



LUND UNIVERSITY

Adaptation of Force Control Parameters in Robotic Assembly

Stolt, Andreas; Linderöth, Magnus; Robertsson, Anders; Johansson, Rolf

Published in:
10th IFAC Symposium on Robot Control

DOI:
[10.3182/20120905-3-HR-2030.00033](https://doi.org/10.3182/20120905-3-HR-2030.00033)

2012

[Link to publication](#)

Citation for published version (APA):
Stolt, A., Linderöth, M., Robertsson, A., & Johansson, R. (2012). Adaptation of Force Control Parameters in Robotic Assembly. In *10th IFAC Symposium on Robot Control* (Vol. 45, pp. 561-566). (IFAC Proceedings Volumes; Vol. 45, No. 22). <https://doi.org/10.3182/20120905-3-HR-2030.00033>

Total number of authors:
4

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Adaptation of Force Control Parameters in Robotic Assembly^{*}

Andreas Stolt^{*} Magnus Linderoth^{*} Anders Robertsson^{*}
Rolf Johansson^{*}

^{*} *Department of Automatic Control, LTH, Lund University, Sweden
(corresponding author e-mail: andreas.stolt@control.lth.se).*

Abstract: Industrial robots are usually programmed to follow desired trajectories, and are very good at position-controlled tasks. New applications, however, often require physical contact between the robot and its environment, and then the position control accuracy is generally not sufficient. Force control is a suitable alternative. The environment is often stiff, and then it is crucial to design appropriate force controllers, which is non-trivial for a robot programmer. This paper presents an adaptive algorithm for choosing force control parameters, based on identification of a contact model. The algorithm is experimentally verified in an assembly task with an industrial robot.

Keywords: Industrial robots, robot control, force control, adaptive control, assembly

1. INTRODUCTION

The robotic applications of today often require physical contact operations between the robot and its environment. Traditionally the solution has been using position controlled robots together with fixtures to achieve the desired accuracy. When the task contains uncertainties, such as part variations, external sensing might be needed. One way to incorporate sensors and specify general tasks is to use the iTaSC framework (De Schutter et al., 2007) (instantaneous Task Specification using Constraints). In (Stolt et al., 2011), it was described how this framework was used in an assembly of an emergency stop button.

Further, there is a need to make it easy for robot operators to specify tasks, especially when external sensing is used. One such example is force controlled assembly. Force sensing is beneficial in these tasks, as it increases the robustness towards uncertainties, e.g., caused by inaccurate gripping, compared to for instance a position controlled implementation. The environment is often stiff, which makes it crucial to design appropriate force controllers. This is a non-trivial task that may be hard for the task programmer. One solution to this problem is to offer a self-tuning mechanism, making the force controllers adaptive.

In this paper, the problem of robotic assembly based on force sensing only is addressed. An adaptive algorithm for choosing force control parameters in a pre-defined controller structure is presented. A contact model is identified, and it is used to tune the force controller. The approach is finally integrated in a sub-assembly of an emergency stop button, previously described in (Stolt et al., 2011). A switch has to be snapped into place in a bottom box. The experimental implementation is based on the iTaSC framework and is made with an industrial robot with

a wrist-mounted 6 degrees-of-freedom force/torque sensor, see Fig. 1.

Identification of contact model parameters has previously been considered by many researchers. In (Erickson et al., 2003), four different methods for estimating the environment contact model were described and experimentally verified. One of the methods was originally presented in (Love and Book, 1995), which describes how the parameters in an impedance controller can be chosen when using contact model parameters estimated with Recursive Least Squares (RLS). A comparison of different algorithms for real-time identification of contact model parameters are described in (Haddadi and Hashtrudi-Zaad, 2008), among them RLS. In (Roy and Whitcomb, 2002), an adaptive force controller is presented, it is based on an estimate of the contact stiffness. A similar approach is presented in (Kröger et al., 2004), which considers adaptive force controllers within the Task Frame Formalism. In (Mallapragada et al., 2006) estimates of contact stiffness and damping are used together with an artificial neural network based gain scheduler for a PI force controller.

