J. De Schutter
H. Van Brussel
Department of Mechanical Engineering
Katholieke Universiteit Leuven
Celestijnenlaan 300B
B-3030 Heverlee, Belgium

Compliant Robot Motion II. A Control Approach Based on External Control Loops

Abstract

A control approach for the execution of robot tasks in contact with the environment is worked out. The input to the controller consists of the task specification described in part I. The control approach is based on external force and tracking loops, which are closed around the robot positioning system. The position control loops tend to decouple and linearize the complex robot dynamics, and therefore they present to the external controllers a system which is easy to model and easy to control. Design and properties of external control loops are discussed in great detail. In particular, the role of a passive compliance with respect to task execution speed and disturbance rejection is analyzed both qualitatively and quantitatively. The resulting compliant motion controller has been tested experimentally, and proved to be very robust and to yield the theoretically expected performance.

1. Introduction

The term "compliant motion" refers to manipulation tasks which involve contact between manipulator and environment, and during the execution of which the end-effector trajectory is modified by the occurring contact forces. Examples are peg-into-hole assembly, following a contour or a surface, opening a door, placing a workpiece against another one, etc. Two kinds of compliance are distinguished. The control system may be programmed to react to force sensor inputs (active compliance), or the contact forces themselves generate the trajectory modifications due to a passive compliance present in the manipulator structure or in the servo. The properties of both active and passive compliance have been discussed thoroughly (Mason 1983; Simons 1980). Many implementations combine the advantages of passive and active compliance (Goto 1980; Simons 1980). This paper is primarily related to active compliance. However, the role of a passive compliance in active compliant motion control is analyzed in depth.

Active force feedback has three important aspects: sensor design, task specification or programming, and task execution or control. Force sensor design does not belong to the subjects treated here (see Van Brussel 1985). The other two aspects deal with integration of the sensor into the robot programming and control system. In this respect, a task specification formalism has been worked out for the programming level in part I of this paper (see pp. 3–17). This formalism produces a task description which contains all the information required in order to allow an entirely automatic execution of the task. In this second part a control strategy is developed which accepts this task description as its input: it allows us to control forces in some directions and velocities in other directions of an operational task frame; it accepts the specification of tracking directions, end-effector and task-frame motion constraints, feedforward velocity information, and task termination conditions. The proposed approach is based on external loops closed around the

The International Journal of Robotics Research,
Vol. 7, No. 4, August 1988,

Fig. 1. Peg-into-hole assembly.

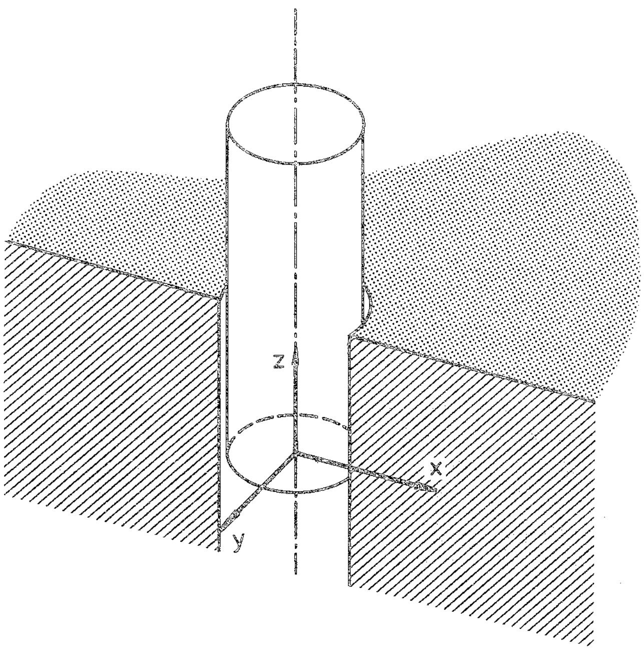

robot positioning system. However, a large part of the procedure, as well as many control properties, can be generalized to other control approaches (De Schutter 1987).

Control of active compliant motion is a complex problem, since the task frame directions do not correspond one-to-one to the joint-space degrees of freedom. As a result, every joint contributes in general both to the execution of motions in the position-controlled and in the force-controlled directions of the task frame, while in the end, whatever the particular control implementation, only a single torque is applied at every joint. Whitney (1987) gives a historic perspective and state of the art in robot force control.

Explicit feedback is the first control technique discussed below. Up to the present it is also the most frequently applied one. Explicit feedback is based on the idea of generalized stiffness or generalized damping: a linear function (represented by a constant matrix) relates the measured force component to desired position increments or desired velocities. These position increments or velocities add to the preprogrammed end-effector position or velocity trajectory. So, in fact, the force-control loops are external loops closed around the joint position or velocity-control loops.

A well-known application of explicit feedback is a peg-into-hole assembly. Define the task frame near the tip of the peg as in Fig. 1. This frame is to follow a trajectory along the z axis while complying with forces along the x and y axes, and to torques around those axes. The stiffness matrix required to perform this task is

$$\mathbb{K}_{ct} = \begin{bmatrix} k_{soft} & 0 & 0 & 0 & 0 & 0 \\ 0 & k_{soft} & 0 & 0 & 0 & 0 \\ 0 & 0 & k_{hard} & 0 & 0 & 0 \\ 0 & 0 & 0 & k_{soft} & 0 & 0 \\ 0 & 0 & 0 & 0 & k_{soft} & 0 \\ 0 & 0 & 0 & 0 & 0 & k_{hard} \end{bmatrix}. \quad (1)$$

In explicit feedback all task frame directions are involved in any part of the control loop; $\mathbb{K}_{ct}$ is a six-by-six matrix. There is no formal distinction between force- and position-controlled directions, as with the control methods discussed below. Yet it is clear that the stiffness matrix (1) hides the functional specification for the peg-into-hole problem mentioned in part I, and implicitly specifies force (k_{soft}) and position (k_{hard}) directions.

However, apart from specifying the relative magnitude of the matrix coefficients, the programmer is also responsible for their absolute magnitude, which has to be interpreted as a servo gain. This coupling between functional specification of the compliant motion mechanism and control implementation constitutes an important drawback of explicit feedback.

Most papers deal with the functional specification aspects of the feedback matrix, but few consider the control aspects. Whitney (1977) focuses on control aspects by investigating a one-dimensional explicit feedback scheme. It is not obvious, however, how to extend this stability analysis to the multidimensional case. Whitney finds that the allowable force feedback gain is inversely proportional to the contact stiffness. A corollary of this statement is that a certain degree of structural or end-effector compliance is always required in order to stabilize the force loop. This is generally indicated as another disadvantage of this type of force control: compliant manipulators or end-effectors handling heavy payloads at high speeds and accelerations (during transfer motions, for example) are subject to considerable position inaccuracies due to static deflections and dynamic deformations at low frequency. However, as we will point out, the need for a passive compliance in the control loop is a general property shared by all compliant motion control approaches (De Schutter 1987).

The free-joint method (Paul 1972) is one of the initial attempts to generate the control torques directly from the sensor readings, thereby avoiding the need for an internal position or velocity loop. For each force-controlled direction, the control system selects a single actuator and a single force component of a sensing wrist. Each associated pair of actuator-force component forms a force servo, which has to approximate the desired compliance in the force-controlled direction of the task frame as nearly as possible. The remaining actuators are position-controlled.

Obviously, since in general the task-frame directions do not coincide with the joint-space degrees of freedom or the wrist sensor directions,

- The obtained compliance in the force directions only approximates the desired compliance.
- The force-servoed joints also cause a compliance in the position directions of the task frame.

Evidently, deviations from the desired position in these directions have to be avoided. Therefore, in order to eliminate accumulated position errors introduced by force-servoed joints, Paul and Shimano (1976) feed appropriate corrections back to the position-servoed joints in a slower loop.

Wu (1980) and Paul (1983) discuss design and stability analysis of a one-dimensional-joint torque-control loop. It is not obvious, however, how to extend this analysis to the multidimensional case. Wu (1980) also presents experimental results from a single-degree-of-freedom manipulator: steady-state loading experiments show good results, but collisions with the environment show large overshoots of the desired force, followed by oscillatory, though quick, recovery.

