A New Actuation Concept for Human-Friendly Robot Design
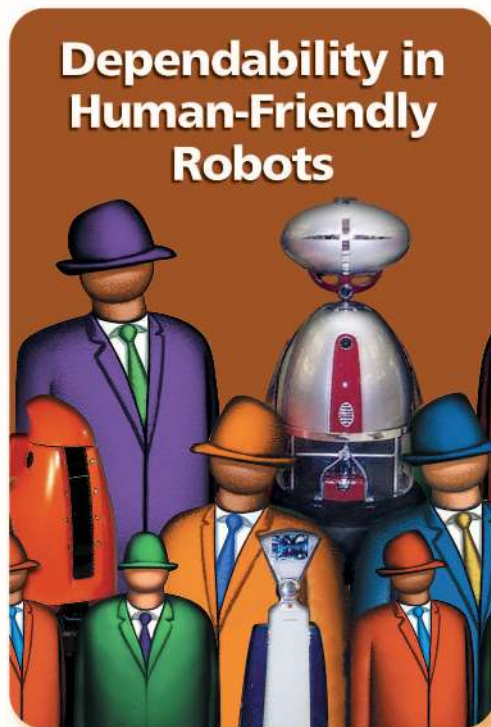
Playing It Safe

**BY MICHAEL ZINN, OUSSAMA KHATIB,
BERNARD ROTH, AND J. KENNETH SALISBURY**

Recently, there has been increasing interest in the emerging field of human-centered robotics. This field focuses on applications such as medical robotics and service robotics, which require close interaction between robotic manipulation systems and human beings, including direct human-manipulator contact. As a result, human-centered robotic systems must consider the requirements of safety in addition to the traditional metrics of performance. To achieve safety we must employ multiple strategies involving all aspects of manipulator design, including the mechanical, electrical, and software architectures. Immediate improvement can often be realized with the use of electronic hardware and software safety mechanisms that intelligently monitor and control manipulator operations. Additional improvements can be realized in the mechanical design. The elimination of pinch points and sharp edges can eliminate the potential for laceration or abrasion injuries. However, the most serious hazard present when working in close proximity with robotic manipulators is the potential for large impact loads, which can result in serious injury or death. 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-degree-of-freedom mass-spring model can be used to predict head accelerations that would occur during an uncontrolled impact. In combination with the HIC index, 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 1 m/s produces a maximum HIC greater than 500, more than enough to cause injury (see Figure 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.]

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. (For the PUMA robot, the required thickness of a compliant cover is more than 5 in, assuming an impact velocity of 1 m/s and an allowable maximum HIC index of 100.) Clearly, adding large amounts of compliant covering is impractical and does not address the root cause of high-impact loads; namely, the large

©1998 ARTVILLE, LLC

effective inertia of most modern robotic arms. Thus, the mechanical characteristics of a robotic system are the limiting factor in improving overall safety. The solution to reducing the effective impedance, and thus improving safety, is to build a lightweight, low-inertia manipulator.

Some types of robotic manipulators—notably those utilizing compliant actuation, such as pneumatic actuators, or those employing compliant drive trains, such as a cable-driven manipulators—do not produce the large-impact loads associated with high-impedance designs. 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 prevents 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 the first mode frequency. 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.

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 [11] and series elastic actuation [9]. 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 actuation approach (DM^2)—has been proposed [12].

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 [2], [5], [11].

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 (German Aerospace Center, Deutsches Zentrum fur Luft–und Raumfhart) lightweight arm design (see Figure 2) [3].
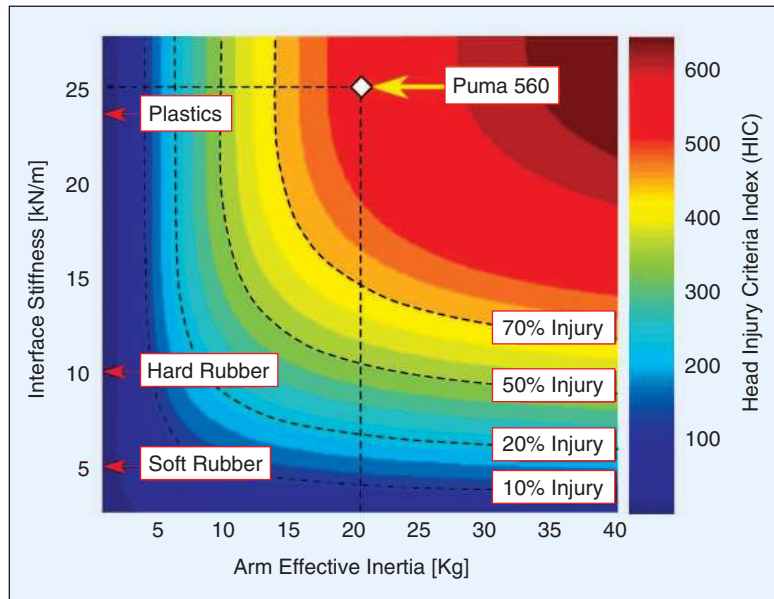

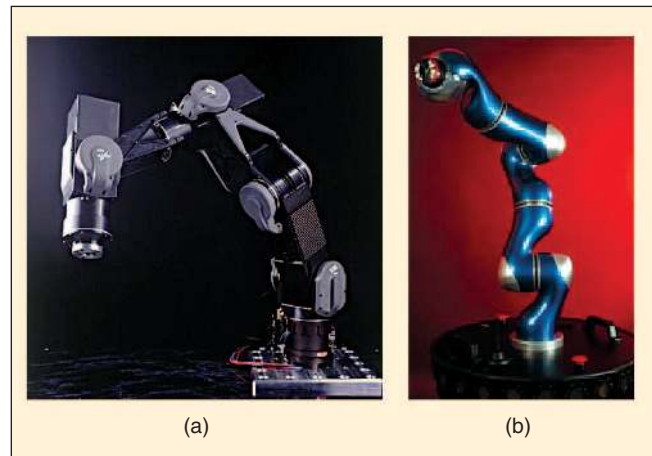**Figure 1.** *HIC as a function of effective inertia and interface stiffness.*