An approach to identification of a contact model with multiple contact points is given in (Weber et al., 2006). The geometry is assumed to be known and this makes it possible to calcu-




Fig. 1. The ABB IRB140 robot used in the experiments. The setup for the assembly scenario can be seen in the lower left part of the photo (detailed view in Fig. 4).

^{*} The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 – Challenge 2 – Cognitive Systems, Interaction, Robotics – under grant agreement No 230902 - ROSETTA. This document reflects only the author's views and the European Community is not liable for any use that may be made of the information contained herein.

late the contact locations; the results presented are based on simulations. An extension of the results provided in (Weber et al., 2006) is (Verscheure et al., 2010), which also considers geometric uncertainties and presents experimental results.

A method for designing force controllers when given environment stiffness by the robot user is presented in (Natale et al., 2000). An industrial robot with a position controlled interface is assumed, and the robot dynamics are taken into consideration when doing the controller design.

2. MODELING

2.1 Contact model

The environment is modeled to consist of a spring and a damper, according to Fig. 3. Thus, interacting with the environment gives the reaction force, F , given by

$$F = K_{env}(x_{env} - x) - D_{env}\dot{x} \quad (1)$$

The stiffness of the environment is denoted by K_{env} , the damping by D_{env} , and the location of the unloaded environment by x_{env} . The environment is further assumed to be decoupled, such that there is one relation (1) for each Cartesian direction.

The environment perceived by the robot will not equal the actual environment, it will rather be a combination of the stiffness and damping properties of the tool attached to the robot, the robot itself, and the actual contact. Hence, a stiff environment might be perceived as a soft one if the tool on the robot is soft. Further on, if the tool has different stiffness properties in different directions, even an isotropic contact material will be perceived to have different stiffness in different directions.

2.2 Adaptation algorithm

The algorithm chosen is the Recursive Least Squares (RLS) method. The contact model (1) is nonlinear, because of the product $K_{env}x_{env}$. This product can, however, be seen as a separate model parameter, and then the model is linear. It can be cast in regressor form according to

$$y = \varphi^T \theta, \quad \begin{cases} y = F \\ \varphi = [-x \ -\dot{x} \ 1]^T \\ \theta = [K_{env} \ D_{env} \ K_{env}x_{env}]^T \end{cases} \quad (2)$$

The RLS algorithm is given as (Johansson, 1993)

$$\begin{cases} \hat{\theta}_k = \hat{\theta}_{k-1} + P_k \varphi_k \varepsilon_k \\ \varepsilon_k = y_k - \varphi_k^T \hat{\theta}_{k-1} \\ P_k = \frac{1}{\lambda} \left(P_{k-1} - \frac{P_{k-1} \varphi_k \varphi_k^T P_{k-1}}{\lambda + \varphi_k^T P_{k-1} \varphi_k} \right) \end{cases} \quad (3)$$

The *forgetting factor* λ can be used to cope with time varying parameters by setting it to a value less than 1. A value of $\lambda = 1$ gives the usual least-squares solution. The initial value of the adaptation gain matrix P has to be chosen, and its magnitude is usually chosen to be large to get a fast convergence to the true values of the estimated parameters.

Each force controlled direction will have nominal parameters, likely not well tuned. These will be used during the estimation of the contact parameters. To assure that the input signals to the estimator are persistently exciting (Johansson, 1993), the




Fig. 2. Illustration of pole placement design (8).

Fig. 3. Contact model.

force reference is set to a sufficiently exciting signal, e.g., a square wave. As the covariance of the estimate is decreased (proportional to the P matrix), the controller parameters are updated based on the contact model parameter estimates. Once the covariance is considered to be low enough, this adaptation phase is finished.

