A New Actuation Approach for Human Friendly Robot Design

Michael Zinn
Bernard Roth

Design Division, Department of Mechanical Engineering
Stanford University
Stanford, CA 94305, USA
zinn@robotics.stanford.edu
roth@robotics.stanford.edu

Oussama Khatib
J. Kenneth Salisbury

Robotics Laboratory, Department of Computer Science
Stanford University
Stanford, CA 94305, USA
ok@robotics.stanford.edu
jks@robotics.stanford.edu

Abstract

In recent years, many successful robotic manipulator designs have been introduced. However, there remains the challenge of designing a manipulator that possesses the inherent safety characteristics necessary for human-centered robotics. In this paper, we present a new actuation approach that has the requisite characteristics for inherent safety while maintaining the performance expected of modern designs. By drastically reducing the effective impedance of the manipulator while maintaining high-frequency torque capability, we show that the competing design requirements of performance and safety can be successfully integrated into a single manipulation system.

KEY WORDS—actuation, human-friendly, safety, robotics, control

1. Introduction

In recent years, there has been great interest generated in the emerging field of human-centered robotics (Giralt and Corke 2001). Human-centered robotics involves the close interaction between robotic manipulation systems and human beings, including direct human–manipulator contact. In such applications, traditional figures of merit, such as bandwidth, maximum force and torque capability, and reachable workspace, do not fully encompass the range of metrics which define the requirements of such systems. Specifically, human-centered robotic systems must consider the requirements of safety in addition to the traditional metrics of performance. The question arises as to whether it is possible to successfully integrate the competing requirements of safety and performance in a single system. To answer this question we must first understand why some robotic systems are unsafe and, alternatively, why some systems have low performance.

1.1. Why Are Some Manipulators Unsafe?

Manipulator safety is dependent on a manipulator's mechanical, electrical, and software design characteristics. However, the biggest danger present when working in close proximity with robotic manipulators is the potential for large impact loads resulting from the large effective inertia (or more generally effective impedance) of many robotic manipulators.

To evaluate the potential for serious injury due to impact we can make use of an empirical formula developed by the automotive industry to correlate head acceleration to injury severity known as the head injury criteria (HIC). A simple two-degrees-of-freedom (2-DOF) mass–spring model can be used to predict head accelerations that would occur during an uncontrolled impact. In combination with the HIC index, the predicted accelerations are used to estimate the likelihood of serious injury occurring during an impact between a robotic manipulator and a human. For the PUMA 560, an impact velocity of one meter per second produces a maximum HIC greater than 500, more than enough to cause injury[1] (see Figure 1).

The International Journal of Robotics Research
Vol. 23, No. 4–5, April–May 2004, pp. 379-398,
DOI: 10.1177/0278364904042193

1. The HIC index is correlated with the maximum abbreviated injury scale (MAIS) to provide a mapping from the calculated HIC values to the likelihood of an occurrence of a specific injury severity level. In Figure 1, HIC values and the corresponding likelihood of a concussive injury (or greater) are shown.

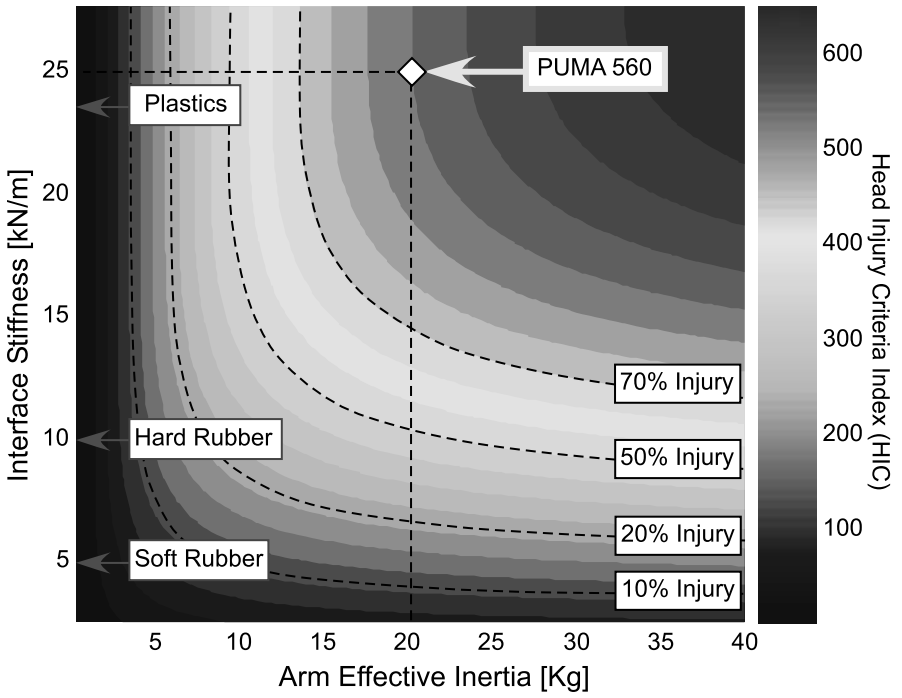

Fig. 1. Head injury criteria as a function of effective inertia and interface stiffness.

As seen in Figure 1, the addition of a compliant covering can reduce impact loading by an order of magnitude or more. However, the amount of compliant material required to reduce impact loads to a safe level can be substantial.[2] Clearly, adding large amounts of compliant covering is impractical and does not address the root cause of high impact loads: namely, the large effective inertia of most modern robotic arms. This hazard can be somewhat mitigated with the use of software and sensor architectures which monitor and interrupt potential anomalies, and thus reduce the chance of uncontrolled impact. However, even the most robust system is subject to unpredictable behavior as a result of electrical, sensor, or software faults. Thus, the mechanical characteristics of a robotic system are the limiting factor in improving overall safety (Zinn et al. 2002).

If inherent safety is to be achieved, we must design manipulators that have naturally low impedance. Unfortunately, most modern robotic manipulators have high effective impedance stemming from their requirements for high performance. The payload requirements and high-bandwidth control necessitate the use of high-inertia gear-head actuators and stiff, bulky structure which drive up the weight and impedance of these systems to unsafe levels.

2. For the PUMA robot, the thickness of a compliant cover required is more than five inches, assuming an impact velocity of one meter per second and an allowable maximum HIC index of 100.

1.2. Why Do Some Manipulators Have Low Performance?

Some types of robotic manipulators, notably those utilizing compliant actuation, such as pneumatic actuators, or those employing compliant drive trains, such as cable driven manipulators, do not produce the large impact loads associated with high impedance designs. We can understand this by examining a simple mass–spring model of an actuator link system with drive-train compliance (see Figure 2(a)).

At low frequencies, the effective impedance at the link can be approximated as the sum of the link's and reflected actuator's impedance (see Figure 2(b)). However, at high frequencies, which produce the bulk of impact load energy, the effective impedance is reduced to the link inertia only (see Figure 2(c)). For many manipulator systems, the actuator reflected inertia, with the N^2 amplification due to gear reduction, is much larger than the link inertia. The attenuation of the actuator's reflected inertia through the compliant drive train can significantly reduce impact loads, improving safety characteristics.

While a compliant actuator or drive train can enhance safety characteristics, the performance of such systems is limited. The flexible modes of the compliant system prevent control bandwidths greater than about one-third of the fundamental resonant frequency. In addition, attenuation of flexible mode oscillations excited by disturbances can be difficult to achieve. This results from the phase delay introduced above

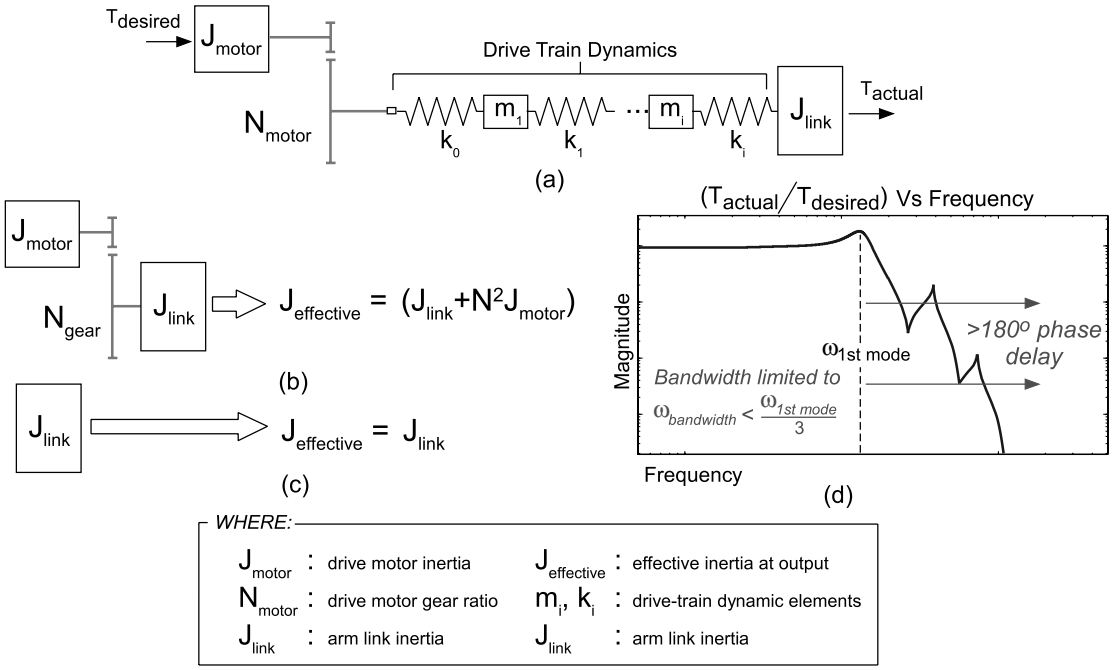

Fig. 2. (a) Robotic manipulator compliant drive-train mass–spring model. (b) Low-frequency effective inertia approximation. (c) High-frequency effective inertia approximation. (d) Open-loop $T_{actual}/T_{desired}$ magnitude versus frequency.

the first mode frequency (see Figure 2(d)). With the resonant frequencies of many cable driven manipulators in the range of 10 Hz or less, high-performance control of such systems is difficult, if not impossible.

1.3. Actuator Characteristics: Obstacle Toward Achieving Safety and Performance

So why is it so difficult to simultaneously achieve safety and performance characteristics in a single manipulator design? The limitations of current actuation technology and the manner in which these actuators are used in manipulator designs are to blame. To understand why, we must examine the characteristics of existing actuation technology. Currently, only electromagnetic, hydraulic, and pneumatic actuators have the power and torque capabilities required for robotic manipulation tasks. Unfortunately, all of these actuation methods have serious deficiencies, limiting their inherent safety and/or performance characteristics.

Hydraulic actuators, which have the highest torque and power density characteristics of any of the actuation methods, are capable of performing tasks which involve the application of thousands of Newton-meters of torque and many kilowatts of power output. However, their very high output stiffness characteristics, which make the hydraulic actuator essentially a pure position source, can render it very dangerous. The output impedance, as compared to the driven manipulator and environment, is virtually infinite, generating very high impact loads during collisions. Typically these actuators are employed at the joint or through a rigid linkage, further increasing the effective inertia of the manipulator. Thus, manipulators that employ hydraulic actuators have very poor inherent safety characteristics.

