By

Benjamin Conrad

A dissertation submitted towards the requirements for the degree of

Doctor of Philosophy
(Mechanical Engineering)

at the

UNIVERSITY OF WISCONSIN - MADISON

2016

Date of oral examination: November 21, 2016

The preliminary is approved by the following members of the Final Oral Committee:

Michael R. Zinn, Associate Professor, Mechanical Engineering

Darryl G. Thelen, Professor, Mechanical Engineering

Heidi Ploeg, Associate Professor, Mechanical Engineering

Peter G. Adamczyk, Assistant Professor, Mechanical Engineering

Matthew S. Allen, Associate Professor, Engineering Physics

Abstract

Continuum manipulator compliance enables operation in delicate environments at the cost of challenging actuation and control. In the case of catheter ablation of atrial fibrillation, the compliance of the continuum backbone lends an inherent safety to the device. This inherent safety frustrates attempts at precise, accurate, and fast control, limiting these devices to simple, static positioning tasks. This dissertation develops Interleaved Continuum-Rigid Manipulation, by which the hysteretic nonlinearities encountered in tendon-actuated continuum manipulators are compensated by discrete rigid joints located between continuum sections. The rigid joints introduce actuation redundancy, which an interleaved controller may use to avoid continuum nonlinearities and dynamic excitations, or to prefer particular configurations that may improve task accuracy, permit greater end-effector forces, or avoid environment obstacles. Two experimental systems explore the potential of these joints to 1) correct for actuation nonlinearities and enhance manipulator performance and 2) increase the manipulator's dexterous workspace. These experiments expose important design and control observations that were not apparent in the general robotic and continuum literature.

Acknowledgments

My ten years in Madison have been great; I have been here longer, and have done, experienced, and learned more than I ever expected. Through these years, the love, support, and interest of my parents, Joe and Connie, and brother Dan has been constant and essential. A number of people chose to make this experience and my growth possible; I could not have done this without the support, encouragement, and patience of my advisor Professor Mike Zinn; the Catholic community that Fr. Eric Nielsen has built (and is now building) at St. Paul's; and the enthusiasm, interest, and simple presence of my friends and current and former lab mates.

Contents

1	Cor	nplian	t, Dexterous Robotic Manipulation	1
	1.1	Minim	nally-Invasive Surgery	2
		1.1.1	Performing Cardiac Ablation	5
		1.1.2	Operating In the Heart	6
		1.1.3	Operating in the <i>Beating Heart</i>	8
	1.2	Contin	nuum Manipulators	10
		1.2.1	The Physics of Tendon-Actuation	14
		1.2.2	Accurate Modeling is Difficult	16
		1.2.3	Design Yields Only Modest Improvements	17
		1.2.4	Advanced Control Can Perhaps Double Performance	18
2	Inte	erleave	d Continuum-Rigid Manipulation	22
	2.1	Possib	pilities of Interleaved Continuum Rigid Manipulation	23
3	An	Interle	eaved Rigid Joint and Continuum Segment	25
•	1111			
	3.1	Overv	iew of the Single DOF Testbed	25
	3.2	A Sing	gle-Input, Multiple-Output Controller	27
	3.3	Single	DOF Performance	30

	3.4	Single DOF Discussion	36
4	A (Clinically-Approximate Five DOF Manipulator	37
	4.1	The Mechanical Design	38
	4.2	Kinematic Modeling and Parameter Spaces	41
	4.3	An Adaptive Closed Loop Controller	44
		4.3.1 Bucketing	46
		4.3.2 Distributing the Buckets	48
		4.3.3 Minimizing the Kinematic Error in a Bucket	49
	4.4	5D Evaluation Overview	51
	4.5	Learning a Half-Cube	51
		4.5.1 Convergence to the Command	52
		4.5.2 Open Loop Variations	52
		4.5.3 Closed Loop Variations	56
		4.5.4 What Are the Consensus Kinematic Parameters?	60
	4.6	Moving Within a Known Space	62
	4.7	Performance	64
	4.8	In Summary	68
5	In (Conclusion	69
Bi	Bibliography 7		
A	Con	stant Curvature Catheter Geometry	83
В	5D	Joint Performance	85
	B.1	FRFs Indicate Control Bandwidth	85

B.2	Hysteresis Loops	86
В.3	Nonlinearities in the Base Roll	89
	B.3.1 Nonlinearities in the Rigid Joints	90
	B.3.2 Nonlinearities in the Catheter Articulation	91
C Init	ial Experiments with Closed Loop Control	93
C.1	5D Forward Kinematics	93
C.2	A Multiple DOF Controller	94
	C.2.1 Tracking a Virtual Target	97
	C.2.2 On the importance of output compliance	103
D Prio	or Publications	104

List of Figures

1.1	Minimally Invasive Laparoscopic Surgery	3
1.2	Elements of cardiac cateher interventions	7
1.3	Working in the left atrium	Ĉ
1.4	Manipulators in the literature	12
1.5	Tendon actuation	15
2.1	An Interleaved Continuum-Rigid Manipulator	22
3.1	Single DOF Testbed	26
3.2	Single DOF Controller	28
3.3	Single DOF Frequency Response	30
3.4	Step Response At Low Articulation	32
3.5	Step Response At High Articulation	33
3.6	Varying Rigid Joint Saturation	34
3.7	Step Response Under Increasing Saturation	34
3.8	Frequency Response Of Saturated System	35
4.1	An 5 DOF Interleaved Manipulator	38
4 2	Kinematic Frames	49

4.3	An adaptive closed loop controller	46
4.4	Initial Iteration of the Open Loop Variations	53
4.5	Open/Few/Intuitive Convergence	54
4.6	Open/Many/Intuitive Convergence	55
4.7	Open/Many/Minimal Convergence	55
4.8	Open/Few/Minimal Convergence	56
4.9	Closed/Few/Intuitive Convergence	57
4.10	Closed/Many/Intuitive Convergence	58
4.11	Closed/Many/Minimal Convergence	59
4.12	Closed/Few/Minimal Convergence	59
4.13	Accumulated Closed Loop Error	60
4.14	Non-Repetitive Training Shapes	63
4.15	Closed Loop Point Moves	63
4.16	Tracking a Sprial	65
4.17	Recasting Kinematic Parameters Across Buckets	66
4.18	Effects of Recasting in Joint and Task Space	67
A.1	Constant Curvature Catheter Geometry	84
B.1	Joint Frequency Responses	86
B.2	Base, Pitch, and Roll Hysteresis	87
В.3	Catheter Hysteresis	89
B.4	Base Roll Nonlinearities	90
B.5	Axial Alignment of the Pitch	91
В6	Non-Constant Curvature	92

	٠	٠	
17	1	1	

C.1	A Multi-DOF Interleaved Controller	95
C.2	Joint Space Convergence And Tracking	99
C.3	5D Task-Space Tracking	100
C.4	Indications Of Cable Stretch	102

List of Tables

4.1	Measured FK Parameters	43
4.2	Experimental controller configurations	52
4.3	Final Intuitive Kinematics, 10th	61
4.4	Final Intuitive Kinematics, 40th	61
4.5	Final Minimal Kinematics, 10th	62
4.6	Final Minimal Kinematics, 40th	62
C.1	Measured And Estimated FK Parameters	.01

1. Compliant, Dexterous Robotic Manipulation

Drive-to-end-effector stiffness is a fundamental attribute of any robotic manipulator, as it captures actuator, transmission, joint, and link stiffness. Maximizing stiffness, while maintaining the desired dynamic performance, naturally leads to manipulators composed of rigid links and revolute or prismatic joints with all degrees of freedom (DOF) concentrated at defined joints whose displacements are easily measured. Stiff joint and link construction greatly simplifies the manipulator modeling and control, to the point that the vast majority of industrial manipulators are operated open loop. Stiff manipulators do not accept disturbances gracefully, leading most to be enclosed in safety cages. The safety cage limits the robot's environment to only those elements that have been modeled and with which it must interact, say to a supply of widgets, a widget receptacle, and the operations to convey a widget to the receptacle. Across industries, robots operating in small environments – where few of the world's complexities penetrate the safety cage – have reached maturity and ongoing innovation comes from fundamental actuator, material, and fabrication advances.

In contrast to stiff, rigid link manipulators, continuum manipulators permit bending along their body due to any applied force – actuator or otherwise. To suffer intentional and inadvertent contacts requires flexibility along the manipulator; this inherent, physical compliance reduces the likelihood of external- or self-damage during contact at the cost of a fundamentally underactuated

¹ Without task feedback. Actuator encoders measure the actuator position but are ignorant of any transmission and link compliance after that measurement. Similarly joint encoders measure transmission, but not link, compliance. Estimating the task position from these measurements requires assuming the effect of link and transmission compliance is negligible with respect to the task magnitude. If only the drives are controlled, then there is no control loop on the actual task position. High link and joint stiffness permits this assumption for the majority of industrial manipulators.

manipulator with essentially infinite DOFs. It should not surprise, then, that the design techniques and control algorithms that produce safety-caged robots are insufficient when operating outside the cage, let alone inside a human.

Robotic cardiac systems operate within a patient's beating heart, accessing tissue with substantially less invasivity than open-heart procedures. These systems are the primary focus of this dissertation, but my observations may generally inform applications across the stiffness spectrum. Fundamentally, there are classes of robotic systems which must safely interact with their environment and for which decreasing link stiffness is the preferred design solution. In these applications, manipulator compliance is essential, and all remaining design freedoms are spent improving accuracy, response speed, and force capability. This dissertation introduces Interleaved Continuum-Rigid Manipulation as a concept to enhance the capabilities and performance of continuum manipulator systems without compromising their essential compliance. Throughout, 'performance' is used to generally mean some combination of accuracy, precision, and response speed.

After describing the motivating application and surveying extant continuum manipulator design, modeling, and control efforts, I introduce interleaved continuum-rigid manipulation, explore some of the new design freedoms, and present results from two experimental interleaved systems.

1.1 Minimally-Invasive Surgery

Operating on and in the human body is a particularly challenging, and potentially rewarding, task, as few things are more valued than individual and familial health. Minimally invasive surgery (MIS) is a relatively recent surgical development that attempts to limit the incidental trauma to healthy tissue in treating various diseases. In short, "It is clear... that almost all advantages of MIS are in the interest of the patient, while surgeons are burdened with almost all disadvantages" [1]. This partition has created strong incentives to develop better tools for MIS and to deploy therapies more



Figure 1.1: a) Manual laparoscopic surgery. The abdomen is inflated to create a free working-space for instruments inserted through small incisions (keyholes). b) Intuitive Surgical's da Vinci system. The surgeon manipulates a haptic interface while viewing a 3D image from the laparoscope arm. The robot has four arms holding interchangeable tools over the draped patient. These arms are redundantly actuated to provide a remote center-of-rotation about each's keyhole and translation through it. c) Each tool consists of a long, rigid section with d) a dexterous end-effector for manipulating tissue.

widely, resulting in the creation of laparoscopic, endoscopic, and catheter (flexible tube) systems.

Laparoscopic surgeries are the most developed, due substantially to the ability to inflate the abdomen and create an open cavity in which to operate. Their tools are characterized by a long, rigid,
small diameter shaft through which the surgeon's hand movements are communicated to miniature implements located inside the patient. Through these tools the surgeon performs essentially
the same therapy as in invasive, open-to-the-sky surgery without the attendant breastbone-to-navel
healthy tissue damage. These procedures were the first to be roboticized commercially, most prominently in Intuitive Surgical's da Vinci, and since only the tools and workflow were changing many
other minimally-invasive surgeries have been adapted to the da Vinci [2, 3]. Robotic "Systems
aim for an enhancement of the skills of the surgeon to conduct medical procedures on the same
scale, therefore not changing the existing paradigms but rather supporting them" [3]. "Surgical
robots help enhance patient outcome by diminishing the surgeon's movement tremors, expanding
range of motion, and allowing for greater flexibility, accuracy, and smoother actions, especially in
small confined anatomical locations" [1]. The surgeon operates the da Vinci (figure 1.1) through

an interface resembling manual surgical tools, having forceps-like handles attached to 5 degree-of-freedom haptic controllers that can render tool forces, manipulator joint limits, and workspace keep-out zones, in a teleoperated fashion. These virtual tools may be easier to operate than their manual, laparoscopic counterparts because the system automates the laparoscopic tool's passage through the abdomen aperture. It has taken some time for these robotic adaptations to equal or surpass [2] the safety and efficacy of the manual minimally-invasive intervention, and the jury is still deliberating on the value-per-patient [4,5].

As laparoscopic surgeries require 1-2cm incisions through the abdominal wall for each of 3-4 instruments, they can be made less-invasive by inserting multiple instruments through a single access. Interventions are more difficult to perform because the tools are smaller, and hence weaker, and have reduced workspaces-both per tool and in coordinated manipulation-while the surgeon's view is essentially limited to a single, low-depth perspective. Endoscopes were first developed to provide visual inspection of internal anatomy, but have been expanded to feature smaller versions of common laparoscopic instruments. They consist of a 10-20mm diameter, flexible, articulated tube which features an optical path and multiple channels for various, sometimes interchangeable instruments. The endoscope's flexibility allows navigation of larger anatomical passages but "There are a number of challenges posed by the flexible endoscope..." including "...a lack of stability, triangulation of instruments for adequate tissue manipulation and inadequate force transmission to perform accurate microsurgery" ([6], see also [7]). As with laparoscopy, elements of these systems are being roboticized to reduce surgeon workload and strain. Whereas the da Vinci removes the surgeon from the patient to a dedicated workstation, many of today's commercial robotic endoscopes assist only the control of the outer endoscope tube and/or camera orientation, leaving the manipulators under the manual control of the bedside surgeon [6]. There are many interesting endoscope-like systems under research, though the most ambitious are systems from [8] and [9] that enter through a single portal and are then assembled into a larger, more capable systems inside the patient.

The laparoscopic and endoscopic approaches mentioned thus far are substantially stiffer than the tissue they manipulate and interact with, occasionally giving rise to incidental tissue damage. This can be mitigated by controlling the entire configuration of the instrument, as is the case in Intuitive Surgical's NeoGuide colonoscope [6] and [10, 11]'s highly articulated robotic probe, or avoided through compliant design, as in cardiac catheters. Cardiac catheters are an example of continuum manipulators, having a continuous backbone constructed from a variety of soft thermoplastics that, even when actuated, remain approximately as compliant as the surrounding tissue and are therefore safe by design. This inherent safety substantially constrains device design and thereby limits system performance; this trade is only worthwhile for devices operating in the vasculature, bronchi, or brain, where, in the event of a mechanical or other system failure, they may be safely removed from the patient through their entrance path without damaging the traversed tissue. The compliance has uses: a cardiac catheter safely traverses the vasculature by conforming to each patient's unique vascular structure and on reaching the heart deforms with the heartbeat and respiration, maintaining tissue contact when a more stiff device might either lose contact with or puncture the myocardium. The inherent safety results from the catheter's low bending stiffness and hence this measure drives catheter material selection and device design.

1.1.1 Performing Cardiac Ablation

The motive application for this work is the treatment of atrial fibrillation (AF) via cardiac catheter systems, but as we are concerned with increasing the capability of the tool the conclusions can be applied to other MIS operations in the heart and continuum manipulation more broadly. I will, though, use AF to determine some guiding workspace and performance requirements to root the

concept in reality. AF is characterized by the rapid, erratic beating of the heart due to errant heart muscle signaling and muscle activation (see [12–14] for longer discussions of the intervention and current status). The discordant beating results in decreased pumping efficiency, greater organ and system stress, and reduced quality of life for the 1-2% of the affected US population; the occurrence increases with age, afflicting 10% of those over 85 [12].

Some patients can be managed by pharmaceuticals, but their limited efficacy and side-effects coupled with age-related progression leads most patients to seek physical interruption of the errant activation [12,14]. As open heart surgery is highly invasive – making it both expensive to perform and recover from – researchers developed minimally-invasive cardiac catheter systems which require a single incision to access the heart. Figure 1.2 shows the basic elements of the system: a control handle external to the patient; a long, flexible, guiding sheath; and an articulating ablation catheter. The catheter is inserted through a small incision into the femoral vein, which it traverses into the right atrium, across the septum, and into the left atrium. Treatment typically involves ablating (burning) tissue around each of the pulmonary vein ostia through physical contact between a radiofrequency emitter and the tissue, whereby the resulting lesions are less conductive than the surrounding tissue. The resulting lesion must be both continuous around each of the ostia and transmural (spanning the wall thickness). As the energy delivered is a function of the probe contact area, force, and duration, treatment success and procedure time depend greatly on catheter controllability and surgeon skill. Care must also be taken to avoid perforating the heart wall which can result in potentially fatal complications such as cardiac tampenade or fistula [15, 16].

1.1.2 Operating In the Heart

Figure 1.3 shows a section of a CT-scanned left atrium along a plane between the superior pulmonary veins and approximately the center of the mitral valve. The section compares well with

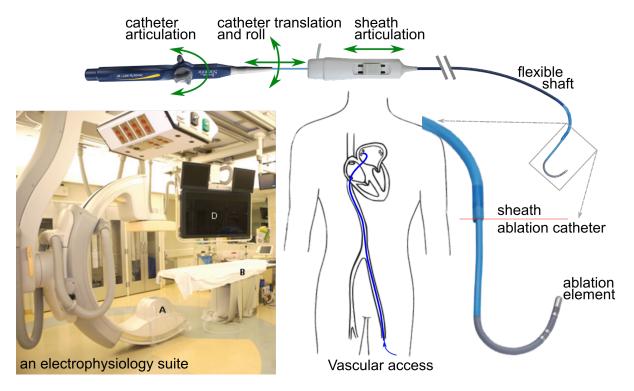


Figure 1.2: Elements of a cardiac catheter intervention. At left, an electrophysiology lab containing A) a CT scanner, B) patient bed, and D) a combination CT/ultrasound/electrocardiogram monitor from [17]. The catheter system (top) enters (anatomy, middle) the femoral vein and navigates the inferior vena cava into the right atrium, across the septum and into the left atrium. This typical manual system has an ablation catheter transiting a deflectable guide sheath into the heart. The guide sheath can be likened to a conduit that allows catheters to be exchanged without needing to re-navigate the vein. The exit from the guide sheath is typically just inside the left atrium from the septum. The ablation catheter articulates by the paddle-knob, translates and rotates with respect to the guiding sheath. The device images were drawn from http://www.bostonscientific.com/en-US/medical-specialties/electrophysiology/products.html

the measurements reported in [18, 19] which range between 27-40mm in diameter and 22-58cm³ in volume, as estimated from x-ray. To access the affected tissue cardiac catheters require three DOFs, commonly consisting of translation, rotation, and articulation, or translation and two-axis articulation. These are with respect to some stationary introducer sheath which forms a conduit between the outside world and the transseptal puncture. Translation affects the length exposed from the distal end of the sheath which commonly, as in Hansen Medical's Magellan Robotic Catheter, alters the articulating length of the catheter [20]. To navigate the vasculature, maximal catheter diameter is limited to approximately 5mm (15 fr).

The figure shows the silhouette of a 3mm catheter that can translate and rotate with respect to a stationary sheath and articulate away from the sheath axis via a tendon routed along the periphery of the articulating section. Having appreciated the catheter compliance in traversing the vasculature to reach the atrium, the compliance is again used to increase lesion uniformity and reduce the chance of perforation. The left pane in figure 1.3 shows two notional catheters ablating the upper wall of the left superior pulmonary vein, one slightly articulated to contact the tissue with the tip of the ablation electrode, and the lighter blue, more highly articulated contacting tissue with the side of the electrode. The ablation electrode is typically a cylinder; orienting the electrode to maximize contact along the side of the cylinder, as in the latter, lighter blue configuration, leads to more electrode-tissue contact and better energy transfer [21–23]. This larger contact area further distributes the contact force over a larger tissue area, decreasing the point load and reducing the likelihood of heart wall perforation [24].

1.1.3 Operating in the *Beating* Heart

Patients are typically conscious-sedated during a cardiac catheter procedure but they continue to respirate at approximately 0.5Hz and their heart continues to beat at a normal sinus rate of 1-2Hz [12,25–27]. [22,28–31] describe gating the therapy with the cardiac cycle so as to ablate when the tissue and catheter are in the same position as the previous cycle. Since the primary regions of interest in AF are the pulmonary vein ostia which typically lie across the atrium from the mitral valve, motion of the atrial roof is primarily towards and away from the catheter as it addresses the tissue. This tissue may reach speeds of 16cm/s as it cycles at 1-2Hz [32] while translating approximately 5mm per cycle [25]. Limitations on sensing and non-intuitive control interfaces make this a challenging control problem, is cardiac tissue interaction possible?

Today, cardiac catheter surgeons create lesions through a point-by-point procedure. The rec-

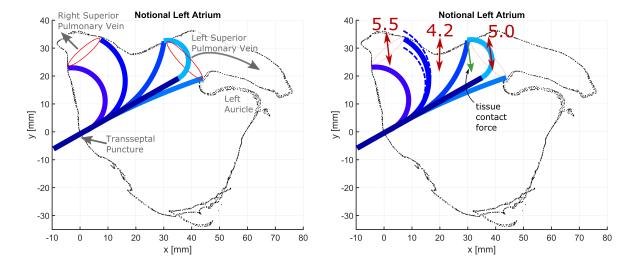


Figure 1.3: A planar section of a typical left atrium intersecting the superior pulmonary veins. Left. Atrial fibrillation is commonly treated through ablation of the pulmonary vein ostia (outlets from the atrium) which would resemble the red ellipses if projected into the plane. A three DOF (translation, rotation, articulation) ablation catheter is shown in four separate poses to indicate the configurations attained during therapy. Right. Average atrial wall movement per cycle in [mm] with respect to a stationary reference [25]; directions are approximate. The catheter's low bending stiffness, relative to its axial stiffness, permits bending due to the tissue contact force (blue-dashed). Between the two configurations addressing the upper left superior pulmonary vein, greater articulation shifts the catheter contact reaction force from primarily axial to bending stiffness, permitting more compliance with the tissue movement and reducing likelihood of perforation. Thanks to David Rutkowski and Dr. Alejandro Roldan for the heart section.

ommended force/time is 400g-s; with surgeons typically applying 20 gram-force, tissue contact lasts 20s [21,33]. The heart tissue is moving throughout this contact, so there often occurs relative motion and varying contact force between the probe and targeted tissue, leading to a more dispersed and uncontrolled therapy than intended [22,25,28,30,34]. Using contact force sensing catheters allows for the capture of some of this contact information, which is then used by surgeons to revisit sites with likely compromised lesions. The relative motion changes with the location of the targeted tissue; some approximate directions are indicated in figure 1.3, right. As the ablation must encircle the pulmonary veins, their differing locations and three-dimensional surfaces will lead to occasions where the tissue is sliding past, as well as moving towards and receding, from the catheter. These directions change with respiration, patient posture, and between patients. This motion explains

why many surgeons prefer to orient the ablation probe parallel to the tissue; more probe contact area delivers the energy more evenly, while catheter relaxation allows it to maintain contact even under modest tissue movement [29], as in the dashed poses of figure 1.3, right.

