Proc. IEEE Int. Conf. on Robotics and Automation, Atlanta, GA, Vol. 3, pp. 410-417, May, 1993.

TRICLOPS: A High-Performance Trinocular Active Vision System

Albert J. Wavering, John C. Fiala, Karen J. Roberts, and Ronald Lumia

Robot Systems Division
National Institute of Standards and Technology
Gaithersburg, MD 20899

Abstract

The design, performance, and application of The Real-time, Intelligently ControLled, Optical Positioning System (TRICLOPS) are described in this paper. TRICLOPS is a multiresolution trinocular camera-pointing system which provides a center wide-angle view camera and two higher-resolution vergence cameras. It is a four degree-of-freedom direct-drive system which exhibits dynamic performance comparable to the human visual system. The vergence degrees of freedom can achieve peak velocities in excess of 30 rad/s, and peak accelerations of 1100 rad/s^2. In an example of visual tracking, TRICLOPS is shown to be capable of following a ball moving at 3 m/s, with rotational velocities of the vergence axes in excess of 6 rad/s.

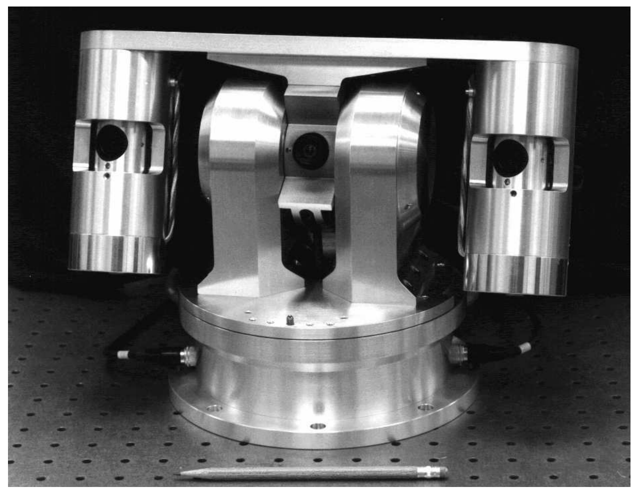

Fig. 1. TRICLOPS active vision system.

1: Introduction

There has been a great deal of interest recently in using movable camera systems for robotic vision applications. The motivations for using such an active or “animate” vision system include the ability to direct sensory and computational resources to regions of interest, and the simplification of a variety of vision computations. These and other motivations for using actively-controlled camera systems have been thoroughly discussed in the literature [4][7][8][20]. In order to support active vision research at the National Institute of Standards and Technology (NIST), a new trinocular robot head, The Real-time, Intelligently ControLled, Optical Positioning System (TRICLOPS), was designed and built. A photo of TRICLOPS is shown in Figure 1. This paper discusses the design features and performance characteristics of TRICLOPS, as well as some experimental results which have been achieved to date.

One of the earliest active vision systems was built at the University of Pennsylvania [15]. The system has a pair of movable cameras with computer-controlled zoom, focus, and aperture along with ten variable-intensity lamps. The positioning mechanism of the system provides two translational degrees of freedom, pan and tilt of both cameras, and coupled symmetric vergence.

At other institutions, such as the University of Illinois, researchers have assembled dynamic stereo camera systems using commercial motion control components, such as rotational and translational stages [1]. Such systems can provide good repeatability and relatively high-speed motions. However, because they are constructed from general-purpose positioning units, they cannot easily be configured to meet special requirements (such as the ability to position the focal point of the camera lenses at the intersection of perpendicular motion axes).

A different approach to aiming cameras is used in the MIT Vision Machine [18]. The Vision Machine consists of two cameras rigidly mounted on a movable platform. Each of the cameras is equipped with a motorized zoom lens, allowing control of the iris, focus, and focal length. Because of the size and weight of the lenses, camera movement is achieved indirectly by pivoting a front surface mirror mounted in front of each lens.

Binocular active vision systems have also been built at the University of Rochester and Harvard University. The Rochester head has independent vergence axes and a coupled tilt axis, and achieves camera rotational velocities of 5.2 rad/s and positioning accuracy of 2.4 mrad [10]. The

Harvard Head is a mobile binocular camera system with computer-controlled aperture and focus, and three mechanical degrees of freedom (pan, tilt, and antisymmetric vergence [12].

The head-eye system developed at the Royal Institute of Technology (the KTH head) in Sweden provides individual pan and tilt control for each of the two cameras, along with a two degree-of-freedom neck and computer-controlled baseline separation [17]. The system also has zoom, aperture, and focus control. The resolution on eye and neck axes is 126 µrad, and maximum rotational velocities are 3.14 rad/s. An interesting feature of the KTH head is a mechanism which automatically compensates for shifts in the position of the lens centers caused by changes in the lens focal length and focusing distance.

Another recent development in the area of active vision hardware is the Spherical Pointing Motor [9]. This device is a small, lightweight pan-tilt mechanism which can be used to move a small camera. It is controlled open-loop and is capable of rotational velocities of up to 10.5 rad/s. Accuracies of 2.6 mrad are achieved. These devices show promise for future small, low cost active vision systems.

