INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 20 Motion Control, Robotics, and Autonomous Systems

CHAPTER 10

Robot Conveyor Picking Simulator for Late Camera Reports and Moving-Target Interception

Pages 235–257

Edited by John D. Hedengren

How to cite this chapter. "Robot Conveyor Picking Simulator for Late Camera Reports and Moving-Target Interception," in *Motion Control, Robotics, and Autonomous Systems*, John D. Hedengren, Ed., vol. 20 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 10, pp. 235–257.

The page numbers in this file are those of the printed volume.

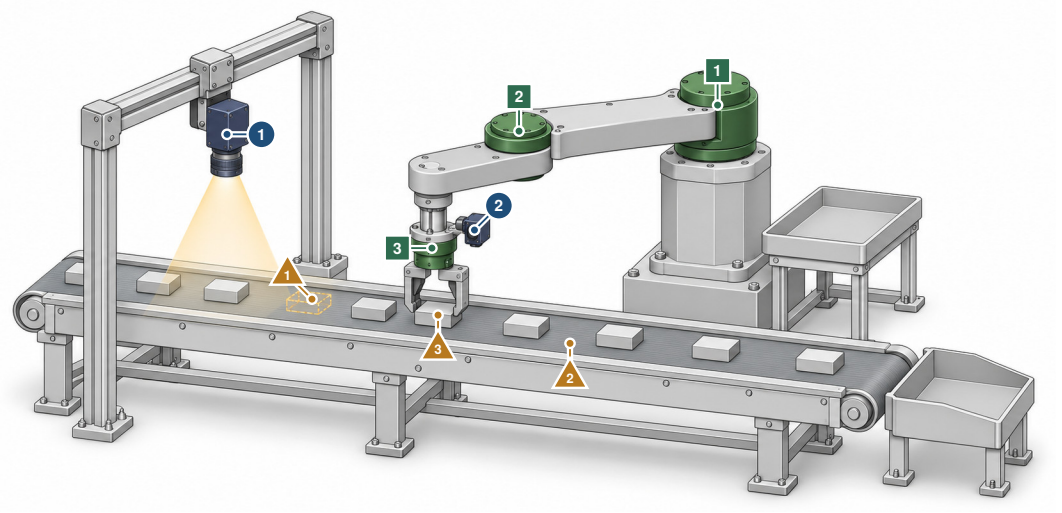
The app for this chapter
https://apmonitor.com/apps/robot-conveyor-picking/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v20-motion-robotics/#robot-conveyor-picking

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 10

Robot Conveyor Picking Simulator for Late Camera Reports and Moving-Target Interception



The controller moves 1 shoulder motor, one torque 2 elbow motor, one torque 3 gripper, which closes and lets go

The controller reads 1 camera upstream, which reports each box late 2 switch that tells a box is held

What disturbs it 1 where the camera last saw a box 2 belt speed 3 mass of the box

Notice that the camera watches the belt upstream of the arm, so every report describes a box that has already moved on, and the gripper (a vacuum cup in the chapter, drawn here with jaws) must close where the box will be. Illustration, not to scale.

IN PLAIN TERMS Catching Boxes on a Moving Belt

The main idea. A robot arm picks small boxes off a moving belt and sets them in a tray. A camera upstream finds each box, but the box has moved on by the time the arm gets there.

The challenge. The arm has to grab each box where it will be, and its hand has to move with the belt as it closes. It must also carry the box gently, or its grip lets go.

Why it matters. On a fast belt, an arm that aims at old reports or waits standing still puts no box in the tray.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment and the fact on early robots (Section 10.2), the try-it segments of scenarios 4 and 2 (Section 10.4), the fact on lost frames (Section 10.6), and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured scenarios (Section 10.6), Technology in Practice (Section 10.7), Limitations (Section 10.8) and the open challenge

The control question A camera reports each product on a moving belt after it has moved on, and a grasp takes only if the gripper is in the right place, nearly still relative to the product and able to hold it, all at one instant. How does a picking arm meet a product where it will be and at the speed it is going, with limited torque and grip?

Learning objectives

- Compute the distance a report-aimed gripper closes behind its product from the belt speed, the latency and the time to the grasp.
- Explain why a gripper that waits at rest fails above the relative speed tolerance, and how arriving at belt velocity removes that failure.
- Bound the carry acceleration by the grip force and the product mass, and explain why the usable limit sits below the bound.
- Predict how a payload left out of the controller's arm model breaks the torque budget and the tracking.
- Configure the cell for a mixed production run and justify each change by the grasp budget it restores.

Prerequisites Sections 1.4.2 (arm dynamics and the payload term), 1.4.4 (quintic profiles and feedforward), 1.4.6 (the age of a report), 1.6.1 and 1.6.4 (running and reading a scenario).

The app https://apmonitor.com/apps/robot-conveyor-picking/

See also Volume 5, Chapter 13 applies each late acoustic fix at the instant it describes, Volume 5, Chapter 7 follows material that has moved on past a slow analyzer, and Volume 5, Chapter 8 closes a robot loop through a delayed link.

10.1 Introduction

Many machines act on an object they measured a moment earlier. A visual servo closes a robot loop through a camera, a hand-eye system that grasps a moving part has to predict where the part will be, and a state estimator that receives a reading late has to incorporate it although it describes an earlier state [1–3]. In each case the measurement says where the object was, the action takes effect where the object will be, and the gap between the two is set by the latency of the sensor, the time the actuator needs to get there and the time it needs to act.

Conveyor picking makes that gap concrete. A camera cannot watch the part of the belt where the arm works, so it looks upstream and reports each product after capture, transfer and recognition. The product rides on while the arm moves to it and while the gripper closes, and the grasp takes only if the gripper is in the right place, nearly still relative to the product and able to hold it, all at one instant. The simulator targets two beliefs. One holds that camera latency is the whole of the timing error. With no latency at all, a gripper sent to the camera's last report on a 0.15 m/s belt still closes 63.0 mm behind the product it was sent for, because that report is 0.42 s old when the grasp completes. The other holds that an accurate position is enough. An arm that waits at rest where a product will be, 0.96 mm from it on average, fails all 60 of its grasps on a 0.30 m/s belt, because the product slides under the gripper at the belt speed.

The class of problem is general. Choosing the point on a predicted path at which to meet a moving object is an optimization of its own [4], prediction and replanning for interception have been developed and tested for objects moving in industrial settings [5], and a Kalman filter can incorporate a measurement that arrives late [3], as the underwater navigation filter of Volume 5, Chapter 13 does with its acoustic fixes. A flying shear that cuts a moving strip, a labeler matched to a product stream and a radar track carried between sweeps have the same shape, and so does a continuous line whose analyzer describes material that has moved on, as in Volume 5, Chapter 7. In each, a measurement is moved to the moment of action by a model of what happened in between, and the actuator must arrive at the object's speed as well as its position.

The behavior matters industrially. Robots pack products in the confectionery industry [6], and on a line that never stops, the number of products an arm catches depends on the order in which it takes them [7] and on how several arms along one belt share the stream [8]. Model-based servos also make

the payload part of the control problem, because the torque a move needs depends on what the arm carries [9].

This chapter treats a conveyor picking simulator in which the camera latency, the belt speed, the product mass, the grip and the arrival pattern can be changed and their consequences read as products placed, missed, failed and dropped. Conveyor tracking, quintic trajectories, computed-torque feedforward and earliest-deadline scheduling are standard methods. The simulator gives a model complete enough to rebuild, in which the controller reads only its own sensors, the products and the camera draw from separate random streams so that two controller settings meet the same products, and the planner checks torque with the controller's own arm model, so that a payload the model omits is invisible to the check. Each grasp condition reduces to a bound that holds before any tuning. An aim that ignores the closing interval misses by at least the belt speed times the closing time, a gripper that waits at rest meets the product at the belt speed, and the grip caps the carry acceleration at the holding force divided by the mass. Eight cases and a graded run expose these bounds one at a time.

Section 10.2 traces how robots came to pick from moving belts. Section 10.3 gives the arm, belt, camera and grasp model, and Section 10.4 turns the grasp into four budgets. Section 10.5 describes the controller layer by layer, Section 10.6 reports the measured scenarios and the production run, Section 10.7 names the devices of real picking cells, and Sections 10.8 and 10.9 state the limitations and the conclusions.

10.2 Picking Products From a Moving Belt

The arm simulated here is a SCARA, a selective compliance assembly robot arm, whose origin the history segment below tells. Both of its joints turn about vertical axes, so the arm sweeps across a plane while gravity loads neither joint. The Delta parallel robot was designed by a research team at the École Polytechnique Fédérale de Lausanne, and a complete model of it, from the kinematics to the inverse dynamics, was written to need few arithmetic and trigonometric operations [10]. Robots for packaging in the confectionery industry, where chocolates are handled, have been reviewed in their own right [6].

> ⧖ HISTORY **An Arm Like a Folding Screen (1978 to 1981)**
>
> The selective compliance assembly robot arm was invented at the University of Yamanashi for assembly work. Its arm gives easily in the horizontal plane and stays stiff in the vertical direction, so a peg slightly out of line can slide into its hole, and its links fold about vertical axes like the panels of a Japanese folding screen. A research group of five companies, formed in April 1978, made the first prototype possible. It was built that October with a servo motor and an optical encoder at each joint, a second prototype followed in 1980, and in 1981 some of the partners began to market their own versions [11].

A moving product first has to be found. A hand-eye system that tracked a moving object with cameras, predicted its motion and closed a gripper on it was reported in 1993 [2], and a tutorial on visual servo control organized the ways a camera can close a robot's loop, with the camera fixed in the workcell or carried on the arm [1]. Interception then became a planning problem in its own right. The meeting point on a predicted path can be selected optimally [4], and prediction, planning and execution were combined into interception strategies tested experimentally for objects moving in industrial settings [5]. A measurement that reaches an estimator late can be incorporated into a discrete-time Kalman filter although it describes an earlier state [3], which is the estimation side of the same timing problem.

A stream of products added the question of order. When items on a belt arrive faster than one robot can take them, the order of the picks changes how many are caught, and online scheduling algorithms were proposed to improve pick-and-place performance on a moving conveyor [7]. Several robots along one conveyor must then coordinate which items each one picks [8].

