DYNAMIC ISSUES IN ROBOT VISUAL-SERVO SYSTEMS

Peter I. Corke
CSIRO Division of Manufacturing Technology
P.O. Box 883, Kenmore, 4069
AUSTRALIA.
pic@brb.dmt.csiro.au

Abstract

This paper poses a number of questions related to the performance and structure of closed-loop visual control, or visual servo, systems. While the fundamentals of visual servo control are well known and systems have been demonstrated for many years, the achieved performance is far less than could be expected. In particular the questions discussed relate to fundamental control system structure, performance metrics, compensator design and the choice of feedback versus feedforward control.

1 Introduction

The use of visual sensors with robot manipulators has a history of over two decades, dating back to early work with block worlds. However the performance of modern closed-loop visual control, or visual-servo, systems is significantly poorer than would be expected from a control systems point of view. In this paper it is argued that the reason for such poor performance is that the dynamics of the system, vision and robot, have been ignored. To this end, a distinction is introduced between *visual servo kinematics* and *visual servo dynamics*. The former, essentially a "pure computer-vision" problem, is well addressed by the literature, while the latter has been largely ignored. Investigating dynamics, rather than just the kinematics, of visual servo systems leads to advantages beyond the obvious ones of tracking objects whose pose is unknown and possibly time varying. Visual servoing admits the possibility of "closing the loop" on the dynamics of the actuator, transmission and structural compliance, allowing superior control of end-point dynamics than can be achieved using axis sensing alone.

Related research [1] introduces the concept of *visual dynamic control* which is concerned with dynamic effects due to the manipulator and machine vision sensor which limit performance and must be explicitly addressed in order to achieve high-performance control. That research takes a very different approach to that done by others in the field, in emphasising temporal and dynamic issues over image processing. "Simplistic computer vision" techniques and a specialised hardware architecture combine to provide a high sample-rate and minimum latency vision system. This is an ideal testbed with which to bring closed-loop dynamic problems to the fore, conduct modelling and experiment with various control strategies.

One problem that dogs any discussion of visual servo performance is the lack of agreed, quantitative, performance metrics. It is actually relatively simple to build a visual closed-loop robot system which will converge on a target, if the controller is sufficiently detuned to ensure stability. Most researchers propose the use of simple proportional controllers but it can be shown that these have two undesirable properties. Firstly, feedback-only controllers have a loop-gain limit due to the open-loop delay, and this limits the achievable closed-loop bandwidth. Secondly, simple feedback controllers have significant phase lag characteristics which lead to poor tracking. Much research has focussed on solving the computer-vision part of the problem, but the control or dynamics problem is generic but not adequately investigated.

The remainder of this section discusses existing research on visual servoing from a control systems perspective, and then describes the experimental platform on which the research described in this paper was conducted. Section 2 then poses a number of questions which seem to have been overlooked by the research community, and offers some observations and partial answers. Finally Section 3 provides some conclusions.

1.1 Research to date

Research in visual servoing has, to date, focussed largely on the visual kinematic issues such as estimating object pose from n points, or analysing the relationship between Cartesian motion and the corresponding image plane motion as encapsulated in an image Jacobian matrix. These computer vision problems underpin, respectively, the *position-based* and *image-based* approaches to visual servoing [2]. Both approaches are well known and have been demonstrated experimentally under laboratory conditions. It could perhaps be argued that the problem is solved for all practical purposes, and the computational costs of these two approaches are comparable and readily achieved. Control of a 6-DOF robot was demonstrated by Ganapathy [3] over a decade ago using a closed-form solution that executed in only 4 ms on a 68000 processor. An iterative approach has even been patented [4]. The 3-D pose estimation problem can be significantly simplified by using non-anthropomorphic direct 3-D sensors — those based on structured lighting are now compact and fast enough to use for visual servoing [5].

Recently, however, there has been a considerable growth in the literature which has been fuelled by personal computing power crossing the threshold which allows analysis of scenes at a sufficient rate to "servo" a robot manipulator. Prior to this, researchers required specialised and expensive pipelined pixel processing hardware.

An obvious characteristic of many reported visual servo systems is that they appear slow and "shaky", and exhibit significant lag with respect to the motion of the tracked object. This can be seen in video tape presentations and in the time axis scaling for many published results. Reported results and comments such as "*slight instability*" are clear indications of closed-loop systems close to their stability limit. Another common comment is of "*noticeable lag*" in the observed performance, which is indicative of a poorly designed control system, inappropriate control architecture or both. To the casual observer there are obvious similarities between the performance of such systems and that of early force controlled robot systems. The force control community learned, long ago, the importance of dynamic modelling and how the dynamics of the manipulator, sensor and environment impact on closed-loop stability [6]. In this paper it is argued that the visual-servo community has yet to learn this lesson.

