Controller design for networks of switching servers with setup times

Erjen Lefeber

Eindhoven University of Technology
Department of Mechanical Engineering

Rendezvous for Mathematics and Computers in Process Engineering

Modeling, optimization and control at different levels

February 19, 2010, TU/e Helix

Acknowledgements This work was supported by the Netherlands Organization for Scientific Research (NWO-VIDI grant 639.072.072).

Inspired by discussions with:

- Varvara Feoktistova, Alexey Matveev (St. Petersburg)
- Jan van der Wal, Josine Bruin
- Stefan Lämmer (TU Dresden)
- Gideon Weiss, Yoni Nazarathy (Haifa)

Motivation

Problem

Problem

How to control these networks?

Decisions: When to switch, and to which job-type

Goals: Minimal number of jobs, minimal flow time

Current approach

Start from policy, analyze resulting dynamics

Kumar, Seidman (1990)

Problem

Problem

How to control these networks?

Decisions: When to switch, and to which job-type

Goals: Minimal number of jobs, minimal flow time

Current approach

Start from policy, analyze resulting dynamics

Kumar, Seidman (1990)

Problem

Problem

How to control these networks?

Decisions: When to switch, and to which job-type

Goals: Minimal number of jobs, minimal flow time

Current approach

Start from policy, analyze resulting dynamics

Kumar, Seidman (1990)

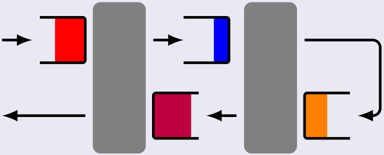

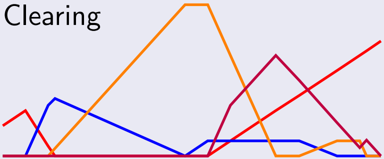

Problem

Current status (after two decades)

Several policies exist that guarantee stability of the network

Remark

Stability is only a prerequisite for a good policy

Open issues

- Do existing policies yield satisfactory network performance?
- How to obtain pre-specified network behavior?

Main subject of study (modest)

Fixed, deterministic flow networks (not evolving, constant inflow)

Problem

Current status (after two decades)

Several policies exist that guarantee stability of the network

Remark

Stability is only a prerequisite for a good policy

Open issues

- Do existing policies yield satisfactory network performance?
- How to obtain pre-specified network behavior?

Main subject of study (modest)

Fixed, deterministic flow networks (not evolving, constant inflow)

Problem

Current status (after two decades)

Several policies exist that guarantee stability of the network

Remark

Stability is only a prerequisite for a good policy

Open issues

- Do existing policies yield satisfactory network performance?
- How to obtain pre-specified network behavior?

Main subject of study (modest)

Fixed, deterministic flow networks (not evolving, constant inflow)

Problem

Current status (after two decades)

Several policies exist that guarantee stability of the network

Remark

Stability is only a prerequisite for a good policy

Open issues

- Do existing policies yield satisfactory network performance?
- How to obtain pre-specified network behavior?

Main subject of study (modest)

Fixed, deterministic flow networks (not evolving, constant inflow)

Approach

Approach

Use ideas/concepts from control theory

Approach

Notions from control theory

1. Generate feasible reference trajectory
2. Design (static) state feedback controller
3. Design observer
4. Design (dynamic) output feedback controller

Parallels with this problem

1. Determine desired system behavior
2. Derive non-distributed/centralized controller
3. Can state be reconstructed?
4. Derive distributed/decentralized controller

Approach

Notions from control theory

1. Generate feasible reference trajectory
2. Design (static) state feedback controller
3. Design observer
4. Design (dynamic) output feedback controller

Parallels with this problem

1. Determine desired system behavior
2. Derive non-distributed/centralized controller
3. Can state be reconstructed?
4. Derive distributed/decentralized controller

Example 1: Single machine

Single machine

$$\sigma_{12} = 3, \sigma_{21} = 1$$

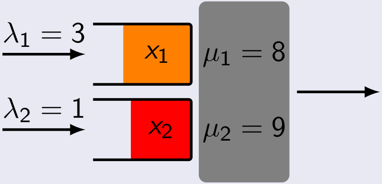

Objective

Minimize:

$$\limsup_{t\to\infty} \frac{1}{t}\int_0^t x_1(\tau) + x_2(\tau)\,\mathrm{d}\tau \qquad \text{or} \qquad \frac{1}{T}\int_0^T x_1(\tau) + x_2(\tau)\,\mathrm{d}\tau$$