Although the above systems represent significant advances in the development of hardware for active vision, no single design incorporated all of the features we desired. These include human-like dynamic performance and range of motion, independent control of vergence degrees of freedom for asymmetrical vergence, ability to position vergence lens focal points at the intersection of vergence and tilt axes, multiresolution imaging capability, high positioning resolution, repeatability, and accuracy, modular design, the use of industrial-quality components, and the ability to withstand the rigors of experimentation. In addition, a completely open controller based on the NASA/NIST Standard Reference Model (NASREM) Architecture [2] was desired to provide total freedom in modifying system parameters and algorithms. In meeting these requirements, TRICLOPS represents yet another step in the evolution of high-performance camera-pointing mechanisms.

The next section describes the details of the mechanical design of TRICLOPS. The control system and performance characteristics of TRICLOPS are discussed in Section 3. Finally, the results of visual tracking experiments are presented in Section 4.

2: TRICLOPS design

One of the primary design goals for TRICLOPS was to build a system that has very high dynamic performance—on a par with the human oculomotor system. There are several reasons for putting a large emphasis on being able to move the cameras very quickly. First, one of the primary advantages of active vision systems is to be able to use camera redirection to look at widely separated areas of interest at fairly high resolution. As is the case with the mammalian visual system, it is desirable to be able to move from one area of interest to the next as fast as possible, so that the time spent redirecting attention is minimized and the time spent acquiring useful image data is maximized.

In addition, it was desired to build a device for which mechanical performance would not be a limiting factor in the development of algorithms for visual tracking and autonomous information-gathering tasks. Although the current state of image processing results in latencies which are typically on the order of hundreds of milliseconds, rapid progress is being made. We would like to be able to take full advantage of anticipated improvements in image processing hardware, which will continue to increase update rates and reduce latency.

The goal of maximum dynamic performance leads to the selection of "micro-miniature" CCD cameras, with manually-adjustable focus, aperture, and focal length, since motorized lenses tend to be quite bulky. This decision also reduces the complexity of camera calibration considerably, since optical parameters do not change while the device is in use. Also, using manually-adjustable lenses simplifies achievement of the goal of placing the focal point of the lens near the center of rotation. The focal point of motorized zoom lenses shifts a great deal as the focal length is changed, and, although it is possible to compensate for this by moving the camera and lens as the zoom is adjusted [17], it results in considerable additional complexity. Many useful tasks can be performed with fixed imaging parameters. However, should computer control over these functions be desired, it is possible to mount motorized lenses to the TRICLOPS cameras. Of course, this would involve compromises in other areas, particularly in terms of dynamic performance and rotation about the focal point.

Another desirable capability is to be able to view the world at different resolutions simultaneously. This is a characteristic feature of human vision, which exhibits a resolution ratio of about 10:1 between the fovea centralis and the periphery (at 0.17 rad eccentricity) [11]. Although several previous active vision systems have employed zoom lenses for varying the image resolution, none have incorporated an additional camera(s) for simultaneous multiresolution sensing. There is a significant (and obvious) advantage to using a separate camera for this purpose, as opposed to using motorized zoom lenses—the wide-angle view is always available to alert the system to new agents which may enter the scene at any time.

To provide simultaneous multiresolution (or *foveal-peripheral*) capability, TRICLOPS has been designed with an integral third center camera which has wide angle (3-4 mm focal length) lens. This color camera serves the purpose of identifying features of interest in the wide field of view.

The center wide-angle lens is used in conjunction with 15-24 mm lenses on the vergence cameras, to provide a maximum resolution ratio of 6.17:1. The system is currently configured with a 4 mm center lens (1.35 rad field-of-view) and 15 mm vergence lenses (0.42 rad field-of-view).

As shown in Figure 7, TRICLOPS has four mechanical degrees of freedom. The four axes are arranged in the following kinematic configuration: 1) pan (or base rotation) about a vertical axis through the center of the base, 2) tilt about a horizontal line that intersects the base rotation axis, and 3) left and right vergence axes which intersect and are perpendicular to the tilt axis (Figure 7). The three mounting locations for cameras are also shown, one at the intersection of the pan and tilt axes, and two more at the intersections of the tilt and vergence axes.

The dynamic performance, accuracy, and other requirements are achieved with a direct-drive design. In many ways, a robot head is an ideal application for a direct-drive system. There is a requirement for large accelerations, low friction, and minimal transmission errors. These are three of the primary characteristics of systems which use motors and feedback devices mounted directly to the axes of motion. Transmission compliance and backlash, which can cause inaccuracy and oscillations, are eliminated. Also, the large gravity torques which cause limitations with direct-drive robot *arms* do not exist for a robot *head* with the configuration shown; the loads on the axes are almost completely inertial. In addition, the use of frameless DC motors and brushless, frameless resolvers provides flexibility in the routing of camera and control wiring. Much of the wiring in TRICLOPS is routed through the centers of these components to minimize external wiring. This results in well-behaved cable motions during high-speed movements and reduces the likelihood of pulling out cables.

