Proceedings of the American Institute of Aeronautics and Astronautics, Arlington, VA, September 26 - 29, 2005.

Intelligent Systems for Construction and Long Range Exploration on Planetary Surfaces

James S. Albus*
National Institute of Standards and Technology, Gaithersburg, Maryland, 20899

For efficient and effective operations on Lunar and/or Martian surfaces, it will be necessary to develop robots with the ability to successfully perform extended missions of exploration and complex tasks of construction without imposing a large workload on astronauts or requiring high bandwidth, rapid turn-around communications with a ground station. For long-range exploration covering distances of hundreds of kilometers while performing experiments and collecting scientific data along the way, robot vehicle speeds of 2 to 5 meters per second will be required. For construction of habitat and base facilities on lunar or planetary surfaces, ultra-light machines with the capability to lift and position large heavy structures and to manipulate parts and tools will be required. To control such machines in performing complex tasks of exploration or construction, a software architecture is required that can integrate sensing, perception, world modeling, planning, and control into a cognitive system that can understand its situation and behave appropriately under a wide variety of complex situations.

I. Introduction

Transportation of equipment to the Moon or Mars, landing safely on the surface, and deploying systems for construction of base facilities and long-range exploration of Lunar and Martian surfaces is a complex and costly enterprise. Weight and volume delivered to planetary surfaces must be minimized.

However, construction of base facilities will require large robots with heavy lift capacity, and long-range exploration will require vehicles capable of much higher speeds than current exploration robots. Current Mars rovers have an average speed of about 30 m/h, i.e., less than 1 cm/s. The Lunar and Martian surfaces are large. The circumference of the Moon is almost 11 000 km (6835 miles) and its surface area is roughly 37 million km^2 (21 million mi^2). The circumference of Mars is more than 21 000 km (13 048 miles) and its surface area is about 145 million km^2 (90 million mi^2). To explore any meaningful percentage of the Lunar or Martian surfaces will require vehicles capable of traveling many kilometers per day. Cruising speed over relatively smooth terrain should average at least 10 km/h, or roughly 3 m/s. To safely achieve such speeds, the design philosophy of planetary rovers must be modified. Vehicles with large wheels and a long wheelbase will be required.

To simultaneously achieve large size and light weight, one solution is to design vehicles that use inflatable structures for compression members and lightweight high-strength cables for tension. Inflatable structures and cables have a further advantage in that they can be packaged in a small volume and encapsulated in airbags for landing. Upon arrival, they can be deployed for operation by simply inflating them. Once inflated, they can be filled with lightweight foam to maintain structural integrity in case of puncture.

To support compression, the most efficient structure is a thin-walled cylinder with length-to-diameter ratio adequate to resist Euler buckling (i.e., length to diameter ratio less than 50.) To minimize weight, the thickness of the cylinder walls should be minimized. This, of course, creates a problem of crushing. The best way to prevent crushing of thin cylinder walls is to inflate the cylinder with sufficient pressure to assure that the cylinder walls always remain in tension.

---

* Senior NIST Fellow, National Institute of Standards and Technology, Gaithersburg, MD, Bldg 220, Rm B-124, AIAA non member.

II. An Exploration Robot Design Concept

Figure 1 illustrates a conceptual design for an exploration robot made from inflated tubes. The basic structure is a tetrahedron with three wheels. The wheels are 1 m in diameter with a wheelbase of 3 m. All wheels are powered by electric motors. The rear wheel provides steering. Mounted above the rear wheel is a directional antenna for long distance communication. The apex of the tetrahedron supports a sensor suite mounted on a pan/tilt head. These sensors are designed for long-range surveillance and mobility path planning.

The main structural members of the chassis are 20 cm (8 in) diameter inflated tubes made from stiff synthetic fabric. A cylinder of diameter d inflated to a pressure of p will support a load of $p\pi d^2/4$. Thus, when inflated to a pressure of about 69 kPa (10 psi,) each 20 cm (8 in) tube can support a compressive load of about 2334 N (502 lbf.)

The tetrahedral chassis supports a work platform that is a smaller inflated tetrahedral volume. The work platform is suspended by eight cables from three attachment points on the top face of the work platform. These attachment points are connected by eight cables to eight winches located at the vertices of the tetrahedral chassis. Movement of the work platform relative to the tetrahedral chassis is enabled by eight winches that control the length and/or tension of the eight cables. The lengths of any six of the cables are sufficient to define the position and orientation of the work platform. The remaining two cables are controlled in tension to prevent any cable from going slack. A tool for grinding and analyzing selected rocks is shown attached to the front face of the work platform. A sensor suite for monitoring the tool with respect to a work sample is mounted above the tool.