Example 1: Single machine

Single machine

$$\sigma_{12} = 3, \sigma_{21} = 1$$

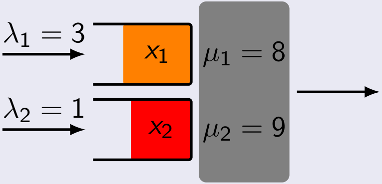

Objective

Minimize:

$$\limsup_{t\to\infty} \frac{1}{t}\int_0^t x_1(\tau) + x_2(\tau)\,\mathrm{d}\,\tau \qquad \text{or} \qquad \frac{1}{T}\int_0^T x_1(\tau) + x_2(\tau)\,\mathrm{d}\,\tau$$

Desired behavior (Problem I)

Single machine

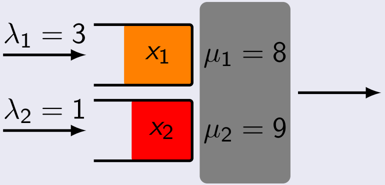

Desired behavior

Remarks

- Many existing policies assume non-idling a-priori
- Slow-mode optimal if $\lambda_1(\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2}) - (\lambda_1 - \lambda_2)(1 - \frac{\lambda_2}{\mu_2}) < 0$.
- Trade-off in wasting capacity: idle $\Leftrightarrow$ switch more often

Desired behavior (Problem I)

Single machine

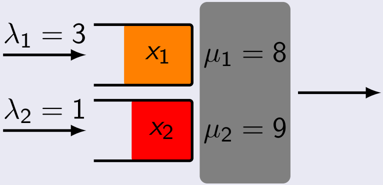

Desired behavior

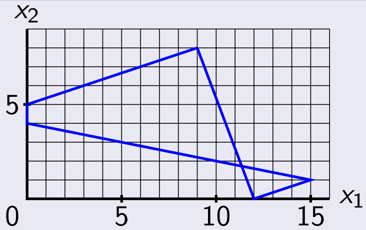

Remarks

- Many existing policies assume non-idling a-priori
- Slow-mode optimal if $\lambda_1(\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2}) - (\lambda_1 - \lambda_2)(1 - \frac{\lambda_2}{\mu_2}) < 0$.
- Trade-off in wasting capacity: idle $\Leftrightarrow$ switch more often

TU/e
NWO

Desired behavior (Problem I)

Single machine

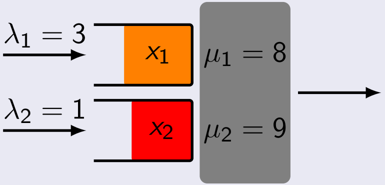

Desired behavior

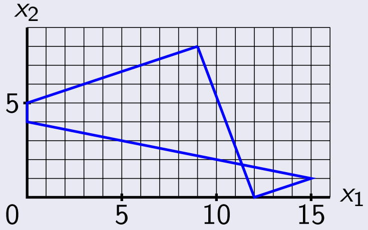

Remarks

- Many existing policies assume non-idling a-priori
- Slow-mode optimal if $\lambda_1(\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2}) - (\lambda_1 - \lambda_2)(1 - \frac{\lambda_2}{\mu_2}) < 0$.
- Trade-off in wasting capacity: idle $\Leftrightarrow$ switch more often