The resolvers used for position feedback provide absolute position information, which eliminates the need to perform any homing or other initialization procedures each time the device is powered up (as would be required with incremental optical encoders, for example). The resolvers are used with 16-bit resolver-to-digital (R/D) converters, resulting in a position sensing resolution of 96 μrad.

TRICLOPS is built up of three independent, detachable modules (see Figure 7). These are the pan/tilt module (which forms the base), and two vergence modules. These modules can be used separately or together as a complete system. The vergence modules are connected to the pan/tilt module by a support block and top connector bar. The baseline separation between vergence modules is determined by the length of this bar. The current (and minimum) baseline separation is 0.2794 m.

All axes have detents which may be used with retractable ball spring plungers to provide static positioning of the shaft at known fixed locations. This is useful for determining the zero offset of the position transducer, and for other calibration procedures. Another important feature of TRICLOPS is that all axes have cushioned hard stops which prevent damage to the system when experimenting with new (and potentially unstable!) algorithms. Each of the TRICLOPS cameras is mounted in a Delrin tube which has screws for fine-adjustment of axial position and roll.

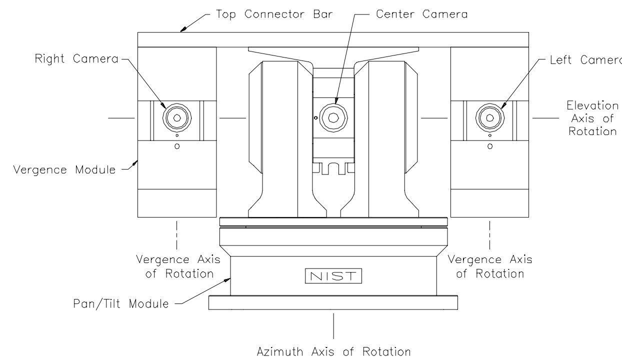

Fig. 2. TRICLOPS degrees of freedom.

The performance specifications which have been achieved with TRICLOPS are summarized in Table 1. As a point of comparison, rapid (saccadic) eye movements in humans can reach peak velocities of 10.5 rad/s and accelerations of 611 rad/s^2 [22]. It is seen that TRICLOPS is capable of motions which are comparable to those of the human

Table 1. TRICLOPS Specifications.

Range of motion:	
Pan	+/-1.68 rad (+/-96.3 deg)
Tilt	0.48 rad down, 1.14 rad up (+27.5 deg, -65.3 deg)
Vergence	+/-0.77 rad (+/-44 deg)
Peak acceleration:	
Pan	70 rad/s^2 (4010 deg/s^2)
Tilt	320 rad/s^2 (18,300 deg/s^2)
Vergence	1100 rad/s^2 (63,000 deg/s^2)
Peak velocity:	
Pan	11.5 rad/s (660 deg/s)
Tilt	17.5 rad/s (1000 deg/s)
Vergence	32 rad/s (1830 deg/s)
Axis position sensing resolution:	96 μrad (0.0055 deg)*
Positioning repeatability:	+/-96 μrad (0.0055 deg)
Interocular distance:	279.4 mm (11.0 in)**
Vergence-center resolution ratio	6.17:1 (max) 3.2:1 (current)
Overall width:	349 mm (13.75 in)
Overall height:	248 mm (9.75 in)
Approximate weight:	17.3 kg (38 lb)

* With 16-bit resolver-to-digital converter

** Distance is determined by length of connector bar—longer distances are possible oculomotor system in terms of peak velocity and acceleration. The range of motion for each of the TRICLOPS degrees of freedom is also comparable to corresponding human motion.

3: Control and image processing system

A multiprocessing controller based on the NASREM Architecture [2] has been constructed to control TRICLOPS and to perform sensory processing and world modeling functions on camera image data. The NASREM Architecture is a hierarchical control system architecture which divides the sensing and control problem into discrete levels. At the lowest level are the servo control loops for the device actuators; at higher levels, groups of sensors and actuators are coordinated as equipment subsystems. The NASREM hierarchical control system provides a modular, easily-modifiable infrastructure for implementing and testing active vision algorithms.

The processes and operation of the TRICLOPS motion control and image processing systems are described in [14]; a brief summary is presented here. The full potential of an active vision system cannot be realized without real-time image processing capability. In our lab, this capability is provided by a Pipelined Image Processing Engine (PIPE)[1] [6] for low-level image processing, along with a multiple-cpu VME system for higher-level sensory processing and world modeling (Figure 7). For controlling the motions of the TRICLOPS device, the architecture consists of a Primitive Level for producing coordinated motions of all axes, and a Servo Level for servoing the axes to commanded joint positions. The Primitive and Servo processes are implemented on five processor boards in a separate VME system, also shown in Figure 7.

The Primitive Level provides a number of different trajectory algorithms for performing high-speed saccades and for testing purposes. There are also vision-based tracking algorithms for tracking object features. The Primitive Level is capable of accepting commands from an operator through a user interface, or from another control system. This allows the device to either be used as part of an integrated robot system or as a stand-alone device.