Exploiting the catheter compliance through parallel alignment, gating the ablation with heart-beat, and using contact force-sensing catheters can improve one-year success rates from 20-55% to 85% [33], but fundamentally limited device controllability and therapy information lead to poorer outcomes and longer procedure times [35]. Improving these and enabling new, more dexterous procedures requires understanding what limits today's ablation catheters and (re)considering continuum manipulator designs in general.

1.2 Continuum Manipulators

Discrete revolute and prismatic joints are found in all manner of industrial and academic robotic manipulators because they are easily fabricated and modeled. In contrast, continuum manipulators are underactuated by design so that both intentional and inadvertent contacts with the environment permit some deformation to the contacting body. In considering the roles of continuum manipulators and discrete joints, [36] observes that in the animal kingdom, rigid joints are found in applications needing precision with flexible, continuum segments occurring in inherently unstructured and varied tasks. The Clemson OctArm exploits this ability to conform to the environment in showing how such an elephant trunk-inspired continuum manipulator readily adapts to grasp unknown objects. Those same objects may each require a different, custom robotic end-effector while the continuum backbone and actuation naturally redistribute contact forces in response to the object's extent [37].

Some rigidly-constructed manipulators attain compliant properties through various joint torque control architectures, variable stiffness actuators, and compliant coverings (skins); Townsend and

Salisbury's Whole Arm Manipulation and ReThink Robotic's Baxter being well-known examples [38,39]. These efforts are generally not equivalent to an underactuated continuum manipulator because their motion is fundamentally constrained by their fixed degrees of freedom while the continuum manipulator has infinitely-many DOF; a manipulator cannot admit motion in directions outside its workspace. This is particularly true in the case of incidental contact along the manipulator, where highly redundant arms, viewed at the end effector, have only a few complaint DOF approaching the shoulder and therefore minimal ability to conform to contact along their bodies. Further, their apparent compliance disappears on power or system failure.

In interventional cardiology, continued therapy maturation and risk reduction may lead to acceptance of platforms that appear too rigid today. This is to say that compliance is valued when operating in the unknown, but greater knowledge of the patient's vasculature, attained through improved sensing techniques, may reduce these unknowns, lessening the associated design constraints and permitting higher-performance designs. My efforts are encouraged by recognizing that: 1) efficacy precedes safety, and 2) that from an engineering point of view cardiac device 'safety' is best, and finally, defined by those systems that have cleared medical society and Food and Drug Administration review.

Continuum manipulators differ primarily in backbone design and actuation mechanism with many interesting permutations. The first continuum manipulator by Hirose consisted of a series of co-actuated rigid joints which together formed a continuum structure [40]. Manipulators of this quasi-continuum design enjoy expansive design freedom and in turn can boast of minimal nonlinear joint effects, but are quickly challenged in providing small radius of curvature while maintaining force capability (large angular displacements lead to lots of joint-clearance space and minimal load-bearing area). One novel development is the Highly Articulated Redundant Probe from [10] (see figure 1.4) whose backbone consists of two concentric cylinders, each composed of segments which

can be locked in a particular actuated position while permitting passage of unlocked members, with the result that the manipulator can follow the leader (along the path of its tip). Despite its novelty, this concept creates a rigid structure that can avoid static obstacles during insertion but thereafter is rigid to disturbances, such as patient movement. The pre-curved super-elastic concentric tube

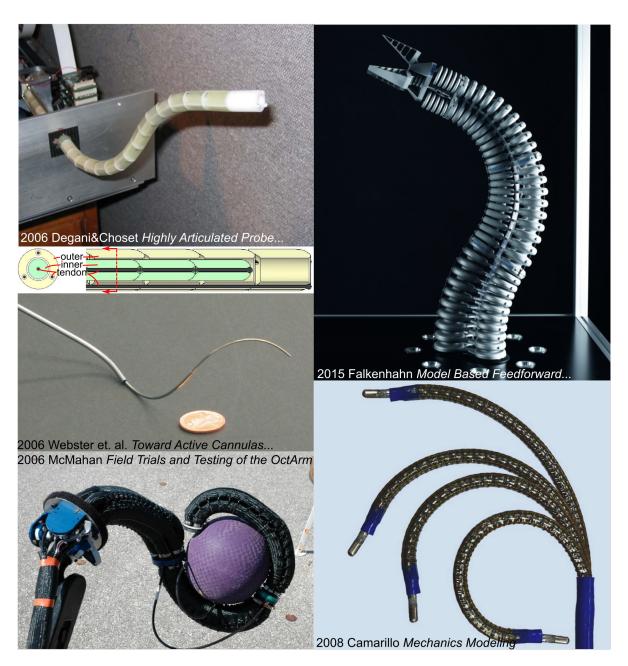


Figure 1.4: Left column from top: [10]'s Highly Articulated Probe; [41]'s active cannulas; Clemson's OctArm pictured in [37]; [42] uses Festo, Ag's Bionic Handling Assistant; Hansen Medical's Artisan catheter as pictured in [43]

manipulators introduced separately by [41] and [44] (figure 1.4) are currently the simplest and most easily scaled continuum manipulators, making them attractive for a variety of small-aperture, minimally invasive surgical applications. Ongoing research is maturing the platform and some, preliminary in-vivo tests have occurred, but they are fundamentally limited in a few regards: the tip force is bounded by the bending stiffness of the concentric tubes, substantial friction may develop between tubes which is very difficult to model and predict, outside disturbances (forces/moments and thermal changes) may release stored elastic energy, and difficulty in sensing 'joint' positions as these result from the difficult-to-model interaction between tubes [45]. These limits may eventually be tamed and I look forward to seeing progress, especially in light of the typical advance in design sophistication between early research catheters and commercial systems.

Clemson's elephant-trunk OctArm and Festo's Bionic Handling Assistant [37, 42] (figure 1.4) consist of pneumatic bellows grouped in parallel to form three-DOF segments (two-axis bending and extension) and serially-arranged segments to form the complete manipulator. Separately, the compliant pressure vessels are the actuators, and when grouped form the compliant, continuum backbone. This concept has been realized at smaller scales pneumatically in [46] and hydraulically by [47]. The hydraulic approach is interesting for biological applications where biocompatible working fluids would be essential in case of malfunction (with air leading to very difficult-to-dislodge and easily fatal air embolisms). These approaches appear to be limited by difficulties in fabrication (especially at small scales), the compressibility of the working fluid and/or delivery lines, and the immense difficulty in modeling the behavior of coupled compliant pressure vessels, both statically and in any dynamic sense.

This work is chiefly concerned with one common class of continuum manipulators, those that are tendon-actuated, like the cardiac ablation catheter pictured in figure 1.4. Figure 1.5 illustrates a stout device in cross-section and articulation. The primary structure is a machined/molded/extruded

elastic backbone that is long with respect to its length such that it has low bending stiffness while resisting axial compression and torsion. A central lumen provides a toolpath while actuation tendons are routed along the periphery through evenly-spaced actuation lumens. Terminating the tendons at the distal end and applying different tensions at the manipulator's proximal end induces bending. One, three, or four tendons may be used per articulating segment to effect bending.

These tendon-actuated manipulators have seen substantial commercial success within minimally invasive surgical systems which safely traverse the vasculature to deploy stents, clear blockages, sample tissue, image vascular and cardiac structures, measure electrical propagation across heart muscle, and to ablate heart tissue from the interior. These interventions are possible with today's manipulators but all would be enhanced by increased dynamic performance. Increased dynamic performance may enable more complex interventions (say mitral valve repair) that today require open-heart or laparoscopic surgeries.

Despite the efficacy of MIS interventions and the commercial success of several companies, fundamental challenges remain; efforts to improve tendon-actuated continuum manipulator performance have focused on better understanding and modeling system kinematics and dynamics, improving open-loop performance via design, and closed-loop control, with each avenue benefiting from improvements in the former.

1.2.1 The Physics of Tendon-Actuation

Tendon-actuated catheters are highly hysteretic due to the sliding interface between the control tendon and the guiding lumen. Consider figure 1.5. A tendon displacement applied at the base of the catheter causes the tendon to slide against the lumen wall. Since the tendon is fixed to the tip, the amount of relative sliding – between initially co-located points on the tendon and wall – decreases to zero from the base to the tip. If friction is neglected, the forces between the

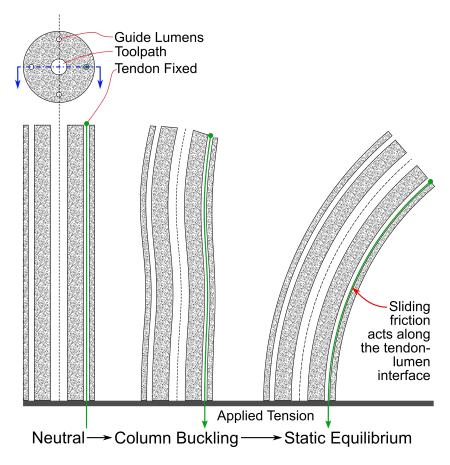


Figure 1.5: Applying a tension to tendon-actuated continuum structure. From left. A cross-section of a continuum structure having four equidistant, peripheral tendon guide lumens and a central toolpath. The relative diameters vary substantially across systems, this illustration is stout for spacing. One tendon is shown in green, fixed at the tip. Applying a tension to this tendon, middle, immediately induces buckling in the structure. Eventually the structure will reach a static equilibrium, right. In the absence of friction the body assumes a constant curvature.

tendon and wall will eventually achieve static equilibrium at some articulation angle, as defined in Appendix A. Sliding friction complicates this static equilibrium. Consider Coulomb friction acting between the guide and tendon where, as is common, the coefficient of static friction is greater than the dynamic coefficient. This inequality gives rise to the stick/slip phenomenon, in which larger forces are required to set an object in motion than to continue that motion. Applied to the tendon/guide sliding, the relative motion will often cease at some point before the tip. Distal to this point there is no relative movement between the tendon and guiding lumen, resulting in no further articulation. Successive tension inputs then lead to a variable stiction point and a manipulator

with variable curvature. Ignoring axial compression, the tendon displacement will lead to equal equilibrium articulation angles between the friction and friction-less cases, but the tip positions will differ according the singular or variable structure curvature. This discrepancy is the origin of the catheter's hysteresis, as it implies that the previous loading influences the static equilibria. External disturbances can upset the static equilibrium, thwarting control efforts.

1.2.2 Accurate Modeling is Difficult

The literature has attempted to model the elements of interaction between the elastic backbone and articulating tendon. [43] was the first to closely consider tendon actuation of a continuum manipulator; they assumed constant curvature which they justified by considering the internal force balance on an element – to the exclusion of friction. This model was validated via a catheter produced by Hansen Medical, Inc. which exhibited low hysteresis. Many efforts have made this same frictionless assumption and used some other arc parameterization to 'model' the catheter, but arrive at essentially identical results with no increase in predictive power. One useful result is [48]'s virtual power approach, which allowed the inclusion of a viscous friction. Their experimental platform had a continuous backbone but discrete tendon guides; these concentrate the tendon friction to small locations and, at least presently, limit the result's applicability. A recent effort [49] explicitly modeled Coulomb friction between the tendon and guide lumen. The resulting, single DOF, closed-form, quasi-static mechanics model demonstrates better agreement with observed behavior but does not consider the transition from dynamic friction to static (in the case of the Coulomb model), tendon stretch, or the effect of using linear-elasticity for what is clearly an elastic body.

[50] considered friction in a dynamic, lumped-parameter numerical model of a four tendon catheter whose articulating length could be changed by extension from a rigid proximal sheath. The numerical approach allowed consideration of a Dahl friction model and tendon stretch, resulting in very good prediction of both static and dynamic catheter configurations. The price of this accuracy is lengthy computation time which would be difficult to achieve in realtime, and particularly so under model parameter uncertainty. The omission of higher-order dynamic terms (Coriolis and centripetal accelerations) would additionally limit accuracy at higher control bandwidths.

The literature has not yet considered: flattening of the guide lumen with articulation, which can lead to pinching of the tendon; the formation and effects of debris from the sliding friction; or the temperature dependence of any of these mechanisms or interactions, especially at increased control bandwidths. The modeling challenges justifies spending substantial effort reducing or designing-out these nonlinearities (as was done in [43]).

1.2.3 Design Yields Only Modest Improvements

Many attempts have been made to avoid or reduce hysteresis through design, often by fundamentally deviating from tendon-actuated designs. [43] and [51] record impressive performance with a catheter from Hansen Medical. Their careful design and construction appears to have substantially avoided the expected hysteresis, while controlling tendon tension rather than pull distance reduces the effect of tendon compliance proximal to the catheter.

Across the literature, uncertainty about the tendon actuation mechanism limits rigorous documentation of the manipulator design, as the relevant physical, material, and process parameters have not been identified. Professional catheter designers are presumably incented against public discussion, while independent characterization of new or used commercial catheters is limited by their unit cost, or their generally fast wear-rate and post-procedure biohazards. I can observe that design can incrementally improve tendon-actuated catheters but has not yet delivered the desired order-of-magnitude performance improvements that, for instance, would enable teleoperation at

approximately 20Hz.

The alternatives to tendon-actuation described above are unexplored in general and in cardiac applications in particular, but it is entirely likely that they will encounter challenges of similar magnitude as the tendon-actuated. Though efficacy precedes safety it appears further work is required to de-risk such an alternative device.

1.2.4 Advanced Control Can Perhaps Double Performance

The modest improvements achieved through improved modeling and design suggest similarly limited performance increases from advanced control architectures. Note that the present interest is in architectures that increase the task performance (accuracy, precision, and/or response speed); questions related to movement or procedure planning and surgeon interfacing lie beyond this interest.

Closest to application are efforts that model system nonlinearities generally, that is considering only a resulting behavior (like a hysteresis loop) rather than the particulars of tendon actuation. [31] develops a system that manipulates a commercial, manual catheter and compensates for the patient's heartbeat by estimating the catheter articulation delay and articulation error (both resulting from catheter and tendon compliance). Their system is driven by a sinusoid at the measured heart rate and modified according to these estimated errors. While desiring to operate at 1Hz, their commercial catheter(s) failed at 0.65Hz, limiting their excitation of device dynamics. This effort may be thought of as a form of model-based control, where the model is the bias (articulation delay) and slope (articulation error) between the catheter's articulation knob input and the tip articulation. This effort did not include any physics-based catheter model, limiting its applicability beyond stably-periodic driving motions and targets. [52] developed a feedforward guidewire position controller that compensates for heart motion. Their system is oriented along the line of motion,

allowing the use of a single DOF guidewire translating with respect to an statically-articulated sheath. The modeled sliding friction and backlash were then used in a feedforward controller to compensate the desired command against these losses.

Using one of the previously-described modeling approaches allows control development within the model's limitations. For instance, assuming constant curvature in a multi-segment continuum manipulator allows [53] to develop a sliding mode controller and show Nyquist stability, but dynamic performance and comparison to their numerical simulations are omitted. [42] develops a model-based feedforward controller for Festo's Bionic Handling Assistant which uses a dynamic model of the coupled bellows to decouple their interactions, linearizing section length control and demonstrating good open-loop tracking. As this is a recent effort it will be interesting to see whether their results extend across the workspace (beyond the demonstrated small actuations) and a fuller exploration of the dynamic performance. It may be that the Bionic Handling Assistant's design and fabrication are sufficiently advanced to mitigate nonlinear effects and allow accurate and precise performance. Though demonstrating a continuum manipulator controller, this result is pneumatic and of limited immediate applicability.

The described challenges in dynamic modeling have motivated some researchers to consider controller designs that, in some way, learn the manipulator nonlinearities. [54] included a 2-layer, 15 node neural network in the feedforward path of a controller for the Clemson OctArm; despite apparently positive results this controller is not used in later works [55]. [56] develops two types of neural networks and separately trains these on the inverse and forward kinematics of Festo's 2-section Compact Bionic Handling Assistant, showing good performance. [57] assumes even less about the manipulator by learning the forward kinematics of Festo's Bionic Handling Assistant online through goal-babbling, which reduces the task error from 30cm to approximately 2cm over a 5-hour training period. [58] apply a neural network to learn the static relationship between tendon

force and tip position in a conical, tentacle-like tendon-actuated manipulator segment. On a large experimental dataset of tip locations and tendon forces, across tip locations they demonstrate that the neural network controller more accurately predicts the tendon force than an Euler-Bernoulli beam model, especially under parameter errors. These results illustrate that a sufficiently large and representative training set can teach control of a continuum manipulator. These methods are fundamentally limited in their inability to consider the actuation history²; some take the instantaneous joint velocity as a model input, but this is incapable of describing phenomena like hysteresis.

The preceding methods essentially attempted to improve the forward, open-loop path by better modeling or learning of the manipulator's characteristics off-line; these are complemented by efforts utilizing feedback in a closed-loop controller. Due to the described hysteresis and underactuated nature, tendon displacements measured proximal to the catheter are insufficient to determine manipulator configuration, and tip location in particular. A task sensor is required; incorporating it into a real-time controller leads to a noncollocated system where the position- or force-controlled drives are separated from the sensor by the backbone compliance, friction hysteresis, and other nonlinearities. Noncollocated systems are characterized by a pole near the imaginary axis; they are difficult to control because the controller must ensure that the pole's locus does not cross the axis into oscillatory motion – ensured by conservative gains and reducing parameter errors – and provide damping if it does [60,61]. The challenge of accurately modeling continuum manipulators has limited analytical explorations and explanations of fundamental system characteristics like mode shapes and noncollocated sensing.

The limited nature of many experiments challenges control architecture comparisons, but to my knowledge the greatest performance improvement arose from modal control [62]. In this, two pose sensors were used to measure the first and second vibrational modes, damping the first mode

²Repetitive or iterative learning control have apparently not been applied to the control of continuum manipulators, but some less-periodic variant may be able to learn the temporal nonlinearities [59].

through an approximate dynamic model and observer. This architecture increased the closed loop bandwidth from 1.1Hz to 2.6Hz for, again, an incremental performance improvement.

Physically-motivated modeling, careful design and fabrication, and advanced control have all resulted in only incremental improvements in response speed and accuracy. These many, good efforts suggest that the available design freedoms are too constraining.

2. Interleaved Continuum-Rigid Manipulation

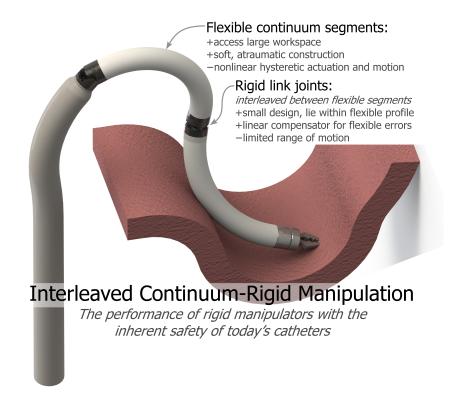


Figure 2.1: A notional interleaved continuum-rigid manipulator addressing tissue.

Reflection on the many challenges of tendon-actuated continuum manipulators motivated a re-evaluation of the fundamental manipulator design. In the absence of theoretical advances, the traditional tendon-actuated architecture is too constraining; only by changing the architecture can we hope to realize better performance, with the design burden to ensure that we maintain the safety profile of current cardiac devices. Recognizing the limitations of continuum manipulators in the literature, while still desiring to perform more advanced tasks, Prof. Zinn and I introduced Interleaved Continuum-Rigid Manipulation which seeks to combine the inherent safety and large workspace access of flexible catheter segments with the precision and speed of traditional revolute and prismatic (rigid link) joints. As will be developed further, these rigid joints lie within the profile of the device and are small relative to the length of flexible segments. Several attributes, summarized in figure 2.1, are immediately apparent: the flexible continuum segments access a large workspace while compliantly interacting with tissue, and the rigid joints are easily-modeled and -controlled with the potential to compensate for flexible segment nonlinearities. Considering any pairing of rigid and flexible joints from an interleaved manipulator, they will generally not be completely redundant actuators, instead having some overlap in their manipulability but differing in their ranges of motion, precision, speed, etc. We expect that the rigid joints will generally be of greater precision, have less actuation range, and possibly be faster than the flexible segment actuators. This concept differs from laparoscopic and endoscopic systems in the flexible segment's compliance and actuation, but can draw inspiration from their rigid joint and end-effector designs.

2.1 Possibilities of Interleaved Continuum Rigid Manipulation

The additional redundancy provided by the rigid joints can be utilized in a number of ways

- the rigid joints might compensate for flexible segment errors, as shown in chapter 3;
- the added redundancy may expand the dexterous workspace and permit more complex tasks, as in chapter 4;
- the controller can simply stop using the flexible segments, preferring the rigid joints for AC and leaving the flexible segments for gross positioning (DC), as in macro/micro approaches [63];

- the rigid joints may enable motions that minimize flexible segment excitation by exploiting the reaction null space [64];
- the extra DOF and precision could be used to manage contact via impedance shaping and control [65, 66];
- the rigid joints might be used for motion compensation as in [31, 52, 67];

The burden is to design these rigid joints so that they lie within the profile of the flexible segments and have a minimal effect on system safety. A number of designs are apparent, but their attributes are of later concern as we must first show that the design freedoms added by Interleaved Continuum-Rigid Manipulation can actually be exploited to increase performance. We do this through two systems: the first, single DOF testbed demonstrates possible performance gains while the second, 5 DOF prototype showcases rigid joints integrated into the manipulator profile and the increased dexterous workspace.