TU/e NWO

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

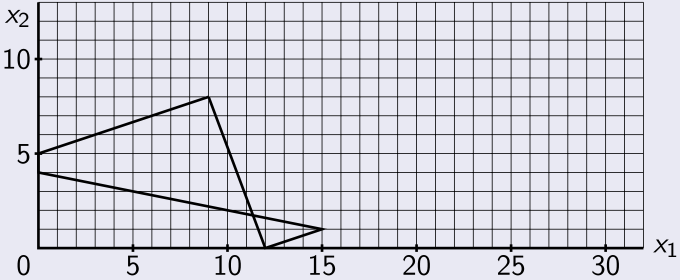

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

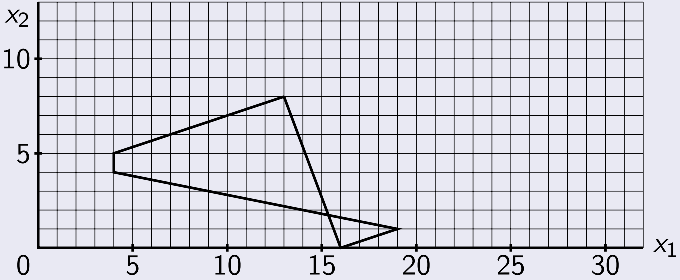

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

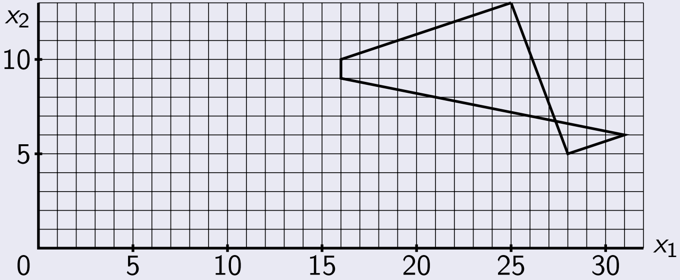

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time ⇒ system settles down at constant energy level

Lyapunov function candidate

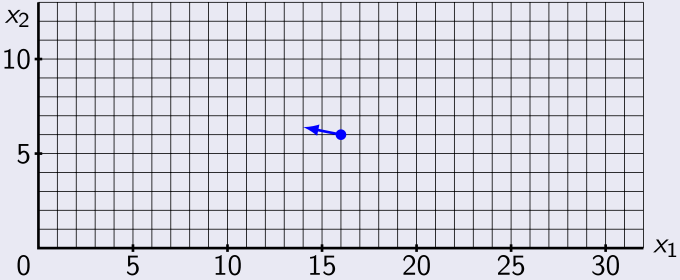

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time ⇒ system settles down at constant energy level

Lyapunov function candidate

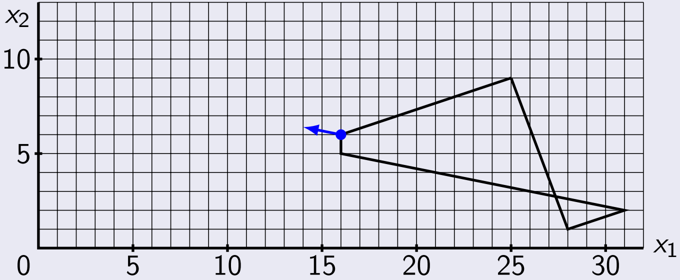

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

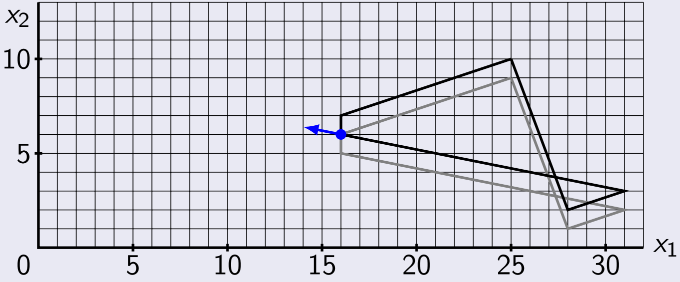

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

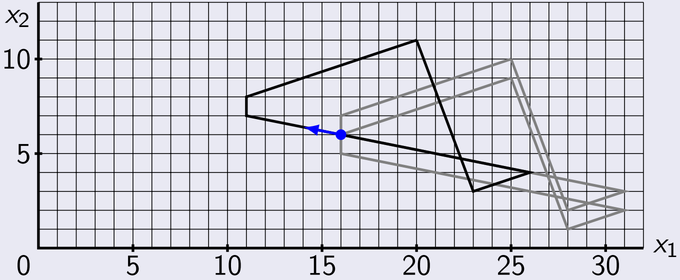

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

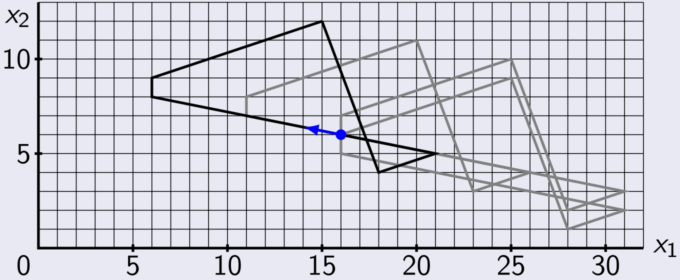

Controller design (Problem II)