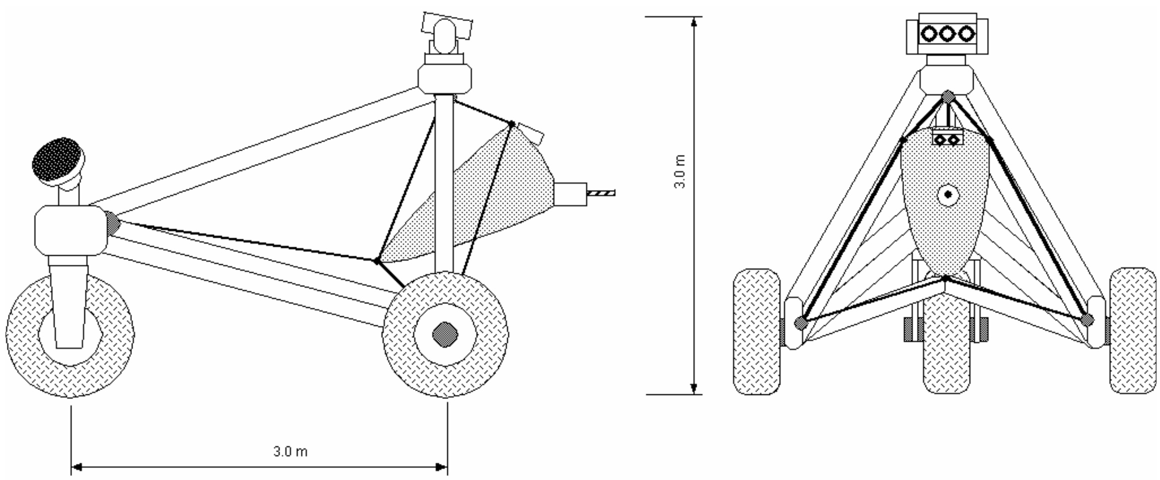

(a) (b)

Figure 1. A conceptual drawing of a long-range exploration robot with inflatable chassis and wheels. *(a) side view, (b) front view. A work platform with a coring tool is suspended by eight cables whose lengths and/or tension are controlled by eight winches mounted at the vertices of the chassis. Six of the cables are controlled in length and the remaining two are controlled in tension. This enables the work platform to be positioned and oriented in six degrees of freedom with all cables maintained in tension.*

Three of the eight winches are located at the apex of the chassis. These control the vertical position and the pitch and roll orientation of the work platform. Two winches located near one of the front wheels control the side-to-side position and the roll orientation of the work platform. A single winch located over the rear wheel controls the forward-backward position. Two winches located near the left front wheel provide tension necessary to keep all cables in positive tension. This arrangement allows the work platform to be controlled in all six degrees of freedom. As shown in Figure 2a, the work platform can rotate more than 90 degrees in pitch so that a tool mounted on the front of the work platform can address rocks either on a vertical surface, or on the ground. The remaining degrees of freedom allow the tool to be oriented normal to the face of selected rocks, either on the ground or a vertical rock face. It also allows the work platform to be moved so as to shift the center of gravity of the vehicle during traversal of steep slopes.

This cable arrangement enables a closed form solution to the forward and inverse kinematic transformations between the chassis and the work platform. The rear attachment point is connected to four winches, one at each vertex of the chassis. The length of any three of these cables defines the position of the rear attachment point relative to the chassis. The fourth cable maintains tension in the other three. The distance between the rear attachment point and the right-front attachment point, plus the lengths of two cables, one from the apex and the other from the right-front wheel, provides the position of the right-front attachment point. Finally, the distance between the left-front attachment point and the other two attachment points, plus a third cable from the apex determines the location of the left-front attachment point. The remaining cable maintains tension in all front cables.

Cables between the vertices of the vehicle chassis and the work platform allow the platform and its attached tools or grippers to be positioned and oriented throughout a large and useful work volume. When the work platform is positioned horizontally, as in Figure 1, tools and grippers attached to it can reach well beyond the front of the robot to perform work on a vertical face in front of the robot. By changing the cable lengths, the work platform can be rotated 90 degrees down so that the tool or gripper attached to the work platform can address the ground beneath the robot as shown in Figure 2a. Thus, the robot can work on a vertical face in front of it, or on the ground beneath it.

