

Reality Browsing: Using Information Interaction and Robotic Autonomy for Planetary Exploration

Peter W. Coppin¹, Michael D. Wagner² and Scott Thayer²

¹*STUDIO for Creative Inquiry / Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213*

²*Field Robotics Center, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213
(412) 268-1565, coppin@cmu.edu*

Abstract. Reality browsing is a framework that enables distributed control of a team of planetary robots. In it, prioritized user queries are serviced in a hierarchical data structure consisting of an Internet-accessible world model, data archives on the remote robots and finally a multiple-robot planner that coordinates query-directed searches. This paper introduces the reality browser concept and outlines important research issues required for implementation.

INTRODUCTION

Recently, a NASA goal of distributed spacecraft control has expanded the network of scientists that can participate in planetary exploration missions (Backes, 2000). As planetary rover missions increase in duration from weeks to months, it becomes critical for scientists to collaborate from around the world. Additionally, with the advent of the Mars Pathfinder mission the Internet has been increasingly used as a means of public participation in these missions. This paper presents a framework for interaction with planetary exploration missions that seamlessly supports scientific investigation, education and public participation through the use of *reality browsing*, a search framework that gathers scientific data in immersive models of remote sites, onboard autonomous robots and in the planetary environment around these robots. A reality browser will consist of updated, 3D representations of remote sites coupled with an infrastructure for the communication of new search criteria to remote autonomous explorers. The paradigm traditional browser interface paradigm will be extended beyond searching online databases and into the physical realm of remote exploration. One or more autonomous robots will be directed by the science goals of users. A planner will optimize scientific return while meeting mission resource constraints and user priority levels. This paper details the reality browser concept and its enabling technologies.

CONTROL AND INTERACTION MODES

A *reality browser* enables comprehensive interaction with robot missions for scientists and lay persons alike. Users will request new scientific information to meet their search criteria, for instance, images of sediment deposits or igneous rocks. While this approach is somewhat analogous to a web-based search engine, queries will not be limited to plain text, but based on multiple sensory realms such as color, size or spectral signatures.

Searches are performed at various levels of data repositories (Fig. 1), beginning with the world model available on the Internet. This mode is relatively easy to perform and does not require control of physical devices. The amount of information in the world model increases as the mission progresses; therefore as planetary exploration missions increase in duration, world models will become very rich. The user is presented with the results of this search and may choose to perform searches deeper into the system – into the memory onboard the robots operating in the remote environment. The term “robot” here includes orbiters, rovers and any other vehicles. This type of search will result in a downlink from robots with appropriate query results, which can be time consuming and costly as bandwidth is limited on systems such as NASA’s Deep Space Network. Two types of searches may then be requested: *increase*

data detail or increase search coverage. A command to increase data detail entails gathering new sensor readings of different type or higher resolution. Increasing data coverage requires searching additional area to detect previously unknown data points.




FIGURE 1. The Reality Browser Control Architecture Gathers Information Requests from Multiple User Interfaces (A). Queries are first serviced in the world model (B). If further searches seek to engage in risky or costly tasks, approval from a human user is required (C). If the desired information is unavailable in the robot memory domain, robots must execute plans to search the world around them (D). Agents send any search results to the user interface (E), until new searches are commanded.

Users may command deeper searches into the physical domain. This is the most costly type of search, as it involves operating robots. Robots will generally be in different locations and be equipped with different sensor suites and mobility platforms. A multiple-robot planner oversees the robot team, assigning search tasks to individual units to achieve search goals. Individual robots have sufficient autonomy to allow recognition of objects matching users' queries.

DATA STORAGE

Search results may source from one of three locations: the world model, robot memory or the physical domain.

- **World Model.** Here sensory data from the remote robots has been transmitted back to Earth and integrated into a world model that is Internet-accessible. This world model fuses multiple types of sensory data in a spatio-temporal map. Multiple sensing modalities are represented in the world model, satellite, panoramic and microscopic imagery, spectrometer readings and so on. Therefore, presenting these data to a reality browser user does not require long-distance communication or physical action.
- **Robot Memory.** Each robot has its own local archive of sensory information. These data have been acquired by an individual robot but have not yet been transmitted back to Earth. When robot memory is searched, requests are transmitted to the entire robot team. Those robots with relevant information reply, those without take no action. The software onboard each robot is responsible for discriminating relevant information to minimize interplanetary communication requirements.
- **Physical Domain.** Data that has not yet been sensed is not represented locally on any robot or in the world model. Acquiring these data for presentation to a user therefore requires a plan of robotic action. Here data understanding processes must contend with more limited data and lead to optimized plans in real-time. This ability is referred to as *science autonomy* and is comprised of classification and planning components (Apostolopoulos, 2000).