2.3 Force controller

The force control used in the assembly framework is decoupled impedance control (Hogan, 1985) for each Cartesian direction x . This setup makes it possible to perform force control in some directions and, e.g., position control in others. The control law used for the impedance controller is given by (4).

$$\ddot{x}_{des} = \frac{1}{M} (F_x - F_{x,ref} - D\dot{x}_{des}) \quad (4)$$

The direction controlled is denoted by x and its desired behavior by x_{des} (the control signal), F_x and $F_{x,ref}$ denote the force and the force reference in the direction of x , respectively. The parameter M is the virtual mass and D the virtual damping of the impedance the direction is controlled to behave like. No position reference is used, as it is only interesting to control the force.

To make the controller safe to use, the maximum output velocity (\dot{x}_{des}) is limited. This limitation is made in such a way that no wind-up problems occur. The switching between different control modes, e.g., from position to force control, is made by bumpless transfer, i.e., the new controller is initially set to have the same control signal as the previous controller.

2.4 Choice of force control parameters

The parameters are chosen according to a pole placement design, of the poles in the transfer function from the force reference, F_{ref} , to the measured force, F . The controller (4) together with the contact model (1), where the location of contact is ignored, gives

$$\begin{cases} \ddot{x} = \frac{1}{M} (F - F_{ref} - D\dot{x}) \\ F = -K_{env}x - D_{env}\dot{x} \end{cases} \quad (5)$$

In (5) the assumption of an ideal velocity controlled robot is made, i.e., $\dot{x} = \dot{x}_{des}$. The time domain equations can be transformed to the frequency domain by the Laplace transform, giving

$$\begin{cases} s^2 X(s) = \frac{1}{M} (F(s) - F_{ref}(s) - DsX(s)) \\ F(s) = -K_{env}X(s) - D_{env}sX(s) \end{cases} \quad (6)$$

By eliminating $X(s)$ in the above equations the following relation between F_{ref} and F is achieved

$$F(s) = \frac{\frac{K_{env}}{M} + \frac{D}{M}s}{s^2 + \frac{D_{env}+D}{M}s + \frac{K_{env}}{M}} F_{ref}(s) \quad (7)$$

Hence, the measured force is related to the force reference by a second-order linear time-invariant dynamical system. A stable pole placement design for such a system can be parameterized according to Fig. 2, which gives the denominator polynomial

$$s^2 + 2\zeta\omega s + \omega^2 \quad (8)$$

Comparison of the coefficients of the denominator in (7) and the specification polynomial in (8) gives that the force control parameters should be chosen as (estimated values of contact stiffness and damping should be used)

$$M = \frac{K_{env}}{\omega^2}, \quad D = \frac{2\zeta K_{env}}{\omega} - D_{env} \quad (9)$$

The actual force controllers are implemented in discrete time, handled by discretization of the control law (4) (the sampling period used was 4 ms). The largest approximation is the assumption of neglected robot dynamics in the realization of the control law (4). This will only be approximately true up to a certain bandwidth, and the stability margins will depend on unmodeled dynamics, e.g., robot stiffness dynamics and time delays originating from sensor processing. The bandwidth of the force controller, ω , will thus have to be chosen with these considerations taken into account.

2.5 Torque control parameters

Torque control during assembly operations often means two or more point contacts. A change in the torque reference will therefore change the measured force, as there is a coupling between the measured force and torque. Usually the contact material for all contacts is approximately the same, which means that the same contact model that was identified during the first phase with only one contact can be reused. The remaining uncertainty is about the location of the second contact relative to the first, and this can be estimated, e.g., with an RLS estimator. Once the location of the contact is estimated, the formulas for controller parameters in Sec. 2.4 can once again be used, with the stiffness $\hat{K}\hat{L}$ and the damping $\hat{D}\hat{L}$, where \hat{L} is the estimated distance between the two contact points.

2.6 Alternative specification of torque control

When performing assembly operations with two-point contacts it is not always easy to choose appropriate set point values for the force and the torque controllers. An alternative is to instead control the force in each contact. The estimation outlined in Sec. 2.5 gives the required information about the relative location of the contacts, i.e., the distance between them. This makes it possible to calculate the force originating from each contact, and transform a reference on the forces in each contact to an equivalent force and torque.

