lead to a predicted voltage constraint violation. This predicted violation activates the power limiting algorithm so that the terminal voltage can be kept above the minimum limit of 3.2 V. As the cell temperature approaches the maximum temperature limit, 45°C in this case, the power capability begins instead to be determined by the maximum temperature limit. This helps the battery temperature to be kept at the limit. Finally, the SOC constraint becomes active and the battery eventually turns off. This performance highlights that the proposed method can estimate the power capability accounting for thermal and electrical constraints. Thus, safe and reliable operation of the battery is achievable.

ONLINE ENERGY TRACKING

The previous two sections have discussed models for the thermo-electric dynamics of the battery and the interactions between the tracks and terrain. These models are useful to plan a mission as seen in the Energy Efficient Coverage Planning section, or develop the power management techniques of the Battery Power Management sec-

tion. However, there may be a difference between the conditions during the execution of the mission, and those assumed in the planning stage. A method to track the available energy online and predict potential mission failures due to energy limitations is needed. This does not necessarily negate the simulation results. In fact, the method highlighted in this section combines prior knowledge from simulations (and/or prior experiments) with real-time data collected as the mission is run to predict whether it can be completed with the remaining energy. If failure is predicted, the Energy Efficient Coverage Planning task can be re-visited and the coverage mission can be re-planned taking the remaining available energy into account.

In this approach, a Bayesian regression model is used to predict mission power when prior knowledge of road segments is available. A road segment has a consistent average grade and surface condition. The model parameters are recursively updated based on real-time measurements of the robot velocity and energy consumption. The updated model is used to predict the future power consumption by leveraging an experimentally validated, linearized vehicle longitudinal dynamics model. The probability of accomplishing the mission can be adaptively estimated during its execution. Details of the approach are given in references 9 through 11. In the example mission of Figure 5, the Bayesian approach outperforms that of the traditional linear regression, which ignores prior knowledge and can under or over estimate the mission energy requirement.

ABOUT THE AUTHORS

Tulga Ersal is an Assistant Research Scientist in the Department of Mechanical Engineering at the University of Michigan. His expertise is in modeling, identification, simulation, and control of system dynamics, including hardware-in-the-loop simulation, with applications to batteries and vehicle powertrains among others.

Youngki Kim is a Ph.D Student in the Department of Mechanical Engineering at the University of Michigan. His research interests include design and control of hybrid electric vehicles and lithium-ion battery modeling and estimation.

John Broderick is a graduate student in the

Department of Electrical Engineering at the University of Michigan. His research interests include modeling and control of ground robots to increase reliability.

Tianyou Guo is a graduate student in the Department of Mechanical Engineering at the University of Michigan. His research interests include vehicle dynamics modeling and control, robotics, terramechanics and sizing and design of hybrid vehicles. His current research focuses on power modeling and control of small unmanned ground vehicles.

Amir Sadrpour is a graduate student in the Department of Industrial and Operations

Engineering at the University of Michigan. His research focuses on energy-based mission reliability assessment for unmanned ground vehicles.

Anna Stefanopoulou is a Professor in the Department of Mechanical Engineering at the University of Michigan and the Director of the Automotive Research Center. Her research interests include estimation and control of internal combustion engines and electrochemical processes such as fuel cells and batteries.

Jason Siegel is an Assistant Research Scientist in the Department of Mechanical Engineering at the University of Michigan. His research interests cover electrochemical energy storage

INTEGRATED CASE STUDY

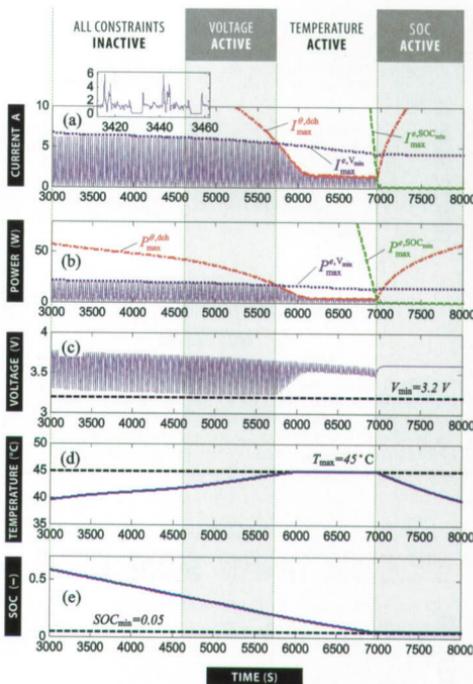
We combined the models and algorithms described above into a single simulation framework and performed a case study to highlight the importance of an integrated solution to the problem of making ground robots more energy and power aware.

