Telepresence across the Ocean

Ioannis Rekleitis, Gregory Dudek, Yasmina Schoueri, Philippe Giguere, and Junaed Sattar
School of Computer Science, McGill University. Montreal, QC, Canada.
[yiannis,dudek]@cim.mcgill.ca, [yasyas,junaed,philg]@cim.mcgill.ca

Abstract

We describe the development and deployment of a system for long-distance remote observation of robotic operations. The system we have developed is targeted to exploration, multi-participant interaction, and tele-learning. In particular, we used this system with a robot deployed in an underwater environment in order to produce interactive webcasts of scientific material. The system used a combination of robotic and networking technologies and was deployed and evaluated in a context where students in a classroom were able to observe and participate to a limited degree in the operation of a distant robot being used for environmental assessment.

1 Introduction

Being able to monitor robotic experiments at a distance is quite important, particularly when the robots are being operated in environments that are remote, hostile, or difficult to access by a human. Since robotics technologies are naturally suited to inaccessible or hostile environments, the combination of robotics and web-casting is a perfect match. For example, mobile robots have been successfully used to acquire data from the surface of Mars which was then almost immediately distributed over the Internet [11]. They have also been used to acquire video images from famed shipwrecks such as the Titanic [6], or were used to monitor the conditions near volcanoes ([14],[1]). In all the above scenarios, a small set of experts are normally located next to the console operating the robot and, in most cases, act as the direct operator of the robot. On the other hand, a large body of data consumers is usually located much further away.

We are investigating a different operational scenario where a small team with a mobile robot is dispatched to a remote location. The task of the team is to relay the information in real-time to one or more "customers" located at a distance from the actual deployment site (in our trial, this distance was of intercontinental scale). A context in which such a scenario would be applicable would be, for example, where students in different classrooms would be able to observe remote habitats such as coral reefs or rainforests. The students and/or the educator can interact with the remote team, potentially requesting the robot to visit or focus on areas of particular interest. In this paper, we report the results of precisely such an exercise where the students were located on a different continent from the experiment.

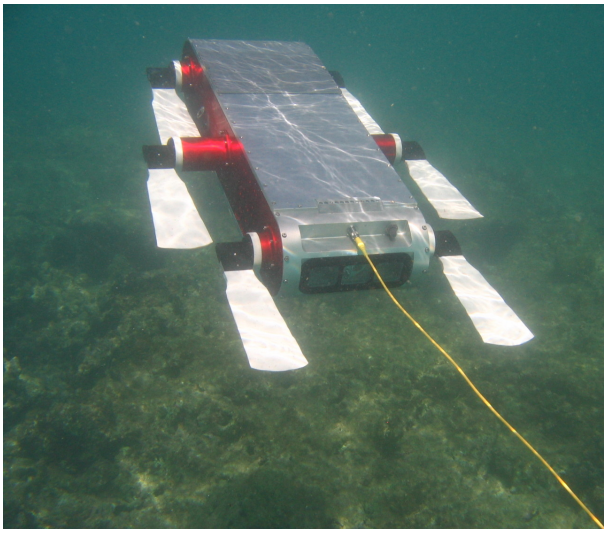

Figure 1. *Ramius*, **a member of the AQUA family of robots, over the reef in Barbados.**

A different, and more critical, scenario is in the context of search and rescue operations. A first responders team equipped with a search and rescue robot is trying to access a remote location. The cognitive load of operating the robot is reportedly high enough that it does not allow an operator to also analyze the information coming from the robot [9]. In turn, help from experts which are located in different parts of the world can be recruited during such operations. By providing critical data from the robot across the Internet, these experts would have near real-time access and would be capable to assist the operator on site to deploy the robot in the most promising areas.