3. An Interleaved Rigid Joint and Continuum Segment

The objective of this first testbed is to determine whether a simple, interleaved manipulator can increase performance over an equivalent continuum manipulator. Approaching this first testbed we realized that a simple, single degree-of-freedom (DOF) system would illumine the many mechanical and control design decisions necessary in realizing a clinical interleaved manipulator. Considering only a single DOF greatly simplified the design, construction, analysis, and interpretation, leading to a quick validation of the potential of interleaved manipulation.

We presented initial results of this single DOF testbed in [68], but that experiment was sensorlimited to being near the zero articulation. After describing the testbed and redundant controller we present results near zero articulation and at 220° with an unexpected observation of the effect of rigid joint saturation.

3.1 Overview of the Single DOF Testbed

The single DOF manipulator consists of a single articulating flexible segment and a proximal revolute joint arranged redundantly. The flexible-segment is a molded 6.4mm diameter urethane body with four tendon guide lumens and a central 3.2mm instrument lumen. This catheter is articulated using a pair of opposing control tendons (Honeywell Spectra fiber, 0.23mm diameter) anchored in the tip and pulled by a pair of 90W DC motors (#273754 Maxon Motor AG) with 5mm diameter pulleys. The tendon position controller (motor encoder) has a closed-loop bandwidth of

approximately 35Hz. Actuation of the tendons causes the flexible segment to articulate within a vertical plane.

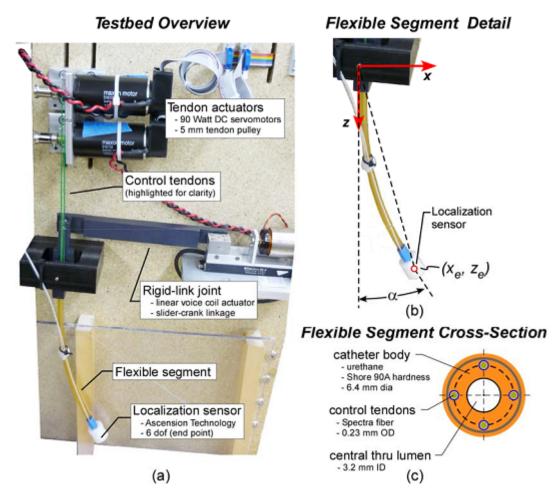


Figure 3.1: Single DOF testbed. a) the flexible segment is actuated by two motors and the rigid joint by a linkage connected to a voice coil motor in a plane substantially along the width of the page. b) testbed coordinate system. c) flexible segment cross-section.

The rigid joint motion provides rotation about a pivot axis located at the base of the flexible segment which is perpendicular to the flexible segment actuation plane. The flexible segment control tendons intersect the rotation axis of the rigid joint to eliminate coupling between the flexible and rigid motions. The rigid joint rotation is accomplished through a slider-crank mechanism actuated by a voice-coil motor (BEI Kimco Magnetics) with approximately 6 mm of travel which results in

±16° of manipulator articulation (when the flexible segment is unarticulated). Catheter tip motion is acquired with an Ascension trakStar magnetic pose sensor which provides a globally-referenced measurement of the catheter's tip pose at approximately 300Hz. The controller is implemented using Matlab xPC 2009a (Mathworks, U.S.A) on an Intel Pentium 4 at 3.0GHz.

3.2 A Single-Input, Multiple-Output Controller

One of the central challenges of the interleaved approach is formulating an effective control strategy. There have been many formal methods developed for multiple-input-multiple-output control system design including μ -synthesis, H_{∞} , and, more recently, design approaches developed for dual-input-single-output systems like the PQ approach [69]. In this testbed we chose to pair a typical tendon-actuated flexible segment with a high-precision and -speed, limited range rigid joint, as depicted in figure 3.1. These joint differences suggest a parallel control structure which explicitly partitions the task error signal, Δx , into high and low frequency components. In the context of the overall control structure proposed given in figure 3.2, the flexible segment and rigid-link task-space controllers $(D_f(s))$ and $D_r(s)$ perform this partitioning function while helping to shape the actuator closed-loop dynamics. Additionally, a task-space loop compensation block, $D_l(s)$ is included to compensate the additional dynamics that result from the parallel path summation of the rigid and flexible segment control signals. The high/low frequency partitioning is motivated by the desire to limit the motion of the limited-stroke rigid joints while correcting for motion errors that result from the slower-responding flexible segments.

The flexible segment control includes a feedforward inverse kinematics block which converts the desired task space configuration to flexible segment joint commands (segment curvatures). As shown in figure 3.2, the task space control signal is transformed to joint space motion commands via the flexible segment and rigid joint Jacobians, J_f and J_r , under the assumption that the task

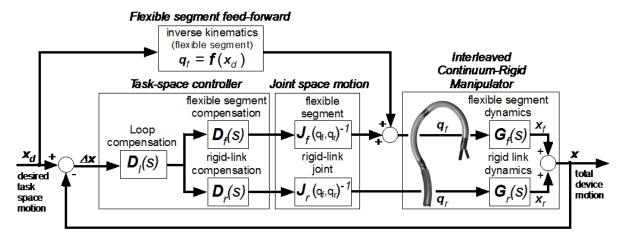


Figure 3.2: Overview of a frequency-partitioning, single-input multiple-output controller based on [69].

space error is small. It should also be noted that both the flexible segment Jacobian J_f and rigid joint Jacobian J_r are functions of the manipulator's configuration. As a result, knowledge of the configuration is required – either through estimation or direct measurement. While the rigid joint position will likely closely track the desired motion, the flexible segment is expected to have significant error and thus direct measurement of its motion is required to properly form the Jacobian for both the flexible segment and rigid joint.

While the specific structure of the compensation blocks $(D_f(s), D_r(s), and D_l(s))$ in figure 3.2) vary depending on the system dynamics and the performance goals, there are general considerations for both the flexible segment and rigid joint control that have bearing on the compensator design. In general, robotic catheter systems regulate control tendon motion with a high-gain position controller that acts on the control tendon actuator positions. This is done to improve disturbance rejection by increasing the static stiffness at the control tendon output and to improve stability margins. In addition, the low torque density of electromagnetic actuators (used almost exclusively in the type of cardiac interventional catheters under consideration here) generally requires the use of a gear reducer. The resulting increase in reflected inertia and friction amplification makes tension control

difficult to implement in a robust manner. As such, it is assumed that the local joint controllers (i.e. tendon position and rigid joint position) are designed to have significantly faster closed-loop dynamics than those of the overall closed-loop interleaved manipulator. When considering the design of the compensation we can assume that the control inputs to the interleaved manipulator are given in terms of joint displacements (q_f and q_r).

To compensate for steady state flexible segment motion errors, integral control (and variants thereof such as lag compensation) has been successfully applied in the reduction of catheter kinematic errors, e.g. [70]. In the context of the interleaved control structure proposed here, the flexible segment compensation block $D_f(s)$ can assume a similar integral-like control structure (K_i/s) . As described in [69], the ratio of joint control compensators $(D_f/r(s) = D_f(s)/D_r(s))$ can be used to examine the frequency partitioning characteristics of the chosen compensator design. Assuming that the magnitude of $D_f/r(s)$ decreases with increasing frequency, the crossover frequency ω_f/r of $D_f/r(s)$ is the point where the magnitude of $D_f(s)$ and $D_r(s)$ are equal and thus the frequency at which the low and high frequency partitioning of the control input occurs. In addition, as described in [69], the phase of $D_f/r(s)$ at the crossover is a representation of the constructive interference between the flexible segment and rigid joint control action. In this paper, where $D_f(s)$ was given an integral-like control structure, a suitable choice for the rigid joint controller could be unity gain $(D_r(s) = 1)$.

As described earlier, the purpose of the overall loop compensation, $D_l(s)$, is to help shape the open-loop system frequency response such that the closed-loop stability margins and performance are satisfactory. In the simplified case where $D_f(s)$ is a simple integral controller and $D_r(s)$ is a unity gain, and when we assume that the rigid-link and flexible segment system plant transfer functions have constant gain and no phase distortion at low frequencies (i.e. $G_f(s) = G_r(s) = 1$), a reasonable choice for $D_l(s)$ could be a simple integral compensator – where the gain is adjusted to

obtain a stable parallel system, presumably with a closed-loop bandwidth which is greater than the crossover frequency of $D_f/r(s)$ (i.e. the partitioning frequency). The control partitioning can be seen by examining the frequency response of the closed-loop system of the simplified case described above. As shown in figure 3.3, the control signals associated with the flexible segment and rigid-link response are partitioned at the crossover frequency $\omega_{f/r}$, the point at which their respective magnitudes are equal. In addition, the high-frequency content from the rigid-link response results in a combined system closed-loop bandwidth ω_{CL} that is well above the crossover frequency $\omega_{f/r}$.

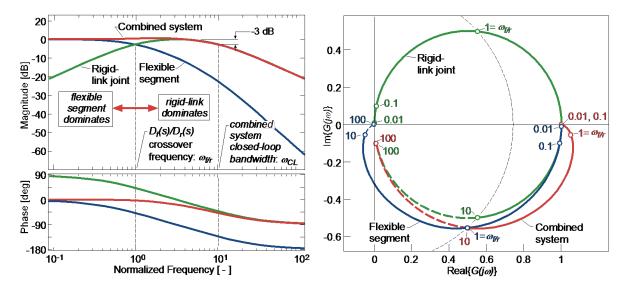


Figure 3.3: Closed-loop frequency response of a one DOF simplified interleaved system. Left. Magnitude and phase as a function of frequency where the contribution of the flexible segment actuation, rigid-link actuation, and the combination of the two are shown; Right. Frequency response shown as real and imaginary part of response.

3.3 Single DOF Performance

To evaluate performance the responses of the interleaved system and a system consisting of only a flexible-segment were compared. The flexible segment-only system was formed by preventing rigid joint motion while using the same flexible segment as the interleaved system. To provide a clear comparison of the control behavior of each system, the flexible segment feedforward term, shown in figure 3.2, was set equal to zero in both the interleaved and flexible-segment only implementations.

In the context of the general PQ design described above, the frequency partitioning between the flexible and rigid joints was adjusted via the flexible segment and rigid-link compensation blocks, $D_f(s)$ and $D_r(s)$, respectively. $D_r(s)$ was set equal to 1.0 while $D_f(s)$ was assigned a simple integral control structure (i.e. $D_f(s) = K/s$). The integral gain of $D_f(s)$ was adjusted to have a crossover frequency of 0.05Hz, below which the flexible segment primarily acts on the task error and above which the rigid-link primarily acts on the task error. The choice of this partitioning frequency is essentially limited by the stability of the flexible segment, which, in addition to the catheter and actuator characteristics, is a function of tendon compliance and friction developed throughout the catheter body. The partitioning frequency must therefore be conservative to allow for variable friction and compliance during actuation; the experimental partitioning frequency of 0.05Hz is the result of this balance. The task-space loop compensation block, $D_l(s)$, is given as an integral controller to eliminate steady-state errors due to internal device friction and kinematic modeling errors. The integral gain of the task-space compensation block was adjusted upward until signs of instability were observed, resulting in an overall system open-loop compensated crossover frequency of approximately 0.6Hz. This crossover frequency is chosen to be below the catheter's first mode of 1.8Hz and maintain sufficient stability margins.

The performance of the two systems was evaluated with a simple step response with the task defined by the catheter tip articulation noted as α in figure 3.1(b). In the first experiment, the manipulator was positioned approximately in the center of its workspace (vertical) and biased slightly positively (11°) to eliminate any effects of control tendon slack. A small effective articulation motion step input command (17°) was applied and the position control performance was measured. Figure 3.4 shows the results of this first, low articulation experiment. The left pane of figure 3.4

demonstrates that the response of the interleaved system is approximately four times faster than the flexible-segment only closed loop system – due primarily to the ability of the rigid joint to effect changes in articulation faster than the more compliant flexible-segment control tendons allow.

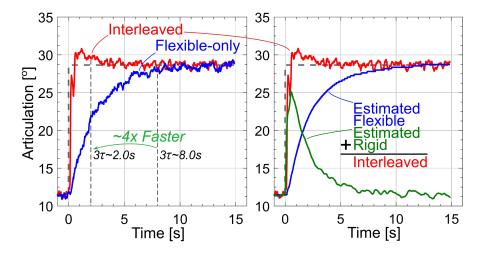


Figure 3.4: Experimental response of the interleaved experimental testbed to a commanded step end-point articulation of 17 degrees from an initial articulation of 11.5°. Left. Response of interleaved system (red) and response of flexible-segment only manipulator (blue); Right. Total response of the interleaved system (red, same as on the left), estimated flexible segment contribution (blue), and estimated rigid-link contribution (green).

To gain a better understanding of the interleaved response, the right pane of figure 3.4 shows a projection of actuator encoder data into the task space to estimate the contribution of each actuator to the total tip articulation. Note that this projection does not include catheter dynamics and therefore has some inherent error. Nevertheless, figure 3.4 clearly shows the effect of the frequency partitioning between the slower flexible segment and faster rigid joint controllers. The response of the rigid joint actuator is almost immediate, reacting to the high-frequency content contained in the step input command. As the slower flexible segment actuator motion increases, the rigid link joint actuator motion decreases in magnitude, returning it to the center of its actuation range. The summation of the two results in a faster response as compared to the flexible segment alone.

In a second experiment, the same step command was applied (17°) with the manipulator initially positioned with an articulation of 218° (i.e. in a U-shaped initial configuration). As seen in figure 3.5, left, the response of the interleaved system suffers from substantial overshoot due primarily to the saturation of the rigid joint. As before, figure 3.5, right, projects the actuator encoder data into the task space to estimate the contribution of each actuator to the total tip articulation. The obvious saturation of the rigid link joint leads the flexible segment to overshoot the command with a corresponding effect on the overall system response.

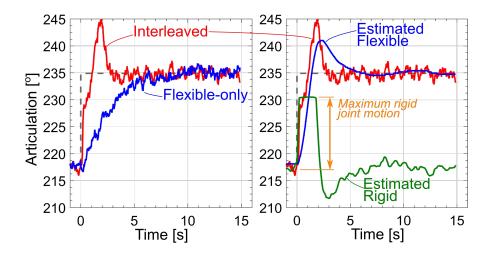


Figure 3.5: Experimental response of the interleaved experimental testbed starting from a large initial end-point articulation. Response of the system to a commanded step end-point articulation of 17° from an initial end-point articulation of 218°. Left – Response of interleaved system (red) and response of flexible-segment only manipulator (blue). Right – response of interleaved system (red), estimated flexible segment contribution (blue), and estimated rigid-link contribution (green). The prominent overshoot is due to plant dynamics that change with configuration: as in figure A.1 applying the same tendon displacement increment in low- and highly-articulated configurations gives greater articulation increment at higher articulations. Increasing articulation also decreases the rigid joint-to-tip distance, decreasing the rigid joint influence. These opposed characteristics lead to greater flexible segment articulation, and less rigid joint articulation, than in the low articulation of figure 3.4 with the resulting rigid joint saturation and tip overshoot.

To understand the source of the response overshoot and explore the effects of rigid joint saturation on overall system stability, a series of similar step response experiments were carried out where limitations on the rigid joint range of motion were emulated by restricting commands to the rigid-link position controller, as in figure 3.6. The experimental step response for decreasing rigid joint range of motion is shown in figure 3.7, indicating that as the rigid joint range of motion is reduced, the response becomes more oscillatory. In the extreme case, where the rigid joint motion is set equal to zero, the system becomes unstable.

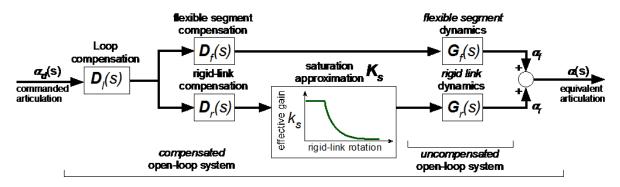


Figure 3.6: Rigid-link saturation is approximated with a variable gain, K_s as described by [71]

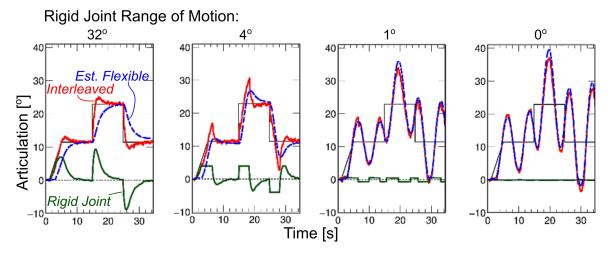


Figure 3.7: Beginning at zero initial articulation, a two-staircase pattern shows the effect of rigid joint saturation. From left: full rigid joint range-of-motion demonstrates performance similar to that in figure 3.4; decreasing to 4° shows lengthier convergence and large overshoot, reminiscent of figure 3.5; 1° is marginally stable; and 0° rigid joint actuation range is unstable.

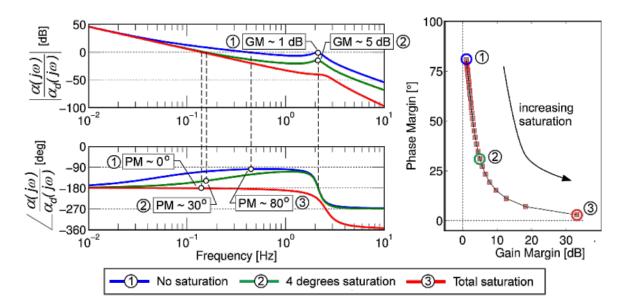


Figure 3.8: Interleaved system stability margins as a function of joint saturation levels. The degradation of the overall system stability (as evidenced by the reduction in phase margin) is shown in the bode plot (left) and the supporting phase-gain plot (right).

The reduction in system stability can be seen in system's stability margins as a function of joint saturation levels. To do this, we constructed a model of the compensated open-loop system which is consistent with experimental frequency response measurements and examine in figure 3.3 the effects of rigid joint saturation by varying the joint gain (as an approximation to saturation [71]). The overall system loop stability as a function of rigid joint saturation can be determined by evaluating the phase and gain margin of the compensated open-loop system as the saturation gain, K_s , is varied from 1.0 (no saturation) to 0.0 (total saturation).

As shown in figure 3.8, the nominal compensated open loop interleaved system (with no saturation) has closed-loop phase and gain margins of approximately 80° and 0.6dB, respectively. The nominal system stability margins reflect the design objective where by the highest possible closed-loop bandwidth was sought. Here the closed loop bandwidth is restricted by the first flexible mode of the flexible segment. The lightly damped mode is responsible for the low gain margin while the integral control action results in a robust phase margin in excess of 80°. As the rigid

link joint saturation increases (approximated by a reduction in K_s), the resulting system phase margin decreases, resulting in a closed loop response which is lightly damped. At the extreme, where the saturation is total (no rigid joint motion), the system's phase margin approaches zero, resulting in a highly oscillatory response. In practice, unmodeled system dynamics would lead to an unstable system. As such, the effects of saturation should be carefully considered in the design of any interleaved system.

3.4 Single DOF Discussion

This single DOF testbed and frequency-partitioned controller demonstrate the fundamental properties of an interleaved manipulator: joint characteristics (accuracy, range, and speed) can be combined into a single manipulator, overall system performance may be increased, and a controller aware of joint characteristics and constraints would be expected to reduce performance to avoid saturation.

The described controller demonstrates one potential method of realizing faster response speed and greater accuracy in an interleaved manipulator. The variable redundancies present in higher dimensions limit its direct application to practical manipulators but the performance improvements demonstrate the benefits of a frequency-partitioning approach. The gains described are also modest, lying on the same scale as those realized under modal control due to the low first vibrational mode of the flexible segment. This limitation is fundamental to essentially all closed-loop flexible segment control schemes, though future work will reexamine this aspect.

4. A Clinically-Approximate Five DOF Manipulator

Having shown the potential for an interleaved manipulator to realize better performance than an equivalent continuum manipulator, we next designed a five degree-of-freedom (DOF) manipulator to expose joint and control design challenges. Where the single DOF system used a rigid joint to explicitly compensate the flexible segment's error, the multi-DOF manipulator pictured in figure 4.1 accesses a large workspace via two flexible segments (proximal shaft and catheter) with two rigid joints (distal pitch and roll) lying within the profile of proximal elements. The prototype traverses a flexible vasculature model to the notional patient's heart, where the rigid joints and catheter operate, with a virtual tip-to-target distance comprising the fifth joint. Four pin supports anchor and shape the vasculature model to generally mimic the path from the femoral vein to the right atrium, as depicted in figure 1.2. The pin supports and tube compliance bend according to external (organ) forces and with rotations of the proximal shaft, capturing to some degree the variable transmission properties encountered in cardiac catheter surgeries.