Data access time varies dramatically based on location. An analogy to modern computer caching systems can be drawn here; methods from this field of research may apply minimizing or optimizing accesses to costly data storage sources.

USER PRIORITIES

A wide variety of users can benefit from a reality browser interface, from the general public to students to scientists and principal investigators. With limited-duration and resource missions, a strict set of priorities must be enforced to ensure that mission-critical users, such as principle investigators, may retrieve the data whenever necessary.

Our approach defines priorities by access data location. Low-priority users may only query information in the world model, costing the mission few resources other than web server CPU cycles. Medium-priority users may search data located on the robots' archives. These users can therefore request that information be transmitted back to Earth, which can be an expensive and time-consuming proposition. Users of the highest priority can make requests that result in robotic action, for instance, deployment of sensors on a known object to a search for new objects – the most resource expensive and risky operation.

PLANNING ACTIONS FOR THE PHYSICAL DOMAIN

A team of cooperating robots with heterogeneous sensor modalities best carries out a wide range of queries from many users. For instance, an orbiter may be most useful for identifying large-scale regions of interest (dried lake beds, alluvial fans, etc.) while clearly rovers are most effective at *in situ*, close-range sensing.

Queries that require interplanetary communication or robotic action must be translated into feasible robot command schedules. The reality browser system must therefore contain a multiple-robot planner that optimizes the number of serviceable queries over any given amount of time. For instance, it could schedule high-resolution imagery of a sediment layer on a canyon wall (a command to increase detail) just before a search for sedimentary rocks in a dried riverbed inside that canyon (a command to increase coverage). Additionally, robots will have limited time, energy, communications and data storage resources, so some queries may not be serviceable due to conflicts with higher-priority users.

Operationally, human mission engineers will verify the output of the planner. This will be feasible because while plans could represent the goals of an immense number of users, the plans themselves should much more distilled and therefore be verifiable by a team of people.

RESEARCH ISSUES

Implementation of a reality browser framework will depend on wide-ranging research efforts, spanning the fields of human/computer interaction, multiple-robot planning and robot autonomy. Four issues in particular must be resolved: representing queries to the system, enabling users to interact with the world model, creating plans for multiple robots and imbuing robots with sufficient autonomy to perform search tasks.

Representing Queries

Users may make very disparate requests such as looking for rock types, weathering features, terrain types and so on. Criteria that characterize items of interest may be very specific and range across multiple sensing modalities. State of the art in robotic exploration involves careful, manual encoding of these criteria into some type of machine learning framework. A reality browser framework must be far more flexible and easy to use, yet complete enough to allow a robot to probably and correctly learn what the user is attempting to find. Two approaches may be taken to address this issue.

In terms of machine learning, the task at hand involves generating a hypothesis that discriminates relevant from irrelevant data. Many machine-learning algorithms depend upon comprehensive examples and / or a priori information from expert sources to generate hypotheses. Generally, system developers train a machine learning systems. However, in the reality browser the user plays this role. Research must investigate how to make this task as intuitive as possible. Providing appropriate training examples may be a time consuming process for the user, who may prefer to explicitly express criteria themselves.

Defining query representation is also tied to the classification method utilized by the robotic team. A promising approach lies with Bayes network classification. Bayes networks support the encoding of prior information in an intuitive, graphical structure that may provide a natural query representation. Other advantages of Bayes networks for this purpose are given in (Pedersen, 2000):

- **Learning from examples.** This allows the user to both specify prior knowledge and provide specific training examples, for instance from a new area of planetary terrain that may contain previously unknown geology. Classification can therefore become more effective as a mission progresses. However, an intolerably large number of training examples may be required to train the system.
- **Using compound evidence.** Both individually and as a team, robots will have multiple sensing modalities. It is important to utilize these modalities collectively to increase performance.
- **Handling noisy or incomplete data.** This is obviously an important feature when used with autonomously, robotically collected data.