Main idea

Lyapunov: if energy is decreasing all the time $\Rightarrow$ system settles down at constant energy level

Lyapunov function candidate

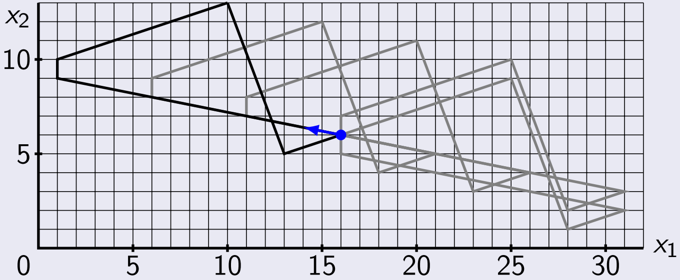

Controller design

Lyapunov function candidate

The smallest additional mean amount of work from all feasible curves for state (work: $x_1/\mu_1 + x_2/\mu_2$).

Phase plane

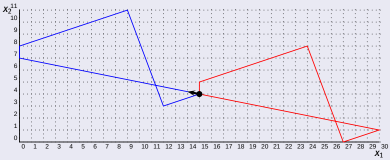

Time evolution work

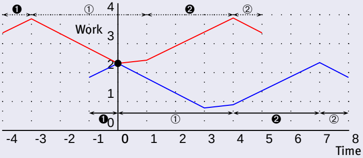

Controller design

Let Lyapunov function candidate decrease as quickly as possible

Controller design

Lyapunov function candidate

The smallest additional mean amount of work from all feasible curves for state (work: $x_1/\mu_1 + x_2/\mu_2$).

Phase plane

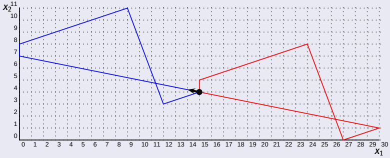

Time evolution work

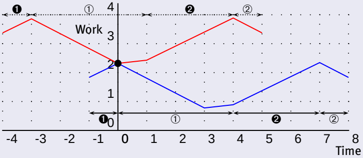

Controller design

Let Lyapunov function candidate decrease as quickly as possible

Controller design (Result)

Single machine

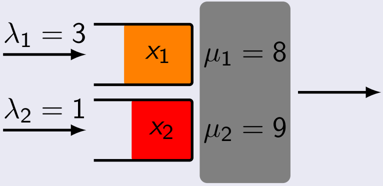

Desired behavior

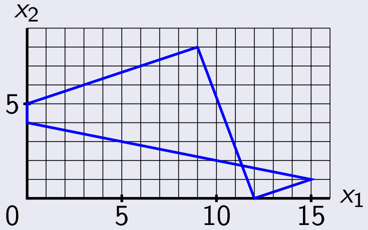

Resulting Controller, cf. [Lefeber, Rooda (2006)]

- When serving type 1:
  1. empty buffer
  2. serve until $x_2 \geq 5$
  3. switch to type 2
- When serving type 2:
  1. empty buffer
  2. serve until $x_1 \geq 12$
  3. switch to type 1

Controller design (Result)

Single machine

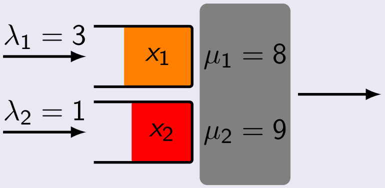

Desired behavior

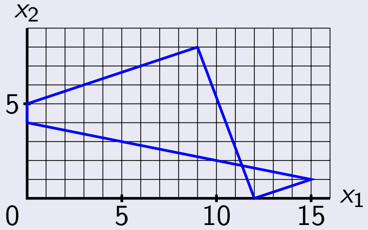

Resulting Controller, cf. [Lefeber, Rooda (2006)]

- When serving type 1:
  1. empty buffer
  2. serve until $x_2 \geq 5$
  3. switch to type 2
- When serving type 2:
  1. empty buffer
  2. serve until $x_1 \geq 12$
  3. switch to type 1

TU/e
NWO

Controller design (Result)

