HAPTIC EMULATION OF HARD SURFACES WITH APPLICATIONS TO
ORTHOPAEDIC SURGERY

by

NIKOLAI ANTHONY HUNGR

B.A.Sc., The University of British Columbia, 2003

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE

in

THE FACULTY OF GRADUATE STUDIES

(Mechanical Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

March 2008

ABSTRACT

A generally accepted goal in orthopaedic surgery today is to maximize conservation of tissue and reduce tissue damage. Bone-conserving implants have bone-mating surfaces that reproduce the natural curvature of bone structures, requiring less bone removal. No small, reliable, inexpensive and universal bone sculpting technique currently exists, however, that can both create and accurately align such complex surfaces. The goal of this thesis was to develop a haptic hard surface emulation mechanism that could be applied to curvilinear bone sculpting using a surgical robot. A novel dynamic physical constraint concept was developed that is able to emulate realistic hard constraints, smooth surface following, and realistic surface rigidity, while allowing complete freedom of motion away from the constraints. The concept was verified through the construction of a two-link manipulator prototype. Tests were run on nine users that involved each user tracing out five different virtual surfaces on a drawing surface using the prototype. The primary purposes of prototype testing were to obtain subjective data on how effectively the dynamic physical constraint concept simulates simple surfaces, to assess how it reacts to typical user interactions and to identify any unexpected behaviour. Users were 100% satisfied with the prototype's ability to emulate realistic and stiff hard surfaces and with its ease of manipulation. The amount of incursion into each of the virtual surfaces by all the users was measured to assess the precision of the system with the goal of deciding whether this new haptic concept should be further developed specifically for precision applications such as surgery. For curvilinear surfaces, 90% of the cumulative distribution of the measured data was less than 2mm, while for linear surfaces it was less than 6mm. Four behavioural effects were noticed: lateral deflection, reverse 'stickiness', hysteresis and instability in certain areas. These effects were studied in detail to determine how to either eliminate them or to minimize them through system design optimization. A computer simulation was also used to model the behaviour of the prototype and to gain further understanding of these effects. These analyses showed that the concept can be successfully used in curvilinear bone sculpting.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGEMENTS

I would like to extend my deepest thanks to my supervisors Antony Hodgson and Christopher Plaskos who guided me, taught me and gave me a glimpse into the exciting world of surgical robotics. I would also like to thank them for the countless hours spent discussing on Skype across two continents.

I am also very grateful to Philippe Cinquin and Jocelyne Troccaz for their warm welcome and full support during my stay in France. It was a privilege to be a part of the working and social environment of the GMCAO-TIMC lab.

A special thank you to Anne-Thérèse Bourreau for her unfaltering help and support and her patience in answering my nagging questions when things got hard. Thanks also to Nabil Zemiti for his essential help in the latter stages of my thesis.

Thank you to NSERC, ICICS and Praxim for providing the funding for my work and for giving me this fantastic opportunity.

I would also like to thank my parents for having gotten me into this mess in the first place and for their continued support, even with 8000 km between us.

And finally, and most importantly, I would like to thank Pilar González Pérez for her unfaltering support and encouragement, and for inspiring in me some of her spirit and energy to lead this project to where it has come.

1.0 HAPTIC EMULATION OF A HARD SURFACE WITH APPLICATIONS TO ORTHOPAEDIC SURGERY

The work presented in this thesis stems from a previous Doctoral project completed in 2005 by Christopher Plaskos (Plaskos, 2005) at l'Université Joseph Fourier in Grenoble, France. In this project, Plaskos presented a miniature robotic bone cutting guide designed to improve femoral implant positioning and fit in total knee arthroplasty (TKA). The robot, named Praxiteles, is mounted on the medial side of the femur. It positions a swinging milling guide to allow the surgeon to accurately mill in the five respective planes of a conventional distal femoral implant. The benefits of this system over manual techniques and existing robotic techniques include increased cutting precision, reduced risk due to retention of manual control of the milling process itself, its small size and a reduction in number of instruments needed in the operating room. The system is currently being commercialized by a French company called Praxim, with the first clinical trials having taken place in December 2007.

With the growing trend within the surgical community towards tissue conservation, Praxim has expressed interest in expanding the Praxiteles system's capabilities to more general bone sculpting. Specific future goals would be to use the system for curvilinear, bone-conserving TKA implants and even in other surgical procedures, such as hip arthroplasty. To realize these goals, while retaining the benefits of the current Praxiteles robot, it is necessary to determine a suitable method of emulating a curvilinear surface.

This Masters thesis focuses on the development of a novel haptic method of emulating a hard, curvilinear surface, and its application to the current Praxiteles robot in the context of bone conserving implants. Chapter 1 provides background information on hard surface emulation and outlines the ideal properties towards which this class of haptic systems strives.

1.1   Introduction to hard surface emulation

In the context of this thesis, hard surface emulation is the ability of a manipulator to simulate a rigid virtual surface of relatively arbitrary shape. As described by Rosenburg (1992) in his paper on virtual fixtures, the use of such a surface implies projecting a virtual interface into a separate environment or workspace. As such, a user is able to move freely within this workspace until they come into contact with the interface. The most basic application of this concept is the use of a physical template, such as a ruler on a piece of paper. A pen can trace any path on the paper workspace only until it comes in contact with the ruler. Three important properties can be extracted from this example that play a major role in hard surface emulation or haptic force feedback systems (Salisbury et al., 1995):

- when not in contact with the surface, the user has complete freedom of motion;

- the stiffness of the surface is, for all practical purposes, infinitely high, regardless of the speed or direction of the user's motion;

- the surface can be traced smoothly and exactly.

To date, virtually all haptic research on stiff wall emulation has focused on impedance or admittance-generating algorithms, which have control schematics similar to the generalized representation presented by Hayward and MacLean (2007) and reproduced here in Figure 1-1. These algorithms are used to determine the forces or displacements required by the haptic architecture to emulate the virtual environment (ex. Colgate et al., 1993; Massie and Salisbury, 1994; Salisbury et al., 1995; Salcudean and Vlaar, 1997; Mark et al., 1996; Kuchenbecker et al., 2006). Hard surfaces are typically approximated as a spring of given stiffness and hence require actual penetration of the virtual surface to activate the restoring forces. With current haptic devices, however, it is difficult to provide sufficient stiffness to convincingly recreate hard surface contacts without inducing contact instability (Salisbury et al., 1995; Cavusoglu et al., 2002; Kuchenbecker

et al., 2006). Maximum stiffness values for common haptic devices are listed in Table 1-1. Lawrence and Chapel (1994) report that stiffnesses greater than 10kNm^{-1} would be difficult for humans to differentiate between, which shows that a large number of haptic devices are far from emulating truly realistic stiff walls. In comparison, true hard collision stiffnesses can be as high as 1000 kNm^{-1} (Kuchenbecker et al., 2006; Hayward and MacLean, 2007).
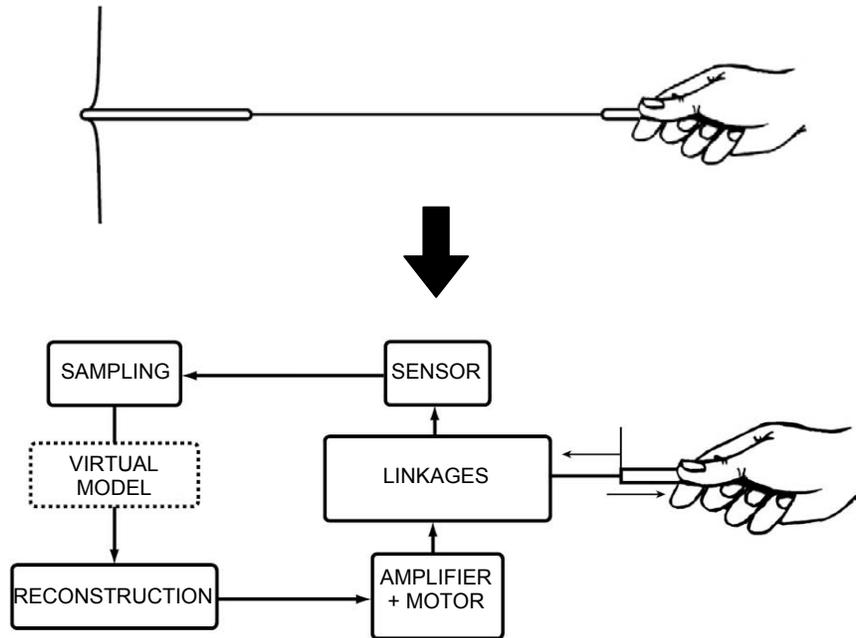
Figure 1-1: *Generalized control schematic of a typical haptic force feedback system, where the virtual model determines the impedance required by the system in order to simulate the virtual environment (from Hayward and MacLean, 2007).*

All the force feedback systems mentioned above are used in the context of teleoperation. In other words, the template (the haptic environment) is separated in space from the workspace (the virtual and/or physical environments). The user, in fact, only interacts physically with the haptic architecture. Hence, it is sufficient for the device to provide a "sensation" of stiffness, which is distinct from providing true precision within the haptic environment itself.

Haptic Device		Maximum Stiffness (kNm^{-1})
Haption Virtuose		2
MPB Freedom		2
Sensable Phantom		2
Sensable Phantom with custom amplifiers		2.7
Mimic Mantis		5.5

Force Dimension Omega		14.5
Moog FSC Robotics HapticMASTER		50
Minimum stiffness required for realistic perception of hardness (Lawrence and Chapel, 1994)		10
True metal-metal contact stiffness (Kuchenbecker et al., 2006; Hayward and MacLean, 2007)		1000

Table 1-1: *Summary of commercially available haptic devices with their respective published maximum stiffnesses.*

In some applications, most notably computer-assisted surgical procedures, there are good reasons for wanting to superimpose the two environments (the haptic and the physical) so as to, for example, ensure that a surgeon does not accidentally move into a dangerous region while handling a cutting tool. This superimposition, however, introduces additional difficulties. In such a situation, it is essential that the haptic device not only

provides a convincing feel during contact with a virtual surface, but it must accomplish it with sufficient precision while simultaneously causing minimal interference with the task due to its presence within the physical workspace. The haptic device must also physically interact with the workspace, so it must be able to distinguish endpoint interactions from user interactions.

Much less research has been carried out in this latter aspect of hard surface emulation. Its importance is significant, however, and will be demonstrated in the following section through examples. The main results of previous research in this domain will be described in Section 1.3. The advantages and drawbacks of the systems developed by this research will then be used to form a list of design constraints for the development of an ideal hard surface emulation device.

1.2   *Applications of hard surface emulation*

Hard surface emulation in the context of assisting users to perform precision motion control tasks has a wide array of potential applications, ranging from large industrial part handling tasks to surgical procedures. Humans are not endowed with the high repeatability, precision or stability of robots. They are, however much better adapted to decision making and strategic planning in variable environments and in controlling physical interactions, such as those involved in using various tools. Haptic interfaces can be used to merge these distinct abilities.

Beginning with the most simple application, a haptic interface can be used to provide directional guidance to a human user. In other words, it could restrict a user's motion along a specific path, allowing him or her to push or pull an object without being concerned with steering it. An industrial example of this is the installation of automobile

doors on an assembly line, as described in Paula (1997). The advantage of haptic guidance in this case specifically benefits the health of the user by preventing unnecessary physical strain arising from the need to steer the heavy door. Another example, on the other size extreme, is needle guidance during pericardial puncture procedures (Schneider and Troccaz, 2001). Here, the haptic interface is used to precisely guide a needle along a predetermined trajectory so as to prevent accidental damage of the heart during the procedure. A similar surgical assistance system can be envisaged for drill guidance in pedicle screw placement in the spine.

A more complex application of hard surface emulation is when it is used to limit the user's motion to a virtual three-dimensional region. The system allows the user to move freely within the region until he or she comes in contact with the boundaries. As an extension to the industrial application mentioned above, such a system would be useful to avoid collisions when handling fragile parts within a constrained environment. In a surgical setting, it could be used to prevent a tool from damaging soft tissue surrounding the operating environment. More specifically, within the domain of orthopaedic surgery, significant research has been carried out in haptic interface design for bone sculpting applications (Ho et al., 1995; Shoham et al., 2003; Brisson et al., 2004; Roche, 2006). In this case, the haptic device is used to provide cutting precision in a manner analogous to using a physical template, while still allowing the surgeon freedom of motion along the surface of the cut.

The industrial and surgical examples presented above show the relevance of hard surface emulation. The techniques developed to date still have significant drawbacks and have therefore not found wide acceptance in the surgical domain. The following section will describe the most important research in hard surface emulation for user assistance and motion control and will outline the advantages and drawbacks of various previously-developed systems.

1.3 Existing hard surface emulation techniques

Haptic interface design consists of two distinct aspects that are to some extent mutually independent: the mechanical architecture and the control method used.

There are many existing mechanical architectures which are strongly related to those used for robots: serial manipulators, parallel manipulators and cable-driven designs. What typically differentiates haptic interfaces from robots, however, is how the joints are controlled: robot joints are almost always driven, whereas haptic joints can be either driven or restrained. As mentioned in Section 2.1, the majority of existing haptic designs have been based on the concept of generating the desired impedance through controlling motor torque. The typical concept is to use sensors to determine when penetration of a virtual surface has occurred and to respond by applying a restoring force through the joint motors to push the user back out. Due to hardware and software limitations, such as response lag time, joint backlash, structural flex, sensor noise, etc, systems based on this concept are not capable of rendering truly hard surfaces or of handling large or sustained user forces without causing instability or lack of precision (Massie and Salisbury, 1994; Kuchenbecker et al., 2006; Cavusoglu et al., 2002; Hayward and MacLean, 2007) These requirements are, however, essential if we wish to use haptic force feedback in guidance and region-restriction tasks. This need has led to the development of a variety of new concepts for emulating hard surfaces based both on modified control algorithms as well as on new mechanical concepts.

1.3.1   *Active Constraint Concept (Acrobot)*

Perhaps the most successful research, as far as commercial development is concerned, is the active constraint concept developed by Harris et al. (2004) and described in Ho et al. (1995). It has been applied in the form of a haptic robot, named Acrobot, and used in cutting tool guidance during total and unicompartmental knee arthroplasties. The concept is illustrated in the block diagram shown in Figure 1-2. The inputs to the algorithm are the user's applied force as a vector and the current location of the robot. The Boundary Controller uses this data to adjust the proportional and derivative gains of the position controller (i.e. impedance). It also adjusts the robot's desired position based on its current location. The principle behind the concept is illustrated in Figure 1-3, which shows three separate regions in space. Region III is the unsafe region, into which the tool must not penetrate. Region I is the safe region, in which the user may move at will. Region II is a region of increasing robot stiffness at the boundary of the two other regions. A more detailed description of how the algorithm functions follows:
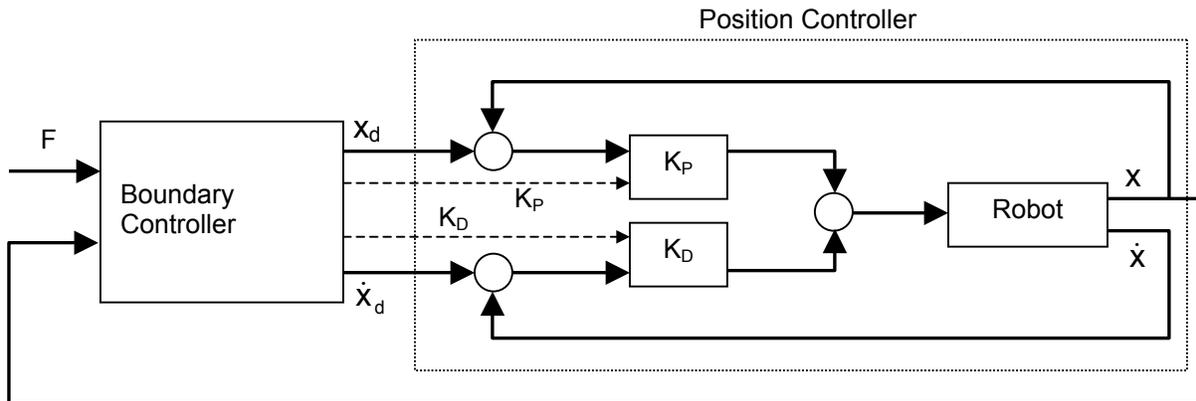
Figure 1-2: *Generalized block diagram representing the active constraint concept used in the Acrobot System (The Acrobot Co. Ltd., London, UK) (Harris et al, 2004).*
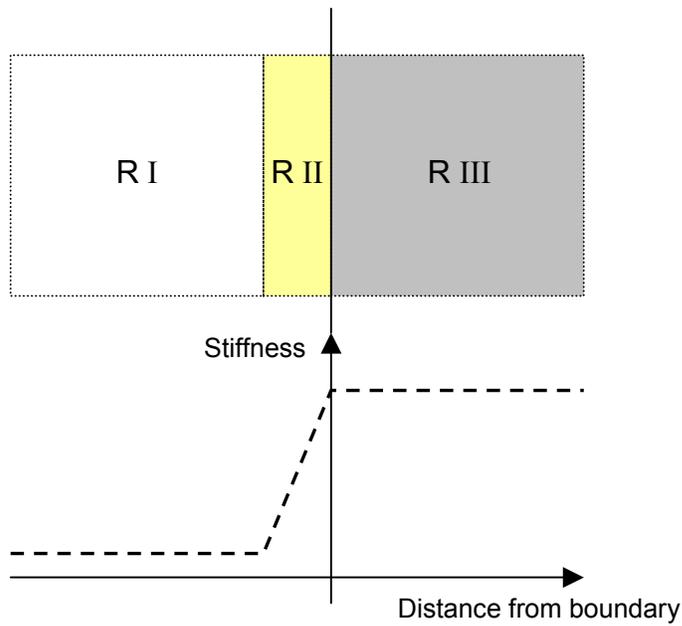
Figure 1-3: *Active constraint principle showing the three regions of distinct robot stiffness (Harris et al, 2004).*

- Region I: the robot impedance (K_P and K_D) is low, making stiffness and damping low and admittance high, to increase the sensitivity of the robot. The position demand is set to the current position.

- Region II: the robot impedance increases and the admittance decreases based on the current location of the tool with respect to the unsafe boundary. Position demand is still set to the current position.

- Region III: the impedance is set very high, while the position demand is set to the nearest point on the boundary in order to push the robot out of the unsafe region.

The purpose of Region II is to provide a smoother transition between the free motion region and the restricted region, thus preventing instability and decreasing the possibility of surface penetration due to delays in the control loop. The drawbacks are springiness at the boundary, vibrating motion at an inclined boundary and restricted motion along the boundary due to the increased impedance in this region (Ho et al., 1995). The latter issue

10

is resolved by applying stiffness only in a direction perpendicular to the surface. In a more practical manner, the stiffness of the robot is based on its distance from the surface along the line of action of the user's force (see Figure 1-4).
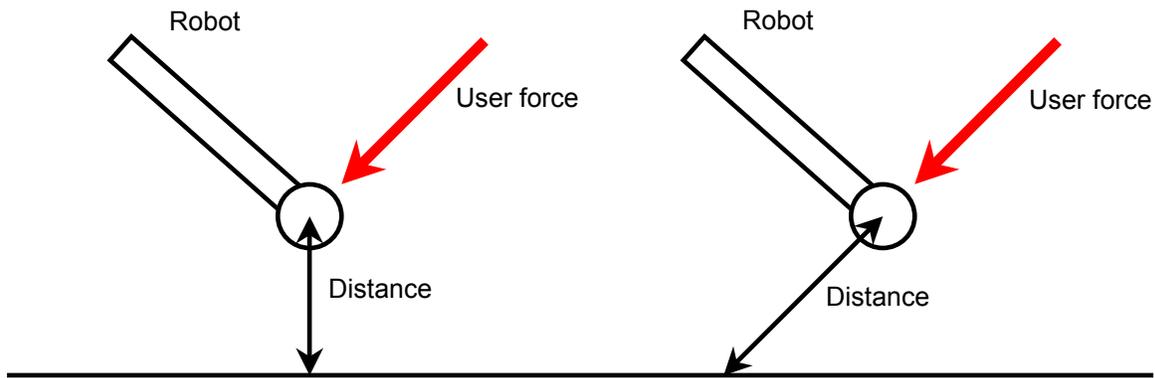
Figure 1-4: *Schematic showing the difference between measuring the distance of the Acrobot's tool tip perpendicular to the surface (left) and along the line of action of the user's force (right). The latter allows motion along the surface since, the smaller the component of the user's force perpendicular to the surface, the less impedance is applied to the robot.*

The true "feel" of the active constraint concept, in terms of hard surface emulation, was not reported in any of the publications found on the subject. However the several documented clinical trials (Jakopec et al., 2001; Jakopec et al., 2003; Rodriguez et al., 2005; Cobb et al., 2006; Davies et al., 2007) suggest that the feel is satisfactory. The precision of the concept is also reported to be successful in the case of planar and cylindrical surfaces (Cobb et al., 2006; Davies et al., 2007). Note that the term "successful" is situation-dependent, and that no numerical indices have been presented to quantify the precision of the technique. An estimate of precision, however, can be made based on the reported 2° leg-alignment precision of the system in unicompartmental knee arthroplasty (Rodriguez et al., 2005; Cobb et al., 2006). Assuming a medial-lateral knee width of 60mm and a worst-case scenario of rotation about the lateral or medial side of the knee, the system's maximum error would be on the order of 2mm. Some practical drawbacks to the design are that a force transducer is required on the interface between

the user and the device. Additionally, the structural architecture and motors must be able to provide sufficient impedance to the user, requiring large parts. This also creates significant friction in the system, requiring motion assistance from the robot to emulate uninhibited motion in the free region (Ho et al., 1995; Harris et al., 2004). All of this results in a relatively costly robot.

1.3.2 *Continually Variable Transmission (Cobot)*

In order to implement a truly 'hard' constraint, different concepts have been tried that do not involve using drive motors to oppose motion across the constraint surface. One of the most promising of these has been called a 'Cobot' and is based on the principle of orienting a friction constraint normal to the constraint surface. Cobots are based on the continually variable transmission (CVT) concept that was first applied to haptic interfaces by Colgate and Peshkin (1999). A CVT device is strictly defined as one having a continuous range of transmission ratios, independent of the amount of torque being applied to it (Kim et al., 2002). The most simple form of haptic CVT is a unicycle, as described in Colgate et al. (1996). A unicycle has two degrees of freedom: linear motion parallel to its wheel and rotational motion about its vertical axis. This means it can move backwards or forwards, while being steered in any direction. The haptic concept involves the user pushing the unicycle in any direction, while a motor steers it in the controlled direction. To implement a surface or path constraint, the motor simply steers the wheel in a direction tangent to the surface. Figure 1-5 shows three haptic interface prototypes based on CVT's in increasing order of degrees of freedom (DOF). The 6 DOF prototype is controlled by spherical CVT's, as shown in Figure 1-5d. The steering rollers, R1 and R2, define the centre sphere's axis of rotation A. Depending on the position of this axis, the output shafts, S1 and S2, rotate at independent velocities. In the case shown in the figure, A is nearly normal to S1 and parallel to S2. Thus, S1 would rotate very slowly while S2 would rotate quickly. The relationship between these angular velocities can be taken

<div align="center">(a)</div>
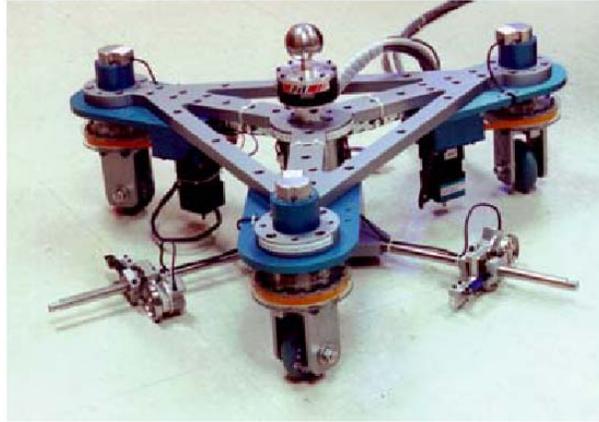
<div align="center">(b)</div>

<div align="center">(c)</div>
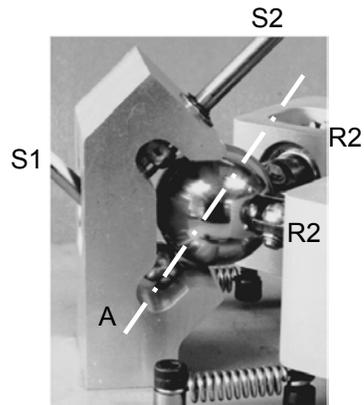
<div align="center">(d)</div>

Figure 1-5: *Prototypes of haptic interfaces based on continuously variable transmissions. (a) Two degree of freedom unicycle (Peshkin et al, 2001). (b) Three degree of freedom tricycle (Peshkin et al, 2001). (c) Manipulator with three revolute joints each powered at its base by a spherical CVT, as shown in (d) (Peshkin and Colgate, 1999).*

advantage of to create rotational motion between two arms of a manipulator (Peshkin et al., 1996).

The primary advantage of this concept is that it relies on real frictional constraint rather than on motor impedances to create surface constraints. This eliminates all sense of "springiness" at constraint boundaries, while within safe regions, motion is completely unconstrained (if force feedback is incorporated). The concept can also be used both for path constraints and for region constraints. The drawbacks, however, are that the concept necessarily requires force sensors to keep track of user intentions. Depending on the task, it can also result in rather bulky architectures with large amounts of inertia. More importantly, the inherent characteristic of the design in which the wheel steers continuously rather than in discrete steps, causes specific undesirable behaviours in two situations: when the user pushes the device from rest and when the device approaches a boundary at a high angle. In the first case, the device "hesitates" while the wheel completes its steering operation before moving in the user's direction (Colgate et al., 1996). In the second case, the device actually penetrates the surface and steers the user off to the side while the wheel once again attempts to complete its steering operation (Colgate et al., 1996; Moore, 2003). The specific amount of penetration normalized to the approach speed was not described. However, Colgate et al. (1996) report a penetration of 10mm, while Moore et al. (2003) report penetration of up to about 40mm, both resulting from normal human-activated approach speeds.

1.3.3   *PADyC Freewheel Mechanism*

Similarly to the continuously variable transmission concept, the PADyC freewheel mechanism is mechanically based. The concept, developed by Troccaz and Delnondedieu (1996) is based on a double freewheel and motor combination that allows passive motion within a set of dynamic constraints. A generalized diagram of the concept

14

is shown in Figure 1-6. The diagram shows a cross-section of two shafts (red and gray), separated by a freewheel. The outer red shaft is able to rotate freely in the counter-clockwise direction. However, it is constrained by the freewheel to the rotational speed of the gray shaft in the clockwise direction. By swapping the roles of the shafts, the reverse is true as well. This concept therefore allows the control of relative motion between two serial manipulator arms, each connected to one of the shafts.

It is evident that the constraints are enforced mechanically. Motors are only used to determine the rotational velocity of the constraints and have no ability to actually drive the mechanism. This passive design makes it inherently safer than active concepts (Schneider et al., 2000; Schneider and Troccaz, 2001; Troccaz and Delnondedieu, 1996). Two reported drawbacks with the design include low stiffness of the system (which may have been caused by the particular system used for testing, rather than by the freewheel concept as such: Schneider and Troccaz, 2001), and jagged motion in certain regions during path or surface following (Delnondedieu, 1997; Colgate et al., 1996). Specific data on the amount of penetration of the virtual surface was once again not presented. However, based on measurements of mechanical flex and hardware response times found in Delnondedieu, 1997, we can make an approximate estimate of 4mm.
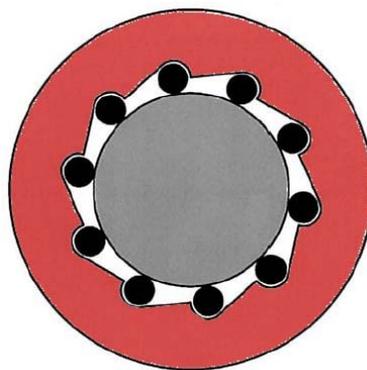
Figure 1-6: *PADyC freewheel concept (Troccaz, 1996).*

1.3.4   PTER Dissipative Passive Concept

The last concept to be described in detail in this chapter is the dissipative passive design developed at the Georgia Institute of Technology (Book et al., 1996). This concept involves the use of clutches and brakes to regulate the relative rotational velocity between two manipulator links. A manipulator based on this scheme is shown in Figure 1-7, while Figure 1-8 illustrates the general principle. Activating the clutch in the direct joint coupling mechanism couples motion of the two links in the same direction. Activating the clutch in the inverting joint coupling mechanism couples motion of the two links in opposite directions. The brakes on each side can be used to stop rotation of either link. The various control methods that can be used to manipulate this system in path-following applications are described in Book et al. (1996), Davis and Book (1997) and Swanson and Book (2003).
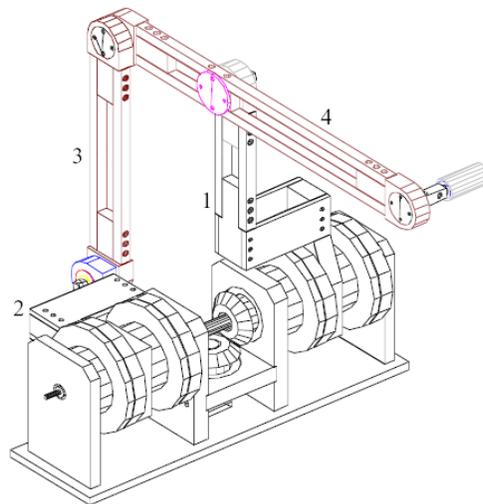
Figure 1-7: *PTER dissipative passive device. Links 1 and 2 are connected to by a coincident axis and their relative motion is controlled by the clutch and brake system at its base (Book et al., 1996).*
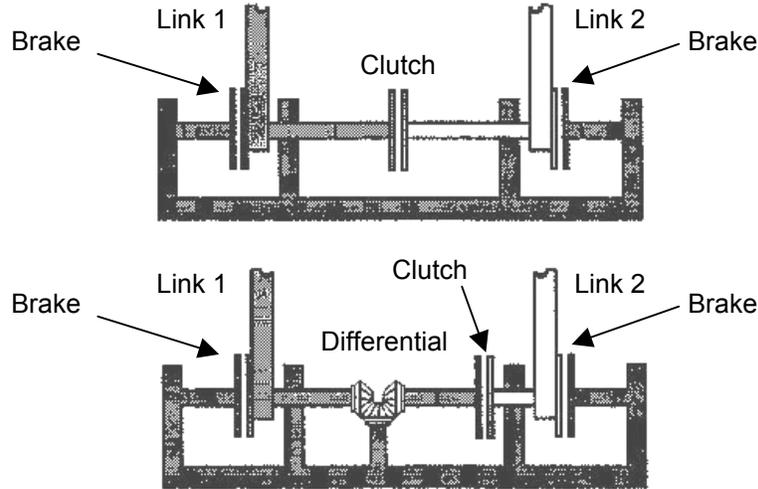
Figure 1-8: *Direct (top) and inverting (bottom) joint coupling mechanisms, forming the basis of the PTER dissipative passive concept (Book et al., 1996).*

Once again, the passivity of this concept is an advantage for safety purposes. No analyses of the "feel" of the device in hard surface emulation were found. However, it can be deduced from the fact that brakes and clutches can provide high frictional constraints over very short periods that such a device could present a convincing impression of a hard virtual surface. Results in Swanson and Book (2003) show promise in surface emulation tasks as far as precision is concerned; however, depth of penetration is still an issue. Path-following also seems to be a challenge, as evidenced in Book et al. (1996) and Swanson and Book (2003), with surface penetration errors of up to 10mm. The feel of the system in path-following tasks was not reported, however, from the path plots shown in Book et al. (1996), smoothness may be an issue.

1.4 Design objectives for an ideal hard surface emulation device

The examples presented in the previous section show that research in haptic stiff wall emulation is still an imperfectly developed field. The many innovative concepts succeed to a certain degree, but no concept to date fully emulates truly realistic rigid guiding surfaces. Figure 1-9 shows a plot of how existing techniques, including those described above, succeed in emulating hard collisions while allowing for smooth surface tracing. The region in the upper right corner of the plot remains largely unpopulated and was the goal of this thesis.
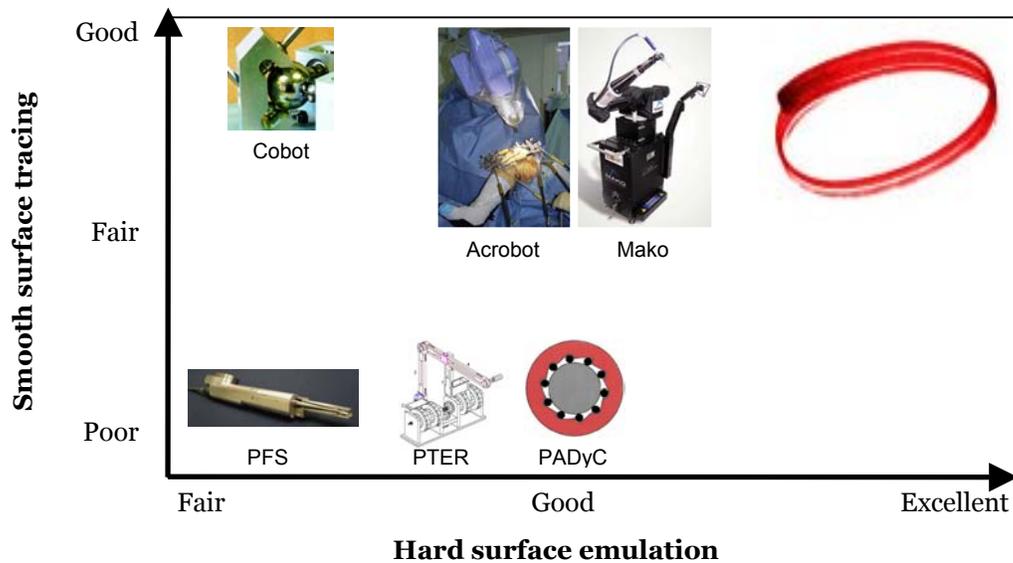
Figure 1-9: *Plot of the distribution of primary existing haptic techniques with respect to their ability to emulate hard collisions while allowing for smooth surface tracing. The red circle marks the goal of this thesis. The Mako Haptic Guidance System and the Precision Freehand Sculptor (PFS) are described in Section 5.4.*

As in any design, it is important to formulate precisely which aspects of device behaviour must be addressed, hence allowing each concept to be developed with specific goals essential to the particular situation encountered in actual use. A list of design constraints

towards which an ideal hard surface emulation device would strive will now be presented. These objectives are based on the results of the above research review, while keeping in mind the basic properties of a real physical template, as described in Section 1.1.

- **Realistic surface collision:** collision with the virtual surface should be as realistic as possible. Hence, minimal or no penetration of the surface should occur upon contact, regardless of the approach velocity and applied force.

- **Realistic surface rigidity:** a constant applied force by the user on the surface should not allow any detectable penetration of the surface or any motion of the surface (i.e. no springiness).

- **Unimpeded surface departure:** the action of pulling away from the surface should not result in any feeling of stickiness or impulse, regardless of the departure acceleration, initial velocity, or initial applied force.

- **Smooth and precise surface tracing:** intents by the user to trace the surface in any direction should result in unimpeded motion, as described in the next objective. No over- or under-penetration of the surface should occur, regardless of the speed of motion and applied force (i.e. no hysteresis or instability).

- **Unimpeded motion freedom away from surface:** when not in contact with the virtual surface, user motion should be completely unimpeded with minimal apparent friction from the device. Gravity effects are situation-dependent and should be considered separately. Ideally, the haptic system should be transparent to the user, as though it did not exist. By strictly adhering to this, gravity effects should not be compensated for, as is the case with any freely-held tool. However, depending on the situation, it may be useful to consider such compensation.

Quantification of these constraints is possible, however, since they are based on the abstract concept of "feeling," it is difficult to specify what values are acceptable in a

given application. Since this concept is individual and situation-dependent, quantification of these constraints will only be developed further in Chapter 5.0 in the specific context of bone sculpting. A generalized hard surface emulation device should therefore be evaluated in comparison to a realistic situation in a qualitative manner. Normalized measurements of penetration, stiffness, stability and friction should be made and then applied to specific situations in which the system could be used, rather than to the general case.

1.5 Overview of following chapters

Chapter 2 describes a novel concept of hard surface emulation which comes close to meeting most of the above ideal design objectives, while introducing some singularities and behavioural particularities inherent to the concept. Chapter 3 evaluates the concept more precisely through the use of a prototype. Chapter 4 presents a simulation designed to understand the behaviours observed in the experiments of Chapter 3. Chapter 5 concludes with a description of how it can be implemented and optimized specifically in the domain of curvilinear bone sculpting for bone conserving distal femoral implants.

2.0 DYNAMIC PHYSICAL CONSTRAINT CONCEPT

The previous chapter outlined a set of design objectives for an ideal hard surface emulation device. I have developed a novel haptic concept that successfully fulfills these objectives. The concept is simple enough to be easily implemented in a variety of manipulator designs. This chapter describes the concept and how it behaves relative to the other dynamic constraint concepts covered in the previous chapter.