This way of specifying the force and the torque during a two-point contact assembly operation will simplify the procedure for the user. The easiest way to implement it is to transform the two-force reference from the user, to a force and torque reference, and keeping separate control of force and torque. The user should, however, be presented with measurements transformed into forces from two contacts.




Fig. 4. Illustration of the coordinate frames used in the assembly task.




Fig. 5. A part of the state machine implementing the assembly sequence. The parameter *adapt* decides whether the adaptation phase should be entered or not.

2.7 Assembly task modeling

The task has been modeled with the iTaSC framework (De Schutter et al., 2007). Two frames have been used to describe the assembly operation, shown in Fig. 4.

- Frame *f1* is attached to the bottom box.
- Frame *f2* is attached to one end of the switch that is gripped by the robot.

The outputs chosen to be controlled are the three Cartesian translations in *f1* and the orientation parameterized by three Euler angles in *f2*.

2.8 Assembly strategy

The location of the bottom box is not assumed to be known exactly. It is therefore not possible to use position control to assemble the switch. A sequence of search motions is used to find the slot in the bottom box that the switch should be assembled into. The assembly sequence can be summarized as:

- (1) Move to start position
- (2) Search for contact in *f1* *z*-direction
- (3) Search for contact in *f1* *y*-direction
- (4) Search for contact in *f1* *x*-direction
- (5) Search for contact in *f2* *x*-rotation direction
- (6) Search for contact in *f2* *z*-rotation direction
- (7) Push down switch (rotation around *f2* *x*-axis)
- (8) Move robot away

Each contact will be force controlled once it has been established. This way the assembly sequence will be robust against uncertainties in the starting position and variations in the involved parts.

The adaptation strategy described in Sec. 2 should only be used when the force control parameters are not well tuned, i.e., usually the first time the assembly is performed or when something has changed, e.g., at the use of a new gripper. The adaptation phases can be considered as separate states in between the nominal ones, see a part of the state machine implementing the sequence in Fig. 5.

3. EXPERIMENTAL RESULTS

All experiments in this section were performed using an ABB IRB140 industrial robot, together with the IRC5 control system. This was extended with an open robot control system, enabling

external sensor integration by modification of the references to the internal servo controller (Blomdell et al., 2005, 2010).

3.1 Contact with different materials

An experiment where contact was made with three different environments was used to test that the adaptation gave the desired performance. An initial search towards the environment was made until a contact force was detected. A force controller was then started with poorly tuned parameters, i.e. a default initial setting, and the adaptive algorithm was initiated. The given force reference was a square wave, and the forgetting factor λ was chosen to be 1, the environment was not assumed to vary over time. Once the covariance of the contact model parameter estimates became low, the force control parameters were updated. A bandwidth of $\omega = 5$ [rad/s] and a relative damping $\zeta = 0.8$ was chosen for the controller.

Data from the experiment is shown in Fig. 6. In the top-most diagrams the environment was a soft plastic foam. The fact that the material was soft can be seen in the force response when contact was made, as the force is slowly built up. The nominal force control parameters were used in the first period of the reference signal, and the parameters were so poorly tuned that hardly anything happened. When the estimated contact parameters were used, however, the reference was satisfactorily tracked. The estimates of the contact stiffness and the damping can be seen to converge in less than 5 seconds.

The second environment used was a mouse pad, displayed in the middle diagrams in Fig. 6. This material was stiffer, but both the control and estimation behavior was similar to the first case. The last environment was a table surface, displayed in the bottom diagrams in Fig. 6. The initial force transient shows that this environment clearly was the stiffest. When the estimated contact parameters were used, the resulting control performance was worse than in the two previous cases. This was probably caused by that the assumptions made when deriving the control parameters were not completely valid for the chosen control bandwidth and the stiffness of the contact material. Even though the performance is worse than for the previous environments, it is acceptable in regular assembly tasks.

The estimate of the stiffness starts with a large transient for all materials, which is caused by the choice of a large initial covariance. Choosing it smaller, however, would lead to slower convergence for the parameter estimates.