Raibert and Craig (1981) present a hybrid force/position control scheme for manipulators. However, since Mason (1983) defines a hybrid controller in more general terms as *any* controller based on the division into force and position directions, we will refer to Raibert and Craig's method as the parallel force/position control scheme.

Two complementary and parallel sets of feedback loops (position and force) control a common plant, being the manipulator in contact with the environment. Each task-frame direction is controlled by only one loop, whereas both sets of loops cooperate to control each manipulator joint. Sensing, control, and actuation take place in three different coordinate systems. Although control takes place in the task frame, in general every sensor force component and every actuator participates in each separate control loop. Hence, the parallel force/position scheme is an extension of, and an improvement over, the free-joint method.

The general parallel force/position control method certainly is a very valuable, useful, and consistent concept. However, the two-dimensional implementation presented by Raibert and Craig does not exploit the hybrid idea to its full extent. The difficulties they encounter in obtaining the control parameters and in extending the experiments to multidimensional manipulator and task configurations arise from the fact that the control laws are implemented in joint space instead of in the task frame. The joint-space nature of the control laws results in coupling between position- and force-control loops. This coupling complicates the design and the stability analysis of the loops. In fact, the control parameters have to be obtained by trial and error, and the determination of satisfactory control matrices may therefore be very difficult, if not impossible, in the six-dimensional case.

Khatib (1987) recognizes this problem and solves it by adding parallel force loops to his operational space dynamic position-control scheme (Khatib 1985). Based on a dynamic model of the robot arm, expressed in task-oriented (i.e., operational) coordinates, dynamically decoupled behavior is achieved for all operational directions, i.e., for all position- and force-controlled task-frame directions.

Only Khatib's approach allows the design and stability analysis of the force loops in the multidimensional case. In addition, in all methods discussed above, little attention is paid to the handling of the approach phase, i.e., the transition phase in which contact with the environment is established starting from motion in free space.

This paper presents a method for compliant motion control which has in common with the free-joint

Fig. 2. One-dimensional external force control: general scheme.

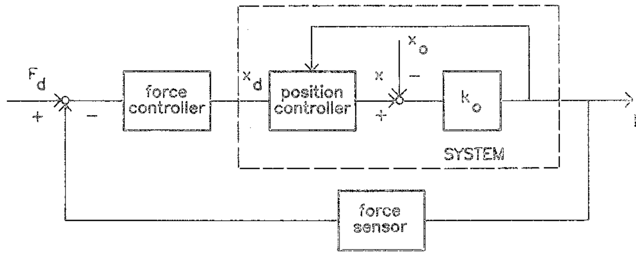

method and with the parallel force/position method that it accepts its input from a functional specification method based on Mason's theory. However, the final control implementation resembles most the explicit feedback method, because it also closes force-control loops around position-control loops.

In this paper:

The force control loops are optimized, allowing general control laws instead of only proportional feedback.
Stability is analyzed in the multidimensional case.
The control method yields good results for the approach phase.

Hirzinger's method (1983; 1985) is based on the same approach.

2. Design of a One-Dimensional External Force-Control Loop

Figure 2 shows the general one-dimensional external force-control scheme. The error between desired force F_d and actual force F is fed into the force controller, which generates position commands x_d. The difference between actual robot position x and the position of the environment x_0 causes a contact force via stiffness k_0. This contact force acts as a disturbance on the position-control loop. The sensor dynamics may be neglected if the structural resonance frequency introduced by the force sensor lies well enough above the position-loop bandwidth. We suppose here that this condition is satisfied.

The effect of the contact force on the position loop may also be neglected in many modern industrial circumstances because

1. The stiffness k_0 is usually small compared to the servo stiffness.
2. The irreversibility of most joint drive systems prevents the passing of end-effector force back to previous links in the kinematic chain.

However, the evolution of industrial manipulators shows a trend toward lightweight constructions, which will be apt to exert controlled forces that make use of full actuator capability. Evidently, in such a case the contact force is no longer negligible. Therefore, we assume that the contact force can either be neglected or is compensated for by direct feedback of the actual measured force to the actuator drive torque.

Both assumptions reduce the general scheme of Fig. 2 to Fig. 3. The design problem consists in finding a suitable force-control law $g_c(s)$, given the closed position-loop transfer function $h(s)$ (with dc gain equal to 1) and the contact stiffness k_0. The response of this control system to an F_d or x_0 input is given by

$$F(s) = \frac{k_0 g_c(s)h(s)}{1 + k_0 g_c(s)h(s)} F_d(s) - \frac{k_0}{1 + k_0 g_c(s)h(s)} x_0(s). \quad (2)$$

For step inputs this results in a steady-state force error if there is no integration in the forward path. Therefore, $g_c(s)$ has to contain a factor $1/s$. On the other hand, $g_c(s)$ also has to contain a factor k_0^{-1} in order to compensate for the stiffness k_0.

The presence of integration in the forward path allows us to define static velocity error coefficients with respect to each of the independent inputs F_d and x_0. For a ramp input $x_0(t) = v_0 t$, the steady-state force error equals

$$\Delta F_{ss} = \frac{1}{\lim_{s \to 0} s g_c(s)h(s)} v_0, \quad (3)$$

and the static velocity error coefficient K_{vx} is defined by

$$K_{vx} = \lim_{s \to 0} s g_c(s)h(s), \quad (4)$$

Fig. 3. One-dimensional external force control: simplified scheme.

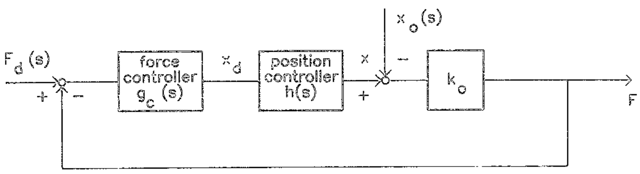

yielding

$$\Delta F_{ss} = K_{vx}^{-1} v_0. \tag{5}$$

The value of K_{vx} is an important property of a force-control law, as we show later. With $g_c(s)$ proportional to k_0^{-1}, one concludes from (4) that K_{vx} is also proportional to the contact compliance k_0^{-1}.

If the transfer function of the internal position-control system is known, the external force-control law can be designed with standard techniques. In the sequel, several approaches are worked out, assuming $h(s)$ is linear and second order:

$$h(s) = \frac{\omega_p^2}{s^2 + 2\zeta_p \omega_p s + \omega_p^2}, \tag{6}$$

where ω_p and ζ_p are the position-control bandwidth and damping ratio, respectively. Typically

$$0.7 < \zeta_p < 1. \tag{7}$$

2.1. Integral Control

Purely integral control, i.e.,

$$g_c(s) = k_0^{-1} k_{fi} \left(\frac{1}{s}\right), \tag{8}$$

is the simplest method and yields good accuracy (no steady-state error). This method resembles the explicit feedback method (generalized damper concept). Therefore, the results presented below give an idea about the (maximum) capabilities of explicit feedback.

The optimal gain k_{fi} is determined by standard control design techniques, yielding

$$k_{fi,\text{design}} \simeq 0.5\omega_p. \tag{9}$$

The resulting bandwidth of the closed force loop is about half the position-loop bandwidth.

Following equation (4), the velocity error coefficient becomes

$$K_{vx} = k_0^{-1} k_{fi} \simeq 0.5 k_0^{-1} \omega_p. \tag{10}$$

2.2. Proportional-Plus-Integral Control

The bandwidth of the closed force loop doubles by adding a proportional control term:

$$g_c(s) = k_0^{-1} \left(k_{fp} + \frac{k_{fi}}{s}\right). \tag{11}$$

Maximum bandwidth results by choosing the break frequency of the control law equal to ω_p:

$$\frac{k_{fi}}{k_{fp}} = \omega_p. \tag{12}$$

Again, the control gains are determined by standard control design techniques, yielding

$$k_{fp,\text{design}} \simeq 1, \qquad k_{fi,\text{design}} \simeq \omega_p. \tag{13}$$

The velocity error coefficient K_{vx} becomes

$$K_{vx} = k_{fi} k_0^{-1} \simeq k_0^{-1} \omega_p. \tag{14}$$

2.3. State Feedback

The state-space approach requires a state-space model of the open-loop system. The physical system has transfer function $h(s)k_0$, but it is very convenient to include the required integrating control factor $1/s$ into the theoretical model. The state-space model for the system then becomes third order.