The case study started with defining an area to cover. The coverage planning algorithm generated a path and planned the most energy efficient velocity trajectory to follow along that path. A trajectory tracking controller then drove the robot dynamics model along this path while the terrain-mechanics model predicted the loads due to the interaction between the tracks and terrain and the battery dynamics model predicted the thermo-electric implications of this trajectory for the battery. Two simulations were run.

The first simulation represented the pre-mission analysis and showed that if the terrain type is dry sand and the mission starts with a fully charged battery at 35°C ambient temperature, then the mission can be completed successfully with 60% remaining SOC and 40°C final battery temperature, 20°C below the assumed maximum allowed limit of 60°C.

The second simulation represented the actual mission scenario, where the mission actually starts with 70% SOC in the battery, which is different from what was assumed for the pre-mission analysis, but would still be sufficient to finish the mission, if the soil type was dry sand throughout the entire area. However, the actual mission scenario also assumed that the terrain type switched from dry sand to sandy loam approximately 1/3 of the way into the mission. When this simulation was run, shortly after the terrain type switched, the mission energy prediction algorithm correctly predicted failure; i.e., if the simulation had continued as is, the battery would have run out of energy before the mission was completed due to the increased power requirements for the sandy loam type of terrain. When failure was predicted, the coverage planning

FIGURE 4
Performance of power capability estimation method during repeated operations at 30°C ambient temperature and natural convection (6 W/m²/K): (a) current; (b) power; (c) voltage; (d) temperature; (e) SOC



algorithm was re-run to cover as much area as possible with the remaining available energy. The simulation then continued with the updated path and velocity trajectory, and showed another implication of the change to a more power-demanding terrain type; namely, the battery reaching its temperature limit of 60°C. When this happened, the power management algorithm described above prevented the battery temperature from exceeding 60°C

and conversion for automotive applications.

Dawn Tilbury is a Professor at the University of Michigan's Mechanical Engineering Department. Her research includes distributed control of mechanical systems with network communication, logic control of manufacturing systems, reliability of ground robotics, and dynamic systems modeling of physiological systems.

Ella Atkins is an Associate Professor in the Aerospace Engineering Department at the University of Michigan. Her research focuses on the safe operation of robotic systems through the management of anomalies, concentrating on aerial vehicles.

Hueli Peng is a Professor in the Department

of Mechanical Engineering at the University of Michigan. His research interests include adaptive control and optimal control, with emphasis on their applications to vehicular and transportation systems. His current research focuses include design and control of hybrid vehicles and vehicle active safety systems. He is currently the U.S. Director of the DOE sponsored Clean Energy Research Center-Clean Vehicle Consortium.

Jionghua (Judy) Jin is a Professor in the Department of Industrial and Operations Engineering at the University of Michigan. Her research interests are primarily in the areas of industrial statistics and quality

engineering. Most recently, her research has focused on data fusion for complex system modeling, design innovation, and performance improvement through optimized decision making, with applications to various automotive and semiconductor manufacturing processes, transportation, and human decision support systems.

A. Galip Ulsoy is a C.D. Mote, Jr. Distinguished University Professor of Mechanical Engineering and the W.C. Ford Professor of Manufacturing at the University of Michigan. His research interests focus on dynamic modeling, analysis, and control of mechanical systems.

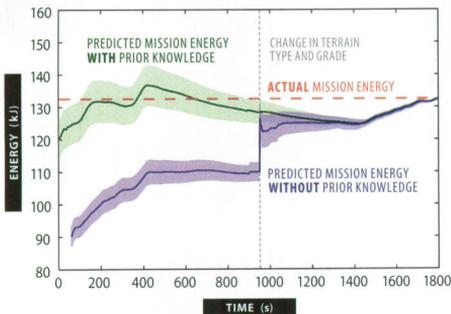


FIGURE 5 Mission energy predictions with prediction confidence intervals (shown as shaded regions) using the Bayesian (green) and linear regression (blue) approaches. Unlike the Bayesian approach, the linear regression approach ignores the prior knowledge and initially underestimates the mission energy requirement.

