PAPER

Trends in ROV Development

AUTHOR
Steve Cohan
Schilling Robotics, Inc.

ABSTRACT

This paper describes three new directions in the development of ROV control systems: exploiting the capabilities of digital video, using model-based control techniques for ROV operations, and providing ROV systems with sophisticated remote diagnostics. Goals of these trends are to increase the value of ROV operations end products (such as video), to increase efficiency of operations by adding automation, and to increase productive time by bringing outside resources to the ROV for maintenance and troubleshooting.

Advances in digital video can enhance the value of video imagery by facilitating the combination of video with other data, improving utility, and increasing the range of analyses and products that can be created from video. Current digital video technology enables systems to treat video in a manner close to the transmission and processing of other sensor data, reducing the number of components and interfaces between the camera and the video imagery's end use. Model-based control has the potential to relieve pilots from managing all tasks concerning low-level motion, coordination, and control, enabling pilots to concentrate on higher-level task planning and execution. Applying software techniques designed specifically for managing data flow in distributed control systems can give remote experts access to all the tools and information necessary to assist with worksite maintenance and troubleshooting.

Introduction

Recent trends in offshore industries (primarily offshore oil and gas) are currently driving changes in how remotely operated vehicles (ROVs) are used. In particular, the movement of offshore operations into deeper water and the movement of processing functions to the seafloor have required ROVs to become more efficient and better able to perform a wide range of tasks. These requirements are in turn driving changes in ROV technology.

This paper discusses three trends in ROV technology that are now feasible because innovations (such as Ethernet communications and sophisticated data gathering and data modeling techniques) have been adapted from use in other industries to applications in subsea intervention. These trends are:

- Treating video imagery as simply another form of data that can be searched, organized, and merged with other data in sophisticated digital asset management systems
- Using model-based control to plan, simulate, and then automatically execute difficult ROV intervention tasks
- Greatly expanding the data available to diagnostic experts who perform remote troubleshooting for ROV field repairs

The means to implement these tools and techniques comes not from the subsea intervention industry but from other related fields. These other fields share with ROV operations the fact that by applying advanced hardware and software to increasingly automate operations, gains are realized in productivity and in the economic value of products and services.

Digital Video

Underwater imagery is a key output for most ROV operations. In particular, video is almost always critical to operations, and in some cases, it is the ROV operator's deliverable product. Video's utility increases substantially when it is combined with other data (such as time and position information) because associating such data with video images enables users to search, organize, and deliver video that is packaged to optimize its utility to customers. For example, an ROV operator's customer may be interested in observing all sections of a surveyed pipeline in which the pipeline location deviates from its previously surveyed route by more than a given number of meters. Being able to associate "built-in" position data with the video would make it easy to deliver this product.

Taking this reasoning one step further, video imagery can be considered as simply another form of data, akin to data from the multitude of other sensors aboard an ROV. Ideally, video could be combined, visualized, and merged with other data (via processes that are transparent to the operator), providing a rich visual window into an operation's state.

Currently, however, video images are seldom transparently combined with other data since conventional analog video formats do not mix directly with sensor data received over ROV telemetry networks. Ironically, the video data is generally traveling over the same media, but current techniques for managing and processing video data complicate its combination with sensor data. To combine video and sensor data, users must supply and interconnect additional components and interfaces for data, as illustrated in Figure 1. The figure shows that the video switcher routes some of the feeds into an overlay unit, which in turn receives sensor data such as depth and heading over a separate data channel. The overlay output is then routed to other destinations, including a tiling unit that combines several video feeds into a single display. The tiling unit also chromakeys a "heads-up display" background over the combined video display, using sensor data obtained over yet another channel. Such a system also needs additional control signals to manage the switcher, overlay unit, and tiler.

As illustrated in Figure 2, digital video has the potential to move freely between digitally enabled devices without inter-

FIGURE 1

Combining video and sensor data currently requires additional components and interfaces.