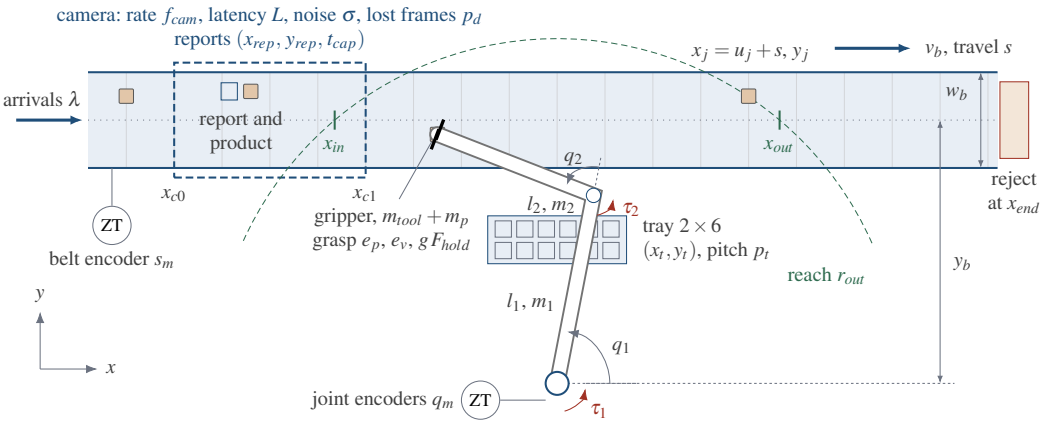


Figure 10.1 The picking cell seen from above, drawn to scale. The shoulder is the origin, the upper arm of length l_1 turns through q_1 from the belt direction and the forearm of length l_2 through q_2 from the upper arm, driven by the joint torques τ_1 and τ_2. Products arrive at rate λ and ride the belt at $x_j = u_j + s$. The camera reports those inside $[x_{c0}, x_{c1}]$ after a latency L, so each report (hollow square) lies behind its product, and the planner meets the products inside the reach r_{out}, between x_{in} and x_{out} on the centerline. A product carried beyond $l_1 + l_2$ is missed, and one that reaches x_{end} goes to the reject bin. The part-present switch sits in the gripper.

> **DID YOU KNOW? Fewer Than Forty Thousand Robots**
>
> At the end of 1981 there were fewer than 40,000 industrial robots in the world, and by 1992 there were 560,000. The first industrial use of the new horizontal arm put eight of them to work assembling audio circuit boards [11].

The joint servos came from model-based control of fast arms. Resolved acceleration control computes the joint torques that produce a commanded acceleration from the manipulator dynamics and corrects the remaining error by feedback [12]. Real-time experiments on a direct-drive arm sampled every 2 ms later compared computed-torque control, which uses the dynamics inside the feedback loop to linearize and decouple the joints, with feedforward compensation, which adds the model torque in the feedforward path, and showed that the off-diagonal terms of the manipulator inertia matrix belong in the torque computation [13]. Both forms depend on the load, whose mass and center of mass can be estimated from general motion of the arm, while its moments of inertia are harder to obtain [9]. A vacuum cup holds its load by astrictive force, one of the forms of prehension that gripper design classifies [14], and that force has to exceed the inertial load of the handling motion. Minimum-time motion along a specified path with known torque limits was solved in 1985 [15], and trajectories generated online from whatever state the arm is in let a controller react at once to a new measurement [16]. The application of this chapter assembles these pieces at teaching scale.

10.3 The Picking Cell

The cell is drawn from above in Figure 10.1. The shoulder of the arm is the origin, the belt runs along $+x$ with its centerline at $y_b = 0.55\,\text{m}$, the camera watches the belt between $x_{c0} = -0.80\,\text{m}$ and $x_{c1} = -0.40\,\text{m}$, upstream of the arm, and a tray of twelve cavities lies between the shoulder and the belt. The six integrated states are the joint angles q_1 and q_2, their rates ω_1 and ω_2, the belt travel s and the belt speed v_b. Everything discrete, which covers the products, the camera frames and their delivery, the gripper, the slip test and the tray, is updated after each integration step and never inside the derivative function.

Table 10.1 Classification of the Signals in the Picking Cell

Symbol	Role	Default	Range
q_1, q_2	CV, joint angles	ready pose	$\lvert q_2\rvert \le 150^\circ$
ω_1, ω_2	states, joint rates	0	—
s, v_b	states, belt travel and speed	0, v_{set}	—
τ_1, τ_2	MV, joint torques	0	± 40, ± 17 N m
—	MV, gripper command	open	open, close
v_{set}	DV, belt speed setting	0.15 m/s	0.05–0.50
λ	DV, arrival rate	24 min^{-1}	0–120
L	DV, camera latency	0.15 s	0–0.5
p_d	DV, lost frames	5 %	0–90
σ	DV, camera noise (1σ)	2 mm	0–10
m_p	DV, product mass	0.10 kg	0.02–2.0
g	DV, grip strength	100 %	10–100
—	DV, extra product, burst of 12	—	events
—	DV, 3 s blackout, 0.6 rad/s bump	—	events
q_m	measured, joint encoders	17 bit	—
s_m	measured, belt encoder	0.1 mm	—
x_{rep}, y_{rep}, t_{cap}	measured, camera report	15 Hz	—
—	measured, part present	off	on, off
—	outputs, product counts	0	—

The signals are classified in Table 10.1. The two joint angles are the controlled variables of the servo loops, whose references come from the planner, while the outcome of the task is scored as products placed, missed, failed and dropped.

10.3.1 Arm Dynamics

Both joints turn about vertical axes, so gravity does no work on the links and the rigid-body dynamics take the form of Equation (10.1), the arm balance of Equation (1.2) in Chapter 1.

$$M(q)\,\ddot{q} + C(q,\dot{q})\,\dot{q} + F(\dot{q}) = \tau, \qquad |\tau_i| \le \tau_{i,\max} \tag{10.1}$$

where $q = [q_1, q_2]^T$ holds the shoulder angle, measured from the belt direction, and the elbow angle, measured from the upper arm, both counterclockwise seen from above, τ holds the two joint torques and F is the joint friction. The upper arm and the forearm are uniform rods of lengths l_1 and l_2 and masses m_1 and m_2, and a point mass $m_{tip} = m_{tool} + m_p$ at the tip stands for the gripper and its vertical quill plus any product of mass m_p being held. With $l_{ci} = l_i/2$ and $I_i = m_i l_i^2/12$, the inertia coefficients are given by Equation (10.2).

$$\begin{aligned} a_1 &= m_1 l_{c1}^2 + I_1 + (m_2 + m_{tip})\, l_1^2 + J_{1m} \\ a_2 &= m_2 l_{c2}^2 + I_2 + m_{tip}\, l_2^2 \\ a_3 &= (m_2 l_{c2} + m_{tip}\, l_2)\, l_1 \end{aligned} \tag{10.2}$$

where J_{1m} and J_{2m} are the rotor inertias of the two motors reflected through their gears. The inertia matrix, the velocity terms and the friction follow as Equation (10.3).

$$\begin{aligned} M &= \begin{bmatrix} a_1 + a_2 + 2a_3\cos q_2 & a_2 + a_3\cos q_2 \\ a_2 + a_3\cos q_2 & a_2 + J_{2m} \end{bmatrix} \\ C\dot{q} &= a_3 \sin q_2 \begin{bmatrix} -(2\dot{q}_1\dot{q}_2 + \dot{q}_2^2) \\ \dot{q}_1^2 \end{bmatrix} \\ F_i &= b_i\,\dot{q}_i + f_{ci}\tanh(\dot{q}_i/0.05) \end{aligned} \tag{10.3}$$

where b_i and f_{ci} are viscous and Coulomb friction coefficients and the tanh smooths the Coulomb term over 0.05 rad/s. Every row carries units of N m. The elbow motor's gyroscopic coupling is left out. The tip position and velocity are given by Equation (10.4).

$$p_{tip} = \begin{bmatrix} l_1 \cos q_1 + l_2 \cos(q_1+q_2) \\ l_1 \sin q_1 + l_2 \sin(q_1+q_2) \end{bmatrix}, \qquad v_{tip} = J(q)\,\dot{q} \tag{10.4}$$

where J is the tip Jacobian, and $a_{tip} = J\ddot{q} + \dot{J}\dot{q}$. A held product adds exactly $m_p J^T J$ to M, which one of the application's verification checks confirms to within 10^{-12}. At the ready pose, $q_2 = 62.6°$, the empty arm has $M_{11} = 0.9996\,\mathrm{kg\,m^2}$ and $M_{22} = 0.1579\,\mathrm{kg\,m^2}$, and a 1.5 kg case raises them to 1.6167 and $0.3417\,\mathrm{kg\,m^2}$. The elbow stops dead at $|q_2| = q_{2,\max} = 150°$, where its rate into the stop is set to zero and an end-stop event is marked. The planner works inside $r_{out} = l_1 + l_2 - 0.03\,\mathrm{m} = 0.72\,\mathrm{m}$, short of the straight arm so that the elbow never has to straighten, and at the elbow limit the inner radius of the reach is 0.20 m.

10.3.2 Belt, Products and Camera

The belt drive follows its speed setting through a first-order response, and each product rides at a fixed belt coordinate, Equation (10.5).

$$\dot{s} = v_b, \qquad \tau_b \dot{v}_b = v_{set} - v_b, \qquad x_j = u_j + s \tag{10.5}$$

where $\tau_b = 0.5\,\mathrm{s}$ and u_j is the belt coordinate of product j. Products enter at $x = -1.45\,\mathrm{m}$ as a Poisson stream of rate λ per minute, never closer than 0.07 m along the belt, with a lateral position y_j uniform within $y_b \pm 0.07\,\mathrm{m}$ and the mass m_p in force when they enter. At reset the belt is preloaded upstream of the camera, the first product at $x = -0.95\,\mathrm{m}$. A product is missed when it passes $x > 0$ beyond the full reach $l_1 + l_2 = 0.75\,\mathrm{m}$, and it leaves the belt at $x_{end} = 0.95\,\mathrm{m}$.

Every $1/f_{cam}$ seconds the camera takes a frame at time t_k. The whole frame is lost with probability p_d, or while a blackout lasts, and otherwise every product inside $x_{c0} \le x_j(t_k) \le x_{c1}$ is reported as shown in Equation (10.6).

$$(x_{rep}, y_{rep}, t_{cap}) = \big(x_j(t_k) + n_x,\; y_j + n_y,\; t_k\big) \tag{10.6}$$

where n_x and n_y are independent normal deviates of standard deviation σ. The report is posted at the first integration step that ends at or after $t_k + L$, the controller reads it at its next scan, and it carries no identity and no velocity. The belt encoder counts s_m in steps of $\Delta s = 0.1\,\mathrm{mm}$, and the joint encoders count q_m in steps of $2\pi/131072 = 4.79 \times 10^{-5}\,\mathrm{rad}$, about 0.036 mm at the full reach. The arrivals and the camera draw from two separate random streams, both seeded from the scenario seed at reset, so a reset replays the same products, noise and lost frames, and a change of controller setting meets exactly the same products. Reset and reload of a scenario behave as Section 1.6.1 describes.