Theoretically, the closed-loop poles may be chosen arbitrarily, and therefore, the bandwidth of the closed loop can be made arbitrarily large. However, in real applications, this bandwidth should be limited because

1. Physical actuators have limited characteristics, and their saturation causes unmodeled nonlinearities in the control loop.
2. The controller becomes too sensitive to disturbance inputs.
3. The system model is valid only within a bounded frequency range.

Therefore, the absolute location of the optimal desired closed-loop poles is determined by experiment. Their relative location, however, is determined by choosing, e.g., a desired step response.

As with the previous control laws, a static velocity error coefficient is defined. Again, K_{vx} is found to be proportional to the contact compliance k_0^{-1}.

2.4. Conclusion

Three force-control laws have been proposed. Theoretically, state-space control yields the largest bandwidth; however, practical experiments will determine whether this theoretical advantage can be exploited. On the other hand, state-space control is much more complex than integral or proportional-plus-integral (PI) control, and requires a higher sampling rate (in case an estimator is used). An important result is that, even with simple PI control, the external force loop can be made as fast as the internal position-control loop.

3. Properties of External Force Control

3.1. Approach

Compliant motion with force feedback assumes physical contact between end-effector and workpiece or environment. Most references pay little or no attention to the control of the transition phases between motion in free space and motion in contact with the environment. They especially underestimate the importance and the specific problems associated with the approach phase. Many publications implicitly make the following assumptions:

1. The approach phase is executed under position control: the robot moves at reduced speed while the force readings are continuously monitored; a conditional stop occurs upon detection of a predetermined force level. (This control strategy explains the term "guarded move" (Will 1975).) However, a design rule for the approach speed as a function of the allowed collision force is nowhere given.
2. The move out of contact is purely position-controlled.

In our view, identical control schemes apply to both motion in contact and the approach phase; from a control point of view any distinction between both phases is rather artificial.

A constant desired force is applied to the force controller. In the no-contact situation the actual force remains zero, and a steady-state approach velocity results:

$$\dot{x}_{ss} = \dot{x}_{d,ss} = K_{vx}\Delta F = K_{vx}F_d. \qquad (15)$$

Hence, the approach speed is determined by the velocity error coefficient (including the compliance value) and the desired force.

If steady state is reached in the previous phase, the end-effector collides with the environment at a velocity given by (15). Many simulations and practical tests (cf. Fig. 10B) have shown stable force responses with acceptable overshooting of the specified force (typically 10 to 20%).

In conclusion, this method yields a simple and acceptable solution to the approach problem because

During the no-contact phase, the approach speed remains automatically limited.
An acceptable force response results after collision with the environment.

3.2. Tracking Capability

In case of relative motion between robot and workpiece, $x_0(t)$ becomes a dynamic (and usually un-

Fig. 4. External force control in the presence of disturbances. A. Position disturbance. B. Force disturbance in the actuation system.

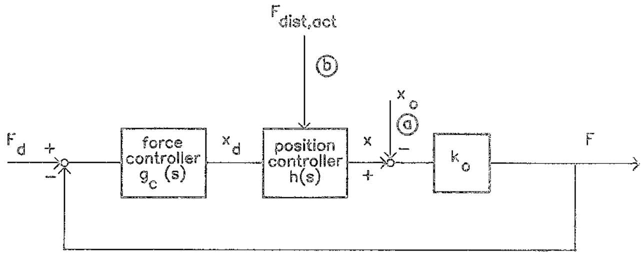

known) disturbance input to the force-control system. Tracking errors (in terms of force) result. For a ramp input ($x_0 = v_0 t$), steady-state is obtained with force error:

$$\Delta F_{ss} = K_{vx}^{-1} v_0. \quad (16)$$

A large velocity error coefficient K_{vx} yields small tracking errors, or, as a corollary, large velocities v_0 are allowed before contact between robot and environment is lost (given a desired force). As a matter of fact, contact is lost if

$$v_0 > K_{vx} F_d. \quad (17)$$

3.3. Disturbance Rejection

Two types of external disturbances are considered (Fig. 4).

3.3.1. Position Disturbance, $x_0(t)$

The previous section analyzed the steady-state force error resulting from an x_0 ramp input. As for the dynamics, the force error has the same dynamics for both a desired force input F_d and for an x_0 input:

$$\Delta F(s) = \frac{1}{1 + k_0 g_c(s) h(s)} (F_d(s) + k_0 x_0(s)). \quad (18)$$

This leads to the following conclusions:

1. A well-designed force-control law automatically possesses a satisfactory position-disturbance rejection (stable and fast recovery, no steady-state error).
2. The maximum force error is determined by the contact stiffness k_0. The more compliance is provided, the stronger the disturbance is rejected.

3.3.2. Force Disturbance in the Actuation System, $F_{dist,act}$

Force disturbances acting on the position loop (e.g., static friction in the actuation) are suppressed by the position controller before affecting the force loop. Assuming a PD-position control scheme, the resulting contact force error is given by

$$\Delta F(s) = \frac{1}{1 + k_0 g_c(s) h(s)} k_0 \left(\frac{h(s)}{k_{pp}} \right) F_{dist,act}(s), \quad (19)$$

where k_{pp} represents the position feedback gain.

This leads to the following conclusions:

1. The force disturbance is filtered by the position-loop transfer function $h(s)$, and is reduced by the relative stiffness k_0/k_{pp}. Again, the more compliance is provided, the stronger the disturbance is rejected.
2. No steady-state error results for a step disturbance.

3.4. Using Force Control in a Velocity Direction

Instead of specifying a constant desired velocity $\dot{x}_d$, a constant force,

$$F_d = K_{vx}^{-1} \dot{x}_d \quad (20)$$

can be specified in a velocity direction. Using force control in a velocity direction has advantages in some cases:

1. If an unexpected collision occurs, the collision force will be limited to F_d (apart from some limited overshoot).
2. Since the motion ends automatically after reaching a force F_d, this control method can be applied for a velocity task with a force terminal condition.

3.5. Adding Feedforward Velocity Information

If x_0 is a known function of time, this information can be introduced into the controller as feedforward position or velocity information (Fig. 5) in order to reduce the tracking error. In a more advanced control scheme, the motion of the workpiece can be estimated from the force readings, based on a model of $h(s)$, and a motion model for the workpiece.

3.6. Influence of the Position Resolution

Practical systems have limited position measurement accuracy. This results in force limit cycles when a constant-contact force is desired in a static situation (i.e., no robot motion is involved). Their peak-to-peak amplitude depends on five factors:

$$\Delta F_{\text{p-to-p}} = (\gamma\delta\epsilon)k_0\,\Delta x_{\text{res}} \tag{21}$$

where

- Δx_{res} is the resolution of the position measurement.
- γ is a factor determined by the relative location within one position increment of the desired robot position which corresponds to the desired force ($\gamma < 1$).
- δ, ϵ are attenuation factors determined by the characteristics of the position-control loop and the force-control loop, respectively ($\delta < 1, \epsilon < 1$).

As a general rule, $k_0\,\Delta x_{\text{res}}$ therefore represents an upper bound for the amplitude of the limit cycles. The limit cycles do not occur in dynamic situations (cf. Fig. 12).

3.7. Role of Stiffness k_0

The stiffness k_0 influences the behavior of a force-control loop in many ways.

1. The influence of the contact force on the position loop is determined by the relative stiffness (k_0 compared with the stiffness of the position servo loop).