Figure 2. *DLR light-weight robot. (a) DLR II. (b) DLR III.*
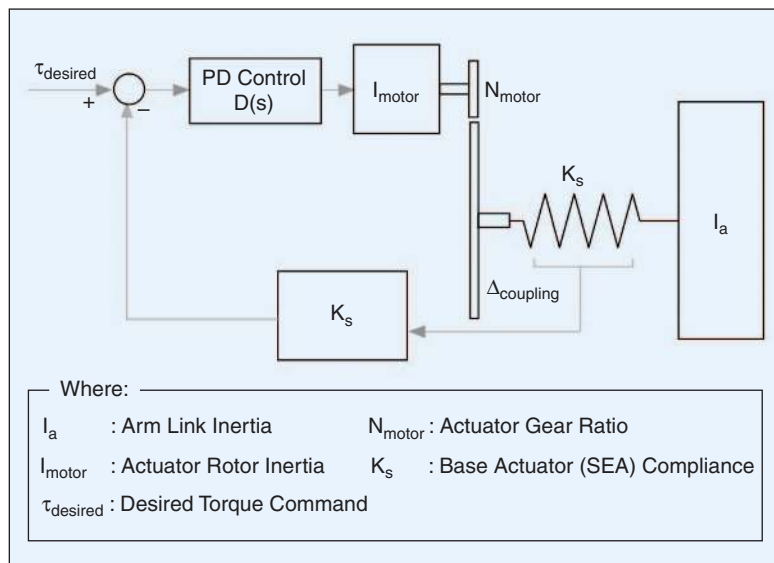

Figure 3. *SEA topology.*

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 magnitude 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.

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 [9], [10]. The SEA approach seeks to mitigate the limitations of high-impedance actuators, such as conventional gear-head electromagnetic or hydraulic actua-

tors, 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 a linear feedback system is implemented to regulate the output torque of the actuator-spring system (see Figure 3).

The main advantage of the SEA topology is that it provides low output impedance across the frequency spectrum. As shown in [9] and [10], the SEA topology reduces the output impedance of the SEA actuator in proportion with the stiffness of the elastic coupling. 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 $10x$–$100x$ 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 the previous section, which has good low-frequency impedance but suffers 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 3). 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 proportional-derivative (PD) controller. The resulting system is stable and possess low 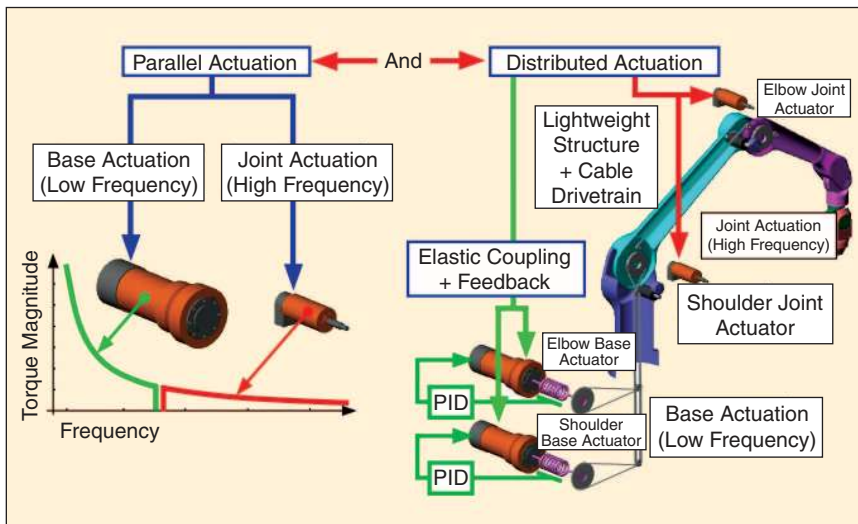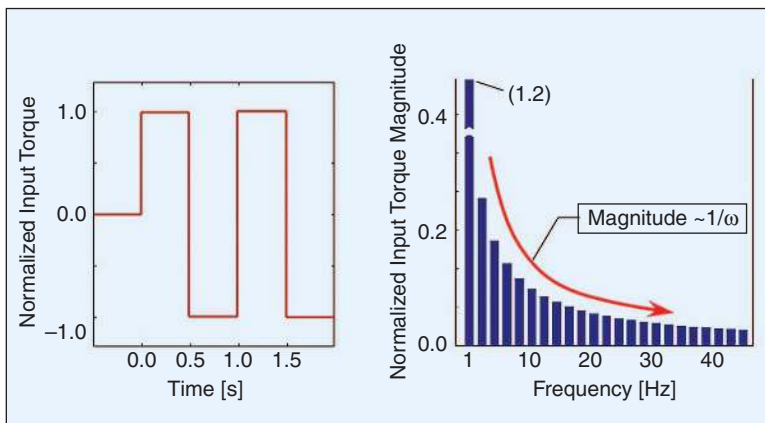impedance over a wide frequency range. In contrast, the coupling between the actuation and driven link is much stiffer for the joint-torque-control approach. Implementation of PD control, in this case, is difficult and requires that the control gains be kept low to maintain stability. As a result, alternative control schemes have been implemented, including proportional-integral (PI) control [11] and full-state feedback [3].

There are tradeoffs with using the SEA actuators. Due to the velocity and torque saturation of the SEA actuator, the maximum output torque above the open loop mode of the system (SEA open loop mode is the unforced coupled motion of actuator and manipulator link inertias through the compliant coupling) falls off as $1/\omega$ regardless

Figure 4. DM² actuation approach.