10.3.3 Grasp, Carry and Placement

The gripper closes in $t_{close} = 0.10\,\mathrm{s}$, a time that also stands for the quill's dip to the belt, and opens in $t_{rel} = 0.08\,\mathrm{s}$. When it finishes closing, at time t_g, the grasp takes only if the three conditions of Equation (10.7) hold for the nearest product on the belt.

$$\begin{aligned} \|p_{tip} - p_j\| &\le e_p \\ \|v_{tip} - [v_b, 0]^T\| &\le e_v \\ m_p \|a_{tip}\| &\le g\,F_{hold} \end{aligned} \tag{10.7}$$

where $e_p = 8\,\mathrm{mm}$, $e_v = 0.15\,\mathrm{m/s}$, $F_{hold} = 10\,\mathrm{N}$ is the holding force at full grip and g is the grip strength. A failed grasp is charged to the first condition that fails, in the order position, speed and grip.

A successful catch is inelastic and conserves generalized momentum, as shown in Equation (10.8).

$$M(q; m_p)\,\dot{q}^{+} = M(q; 0)\,\dot{q}^{-} + m_p J^{T}(q) \begin{bmatrix} v_b \\ 0 \end{bmatrix} \tag{10.8}$$

where $M(q; m)$ is the inertia matrix with mass m held. A gripper already moving at belt velocity therefore catches the product without any change in its joint rates, which a verification check confirms to within 10^{-12}. While a product is held it slips out, and counts as a drop, whenever $m_p \|a_{tip}\| > gF_{hold}$. On release the product is placed if the tip is within 12 mm of the nearest tray cavity and that cavity is empty, it is put back on the belt if the tip is over the belt, and otherwise it is dropped. The tray holds 2×6 cavities at a 45 mm pitch, and the cell program replaces it at once by an empty one after every twelfth release over it, placed or not, so a missed cavity cannot put the program's count out of step with the tray.

10.3.4 Parameters, Numerics and Scoring

The parameters are listed in Table 10.2. The Show Model dialog offers a slider for every parameter marked adjustable, the control panel holds the tuning and the options, and the parameters marked hidden set the geometry of the cell. The six states advance by fourth-order Runge-Kutta integration on a fixed 1 ms step, the controller runs every 2 ms, and the historian logs at 50 Hz, in the step order of Section 1.6.1, with the end stop and the discrete events handled after the integration. This chapter writes Δt for the integration step and h for the control period, the reverse of the symbols in that section. Under a fixed torque history, halving the step to 0.5 ms changes the joint angles after 1 s by 5.2×10^{-8} rad, and the next halving shrinks that difference by a factor of 14.1, close to the 16 expected of a fourth-order method. Case 6 run for 60 s at 0.5 ms reproduces the placed, missed, failed and dropped counts and all 11 torque-limit hits, with the energy within 2.3 J of 1397.4 J and the tracking error at 6.190 against 6.202 mm rms.

The scorecard counts products placed, missed, failed and dropped, together with the picks per minute, the torque-limit hits, the shaft work $\sum_i \int |\tau_i \omega_i|\, dt$ with braking counted as spent, and the rms distance of the tip from the tip position of its reference in automatic mode. A torque-limit hit is counted once per episode, and a new episode begins only after 50 ms clear of the limit. These scores use the true state, and the controller never reads them.

10.3.5 Assumptions

A planar arm. Both joints turn about vertical axes and the quill's dip to the belt is folded into t_{close} and t_{rel}. A real cell also plans that vertical axis, and its settling adds to the cycle.

Rigid links and a point payload. Real arms flex at tens of m/s^2, and a carton has rotational inertia of its own, which is the part of a load that is hardest to estimate [9].

Ideal drives up to a limit. Each drive delivers the commanded torque up to $\tau_{i,\max}$, and the speed ratings enter only as planner limits. A real motor loses torque with speed, and its current loop and gearbox add lag and backlash.

One clock and exact time stamps. The camera and the encoder share a clock and every frame is stamped exactly. On a real line, clock offset and trigger jitter between the two add a position error proportional to the belt speed.

An ideal recognizer. The camera reports positions only, with Gaussian noise, and finds every product in its view unless the whole frame is lost. Misrecognition, overlapping products and occlusion by the arm are absent.

A threshold grasp. Equation (10.7) either holds or not, and a slip loses the product at once. Real vacuum cups creep and skid first, which turns the three thresholds into a probability.

An exact controller model. The controller's arm model matches the plant apart from the payload switch. An identified model would leave some tracking error that only feedback removes.

Table 10.2 Model, Controller and Numerical Constants. Class A Is Adjustable in the Interface, H Is a Hidden Parameter and C Is Fixed in the Code

Symbol	Meaning	Value	Class
Arm and drives			
l_1, l_2	link lengths	0.40, 0.35 m	A
m_1, m_2	link masses	4.0, 2.0 kg	A
m_{tool}	gripper and quill at the tip	0.5 kg	A
$q_{2,\max}$	elbow travel	150°	A
$\tau_{1,\max}$, $\tau_{2,\max}$	torque limits	40, 17 N m	A
$\omega_{1,\max}$, $\omega_{2,\max}$	speed ratings (planner)	7.0, 12.0 rad/s	A
J_{1m}, J_{2m}	reflected rotor inertias	0.05, 0.015 kg m^2	A
b_1, b_2	viscous friction	0.6, 0.3 N m s/rad	A
f_{c1}, f_{c2}	Coulomb friction	0.6, 0.3 N m	A
—	joint encoder count	$2\pi/131072$ rad	H
Line and camera			
y_b, w_b	belt centerline, width	0.55, 0.20 m	H
τ_b	belt drive time constant	0.5 s	A
Δs	belt encoder count	0.1 mm	A
—	entry point, reject bin	−1.45, 0.95 m	H
—	minimum spacing on the belt	0.07 m	A
f_{cam}	camera frame rate	15 Hz	A
x_{c0}, x_{c1}	camera field of view	−0.80, –0.40 m	H
Gripper and tray			
t_{close}, t_{rel}	closing and release times	0.10, 0.08 s	A
e_p, e_v	grasp tolerances	8 mm, 0.15 m/s	A
F_{hold}	holding force at full grip	10 N	A
x_t, y_t, p_t	tray center and pitch	0, 0.30, 0.045 m	H
—	placement tolerance	12 mm	C
Controller			
K_{p1}, K_{p2}	position gains	400, 100 N m/rad	A
K_{d1}, K_{d2}	rate gains	25, 6 N m s/rad	A
$\alpha_{1,\max}$, $\alpha_{2,\max}$	acceleration limits	40, 80 rad/s^2	A
a_{carry}	carry acceleration limit	30 m/s^2	A
v_a	assumed belt speed	0.15 m/s	A
—	torque share for planning	90 %	C
τ_f	joint rate filter	1 ms	C
—	replan and schedule period	20 ms	C
Numerics			
Δt, h	integration step, control period	1, 2 ms	C
—	historian rate	50 Hz	C

10.4 Four Budgets for One Grasp

Equation (10.7) is what makes conveyor picking different from trajectory tracking. The three conditions must hold at one instant, each is met by a different part of the controller, and each yields a bound that decides feasibility before any gain is tuned. A fourth budget, the drive torque, is where the payload enters.

10.4.1 Position and the Age of a Report

A report describes the product at its capture time. An estimator that aims at the report as if it were current sends the gripper to where the product was, and when the grasp completes the product has moved on by the distance of Equation (10.9), the error vL of Equation (1.11) with the age carried on to the grasp.

$$d = v_b\,(t_g - t_{cap}) = v_b\,(L + T_d) \tag{10.9}$$

where T_d is the time from the delivery of the report to the completed grasp. T_d contains at least the closing time, so even an instantaneous camera leaves $d \geq v_b t_{close}$. With $e_p = 8\,\text{mm}$ and $t_{close} = 0.10\,\text{s}$, an aim that does not predict the closing interval fails above $e_p / t_{close} = 0.08\,\text{m/s}$ whatever the latency. In case 4 on a 0.15 m/s belt, $T_d = 0.42\,\text{s}$ on average, and over latencies from 0 to 0.5 s the distance from the gripper to the product it was sent for follows $d = 64.4 + 148.2L$ in mm with no residual above 3.2 mm, a slope 1.2 % below the belt speed, as shown in Figure 10.2. The two points that sit 3 mm above the line, at 0.15 and 0.20 s, are runs in which one grasp of 42 lands on a following product, which shifts the later attempts and raises their mean T_d to 0.45 s. The remedy is the time stamp itself. Moved on by the belt travel since t_{cap} and predicted to t_g, the estimate places 24 of 25 products at every latency in that range, with a mean position error at the grasp between 0.81 and 1.06 mm.

> **TRY IT IN THE APP** **4 · Camera Delay**
>
> Load the case and press Run. The controller aims at the camera's last report, which arrives 300 ms late. Watch where the gripper closes and where its product is by then: the gripper closes on empty belt behind it, and none of 42 grasps in the first minute take. Press Reset, set Camera latency to 0 and run again. The gripper still closes behind its product.

10.4.2 Relative Speed at the Grasp

A gripper that arrives early and waits at rest meets the product at a relative speed equal to the belt speed, so a rest intercept is feasible only while $v_b < e_v$. Case 2 at 30 arrivals a minute shows a sharp edge, with the belt speed set before the run so that the belt starts at that speed. Waiting at rest places every product that reaches the arm, with no failed grasp, up to 0.14 m/s, fails 12 grasps at 0.15 m/s, where the relative speed at the grasp spans 0.1476 to 0.1509 m/s about the tolerance, and places none from 0.16 m/s upward. Across the sweep the mean relative speed is 0.98 to 1.002 times the belt speed. Arriving at belt velocity instead places 23 to 30 products at every speed from 0.05 to 0.50 m/s, with a mean relative speed of 0.0024 to 0.0032 m/s at the grasp.

> **TRY IT IN THE APP** **2 · Higher Belt Speed**
>
> Run the case as loaded. The arm arrives early and waits at rest, and every grasp fails on speed as the product slides under the gripper at the belt's 0.30 m/s. Press Reset, switch Meet the chocolate to Move with it and run again. The gripper now follows each product while it closes, and 30 are placed in the first minute with no failure.