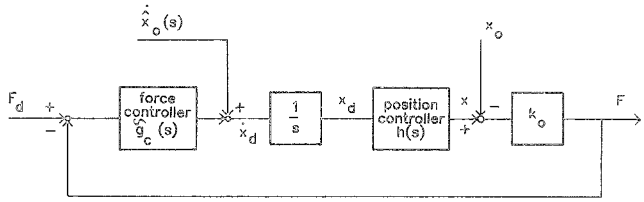

Fig. 5. Adding feedforward velocity information.

2. k_0 directly affects the force resolution and accuracy in the presence of finite position resolution (21).
3. k_0 indirectly affects the obtainable bandwidth of the force-control loop. Indeed, the controller output is proportional to k_0^{-1}. Therefore, the larger k_0^{-1}, the more likely the actuators get saturated, and a decrease of bandwidth becomes necessary.
4. The static velocity error coefficient K_{vx} is proportional to k_0^{-1}, and therefore the execution speed of a compliant motion task is proportional to the degree of passive compliance.
5. The disturbance rejection is proportional to k_0^{-1}.
6. The larger k_0^{-1}, the more deviation exists from the ideal world with infinitely stiff objects. This has consequences in the multidimensional case, particularly when controlling tracking directions.

In conclusion, high structural flexibility is disadvantageous with respect to motions in free space. However, it offers several advantages for compliant motion control, such as a high speed of execution, a high disturbance rejection, and a high force resolution. The first two properties can be generalized to other force-control methods (De Schutter 1987). On the other hand, too much flexibility causes a reduction of bandwidth due to actuator saturations, and disturbs control of tracking directions due to a deviation from the rigid environment hypothesis.

3.8. Applying External Control to Other Types of Feedback

Similar external control schemes apply to any type of sensor feedback where $\Delta x = x - x_0$ is deducible from the measurement. Examples include vision tracking,

Fig. 6. Conceptual organization of multidimensional external force control.

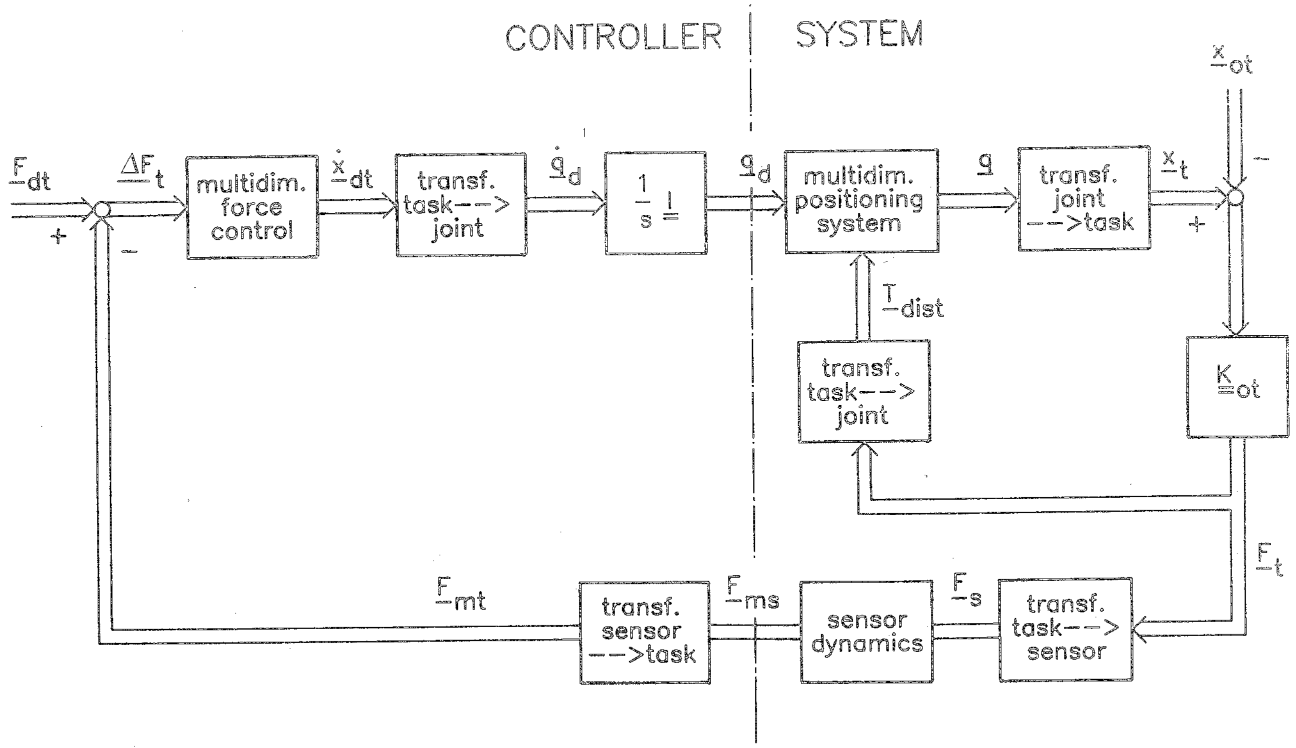

proximity sensing, and tracking directions in multidimensional compliant motion. Clearly, in case of noncontact sensors there is no external force which disturbs the position loop.

4. Control of Multidimensional Compliant Motion Tasks

4.1. Multidimensional Force Control

Figure 6 shows the conceptual organization of multidimensional external force control in case

1. The position control law is implemented in joint space (with coordinates $\mathbf{q}$).
2. The forces are measured at the robot wrist.
3. The stiffness at the contact point is described by matrix $\mathbf{K}_{ot}$.

In order to reduce the complexity, we make the following assumptions:

1. The influence of contact forces on the position loops is supposed to be negligible or compensated for, as in the one-dimensional case.
2. The sensor dynamics are neglected.
3. The position loops are supposed to be dynamically decoupled in the task space; i.e.,

$$\mathbf{x}_t(s) = \begin{bmatrix} h_{xt}(s) & 0 & 0 & 0 & 0 & 0 \\ 0 & h_{yt}(s) & 0 & 0 & 0 & 0 \\ 0 & 0 & h_{zt}(s) & 0 & 0 & 0 \\ 0 & 0 & 0 & h_{\alpha xt}(s) & 0 & 0 \\ 0 & 0 & 0 & 0 & h_{\alpha yt}(s) & 0 \\ 0 & 0 & 0 & 0 & 0 & h_{\alpha zt}(s) \end{bmatrix} \mathbf{x}_{dt}(s). \quad (22)$$

This third assumption is perfectly satisfied if the position-control law is based on nonlinear decoupling in the task space (Khatib 1985). On the other hand, the assumption is approximately valid (after linearization) if the position control is designed using nonlinear

Fig. 7. Simplified multidimensional force-control scheme using external loops.

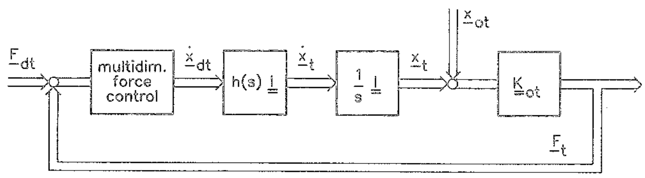

decoupling in joint space and if the dynamics of all decoupled joint motions are chosen to be identical.

Finally, this assumption offers an acceptable working model for many industrial manipulators provided with decentralized controllers as long as the bandwidth and damping ratio of all independent joint-control systems are selected to be almost identical. Indeed, for many industrial manipulators the variation of dynamics and the coupling between the joint-control systems remains sufficiently limited.

Hence, (22) is approximately satisfied with

$$h_{xt}(s) \simeq h_{yt}(s) \simeq \cdots \simeq h_{\alpha zt}(s). \qquad (23)$$

These three assumptions reduce the multidimensional force-control scheme to the scheme of Fig. 7. In this figure the remaining coupling between the force-control loops stems from the stiffness matrix $\mathbb{K}_{ot}$. Introducing its inverse in the multidimensional force control law eliminates this coupling. The transfer matrix of the open-loop system then becomes

$$\mathbb{F}_t(s) = \left(\frac{h(s)}{s}\right)\mathbb{I}\mathbb{x}_{dt}(s), \qquad (24)$$

where $\mathbb{I}$ represents a six-dimensional unity matrix. Six decoupled and identical systems result, which are controlled using independent one-dimensional force controllers.