Figure 5. Torque versus frequency: 1-Hz² wave.

of the control loop controller bandwidth [10]. This behavior is an open-loop characteristic of the SEA actuator 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- and open-loop impedance characteristics.

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 actuator. This is not a major consideration for manipulation systems which do not require fast dynamics, such as walking robots for which the SEAs 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 SEAs are prohibitive. Other approaches have been proposed, such as the use of a nonlinear elastic coupling, whose compliance can be changed through coactivated actuators [1]. Unfortunately, the bandwidth limitations affecting the series elastic actuator, while mitigated somewhat by the variable compliance, is still a limiting factor in performance.

New Actuation Approach: Distributed Macro-Mini Actuation

To address the limitations of current actuation technology, we have proposed a new approach that seeks to relocate the major source of actuation effort from the joint to the base of the manipulator [12]. This can substantially reduce the effective inertia of the overall manipulator 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 summarized in Figure 4.

The first element of the actuation approach is to divide the torque generation into separate low- and high-frequency parallel actuators. The effectiveness of this approach can be seen clearly when one considers 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 5). This torque versus frequency profile is ideally fit using a large-output, low-frequency actuator coupled with a high-frequency servomotor.

In order for the DM^2 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 nondestructively 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 low-frequency actuation, we achieve low impedance by using a series elastic actuator (see the section "Series Elastic Actuation"). Because the DM^2 approach does not require that the base actuator be capable of supplying high-frequency torques, the bandwidth limitations of SEA actuators do not pose a difficulty.
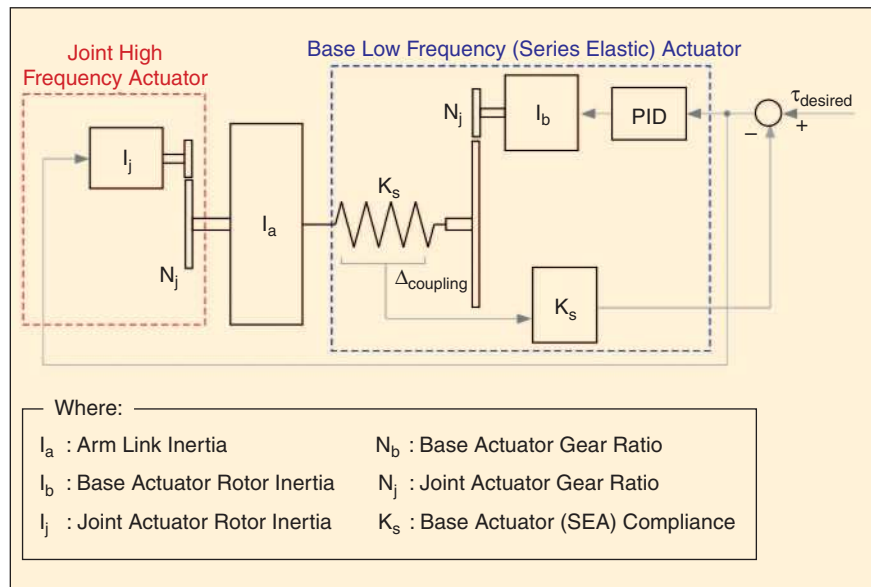

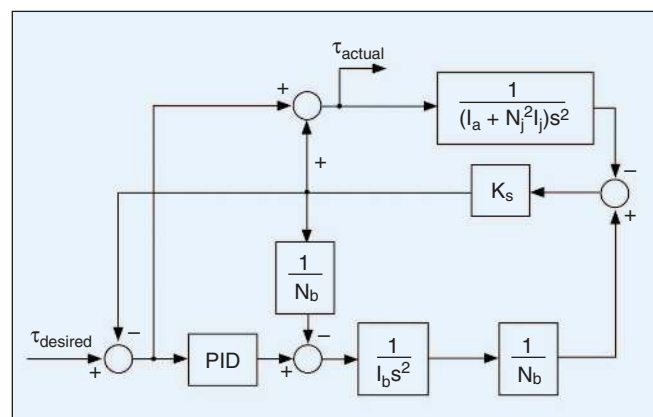**Figure 6.** DM^2 actuation and control topology (single DOF).

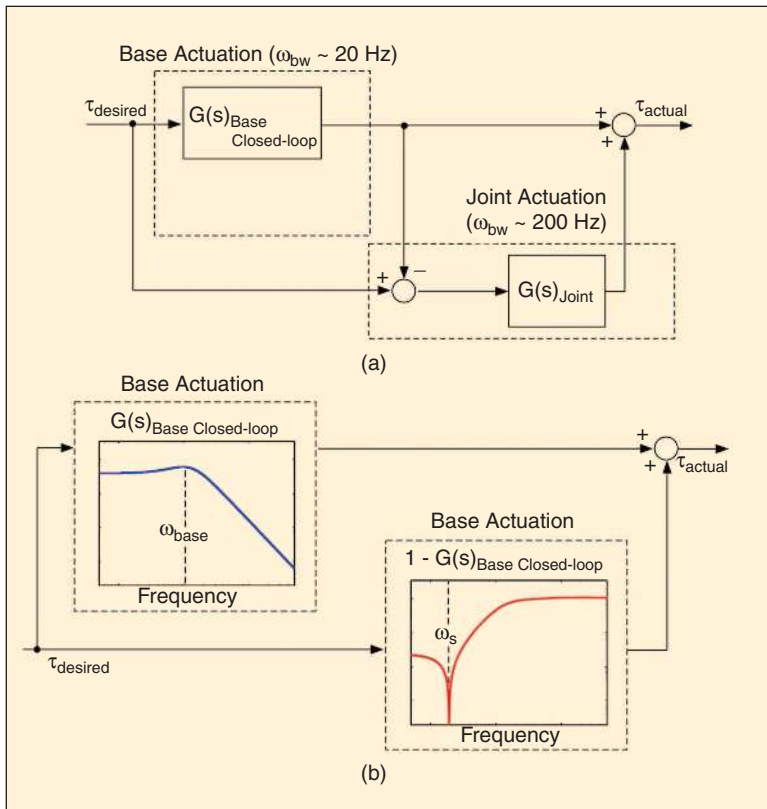**Figure 7.** DM^2 actuation and control block diagram representation (single DOF).