2.1 Description of the dynamic physical constraint concept

In developing the idea for the dynamic physical constraint concept, the properties of a true physical template, such as a ruler, were considered as the ideal. Although it is able to emulate true hard surface contact, the problem with such a template is that its shape cannot be modified. This led to the idea of a moving physical constraint that can be adjusted depending on the user's current position in space. I call this idea a 'dynamic physical constraint'.

The simplest form of this concept is a one-dimensional constraint that allows motion in one direction but not in the other. An extension of this, in which the constraint can be used to prevent motion in one direction while allowing it in the perpendicular direction, is shown in Figure 2-1(a). This passive manipulator has one rotational degree of freedom (R configuration). The end of link 1 is constrained to move along a circular path about the fixed centre point. It passively positions itself along the line connecting the user's current position and the centre point, creating a physical barrier which the user collides with when he or she moves towards the centre.

A further extension to this form is shown in Figure 2-1(b). The concept is identical, except a second degree of freedom is added by replacing L1 with a parallel prismatic joint (RP configuration). Given a desired virtual surface shape, the prismatic joint adjusts its length radially, based on the manipulator's position about the centre, to ensure the stopper is always on the surface. In this case, the manipulator is semi-active since the rotational degree of freedom is passive, while the translational degree of freedom is active.
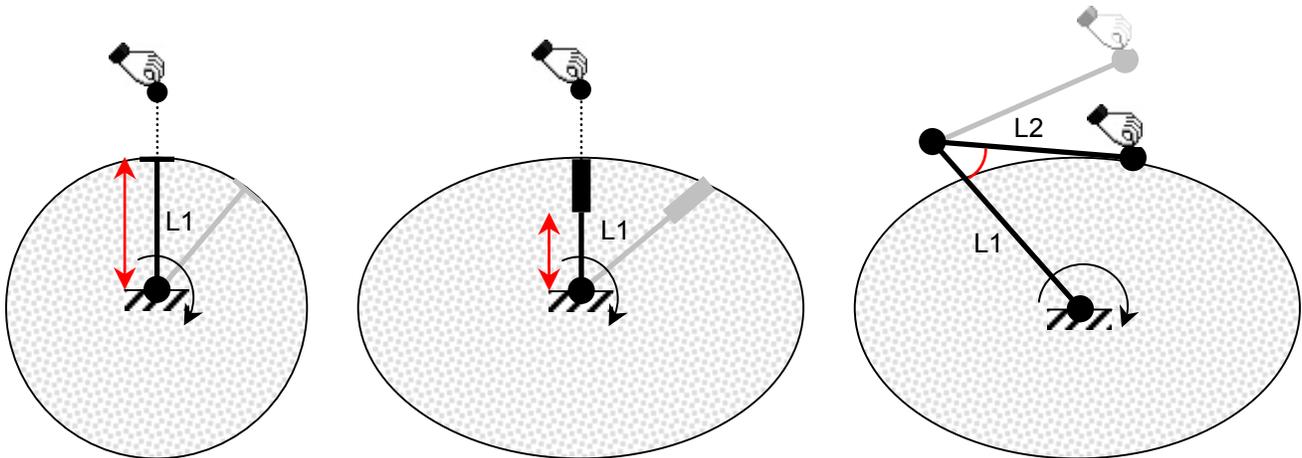
Figure 2-1: *Three different two-dimensional implementations of the dynamic physical constraint concept, from simplest to more complex. (a) L1 is fixed in length and rotates about the attachment point to always remain in line with the user. (b) Same as the previous implementation, except that L1 can vary in length, allowing different constraints at different values of θ. (c) A two-link implementation, where the constraint is located at the elbow, limiting the amount to which the elbow can bend. The red lines depict the respective constrained dimensions.*

A different two-dimensional implementation of the dynamic physical constraint concept is shown in Figure 2-1(c). Once again, the first degree of freedom is revolute about the fixed centre point. The second degree of freedom is also revolute, resulting in a common two-link rotational manipulator (RR configuration). The dynamic physical constraint, in this case, is applied at the manipulator's elbow, allowing free rotation of L2 away from it, but obstructing motion towards it when the free outer link contacts the stopper. In other words, the user can apply scissor-like motion to the manipulator until contact with the angular constraint is achieved. This configuration is semi-active because all degrees of freedom are passive, while the revolute physical constraint actively adjusts itself depending on the current position of L1. A particular advantage of this configuration is that a rotational constraint is very easily implemented using an electric motor. Its workspace is also large and adjustable.

The configurations presented above are by no means exhaustive. PP, PR, and of course three or more degree of freedom configurations can be envisioned. Movable stoppers can also be implemented to prevent movement in the opposite direction, if desired. These additional stoppers can either be actuated independently or from the same motor actuating the primary stopper.

2.2 Comparison to design objectives

This section describes how the dynamic physical constraint concept fulfills the objectives outlined in Section 1.4. The primary accomplishment of this concept is that it simulates realistic hard surface contact because of the true physical contact between the user and the dynamic constraint. When not in contact with the virtual surface, the user can move the manipulator with complete freedom, without the need to backdrive motors. Additionally, no force sensors are necessary to predict the user's intent. The control loop relies solely on simple mechanical position sensing of the manipulator. The concept also allows the user to apply a constant force against the virtual surface without jeopardizing the surface rigidity and stability. Finally, the virtual surface can be traced without stiffness caused by the manipulator and with the same rigidity as upon initial contact. These properties show that the concept successfully fulfills the objectives of an ideal hard surface emulation device.

2.3 Limitations of concept

Although the dynamic physical constraint concept largely satisfies the objectives outlined previously, it can exhibit certain behaviours which deviate from the ideal and which may

RP Configuration	RR Configuration
(a)	(b)
• S: any surface • F⊥S and F→S	
(c)	(d)
• S: any surface where the angle (α) between its tangent and L2 is < 90°. • F: any force that causes a negative moment about (0,0), and that has a component tangent to the surface and towards (0,0).	
	(e)
	• L2⊥S
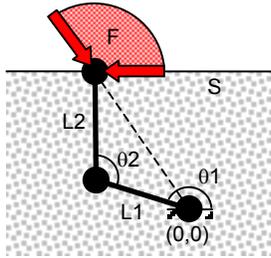
Table 2-1: *Response limitations of the dynamic physical constraint concept in the RP and RR configurations. (a) and (b) A force with a component towards the surface will cause lateral deflection. (c) and (d) The forces as shown will cause penetration of the surface. (e) A force applied in the sector shown will bring the end-effector out of reach of the surface, resulting in penetration. S = the virtual surface, F = the force applied by the user, (0,0) = robot attachment point.*

therefore limit its usefulness under particular conditions. Table 2-1 shows a list of these for the RP and RR configurations described in Section 2.1.

2.3.1 *Lateral deflection*

The first, and most significant behaviour ((a) and (b)) is a property of the concept itself and so manifests itself regardless of the configuration used. In the situations shown in the table, the user would ideally expect the robot to resist his or her force and not allow any movement. However, the applied force F creates a moment about the attachment point and, since $\theta 1$ is free, causes the robot to rotate. As this rotation is detected, the stopper is actively readjusted to limit $\theta 2$ and prevent incursion into the surface; the net effect is that the end effector tracks laterally a small distance along the surface before coming to rest. This lateral deflection effect grows in strength the more parallel the surface's tangent is to the line connecting the end effector and the attachment point, and the greater the moment arm of the applied force is.

The amount of lateral deflection is dependent on the user's interaction with the mechanism. If the user's applied force remains vertical throughout the deflection, then the end effector will continue to slide laterally. If the user exhibits lateral stiffness, however, a horizontal force component will build as the end effector is deflected, causing the deflective motion to stop at a point of equilibrium. This effect is called the equilibrium point effect will be described further in Section 4.2.1.

2.3.2 *Inability to create constraint (generic)*

The second limitation inherent to both configurations is shown in (c) and (d). In this case, the constraint would not be able to prevent penetration into the surface, since the constraint only applies radially towards the attachment point (an inverse constraint could

be envisioned to counter this issue). Additionally, in case (d), although the user wants to move down the surface, he or she is restricted by the physical constraint. This results in a clockwise moment about the centre, causing the constraint to push the end effector radially outwards and further into the surface. In the special case when the attachment point is located on the surface, the user would not be able to move towards the centre.

2.3.3 Inability to create constraint (configuration-dependent)

The third limitation shown in Table 2-1 is specific to the RR configuration and occurs when L2 is perpendicular to the tangent of the constraint surface. Any force that causes a counter-clockwise moment about the centre brings the surface out of reach of the robot. Additionally, in the configuration shown, the end-effector motion which would result from a rotation of link 2 would be parallel to the constraint surface and therefore would be ineffective in enforcing this constraint.

2.4 Design conditions

Table 2-2 shows a set of conditions that are required in the design of the device in order to minimize or avoid the problems described in the previous section. This discussion is restricted to designs where the physical constraint limits movements towards the centre of the mechanism. The inverse of these conditions can be applied if the constraint is used for movements in the opposite direction.
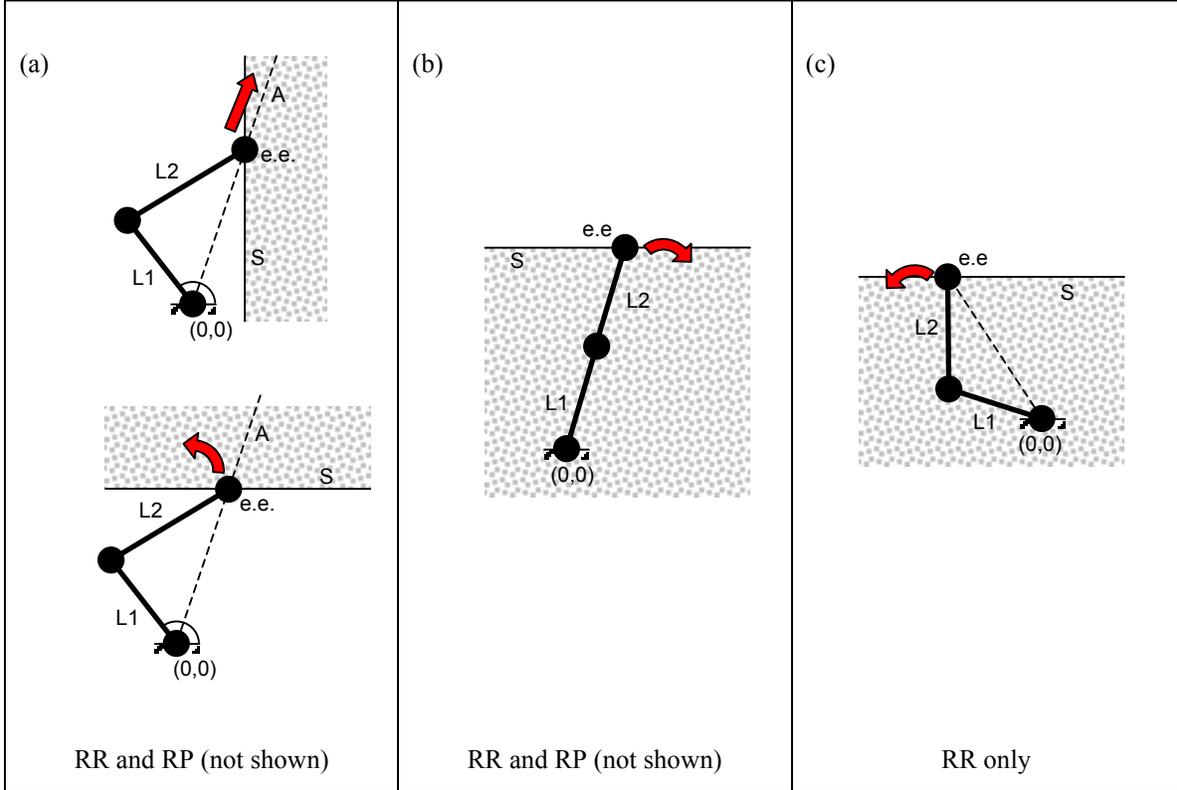
Table 2-2: *Conditions in which the dynamic physical constraint device must not be placed, in order to prevent situations where the constraint cannot be enforced. The RP configuration is not illustrated. e.e. = end effector.*

Outward orientation: The first condition (a) is valid for both the RR and RP configurations and states that the extension of the line (A) connecting the end effector and the robot's attachment point must be directed outwards through the constraint surface. Since A is the direction along which the robot expands, either passive motion or motion caused by the active constraint would cause penetration into the surface.

Surface within joint limits and away from singularities: The second condition (b) is also valid for both configurations and states that the robot must not be placed into a fully extended position while in contact with the surface. This would put parts of the surface out of reach of the robot and, in the case of the singular configuration of the RR design,

would prevent the device from generating an extension that would return the endpoint to the surface, if the endpoint were moved inwards.

Alignment condition: The third condition (c) relates to the problem described in section 3.3.3 and applies only to the RR configuration. Link 2 must never be perpendicular to the surface because a counter-clockwise rotation of the second link can cause motion across the constraint surface that cannot be resisted by repositioning of the dynamic constraint. More generally, this condition specifies that with any free joints assumed to be fixed, movement of the dynamic constraint must be able to produce a component of displacement perpendicular to and directed outwards from the constraint surface in order to be able to be effective in preventing incursion through the surface.

If these design conditions are satisfied, the only issue that a user will have to deal with is the induced lateral deflection. However, convex surfaces that wrap around the centre point would have tangents which are predominantly nearly perpendicular to the radial line connecting the centre and the end effector, so the induced lateral deflections are likely to be small. Even in cases where the induced deflections are larger, it is possible that this effect would be easily learned and anticipated by human operators and hence, easily overcome.

2.5   Conclusions

This chapter has presented a novel dynamic physical constraint concept for haptic hard surface emulation. The concept was described in general terms, and two specific robot architectures were presented that could be used to create a two-dimensional haptic interface. A list of design conditions and guidelines was then developed based on an analysis of the concept's behaviour and limitations, but it is difficult to tell analytically

whether or not a user will be able to use the device effectively and to deal with the lateral deflection issue. To validate and test the overall concept, primarily for its feel and precision, a prototype was built for user tests. This prototype is described and analysed in the following chapter.

3.0 PROTOTYPE DESCRIPTION AND TESTING

Generally, the purpose of building a prototype is to validate a conceptual design. In this project, a prototype was necessary for three particular reasons: to evaluate how the dynamic physical constraint concept feels as a haptic hard surface emulation device; to evaluate how it performs in potentially problematic situations, such as near singularities; and to evaluate what precision can be achieved as a function of movement speed. This chapter describes the prototype in detail and how it was tested.

3.1 Introduction

The prototype was developed by modifying the original Praxiteles robot mentioned in Chapter 1. Praxiteles was used for a number of reasons. The first reason is that this project was initiated by Praxim, who was interested in expanding Praxiteles' capabilities. The convenience of using an existing architecture allowed the project to focus purely on the development of a novel haptic concept rather than on the technicalities of developing a functioning manipulator. I was also able to keep manufacturing complexity and costs low by reusing existing components.

Regardless of these initial reasons, Praxiteles, in fact, has particular advantages for this project. It is designed as a modular robot: its first output axle can be easily mounted to any surface, while a variety of extensions can be plugged onto its second output axel. It is driven by brushless DC servomotors coupled with backdriveable harmonic drives. The motors are supplied with Hall sensors, which enable the angular position of each axis to be directly measured. It is also small and simple, which is ideal for the eventual goal of using the system in computer assisted surgery.

The original Praxiteles is shown in Figure 3-1 mounted to the distal femur during a cadaver trial. It has two revolute joints, labelled Axis 1 and Axis 2. The robot is attached to the bone at its first axis using a fixation piece and two bicortical pins. A rotational milling guide is attached to its second axis. This extension holds the hand mill, allowing axial and pivoting motions so that the mill can be moved in the cutting plane determined by the robot. Optical markers are used for calibration and intraoperative navigated computer assistance. My prototype was based primarily on the Praxiteles base and the motor unit. The following section describes the prototype in detail.
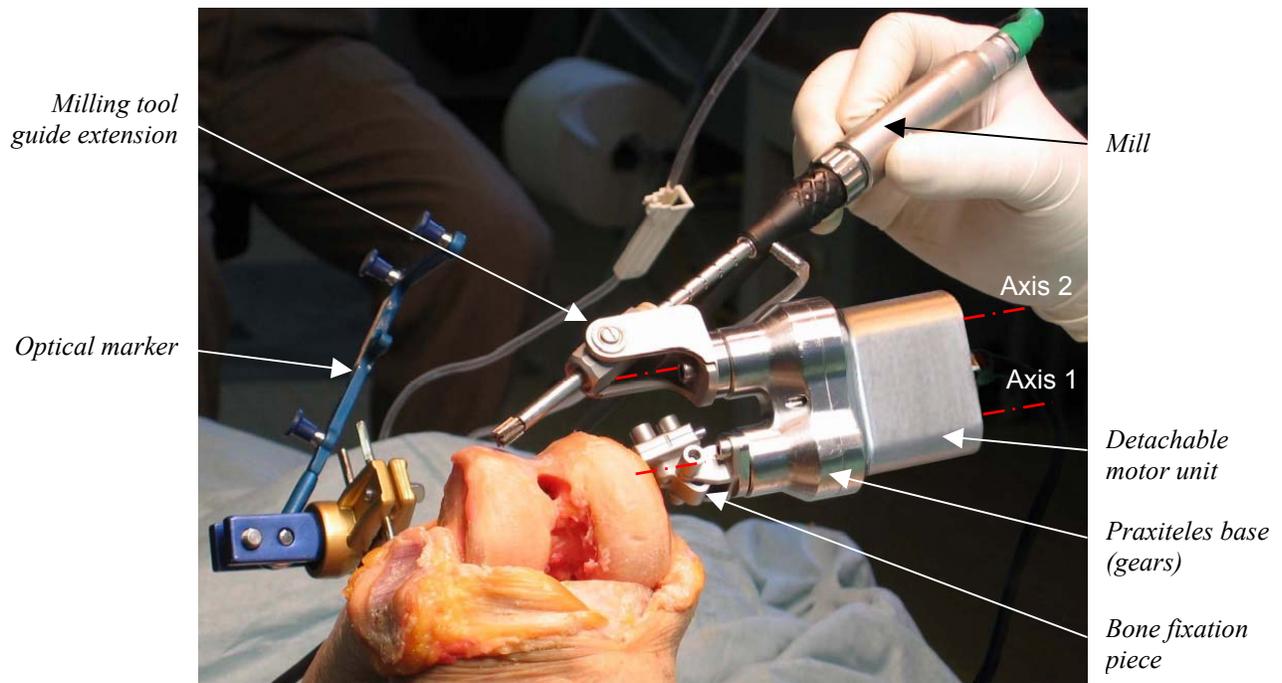
Figure 3-1: *The original Praxiteles shown in a distal femoral bone milling cadaver test (photo credit: Christopher Plaskos).*

3.2 Prototype description

The prototype and its control box are shown in Figure 3-2. The prototype is a two-link RR manipulator. It is rigidly mounted to a flat drawing surface at axis 1. The first link consists of the original Praxiteles base and motor unit. It can rotate freely about axis 1. The motor unit consists of two brushless DC servomotors with linear Hall sensors (Faulhaber BL 2036B). The second link and the dynamic physical constraint were custom manufactured for this prototype. The second link can rotate freely about axis 2 until it comes in contact with the physical constraint which is actively driven about this axis by one of the Praxiteles motors (see Appendix A for design drawings). This motor is connected through a 1:100 harmonic drive located inside the Praxiteles base. The other motor, connected to axis 1, is only used for its Hall sensor and remains unpowered at all

Figure labels (top image): Axis 1, Axis 2, Pen (end effector), Praxiteles base (link 1 - free), Dynamic physical constraint (driven), Fixation mount, Link 2 (free)

Figure labels (bottom image): Power supply, Output to robot, Emergency cut-off switch, Serial port input from computer, Faulhaber motion controllers
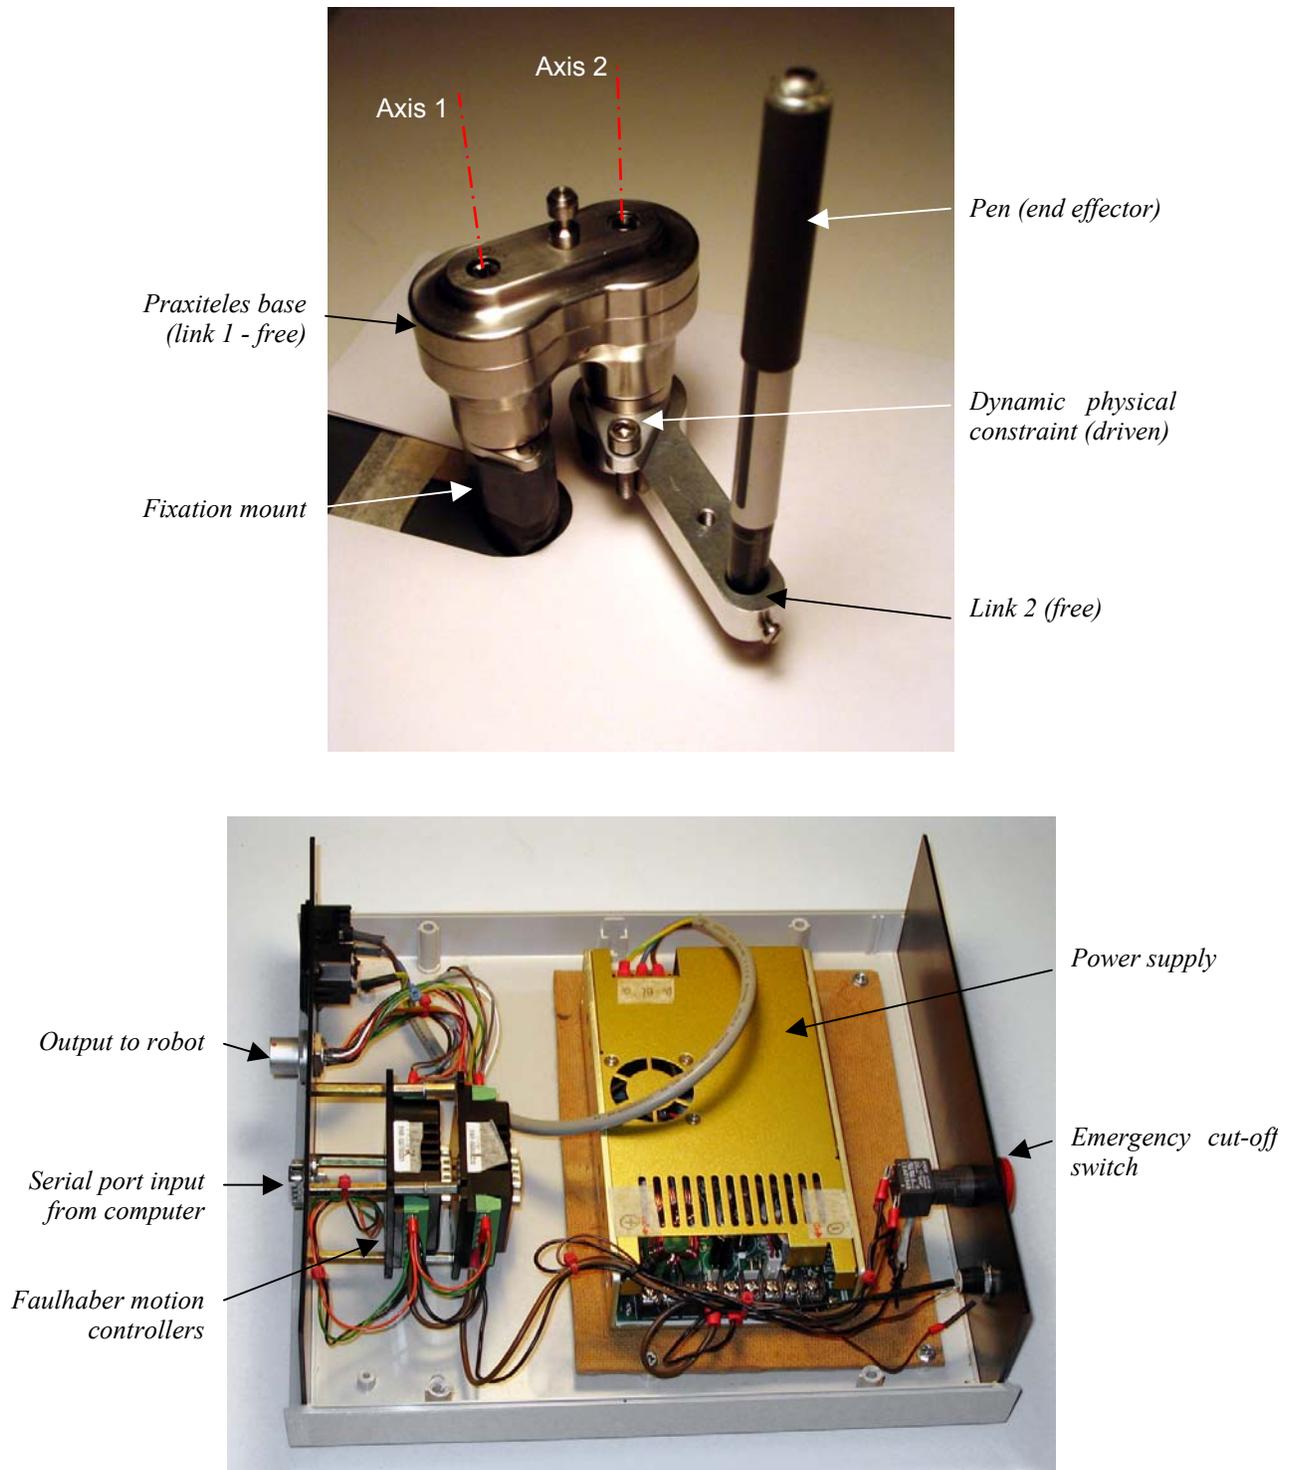
Figure 3-2: *(a) The prototype used to test the dynamic physical constraint concept. The motor unit is not shown, and plugs into the top of the Praxiteles base. (b) The electronics used to control the prototype.*
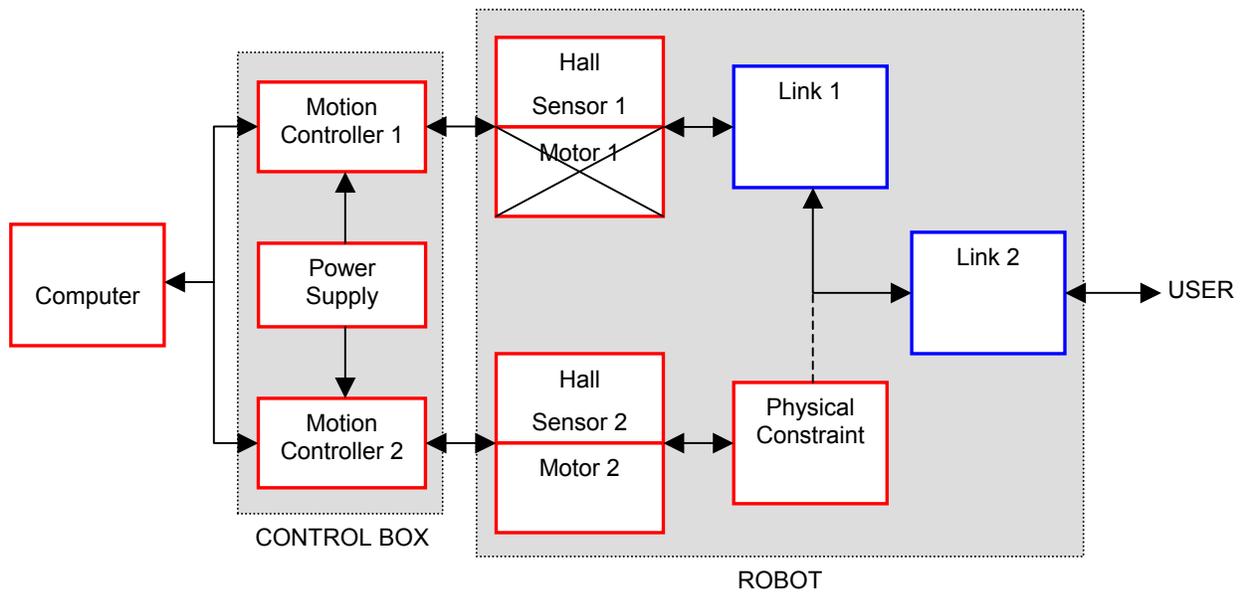
Figure 3-3: *Block diagram showing the general arrangement of all the components in the prototype. Motor 1 is always unpowered and is used only for its Hall sensor. The blue components are passive and the red components are active. The dotted line signifies that interaction between the physical constraint and link 2 only occurs when they are in contact with each other.*

times. The harmonic drive connected to this motor was removed and replaced by a direct drive axle which dramatically reduces the friction about this axis. The manufacturer claims that the Hall sensors have half a degree of precision, which was deemed sufficient for the use of this prototype. A pen is mounted at the robot's end effector, and traces the user's motions on the drawing surface shown.

The control box consists of a power supply and two motion controllers (Faulhaber MCBL 3006 S) connected in series (one for each motor). The motion controllers manage the input and output to and from the motors and Hall sensors and communicate with the computer through a serial connection.

A general block diagram of the whole system is shown in Figure 3-3. The user moves the end effector located on link 2. The motion of link 2, in turn, affects the motion of link 1. Interaction of the links with the physical constraint occurs only when they are in contact. During these interactions, Hall sensor 1 tracks the position of link 1 about axis 1 and sends the information continuously through motion controller 1 to the computer. The computer then uses this information to decide where the physical constraint should be positioned and sends this through motion controller 2 to motor 2. When link 2 is not in contact with the physical constraint, the two robot links act passively and independently from the active constraint.

3.3 Control program

The dynamic constraint is controlled by a laptop computer (AMD Athlon 64 Processor, 797MHz, 512 MB RAM, 115200 Baud RS-232 connection). Figure 3-4 shows a block diagram of the control program used. The program was written in Visual Basic 6.0 using an MSComm control for serial communication. It relies on continuous readings of the position of the first axis. Based on this position, it determines the respective position in which the physical constraint must be placed to prevent incursion of the end effector into the virtual surface. This is done using a modified inverse kinematics calculation where the known variables are the position of the first joint, the equation of the virtual surface and the length of the robot links, and the unknown is the position of the end effector. The actual implementation of these calculations can be done in two ways: the modified inverse kinematics can be calculated online using numerical methods to solve the non-linear problem, or basic inverse kinematics can be used beforehand to create a lookup table of joint 1 and 2 positions spanning the entire surface. The latter method then uses a generic binary search algorithm and interpolation to find the matching joint 2 position. These methods are explained in detail in the Appendix B.

36
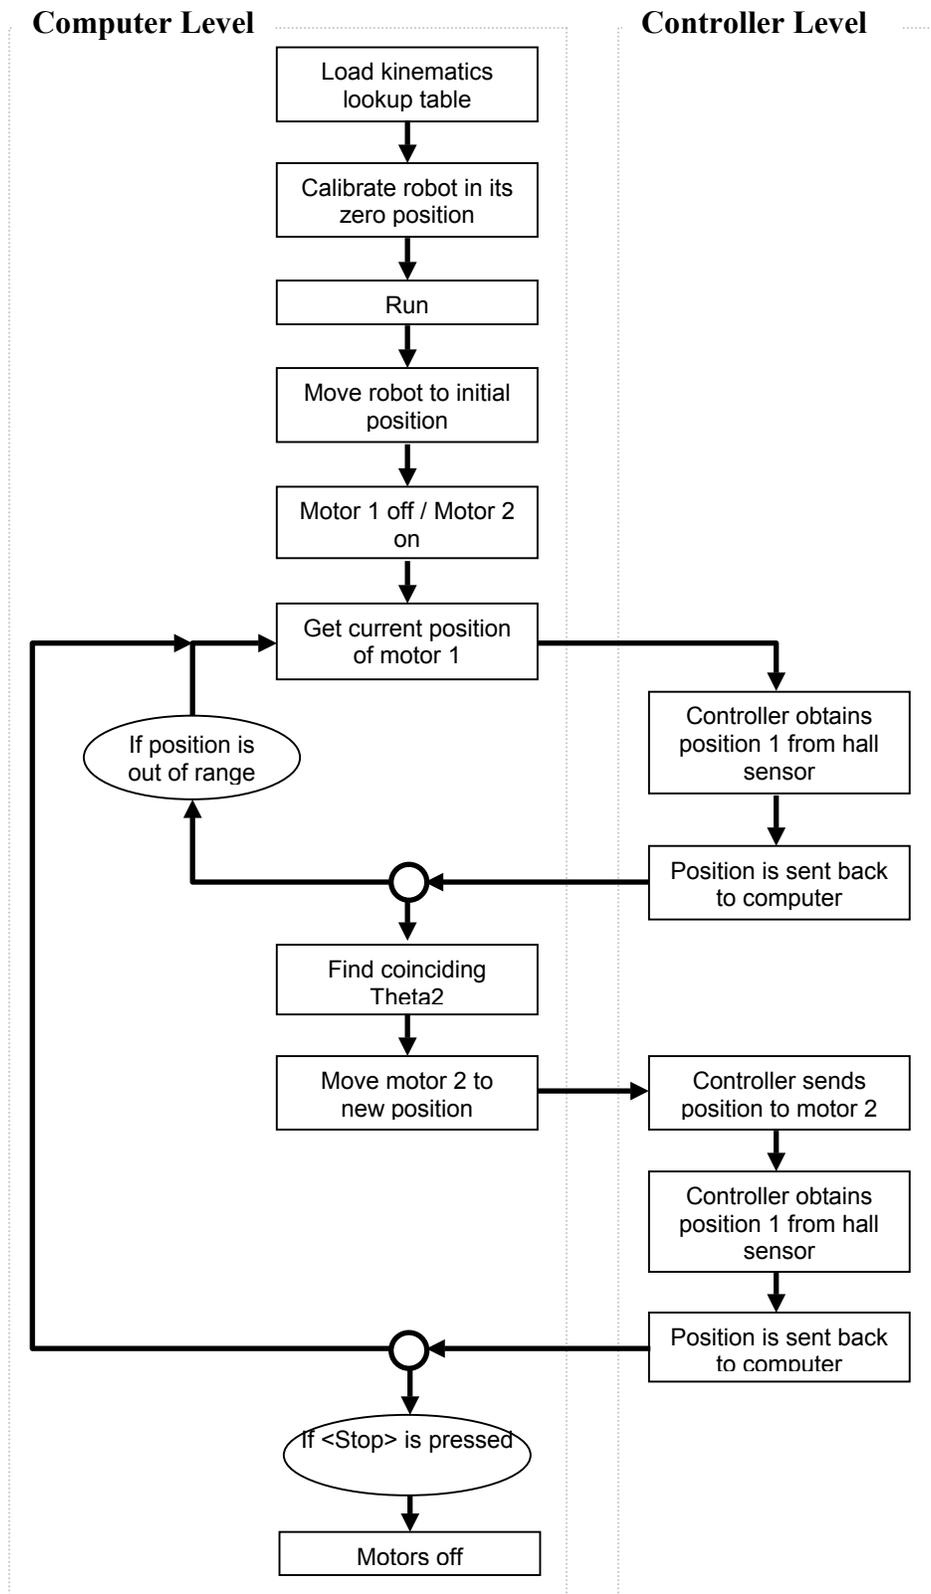
Figure 3-4: *Block diagram of prototype control program.*

Since the efficiency of the control loop is essential for the rapid response of the system to the user's motions, the execution speeds of the two methods described above were compared. Each method was used to find the coinciding joint 2 locations for 150 different joint 1 positions, using an elliptical surface. The speed and number of iterations required to find each position were recorded, using the Windows system high-resolution performance counter. The online numerical method, using Halley's root finding method, required, on average, 22µs and 4 iterations to converge to a solution. The kinematic lookup table method, using 2021 rows, required, on average, 9µs and 12 iterations to converge to a solution. The difference in speed is due to the time-consuming mathematical functions used in the numerical method, compared to the simple integer division used in the binary search routine.

If we consider an ultimate desired control loop frequency of 1kHz as being sufficient to manage human motions, we can see that the completion times of these routines are about 2% or less of the overall desired loop times. Therefore, in terms of speed, we can conclude that either would be satisfactory. The lookup table, however, has the advantage of being able to be used for complex digitized surfaces that cannot be easily defined by an equation.

3.4   Prototype testing

The primary purposes of prototype testing were (1) to obtain subjective data on how effectively the dynamic physical constraint concept simulates simple surfaces and how it reacts to typical user interactions and (2) to verify my theoretical analysis, identify any unexpected behaviour, and quantify the extent of incursions into the constraint surfaces.

3.4.1 Testing Method

The prototype was tested using a variety of virtual surfaces, as shown in Figure 3-5.

1. The ellipse (major axis = 80mm, minor axis = 50mm) is the simplest surface, as it is only slightly more complex than a circle, requiring minimal motion of the dynamic constraint as the end effector traverses the length of the constraint surface.

2. The horizontal line (y = 60mm, -34mm $\leq$ x $\leq$ 80mm) is more challenging to execute because it requires more motion of the physical constraint as the extremities of the line are approached, but it also allows simple and effective measurement of the precision of the system.

3. The hybrid circle (radius = 40mm)/ellipse (major axis = 80mm, minor axis = 50mm) surface is a modified version of the ellipse and has a tighter radius of curvature; this shape was chosen because it corresponds to a particular implant design.