Vision is admittedly a "difficult" sensor to use for closed-loop control due to its low sample rate, high latency (often one or more sample intervals) and low spatial resolution. However most reports demonstrate a closed-loop bandwidth which is only a fraction of that which would be expected. Evidence of stability problems, or tracking lag, has tended to be ignored or blamed on the low sample rate or low-bandwidth communications link between the vision system and robot. Unfortunately, rectifying these "deficiencies" will not solve a problem whose roots are in the fundamentals of the system dynamics. The evidence points to an underlying lack of control systems knowledge in this research community. This "dynamics problem" is generic, that is, its existence is independent of the type of image processing or scene interpretation performed. Its existence is due to the temporal characteristics of the camera and vision system, and the robot and its controller.

Problems of delay in visual servo systems were first noted over 15 years ago [7] but are not always taken into consideration. The sources of delay include the charge integration time within the camera, serial pixel transport from the camera to the vision system, and the computation time for feature extraction. If the target motion is constant then prediction can be used to compensate for the latency, but combined with low sample rate leads to poor disturbance rejection and long reaction time to target "maneuvers". Grasping objects on a conveyor belt is an ideal, though non-challenging, application for prediction since the target velocity is constant.

1.2 Experimental system

The research described [1] takes a different approach from others in the field, by emphasising temporal and dynamic issues over image processing. "Simplistic" computer vision techniques and a specialised hardware architecture combine to provide a high sample-rate and minimum latency vision system. This is an ideal testbed with which to bring closed-loop dynamic problems to the fore, conduct modelling and experiment with various control strategies. A block diagram of a 1-DOF system is given in Figure 1. The camera lens is modelled as a simple gain, K_{lens}, which, due to perspective, is a function of target distance. The system is multi-rate due to the different sample intervals for the robot's servo system and the camera's field rate. Manipulator structural dynamic effects are represented by $G_{struct}(s)$, which

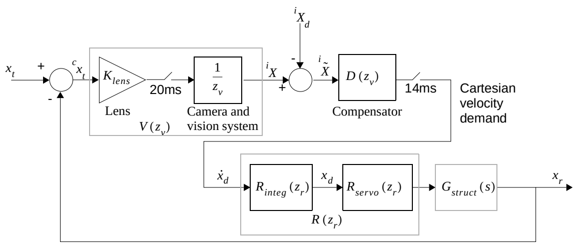

Figure 1: Block diagram of the 1-DOF visual feedback system. x_t is the world coordinate location of the target, and ${}^iX_d$ is the desired location of the target on the image plane. Note the two samplers in this system.

is assumed, for now, to be unity.

The performance of this system, with proportional control $D = K_p$, has been measured and Figure 2 shows a step response to a visual square wave excitation. The 'early' and 'late' response are related to when the visual stimulus occurs with respect to the system's samplers. The closed-loop frequency response is given in Figure 3 and shows evidence of a low frequency pole at approximately 2 Hz, and time delay. The phase is linear with respect to frequency and indicates a delay of 58.0 ms.

To simplify analysis, the system may be approximated by a single-rate model operating at the visual sampling interval of 20 ms. The open-loop transfer function may be written as

$$V(z)D(z)R(z) = \frac{K}{z^2(z-1)} \qquad (1)$$

where

$$K = \frac{T_v}{T_r} K_p K_{lens} \qquad (2)$$

and T_v and T_r are the vision and robot servo sample intervals respectively. The overall delay is Ord(denominator) − Ord(numerator), in this case three sample periods or 60 ms which agrees with previous observations of latency and the step response of this model.

The single-rate open-loop model (1) can be used for stability analysis as well as state-estimator and controller design using classical digital control theory.

The root-locus diagram in Figure 4 shows that for the known gain setting, $K = 0.132$, the poles are located at $z = 0.789$, 0.528 and -0.316. The dominant pole corresponds to a frequency of 1.9 Hz which is in good agreement with the estimate of 2 Hz made earlier. This pole is very sensitive to changes in loop gain such as those

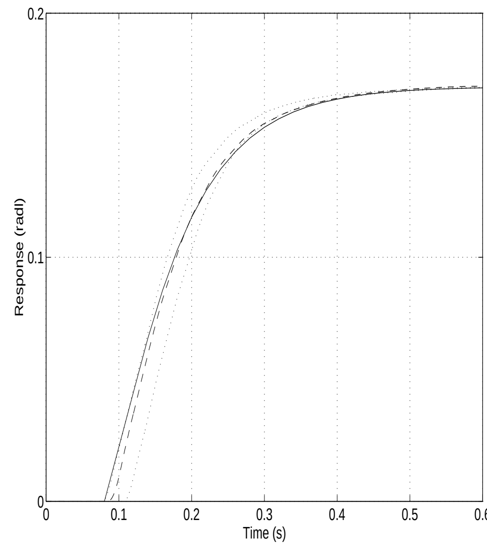

Figure 2: Comparison of measured and simulated step response. Measured early and late step responses (dotted), single-rate model (solid), and the multi-rate model (dashed).