Pneumatic actuators, on the other hand, can be made very compliant. Due to the near zero inductance of the compressible gas, their output impedance is low over a wide frequency range, reducing uncontrolled impact loads to potentially safe levels. However, pneumatic actuators have very low-bandwidth capabilities. Even when pressure control is implemented (as opposed to conventional flow control), control bandwidths are limited to less than 20 Hz, which is insufficient for high-performance tasks (Hollerbach, Hunter, and Ballantyne 1991). Making matters worse, the slow bandwidth capabilities render the large amount of stored potential energy in the compressible gas a serious hazard. Thus, while the natural compliance of pneumatic actuation reduces its effective inertia, its low-bandwidth characteristics limit the performance characteristics of manipulators which use them (for the same reasons described in Section 1.2).

Primarily as a result of the limitations of pneumatic and hydraulic actuators, many current human-centered research efforts use manipulation devices that employ electromagnetic actuation as their primary torque source. The primary limitation of electromagnetic motors is their relatively low torque

and power density. The use of electromagnetic motors without a torque magnifying reducer is limited to direct drive systems that must employ large direct-current (DC) torque motors which are heavy and inefficient. To increase the torque output to useful levels, gear reducers are almost universally employed when using electromagnetic actuators. Unfortunately, the increase in torque and power density that results must be traded against the large increase in reflected inertia which increases with the square of the gear reduction. Reduction ratios employed in most systems more than double the effective inertia of the manipulator, trading safety for improved performance.

2. New Actuation Approaches

New actuation approaches have been developed to overcome the safety and performance limitations of existing systems. Chief among these are the joint torque control approach (Vischer and Khatib 1995) and series elastic actuation (Pratt and Williamson 1995). However, for reasons discussed in the following sections, these approaches do not simultaneously achieve the characteristics necessary for both safety and performance. To address these limitations and create a unified high-performance and safe robotic manipulator, a new actuation approach, referred to as the distributed macro–mini (DM^2) actuation approach, has been proposed (Zinn et al. 2002).

2.1. Joint Torque Controlled Actuation

Joint torque control was developed to eliminate the deleterious effects of nonlinearities and friction inherent in the actuator–transmission systems generally found in industrial robots. Initial implementations were successful in substantially reducing joint friction effects but wide joint actuation bandwidth was difficult to achieve without actually reducing the friction and nonlinearities in the actuator–transmission system (Holmberg, Dickert, and Khatib 1992; Vischer and Khatib 1995; Hirzinger et al. 2001).

In response, joint torque control systems employ high-performance actuator and transmission designs with integrated torque sensors to achieve the performance levels desired. Perhaps the most successful of these has been the new DLR lightweight arm design; see Figure 3 (Hirzinger et al. 2002). The implementation of joint torque control allows for near-zero low-frequency impedance, which gives the DLR arm excellent force control characteristics. However, above the control bandwidth, joint torque control is ineffective at reducing the impedance of the manipulator. The open-loop characteristics of the manipulator and reflected actuator inertia dominate. Thus, the magnitudes of impact loads, which are determined by the high-frequency impedance of the contacting surfaces, are not attenuated.

While the joint torque control has been successful in improving the force and impedance control of robotic manipulators, their fundamental open-loop characteristics make inherent safety difficult to achieve and thus do not satisfy the human-centered robotic requirements of both performance and safety.

2.2. Series Elastic Actuation

Recently, a class of actuators, known as series elastic actuators (SEAs), has been developed to address the problems of high impedance actuators (Pratt and Williamson 1995; Robinson 2000). The SEA approach seeks to mitigate the limitations of high impedance actuators, such as conventional gear-head electromagnetic or hydraulic actuators, by placing an elastic element between the output of the actuator and the robotic link. The elastic element limits the high-frequency impedance of the actuator to the stiffness of the elastic coupling. To limit the low-frequency impedance, and thus transform the actuator into an approximate pure torque source, a linear feedback system is implemented to regulate the output torque of the actuator–spring system (see Figure 4).

The main advantage of the SEA topology is that it provides low output impedance across the frequency spectrum. As shown in Pratt and Williamson (1995) and Robinson (2000), the SEA topology reduces the output impedance of the SEA in proportion with the stiffness of the elastic coupling, as shown in the output impedance transfer function in eq. (1):

$$\frac{F(s)}{X(s)} = \frac{s^2(N_{motor})^2 I_{motor}}{(s^2(N_{motor})^2 I_{motor}/K_s) + 1 + N_{motor} D(s)}. \quad (1)$$

At frequencies below the closed-loop bandwidth of the SEA controller, the output impedance is reduced as a function of the control gains. Impedance reduction of $10\times$ to $100\times$ is common, and is only limited by the maximum obtainable bandwidth. At frequencies above the closed-loop bandwidth, the output impedance reduces to the stiffness of the elastic coupling. This is in contrast to other approaches, such as joint torque control discussed in Section 2.1, which have good low-frequency impedance but suffer from large high-frequency impedance.

It is interesting to note the similarities between the SEA and joint torque control approaches. The topology of joint torque control is identical to that of the SEA approach (shown in Figure 4). The difference between the two approaches lies in their differing control approaches, which are driven by their very different open-loop characteristics. As described earlier, series elastic actuation has a compliant coupling between the actuator and driven link, the stiffness value of which is chosen so that the open-loop mode of the system is well below the obtainable closed-loop bandwidth of the SEA control. As a result of the low stiffness compliance, the open-loop gain is very low, which allows for the use of a simple, high-gain PD controller. The resulting system is stable and possesses low impedance over a wide frequency range. In contrast, with the joint torque control approach, the coupling between the

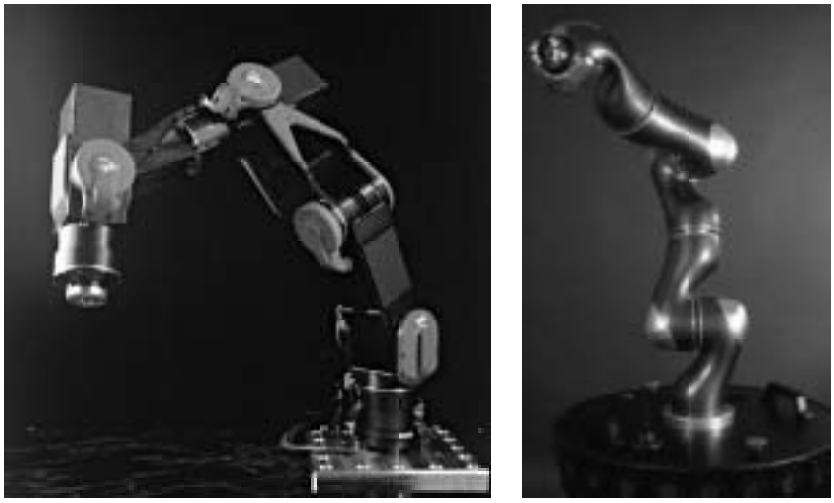

Fig. 3. DLR lightweight robot: (a) DLR II; (b) DLR III.

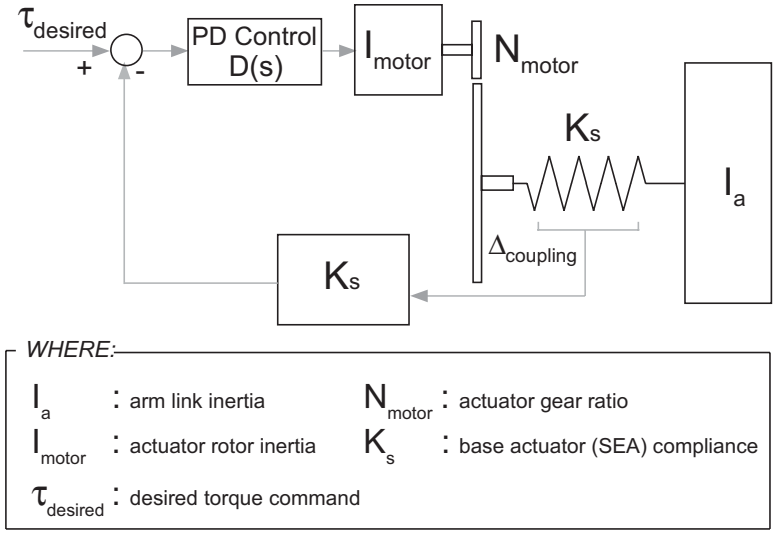

Fig. 4. Series elastic actuation topology.

actuation and driven link is much stiffer. Implementation of PD control is difficult and requires that the control gains be kept low to maintain stability. As a result, alternative control schemes have been implemented, including PI control (Vischer and Khatib 1995) and full-state feedback (Hirzinger et al. 2002).

There are trade-offs with using the SEAs. Due to velocity and torque saturation of the SEA, the maximum output torque above the open-loop mode of the system[3] falls off as $1/\omega$ regardless of the control loop controller bandwidth (Robinson 2000). This behavior is an open-loop characteristic of the SEA topology and represents a fundamental physical limitation of the actuator. The choice of the elastic coupling stiffness (in relation to the manipulator and motor reflected inertia) determines the open-loop mode frequency. A stiffer coupling improves the high-frequency torque performance but adversely affects the desirable closed-loop and open-loop impedance characteristics.

The use of a compliant coupling and the closed-loop control of the SEA output torque limits the bandwidth of any

3. SEA open-loop mode: unforced coupled motion of actuator and manipulator link inertias through the compliant coupling.

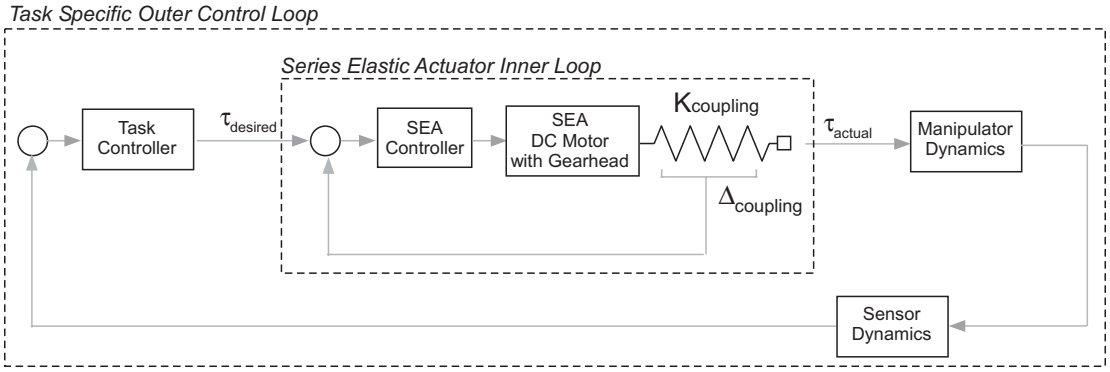

Fig. 5. Nested series elastic actuation and outer task control loops.

task which relies on a series elastic actuator as its only torque source. This limitation derives from the use of the SEA closed-loop system within a larger, task-specific control loop. As shown in Figure 5, the design and resulting stability of the task-specific control loop is dependent on the interaction between the inner SEA closed-loop system and the outer task-specific control loop. If the outer loop bandwidth approaches the bandwidth of the inner loop, instability is likely to occur. As a result, the task specific control loop cannot be closed at a rate faster than the inner loop.