The Servo Level processes are executed at a 2 KHz loop rate, which allows a closed-loop position bandwidth of 50 Hz to be achieved for the vergence axes with a standard PD algorithm. The closed loop bandwidths of the pan and tilt axes for this algorithm are about 4 Hz and 12-13 Hz, respectively. The Servo Level also provides other algorithms, such as PID control with velocity and acceleration feedforward. Servo parameters, such as servo gains, can be changed every cycle if desired.

All interprocess communication of command, status, and feedback information occurs via VME common memory interface buffers. An external host computer can therefore obtain position and velocity feedback and provide commands to the system at either trajectory generation or servo levels.

The joint trajectory following performance achieved by the system is demonstrated in Figure 7. This plot shows the desired and actual joint positions for quintic polynomial trajectories that move each joint at close to its maximum velocity. To obtain this data, the trajectory generation algorithm was modified to allow different motion times for each degree of freedom. The motions for all axes were performed and recorded simultaneously. For the vergence joint, the motion is from -0.75 rad to 0.75 rad in 0.091 s. This requires a peak acceleration of over 1000 rad/s^2 and a peak velocity of more than 30 rad/s. For the tilt axis, the motion is from -1.10 rad to 0.45 rad in 0.167 s. The peak commanded velocity and accelerations for this motion are 17.4 rad/s and 320 rad/s^2. The base pan motion of Figure 7 goes from -1.57 rad to 1.57 rad in 0.523 s. This motion requires maximum velocities and accelerations of 11.3 rad/s and 66 rad/s^2. The commanded and feedback velocities for these motions are shown in Figure 7.

As the plots show, the joint feedback positions follow the desired trajectories very closely. This degree of performance is the result of the high bandwidth of the system, as well as the use of velocity and acceleration feedforward terms. These motions were performed using a PID servo algorithm. The integral term counteracts the effects of cable-twist and friction torques which would otherwise cause steady-state position errors.

4: Tracking experiments

One of the most basic capabilities which may be provided by an active vision system is to track an object, using visual information to direct the gaze and keep the object in the field of view of the cameras. This section will briefly describe one of the visual tracking algorithms which has been implemented on TRICLOPS, and will summarize the results which have been obtained to date. Further details may be found in [14].

In the current tracking experiments, the emphasis has been placed on maximizing the tracking performance in terms of speed (bandwidth). As such, the visual processing has been simplified to the minimum required functionality to reduce the image processing delays as much as possible. Consequently, the image processing used for these tracking experiments consists of computing the centroid of a thresh-

1. Commercial equipment and materials are identified in this paper in order to adequately specify the experimental procedure. Such Identification does not imply recommendation or endorsement by NIST, nor does it imply that the materials or equipment identified are necessarily the best for the purpose.