10.4.3 Holding Force During the Carry

Once a product is held, the only force that can pull it off the cup is its own acceleration, so the carry must keep $\|a_{tip}\| \leq g F_{hold} / m_p$. For 0.10 kg at 15 % grip the bound is $15\,\text{m/s}^2$, and for the 0.6 kg boxes of the challenge at 60 % it is $10\,\text{m/s}^2$. The carry limit a_{carry} bounds the planned acceleration, while the real one also carries the tracking error, so the usable limit sits below the bound. In case 7 the largest force any carry needs is 1.445 N at a $14\,\text{m/s}^2$ limit and 1.230 N at $12\,\text{m/s}^2$, 3.2 and 2.5 % above the planned value, and a $15\,\text{m/s}^2$ limit drops 20 of 24 products.

10.4.4 Torque and the Payload

The payload also enters the torque budget, because it adds $m_p J^T J$ to M. A 1.5 kg case raises M_{11} at the ready pose by 62 % and M_{22} by 116 %. The planner checks every trajectory against 90 % of each drive with the controller's own arm model, so a model that leaves the case out approves carries that the drives cannot deliver.

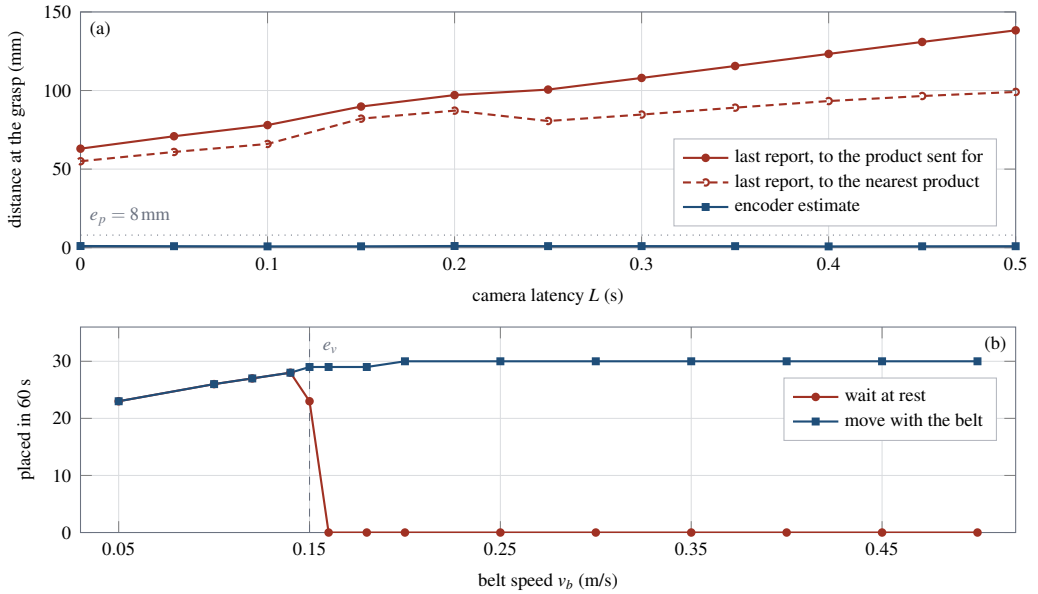


Figure 10.2 Grasp outcome against camera latency and belt speed, measured on the application. (a) Case 4 at 0.15 m/s. Aimed at the last report, the gripper closes behind the product it was sent for (solid, filled circles) by $64.4 + 148.2L$ in mm. The application's own grasp error, the distance to the nearest product (dashed, open circles), grows more slowly because a following product is sometimes nearer. Moved on by the encoder (squares), the error stays between 0.81 and 1.06 mm and 24 of 25 products are placed at every latency. (b) Case 2 at 30 arrivals a minute, the belt speed set before each run. Waiting at rest (circles) places nothing once the belt runs faster than the relative speed tolerance e_v (dashed line), while arriving at belt velocity (squares) places 23 to 30 at every speed.

10.5 Implemented Controller

The controller reads only its own sensors, which are the joint encoders, the belt encoder, the delivered camera reports and the part-present switch in the gripper. A verification check wraps the simulated world so that any other read by the control law fails, and it passes over four scenarios. The controller is built in four layers, as shown in Figure 10.3. None of them is a PID loop. The application does not use the velocity-form PID of Section 1.6.2, and with no integral action anywhere there is nothing for anti-windup to act on.

10.5.1 Sensors

The joint rates are a filtered backward difference of the encoder angles, $\hat{\omega}_k = \mathrm{e}^{-h/\tau_f}\,\hat{\omega}_{k-1} + (1 - \mathrm{e}^{-h/\tau_f})(q_{m,k} - q_{m,k-1})/h$, with $h = 2$ ms and $\tau_f = 1$ ms. The belt speed $\hat{v}_b$ is the slope of the belt encoder count over the last 50 ms, passed through a first-order filter of 0.1 s. Every belt encoder count is also pushed, with its time, into a history of 1024 samples at the control period, which covers 2.05 s.

10.5.2 Estimator and Tracks

Three estimates of a product's position are offered, Equation (10.10).

$$\hat{x}(t) = \begin{cases} x_{rep} & \text{last report} \\ x_{rep} + v_a\,(t - t_{cap}) & \text{assumed speed} \\ x_{rep} + s_m(t) - s_m(t_{cap}) & \text{encoder} \end{cases} \tag{10.10}$$

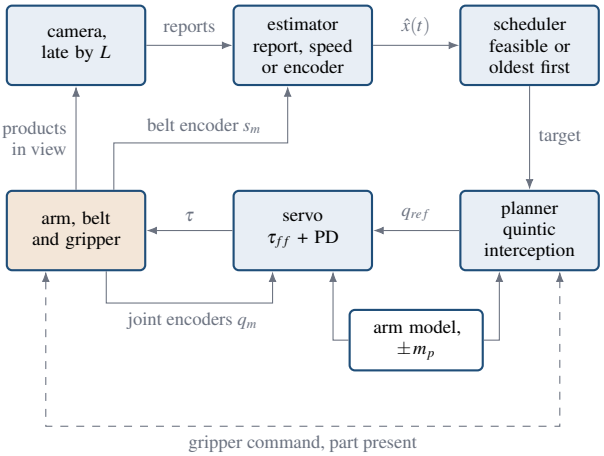


Figure 10.3 Layers of the picking controller. The estimator turns late camera reports and the belt encoder count into positions at any time, the scheduler chooses the next product, the planner builds a minimum-time quintic that meets it, and the servo follows that reference with computed-torque feedforward and a PD correction. One arm model, with or without the payload m_p, serves both the feedforward and the planner's torque check. The task sequence in the planner commands the gripper and reads the part-present switch.

where $s_m(t_{cap})$ is read from the encoder history by linear interpolation and v_a is a typed-in belt speed. A stored history serves the same purpose in Chapter 9, which forms each innovation of its pose filter against the pose stored for the moment of the measurement. The last two hold each product as a track in belt coordinates, $u = x_{rep} - s_m(t_{cap})$ or $u = x_{rep} - v_a t_{cap}$, where a product does not move, and average successive reports as $u \leftarrow u + k_n(u_{rep} - u)$ with $k_n = \max(1/n, 0.25)$ after the nth report. The average is exact for the first four reports and exponential after that, and its stationary variance under independent noise is $0.25/(2-0.25) = 1/7$ of the camera's. The 2 mm camera noise therefore becomes 0.76 mm per axis and a mean distance of 0.95 mm, against the 0.99 mm measured in case 1. Predicting to a later time t_a, the encoder track adds $\hat{v}_b(t_a - t)$. A report joins the track whose position at t_{cap} lies within 30 mm along the belt and 25 mm across it, or it starts a new track. The last-report estimator matches a report to a track when the report lies between 20 mm upstream and 80 mm downstream of that track's previous report. A change of estimator or of v_a restarts every track and abandons an approach in progress.

10.5.3 Scheduler

While the arm is free, the scheduler runs every 20 ms over the tracks not yet picked and tried fewer than twice. Feasible first plans an interception for each candidate, drops any it cannot meet inside the reach before the grasp finishes, and takes the one with the earliest deadline, the time left before the product passes $x_{out} - \hat{v}(t_{close} + 0.03\,\mathrm{s})$, where $\hat{v}$ is the belt speed the estimator uses ($\hat{v}_b$ or v_a). Oldest first takes the lowest track number whether it can be caught or not. With the last-report estimator the scheduler takes the most downstream report inside the reach.

10.5.4 Planner

Every move is one quintic polynomial per joint, from the present reference state of angle, rate and acceleration to a goal state, over one duration T shared by both joints. The reference is therefore continuous in acceleration across replans, and the feedforward torque never steps [16, 17]. A rest-to-rest move of Δq peaks at the rate $1.875\,\Delta q/T$ and the acceleration $(10/\sqrt{3})\,\Delta q/T^2 = 5.774\,\Delta q/T^2$ of Section 1.4.4, which gives the starting duration of Equation (10.11), the quintic duration of

Equation (1.5) for each joint with two more bounds.

$$T_0 = \max_i \left\{ 0.12,\ \sqrt{\frac{5.774\,|\Delta q_i|}{\alpha_{i,\max}}},\ \frac{1.875\,|\Delta q_i|}{\omega_{i,\max}},\ \frac{|\Delta \dot{q}_i|}{\alpha_{i,\max}} \right\} \tag{10.11}$$

where T_0 is in seconds. T is then rescaled, at most ten times, by the largest of the ratios

$$\sqrt{\frac{\max|\ddot{q}_i|}{\alpha_{i,\max}}},\quad \frac{\max|\dot{q}_i|}{\omega_{i,\max}},\quad \sqrt{\frac{\max|\hat{\tau}_i|}{0.9\,\tau_{i,\max}}},\quad \sqrt{\frac{\max\|a_{tip}\|}{a_{carry}}}$$

the last only while a product is held, each peak taken from 21 samples of the polynomial, until the tightest ratio lies between 0.97 and 1, and never beyond $8T_0$. The torque $\hat{\tau}$ comes from the controller's arm model, which is how a payload left out of that model escapes the check.