4. The sine wave (amplitude 5mm, y = 60mm, -34mm $\leq$ x $\leq$ 80mm) is designed to test the ability of the user to navigate a profile with relatively small details superimposed on a larger profile.

5. Finally, the tri-planar surface (left plane: slope = 1, y-intercept = 100mm, middle plane: slope = 0, y-intercept = 74mm, right plane, slope = -1, y-intercept = 100mm) is a simplified version of the five-cut pattern usually used for femoral implants. It was chosen to test the system's ability to handle short straight lines and corners.

The first four surfaces were tested on nine different users. Two users were female and the rest were male, and all users were within the age of 25 and 35. The users had no prior knowledge of the system. They were asked to use the prototype to trace out and describe verbally the four surfaces, whose shapes were initially unknown to them. They were given three rules: (1) to keep the prototype's end-effector within specific boundaries pre-

(a) Ellipse

(b) Horizontal line

(c) Hybrid circle/ellipse

(d) Horizontal sine-wave

(e) Tri-planar
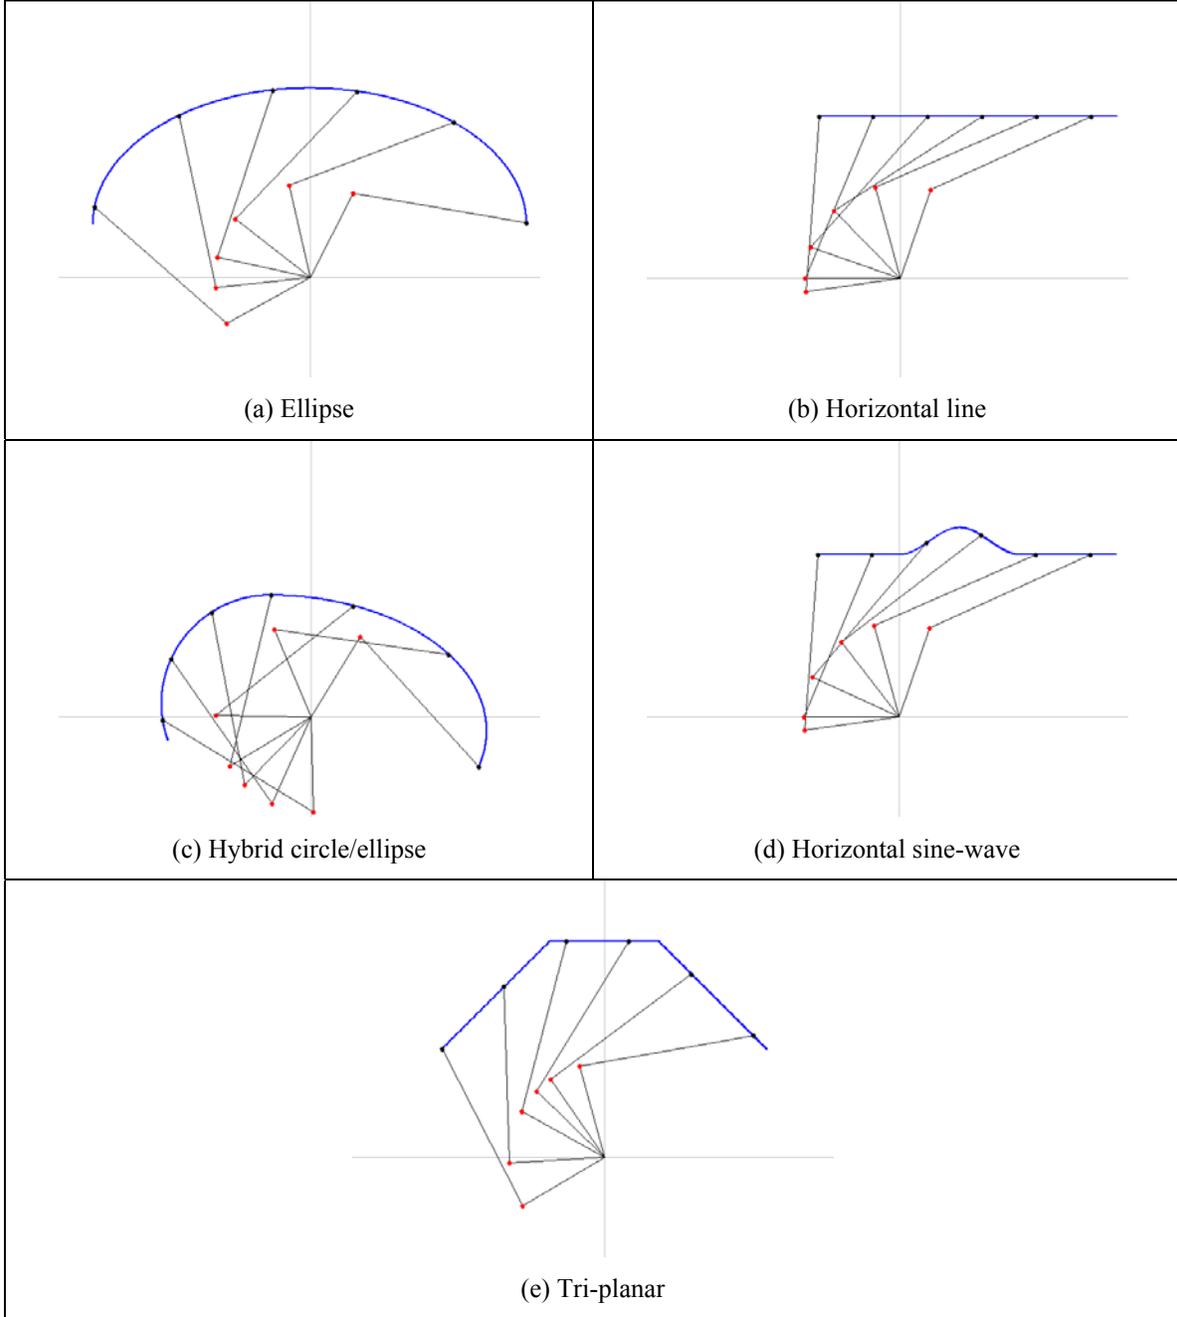
Figure 3-5: *Virtual surfaces used to test the prototype. The surface is drawn in blue, with the robot sketched at a number of locations in black.*
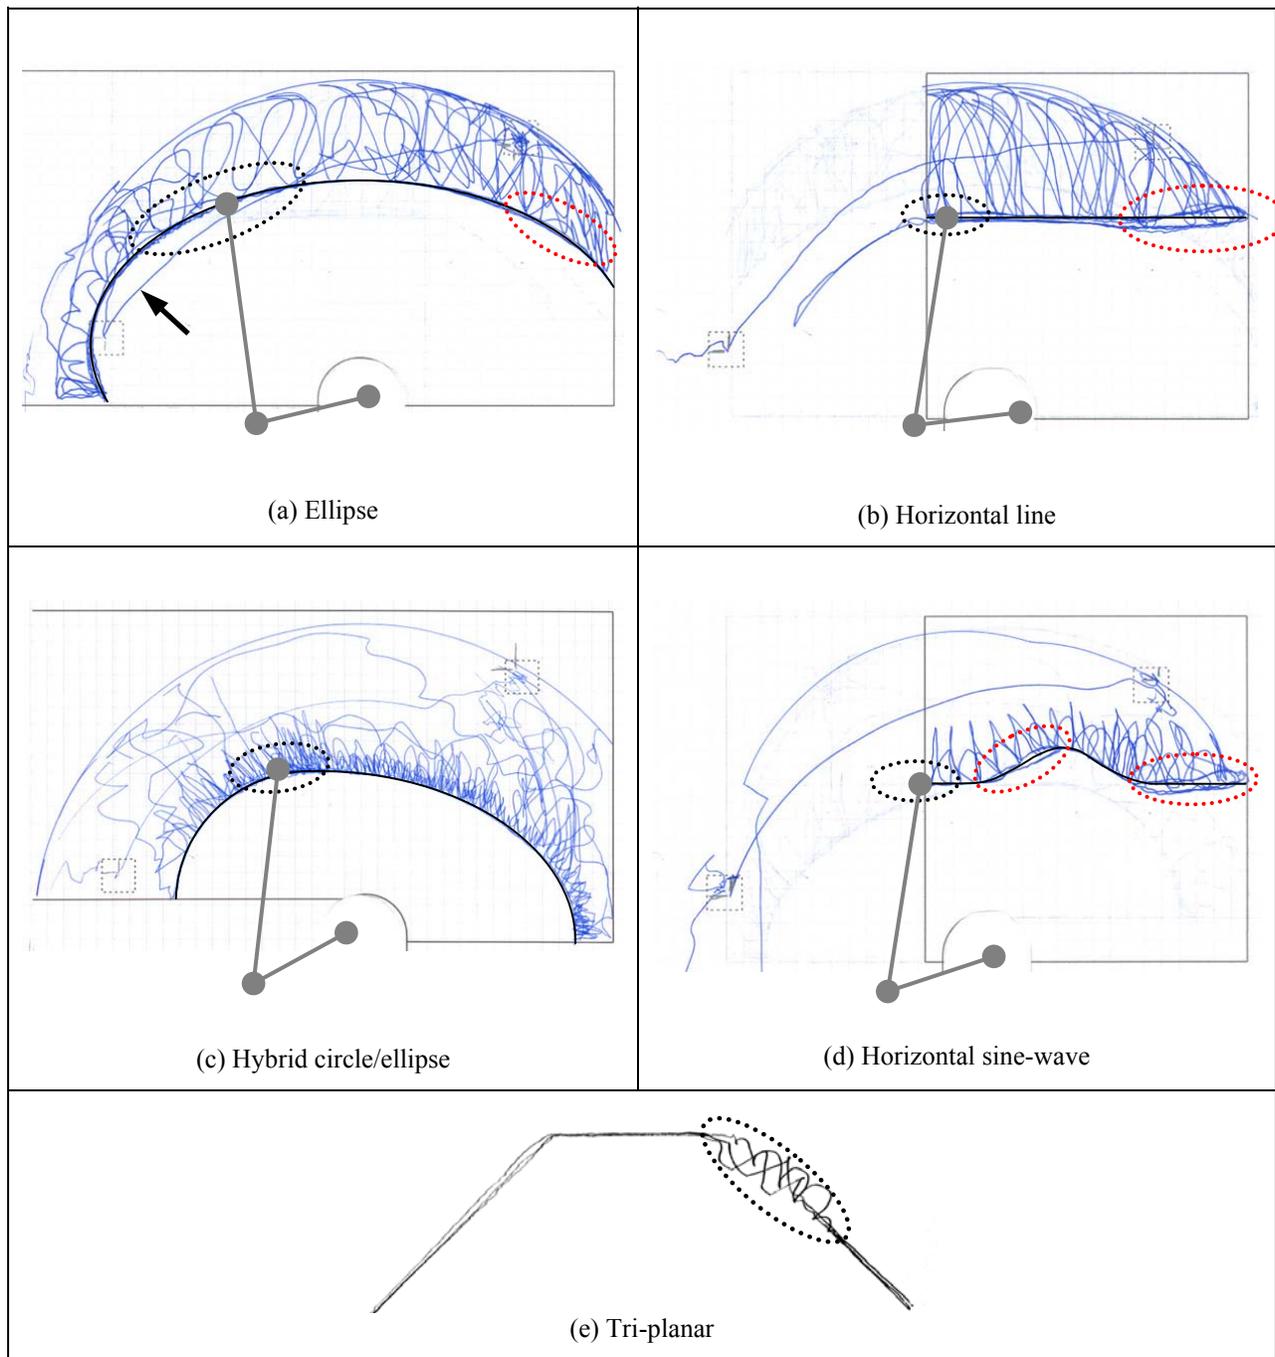
Figure 3-6: *Sample, representative scans of five drawing surfaces used in the prototype testing (see Appendix C for scans of all data). Regions of instability are circled in black, and regions of hysteresis are circled in red. The approximate location of the robot links is drawn in gray at the regions of instability. The line marked by a black arrow was done during the calibration procedure, before the robot was turned on, and should be ignored.*

41

drawn on the drawing surface, (2) not to invert the elbow such that the second link would no longer be within reach of the dynamic constraint and (3) to restrain from moving the prototype unreasonably fast. All other motions were permitted. Prior to each attempt, the location of the drawing surface (a sheet of graph paper) was calibrated with respect to the prototype by marking and recording the position of two distinct points. During the trials, observations were made with respect to user reactions, comments and ability to handle the prototype. Afterwards, the resulting motion traces drawn by the pen on the drawing surface were scanned, digitized, calibrated and compared to the original surfaces to determine the extent of surface penetration.

The fifth, tri-planar surface was tested randomly by a number of users without collecting data. Observations were made on the robot and user's behaviours.
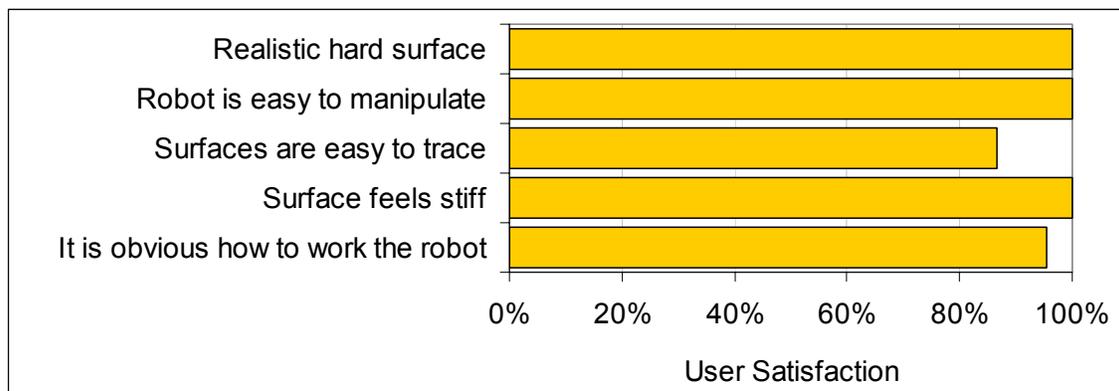
Figure 3-7: *Average user satisfaction with respect to the design objectives outlined in Section 1.4 during prototype testing.*

3.4.2 Test observations

Figure 3-6 shows a sample scan of each of the five surfaces created by the users. These sample scans were chosen for their clarity in presenting the various behavioural effects

and for being most representative of the average trace. Digitized versions of all the user scans can be seen in Appendix C. Lines that are outside the surface boundaries or that are obviously cutting through the surfaces were done during calibration and should be ignored. Figure 3-7 shows the average user satisfaction with respect to the specific design objectives outlined in Section 1.4. The following is a list of general observations made during and after the testing with regards to general user reactions:

- all surfaces were successfully traced and identified by the users.

- all users were thoroughly convinced of hitting a hard surface. Most users were additionally satisfied by the "click" sound emitted upon contact.

- users reported no difficulty in moving the end effector.

- users reported no difficulty in tracing the surfaces once found.

- all users were surprised by the stiffness and strength of the virtual hard surface.

- a rapid learning curve was observed in the users' confidence and ability to find and trace out the surface: after the first surface, users were able to follow the surfaces almost immediately by finding a contact point and then tracing along it.

- some users' motions were affected by the two-link architecture of the robot: their motion attempted to follow the rotation of the second link.

Overall, user reactions to the system were very positive. The simplicity and realistic feel of the system was a major factor in their feedback.

Along with user reactions, observations were made on the prototype's behaviour. The following is a list of these observations:

- the lateral deflection effect, described in Section 2.3.1, was only noticeable at the extreme right of the linear surfaces. The robot would take control of the user's motions and push them to the right in an unexpected manner. As

predicted, however, after several tries, users were able to control this behaviour and did not find it disturbing.

- some users noticed "stickiness" when moving to the left at the extreme right of the linear surfaces.

- during tracing of the surface, regions of hysteresis were noticed primarily at the rightmost extreme of the linear surfaces (Figure 3-6 (b) and (d)). The pen would draw a counter-clockwise loop at these locations (ie, there would be a significant incursion during the rightward motion, followed by a return closer to the intended virtual surface). The left side of the sine wave was affected in the same manner. The ellipse displayed minor hysteresis on its right side, while the hybrid surface displayed essentially none.

- during tracing of the surface, regions of instability were noticed with the elliptical surface and less so with the hybrid surface (see Figure 3-6 (a) and (c)).

The tri-planar surface, tested separately, showed some different responses. The corners were very accurately emulated with virtually no rounding occurring, although users typically produced a slight overshoot as they moved around a corner; if one imagines this occurring when using a cutting tool, the extra material could be easily removed with a second pass in the reverse direction. The lateral deflection effect was hardly noticed anywhere along the surface. Significant instability, however, was noticed just to the right of the second corner.

In order to successfully implement the dynamic physical constraint concept for specific tasks, it is important to understand all of these behavioural characteristics. The following chapter attempts to explain them and suggests ways of eliminating or reducing them.

4.0 ANALYSIS OF TEST RESULTS

The behavioural observations described in the previous Chapter will be explained in the following sections. Specifically, I will discuss the causes of the lateral deflection, reverse stickiness, hysteresis and instability behaviours. Suggestions will also be made for addressing these issues. I will begin, however, by presenting an analysis of the prototype's ability to prevent incursions across the constraint surface.

4.1 Penetration analysis

The amount of surface penetration that occurred with this prototype was determined from the scans of the user motion traces. Each trace had two calibration points marked on them which were measured in both paper and robot coordinates. The scans of the motion traces were digitized by hand using a digitization program at high zoom levels (note that only the inner lines representing the trace of the virtual surface were digitized. The other lines in which the user was not in contact with the surface were ignored). A scaled homogenous transformation between the paper and robot reference frames was then applied (see Figure 4-1) in order for the digitized points to be overlaid on a plot of the theoretical constraint surface. These plots can be seen in the Appendix C.
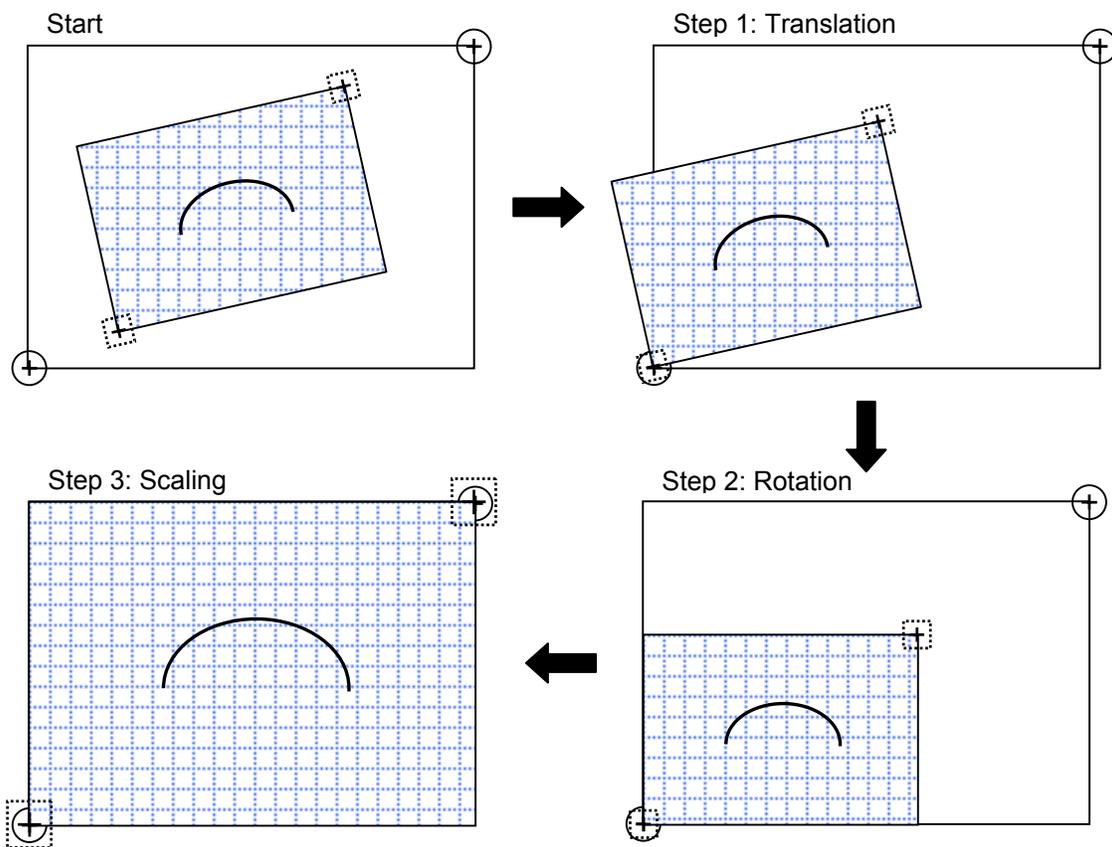
Figure 4-1: *Scaled homogeneous transformation used to calibrate the digitized scan of the user's motion trace (blue graph paper) with respect to the robot's coordinates (large white square).*

The distance of each calibrated digitized point from the intended surface, along a line normal to the surface, was calculated to determine the amount of penetration into the virtual surface. A summary of the results is shown in Figure 4-2. The cumulative distribution of this data is shown in Figure 4-3. The surface that was most accurately traced was the hybrid circle/ellipse surface, followed closely by the elliptical surface. The average penetration for the hybrid surface was 0.83mm (SD±0.46) and for the elliptical surface was 1.01mm (SD±0.59). From the cumulative distribution plot, we can see that 90% of the measurements for these two convex curvilinear surfaces had less than 2mm of incursion for all the users. The sine wave and the horizontal line were significantly less accurate with penetrations of up to 10mm occurring and 90% of the measurements being less than 6mm. The primary region of penetration for these surfaces was at their extreme right, where the robot was nearing its fully stretched out singular configuration. In this region, the robot was highly susceptible to the lateral deflection behaviour, typically producing a sudden lateral acceleration of the end effector when an incursion occurred. The faster the end effector moved, the greater the depth of incursion. The sine wave surface also saw significant penetration on the left side of the wave.

Sources of error in these precision measurements include the following:

- Position sensor errors: the hall sensor on the first joint had no gear reduction, so the half a degree of error in the Hall sensor readings could have caused errors in the measurement of the two calibration points in the robot coordinates. Half a degree translated into maximum error at the end effector (using the worst case scenario of the fully stretched out robot with link lengths of 35 and 65 mm) is 0.87 mm. The robot, however, was never in such a stretched out position, so a more realistic error would be on the order of 0.5mm (based on a radial distance of 60mm). Note that the half a degree sensor error was reported by the manufacturer as being pessimistic. The sensor error on the first joint can be considered insignificant due to the 1:100 gear reduction.

- The backlash in the harmonic drives is on the order of 0.1° (Plaskos, 2005).
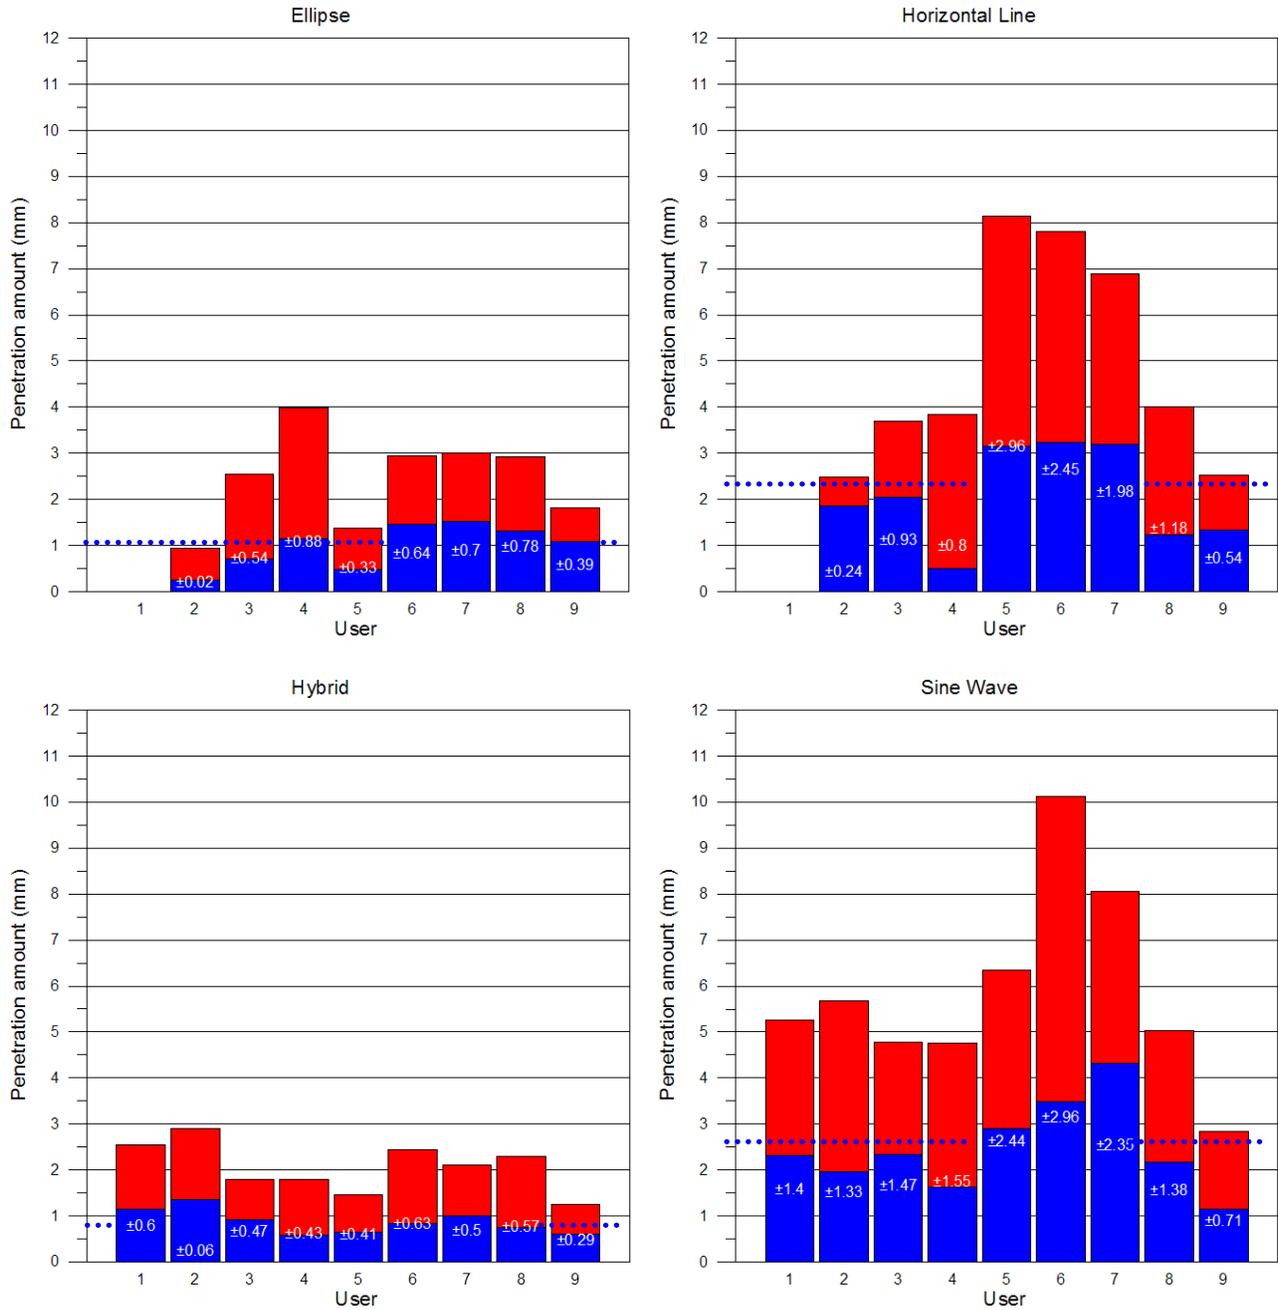
Figure 4-2: *Summary of the amount of surface penetration experienced by the nine users during the experiment, for each of the four surfaces. The red bars show the maximum amount of penetration along the entire surface. The blue bars show the mean amount of penetration along the entire length of the surface, with standard deviations written in white. The overall weighted averages, marked with a dotted blue line, were: 1.01 (SD±0.59) for the ellipse, 2.23 (SD±1.74) for the horizontal line, 0.83 (SD±0.46) for the hybrid and 2.61 (SD±1.96) for the sine wave.*
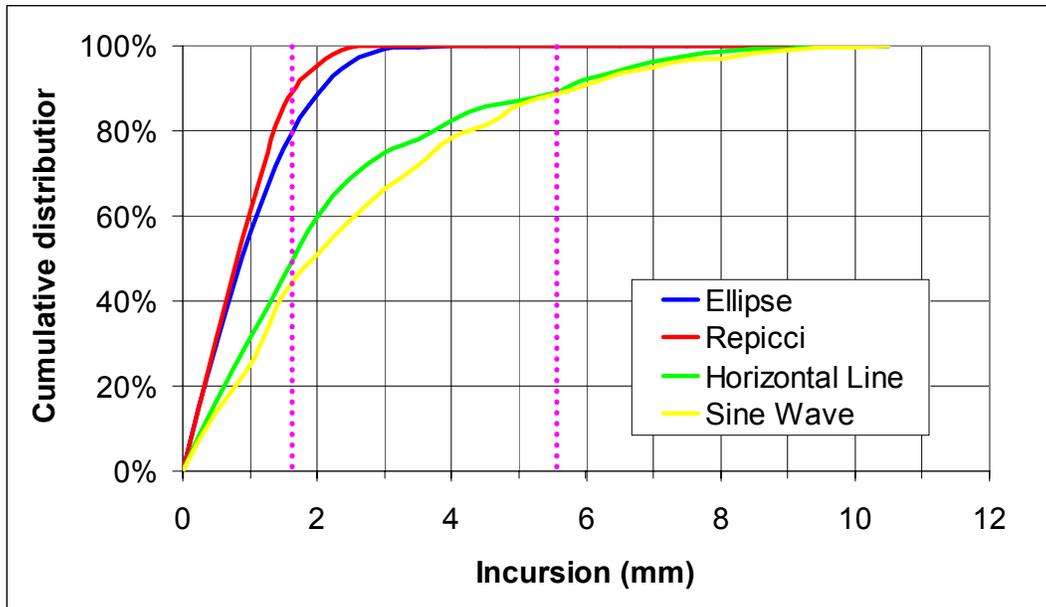
Figure 4-3: *Cumulative distribution of the surface penetration measurements for all nine users during the experiment, for each of the four surfaces. The dotted lines mark the approximate 90% cumulative distribution for the convex ellipse and Repicci surfaces and the linear horizontal and sine wave surfaces.*

- Hand measurement errors: the calibration points were measured in paper coordinates by hand using precision calipers, so errors could have been on the order of about 1 to 2 mm.

- Joint looseness: the prototype joints have minimal looseness in the vertical direction. In the fully stretched out position, there is about 1mm of vertical play. Since the prototype is 1cm above the drawing surface, this would cause insignificant horizontal error. The pen also had minimal play in its seating. This was minimized by the low position of the robot and also by instructing the users to hold the prototype by link 2, rather than by the pen.

Combining the above errors would give us a total measurement error of:

$$E_{total} = \pm\sqrt{0.5^2 + 0.1^2 + 2^2} = \pm 2.1 mm$$

49

The objective of this precision analysis was to get a general idea of the order of precision of the system with the goal of deciding whether this new haptic concept should be further developed specifically for precision applications such as surgery. To obtain an implant frontal plane alignment accuracy within 3°, as specified in Plaskos (2005), and assuming an approximate medial-lateral knee width of 60mm, the precision of the system would ideally be within 1.5mm. The system was designed in the simplest manner possible for concept validation purposes. From the beginning, therefore, the electronics, software, and mechanics of the system were not optimized for accuracy, communication speeds and true real time reaction speeds. The majority of the reported errors can be dealt with by redesign and optimization. For the simplicity of the prototype, therefore, these results are very encouraging.

4.2   Lateral deflection analysis

The lateral deflection behaviour was predicted in the design validation phase described in Section 2.3.1. This behaviour is characteristic of the dynamic physical constraint concept and cannot be avoided. It can, however be limited by the optimization of design parameters. To determine what these parameters are, it is important to understand and be able to simulate the severity of the issue in various situations.

4.2.1   Equilibrium point concept

The deflection response can be explained using the equilibrium point concept described by Hogan (1985). The concept is illustrated in Figure 4-4 in the context of my prototype and refers to the idea that activated muscles exhibit a spring-like behaviour so that, for a
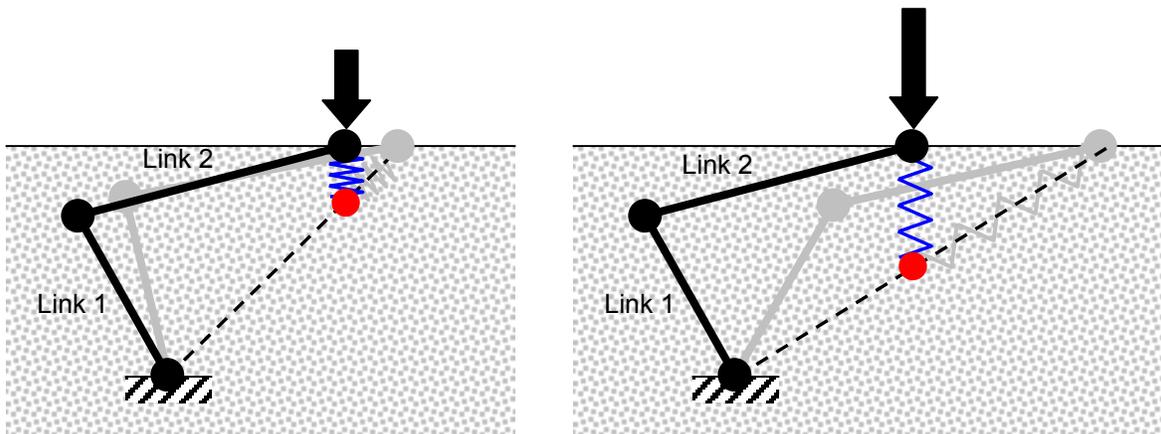
Figure 4-4: *Illustration of the virtual equilibrium point concept. The equilibrium point is drawn in red, and its depth from the surface depends on the amount of force applied by the user. The model of the mechanical interaction between the user and the surface is represented by the blue spring. The user's lateral stiffness is represented by the gray spring. The intersection point between the dotted line and the surface is where the end effector will be deflected to if the equilibrium point remains at rest and all friction in the system is negligible. The gray solid lines represent the final deflected position of the robot.*

constant level of activation, muscles will automatically generate a restoring force when the limb is perturbed from the equilibrium position defined by the activation state. This equilibrium position can be imagined as the location of a 'ghost' or virtual arm where it is not subjected to any interaction forces. In the absence of interaction forces, the real and virtual arms will coincide, but the real arm can be displaced from the virtual arm by applying interaction forces. For example, when a person pushes against a flat surface (such as a table or, in my case, the virtual surface emulated by the physical constraint), we can imagine a virtual point being cast in the direction of the force to a point below the surface. This point is the virtual location where the system would be in equilibrium if the surface contact were suddenly removed. If we model the muscle behaviour as a spring-damper system, the location of the equilibrium point relative to the surface under static conditions would depend on the spring constant and the force applied by the user. Therefore, the harder the user pushes against the surface, the deeper the equilibrium point would be, assuming a relatively constant stiffness.

4.2.2 *Application of the equilibrium concept to the lateral deflection analysis*

The relevance of this concept to my analysis of the prototype's lateral deflection behaviour is that if we know the location of the equilibrium point, we can determine how much deflection will occur. This is done by simply extending the line connecting the robot's attachment point and the equilibrium point until it reaches the surface. To explain this, imagine that the end effector is located within the surface at the virtual equilibrium point, frozen in time. In this situation, the physical constraint at the robot's elbow joint will want to push the end effector back to the surface by expanding the angle between the two links. This expansion will cause motion radially outward, in other words, along the line connecting the attachment point and the equilibrium point. The end-effector comes to rest along this line because the force due to displacement of the end-effector away from the equilibrium position produces no net moment about the robot's grounded pivot point.

Using this concept, the amount of lateral deflection experienced by the user depends on the equilibrium point location. We would therefore predict the following:

- the harder the user pushes against the surface, the deeper the equilibrium point is and the larger the lateral deflection.

- the closer the direction of the user's force coincides with the radial distance of the end effector from the attachment point, the less deflection there is. When the direction does coincide, no deflection occurs.

- the shape of the surface affects the amount of deflection felt.