The scarcity and immaturity of powerful actuators and encoders at the <6mm diameter scale motivates remote actuation of the rigid joints. Remote actuation locates the drives external to the patient, allowing unconstrained servo selection at the cost of longer transmission elements that accumulate errors which vary with vasculature configuration. In this prototype large reductions in the distal rigid joints minimize these errors by the reduction ratio, offering a trade between rigid joint extent and performance (an increased gear reduction increases precision and accuracy

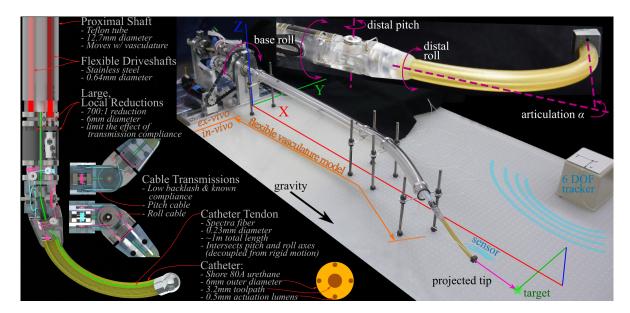


Figure 4.1: A 5D interleaved continuum-rigid manipulator. The background shows the overall manipulator, the left the joint design, and the upper right actual joints. The rigid joints are actuated by flexible stainless-steel (piano wire) driveshafts. These drive miniature planetary gearheads which then actuate cable assemblies to manipulate the rigid pitch and roll joints. As the roll joint is distal to the pitch, its cable assembly is communicated across the pitch joint by wrapping a transfer pulley. This wrapping ensures constant cable length but introduces coupling between the pitch and roll joints. The catheter has the same cross-section as that in the single DOF testbed.

while decreasing speed). In critical dimensions the prototype of figure 4.1 is 2-3 times larger than the maximal clinical device because further minimization would have substantially increased cost, complicated fabrication and assembly, and limited system inspection and performance evaluation.

Initial results of this prototype were shown in [72] and with improved rigid joints and a new controller in [73] and [74]. Improving on these, this chapter details the manipulator design and construction, presents an adaptive multi-DOF controller, and performs experiments that demonstrate characteristics of the manipulator, controller, and overall approach.

4.1 The Mechanical Design

The proximal roll joint is driven by a brushed 90W DC gear motor (Maxon Motor AG, #273754) having an 18:1 reduction (#166159) and a 500 count/turn encoder (#110514). This motor drives

a timing belt with a further 3.2:1 pulley reduction to which the proximal shaft connects. The proximal shaft is an 850mm long, 12.7mm outer diameter Teflon tube with a 1.5mm wall thickness, chosen for its torsional rigidity and bending stiffness.

The distal pitch rigid joint is driven by a brushed 90W DC motor (Maxon #273754) with a 1024 counts/turn encoder (#225787) connected to a 0.64mm diameter by 890mm long flexible stainless steel driveshaft. The distal end of this shaft drives a 699.5:1 planetary gearhead (Precision Microdrives #206-108) whose output connects to a cable transmission having a 2.6:1 reduction which drives the pitch joint. The distal roll rigid joint is identical to the distal pitch, with the exception of having a final cable reduction of 2.3:1. Both cables are made from 25N, 0.15mm diameter, braided, and Teflon coated Honeywell Spectra cable, sold as PowerPro fishing line. Figure 4.1 also details the rigid joint designs and shows how the catheter tendon intersects each of the proximal joint axes to decouple articulation from joint actuation. The joints themselves consist of four 3D-printed parts detailed in figure 4.1 riding on miniature ball bearings and printed via 3D Systems Viper stereolithography system at 0.05mm resolution. The pitch joint range of motion is ±50° and the distal rolls through ±200°.

The 110mm by 6.25mm distal catheter is actuated by an 11W DC gear motor (Maxon #222050) with a 109:1 planetary gearhead reduction (#143987) and a 512 count/turn encoder (#201940). This gearhead turns a 12mm diameter pulley to pull an approximately 1350mm long tendon made of Teflon-coated, 0.23mm Honeywell Spectra woven fiber (sold as PowerPro Braided Fishing Line). Along the length of the proximal section, three thin-walled 3mm Teflon tubes separate and loosely guide the two rigid joint driveshafts and the catheter tendon. The catheter is cast from Smooth-On PMC-780 urethane; figure 4.1 details the arrangement and sizes of the interior lumens. The Teflon-coated tendon exhibits substantially lower sliding friction against the urethane guiding lumen than typical in clinical catheters, due partly to the Teflon/urethane versus stainless steel/polyimide inter-

faces, but more substantially to the larger tendon/lumen clearance (0.23/0.50mm) versus that seen in clinical catheters (0.25/0.31mm according to [75]). The first vibrational mode of this catheter is approximately 1.5Hz measured at zero articulation and with all other joints undeflected and static. The motor position control loops are closed substantially above the system task bandwidth (30Hz > 10Hz) via a servo controller implemented in Matlab xPC 2009 (The Mathworks, Inc.). Manipulator configuration is again measured by an Ascension trakSTAR 6 DOF electromagnetic pose sensor mounted at the catheter tip operating at 300Hz, with the fifth, virtual tip-to-target distance projected along the sensor body.

The large, local reductions of the rigid joints are the defining aspect of this design; they determine the outer diameter (13mm) of the rigid joints and permit accurate rigid joint actuation even under substantial proximal segment movement. The small diameter of the flexible driveshafts accept these movements at the cost of shaft windup prior to the reduction, approximately 50° per N-mm of shaft torque according to linear elasticity, but only 0.07°/N-mm after the reduction. The shaft windup and large reduction introduce additional delay in the rigid joints which can lead to limit cycling; ignoring dynamic loads this is acceptable as long as the rigid joints are a few times faster than the slowest flexible segment. Output cable compliance is more detrimental as it leads directly to static and dynamic rigid joint errors; the Kevlar cables were tensioned to approximately 2N such that no slack or joint play were noticeable and fixed in place.

As seen in figure 4.1 the joints clearly lie within the profile of the proximal segment and the transition from the catheter is smooth, satisfying one of the basic rigid joint requirements. It is harder to quantify, and hence argue, the impact of the rigid joints on the manipulator safety. The joints, and particularly the pitch, can be designed with the ability to breakaway if some mechanical malfunction occurred necessitating the removal of the manipulator with an actuated pitch. Moreover,

¹Revolute joints, for instance, may be designed with breakaway mechanisms, wherein an actuated joint can come out of its nominal center of rotation to return to its un-actuated, safe configuration, much as various biological joints

as concerning as mechanical malfunction may be to the design engineer, a variety of studies have tabulated procedure complications for commercial robotic MIS systems (laparoscopic, endoscopic, and catheter) and broadly find these systems are not suffering mechanical failures [76–78]; the internal quality control systems from best-practice engineering and manufacturing appear sufficient to avoid these incidents [4].

Rigid joints can traverse the septum if there is sufficient compliance proximal and distal to them to allow the rigid elements to complete the motion. What matters is not the total extent of the rigid joints but that the rigid portion with the proximal and distal flexible segment compliances be sufficient to bend to the contacts. Simply put, longer rigid elements require greater proximal and flexible compliance up to the point that the elements span three contacts. Safety should be considered more deeply, though in the context of an application and only after we have determined some of the possibilities. Now, let's develop the manipulator model.

4.2 Kinematic Modeling and Parameter Spaces

In modeling the manipulator, we seek to relate the joint-space inputs (motor rotations) to the task-space outputs (projected tip positions). Broadly speaking, this input/output modeling may be performed by 1) rigorously investigating and describing every physical phenomenon affecting the output or 2) creating a simplified description that may be adapted to observations as needed. The first approach is essentially that of most robotic manipulators, but its application is challenged by the difficulty in modeling the catheter (as described in 1.2.2), flexible vasculature, and the other nonlinear effects. This leaves us with the second approach, requiring the formulation of a model that may be adapted to the expected nonlinearities (several of which are described in appendix B).

We seek a kinematic, as opposed to a dynamic, model because dynamic motions are challenging can be extended beyond their typical limits.

to separate from nonlinear effects in the observed task-space motion, and because a dynamic model is inherently more complex than a kinematic model and would therefore require adapting more parameters to approximate the observed motion. In choosing this simplified kinematic model, we should expect that it will become less accurate at greater joint velocities.

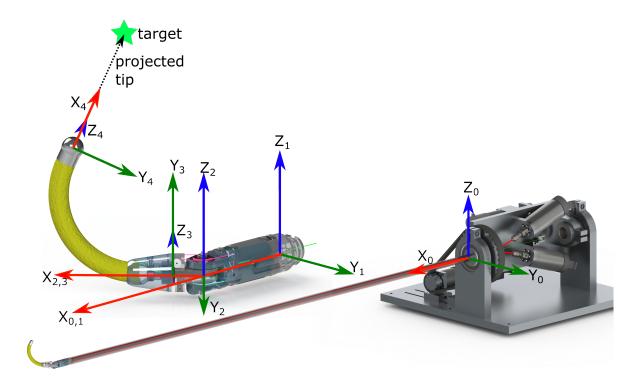


Figure 4.2: Kinematic frames.

The model is represented through homogeneous transformation matrices between the five principal frames noted in figure 4.2. Many unmodeled effects may occur within these basic frames, some of which are documented in appendix B, so the model includes additional rotations and translations that may be adapted to incorporate nonlinearities as they are encountered. The homogeneous

transformation matrices between these frames are

$$H_0^1 = T_x(t_{x01})T_y(t_{y01})T_z(t_{z01})R_y(r_{y01})R_z(r_{z01}) * R_x(q_1)$$

$$H_1^2 = T_x(L_{prox}) * R_z(q_2)$$

$$H_2^3 = T_x(L_{ptch}) * R_x(q_3)$$

$$\alpha = (k_p q_4)^{k_e} \qquad r_{cath} = \frac{L_{cath}}{\alpha}$$

$$H_3^4 = T_x(L_{roll})R_y(r_{y34})R_z(r_{z34}) * T_x(r_{cath}\sin(\alpha))T_y(r_{cath}(1-\cos(\alpha)))R_z(\alpha)$$

$$H_4^5 = R_y(r_{y45}) * T_x(q_5)$$

where t_* and r_* denote translations and rotations between the indicated frames that may be locally adapted; k_p is a gain on the catheter articulation angle and k_e is a variable exponent (see appendix B); and $L_{prox} = 25.5$, $L_{ptch} = 11.0$, and $L_{roll} = 12.5$ mm.

These individual transformations may be multiplied to find the predicted target location from kinematic parameters and joint angles, $H_5^0(kn,q) = H_1^0 H_2^1 H_3^2 H_4^3 H_5^4$, with the set of relations referred to as the manipulator's forward kinematics (FK). More will be said in later sections, but for now the as-measured kinematic parameters are

	measured		'intuitive'	'minimal'
t_{x01}	773	mm	constant	variable
t_{y01}	-103	mm	constant	variable
t_{z01}	0	mm	constant	variable
r_{y01}	0	0	constant	constant
r_{z01}	-28	0	constant	constant
r_{y34}	0	0	variable	constant
r_{z34}	0	0	variable	constant
k_p	1	[-]	variable	constant
k_e	1	[-]	variable	constant
L_{cath}	95	mm	variable	constant
r_{y45}	0	0	constant	constant

Table 4.1: Measured forward kinematic parameters and named sets for section 4.5.

4.3 An Adaptive Closed Loop Controller

From the initial experiments in appendix C and from observation of the nonlinearities in appendix B, we expect the manipulator performance to vary across the workspace and we seek control techniques that permit spatial adaptation. A classical technique is that of gain scheduling, whereby the workspace is divided into domains having differing local task controller gains that produce better behavior within each domain. As Slotine and Li [79, p.205] summarize, "The main problem with gain scheduling is that [it] has only limited theoretical guarantees of stability in nonlinear operation, but uses some loose practical guidelines such as 'the scheduling variables should change slowly' and 'the scheduling variables should capture the plant's nonlinearities'." The latter aspects are broadly the reasoning behind the intuitive set of kinematic parameters chosen above.

An alternative to discriminating by task position is to encode the system's expected parameter variations in a framework that estimates parameter values based on recent history. This includes the many works in model adaptive control and those employing forms of estimation. One common example is Kalman estimation, where some assumed statistical model of the parameter variations is used to estimate the parameters in a way that is consistent with the observations and expected variance. Since these approaches are only as useful as the assumed model or parameter variance are true, we can use many samples and frame the parameter estimation as an overdetermined optimization.

The interleaved manipulator encounters both spatial and historical nonlinearities, as well as their combination, so we should like to combine the parameter tuning aspects of the model adaptive and state estimating controllers with the location awareness of scheduled approaches. The general strategy, then, is to allocate sets of kinematic parameters across the workspace and estimate their values from recent motions proximate to each.

Figure 4.3 presents this general, adaptive parameter closed loop controller for the 5DOF manipulator. The task space controller is quite similar to that described in section C, but it is made much more accurate by the addition of the kinematic parameter estimation. Task space is described by a position and pointing vector $x = \{x, y, z, u_x, u_y, u_z\}^T$ having five independent variables that fully specify the 5DOF manipulator's position. The task space error is formed between the desired position $x^* = \{x, y, z\}^T$ and $\hat{x} = \{x, y, z\}^T$, which is the projection along the measured tip's pointing vector a distance equal to the commanded tip-to-target distance (such that the virtual tip-to-target distance joint has no error). The resulting xyz error vector is integrated, gained, and added to the desired task, while the pointing vector is taken directly as desired (due to the non-integrability of unit vectors). The inverse kinematics goal is

$$x_{goal} = \begin{cases} \left\{ x \\ y \\ z \right\}^* + \left\{ K_{ix}/s \\ K_{iy}/s \\ K_{iz}/s \right\} \left(\left\{ x \\ y \\ z \right\}^* - \left\{ x \\ y \\ z \right\} \right) \\ u_x^* \\ u_y^* \\ u_z^* \end{cases}.$$

In contrast to the prior work, this is an exactly-determined optimization problem with a unique global minimum. The goal position is used (dark green lines) to find a set of kinematic parameters estimated from prior motions (detailed next). Having a goal and local kinematics, the problem is packaged² for the NLOpt nonlinear optimization library [80] and quickly found by the local, derivative-free Subplex algorithm of [81]. As the manipulator executes the solution via the Matlab xPC servo controller, the Ascension trakSTAR measures the tip pose and sends the most recent 300Hz sample to LabVIEW over USB. These commanded joint positions and measured poses are

²Code available at https://github.com/bc0n/ICRMKinematics

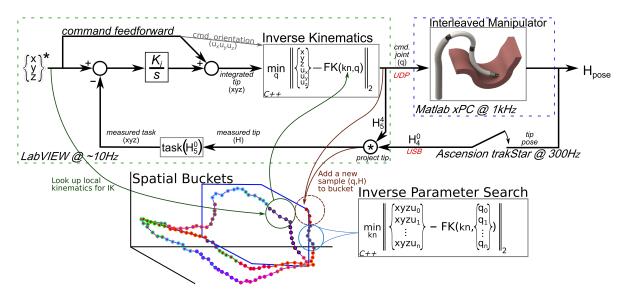


Figure 4.3: An adaptive closed loop controller. The task space integral controller pursues the current command using kinematic parameters from the nearest bucket (green arrows). The inverse kinematics use these locally-adapted parameters to find the set of joint positions that best accomplish the desired position (xyz) and orientation (tip pointing vector \vec{u}). The resulting joint positions q are sent to the manipulator's Matlab xPC-based servocontroller and the resulting pose H is measured by an Ascension trakSTAR electromagnetic tracker. Each set of commanded joint positions and measured poses is added (brown arrows) to the nearest bucket (typically this is the measured xyz task space position, other schemes are described in the text). When a bucket has sufficiently many spatially disparate samples, an optimization begins to find a new set of kinematic parameters that best relate the commanded joint positions to the measured manipulator pose. In this optimization, the cost function is the norm of the differences between the measured positions and tip pointing vectors and the position and pointing predicted by the forward kinematics with the test kinematic parameter set.

collected (brown lines) into buckets and used to further adapt the kinematic parameters to the non-linear features in each bucket. This outer loop is implemented in LabVIEW (National Instruments) and runs at approximately 10Hz on an Intel Core 2 Duo at 3GHz.

The next subsections detail the kinematic parameter estimation through spatially-distributed buckets.

4.3.1 Bucketing

Kinematic parameter estimation is performed on large sets of samples, each sample consisting of the five joint angles and the resulting, measured tip position. These samples are grouped into a distribution of buckets, figure 4.3, with each bucket containing samples that traveled through its domain and each estimating the best kinematic parameters for its contents. Buckets are strewn throughout the workspace to permit the local adaptation of kinematic parameters to the manipulator's behavior within each bucket.

The controller solves a bucket only when the number of points in the bucket are greater than the number of kinematic parameters to be found, but what should be an overdetermined problem with a unique solution may, depending on the configuration, kinematic model, and parameter limits, have redundant solutions. Uniquely solving the kinematic parameters desires some broad experience across the workspace, but this may span nonlinear features that are too complex to be approximated with the kinematic model. Depending on the bucket extent and the number and spatial distribution of samples, the kinematic parameters may vary greatly between optimizations on the same bucket as new samples are added, and they may differ substantially from those of neighboring buckets.

The choice of kinematic model is important, as are the parameter limits. In this work the parameter limits are held constant across all buckets, but they are a prime candidate for encouraging consensus across neighboring buckets. Some ability to seek consensus is practically important, since when a trajectory crosses from one bucket to the next the kinematic parameters will change, as seen later in section 4.7; the trajectory will remain smooth in task space, joint space will be discontinuous. This is one negative effect of the bucketing approach, that the kinematic parameters used in the inverse kinematics are not continuous between buckets. As the nonlinearities are captured by the buckets are locally smooth, the parameters would ideally reside on some surface which is anchored by each bucket's kinematic parameter estimate while also varying to meet the neighboring buckets, with the variance a function of the estimate's sample population and dissension. Alas.

Limited computational resources require that we establish an upper limit on the maximum

number of samples from one bucket that we use in the inverse parameter search. For well-traveled buckets, the most representative set of samples will have the maximal joint-distance, that is we want the n samples that individually maximize their distance from all other points. This would be an $O(n^3)$ search (like the all-pairs, shortest path problem) and while relatively simple to implement, the present controller simply chooses up to 1000 samples spread evenly across time for a moderately disparate set.

4.3.2 Distributing the Buckets

The buckets may be distributed throughout space according to several logics; given our desire to adapt the kinematics to locally approximate nonlinear errors, how do we determine 'local', especially when the nonlinearities are not known in advance? Several strategies are apparent and have been tested, but they have not been explored theoretically.

- We could evenly distribute n buckets across the task space and allocate samples by the measured task position. Increasing n improves the accuracy of each bucket's kinematic estimate, but increases computation time (as more buckets require more inverse parameter searching) and increases the likelihood of overfitting from the smaller number and more compact spatial distribution of the samples in each bucket, thereby decreasing the bucket-to-bucket consensus. If the task permits redundant motion, though, this scheme may result in buckets with samples from multiple, redundant configurations, leading to kinematic parameter estimates that are insufficiently local to any single configuration.
- Alternately, the buckets can be distributed in joint space, since some of the nonlinearities we seek to approximate are strong functions of joint actuation. But truly capturing the nonlinearities requires buckets on the same scale as the nonlinearities and would likely require extensive training to fully capture.

- Cases with nonstationary hysteretics may require a path- and velocity-based definition, and again the problem size is large.
- We might allocate buckets on-the-fly, according to the distance from the previous bucket or
 on observing kinematic errors above a particular threshold, akin to a depth-first search in
 kinematic space.
- Rather than tying sets of kinematic parameters to individual buckets, they may instead be classified by their expected variability and made functions of the manipulator's configuration.

The first two approaches were tested during development; with no apparent benefit to the joint space, the task space grid is used in the coming experiments. As implemented, the parameters are divided into constant and variable sets (see table 4.1), where the constant set are estimated with samples from all buckets (everywhere we've traveled) and only the variable are locally estimated. This might be generalized by allowing individual parameters to vary over different spatial and/or temporal scales, though the estimation algorithm becomes more challenging to conceive and implement.

4.3.3 Minimizing the Kinematic Error in a Bucket

Finding the set of kinematic parameters that best explains a bucket's samples requires solving

$$\min_{kn} \left\| \sum_{i} \overline{xyzu_i} - FK(kn, \overline{q_i}) \right\|_2$$

on the *i* samples selected from the bucket, where $\overline{xyzu_i}$ is a vector of the measured tip position and pointing angle and $\overrightarrow{q_i}$ is the corresponding vector of commanded joint angles. This optimization is implemented³ in the same C++ library as the IK and also uses the NLOpt nonlinear optimization library [80] and the local, derivative-free Subplex algorithm of [81].

³Code available at https://github.com/bc0n/ICRMKinematics

The optimization proceeds smoothly when the selected parameters of the kinematic model are non-redundant and the samples spatially disparate, and slowly when not. From a physical perspective, some of the kinematic parameters, like the 0-1 translations which locate the first frame or the tip pointing angle r_{y45} , should be constant and only solved when considering the entirety of motion. This sort of reasoning lead to the 'intuitive' set of kinematic parameters in table 4.1, while the 'minimal' set includes only those strictly necessary to achieve good task performance (xyz, since the u pointing vector is not used in the task controller). When the kinematic parameters are separated into constant/variable sets, the optimization runs first on the set of samples drawn from all buckets (with the variable parameters held at their default values) and then on each individual bucket with the just-found constant kinematic parameters held constant. This framing reduces the complexity of the kinematic optimization, and the gains can be massive: solving all 11 parameters on a bucket may take ten minutes, solving the 5 intuitive parameters a minute, and the 3 minimal parameters a second, all with equivalent residual values.

But the optimization residual is not the only metric of quality. The full parameter set will typically overfit each bucket, leading to large discontinuities in the kinematic estimates of adjacent buckets. This occurs even with well-populated buckets and spatially-disparate samples; conceptually, the only cure to overfitting is having samples that exhibit motion in each of the kinematic parameter dimensions. The intuitive set has configuration-dependent kinematic redundancies, such as between r_{y34} and k_p at low articulation, but the constant parameters ensure broad similarity across buckets. The minimal set avoids the redundancies, but it learns very little about the manipulator and *should* be a poor guide to maneuvers in unexplored areas of the host bucket and those beyond.

4.4 5D Evaluation Overview

Beginning this work we desired to increase the accuracy and execution speed of continuum manipulators. Towards this, section 4.5 looks at the manipulator's accuracy when executing half of a 3D cube over a number of iterations under open and closed loop controller variations. The next experiment, section 4.6, tests the quality of the learned parameters through non-repetitive training and point-to-point motions. Section 4.7 considers benefits of learning on the trajectory tracking accuracy and speed, while illustrating the fundamental challenge of non-collocated sensing.