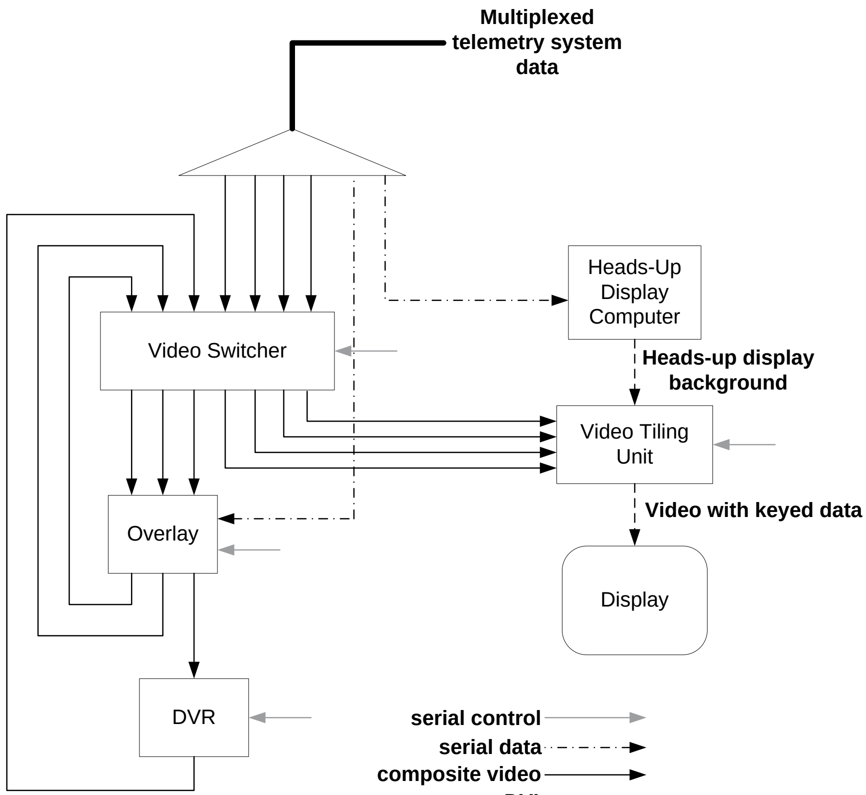

mediate conversions to analog form and without traversing extra physical interfaces with their inherent requirements for cabling, power, and configuration. In the system shown in Figure 2, all devices that produce or consume video interface to an Ethernet-based local area network (LAN) (Stanley, 2007). The system's other sensor data and state data traverse the LAN and are available to all entities that perform data processing, display, and storage functions.

As illustrated in Figure 3, devices that do not typically include native Ethernet interfaces, such as displays and cameras (other than IP cameras), can be augmented with electronics and software that transform between the device's native format and compressed packetized video.

Techniques for digital video compression, transmission and recording are technically well established and commercially available (Axis Communications, 2007; Force Incorporated, 2005). These techniques and standards make it feasible to have sophisticated digital asset management systems that provide their users with powerful tools to associate different types of data; to store, catalog, and retrieve large quantities of video content; and to distribute it to their customers (Virage Autonomy Systems, 2007; Artesia Digital Media Group, 2007). Such systems are the state of the art for broadcast, security, and video production (VisualSoft Ltd., 2004). They eliminate difficulties such as managing large volumes of physical media (for example, tapes and disks) and provide much more efficient means to retrieve content of interest via nonlinear editing capabilities. *Video analytics* encompasses an even more advanced set of techniques in which specialized algorithms extract very specific information from a video data set by analyzing trends and patterns (Siemens

FIGURE 2

HD video travels with other data over an Ethernet system.

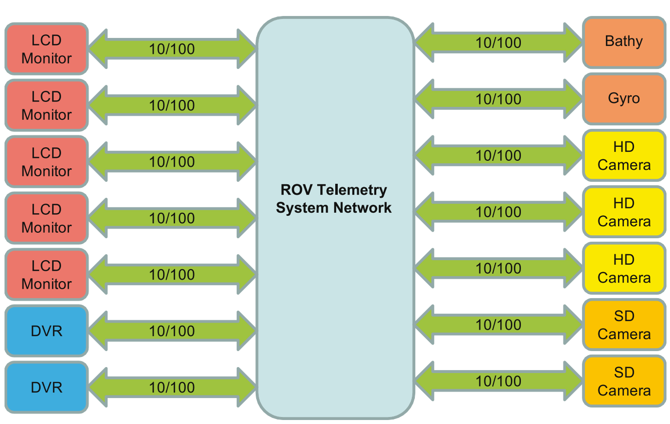

FIGURE 3

Devices that do not typically include native Ethernet interfaces can be equipped to transform between the device's native format and compressed packetized video.