- if the direction of the user's force angles away from the robot's attachment point, then the equilibrium point also moves away and the amount of deflection increases (and vice versa).
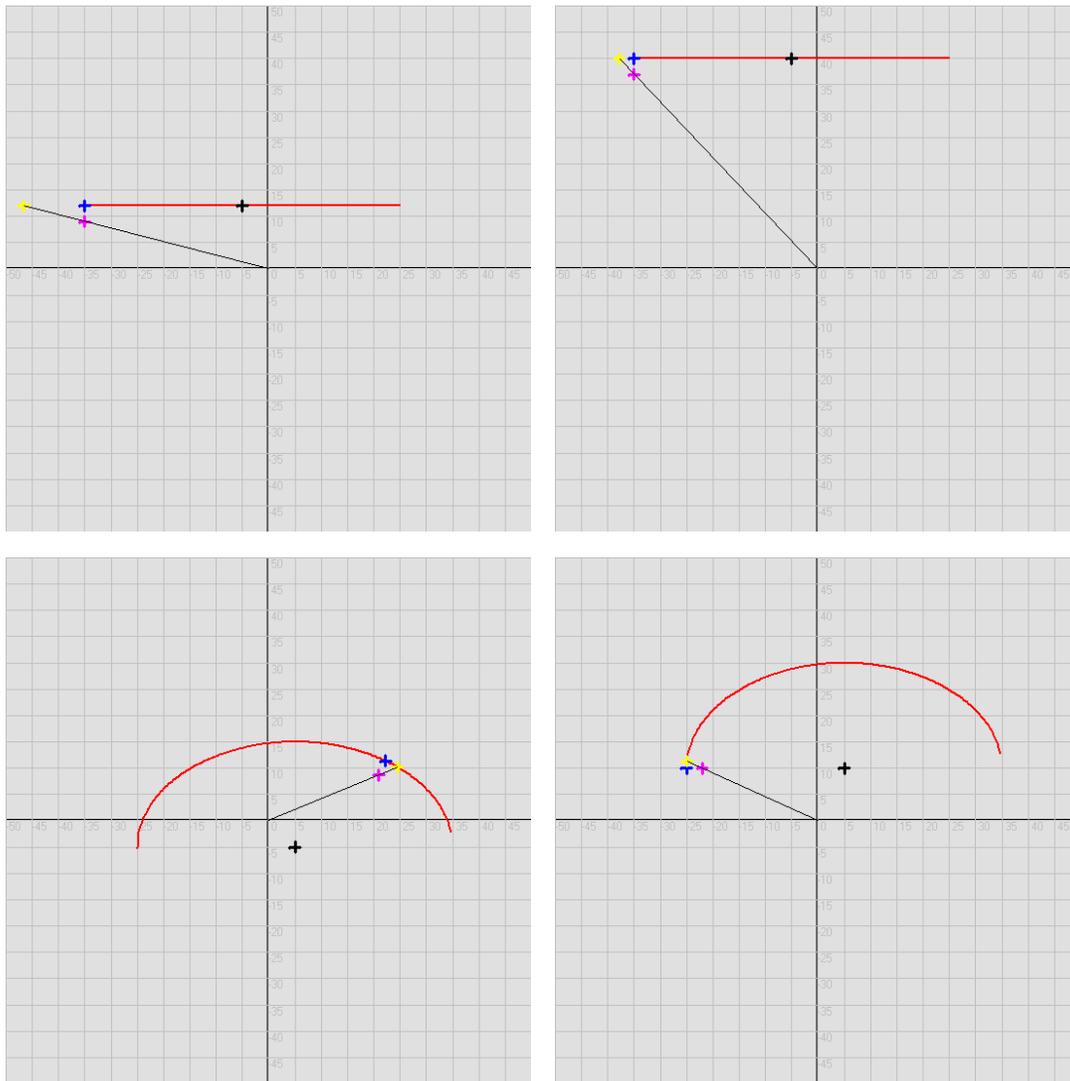
Figure 4-5: *Diagrams showing the maximum lateral deflection along a sample linear and an elliptical surface due to the equilibrium point effect. At different robot attachment points, the location and amount of this maximum varies. The surface is shown in red. The robot attachment point is at the origin of the graphs. The magenta cross shows the equilibrium point, located 3mm inwards and normal to the surface. The blue cross is the point on the surface at which the user's force (represented by the equilibrium point) was applied, and the yellow cross marks the point where the user would be deflected to. The black cross marks the radial centre of the surface.*

It is interesting to note that this concept is independent of the prototype's architecture. In fact, it can be seen that the RP configuration, described in Section 2.1, would react in

53

exactly the same manner. Therefore, the only design parameter which can be optimized to minimize this behaviour for a given constraint surface is the location of the robot's attachment point relative to the surface.

4.2.3   Optimization for minimum lateral deflection

Optimization of the location of the attachment point depends on the shape of the surface being emulated and the distance past the surface where the equilibrium point lies. For every different shape, the optimal location at which the maximum lateral deflection across the entire surface is minimized varies. In order to determine the characteristics of this variation, a program was written to calculate the maximum deflection that would occur across the surface for a matrix of attachment point locations. An example of how two different robot attachment points result in distinct maximum deflections can be seen in Figure 4-5. The program is based on the equilibrium point concept and produces two-dimensional colour plots showing how much the location of the attachment points affects the deflection behaviour. Some sample plots are shown in Figure 4-6 for an equilibrium point penetration distance of 3mm. Based on these plots and the geometric nature of the optimization procedure, we can conclude the following:

- for convex surfaces, the least amount of deflection across the entire length of the surface occurs when the attachment point is at the mirrored radial centre of the surface (the concept of a mirrored radial centre is described in Herburt et al. (2005)).

- the deflection distribution across a surface is minimally dependent on the size of the surface. In other words, concentric surfaces have essentially the same deflection distribution for any given attachment point location. Because of this, the x and y axes of the plots in Figure 4-6 are normalized with respect to surface size.

54

Circular (b=a)

Elliptical (b = 0.8a)

Elliptical (b = 0.5a)

Elliptical (b = 1.5a)

Quarter ellipse

Horizontal line
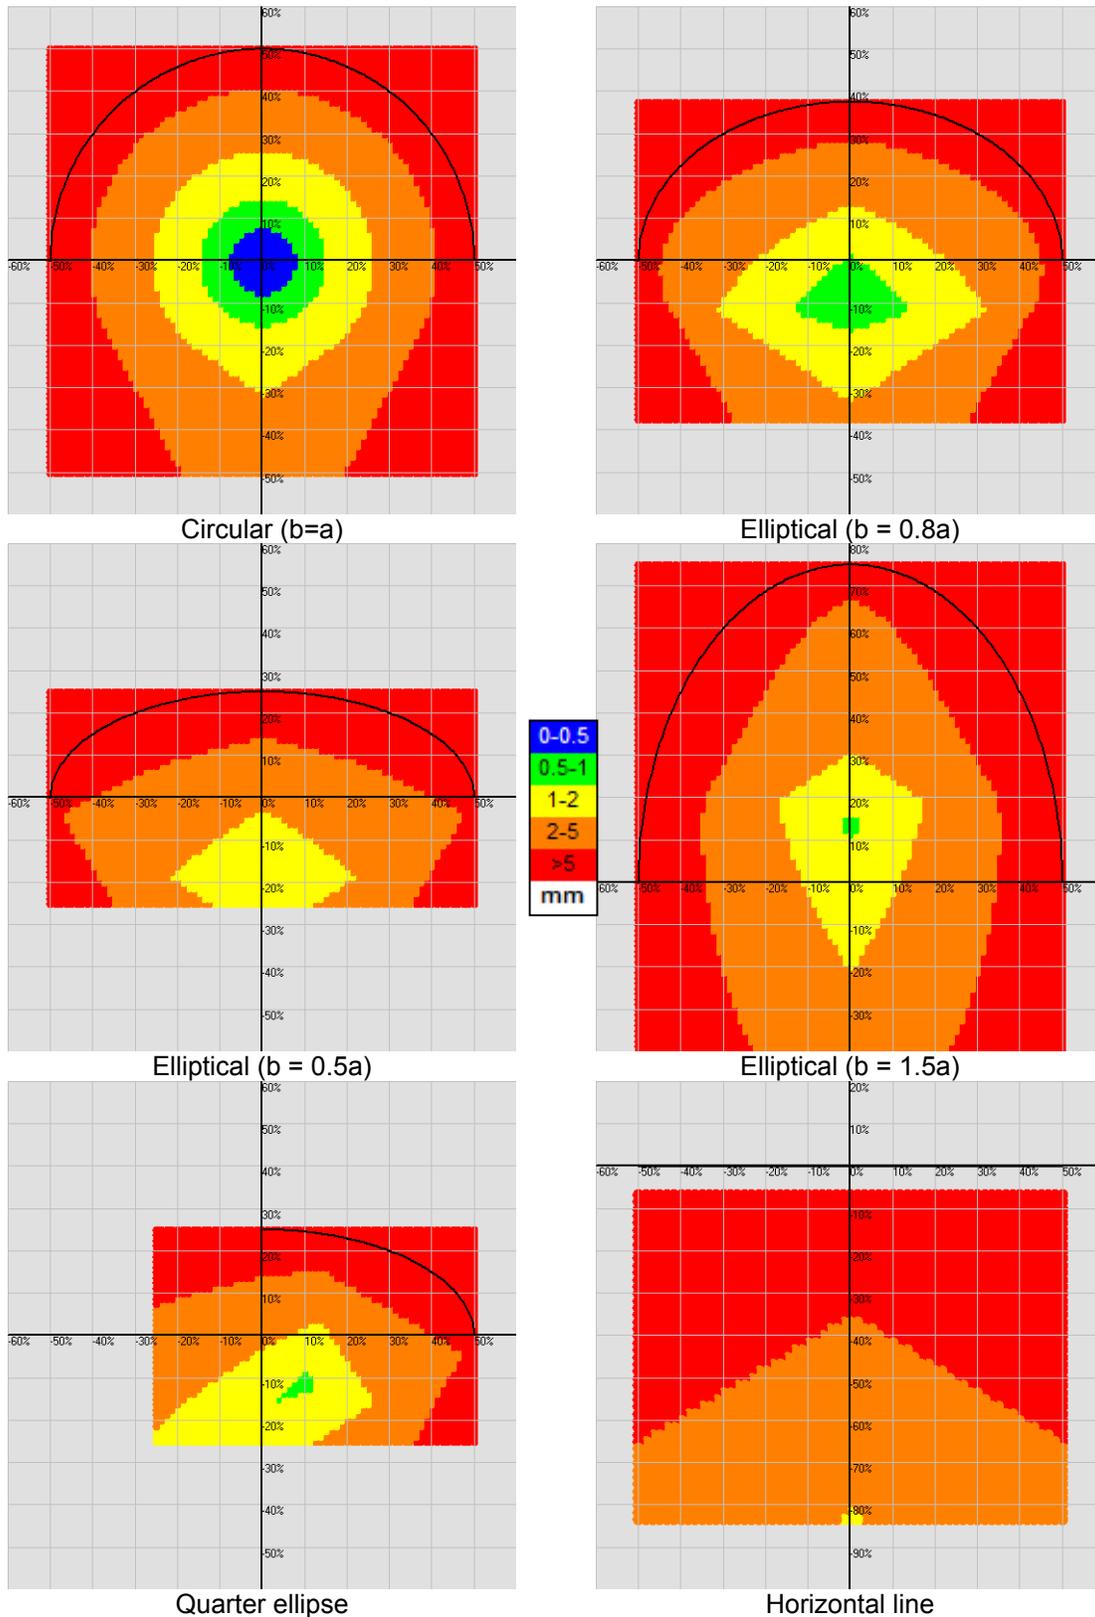
Figure 4-6: *Attachment point optimization plots for various surface shapes. The colours show the amount of maximum lateral deflection along the length of the surface for the various attachment points.*

- a circular surface is the most basic, since no deflection occurs at all when the attachment point is located on its radial centre (the system, in fact, turns into a one degree of freedom pendulum-like manipulator).

- the colour map of a hemi-elliptical surface can be obtained by stretching a hemi-circular shape.

- with linear surfaces, the further the attachment point, the smaller the maximum deflection.

- the colour maps of composite surfaces are the result of the addition of the separate components.

These generalized trends can be applied to any surface to approximate the optimal location of the robot's attachment point. Most convex surfaces are minimally affected when the robot is attached in the vicinity of their mirrored radial centre, having maximum deflections approximately equal to or less than the amount of equilibrium point penetration.

4.3   Reverse stickiness analysis

The stickiness that users felt at the extreme right of the linear surfaces, when trying to reverse their motion back towards the centre, can be described as being the opposite effect of the lateral deflection behaviour. In fact, the stickiness comes from the fact that the user must overcome this deflection effect. It is not possible to overcome it by force, since increasing the force conformingly increases the amount of deflection. The user must, in fact, change the direction of his or her applied force, such that a counter-clockwise moment is created about the attachment point, as shown in Figure 4-7. At the far right of a linear horizontal surface, this would involve applying a force at an angle that

is unintuitively acute with respect to the surface, hence the stickiness. With an RP configuration, the effect would be identical. As with the lateral deflection behaviour, this effect was easily learned by the users, and presented no significant difficulty to overcoming it. Since this effect is related directly to the lateral deflection issue, it is optimized when the lateral deflection issue is optimized.
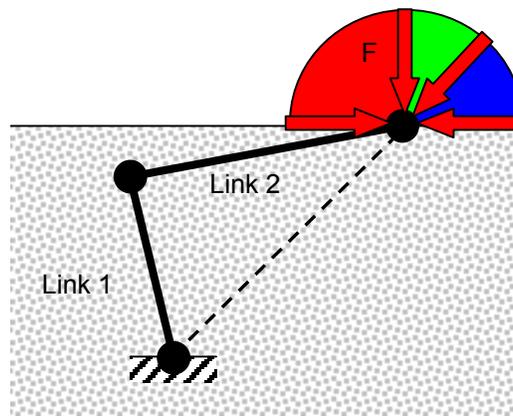
Figure 4-7: *Illustration of the stickiness effect. A force applied in the direction delimited by the red region will cause clockwise rotation of the robot about its base and hence motion to the right (lateral deflection effect). A force applied in a direction delimited by the green region, would also result in lateral deflection to the right even though the user is expecting motion to the left. This is the region in which stickiness is felt, since the user fights against the unexpected lateral deflection effect. The blue region indicates the region in which the user's applied force results in expected motion to the left.*

4.4 Hysteresis analysis

A more significant problem encountered in the behaviour of the prototype was hysteresis. Hysteresis in a mechanical system occurs when the output of the system varies depending on the path taken rather than on its current state.

4.4.1 Hypothesis

The prototype showed varied amounts of hysteresis for the different surfaces. The most apparent case was at the right side of the horizontal line (Figure 3-6(b)) where back and forth motion by the user caused the end-effector to move along a counter-clockwise circular path. This was likely caused by the response delay of the system. The faster the user traced the surface towards the right, the less time the dynamic constraint had to react to push the user back onto the surface. In the opposite direction, the physical constraint could not retreat quickly enough. Figure 4-8 shows a schematic of this effect. This behaviour was likely the cause of the prototype's lack of precision in these areas.
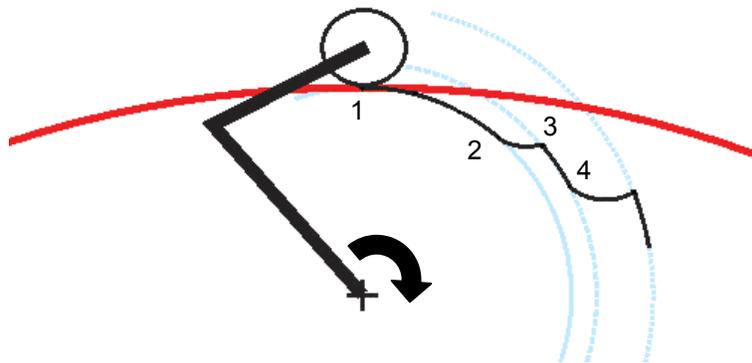
Figure 4-8: *Generalized schematic of the hysteresis problem. The virtual surface is shown in red, while the actual path followed by the end effector is shown in black. At point 1, the system is rotating to the right. At point 2, the system realizes it has penetrated the surface and it moves the constraint accordingly. The system keeps rotation, however, so when it stops moving, it has not actually reached the surface (point 3). At point 4, it once again realizes that it is within the surface and the same procedure repeats itself until the user slows down enough to give the system enough time to reach the surface.*

4.4.2 Hysteresis tests

A simple test was done to ensure that this hysteresis was in fact due to the robot's reaction speed. For this test, the harmonic drive on the first axis of the original Praxiteles

(a) Ellipse

(b) Horizontal line

(c) Hybrid circle/ellipse

(d) Horizontal sine-wave

(e) Four-planar

(f)
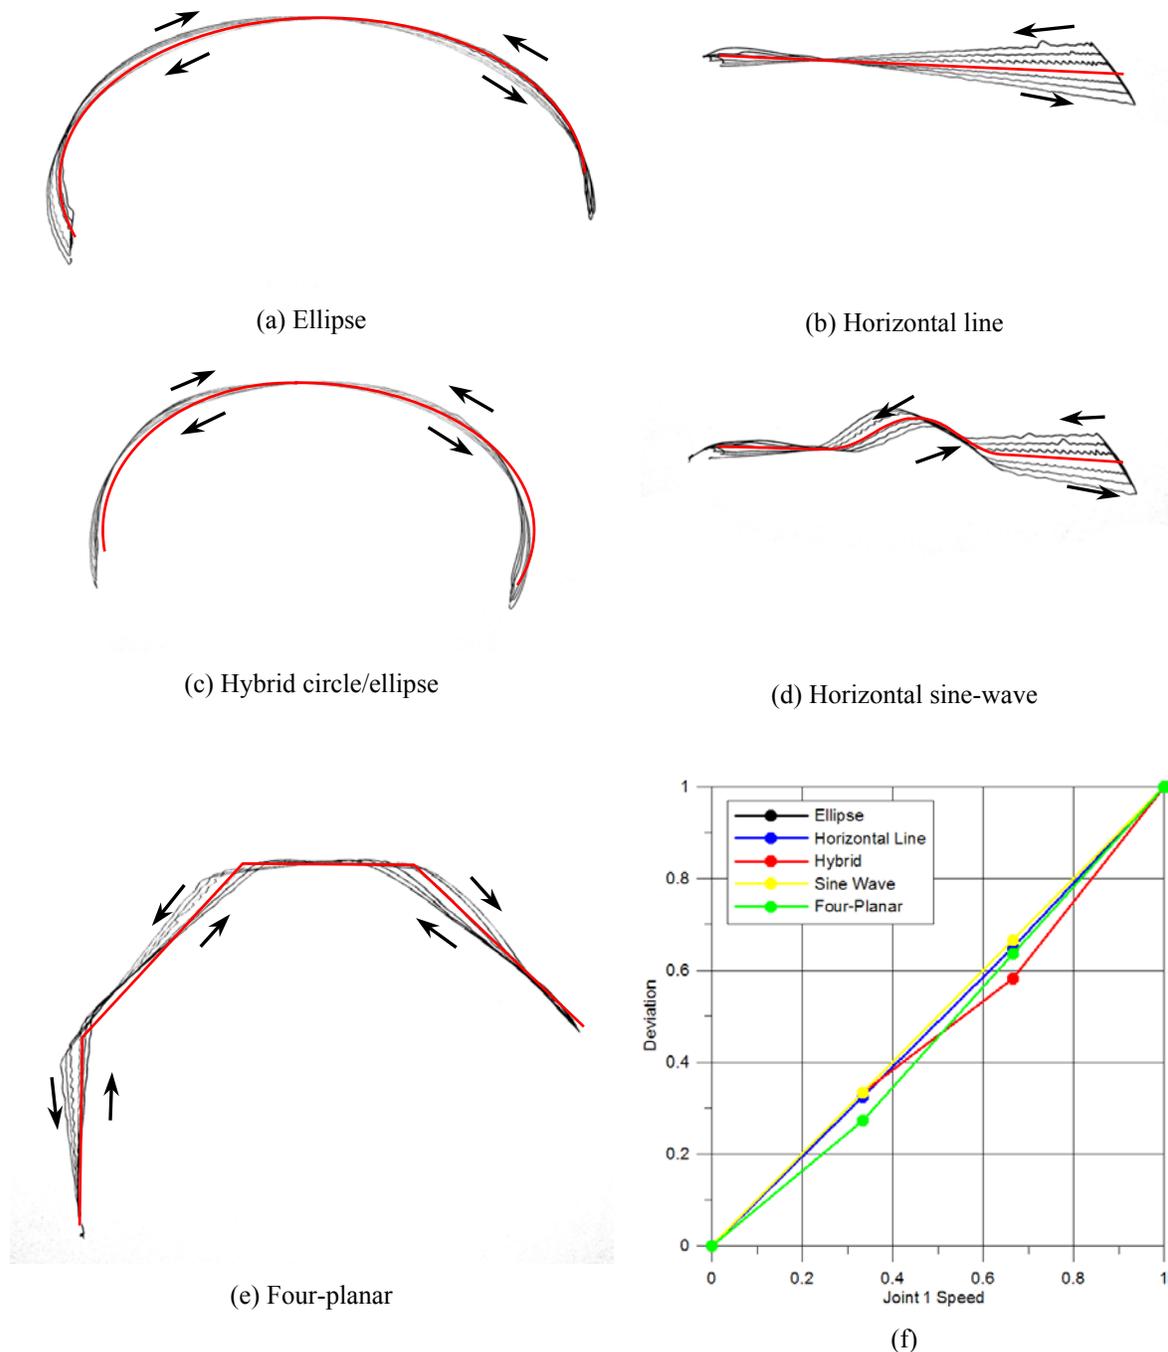
Figure 4-9: *Scanned traces of the hysteresis tests, showing how hysteresis increases with increasing operation speed. The faster motions tend to be the most deviated from the reference surface (shown in red). The step-like motion described in Figure 4-8, is large enough to be visible in some areas. Note that the uneven bumps in some of the traces are likely due to mechanical misalignment between the robot's motor axes and the gear axes. (f) Normalized plot showing how the amount of deviation of the end effector from the intended surface varies linearly with the rotational speed of the robot's first joint. Note that the Four-planar surface is simply an extension of the tri-planar surface used in the user testing.*

base was re-installed. In addition, the second link was fixed to the physical constraint. This made it possible to run the robot automatically, with no user interactions. The same control algorithm was used to position the physical constraint as in Section 3.3, but the first joint was powered at a constant velocity such that the robot traced the virtual surface from one end to the other. The robot was run at three different joint 1 speeds for each surface: 5rpm, 10rpm and 15rpm. The pen traced each run on the same drawing surface.

The resulting traces can be seen in Figure 4-9. Six lines can be seen on each trace. The lower three are the traces created during clockwise motion of the robot, while the upper three are the corresponding counter-clockwise traces. The two middle lines, therefore, were drawn at 5rpm. The next ones outward from the centre were drawn at 10rpm, and the two outermost lines were drawn at 15rpm. This can be plainly seen in (b). The true surface is in the middle of all the six lines. The areas with no or little hysteresis occurred where the physical constraint did not have to move to keep the end effector on the surface. The plot in (f) shows that with increasing speed of the first axis, the hysteresis increases linearly for all of the surfaces.

4.4.3   Sources of hysteresis

It is evident from these results that the hysteresis is caused by the finite response time of the robot. There are two likely sources of this. The first is that the motor driving the physical constraint has a finite maximum velocity and acceleration at which it can move from one point to another. These values were set to the maximum recommended values, which are: 490rpm and 1190rad/s^2 (including the gear reduction). The motor, which is governed by a PID controller built into the motion controllers, was also tuned to react as accurately and rapidly as possible to a step response. During the tuning procedure, a small amount of overshoot was allowed in order to minimize the rise time. The motor's

rise time for a step of 1 degree (including gear reduction) was measured to be approximately 50ms.

The second source of this finite response time is the communication speed between the motor and the computer. A timeline of the entire communication procedure is shown in Figure 4-10. The timeline begins with the moment the computer program queries for the current position of the first axis. The request is sent to a transmission buffer which sends the message through the serial port to the motion controller. The controller returns the position of the first axis in a similar manner. The calculation of the required position of the physical constraint is insignificant and is sent as a move command to the motor. Once received by the motion controller, the move command must be processed and its successful receipt must be confirmed back to the program before actually putting the motor into motion. This is the procedure starting at about 6.5ms and it is unknown why this procedure has such a long delay. The motor only starts moving 9.5ms after the initial move command is sent at the 5ms marker. This is a significant portion of the total time delay.

One entire communication loop takes 18ms, as measured in real time on the computer using the operating system's high resolution performance counter. This, in fact, is only a fraction of the rise time required by the motor as reported above, but is still significant enough to have an effect. However, it is evident that to decrease the hysteresis of the system, the speed of the motor would first have to be increased.
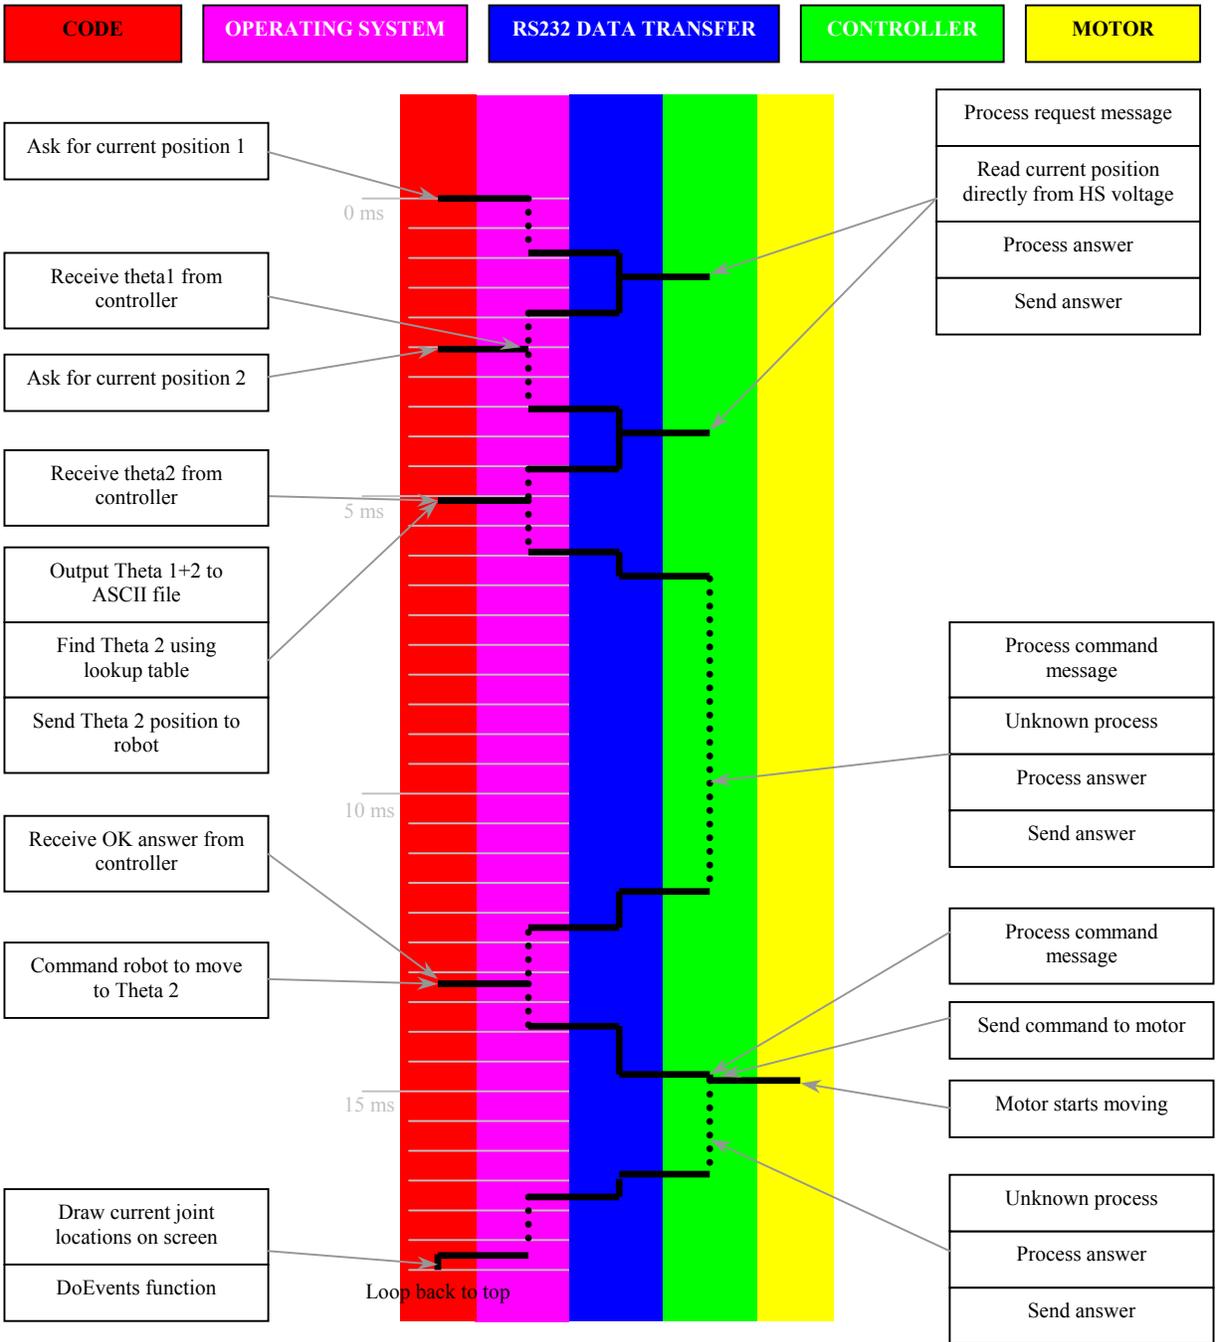
Figure 4-10: *Timeline of the entire communication process between the computer and the robot. The coloured regions represent the software and hardware mediums through which the procedure must travel. Dotted lines represent periods of time whose length is known, but the details of what occurs during these intervals is unknown. All processes are, in general, invariable, unless other uncontrollable processes, such as operating system priority tasks, take effect.*

4.5 Instability analysis

Instability was the most significant problem encountered. The instability was in the form of unexpected vibration in specific regions while tracing along the surface, or while trying to hold the end effector still within those regions. The largest instabilities were seen adjacent to the top right corner of the tri-planar surface (see Figure 4-11) and had a maximum amplitude on the order of 10mm.
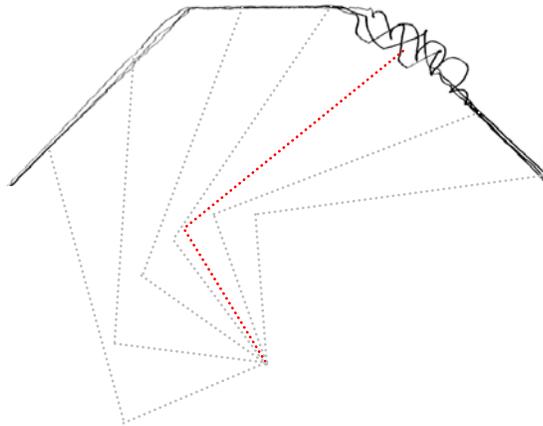
Figure 4-11: *Scan of a sample trace of the tri-planar surface, showing large instability adjacent to the upper right corner. The motion of the robot links along the surface is shown in gray. The red line shows how link 2 is nearly perpendicular to the surface in the unstable region.*

4.5.1 Hypothesis

It was noticed that for all the tested surfaces, the instabilities only occurred when the second link was within 15° from the perpendicular to the surface (see Figure 3-6). As described in Section 3.3, the position of the constraint (let us call it θ_c), is a function of the position of the first link (θ_1). This means that as the user approaches the point on the surface at which the second link becomes perpendicular to the surface, the rotational

speed (let us call it $d\theta_c$) required by the dynamic constraint to prevent incursion goes to infinity.

An illustration of my hypothesis for the instabilities is shown in Figure 4-12. A small incursion produces a modest change in θ_1, but since $d\theta_c/d\theta_1$ is relatively large, this produces a large change in the commanded value of θ_1. This, in turn, produces a strong reaction extending the arm, which may knock the end effector off the surface. Since the user's equilibrium point is still under the surface however, the end effector is once again pushed into the surface, and the cycle repeats.
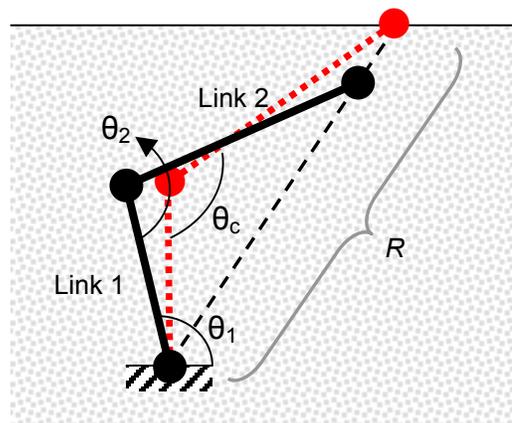
Figure 4-12: *Schematic of a different control strategy that could be used to solve the instability problem in the prototype. It relies on the measurement of both θ_1 and θ_2. Then, using forward kinematics and geometry, R, the radial distance from the base to the virtual surface, can be calculated and used to position the constraint θ_c.*

The above hypothesis has focussed on the RR configuration of the prototype. It would, however, be useful to know whether this instability behaviour is a general property of the dynamic physical constraint concept, or if it is configuration-dependent. I hypothesize that an RP configuration would likely not exhibit this behaviour since the perpendicularity effect would not occur anywhere along the surface. This could be confirmed by building a prototype using this configuration.

4.5.2 Prediction of unstable regions

To help verify my hypothesis that instability occurs in regions where the second link is essentially perpendicular to the surface, a program was written to predict where these regions are located, and their likely intensity. The program uses inverse kinematics to calculate the ratio of $d\theta_c/d\theta_1$ at various points along the surface. Figure 4-13 shows the plots output by the program. Note that all the regions of highest $d\theta_c/d\theta_1$ ratios match exactly the regions of instability on the scans in Figure 3-6. Note also that the maximum ratios for the two surfaces that have planar components are the highest, while the ellipse and hybrid surfaces are significantly lower, which also coincides with the prototype results. This program can be used to optimize the link lengths in order to minimize the intensity of these regions. This will be discussed in the specific context of bone sculpting surfaces in Section 5.9.
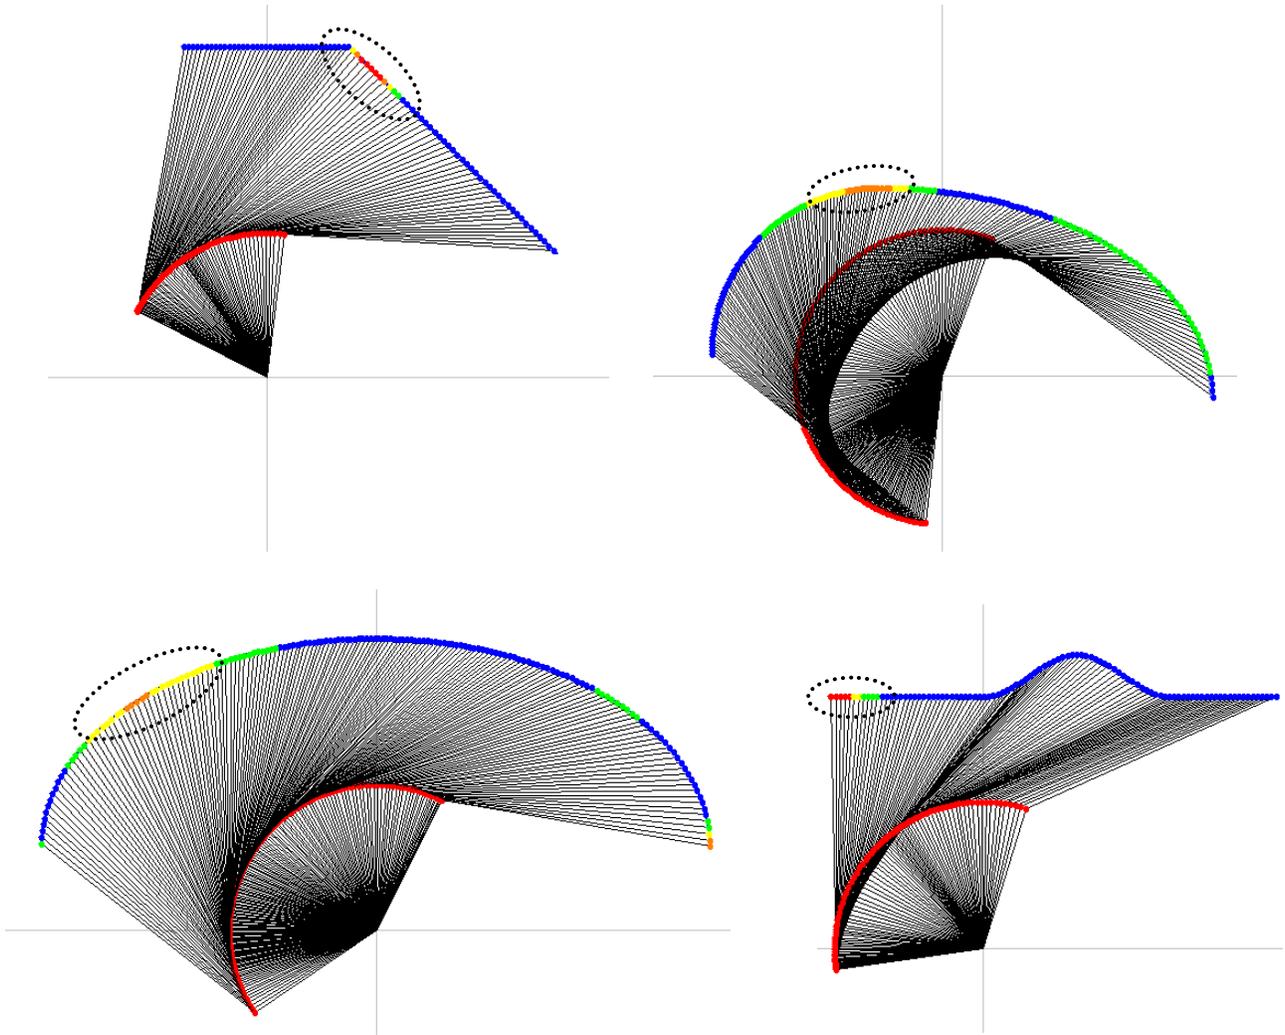
Figure 4-13: *Plots output by the instability prediction program, showing the evolution of the robot's motion along the length of the four different surfaces that exhibited instability in the prototype. The surface has been overlaid with colour intensities, showing the regions of maximum $d\theta_c/d\theta_1$ ratios. The intensity scale ranges from blue (lowest ratio) to red (highest ratio). The maximum velocity ratio for: the planar surface = 9.9, the hybrid surface = 1.0, the elliptical surface = 1.8, and the sinusoidal surface = 9.9.*

4.5.3 Prototype simulation

In an attempt to further explore the physical reasons for the instability issues, mathematical models of both the RP and RR configurations were derived. These models allowed me to simulate the behaviour of the system during surface tracing motions.