Tasks such as position control and end-effector impedance control are limited to a bandwidth that is significantly below the closed-loop bandwidth of the SEA. This is not a major consideration for manipulation systems which do not require fast dynamics, such as walking robots, for which the series elastic actuators were originally developed. However, for tasks requiring high-bandwidth control, such as high-speed trajectory tracking or high-frequency disturbance rejection, the limitations of the series elastic actuators are prohibitive. Other approaches have been proposed, such as the use of a nonlinear elastic coupling whose compliance can be changed through co-activated actuators (Bicchi, Rizzini, and Tonietti 2001). Unfortunately, the bandwidth limitations affecting the series elastic actuator, while mitigated somewhat by the variable compliance, are still a limiting factor in performance.

3. Distributed Macro–Mini Actuation Approach

Recently, a new actuation approach, referred to as the DM^2 actuation approach, has been developed to overcome the safety limitations of joint torque control and the performance limitations of series elastic actuation (Zinn et al. 2002). As the name implies, the DM^2 actuation approach employs a pair of actuators, connected in parallel and distributed to different locations on the manipulator. The effective inertia of the overall manipulator is substantially reduced by isolating the reflected inertia of the actuator while greatly reducing the overall weight of the manipulator. Performance is maintained with small actuators collocated with the joints. Our approach partitions the torque generation into low- and high-frequency components and distributes these components to the arm location where they are most effective. The overall approach is shown in Figure 6.

The first part of the DM^2 actuation approach is to divide the torque generation into separate low- and high-frequency actuators whose torque sum in parallel. The effectiveness of this approach can be seen clearly when we consider that most manipulation tasks involve position or force control which are dominated by low-frequency trajectory tracking or DC load torques. High-frequency torques are almost exclusively used for disturbance rejection. Even haptic device torque profiles, which might require rapid changes approximating a square wave input, have a torque magnitude versus frequency curve that falls off with increasing frequency by $1/\omega$ (see Figure 7). This partition is even more compelling when we consider power requirements versus frequency. Using the square wave example above, power versus frequency falls off with $1/\omega^2$. This power versus frequency profile is ideally fit using a large output, low-frequency actuator coupled with a high-frequency small torque motor.

In order for the DM^2 actuation approach to work properly, both the high- and low-frequency actuators must have zero or near-zero impedance. This is due to the fact that during power transfer the actuator torques will add non-destructively only if their respective impedance is zero. In particular, each actuator must not have significant impedance within the frequency range of the opposing actuator. Only if this condition is true will the DM^2 concept work. For the high-frequency actuation, very low impedance is achieved by using a small low-inertia torque motor connected to the manipulator through a low-friction, low-reduction cable transmission. For the

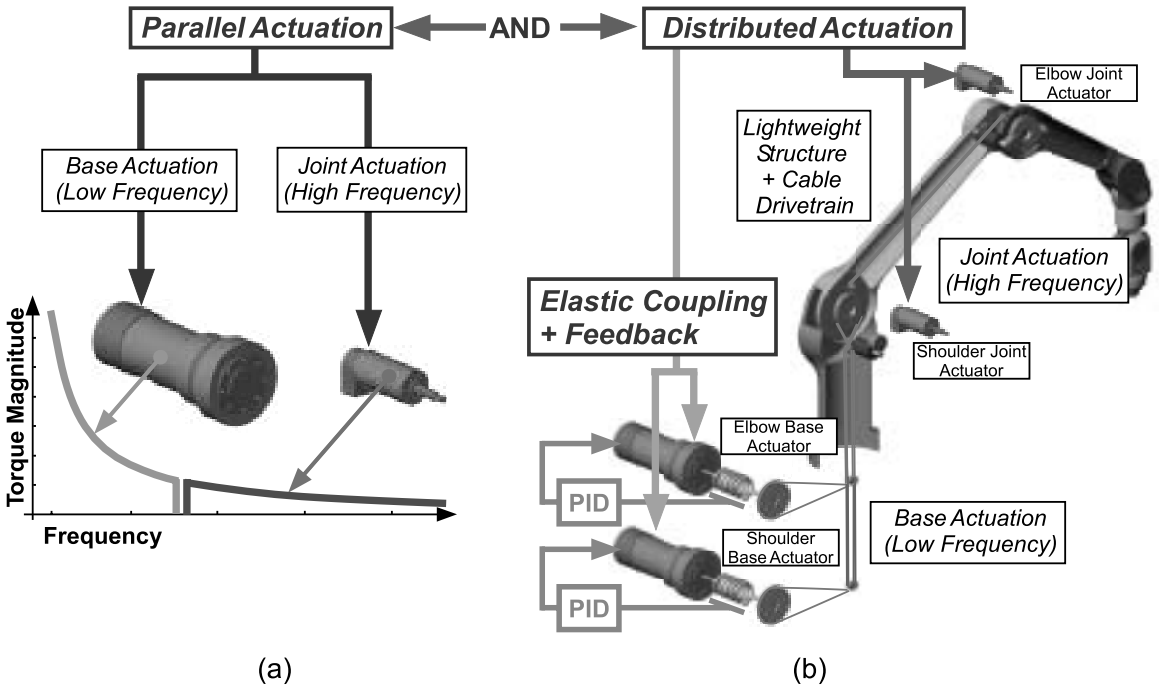

Fig. 6. DM^2 actuation approach. (a) Partition of torque into low- and high-frequency (parallel) components. (b) Distributed actuation: large, low-frequency actuators are located at the base; small, high-frequency actuators are located at the joints.

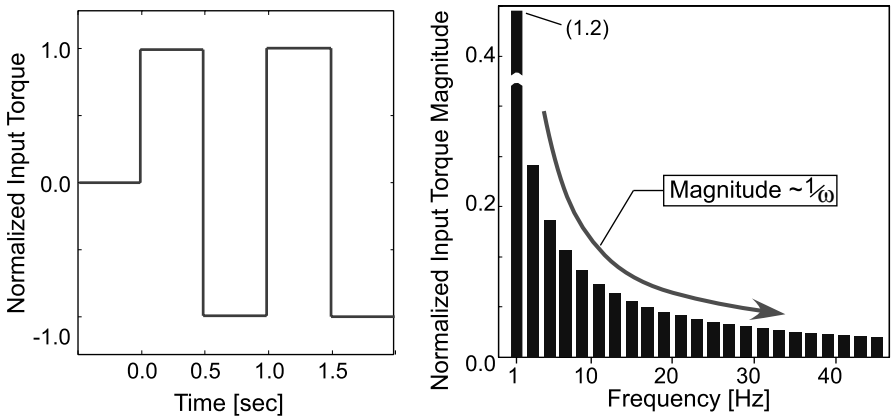

Fig. 7. Torque versus frequency: 1 Hz square wave.

low-frequency actuation, we achieve low impedance by using a series elastic actuator (see Section 2.2). Because the DM^2 actuation approach does not require that the base actuator be capable of supplying high-frequency torques, the bandwidth limitations of SEAs do not pose a difficulty.

The second part of the DM^2 actuation approach, which differs from previous attempts at coupled actuation (Morrell 1996), is to distribute the low- and high-frequency actuators to locations on the manipulator, where their effect on contact impedance is minimized while their contribution to control bandwidth is maximized. This is achieved by locating the low-frequency series elastic actuator remotely from the actuated joint. This is particularly advantageous as the low-frequency components of most manipulation tasks are considerably larger in magnitude than the high-frequency components, and consequently require a relatively large actuator. Locating the large SEA at the base significantly reduces the weight and inertia of the manipulator. The high-frequency ac-

tuators are located at the manipulator joints and connected through a stiff, low-friction transmission, providing the high-frequency torque components that the low-frequency base actuators cannot. The high-frequency torque actuator must be connected to the joint inertia through a connection which produces a high primary mode vibration frequency. By locating the actuator at the joint and by using a low-inertia servomotor, we can achieve this high-bandwidth connection with a minimum amount of weight and complexity.

The DM^2 actuation approach is analogous to the design of robotic manipulators for use in zero gravity. Under such conditions, gravity induced torques do not exist. Joint actuators provide torques related only to the task, such as trajectory tracking and disturbance rejection, both of which are primarily medium to high frequency in content. We achieve the zero gravity analogy by compensating for gravity torques and low-frequency torques using the low-frequency actuators located at the base of the manipulator. With the effects of gravity and low-frequency torques compensated, joint torque requirements become similar to those encountered by a zero gravity robotic manipulator. However, unlike robotic manipulators designed for space applications, the DM^2 joint actuators do not require a large gear reducer to achieve the required torque and power densities.

3.1. DM^2 Actuation Control Approach

Perhaps the most challenging aspect of a DM^2 implementation is the development of a control approach which leverages the characteristics of the parallel actuator structure while dealing with the unique control challenges associated with the use of low-impedance actuation.

At the joint level, the DM^2 actuation approach is essentially a dual-input single-output system. The redundant actuators provide an additional degree of freedom which can be used in optimizing system performance while minimizing actuation effort. For example, in the case of trajectory tracking, we can use LQR control techniques to obtain an optimum control law based on minimizing control effort and tracking error. The low- and high-frequency actuation effort partitioning can be accomplished in a similar manner. However, this type of control structure is specific to a given task, in this case to trajectory tracking, and does not provide a black-box interface to the actuation similar to the use of a single actuator. In particular, for applications involving a number of different control modes, such as free-space motion with contact transitions, or for applications requiring a low-impedance torque source, such as haptics or tele-robotic master devices, we desire an actuation control scheme which allows the use of the parallel actuation system as a single torque source.

3.1.1. Near-Perfect Torque Source

As such, our control approach seeks to exploit the DM^2 actuation's unique characteristics to construct a near-perfect torque source. The characteristics of a perfect torque source, consisting of zero output impedance and infinite control bandwidth, would enable a manipulator to possess the characteristics necessary for both inherent safety and high-performance tasks. While a perfect torque source is impossible to achieve, a near-perfect torque source, with low output impedance relative to the driving load and high-bandwidth torque capability offers many of the same advantages.

Figures 8 and 9 show a physical schematic diagram of the control structure and an equivalent block diagram representation, respectively. The transfer function of the control structure shown in Figure 9 has unity gain and zero phase over all frequencies ($T_{actual}(s)/T_{desired}(s) = 1$). A simplified representation, shown in Figure 10, demonstrates how the control approach utilizes the low-frequency base actuator's low pass filter characteristics to partition the control torques into low- and high-frequency components.