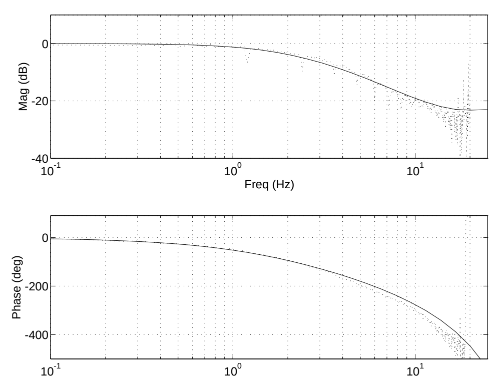

Figure 3: Bode plot of single-rate model closed-loop transfer function (solid) with measured frequency response (dotted) overlaid. Note the considerable phase error for moderate frequency demands, eg. 1 Hz. ($K_p = 1.3 \times 10^{-4}$, $K_{lens} = 708$)

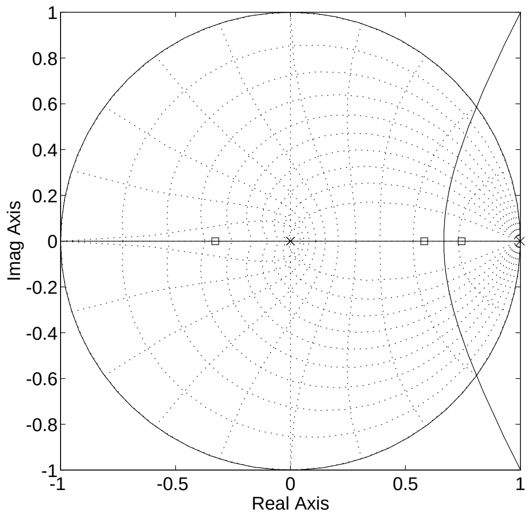

Figure 4: Root-locus of single-rate model. Closed-loop poles for gain $K_p = 1.3 \times 10^{-4}$ are marked ($K_{lens} = 708$).

caused by variation in target distance. Figure 3 shows the Bode plot of the single rate model overlaid on the experimentally determined frequency response function from above.

2 Some questions

This section poses a number of questions which need to be addressed in order to achieve high performance visual closed-loop control.

2.1 Is image feature trajectory generation necessary?

Most robot controllers accept absolute position setpoint commands but an end-effector mounted vision sensor provides relative position information. In general, due to limited torque capabilities, the robot will be unable to reach an arbitrary target position within one setpoint interval and thus requires some "planned" motion over many intervals. In robotics this is the well-known trajectory generation problem, but that approach has a number of disadvantages in this situation, since a moving target requires continual path adjustment [8] while also maintaining continuity of position and its time derivatives.

A far simpler approach is to consider the positioning task as a control problem:

> Perceived error generates a velocity demand which moves the robot toward the target.

This results automatically in target-following behaviour, but without the complexities of an explicit trajectory generator. This behaviour is analogous to the way we control cars and boats within an unstructured environment: we do not compute a sequence of precise spatial locations, but rather a velocity which is continually corrected on the basis of sensory inputs until the goal is achieved. Stated another way, we consider visual control as a *steering problem.*

2.2 What is the most appropriate axis control method to employ in a visual servo system?

At a very fundamental level a robot link is moved by an actuator that applies a torque or force. However it is common to provide local axis-level control loops which create the abstraction of a link whose position or velocity may be specified. Most industrial robot applications involve positioning tasks, bringing the tool to a known point in the workspace, and this requires position controlled axes. When the robot is in contact with a surface some positional degrees of freedom are lost and the robot must perform compliant motion requiring some, or all, joints to be torque (or force) controlled. In some teleoperation applications it is desirable to consider axes as being velocity controlled. Robot joints can thus be considered as torque, velocity or position "sources", and the choice is dependant upon the robot's application.

The argument in Section 2.1 would indicate that axis velocity control was the "natural" choice for closed-loop visual control. Surprisingly the majority of reported experimental systems implement a visual feedback loop "on top of" the robot's existing position-control loops. This has occurred, probably, for the pragmatic reason that most robot controllers present an abstraction of the robot as a position-controlled device. However some researchers in visual servoing who have "built their own" robot controllers at the axis level have all chosen to implement axis-position control [9–11]. The redundant levels of control add to system complexity and may impact on closed-loop performance [12]. Exceptions to this approach are the citrus-picking robot [13] which closes the vision loop about a hydraulic actuator which is a "natural" velocity source, and also [14, 15]. Weiss [2] proposed closing a visual-feature control loop around a torque servo but conducted only simulation studies.

Sending incremental motion commands to a robot position controller is essentially velocity control — the controller is integrating the increments to create an absolute coordinate as a setpoint for the position servos. From a control systems perspective, however, there are some subtle differences between a velocity controlled axis and an incremental position controlled axis. In particular the inner position loop will generally be very "tight", that is high-gain, in order to minimize position error. For an ideal robot this is not a problem, but it will exacerbate resonance problems due to structural and/or drive compliance. It has been shown [16] that for the case of significant structural compliance it is preferable to use a "soft" velocity loop in order to increase system damping. This is particularly desirable in the non-colocated control situation of an eye-in-hand camera.