Figure 8. (a) DM² 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.

quency 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 actuator at the base significantly reduces the weight and inertia of the manipulator. The high-frequency actuators 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 that 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 DM² 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
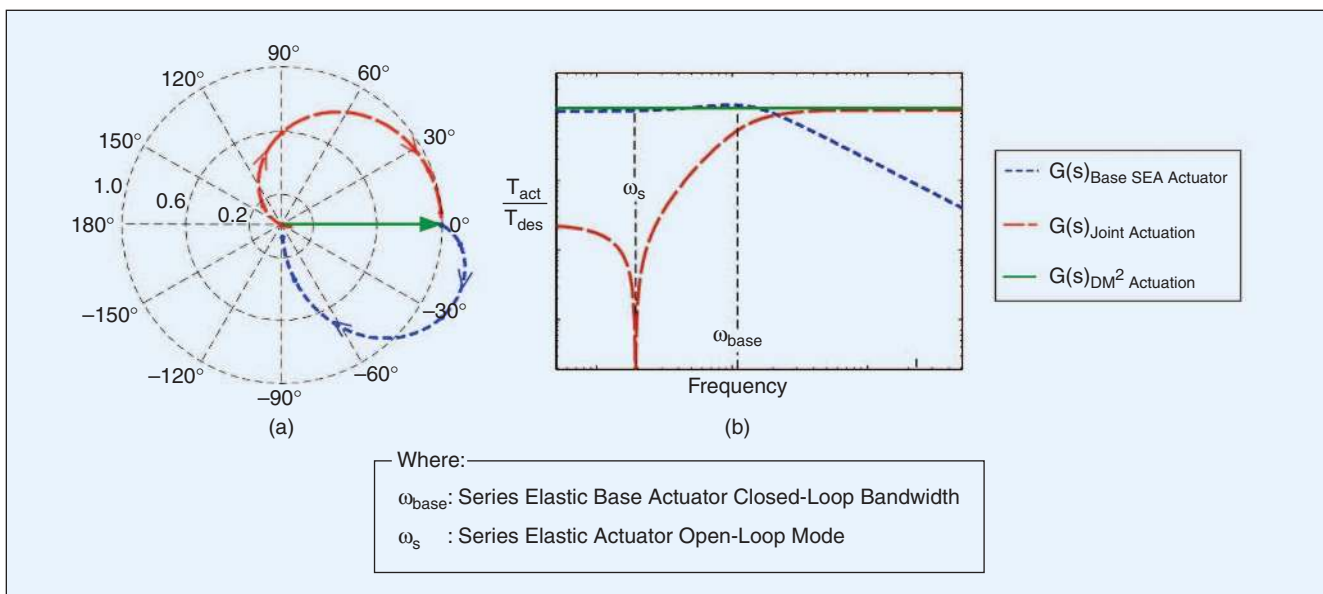
The second part of the DM² actuation approach, which differs from previous attempts at coupled actuation—the most notable of which is the parallel coupled macro–mini actuations approach [8]—is to distribute the low- and high-fre-

Figure 9. (a) Perfect torque source: base, joint, and combined DM² actuator torque magnitude versus phase polar plot. (b) Near-perfect torque source: base, joint, and combined DM² actuator torque magnitude versus frequency.

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. Thus, the impedance of the DM^2 approach is superior to that of current space robotic manipulators.

Actuation Control Approach

Perhaps the most challenging aspect of a DM^2 implementation is the development of a control approach that 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 approach is essentially a dual-input, single-output system. The redundant actuators provide an additional degree of freedom that 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 telerobotic master devices, we desire an actuation control scheme that allows the use of the parallel actuation system as a single torque source.

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.

A physical schematic of the control structure along with an equivalent block diagram representation are shown in Figures 6 and 7, respectively. The transfer function of the control structure shown in Figure 7 has unity gain and zero phase over all frequencies ($\tau_{actual}(s)/\tau_{desired}(s) = 1$).