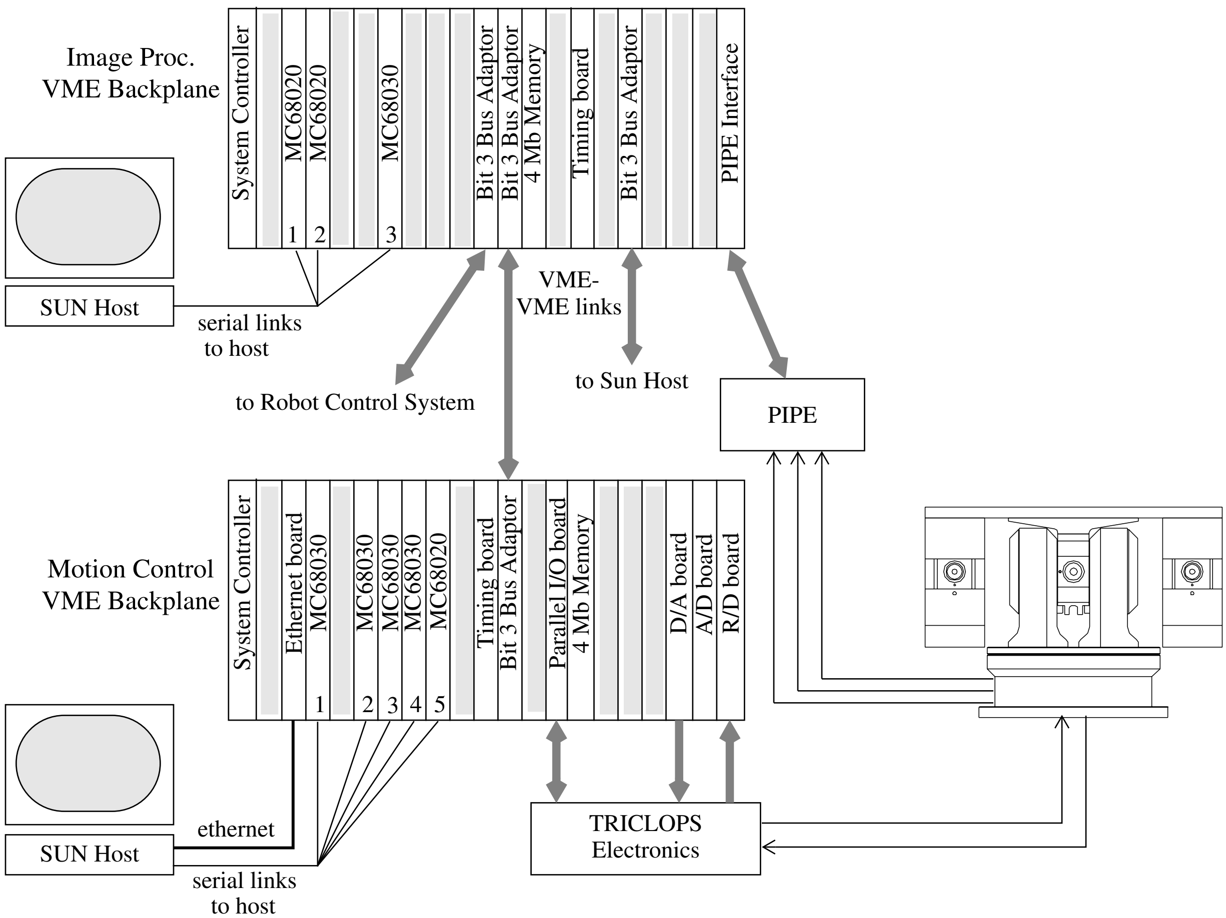

Fig. 3. TRICLOPS control and image processing hardware.

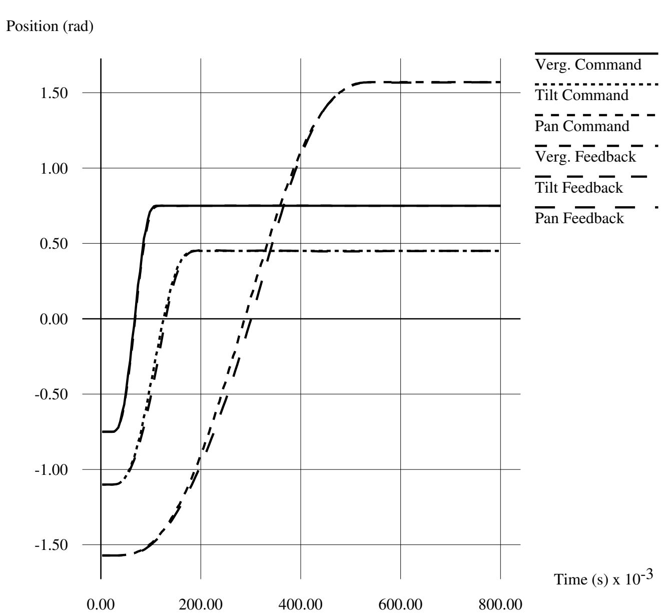

Fig. 4. High-speed trajectory following.

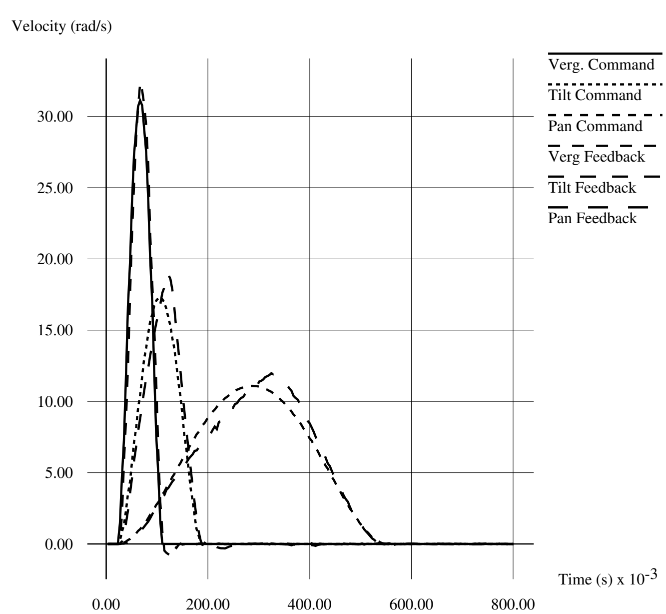

Fig. 5. Velocities during high-speed motions.

olded intensity image of a single object which presents a high-contrast image to both the left and right vergence cameras. For these experiments, only the two vergence cameras are used. The cameras are calibrated using the method proposed by Tsai [21].

In the current approach to tracking, 3D world position-based motion modeling of the target is used with prediction to generate goal fixation points. A block diagram of the tracking control system is presented in Figure 6. New centroid data are obtained from the world model every 1/30 s. The cameras are used in non-interlaced mode, and only information from the odd fields is used. Delayed joint positions which correspond in time to the centroid data are also read in at this time. The amount of delay is equivalent to the image acquisition and processing latency (about 0.084 s for the processing currently used for visual tracking). The delayed positions are supplied by the Servo processes along with the most recent feedback data. The centroid information is corrected for lens distortion, and the corrected centroids are used along with the delayed joint positions to compute the position of the object with respect to the TRICLOPS world reference frame using a triangulation algorithm.