Single machine

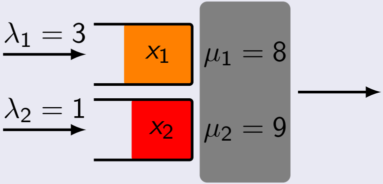

Desired behavior

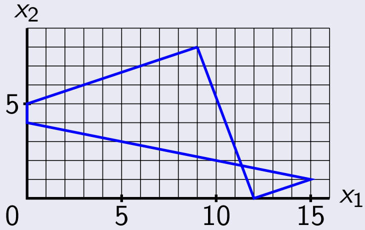

Resulting Controller, cf. [Lefeber, Rooda (2006)]

- When serving type 1:
  1. empty buffer
  2. serve until $x_2 \geq 5$
  3. switch to type 2
- When serving type 2:
  1. empty buffer
  2. serve until $x_1 \geq 12$
  3. switch to type 1

Recap

Notions from control theory

1. Generate feasible reference trajectory
2. Design (static) state feedback controller
3. Design observer
4. Design (dynamic) output feedback controller

Parallels with this problem

1. Determine desired system behavior
2. Derive non-distributed/centralized controller
3. Can state be reconstructed?
4. Derive distributed/decentralized controller

Example 2: Kumar-Seidman case

Transactions on Automatic Control, Vol 35, No 3, March 1990

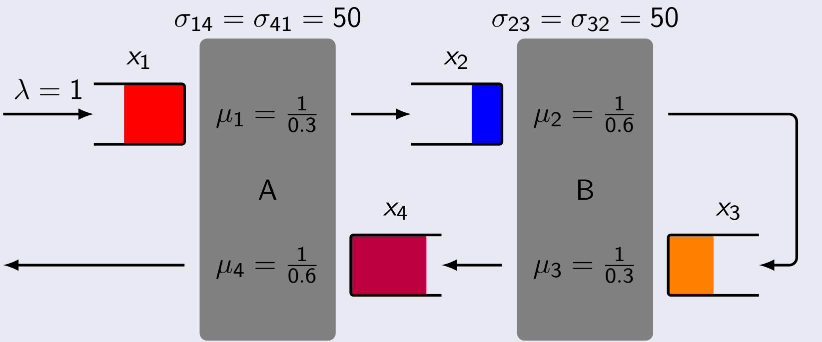

Observation

Sufficient capacity (consider period of at least 1000).

Desired behavior

Desired behavior

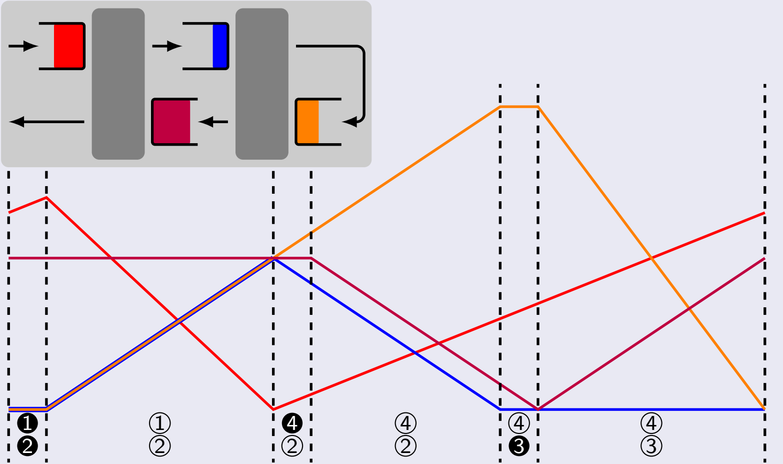

Resulting controller

Network

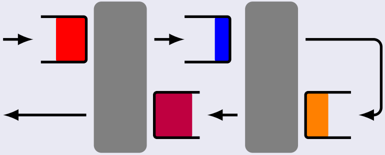

Desired behavior

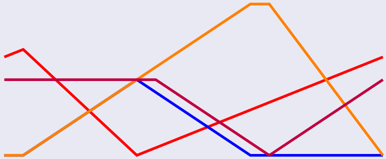

Resulting controller

Mode (1,2): to (4,2) when both $x_1 = 0$ and $x_2 + x_3 \geq 1000$

Mode (4,2): to (4,3) when both $x_2 = 0$ and $x_4 \leq 83\frac{1}{3}$

Mode (4,3): to (1,2) when $x_3 = 0$

Remark:

