# Decentralized Control

**Sanjay Lall**
**Stanford University**

Connections II
Caltech, August 17, 2006

# Terrestrial Planet Finder

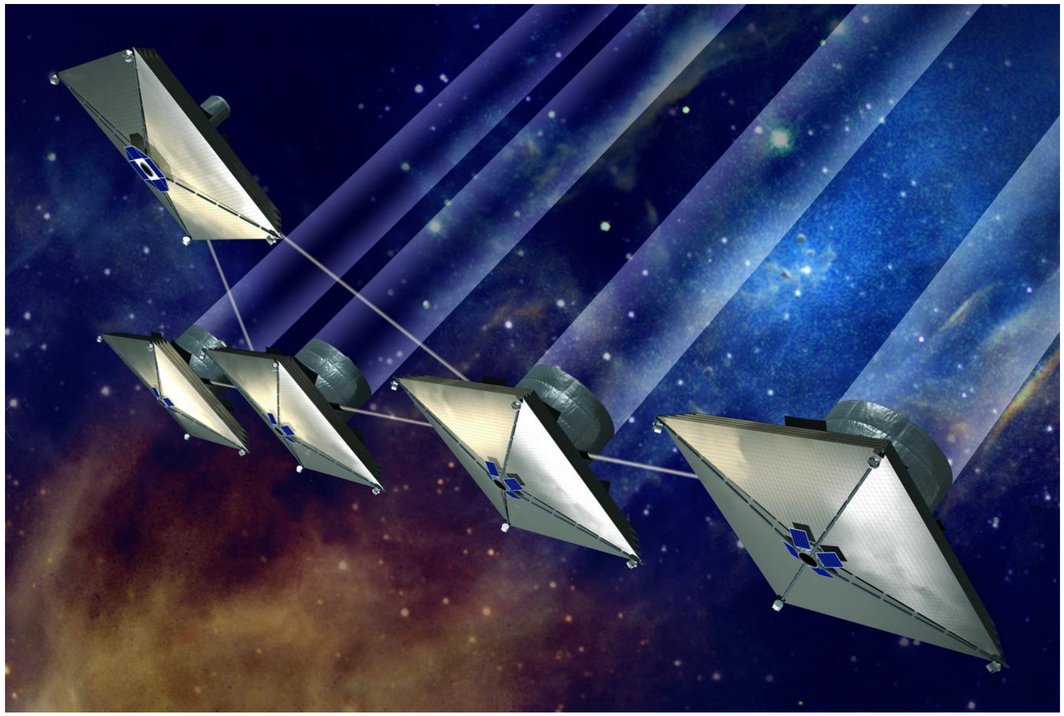

Image: Jet Propulsion Laboratory http://planetquest.jpl.nasa.gov/gallery/frequentImages.html

- Interferometry with infrared gathered by multiple telescopes
- Decentralized control to achieve precision formation flying

# Example: Team Decision

Minimize the cost

$$\mathrm{E}\big((u_1 + u_2 - x)^2\big)$$

- $x$ is unknown and random
- Player $1$ measures $y_1 = A_1 x + w_1$ and chooses $u_1$
- Similarly for player $2$

# Example: Team Decision

The cost is

$$\int \big(\gamma_1(y_1) + \gamma_2(y_2) - x\big)^2 \, d\mu$$

- Convex in $\gamma$. Pair of KKT conditions, one is

$$\gamma_1(y_1) = \mathrm{E}(x \,|\, y_1) - \mathrm{E}\big(\gamma_2(y_2) \,|\, y_1\big)$$

- Called *second guessing*

# Quadratic case

The cost is

$$\int c\big(\gamma_1(y_1), \gamma_2(y_2), x\big)\, d\mu$$

- $c$ is homogeneous positive quadratic, measure is Gaussian
- Optimal cost achieved by *linear* functions $\gamma_i$
- Linearity reduces cost to quadratic in coefficients of $\gamma_i$

# Quadratic case

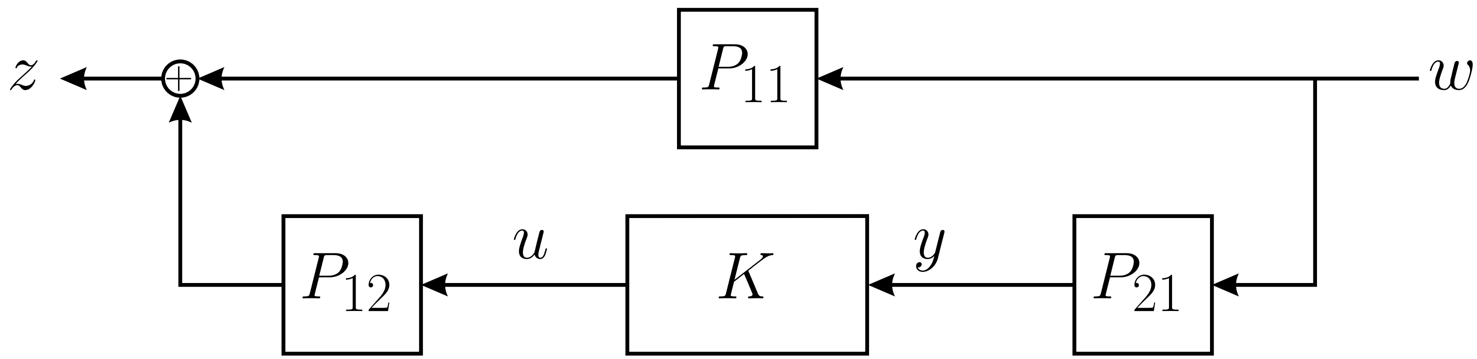


- $z = \big(P_{11} + P_{12} K P_{21}\big) w$
- Subject to sparsity constraints on $K$, minimize