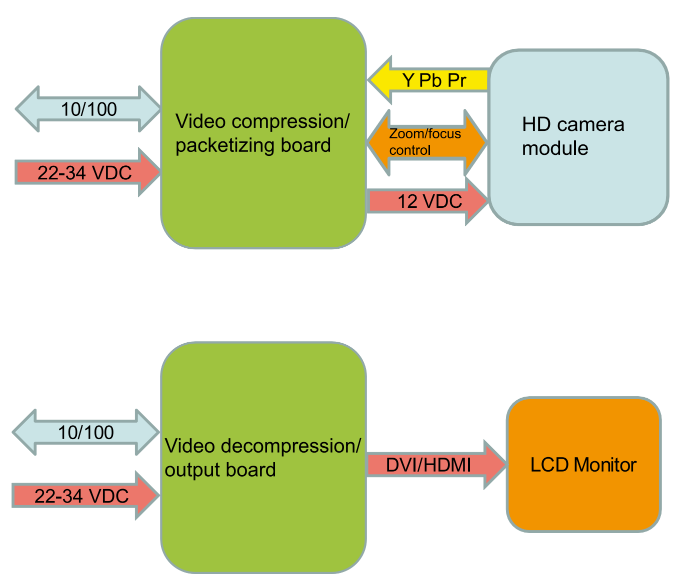

Switzerland, 2007). While many of today's commercial offerings in video analytics are geared toward security applications, ROV operations-oriented tasks that require video capture are also candidates for automation through video analytics.

Important digital video attributes for ROV use include bandwidth requirements, time latency, and storage requirements. Because bandwidth in general (and in ROV telemetry systems in particular) is a finite resource, techniques that are bandwidth-hungry drive tradeoffs in cost and complexity. Low-latency displays are imperative for teleoperation since ROV maneuvers and manipulator operations are degraded if images are delayed by more than approximately 100 ms. The H.264 digital video compression standard in particular offers streaming high-definition (HD) video that has sufficient compression to make its network bandwidth and storage requirements reasonable for ROV applications. There are now product offerings on the market designed to transmit HD video (1920 X 1080, interlaced or progressive) over Ethernet networks with both low latency (less than 100 ms) and low bandwidth (less than 30 mbps) (4i2i Communications, 2007; HaiVision Systems, 2006). The H.264 standard's popularity also ensures that a multitude of devices and software will be available to display, manage, and deliver the video.

In the near term, it will be practical to deliver high-quality HD video systems that have a reasonable cost, that meet ROV telemetry system requirements, and that can combine data in ways unconstrained by conventional video management techniques. Such systems will let ROV operators further enhance the value of their end products by using content management tools that are now standard in the broadcast and video production industries.

Model-Based Control

Today's computer-aided design and engineering systems provide detailed geometric (and sometimes physical) models of components and systems, and these data are typically used beyond the design stage in manufacturing, quality assurance, and technical documentation. Such models can also be used in various ways by organizations and people who install, commission, and operate subsea installations and equipment.

Planning for intervention and operations is one use for such models. Using simulation software, inspections and interventions can be designed and rehearsed, yielding useful products like work flow analyses, contingency plans, and training scenarios for the pilots and technicians who will perform the tasks.

Task execution is a natural progression from task planning, and if the steps for a task have been simulated, refined, and then stored in a manner that is interpretable by the system executing the task, the task can be automatically executed by interpreting and replaying those steps. The user's input is needed to define the environment and the task (modeling), and to verify that what has been modeled will be effective (simulation). Automated systems then have the potential to either assist in or autonomously perform task execution.