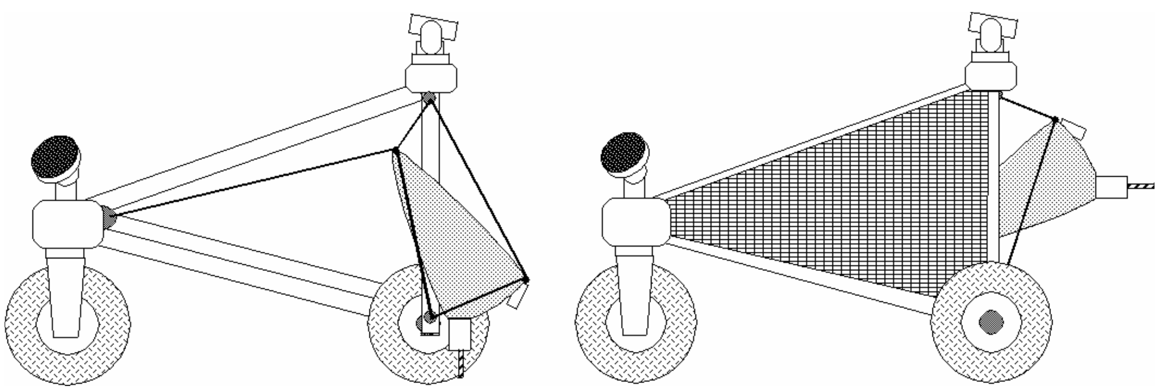

(a) (b)

Figure 2. Long range exploration robot. *(a) A cut away view of a long-range exploration robot with the work platform and coring tool rotated into a vertical position for operating on rock samples lying on the ground. (b) A long-range exploration robot with solar cells installed on its sides.*

The upper rear surfaces of the tetrahedral chassis can be covered with solar cells as shown in Figure 2b to generate electricity to power the robot.

III. Construction Robot Design Concepts

For building and maintaining base facilities, large construction robots will also be required for lifting, transporting, and positioning large heavy objects such as pressurized habitats for astronauts, or facilities for processing in-situ material. Large construction robots will also be required for mining and transporting ore for extraction of chemicals required for fuel and building materials. Objects to be manipulated may be several meters in size and have a mass of hundreds of kilograms.

A conceptual drawing of a large construction robot for lifting and positioning large heavy objects is shown in Figure 3. The main structural members are 50 cm (20 inch) diameter inflated tubes 10 m (33 ft) in length made from stiff synthetic fabric. When inflated to a pressure of about 69 kPa (10 psi), each tube can support a compressive load of about 13 966 N (3140 lbf) Each of the tubes is supported by a two wheeled (or two tracked) vehicle for mobility.

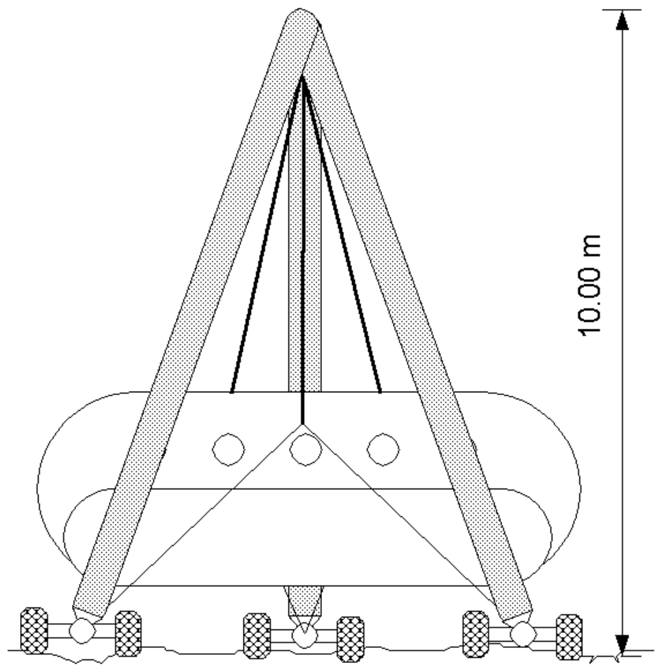

Figure 3. A conceptual drawing of a large construction robot manipulating a pressurized habitat module. *The robot consists of three inflated cylinders 50 cm in diameter and 10 m long. Each of the three cylinders is supported by a two-wheeled (or two tracked) crawler vehicle. Three heavy winches at the apex control three heavy lifting cables to lift and position the load in roll and pitch. Three lighter winches control three tag lines to provide fine positioning in x, y, and yaw. A fourth tag line winch maintains tension in all cables.*