$$\mathrm{E}\|z\|^2 = \|P_{11} + P_{12} K P_{21}\|_2^2$$

# Finite State Systems

with Randy Cogill, Stanford

# The Team Decision Problem

The team decision problem is to pick $\gamma_1$ and $\gamma_2$ to maximize

$$\sum_{y_1,y_2} r\big(y_1, y_2, \gamma_1(y_1), \gamma_2(y_2)\big)$$

- Players measure $y_1, y_2$ and take actions $u_1, u_2$.
- $r(y_1, y_2, u_1, u_2)$ is the reward
- $Y_i$ and $U_i$ are finite sets
- $\gamma_i : Y_i \to U_i$ is the policy for player $i$

# The Team Decision Problem

The team decision problem is to pick $\gamma_1$ and $\gamma_2$ to maximize

$$\sum_{y_1,y_2} r\big(y_1, y_2, \gamma_1(y_1), \gamma_2(y_2)\big)$$

- NP-hard, therefore
  - Is this problem *approximable?*
  - Can we exactly solve *restricted instances?*
- *Person-by-person* optimal solutions may be arbitrarily bad.

# Team Decision Problems

The team decision problem is to pick $\gamma_1$ and $\gamma_2$ to maximize

$$J(\gamma_1, \gamma_2) = \sum_{y_1, y_2} r\big(y_1, y_2, \gamma_1(y_1), \gamma_2(y_2)\big)$$

- Is this problem *$\varepsilon$-approximable*, i.e., can we efficiently find $\gamma_i$ so that

$$J_{\text{opt}} \leq \varepsilon J(\gamma_1, \gamma_2)$$

- Clearly $\varepsilon > 1$; we find an algorithm with

$$\varepsilon = \min\{|U_1|, |U_2|\}$$

# Approximation Algorithm

There is an algorithm that achieves $\varepsilon = |\mathcal{U}_2|$ for

$$J_{\text{opt}} \leq \varepsilon J(\gamma_1, \gamma_2)$$

The algorithm: choose the strategy $\gamma_1^{\text{approx}}$ for *player one* to maximize

$$\sum_{u_2, y_1, y_2} r(\gamma_1(y_1), u_2, y_1, y_2)$$

Then choose $\gamma_2$ to maximize

$$\sum_{y_1, y_2} r\big(\gamma_1^{\text{approx}}(y_1), \gamma_2(y_2), y_1, y_2\big)$$

# Multimodular Reward

A function $f : \mathbb{Z}^n \to \mathbb{R}$ is called *multimodular* if

there exists a concave function $g$ such that

$$f(x) = g(x) \text{ for all } x \in \mathbb{Z}^n$$

- Remains NP-hard, even with multimodular reward and two players.
- We have an algorithm which achieves approximation ratio $\varepsilon = 3$.
- Any number of actions and measurements.

# Summary: Finite State Systems

The team decision problem is approximable

- Approximation ratio $\varepsilon = \min\{|U_1|, |U_2|\}$
- For multimodular rewards, $\varepsilon = 3$
- Best possible $\varepsilon$ is unknown
- Decentralization adds *bilinear constraint* to usual LP
- Duality theory for these problems