In [1] the three axis control modes were compared for a simple linear discrete-time visual servo simulation model with proportional control. Torque control resulted in the slowest closed-loop poles, and position control the fastest, but with a steady state error to a step or higher order target motion. The addition of an extra integrator remedies this, but as discussed above, the structure is now essentially velocity controlled. However when a full non-linear motor model was simulated, the torque-mode controller did not behave as expected — simulation resulted in very lightly damped oscillatory responses. This was found to be due to a combination of non-linear friction, integral action and the relatively low sample rate. The unmodeled dynamics, particularly stick/slip friction, are complex and have time constants that are shorter than the visual sample interval. It is a "rule of thumb" that the sample rate should be 4 to 20 times the natural frequency of the modes to be controlled. Improved torque-mode control would require one or more of the following: a higher sample rate, friction measurement and feedforward, adaptive, or robust control. A more straightforward option is the use of axis velocity control. Velocity feedback is effective in linearising the axis dynamics and eliminating much of the parameter variation. Parameter sensitivity studies showed robustness with changes of up to $\pm 25\%$.

2.3 How should visual servo performance be evaluated?

Performance of the visual servo system can be expressed in terms of time or frequency domain metrics. In the frequency domain, performance is related to the closed-loop bandwidth. While "high closed-loop bandwidth" is desirable to reduce error it can also lead to a system that is sensitive to noise and unmodeled dynamics. Additionally, the notion of "bandwidth" is not straightforward since it implies both magnitude (3 dB down) and phase ($45°$ lag) characteristics. Time delay introduces a linear increase in phase with frequency, and the Bode plot of Figure 3 indicates $45°$ lag occurs at less than 1 Hz. For tracking applications phase is a particularly meaningful performance measure — for instance the citrus picking robot's tracking specification [13] was in terms of phase for a particular fruit swinging frequency.

Common time domain metrics include settling time, overshoot and polynomial signal following errors. However many published visual servo results should be subject to close scrutiny about what they really demonstrate. Any closed-loop system of Type 1 or greater will have zero steady state error to a step demand. Since, as described above, most visual servo systems incorporate at least one integrator (by virtue of incremental positioning), zero settling error is to be expected. In fact settling error is the *least challenging* test of control performance. For example, the feedback system introduced above has a fast square wave response (200 ms) but simulated tracking of sinusoidal motion in Figure 5 shows a robot lag of approximately $30°$ and the centroid error peaking at over 80 pixels.

A better metric for tracking performance would be related to tracking error such as peak-to-peak error, or

$$\eta = \sqrt{\frac{\displaystyle\int_{t_1}^{t_2} \left({}^iX(t) - {}^iX_d(t)\right)^2 dt}{t_2 - t_1}} \tag{3}$$

which is the RMS pixel error value over the time interval $[t_1, t_2]$. The choice of appropriate target motion with which to evaluate (3) depends upon the application. For instance, if tracking a falling or flying object it would be appropriate to use a parabolic input. A sinusoid is particularly challenging since it is a non-polynomial with significant and persistent acceleration that can be created readily in the laboratory using a pendulum or turntable.

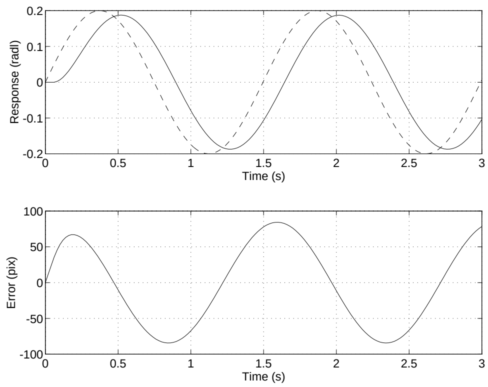

Figure 5: Simulated tracking performance of visual feedback controller, where the target is moving sinusoidally with a period of 1.5 s (0.67 Hz). Note the considerable lag of the response (solid) to the demand (dashed).

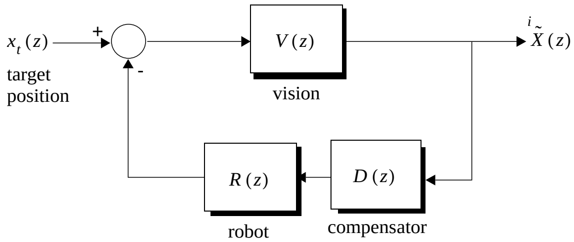

Figure 6: Block diagram showing target position as input and image-plane error as output.