Because of the image processing delay, the computed object position represents the position at time t-0.084 s, where t is the current time. If this position is used directly as the goal for tracking, very poor performance will result. The need for prediction to compensate for image processing delays in real-time visual servoing tasks is well known (see [3][10][19]). In contrast with the α–β–γ filters used for prediction in [3] and [10], the prediction used in the current experiments is based on performing least-squares fits of low-order polynomial equations of time to a limited sequence of object position data. These object motion functions are then evaluated at a time approximately equal to the estimated processing delay. Cubic polynomials which are updated using the 12 most recent object position estimates have worked well in the current experiments. The object motion functions are defined such that the result of the equation when evaluated at t = 0 gives the most recent object position estimate. This allows the pseudoinverse of the functions of time matrix to be constant (and able to be computed a priori). The least squares estimate can then be updated every centroid data cycle with a single matrix multiply.

After determining the predicted position of the object with respect to TRICLOPS, the desired position of the pan axis is determined. Several different means of distributing motion between the pan and vergence axes have been implemented and evaluated. These include not moving the base at all, moving the base to maintain symmetrical vergence, and moving the base some fixed percentage of the angle required for symmetrical vergence. One of the most effective methods consists of commanding the base to increment the command position by some percentage of the distance between the angle required for symmetrical vergence and the current command position. This has the effect of gradually moving the pan axis toward the direction of the object. As a result, vergence asymmetry is reduced and there is no “preferred position” of the base. The responsiveness of the base is determined by the size of the percentage parameter. A value of 0.1 is a good compromise between moving the base too sluggishly and putting too much of the tracking burden on the base.

Finally, the inverse kinematic equations are solved (with the desired base angle as an input parameter) to yield the desired joint angles. The desired joint angles are commanded to the Servo level, which executes the high-bandwidth PD servo algorithm to achieve them. New joint angle goals are computed on the same 30 Hz timing cycle as image centroids. These position commands are interpolated by the Servo Level into 0.0005 s subcommands. This interpolation greatly smooths the tracking motion, but it also incurs some additional delay which must be accounted for in the prediction of object position. This delay, combined with the image processing delay, results in a total delay of about 0.117 s. However, the prediction procedure itself introduces a slight frequency-dependent phase lead, such that using a time of 0.090 s in the prediction equations yields better overall tracking than using 0.117 s.

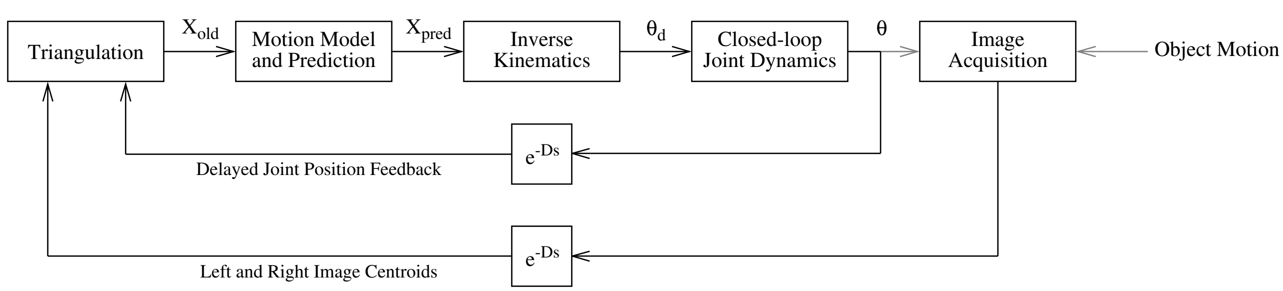

Fig. 6. Block Diagram of Tracking Algorithm.

Several experiments have been performed to evaluate the tracking performance attained using the above techniques with TRICLOPS. In one test, a 3.5 cm diameter white ball attached to the end of a 0.44 m rod is rotated in a horizontal circle whose center is 1.1 m in front of TRICLOPS. Figure 7 shows the set-up for this experiment.

The accuracy of the position estimates determined by TRICLOPS, with and without prediction, for a ball orbital velocity of about 6.9 rad/s (tangential velocity = 3 m/s) is shown in Figure 7. The solid line in the figure indicates the actual path of the ball during a single revolution, as determined by careful manual measurement. The smooth dotted line close to the actual path is the estimated path of the ball without prediction. It is seen that the spatial error in the estimated path without prediction is quite small. However, because this estimate does not include prediction, TRICLOPS can only track the ball up to an orbital velocity of 1.8 rad/s using this estimate directly.