# Feedback

- Cost function

$$\mathrm{E}\Big(k^2(u_1 - x)^2 + (u_1 - u_2)^2\Big)$$

- Measurement $2$ depends on action of player $1$

$$y_1 = x \qquad y_2 = u_1 + w_2$$

Nonlinear $\gamma_i$ perform better than any linear $\gamma_i$

New result by Rotkowitz: worst-case problem has a linear optimal solution

# No Feedback

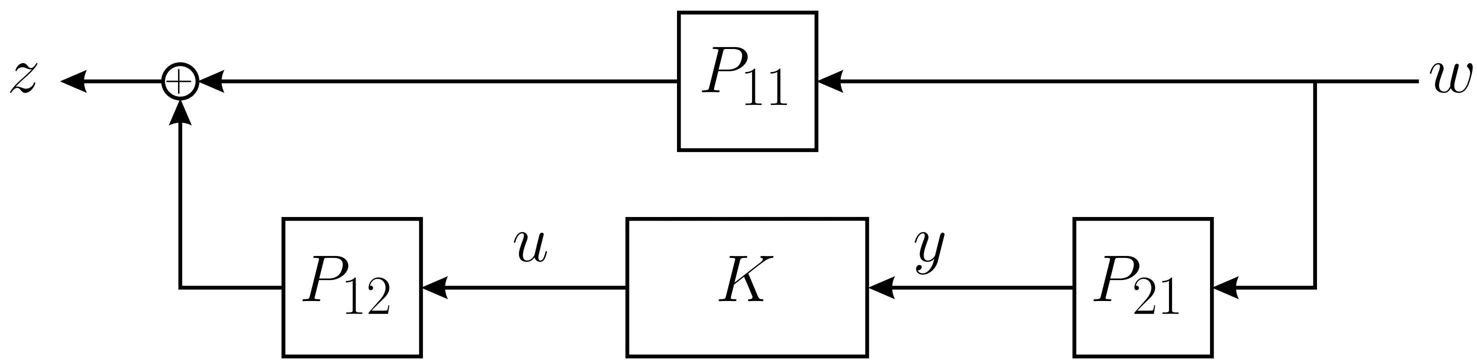

# Feedback

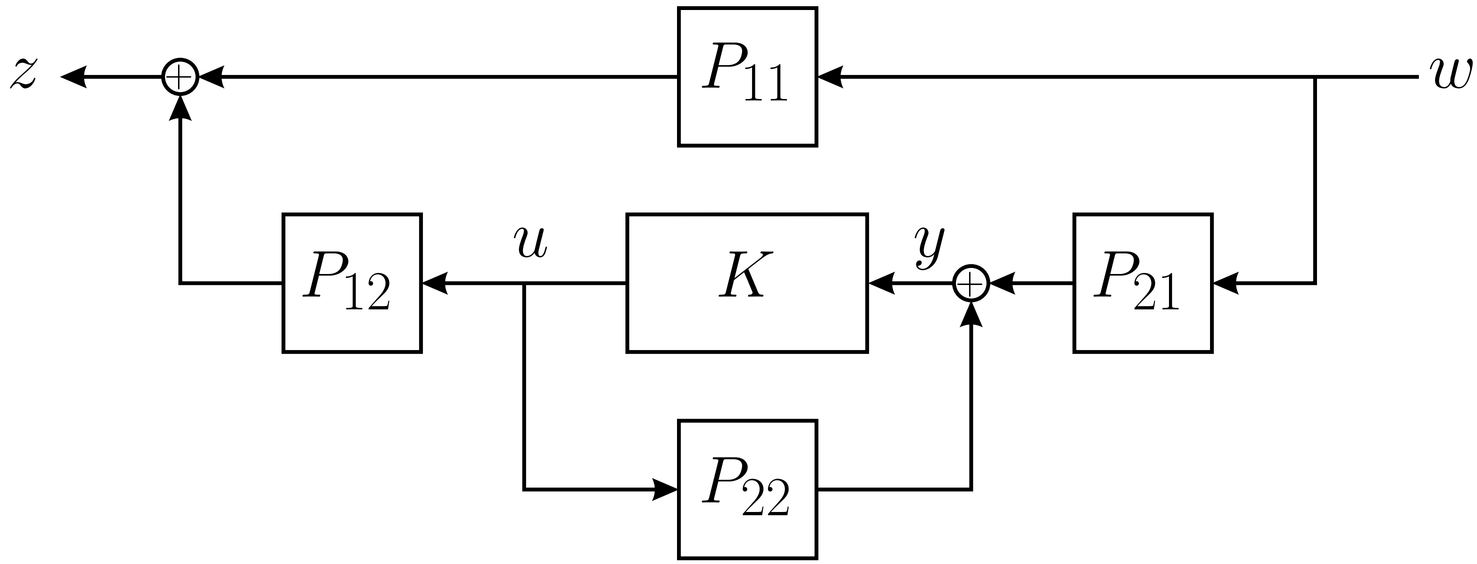


- $z = \big(P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21}\big)w$
- Subject to sparsity constraints on $K$, minimize

$$\mathrm{E}\|z\|^2 = \|P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21}\|_2^2$$

# Witsenhausen Counterexample

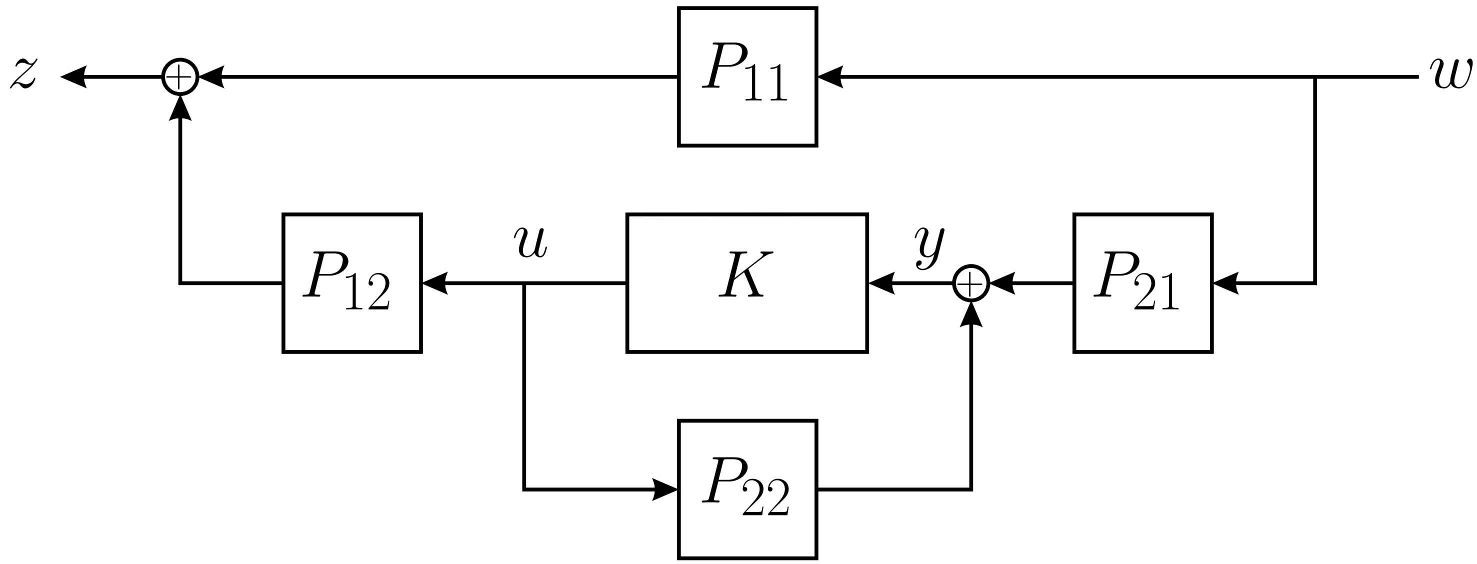


$$P_{11} = \begin{bmatrix} -k & 0 \\ 0 & 0 \end{bmatrix} \qquad P_{12} = \begin{bmatrix} k & 0 \\ 1 & -1 \end{bmatrix} \qquad P_{21} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad P_{22} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