4.5.3.1   RP model

The RP model was based on the model of an inverted bar pendulum with variable length. A free body diagram is shown in Figure 4-14(a). The equations of motion for this system are as follows:

$$\underline{\Sigma M = I\ddot{\theta}(t)}:$$

$$\ddot{\theta}(t) = \frac{1}{I}\left(G_x(t)\frac{R(t)}{2}\sin\theta(t) - G_y(t)\frac{R(t)}{2}\cos\theta(t) + F_{ex}(t)\frac{R(t)}{2}\sin\theta(t) - F_{ey}(t)\frac{R(t)}{2}\cos\theta(t)\right)$$

$$\underline{\Sigma F_x = m\ddot{x}_p(t)}:$$

$$G_x(t) - F_{ex}(t) = m\ddot{x}_p(t)$$

$$\underline{\Sigma F_y = m\ddot{y}_p(t)}:$$

$$G_y(t) + F_{ey}(t) = m\ddot{y}_p(t)$$

where, $\ddot{x}_p(t)$ and $\ddot{y}_p(t)$ are calculated from geometry. A block diagram of the control loop used to simulate this dynamic model is shown in Figure 4-14(b). In this control loop, an equilibrium point is fed into the Equilibrium Force block, in which the force applied to the robot by the user is calculated using a spring-damper approximation, as

(a)

(b)
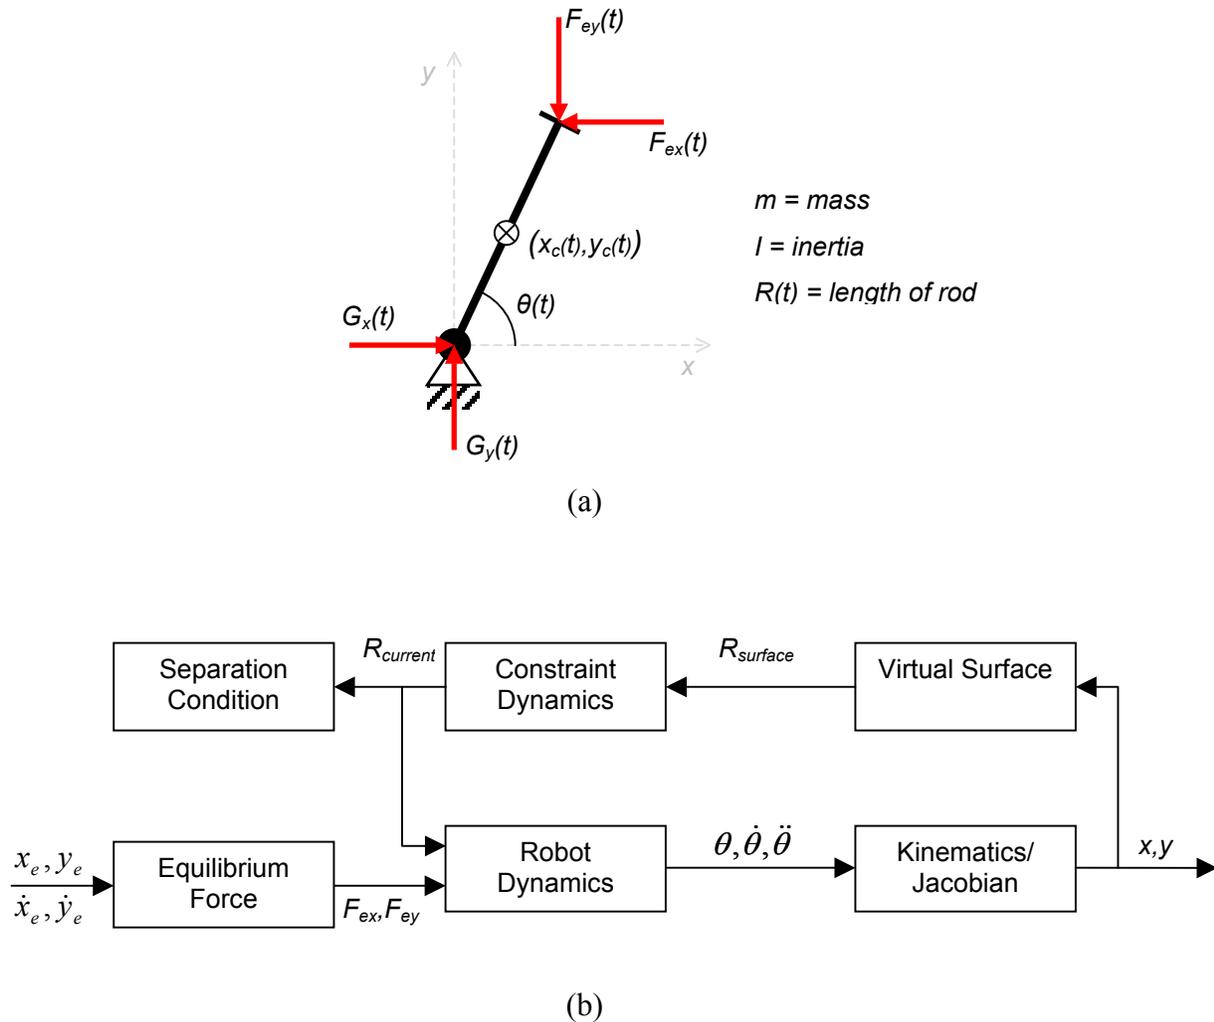
Figure 4-14: *(a) Free body diagram of the RP model. F_e is the force applied by the user and is determined by the equilibrium point location. G is the reaction force from the ground. The length of the rod is variable with time. The inertia of the system is assumed to be constant. The centre of mass (x_c, y_c) of the rod was assumed to be in the middle of the rod. Gravity is ignored because the robot is in a horizontal plane. (b) Block diagram of the control loop used in the simulation.*

described in Section 4.2.1. This force is fed into the robot dynamics, from which the current angular position of the robot about its base can be calculated. A kinematics block outputs the current end effector location. Based on this current location, whether it is on the virtual surface or not, the Virtual Surface block calculates the length of the target radial constraint at this position (ie. the radial distance from the base of the robot to the virtual surface). This is set as the target radius to which the robot must drive the current

end-effector location. The Constraint Dynamics block simulates the motor delays in readjusting the length of the constraint, by applying a PID control loop based on a mass-spring-damper-type model and using the error between the radial distance of target surface and the radial distance of the current end effector position. At this point, the radial acceleration of the end effector is compared to the acceleration of the moving constraint to check if separation between the two has occurred. If so, the simulation is stopped.

In this model, the inertial components of the system are significantly simplified and are held constant, even though the robot changes in length with time. Friction is also ignored. These elements were assumed to be fairly minimal and would most likely not affect the simulation with respect to the instability, hysteretic and deflection elements that I was attempting to simulate.

4.5.3.2   *RR model*

The RR model was based on a point mass pendulum, rather than on a bar pendulum, and is shown in Figure 4-15(a). A point mass model was chosen for simplicity, since the inertial components of the two links would add unnecessary complexity. The dynamics of this model function very similarly to the RP model, except that the inertial component is concentrated in a point mass, rather than a rod. The main difference comes from the definition of the constraint, which in this model is revolute and is defined by the angle θ_2 that, in turn, defines the constrained radial distance between the base of the robot and the virtual surface R.
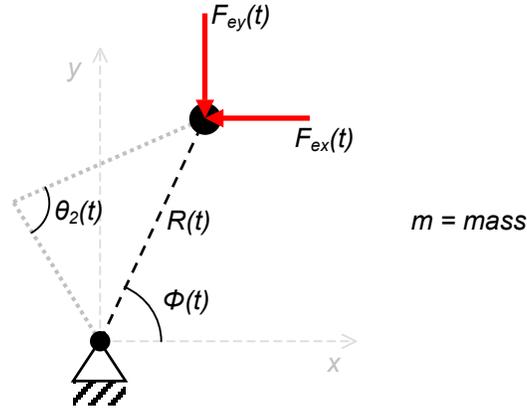
(a)
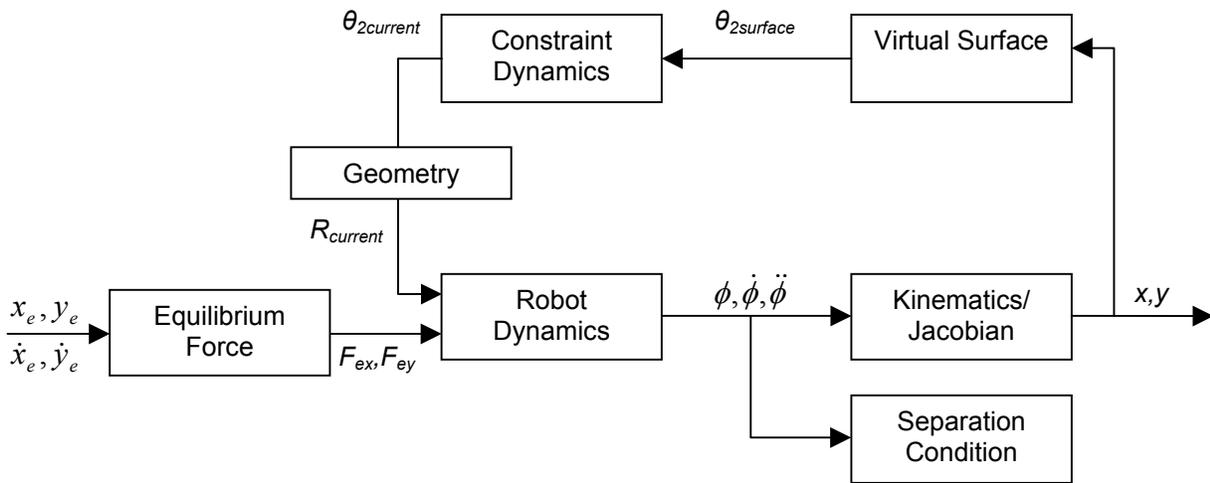
(b)

Figure 4-15: *(a) Free body diagram of the RR model. F_e is the force applied by the user and is determined by the equilibrium point location. Gravity is ignored because the robot is in a horizontal plane. (b) Block diagram of the control loop used in the simulation.*

The dynamic equation for the system is:

$$\underline{\Sigma M = mR^2(t)\ddot{\phi}(t)}:$$

$$\left(-F_{ex}(t)\sin\phi(t) + F_{ey}(t)\cos\phi(t)\right) \cdot R(t) = mR^2(t)\ddot{\phi}(t)$$

In the block diagram of Figure 4-15(b), the Virtual Surface block outputs the target angle of the revolute constraint, θ_2 . A PID block similar to the one used in the RR model is

70

then used to output the current angle of the second joint. This angle is then input into the robot dynamics after geometrically converting it to the radial distance R.

Robot Parameters		Equilibrium Parameters		Constraint Dynamics Parameters		Separation Parameters	
$m =$	0.05 kg	$k_h =$	100 N/m	k_c (RR) =	0.209 Nm/rad	$m_{arm} =$	1 kg
I (RP) =	0.003 kgm^2	$k_v =$	100 N/m	d_c (RR) =	0.013 Nms/rad		
Link 1 length (RR) =	0.035 m	$d_h =$	5 Ns/m	k_c (RP) =	49.388 N/m		
Link 2 length (RR) =	0.065 m	$d_v =$	5 Ns/m	d_c (RP) =	3.143 Ns/m		

Table 4-1: *Parameters used in the RR and RP simulations. m = mass of the robot. I = inertia of the RP pendulum. k_h, d_h, k_v, d_v = horizontal and vertical spring constants and damping coefficients in the equilibrium point model. k_c, d_c = spring constants and damping coefficients used in the PID constraint dynamics model. m_{arm} = approximate mass of the user's arm.*

4.5.3.3   *Implementation and model parameters*

The simulation of the RP and RR models was done in Simulink R2007a (The Mathworks, Inc.). The parameters that were used in the simulation are listed in Table 4-1. The robot inertial parameters were estimated and the link lengths were measured. The spring constants of the equilibrium point model were chosen based on approximate measurements made on the human arm's ability to resist a sudden step force. These measurements were made by dropping a 1kg mass onto a blindfolded subject's arm, flexed at 90°, from just above the hand, and approximately measuring how far the arm dropped. These were found to be on the order of 100N/m in both horizontal and vertical directions and fall within ranges reported in the literature (Mussa-Ivaldi et al., 1985; Rahman et al., 1999). The human arm was assumed to be critically damped and the damping coefficients were chosen correspondingly. The constraint parameters were based on the measured 90% rise time of the motor, as described in Section 4.4.3, and assuming critical damping. The mass of the human was used to determine whether its

acceleration was greater than that of the end effector. This was the condition used for separation of the user's arm from the robot.

4.5.4   Simulation results

The results of the simulations are shown in Figure 4-16 for linear and elliptical surfaces with various equilibrium point paths. The plots show that the simulations successfully emulate the hysteretic behaviour of the prototype. Increasing the rise time of the constraint model or the velocity of the equilibrium point motion increased the amount of hysteresis, in both the RP and RR configurations. This supports my hypothesis stating that hysteresis is caused by the system's response lag. It can also be seen that the amount of hysteresis is the same for both the RP and RR configurations, which makes sense, since fundamentally, the two configurations depend on the same radial distance between the base of the robot and the virtual surface.

The lateral deflection behaviour was also successfully emulated, as shown in Figure 4-16 (e). The same amount of deflection occurred for the RP configuration as for the RR configuration.

The simulations did not, however, show any instabilities such as those experienced with the prototype. Variations in the model parameters, including overdamping, underdamping, increasing the rise time of the constraint model as well as the equilibrium point model, and increasing the speed of the equilibrium point did not result in any instabilities in any of the cases tested. If any instabilities occurred, they were across the
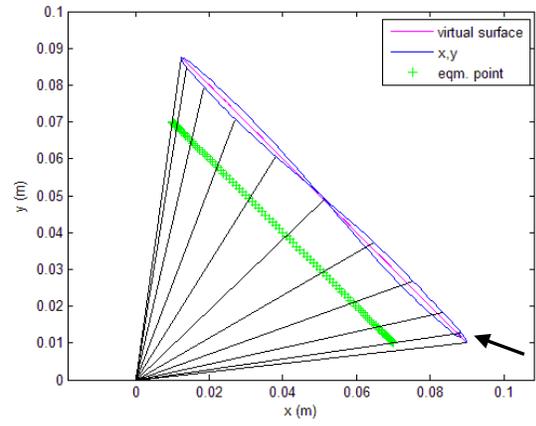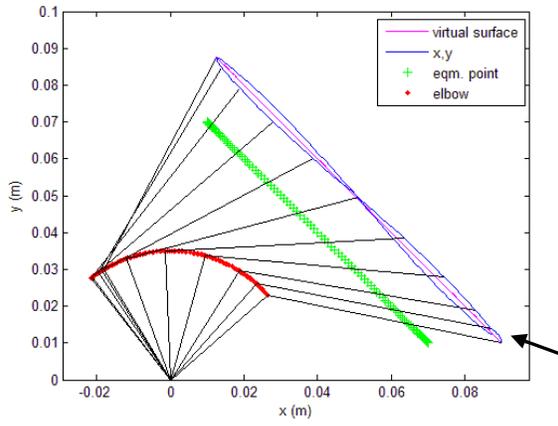
(a)
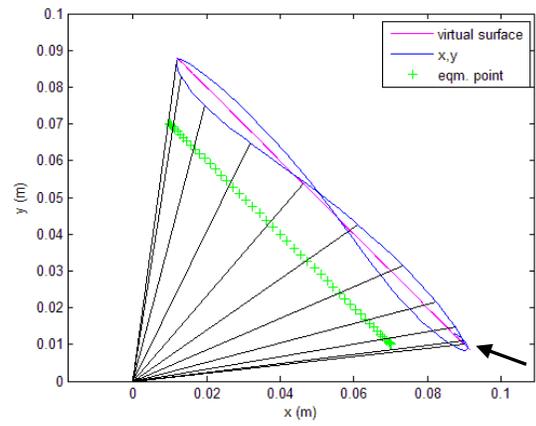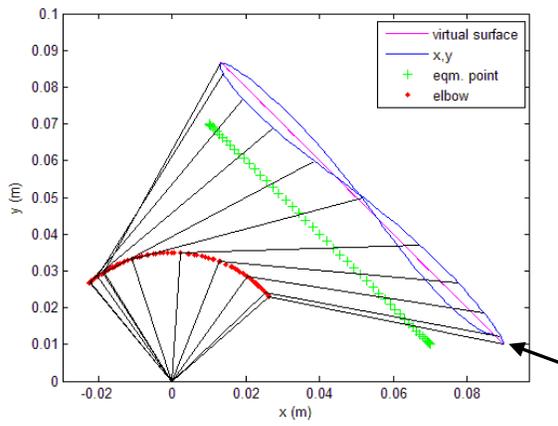
(b)
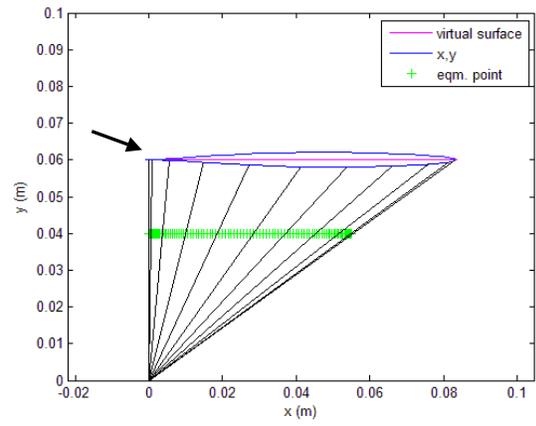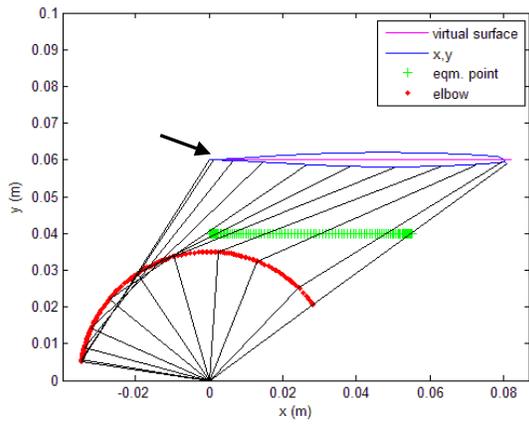
(c)
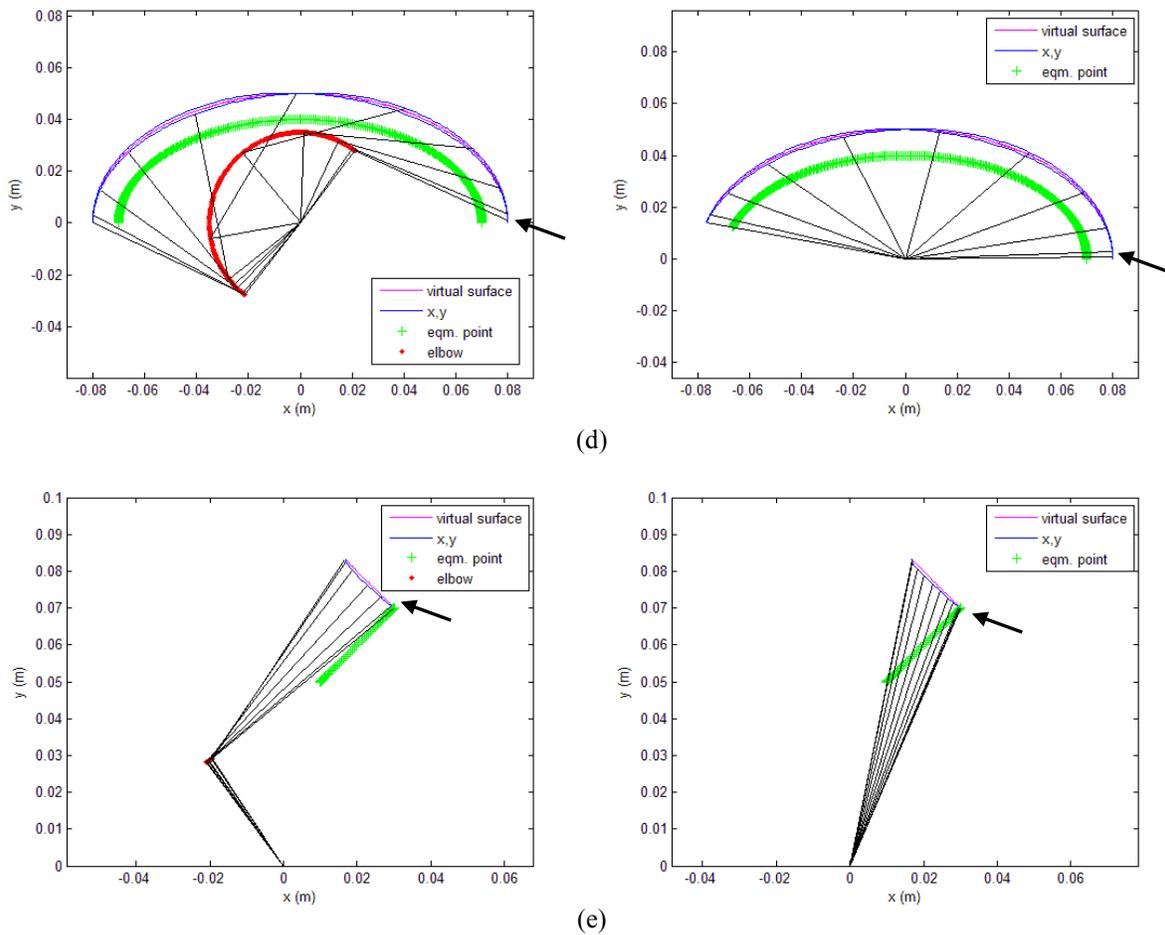
Figure 4-16: *Plots showing the motion of the simulated RR (left plots) and RP (right plots) robot models for various surface types and equilibrium point paths. The beginnings of the trajectories are marked by an arrow. All equilibrium point trajectories followed a sinusoidal velocity profile, starting and ending at rest. (a) Diagonal line emulating the right diagonal surface of the four-planar surface (total trajectory time 2s). (b) Same surface as in (a) but at double the equilibrium point speed (total trajectory time 4s). (c) Horizontal line (total trajectory time 4s). (d) Ellipse (total trajectory time 8s). (e) Same surface as in (a) but with an equilibrium point trajectory perpendicular to the surface and only in one direction (total trajectory time 1s).*

entire length of the surface, not just in specific regions. This was to be expected, since underdamping of the system would cause bounce. These results mean that none of the elements and interactions as they are represented in the model contribute to the instabilities seen in the prototype. Comparing the model to the prototype's control loop, the main difference comes from the fact that θ_c in the model is defined by the radial

74

distance R between the base of the robot and the point on the virtual surface to which the end effector is being constrained. In the prototype, however, θ_c is defined by θ_1. As in my hypothesis, this dependence on the first joint makes the second joint very sensitive in regions where the ratio of $d\theta_c/d\theta_1$ is large. Since this sensitivity does not exist in the model, the instabilities are not to be expected.
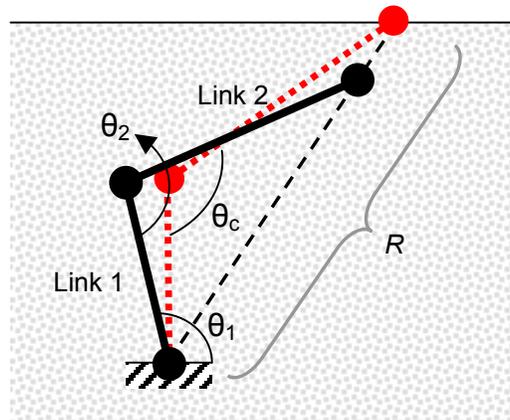
Figure 4-17: *Schematic of a different control strategy that could be used to solve the instability problem in the prototype. It relies on the measurement of both θ_1 and θ_2. Then, using forward kinematics and geometry, R, the radial distance from the base to the virtual surface, can be calculated and used to position the constraint θ_c.*

4.5.5 *Proposed solution to the instabilities*

Based on the analysis above, I can propose a solution to the instabilities by changing the way the prototype is controlled. Referring to Figure 4-17, this would involve measuring the location of the free-moving second link (θ_2). Applying forward kinematics, the current position of the end effector (whether in contact with the virtual surface or not) can be calculated online. Then, using geometry, the radial distance between the robot base and the virtual surface (R) can be calculated and used to position the constraint θ_c. This control method would eliminate the sensitivity of the constraint to changes in θ_1. To test the effectiveness of this solution would require placing an encoder on the second link. In terms of further analysis of the instabilities in the system, the next step would be to try to

explain them using ideas from control theory, such as root locus diagrams. Unfortunately, neither could be done within the time constraints of the thesis, but they will be part of the future work on this project.

4.6   Analysis conclusions

In summary, I can recommend the following solutions to the four major behavioural issues displayed by the prototype:

- Lateral deflection and reverse stickiness: optimize the location of the robot's fixation point with respect to the virtual surface.

- Hysteresis: decrease the response time of the system to the user's actions by increasing the hardware communication speeds, by using a motor with a higher saturation speed and by reducing the gear ratio on the motor shaft.

- Instability: either optimize the link lengths for a set of specific virtual surfaces, such that the second link never approaches the perpendicular of the surface, or place an encoder on the second link and apply the control procedure proposed in Section 4.5.5.

In the previous two chapters, I presented a simple prototype used to test the dynamic physical constraint concept. Prototype testing showed that the concept is indeed an effective method of emulating a realistic hard surface, despite exhibiting a few behavioural complications that can either be addressed by design or easily learned and accommodated by the user. The overwhelmingly realistic feel and reasonable precision of this very simple prototype in my view justifies its further development and application to a specific task in surgery, namely bone sculpting for arthroplasty procedures.

5.0 IMPLEMENTATION OF THE DYNAMIC PHYSICAL CONSTRAINT CONCEPT TO SURGICAL BONE SCULPTING

Bone sculpting is the surgical procedure of shaping bone surfaces in preparation for the placement of orthopaedic implants. The typical tools used in such a procedure are an oscillating saw or a mill. The benefit of the former is its ability to do rapid planar cuts, and that of the latter is to do more accurate and complex cuts. The vast majority of commercial implants have a bone-mating surface whose profile is based on a single standard geometric shape, or a set of these shapes. Typical shapes include planes, partial cylinders or partial spheres. One of the main reasons for the use of these simple shapes is the ease with which they can be replicated on the mating bone surface. Surfaces that reproduce the natural shape of bone structures are currently too complicated to sculpt. In fact, no simple, reliable, inexpensive and universal bone sculpting technique currently exists that would justify the development of more complex implant-to-bone mating surfaces.

This chapter applies the new dynamic physical constraint hard surface emulation concept specifically to the design of a miniature distal femoral bone sculpting mechanism. The importance of bone sculpting is first explained, followed by an analysis of existing bone sculpting systems. The concept is then described within the setting of bone sculpting, including the optimization of the prototype described in the previous chapter.

5.1   The importance of bone sculpting

A generally accepted trend in orthopaedic surgery today is the conservation of tissue and the reduction of tissue damage during operations. Minimally invasive surgery (MIS) is the most familiar example related to this trend. MIS involves carrying out an operation using an exposure smaller than in standard procedures. Another example, more directly related to orthopaedic surgery, is the conservation of bone during the placement of implants. The primary reason for conserving bone is that implants do not last forever and revision procedures are made easier if there is more bone stock available to work with at the time of the revision. Proof of the current interest in this trend can be seen by an examination of product brochures and publicity campaigns for the majority of the currently available orthopaedic implants, where the term "bone-conserving" is commonplace.

This trend towards reducing the impact on existing tissue has, in fact, gained sufficient popularity as to prompt the organization of the first Consensus Conference on Tissue Sparing Surgery, which was held in June 2006 in Rapallo, Italy. The principal organizer of the conference introduced a list of rules that govern tissue-sparing procedures (Pipino, 2006). Two of the rules are of specific interest to the present discussion:

- A prosthesis integrates in the joint and does not substitute it.

- Mini-prostheses are used to maximize the conservation of bone stock.

The first rule implies that a prosthesis should be designed in a manner that is the least invasive towards the natural tissue as possible. An extension of this leads to the second rule which states specifically that the smallest possible prostheses should be used, while considering stability, adhesion and distribution of bone stress. To date, there is very little documented quantitative support of the benefits or drawbacks to such bone-conserving implants. In the year and a half since the conference, however, several authors have begun addressing the issue in the literature. Although the existing papers are primarily

based on opinion, rather than fact, they show the emerging importance of the tissue-sparing principle in modern orthopaedic surgery.

Pipino (2006) describes the natural evolution of tissue-sparing philosophy as being driven by the goals of reducing surgical aggression and the optimization of patient postoperative recovery. Rossi et al. (2007), while supporting the movement, take a more cautionary stance and warn that smaller may not always be better, depending on each particular patient and situation. Prostheses should first and foremost ensure the elimination of pain and the restoration of normal anatomical function. In a separate discussion, Engh (2007) approaches the subject from the point of view of its benefits for younger patients. Focussing on knee implants, he states that younger patients with highly active lives require more specialized approaches to knee surgery, including customized instruments and implants specific to their unique anatomy. He states further that, in the future, knee arthroplasties will involve a significant reduction in bone removal: current TKA implants remove approximately 20mm of bone, including both the femoral and tibial implants. Potential reduction of their size is limited by stress concentrations at the junctions of the planar surfaces that make up their bone-mating underside. To increase bone conservation, he implies that there is a need for the development of more complex, non-planar resurfacing implants.

As mentioned above, a major reason for the use of bone-conserving implants is in anticipation of the possible need for revision surgery. The removal of a primary implant can cause excessive bone loss, causing a significant increase in the complexity of the revision surgery. This complexity includes difficulties in alignment, bone fixation, soft tissue balancing, implant support and, to a lesser extent, increased potential for bone fracture (Haas et al., 2000; Whiteside, 2004; Backstein et al., 2006). Conserving more bone in the initial operation would, therefore, naturally be of benefit in case of the need for a revision. This is especially important for young and active patients who are more likely to require revision surgery later in their lives (Robertsson et al., 2001; Rand et al., 2003; Gioe et al., 2004; Eickmann, 2006; Collier et al., 2006). The number of hip and

knee replacements in male and female patients under the age of 65 saw a 10 year average increase of 22% in 2004-2005 in Canada (Canadian Institute for Health Information (CIHI) report, 2006).

As the above discussion implies, the design of bone-conserving implants requires the ability to create more complex, anatomically-similar bone-mating surfaces. From an engineering standpoint, the use of such curvilinear surfaces would have the additional benefit of reducing the stress concentrations found at the corners of planar surfaces. In the past, these have been found to cause implant failures (Wada et al., 1997; Swartz et al., 2001; Panousis et al., 2004; Campbell et al., 1998), as shown in Figure 5-1. The risk of such failures limits the amount by which the thickness of these planar implants can be reduced, and consequently the amount of bone conservation.
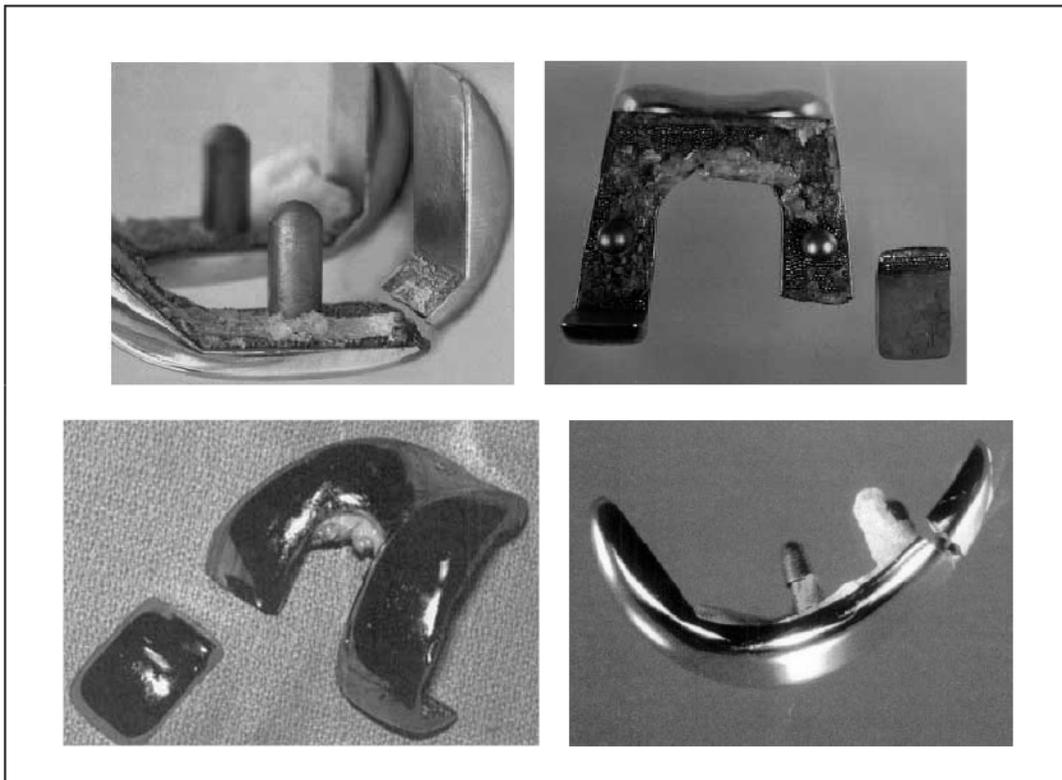
Figure 5-1: *Examples of implants fractured at regions of high stress-concentration. (Wada et al., 1997; Swartz et al., 2001; Panousis et al., 2004; Campbell et al., 1998).*

Although the need and interest is certainly present, the intra-operative complexity of sculpting such surfaces has so far deterred most advances in research and development in this field. A possible additional deterrent to a modified approach is the natural reservation of surgeons to accept new technology, especially robotic and computer-assisted technology (Breisch, 1998; Davies, 2000; Bargar, 2007). Such resistance can be overcome only if newly-developed devices are able to closely replicate the freedom and agility of instruments which surgeons are currently accustomed to using, while adding the precision of robotics. The successful development of bone-sculpting tools would then likely spur increased interest in the development and acceptance of bone-conserving orthopaedic implants.

5.2 Why focus on distal femoral bone sculpting?

Bone sculpting is a broad topic in orthopaedic surgery and finds its use in many procedures. Hip and knee arthroplasty are of primary importance because of how frequently they are performed (according the CIHI, 33,590 knee replacements were reported in 2004-2005 in Canada, 30% more than hip replacements and with a 10 year change of 125%). The work described in this chapter is focussed specifically on distal femoral bone sculpting for unicompartmental knee arthroplasty (UKA) because the anatomy of the knee allows relatively easy access and the natural shape of the distal femoral head is well adapted to curvilinear shapes. Additionally, increasing interest can be seen in UKA procedures for reasons such as tissue (bone and ligament) and physiological conservation, faster recovery, increased range of motion and low cost compared to TKA procedures (Riaz and Maistrelli, 2005).

Non-planar femoral UKA implants already exist commercially, although their non-linearities are typically only two-dimensional (ie, the prepared bone surface is a circular

or elliptical cylinder or some other compound cylindrical shape). My collaborating company, Praxim, feels that there is significant value in initially solving a two-dimensional problem for the new concept as it would be a relatively simple way of determining the strengths and limitations of the concept in a surgical context, while also ensuring that existing techniques exist as points of comparison and validation. It would also provide a good starting point as to the types and sizes of surfaces the mechanism should be capable of handling. The concepts developed from this two-dimensional problem could then be naturally extended to other, more complex problems. For example, total knee and total hip arthroplasties are procedures where three-dimensional non-linearities are unavoidable and additional complexities such as constrained access, small implant sizes and challenging mechanism fixation, come into play. The potential for future expansion of the concept to more complex situations will be discussed in the concluding chapter.