A second method by which queries could be represented is with a sort of grammar of explicit criteria. A descriptive search engine interface can be understood by analogy to text based search engines. Boolean functions placed in additive or subtractive sequence provide a description of the sought objects or features. On the web, the descriptors are text-based; in the reality search engine the descriptors are visual. The grammar will enable scientists to describe features/ objects in relation to size, sharpness, color, orientation and more. For example, if searching for meteorites, a user could specify a search for objects less than 10 cm, and a sharpness or color of a given value, spectral characteristics and more (Fig. 2).




FIGURE 2. Simplified Search Criteria to Find a Characteristic Small, Dark and Round Object such as a Meteorite.

Information Interaction

A reality browser user interface will follow information interaction principles outlined by the EventScope project at Carnegie Mellon University (see <http://www.eventscope.org>), and its predecessor, the Big Signal Project. Data retrieved for searches must be presented in a manner that enables comparison to known concepts and an understanding of spatial and temporal relationships to promote learning and research. Therefore, rather than simply returning simple data points, the reality browser will update the world model with the new data and present a view of the model focused on the relevant results. The primary aspects of such an interface include:

- **Spatial / Chronological Maps.** Data are gathered by robots will be assembled into both spatial maps and associated timelines. These evolve throughout the mission, allowing users to uncover relationships in spatial dimensions while maintaining temporal relationships. Maps will come from three-dimensional range information from rovers such as (Maimone, 1998; Stoker, 1999), but other important sources include high-resolution panoramic imagery.
- **Context Libraries.** The “rock library” of Big Signal Antarctica 2000 (Coppin, 2000) allowed users with little knowledge of geology to understand meteorite and rock data retrieved by an autonomous robot in Antarctica. It provided examples of various rocks that could be found in these regions, allowing middle school students to classify new rocks found. Analogous *context libraries* will be available to provide background information on future planetary exploration missions. Each mission’s context library blends prior knowledge with a user interface designed to make new discoveries for themselves. Flexible, user-modifiable context libraries may prove an important source for training examples for queries. Finally, more knowledgeable users may also use the context libraries as a reference when explaining concepts to students or the public, for instance, in the form of school curricula.

- **Flags.** Flags on the two- or three-dimensional maps described above will designate the presence of data types associated with points in space. Flags also enable collaboration between users, who will place flags with comments or links to other information sources, such as pages on the World Wide Web.
- **Multiple Scales of Resolution.** Users should have access to many different scales from satellite imagery to microscopic. Tying small-scale features with large-scale features is a critical aspect of geologic inquiry.

Automated Planning of Multiple Robots

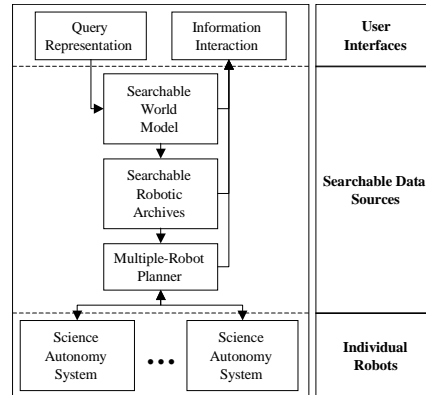


FIGURE 3. Information and Command Flow in the Reality Browser Framework.

Current automated mission planning tools such as ASPEN (Backes, 1999) are intended for use with single robots. Consequently, they contain path planners and databases embedded with sequence generation tools. Because the framework presented here includes the automatic planning of multiple robots, a monolithic planner must be split into a multiple-robot planner and task autonomy control systems onboard each robot, referred to as *science autonomy systems* (see section below). This planner fits into the reality browser architecture as shown in Figure 3. The multiple-robot planner takes as input a potentially large list of prioritized queries and then generates a sequence of tasks for each robot to execute. An important role of the multiple-robot planner is to service similar or overlapping queries with minimal expenditure of robot resources. It must also realize when queries cannot be serviced and report this to the appropriate users.

It is hypothesized that electoral models for structuring users' access to power will lead to more effective multiple robot control schemes than economic models (e.g., laissez-faire capitalism). This is because, as long as specialized robots continue to be custom-made instead of mass-produced, the supply of robots designed to service a given user's query cannot increase to meet the growing demand of users quickly enough to accomplish all tasks in a limited time.