# Linear Decentralized Control

Michael Rotkowitz, ANU

# Perfectly Decentralized Control

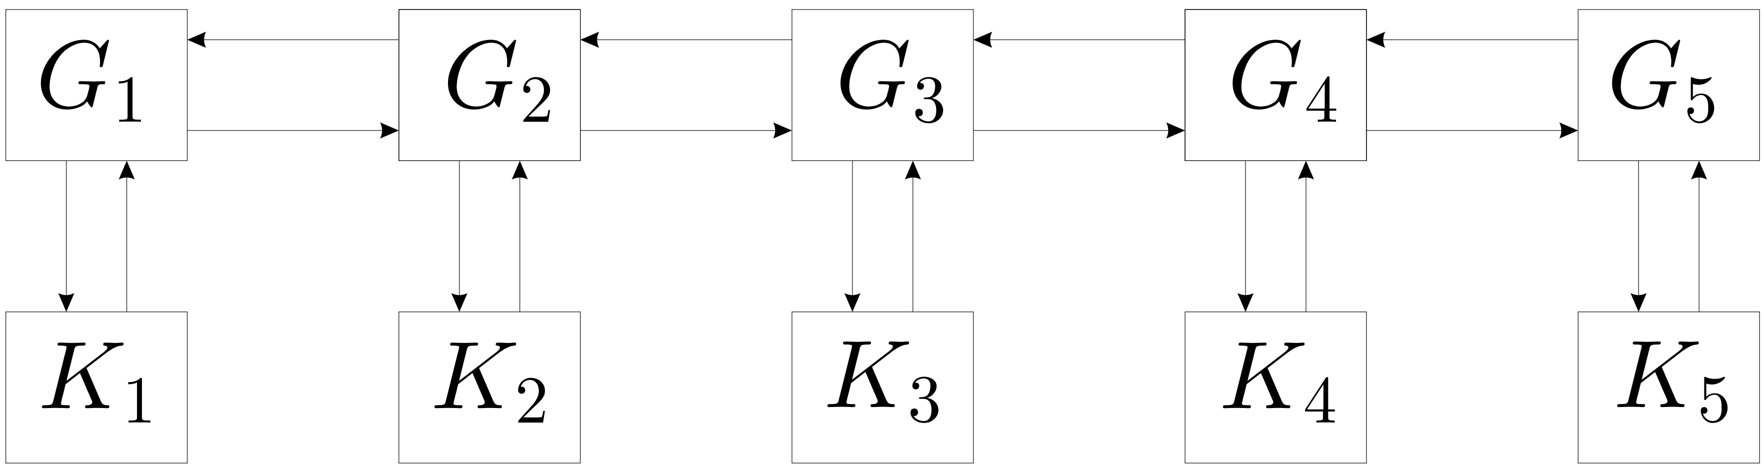


Control design problem is to find $K$ which is block diagonal, whose diagonal blocks are 5 separate controllers.

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \end{bmatrix} = \begin{bmatrix} K_1 & 0 & 0 & 0 & 0 \\ 0 & K_2 & 0 & 0 & 0 \\ 0 & 0 & K_3 & 0 & 0 \\ 0 & 0 & 0 & K_4 & 0 \\ 0 & 0 & 0 & 0 & K_5 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \end{bmatrix}$$