This large construction robot would typically be assisted by one or more smaller robots of the design shown in Figures 1 and 2. These smaller robots would be outfitted with grippers and small tools to make electrical and hydraulic connections, open and close valves, and perform joining operations, such as running nuts, closing snaps, and attaching and removing lifting cables to and from heavy loads.

The large construction robot could also be fitted with a work platform so that it could manipulate heavy tools such as a trenching machine, jackhammer, backhoe, or clamshell for mining ore.

Our lab at the National Institute of Standards and Technology (NIST) has had considerable experience with large robots using cable suspended work platforms such as are illustrated in Figures 1 – 3. Laboratory versions of the NIST RoboCrane[1] developed by NIST for DARPA, Navy, and Marine Corps have demonstrated capabilities for welding, grinding, sawing, and milling operations; as well as lifting, positioning, and fixturing large heavy objects for assembly. RoboCrane technology commercialized for the U.S. Air Force has demonstrated the capability of lightweight cable-driven robots to lift and precisely position work platforms, and manipulate tools for paint stripping, inspection, repair, and repainting large aircraft.

IV. Control Issues

Round trip communication time delays between the Earth and Mars (about 20 min), and even between Earth and the Moon (about 3 s) impose severe limitations on performance of robot tasks that require significant operator interaction. The current approach for exploration on Mars is to limit the speed of locomotion to very slow speeds (about 1 cm/s.) This is not feasible for long-range exploration of hundreds or thousands of kilometers.

Future exploration missions will require that robots be capable of performing lengthy traverses and executing complex exploration tasks without direct human involvement. To support average speeds of 5km/h to 10 km/h (50 to 100 km per 10 h day), intelligent controllers must be provided that support autonomous mobility such that human intervention is seldom required. Future exploration robots should be able to collect data and transmit results more or less continuously, without waiting for instructions from a human operator between each task segment.

To achieve this level of autonomy, planetary exploration and construction robots must be endowed with considerable on-board intelligence. They must have sensors and sensory processing algorithms capable of perceiving the geometry of the local environment, and building geometric models of the terrain and objects with which that the robot must interact. They must be able to perceive and understand their environment, and have cognitive capabilities that allow them to reason about their situation, and plan behavior that is both effective and efficient. They must have sufficient knowledge, skills, and abilities to perform lengthy and complex tasks of locomotion and manipulation without direct supervision.

Exploration robots must have the ability to plan safe paths through difficult terrain, and to select trajectories that minimize energy and maximize exposure of solar panels to solar radiation. They must be able to recognize geological features of interest, and to choose optimal ways to approach and analyze those features. Construction and maintenance robots must be capable of autonomously planning and executing digging and grading tasks necessary for site preparation and mining of ore. They must have sufficient knowledge, skills, and abilities to autonomously perform lifting, positioning, mating, and joining tasks required for assembly and repair of base facilities. Robot control systems must be designed with sufficient intelligence to operate for lengthy periods without human intervention. Interventions by human operators must be infrequent and brief.

V. The 4D/RCS Reference Model Architecture

4D/RCS is one example of a reference model architecture that can be used to build intelligent control system with these capabilities. 4D/RCS is the most recent version of the Real-time Control System (RCS) that has evolved over the past 30 years at NIST and elsewhere for the design of controllers for a wide variety of intelligent systems.[2,3] Applications include industrial and laboratory robots, machine tools, coordinate measuring machines, automated factories, general mail and stamp distribution systems, automated mining systems, undersea vehicles, autonomous ground vehicles, and the NIST RoboCrane. 4D/RCS combines the RCS architecture with the 4-D approach for high speed autonomous driving developed by Dickmanns[4] at the Universitat der Bundeswehr in Munich. 4D/RCS was developed for the Army Research Laboratory Demo III program[5] and has subsequently been adopted by the Army Tank and Automotive Research and Development Engineering Center for their Vetronics Technology Integration Program, and by the Army Future Combat System (FCS) Program for the Autonomous Navigation System that will be deployed on all FCS ground vehicles, both manned and unmanned. 4D/RCS provides a theoretical foundation for designing, engineering, integrating, and testing software for intelligent unmanned vehicle systems.