4.2. Control of General Compliant Motion Tasks

Similar one-dimensional external controllers are also provided for the tracking directions in order to eliminate the orientation errors. These orientation errors are estimated based on contact forces measured in velocity directions, or actual velocities measured in force directions (see part I). In the velocity directions, the desired velocity is already specified, and no external control is needed.

Finally, variable task frames, end-effector (EE) and task-frame (TF) motion constraints, feedforward velocity (ff-vel), and task termination conditions can be introduced as explained in part I, yielding the flowchart of Fig. 8. In this flowchart, after applying the one-dimensional external control laws, the resulting desired velocities in force (f), tracking (tr), and velocity (v) directions are assembled to constitute the desired velocity vector of both the end-effector and the task frame. Then, when applicable, the optional motion constraints are performed on the end-effector velocity vector and the task-frame velocity vector; the optional feedforward velocity is added to both end-effector and task-frame velocities, and, if required, the relation between the task frame and its reference frame is updated. Finally, the end-effector velocity is transformed to joint space, integrated, and output to the joint position controllers.

5. Applications

The task specification formalism described in part I and the control strategy presented above have successfully been applied to an experimental test setup consisting of

- An electrohydraulic Cincinnati-T3 robot (six d.o.f., payload of 45 kg, accuracy of 1.2 mm).
- A six-dimensional force/torque sensor based on strain gauge measurements and described in Van Brussel (1985).
- A six-dimensional compliant structure, with known stiffness matrix, which is mounted between sensor and end-effector in order to provide a sufficient degree of passive compliance.
- An HP1000-A700 minicomputer in which the control flowchart of Fig. 8 is implemented at a sampling rate of 50 Hz.

In order to design the external control loops, we identify the position-loop dynamics $h(s)$. Using step responses in each of the operational directions, we find