3.2 Adaptation in an assembly sequence

The adaptation strategy was used to tune the force control parameters in an assembly sequence, experimental data is displayed in Figs. 7-10. Force data from the beginning of the sequence is shown in Fig. 7. State 2 was the search motion in $f1$ z -direction, and the adaptation of the force control parameters for the z -coordinate was started in state 2.5, when contact was detected. The initial parameters were poor, which the large initial force transient shows. On the other hand, the transient gave good excitation for the estimation algorithm. Initially, the force reference in state 2.5 was a sinusoid, to get a reference that would not be too hard for the poor controller to follow. Once the covariance of the estimate decreased below a threshold, the reference was switched to a square wave, to get more excitation. In order not to disturb the estimation algorithm, all other output directions were controlled to keep their current position

during the adaptation phase. This phase was finished once the covariance decreased below a second threshold.

Search motions and adaptation in the $f2$ y - and x -directions then follow. Here it can be noted that the initial transients are much lower than for the z -direction and that the adaptation phases lasts somewhat longer. State 5 was the rotational search around the $f2$ x -axis, where the forces were controlled to be constant to keep the contact.

The identified contact model parameters and the norm of the P -matrix (a measure of the size of the covariance) are shown in Fig. 8. It can be seen that the contact in the z -direction was considerably stiffer than in the other directions. The contact material itself had approximately the same properties in all directions, but the gripper and the switch was much stiffer in the z -direction than in the others. The slower convergence for the estimation in the y -direction can clearly be seen in the plot of the norm of P . Occasionally the algorithm gave unreasonable estimates, such as negative parameters, and to avoid problems with this the algorithm was supervised. The values used for choosing force controllers were projected into allowed intervals, see e.g., the damping parameter around $t = 14$ [s]. The estimation of the contact location, x_{env} in (1), is not shown because it is not relevant for the assembly sequence, but it also converged to a reasonable value for each contact model.

The search speeds in the assembly sequence had to be slow to handle the initial force control parameters, see e.g., the transient in the z -force in Fig. 7 at $t = 4$ [s]. Once the control parameters had been tuned, it was possible to increase all search speeds.

The data shown in Figs. 7 and 8 have only been a one-point contact. The two-point contact was made in the second part of the assembly, see experimental data in Figs. 9 and 10. The adaptation for the torque controller (around $f2$ x -axis) started when the two-point contact was detected, i.e., when state 5.5 was entered. The resulting controller, active in the end of the adaptation phase, shows some overshoots when the reference is a square wave. This means that better reference tracking probably can be achieved by decreasing the control bandwidth, but this is not good for the performance in the assembly sequence, where it needs to react fast to disturbances caused by movements in other directions not to lose contact. The following action in the assembly sequence was to find the slot with the second contact point, by a search around the $f2$ z -axis. Once found, detected by a large z -torque, the switch was pushed down until it was correctly inserted. Finally, the whole assembly was lifted to show that the sequence had finished.

The estimation of the distance between the two contact points is shown in Fig. 10. The estimate initially varies, and even becomes negative, which is handled by the previously mentioned supervision of the algorithm. A negative distance is further considered to be more of an issue than a negative damping parameter, so the P -matrix was also reset to a larger magnitude to restart the estimation. The estimate finally converges to approximately 34 [mm], which is within 1 [mm] from the true value.

3.3 Alternative specification of torque reference

The approach where the user specifies two forces instead of one force and one torque in a two-point contact situation (Sec. 2.6) was implemented in the assembly sequence. The only relevant state in the sequence was state 6, and experimental data from




Fig. 6. Experimental data from an experiment where contact was made with different environments. The top diagrams show contact with soft plastic foam, the middle diagrams contact with a mouse pad, and the bottom diagrams contact with a table surface. The left diagrams show the measured force in blue and the force reference in black. The right diagrams show estimated stiffness and damping.