The 4D/RCS reference model architecture consists of a multi-layered multi-resolutional hierarchy of computational nodes each containing elements of sensory processing (SP), world modeling (WM), value judgment (VJ), and behavior generation (BG). A block diagram of the 4D/RCS architecture is shown in Figure 4. At the lower levels, these nodes generate goal-seeking reactive behavior. At higher levels, they enable goal-defining deliberative behavior. Throughout the hierarchy, interaction between SP, WM, VJ, and BG give rise to perception, world modeling, decision making, planning, and reasoning.

At low levels, range in space and time is short and resolution is high. At high levels, distance and time are long and resolution is low. This enables high-precision fast-action response over short intervals of time and space at low levels, while long-range plans and abstract concepts are being formulated over broad regions of time and space at high levels.

4D/RCS closes feedback loops at every level. SP processes focus attention (i.e., window regions of space or time), group (i.e., segment regions into entities), compute entity attributes, estimate entity state, and assign entities to classes at every level. WM processes maintain a rich and dynamic database of knowledge about the world in the form of images, maps, entities, events, and relationships at every level. Other WM processes use that knowledge to generate estimates and predictions that support perception, reasoning, and planning at every level.

The 4D/RCS nodes have internal structure such as shown in Figure 5. Within each node there typically are four functional elements or processes: 1) behavior generation, 2) world modeling, 3) sensory processing, and 4) value judgment. There is also a knowledge database that represents the node's best estimate of the state of the world at the range and resolution that are appropriate for the behavioral decisions that are the responsibility of that node. A methodology for designing 4D/RCS computational nodes and establishing communications between them has been published.