We have developed a system that evaluates the feasibility of such remote robot interaction based on the deploy-

978-0-7695-4040-5/10 $26.00 © 2010 IEEE
DOI 10.1109/CRV.2010.41

ment of an underwater vehicle in the Caribbean Sea, as illustrated in Fig. 1. AQUA is specifically designed for surveying operations, for example monitoring the conditions of a coral reef or observing fish behavior. During this deployment, the robot performed a variety of operations such as following a pre-set depth profile while swimming, and swimming following different trajectories. Other experiments included using an unsupervised learning technique that confined the robot motion above a coral reef, using visual appearances [7].

During the experimental sessions most relevant to this paper, the robot was deployed over a reef structure near the shore. The robot was monitored by a local operator, while visual information as well as the state of the robot was collected by a second laptop. This laptop broadcasted the collected information in a coherent manner as a web service. At the same time, a graduate student from our lab was giving a guest lecture at a high school in Montreal. During the lecture, the student had access to the information broadcasted from the experimental site. Students from the biology and the robotics class attended the lecture. This lecture combined information on marine biology, the importance of coral reefs, and also topics on computer science and robotics. This single presentation introduced students to underwater robotics and increased their motivation to enter a robotics competition. At the same time, our experiments demonstrated the ability to monitor robotic experiments at remote sites over the Internet using available off-the-shelf web-based technology. This resulted in highly reduced costs as opposed to custom made setups traditionally used by NASA and other research groups [15].

Our sample application is the provision of real-time interactive web-casting in an educational context. One factor in such an application is the need to provide data on schedule irrespective of Internet service interruptions. Even though we were unable to preclude the possibility of system failures (including loss of Internet connectivity for the entire island we worked from), it was important for the broadcast to proceed on time. We addressed it by partitioning the system between a remote server group located with the robot, and a local proxy server located nearby the clients (where nearby is expressed with respect to Internet connectivity, and may still be rather distant in geographic terms).

The next section presents related work. The experimental setup is described in Section 3. Section 4 presents our experiments of offering a window to the reef in Barbados to students located in a high school in Montreal. Finally, conclusions and future work are discussed in Section 5.

2 Related Work

The remote monitoring and operation of robotic systems has many applications. Of particular interests are the deployment of robots in remote environments for medical purposes, or for search and rescue operations. The role of teleoperation and telepresence is quite important in many applications [4].

As noted earlier, the further frontier where robots operate and send back data from, is the planet Mars. Future missions are going to maintain human telepresence on another planet for the years to come with the "Mars Science Laboratory" (MSL) [27] and ESA's ExoMars [25]. At the Canadian Space Agency (CSA), the remote operation of robotic systems for on-orbit-servicing of satellites has been extensively studied [18]. A dual arm manipulator has been teleoperated successfully across the Atlantic ocean to perform the capture of a tumbling satellite mock-up. The above scenario is currently extended to remotely monitor and operate the robotic system from the International Space Station (ISS). Further plans also include the operation of a planetary rover operating at CSA's Mars Emulation terrain from ISS [12].

In the area of search and rescue robotics, a lot of effort has been spent in the human-robot interaction. In addition to the individual approaches, competitions designed to bring the different researchers together have also been organized [10].

Telepresence underwater was considered as early as 1995, when a team from NASA performed experiments in Antarctica [23], but without broadcasting the information further than the operators station. The Aquarius research station operated by the NOAA has executed a series of webcasts from deep undersea, but these appear to have been unidirectional feeds without interactive control or dialogue [1].

Pioneering work by a group of Italian researchers is described in [2, 3, 26]. Different experiments were performed by the underwater vehicle Romeo first controlled on site and then teleoperated over a satellite link. It is worth noting that special hardware was required to establish a satellite connection during operations.

Robotic testbeds have also been used in efforts to make education more engaging. NASA's Site of Remote Sensing project introduced telerobotics to high-school students [17], and University of Essex used a web based interface for the control of a mobile robot [28, 24]. The JASON project was built to give teachers access to scientific expeditions[2]. Finally, manipulators have been used by the University of Verona as a teaching tool for computer science students [17].