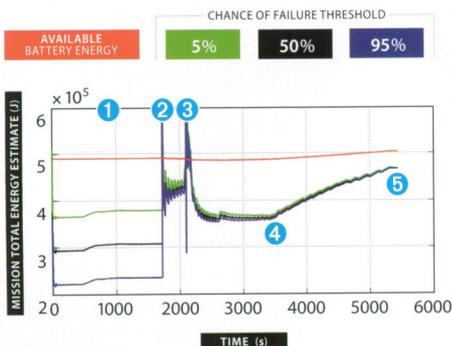


FIGURE 6 Tracking mission energy during a simulated coverage task mission execution.

- 1 Mission starts as predicted. Robot has sufficient energy to complete the mission.
- 2 Terrain type changes unexpectedly. Predicted energy requirement starts increasing.
- 3 Failure is predicted. Mission is re-planned with the remaining available energy.
- 4 Battery temperature reaches the maximum allowable limit. Battery power is constrained to regulate temperature. Lower speeds lead to inefficient motor operation. Mission energy prediction starts increasing.
- 5 Mission is completed without overheating and before the battery runs out of charge.

by limiting the power that can be drawn from the battery. This, however, also prevented the robot from following the optimal velocity trajectory and caused it to operate under inefficient conditions. Inefficient operation caused the mission energy requirement predictions to gradually increase. Nevertheless, the mission was completed before another failure was predicted. **Figure 6** shows the mission energy requirement predictions during this second simulation.

CONCLUSIONS

Using a coverage problem as an example, this article shows that managing the battery and the mission properly is critical for ground robots to successfully complete given tasks and make maximum use of their capabilities. To this end, the problems of energy efficient coverage planning, predicting the locomotion power requirements, controlling the battery power with thermal and electrical constraints, and tracking the mission energy requirements online based on a combination of prior knowledge and real-time data are all tightly connected to each other. Therefore, the best answers to the question of how a ground robot should most effectively utilize its battery are more likely to come from such integrated solutions. The collaboration presented here is a first demonstration of such an integration. ■

REFERENCES

- 1 D. M. Tilbury and A. G. Ulsoy, "A New Breed of Robots That Drive Themselves," *Mechanical Engineering Magazine*, vol. 133, no. 2, pp. 28-33, 2011.
- 2 H. Choset, "Coverage for robotics—a survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, no. 1-4, pp. 113-26, 2001.
- 3 J. A. Broderick, D. M. Tilbury, and E. M. Atkins, "Optimal coverage trajectories for a UGV with tradeoffs for energy and time," *Autonomous Robots*, pp. 1-15, in press.
- 4 H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, no. 3, pp. 247-253, 2000.
- 5 J. Y. Wong and C. F. Chiang, "A general theory for skid steering of tracked vehicles on firm ground," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 215, no. 3, pp. 343-355, 2001.
- 6 T. Guo and H. Peng, "A simplified skid-steering model for torque and power analysis of tracked small unmanned ground vehicles," *American Control Conference*, Washington, DC, pp. 1106-1111, 2013.
- 7 D. Andrea, *Battery management systems for large lithium-ion battery packs*, Norwood, MA: Artech House, 2010.
- 8 Y. Kim, S. Mohan, J. B. Siegel, and A. G. Stefanopoulou, "Maximum power estimation of lithium-ion batteries accounting for thermal and electrical constraints," *ASME Dynamic Systems and Control Conference*, 2013.
- 9 A. Sadrpour, J. Jin, and A. G. Ulsoy, "Experimental Validation of Mission Energy Prediction Model for Unmanned Ground Vehicles," *American Control Conference*, Washington, DC, 2013.
- 10 A. Sadrpour, J. Jin, and A. G. Ulsoy, "Mission energy prediction for unmanned ground vehicles using real-time measurements and prior knowledge," *Journal of Field Robotics*, vol. 30, no. 3, pp. 399-414, 2013.
- 11 A. Sadrpour, A. G. Ulsoy, and J. Jin, "Real-Time Energy-Efficient Path Planning for Unmanned Ground Vehicles Using Mission Prior Knowledge," *International Journal of Vehicle Autonomous Systems*, in press.