- Non-distributed/centralized controller

Proof

Monodromy operator

x_i^k: buffer contents at k^{th} start of mode (1,2). For $k > 2$:

$$\begin{aligned} x_1^{k+1} &= 50 + \frac{3}{7}(x_1^k + 50) + \max\left(\frac{3}{7}(x_1^k + 50), \frac{3}{5}x_4^k\right) \\ x_2^{k+1} &= 0 \qquad x_3^{k+1} = 0 \qquad x_4^{k+1} = \frac{5}{7}(x_1^k + 50) \end{aligned} \tag{1}$$

Observation

With $y_1^k = (x_1^k - 650)/7$, $y_4^k = (x_4^k - 500)/5$ we get from (1):

$$0 \leq \max(y_1^{k+2}, y_4^{k+2}) \leq \frac{6}{7}\max(y_1^k, y_4^k)$$

So system converges to fixed point $(650, 0, 0, 500)$.

TU/e NWO

Observability

Network

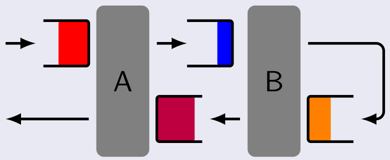

Assumptions

- Clearing policy used for machine B
- At $t = t_1$: ③ starts
- At $t = t_2 > t_1$: ③ stops

System state can be reconstructed at machine A

- $x_3(t_2) = 0$, and $0.3(t_2 - t_1) = x_3(t_1) = x_3(t_1 - 50)$
- $x_2(t_1 - 50) = 0$, and $x_2(t_2) = \int_{t_1-50}^{t_2} u_1(\tau)\,\mathrm{d}\tau$

Observation

Observability determined by network topology

Observability

Network

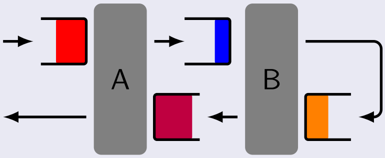

Assumptions

- Clearing policy used for machine B
- At $t = t_1$: ③ starts
- At $t = t_2 > t_1$: ③ stops

System state can be reconstructed at machine A

- $x_3(t_2) = 0$, and $0.3(t_2 - t_1) = x_3(t_1) = x_3(t_1 - 50)$
- $x_2(t_1 - 50) = 0$, and $x_2(t_2) = \int_{t_1-50}^{t_2} u_1(\tau)\,\mathrm{d}\,\tau$

Observation

Observability determined by network topology

Observability

Network

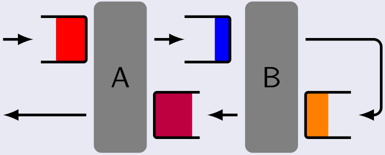

Assumptions

- Clearing policy used for machine B
- At $t = t_1$: ③ starts
- At $t = t_2 > t_1$: ③ stops

System state can be reconstructed at machine A

- $x_3(t_2) = 0$, and $0.3(t_2 - t_1) = x_3(t_1) = x_3(t_1 - 50)$
- $x_2(t_1 - 50) = 0$, and $x_2(t_2) = \int_{t_1-50}^{t_2} u_1(\tau)\,\mathrm{d}\,\tau$

Observation

Observability determined by network topology

Network

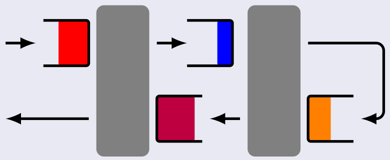

er,

Desired behavior

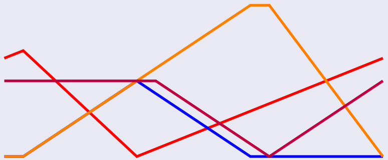

Distributed controller

Serving 1: Serve at least 1000 jobs until $x_1 = 0$, then switch. Let $\bar{x}_1$ be nr of jobs served.

Serving 2: Serve at least 1000 jobs until $x_2 = 0$, then switch.

Serving 4: Let $\bar{x}_4$ be nr of jobs in Buffer 4 after setup. Serve $\bar{x}_4 + \frac{1}{2}\bar{x}_1$ jobs, then switch.

Serving 3: Empty buffer, then switch.

Simulation results Initial condition $(1000, 1000, 1000, 1000)$.
Deterministic/Exponential service times, setup times.