Fig. 7. Experimental data from the beginning of the assembly sequence. The top diagram shows the state sequence, with state numbers defined in Sec. 2.8. The remaining diagrams show the measured force (blue) and the force reference (black) for the coordinate directions in frame $f1$. The reference is only shown when the coordinate is force controlled. The adaptation phases are marked with vertical green lines.




Fig. 8. Experimental data from the beginning of the assembly sequence. The top diagram shows the state sequence, the second the stiffness parameter estimate, the third the damping parameter estimate, and the last diagram the norm of the P -matrix. The beginning of each adaptation phase is marked with a green line. The first phase is for the parameters in the z -direction, the second in the y -direction, and the third in the x -direction.




Fig. 9. Experimental data from the final part of the assembly sequence. The top diagram shows the state sequence. The second diagram shows the measured force (solid lines) and the force reference (dashed lines) for the coordinate directions in frame $f1$. The third diagram shows the measured torque (solid lines) and the torque reference (dashed lines) around the coordinate axis of frame $f2$. The adaptation phase is marked with vertical green lines. Only the torque around the x -axis is controlled in the adaptation phase.




Fig. 10. Experimental data from the final part of the assembly sequence. The top diagram shows the state sequence, the second diagram the distance parameter estimate (between the contact points), and the last diagram the norm of the P -matrix. The adaptation phase is marked with green lines.




Fig. 11. Experimental data from state 6 in the assembly sequence. The top diagram shows the z -force, the middle diagram the x -torque, and the bottom diagram the estimated equivalent forces acting on the two ends of the switch. Measured force/torque are shown with solid lines and references with dashed lines.

this state is shown in Fig. 11. The first contact point was the end of the switch that first made contact, and the force reference for this point was set to 5 [N], enough to not lose contact. It was desired that the other end of the switch slides down into the slot, and the reference was therefore a larger force, here 15 [N]. By using the identified distance between the two contact points,

the given specification was translated to a force and a torque reference, see the two top diagrams in Fig. 11.

The control performance is good in the beginning of the time slot shown in Fig. 11. The references are lost in the end, and this was caused by that the switch slid down into the slot,

i.e., contact was lost for the second point. This can also be seen in the torque diagram, as the torque approaches 0. This behavior is an indication of a successful assembly.

4. DISCUSSION

The adaptation algorithm described in Sec. 2 was successfully implemented on an industrial robot system. The achieved performance is satisfactory, both for soft and stiff contacts, and it can be used to free the user from the tedious work of tuning the force controllers manually. Some performance degradation for stiff contacts is present that is not foreseen by the design procedure. This is caused by a too coarse approximation of the robot dynamics, by making the assumption of an ideal velocity controlled robot. To get a better control design, which considers the limitations of the robot system, also the robot dynamics has to be modeled.

An option that might enhance the control performance is to resort to an optimal controller, e.g., an LQG or H_∞ -controller. But this means that the impedance control structure has to be abandoned, which might not be desirable. The impedance control parameters have a physical interpretation that might be valuable, e.g., in an error situation.

The adaptation is currently implemented as separate states in the controlling state machine. A dedicated excitation signal is used to assure that input data to the estimation algorithm is sufficiently exciting. A further development could be to run the adaptation algorithm in each assembly operation without an excitation signal.

The contact locations have been estimated during the assembly sequence, but this information has so far not been used. One way to use it is to decrease the search times, by starting the search motions closer to the identified contact locations, after a number of successful assembly operations have been performed.

The method of using two forces instead of one force and one torque in a two-point contact scenario simplifies the task specification for the user, as the coupling between the force and the torque can be ignored. Generalizing the strategy to more than two-point contacts is hard, as the conversion from force and torque measurements to multiple forces is very hard or even impossible to solve.

To the best of the authors knowledge, a similar approach has not been previously presented within assembly. Adaptive force control with comparable results has been performed, e.g., in (Roy and Whitcomb, 2002) and (Kröger et al., 2004). They both show similar results for corresponding contact stiffnesses, but this paper also considers significantly stiffer contact environments. A stiffness of over 100 [N/mm] was estimated in Fig. 6, compared to a stiffness around 20 [N/mm] in (Kröger et al., 2004) and below 1 [N/mm] in (Roy and Whitcomb, 2002).