# Communicating Controllers

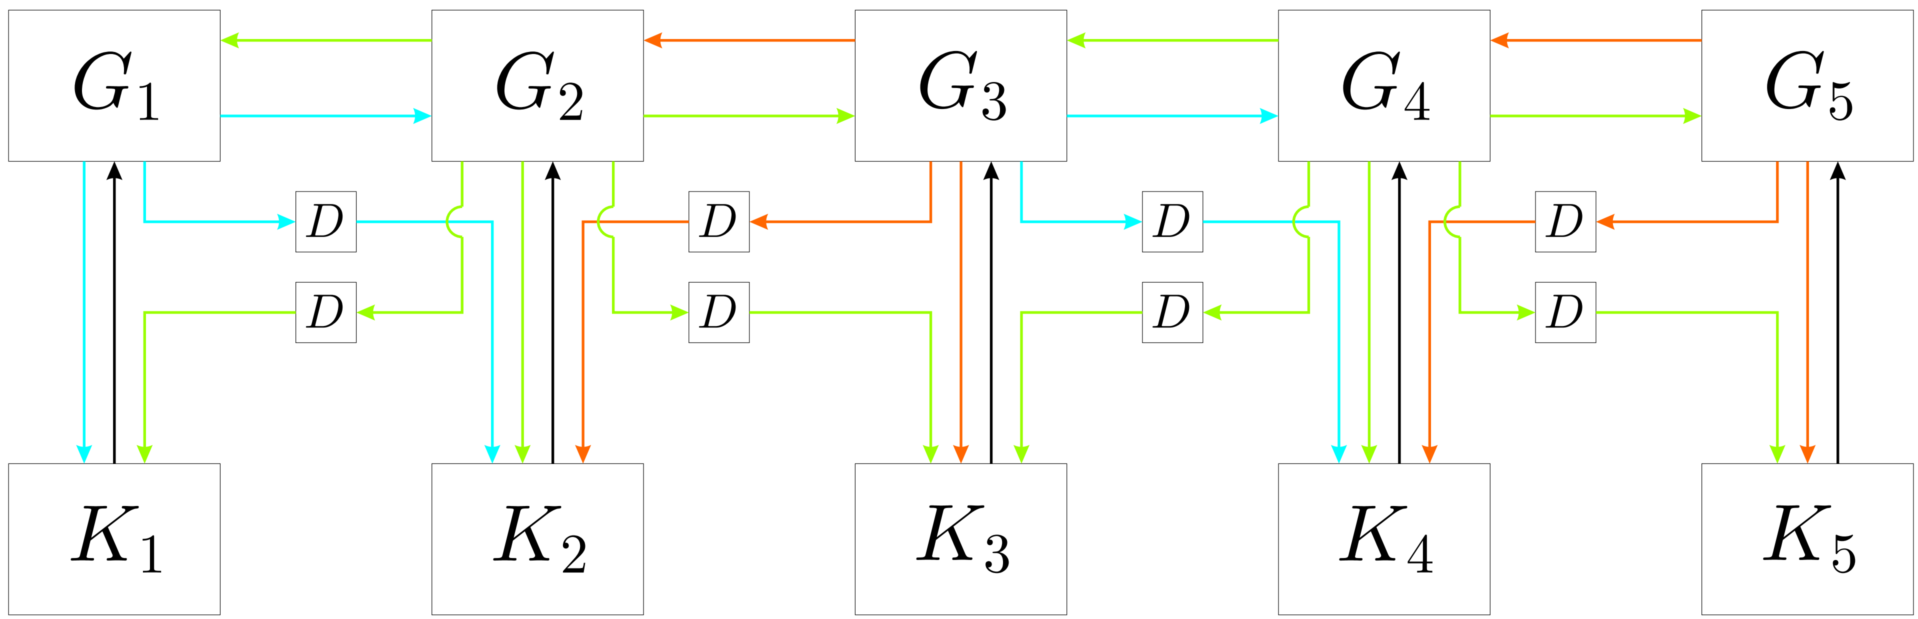


Control design problem is to find $K$ of the form

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \end{bmatrix} = \begin{bmatrix} K_{11} & DK_{12} & 0 & 0 & 0 \\ DK_{21} & K_{22} & DK_{23} & 0 & 0 \\ 0 & DK_{32} & K_{33} & DK_{34} & 0 \\ 0 & 0 & DK_{43} & K_{44} & DK_{45} \\ 0 & 0 & 0 & DK_{54} & K_{55} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \end{bmatrix}$$

# Example: TCP Congestion Control

- Feedback model of TCP over multiple links.
- Doyle, Kelly, Low, Paganini, Srikant.
- Decentralized control design.
- Can be used to analyze the effects of heterogeneous TCP variants.
- Figure: Paganini et al, 2000.

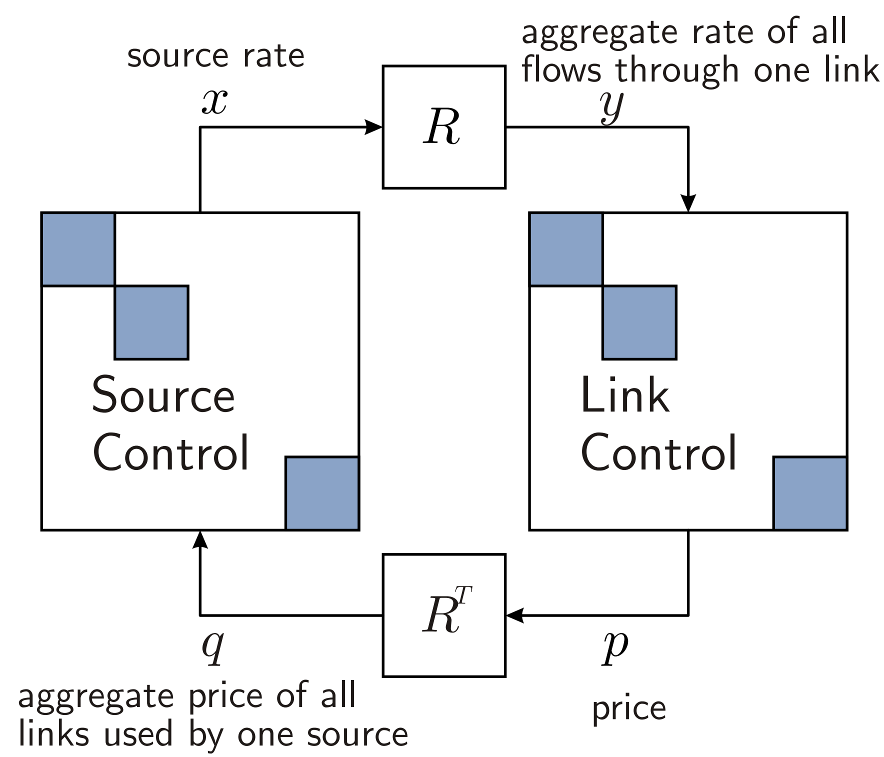

# Sparsity Constraints

We'd like to find $K$ satisfying sparsity constraints (and others)