By using the actual measured torque output from the low-frequency base actuators in combination with the desired torque, we automatically compensate for the non-ideal behavior of the base actuators. Assuming that the smaller joint actuators can produce this torque, the sum of the combined torques is a perfect realization of the desired torque. The frequency partitioning can be clearly seen if we rearrange the structure in Figure 10(a) into a pure parallel structure, as shown in Figure 10(b). As seen in Figure 10(b), the base actuator's transfer function falls off above its closed-loop bandwidth, w_{base}, while the equivalent joint actuator's transfer function approximates a double lead filter, which adds phase to the combined system above the open-loop mode frequency, w_s, and attenuates the DC and low-frequency components commanded to the high-frequency actuator.

The combined actuator control structure creates a perfect torque source in the linear sense, where the torques sum to unity magnitude and zero phase, as seen in Figures 11(a) and (b). Thus, by using the simple control structure described above, we can create a unified actuator with the desirable characteristics of low impedance (necessary for inherent safety) and high-bandwidth torque control (necessary for high performance).

3.1.2. Manipulation Control

The DM^2 control structure allows for straightforward implementation of the DM^2 actuation approach in multi-DOF manipulators system. Assuming that the assumptions of a near-perfect torque source hold, the DM^2 actuation approach is particularly well suited to control methods, such as operational control (Khatib 1987), which assume that the control torques are directly applied to the joint with little or no unmodeled disturbances from sources such as actuator friction or reflected inertia.

The perfect torque source structure breaks down when the assumptions of the model shown in Figures 8 and 9 are no

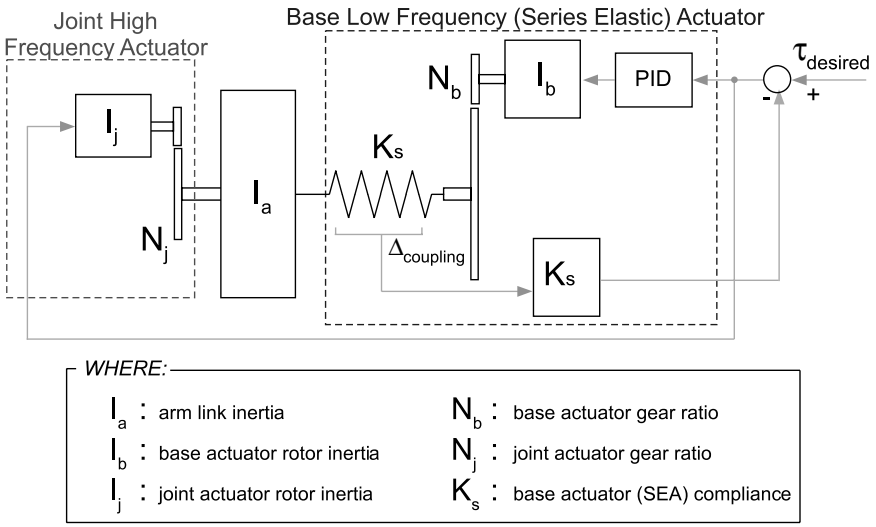

Fig. 8. DM^2 actuation and control topology (single degree of freedom).

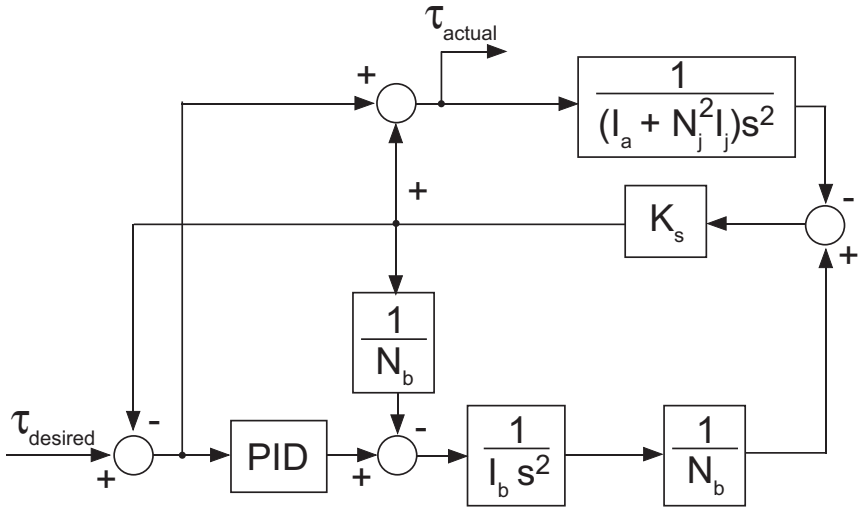

Fig. 9. DM^2 actuation and control block diagram representation (single degree of freedom).

longer valid. The main challenge in implementing the control scheme is in identifying and avoiding the situations where this ideal model breaks down.

3.1.3. Effects of Actuator Saturation

One significant deviation from the ideal model occurs when one of the DM^2 parallel actuators saturates. DM^2 actuator torque saturation represents the threshold above which the joint actuator can no longer compensate for the phase and magnitude error of the low-frequency base actuator. Commanded torques which force the high-frequency joint actuator to saturate will cause both magnitude errors and phase lag to occur, invalidating the perfect torque source characteristics of the combined parallel actuation. This effect is illustrated in Figures 12(a) and (b).

In Figures 12(a) and (b), the frequency response of the base series elastic actuator, the joint actuator, and the combined DM^2 actuator is shown on a polar plot of magnitude versus frequency (Figure 12(a)) and as a bode plot (Figure 12(b)). The effect of saturation can be seen as both magnitude and phase errors in the resulting parallel actuation response. As the

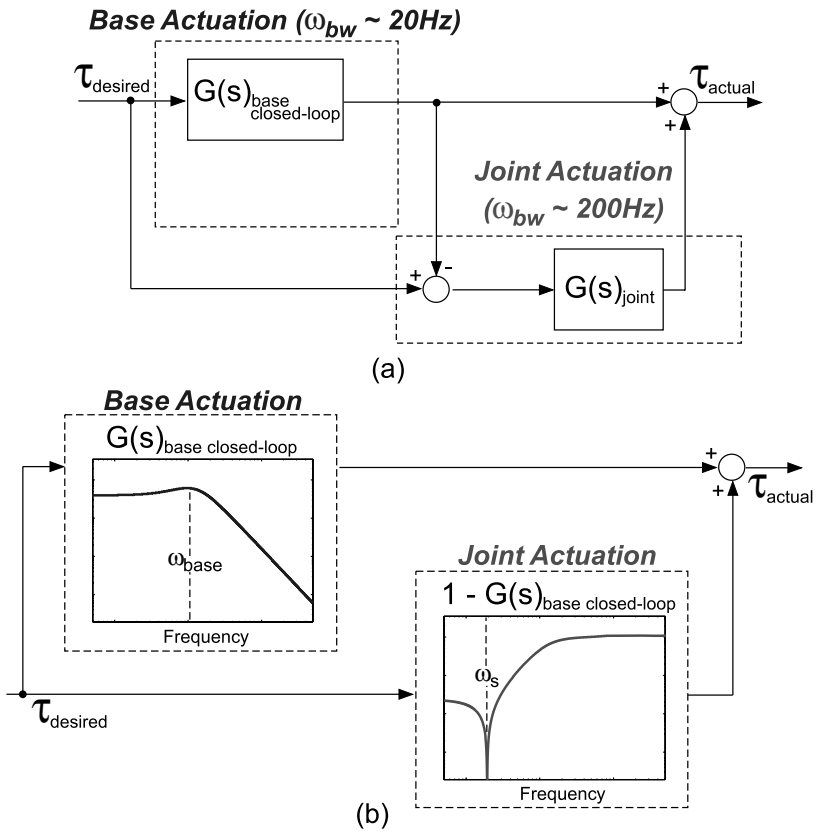

Fig. 10. (a) DM^2 actuation control structure ($G(s)_{base-closed-loop}$, base actuator closed-loop transfer function; $G(s)_{joint}$, joint actuator transfer function). (b) Equivalent parallel structure.

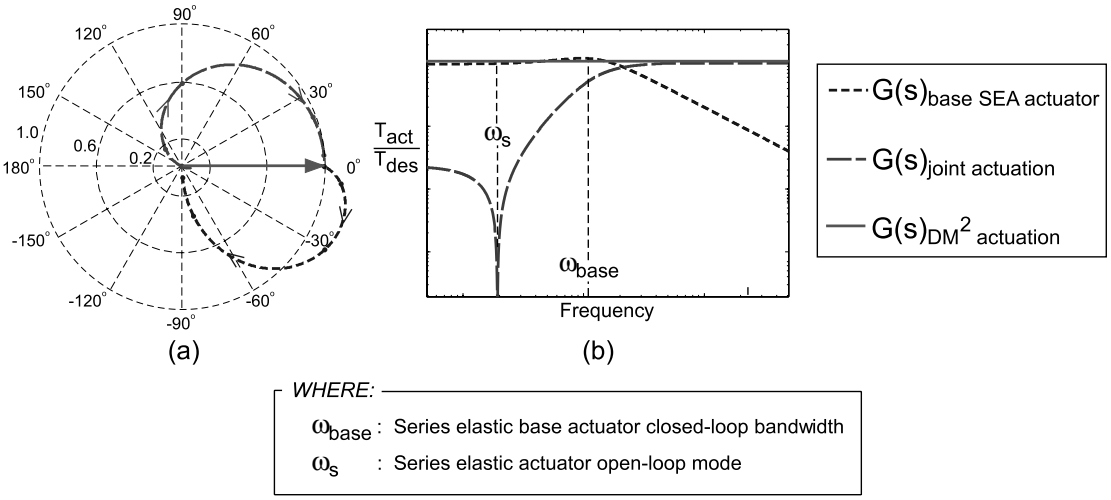

Fig. 11. (a) Perfect torque source: base, joint, and combined DM^2 actuator torque magnitude versus phase polar plot. (b) Near-perfect torque source: base, joint, and combined DM^2 actuator torque magnitude versus frequency.

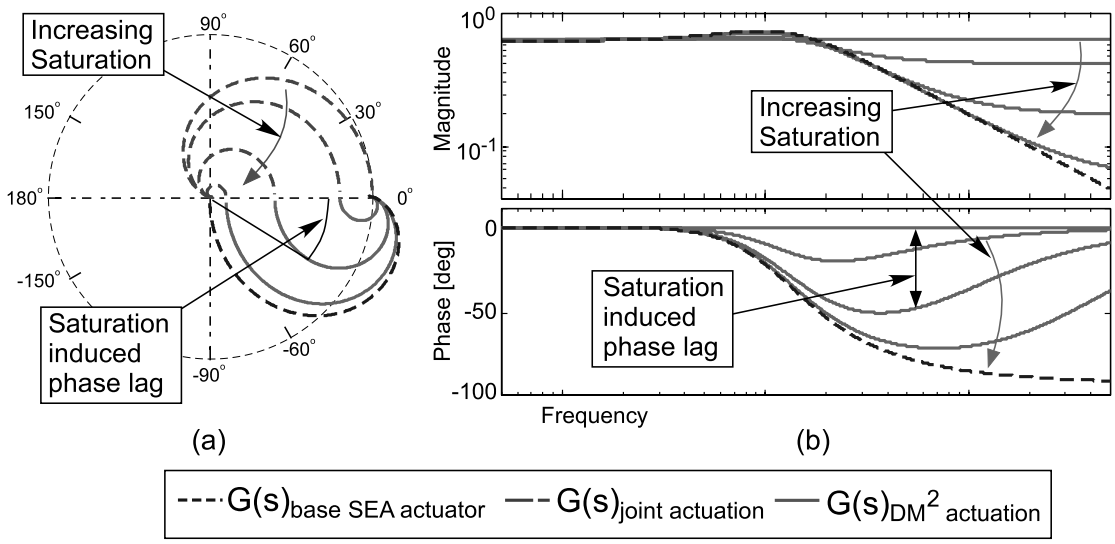

Fig. 12. Breakdown of perfect torque source due to saturation. (a) Base, joint (with saturation), and resulting DM^2 actuator torque magnitude versus phase polar plot. (b) Bode plot of DM^2 actuator torque with joint actuator saturation.

joint actuator approaches complete saturation, the combined parallel actuator's response approaches that of the single base series elastic actuator with its lower bandwidth constraints. This is particularly problematic in that a task control loop, such as position tracking, which under normal conditions is stable, can become unstable as a result of a torque command which exceeds the capabilities of the smaller joint actuator.