A simplified representation, shown in Figure 8, 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 nonideal behavior of the base actuators. Assuming that the smaller joint actuators can produce this torque, the combined torques' sum is a perfect
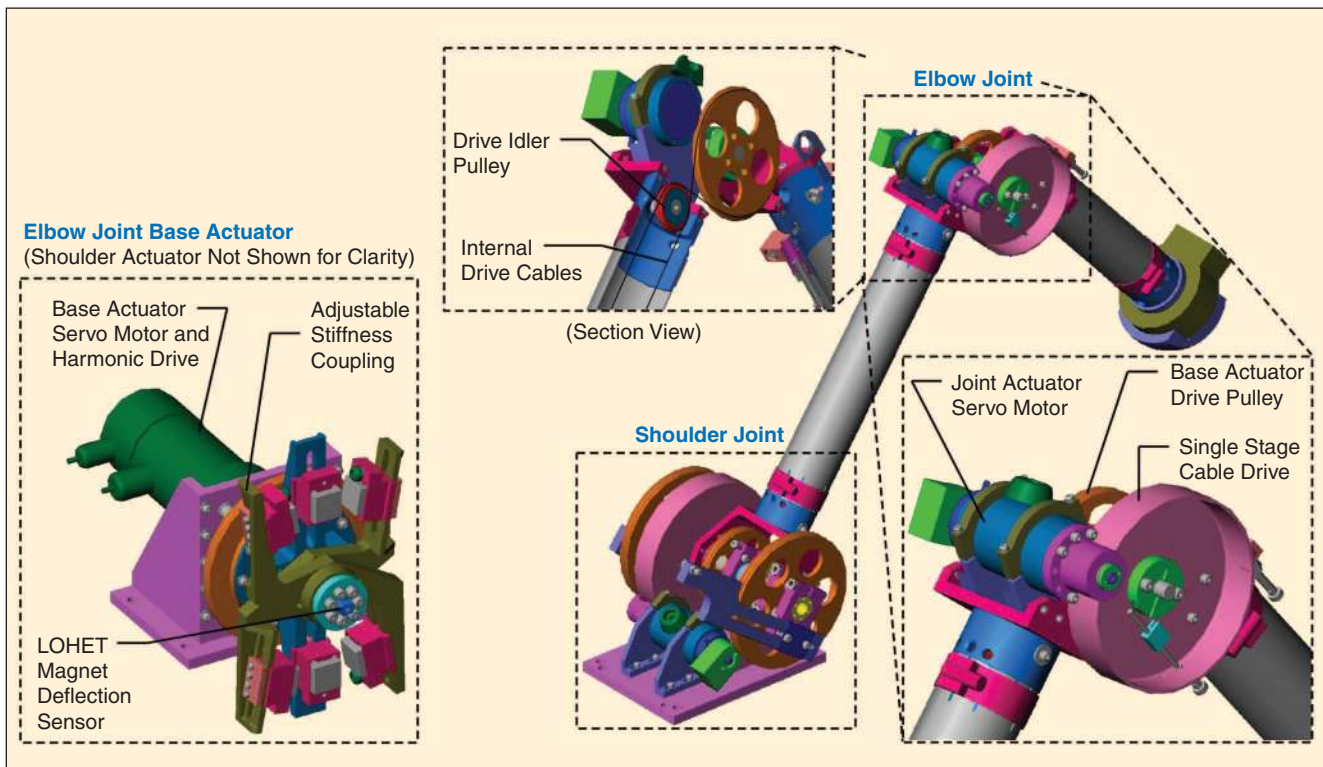

Figure 10. *Two-axis DM^2 prototype.*

realization of the desired torque. The frequency partitioning can be clearly seen if we rearrange the structure in Figure 8(a) into a pure parallel structure, as shown in Figure 8(b). As seen in Figure 8(b), the base actuator's transfer function falls off above its closed-loop bandwidth, ω_{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, ω_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 Figure 9(a) and (b). Thus, by using the simple control structure described above, we can create a unified actuator with the desirable characteristics of the low impedance necessary for inherent safety and the high bandwidth torque control necessary for high performance.

Manipulation Control

The DM^2 control structure allows for straightforward implementation of the DM^2 approach in multidegree-of-freedom manipulator systems. Assuming that the assumptions of a near-perfect torque source hold, the DM^2 approach is particularly well suited to control methods, such as operational space control [6], 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 6 and 7 are no longer valid. The main challenge in implementing the control scheme is in identifying and avoiding the situations where this ideal model breaks down (see [13] for more details).

Promising Results: Safety and Performance

To demonstrate the effectiveness of the DM^2 approach, we have designed and built a two-axis prototype robotic arm that incorporates the important characteristics of the DM^2 approach. The overall design approach is shown in Figure 10.

Preliminary experimental and simulation results have demonstrated the effectiveness of the DM^2 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 10), the
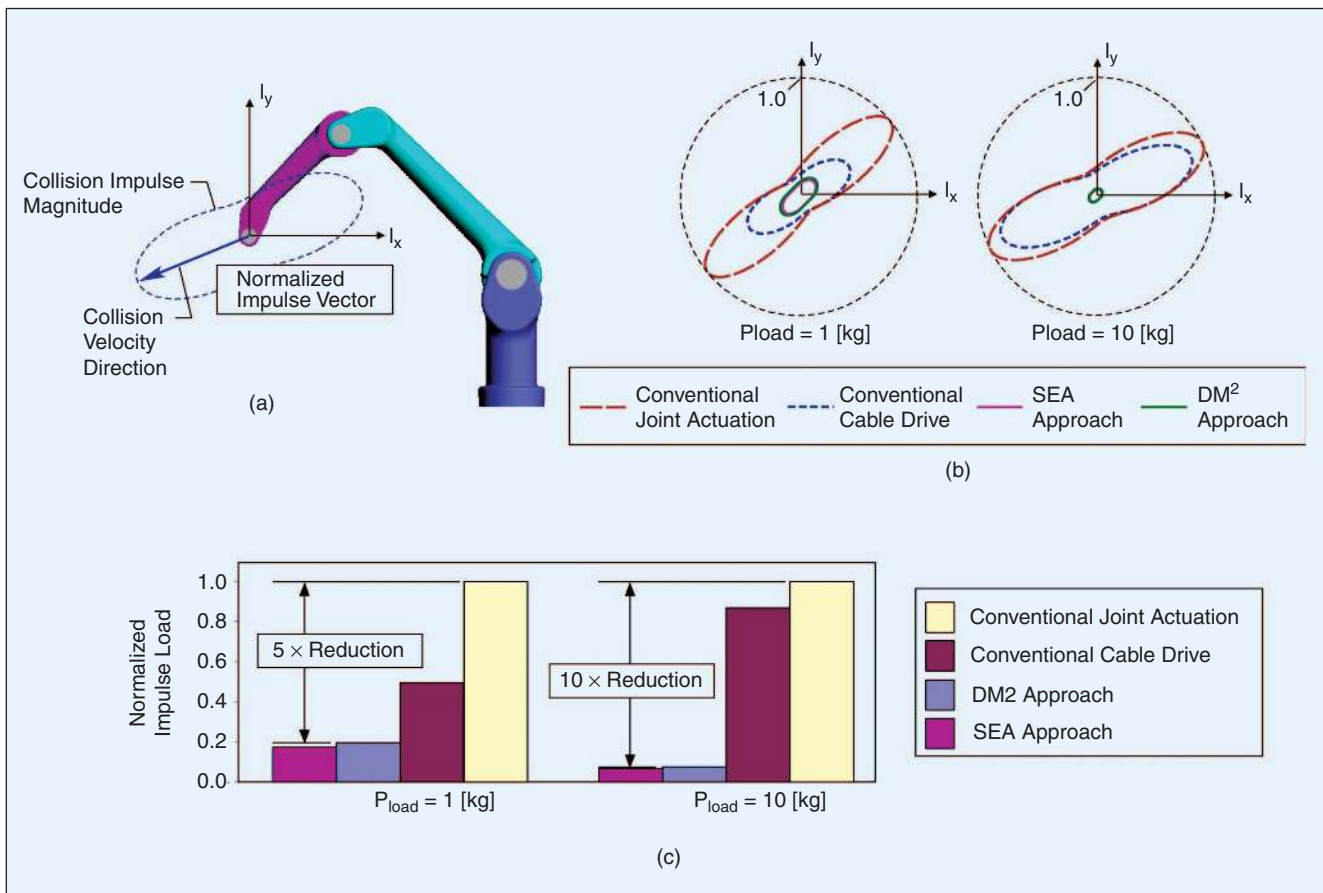