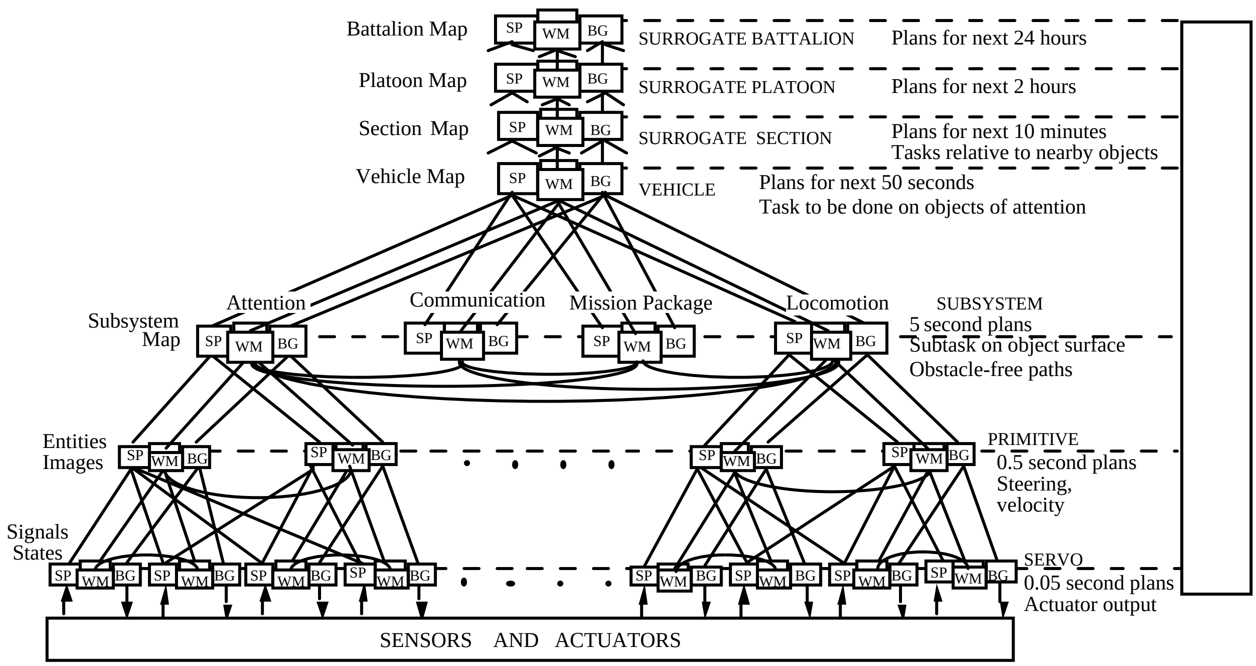

Figure 4. A 4D/RCS reference model architecture for an individual vehicle. *Processing nodes are organized such that the BG modules form a command tree. Information in the knowledge database is shared between WM modules in nodes above, below, and at the same level within the same sub tree. On the right, are examples of the functional characteristics of the BG modules at each level. On the left, are examples of the data types and maps maintained by the WM in the knowledge database at each level. Sensory data paths flowing up the hierarchy typically form a graph, not a tree. VJ modules are hidden behind WM modules. An operator interface provides input to, and output from, modules in every node. A feedback control loop is typically closed through every node.*

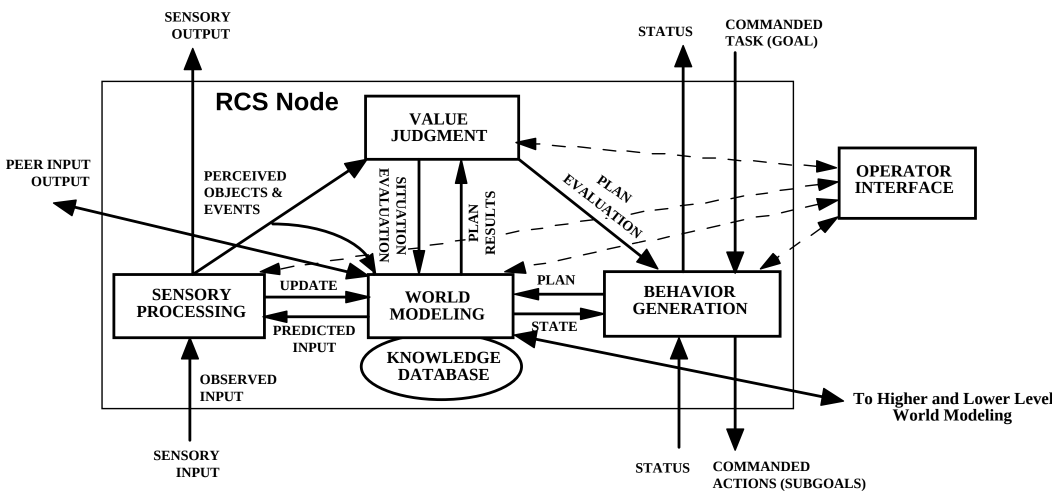

Figure 5. Internal structure of a typical 4D/RCS node. *The functional elements within a RCS node are behavior generation, sensory processing, world modeling, and value judgment. These are supported by a knowledge database. Each functional element in the node may have an operator interface. The connections to the Operator Interface enable a human operator to input commands, to override or modify system behavior, to perform various types of teleoperation, to switch control modes (e.g., automatic, teleoperation, single step, pause,) and to observe the values of state variables, images, maps, and entity attributes. The Operator Interface can also be used for programming, debugging, and maintenance.*

The 4D/RCS is a hybrid architecture that is both deliberative (with the hierarchical ability to plan) and reactive (with the reflexive ability to respond rapidly to exigencies). At every level within the BG hierarchy, planning processes receive goals and priorities from superiors and deliberatively decomposes those goals into subgoals with priorities and timing requirements for subordinates at levels below. Also at every level, executor processes close a feedback loop that produces reactive behavior.

At every level within the SP hierarchy, perception processes focus attention, segment, filter, and classify information derived from subordinate levels. Events are detected, objects recognized, situations analyzed, and status reported to superiors at the next higher level.

At every level, sensory processing and behavior generation processes have access to a model of the world that is resident in a knowledge database. This world model enables the intelligent system to analyze the past, plan for the future, and perceive sensory information in the context of expectations and task requirements.

At every level, a set of cost functions enable value judgments that determine priorities that guide intelligent decision making, planning, and situation analysis. This provides a robust form of value driven behavior that can function effectively in an environment filled with uncertainties and unexpected events.