2.4 What performance should be expected?

The camera frame rate is the sample rate of a visual servo system and this is a function of the standard video formats typically employed (CCIR or RS170). Most high-performance systems reported achieve field-rate processing at best, giving a sample rate of 50 or 60 Hz, which is a major performance limiting factor. Franklin [17] suggests that the sample rate of a digital control system be between 4 and 20 times the desired closed-loop bandwidth. Such a rule implies a closed-loop bandwidth of between 2.5 Hz and 12 Hz. Corresponding response times would be in the range 80 to 400 ms, which the experimental system achieves.

2.5 What are appropriate control laws to achieve high closed-loop performance?

In control system terms the visual servo system is a regulator (trying to maintain a constant image plane target centroid) and the target motion is a non-measurable disturbance input. Tracking performance can be expressed as a transfer function for image plane error response to target input, as shown in Figure 6

$$\frac{^i\tilde{X}(z)}{x_t(z)} = \frac{V(z)}{1+V(z)R(z)D(z)} \quad (4)$$

$$= \frac{K_{lens} z(z-1) D_D(z)}{z^2(z-1)D_D(z) + K_{lens} T D_N(z)} \quad (5)$$

where $^i\tilde{X} = {^iX} - {^iX_d}$ is the image plane error and ^iX_d is assumed, for convenience, to be zero.

The steady-state value of the image plane error for a given target motion, $x_t(z)$, can be found by the final-value theorem. For proportional control, and substituting previously identified models for $V(z)$ and $R(z)$, the steady-state error is

$$^i\tilde{X}(\infty) = \lim_{z \to 1} (z-1) \frac{K_{lens} z(z-1)}{z^2(z-1) + K_p K_{lens}} x_t(z) \quad (6)$$

which may be evaluated with the target motion chosen as one of

$$x_t(z) = \begin{cases} \frac{z}{z-1} & \text{for a step input} \\ \frac{Tz}{(z-1)^2} & \text{for a ramp input} \\ \frac{T^2}{2}\frac{z(z+1)}{(z-1)^3} & \text{for a parabolic input} \end{cases} \quad (7)$$

For the steady-state tracking error to be zero, the numerator of (6) must cancel the poles at $z = 1$ of $x_t(z)$ and retain a factor of $(z-1)$. This numerator comprises the poles of the robot and compensator transfer functions — the latter may be selected to give the desired steady-state error response by inclusion of an appropriate number of integrator terms, or $(z-1)$ factors. The number of such open-loop integrators is referred to as the *Type* of the system in classical control literature. Equation (6) has one open-loop integrator and is thus of Type 1 giving zero steady-state error to a step input. A ramp input however will result in a finite error of T/K_p pixels which can be seen as the substantial lag in Figure 5. Also note that the numerator of the closed-loop transfer function (5) has a zero at $z = 1$ which acts as a differentiator. From (5) compensator poles will appear as additional closed-loop zeros.

Common approaches to achieving improved tracking performance are:

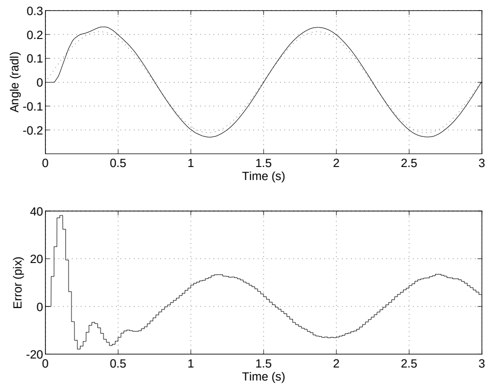

Figure 7: Simulation of pole-placement fixation controller, showing response (solid) and demand (dotted). Closed-loop poles at $z = 0.6 \pm 0.4j, 0.4$.

1. Increasing the loop gain, K_p, which minimises the magnitude of ramp-following error. However for this system, as shown in the root-locus plot of Figure 4, this approach is limited, since the poles leave the unit circle as gain is increased. It is not possible to achieve the stable high-gain control required for accurate tracking.

2. Increasing the Type of the system, by adding additional open-loop integrators.

3. Introducing feedforward of the signal to be tracked. This is an important issue in machine-tool control and has been discussed, for instance, by Tomizuka [18]. This approach is discussed in Section 2.7.

A number of control laws including PID, Smith's predictor, pole-placement, and LQG, were synthesised and evaluated in both simulation and experiment. In order to achieve good sinusoid tracking performance, it was found necessary to increase the Type of the open-loop system to 2. However the location of the closed-loop poles is fundamentally constrained if a stable compensator is required. As closed-loop bandwidth is increased the robustness with respect to plant parameter variation, in particular saturation, is greatly reduced.

Figure 7 shows the simulated closed-loop response for sinusoidal target motion, a detailed non-linear system model, and a pole-placement controller. The peak-to-peak error in pixel space is approximately 28 pixels — clearly a great improvement over simple proportional control — and the best performance achieved in this work. However a number of characteristics of the response deserve comment:

- The error is greatest at the peaks of the sinusoid, which correspond to greatest acceleration. The Type 2 controller has zero error following a constant velocity target (which the slopes of the sinusoid approximate), but will have finite error for an accelerating target. That performance is dependent on the closed-loop pole locations which cannot be made arbitrarily fast.