Figure 11. Comparison of impulse load due to impact for various actuation concepts. (a) Normalized impulse vector: impulse due to collision of manipulator end effector with rigid object. Impulse magnitude changes with angle due to variation of end-effector effective inertia as a function of impact direction. (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.

effective joint inertia was reduced by almost a factor of ten. We can use the effective inertia, graphically illustrated as a belted ellipsoid [7], to calculate the impulse due to impact at any point on the manipulator. To demonstrate the effectiveness of the DM^2 approach in reducing impact loads, Figure 11 shows the normalized impact impulse for two cases of end-point load (P_{load}) for a two-degree-of-freedom 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 both the DM^2 and SEA approaches. Because the DM^2 joint actuators are lightweight and have very small reflected inertia, their contribution to the effective inertia is minimal. As a result, the improved performance that results from the addition of the small joint actuators does not compromise the safety levels of the DM^2 as compared to the SEA actuated manipulator. While this is just an illustrative example, we see that in combination with a lightweight structure and compliant covering, this new actuation approach can be used to design a manipulator that reduces impact loads substantially, thus ensuring inherent safety.

In addition to safety, the DM^2 approach, with the introduction of the high-frequency joint actuator and implementation of the control approach described in the section "DM^2 Actuation Control Approach," has been shown experimentally to improve manipulator performance. As shown in Figure 12, open-loop end-effector force control with the DM^2 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 approach. Initial experiments demonstrated a position control bandwidth of approximately 5 Hz as com-
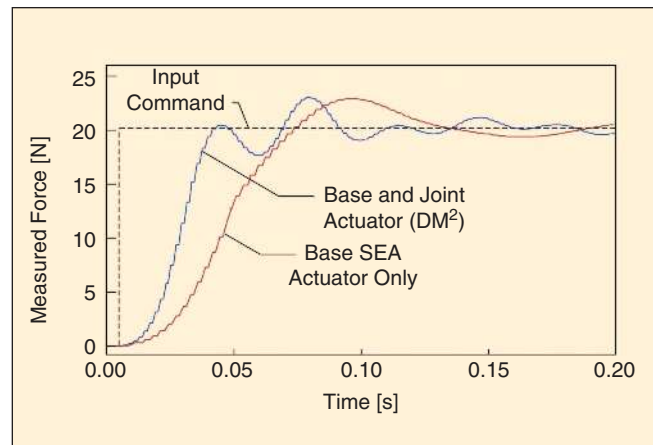

Figure 12. *Open-loop end-effector force (step) response.*

pared to a 2-Hz bandwidth using the base actuator alone (see Figure 13), reducing the position tracking error by more than a factor of ten. 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 two-DM^2 axis testbed, we performed end-effector position tracking control experiments along a 15-cm linear path at cycle rates of .25 Hz, 1.0 Hz, and 2.0 Hz. The results of the experiments, which contrast the DM^2 actuated and base (SEA) actuated performance, are shown in Figure 14. 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-Hz and 2.0-Hz rate experiments, significant phase and amplitude distortion were observed.
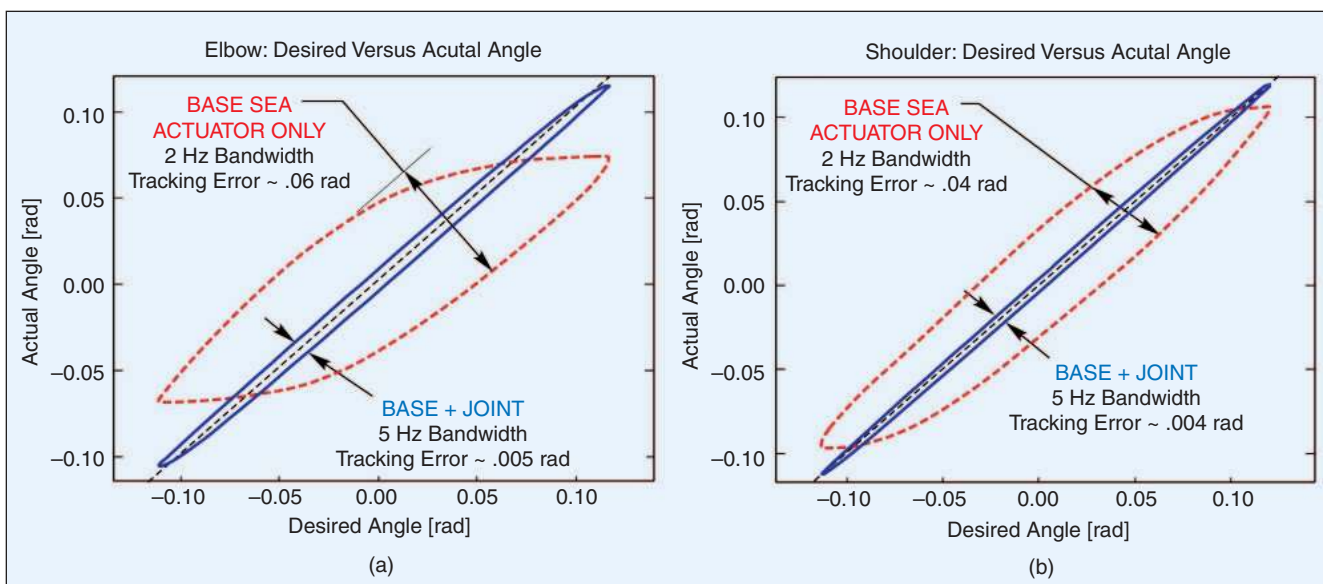