The goal is where the product will be when the grasp takes, which depends on T, which depends on the goal. Up to eight fixed-point iterations settle Equation (10.12), because the product moves slowly compared with the arm [4]. A table tennis ball moves fast compared with the arm, and Chapter 11 predicts its crossing point and swings within the time left.

$$T = T_{\min}\big(q_{ref} \to q(\hat{p}(t+T+\Delta))\big) \tag{10.12}$$

where $\hat{p}$ is the predicted product position and Δ is the delay from arrival to grasp. Waiting at rest, the arm arrives 0.03 s early and $\Delta = 0.03 + t_{close}$, so the product reaches the stationary gripper as the grasp takes. Moving with the belt, $\Delta = 0$, the goal carries the joint rates $\dot{q} = J^{-1}[\hat{v}, 0]^T$, with $\hat{v}$ the estimator's belt speed as in the scheduler, and the accelerations that keep the tip on a straight line, and while the gripper closes the reference follows the product along the belt at $\hat{v}$, stopping at the edge of the reach. If the product has not yet entered the reach, the arrival is delayed until it does. An arrival more than 0.25 s beyond the minimum time is made in two moves, first at minimum time to a staging point under the meeting point, or 0.12 s of belt travel behind it when moving with the belt, then a final approach launched when the time left falls to its own minimum plus 40 ms. In case 1, 13 of 19 approaches were made this way. Every 20 ms the goal is refined toward the latest estimate at the same arrival time, except in the last 60 ms, and a correction that would break a limit is absorbed by arriving later. A final approach more than twice over a limit is abandoned and the arm holds, which happens only at speed ratings too low to follow the belt.

10.5.5 Servo

Each joint follows its reference with the law of Equation (10.13), executed every $h = 2$ ms.

$$\tau_i = \tau_{ff,i} + K_{p,i}\,(q_{ref,i} - q_{m,i}) + K_{d,i}\,\big(\dot{q}_{ref,i} - \delta\,\ddot{q}_{ref,i} - \hat{\omega}_i\big) \tag{10.13}$$

where $\tau_{ff} = \hat{M}\ddot{q}^* + \hat{C}\dot{q}^* + \hat{F}$ is the inverse dynamics of the controller's arm model, the reference feedforward of Section 1.4.3 written for the arm, friction included, along the reference q^* half a control period ahead, because the torque computed now is held for the whole period, and along the reference itself while the gripper closes. The shift $\delta = h/2 + \tau_f = 2$ ms moves the reference rate back to the instant that the filtered backward difference $\hat{\omega}_i$ describes. The arm model either knows only the arm or also counts the product, at its recipe mass, while the part-present switch is on. The command is clipped at $\pm\tau_{i,\max}$, the only output limit, with no rate limit on the torque. The defaults are $K_{p1} = 400$ N m/rad, $K_{d1} = 25$ N m s/rad, $K_{p2} = 100$ N m/rad and $K_{d2} = 6$ N m s/rad. Against the diagonal inertias at the ready pose they give a shoulder loop of 20.0 rad/s at a damping ratio of 0.63 and an elbow loop of 25.2 rad/s at 0.75, and a 1.5 kg case lowers these to 15.7 rad/s at 0.49 and 17.1 rad/s at 0.51. In the convention of Section 1.3 each loop is reverse acting ($a = +1$), because an angle short of its reference receives a positive torque and a positive torque accelerates the joint forward, as the application's sign checks confirm. The application gives the loops no action label. There is no integral term, because with no gravity on the links there is no steady torque for one to find, and the feedforward already supplies the inertia and friction torques the motion needs.

Table 10.3 Scenario Sequence and the Changes Each Makes to the Base Line

#	Change	Teaching point
1	0.10 m/s, 20 min^{-1}	the cell working, report against product
2	0.30 m/s, 30 min^{-1}	speed at the grasp
3	0.20 to 0.30 m/s at 20 s, assumed 0.20 m/s, move with it	a typed-in speed against the encoder
4	$L = 0.30$ s, last report, move with it	the age of a report
5	0.25 m/s, 30 min^{-1}, 50 % lost, 3 s blackout at 27.9 s, move with it	one look is enough, none is not
6	1.5 kg, $F_{hold} = 60$ N, move with it	the payload in the model
7	$g = 15$ %, move with it	carry acceleration against grip
8	0.25 m/s, 20 min^{-1}, 12 more at 10 s, oldest first, move with it	scheduling and the ceiling
9	100 boxes of 0.6 kg at 36 min^{-1}, 0.22 to 0.30 m/s at 60 s, $L = 0.2$ s, 10 % lost, $g = 60$ %, assumed 0.22 m/s	graded production run

10.5.6 Manual Mode and Transfer

In manual mode the operator sets τ_1 within ±40 N m and τ_2 within ±17 N m, and the arrow keys add ±3 N m and ±1 N m while held. The scheduler, planner and gripper stay idle, the estimator keeps tracking, and the reference follows the measured arm. A return to automatic mode plans a stop from the reference of the last manual scan, so the initial servo error is only the arm's travel over one control period, 2.7 mrad for a shoulder turning at 1.35 rad/s. In automatic mode the manual torques follow the clipped command after every scan, so a switch to manual starts from the last automatic torque. Manual mode exposes the inertial coupling of Equation (10.3). At the ready pose, 1 N m at the shoulder accelerates the shoulder at 1.57 rad/s^2 and the elbow at −2.39 rad/s^2, while 1 N m at the elbow accelerates the elbow at 9.95 rad/s^2 and the shoulder at −2.39 rad/s^2.

10.5.7 Plant, Controller and Analysis Models

Equations (10.1) to (10.8), with the true payload, friction and limits, make up the simulation model. The controller's model is the same arm dynamics, exact apart from the payload switch and used by both the feedforward and the torque check, together with one of the three position estimates of Equation (10.10). The analysis model is the set of budgets of Section 10.4 and the quintic peak coefficients, which this chapter uses to read the measurements and which appear nowhere in the source. No model is identified from data.

10.6 Scenarios and Measured Results

Table 10.3 lists the eight cases and the challenge. Every case starts from the same line, a 0.15 m/s belt, 24 arrivals a minute, 150 ms of latency, 5 % of frames lost, 2 mm of noise, 0.10 kg products and full grip, with the encoder estimator, waiting at rest, the arm-only model, feasible-first scheduling and a carry limit of 30 m/s^2. The values in Table 10.4 were measured by running the application's own engine at its default seed over the first 60 s unless another span is named, with each change made before the run starts. Position errors are means over the grasps concerned, the failed ones where a case fails, and estimate errors are means over the logged samples in which a target is assigned. Over six other seeds the counts move by a few products, and every comparison below keeps its direction except where the text says otherwise. Figure 10.4 draws the products placed as bars.

10.6.1 The Cell Working, and the Age of a Report

On the slow belt of case 1 the cell places 18 of the 24 products that arrive in the first minute, and the other six are still on their way. No product is missed and no grasp fails. While the arm goes for a product, the camera's last report is 20.2 mm behind it on average and the encoder estimate is 0.99 mm from it. Switched to the last report, the gripper closes 46.8 mm from the nearest product on average

Table 10.4 Measured Results

Case and condition	Result
1, encoder, at rest	18 of 24 placed, none missed or failed; estimate 0.99 mm, last report 20.2 mm
1, last report	0 placed, 32 failed, 17 missed; 46.8 mm from the nearest product
2, at rest	0 placed, 60 failed on speed at 0.3005 m/s; position 0.96 mm
2, move with it	30 placed, none failed; 0.0027 m/s; 0.61 mm rms
2, move with it, feedforward off	30 placed, 15 failed, 1 hit; 16.1 mm rms
3, assumed speed, first 20 s	9 of 11 placed, none failed
3, assumed speed, 60 s	10 placed, 40 failed after the change, 14 missed; 99.7 and 120.4 mm
3, encoder	24 of 25 placed, none failed or missed
4, last report, $L = 0.30$ s	0 placed, 42 failed, 23 missed; 84.7 and 108.0 mm
4, last report, $L = 0$	0 placed, 42 failed; 55.0 and 63.0 mm
4, encoder, $L = 0.30$ s	24 of 25 placed, none failed; 0.99 mm
5, 50 % lost	27 placed, 3 missed, all 3 never seen; 1.28 mm
5, none lost	27 placed, 3 missed; 1.07 mm
6, arm only	24 placed, 11 hits; 47.1 and 24.3 N m; 6.20 mm rms, 43.7 mm worst
6, arm + payload	24 placed, 0 hits; 0.26 mm rms, 0.75 mm worst
6, arm only, $\alpha_{\max}$ 25 and 50	24 placed, 0 hits; 4.93 mm rms
7, $a_{carry} = 30$	0 placed, 24 of 25 dropped
7, $a_{carry} = 15$	4 placed, 20 dropped
7, $a_{carry} = 14$	24 placed, 0 dropped; peak 1.445 N
7, $a_{carry} = 12$	24 placed, 0 dropped; peak 1.230 N
8, oldest first	22 placed, 8 missed, 16 failed; 145.6 mm
8, feasible first	25 placed, 5 missed, none failed
2, 120 min^{-1}, move with it, 90 s	105 placed, 70.0 min^{-1}; 46 of 161 missed
9, as shipped	0 placed, 224 failed, 100 missed, 1 hit
9, move with it	0 placed, 35 dropped, 158 failed
9, and $a_{carry} = 8$	35 placed, 34 of them before 60 s; 158 failed
9, and encoder	98 placed, 2 missed, 0 failed, 0 dropped, 0 hits; 1.97 mm rms
9, and arm + payload	98 placed, 2 missed, 0, 0, 0; 0.25 mm rms; ends at 180.2 s

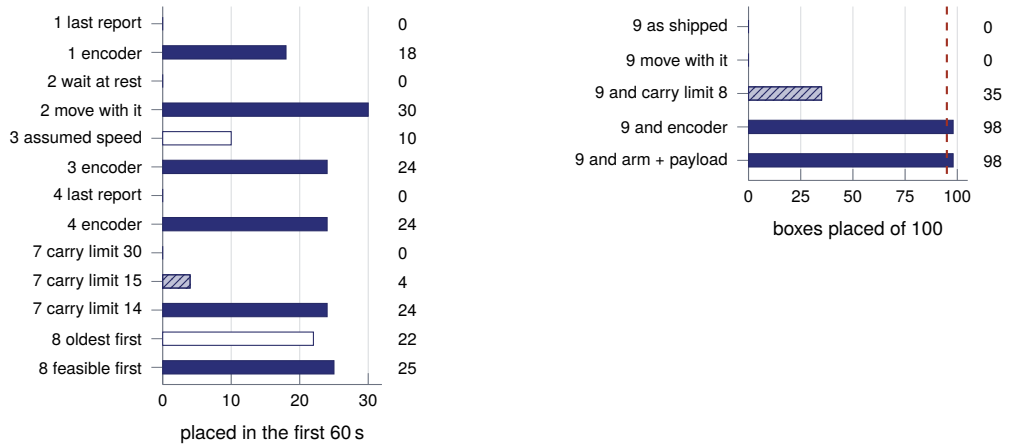


Figure 10.4 The products placed in Table 10.4 drawn as bars, labeled by case and setting. Left, the first minute of cases 1 to 4, 7 and 8: white bars are the setting that fails and solid bars the one that works, and the hatched bar is the carry limit equal to the bound the grip suggests. Right, the production run as each change is added to the one before, with the 95 boxes it asks for as a dashed line and the hatched bar placing only part of the run. Notice how far a single setting moves each count, and that the run is met only once the encoder estimator joins the first two changes.

and all 32 grasps fail. The latency contributes 15 mm of that, and the rest is the belt moving while the arm travels and the gripper closes, as Equation (10.9) predicts.