- There is considerable delay at $t = 0$ before the robot commences moving, due to delay in the vision system and controller.

- Once again this is a Type 2 system and the closed-loop response will include a double differentiator. This effect is evident in the initial transient shown in the simulated closed-loop response of Figure 7.

The closed-loop poles cannot be made arbitrarily fast and are constrained by the practical requirement for compensator stability. Linear discrete-time simulations indicate no difficulty with an unstable compensator, but non-linear simulations show that the closed-loop system will be unstable. This is believed to be due to non-linearities in the plant which change the apparent dynamics at saturation which is exacerbated by high controller gains. The addition of integrators, while raising the system Type, places stringent constraints on closed-loop and estimator pole locations.

The Smith's predictor also warrants comment since it is commonly cited by the active vision community. The essential characteristic of this predictor is cancellation of open-loop poles. In linear simulation this approach works well, but in non-linear simulation or experiment it lacks robustness. Stability can be achieved but only by choosing "slow" closed-loop poles.

2.6 What are the limitations to closed-loop performance?

For a feedback system such as described so far the limiting factor is the low loop gain necessary to ensure stability in the presence of open-loop delay. The architecture of the system, and in particular communications between the vision and robot control subsystems should be designed in

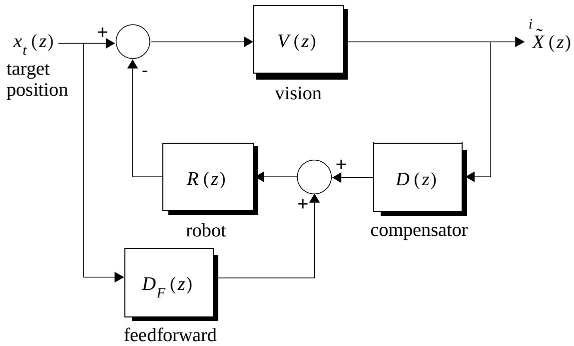

Figure 8: Block diagram of visual servo with feedback and feedforward compensation.

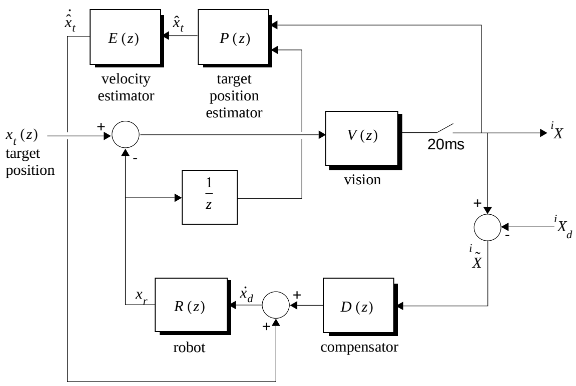

Figure 9: Block diagram of velocity feedforward and feedback control structure as implemented. $P(z)$ estimates target bearing.

order to minimize this delay. Prediction can be used to "cover" system delay, but at the expense of poor disturbance rejection with respect to unmodeled target motions. More sophisticated compensators can be employed, but they remain fundamentally constrained by the requirements for compensator stability and robustness.

2.7 Why are most visual servo systems based only on feedback control?

The term "visual servo" is almost synonymous with "visual feedback" but this need not be so. The design constraints inherent in feedback-only control lead naturally to the consideration of feedforward control strategies, since additional design degrees of freedom are available by manipulating the system zeros as well as the poles. The feedforward quantity, target position or velocity in this case, is not directly measurable but can be estimated. The feedforward term is shown in Figure

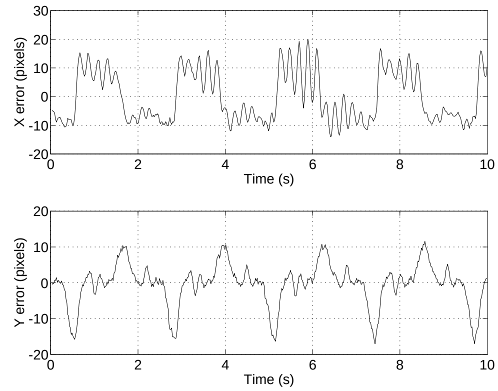

Figure 10: Measured tracking performance, centroid error, for target on turntable revolving at 4.2 rad/s.

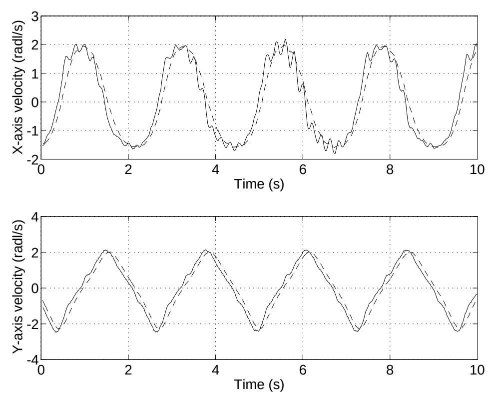

Figure 11: Measured tracking v velocity for target on turntable revolving at 4.2 rad/s, showing axis velocity (solid) and estimated target velocity (dashed).