Figure 13. *Comparison of position tracking performance using base actuation only with combined base and joint actuation (DM^2).*
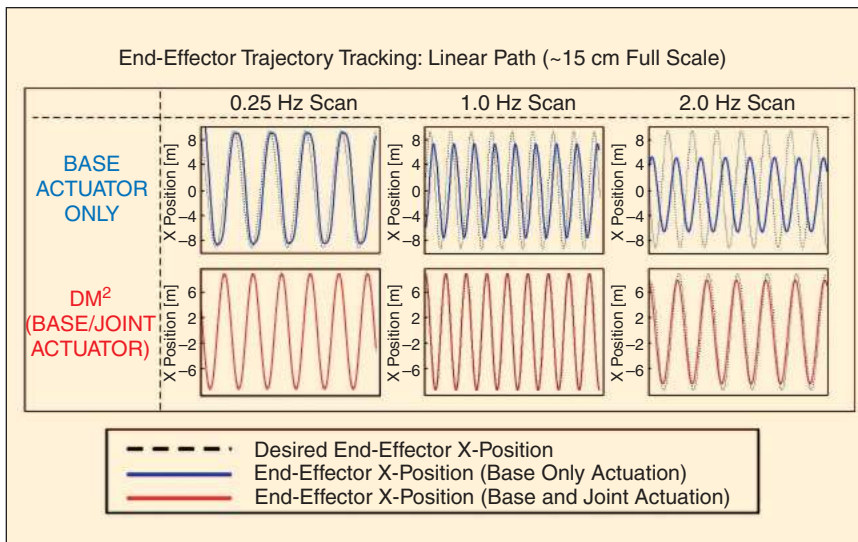
Figure 14. *End-effector position tracking control experimental results.*
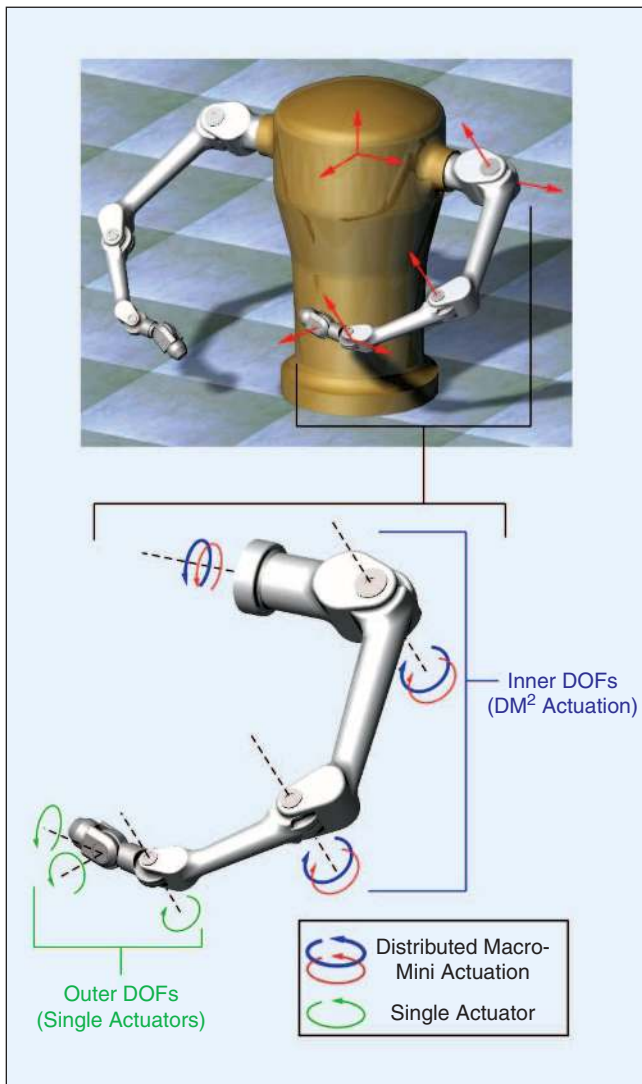

Figure 15. *Implementation of DM² actuation for multi-DOF manipulator.*

Distributed Macro-Mini Implementation

Finally, a few words should be said about the implementation of a DM^2 actuated robotic system. The DM^2 approach is essentially a tradeoff between safety, performance, and design complexity. However, this design trade is not necessarily a zero-sum game.

The primary reason for the introduction of our new actuation approach were to reduce contact impedance and 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 [7], the dynamics of a redundant manipulator is bounded by the dynamics of the outermost degrees of freedom that span the task space. In the case of a redundant manipulation system, such as a dual manipulator–mobile base's system depicted in Figure 15, 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 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 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.