Distributed controller

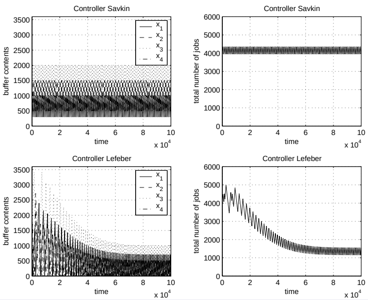

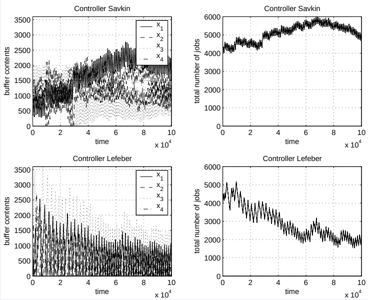

Conclusions

New approach

1. Determine desired system behavior (trajectory generation)
2. Derive non-distributed/centralized controller (state feedback)
3. Derive distributed/decentralized controller (output feedback)

Advantage

All three problems can be considered separately

Centralized control

Approach can deal with

- Arbitrary networks
- Finite buffers
- Transportation delays

Decentralized control

- Observer based approach results in new, tailor-made controllers that perform better

Conclusions

New approach

1. Determine desired system behavior (trajectory generation)
2. Derive non-distributed/centralized controller (state feedback)
3. Derive distributed/decentralized controller (output feedback)

Advantage

All three problems can be considered separately

Centralized control

Approach can deal with

- Arbitrary networks
- Finite buffers
- Transportation delays

Decentralized control

- Observer based approach results in new, tailor-made controllers that perform better

Conclusions

New approach

1. Determine desired system behavior (trajectory generation)
2. Derive non-distributed/centralized controller (state feedback)
3. Derive distributed/decentralized controller (output feedback)

Advantage

All three problems can be considered separately

Centralized control

Approach can deal with

- Arbitrary networks
- Finite buffers
- Transportation delays

Decentralized control

- Observer based approach results in new, tailor-made controllers that perform better

Conclusions

New approach

1. Determine desired system behavior (trajectory generation)
2. Derive non-distributed/centralized controller (state feedback)
3. Derive distributed/decentralized controller (output feedback)

Advantage

All three problems can be considered separately

Centralized control

Approach can deal with

- Arbitrary networks
- Finite buffers
- Transportation delays

Decentralized control

- Observer based approach results in new, tailor-made controllers that perform better

For further reading

E. Lefeber and J.E. Rooda.
Controller design of switched linear systems with setups.
Physica A, 363(1):48–61, April 2006.

E. Lefeber, J.E. Rooda.
Controller Design for Flow Networks of Switched Servers with Setup Times: the Kumar-Seidman Case as an Illustrative Example.
Asian Journal of Control, 10(1), 55-66, 2008.

Background: Control theory

System dynamics (linear)

$$\dot{x}(t) = Ax(t) + Bu(t) \qquad x \in R^n, u \in R^k$$
$$y(t) = Cx(t) \qquad y \in R^m$$

where $u(\cdot)$ is a function to be designed.

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller $u(\cdot)$, such that

$$\lim_{t\to\infty} \|x(t) - x_r(t)\| = 0.$$

TU/e
NWO

Background: Control theory

System dynamics (linear)

$$\dot{x}(t) = Ax(t) + Bu(t) \qquad x \in R^n, u \in R^k$$
$$y(t) = Cx(t) \qquad y \in R^m$$

where $u(\cdot)$ is a function to be designed.

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller $u(\cdot)$, such that

$$\lim_{t \to \infty} \|x(t) - x_r(t)\| = 0.$$

Background: Control theory

System dynamics (linear)

$$\dot{x}(t) = Ax(t) + Bu(t) \qquad x \in R^n, u \in R^k$$
$$y(t) = Cx(t) \qquad y \in R^m$$

where $u(\cdot)$ is a function to be designed.

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller $u(\cdot)$, such that

$$\lim_{t\to\infty} \|x(t) - x_r(t)\| = 0.$$

Background: Control theory, Example tracking control

Controller

$$u = u_r + K(x - x_r)$$

Error dynamics

Define $e = x - x_r$, then:

$$\dot{e} = Ax + B(u_r + Ke) - (Ax_r + Bu_r) = (A + BK)e$$

Make sure that K is such that eigenvalues of $A + BK$ are in left half of complex plane.