4.5 Learning a Half-Cube

As a first experiment, figure 4.4 shows the manipulator executing a six-sided path around the edges of a 100mm side-length half-cube, located at approximately $q = \{0, 0, 0, 1, 200\}$ and well within the manipulator's workspace. The objective of this experiment is to demonstrate the controller of section 4.3 and how it adapts through eight controller variations. The controllers in table 4.2 differ according to: whether the task space loop is open or closed (see figure 4.3), the bucket grid spacing of 30 (few) or 6mm (many), and if the locally-adapted kinematics are from the intuitive or minimal sets defined in table 4.1 and described in section 4.3.

This experiment implements the controller of section 4.3 as a serial process: first, the IK uses the kinematic parameters of the nearest bucket (or as-measured if this is the first iteration) to find the joint angles for each of the commanded points; second, the set of joint angles are executed by the manipulator and the tip position is recorded; third, the commanded joint angles and measured positions are sorted into the buckets, the constant parameter optimization is performed on samples drawn from all buckets, and then the varying parameters are found on the samples in each bucket. When a new estimate of the constant parameters arrives, all of the existing buckets are re-solved

from this new basis to maintain consistency between buckets.

code	task loop	number of buckets	kinematic set
OFI	open	few	intuitive
OFM	open	few	minimal
OMI	open	many	intuitive
OMM	open	many	minimal
CFI	closed	few	intuitive
CFM	closed	few	minimal
CMI	closed	many	intuitive
CMM	closed	many	minimal

Table 4.2: Experimental controller configurations.

4.5.1 Convergence to the Command

The experiment begins in figure 4.4 from the neutral position with the as-measured kinematic parameters (wireframe) and attempts to move the projected tip through the commanded points (black). On the initial iteration the four open loop controllers all use the as-measured kinematics (no refinement yet), compute the same joint angles, and, on execution, trace virtually identical paths. The misplaced, distorted, and tilted result indicates the controller's challenge going forward.

4.5.2 Open Loop Variations

To begin, consider the convergence of the Open-loop, Few buckets, Intuitive parameter (OFI) variation in figure 4.5. Since this the task loop is open, the convergence from blue to red is entirely due to improvements in the kinematic parameter estimates. The task error plot is the Euclidean norm of the difference between the commanded point and the projected tip. To form a reasonable expectation of the task performance consider the approximately 1.5° angular repeatability seen in the hysteresis plots of section B and an average tip-to-target distance of 200mm; in the absence of kinematic error we would expect $\sin(1.5) * 200 = 5$ mm of task error.

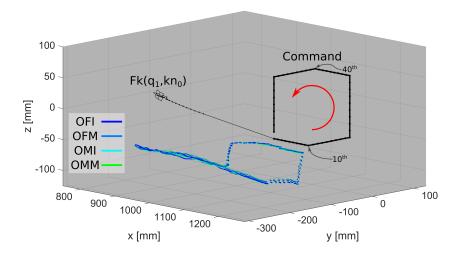


Figure 4.4: Initial iteration of the open loop variations. The wireframe manipulator is true to size and drawn from the as-measured kinematics configured to point at the first target.

The third graph looks at the number of samples in the bucket used to approach the 40th commanded point and how that bucket's kinematic parameters evolve with iteration. Since the kinematic space is hard to directly visualize, the graph plots the kinematic objective on the bucket's contents from each of the iterations. The darkest blue line results from evaluating the initial, asmeasured kinematics on the chosen bucket's samples after each iteration. Similarly, the yellow line indicates how well the kinematics estimated on the bucket from iteration ten perform over all of the iterations. As the bucket's population grows with each iteration, these lines show how the kinematic estimate is improving. The residual will typically decrease with iteration and from blue to red, as seen, but equally important is the later flatness, indicating that the kinematic parameters have converged.

The failure to achieve the upper corner of points is due to the controller getting 'stuck' with one particular set of sub-optimal parameters and no ability to improve them. To see this, consider the interplay between the inverse kinematics and the inverse kinematics estimation: the IK is given a task and uses the kinematics from the nearest, visited bucket. In one variation of the stuck

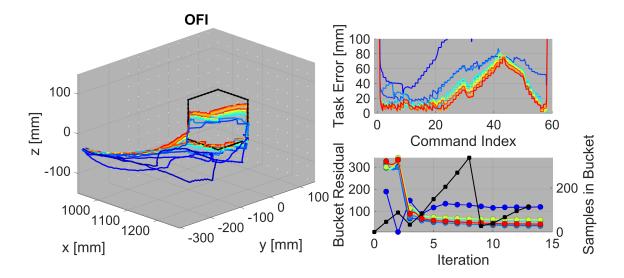


Figure 4.5: Convergence of the open-loop, few-bucketed, intuitive kinematic parameter controller variation. The jet colormap (blue-green-orange-red) shows the refinement with iteration in 3D on the left and in task error on the top right. The background grid of points indicate the buckets' extent and distribution. The bottom right graph shows the bucket's kinematic parameter convergence and the number of samples versus iteration (black). See the text for further description.

bucket, the IK's solution does not lead to motion through the nearest bucket or through some new, closer bucket and, having learned nothing, the cycle repeats. The other stuck variation is when the manipulator travels through the nearest bucket, but that motion is very similar to prior runs through that bucket, to again result in no learning and repetition of the cycle. So, the controller does well along the lower two sides but is never able to converge to the upper corner because it is unable to travel through the regions with error.

Increasing the bucket density from every 30mm to every 6mm, open/many/intuitive in figure 4.6, should allow the bucket estimates to more accurately represent their contents and may lead to more accurate IK solutions. The increased locality of the buckets and potential accuracy is counteracted by the fewer and more spatially-compact samples in each bucket. Smaller buckets with fewer samples increases the chance of overfitting; for example the aqua trajectory shows a very large diversion from the command between indices 10 and 14. Increased bucket density does

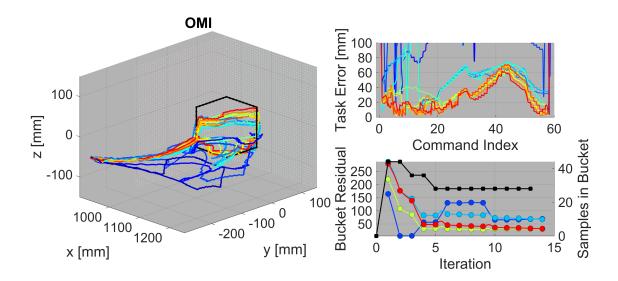


Figure 4.6: Convergence of the open-loop, many-bucketed, intuitive kinematic parameter controller variation. See the text for discussion.

not substantially help the open loop convergence, as the manipulator again gets stuck along the upper sides.

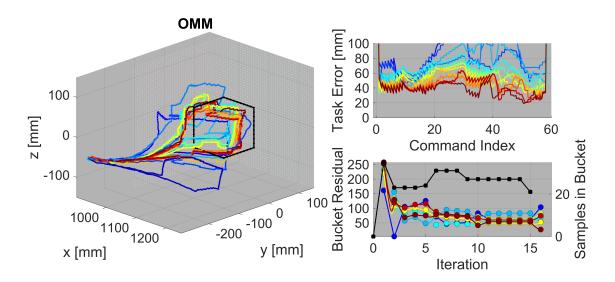


Figure 4.7: Convergence of the open-loop, many-bucketed, minimal kinematic parameter controller variation. See the text for discussion.

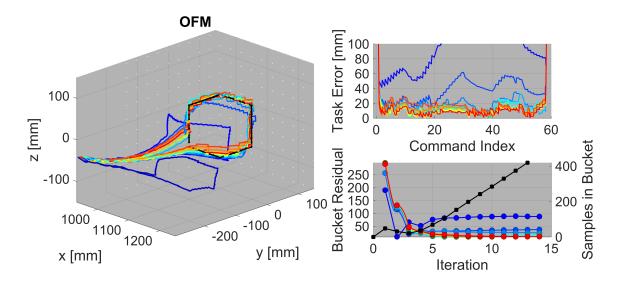


Figure 4.8: Convergence of the open-loop, few-bucketed, minimal kinematic parameter controller variation. See the text for discussion.

Switching to the minimal parameter set with many buckets, open/many/minimal in figure 4.7, shows that the minimal representation can also get stuck, though in a different area than the intuitive set. Comparing this to the open/few/minimal variation of figure 4.8, shows that good performance is actually possible open loop. Whether the open loop controller get stuck is, then, primarily a function of the bucket density and where it gets stuck depends on the kinematic representation.

4.5.3 Closed Loop Variations

Closing the task loop is one obvious and effective way to encourage the manipulator to converge to the command and prevent stuck kinematic estimates. To promote the maximum amount of learning, the controller's integrator is reset after achieving a given command. Pursuit of the next command, then, begins from the same position as in the open loop controllers (compare the darkest blue of figures 4.5 and 4.9) and the figure shows large sweeps from the previous command, out to

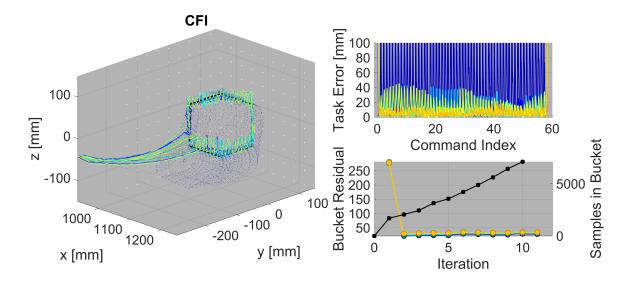


Figure 4.9: Convergence of the closed-loop, few-bucketed, intuitive kinematic parameter controller variation. The first iteration is voluminous and rendered as very tiny points to show the subsequent iterations. See the text for further discussion.

the open loop solution, and then closed loop motion back inward.

Resetting the integrator is atypical, but it serves to create large motions through many buckets and it more clearly shows the effect of learning over iterations. In figure 4.9, the peak heights are the error after the move to the open loop positions, with the subsequent valley the result of the closed loop convergence to the command. The benefit of the parameter adaptation is the decreasing peak heights from blue to red, that the closed loop controller has to do less work to achieve the desired command.

Between the few- and many-bucketed controllers in figures 4.9 and 4.10, the increased number of buckets leads to faster, more accurate learning. In the limit that there is one bucket for every commanded point, the parameters of each bucket can be viewed as a point-specific gain and likened to the classical gain scheduling approach. As before, though, increased bucket density can lead to overfitting and the estimates are more susceptible to sensor noise, so in that limit there may need to be some heuristics to guide the optimization

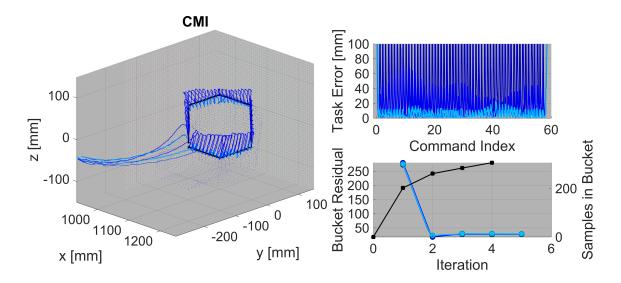


Figure 4.10: Convergence of the closed-loop, many-bucketed, intuitive kinematic parameter controller variation. The first iteration is voluminous and rendered as very tiny points to show the subsequent iterations. See the text for further discussion.

Looking at the closed loop minimal variations, figures 4.11 and 4.12, the integral resets again deposit a surplus of samples into both the few- and many-bucketed variations, which the minimal kinematic representation learns very quickly. Only slight refinements in the parameters occur after the first closed loop iteration, leading to very accurate performance.

Which of the closed loop controllers is best, since all achieve the commands? One simple metric is to accumulate the task error with time; this includes the open loop accuracy (feedforward term) and how quickly the integral controller eliminates that error. This metric is computed for the four closed loop variations on the final trip around the half cube in figure 4.13. It suggests that more buckets lead to faster, more efficient convergence to the command. The minimal kinematic representations enables slightly better performance than the intuitive sets, implying that attempts to provide a physically-informed representation may not be necessary or beneficial.

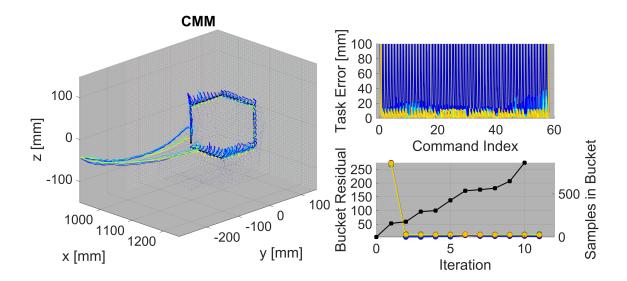


Figure 4.11: Convergence of the closed-loop, many-bucketed, minimal kinematic parameter controller variation. The first iteration is voluminous and rendered as very tiny points to show the subsequent iterations. See the text for further discussion.

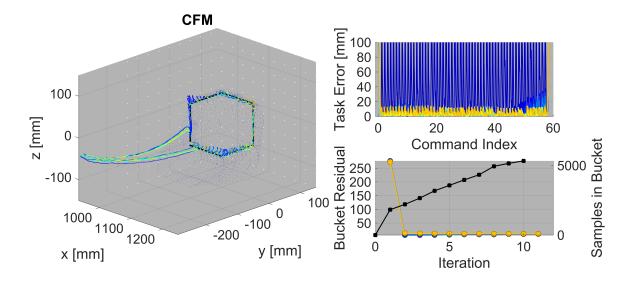


Figure 4.12: Convergence of the closed-loop, few-bucketed, minimal kinematic parameter controller variation. The first iteration is voluminous and rendered as very tiny points to show the subsequent iterations. See the text for further discussion.

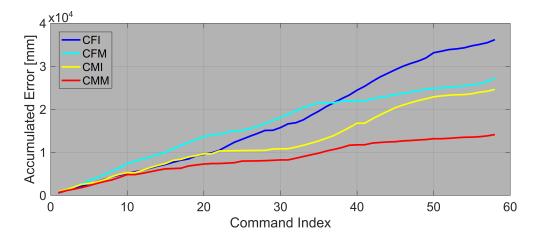


Figure 4.13: Since the task errors can be very noisy, accumulating them with time gives a clearer indication of controller performance.

4.5.4 What Are the Consensus Kinematic Parameters?

Let us close this first experiment by asking what are the true kinematic parameters, and have any of the controllers reached them? For simplicity consider the 10^{th} and 40^{th} commanded points, the lower and upper corners of the half-square as indicated in figure 4.4. In the open loop controllers, the lower corner converged quickly while the upper often led to stuck behavior, so these points will give an indication of the well-learned and stuck conditions. The bold column labels in tables 4.3-4.6 recall which parameters were allowed to vary between buckets and which were held constant across all buckets (table 4.1).

First, none of the controllers affirmed the as-measured kinematics, straying by approximately 100mm. Of the intuitive controllers the closed loops' agreement on the constant parameters is encouraging, while the differing estimates of the open loop controllers is not surprising given their significantly fewer and spatially-compact samples. It is interesting to see that shorter catheters were preferred, though it's too much to say that the controllers are detecting the actual shortening of the catheter from to the applied tension. Local adaptations in the other parameters are plausible; for instance, the switching of r_{y34} between the corners may be in response to out of plane catheter

bending from changing gravitational loads of the tip's sensor, laser, and their cables.

The final minimal kinematic parameters of tables 4.5 and 4.6 are harder to interpret, mainly due to the oddity of holding distal features constant while varying proximal ones in a serial manipulator. That aside, the closed loop controllers again agree on the gross translation of the manipulator for the upper and lower corners. The majority of the minimal's constant parameters are not physically plausible – it would be hard to not notice the -57° error in r_{z01} – but these concerns are subsumed into the variable translations. Given a sufficiently large parameter range on the variable set, the minimal representation does not need to solve the constant parameters.

	t_{x01}	t_{y01}	t_{z01}	r_{y01}	r_{z01}	r_{y34}	$\mathbf{r_{z34}}$	$\mathbf{k_p}$	$\mathbf{k_e}$	${ m L_{cath}}$	r_{y45}	$n_{samples}$
measured	773	-103	0	0	-28	0	0	1	1	95	0	
OFI	831.0	-177.7	20.0	29	-52	-0.0	-5.7	1.2	0.8	80.0	-12	16,368
OMI	819.8	-181.7	-21.2	12	-46	-0.0	-5.7	1.0	1.5	80.0	-5.7	19,194
CFI	803.0	-167.2	-38.3	17	-52	-12	17	0.8	1.2	82.4	5.7	122,275
CMI	804.2	-171.8	-40.5	17	-52	-12	0.0	1.1	0.8	91.7	5.7	75,136

Table 4.3: Final kinematic estimates of the intuitive set for the bucket nearest to the 10^{th} commanded point (lower corner). The CMI variation has fewer samples than the CFI from completing only 5 circuits of the half-cube.

	t_{x01}	t_{y01}	t_{z01}	r_{y01}	r_{z01}	$\mathbf{r_{y34}}$	${f r_{z34}}$	${ m k_p}$	$\mathbf{k_e}$	${ m L_{cath}}$	$ m r_{y45}$	$n_{samples}$
measured	773	-103	0	0	-28	0	0	1	1	95	0	
OFI	831.0	-177.7	20.0	29	-52	17	-17	1.2	0.8	80.0	-12	16,368
OMI	819.8	-181.7	-21.2	12	-46	17	-17	1.2	0.8	80.0	-5.7	19,194
CFI	803.0	-167.2	-38.3	17	-52	17	-0.0	0.9	1.2	80.0	5.7	$122,\!275$
CMI	804.2	-171.8	-40.5	17	-52	17	-12	0.9	1.5	92.8	5.7	$75,\!136$

Table 4.4: Final kinematic estimates of the intuitive set for the bucket nearest to the 40^{th} commanded point (upper corner).

It is difficult to learn much about the manipulator from the kinematic parameter estimates of the minimal set; its ease of solution and accuracy are attractive but it is not obvious whether the

	$\mathbf{t_{x01}}$	$\mathbf{t_{y01}}$	$\mathbf{t_{z01}}$	r_{y01}	r_{z01}	r_{y34}	r_{z34}	k_p	k_e	L_{cath}	r_{y45}	$n_{samples}$
measured	773	-103	0	0	-28	0	0	1	1	95	0	
OFM	753.6	-93.5	-5.6	29	-57	-17	-12	1.2	0.8	126.5	-5.7	$16,\!275$
OMM	750.0	-97.9	-19.5	17	-57	-5.7	5.7	0.9	1.0	149.8	-12	18,740
CFM	786.2	-92.0	20.0	29	-57	-12	-5.7	1.2	0.8	95.8	-0.0	110,169
CMM	792.1	-93.6	20.0	29	-57	-12	-12	1.2	0.8	92.6	-0.0	110,876

Table 4.5: Final kinematic estimates of the minimal set for the bucket nearest to the 10^{th} commanded point (lower corner).

	$\mathbf{t_{x01}}$	$\mathbf{t_{y01}}$	$\mathbf{t_{z01}}$	r_{y01}	r_{z01}	r_{y34}	r_{z34}	k_p	k_e	L_{cath}	r_{y45}	$n_{samples}$
measured	773	-103	0	0	-28	0	0	1	1	95	0	
OFM	750.2	-120.7	-66.4	29	-57	-17	-12	1.2	0.8	126.5	-5.7	16,275
OMM	750.0	-129.8	-98.0	17	-57	-5.7	5.7	0.9	1.0	149.8	-12	18,740
CFM	761.0	-109.9	-50.6	29	-57	-12	-5.7	1.2	0.8	95.8	-0.0	110,169
CMM	763.1	-113.1	-46.0	29	-57	-12	-12	1.2	0.8	92.6	-0.0	110,876

Table 4.6: Final kinematic estimates of the intuitive set for the bucket nearest to the 40^{th} commanded point (upper corner).

adaptations are useful beyond the half-cube trajectory. The next experiment tests this by following several, non-repeating shapes and evaluating the benefit of adaptation.

4.6 Moving Within a Known Space

The previous experiment attempted the same trajectory multiple times and gradually adapted the local kinematics to improve open and closed loop performance, but it is not clear if the controller has learned anything beyond that exact motion. In this section the controller trains on a series of planar spirals, figure 4.14, and performs a series non-repetitive motions within the learned space. None of the trajectories are exactly repeated and the evaluation motions are point-to-point across the learned space.

While we expect the training to lead to better open loop performance, the more interesting question for figure 4.15 is if the training makes the closed loop motion more direct or efficient. For this we will again use the accumulated task error, since it incorporates both the feedforward error

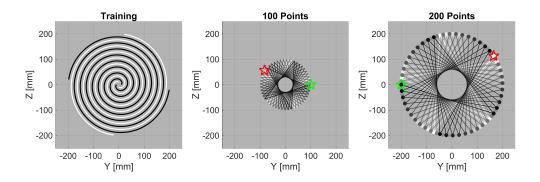


Figure 4.14: Left. Four 500 point training spirals from 10 to 20mm. Center. Point-to-point motion over a 100mm radius circle, from the green to red stars counter-clockwise. Right. 200mm radius circle for the same.

and how long that error resists the integral controller.

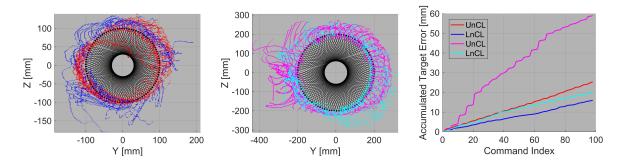


Figure 4.15: Left. 100 points are distributed around a 100mm radius circle wholly within the learned space; from the current point, the next is 40 away, with the trajectory proceeding counterclockwise through approximately 40 cycles. Center. 100 points distributed at a radius of 200mm, or along the learned space's outer edge. Proceeding as previously, proximity to the edge and the greater travel distance lead to greater error. Right. Accumulated error for the un-learned, closed loop (UnCL) and learned, closed loop, many-bucketed, minimal parameter set (LnCL) controllers. Less error and error slope indicates better relative performance.