3 Experimental Setup

A number of key components were essential for this experiment to succeed. In this section, we provide a short

[1] http://www.livingoceansfoundation.org

[2] http://www.jason.org/

description of these components. We start first by presenting the amphibious robotic platform (AQUA), the central part of this experiment. We then discuss the other essential hardware component: the networking elements used to relay in real-time the information. The communication protocols and software architecture adopted to expose the state of the robot to the outside world is presented last.

3.1 AQUA Robotics Platform

The vehicle used in the experiments, nicknamed *Ramius*, (seen in Fig. 2) is a hexapod robot specifically designed for amphibious locomotion. It is part of the family of the amphibious robots named AQUA [5]. The platform itself was adapted from the successful RHex platform [19]. This family of robots has been used extensively in field trials, notably in visual-servoing tasks described in [22],[20]. The robot is capable of operating in different autonomy scenarios from full tele-operation to tetherless autonomy [21].

Figure 2. Picture of the swimming robot *Ramius* while being deployed in Barbados.

Every robot of the AQUA family is equipped with six limbs. Each limb is moved using a single electrical motor, greatly enhancing the robustness of the platform while simplifying the design. The limbs themselves are flippers, and thrust is generated by moving them rapidly. The location and orientation of the six flippers is such that thrust underwater can be generated in five degrees of freedom: pitch, roll, yaw, heave and surge.

Two PC/104 single-board computers, one a 300 MHz Pentium-equivalent running QNX and the other a 1400 MHz Pentium-M running GNU/Linux, are used for on-board computation. Communication to a remote operator laptop is done over Firewire transmitted via a fiber-optic tether.

3.2 AQUA Stability Augmentation System

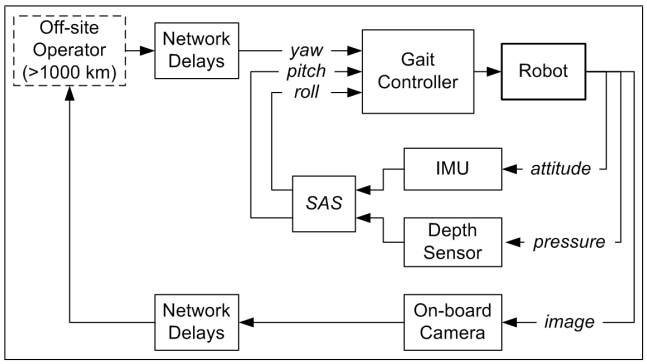

Figure 3. Block diagram of control scheme for the stability augmentation system. The dashed box represents a hypothetical off-site operator, trying to remotely change the heading of the robot.

A proportional-derivative linear controller [16] in the Stability Augmentation System (SAS) was used to maintain the pitch and roll angles of the robot. The SAS modified the swimming pattern of the flippers to generate proper pitch or roll correcting moments. A simple autopilot maintained the depth of the robot $1 \pm 0.2\ m$, using small pitching corrections executed by the SAS. The yaw angle was left uncontrolled, so an operator could modify the heading of the vehicle by issuing turning commands. Fig. 3 shows the block diagram of the control architecture.

There was a significant cognitive load reduction by having this autopilot stabilizing the robot. It indeed relieved the operator from maintaining the critical pitch and roll angles; both are needed to maintain the vertical orientation of the robot. The image captured by the on-board cameras are therefore easier to interpret, since the vertical orientation is known by the user and stable over time.