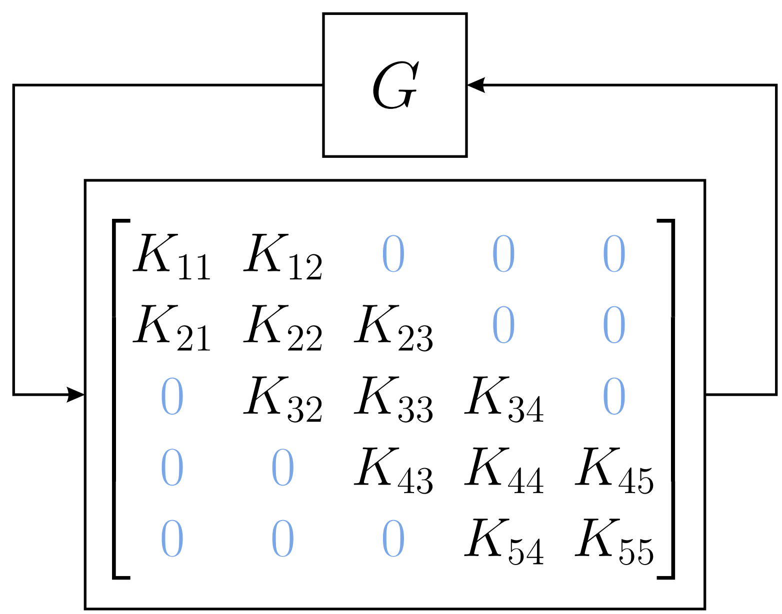


- *Quadratic invariance* property characterizes solvable problems
- Radner, Witsenhausen, Ho and Chu, Fan, Speyer and Jaensch, Voulgaris, Bamieh, Tsitsiklis, Blondel, Rantzer

# Problem Formulation

The set of $K$ with a given decentralization constraint is a subspace $S$, called the *information constraint*.

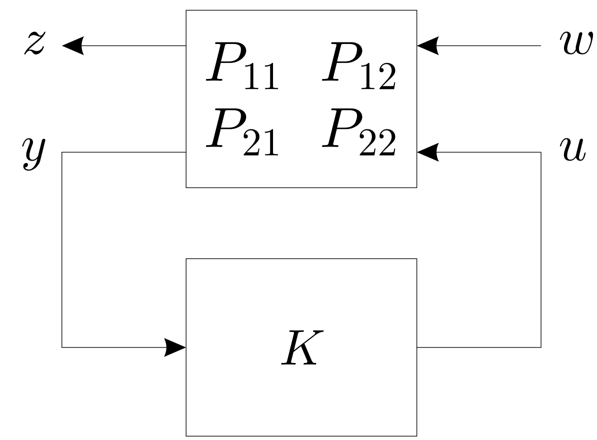


We would like to solve

$$\begin{aligned} \text{minimize} \quad & \|P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21}\| \\ \text{subject to} \quad & K \text{ is stabilizing} \\ & K \in S \end{aligned}$$

Many physical interpretations. Easy without the constraint that $K \in S$

# Information Constrained Problems

Use the change of variables (when $P$ is stable)

$$R = -K(I - P_{22}K)^{-1}$$

gives equivalent problem

$$\begin{array}{ll} \text{minimize} & \|P_{11} - P_{12}RP_{21}\| \\ \text{subject to} & R \text{ is stable} \\ & R(I - P_{22}R)^{-1} \in S \end{array}$$

- Subspace constraint on $K$ is not convex in $R$
- If $R$ is stable, then $K$ is stabilizing.

# Quadratic Invariance

Suppose $S$ is a subspace. Given a linear system $G$, the subspace $S$ is called *quadratically invariant* under $G$ if

$$KGK \in S \qquad \text{for all } K \in S$$

## Main Result

The subspace $S$ is quadratically invariant under $G$ if and only if

$$K \in S \iff K(I - GK)^{-1} \in S$$

# Theorem

- $\mathcal{U}$ and $\mathcal{Y}$ are Banach spaces
- $G : \mathcal{U} \to \mathcal{Y}$ is compact
- $S \subset L(\mathcal{Y}, \mathcal{U})$ is a closed subspace
- Let $M = \{ K \in L(\mathcal{Y}, \mathcal{U}) \; ; \; (I - GK) \text{ is invertible} \}$
- Let $h(K) = -K(I - GK)^{-1}$

Then the subspace $S$ is quadratically invariant under $G$ if and only if

$$h(S \cap M) = S \cap M$$

# Quadratic Invariance

If $S$ is quadratically invariant, we find the optimal controller by solving

$$\begin{array}{ll} \text{minimize} & \|P_{11} - P_{12}RP_{21}\| \\ \text{subject to} & R \text{ is stable} \\ & R \in S \end{array}$$

- Then optimal controller is $K = -R(I - P_{22}R)^{-1}$
- Holds for continuous-time and discrete-time systems.
- New results cover unstable systems, nonlinear systems.

# Proof Outline

If $S$ is quadratically invariant, then for all $K$ in $S$ and all $n$ we have

$$K(GK)^n \in S$$

because

$$\begin{aligned} 2K(GK)^{n+1} = \big(K + K(GK)^n\big)G\big(K + K(GK)^n\big) \\ - KGK - K(GK)^n GK(GK)^n \end{aligned}$$

# Proof Outline

For $\lambda \in \mathbb{C}$ sufficiently large

$$K(I - GK/\lambda)^{-1} = K + \frac{KGK}{\lambda} + \frac{KGKGK}{\lambda^2} + \dots$$

since $K(GK)^n \in S$ we have that

$$K(\lambda I - GK)^{-1} \in S$$

Now use analyticity of resolvent to show this holds for $\lambda$ in the unbounded connected component of the resolvent set.