8, and the closed-loop transfer function becomes

$$\frac{{}^i\tilde{X}}{x_t} = \frac{V(z)\,(1 - R(z)D_F(z))}{1 + V(z)R(z)D(z)} \tag{8}$$

which has the same denominator as (5) but an extra factor in the numerator. Clearly if $D_F(z) \equiv R^{-1}(z)$ the tracking error would be zero, but such a control strategy is not realizable since it requires (possibly future) knowledge of the target position which is not directly measurable.

Such a controller has been implemented, see Figure 9, based on estimated target velocity feedforward and with high-performance digital axis velocity-control loops. In the experiment the camera has 2-DOF rotational motion and is fixated on an object revolving on a turntable. The resulting centroid tracking error is shown in Figure 10. Figure 11 compares the measured and estimated target velocity. Since the majority of the control effort is due to the feedforward path, lower feedback gains can be used thus minimising problems due to latency. The use of a feedforward term has been proposed previously by De Schutter [19] in the context of robotic force control.

2.8 What other advantages are offered by "visual dynamic" control?

The analysis of visual servo dynamics leads to advantages beyond the obvious ones of improved tracking performance. Visual servoing admits the possibility of "closing the loop" on the dynamics of the actuator, transmission and structural compliance, allowing superior control of end-point dynamics than can be achieved using axis sensing alone.

Robot speed, or cycle time, is a critical factor in the economic justification of a robot. Although machines capable of extremely high tool-tip accelerations now exist, the overall cycle time is dependent upon other factors such as settling time and overshoot. High speed and acceleration are often achieved at considerable cost since effects such as rigid-body dynamics, non-linear friction and link and transmission flexibility [20] become significant. To achieve precise end-point control using joint position sensors the robot must be engineered to minimize these effects by employing gearless transmissions and massive links to minimise friction and compliance. End-point relative position measurement admits direct end-point, rather than joint-angle, control. With appropriate control strategies this may result in reduced overshoot and thus lower cycle time.

Figure 12 shows the measured end-effector motion of a Puma robot (under native position controller with arm fully outstretched and horizontal) after a high velocity motion. The end-effector has overshot by around 12 mm. Although not visible in the figure, the camera acceleration peaks at 4 g and oscillates at around 19 Hz. This mode is due to torsion in the vertical drive shaft.

A computed torque controller, with friction feedforward, runs at 500 Hz

$$\tau_d = \mathbf{M}\left\{\ddot{\theta}_d + K_p(\theta_d - \theta) + K_v(\dot{\theta}_d - \hat{\dot{\theta}})\right\} + \mathbf{F}(\hat{\dot{\theta}}) \quad (9)$$

This is overlaid with a variable structure control strategy which moves the arm under open-loop trajectory generator control to the vicinity of the target, and once found incorporates visual control for fine positioning

$$\dot{\theta}_d = {}^iK_p({}^iX_d - {}^iX) + {}^iK_v{}^i\hat{\dot{X}} \quad (10)$$

Experimental results with this controller are shown in Figure 13. The end-effector has followed an almost critically damped path to the target which was displaced approximately 20 mm from the destination of the trajectory generator. The target came into view at 0.32 s and the transition in control strategies occured four field times later at approximately 0.40 s. The end-effector has settled to within 1 pixel or 0.78 mm of the target whose position was unknown in advance.

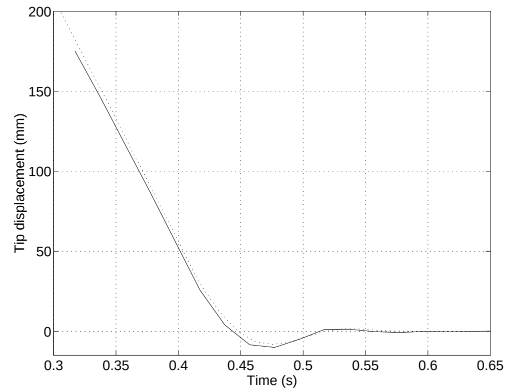

Figure 12: Measured tip displacement determined from axis sensing (dotted) and end-effector mounted camera (solid).

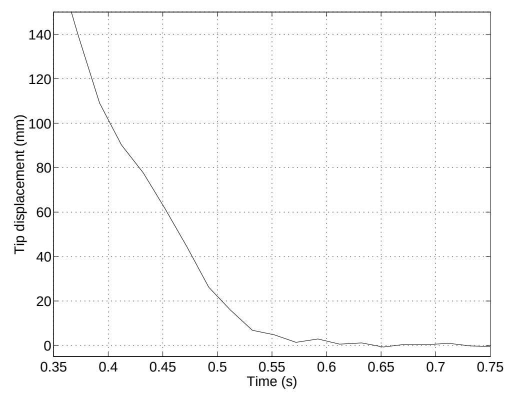

Figure 13: Measured tip displacement determined from end-effector mounted camera for hybrid visual control strategy.