It will also prevent a potentially dangerous condition called pilot-induced oscillation. These are undesired oscillations that are generated when a pilot tries to correct the attitude of a vehicle, but in such a way that his input overcorrects, thus resulting in oscillations of increasing amplitude [13]. They occur for vehicles that are very responsive, and when significant delays are present in the human pilot inputs. Considering that our vehicle have pitching and rolling rates close to $30^o/sec$ and $90^o/sec$ respectively, the introduction of long-distance network delays (over a second

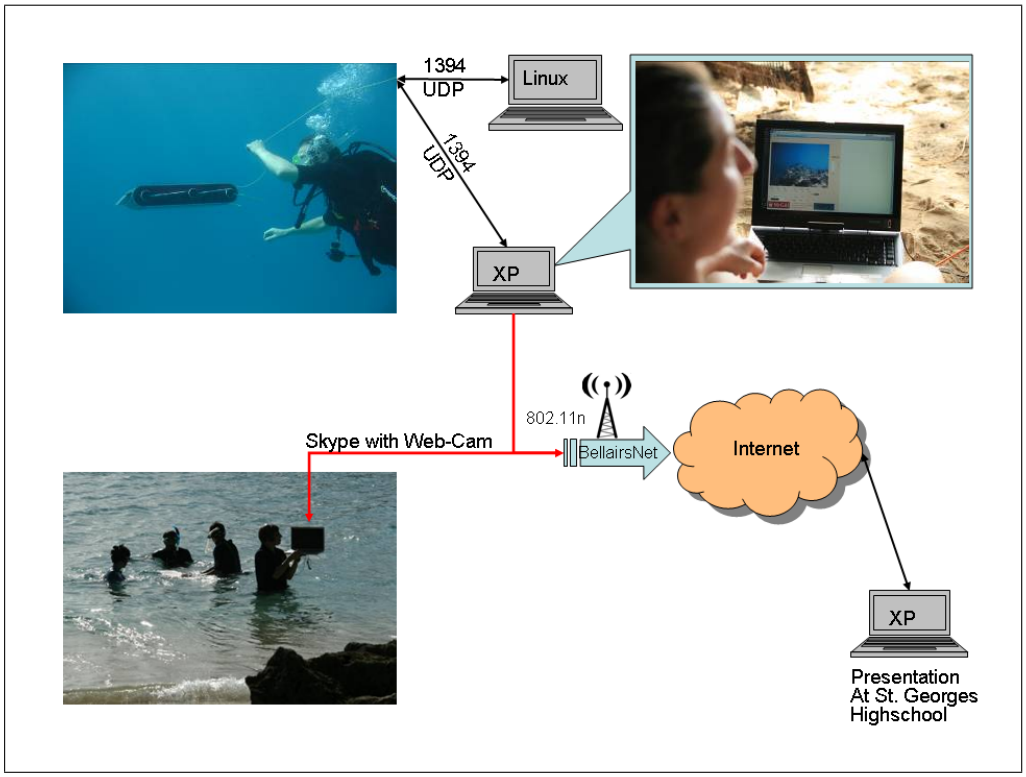

Figure 4. The different components of the experimental setup.

if satellite hops are present) would make stabilizing these two axis by a remote operator impossible [8].

3.3 Software Architecture: Proxy-based streaming

Our delivery system is targeted to two classes of clients: a primary user who needs maximum performance streaming content, and a set of additional secondary clients who only want partial data. In a presentation context, these correspond to a lecturer or classroom teacher, and a set of students who may also be accessing the content.

Since the actual robot and the associate interface may be geographically remote, a reliable and fast connection cannot be relied upon. Furthermore, since the remote site may have unpredictable Internet connectivity constraints, it may not even be convenient for the clients to connect directly to the robot server system. This concern is further exacerbated by operational security considerations.

For all these reasons, and in order to provide a responsive user interface to both types of client we have employed a remote proxy server to provide web-based data to the clients. Our approach uses 4 subsystems.

1. A Microsoft Robotics Developers Studio Services (MRDS) [3] host which interacts directly with the robot and serves live real-time data.
2. A secondary host located near the robot which pulls data from MRDS and supplements it with additional feeds to include media other than what is coming directly from the robot. This secondary host pushes data to a proxy server (ideally located near the clients).
3. A proxy server that caches the data from the remote location and also provides a backup feed in the case of a loss of connectivity.
4. A web-based client used to view the content. This can be deployed on any standard web browser.