Summary

We have presented a new actuation concept for human-friendly robot design, referred to as DM^2. The new concept of DM^2 was demonstrated on a two-degree-of-freedom 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 approach.

Acknowledgments

Keywords

Human-friendly, low impedance, safety actuation, parallel actuation, distributed actuation.

References

[1] A. Bicchi, L. Rizzini, and G. Tonietti, "Compliant design for intrinsic safety: General issues and preliminary design," in *Proc. Int. Conf. Intell. Robots Syst.,* Maui, HI, 2001, pp. 1864–1869.

[2] G. Hirzinger, A. Albu-Schäffer, M. Hähnle, I. Schaefer, and N. Sporer, "A new generation of torque controlled light-weight robots," in *Proc. Int. Conf. Robotics Automation*, Seoul, Korea, 2001, pp. 3356–3363.

[3] G. Hirzinger, N. Sporer, A. Albu-Schaffer, M. Hahnle, and A. Pascucci, "DLR's torque-controlled light weight robot III—Are we reaching the technological limits now?" in *Proc. Int. Conf. Robotics Automation*, 2002, pp. 1710–1716.

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

[5] R. Holmberg, S. Dickert, and O. Khatib, "A new actuation system for high-performance torque-controlled manipulators," in *Proc. 9th CISM-IFToMM Symp. Theory Practice Robots Manipulators*, Udine, Italy, Sept. 1992, pp. 285–292.

[6] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE J. Robot. Automat.*, vol. RA-3, pp. 43–53, Feb. 1987.

[7] O. Khatib, "Inertial properties in robotic manipulation: An object-level framework," *Int. J. Robotics Res.*, vol. 14, no. 1, pp. 19–36, Feb. 1995.

[8] J.B. Morrel, "Parallel coupled micro–macro actuators," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1996.

[9] G. Pratt and M. Williamson, "Series elastic actuators," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, vol. 1, Pittsburgh, PA, 1995, pp. 399–406.

[10] D. Robinson, "Design and analysis of series elasticity in closed-loop actuator force control," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, June 2000.

[11] D. Vischer and O. Khatib, "Design and development of high-performance torque-controlled joints," *IEEE Trans. Robot. Automat.*, vol. 11, pp. 537–544, Aug. 1995.

[12] M. Zinn, O. Khatib, B. Roth, and J.K. Salisbury, "A new actuation approach for human friendly robot design," in *Experimental Robotics VIII, Springer Tracts in Advanced Robotics,* B. Siciliano and P. Dario, Eds. Berlin: Spinger-Verlag, 2002.

[13] M. Zinn, O. Khatib, B. Roth, and J.K. Salisbury, "Actuation methods for human-centered robotics and associated control challenges," in *Control Problems in Robotics*, A. Bicchi, H. Christensen, and D. Prattichizzo, Eds. Berlin, Germany: Springer-Verlag, 2003.

Michael Zinn received his M.S. and B.S. degree from M.I.T. in 1988 and 1987, respectively, and is currently completing his Ph.D. in mechanical engineering at Stanford University. He has extensive experience in mechanical and controls design with more than 15 years experience designing complex electro-mechanical systems. His current research interests include human–centered robotics and medical device robotics.

Oussama Khatib received his Ph.D. in 1980 from Sup'Aero, Toulouse, France. His current research is in human-centered robotics, human-friendly robot design, dynamic simulations, and haptic interactions. His exploration in this research ranges from the autonomous ability of a robot to cooperate with a human to the haptic interaction of a user with an animated character or a surgical instrument. Prof. Khatib is the president of the International Foundation of Robotics Research, IFRR, and co-editor of the Springer Tracts in Advanced Robotics. He is a "Distinguished Lecturer" of the IEEE and a recipient of the JARA Award.

Bernard Roth has been a faculty member of Stanford University's Mechanical Engineering Department since 1962. He teaches courses in design, robotics, kinematics, and the 24 societal and personal aspects of technology. He is an active researcher with over 150 publications in the areas of robotics, kinematics, and design. He and his students have been pioneers in the design and construction of robotic device and in the development of a rational theory for robot design and control. He has served as president of the International Federation for the Theory of Machines and Mechanisms and as chairman of the Design Engineering Division of the American Society of Mechanical Engineers. He is a director of the International Federation for Robotics Research. Dr. Roth is the recipient of many awards for his research and teaching, including The Joseph F. Engleberger Award for Robotics, the Machine Design Award, the Melville Medal, five Best Paper Awards from the American Society of Mechanical Engineers, and two Japanese Society for the Promotion of Science Awards.

J. Kenneth Salisbury received his Ph.D. from Stanford University in mechanical engineering in 1982. At MIT from 1982-1999, he served as principal research scientist in mechanical engineering and as a member of the Artificial Intelligence Laboratory. Some of the projects with which he has been associated include the Stanford-JPL Robot Hand, the JPL Force Reflecting Hand Controller, the MIT-WAM arm, and the Black Falcon Surgical Robot. His work with haptic interface technology led to the founding of SensAble Technologies Inc., producers of the PHANTOM haptic interface and 3D FreeForm software. In 1997, he joined the staff at Intuitive Surgical, in Mountain View, California, where his efforts have focused on the development of telerobotic systems for the operating room. In 1999 he joined the faculty at Stanford in the Departments of Computer Science and Surgery, where his research focuses on human centered robotics, cooperative haptics, and surgical simulation. He currently serves on the National Science Foundation's Advisory Council for Robotics and Human Augmentation as scientific advisor to Intuitive Surgical, Inc. and as technical advisor to Robotic Ventures, Inc.

Address for Correspondence: Michael Zinn, Design Division, Department of Mechanical Engineering, Stanford University, Stanford, CA. E-mail: zinn@robotics.stanford.edu.