As a result, the controller must prevent saturation of the joint actuator from occurring. This can be accomplished by simply limiting the control input. This approach is taken when there is no outer task closed loop, such as with simple haptic rendering where the desired torque is a function of manipulator position alone and no effort is made to compensate the system output. In the case when there is a control loop wrapped around the DM^2 actuation, the control gains must be reduced or the control input limited to avoid saturation. As such, the avoidance of saturation poses a limit on the ultimate performance of a DM^2 actuated manipulator. Fortunately, as discussed in Section 3, the torque requirement of the high-frequency joint actuator is substantially less than the low-frequency base actuators. Thus, avoidance of joint actuator saturation can be achieved with proper sizing of the joint actuator with respect to the given manipulator tasks.

3.1.4. Manipulator Control With Low Impedance Actuation

Another deviation from the ideal model, which can have a significant effect on performance, is the existence of compliance in the drive train between the manipulator link and the joint actuator. While the joint actuator has a relatively stiff single-stage transmission design, some level of compliance is unavoidable. The drive-train compliance in combination with the low reflected inertia of the joint actuator produces low-frequency oscillations which can limit closed-loop performance. This effect can be understood by augmenting the DM^2 model in Figure 8 to include joint actuator drive-train compliance (see Figures 13 and 14).

The transfer function of the system shown in Figure 13 no longer represents a pure torque source. The addition of the oscillatory pole, due to the drive-train compliance, is clearly visible on the bode plot on the system transfer function in Figure 15. The presence of the flexible mode is of no surprise and, at first glance, does not seem to pose a significant problem. The relatively high frequency of the oscillatory mode would suggest that the proper choice of gains will provide adequate performance while avoiding excitation of the oscillatory pole.

However, as a result of the low reflected inertia of the DM^2 actuation approach, the ratio of joint actuator's reflected inertia to driven link inertia is very small, typically 1 : 10 or less. This large mismatch between actuator and drive-train inertia can cause serious low-frequency oscillations to occur when position or velocity feedback is introduced, limiting the maximum task control loop bandwidth achievable. For successful implementation of the DM^2 actuation approach, it is important to fully understand this phenomenon and to develop techniques to address and compensate for its effects.

We can more clearly understand this phenomenon using a simplified model of the DM^2 system, which includes the drive-train compliance but ignores the coupling with the low-frequency base actuator. Figure 16(a) and eq. (2) show the assumed model and its uncompensated open-loop transfer function.

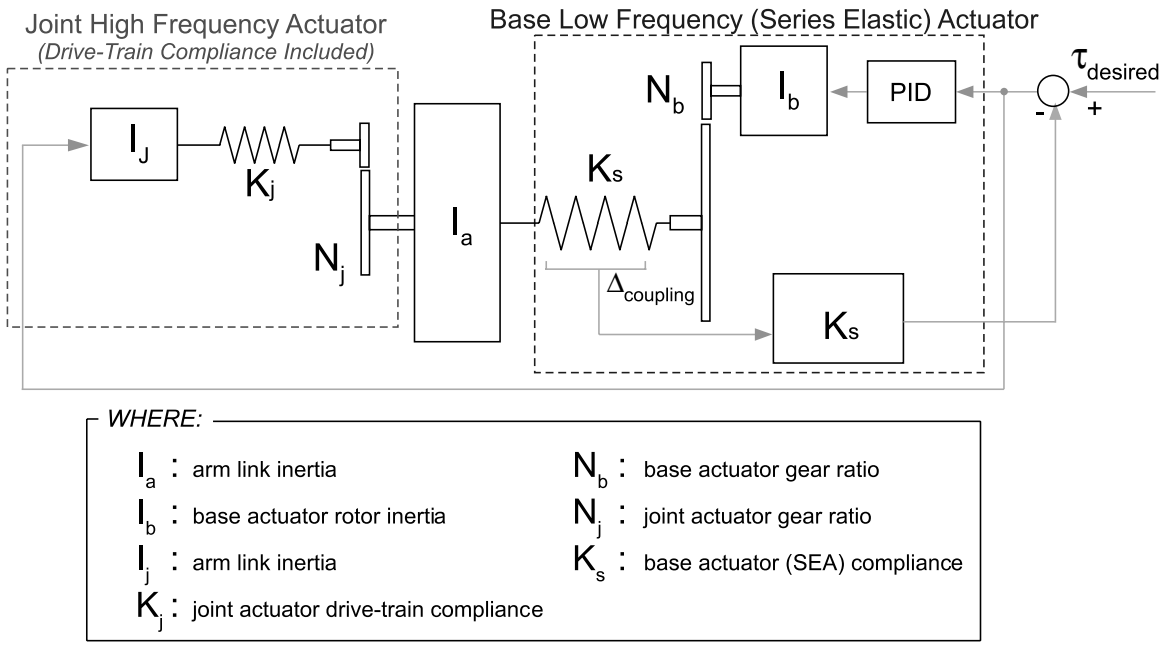

Fig. 13. Augmented DM^2 actuation and control topology (single degree of freedom). Joint actuator drive-train compliance included.

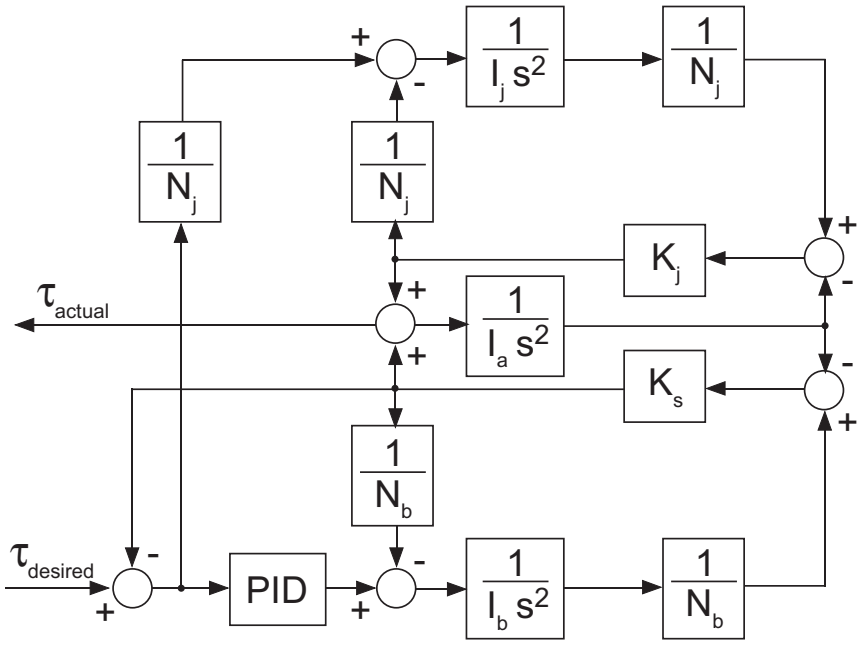

Fig. 14. Augmented DM^2 actuation and control block diagram representation (single degree of freedom). Joint actuator drive-train compliance included.

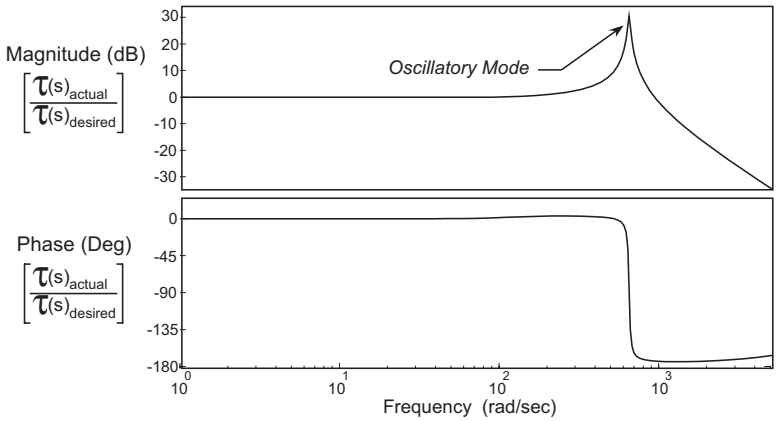

Fig. 15. Bode plot of augmented DM^2 actuation and control block diagram representation (single degree of freedom). Joint actuator drive-train compliance included.

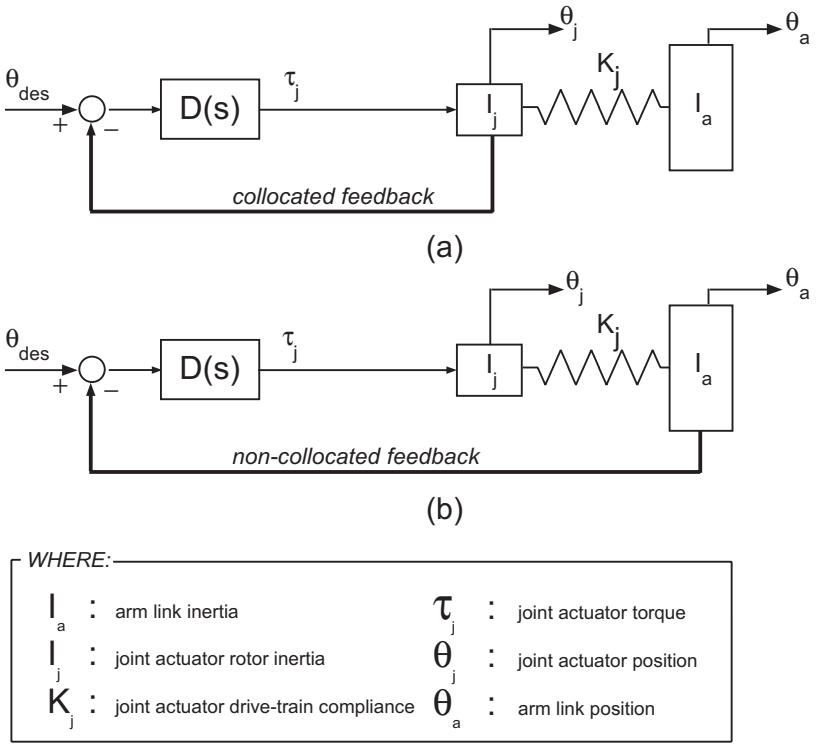

Fig. 16. Spring–mass model of joint actuator and driven link inertias: (a) collocated control; (b) non-collocated control.

$$\frac{\theta_j(s)}{\tau_j(s)} = \frac{s^2 I_a + K_j}{s^2(s^2 I_a I_j + K_j(I_a + I_j))}. \quad (2)$$

In many servo-systems, including robotics, the actuator and link inertias are matched or nearly matched to achieve optimum power and acceleration transfer from motor to load. In this situation, the poles and zeros of the transfer function, given by eq. (3), are approximately equal in frequency.

$$\omega_{zero} = \sqrt{\frac{K_j}{I_a}} \quad \text{and} \quad \omega_{pole} = \sqrt{\frac{K_j(I_j + I_a)}{I_a I_j}}. \quad (3)$$

However, in a system employing low impedance actuation, the zero frequency can be an order of magnitude below the frequency of the flexible mode pole. This large separation amplifies the flexible mode peak by a factor approximately equal to the ratio of drive link to motor inertias (see Figure 17).