5.3 *Existing curvilinear implant bone-mating surfaces*

Several implants with non-planar bone-mating surfaces exist commercially in UKA. These, along with a typical tri-planar implant, are shown in Appendix D. The reported dimensions are based on approximations obtained from available images and implant sizing charts and give a general idea of their sizes. The Biomet Oxford implant is supplied with appropriately-sized concave spherical milling bits that are used to mill the bone surface to the appropriate shape. The Wright Advance implant uses a metal template and a square end mill to prepare the bone surface. The planar Zimmer M/G implant uses a slotted template to guide cuts made with an oscillating saw. All these techniques are efficient and require minimal guess work by the surgeon. They, however, each require their own set of specific tools to match a variety of implant sizes and require modifications to be used in conjunction with computer-assisted surgical systems.

The more complex Biomet Repicci implant is designed to imitate the natural anatomical shape of the femoral condyle, and can be approximated by a combination of circular and elliptical shapes. The surgical procedure for preparing the bone surface for this implant involves a trial-and-error procedure using a hand-held burring tool. The amount of bone surface removed during this procedure is small. However, one can imagine the difficulty of using hand-eye coordination to mill the entire articular surface of the knee for a similar TKA implant. In fact, true bone-conserving implants would essentially be resurfacing implants that only replace the damaged articular cartilage to a maximum depth of about 2mm, as in the Repicci implant. This would eliminate the possibility of using physical templates or shaped milling tools such as those described above.

5.4   Existing bone-sculpting systems

The use of robotics in orthopaedic surgery has received significant attention for the last two decades. The first successful robots to be developed, such as Robodoc and Caspar, were autonomous robots that performed surgical cuts based on a pre-operative plan. These robots have been classified as "active" robots due to their powered joints. More recently, however, with the development of new haptic techniques such as those described in Chapter 1.0 , a different class of robots has emerged, called semi-active or synergistic robots. These robots work with the surgeon to provide accuracy and safety. The surgical procedure remains in the hands of the surgeon, but the robot either acts as a cutting constraint or provides motion guidance. A further class of robots, that could be called active-passive robots, autonomously place a passive cutting guide in pre-calculated positions.

Some of the primary bone-sculpting robots developed to date are shown in Figure 5-2. A list of some comparative elements of each robot is shown in Table 5-1, to help evaluate the state-of-the-art. The majority have been designed for knee arthroplasty. Most have

been clinically approved, which is a good sign that shows that acceptance of robotics in the operating room (OR) is growing. Clinical approval, however, does not guarantee success. Initial cost is a significant drawback of these systems, since their ultimate cost-effectiveness over time is difficult to establish (Bargar, 2007). Although most system costs are either undetermined as of yet, or not publicized, the majority shown in Table 5-1 can be estimated (based on their size and complexity) as likely being greater than half a million US dollars. Size is also an important factor, as the majority of surgical robots are large, free-standing machines that are difficult to move and take up significant OR space. Another related aspect is the complexity of the design which is often imposing and contains many hidden parts and functions. As with any mechanical concept, greater complexity provides more opportunities for potential malfunction. The intrinsic safety of the concepts listed in the table was evaluated by listing those aspects of the system over which the surgeon has direct control, as opposed to those that are automatic. System complexity was roughly quantified by the approximate number of mechanical parts making up the entire system (taken primarily from patent documents). The intrinsic safety of the concepts listed in the table was evaluated based on whether the following two aspects of the system are controlled directly by the surgeon or electronically: the tool tip motion and the restriction of tool tip motion to a safe region or a specified path. In general, autonomous robots have both aspects controlled electronically, semi-active robots restrict tool motion electronically while cutting motions are generated by the user, and active-passive robots use mechanical motion restriction with all cutting done by the user. As is standard in FDA and CE marking approvals, the greater the electronic role in the system, the greater the risk. The cutting accuracies were obtained or calculated from various sources to give an idea of what kind of precision has been attained to date using these robots.
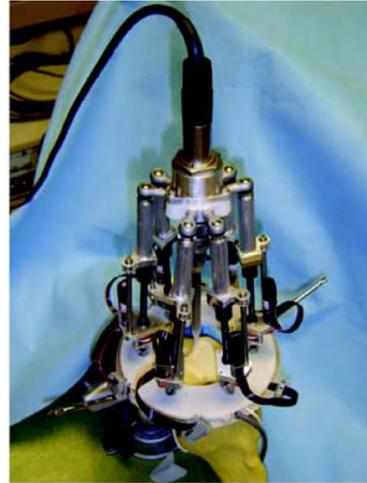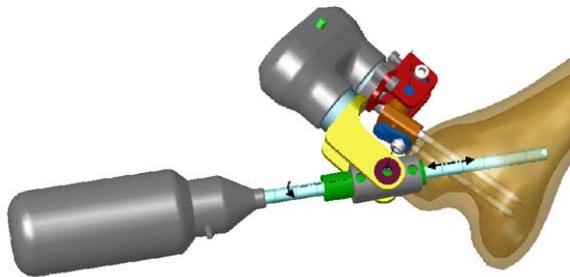
Figure 5-2: *Existing bone-sculpting robots. First row: active robots. Second row: semi-active robots. Third row: active-passive robots.*

ROBOT NAME	Robodoc	MBARS	ACROBOT	PFS	Mako HGS	Praxiteles	BRIGIT
Robot class	Active	Active	Semi-Active	Semi-Active	Semi-Active	Active-Passive	Active-Passive
Description	Autonomous SCARA-type robot that performs cuts based on pre-operative plans.	Parallel bone-mounted mini robot for autonomous milling.	Backdriveable robot that uses active constraints to limit motion to a specific region.	Hand-held milling device that retracts its blade when it exits a predefined cutting region.	Force-feedback haptic arm robot used to control cutting within a small incision.	Bone-mounted mini robot that autonomously positions a cutting guide.	Articulated robot that autonomously positions a cutting guide.
Type of Surgery	THA, TKA	PFA	UKA, TKA	TKA, L	UKA	TKA	TKA
Surface Type	Any	Any	Any	Any	Any	Planar	Planar
Current Stage of development	Clinical (since 1992)	Research	Clinical (since 2001)	Research	Clinical (since 2006)	Clinical (since 2007)	Awaiting clinical approval
Cost (US $)	$600,000	Unknown	>>	Unknown	>>	$220,000	>>
Size (cm^3)	>>	3000	20,000	100	70,000	80	20,000
Operation time (min.)	90 ^{D}	Unpublished	104	Unpublished	Unpublished	Unpublished	Unpublished
Complexity(# parts)	>>	100	>>	20	100	20	>>
Cutting Accuracy (mm)	< 1	Unpublished	< 2	1.5	Unpublished	0.5	< 1
Intrinsic safety of concept	*	*	**	***	**	***	***

Table 5-1: *Comparison of existing bone-sculpting robots. PFS = Precision Freehand Sculptor , HGS = Haptic Guidance System, THA = Total Hip Arthroplasty, TKA = Total Knee Arthroplasty, UKA = Unicompartmental Knee Arthroplasty, PFA = Patellofemoral Arthroplasty, L = Laminectomy. The symbol ">>" means the value is unknown but large compared to the others. For cost, ">>" means greater than half a million US dollars.*

	Advantages	Disadvantages
Robodoc	- precise - any surface possible - can be used in any part of body	- fully automated, therefore surgeon is left out of control loop - large - expensive - requires rigid fixation of limb
MBARS	- small - bone-mounted - stiff, due to parallel architecture	- fully automated, therefore surgeon is left out of control loop - not viable for semi-active control - complex parallel architecture - force sensors required for active control
ACROBOT	- any surface possible - precise - can be used in any part of body - can change rigidity of contact	- large - requires rigid fixation of limb - simulates an elastic, rather than a rigid surface - relies on motors to restrain motion - requires force feedback - requires large motors to counter user forces
Precision Freehand Sculptor	- small - hand-held - can be used for MIS - can be used in any part of body	- makes rough cuts - precision dependent on optical tracking and cutter retraction speed
Mako Haptic Guidance System	- any surface possible - can be used in any part of body	- mechanical rigidity difficult to obtain - relies on motors to restrain motion - requires rigid fixation of limb - relatively large - expensive
BRIGIT	- precise - can be used in any part of body	- limited surfaces shapes possible - large - requires rigid fixation of limb - expensive
Praxiteles	- precise - small - inexpensive - bone-mounted - simulates hard surface - no force sensors needed	- limited surface shapes possible

Table 5-2: *Advantages and disadvantages of existing bone-sculpting robots.*

Table 5-2 summarizes the specific advantages and disadvantages of the various systems, based on the elements in Table 5-1. The following section uses the results presented in both tables to create a list of objectives for the design of a haptic bone sculpting system.

5.5 Design objectives for bone-sculpting

The goal of this chapter is to show how the dynamic physical constraint hard surface emulation concept was used to create a semi-active bone-mounted sculpting tool. The design objectives outlined below were used to assess the new system.

- **Preservation of the Praxiteles architecture:** since the project was initiated by Praxim's interest in expanding the Praxiteles architecture to curvilinear bone sculpting, an important aspect of the project was the conservation of this simple architecture.

- **Semi-active design:** to ensure the inherent safety of the robot, it was important that the cutting tool motion be directly controlled by the surgeon. The robot would, however, provide motion constraint in the form of a virtual template.

- **Freedom of motion:** when not in contact with the virtual template, the user should be virtually unimpeded with minimal apparent friction from the device. The system should ideally not require force feedback or friction and gravity compensation.

- **Precision:** the simulation of the virtual template should be similar in rigidity to that of a physical template. This is both for the realistic "feel" of the system but more importantly for the precision of the system. Collision, constant applied forces, and tracing of the surface should all be done with a
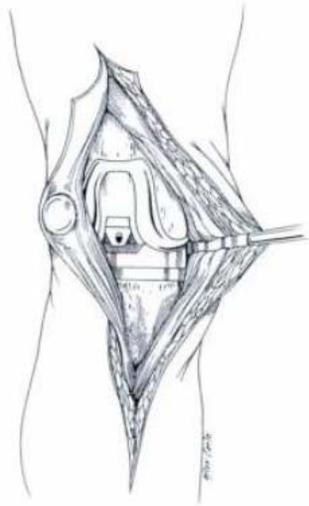
TKA Standard Medial
Parapatellar Approach
(Callaghan, 2003)

TKA MIS Subvastus Approach
(Pagnano and Meneghini, 2006)
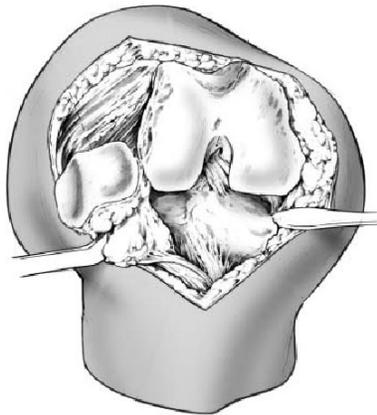
TKA MIS Midvastus Approach
(Hube et al., 2002)
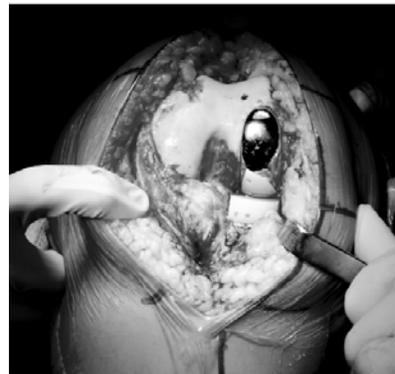
UKA Medial Parapatellar
Approach (Sah and Scott, 2007)

Figure 5-3: *The most common approaches to surgical exposure of the knee for TKA and UKA.*

maximum over- and under-penetration limited to 1.0 mm, as is the case in the best of the systems described in Section 5.4.

- **Smooth surface following:** tracing of the virtual template should be done with the same freedom of motion as in the third constraint above.

- **Minimum obstruction:** the design should not obstruct the surgeon's view and access to the surgical site and must be able to function within the space defined by a typical surgical exposure. Typical surgical exposures of the knee during TKA and UKA are shown in Figure 5-3. Exposures vary from a 25-30 cm incision in the standard TKA medial parapatellar approach to an 8-10 cm incision in the UKA medial parapatellar approach. In UKA, access to the femur is limited to either the medial or the lateral sides depending on which condyle is being replaced.

5.6 *Description of the dynamic physical constraint bone-sculpting system*

The proposed system architecture is essentially based on the prototype described in Chapter 3.0 and is based on a two-link manipulator arm. A CAD drawing of the system is shown in Figure 5-4(c).   In brief, it consists of the motor and gear reduction components of the original Praxiteles (Figure 5-4(a)) as the first link and a rotational physical constraint assembly as the second link. The robot would be mounted like the original Praxiteles, on the medial or lateral side of the femur, depending on the condyle involved in the operation, as shown in Figure 5-4(b).

Two distinct milling configurations can be envisioned and are graphically represented in Figure 5-5. The first is a peripheral milling configuration in which a cylindrical milling bit is oriented normal to the sagittal plane. The second is a top-milling configuration in which a ball-mill is oriented in the sagittal plane. The former has the advantage of
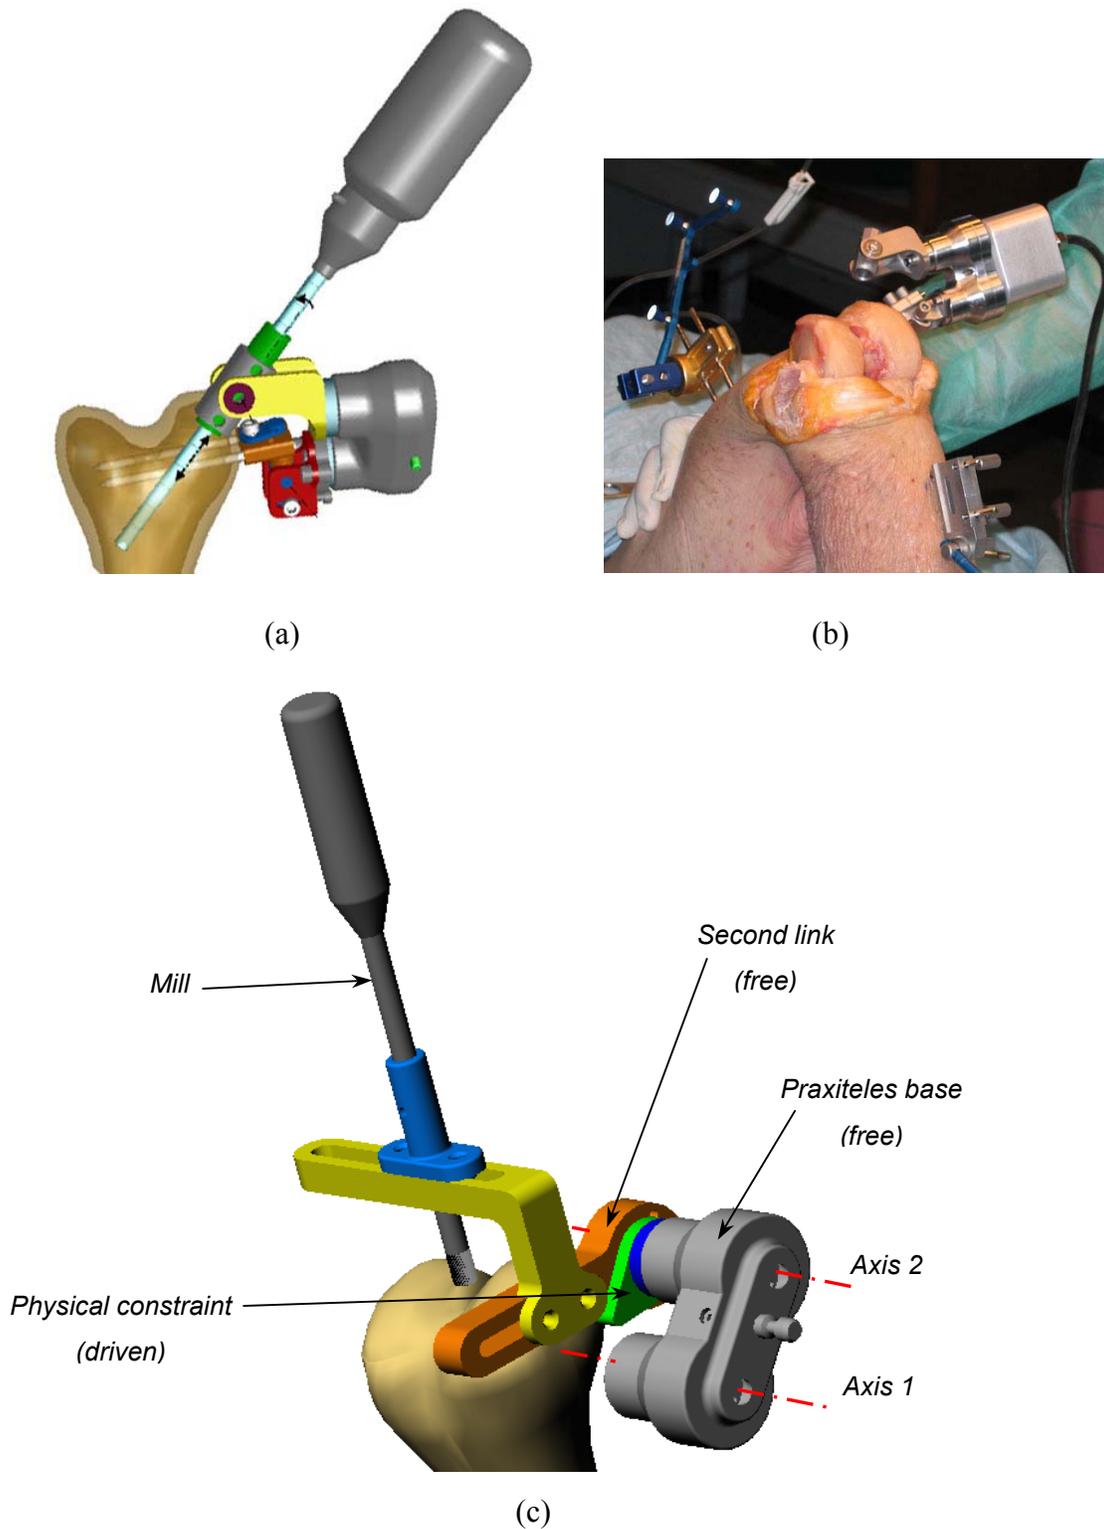
Figure 5-4: *(a) A CAD drawing of the original Praxiteles. (b) A photograph of the original Praxiteles system mounted on the lateral side of the distal femur during cadaver trials. (c) A CAD drawing of the proposed bone sculpting system. Axis 1 denotes the axis about which the robot is attached to the bone (the attachment mechanism would be similar to those shown in (a) and (b)). Axis 2 denotes the axis about which the physical constraint is activated.*

allowing a cylindrical surface to be cut in a single pass. Unless a shaped cutter is used, it is, however, limited to two-dimensionally varying surfaces. Since access is from the side, the peripheral milling configuration is less invasive than the top-milling configuration which requires full bone exposure from the top. The former also has the advantage of decreasing the bending moment applied by the user on the base of the robot where it is fixed to the bone (see Figure 5-5).

Peripheral milling
configuration

Top-milling
configuration
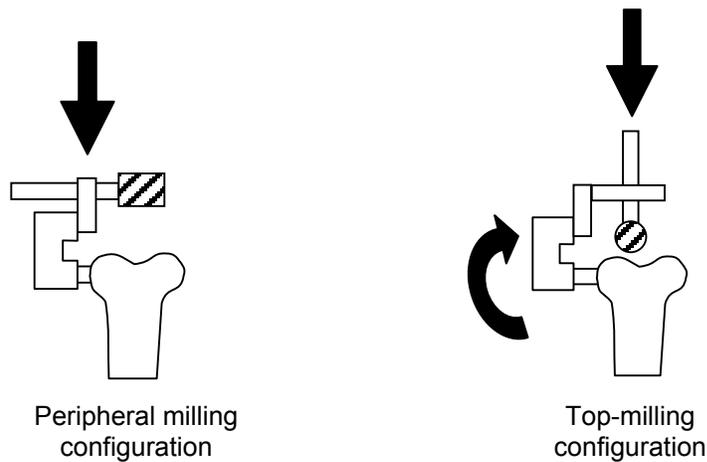
Figure 5-5: *Two possible milling configurations, showing the effect of user forces on the tool.*

Although the peripheral milling configuration appears to be advantageous, it was decided that the top-milling configuration was more suitable for initial evaluation. The major advantage of the top-milling configuration is that it allows for multi-pass contoured cutting of three-dimensionally varying surfaces. It also decreases the interference of the robot architecture with the surgeon's milling motions, since the mill is further away from the robot architecture. The minimal invasiveness of the peripheral configuration was also not deemed to be a significant advantage, seeing as the typical surgical exposures during knee arthroplasties use top rather than side approaches (Figure 5-3). It was felt that the versatility of the top-milling configuration outweighed the advantages of the peripheral milling configuration.

Based on this decision, the mill would be fixed to the second link through an adjustable connector (yellow piece in Figure 5-4(b)) in order to orient it in the sagittal plane. Its location in the medio-lateral direction with respect to the bone, would need to be monitored in real time using a navigation system. This location would determine which two-dimensional virtual contour would be implemented in the control software.

Finally, it should be noted that rigid fixation and precise calibration are both important aspects that would affect the overall sources of error in the system. These issues, however, are not the focus of this thesis and would ideally be solved using methods already developed for the original Praxiteles.

5.7 Design assessment

The bone sculpting system can be assessed by examining the properties displayed by the prototype, as described in Chapter 3.0 :

- The effect of a hard virtual surface was emulated very successfully.

- The stiffness of the system was truly realistic with a satisfying "click" at the moment of contact which increased the effect of a hard constraint.

- When not in contact with the virtual surface, motion appeared frictionless.

- Surface tracing was also virtually frictionless, smooth and satisfactory.

- The small size and simplicity of the system ensures minimal obstruction of user motions.

The surfaces used during prototype testing were all scaled-up versions of the existing UKA surfaces listed in Section 5.4. Scaled surfaces were used because the prototype did not allow for realistic sizes to be tested due to interference between the physical

constraint and the robot base. The use of larger surfaces magnified the motions of the robot, making it easier to identify the concept's issues. However, as shown in the analysis of the prototype in Chapter 4.0 , the lateral deflection, hysteresis and instability issues are all size and motion-dependent. In other words, the tighter the surface wraps around the robot's attachment point, the smaller the motions and the smaller the effect these issues have on the system.

The design currently successfully meets all the design objectives except for precision. Since the precision of the prototype was also motion dependent and resulted from insufficiently rapid hardware response time, it was felt that in a slow and controlled milling situation, and using appropriate hardware, the precision of the system could be dramatically increased. This miniature bone-mounted device, therefore, seems to be a feasible and potentially advantageous method for distal femoral bone-sculpting.

5.8   *Optimization of attachment location*

As described in Chapter 4.0 , the amount of lateral deflection due to the equilibrium point effect is influenced by the attachment location of the first axis of the robot with respect to the surface being emulated. The optimization procedure described in Section 4.2.3 was used to determine an optimal fixation position on the bone for the circular and Repicci unicondylar implant shapes listed in Appendix D. The results are shown in Figure 5-6 in the form of colour maps. These maps show how the maximum lateral deflection across the length of the surface varies with respect to the attachment point for a number of equilibrium point surface penetrations. Since geometric scaling of the implant does not affect the deflection calculations, the x- and y-axes are normalized with respect to the surface's antero-posterior (AP) width. Therefore the results shown apply to all implant sizes. The effect of the amount of equilibrium point penetration on the amount of
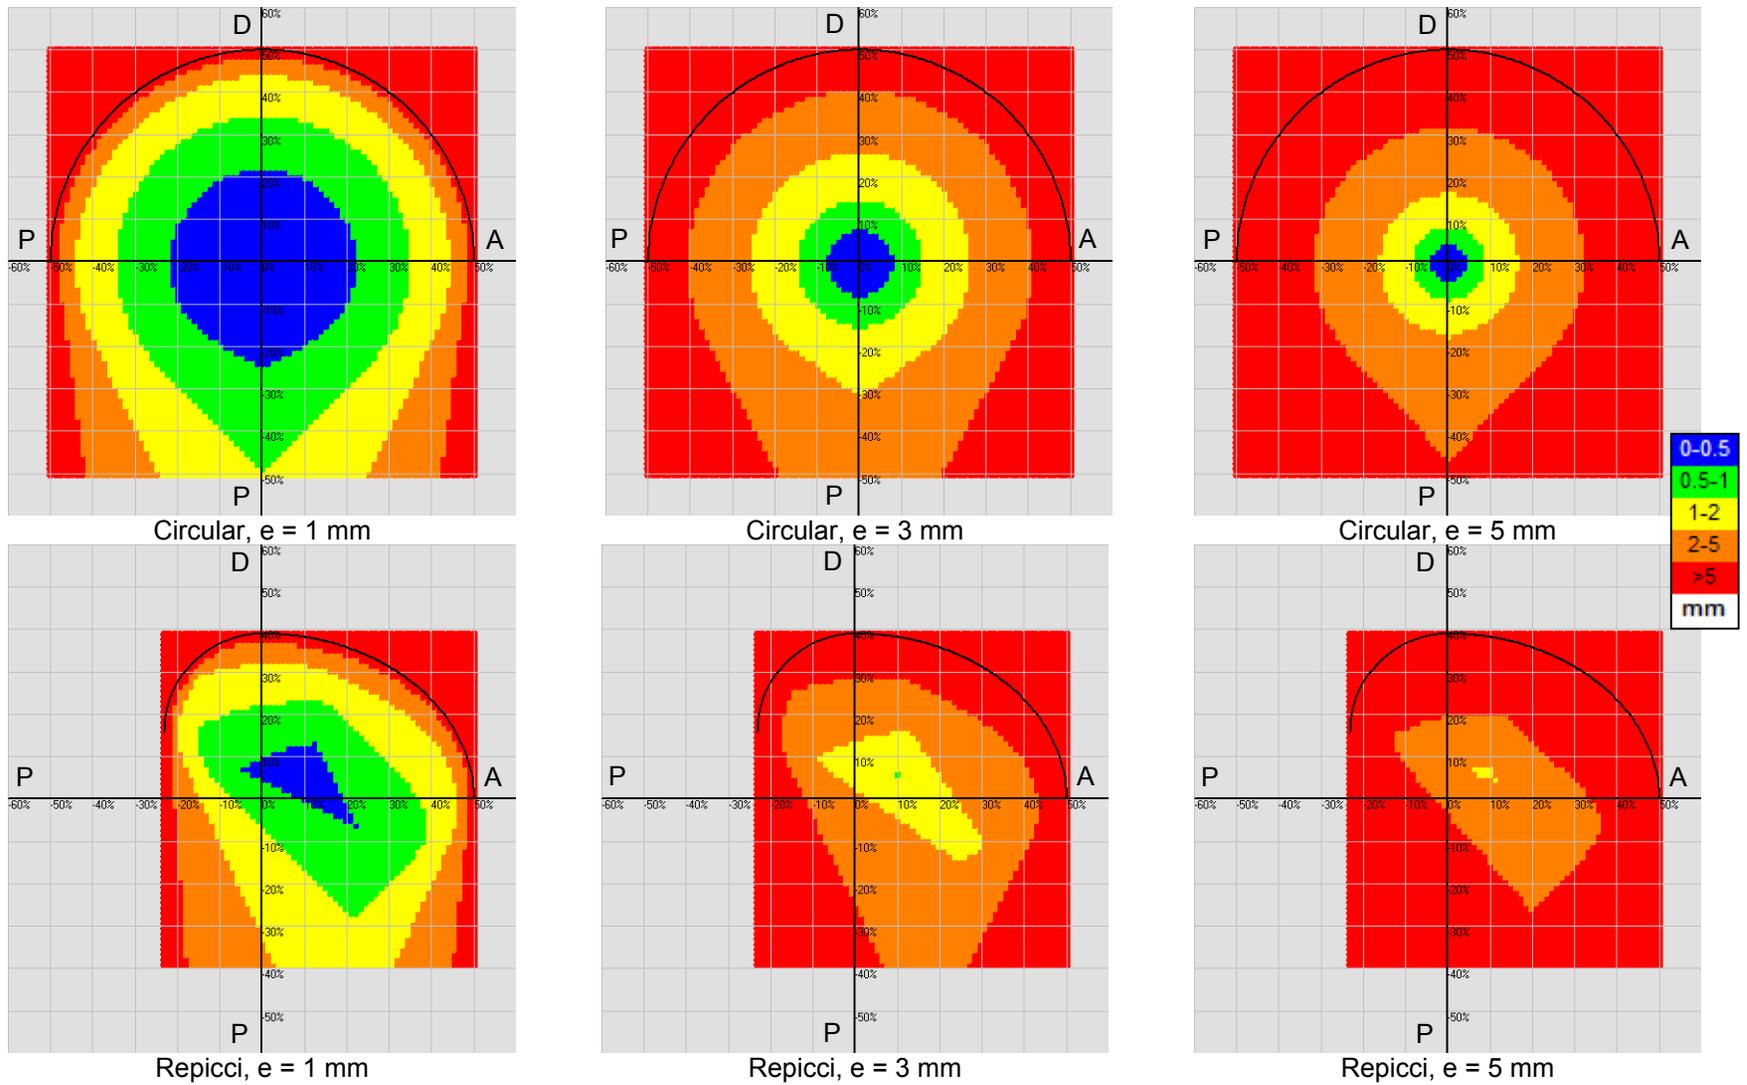
Figure 5-6: *Colour maps showing the amount of maximum lateral deflection caused by different amounts of equilibrium point penetration for circular and Repicci-type implant-to-bone surfaces. x- and y- axes are normalized to the percentage of the AP width of the surface. Note that for simplicity, the AP width of the Repicci-type surfaces was taken as the width of the major axis of the elliptical portion of the surface. A = Anterior, P = Posterior / Proximal, D = Distal.*

95

maximum deflection is approximately linear, which coincides with the linear geometry used in the equilibrium point analysis.

For circular-shaped unicondylar implants, the ideal attachment point is the surface's centre of curvature. It would be very difficult, however, for the surgeon to mount the robot accurately enough at this exact location. However, based on the colour maps, it can be seen that attachment points as far as 20% of the AP distance away from the centre result in a maximum deflection of less than the amount of equilibrium point penetration. An appropriate range of attachment points, therefore, can be considered as being within the colour contours that are less than the amount of equilibrium penetration. Similarly, for Repicci-type implants, the same region aligned about its mirrored radial centre is acceptable. This gives the surgeon a region of 40% of the AP width within which to mount the robot. This is sufficiently large to use eye-balling or manual measurements without causing significant problems with the robot's behaviour. It also gives the surgeon flexibility to choose a mounting location based on bone quality.

5.9   *Optimization of link lengths*

It is important to optimize the lengths of the robot's links to ensure their suitability for distal femoral bone sculpting. Optimizing the lengths requires the consideration of multiple variables.
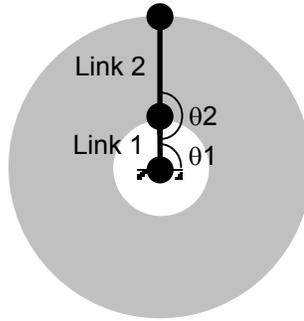
Figure 5-7: *Illustration showing the doughnut-shaped workspace (gray region) of the prototype.*

5.9.1   Optimizing for robot workspace

The first variable is ensuring that the robot workspace allows for all possible surface sizes. The two-dimensional manipulator prototype described in Chapter 3.0   has a ratio of link lengths of approximately one to two (link 1 : link 2). Its workspace is, therefore, doughnut shaped, with an unreachable circular region of radius equal to the length of link 1 and centred about the attachment point, as shown in Figure 5-7. This is an evident drawback since nearly all the surfaces listed in Appendix D have a radius smaller than the 3.5cm length of the Praxiteles base. From Appendix D, the smallest surface radius is 1.4cm. This is already a small radius for adult femoral sizes, however it is important to apply a factor of safety to be absolutely sure that it can be reached. A minimum radius of 1cm was, therefore, chosen. The difference in length between the two links should not exceed this.
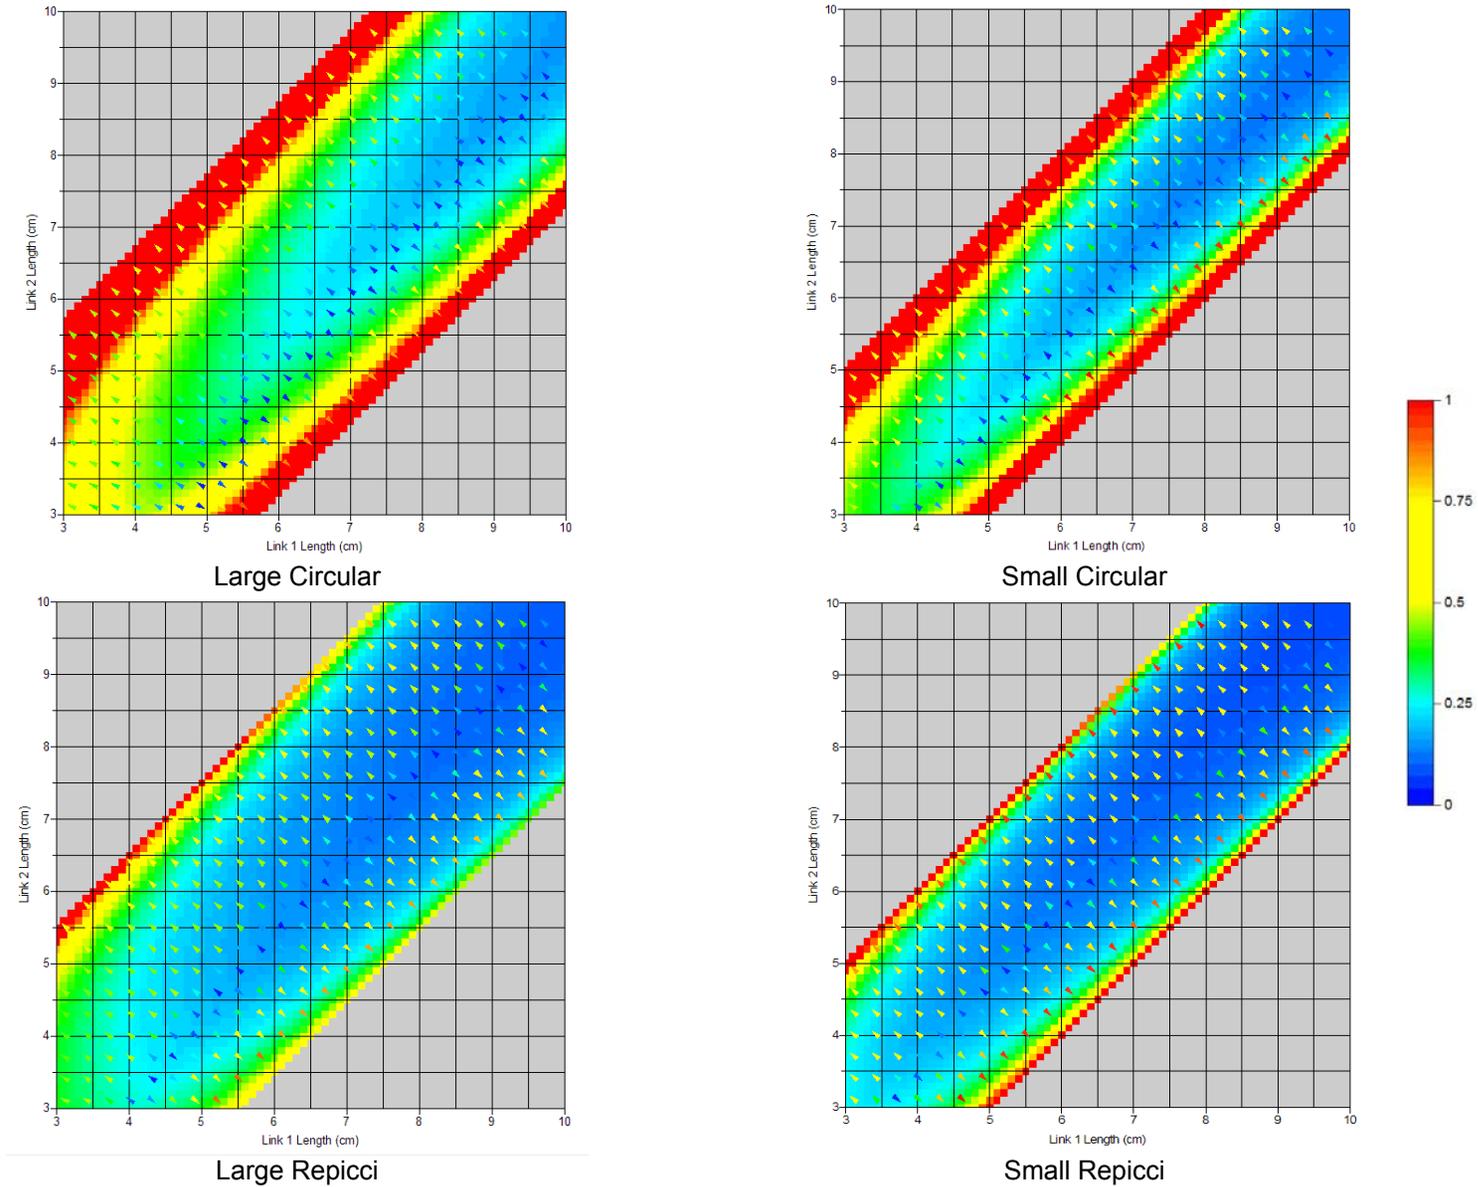
Figure 5-8: *Colour maps showing the distribution of stability (base colour maps) and average approach angles as described in Section 5.9.3(vector map overlays) for various link lengths Red regions are more likely to be affected by instability.*

5.9.2 Optimizing for surface instability

The second issue is ensuring that the chance of surface instability occurring is reduced for the given surfaces, even if the new control method proposed in Section 4.5.5 is applied. This can be done by optimizing the ratio of link lengths, as described in Section 4.5.2. Colour maps showing the distribution of the maximum $d\theta_c/d\theta_1$ ratio for a grid of link lengths and for large and small sizes of the circular and Repicci surfaces are shown in Figure 5-8. These show that for all sizes and shapes, the optimal link length ratio tends to be just above 1:1, with increasing link lengths resulting in improved stability.