Compactness of $G$ gives a countable spectrum

# Example: Advance Warning

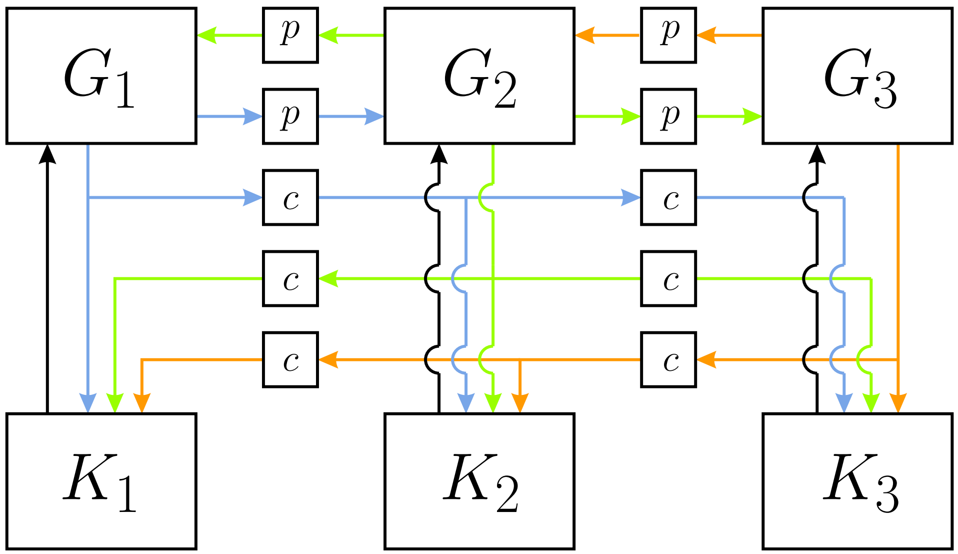


- Quadratic invariance if

$$c \leq p$$

- For vehicle formations

$$1/p \;\propto\; \text{speed of sound}$$

- $p$ is the *propagation delay*

- $c$ is the *communication delay*

- Similar results in stochastic team theory.

# Quadratically Invariant Constraints

The following problems have quadratically invariant constraints

- Strings and arrays
- Group symmetries
- Hermitian symmetries
- One-step-delayed problem
- Funnel-causality (Bamieh and Voulgaris, 2002)
- Partially-nested (Ho and Chu, 1972)

# Control over Networks

We have a network of plants connected by delays

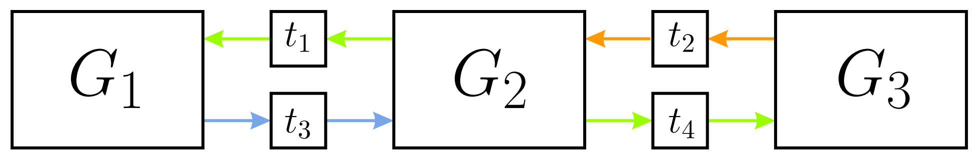


and an associated weighted directed graph

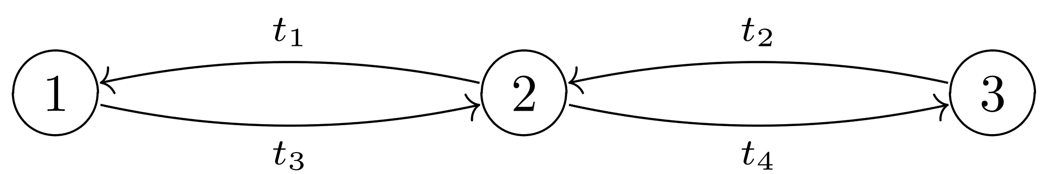


- $p_{ij}$ is the shortest path length from plant $i$ to plant $j$.
- $c_{ij}$ is the delay before controller $i$ receives sensor data from plant $j$
- Quadratic invariance if $c_{ij}$ satisfy triangle inequality and $c_{ij} \leq p_{ij}$

# Two-Dimensional Lattice

- Controllers communicate on lattice
- Dynamics propagate in space
- Delay proportional to distance
- Quadratically invariant iff