Many processes in various disciplines already use the details in a model to execute tasks in the physical world. For example, rapid prototyping technologies for creating prototype machined parts are widely available (Castle Island, 2007). These systems typically "grow" a model by building up two-dimensional cross-sections of a computer-assisted design (CAD) model (Figure 4).

For many years, computer-assisted surgery has used models derived from medical images to plan and then execute complex neurologic and orthopedic surgical procedures. More recently, these techniques have facilitated minimally invasive spinal fusion

FIGURE 4

Stereolithography system for rapid prototyping using computer models. Photo of the Viper™ Pro SLA® system courtesy of 3D Systems Corporation.

FIGURE 5

Workflow for computer-assisted surgery.

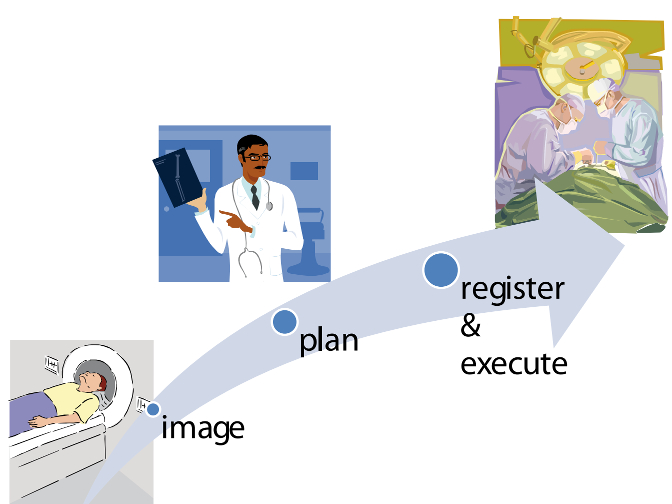

and joint replacement procedures (Nolte, 1999). In the operating room, the surgeon "registers" the patient's anatomy with the image data (that is, he or she matches up particular points on the patient's body with the same points on the model of the patient's body) as shown in Figure 5. In some cases, the procedure can then be executed autonomously by a robotic device (Taylor, 1994). Alternatively, the surgeon can use an image guidance system in which software provides navigational displays that guide the surgeon in placing and orienting instruments relative to the patient's body.

Model-based control conceptually maps over to ROV tasks for operations such as hot stab placement since this task requires proper position and orientation of the hot stab tool relative to the hot stab receptacle. As this task is currently performed by ROV operators, the lack of three-dimensional information and the large number of degrees of freedom can make it difficult to move the ROV and manipulator to the proper location. The operator must determine how to best partition the task among these many degrees of freedom (ROV motion, manipulator, camera views), which can be a difficult process requiring skill and experience.

Augmenting the hot stab operation with model-based control, however, has the potential to simplify the operation. Ideally, an existing model of the installation would be the starting point. This model would then be registered with the actual installation's position and orientation, allowing the pilot to specify locations of interest on the model. At this point, the ROV control system could assist the pilot in moving to these locations. If the ROV were equipped with DP and/or navigation capability, it could move autonomously to the desired locations so that the manipulator workspace would be well placed for the hot stab.

It is possible to envision tasks that require a manipulator-mounted sensor to be moved along a path that is defined relative to the modeled installation, such as inspecting a structural joint in nondestructive testing. By using a model of the component, the manipulator control system could move the tool in the required spatial path while maintaining proper distance from the surface under inspection.

A key requirement for successful model-based control is that objects be mapped into the relevant coordinate system for control with sufficient accuracy for the task to be executed. Accurate mapping in turn depends on accuracy of the data input to the registration procedure, i.e., the procedure that calculates the transformation between the coordinate system where the work is planned (the "model" coordinate system) and an available and useful coordinate system at the worksite. One technique for finding this transformation is to obtain the coordinates of a set of "fiducial" points (that is, points used as a standard reference) in both coordinate systems. Once these fiducial points have been determined, planners can calculate the transformation between these coordinate systems using either direct geometric methods or "least squares" methods. With the latter method, adding extra fiducial points reduces the error and therefore increases the accuracy of the transformation (Umeyama, 1991). While the fiducial technique is mathematically simple and has an analytical solution, it can be difficult in the subsea environment to find a good set of fiducial points that are easily measured.