5.9.3 Optimizing for typical approach angles

A third issue concerns the relative velocities of the two links as the end-effector nears the surface. During the experimental trials described in Section 3.4.2, as well as during several cadaver trials of the original Praxiteles system, it was observed that users apply motions generally either parallel or perpendicular to the milling surface. Motions parallel to the surface are surface tracing motions. Motions perpendicular to the milling surface, however, imply collision of the second link with the physical constraint. If the physical constraint is in motion before contact, this could cause an unnatural bump upon contact. Since the speed at which the physical constraint is moving depends on the speed of rotation of the first joint, the strength of this bump can be generalized to the amount by which the first link rotates during a perpendicular trajectory towards the surface. This is dependent on the link lengths.

Figure 5-9(a) and (b) show a time-lapse plot of the robot's motion in two cases. In case (a), the majority of the motion comes from the second link, so minimal bump is felt by the user upon contact with the physical constraint. In case (b), however, the majority of the motion comes from the first link, resulting in an unexpected impulse from the
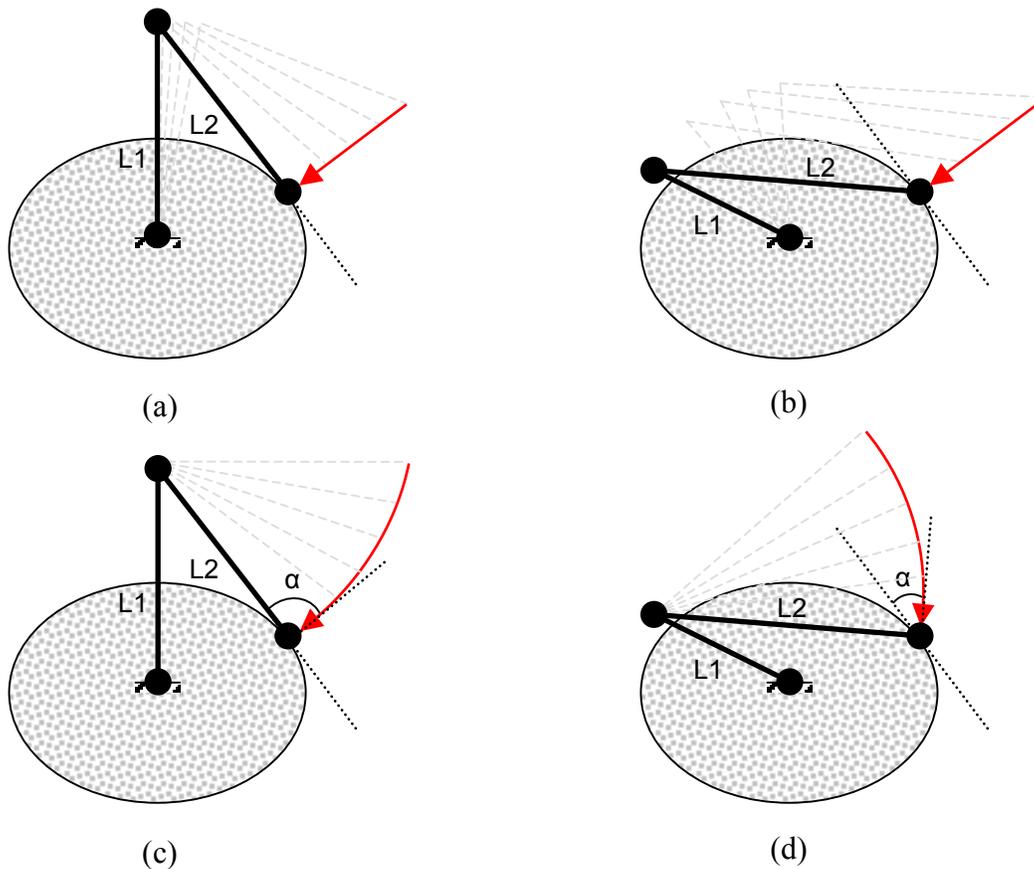
Figure 5-9: *Illustration of how the link lengths affect the relative motion of each link during an approach perpendicular to an arbitrary surface. (a) When the link lengths are approximately equal, the amount of motion of the first link is minimal during a perpendicular trajectory. (b) When the first link is shorter than the second link, the amount of motion of the first link is greater during a perpendicular trajectory. (c) If the user rotates the second link about the elbow, the approach angle (α) is perpendicular, resulting in a convincing collision with the surface. (d) Rotation of the second link about the elbow results in an acute approach angle (α).*

physical constraint as it correspondingly re-adjusts its position based on the motion of the first link.

In order to facilitate the optimization of the link lengths to decrease the amount of motion of the first link for the variety of bone-sculpting surfaces, cases (a) and (b) can be viewed in a simpler manner. The ideal case for any combination of link lengths would be if the

first link would remain motionless during a perpendicular approach angle, since this would result in no unnatural impulse upon contact. If the first link is, therefore, held still in both the cases presented in (a) and (b), we can see in Figure 5-9(c) and (d), that rotating the second link in the first case results in a perpendicular approach angle at the surface (good), while in the second case it results in an acute approach angle (bad). This simplified analysis was used to determine the optimal link length ratio for the circular and Repicci surfaces by determining the average link 2 approach angle along the entire length of each surface for various link lengths. The results, scaled to use the same colour scale as the instability results, are shown in the colour maps of Figure 5-8 as vector map overlays. The bluer the colour, the higher the average approach angle across the surface and the better the link length ratio with respect to this property.

5.9.4   *Optimizing for hardware constraints*

Other aspects that are important to consider in the optimization of the robot are hardware constraints, such as motor and sensor sizes, interference of parts, such as between the dynamic constraint and the base of the robot, and of course encumbrance of the surgeon's motions due to large sizes and inertias. The current length of the prototype's first link should be considered a minimum since the hardware is already very closely spaced. This length should, however, not be significantly longer, to allow for freedom of rotation without interference with the surgical environment. An appropriate maximum is double the current prototype's link 1 length.

5.9.5   *Summary of optimization*

A summary of the results of the link length optimization is as follows:

- Optimizing for robot workspace: the difference in length between the two links should not exceed 1cm.

- Optimizing for surface instability: the ratio between the links should be just above 1:1, and the longer the better.

- Optimizing for typical approach angles: the first link should be 0.5 cm longer than the second link.

- Optimizing for hardware constraints: the length of the first link should be between 3.5 and 7cm.

According to this summary, if it were decided to keep the original Praxiteles base for the first link, then an optimal second link length would be 3cm. If the entire system were to be redesigned, then the optimal lengths would be 5cm and 4.5cm, respectively. These lengths result in a sufficient workspace, low instability, high average approach angles, and are still small enough to prevent encumbrance to the surgeon and surgical environment.

5.10 Bone sculpting conclusions

This chapter presented a potentially very advantageous semi-active solution to distal femoral bone sculpting. The use of the dynamic physical constraint concept allows for the design of a simple, miniature and realistic-feeling virtual template device that can be used to emulate all types of convex surfaces. Such a device could encourage the development of new, bone-conserving implants that mimic the natural curvature of femoral condyles.

6.0 CONCLUSIONS AND FUTURE WORK

This thesis presented a novel solution to the haptic hard surface emulation problem. The solution was based on a dynamic physical constraint concept which was used to replace the standard impedance-type force feedback control loop typically used in haptics. The concept is simple, small, requires only basic position control with no force sensors, and above all creates a truly realistic hard surface effect. Surface contact, surface tracing and free-space motion are all successfully emulated.

6.1   Future Work

6.1.1   Prototype improvements

The prototype developed during this thesis was successful in demonstrating the feasibility of the dynamic physical constraint concept as a haptic hard surface emulation device. However, some improvements are required for its use in applications in which the device interacts with not only the user, but the environment as well.

For applications such as bone sculpting, precision is very important. As discussed in Section 4.4, the precision of the prototype was affected by its hysteretic behaviour. To decrease the amount of hysteresis, the physical constraint would have to be driven at faster speeds. One way of doing this would be to decrease the harmonic drive reduction ratio used in the second joint. A balance would have to be found of providing sufficient strength to counter the user's forces while increasing the harmonic drive's output speed. Using a 1:50 reduction ratio, for example, would double the reaction speed and would potentially significantly reduce the amount of hysteresis. If this were not enough, an additional solution would be to use a stronger and faster motor. A further modification that would increase the precision would be to use an external encoder to locate the first joint, to improve on the accuracy of the Hall sensor being used currently.

As mentioned in Section 4.5.5, some work is still needed to ensure my hypothesis for instability is correct. To do this, the prototype could be modified for the proposed new control strategy. Further control theory analysis could also be done to improve our understanding of the system.

Finally, the workspace of the prototype needs to be better matched to the intended surgical application. Currently, the prototype is unable to emulate small, tight surfaces

because of interference between the physical constraint and the robot's base. Creating a slot in the second link that would allow the constraint to penetrate within the link would make the system more compact. Another, more complicated option would be to redesign the second link to be curved inward rather than straight. That would bring the end effector closer into the centre of the robot.
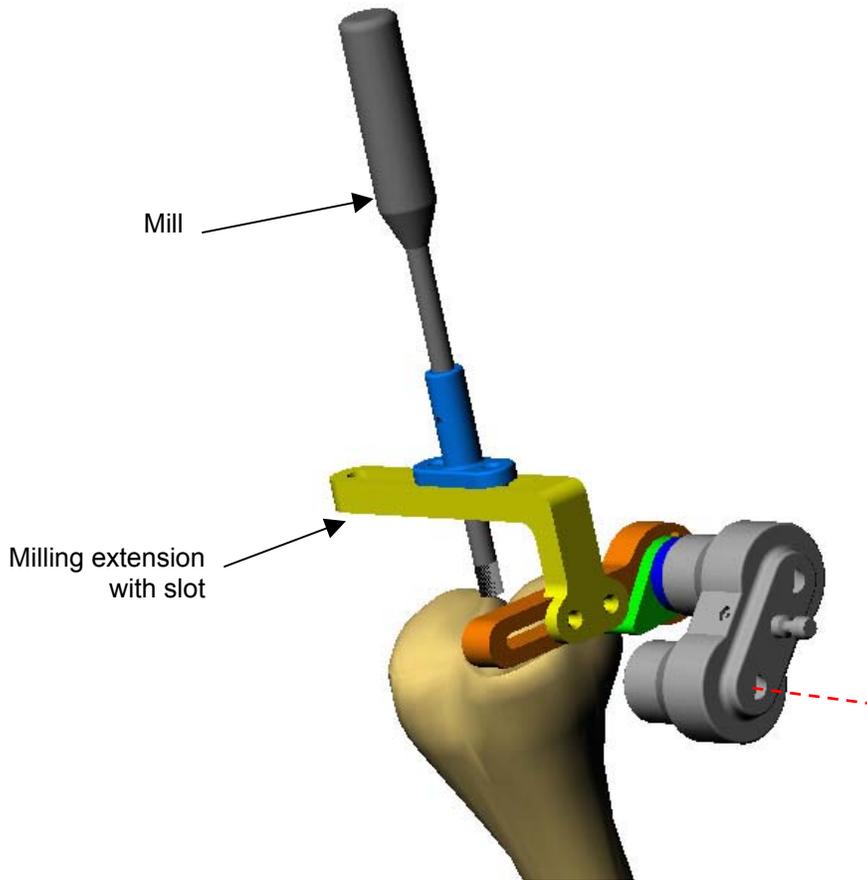

Figure 6-1: *Illustration of how the prototype could be extended for bone milling. The robot is fixed to the bone by the first axis (marked in red). The mill is either fixed in the slot or can slide laterally along it. The mill is also fixed vertically along its axis.*

6.1.2   *Extension of the prototype to bone milling*

The next step in the validation of the dynamic physical constraint concept would be to test the system's behaviour when actually interacting with environmental forces, such as those caused by milling. This could be done quite easily by mounting a milling extension onto the second link and attaching the prototype to the milling surface. A sketch that illustrates this idea can be seen in Figure 6-1. In this case, the milling extension would have a slot running parallel to the robot's two axes. The mill could either be fixed at a set location within this slot, or could be kept loose to enable the milling of a cylindrical surface. In the specific context of bone milling, this extension of the prototype could be used first on sawbones to test its feasibility and to see if there are any unexpected behavioural issues caused by the addition of milling forces. The prototype could then be tried on animal bones to mimic the conditions of a true human femoral milling situation. These tests could also be used to determine the prototype's milling accuracy.

6.2   *Extensions of the dynamic physical constraint concept*

6.2.1   *Architectural variations*

The dynamic physical constraint concept provides a method of constraining joint motion in a realistic manner. There are numerous variations in how the concept can be applied to manipulators. The concept primarily discussed in this thesis consists of a two degree of freedom (DOF) revolute system in which the first degree of freedom is freely mobile, while the second is constrained by a dynamic physical constraint. As mentioned earlier, the constraint does not have to be revolute. A prismatic joint can be envisioned that constrains motion radially (RP configuration). By the same logic, the unconstrained

degree of freedom of a two DOF architecture does not have to be revolute either: this would create a PP configuration. A double constraint can also be envisioned, where motion in both directions would be constrained. In this case, the joint would require two physical constraints overlaid on top of each other that could be driven either jointly, when the distance between the upper and lower constraint remains equal, or separately when the two constraints are variable with respect to each other. In a two DOF architecture, a second constraint (either dynamic or passive) could also be placed on the first joint, limiting its range of motion, allowing the robot to produce constraints in directions that the single constraint fails at. Evidently, there are many possibilities and it can be seen that this concept is very adaptable to many situations.
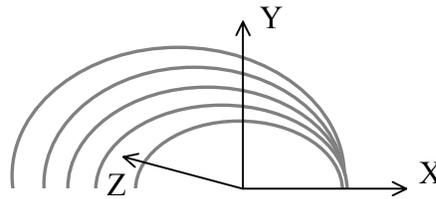

Figure 6-2: *Contouring of a milling surface in the sagittal (YX) plane. Each contour varies in the medio-lateral (Z) direction.*

6.2.2 Extension to three or more dimensions

The extension of the concept to three or more dimension would, of course, be very useful. The bone milling extension described in Section 6.1.2 can, in fact, be extended to allow three dimensional milling in a very simple manner. A linear encoder could be placed on the slot in the extension to link 2. This would be used to track the location of the mill in the medio-lateral direction. The three-dimensional milling surface could be modelled by a set of planar contours in the sagittal plane, as shown in Figure 6-2. Depending on the medio-lateral location of the mill, the physical constraint would adjust itself to the appropriate contour. The benefits of this method are that it requires minimal

modifications to the original prototype. Even though it is a three DOF system, the behaviour of the system would be similar to that of the original two DOF system.
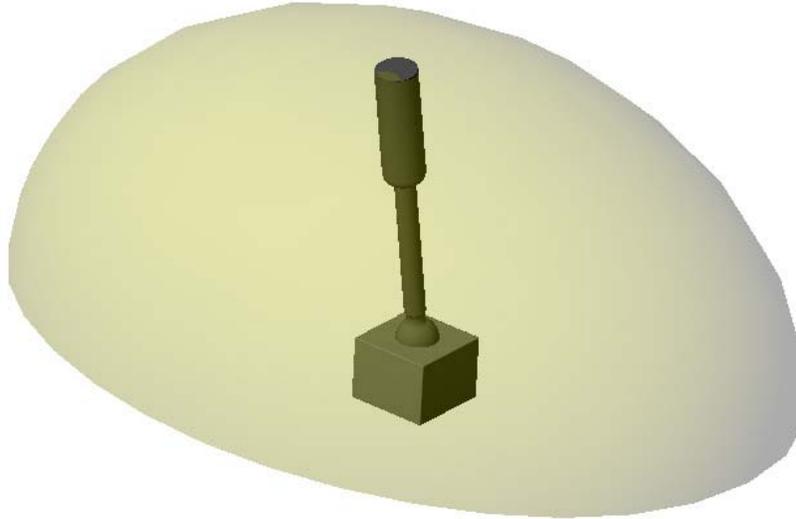

Figure 6-3: *Sketch of a three degree of freedom RRP configuration using a linear dynamic physical constraint, showing a sample three-dimensional curvilinear surface.*

Another architecture that can be envisioned is based on a ball joint (or double revolute joint) and a linear physical constraint (RRP configuration), as illustrated in Figure 6-3. In such a system, the ball joint would be freely mobile and its position in the XZ and YZ planes would be monitored. Based on this position, the linear physical constraint would adjust its radial length preventing incursion into the virtual surface. Many architectures exist to which this concept could be applied.

6.2.3 Other applications

The application of this novel concept is potentially very broad. To the author's knowledge, no design currently exists in the general field of haptic hard surface emulation that provides such realism. The concept could be used in the design of haptic interfaces for anything from desktop virtual environments (such as a three dimensional spatial mouse or joystick) to industrial applications or telemanipulation tasks to provide a true sense of touch through a manipulator.

In the surgical field, the concept could be widely used as well. It could, of course, be used for implant-bone surface preparation in joints other than the knee: hip, elbow, and ankle joints, for example. It could, in fact, be used in any kind of osteotomy for which there is a need for accurate curvilinear bone milling (in the milling of the femoral cavity during hip arthroplasty, for example). One could also imagine other applications in other surgical fields such as ear or brain surgery where it is important to prevent movement into critical and easily damaged structures.

REFERENCES

Backstein D, Safir O and Gross A. "Management of Bone Loss: Structural Grafts in Revision Total Knee Arthroplasty." *Clinical Orthopaedics And Related Research.* 446:104-112, May 2006.

Bargar WL. "Robots in Orthopaedic Surgery: Past, Present, and Future." *Clinical Orthopaedics And Related Research.* 463:31-36, Oct. 2007.

Book W, Charles R, Davis H and Gomes M. "The Concept and Implementation of a Passive Trajectory Enhancing Robot." *Proceedings of the ASME International Mechanical Engineering Conference and Exhibition.* 58:633-638, 1996.

Breisch SL. "Future is here does it work?." *AAOS Bulletin.* Aug. 1998. http://www.aaos.org. Accessed: Jan. 10, 2008

Brisson G, Kanade T, DiGioia A and Jaramaz B. "Precision Freehand Sculpting of Bone." *Proceedings of MICCAI 2004.* pp. 105-112, 2004.

Callaghan JJ. *The Adult Knee.* Lippincott Williams & Wilkins. p. 1106, 2003.

Campbell MD, Duffy GP and Trousdale RT. "Femoral component failure in hybrid total knee arthroplasty." *Clinical Orthopaedics and Related Research.* 356: 58-65, 1998.

Cavusoglu MC, Feygin D and Tendick F. "A Critical Study of the Mechanical and Electrical Properties of the PHANToM Haptic Interface and Improvements for High-Performance Control." *Presence.* 11(6): 555-568, Dec. 2002.

Cobb J, Henckel J, Gomes P, Harris S, Jakopec M, Rodriguez y Baena F, Barrett A and Davies B. "Hands-on robotic unicompartmental knee replacement: A prospective, randomised controlled study of the Acrobot system." *The Journal of Bone and Joint Surgery (British volume).* 88-B:188-197, 2006.

Colgate JE, and Wannasuphoprasit W and Peshkin MA. "Cobots: Robots For Collaboration With Human Operators." *Proceedings of the International Mechanical Engineering Congress & Exposition (IMECE).* 58:433-439, 1996.

Colgate JE, Grafing PE, Stanley MC and Schenkel G. "Implementation of stiff virtual walls in force-reflecting interfaces." *Proceedings of IEEE Virtual Reality Annual International Symposium (VRAIS)*. pp. 202-208, 1993.

Colgate JE, Peshkin MA and Wannasuphoprasit W. "Nonholonomic Haptic Display." *IEEE International Conference on Robotics and Automation*. 1:539-544, 1996.

Collier MB, Eickmann TH, Sukezaki F, McAuley JP and Engh GA. "Patient, Implant, and Alignment Factors Associated With Revision of Medial Compartment Unicondylar Arthroplasty." *The Journal of Arthroplasty*. 21(6 Suppl 2):108-115, Sep. 2006.

Davies B. "A review of robotics in surgery." *Proceedings of the Institution of Mechanical Engineers. Part H, Journal of Engineering in Medicine*. 214(1):129-140, 2000.

Davies B, Rodriguez y Baena FM, Barrett ARW, Gomes MPSF, Harris SJ, Jakopec M and Cobb JP. "Robotic control in knee joint replacement surgery." *Proceedings of the Institution of Mechanical Engineers. Part H, Journal of Engineering in Medicine*. 221(1):71-80, Jan. 2007.

Davis H and Book W. "Torque Control of a Redundantly Actuated Passive Manipulator." *Proceedings of the American Control Conference*. pp. 959-963, Jun. 1997.

Delnondedieu Y. "Un Robot à Sécurité Passive en réponse aux problèmes d'ergonomie et de sécurité en robotique médicale." PhD Thesis, Institut National Polytechnique de Grenoble, France, 1997.

Eickmann TH, Collier MB, Sukezaki F, McAuley JP and Engh GA. "Survival of Medial Unicondylar Arthroplasties Placed by One Surgeon 1984–1998." *Clinical Orthopaedics And Related Research*. 452:143-149, Nov. 2006.

Engh GA. "Advances in knee arthroplasty for younger patients: traditional knee arthroplasty is prologue, the future for knee arthroplasty is prescient." *Orthopedics*. 30(8 Suppl):55-57, Aug. 2007.

Gioe TJ, Killeen KK, Grimm K, Mehle S and Scheltema K. "Why Are Total Knee Replacements Revised? Analysis of Early Revision in a Community Knee Implant Registry." *Clinical Orthopaedics And Related Research*. 428:100-106, Nov. 2004.

Haas SB, Nelson CL and Laskin RS. "Posterior stabilized knee arthroplasty: an assessment of bone resection." *The Knee*. 7(1):25-29, Jan. 2000.

Harris SJ, Davies BL and Jakopec M. "Active-Constraint Robots." United States Patent Application # US 2004/0128026 A1. Jul. 2004.

Hayward V and MacLean KE. "Do It Yourself Haptics, Part-I." *IEEE Robotics and Automation Magazine*. submitted for publication, 2007.

Herburt I, Moszynska M and Peradzynski Z.. "Remarks on Radial Centres of Convex Bodies." *Mathematical Physics, Analysis and Geometry*. 8(2):157-172. May 2005.

Ho SC, Hibberd RD and Davies BL. "Robot Assisted Knee Surgery: Establishing a force control strategy incorporating active motion constraint." *IEEE Engineering In Medicine And Biology*. 14(3):292-300, May/Jun. 1995.

Hogan N. "Impedance control - An approach to manipulation. I - Theory. II - Implementation. III - Applications." *ASME Transactions Journal of Dynamic Systems and Measurement Control B*. 107:1-24, Mar. 1985.

Hube R, Sotereanos NG and Reichel H. "The Midvastus Approach for Total Knee Arthroplasty." *Orthopedics and Traumatology*. 10(3):235-244, Sep. 2002.

Jakopec M, Harris SJ, Rodriguez y Baena F, Gomes P and Davies BL. "The Acrobot system for total knee replacement." *The Industrial Robot*. 30(1):61-66, 2003.

Jakopec M, Harris SJ, Rodriguez y Baena F, Gomes P, Cobb J and Davies BL. "The First Clinical Application of a "Hands-On" Robotic Knee Surgery System." *Computer Aided Surgery*. 6:329-339, 2001.

Kim J, Park FC, Park Y and Shizuo M. "Design and Analysis of a Spherical Continuously Variable Transmission." *Journal of Mechanical Design*. 124:21-29 Mar 2002.

Kuchenbecker KJ, Fiene J and Niemeyer G. "Improving Contact Realism through Event-Based Haptic Feedback." *IEEE Transactions On Visualization And Computer Graphics*. 12(2):219-229, Mar. 2006.

Lawrence DA and Chapel JD. "Performance Trade-offs for Hand Controller Design." *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*. 4:3211-3216, May 1994.

Mark WR, Randolph SC, Finch M, Van Verth JM and Taylor II RM, "Adding Force Feedback To Graphics Systems: Issues And Solutions." *Proceedings SIGGRAPH 96*. pp. 447-452, 1996.

Massie TH and Salisbury JK. "The PHANTOM Haptic Interface: A Device for Probing Virtual Objects." *Proceedings of the ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 1994.

Moore Jr. CA, Peshkin MA, and Colgate JE. "Cobot Implementation of Virtual Paths and 3-D Virtual Surfaces." *IEEE Transactions on Robotics and Automation*. 19(2):347-351, Apr. 2003.

Mussa-Ivaldi FA, Hogan N and Bizzi E. "Neural, Mechanical, and Geometric Factors Subserving Arm Posture in Humans." *The Journal of Neuroscience*. 5(10):2732-2743, Oct. 1985.

Pagnano MW and Meneghini RM. "Minimally Invasive Total Knee Arthroplasty with an Optimized Subvastus Approach." *The Journal of Arthroplasty*. 21(4 Suppl 1):22-26, 2006.

Panousis K, Murnaghan C, Koettig P and Grigoris P. "Fracture of the femoral component of a Brigham unicompartmental knee: a case report." *Knee Surgery, Sports Traumatology, Arthroscopy*. 12:307-310, Jul. 2004.

Paula G. "Cobots for the assembly line." *Mechanical Engineering Magazine Online*. Oct. 1997. http://www.memagazine.org. Accessed: Nov. 11, 2005.

Peshkin MA , Colgate JE and Moore C. "Passive robots and haptic displays based on nonholonomic elements." *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*. pp. 551-556, Apr. 1996.

Pipino F. "Tissue-sparing surgery (T.S.S.) in hip and knee arthroplasty (Editorial)." *Journal of Orthopaedics and Traumatology*. 7(1):33-35, Mar. 2006.

Plaskos C. "Modeling and Design of Robotized Tools and Milling Techniques for Total Knee Arthroplasty." PhD Thesis, Université Joseph Fourier, Grenoble, France, 2005.

Rahman MM, Ikeura R and Mizutani K. "Investigating the Impedance Characteristic of Human Arm for Development of Robots to Cooperate with Human operators." *Proceedings of the 1999 IEEE International Conference on Systems, Man, and Cybernetics*. 2:676-681, 1999.

Rand JA, Trousdale RT, Ilstrup DM and Harmsen WS. "Factors Affecting the Durability of Primary Total Knee Prostheses." *The Journal of Bone and Joint Surgery (American volume)*. 85-A(2):259-265, Feb. 2003.

Riaz S and Maistrelli GL. "Unicompartmental Knee Arthroplasty and its Minimally Invasive Version." *Pakistan Journal of Medical Sciences*. 21(3):361-374, Jul.-Sep. 2005.

Robertsson O, Knutson K, Lewold S and Lidgren L. "The Swedish Knee Arthroplasty Register 1975-1997: an update with special emphasis on 41,223 knees operated on in 1988-1997." *Acta Orthopaedica Scandinavica*. 72(5):503-513, Oct . 2001.

Roche M. "Changing the Way Surgeons Plan and Execute Minimally Invasive Unicompartmental Knee Surgery." *Orthopaedic Product News*. Jul./Aug. 2006. Available: http://www.makosurgical.com.

Rodriguez F, Harris S, Jakopec M, Barrett A, Gomes P, Henckel J, Cobb J and Davies B. "Robotic clinical trials of uni-condylar arthroplasty." *International Journal of Medical Robotics and Computer Assisted Surgery*. 1(4):20-28, 2005.

Rosenberg LB. "Virtual Fixtures: Perceptual Overlays Enhance Operator Performance in Telepresence Tasks." Interim Report, Biodynamics and Biocommunications Division, Wright Patterson Air Force Base, 1992.

Rossi P, Castoldi F, Rossi R, Caranzano F, Baronetti M and Dettoni F. "TSS and traditional surgery in hip and knee replacement." *Journal of Orthopaedics and Traumatology*. 8(3):157-163, Sep. 2007.

Sah AP and Scott RD. "Lateral Unicompartmental Knee Arthroplasty Through a Medial Approach. Study with an Average Five-Year Follow-up." *The Journal of Bone and Joint Surgery (American volume)*. 89(9):1948-1954, Sep. 2007.

Salcudean SE and Vlaar TD. "On the Emulation of Stiff Walls and Static Friction with a Magnetically Levitated Input/Output Device." *ASME Journal of Dynamics, Measurement and Control*. 119:127-132, Mar. 1997.

Salisbury K, Brock D, Massie T, Swarup N and Zilles C. "Haptic rendering: programming touch interaction with virtual objects." *Proceedings of the 1995 symposium on Interactive 3D graphics*. pp. 123-130, 1995.

Schneider O and Troccaz J. "A Six-Degree-of-Freedom Passive Arm with Dynamic Constraints (PADyC) for Cardiac Surgery Application: Preliminary Experiments." *Computer Aided Surgery*. 6:340–351, 2001.

Schneider O, Troccaz J, Chavanon O and Blin D. "PADyC : a Synergistic Robot for Cardiac Puncturing." *Proceedings of the 2000 IEEE International Conference on Robotics & Automation*. pp. 2883-2888, Apr 2000.

Shoham M, Burman M, Zehavi E, Joskowicz L, Batkilin E and Kunicher Y. "Bone-Mounted Miniature Robot for Surgical Procedures: Concept and Clinical Applications." *IEEE Transactions on Robotics and Automation*. 19(5):893-901, Oct. 2003.

Swanson DK, Book WJ. "Path-Following Control for Dissipative Passive Haptic Displays." *Proceedings of the11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'03)*. pp. 101-108, Mar. 2003.

Swarts E, Miller SJ, Keogh CV, Lim G and Beaver RJ. "Fractured Whiteside Ortholoc II Knee Components." *Journal of Arthroplasty*. 16(7): 927-934, Oct. 2001.

Troccaz J and Delnondedieu Y. "Semi-active guiding systems in surgery: a two-dof prototype of the passive arm with dynamic constraints (PADyC)." *Mechatronics*. 6(4):399-421, 1996.

Troccaz J. "Disengageable Free Wheel." United States Patent # 5529159. Jun. 1996.

Wada M, Imura S, Bo A, Baba H and Miyazaki T. "Stress fracture of the femoral component in total knee replacement: a report of 3 cases." *International Orthopaedics*. 21:54-55, 1997.

Whiteside LA. "Ligament Balancing in Revision Total Knee Arthroplasty." *Clinical Orthopaedics And Related Research*. 423:178-185, Jun. 2004.

APPENDIX A - PROTOTYPE DESIGN DRAWINGS

| RevNo | Revision note | | | Date | Signature | Checked |

	Description	Quantité
A	Stopper	1
B	Screw ⌀M3 x 10	4
C	Bearing Cap	1
D	Ball Bearing SKF 61800	2
E	Link 2	1
F	Spacer	1
G	Bolt ⌀M5 x 30	1
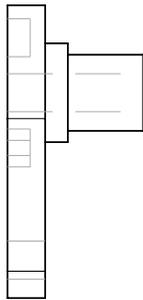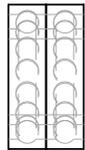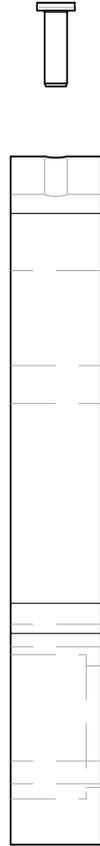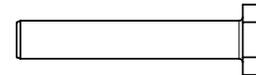
A B C D E F G
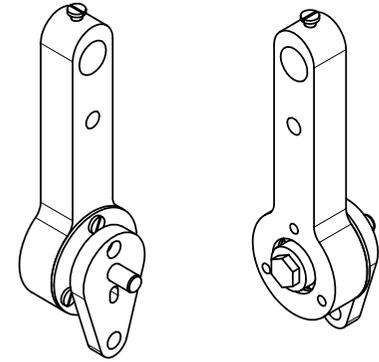
PRAXITELES STOPPER DESIGN ASSEMBLY DRAWING - V1

N. HUNGR

Edition	Sheet
mm	18/09/2007

$\varnothing 13 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$\varnothing 10h6$

$\varnothing 5 \begin{smallmatrix} +0.20 \\ +0 \end{smallmatrix}$

18

8

5

$3 \begin{smallmatrix} +0.20 \\ +0.00 \end{smallmatrix}$

$6.25 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$7.25 \begin{smallmatrix} +0.05 \\ -0.05 \end{smallmatrix}$

$2 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$\varnothing 23$

$22 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$\varnothing M5$

$\varnothing 10$

$\varnothing 5 \begin{smallmatrix} +0.05 \\ +0.00 \end{smallmatrix}$

$\varnothing 3 \begin{smallmatrix} +0.05 \\ +0 \end{smallmatrix}$
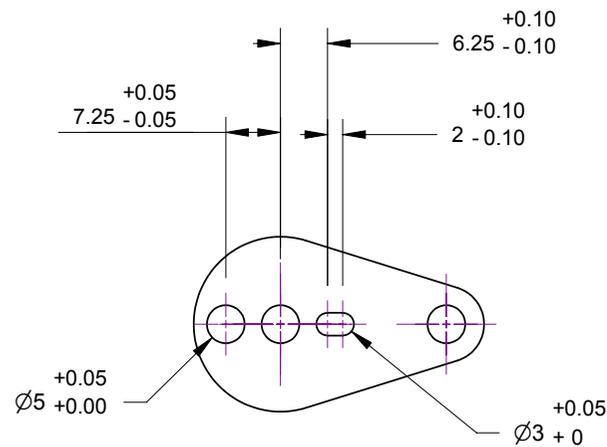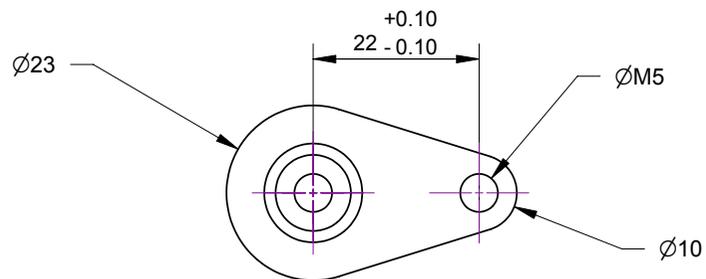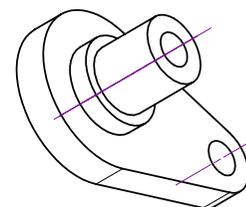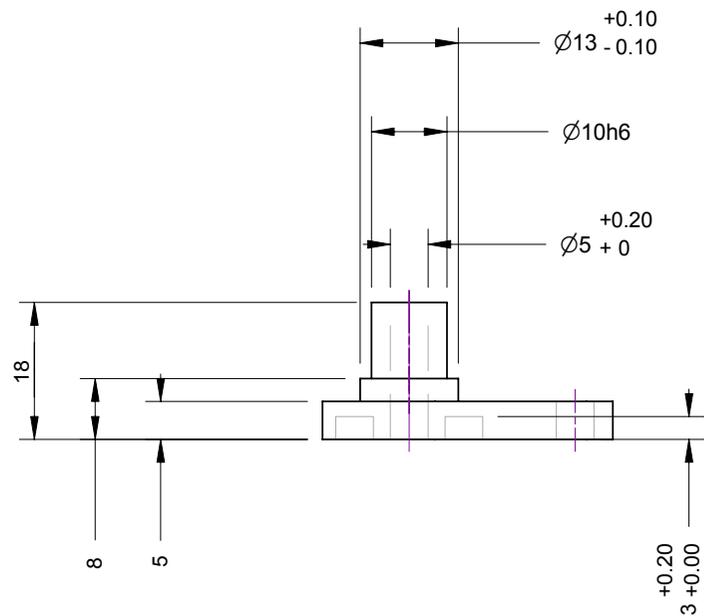
Matière: AL 2017A

PRAXITELES STOPPER - V2

N. HUNGR

Edition
mm

Sheet
19/09/2007

RevNo	Revision note		Date	Signature	Checked

A

B

$\varnothing 16 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$12 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

3 trous Ø3

Ø31

1

C

D
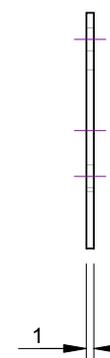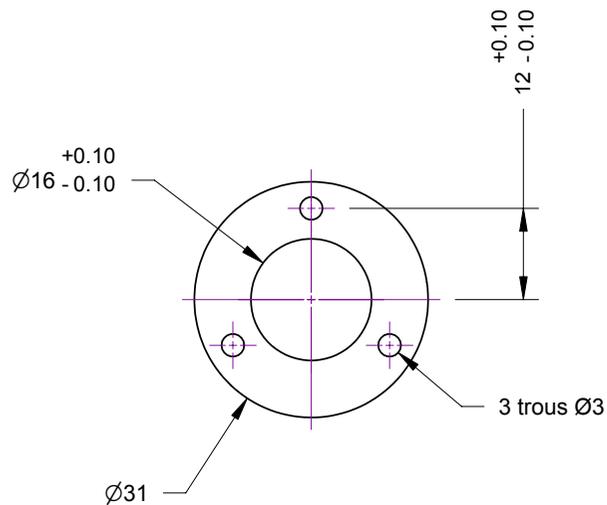
Matière: Acier normal