Science Autonomy

Robots servicing user queries in this architecture must have sufficient task autonomy to discriminate sensor information meeting query criteria from irrelevant information. Serious constraints exist in the form of limited robot resources and communications bandwidth. An intelligent search must therefore be performed, complete with selective sensor deployment and data reduction techniques.

The science autonomy system (SAS) onboard the Nomad robot was developed to perform this type of search, hunting for meteorites in Antarctica (Wagner, 2000). Although the SAS was developed specifically for autonomous meteorite search, it is designed as a control architecture for a more general class of autonomous scientific exploration missions. These missions may require vastly different types of sensors, actuators and data understanding algorithms. Scientific goals in the SAS are framed as classification problems. Objects are classified as belonging to certain known exemplar classes. The goal of the SAS is therefore to classify as many objects in the world as possible. However, different sensors on a robot are deployed with varying time and energy costs. It is therefore beneficial to intelligently select sensor usage because most field robots operate with limited time and energy resources. To achieve intelligent sensor selection, the SAS utilizes the following approaches:

- Deploying sensors must be associated with some kind of a cost such as energy, time or digital storage space.
- The system must provide an estimate of the information that would be gained by deploying all sensors in the system. This quantity, referred to as *information gain*, is weighed against sensor deployment costs to create action plans.

CONCLUSIONS

Public excitement in past NASA missions has been immense despite the lack of any real interaction. World models have been created that have allowed students to experience Mars, but only as a byproduct of exploration directed by a select few people. Implementation of the reality browser concept in future planetary exploration missions would likely lead to increased participation amongst both mission scientists and the public around the world. With increased participation there will certainly be a wealth of new discoveries.

ACKNOWLEDGMENTS

The authors would like to thank the members of the EventScope and Robotic Antarctic Meteorite Search projects whose hard work has inspired the concepts presented here. Additionally we thank Dr. David Crown at the University of Pittsburgh for providing us with numerous invaluable insights.

REFERENCES

- Apostolopoulos D. S., Wagner, M. D., Shamah, B. N., Pedersen, L., Shillcutt, K., and Whittaker, W. L., "Technology and Field Demonstration of Robotic Search for Antarctic Meteorites," *to appear in the International Journal of Robotics Research 2000*.
- Backes, P. G., Rabineau, G., Tao, K. S., and Chien, S., "Automated Planning and Scheduling for Planetary Rover Distributed Operations," in proceedings of *1999 IEEE International Conference on Robotics and Automation*, Detroit, Michigan, 1999, pp. 984-991.
- Backes, P. G., and Norris, J. S., "Mars Rover Mission Distributed Operations," *TMOD Technology and Science Program News*, Issue 19, June 2000.
- Coppin, P., Pell, R., Wagner, M. D., Hayes, J. R., Li, J., Hall, L., Fischer, K., Hirschfield, D., and Whittaker, W. L., "EventScope: Amplifying Human Knowledge and Experience via Intelligent Robotic Systems and Information Interaction," *IEEE International Workshop on Robot-Human Interaction*, Osaka, Japan, September 2000.
- Maimone, M., Matthies, L., Osborn, J., Rollins, E., Teza, J., and Thayer, S., "A Photo-Realistic 3-D Mapping System for Extreme Nuclear Environments: Chernobyl," to appear in proceedings of *IEEE/RSJ International Conference on Intelligent Robotic Systems 1998*.
- Pedersen, L., Apostolopoulos, D., and Whittaker, W. L., "Bayes Networks on Ice: Robotic Search for Antarctic Meteorites," submitted to *Neural Information Processing Symposium 2000*, Denver, Colorado, 2000.
- Stoker, C., Zbinden, E., Blackmon, T., Nguyen, L., "Visualizing Mars Using Virtual Reality: A State of the Art Mapping Tool Used on Mars Pathfinder," presented at the Extraterrestrial Mapping Symposium: Mapping of Mars, ISPRS, Caltech, Pasadena, CA, July 1999.
- Wagner, M. D., Apostolopoulos, D., Shillcutt, K., Shamah, B., Simmons, R., and Whittaker, W., "The Science Autonomy System of the Nomad Robot," submitted to the *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2000.