A more sophisticated method is to collect a relatively large set of points (called a point cloud) from a surface in the work environment that is easily represented mathematically in the model, and then find the coordinate transformation that minimizes the point cloud's distance from the modeled surface. The advantage of this solution is that it does not require that specific points in the work environment be measured, but only a set of points in a region of interest.

Regardless of the method used, coordinates of points in the work environment must be determined. One method is to locate the ROV and all fiducial points using conventional acoustic measurement techniques. A shortcoming of this approach is the variable (and not very high) accuracy inherent in acoustic positioning, which depends on factors such as local environmental conditions (which affect sound velocity) and the accuracy and layout of the transponder array. Significant effort is also required to set up the array, and the number of transponders required and the effort to survey them could add unacceptable cost and complexity.

An alternative approach is using an ROV-based position sensor to collect either fiducial coordinates or point clouds from the work region. For example, a properly calibrated manipulator can serve as a "portable coordinate measuring machine," quickly digitizing multiple points in its workspace. In this case, the collected points are transformed from the manipulator or ROV coordinate frame to the model coordinates. For this approach to work, the ROV/manipulator platform must remain stationary relative to the environment. In theory, once the registration has been obtained, an ROV with DP and/or navigation capability could update the registration as the ROV moves (at least over small distances), eliminating the requirement for the intervention to take place in a completely stationary state.

Another promising approach involves using a measurement technology that does not involve contact (such as multibeam sonar) to digitize the local environment and provide three-dimensional data. This approach would be much more automated and efficient than using a device such as a manipulator. However, various tradeoffs in range, power, and precision and accuracy must be considered in developing a feasible multibeam sonar approach.

Ideally, during task execution under model-based control, pilots would be relieved from low-level path planning decisions and would be able to think of operations in terms of "macro" steps, such as moving along a path or into a receptacle. The payoff for developing, refining, and commercializing model-based control capability is the automation of various ROV intervention tasks, with potential benefits of greater productivity, reduced operator load, and increased capability for complex interventions.

Remote Diagnostics

ROV manufacturers strive to make highly reliable systems that have minimal maintenance requirements. However, as ROV systems become more capable and feature rich, system complexity inevitably increases. Well-designed systems can mask complexity by presenting users with simple interfaces and intuitive workflows. When problems occur, however, it is sometimes necessary to remove the covers and investigate the system's internals.

ROV pilot/technicians frequently have to perform in-field diagnostics and repairs, particularly on mechanical, hydraulic, and electronic systems. Given access to drawings, schematics, and training, they can deal with most failures and can conduct preventive maintenance. However, as systems become more complex (with many components and interfaces), troubleshooting also becomes more complex. When ROV manufacturers deploy new functions, particularly software-intensive ones, pilot/technicians must simultaneously learn these new capabilities, train others, and solve problems in the offshore environment. To succeed in this situation, they need quick access to information and expert advice. Remote diagnostics capability can provide this assistance.

Remote diagnostics in the most basic form could be a telephone or instant messaging connection between a knowledgeable troubleshooter at the manufacturer's facility and the technician at the equipment site who serves as eyes and ears for the expert. For all but the simplest problems, this approach results in a slow, painstaking, error-prone process. To make the process efficient, the expert must have the data and equipment controls available at his or her location.

Various industrial and medical systems now feature remote diagnostics capability (GE Energy, 2007; Delphi Medical Systems, 2007). In addition to diagnosing and repairing malfunctions, these systems gather operational data that enables their manufacturers to recommend preventative maintenance; to suggest upgrades, complementary products, or training based on usage trends; and even to re-supply customers with consumables. Remote support is also widely available for software users, from PC consumers to purchasers of complex business systems.

