

# A Framework for Distributed Rover Control and Three Sample Applications

Steve McGuire  
Autonomy and Robotics Area  
NASA Ames Research Center, MS 269-3  
Moffett Field, CA. 94035-1000

## Abstract

*In order to develop quality control software for multiple robots, a common interface is required. By developing components in a modular fashion with well-defined boundaries, roboticists can write code to program a generic rover, and only require very simple modifications to run on any robot with a properly implemented framework. The proposed framework advances a GenericRover that could be any rover, from Real World Interface's All Terrain Robot Vehicle Jr. series, to the Fido-class rovers from the Jet Propulsion Laboratory, to any other research robot. Using these generic hardware interfaces, software designers and engineers can concentrate on the actual code, and not have to worry about hardware details. In addition to the hardware support framework, 3 sample applications have been developed to demonstrate the flexibility and extensibility of the framework.*

## 1 Introduction

Every robotic rover architecture relies on some distributed mechanism to share rover data and to issue commands. However, these independent architectures are very hardware dependent and are not easily extensible to different rover platforms. The proposed framework allows high-level software to be written to support a reference rover. Through the use of well-defined features and interfaces, high-level controllers need not be concerned with hardware dependencies. To support this framework, a foundation layer is required to transform high-level commands into the required hardware interface for the rover.

The proposed framework was designed to support Real World Interface's (RWI) All Terrain Robot Vehicle (ATRV)-Jr. robot, but could be easily extended to any control interface. Every detail of the rover is hierarchically descended from a basic rover object. Follow-

ing previous work in the group, commands may be sent to the physical hardware via standardized messages; for convenience, the K9 rover messages were used as templates. In the case of the ATRV-Jr., the hardware framework is responsible for translating these high-level messages into Common Object Request Broker Architecture (CORBA) operations. CORBA is used as an object transport by default; the commercially supported control system aboard the ATRV-Jr. communicates via CORBA. By using asynchronous command queues to process rover requests, most framework calls return to their calling function very quickly, generally only requiring a memory copy.

Distributed operation is handled transparently by the CORBA object transports, concurrently allowing multiple disparate clients to issue requests to the rover and observe the state. Much of the work of the base class involves translating robot data to appropriately standardized common output values. For example, the CORBA components for the Inertial Measurement Unit (IMU) provide one integration step of acceleration data; the framework then interpolates the second required integration to provide a rough estimate of position.

## 2 Hardware Support

### 2.1 Hardware to CORBA Interface

Because of the modular design of CORBA, certain components can be replaced with higher-performance implementations. In the ATRV Jr., the only components that are not driven by serial port are the stereo framgrabbers, which connect via two Peripheral Component Interconnect (PCI) bus slots. These framgrabbers are commercially available Hauppauge WinTV NTSC cards that utilize a Brooktree 878 chipset. The original code provided by RWI used a single-threaded, single buffering capture scheme which

as controls on rotation and translation. The results of [5] may be useful to integrate into the base controller, as the sensors available on K9 are very similar to the sensors of the ATRV-Jr.

### 3.2 Resolution

The calibration setup was needed to correct for a problem that was only observed in field tests that related to the difference between a turn command and the actual result of the command. These commands were given by hand and thus were directly observable; the same class of commands is used by the sample tracker described below and thus need to become more accurate. The model of the base drivetrain that is exposed by RWT's software does not expose raw wheel counts, and appeared calibrated for non-ablative surfaces such as solid floors. On such surfaces, turns are much more accurate than for turns on gravel.

A basic calibration is proposed to determine linear coefficients of the rotational and translational velocities to characterize a surface's requirements for accurate motion. Also, because the work of [5] is available for integration into the base controller, a combination approach might give the best results.

## 4 Sample Applications: Utilities

### 4.1 *grover*: A Rover Control Utility

*grover* was designed to exercise all of the possible functions of all of the rover extensibilities. The command line interface is suited to directing the robot through a series of low-level operations. In essence, the purpose of this application is a test and utility driver that can issue specific commands for adjusting the rover's state. For example, this utility was used in the proving tests to issue specific drive commands and to adjust the pantilt head to specific values. Also, *grover* was used to capture imagery used by E. Bandari and E. Ricks of the Autonomy and Robotics Area (ARA) to characterize the vision system aboard the rover.

### 4.2 *panovision*: A Telemetry and Imaging Utility

To demonstrate the flexibility of the framework, and to support other researchers in the Autonomy and Robotics Area, *panovision* was implemented. A commercially available camera was placed into an optics package manufactured by Carnegie Mellon University



Figure 2: CMU Omnicam aboard ATRV-Jr.

(CMU) to provide a full 360 degree field of view. The complete enclosure was then mounted on the ATRV-Jr using a custom manufactured mount. (Figure 2)

The chosen camera is a Sony EVI-370D block camera that has a serial interface to adjust camera parameters and a NTSC video output. By reviewing technical material provided by Sony and modifying source code provided by CMU[4], a library of camera control functions was developed. Some functions include adjustable zoom, focus control, and exposure control.

#### 4.2.1 Telemetry Control

To facilitate other research, rover telemetry had to be captured by interrogating various hardware components within a frame of reference. This hardware requirement was accomplished by using the developed framework. In addition, assuming a proper implementation was available for another rover, only superficial modifications would be necessary to run the image capture and telemetry capture on another platform.