In practice, the secondary host serves to isolate the operational robot system from Internet-related activities. In our application, we also use it to insert and transmit a stream of background commentary to the clients. The proxy server provides low-latency data and can serve a large number of clients. In practice, we found the long-haul connection to suffer from intermittent drop-outs or occasional high latency. Using a local proxy naturally could not improve frame rates, but it did assure that the local client refreshes were rapid and also that they did not cause further congestion or thrashing. In addition, the local proxy was able to supply a pre-prepared static feed in the case of an extended drop-out or loss of service. This provided a critical level of reliability for real-time mission critical use.

In our trial application, the primary and secondary hosts were located on a Caribbean island, while the proxy host

[3] http://msdn.microsoft.com/en-us/robotics/default.aspx

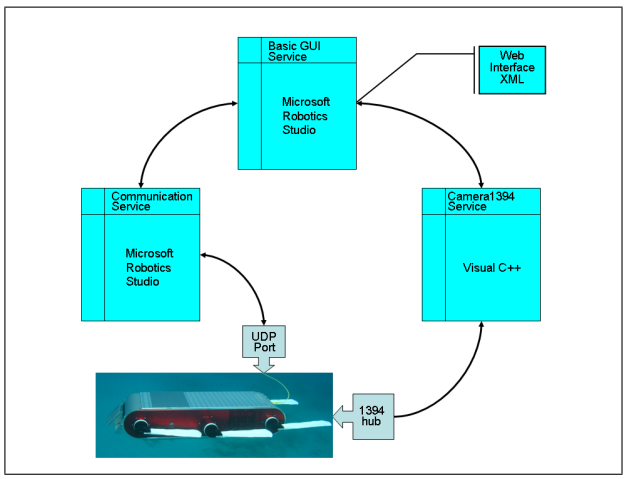

Figure 5. The different services developed using the Microsoft Robotics Studio.

and the clients were both located in the same city in Canada. Figure 4 presents an outline of the connections, especially the primary and secondary hosts.

3.4 Software Architecture: Microsoft Robotics Studio Services

We employed the Microsoft Robotics Developers Studio (MRDS) in order to expose the state and control of the Aqua robot. MRDS is an infrastructure which provides a uniform look and feel interface of sensor information between both built-in and user-added services. At the lower level, we employed a service that polled the robot at 10Hz, requesting the state of the robot via UDP. The collected state information comprised:

- orientation of the robot (yaw, pitch and roll angles) measured by a Microstrain^{TM} 3DM-GX1 Inertial Measurement Unit,
- depth estimate reported by a depth sensor,
- six leg (flipper) angles,
- battery voltage.

The basic service termed *Communication Service* exposes the above mentioned data to any other service that subscribes to it. Currently, the only other service developed is a basic GUI Service which exposes the state of the robot to remote users. An XSLT style-sheet is used to provide the formating of the raw data. A second program in Visual C++ was used to obtain individual camera frames and make them available to the XSLT-based GUI. Figure 5 presents the different services and their interaction, while Fig. 6 shows a snapshot of the web-based interface.

Figure 6. The monitoring GUI developed using the Microsoft Robotics Studio.

4 Experimental Results

For our remote telepresence field trials, a two-part setup was used – one end being the robot and the associated communication infrastructure at the field trials, and the other was a remote presenter equipped with a laptop computer hooked up to the Internet. MRDS offers a built-in, web-based service to present a remote robot console. To test this service, we used an Internet Explorer browser on a laptop with Windows Vista Home Premium installed. With the robot transmitting live videos and telemetry information, the MRDS services collated the data and presented them on a publicly accessible website (which uses proprietary Microsoft HTML extensions, and hence the need to use Internet Explorer for browsing the above-mentioned website).