Obviously, connectivity is required to implement remote diagnostic functionality. Such connectivity is widely available today via satellite links; although at sea, bandwidth availability varies and may not be on par with similar services using land-based networks. Still, available bandwidth as low as DSL or dial-up levels can support services such as Windows Remote Desktop when tuned properly (Russel, 2001). Under this limitation, it is important to use the available bandwidth to transport the most salient information to the expert. Given that ROV control systems have basic Internet connectivity, there are several options for organizing internal system state data so that it is useful to an expert.

One option is to provide a duplicate pilot station at the remote support site. This option is feasible if the control system architecture is sufficiently modular and abstract (that is, if the additional pilot station is simply another instance of a general pilot station and can be driven by available data channels). The expert then has the same view and same set of diagnostic tools that are available to the pilot/technician. This arrangement is useful when the regular operational user interface can provide the required diagnostic information (for example, when sufficient information to solve the problem is available, but the pilot/technician lacks the experience needed to solve the problem quickly). This method can also provide a training opportunity; assuming that time zone differences can be accommodated and voice/video communications are available, the expert can guide the pilot/technician through the troubleshooting process as he is performing it.

A disadvantage of the duplicate pilot station option is that solving the current problem may require data that is not available at the pilot station. The ability to "dig deeper" is critical to maximizing the expert's impact. Software tools that are unavailable at the standard pilot station (such as capabilities for advanced plotting, statistical analysis, and simulation) could be available at the support location and could be used to solve the current problem.

In general, it is best for troubleshooters to have access to all internal state data so that regardless of the problem, all data that might be relevant are available for examination and analysis. In the approach illustrated in Figure 6, a dedicated software function captures and transmits data representing a system's full internal state and communicates this data over a channel dedicated to this purpose. In this fairly typical arrangement, the capture-and-transmit function gathers data from predetermined internal sources and transmits it outside the system. The data's remote user does not have general access to arbitrary data traversing the control system.

FIGURE 6

A system's full internal state is captured to facilitate remote diagnostics.

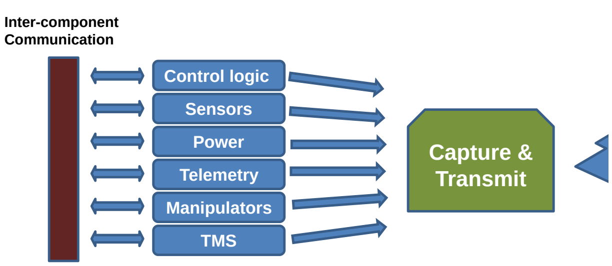

FIGURE 7

The Publish/Subscribe architecture enables data gathering as an inherent part of inter-process communications.

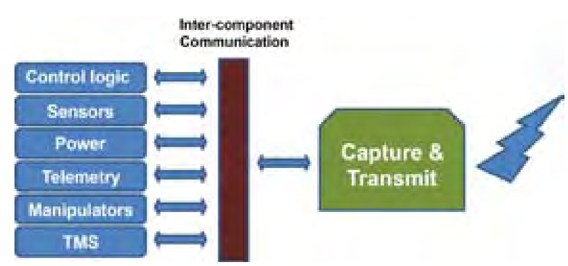

Implementing this approach presents significant challenges, however. Brute-force transmission of all relevant data may become bandwidth-intensive, and the data items of value in any particular situation are not weighted any differently than other data items. In cases of limited bandwidth, it would be useful to have a data capture system in which the expert performing the diagnosis could select the data items to be sent from the remote system.

Another challenge in this architecture is managing change. When new functionalities are deployed, data channel contents must be updated to handle the new data associated with these functions. If this update is not automatic (that is, inherently part of the system update), sustaining full diagnostic functionality becomes a manual maintenance task and thus subject to error.

A more sophisticated approach would enable the expert to query the remote system about which data items are available. The expert could then request these items as needed, adding new data items when required and deleting those that are not useful in a particular situation. Such an approach has several advantages:

- Since it is not necessary to decide *a priori* which data items will be made available, network bandwidth can be used optimally in each situation.
- When necessary, high-priority data can be favored over low-priority data.
- The data's presence in the system makes it inherently available for transmission to the expert conducting the diagnosis.