Where the closed loop experiments in section 4.5 quickly learned to accurately execute the half-cube, the learned controller here does not inspire the same confidence. That is, we may have expected some open loop error but also that the controller would directly drive to the commanded point in a relatively linear fashion, but the learned controller in figure 4.15 mostly resembles a

hairball. The likely explanation is that the nearly-identical learning of the previous experiment included direction-dependent actuation hysteretics that the non-repetitive commands here do not encounter. This may also be an artifact of the planar training shapes, that the controller was denied access to useful knowledge and the estimates suffered as a result. More experiments could draw this out further; we turn instead to the dynamic response of the manipulator.

4.7 Performance

Local adaptation of the kinematic parameters improved the task accuracy of the preceding experiments; let's now look at the performance while tracking a moving target. Where the controllers previously waited for task convergence before tackling the next command, this experiment is configured to be always moving. Good performance follows from accurate open loop (feedforward) steps and aggressive-but-stable integral controller gains.

In the training phase a 10mm error requirement and the same four spirals from the prior experiment were used to learn the same 200mm radius planar volume. The evaluation phase executes a single-cycle spiral through 200, 100, 40, or 10 points. Commanding the points at approximately 1Hz (limited by the IK speed) results in a task space velocity that changes with the number of points and the spiral radius: 200 pts = 1.6:4.7 mm/s, 100 = 3.2:9.4, 40 = 8.1:23, and 10 = 35:91.

Figure 4.16 shows the results of four controller configurations, each having integral controller gains of 0.5 (k_i in figure 4.3). The first, yellow controller (UnCl) is a typical closed loop controller with the as-measured (un-learned) kinematics. It exhibits substantial open loop error, but the integral controller eventually delivers good tracking. Learning the local kinematic parameters (LnOL) improves the open loop error and, critically, this performance is not a function of the command rate (ignoring manipulator dynamics). Using the learned parameters (LnCL) gives better

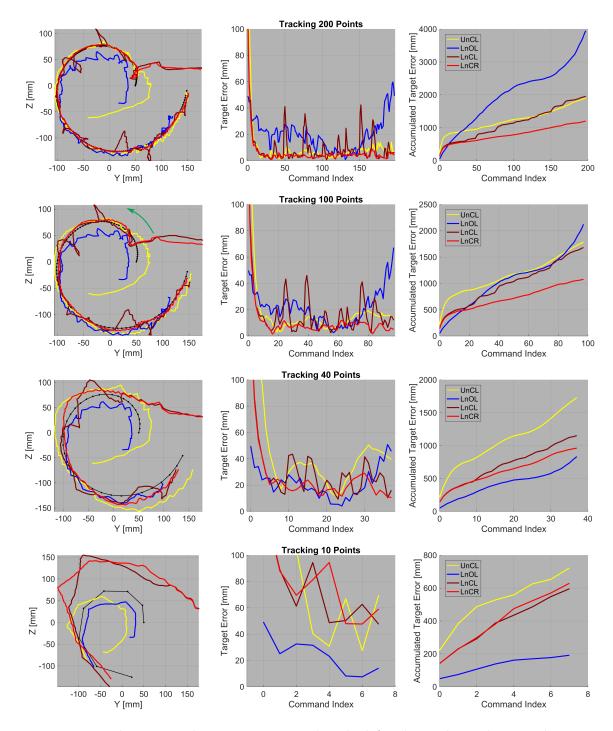


Figure 4.16: Tracking a spiral at increasing speeds. The left column shows the spatial motion, with the center and right providing the tracking error. Four controllers are used, an un-learned kinematics/closed loop (UnCL), a learned/open (LnOL), a learned/closed (LnCL), and a learned/closed with recast integral error (LnCR). (A data error prevented LnCL and LnCR from using the parameters used in the more accurate LnOL trajectory.)

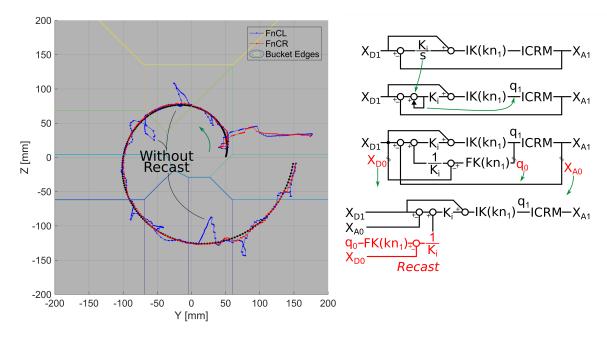


Figure 4.17: Looking more closely at the learned/closed (blue, LnCL) and learned/closed with recast integral error (red, LnCR). The LnCL trajectory exhibits several large flares from the commanded, 200 point trajectory, despite the relative slowness of the command. As shown on the right, crossing buckets requires modifying the accumulated error via the new kinematics.

initial, open loop response⁴ and often less overall tracking error (right column).

There is a penalty in using the learned parameters with an integral controller, as in LnCL, in that the accumulated errors within the integral are relative to the manipulator's previous performance. When the kinematic parameters change from one bucket to the next, figure 4.17, the effect of the accumulated integral will change, often substantially, and the controller will need to work through this change before it can consider the next command. One could reset the integrators whenever the kinematic parameters change and return the controller to the open loop position of that command, but since errors have accumulated we should look at the control in greater detail. The right side of figure 4.17 manipulates the block diagram to expose the accumulated error inside the previous joint command. Using the current kinematics and the previous joint angles allows the accumulated error to be recast from the former kinematic set to the new. This results in substantially smoother

⁴At least nominally; in contrast to the LnOL, the LnCL and LnCR were only minimally trained and therefore show larger than appropriate open loop error.

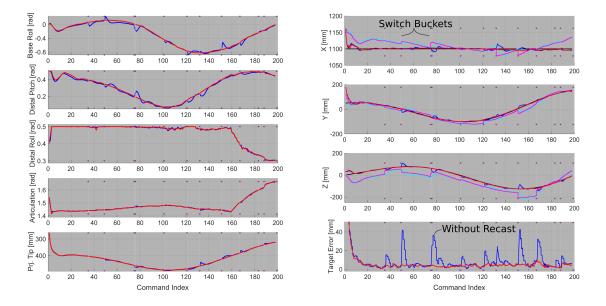


Figure 4.18: Joint, left, and task space, right, for the trajectory in figure 4.17. The learned/closed (LnCL) is in blue and the learned/closed with recast integral error (LnCR) is red. Bucket boundaries are noted by slight vertical lines; as buckets are chosen by the commanded point both controllers experience the same transitions at the same times. In the right pane the accumulated (integral) error is shown in aqua (LnCL) and magenta (LnCR).

performance across bucket edges, as suggested in figure 4.17 and more clearly in joint and task space, figure 4.18.

For example, consider the second transition at command index 49; in task space the recast integral (magenta) crosses the bucket edge and immediately comes to a new value, leaving zero increase in the task error along any axis and no change to the joint angles. In comparison, the unmodified error (aqua) blithely applies its old error to the command and creates a large error that requires time to remove and eventually settles at the recast level. As both controllers indicate that some error is required to achieve the command, this also demonstrates why zeroing the accumulated error would degrade performance. While good task performance follows from accurate kinematic parameters, this recasting is still appropriate when the parameters are dubious. As shown, it is really a means of ensuring that the IK goal and search occur from the same kinematic basis.

The last row in figure 4.16 exhibits the fundamental motivation for this adaptive controller: reducing open loop error is essential to realizing faster performance as all closed loop approaches are limited by the controller's non-collocated dynamics. At slower speeds the closed loop controller can improve overall accuracy, but learning is the only vehicle for improving temporal response.

4.8 In Summary

The 5D interleaved manipulator is complex, but that complexity opens certain avenues to increased task accuracy and performance. Key to both of these are accurate kinematic parameters that locally approximate the manipulator's nonlinear behavior. Surprisingly, crude approximations appear acceptable if they are sufficiently local and span the task space (not directly but as functions of the forward kinematics). And while much work can be done to further estimate, characterize, and avoid manipulator nonlinearities, improving accuracy, increased temporal response is attained only by open loop improvements.

5. In Conclusion

Compliant, minimally invasive treatment of deep tissue is virtually impossible with today's systems because they are incapable of competently and safely interacting with tissue. Awaiting theoretical advances, the traditional tendon-actuated architecture is too constraining; only by changing the architecture can we hope to realize better performance, with the design burden to ensure that we maintain the safety profile of current cardiac devices. In this expanded design space, physical and theoretical limitations may be avoided or delayed through system and control design.

This work demonstrated the potential of Interleaved Continuum-Rigid Manipulation to increase the task accuracy and speed beyond today's minimally invasive systems. The rigid joints introduce actuation redundancy, which an interleaved controller may use to avoid continuum nonlinearities and dynamic excitations, or to prefer particular configurations that may improve task accuracy, permit greater end-effector forces, or avoid environment obstacles. Unrestrained by actuator and joint design, the single degree of freedom system realized faster, more accurate performance through an understanding of the system actuators and the error dynamic. It is conceptually possible to attain this same performance in the 5 degree of freedom system, though the joint transmissions and redundant control are challenging.

In the more general case of accurately and quickly executing commanded trajectories, the 5D system demonstrated the essential nature of the kinematic model to the system performance. The task of implementing the combined kinematic estimator and closed loop controller entails many

considerations, but when well-trained, the bucketing approach consistently achieved open loop errors of 10mm. Closing the loop can further increase this accuracy and ensure performance in unlearned areas, but care must be taken that the controller dynamics do not substantially limit the overall system performance.

To conclude, recall the overall concept in figure 2.1. In both experimental systems, the flexible catheter segments are essential to safely accessing a large workspace. Diligent mechanism and control design allows actuation nonlinearities to be compensated by the limited-extent rigid joints. The added redundancy increases task dexterity and performance, and may allow the complete avoidance of flexible segment actuation nonlinearities. These benefits open a new, unappreciated avenue for increasing continuum manipulator capability and enable operation in challenging environments. Despite this work, there remain many interesting questions in interleaved manipulation precisely because the new design freedoms are liberating.

Bibliography

- [1] Carlos Eduardo Diaz, Roemi Fernandez, Manuel Armada, and Felipe de Jesus Garcia Gutierrez. State of the art in robots used in minimally invasive surgeries. natural orifice transluminal surgery (notes) as a particular case. *Industrial Robot: An International Journal*, 42(6):508–532, 2015.
- [2] M. Jung, P. Morel, L. Buehler, N. C. Buchs, and M. E. Hagen. Robotic general surgery: current practice, evidence, and perspective. *Langenbeck's Archives of Surgery*, 400(3):283–292, 2015.
- [3] Mathias Hoeckelmann, Imre J Rudas, Paolo Fiorini, Frank Kirchner, and Tamas Haidegger. Current capabilities and development potential in surgical robotics. Int J Adv Robot Syst, 12:61, 2015.
- [4] Johannes Bonatti, George Vetrovec, Celia Riga, Oussama Wazni, and Petr Stadler. Robotic technology in cardiovascular medicine. *Nature Reviews Cardiology*, 11(5):266–25, May 2014.
- [5] Dimitrios Stefanidis, Robert D. Fanelli, Ray Price, and William Richardson. Sages guidelines for the introduction of new technology and techniques. Surgical Endoscopy, 28(8):2257–2271, 2014.
- [6] Nisha Patel, Ara Darzi, and Julian Teare. The endoscopy evolution: the superscope era. Front-line gastroenterology, 6:101–107, 2015.

- [7] Ewout A Arkenbout, Paul WJ Henselmans, Filip Jelínek, and Paul Breedveld. A state of the art review and categorization of multi-branched instruments for notes and sils. Surgical endoscopy, 29(6):1281–1296, 2015.
- [8] Tao Shen, Carl A Nelson, Kevin Warburton, and Dmitry Oleynikov. Design and analysis of a novel articulated drive mechanism for multifunctional notes robot. *Journal of mechanisms* and robotics, 7(1), Feb 2015.
- [9] Selene Tognarelli, Marco Salerno, Giuseppe Tortora, Claudio Quaglia, Paolo Dario, Marc Oliver Schurr, and Arianna Menciassi. A miniaturized robotic platform for natural orifice transluminal endoscopic surgery: in vivo validation. Surgical endoscopy, pages 1–8, 2015.
- [10] Amir Degani, Howie Choset, Alon Wolf, and Marco Zenati. Highly articulated robotic probe for minimally invasive surgery. In Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on, pages 4167–4172. IEEE, 2006.
- [11] Amir Degani, Howie Choset, Alon Wolf, Takeyoshi Ota, Marco Zenati, et al. Percutaneous intrapericardial interventions using a highly articulated robotic probe. In *Biomedical Robotics* and *Biomechatronics*, 2006. BioRob 2006. The First IEEE/RAS-EMBS International Conference on, pages 7–12. IEEE, 2006.
- [12] Ziad Issa, John M Miller, and Douglas P Zipes. Clinical Arrhythmology and Electrophysiology: A Companion to Braunwald's Heart Disease: Expert Consult: Online and Print, chapter 15. Elsevier Health Sciences, 2012.

- [13] Jason G. Andrade, Lena Rivard, and Laurent Macle. The past, the present, and the future of cardiac arrhythmia ablation. *Canadian Journal of Cardiology*, 30(12):s431–s441, December 2014.
- [14] Matthew Wright and Sanjiv M. Narayan. Ablation of atrial fibrillation. Trends in Cardiovascular Medicine, 25(5):409–419, July 2015.
- [15] Ziad Issa, John M Miller, and Douglas P Zipes. Clinical Arrhythmology and Electrophysiology: A Companion to Braunwald's Heart Disease: Expert Consult: Online and Print. Elsevier Health Sciences, 2012.
- [16] Alexander Wutzler, Thomas Wolber, Wilhelm Haverkamp, and Leif-Hendrik Boldt. Robotic ablation of atrial fibrillation. *Journal of visualized experiments: JoVE*, 99(99), 2015. Amigo Remote Catheter System.
- [17] Michelle G Glowny and Frederic S Resnic. What to expect during cardiac catheterization.

 Circulation, 125(7):e363–e364, 2012.
- [18] Ramachandran S. Vasan, Martin G. Larson, Daniel Levy, Jane C. Evans, and Emelia J. Benjamin. Distribution and categorization of echocardiographic measurements in relation to reference limits. *Circulation*, 96:1863–1873, 1997.
- [19] Roberto M. Lang, Michelle Bierig, Richard B. Devereux, Frank A. Flachskampf, Elyse Foster, Patricia A. Pellikka, Michael H. Picard, Mary J. Roman, James Seward, Jack Shanewise, Scott Solomon, Kirk T. Spencer, Martin St. John Sutton, and William Stewart. Recommendations for chamber quantification. European Heart Journal - Cardiovascular Imaging, 7(2):79–108, 2006.
- [20] Magellan robotic catheter 10fr, 2015.

- [21] Ziad Issa, John M Miller, and Douglas P Zipes. Clinical Arrhythmology and Electrophysiology: A Companion to Braunwald's Heart Disease: Expert Consult: Online and Print. Elsevier Health Sciences, 2012.
- [22] William W. B. Chik, Michael Anthony Barry, Jim Pouliopoulos, Karen Byth, Christine Midekin, Abhishek Bhaskaran, Gopal Sivagangabalan, Stuart P. Thomas, David L. Ross, Alistair McEwan, Pramesh Kovoor, and Aravinda Thiagalingam. Electrogram gated radiofrequency ablations with duty cycle power delivery negates effects of ablation catheter motion.

 Circulation: Arrhythmia and Electrophysiology, 113, 2014.
- [23] Sven Knecht, Tobias Reichlin, Nikola Pavlovic, Beat Schaer, Stefan Osswald, Christian Sticherling, and Michael Kuhne. Contact force and impedance decrease during ablation depends on catheter location and orientation: insights from pulmonary vein isolation using a contact force-sensing catheter. Journal of Interventional Cardiac Electrophysiology, 43(3):297–306, 2015.
- [24] Boris Schmidt, Roland R. Tilz, Kars Neven, K.R. Julian Chun, Alexander Furnkranz, and Feifan Ouyang. Remote robotic navigation and electroanatomical mapping for ablation of atrial fibrillation. *Circulation: Arrhythmia and Electrophysiology*, 2:120–128, 2009.
- [25] Hanno U. Klemm, Daniel Steven, Christin Johnsen, Rodolfo Ventura, Thomas Rostock, Boris Lutomsky, Tim Risius, Thomas Meinertz, and Stephan Willems. Catheter motion during atrial ablation due to the beating heart and respiration: Impact on accuracy and spatial referencing in three-dimensional mapping. *Heart Rhythm*, 4(5):587–592, 2007.
- [26] Shelten G. Yuen, Daniel T. Kettler, Paul M. Novotny, Richard D. Plowes, and Robert D. Howe. Robotic motion compensation for beating heart intracardiac surgery. *International Journal of Robotics Research*, 28(10):1355–1372, 2009.

- [27] Rogerio Richa, Antonio P.L. Bo, and Philippe Poignet. Beating heart motion prediction for robust visual tracking. In *Robotics and Automation (ICRA)*, 2010 IEEE International Conference on, pages 4579–4584. IEEE, 2010.
- [28] Saurabh Kumar, Joseph B. Morton, Karen Halloran, Steven J. Spence, Geoffrey Lee, Michael C. G. Wong, Peter M. Kistler, and Jonathan M. Kalman. Effect of respiration on catheter-tissue contact force during ablation of atrial arrhythmias. *Heart Rhythm*, 9(7):1041–1047, 2012.
- [29] Waqas Ullah, Ross J. Hunter, Victoria Baker, Mehul B. Dhinoja, Simon Sporton, Mark J. Earley, and Richard J. Schilling. Factors affecting catheter contact in the human left atrium and their impact on ablation efficacy. *Journal of Cardiovascular Electrophysiology*, 26(2):129–136, February 2015.
- [30] Hisao Matsuda, Abdul Shokor Parwani, Philipp Attanasio, Martin Huemer, Alexander Wutzler, Florian Blaschke, Wilhelm Haverkamp, and Leif-Hendrik Boldt. Atrial rhythm influences catheter tissue contact during radiofrequency catheter ablation of atrial fibrillation: comparison of contact force between sinus rhythm and atrial fibrillation. *Heart and vessels*, pages 1–9, 2015.
- [31] Mahta Khoshnam and Rajni V Patel. Robotics-assisted catheter manipulation for improving cardiac ablation efficiency. In *Biomedical Robotics and Biomechatronics* (2014 5th IEEE RAS & EMBS International Conference on, pages 308–313. IEEE, 2014.
- [32] Giovanni Di Salvo, Pio Caso, Rosalia Lo Piccolo, Angela Fusco, Alfonso R. Martiniello, Maria G. Russo, Antonio DOnofrio, Sergio Severino, Paolo Calabr, Giuseppe Pacileo, Nicola Mininni, and Raffaele Calabr. Atrial myocardial deformation properties predict maintenance

- of sinus rhythm after external cardioversion of recent-onset lone atrial fibrillation. *Circulation*, 112(3):387–395, 2005.
- [33] Josef Kautzner, Petr Neuzil, Hendrik Lambert, Petr Peichl, Jan Petru, Robert Cihak, Jan Skoda, Dan Wichterle, Erik Wissner, Aude Yulzari, and Karl-Heinz Kuck. Efficas ii: optimization of catheter contact force improves outcome of pulmonary vein isolation for paroxysmal atrial fibrillation. *Europace*, 17(8):1229–1235, June 2015.
- [34] Hiroshi Nakagawa, Josef Kautzner, Andrea Natale, Petr Peichl, Robert Cihak, Dan Wichterle, Atsushi Ikeda, Pasquale Santangeli, Luigi Di Biase, and Warren M. Jackman. Locations of high contact force during left atrial mapping in atrial fibrillation patients. *Circulation: Arrhythmia and Electrophysiology*, 6:746–753, July 2013.
- [35] Philip Aagaard, Andrea Natale, and Luigi Di Biase. Robotic navigation for catheter ablation: benefits and challenges. Expert review of medical devices, 12(4):457–469, 2015.
- [36] Lara S. Cowan and Ian D. Walker. The importance of continuous and discrete elements in continuum robots. *International Journal of Advanced Robotic Systems*, 10:1–13, 2013.
- [37] William McMahan, V Chitrakaran, M Csencsits, D Dawson, Ian D Walker, Bryan A Jones, M Pritts, D Dienno, M Grissom, and Christopher D Rahn. Field trials and testing of the octarm continuum manipulator. In Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on, pages 2336–2341. IEEE, 2006.
- [38] Kenneth Salisbury, William Townsend, B Ebrman, and David DiPietro. Preliminary design of a whole-arm manipulation system (wams). In Robotics and Automation, 1988. Proceedings., 1988 IEEE International Conference on, pages 254–260. IEEE, 1988.
- [39] baxter: Collaborative robits pioneer, 2015.

- [40] Shigeo Hirose. Biologically inspired robots: snake-like locomotors and manipulators. Oxford University Press, 1993.
- [41] P. Sears and Pierre Dupont. A steerable needle technology using curved concentric tubes. In Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on, pages 2850–2856, Oct 2006.
- [42] Valentin Falkenhahn, Alexander Hildebrandt, Rudiger Neumann, and Oliver Sawodny. Model-based feedforward position control of constant curvature continuum robots using feedback linearization. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 762–767. IEEE, 2015.
- [43] David B. Camarillo, Christopher F. Milne, Christopher R. Carlson, Michael R. Zinn, and J. Kenneth Salisbury. Mechanics modeling of tendon-driven continuum manipulators. *Robotics*, *IEEE Transactions On*, 24(6):1262–1273, 2008.
- [44] Robert J Webster III, Allison M Okamura, and Noah J Cowan. Toward active cannulas: Miniature snake-like surgical robots. In *Intelligent Robots and Systems*, 2006 IEEE/RSJ International Conference on, pages 2857–2863. IEEE, 2006.
- [45] C. Bergeles and Guang-Zhong Yang. From passive tool holders to microsurgeons: Safer, smaller, smarter surgical robots. Biomedical Engineering, IEEE Transactions on, 61(5):1565–1576, May 2014.
- [46] Farahnaz Maghooa, Agostino Stilli, Yohan Noh, Kaspar Althoefer, and Helge A Wurdemann.