While the predicted ball positions in Figure 7 have slightly increased spatial errors, tracking speed is greatly improved due to the compensation for the processing delay. When the prediction is added, TRICLOPS can track the ball at velocities up to 8.6 rad/s.[1] Although not shown in the figure for reasons of clarity, the path actually tracked by TRICLOPS (obtained by applying forward kinematics to the time history of joint positions) follows the commanded path quite closely. For this orbital velocity, the magnitude of the velocity of the vergence axes reaches a maximum of about 6.5 rad/s, and the magnitude of the peak base rotation velocity is about 1.7 rad/s. The ball is also tracked smoothly and stably when unpredictable motions of the ball at the end of a hand-held rod are performed.

1. At speeds this high, there is some overshoot in the predicted path due to the frequency response characteristics of the prediction.

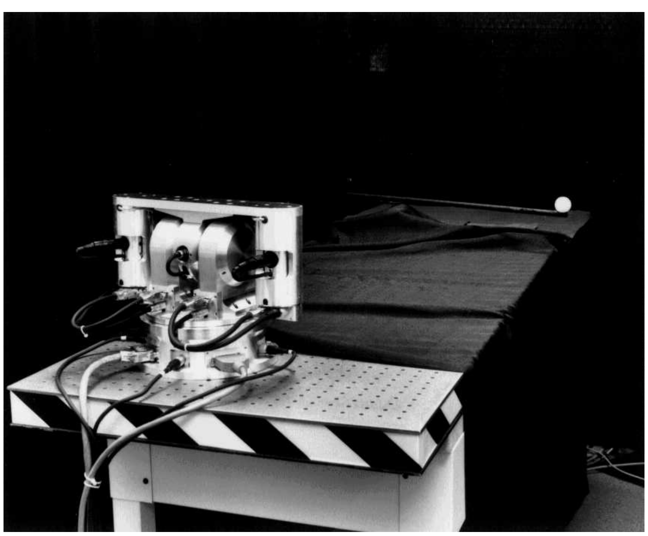

Fig. 7. Experimental set-up for circular tracking.

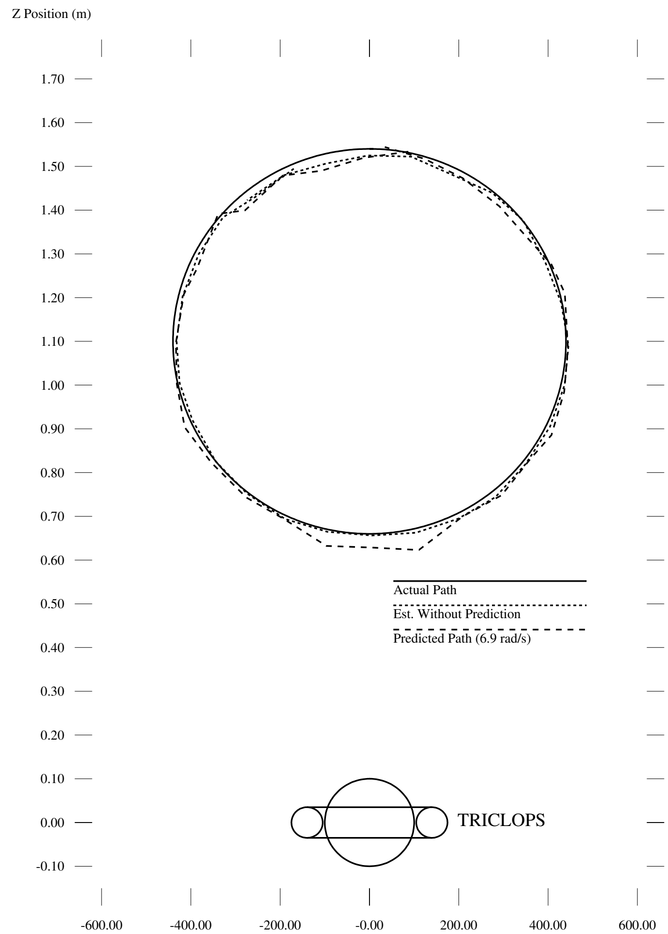

Fig. 8. Estimated ball path with and without prediction.

As would be expected, the geometric error increases with increasing distance from the head. The maximum error of the estimated path without prediction is about 0.025 m at the far side of the circle. This accuracy, particularly in range, is target-dependent, however. The data of Figure 7 were obtained during tracking in ambient light. Accuracies on the order of a few millimeters have been obtained when a special target (for example, an internally-illuminated ball) is used to produce better quality images.

5: Conclusions

Although it has been only about six years since the development of the first animate vision systems patterned after the human eye-head arrangement, much progress has been made. A significant recent development is the commercial availability of micro-miniature CCD cameras, which has enabled the goal of human-like dynamic performance to be achieved.

TRICLOPS is a new direct-drive active vision system which takes advantage of these imaging devices to enable very large rotational velocities and accelerations. TRICLOPS is capable of peak velocities of over 32 rad/s for the vergence degrees of freedom, 17.5 rad/s for the tilt axis, and 11.5 rad/s for pan rotations. The visually-based closed-loop performance of the system has been demonstrated by implementing a simple tracking application. Using triangulation and 3D motion modeling and prediction, TRICLOPS is able to keep a ball in the field of view of the vergence cameras for tracking velocities in excess of 6 rad/s.