Case 4 raises the latency to 300 ms. Aimed at the last report, the gripper closes 84.7 mm from the nearest product and none of 42 grasps take (Table 10.4), and Figure 10.5 shows one of them closing

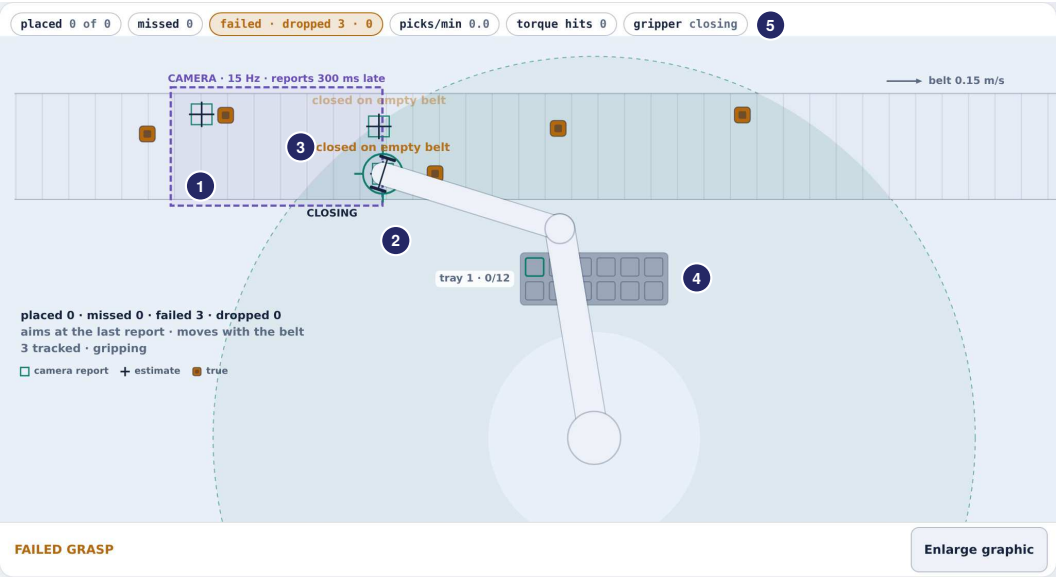


Figure 10.5 The picking application in scenario 4, Camera delay, cropped to the plant view, at $t = 8.20$ s with the settings the scenario loads, so the controller aims at the camera's last report. (1) The 15 Hz camera reports 300 ms late. (2) The gripper is closing inside the ring that marks its aim, over the last report of the product it was sent for, while that product has ridden 98.7 mm further along the 0.15 m/s belt. (3) Both floating labels read closed on empty belt. (4) The tray waits empty, and (5) the badges count three failed grasps, nothing placed and no torque-limit hits.

on empty belt. The application's grasp error is the distance to the nearest product on the belt, and the product each grasp was sent for is 108.0 mm away on average. With the latency set to zero the two distances fall to 55.0 and 63.0 mm and still no grasp takes. The distance to the product sent for falls by 45.0 mm, which is the belt speed times the latency, while the distance to the nearest product falls by only 29.7 mm, and the fit of Figure 10.2 extends the same arithmetic over eleven latencies. With the encoder estimator at 300 ms, 24 of 25 products are placed with no failed grasp and a position error of 0.99 mm, because the encoder count at the report's time stamp comes from the stored history.

🕮 **Pause.** Try Exercise 10.1 before reading on.

10.6.2 Speed at the Grasp

Case 2 runs the belt at 0.30 m/s with the arm still arriving early and waiting at rest. The gripper closes 0.96 mm from each product on average, and all 60 grasps in the first minute fail on relative speed, the product sliding under the gripper at 0.3005 m/s against the 0.15 m/s tolerance. On the slow belt of case 1 the same wait left 0.0994 m/s, inside the limit. Moving with the belt, the relative speed at the grasp falls to 0.0027 m/s and 30 products are placed with no failure. Switching the torque feedforward off in this case raises the tracking error from 0.61 to 16.1 mm rms, costs 15 failed grasps, 11 of them on position and 4 on speed, and brings one torque-limit hit.

🕮 **Pause.** Try Exercise 10.2 before reading on.

10.6.3 A Typed-In Speed Against the Encoder

Case 3 runs the belt at 0.20 m/s with the controller assuming 0.20 m/s, and at 20 s the line speeds up to 0.30 m/s. All 9 products that reach the arm in the first 20 s are placed. After the change the estimate falls behind by the speed error times the age of each report, and by 60 s 40 grasps have failed, closing

99.7 mm from the nearest product and 120.4 mm from the one each was sent for, with 14 products missed. The encoder estimator measures what the belt did rather than what it was told, and over the same minute it places 24 of the 25 products that arrive with no failed grasp.

10.6.4 Lost Frames

Case 5 loses half of all frames on a 0.25 m/s belt, where a product crosses the field of view in 1.6 s, and the camera goes dark for 3 s at 27.9 s. Each product seen is reported 12.0 times on average, against 23.5 with no frames lost. Twenty-seven products are placed and three are missed, and all three crossed the field of view in the dark and were never reported. No grasp fails, and the estimate stays 1.28 mm from its product on average against 1.07 mm with no frames lost, where the count is again 27 placed and 3 missed. Over six other seeds the misses are always the products never reported, and at one seed losing half the frames costs a single miss that the full stream avoids. A product seen once is carried forward by the encoder for as long as it rides the belt, and a product never seen cannot be picked by any estimator.

> ✦ DID YOU KNOW? **Half the Frames, the Same Count**
>
> From the simulator: in case 5, losing half of all camera frames costs no product. The arm places 27 and misses 3 either way, and the estimate moves only from 1.07 mm to 1.28 mm from its product (Table 10.4). Once a product has been seen, the belt encoder carries it.

10.6.5 Payload and Grip

Case 6 changes the product to a 1.5 kg case held by a 60 N cup, with the controller's arm model still knowing only the arm. The planner believes the carries fit inside the drives and they do not. The commands reach 47.1 N m at the shoulder and 24.3 N m at the elbow against limits of 40 and 17 N m, the torque limit is hit 11 times, and the tip lags its reference by 6.20 mm rms, 43.7 mm at worst. With the payload in the model there are no hits and the tracking error is 0.26 mm rms, 0.75 mm at worst. Lowering the acceleration limits to 25 and 50 rad/s^2 instead also stops the hits, at six of seven seeds, while the lag stays at 4.93 mm rms because the feedforward still leaves the case out. Case 7 cuts the grip to 15 % of 10 N. At the default carry limit of 30 m/s^2, 24 of the 25 products that arrive are dropped and the other is still on the belt. At 15 m/s^2, 20 are still dropped, and at 14.5 or 14 m/s^2 all 24 reach the tray. Over seven seeds a 15 m/s^2 limit drops 8 to 23 products and 14 m/s^2 drops none.

🕮 **Pause.** Try Exercise 10.3 before reading on.

10.6.6 Scheduling and the Throughput Ceiling

Case 8 sends a burst of twelve products at the minimum spacing into a stream of 20 a minute on a 0.25 m/s belt. Taking the oldest first, the arm spends 16 grasps on products it could no longer catch, closing 145.6 mm from the nearest one on average, and places 22 while 8 are missed. Feasible first skips what cannot be caught and places 25, with 5 missed and no failed grasp. Over seven seeds feasible first never fails a grasp while oldest first fails 8 to 18, and at one seed both place 31. Pushed with 120 arrivals a minute on a 0.30 m/s belt while moving with it, the arm places 105 products in 90 s, 70.0 min^{-1}, and misses 46 of the 161 that arrive. At that ceiling one pick takes 0.857 s, made up of 0.303 s of approach, 0.100 s of closing, 0.313 s of carry, 0.133 s of placing and release and 0.009 s of waiting.

10.6.7 The Production Run

The challenge sends 100 boxes of 0.6 kg at 36 a minute on a belt that speeds up from 0.22 to 0.30 m/s at 60 s, with 200 ms of latency, one frame in ten lost and the cup at 60 % of 10 N. It is met by placing at least 95 with no drop and no torque-limit hit. As shipped, waiting at rest with an assumed belt speed, it places none and fails 224 grasps and hits the torque limit once, and the run stops at 181.9 s when the last box is decided. The 68 failures before the speed change are on relative speed, and the

156 after it are on position while also moving too fast. Moving with the belt drops 35 boxes, because the full carry limit asks a 6 N grip for more than it holds. An 8 m/s^2 carry limit stops the drops and places 35, 34 of them before the speed change defeats the assumed speed. The encoder estimator brings the count to 98 of 100 with no failed grasp, and putting the payload in the arm model keeps 98 while cutting the tracking error from 1.97 to 0.25 mm rms over the run, which ends when the last box is decided at 180.2 s. The payload model also lowers the largest carry force from 5.46 to 4.89 N, that is from 9.10 to 8.15 m/s^2 against the 10 m/s^2 the grip allows. Over six other seeds the same four changes place 95 to 100 boxes, and the shipped settings place none.