This effect severely limits the achievable closed-loop bandwidth and thus performance in general. The effect can be quite puzzling considering that the flexible mode frequency can be very high (an order of magnitude or more above the open-loop crossover frequency) and still cause excessive oscillations in the closed-loop response. Only when we consider the zero, whose frequency is affected by the much larger drive link inertia, does it become clear why the problem exists.

Another way to analyze the problem is to examine the symmetric root locus (see Figure 18) of the system shown in Figure 16(a). When the ratio of joint motor rotor inertia to arm inertia, I_j/I_a, is close to 1 : 1, the oscillatory poles are drawn toward the transmission zeros as the gain is increased, reducing their residues, which reduces the magnitude of oscillations and allows for larger closed-loop gains. However, when the motor inertia, I_j, is much less than the arm inertia, I_a, the transmission zeros are located too far from the oscillatory poles to have a stabilizing effect and instead attract the dominant second-order poles. This phenomenon can be clearly seen if we look at the symmetric root locus for the transfer function in eq. (2).

As seen in Figure 18, when the motor inertia, I_j, is smaller than the arm inertia, I_a, the optimal control gains drive the dominant poles toward the zeros, indicating that a large amount of control effort would be required to modify the system behavior away from the low-frequency zeros. As a result, achieving high-bandwidth closed-loop control can be difficult and represents the biggest challenge when implementing the DM^2 actuation approach.

3.1.5. Achieving High-Bandwidth Control

The challenge of implementing high-bandwidth control in a DM^2 actuated system has been addressed through the combined implementation of prudent mechanical design techniques, which favorably modify the manipulator's open-loop dynamics, and control augmentation such as filtering and proper actuator–sensor placement.

Modification of the manipulator dynamics primarily involves attempts to increase the stiffness of the coupling between the motor inertia and link inertia. A stiffer coupling will increase both the frequency of the oscillatory poles and the transmission zeros, allowing for a higher crossover frequency. An alternative approach is to intentionally increase the motor's inertia, thereby reducing the frequency of the oscillatory poles to the frequency of the zeros. However, this approach is only useful when the motor and link inertias differ by less than approximately a factor of 2. Otherwise, the required increase in motor inertia is excessively large and severely reduces the acceleration capability of the system. Regardless, in the case of low impedance actuation, a large increase in actuator inertia would substantially increase the reflected inertia of the actuator, adversely affecting its safety characteristics, and thus cannot be considered for human-friendly robotic systems. Thus, mechanical modifications are limited to those that increase the frequency of the transmission zeros, such as stiffening the motor transmission or reducing the driven link inertia

In addition to mechanical modifications and control signal filtering (Ellis and Lorenz 2000), a somewhat surprising method to deal with the low-frequency oscillations associated with low impedance actuation is to change the control topology from collocated to non-collocated control. We can understand this by examining the open-loop transfer function of a simple mass–spring model of an actuator link system which employs non-collocated control. Figure 16(b) and eq. (4) show the assumed model and its associated transfer function.

$$\frac{\theta_a(s)}{\tau_j(s)} = \frac{K_j}{s^2(s^2 I_a I_j + K_j(I_a + I_j))}. \quad (4)$$

At first glance, this seems counter-intuitive since, in most cases, the stabilizing effects of the zeros associated with collocated control are beneficial and allow for more aggressive gains. However, in the case of large inertia mismatch, the collocated control zero is the main cause of the problem. A comparison of peaking amplitude (see Figure 19) shows that for large mismatches the non-collocated control may be better than a collocated approach. Of course, this does not take into account the tendency of the oscillatory poles to become unstable, and special care must be taken to insure their stability, such as the use of a notch filter or a gain stabilizing lag network (Cannon and Schmitz 1984). With this consideration, we can conservatively assume that when using non-collocated control we can achieve a crossover frequency as high as one-fifth of the flexible mode frequency. With this assumption, we can see from Figure 19 that when the joint motor inertia is much less than the arm inertia ($I_j/I_a < 10$) the use of non-collocated control allows for a higher closed-loop bandwidth than collocated control. This, in fact, has been shown to be the case on a two-axis testbed, where the motor link inertia ratios range from 1 : 50 to less than 1 : 100.

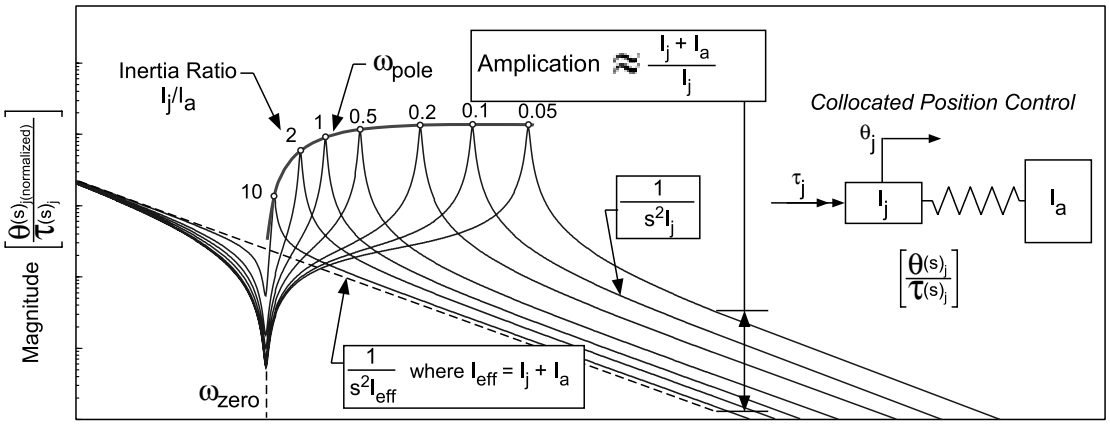

Fig. 17. Open-loop transfer function of collocated motor position control: amplification of oscillatory pole due to mismatched actuator link inertia.

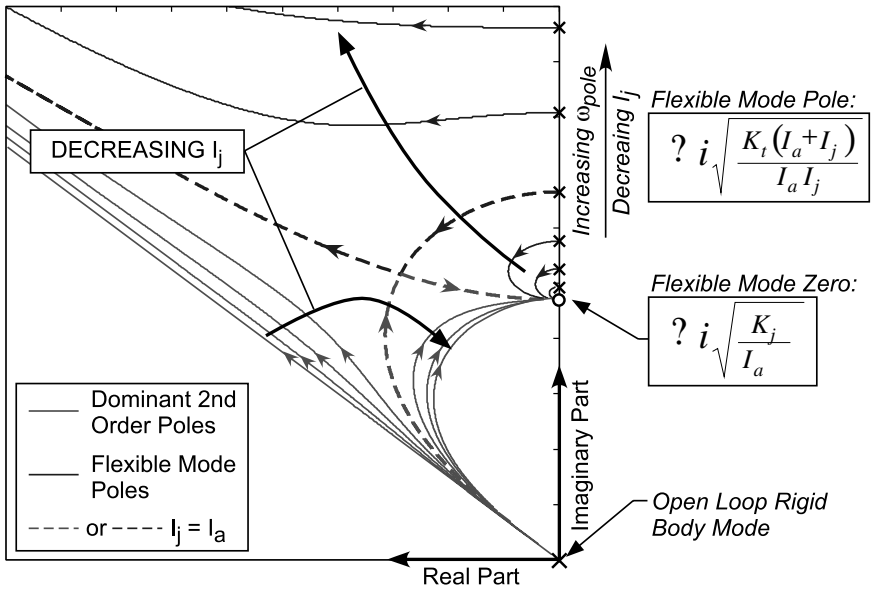

Fig. 18. Symmetric root locus of collocated position control system with shaft compliance.

3.2. Promising Results: Safety and Performance

To demonstrate the effectiveness of the DM^2 actuation approach, we have designed and built a two-axis prototype robotic arm, which incorporates the important characteristics of the DM^2 actuation approach. The overall design approach is shown in Figure 20.

Preliminary experimental and simulation results have demonstrated the effectiveness of the DM^2 actuation approach. The reduction in impact loading by an order of magnitude, as compared to conventional joint actuated manipulators, substantially improves the inherent safety of the manipulator. In the case of a two-axis prototype developed at Stanford (see Figure 20), the effective joint inertia was reduced by almost a factor of 10. We can use the effective inertia, graphically illustrated as a belted ellipsoid (Khatib 1995), to calculate the impulse due to impact at any point on the manipulator. To demonstrate the effectiveness of the DM^2 actuation approach in reducing impact loads, Figure 21 shows the normalized impact impulse for two cases of end-point load (P_{load}) for a 2-DOF planar manipulator. The impact impulse reduction increases rapidly with increasing load, as the required increase in actuator torque capability affects the reflected inertia of the conventional and cable-driven manipulators, while minimally affecting the reflected inertia of the distributed-parallel approach. While this is just an illustrative example, we see that, in combination with a lightweight structure and compliant covering, this new actuation approach can

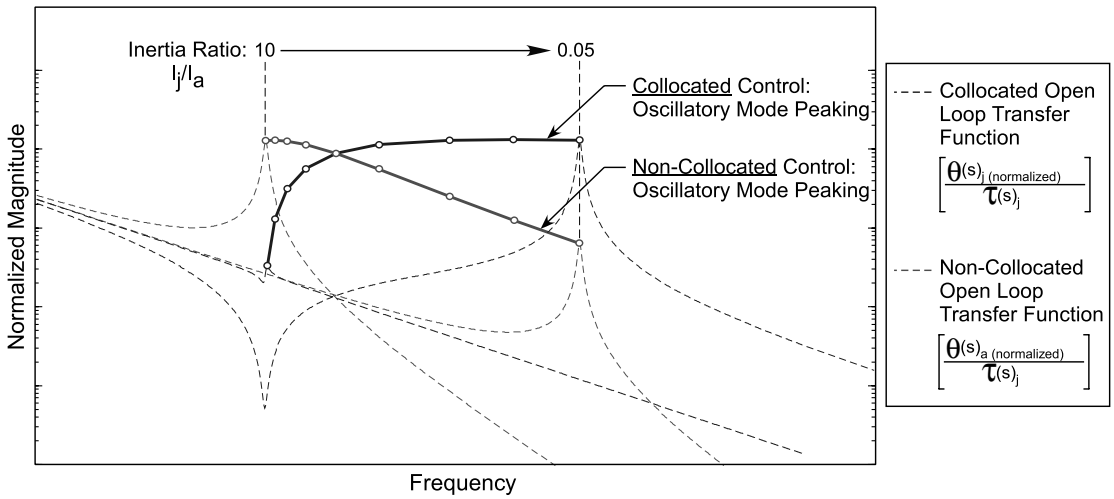

Fig. 19. Variation of peaking amplitude for collocated and non-collocated position control for varying motor to load inertia ratios, I_j/I_a.