 Tendon and pressure actuation for a bio-inspired manipulator based on an antagonistic principle. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 2556–2561. IEEE, 2015.

- [47] Yan Bailly, Yacine Amirat, and Georges Fried. Modeling and control of a continuum style microrobot for endovascular surgery. *Robotics, IEEE Transactions On*, 27(5):1024–1030, 2011.
- [48] W.S. Rone and P. Ben-Tzvi. Continuum robot dynamics utilizing the principle of virtual power. *Robotics, IEEE Transactions on*, 30(1):275–287, Feb 2014.
- [49] Guru Subramani and Michael R Zinn. Tackling friction-an analytical modeling approach to understanding friction in single tendon driven continuum manipulators. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 610–617. IEEE, May 2015.
- [50] Jinwoo Jung, Ryan S. Penning, and Michael R. Zinn. A modeling approach for robotic catheters: effects of nonlinear internal device friction. *Advanced Robotics*, 28(8):557–572, 2014.
- [51] David B. Camarillo, Christopher R. Carlson, and J. Kenneth Salisbury. Task-space control of continuum manipulators with coupled tendon drive, volume 54 of Springer Tracts in Advanced Robotics, chapter Experimental Robotics, pages 271–280. Springer, 2009.
- [52] S.B. Kesner and R.D. Howe. Position control of motion compensation cardiac catheters.

 Robotics, IEEE Transactions on, 27(6):1045–1055, Dec 2011.
- [53] Mircea Ivanescu, Decebal Popescu, and Nirvana Popescu. A decoupled sliding mode control for a continuum arm. *Advanced Robotics*, 29(13):831–845, 2015.
- [54] D. Braganza, D. M. Dawson, Ian D. Walker, and N. Nath. A neural network controller for continuum robots. Robotics, IEEE Transactions On, 23(6):1270–1277, 2007.
- [55] Apoorva D Kapadia, Katelyn E Fry, and Ian D Walker. Empirical investigation of closed-loop control of extensible continuum manipulators. In *Intelligent Robots and Systems (IROS 2014)*, 2014 IEEE/RSJ International Conference on, pages 329–335. IEEE, 2014.

- [56] A Melingui, Rochdi Merzouki, JB Mbede, Coralie Escande, and N Benoudjit. Neural networks based approach for inverse kinematic modeling of a compact bionic handling assistant trunk. In Industrial Electronics (ISIE), 2014 IEEE 23rd International Symposium on, pages 1239–1244. IEEE, 2014.
- [57] M. Rolf and J.J. Steil. Efficient exploratory learning of inverse kinematics on a bionic elephant trunk. Neural Networks and Learning Systems, IEEE Transactions on, 25(6):1147–1160, June 2014.
- [58] Michele Giorelli, Federico Renda, Marcello Calisti, Andrea Arienti, Gabriele Ferri, and Cecilia Laschi. Neural network and jacobian method for solving the inverse statics of a cable-driven soft arm with nonconstant curvature. *Robotics, IEEE Transactions On*, 31(4):823–834, 2015.
- [59] D.A. Bristow, M. Tharayil, and A.G. Alleyne. A survey of iterative learning control. *Control Systems*, *IEEE*, 26(3):96–114, June 2006.
- [60] Robert H. Cannon and Eric Schmitz. Initial experiments on the end-point control of a flexible one-link robot. *Robotics Research, International Journal on*, 3(3):62–75, 1984.
- [61] Andre Preumont. Vibration Control of Active Structures: An Introduction, volume 96 of Solid Mechanics and its Applications. Kluwer Academic Publishers, 2 edition, 2002.
- [62] Ryan Penning and Michael Zinn. A combined modal-joint space control approach for minimally invasive surgical continuum manipulators. *Advanced Robotics*, 28(16):1091–1108, 2014.
- [63] Kourosh Parsa, Jorge Angeles, and Arun K. Misra. Control of macro-micro manipulators revisited. Journal of Dynamic Systems, Measurment, and Control, 127(4):688–699, 2004.
- [64] D. N. Nenchev. Reaction null space of a multibody system with applications in robotics.
 Mechanical Sciences, 4:97–112, 2013.

- [65] T. Wongratanaphisan and M. Cole. Robust impedance control of a flexible structure mounted manipulator performing contact tasks. *Robotics, IEEE Transactions on*, 25(2):445–451, April 2009.
- [66] Loris Roveda, Federico Vicentini, Nicola Pedrocchi, Francesco Braghin, and Lorenzo Molinari Tosatti. Impedance shaping controller for robotic applications involving interacting compliant environments and compliant robot bases. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 2066–2071. IEEE, 2015.
- [67] Angelica Ruszkowski, Omid Mohareri, Sam Lichtenstein, Richard Cook, and Septimiu Salcudean. On the feasibility of heart motion compensation on the davinci surgical robot for coronary artery bypass surgery: Implementation and user studies. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 4432–4439. IEEE, May 2015.
- [68] Benjamin L Conrad, Jinwoo Jung, Ryan S Penning, and Michael R Zinn. Interleaved continuum-rigid manipulation: An augmented approach for robotic minimally-invasive flexible catheter-based procedures. In *Robotics and Automation (ICRA)*, 2013 IEEE International Conference on, pages 718–724. IEEE, 2013.
- [69] S. J. Schroeck, W. C. Messner, and R. J. McNab. On compensator design for linear timeinvariant dual-input single-output systems - mechatronics, ieee/asme transactions on. *Mecha*tronics, IEEE/ASME Transactions on, 6(1):50–57, March 2001.
- [70] Ryan S Penning, Jinwoo Jung, Justin Borgstadt, Nicola J Ferrier, Michael R Zinn, et al. Towards closed loop control of a continuum robotic manipulator for medical applications. In Robotics and Automation (ICRA), 2011 IEEE International Conference on, pages 4822–4827. IEEE, 2011.

- [71] Katsuhiko Ogata. System Dynamics. Prentice Hall, 3rd edition, 1998.
- [72] Benjamin L Conrad and Michael R Zinn. Interleaved continuum-rigid manipulation approach: Development and functional evaluation of a clinical scale manipulator. In *Intelligent Robots* and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on, pages 4290–4296. IEEE, IEEE, September 2014.
- [73] Benjamin Conrad and Michael Zinn. Closed loop task space control of an interleaved continuum-rigid manipulator. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 1743–1750. IEEE, 2015.
- [74] Benjamin L Conrad and Michael R Zinn. Interleaved continuum-rigid manipulation: An approach to increase the capability of minimally-invasive surgical systems. IEEE/ASME Transactions on Mechatronics, 2016.
- [75] David B. Camarillo, Christopher R. Carlson, and J. Kenneth Salisbury. Configuration tracking for continuum manipulators with coupled tendon drive. *Robotics, IEEE Transactions On*, 25(4):798–808, 2009.
- [76] Orhan Agcaoglu, Shamil Aliyev, Halit Eren Taskin, Sricharan Chalikonda, Matthew Walsh, Meagan M. Costedio, Matthew Kroh, Tomasz Rogula, Bipan Chand, Emre Gorgun, Allan Siperstein, and Eren Berber. Malfunction and failure of robotic systems during general surgical procedures. Surgical Endoscopy, 26(12):3580–3583, December 2012.
- [77] Diana C. W. Friedman, Thomas S. Lendvay, and Blake Hannaford. Instrument failures for the da vinci surgical system: a food and drug administration maude database study. Surgical Endoscopy, 27(5):1503–1508, May 2013.

- [78] Homa Alemzadeh, Ravishankar K. Iyer, Zbigniew Kalbarczyk, Nancy Leveson, and Jaishankar Raman. Adverse events in robotic surgery: A retrospective study of 14 years of fda data. In The 50th Annual Meeting of the Society of Thoracic Surgeons, 2015.
- [79] Jean-Jacques E. Slotine and Weiping Li. On the adaptive control of robot manipulators. The International Journal of Robotics Research, 6(3):49–59, 1987.
- [80] Steven G. Johnson. The nlopt nonlinear-optimization package, 2015.
- [81] Thomas H. Rowan. Functional Stability Analysis of Numerical Algorithms. PhD thesis, University of Texas Austin, 1990.
- [82] Paul Zarchan and Howard Musoff. Fundamentals of Kalman Filtering: A Practical Approach, volume 208 of Progress in Astronautics and Aeronautics. American Institute of Aeronautics and Astronautics, 2005.
- [83] D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the lipschitz constant. Journal of Optimization Theory and Applications, 79(1):157–181, 1993.
- [84] Adria Colome and Carme Torras. Closed-loop inverse kinematics for redundant robots: Comparative assessment and two enhancements. *Mechatronics*, *IEEE/ASME Transactions on*, 20(2):944–955, 2015.
- [85] Stefano Chiaverini, Giuseppe Oriolo, and Ian D. Walker. Springer Handbook of Robotics, chapter 11. Springer Science & Business Media, 2008.
- [86] Jorge Nocedal and Stephen Wright. Numerical optimization. Springer Science & Business Media, 2006.

A. Constant Curvature Catheter Geometry

Geometric relations for a constant curvature, 100mm long by 6mm diameter catheter. From the sector-arc length formula $s=\alpha R$, we assume that the backbone does not compress by letting s=L=100mm and find the arc radius $R=L/\alpha$. This is the well-known catheter singularity of $R\to\infty$ as $\alpha\to0$. The tip position is given by

$$x = \frac{L}{\alpha} (1 - \cos \alpha) \tag{A.1}$$

$$y = \frac{L}{\alpha} \sin \alpha \tag{A.2}$$

and the backbone is traced by letting $\theta = 0..\alpha$ in place of α above.

The effect of tendon displacements is found by considering the inner arc

$$\alpha = \frac{L}{R} = \frac{L_{in}}{R - \frac{w}{2}}.$$

Displacing a tendon an increment δ results in a change of articulation $\Delta \alpha$ of

$$\Delta \alpha = \alpha_0 - \alpha_1 = \frac{L_{in}}{R - \frac{w}{2}} - \frac{L_{in} + \delta}{R - \frac{w}{2}} = \frac{\delta}{R}.$$

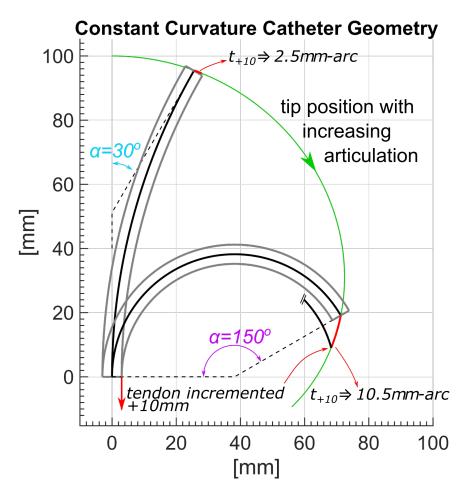


Figure A.1: Two catheters articulated to $\alpha=30^\circ$ and 150°. Articulation angle α is defined as the angle between the tip pointing vector and the neutral position, light blue in the figure. This angle also appears between the tip perpendicular and the base perpendicular, as in the magenta. Increasing articulation results in the spiral shape indicated by the green path. Applying a tendon displacement of 10mm to the two initial poses results in substantially different tip motions; 2.5mm-arc becomes 10.5mm-arc at $\alpha=150^\circ$.

B. 5D Joint Performance

This appendix documents the performance and obvious nonlinearities of the four physical joints. To begin evaluation, let's start by comparing the ideal pitch joint performance to the single DOF testbed from section 3.3; in figure 3.4 the voice coil actuates the catheter tip from 11° to 25° in approximately 1s and with the pitch joint local reduction of 700*2.6:1, an equivalent actuation of the pitch joint would require reaching 4107rpm in one second. This is well within the speed and power envelope of the drive motors, but the shaft windup, cable compliance, and catheter dynamics limit realization in the 5D system.

B.1 FRFs Indicate Control Bandwidth

To more generally capture these effects, figure B.1 gives frequency response functions (FRFs) for each of the drives, showing the relationship between the commanded and measured joint movement as a function of frequency. First, the distal roll, pitch and catheter have approximately the same corner frequency near 1.2Hz, while the base roll is slightly faster at 1.9 Hz. From the start, then, the joints in this 5D manipulator are equally-fast and a frequency partitioning controller like that used in section 3.2 is not appropriate. The catheter response depends on the operating point, with the near-neutral 1rad displaying near unity response up to the corner, while the 2 and 3rad points decrease early. This reduction at higher articulations may be due to the catheter's tendon hysteresis, but more likely is the influence of a tip-mounted 3g laser diode, the Ascension sensor,

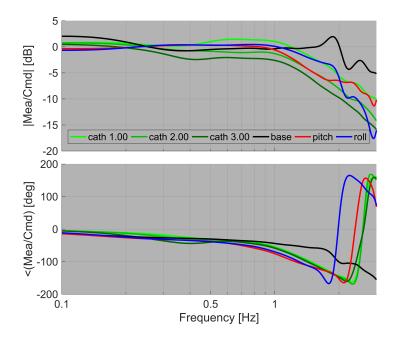


Figure B.1: Frequency response functions for each joint. The joints were driven by 0.1rad amplitude chirps spanning 0.01 to 3Hz and the tip pose measured by the Ascension sensor. Since using the IK to transform the measured pose into joint space would introduce uncertainty, the cleanest measured pointing angle was normalized and used in computing the FRFs, which were then filtered to 10Hz to more clearly show joint characteristics. Articulation operating points of 1, 2, and 3rad were used in the catheter FRF.

and their cables. These disturbances are also responsible for the base roll's resonance at 1.9Hz, where the long, free length of the two cables allowed a slow oscillation to occur. The observed corner frequencies are the result of the device dynamics and the control loop delays; some of these may be reduced with subsystem modeling and refinements to the system architecture,

B.2 Hysteresis Loops

Next, let's look at the hysteresis seen when commanding small (0.2rad) sinusoids at frequencies of 0.1, 0.4, 0.8, and 1Hz to each of the joints. The IK are again avoided by using the measured pointing angles and normalizing the command and measured, though since this is an angular comparison we do not see hysteretic effects in translation (such as the base roll's walking behavior from appendix

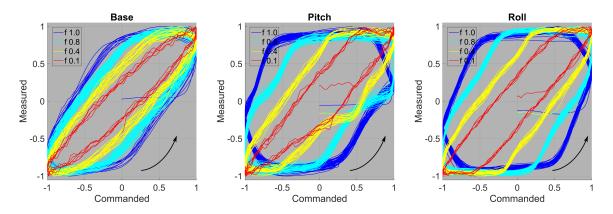


Figure B.2: Commanded joint angle versus measured tip angle while executing a 0.2rad sinusoid at frequencies of 0.1, 0.4, 0.8, and 1Hz. The command and measurement are normalized to -1:1 to avoid using the inverse kinematics

B). Figure B.2 shows the first three joints and figure B.3 the catheter around two operating points.

First, the nearly identical lag in the lowest frequency sinusoid (0.1Hz = red) across the joints is due to controller and sensor delays (as in figure 4.3, a 10ms timed loop iteratively issuing sets of joint angles through a UDP broadcast to the 1kHz servo controller which are then measured at 300Hz and saved by the controller). At least five cycles were performed at each frequency, and the distance between these cycles indicates the manipulator's repeatability: a 0.1[-] wide band while executing a 0.2rad command indicates an angular repeatability of approximately 1.5°. Next, with the slight exception of the pitch joint, the slow hysteresis loops are symmetric and centered on the origin, indicating that the joints behave similarly when driven forward or backward. Considered individually, the widening of the base roll with increasing frequency is likely due to some proximal shaft windup and bending and twisting dynamics in the catheter (the catheter is actuated to 1rad to improve sensor resolution; given the tendon and tip loads, the lowest vibrational mode may be an twisting instead of bending).

The pitch joint is rather complex, so consider first the more typical behavior of the roll joint.

At low frequencies the hysteresis loop is quite square, and dominated by the flexible driveshaft windup. When the joint is driven on a sine, the shaft reverses direction every peak and valley,

necessitating the shaft's unwinding and winding-up in the opposite direction, during which no tip movement occurs. We can use the length of the horizontal sections to directly estimate the shaft windup: the 0.2rad amplitude command was normalized to the range of -1:1, so a 0.3[-] flat segment indicates a windup of $0.3/2*0.2=0.17^{\circ}$ in joint space or 120° on the other side of the reduction (shaft space). As the sinusoid frequency increases the windup becomes larger and the shaft spends more time unwinding and winding than actually turning driving the distal joint. Approaching 1Hz, we may be seeing the shaft's second torsional mode (the first is the (un)winding) and the onset of catheter bending dynamics (as with the base roll, it is articulated to 1rad).

Returning to the pitch joint, it experiences the same flexible driveshaft dynamics as the roll and, since the catheter is straight and unloaded for the pitch test, lower catheter dynamics. Since the roll joint is distal to the pitch, and these joints are coupled, rotating the pitch requires a small, counteracting actuation of the roll to prevent the roll joint from twisting as the pitch rotates (see figure 4.1 for reference). This coupling leads the pitch hysteresis to be a function of the windup of both of the driveshafts. The negative-going-positive behavior is odd; the shaft effects should result in a symmetric plot, so the more likely cause is some disturbance after the reductions, perhaps from the laser and sensor cables or from a manufacturing defect or wear in the joint itself.

Switching to consider the catheter, figure B.3, the same set of sines was performed at 1 and 3rad articulations. We again see the delay at low frequencies, but notice how the loops become increasingly square with frequency. Though skipping a comprehensive analysis, this effect is consistent with tendon/lumen friction: as the sinusoid frequency increases, the tendon forces also increase, and since the friction force is a function of the tendon tension the increasing force can cause the tendon to 'bite' into the lumen and further delay the onset of catheter movement. Notice also that the 3rad plot is not centered on the origin, but rather that the negative lobe never reaches the negative corner. The motion between the highly articulated upper right corner to the less-articulated lower

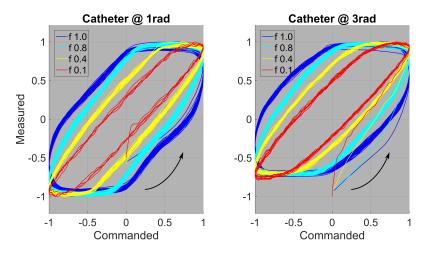


Figure B.3: Commanded joint angle versus measured tip angle while executing a 0.2rad sinusoid at frequencies of 0.1, 0.4, 0.8, and 1Hz from an initial articulation of 1 or 3rad. The command and measurement are normalized to -1:1 to avoid using the inverse kinematics.

left corner (counter-clockwise) is one of relaxation, so as the tendon tension decreases we need to wait for the decreasing tension to propagate along the tendon. As a last comment, the catheter dynamics, if any, will not be very obvious in this plot because articulating the catheter increases the in-plane resonant frequencies while the increasing friction leads to greater damping; the only place that dynamics will be readily observable is during relaxation as an antisymmetric bloat.

This section attempted to document and explain some characteristics of each joint's performance. Each joint has some obvious nonlinearities which the controller will need to considered; in addition to those described, some additional, physical traits are described in Appendix B. In summary, we see that actuation speed will be jointly limited by the catheter and flexible drive shaft dynamics. Motions appear to be quite repeatable, while the absolute accuracy is a function of the workspace and direction of motion.

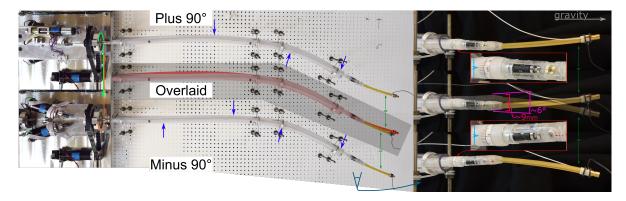


Figure B.4: The primary unmodeled nonlinearities experienced in actuating the base roll joint. First, rotating the base from 90° to -90° leads to movement of the vasculature, shown by overlaying the positive and negative images. The rotation affects where the proximal shaft contacts the vasculature model, as indicated by the blue arrows. As the contacts change so too does the tip position, translating by 9mm and pointing angle by 6° . Gravity loading on the tip sensor, laser diode, and their cables also contributes. Inset in the right details are zooms of the distal end of the proximal shaft, showing a circumference of 43mm with positions of 1mm ($+90^{\circ}$) and 19mm (-90°) indicated. For a rotation of 180° , the distal end rotates only 150° .

B.3 Nonlinearities in the Base Roll

The two flexible segments have the greatest modeling uncertainty. As described in section 4.1, this first, base roll joint is composed of a Teflon tube that suffers from twisting and 'walking' inside the guiding, vinyl vasculature model. Both the twisting and walking are caused by contact with the vasculature model at the two or three regions indicated in figure B.4. Teflon, and thermoplastics in general, exhibits hysteretic effects over differing time spans; since the manipulator is typically parked with zero joint angles there is a residual bend that increases contact forces at some angles. The overlaid images clearly show the walk of approximately 9mm in the global Z direction and a shaft windup of approximately 30° , requiring the kinematics to have a translation in Z and a pointing angle adjustment that both vary with the tip-to-target distance q_1 .

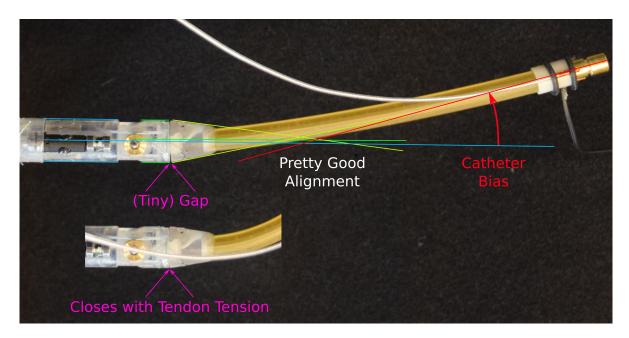


Figure B.5: The pitch is very well aligned in this photo, $< 5^{\circ}$ is probably typical. The initial bias in the catheter is obvious.