As a rule of thumb, at each successively higher level, range in space and time increase by about an order of magnitude, accompanied by an order of magnitude decrease in resolution. For example, plans at the Servo level may have a time horizon of 50 ms, and output from the Servo level may be updated every 5 ms. Plans at the Primitive level may have a time horizon of 500 ms. Output from the Primitive level may be updated on 50 ms intervals, and plans may be recomputed as often as every 50 ms. At the Subsystem level, plans extend to a 5 s time horizon. Output is updated and plans recomputed every 500 ms. At the Vehicle level, plans extend to a 1 min time horizon. Output is updated and plans recomputed every 5 s. At the Section level, plans extend to a 10 min time horizon. At each successively higher level, plans extend further into the future, and are recomputed less frequently. At Platoon, Company, or Battalion, planning horizons may extend hours or even days into the future.

VI. Test Results

Experiments with eXperimental Unmanned Vehicles (XUVs) sponsored by the Army Research Laboratory have demonstrated the ability of robot vehicles with 4D/RCS controllers to navigate autonomously over rough terrain and to negotiate steep slopes and erosion features at an average speed of more than 5 km/h.

A series of tests were conducted during the winter of 2002-2003 on three types of terrain: rolling/arid, rolling/vegetated, and urban.[6] The rolling/arid tests were conducted at Tooele Army Depot, Utah on open terrain with sagebrush up to 1 m high, deep gullies, erosion features, rocks, and breastworks of an abandoned dam with steep slopes. The rolling/vegetated tests were conducted on maneuver ranges at Ft. Indiantown Gap, Pennsylvania through woods of varying density, open terrain with tall grass and brush, dirt roads, creeks, lakes, and trails with constriction points such as bridges. Urban tests were conducted among Army barracks at Ft. Indiantown Gap where there are buildings, phone poles, vending machines, parked cars, dumpsters, and culverts. For the tests, rubble was dumped on the streets, along with abandoned cars, and human mannequins.

Routes for XUV missions were laid out on a terrain map by trained Army scouts using aerial photographs. This information was conveyed to the XUVs in the form of GPS waypoints spaced about 100 m apart. Two test courses were laid out at each test site with different levels of terrain complexity. Test missions of three different lengths, 500, 1000, and 2000 m were defined, and randomized over start points and various other parameters. Two teams ran a total of 144 missions over a period of six days at each of the test sites.

The tests were designed to test the XUVs ability to traverse cross-country and through urban areas. The XUVs operated completely autonomously until they got into trouble and called for help. At that point, an operator was called in to teleoperate the vehicle out of difficulty. During these teleoperations, data was collected on the cause of the difficulty, the type of operator intervention required to extract the XUV, the work load on the operator, and the time required before the XUV was returned to autonomous mode. Typical reasons for calling for help were that the XUV was unable to proceed because of some terrain condition or obstacle (such as soft sand on a steep slope, or dense woods), and was unable to find an acceptable path plan after several attempts at backing up and heading a different direction.

During three major experiments designed to determine the technology readiness of autonomous driving, the Demo III experimental unmanned vehicles were driven a total of 550 km, over rough terrain: 1) in the desert; 2) in the woods, through rolling fields of weeds and tall grass, and on dirt roads and trails; and 3) through an urban environment with narrow streets cluttered with rubble, parked cars, dumpsters, culverts, telephone poles, and manikins. Tests were conducted under various conditions including night, day, clear weather, rain, and falling snow. The unmanned vehicles operated over 90 % of both time and distance without any operator assistance. A detailed report of these experiments has been published[6], along with high resolution ground truth data describing the terrain where the XUVs experienced difficulties.[7] More recent results have demonstrated the ability of XUVs to traverse distances of more than 40 km through difficult terrain, including rolling hills, dirt roads and trails, woods and fields without any human intervention.

It should be noted that all of the above tests were performed in environments devoid of moving objects such as on-coming traffic, pedestrians, or other vehicles. It should also be noted that only the lower three levels of the 4D/RCS architecture were involved in these tests. All of the higher level capabilities were provided by the human operator, and by subject matter experts.