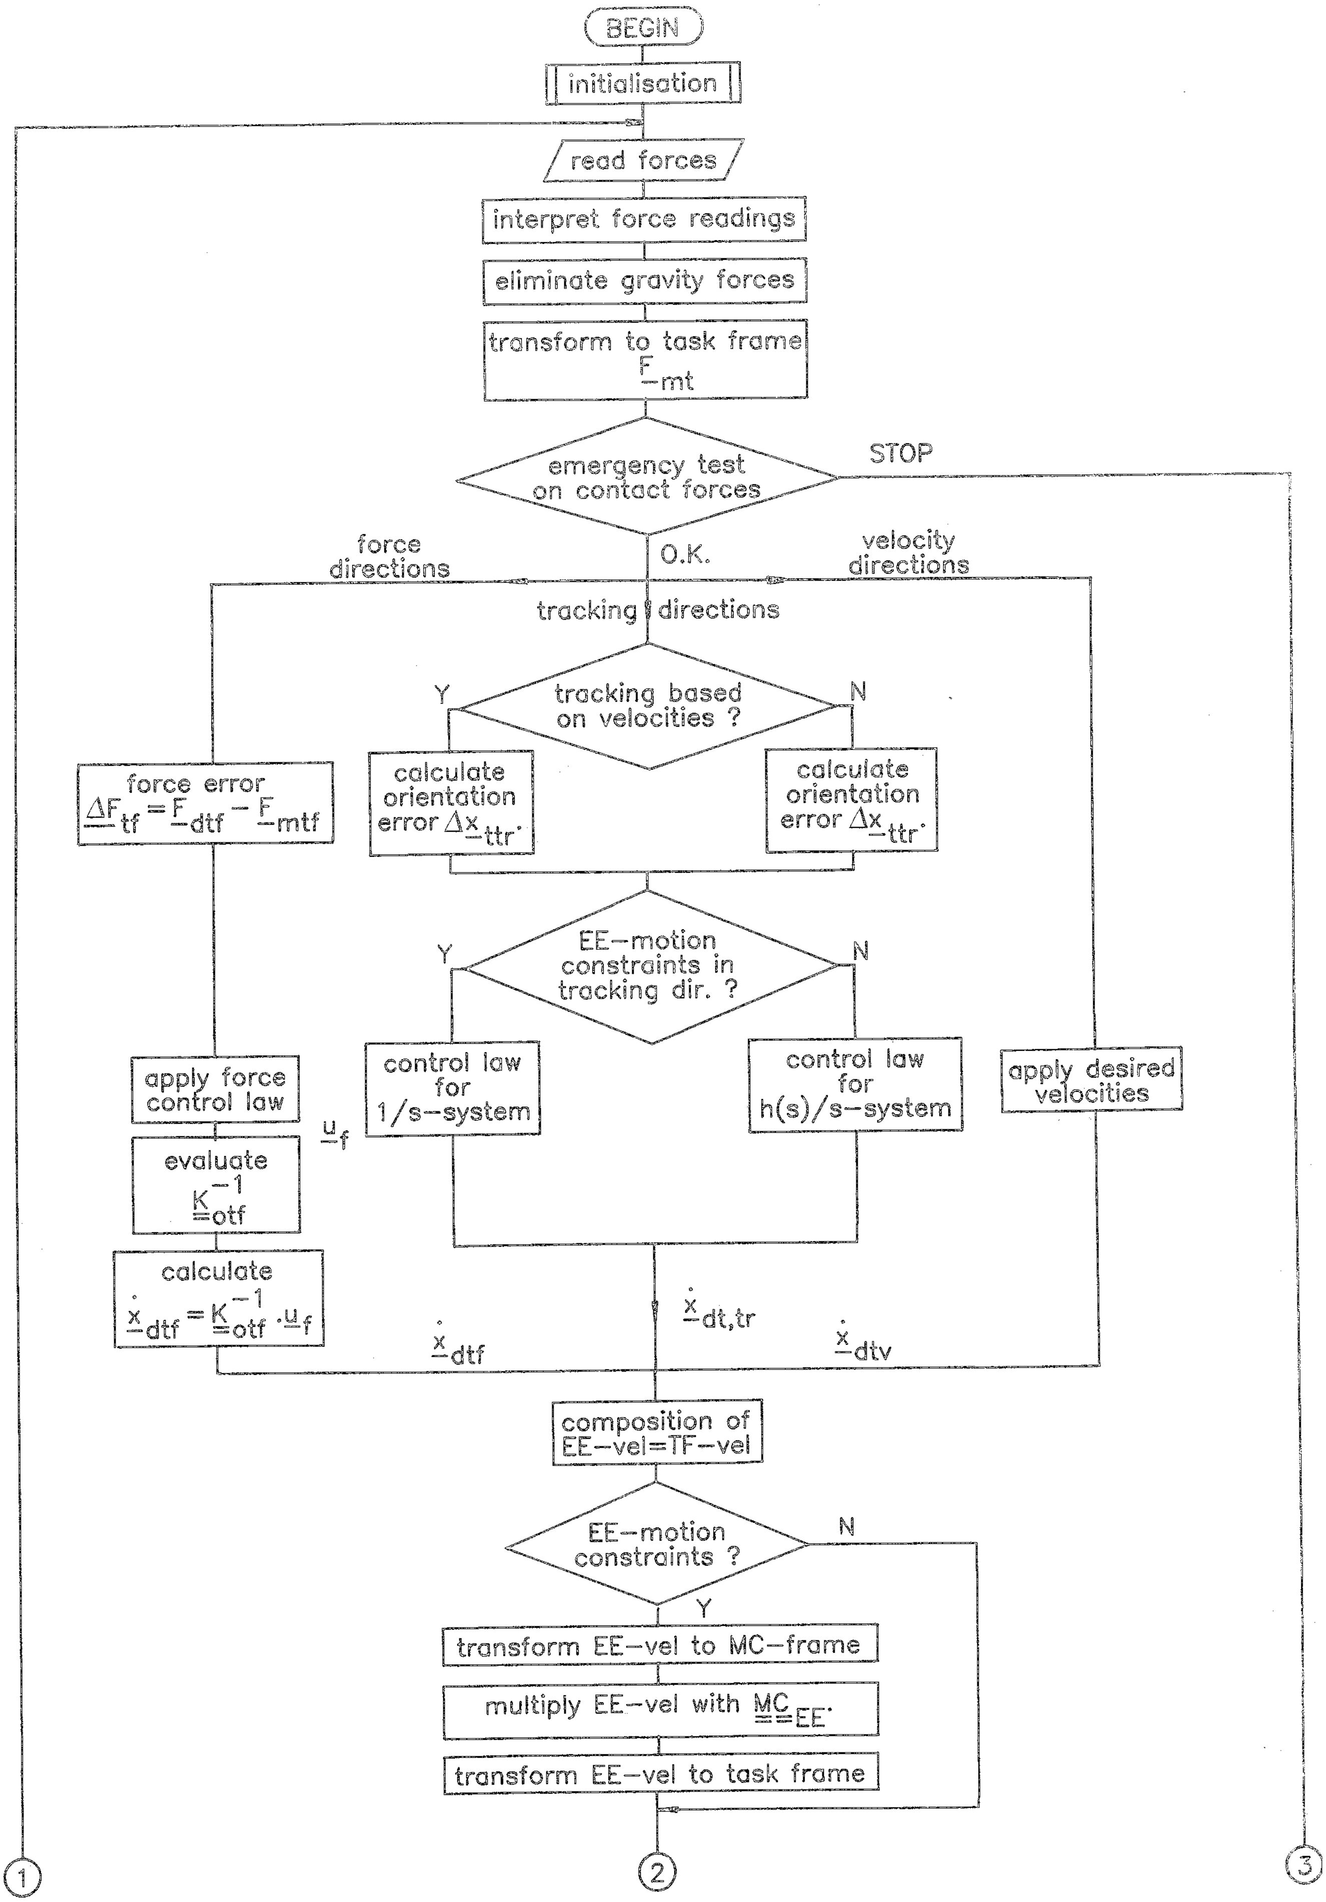

BEGIN
initialisation
read forces
interpret force readings
eliminate gravity forces
transform to task frame
F mt
emergency test on contact forces
STOP
O.K.
force directions
velocity directions
tracking directions
tracking based on velocities ?
Y
N
force error
ΔF tf = F dtf − F mtf
calculate orientation error Δx ttr.
calculate orientation error Δx ttr.
EE−motion constraints in tracking dir. ?
Y
N
apply force control law
u f
evaluate K otf −1
calculate
ẋ dtf = K otf −1 · u f
control law for 1/s−system
control law for h(s)/s−system
apply desired velocities
ẋ dt,tr
ẋ dtf
ẋ dtv
composition of EE−vel=TF−vel
EE−motion constraints ?
N
Y
transform EE−vel to MC−frame
multiply EE−vel with MC EE.
transform EE−vel to task frame
1
2
3

Fig. 8. Flowchart for the implementation of compliant motion control using external controllers.