10.7 Technology in Practice

The cell of Section 10.3 reads joint encoders, a belt encoder, a camera and a part-present switch, and it moves two joint drives and a gripper. Figure 10.6 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

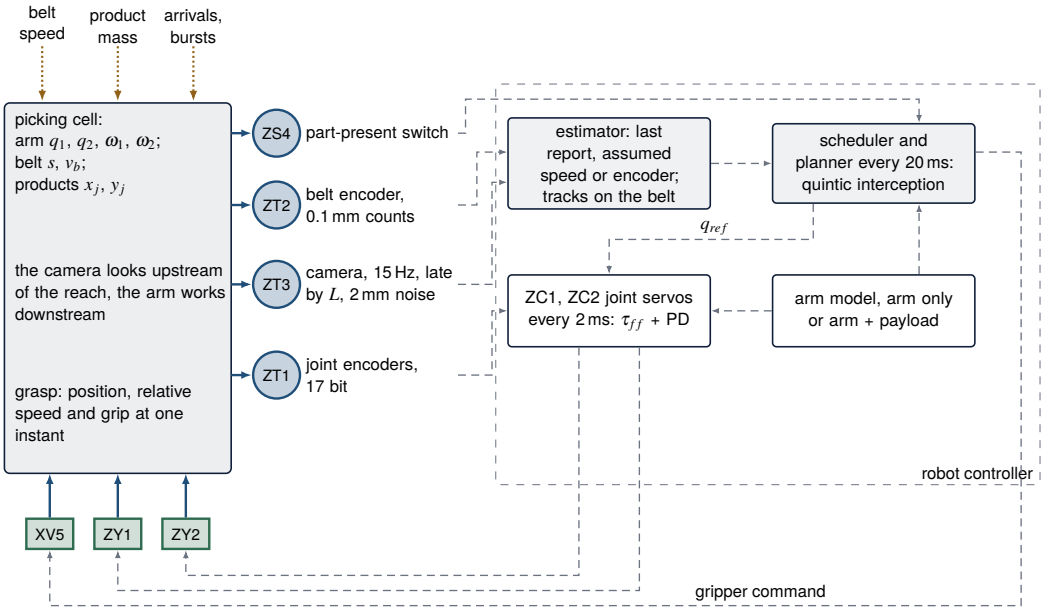
A robot synchronized with a moving conveyor is a conveyor tracking robot, and it needs frequent updates of the desired position on the conveyor. Line tracking is done either with the robot's base moving along the line or with the base fixed beside it, as here [18]. The camera finds the products, and the belt encoder says how far they have gone since. The camera of Equation (10.6), 15 Hz with a latency and 2 mm of noise, stands for such a vision system, and its lost frames for the frames a real one loses to its lighting. The first SCARA prototype measured each joint with an optical encoder [11], and the 17-bit counts of Table 10.1 stand for such encoders.

> **TECHNOLOGY Sensor: Machine Vision Camera**
>
> A vision system is one or more cameras interfaced to guide a robot to locate and identify a part. On a conveyor the camera finds each product and sends its offset to the robot controller, which combines it with the conveyor's speed to track the product until it comes within reach. One published line of this kind picked sachets from random places at 50 a minute to an accuracy of ±3 mm. Lighting is the usual trouble: moving parts, factory lights, the air and daylight from windows can upset a system that worked in the laboratory, and parallax and the match between camera and robot data call for a high mount and a calibration [18]. On a trend, a lighting fault looks like lost frames: the reports stop and the estimate rides on the belt encoder.

What Moves It

Each joint of that prototype was driven by a direct-current servo motor through a strain-wave (harmonic) gear reducer, and vibration in its servo motion led to a motor with tachogenerator feedback in the second [11]. The torque limits of 40 and 17 N m in Equation (10.1) stand for such drives, with the current loop and the gearbox taken as ideal (Section 10.3.5). The gripper of Equation (10.7), closing in 0.10 s and holding up to 10 N at full grip, stands for a vacuum cup, whose holding force must exceed the forces of the carry.

> **TECHNOLOGY Final Element: Vacuum Cup**
>
> A vacuum cup is a pneumatic pickup that holds a part by suction from a venturi or a vacuum pump, and it lets the part go when the vacuum supply is switched off. Cups are round or oval, from 12.7 mm to 101.6 mm across, and made of silicone, other rubbers or polyurethane. Choosing one takes experience and sometimes trial and error, and a hole in the part, a rough surface or a finish that must not be marked can rule a cup out. The force a grip must supply is the sum of the forces from mass, inertia and the Coriolis and centrifugal effects of the arm's motion, all of which change as the robot moves [18]. On a trend, a product lost in the carry is a part-present signal that drops out before the tray.



ZY1, ZY2: shoulder and elbow drives, a servo motor and gear reducer on each joint, limited to 40 and 17 N m. XV5: the vacuum valve of the gripper cup, which closes in 0.10 s and holds up to 10 N at full grip.

Figure 10.6 The loops of the chapter with real devices in place of the schematic's symbols. ZT1 counts the joint angles, ZT2 the belt travel, ZT3 is the camera and ZS4 the part-present switch. In the robot controller (dashed frame) the estimator feeds the scheduler and planner, which sets the joint references q_{ref} of the servos ZC1 and ZC2 and commands the gripper valve XV5. The servos drive ZY1 and ZY2 with torques computed from the arm model. Notice that the camera and the belt encoder meet only in the estimator, where the time stamp of each report lets the encoder count move it on.

Where the Controller Runs

The servo of this chapter runs every 2 ms, and the scheduler and the replanning every 20 ms. In discrete manufacturing a supervisor for each cell may coordinate a machine controller and a robot controller under the plant's supervisory loop, and such loops run with cycle times from milliseconds to minutes [19]. The first SCARA motions were computed by an 8-bit microcomputer, which shaped a motion curve with a deceleration interval twice its acceleration interval to prevent vibration [11]. The structure of this chapter lives in the robot controller: the computed-torque servo is its fast layer, and the estimator, the scheduler and the quintic planner sit in the layer above it.

What Is Modernizing It

Three-dimensional machine vision has lagged two-dimensional cameras because it needs light sources and calibration built into the sensor, and better auto-calibration routines are widening its use. Recognition based on geometry tolerates changes of lighting, focus and shape, and it is key to bin picking and tracking [18]. Lightweight arms with seven or more joints, often with force sensing in each joint, are now on the market, and robot hands are beginning to enter manufacturing [20]. Corrections that restore the rigid-body condition between camera and robot coordinates make the registration of the next segment more accurate [21].

IN PRACTICE **Registering the Camera to the Arm**

Before a vision-guided cell runs, the camera's coordinates are tied to the robot's. The common method measures the same points, called fiducials, in both frames and fits one rigid transformation between them. Noise and bias in those points then shift every target the robot is sent to. In NIST tests on two arms, corrections interpolated from fiducials near each target cut the root-mean-square target error to less than 20 % of its uncorrected value [21]. Lighting is settled on the line itself, because laboratory trials rarely match production [18]. The camera of Equation (10.6) reports in the arm's own coordinates with noise and no bias, so this step is absent from the model.

10.8 Limitations

The arm is planar and rigid with a point payload and ideal drives, so the vertical axis, link flexibility, the torque-speed curve and backlash are all absent. The camera and the encoder share a perfect clock, and recognition never fails inside the field of view except by losing a whole frame. The grasp is a threshold and a slip is total, which makes each failure mode clean but sharper than a real cup. The scheduler looks one product ahead and does not optimize an order of several or share the belt with another arm [8]. The numerical targets are the application's own design choices, not measured industrial performance.

One reporting detail affects how the numbers are read. The grasp error shown in the trends is the distance to the nearest product, which understates the miss when products are closely spaced, and this chapter reports the distance to the product each grasp was sent for beside it. The joint loops report no overshoot or settling time, because their reference moves a little at every sample and a step response is never defined.

OPEN CHALLENGE **Testing a Grasp on a Moving Part**

A NIST roadmap for robot dexterity lists grasping objects while in motion among the capabilities that need test methods. Its first test methods were to assume structured and static conditions, with no picking from a conveyor belt, and it notes that grasping cannot be separated from perceiving the object. Only in the longer term would tests cover cluttered, mixed parts that may be moving [20]. A moving grasp is hard to test because it couples the camera, the estimate, the arm and the gripper at one instant, as Equation (10.7) does. Progress would be a repeatable test that scores a whole picking cell against a stream of moving parts, as the production run of this chapter does for its model.

10.9 Conclusion

A picking robot acts on camera reports that are already old, and the simulator of this chapter makes that age, and what it costs, measurable. Aimed at the last report, a gripper on a 0.15 m/s belt closes behind the product it was sent for by $64.4 + 148.2L$ in mm, so removing the latency still leaves 63.0 mm, while moving each report on by the encoder travel since its time stamp places 24 of 25 products at every latency up to 0.5 s. Waiting at rest in the right place fails every grasp above the 0.15 m/s relative speed the gripper takes, and arriving at belt velocity places 30 of 32 at 0.30 m/s with 0.0027 m/s left at the grasp. A payload left out of the controller's model drives the torque commands to 47.1 and 24.3 N m against limits of 40 and 17, and a 1.5 N grip needs a carry limit of 14.5 m/s^2 or less rather than the 15 m/s^2 its holding force suggests. The arm and the grasp are idealized, and the clock is perfect. A fuller model would add the vertical axis, a torque-speed curve, clock offset between the camera and the encoder, recognition errors and a second arm sharing the belt. The application is freely available practice in meeting a moving object where it will be and at the speed it is going.

KEY POINTS

- A picking arm acts on camera reports that are already old, so it must meet each product where it will be and at the speed it is going.
- Aimed at the last report, a gripper still closes 63.0 mm behind its product with no latency, while the encoder estimate places 24 of 25 at every latency up to 0.5 s.
- Waiting at rest fails every grasp above the gripper's 0.15 m/s, and arriving at belt velocity places 30 of 32 at 0.30 m/s.
- A payload left out of the arm model drives the torques past their limits.

Where else this applies. The same age of a measurement is met in Volume 5, Chapter 13 and Volume 5, Chapter 7.

References

[1] S. Hutchinson, G. D. Hager, and P. I. Corke, "A tutorial on visual servo control," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 651–670, 1996.

[2] P. K. Allen, A. Timcenko, B. Yoshimi, and P. Michelman, "Automated tracking and grasping of a moving object with a robotic hand-eye system," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 2, pp. 152–165, 1993.

[3] T. D. Larsen, N. A. Andersen, O. Ravn, and N. K. Poulsen, "Incorporation of time delayed measurements in a discrete-time Kalman filter," in *Proceedings of the 37th IEEE Conference on Decision and Control*, vol. 4, 1998, pp. 3972–3977.

[4] E. A. Croft, R. G. Fenton, and B. Benhabib, "Optimal rendezvous-point selection for robotic interception of moving objects," *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, vol. 28, no. 2, pp. 192–204, 1998.

[5] D. Hujic, E. A. Croft, G. Zak, R. G. Fenton, J. K. Mills, and B. Benhabib, "The robotic interception of moving objects in industrial settings: strategy development and experiment," *IEEE/ASME Transactions on Mechatronics*, vol. 3, no. 3, pp. 225–239, 1998.

[6] J. S. Dai, "Robotics and automation for packaging in the confectionery industry," in *Robotics and Automation in the Food Industry*. Woodhead Publishing, 2013, pp. 401–419.

[7] R. Mattone, L. Adduci, and A. Wolf, "Online scheduling algorithms for improving performance of pick-and-place operations on a moving conveyor belt," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, vol. 3, 1998, pp. 2099–2105.

[8] H. I. Bozma and M. E. Kalalıoğlu, "Multirobot coordination in pick-and-place tasks on a moving conveyor," *Robotics and Computer-Integrated Manufacturing*, vol. 28, no. 4, pp. 530–538, 2012.