A member of our team in Montreal provided live commentary about the experiments that were being conducted in Barbados to an audience of high-school students; see Fig. 7. A large section of the students are involved in robotics projects, and as a team they take part in different robotic competitions at a provincial and national level. Along with commentary by this team member in Montreal, the audience was also able to directly talk to the team members working in Barbados, via a live audio and video link. Questions by the students were answered by our team members from the seaside, as they worked on conducting their experiments with the Aqua robot. Together with live audio and video feeds, the students were able to see live video coming from the robot's camera during the experiments, and were able

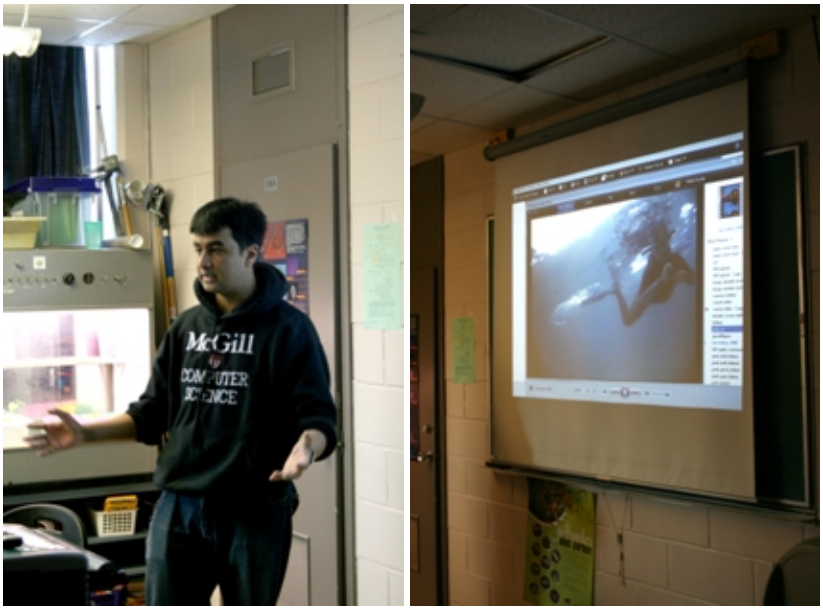

(a) (b)

Figure 7. (a) The presenter in Montreal, in front of the classroom. (b) Picture of the live video feed from the robot, projected on the screen.

to look at different telemetry measurements coming from the robot's internal sensors. The interaction with the team members at the experiment site, coupled with the live data from the experiments gave the students an all-round idea about the intricacies of field robotics and the scientific impact of the experiments, all from sitting in their classroom.

5 Conclusions and Future Work

The telepresence experiments we performed illustrated the feasibility of building and deploying web-based services that present the collected data together with the state of a robot. The employment of web-based technologies enabled the transmission of the real-time data in a coherent manner over the Internet. This enabled us to engage two groups of high school students with diverse interests in marine biology and robotics in an online presentation of our experiments with an underwater robot.

One of the lessons learned was that having participants at different stages in the control and communication chain had definite advantages. It also made for greater flexibility and, of course, robustness to bandwidth fluctuations. Due to the variable latencies and degrees of engagement (and hence impressiveness), the experience of a user who was with the robot, on site, or simply at the remote location was quite different and each had a unique and useful perspective. The experiment emphasized that there are several distinct classes of "telepresence" to be achieved in a context like this. In terms of the experience of the observers, they can feel like they are in the role of the robot, in the role of a companion to the robot, in the role of the robot supervisor on shore, or in the role of a member of the development team located at an arbitrary location. Each of these roles was, to some extent, present in our tests. The tradeoff between them, the nature of each, and their respective advantages seem to be significant in a pedagogical context, but they are difficult to define precisely let alone quantify. In future work, we would like to further examine the effect of each on the final experience.

We are currently working to expand our framework in a variety of ways. We would like to enable selected users to operate the vehicle over the web-interface. This is now possible since the most critical axes, pitch and roll, are maintained by the robot's autopilot system. Moreover, we are looking into interactive technologies to allow remote users to convey to the on-site operator their preferences as to where the robot should investigate. This would allow them to focus their and the robots attention in areas of particular interest.