The Publish-Subscribe architecture is an inter-process communications architecture that supports this flexible data-gathering functionality (Matteucci, 2003). In this architecture, a "broker" manages communications and data flow between different entities and components as illustrated in Figure 7. Producers of data items register them with the broker, along with other information such as data format and the frequency of data updates. Producers then "publish" data to the broker as the data are updated. Data consumers also register with the broker for specific data items of interest to them, providing additional information such as frequency at which updates are desired. After data consumers have registered, the broker transmits fresh data updates to them as updates become available.

Using this Publish-Subscribe data architecture, the "capture and transmit" function in Figures 6 and 7 can ensure that the most relevant parameters are brought over from the remote system (taking available bandwidth into account). Such a system also largely eliminates the need to update the diagnostic software each time that the remote system's software is updated. For example, if a new software release for the remote system includes data items that were not previously gathered, the diagnostic software does not need to be updated explicitly to recognize the new data. Instead, the diagnostic software can simply query the remote system to determine which data items are available, and can select those items for capture and transmission.

This kind of flexibility would provide experts performing troubleshooting with access to the most important data when they need it, allowing faster response and better results for pilot/technicians and ROV end users.

References

4i2i Communications Ltd. 2007. Retrieved 31 Dec. 2007. http://www.4i2i.com.

Artesia Digital Media Group. 1999-2008. Retrieved 31 Dec. 2007. www.artesia.com.

Axis Communications. Digital video compression: reviewing the methodologies and standards to use for video transmission and storage. Retrieved 31 Dec. 2007. http://www.axis.com/documentation/whitepaper/video/compression_standards.htm.

Castle Island Co. 2007. What is rapid prototyping? 16 Nov. 2007. Retrieved 31 Dec. 2007. http://home.att.net/~castleisland/rp_int.htm.

Delphi Medical Systems. 2007. Delphi medical systems introduces customized remote monitoring service for medical devices. Retrieved 31 Dec. 2007. http://www.peakind.com/news/pressReleases/pr480-10182005/.

Force, Incorporated. 2005. Digital fiber optic multichannel V/A/D transport systems. Retrieved 31 Dec. 2007. http://www.fiber-optics.info/articles/Multichannel_Video.htm.

GE Energy. 1997-2007. Precipitator real-time diagnostics. Retrieved 31 Dec. 2007. http://www.gepower.com/prod_serv/serv/particulate_matter/en/prec_real_time/index.htm.

HaiVision Systems Inc. 2006. Retrieved 31 Dec. 2007.http://www.haivision.com.

Matteucci, M. 2003. Publish/subscribe middleware for robotics: requirements and state of the art. Jan. 13, 2003. Technical Report N2003.3, Dipartimento di Elettronica e Informazione – Politecnico di Milano, Milano I.

Nolte, L.P. and Ganz, R. (Editors). 1999. Computer assisted orthopedic surgery, Hogrefe & Huber Publishers.

Russel, C. 2001. Tuning Windows XP Remote Desktop for network bandwidth. Retrieved 18 Mar. 2008. http://www.microsoft.com/windowsxp/using/mobility/learnmore/russel_august27.mspx.

Siemens Switzerland Ltd, Building Technologies Group. 2007. Trends in intelligent video analytics. 29 June 2006. Retrieved 31 Dec. 2007. https://www.buildingtechnologies.siemens.com/press/press_release/subchannel_press_release_2006/trends_in_intelligent_video_analytics.htm.

Stanley, J. and Schilling, T. 2007. Moving to the next phase—automatically. Ocean News & Technology. 13(5):41-42.

Taylor, R.H. et al. 1994. An image-directed robotic system for precise orthopedic surgery. IEEE T Robotic Autom. 10(3):261-275.

Umeyama, S. 1991. Least-squares estimation of transformation parameters between two point patterns. IEEE T Pattern Anal. 13:376-380.

Virage Autonomy Systems. 2007. Retrieved 31 Dec. 2007. http://www.virage.com.

VisualSoft Ltd. 2004. Retrieved 31 Dec. 2007. http://www.visualsoft.ltd.uk/.