[9] C. G. Atkeson, C. H. An, and J. M. Hollerbach, "Estimation of inertial parameters of manipulator loads and links," *The International Journal of Robotics Research*, vol. 5, no. 3, pp. 101–119, 1986.

[10] F. Pierrot, C. Reynaud, and A. Fournier, "DELTA: a simple and efficient parallel robot," *Robotica*, vol. 8, no. 2, pp. 105–109, 1990.

[11] H. Makino, "Development of the SCARA," *Journal of Robotics and Mechatronics*, vol. 26, no. 1, pp. 5–8, 2014.

[12] J. Luh, M. Walker, and R. Paul, "Resolved-acceleration control of mechanical manipulators," *IEEE Transactions on Automatic Control*, vol. 25, no. 3, pp. 468–474, 1980.

[13] P. K. Khosla and T. Kanade, "Experimental evaluation of nonlinear feedback and feedforward control schemes for manipulators," *The International Journal of Robotics Research*, vol. 7, no. 1, pp. 18–28, 1988.

[14] G. J. Monkman, S. Hesse, R. Steinmann, and H. Schunk, *Robot Grippers*. Wiley-VCH, 2006.

[15] J. E. Bobrow, S. Dubowsky, and J. S. Gibson, "Time-optimal control of robotic manipulators along specified paths," *The International Journal of Robotics Research*, vol. 4, no. 3, pp. 3–17, 1985.

[16] T. Kröger and F. M. Wahl, "Online trajectory generation: basic concepts for instantaneous reactions to unforeseen events," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 94–111, 2010.

[17] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2009.

[18] R. Bostelman and J. Falco, "Survey of industrial manipulation technologies for autonomous assembly applications," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7844, 2012.

[19] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[20] J. Falco, J. Marvel, and E. Messina, "A roadmap to progress measurement science in robot dexterity and manipulation," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7993, 2014.

[21] M. Franaszek, G. S. Cheok, K. van Wyk, and J. A. Marvel, "Improving 3D vision-robot registration for assembly tasks," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8300, 2020.

10.10 Exercises

10.1 Concept check · 2 min

Scenario 4 aims the gripper at the camera's last report while the arm moves with the belt at 0.15 m/s. The camera latency is set to zero and nothing else is changed. What happens to the grasps?

(a) They take, because a report with no latency describes the product as it is now

(b) They fail on relative speed, because the gripper waits at rest under the report

(c) They still fail, because the product moves on while the arm travels and the gripper closes

(d) They take for half the products, those reported in the frame just before the approach

10.2 Concept check · 2 min

In scenario 2 the belt runs at 0.30 m/s and the arm arrives early and waits at rest where each product will be. The gripper closes 0.96 mm from the product on average, yet all 60 grasps of the first minute fail. Why?

(a) The camera noise of 2 mm is larger than the position error the gripper can take

(b) The product's own acceleration at the catch exceeds what the holding force can follow

(c) The arm must reach belt speed while the gripper closes, and the drives hit their torque limit

(d) The product slides under the gripper at the belt speed, twice the relative speed it can take

10.3 Concept check · 1 min

In scenario 6 the controller's arm model knows only the arm, and the planner checks every move against 90 % of each drive. Why does the torque limit still get hit while a 1.5 kg case is carried?

(a) The planner computes the torque with the controller's arm model, which leaves the case out

(b) The torque share of 90 % leaves too little for the Coulomb friction of the joints

(c) The servo gains amplify the quantization of the joint encoders into torque spikes

(d) The catch conserves momentum, so every grasp sends a torque step through both joints

10.4 Theory and calculation · 6 min

Use the grasp tolerance $e_p = 8\,\text{mm}$ and the closing time $t_{close} = 0.10\,\text{s}$ of Table 10.2 and Equation (10.9).

(a) An instantaneous camera ($L = 0$) feeds a controller that aims at the report and does not predict the closing interval, so the shortest possible T_d is t_{close}. Find the fastest belt on which a grasp can still meet e_p. *Checkpoint:* 0.08 m/s.

(b) Scenario 4 runs at 0.15 m/s with $L = 0.30\,\text{s}$, and the time from delivery to grasp averages $T_d = 0.42\,\text{s}$. Find the distance from the gripper to the product it was sent for. *Checkpoint:* 108 mm.

(c) A camera with no latency replaces the old one. Find how much of that distance it removes and how much remains. *Checkpoint:* 45 mm removed, 63 mm remaining.

10.5 Theory and calculation · 6 min

At the ready pose ($q_2 = 62.6°$) the empty arm has $M_{11} = 0.9996\,\mathrm{kg\,m^2}$. The shoulder servo has $K_{p1} = 400\,\mathrm{N\,m/rad}$ and $K_{d1} = 25\,\mathrm{N\,m\,s/rad}$. Treat the shoulder as a single inertia M_{11} under the PD law with the feedforward exact, so that its error obeys $M_{11}\ddot{e} + K_{d1}\dot{e} + K_{p1}e = 0$.

(a) Find the natural frequency and the damping ratio of the shoulder loop. *Checkpoint:* 20.0 rad/s and 0.63.

(b) A point mass m_p at the tip adds $m_p r^2$ to M_{11}, where r is the tip's distance from the shoulder. With $l_1 = 0.40\,\mathrm{m}$ and $l_2 = 0.35\,\mathrm{m}$, find r at the ready pose and M_{11} with a 1.5 kg case held. *Checkpoint:* $r = 0.641\,\mathrm{m}$ and $1.617\,\mathrm{kg\,m^2}$.

(c) Find the natural frequency and damping ratio with the case, and compare the relative change of each. *Checkpoint:* 15.7 rad/s and 0.49.

10.6 Theory and calculation · 6 min

The encoder estimator averages the reports of one product in belt coordinates with the gain $k_n = \max(1/n, 0.25)$ after the nth report. Each report carries independent noise of $\sigma = 2\,\mathrm{mm}$ on each axis.

(a) For an exponential average $u \leftarrow u + k(u_{rep} - u)$ with a fixed k, the stationary variance is $k/(2-k)$ times the variance of one report. Find the standard deviation per axis for $k = 0.25$. *Checkpoint:* 0.756 mm.

(b) The mean length of a two-dimensional error whose axes are independent normals of standard deviation s is $s\sqrt{\pi/2}$. Find it and compare it with the 0.99 mm measured in case 1. *Checkpoint:* 0.95 mm.

(c) In scenario 5 the belt runs at 0.25 m/s past the 0.40 m field of view of a 15 Hz camera, and half the frames are lost. Find how many reports a product gets on average with all frames and with half of them. *Checkpoint:* 24 and 12.

10.7 App challenge · 8 min

Load scenario 4, Camera delay. It aims at the last report (Where the chocolate is set to Last report), moves with the belt, and the camera reports 300 ms late. Use the default seed and run each part for 60 s at 2×. Read the scorecard tiles Placed, Failed grasps and Missed at 60 s. Counts are exact at the default seed.

(a) Run the scenario as loaded. *Checkpoint:* 0 placed, 42 failed grasps, 23 missed.

(b) Reload the scenario with the reload button, set Camera latency to 0 ms before Run, and run again. Explain the result with Equation (10.9). *Checkpoint:* 0 placed, 42 failed grasps, 23 missed.

(c) Reload the scenario, set Where the chocolate is to Encoder before Run, keep the 300 ms latency, and run again. *Checkpoint:* 24 placed (Placed reads 24 of 24), no failed grasp, none missed.

10.8 App challenge · 8 min

Load scenario 7, Reduced grip. The cup holds 15 % of 10 N and each product weighs 0.10 kg. Use the default seed, set each value before Run, run for 60 s at 2×, and read the tiles Placed and Drops.

(a) Compute the largest tip acceleration the grip can hold. *Checkpoint:* $15\,\mathrm{m/s^2}$.

(b) Reload the scenario with the reload button, set Carry acceleration limit to $15.0\,\mathrm{m/s^2}$, and run. *Checkpoint:* 4 placed, 20 dropped.

(c) Reload, set the limit to $14.5\,\mathrm{m/s^2}$, and run. Explain why the bound of part (a) is not a usable setting. *Checkpoint:* 24 placed, none dropped.

10.9 App challenge · 9 min

Load the challenge, a production run of 100: 0.6 kg boxes at 36 a minute, a belt that speeds up from 0.22 to 0.30 m/s at 60 s and a grip of 60 % of 10 N. Use the default seed and run at 5× until the run stops itself, about 180 s of simulated time. Reload the scenario with the reload button before each part, make the listed changes before Run, and read the tiles Placed, Drops, Failed grasps and Tracking error. Counts are exact. Give the tracking error within 0.02 mm.

(a) Set Meet the chocolate to Move with it. *Checkpoint:* 0 placed, 35 dropped, 158 failed grasps.

(b) Also set Carry acceleration limit to 8.0 m/s^2. *Checkpoint:* 35 placed, none dropped, 158 failed grasps.

(c) Also set Where the chocolate is to Encoder. *Checkpoint:* 98 placed, none dropped, no failed grasp, 1.97 mm rms.

(d) Also set Controller's arm model to Arm + payload. *Checkpoint:* 98 placed, 0.25 mm rms.

10.10 Application area · 4 min

The chapter assumes that the camera and the belt encoder share one clock. On a real line the trigger time stamped on a frame can be offset from the encoder's clock.

(a) Find the position error a 10 ms offset adds on a 0.30 m/s belt. *Checkpoint:* 3 mm.

(b) The encoder estimator leaves about 1 mm of error at the grasp, and the grasp tolerance is 8 mm. Adding the two errors in the worst case, find the largest clock offset the cell can take at 0.30 m/s. *Checkpoint:* About 23 ms.

10.11 Application area · 3 min

Pushed with 120 arrivals a minute on a 0.30 m/s belt, the arm completes one pick in 0.857 s: 0.303 s of approach, 0.100 s of closing, 0.313 s of carry, 0.133 s of placing and release and 0.009 s of waiting.

(a) Find the ceiling in picks per minute. *Checkpoint:* 70.0 per minute.

(b) A line must clear 120 products a minute with no misses. Find the least number of identical arms, and name what the chapter's scheduler lacks for that layout. *Checkpoint:* Two arms.

10.12 Review the agent · 5 min

An AI assistant reads the scenario 6 results (a 1.5 kg case, Arm only model, 11 torque-limit hits and 6.20 mm rms of tracking error) and writes: "The servo is too soft for the heavy case. Double both position gains to 800 and 200 N m/rad so the joints hold their references more stiffly. The lag will shrink and the torque-limit hits will stop." Find the flaw and say how you would check it in the app.