PRAXITELES BEARING CAP - V1

N. HUNGR

Edition mm	Sheet 17/09/2007

| RevNo | Revision note | | Date | Signature | Checked |

$10 \begin{smallmatrix} -0.10 \\ -0.00 \end{smallmatrix}$

Ø19H7

3 trous de vis ØM3

$Ø16 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

ØM5

$Ø10 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

R5

Ø31

$12 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$45 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

$65 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

67.5

$7.5 \begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$

15

ØM3

6

12

Matière: AL 2017A

PRAXITELES LINK 2 - V2

N. HUNGR

Edition mm

Sheet 19/09/2007

RevNo	Revision note	Date	Signature	Checked

A

B

$\varnothing 5 {}^{+0.20}_{+0.00}$

$\varnothing 13 {}^{+0.10}_{-0.10}$

4

R0.5

C
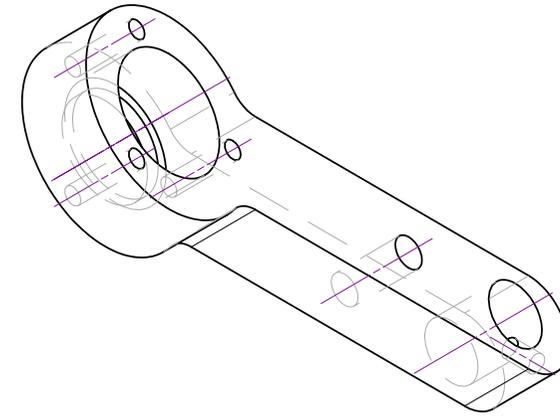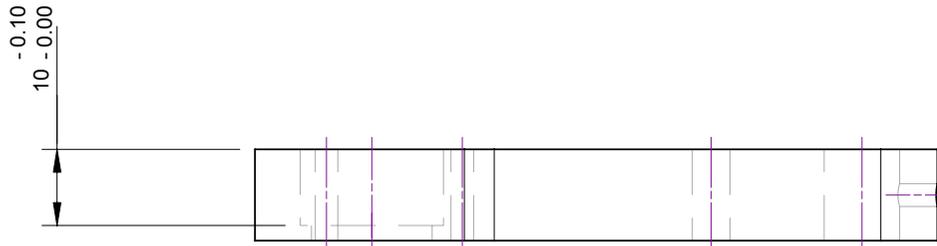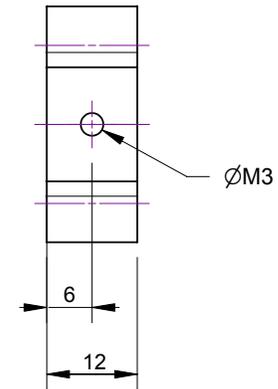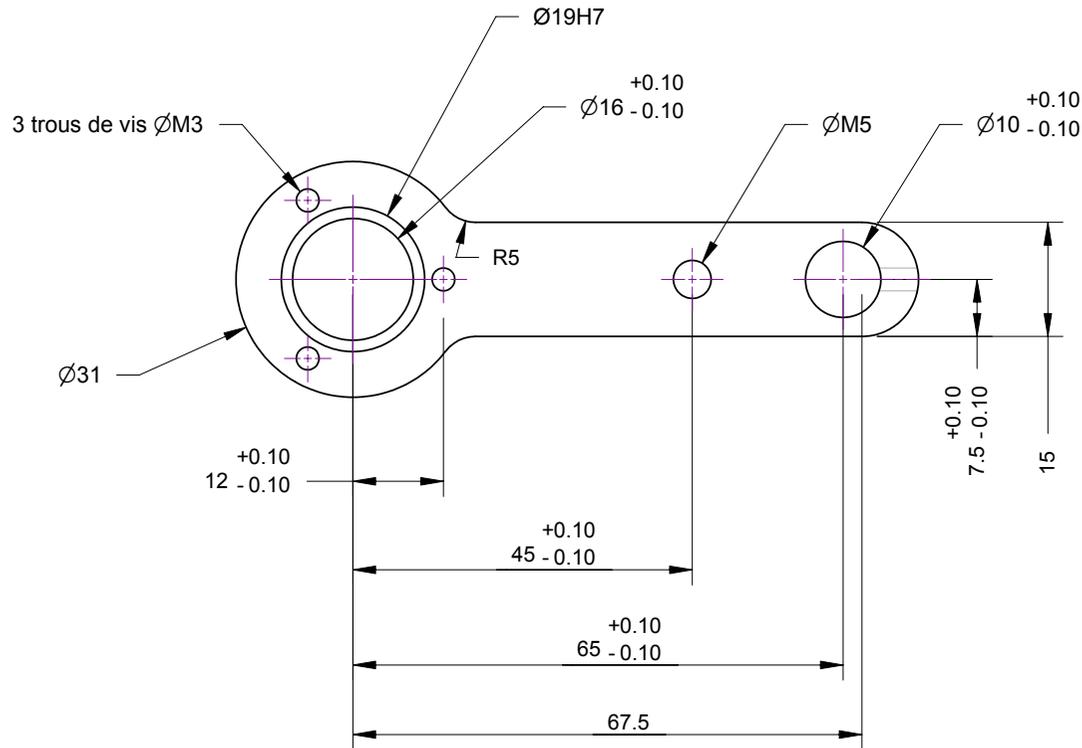
D
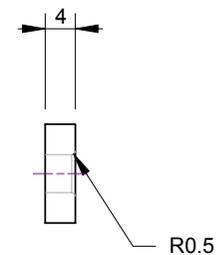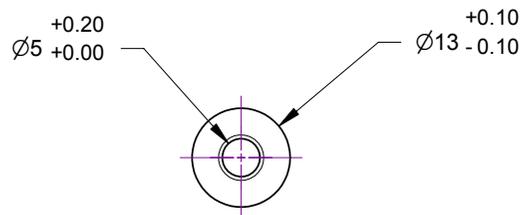
Matière: AL 2017A

PRAXITELES SPACER - V1

N. HUNGR

Edition
mm

Sheet
17/09/2007

APPENDIX B - NUMERICAL AND LOOKUP TABLE METHODS FOR PROTOTYPE CONTROL

Numerical method:

Method layout:

Program outline:

Input:

θ_1, s(x,y), (x,y)$_{initial}$

Calculate:

f(x,y) = d(x,y) - L$_2$

f'(x,y)

Apply Halley's method:

$$x_{n+1} = x_n - \frac{2f(x_n,y_n)f'(x_n,y_n)}{2f'(x_n,y_n)^2 - f(x_n,y_n)f''(x_n,y_n)}$$

Loop until x_n - x_{n+1} < *precision*

Calculate and output:

θ_2
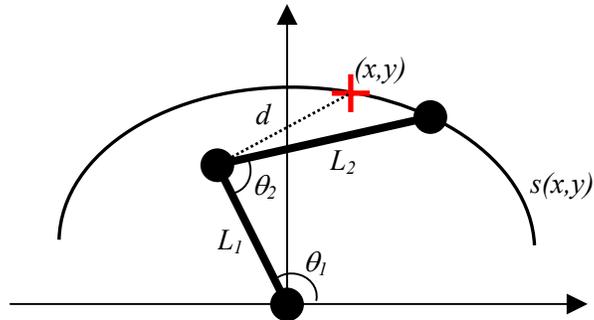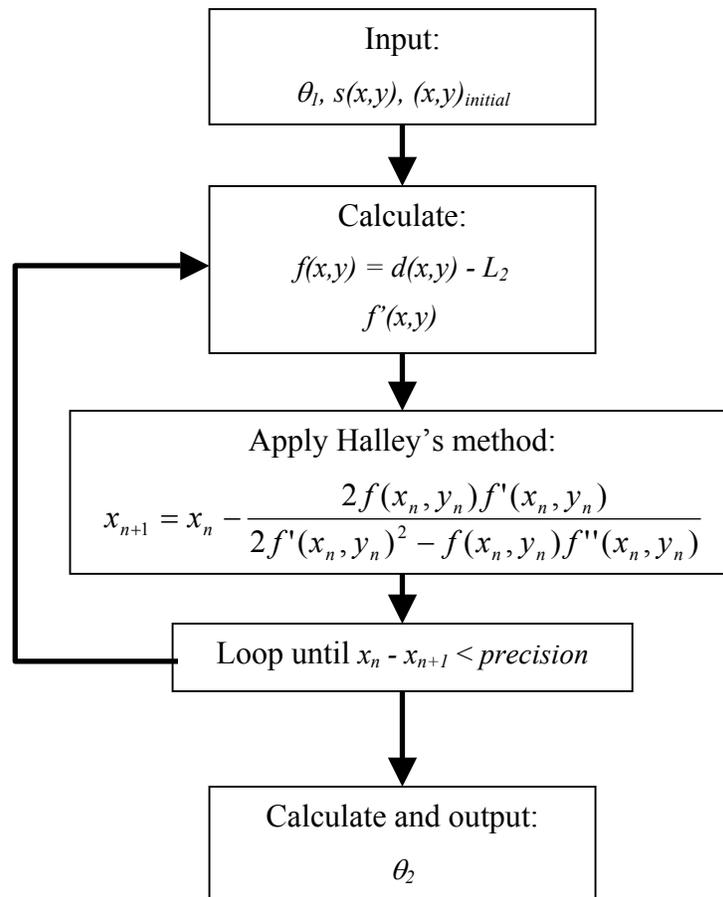
Lookup Table Method:

Program outline:

```
┌──────────────────────────────────────────┐
│                 Input:                     │
│                                            │
│   θ₁, List(θ₁,ᵢ , θ₂,ᵢ), i_low, i_high     │
└──────────────────────────────────────────┘
```

Input:

θ_l, $List(\theta_{1,i}$, $\theta_{2,i})$, i_{low}, i_{high}

Find:

i_{middle}

If: $\theta_l < \theta_{l,imiddle}$

then: $i_{high} = i_{middle}$

else: $i_{low} = i_{middle}$

Loop until $i_{high} - i_{low} = 1$

Interpolate between

$\theta_{2,ihigh}$ and $\theta_{2,ilow}$
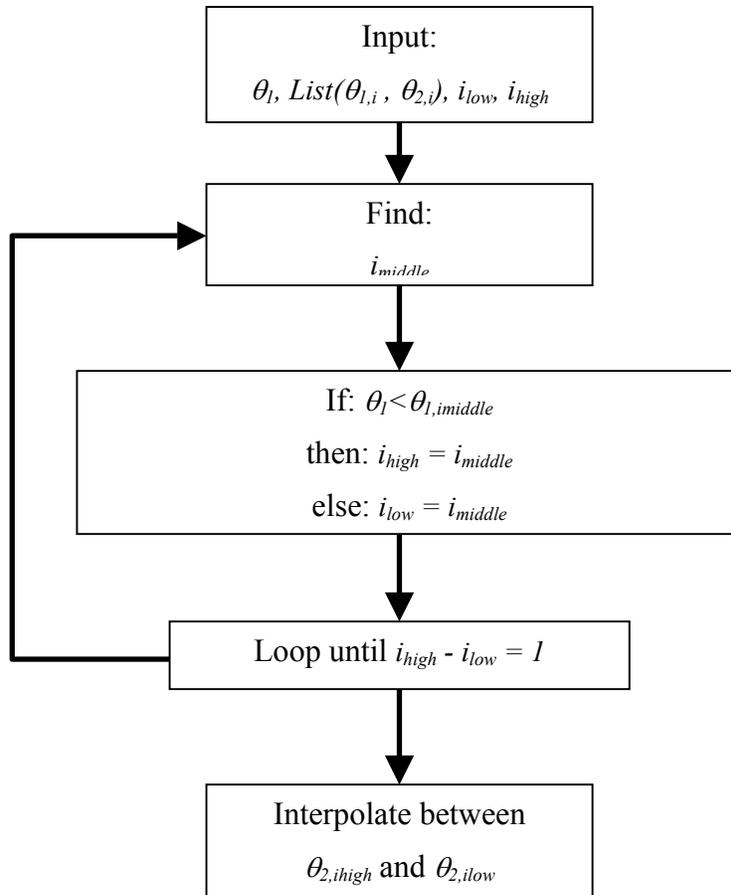
# APPENDIX C - DIGITIZED PROTOTYPE TEST SCANS

Surface 1: Ellipse
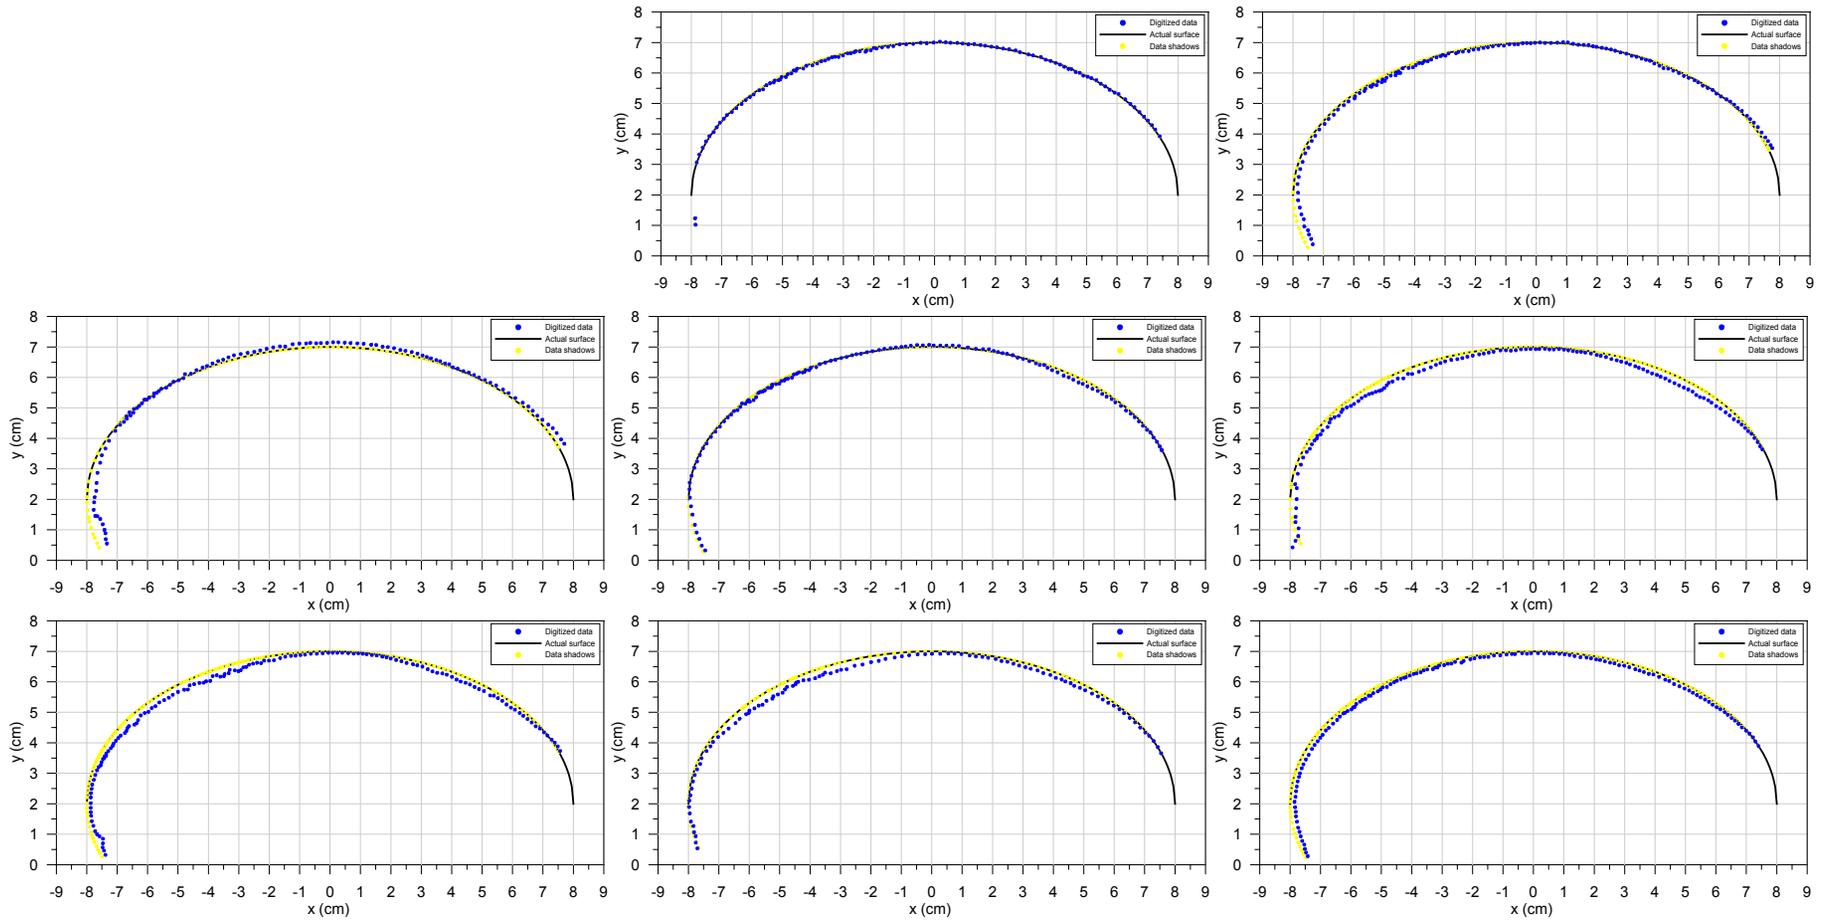
# Surface 2: Horizontal line
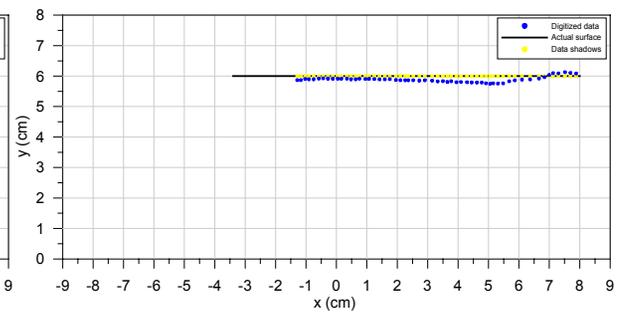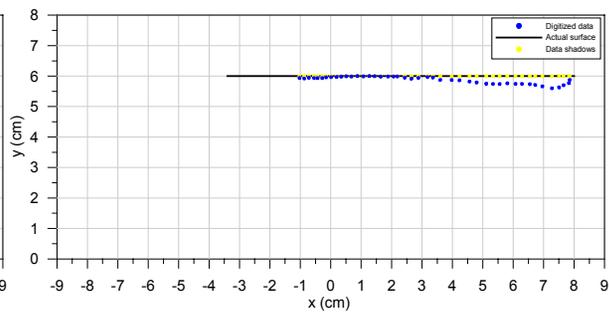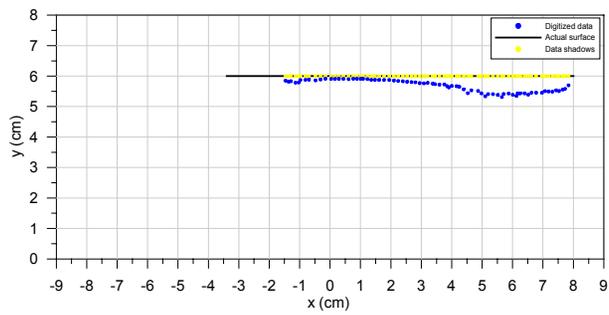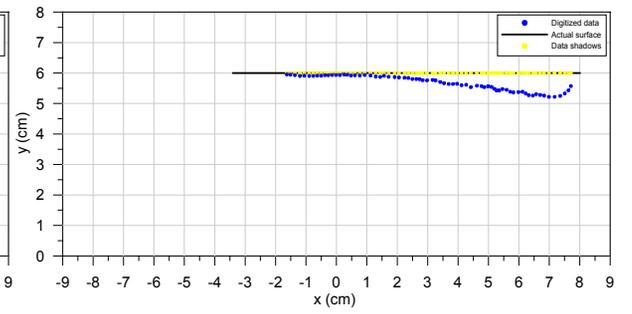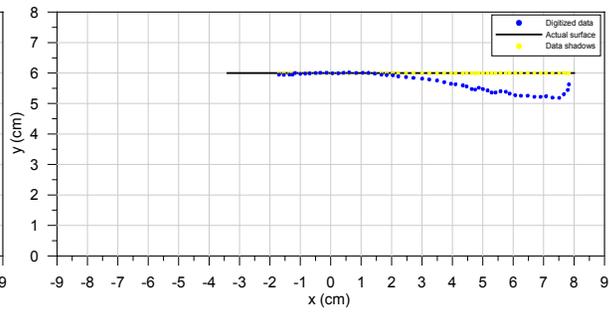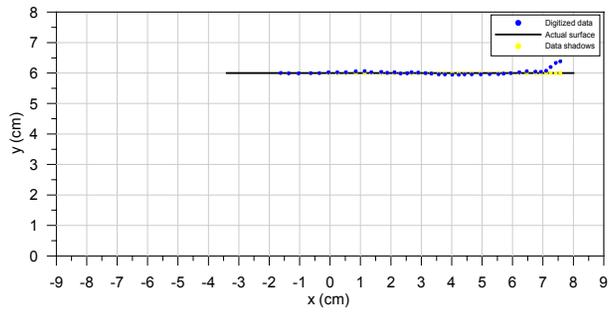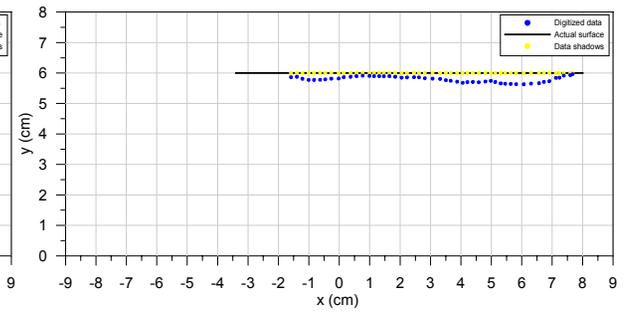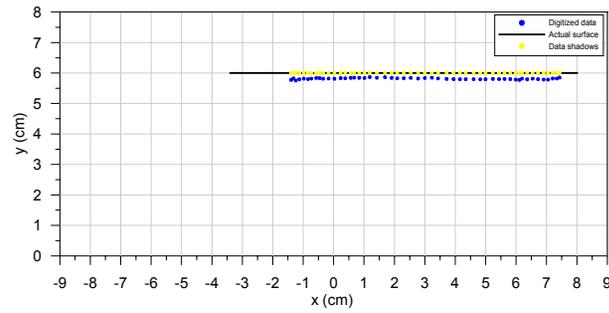
# Surface 3: Hybrid circle/ellipse
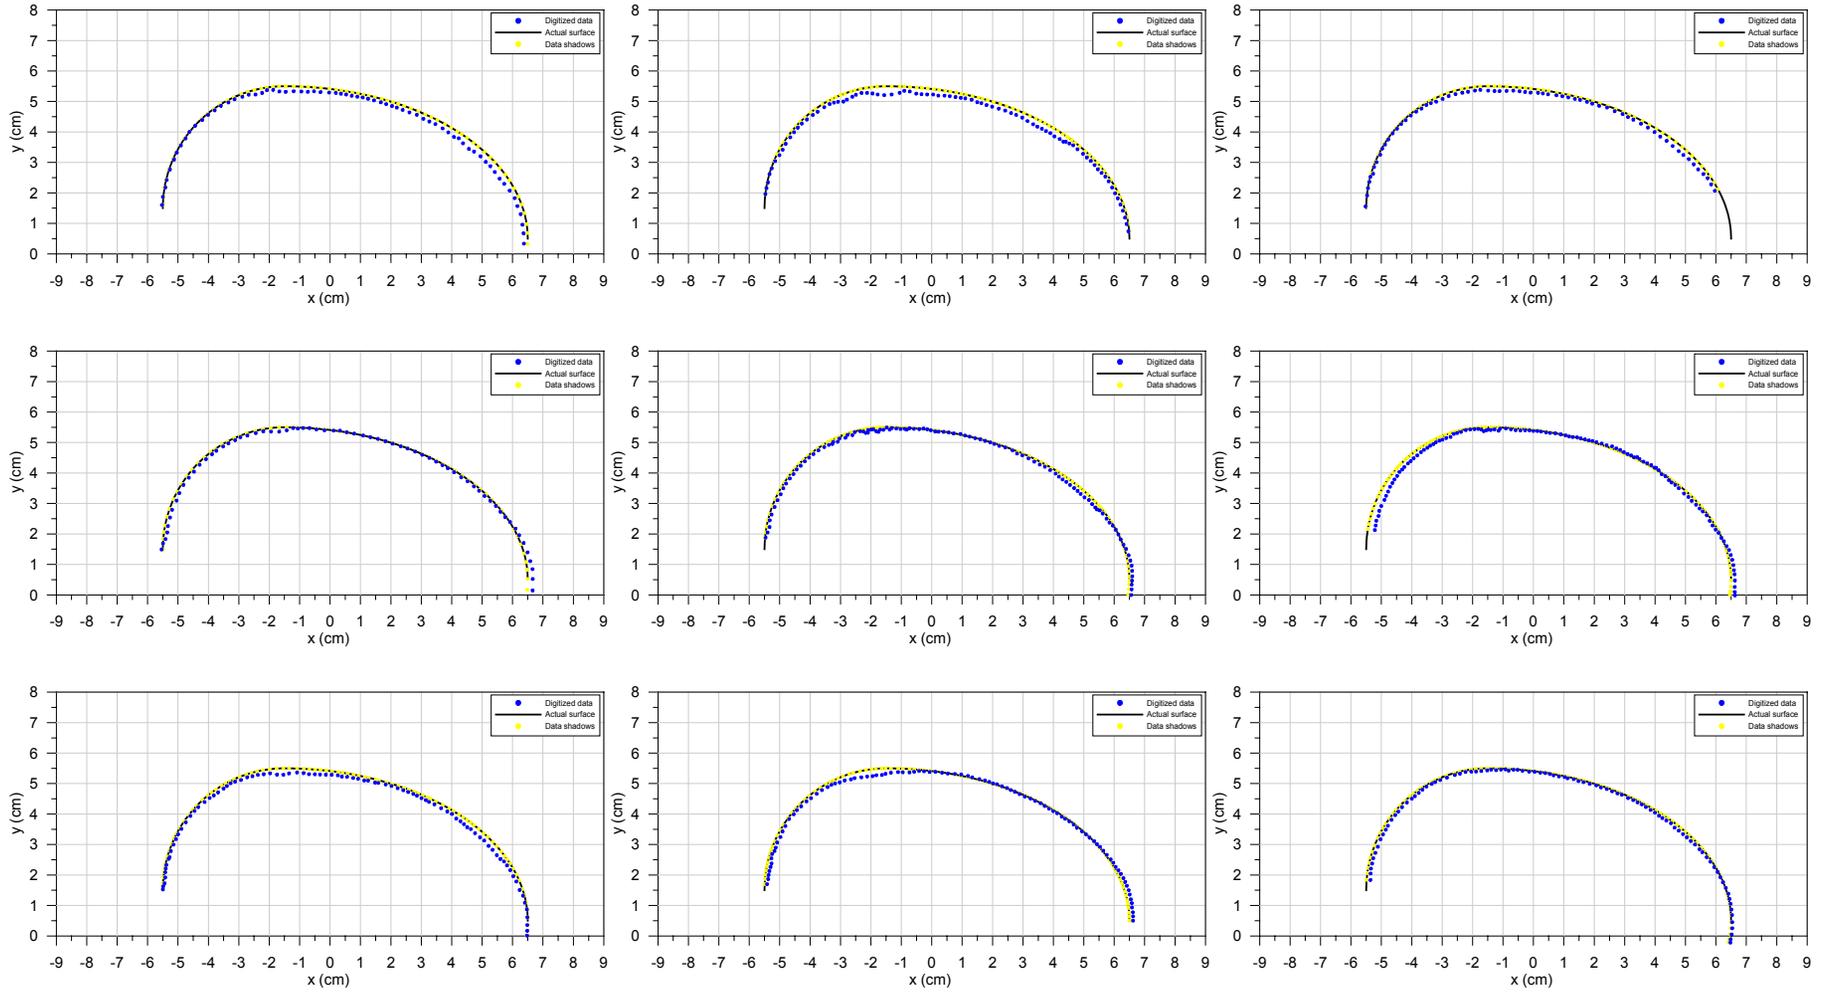
Surface 4: Horizontal sine-wave
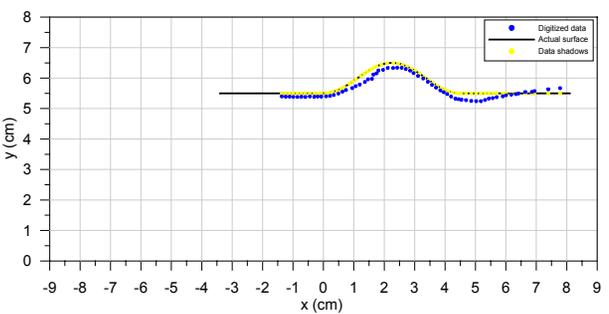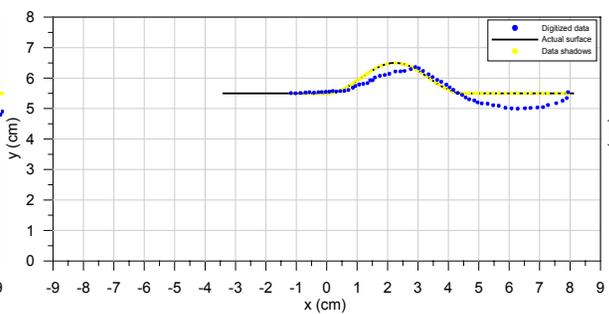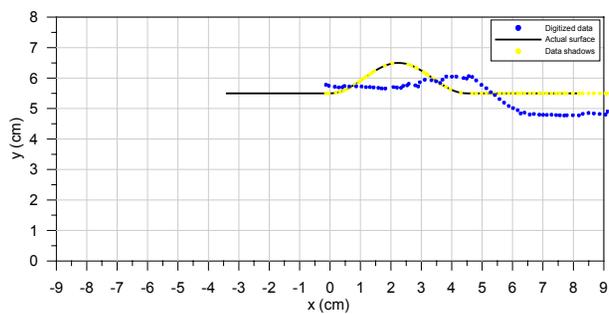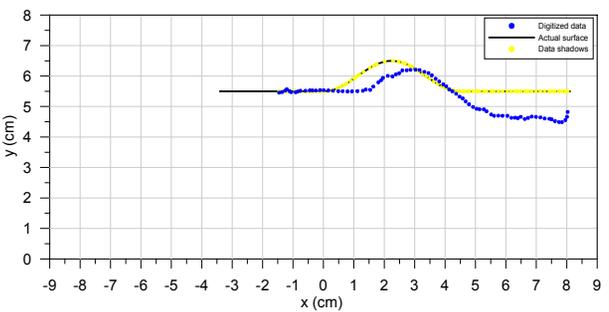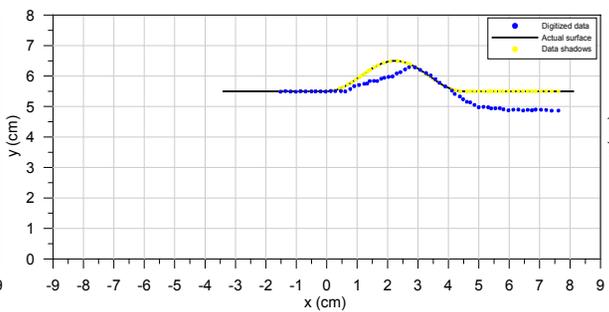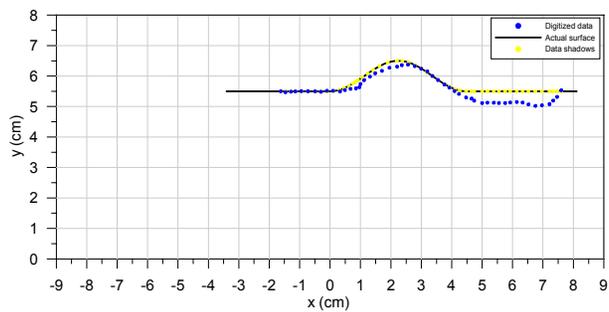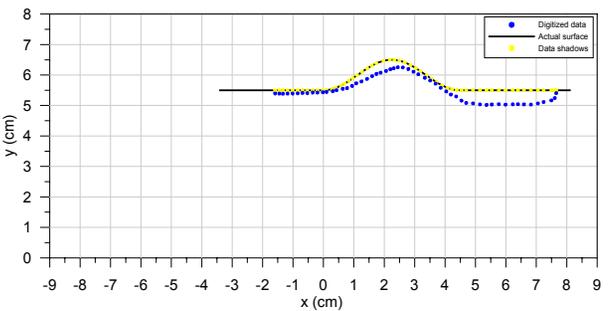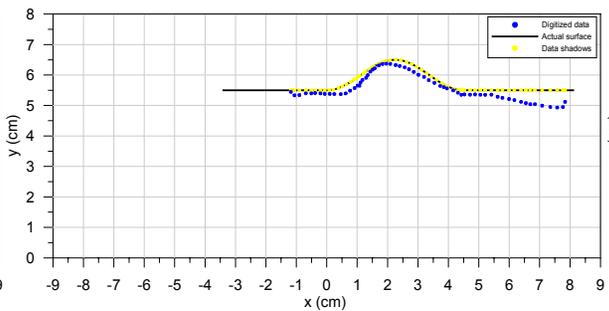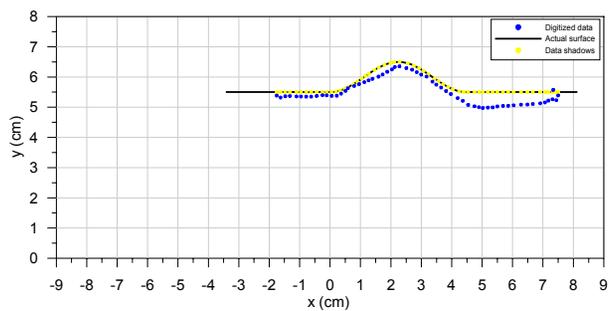
APPENDIX D - UNICOMPARTMENTAL FEMORAL IMPLANT DIMENSIONS

Implant Shape	Implant Name	Sizes	Dimensions (mm)								
			AP Length	Radius of Curvature 1		Origin of Radius 1		Radius of Curvature 2		Origin of Radius 2	
				a	b	x	y	a	b	x	y
Spherical											
	Biomet Oxford	Small	37	19	-	22	22				
		Medium	40	21	-	24	24				
		Large	42	23	-	26	26				
		Extra Large	45	25	-	28	28				
Circular											
	Wright Advance	1	39	18	-	23	24				
		2	43	20	-	26	27				
		3	48	22	-	29	30				
		4	53	25	-	32	33				
Combined Circular-Elliptical											
	Biomet Repicci	102100	45	14	-	26	19	31	24	30	29
		102110	48	15	-	27	20	33	26	32	30
		102112	51	16	-	29	21	35	27	35	32
		102120	54	17	-	31	23	37	29	37	34
		102130	60	19	-	34	25	41	32	41	38
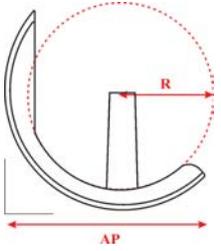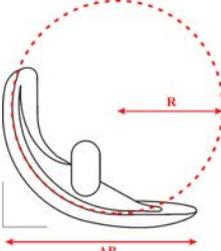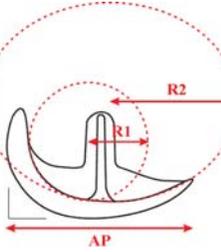
Implant Shape	Implant Name	Sizes	Dimensions (mm)												
			AP Length	L1				L2				L3			
				m	b	xstart	xend	m	b	xstart	xend	m	b	xstart	xend
Planar															
	Zimmer M/G	Small	42	0	5	18	42	-0.18	17	6	18	-4.32	111	5	6
		Small+	44	0	5	19	44	-0.19	18	6	19	-4.53	116	5	6
		Regular	47	0	6	20	47	-0.20	19	7	20	-4.83	124	6	7
		Regular+	50	0	6	22	50	-0.22	20	7	22	-5.14	132	6	7
		Large	52	0	6	22	52	-0.23	21	8	22	-5.35	137	6	8
		Large+	56	0	7	24	56	-0.24	23	8	24	-5.76	148	7	8
		Large++	60	0	7	26	60	-0.26	24	9	26	-6.17	158	7	9
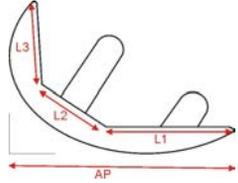
Note: Since the true dimensions are not accessible to the public, the dimensions presented here are approximated from available implant images and sizing charts.