5. CONCLUSIONS

A method for self-tuning of force controllers to use in industrial robots has been described. It was based on identification of a contact model using an RLS algorithm. The force controller considered was an impedance controller and its parameters were chosen according to a pole placement design. The method was implemented on an industrial robot system and used in an assembly task.

REFERENCES

- Blomdell, A., Bolmsjö, G., Brogårdh, T., Cederberg, P., Isaksson, M., Johansson, R., Haage, M., Nilsson, K., Olsson, M., Olsson, T., Robertsson, A., and Wang, J. (2005). Extending an industrial robot controller—Implementation and applications of a fast open sensor interface. *IEEE Robotics & Automation Magazine*, 12(3), 85–94.
- Blomdell, A., Dressler, I., Nilsson, K., and Robertsson, A. (2010). Flexible application development and high-performance motion control based on external sensing and reconfiguration of ABB industrial robot controllers. In *Proc. ICRA 2010 Workshop on Innovative Robot Control Architectures for Demanding (Research) Applications*, 62–66. Anchorage, Alaska, USA.
- De Schutter, J., De Laet, T., Rutgeerts, J., Decré, W., Smits, R., Aertbeliën, E., Claes, K., and Bruyninckx, H. (2007). Constraint-based task specification and estimation for sensor-based robot systems in the presence of geometric uncertainty. *Int. J. Robotics Research*, 26(5), 433.
- Erickson, D., Weber, M., and Sharf, I. (2003). Contact stiffness and damping estimation for robotic systems. *Int. J. Robotics Research*, 22(1), 41.
- Haddadi, A. and Hashtrudi-Zaad, K. (2008). Online contact impedance identification for robotic systems. In *Proc. Int. Conf. Intelligent Robots and Systems (IROS)*, 974–980. Nice, France.
- Hogan, N. (1985). Impedance control: An approach to manipulation: Parts i-iii. *ASME J. Dynamic Systems, Measurement, and Control*, 107, 1–24.
- Johansson, R. (1993). *System Modeling and Identification*. Prentice Hall, Englewood Cliffs, NJ.
- Kröger, T., Finkemeyer, B., Heuck, M., and Wahl, F. (2004). Adaptive implicit hybrid force/pose control of industrial manipulators: Compliant motion experiments. In *Proc. Int. Conf. Intelligent Robots and Systems (IROS)*, 816–821. Sendai, Japan.
- Love, L. and Book, W. (1995). Environment estimation for enhanced impedance control. In *Proc. Int. Conf. Robotics and Automation (ICRA)*, 1854–1859. Nagoya, Japan.
- Mallapragada, V., Erol, D., and Sarker, N. (2006). A new method of force control for unknown environments. In *Proc. Int. Conf. Intelligent Robots and Systems (IROS)*, 4509–4514. Beijing, China.
- Natale, C., Koeppe, R., and Hirzinger, G. (2000). A systematic design procedure of force controllers for industrial robots. *Mechatronics, IEEE/ASME Transactions on*, 5(2), 122–131.
- Roy, J. and Whitcomb, L. (2002). Adaptive force control of position/velocity controlled robots: theory and experiment. *Robotics and Automation, IEEE Transactions on*, 18(2), 121–137.
- Stolt, A., Linderth, M., Robertsson, A., and Johansson, R. (2011). Force Controlled Assembly of Emergency Stop Button. In *Proc. Int. Conf. Robotics and Automation (ICRA)*, 3751–3756. Shanghai, China.
- Verschuere, D., Sharf, I., Bruyninckx, H., Swevers, J., and De Schutter, J. (2010). Identification of contact parameters from stiff multi-point contact robotic operations. *Int. J. Robotics Research*, 29(4), 367–385.
- Weber, E., Patel, K., Ma, O., and Sharf, I. (2006). Identification of contact dynamics model parameters from constrained robotic operations. *ASME J. Dynamic Systems, Measurement, and Control*, 128, 307–318.