Overcoming these limitations is the primary focus of current NIST research. We have analyzed the set of tasks required for driving in normal traffic, and begun to catalog the knowledge required for autonomous driving. The Army Research Lab Robotics Technology Alliance, NIST, and a number of private contractors are currently developing sensors and algorithms for autonomous driving on normal roads and streets, e.g., driving on country roads and city streets with on-coming traffic, negotiating intersections with traffic signals and pedestrians, and maneuvering in and out of parking spaces. Researchers are analyzing the sensing capabilities and perception algorithms necessary to detect and track moving objects under various conditions in normal traffic. Ontologies and tools are being designed for representing objects, events, situations, and relationships that are required for competent driving behavior.[8] ARL sponsored research at NIST has recently been focused on tactical behaviors for teams of real and virtual, manned and unmanned, military ground and air vehicles.

VII. Summary and Conclusions

This paper has presented conceptual designs for two classes of future planetary robots: one for long-range exploration, the other for construction and heavy lifting. Both of these designs are based on inflatable technology. This enables robots that are both large and lightweight, and are easily packaged in a small volume for transportation to, and landing on, lunar or planetary surfaces. The use of cables and computer-controlled winches to build lightweight robots with large working volumes is based on the NIST RoboCrane.

A description is given of the 4D/RCS reference model architecture for intelligent vehicle systems. 4D/RCS incorporates and integrates many different concepts and approaches to robot control systems into a harmonious whole. It is hierarchical but distributed, deliberative yet reactive. It spans the space between the cognitive and reflexive, between planning and feedback control. It bridges the gap between spatial distances ranging from kilometers to millimeters, and between time intervals ranging from months to milliseconds. And it does so in small regular steps, each of which can be easily understood and readily accomplished through well known computational processes.

4D/RCS was developed for the Army Research Lab autonomous ground vehicle program and has been adopted by the Army Future Combat System for Autonomous Navigation Systems. The underlying technology, RCS, is well documented and has been used for a wide variety of applications over the past three decades. Software development tools are commercially available. The 4D/RCS version 2.0 reference document[2] and two introductory textbooks[3,9] are also available as well as a number of other publications. Information about 4D/RCS can be accessed over the Internet at http://www.isd.mel.nist.gov/projects/rcs/

Acknowledgments

This work was partly supported by the Army Research Laboratory, Charles Shoemaker, Program Manager. It was also supported by the Intelligent Systems Division of NIST through funding for Research and Engineering of Intelligent Systems.

References

[1]Albus, J.S., Bostelman, R., and Dagalakis, N.G., "The NIST ROBOCRANE," *Journal of Robotic Systems*, Vol. 10, No. 5, 1993, pp. 709-724.

[2]Albus, J.S., et.al. "4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems," *NISTIR 6910*, National Institute of Standards and Technology, Gaithersburg, MD, 2002.

[3]Albus, J.S. and Meystel, A.M., *Engineering of Mind: An Introduction to the Science of Intelligent Systems*. New York, NY: John Wiley & Sons, Inc., 2001.

[4]Dickmanns, E.D., "A general dynamic vision architecture for UGV and UAV," *Jounral of Applied Intelligence*, Vol. 2, 1992, pp. 251-270.

[5]Shoemaker, C., Bornstein, J., Myers, S., and Brendle, B. "Demo III: Department of Defense testbed for unmanned ground mobility," *Proceedings SPIE Conference on Unmanned Ground Vehicle Technology*, SPIE Vol. 3693, Orlando, FA, April, 1999.

[6]Camden, R., Bodt, B., Schipani, S., Bornstein, J., Runyon, T., French, F., Shoemaker, C., Jacoff, A., and Lytle, A., "Autonomous Mobility Technology Assessment Final Report," Army Research Laboratory Technical Report – 3471, Aberdeen Proving Ground, MD, April 2005.

[7]Witzgall, C., Cheok, G.S., Gilsinn, D.E., "Terrain Characterization from Ground-Based LADAR," *Proceedings of PerMIS '03 Workshop*, National Institute of Standards and Technology, Gaithersburg, MD, 2003

[8]Barbera, A., Albus, J., Messina, E., Schlenoff, C., and Horst, J., "How Task Analysis Can Be Used to Derive and Organize the Knowledge For the Control of Autonomous Vehicles," *The Robotics and Autonomous Systems Journal: Special Issue on the 2004 AAAI Knowledge Representation and Ontologies for Autonomous Systems Spring Symposium*, 2004.

[9]Gazi, V., Moore, M.L., Passino, K.M., Shackleford, W.P., Proctor, F.M., and Albus, J.S., *The RCS Handbook: Tools for Real Time Control Systems Software Development*, John Wiley & Sons, NY, 2001