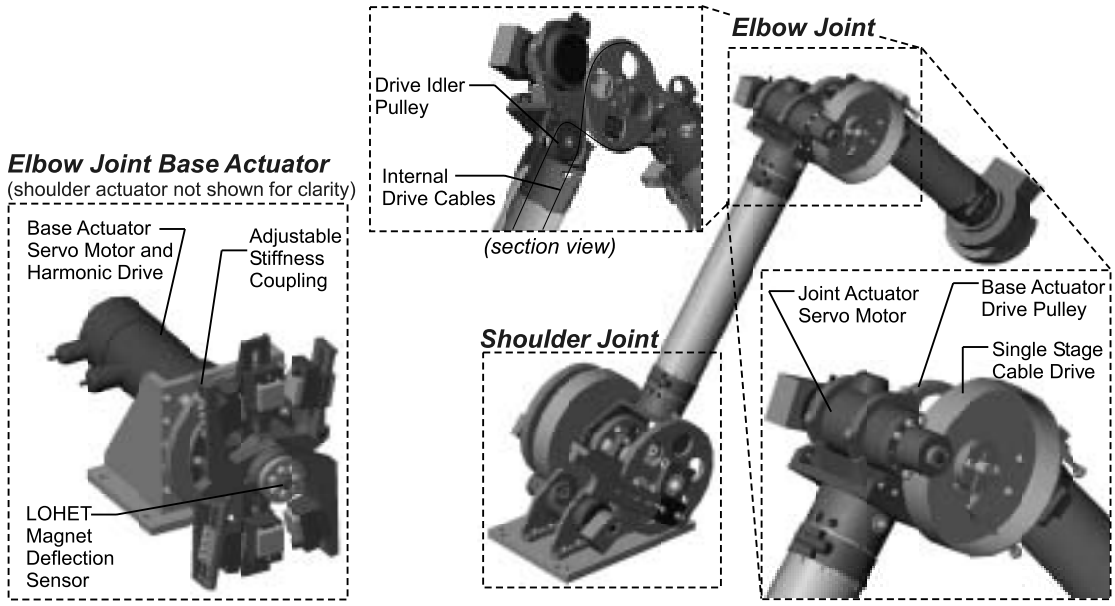

Fig. 20. Two-axis DM^2 prototype.

be used to design a manipulator that reduces impact loads substantially, thus ensuring inherent safety.

In addition to safety, the DM^2 actuation approach, with the introduction of the high-frequency joint actuator and implementation of the control approach described in Section 3.1, has been shown experimentally to improve manipulator performance. As shown in Figure 22, open-loop end-effector force control with the DM^2 actuation approach improves the speed of response over that of the base series elastic actuator alone. Both approaches have very low steady-state error due to their very low output impedance.

Trajectory tracking experiments carried out on the two-axis planar manipulator testbed demonstrate the feasibility of the DM^2 actuation approach. Initial experiments demonstrated a position control bandwidth of approximately 5 Hz as compared to a 2 Hz bandwidth using the base actuator alone (see Figure 23), reducing the position tracking error by more than a factor of 10. The higher achievable closed-loop position bandwidth allows the DM^2 actuated arm to accurately follow trajectories at rates that are not possible with the base actuator alone. Using the DM^2 two-axis testbed, we performed end-effector position tracking control experiments along a 15 cm

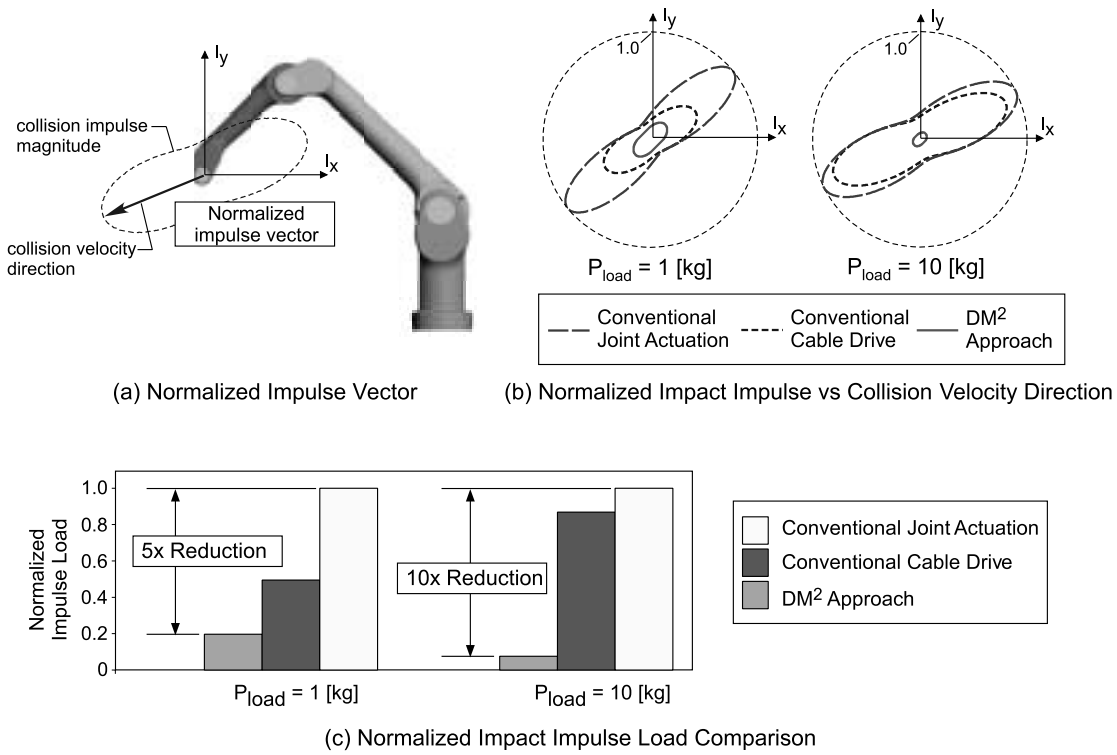

Fig. 21. Comparison of impulse load due to impact for various actuation concepts. (a) Normalized impulse vector. Impulse due to collision of manipulator end-effector with rigid object. Impulse magnitude changes with angle due to variation of end-effector effective inertia as a function of impact direction. (b) Normalized impact impulse versus collision velocity direction for various actuation concepts and values of end-point load (P_{load}). (c) Comparison of normalized impact impulse load for various actuation concepts and values of end-point load (P_{load}). Impulse values are normalized by impact velocity and maximum effective inertia.

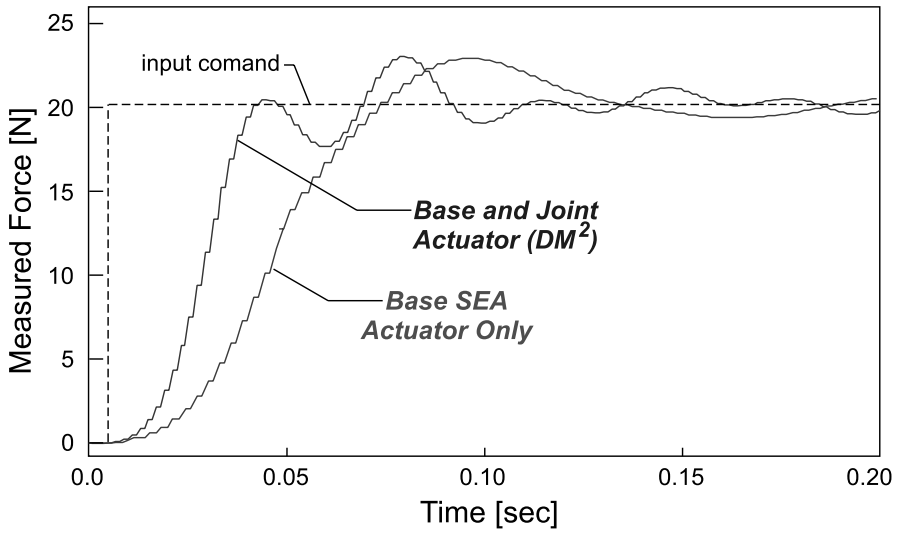

Fig. 22. Open-loop end-effector force (step) response.

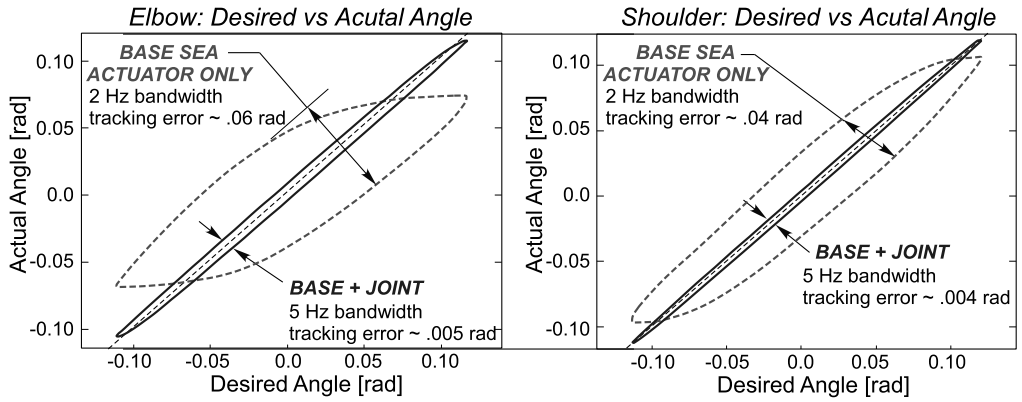

Fig. 23. Comparison of position tracking performance using base actuation only with combined base and joint actuation (DM^2).