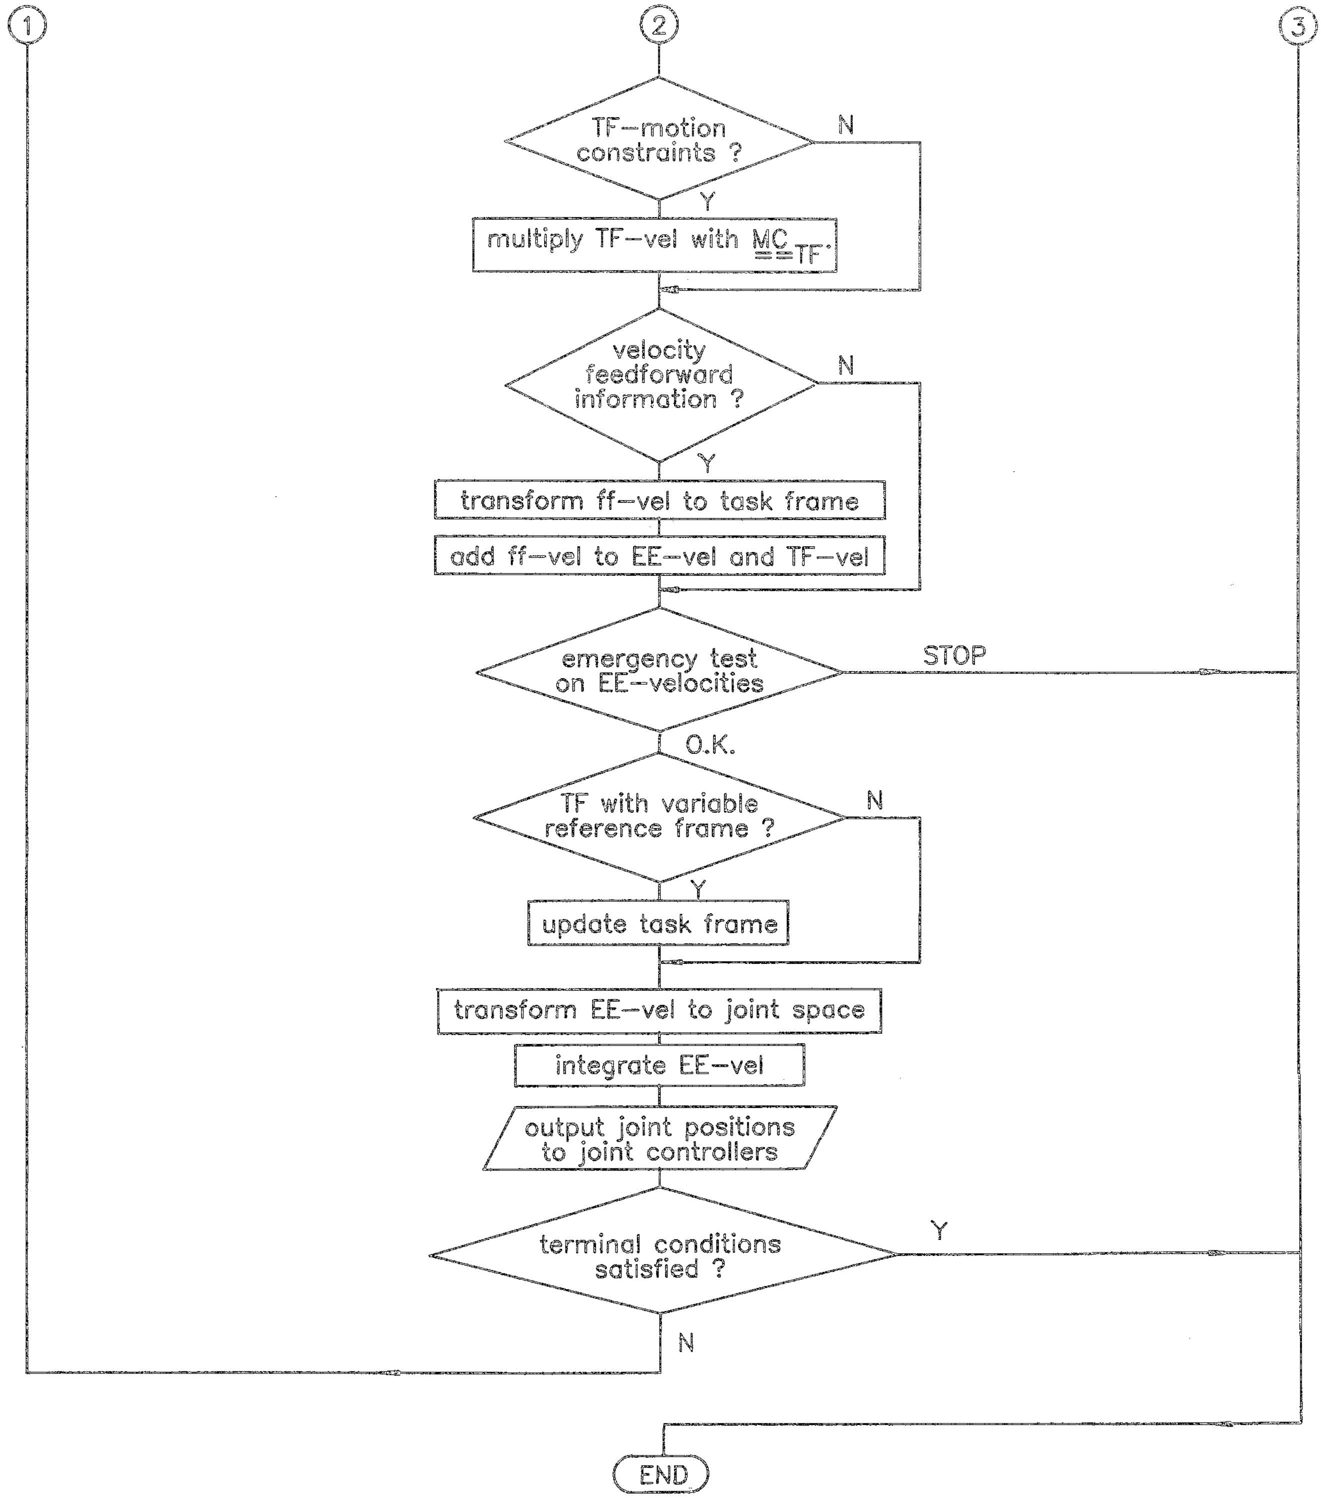

Fig. 9. *Position step response used for the identification of the open-loop system.*

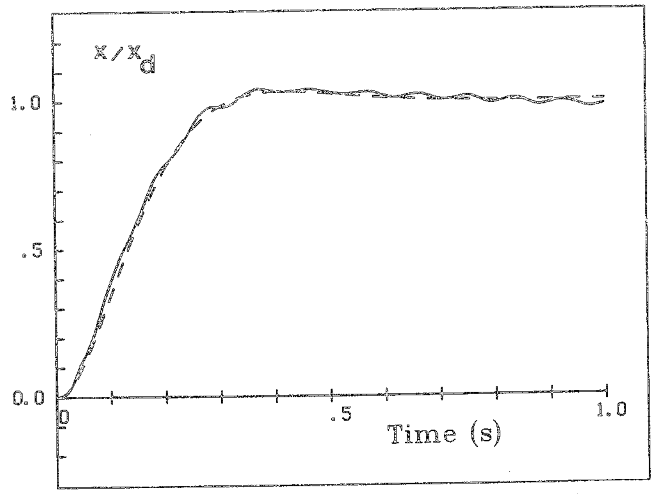

Fig. 10. *PI control. A. Force step response. B. Approach response.*

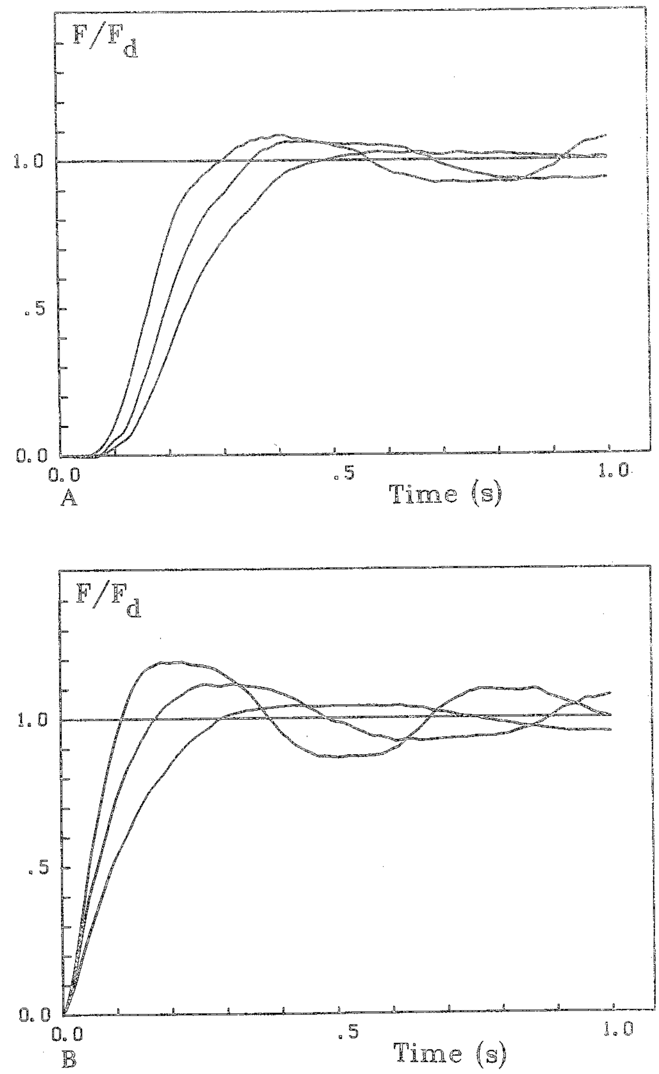

that $h(s)$ is linear and second order (Eq. (6)) with $\zeta_p = 0.75$, $\omega_p = 11$ rad/s, and with negligible dead time (Fig. 9). In addition, a sufficiently strong decoupling exists between the operational space directions.

5.1. Evaluation of the External Force-Control Designs

The control designs of Section 2 are evaluated by applying a constant desired force in one direction of the operational space. Figure 10 shows step responses (i.e., starting from zero contact force) and approach responses (i.e., collisions at steady-state approach speed) for three PI controllers. The fastest one is designed according to Eq. (9). In the other two the control gains are reduced to 71% and 50%, respectively. The desired contact force is 15 N, and the contact stiffness k_0 is about 3 N/mm.

The fastest PI controller has the same response speed as the inner position loop. Indeed, the 90% rise time equals 0.24 s for both a force step (Fig. 10A) and a position step (Fig. 9). A second important result is that this controller results in less than 20% force overshoot after collision with the environment (Fig. 10B). The steady-state approach velocity, given by (15), equals 56 mm/s for the fastest controller and for the given stiffness and desired force.

Similar experiments have been performed using integral and state feedback controllers. As found theoretically, the maximum response speed of integral controllers is about half of the position-loop response speed. As for state feedback, it was impossible to implement stable-state feedback controllers which performed faster than PI control: again, the maximum response speed equals the response speed of the inner position loop. Figure 10 also shows the presence of limit cycles due to the finite position resolution, which is typical for a static situation.

Fig. 11. Workpiece used for two-dimensional contour tracking.

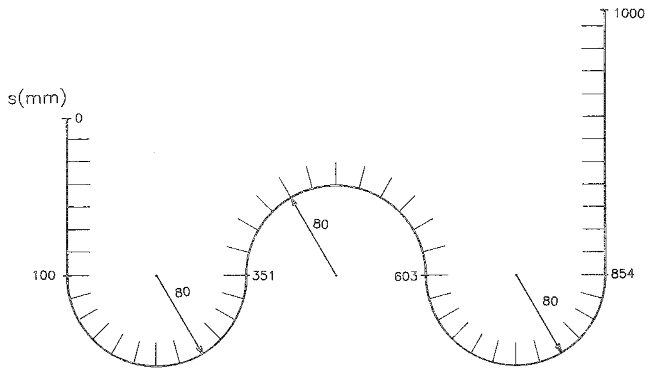

Fig. 12. Two-dimensional contour tracking: force tracking errors. A. V = 10 mm/s. B. V = 20 mm/s.