One of the main limitations encountered was the intermittent and high-latency Internet connection in our contact, which is common place in remote locations and with a constrained budget. The current implementation of the web interface was serving the data from the robot's Firewire camera one image at a time. This resulted in very slow update rates of less than 1Hz. Currently, we are working on a live streaming video implementation, which after the initial connection is established, the camera output is streamed live

through the GUI interface.

The experiments performed illustrated the power of a remotely operated robot as a teaching tool in different fields. Students in biology had a unique opportunity to observe the rich life of a coral reef. For the robotics students, their exposure to a real field testing of an underwater robot gave rise to many questions ranging in topics from low level issues, such as the power system and the motors used, to high level topics, such as sensor choice, visual servoing, and planning algorithms. We are looking forward to build on our experience and broaden our audience to several different groups.

6 Acknowledgments

We would like to thank the rest of the members of the AQUA project: Y. Girdhar, N. Plamondon, B. Vertzijlenberg, A. Topol, M. Jenkin, S. Ibararan, and E. Dumas, that helped with the experiments. We would like also to acknowledge the financial support of Microsoft Research that contributed to the success of this project.

References

[1] D. Apostolopoulos. Locomotion configuration of a robust rappelling robot. In *Proc. of the International Conference on Intelligent Robots and Systems-Volume 3*, page 3280, Washington, DC, USA, 1995. IEEE Computer Society.

[2] R. Bono, G. Bruzzone, M. Caccia, E. Spirandelli, and G. Veruggio. Romeo goes to antarctica [unmanned underwater vehicle]. In *OCEANS '98 Conference Proceedings*, volume 3, pages 1568–1572, Sep 1998.

[3] G. Bruzzone, R. Bono, G. Bruzzone, M. Caccia, M. Cini, P. Coletta, M. Maggiore, E. Spirandelli, and G. Veruggio. Internet-based satellite teleoperation of the ROMEO ROV in Antarctica. In *Proc. of the 10th Mediterranean Conference on Control and Automation*, 2002.

[4] J. Burke, R. Murphy, E. Rogers, V. Lumelsky, and J. Scholtz. Final report for the darpa/nsf interdisciplinary study on human-robot interaction. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 34(2):103 – 112, May 2004.

[5] G. Dudek, M. Jenkin, C. Prahacs, A. Hogue, J. Sattar, P. Giguère, A. German, H. Liu, S. Saunderson, A. Ripsman, S. Simhon, L. A. Torres-Mendez, E. Milios, P. Zhang, and I. Rekleitis. A visually guided swimming robot. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada, August 2005.

[6] R. Eustice, H. Singh, J. Leonard, M. Walter, and R. Ballard. Visually navigating the RMS titanic with SLAM information filters. In *Proc. of the Robotics: Science and Systems (RSS)*, Cambridge, MA, Jun. 2005.

[7] P. Giguere, G. Dudek, C. Prahacs, N. Plamondon, and K. Turgeon. Unsupervised learning of terrain appearance for automated coral reef exploration. In *Proc. of the Canadian Conference on Computer and Robot Vision*, Kelowna, British Columbia, May 2009.

[8] J. J. Corde Lane, C. R. Carignan, B. Sullivan, D. Akin, T. Hunt, and R. Cohen. Effects of time delay on telerobotic control of neutral buoyancy vehicles. In *Proc. of the IEEE International Conference on Robotics and Automation*, volume 3, pages 2874 – 2879, May 2002.

[9] M. W. Kadous, R. K.-M. Sheh, and C. Sammut. Effective user interface design for rescue robotics. In *Proc. of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction (HRI)*, pages 250–257, Salt Lake City, Utah, 2006.

[10] G. Kantor, S. Singh, R. Peterson, D. Rus, A. Das, V. Kumar, G. Pereira, and J. Spletzer. Distributed search and rescue with robot and sensor teams. In *Field and Service Robotics*, pages 327–332. Sage Publications, 2003.