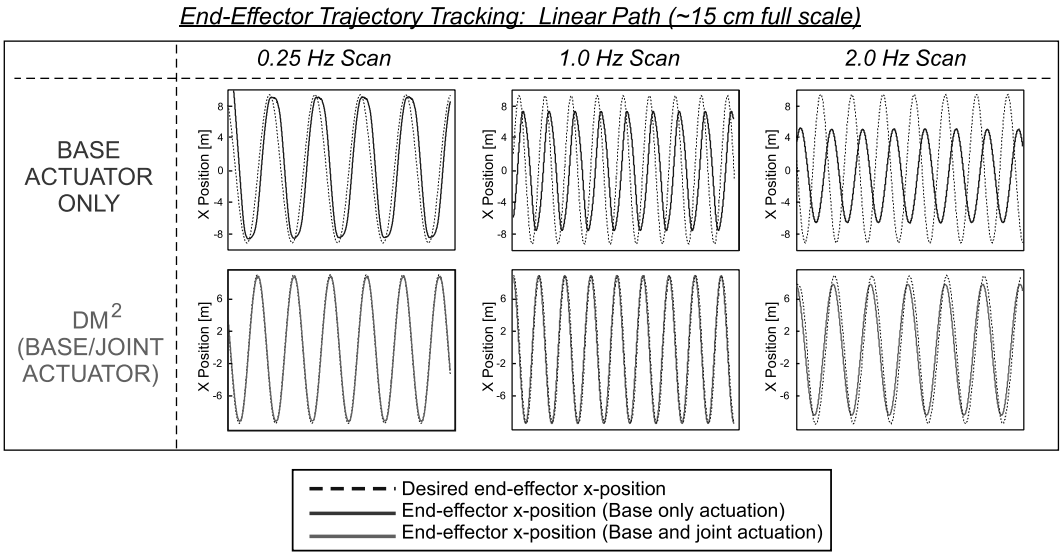

Fig. 24. End-effector position tracking control experimental results.

linear path at cycle rates of 0.25, 1.0, and 2.0 Hz. The results of the experiments, which contrast the DM^2 actuated and base (SEA) actuated performance, are shown in Figure 24. The DM^2 actuated testbed showed good tracking control for all three cases, with only a small amount of amplitude and phase distortion occurring during the 2.0 Hz rate experiment. The same experiment performed using the base actuators alone produced significant tracking error. During the 1.0 and 2.0 Hz rate experiments, significant phase and amplitude distortion were observed.

3.3. Distributed Macro–Mini Implementation

Finally, a few words should be said about the implementation of a DM^2 actuated robotic system. The DM^2 actuation approach is essentially a trade-off between safety, performance, and design complexity. However, this design trade is not necessarily a zero-sum game. Recall that the primary reason for the introduction of our new actuation approach was (1) to reduce contact impedance and (2) to maintain task performance levels. If the task is performed by a manipulator's end-effector, then high-frequency torque and force capabilities need only be provided at the end-effector. As shown in Khatib (1995), the dynamics of a redundant manipulator is bounded by the dynamics of the outermost degrees of freedom which span the task space. In the case of a redundant manipulation system, such as a dual manipulator–mobile base's system depicted in Figure 25, the mobile base degrees of freedom need not employ our new actuation approach to maintain task performance levels which, due to the redundancy of the system, are bounded by the outer six degrees of freedom. Another possible

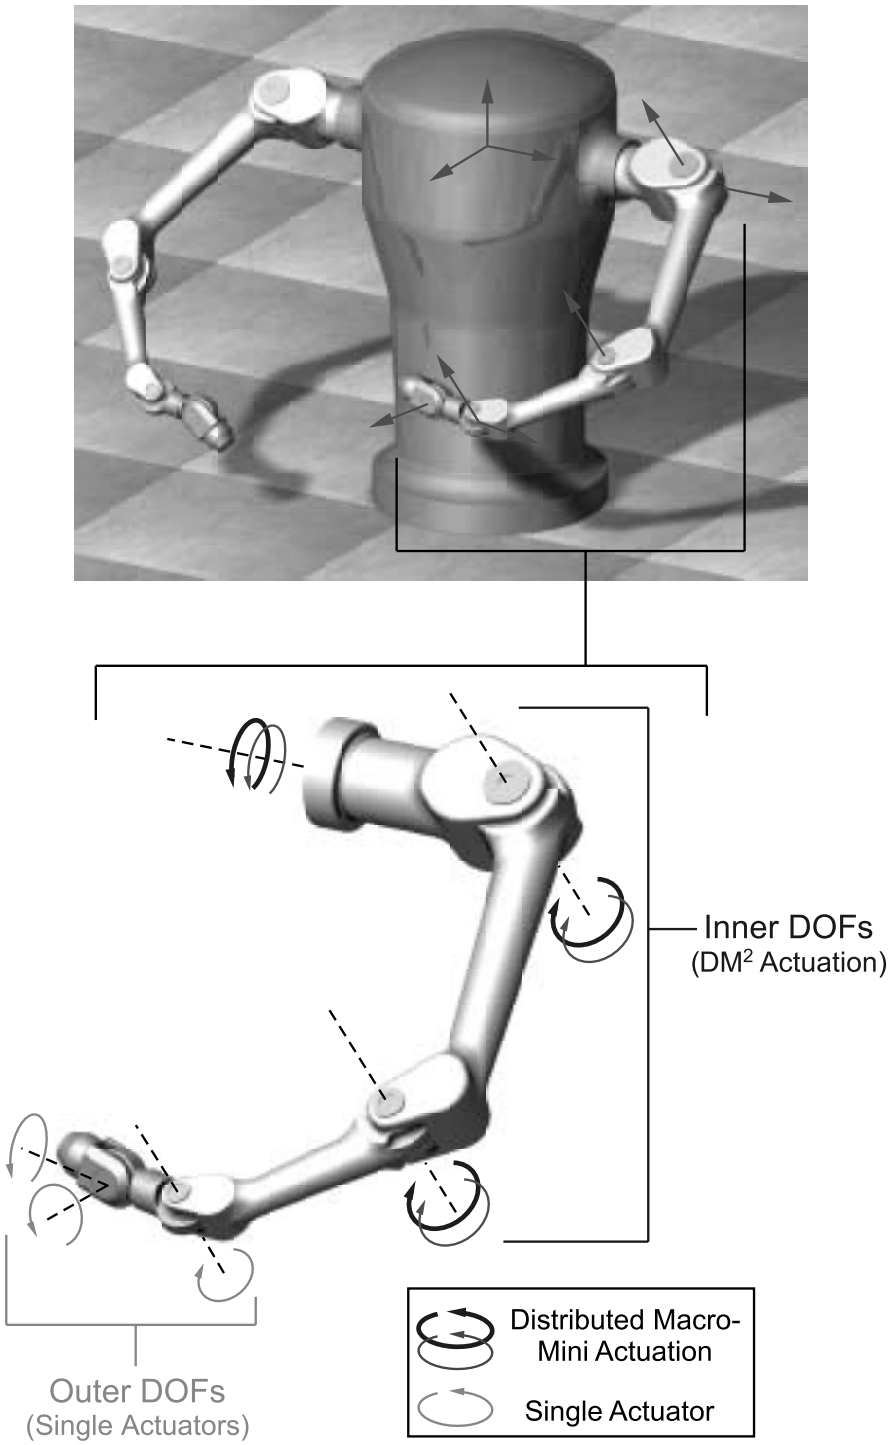

Fig. 25. Implementation of DM^2 actuation for multi-DOF manipulator.

approach is to design the wrist such that required task torques are small, as would be the case for a compact wrist design. In this case, the wrist actuation could be provided by smaller conventional EM actuators. The large DC and low-frequency torques provided by the base actuators of the DM^2 actuation approach would not be required. The higher impedance of the wrist actuators would not compromise safety because impact loads would be limited by the inner three degrees of freedom. Thus, our new human friendly actuation approach can be implemented in a manner which maximizes the safety and performance characteristics, while minimizing the additional complexity associated with its dual actuation approach.

4. Summary

We have presented a new actuation concept for human-friendly robot design, referred to as DM^2 actuation. The new concept (DM^2) was demonstrated on a 2-DOF prototype robot arm that we designed and built to validate our approach. The new actuation approach substantially reduces the impact loads associated with uncontrolled manipulator collision by relocating the major source of actuation effort from the joint to the base of the manipulator. High-frequency torque capability is maintained with the use of small, low-inertia servomotors collocated at the joints. The servomotors, integrated with a low-reduction, low-friction cable transmission, provide the high-frequency torque required for high-performance tasks while not significantly increasing the combined impedance of the manipulator–actuator system. The low output impedance and complete frequency coverage of the new actuation approach allow the combined manipulator system to approximate a pure torque source. This in turn allows for very good open-loop joint torque control over a wide frequency range. Initial experimental and simulation results validate the DM^2 actuation approach.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. EIA-9977717, the support of which is gratefully acknowledged. In addition, we would like to thank our colleagues Gunter Neimeyer and Kenneth Waldron, at Design Division, Department of Mechanical Engineering, Stanford University, Stanford, CA, for their helpful insights and discussion in preparing this paper.

References

Bicchi, A., Rizzini, L., and Tonietti, G. 2001. Compliant design for intrinsic safety: general issues and preliminary design. *Proceedings of the International Conference on Intelligent Robots and Systems*.

Cannon, R. H., and Schmitz, E. 1984. Initial experiments on the end-point control of a flexible one-link robot. *International Journal of Robotics Research* 3(3):62–75.

Ellis, G., and Lorenz, R. D. 2000. Resonant load control methods for industrial servo drives. *Proceedings of IEEE Industry Applications Society*, Rome, Italy.

Giralt, G., and Corke, P., editors. 2001. *Proceedings of the IARP/IEEE-RAS Joint Workshop on Technical Challenge for Dependable Robots in Human Environments*, Seoul, Korea.

Hirzinger, G., Albu-Schäffer, A., Hähnle, M., Schaefer, I., and Sporer, N. 2001. A new generation of torque controlled light-weight robots. *Proceedings of the International Conference on Robotics and Automation*.

Hirzinger, G., Sporer, N., Albu-Schäffer, A., Hähnle, M., and Pascucci, A. 2002. DLR's torque-controlled light weight robot: III. Are we reaching the technological limits now? *Proceedings of the International Conference on Robotics and Automation*, pp. 1710–1716.

Hollerbach, J., Hunter, I., and Ballantyne, J. 1991. *A Comparative Analysis of Actuator Technologies for Robotics*, MIT Press, Cambridge, MA, pp. 299–342.

Holmberg, R., Dickert, S., and Khatib, O. 1992. A new actuation system for high-performance torque-controlled manipulators. *Proceedings of the 9th CISM-IFToMM Symposium on the Theory and Practice of Robots and Manipulators*, Udine, Italy, pp. 285–292.

Khatib, O. 1987. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation* 3(1):43–53.

Khatib, O. 1995. Inertial properties in robotic manipulation: an object-level framework. *International Journal of Robotics Research* 14(1):19–36.

Morrell, J. B. 1996. Parallel coupled micro-macro actuators. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Pratt, G., and Williamson, M. 1995. Series elastic actuators. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 1, pp. 399–406.

Robinson, D. 2000. Design and analysis of series elasticity in closed-loop actuator force control. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Vischer, D., and Khatib, O. 1995. Design and development of high-performance torque-controlled joints. *IEEE Transactions on Robotics and Automation* 11(4):537–544.

Zinn, M., Khatib, O., Roth, B., and Salisbury, J. K. 2002. A new actuation approach for human friendly robot design. *Proceedings of the International Symposium on Experimental Robotics*, Sant'Angelo d'Ischia, Italy.