Remark

The controller design holds for arbitrary reference.

Background: Control theory, Example tracking control

Controller

$$u = u_r + K(x - x_r)$$

Error dynamics

Define $e = x - x_r$, then:

$$\dot{e} = Ax + B(u_r + Ke) - (Ax_r + Bu_r) = (A + BK)e$$

Make sure that K is such that eigenvalues of $A + BK$ are in left half of complex plane.

Remark

The controller design holds for arbitrary reference.

Background: Control theory, Example tracking control

Controller

$$u = u_r + K(x - x_r)$$

Error dynamics

Define $e = x - x_r$, then:

$$\dot{e} = Ax + B(u_r + Ke) - (Ax_r + Bu_r) = (A + BK)e$$

Make sure that K is such that eigenvalues of $A + BK$ are in left half of complex plane.

Remark

The controller design holds for arbitrary reference.

Background: Control theory

System dynamics (linear)

$$\dot{x}(t) = Ax(t) + Bu(t) \qquad x \in R^n, u \in R^k$$
$$y(t) = Cx(t) \qquad y \in R^m$$

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller

Problem III: Observer design

Reconstruct x using only measurement of y

Background: Control theory

System dynamics (linear)

$$\dot{x}(t) = Ax(t) + Bu(t) \qquad x \in R^n, u \in R^k$$
$$y(t) = Cx(t) \qquad y \in R^m$$

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller

Problem III: Observer design

Reconstruct x using only measurement of y

Background: Control theory, Example observer design

Observer

$$\dot{\hat{x}} = A\hat{x} + Bu + L(y - \hat{y})$$
$$\hat{y} = C\hat{x}$$

Observer error dynamics

Define $e = x - \hat{x}$, then

$$\dot{e} = A\hat{x} + Bu + LCe - (Ax + Bu) = (A + LC)e$$

Make sure that L is such that eigenvalues of $A + LC$ are in left half of complex plane.

Background: Control theory

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller assuming x is available for measurement

Problem III: Observer design

Reconstruct x using only measurement of y

Problem IV: Output feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller assuming only y is available for measurement

Background: Control theory

Problem I: Trajectory generation

Determine feasible functions $x_r(t)$, $u_r(t)$.

Problem II: State feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller assuming x is available for measurement

Problem III: Observer design

Reconstruct x using only measurement of y

Problem IV: Output feedback tracking control

Given arbitrary feasible $x_r(t)$, $u_r(t)$, find a controller assuming only y is available for measurement

Background: Control theory, Example tracking control

System dynamics (linear)

$$\dot{x}(t) = Ax(t) + Bu(t) \qquad x \in R^n, u \in R^k$$
$$y(t) = Cx(t) \qquad y \in R^m$$

Dynamic output feedback tracking controller

$$u = u_r + K(\hat{x} - x_r)$$
$$\dot{\hat{x}} = A\hat{x} + Bu(t) + L(y - \hat{y})$$
$$\hat{y} = C\hat{x}$$

where K and L from previous designs can be used.

Adaptive control

System dynamics

$$\dot{x} = ax + u \qquad a \text{ unknown parameter}$$

Controller

$$u = -\hat{a}x - kx \qquad k > 0$$
$$\dot{\hat{a}} = \gamma x^2 \qquad \gamma > 0$$

Result

$$\lim_{t\to\infty} x(t) = 0$$

Furthermore, $\hat{a}(t)$ converges to a constant (not to a!)

Adaptive control

System dynamics

$$\dot{x} = ax + u \qquad a \text{ unknown parameter}$$

Controller

$$u = -\hat{a}x - kx \qquad k > 0$$
$$\dot{\hat{a}} = \gamma x^2 \qquad \gamma > 0$$

Result

$$\lim_{t \to \infty} x(t) = 0$$

Furthermore, $\hat{a}(t)$ converges to a constant (not to a!)

TU/e
NWO

Adaptive control

System dynamics

$$\dot{x} = ax + u \qquad a \text{ unknown parameter}$$

Controller

$$\begin{aligned} u &= -\hat{a}x - kx & \qquad k > 0 \\ \dot{\hat{a}} &= \gamma x^2 & \qquad \gamma > 0 \end{aligned}$$

Result

$$\lim_{t\to\infty} x(t) = 0$$

Furthermore, $\hat{a}(t)$ converges to a constant (not to a!)