[11] M. Maimone, J. Biesiadecki, E. Tunstel, Y. Cheng, and C. Leger. *Intelligence for Space Robotics*, chapter Surface Navigation and Mobility Intelligence on the Mars Exploration Rovers, pages 45–69. TSI press, 2006.

[12] E. Martin, R. L'Archevêque, S. Gemme, I. Rekleitis, and E. Dupuis. The avatar project: Remote robotic operations conducted from the international space station. *IEEE Robotics and Automation Magazine*, 14(4):20–27, Dec. 2008.

[13] D. T. McRuer. Pilot-induced oscillations and human dynamic behavior. Technical Report NASA Contractor Report 4683, National Aeronautics and Space Administration, July 1995.

[14] G. Muscato, G. Nunnari, and S. Guccione. Robots for volcano exploration: A new perspective. In *Proc. of the Eight International Symposium on Robotics with Applications, World Automation Congress (WAC2000)*, Maui, 2000.

[15] C. F. Olson, L. H. Matthies, J. R. Wright, R. Li, and K. Di. Visual terrain mapping for mars exploration. *Computer Vision and Image Understanding*, 105(1):73–85, 2007.

[16] N. Plamondon and M. Nahon. Trajectory tracking controller for an underwater hexapod vehicle. In *Oceans 08 MTS/IEEE*, Quebec City, Canada, September 2008.

[17] M. Portsmore, C. Rogers, P. Lau, and E. Danahy. Remote sensing and tele-robotics for elementary and middle school via the internet. *Computers in Education Journal*, XIV(2):72 75, April - June 2004.

[18] I. Rekleitis, E. Martin, G. Rouleau, R. L'Archevêque, K. Parsa, and E. Dupuis. Autonomous capture of a tumbling satellite. *Journal of Field Robotics, Special Issue on Space Robotics, Part II*, 24(4):275–296, April 2007.

[19] U. Saranli, M. Buehler, and D. E. Koditschek. Rhex: A simple and highly mobile hexapod robot. *International Journal of Robotics Research*, 20(1):616 – 631, Jul. 2001.

[20] J. Sattar and G. Dudek. A boosting approach to visual servo-control of an underwater robot. In *Proc. of the 11th International Symposium on Experimental Robotics, ISER*, Athens, Greece, July 2008.

[21] J. Sattar, G. Dudek, O. Chiu, I. Rekleitis, P. Giguere, A. Mills, N. Plamondon, C. Prahacs, Y. Girdhar, M. Nahon, and J.-P. Lobos. Enabling autonomous capabilities in underwater robotics. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3628 – 3634, Nice, France, 2008.

[22] J. Sattar, P. Giguère, G. Dudek, and C. Prahacs. A visual servoing system for an aquatic swimming robot. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, Canada, August 2005.

[23] C. Stoker, D. Burch, B. H. III, and J. Barry. Antarctic undersea exploration using a robotic submarine with a telepresence user interface. *IEEE Expert*, 10(6):14–23, Dec 1995.

[24] P. W. Tsui and H. Hu. A framework for multi-robot foraging over the internet. In *IEEE International Conference on Industrial Technology*, volume 2, pages 897 – 902, Dec. 2002.

[25] J. Vago. Overview of exomars mission preparation. In *Proc. of the 8th ESA Workshop on Advanced Space Technologies for Robotics & Automation*, Noordwijk, The Netherlands, Nov. 2004.

[26] G. Veruggio. Marine robotics: a global interdisciplinary approach to the scientific, technological and educational aspects. In *5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, Lisbon, Portugal, 5-7 July 2004.

[27] R. Volpe. Rover functional autonomy development for the mars mobile science laboratory. In *Proc. of the IEEE Aerospace Conf.*, Big Sky, MT, USA, 2006.

[28] L. Yu, P. W. Tsui, Q. Zhou, and H. Hu. A web-based telerobotic system for research and education at essex. In *Proc. of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, volume 1, pages 37 – 42, July 2001.