To adjust for pitching terrain, the onboard pitch sensor is utilized. In the framework, a relative tilt is defined to be the tilt angle of the pan-tilt head with respect to the rover body; an apparent tilt is defined to be the angle between the optical axis of the camera and the surrounding terrain, using a planar assumption. This apparent tilt can be calculated by summing the reading from the pitch sensor and the tilt reading from the pan-tilt head. By using the readings from the pitch sensor, the rover can compensate for uneven terrain and correct artificially high or low relative tilt angles.

### 5.3 Finish Conditions

Once the tracker could analyze an image and report on the best possible match of the target, the framework was utilized to center the target in the field of view and to issue a drive command to the base. In order to determine when the rover has actually reached a target, the measure of the apparent tilt angle defined above is used. Presumably, when a target is within range of an on-board instrument, the pan-tilt will be "looking" downward beyond a specific, predefined angle.

### 5.4 Role of Subcapture

Using the subcapture capability described above, the performance of the tracking algorithm was increased by limiting the search space. In every iteration, the rover seeks to minimize the absolute pan angle and keep the target centered in the image plane. Once the rover has met these objectives, the possible search space is limited to one quarter of the size of the initial space to increase tracking performance.

Due to the distributed nature of the vision system, a possible conflict exists where a request to change the subimage may occur, but does not take effect instantaneously. The result of such a conflict is that the returned imagery is from a different section of the image than the specified coordinates requested. To prevent this conflict, the vision object pauses the tracker until the requested image coordinates are available. The pause operation is transparent to the visual servoing client and is provided by the particular framework implementation. This operation is done at the lowest possible level to reduce the communications bandwidth required by imaging operations.

### 5.5 Disparity and Prediction Measure

Because the rover is moving with nearly constant translational speed, a further enhancement of the tracker would be to predict the location of the target in the next frame. By assuming a linear model of movement, the target disparity, or the distance from the target to the center point of the image, is measured and recorded. Using a simple approximation, the next location of the target in the source image is predicted. Due to insufficient resources, this information is calculated and stored for further analysis, but not integrated into the tracker control loop.

### 5.6 Tracker Performance

An operator selects a target using an initial image at the full working resolution of the camera. The tracker loop then drives the rover towards the target, logging an image during every iteration. During full resolution operation, the tracker loop can run at approximately 2 Hz. During subimage operation, the loop speed increases to approximately 5 Hz.

#### 5.6.1 Laboratory

In laboratory tests, the tracker's performance is capable of driving the rover at speeds of up to 20 cm/s. See Figures 3, 4, and 5 for starting, intermediate, and ending points. The operator-designated target is symbolized by a crosshair in 3; the rover's estimates are in 4 and 5. The performance of the tracker in the lab setting is sufficiently good to warrant further investigation; however, the laboratory setting makes use of controlled lighting and a well-defined textured target, the phone directory of the Center.



Figure 3: Lab Test Begin




Figure 3: Field Test Complete

an extensible control and reporting interface for any type of rover that has the same logical functions. For example, most rovers have capabilities such as moving a pan-tilt head, moving the drivetrain and chassis, and retrieving current rover state.

Any software which needs to have access to rover control need only host a framework object, allowing user interface development to become independent from object transport development. Essentially, the framework establishes a local set of objects that proxy requests to the rover, hiding the underlying object transport from the developer. Through the use of such abstractions, any rover could be supported in a transparent fashion.

In addition, by using CORBA over Transmission Control Protocol/Internet Protocol (TCP/IP) as a distributed object transport, any operation can be initiated by any machine that has IP connectivity to the rover. As a side benefit, several CORBA implementations are freely available, negating the direct or hidden licensing costs associated with other distributed object models.

Many further developments are envisioned. For example, a generic simulation rover that conforms to the framework and provides mockup data and telemetry would be useful for testing rover control software offline, in a device independent fashion. Control software under development can support any rover that has basic functions; components that require specialized control such as arm control or other device control can simply extend the framework in the appropriate direction for the appropriate hardware platform. In this case, only rover-specific code would need to be written; core functionality would be provided by the framework and the framework's underlying hardware

implementation.

## Acknowledgments

Many thanks to M. Bualat, R. Washington, M. Deans, A. Wright, L. Edwards, K. Bass, M. Fair, E. Ricks, E. Bandari, and the other members of the Autonomy and Robotics Area of Code IC for supporting this work.

## References

- [1] A. Schiffler, [aschiffler@home.com](mailto:aschiffler@home.com) "libbgrab" <http://www.ferzkopp.net/Software/libbgrab/>
- [2] J. Borenstein, L. Feng, "UMBMark - A Method for Measuring, Comparing, and Correcting Deadreckoning Errors in Mobile Robots", UM-MEAM-94-22, University of Michigan 1994
- [3] D. Wettergreen, H. Thomas, M. Bualat. "Initial Results from Vision Based Control of the Ames Marsokhod Rover". *Proceedings IROS '97*, pp. 1377-1382.
- [4] J. McMahaill, [jmcm@frc.ri.cmu.edu](mailto:jmcm@frc.ri.cmu.edu) "VISCA Camera Interface", Robotics Institute, Carnegie Mellon University, 1999.
- [5] R. Xu, [rxu@andrew.cmu.edu](mailto:rxu@andrew.cmu.edu) "State Estimation On K9", NASA Ames Research Center, 2001.