B.3.1 Nonlinearities in the Rigid Joints

Several effects are likely to occur in both rigid joints: driveshaft frictions and torques vary with configuration, the absence of a local joint encoder leads to centering error, cable compliance after the gearhead output delays actuation, disturbances from loading more distal elements. Not all of these are photogenic, but figure B.5 does find some slight axial play along the distal roll axle. Beyond this nit-picking, the rigid joints appear to be in good form.

B.3.2 Nonlinearities in the Catheter Articulation

Figure B.5 shows the appreciable, residual bias in the catheter. This leads to a deadzone near the neutral axis; the inverse kinematics typically prefers the catheter keep 1rad of articulation, so this deadzone is very rarely encountered. Other nonlinearities were described in section 1.2; see in figure B.6 an example of the 5DOF's non-constant catheter curvature.

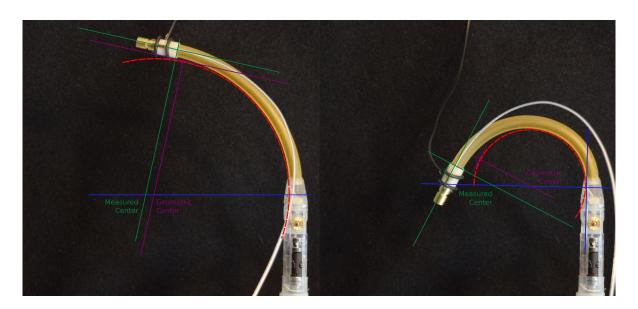


Figure B.6: Friction affects every catheter motion, severity increasing with articulation.

C. Initial Experiments with Closed Loop Control

This section briefly reviews a prior controller for the 5DOF manipulator through a task space closed loop controller.

C.1 5D Forward Kinematics

The global coordinate system is shown in figure 4.1 and located on the flat-holed plane at the face of the bearing clamp block. The forward kinematics are composed of homogeneous transformation matrices of the form ${}_{5}^{0}H = {}_{1}^{0}H_{2}^{1}H_{3}^{2}H_{4}^{3}H_{5}^{4}H$. During actuation of the base roll the flexible vasculature can move, but owing the stiffness of the proximal shaft there is little translation or walking of the proximal shaft about the vasculature exit. The first transformation matrix captures the static displacement through the vasculature model and the base roll joint by

$$_{1}^{0}H = T_{x}(811.1)T_{y}(-157)T_{z}(-8.7)R_{z}(-0.37)R_{x}(q_{p1}),$$

with q_{p1} being the joint position. This base roll frame is attached to the proximal section of the rigid joint yet located at the pitch axis. The distal pitch frame is located at the pitch axis and attached to the distal pitch, having a transformation matrix of

$${}_{2}^{1}H = R_{z}(q_{p2}).$$

The distal roll joint is offset from the pitch axis by approximately 8.2mm, resulting in

$$_{3}^{2}H = T_{x}(8.2)R_{x}(q_{p3}).$$

The catheter begins at this distal roll frame, with the tip located by

$$R = L_c/q_{p4} \tag{C.1}$$

$$x = R\sin(q_{p4}) \tag{C.2}$$

$$y = R(1 - \cos(q_{n4})) \tag{C.3}$$

producing

$$_{4}^{3}H = T_{y}(y)T_{x}(x)R_{z}(q_{p4}).$$

Finally, the projected tip is a simple transformation along the x axis

$$_{5}^{4}H = T_{x}(q_{p5}).$$

C.2 A Multiple DOF Controller

The overall controller topology is shown in figure C.1; the task is to move the projected tip \hat{x} to a position x^* in 3D space, as shown in figure 4.1. The tip positions \hat{x} are formed from the filtered, measured tip pose projected along the tip pointing vector the commanded tip-to-target distance (the virtual tip-to-target distance joint has no error). The resulting task-space error vector is integrated, gained, and added to the desired task. The integral and gain are, then, moving the inverse kinematics (IK) goal in response to manipulator errors. As described in section 1.2.4, fundamental uncertainties about catheter articulation and the plant's variable poles limit the rote design and application of common control methods; a simple integral controller is sufficient for the present purpose but will be revisited in future work. IK convergence to the integrated/gained/summed task command produces a new set of joint angles which are executed by the manipulator. The outer loop is implemented in LabVIEW (National Instruments, Inc.) and runs at 10Hz on an Intel Core 2 Duo at 3GHz. Joint angle commands q_p^* are communicated to the Matlab xPC servo controller

running at 30Hz via UDP, while the Ascension sends the most recent 300Hz sample to LabVIEW over USB.

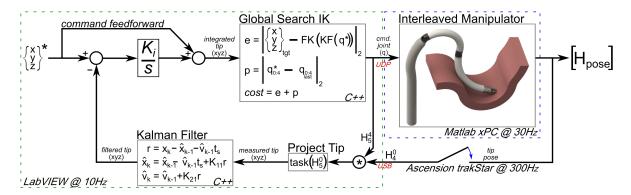


Figure C.1: The multi-DOF task-space controller. The 1D, 2-state polynomial Kalman filter is used individually on the measured xyz tip positions and within the inverse kinematics on the joint angles, with unique filter gains for each [82].

The task-space controller in figure C.1 is so named because the outermost error and integral controller operate on the task vector, the projected tip xyz position. In this design the summation of the integrated and feedforward terms is a virtual xyz position, the 'integrated tip.' The discrepancy between the actual and integrated tip locations is an indication of the magnitude of the unmodeled plant nonlinearities. A common, alternate framing is to move the IK before the outermost error and place another IK along the feedback path to convert the measured position into estimated joint angles. This joint-space topology has two difficulties in a redundant interleaved manipulator: some effort must be taken to ensure that the now-separated IKs seek the same local minima while not compromising their search; and the controller must be prevented from actuating redundant joints in opposing fashions where no task-space improvement is seen. This distinction is only apparent in nonconvex and underspecified (if the task DOF is less than the manipulator DOF such that redundant motions remain while achieving the task) IK implementations.

The IK are implemented in a C/C++ library which calls the NLopt nonlinear optimization

library ([80], version 2.4.2) and particularly the global, derivative-free, constrained, and objective-scaling DIRECT algorithm of [83]. The IK objective is to minimize the sum of the task error and the difference of the four physical joints from the previous solution constrained within the physical joint limits. Note that this objective does not eliminate the manipulator redundancy – because the joint movement minimization 'task' lies within the tip positioning task – but it does introduce a preference between possible solutions. Even when preferring to minimize the difference between successive solutions in the objective, experiments suffered excessive joint chatter due to sensor noise jostling between solutions, motivating filtering of the joint positions. While evaluating this objective, each candidate joint angle vector is Kalman filtered before computing the FK and evaluating the objective; only the final solution is used in the Kalman state update. Including the Kalman filtering inside the objective encourages similarity between the converged and physically-realized tasks without limiting the global joint angle search or delaying sensor feedback in the outer control loop. This solution was sufficient for the present effort but will be considered more deeply in future work.

That future work may also explore the applicability of the many varieties of redundant inverse kinematic methods to a general interleaved manipulator (see [84] for an overview). The application of these methods is difficult due to their assumption that the FK are accurate, which is not generally the case in continuum manipulators (note the discrepancy in desired and integrated task position in figure C.3 below). The incremental joint updates do not always or often result in the expected gains, which can lead to chatter and impair real-world convergence. Moreover, the class of Jacobian-inverse solvers are all simple gradient descent optimizers [84,85], with their slow convergence explained by the general lack of using the Wolfe conditions (sufficient decrease and curvature) to compute the necessary step size, while their local nature guarantees they get trapped in local minima [86]

¹This is a simple two-state polynomial Kalman filter applied separately to each joint angle [82]. Its simplicity allowed the Kalman gains to be computed off-line from the expected variance. The variance can be used to introduce a preference between joints, as it controls the retarding influence on the filter result.

(few papers in the literature detail the entire IK implementation, challenging replication). The explicit use of task Jacobians in various task priority contexts limits the consideration of more task DOFs than manipulator DOFs and, while adding sub-tasks orthogonal to the primary task ensures that there is a global minimum, if operating in free-space there may not be obvious sub-tasks to specify the remaining DOFs [85]. In the presence of FK errors and the absence of IK methods robust to these errors, an interleaved manipulator is better served by a more powerful and general optimization method that avoids these difficulties.

C.2.1 Tracking a Virtual Target

In evaluating this interleaved prototype design and controller, we are chiefly interested in how the controller uses the rigid and flexible joints to minimize error over shifting redundancies. Figures C.2 and C.3 track a virtual circle under increasing the controller gain. The experiment begins in the left panel of figure C.2 by converging from the manipulator start position (approximately [910,-280,91]mm) to the first point in the virtual ellipse at [1050,140,-24]mm. Once the position error norm was below approximately 2mm the virtual circle target began moving through 7000 points about a 200mm diameter circle inclined 30° from horizontal (YZ) at 0.1s/point or 0.9mm/s. Going this slowly avoids excitation of the catheter while helping to clearly show controller-induced errors. No noticeable improvement was observed from independently varying the x/y/z gains on the task error vector in figure 4.3, so $K_i = K_{ix} = K_{iy} = K_{iz} = 0.06$ (red), 0.07 (blue), and 0.09[-] (black) while the joint Kalman filters were adjusted for best performance at each gain.

Consider the gradual convergence in figure C.2. An integral controller on a linear plant will generally demonstrate gradual, asymptotic convergence to the command. Increasing the integral gain will increase the convergence rate until it grows oscillatory. Immediately we can see in the left column that the converging angular error is nonlinear and varies with gain, developing some

apparently oscillatory motion when $K_i = 0.09$. Angular error is the arctangent of the Euclidean error between the target and tip positions and the tip-to-target distance (this is also the fifth, virtual joint position). The angular error does not suggest instability, but notice how increasing the gain led the IK to prefer different base and distal roll orientations than the lesser tunings. Prior to instability we observe divergence in the physical and integrated IK tasks – the error in the IK's forward kinematics is so large that the new joint commands do not reduce the physical error. This also complicates the typical, incremental tuning strategy employed on most controllers as the 0.09 tuning is now operating in a substantially different configuration than the prior tunings, frustrating their comparison. The stability of large gains will also vary across the workspace, a form of input coupling that can be challenging to decouple on a redundant system. The tuning is further complicated by the varying FK across configurations, leading to a tuning that is limited by the most unruly configuration in the trajectory.

The right panel in figure C.2 shows the joint positions and angular error while tracking the circle in figure C.3. As in the initial convergence, the 0.09 tuning attains different configurations than the 0.06 and 0.07. The 0.09's larger gain moves the integrated tip across the xyz workspace more readily, and in doing so switches between optimal configurations more quickly than the lesser tunings. This switching leads to greater error, as the controller is moving the integrated tip, and hence joint solution, faster than the manipulator can move the physical tip. Where the 0.09 initial convergence was not highly oscillatory, it is nearly unstable around segments of the circle and these are what limit the integral controller tuning.

The right panel in figure C.3 shows the measured and integrated x, y, and z tip positions with time. The parameters of the FK described in Appendix C.1 were found by commanding simultaneous sinusoids on each of the joints over their joint ranges and fitting the FK-estimated tip position

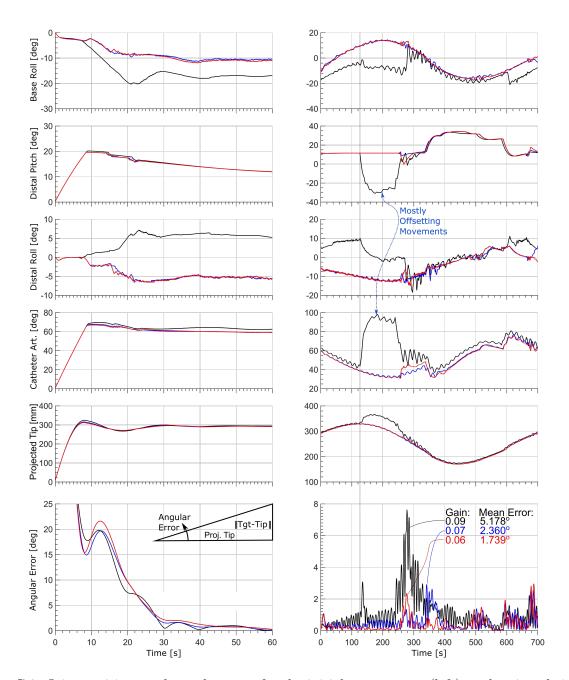


Figure C.2: Joint positions and angular error for the initial convergence (left) to the virtual circle, and tracking of the circle in figure C.3. The color-gain associations given in the lower right graph apply across the figure (and in the next). The initial convergence is aided by a delayed application of the full task error, leading to the smooth curve for the first 7s. The angular error is defined as the arctangent of the target/tip Euclidean error norm and the projected tip distance. Joint limits are $\pm 170^{\circ}$, $\pm 35^{\circ}$, $\pm 60^{\circ}$, $0:280^{\circ}$, and 0:1000mm and are not a concern.

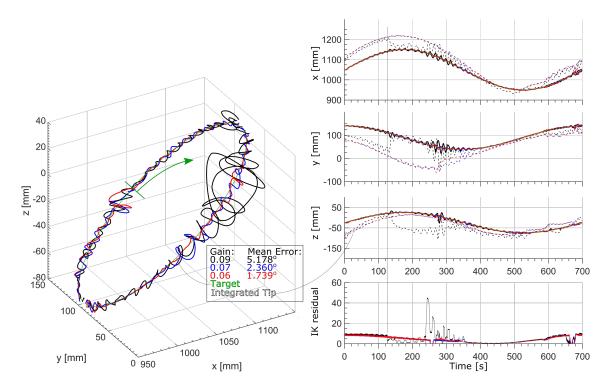


Figure C.3: A rotated, inclined view of the 3D virtual circle and the measured projected tip positions for the three listed gains. The right pane details the xyz positions and IK convergence with time; the associated joint positions and angular error are given in the right pane of figure C.2. Notice also the discrepancies between the integrated (dashed), commanded (green), and measured (black/blue/red) xyz positions which arise from dynamic forward kinematics errors.

against the measured tip position while varying the parameters.² As the FK were implemented in the C/C++ IK library, it was a simple extension to use the same nonlinear optimizer to find the (potentially overfit) FK parameters given in table C.1. The 50-100mm discrepancy between the actual and integrated tip positions evinces substantial errors remaining even after the parameter estimation, but inspection of the integrated tip around t=450s and the good agreement between integrated and measured positions shows that the remaining errors vary with the configuration. The described parameter estimation determined the zeroth order (bias), better agreement requires a higher-order fit; lacking more robust manipulator modeling and at the risk of overfitting, these

²This fitting also reduces the constant bias in the Ascension sensor. Anecdotally, the field is known to vary by a few centimeters across its sensing space. I expect and have observed variations within one centimeter over the dexterous workspace of this prototype – too small and smooth to explain the integrated/commanded discrepancy.

parameters should be estimated on-line.

	measured	fit	
L_{cath}	95	106	mm
0_1R_z	-12	-21	0
${}^{0}_{1}T_{x}$	814	811	mm
$\bar{_1^0}T_y$	-185	-157	mm
$_{1}^{ar{0}}T_{z}^{ar{0}}$	34.7	-8.7	mm
$\frac{\bar{2}}{3}T_x$	8.2	8.2	mm

Table C.1: Measured versus estimated forward kinematic parameters.

The vertical rule in figures C.2 and C.3 at $t \approx 125$ s denotes the onset of an IK solution switch in the 0.09 tuning, as it is followed by large joint movements, momentarily-increased task error, and immediately-decreased IK residual. While the manipulator executes the large joint motions to achieve the solution switch, the task error continues to integrate; depending on the configuration and the time to realize the new solution, the integrated tip may step across another solution boundary and prefer some different, or the previous, solution. Even if the manipulator does achieve the desired solution, the FK errors may prevent the expected task-error minimization from physical realization. In considering convergence to one of multiple solutions, we should ask why there is not a smooth transition between minima, given that there are two free DOFs? This expectation is correct in the absence of joint velocity limits, where even highly-aligned joints (nearly singular) may contribute their remaining, non-overlapping motion to achieving the objective. Joint velocity limits largely prevent this; in the 0.09 joint trajectories of figure C.2 the indicated motions of the distal pitch and catheter articulations are largely offsetting as there is only 10° of distal roll between them. So while the use of a nonlinear, gradient-free optimization avoids issues of ill-conditioned matrix inversion near singularities, the collapsing null-space can still diminish the 5 DOF manipulator to a 4- or 3-DOF, where the 4 DOF solution space may not have a smooth transition between configurations and the 3 DOF only a single minima that shifts with the task coupling.

This same effect is responsible for the notable increases in error in the 0.06 and 0.07 angular errors at t = 280s (0.06) and t = 340s (0.07), and both at t = 680s. In these cases, the smaller gain moves the integrated tip at closer to the same speed as the manipulator, allowing the integrated tip movement to be more closely, and accurately, guided by the real task error. Note that in the experimental data there are no convergence failures (figure C.2, lower right); the drops in the IK residual are merely the realization that a better solution is possible with essentially zero change in the joint angles. This occurs when the error vector is oriented along the current line of motion, such that the integrated tip position is the same as the previous command.

Above all, this simple convergence and circle tracking experiment demonstrates the challenges in tuning the simple controller of figure 4.3, and the substantial importance of accurate FK parameters.

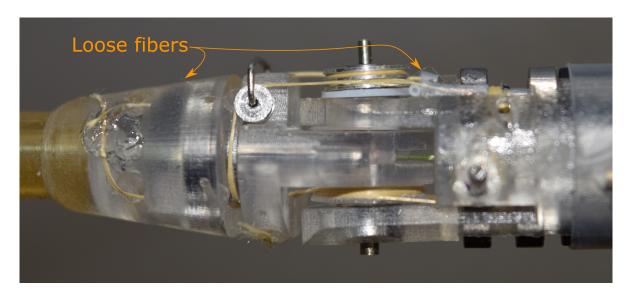


Figure C.4: The loose fibers indicate unraveling of the Kevlar fiber cable, leading to lower joint stiffness.

C.2.2 On the importance of output compliance

The performance of this first prototype is primarily limited by rigid joint cable stretch. Where the flexible shaft windup is minimized by the distal reductions, any stretch in the output cable transmissions is immediately present in the joint actuation. Cable stretch is a function of friction along the cable path, static and dynamic friction in the joints, and the gravity loading of more distal elements. Prior to output cable routing, the joint friction was observed to be very low due to the ball bearings visible in figure 4.1. Friction along the cable path is more difficult to estimate beyond observing that the guiding pulleys (and Teflon conduits on the roll joint) freely rotated during assembly. The gravity loading was reduced by orienting the testbed vertically (see the arrow in figure 4.1) but this has minimal effect in configurations away from the neutral axis.

More important is cable creep, by which the unwoven, multi-stranded Kevlar cables lengthened through fiber movement. This is visible in figure C.4 in the form of loose fibers preceding the transfer and redirection pulleys. A loaded, multi-strand cable is wound such that an increasing axial load draws the fibers into greater mutual contact, with the resulting inter-fiber friction encouraging cable cohesion across pulleys. At this scale loose fiber ends should be expected, but their occurrence near pulleys indicates that fibers can escape the cable (likely during unloading). If fibers can escape, then there can also be axial movement of the fibers, as the same forces that encourage fiber ends to remain within the bundle also limit their axial movement. Under tensile loads, the fibers move to lengthen the cable and reduce the cable strain. The cable lengths or resulting cable/joint compliance were not measured during assembly but, anecdotally, the rigid joints exhibit greater compliance than when assembled. As the pulley grooves determine the maximum cable diameter, the pulleys may be replaced in the future with larger-grooved versions for use with single-strand polyimide cables.

D. Prior Publications

- [1] Benjamin L Conrad, John C Springmann, Lisa A McGill, and Timothy A Shedd. Effectiveness of linear spray cooling in microgravity. In *Space, Propulsion and Energy Sciences International Forum*, pages 67–72, 2009.
- [2] B.L. Conrad and S.T. Sanders. Delayed pulse referencing for cancellation of polychromatic beating noise in grating spectrometry. *Applied Physics B*, 96(4):715–718, 2009.
- [3] Benjamin L. Conrad and Scott T. Sanders. Hrry-spectral and related terms a refined nomenclature for high repetition rate hyperspectral systems. *Journal of Optics and Laser Technology*, 46:134–138, 2012.
- [4] Benjamin L Conrad, Jinwoo Jung, Ryan S Penning, and Michael R Zinn. Interleaved continuum-rigid manipulation: An augmented approach for robotic minimally-invasive flexible catheter-based procedures. In *Robotics and Automation (ICRA)*, 2013 IEEE International Conference on, pages 718–724. IEEE, 2013.
- [5] Benjamin L Conrad and Michael R Zinn. Interleaved continuum-rigid manipulation approach: Development and functional evaluation of a clinical scale manipulator. In *Intelligent Robots and Systems (IROS 2014)*, 2014 IEEE/RSJ International Conference on, pages 4290–4296. IEEE, IEEE, September 2014.

- [6] Benjamin Conrad and Michael Zinn. Closed loop task space control of an interleaved continuumrigid manipulator. In Robotics and Automation (ICRA), 2015 IEEE International Conference on, pages 1743–1750. IEEE, 2015.
- [7] Benjamin Conrad and Michael Zinn. Interleaved Continuum-Rigid Manipulation: An Approach to Increase the Capability of Minimally Invasive Surgical Systems. In *Mechatronics*, *IEEE/ASME Transactions on*, forthcoming.