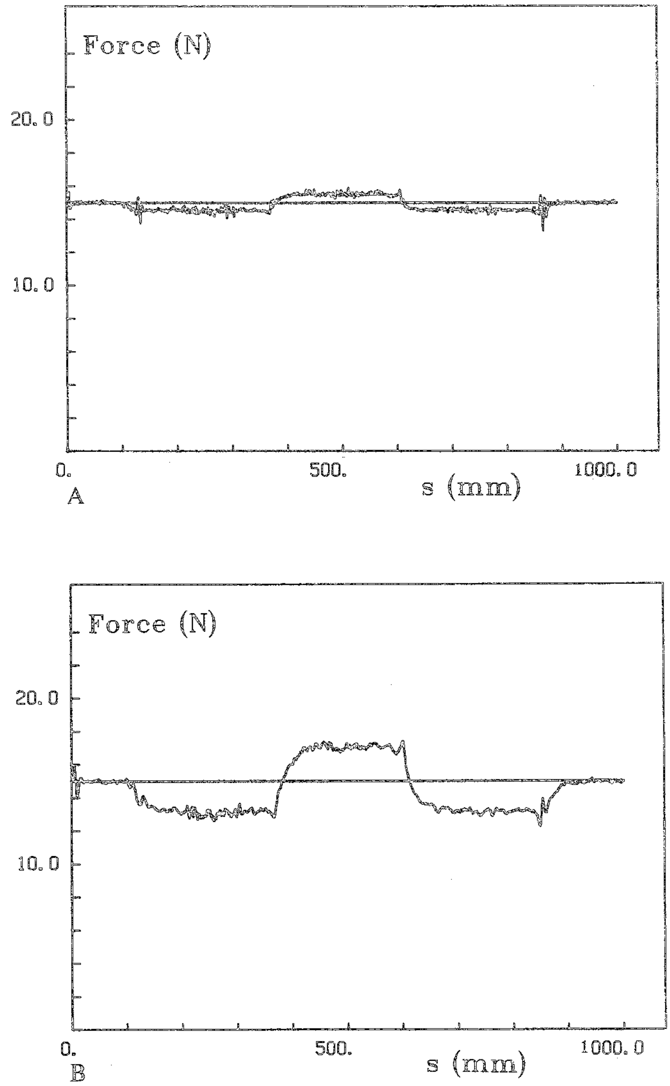

5.2. Two-Dimensional Contour Tracking

Two-dimensional contour tracking involves one force direction (normal to the contour), one velocity direction (tangential to the contour), and one tracking direction (rotation in the plane of the contour). A complete formal description for this task is given in Section 7.4 of part I. The contour (Fig. 11) contains three half circles (radius 80 mm) and two linear parts.

Theoretically, the force tracking error in the circular parts is calculated as

$$\Delta F_{ss} = (K_{vx})_f^{-1}(K_{vx})_{tr}^{-1}\left(\frac{v^2}{r}\right), \qquad (25)$$

where $(K_{vx})_f$ is the velocity coefficient of the force controller, $(K_{vx})_{tr}$ is the velocity coefficient of the tracking controller, v is the tangential velocity, and r is the radius.

A PI controller is selected for the force direction, with $k_{fi} = 8.3\ s^{-1}$ and $k_{fp} = 0.75$, and the contact compliance k_0^{-1} is chosen equal to 0.32 mm/N. This yields a velocity coefficient $(K_{vx})_f = 2.65$ mm/(s · N). On the other hand, a slower, pure I controller suffices for the tracking direction, with $k_{tr,i} = (K_{vx})_{tr} = 1\ s^{-1}$. Equation (25) results in force tracking errors of 0.47 N and 1.93 N for a radius of 80 mm and a tangential velocity v of 10 and 20 mm/s, respectively. Figure 12 confirms these theoretical results. Figure 12 also confirms that no force limit cycles exist in a dynamic situation.

5.3. Other Applications

The control strategy presented in this paper has successfully been applied to a number of compliant motion tasks, such as tracking a three-dimensional welding seam, palletizing of blocks or boxes starting from a partially ordered set, opening a door without knowledge of hinge location or orientation, etc. A very demanding task for the Cincinnati robot consists of a peg-into-hole assembly with diameters of only 30 mm and clearances of 0.1 mm. Successful insertions are guaranteed for orientation errors up to 3°.

6. Conclusion

A control strategy is designed which offers an entirely automatic control solution for every compliant task specifiable, using the formalism presented in part I. This strategy is based on external loops closed around the robot positioning system.

Design and properties of one-dimensional external force loops have been studied in great detail. In particular, external loops

- Offer a simple solution to the approach problem
- Show a high rejection for disturbances in the actuation system
- Require a passive compliance in order to overcome the limited position resolution and to obtain an acceptable execution speed and disturbance rejection
- Have a bandwidth that, in a good design, approaches the position-loop bandwidth

Multidimensional force control and control of general compliant motion is performed by using a set of independent one-dimensional control laws, provided the position loops achieve a sufficient decoupling of the robot dynamics. The general applicability of the task specification formalism as well as the external control strategy has been demonstrated in an experimental test setup built around a Cincinnati-T3 robot.

Favorable conditions are present for implementing compliant motion in a robot programming language (Van Aken 1987):

- The separation between programming and control
- The variety of the potential applications
- The simplicity and the robustness of the control strategy

Acknowledgment

This work was partly sponsored by a research fellowship of the Belgian National Fund for Scientific Research.

References

De Schutter, J. 1987. A study of active compliant motion control methods for rigid manipulators based on a generic scheme. *Proc. IEEE Conf. on Robotics and Automation,* pp. 1060–1065. Washington: IEEE Computer Society Press.

Goto, T., Takeyasa K., and Inoyama, T. 1980. Control algorithm for precision insert operation robots. *IEEE Trans. Syst., Man, Cybernetics* SMC-10(1):19–25.

Hirzinger, G. 1983 (Guadalajara, Mexico). Direct digital control using a force-torque sensor. *Proc. IFAC Symp. on Real Time Digital Control Applications.*

Hirzinger, G., and Landzettel, K. 1985. Sensory feedback structures for robots with supervised learning. *Proc. IEEE Conf. on Robotics and Automation,* pp. 627–635. Washington: IEEE Computer Society Press.

Khatib, O. 1985. The operational space formulation in robot manipulator control. *Proc. 15th Int. Symp. on Industrial Robots,* pp. 165–172. Tokyo: Japanese Industrial Robot Association.

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

Lozano-Perez, T. 1983. Robot programming. *Proc. IEEE* 71(7):821–840.

Paul, R. P. 1972. Modelling, trajectory calculation, and servoing of a computer controlled arm. AIM-177. Stanford, Calif.: Stanford University, Artificial Intelligence Lab.

Paul, R. P., and Shimano, B. 1976 (San Francisco). Compliance and force control. *Proc. Joint Automatic Control Conf.,* pp. 694–699.

Paul, R. P. 1983. *Robot Manipulators, Mathematics, Programming and Control.* Cambridge: MIT Press.

Raibert, M. H., and Craig, J. J. 1981. Hybrid position/force control of manipulators. *ASME J. Dyn. Syst., Meas., Control* 103(2):126–133.

Simons, J. 1980. Force feedback in robot assembly using an active wrist with adaptable compliance. Ph. D. Thesis, Kath. Univ. Leuven (Belgium), Dept. of Mechanical Engineering.

Van Aken, L. et al. 1987 (June 2–3, Stuttgart). LOLA (Leuven off-line language): an enhanced manipulator level off-line robot programming system. *Proc. IFIP Working Conf. on Off-Line Programming of Industrial Robots,* pp. 153–164. Amsterdam: North-Holland.

Van Brussel, H., Belien, H., and Thielemans, H. 1985. Force sensing for advanced robot control. In *Proc. 5th Int. Conf. on Robot Vision and Sensory Control.* ed. N. J. Zimmerman. Bedford (U.K.): IFS.

Whitney, D. E. 1977. Force feedback control of manipulator fine motions. *ASME J. Dyn. Syst., Meas., Control* 99(2):91–97.

Whitney, D. E., 1987. Historic Perspective and State of the Art in Robot Force Control. *Int. J. Robotics Research* 6(1):3–14.

Will, P. M., and Grossman, D. D. 1975. An experimental system for computer controlled mechanical assembly. *IEEE Trans. Comput.* C-24/g:879–888.

Wu, C. H., and Paul, R. P. 1980. Manipulator compliance based on joint torque control. *19th IEEE Conf. on Decision and Control,* pp. 88–94.