TRICLOPS provides concurrent multi-resolution sensing capability by locating an integral third wide-angle view camera between the two vergence cameras. Other unique features of this system include modular construction, the use of industrial-quality components, and the incorporation of design elements which decrease the possibility of damaging the system during experimentation.

We have just begun to explore the potential of TRICLOPS. Future plans include using the center wide-angle view camera in conjunction with algorithms which use color and motion to autonomously direct the attention of the system. We also intend to connect the TRICLOPS control and image processing/world modeling systems with another existing NASREM system used to control a seven degree-of-freedom manipulator [13]. This will provide a very useful platform for studying visually-guided manipulation tasks, such as assembly. TRICLOPS will also be used in the development of active vision applications in inspection, surveillance, and mobile navigation and exploration.

The high-speed motion, multiresolution imaging, and fast visual tracking capabilities of TRICLOPS are demonstrated on an accompanying videotape [16].

References

[1] Abbott, A. L., Ahuja, N. 1990. Active surface reconstruction by integrating focus, vergence, stereo, and camera calibration. In *Proc. Third International Conference on Computer Vision*, Osaka, Japan, pp. 489 - 492.

[2] Albus, J. S., McCain, H. G., Lumia, R. 1987. NASA/NBS Standard Reference Model for Telerobot Control System Architecture (NASREM). National Inst. Standards and Technol., Gaithersburg, MD, Technical Note 1235.

[3] Allen, P., Yoshimi, B., Timcenko, A. Real-time visual servoing. In *Proc. DARPA Image Understanding Workshop*, Pittsburgh, PA, 1990.

[4] Aloimonos, J., Weiss, I., Bandyopadhyay, A. 1987. Active vision. In *Proc. First International Conference on Computer Vision*, pp. 35-54.

[5] Aloimonos, J., Shulman, D. 1989. *Integration of Visual Modules: An Extension of the Marr Paradigm*. Academic Press, Inc.: New York, NY.

[6] Aspex, Inc. 1987. PIPE--An introduction to the PIPE system. New York, NY.

[7] Bajcsy, R. 1988. Active perception. *Proceedings of the IEEE* 76(8): 966-1005.

[8] Ballard, D. H. 1991. Animate vision. *Artificial Intelligence*, 48:57-86.

[9] Bederson, B. B., Wallace, R. S., Schwartz, E. L. 1992. Two miniature pan-tilt devices. In *Proc. IEEE Conference on Robotics and Automation*, pp. 658 - 663.

[10] Brown, C. M. 1989. Kinematic and 3D motion prediction for gaze control. In *Proc. Workshop on Interpretation of 3D Scenes*, Austin, TX, pp. 145-151.

[11] Carpenter, R. H. S. 1988. *Movements of the Eyes*. Pion Limited: London, England.

[12] Clark, J. J., Ferrier, N. J. 1988. Modal control of an attentive vision system. In *Proc. 2nd International Conference on Computer Vision*, pp. 514 - 523.

[13] Fiala, J., Wavering, A., Lumia, R. 1992. A manipulator control testbed: Implementation and applications. In *Proc. AAS Conf. on Guidance and Control*, Keystone, CO.

[14] Fiala, J., Lumia, R., Roberts, K., and Wavering, A. TRICLOPS: A tool for studying active vision. Submitted to the *International Journal of Computer Vision*.

[15] Krotkov E., Fuma, F., Summers, J. 1988. An agile stereo camera system for flexible image acquisition. *IEEE Journal of Robotics and Automation* 4(1): 108 - 113.

[16] Lumia, R. et al 1993. TRICLOPS: A tool for active vision. Video tape. In *Video Proc. IEEE Conference on Robotics and Automation.*

[17] Pahlavan, K., Eklundh, J. O. 1990. A head-eye system for active, purposive computer vision. KTH, Stockholm, Sweden, TRITA-NA-P9031.

[18] Poggio, T., et al. 1988. The MIT vision machine. In *Proc. DARPA Image Understanding Workshop*, pp. 177 - 198.

[19] Robinson, D. A. 1987. Why visuomotor systems don't like negative feedback and how they avoid it. In Michael Arbib and Allen Hanson, eds, *Vision, Brain, and Cooperative Computation*. MIT Press: Cambridge, MA.

[20] Swain, M. J., Stricker, M., eds. 1991. Promising directions in active vision. Report from the NSF Active Vision Workshop, University of Chicago.

[21] Tsai, Roger Y. 1987. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics and Automation* RA-3(4): 323-344.

[22] Wurtz, R. H., M. E. Goldberg, eds. 1989. *The Neurobiology of Saccadic Eye Movements*. Elsevier: New York, NY.

Acknowledgments

We would like to thank our coworkers for their valuable contributions, including Tom Wheatley for building the TRICLOPS electronics, Marilyn Nashman for developing parts of the visual processing software, Adam Jacoff for aiding in the design and construction of TRICLOPS, and Henry Schneiderman for ideas and suggestions relating to prediction and curve-fitting.