3 Conclusions

This paper has discussed several important aspects of visual closed-loop systems that receive little attention in the literature. It has been argued that ignoring the temporal and dynamic characteristics of the system being controlled leads to the poor closed-loop performance demonstrated in many papers. It is also argued that axis velocity control has many advantages for use in visual servo systems. These include a 'natural' interpretation of visual servoing as a steering problem, and improved dynamic performance for manipulators with structural compliance.

The key conclusions are that in order to achieve high-performance visual servoing it is necessary to minimize open-loop latency, have an accurate dynamic model of the system and to employ a feedforward type control strategy. Prediction can be used to overcome latency but at the expense of reduced high frequency disturbance rejection. Open-loop latency is reduced by choice of a suitable control architecture.

An accurate dynamic model is required for control synthesis. Feedback-only controllers have a loop gain limit due to the significant delay in pixel transport and processing. Simple feedback controllers have significant phase lag characteristics which lead to poor tracking. More sophisticated feedback controllers can overcome this but the solution space becomes very constrained and the controllers are not robust with respect to plant parameter variation. Feedforward control results in a robust controller with excellent tracking capability. In addition dynamic visual servoing can be used to "close the loop" on the dynamics of the actuator, transmission and structural compliance, allowing superior control of end-point dynamics than can be achieved using axis sensing alone.

References

[1] P. I. Corke, *High-Performance Visual Closed-Loop Robot Control.* PhD thesis, University of Melbourne, Dept. Mechanical and Manufacturing Engineering, July 1994.

[2] L. Weiss, *Dynamic Visual Servo Control of Robots: an Adaptive Image-Based Approach.* PhD thesis, Carnegie-Mellon University, 1984.

[3] S. Ganapathy, "Real-time motion tracking using a single camera," Technical Memorandum 11358-841105-21-TM, AT&T Bell Laboratories, Nov. 1984.

[4] J.-C. Yuan, F. Keung, and R. MacDonald, "Telerobotic tracker," Patent EP 0 323 681 A1, European Patent Office, Filed 1988.

[5] S. Venkatesan and C. Archibald, "Realtime tracking in five degrees of freedom using two wrist-mounted laser range finders," in *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 2004–2010, 1990.

[6] S. D. Eppinger and W. P. Seering, "Introduction to dynamic models for robotic force control," *IEEE Control Systems Magazine*, vol. 7, no. 2, pp. 48–52, 1987.

[7] J. Hill and W. T. Park, "Real time control of a robot with a mobile camera," in *Proc. 9th ISIR*, (Washington, DC), pp. 233–246, Mar. 1979.

[8] J. Feddema, *Real Time Visual Feedback Control for Hand-Eye Coordinated Robotic Systems.* PhD thesis, Purdue University, 1989.

[9] N. Papanikolopoulos, P. Khosla, and T. Kanade, "Vision and control techniques for robotic visual tracking," in *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 857–864, 1991.

[10] P. Sharkey, I. Reid, P. McLauchlan, and D. Murray, "Real-time control of a reactive stereo head/eye platform," *Proc. 29th CDC*, pp. CO.1.2.1–CO.1.2.5, 1992.

[11] N. Houshangi, "Control of a robotic manipulator to grasp a moving target using vision," in *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 604–609, 1990.

[12] P. Corke and M. Good, "Dynamic effects in high-performance visual servoing," in *Proc. IEEE Int. Conf. Robotics and Automation*, (Nice), pp. 1838–1843, May 1992.

[13] T. Pool, *Motion Control of a Citrus-Picking Robot.* PhD thesis, University of Florida, 1989.

[14] K. Hashimoto, T. Kimoto, T. Ebine, and H. Kimura, "Manipulator control with image-based visual servo," in *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 2267–2272, 1991.

[15] P. Corke and M. Good, "Controller design for high-performance visual servoing," in *Proc. IFAC 12th World Congress*, (Sydney), pp. 9–395 to 9–398, 1993.

[16] M. Roberts, "Control of resonant robotic systems," Master's thesis, University of Newcastle, Australia, Mar. 1991.

[17] G. Franklin and J. Powell, *Digital Control of dynamic systems*. Addison-Wesley, 1980.

[18] M. Tomizuka, "Zero phase error tracking algorithm for digital control," *Journal of Dynamic Systems, Measurement and Control*, vol. 109, pp. 65–68, Mar. 1987.

[19] J. D. Schutter, "Improved force control laws for advanced tracking applications," in *procra*, pp. 1497–1502, 1988.

[20] M. C. Good, L. M. Sweet, and K. L. Strobel, "Dynamic models for control system design of integrated robot and drive systems," *ASME Journal of Dynamic Systems, Measurement and Control*, vol. 107, p. 53, 1985.