$$c \leq \frac{p}{\sqrt{2}}$$

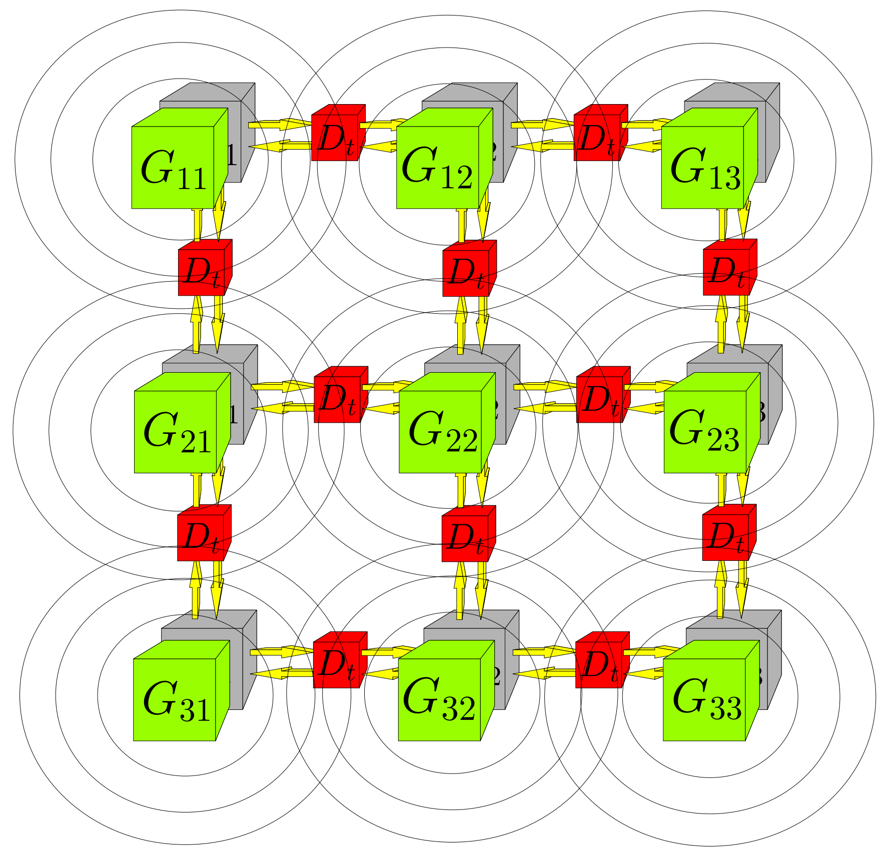

# Unstable Plants

For unstable $P$

- Guess a stable stabilizing controller $K_{\text{nom}} \in S$.
- Use the change of variables

$$K = K_{\text{nom}} + Q(I - GK_{\text{nom}} + GQ)^{-1}(I - GK_{\text{nom}})$$

This results in the equivalent convex problem

$$\begin{array}{ll} \text{minimize} & \|T_1 - T_2 Q T_3\| \\ \text{subject to} & Q \text{ is stable} \\ & Q \in S \end{array}$$

# Summary: Linear Decentralized Control

- Optimal-norm synthesis subject to *quadratically invariant information constraints* is a convex optimization problem.

- Effective computation of $H_2$ and $H_\infty$ optimal controllers

- Many things remain unknown
  - Estimation structure of controllers
  - Dynamic programming approach
  - State-space formulae, minimal order

# Approximation Bounds for Markov Decision Processes

with Randy Cogill, Stanford

# Markov Decision Processes

Given policy $\mu$, find the smallest approximation ratio $\varepsilon > 1$ such that

$$J(\mu) \leq \varepsilon J_{\text{opt}}$$

- States $x_t$ and actions $u_t$, policies $u_t = \mu(x_t)$
- The *average per-period cost* is

$$J(\mu) = \lim_{t \to \infty} \frac{1}{t+1} \sum_{s=0}^{t} \mathrm{E}\big(r(x_s, \mu(x_s))\big)$$

- NP-hard, so look at *heuristic policies* instead of optimal policies

# Approximation Ratio

For a specific policy $\mu$,

- We can compute an upper bound $\alpha$ on $J(\mu)$
- We can compute a lower bound $\beta$ on $J_{\text{opt}}$
- Then we have

$$J(\mu) \leq \frac{\alpha}{\beta} J_{\text{opt}}$$

# Upper Bounds

Markov process with cost

$$J = \lim_{t\to\infty} \frac{1}{t+1} \sum_{s=0}^{t} \mathrm{E}\big(r(x_s)\big)$$

For finite-state, irreducible, aperiodic systems, given any function $h$

$$J \le \sup_z \Big( r(z) + \mathrm{E}\big(h(x_{t+1}) \,|\, x_t = z\big) - h(z) \Big)$$

- Due to Odoni, 1969. Based on weak duality.
- Extensive theory: subharmonic fns, maximum principle, Lyapunov
- Problem in applications: general state spaces require bounded $h$

# General Result

Theorem: suppose $x_0, x_1, \ldots$ is Markov. If there is an $\varepsilon > 0$ such that

$$\sup_z \Big( \mathrm{E}\big( |h(x_{t+1})|^{1+\varepsilon} \,\big|\, x_t = z \big) - |h(z)|^{1+\varepsilon} \Big) < \infty$$

then

$$J \le \sup_z \Big( r(z) + \mathrm{E}\big( h(x_{t+1}) \,\big|\, x_t = z \big) - h(z) \Big)$$

- Possibly uncountable state space
- Bounded $h$ is a special case, usually inapplicable since $r$ unbounded
- Related to Foster's criterion, but gives lower and upper bounds

# Optimal Achievable Cost

Theorem: The following gives a lower bound on $J_{\text{opt}}$

$$J(\mu) \geq \inf_{z,w}\Big(r(z,w) + \mathrm{E}\big(h(x_{t+1}) \,|\, x_t = z, u_t = w\big) - h(z)\Big)$$

for all $\mu$ such that there is an $\varepsilon > 0$ such that

$$\sup_{z}\Big(\mathrm{E}\big(|h(x_{t+1})|^{1+\varepsilon} \,|\, x_t = z, u_t = \mu(z)\big) - |h(z)|^{1+\varepsilon}\Big) < \infty$$

- Possibly uncountable state space, finite action space
- Does not require solving Hamilton-Jacobi equation

# K-Server Problem

Many applications, e.g. vehicle tasks, disk head motion, caching, etc.

- $S$ is a finite set
- $x_i(t) \in S$ is the position of $i$'th server
- $r(t) \in S$ is the request at time $t$
- Control action $u(t) \in \{1, \dots, k\}$

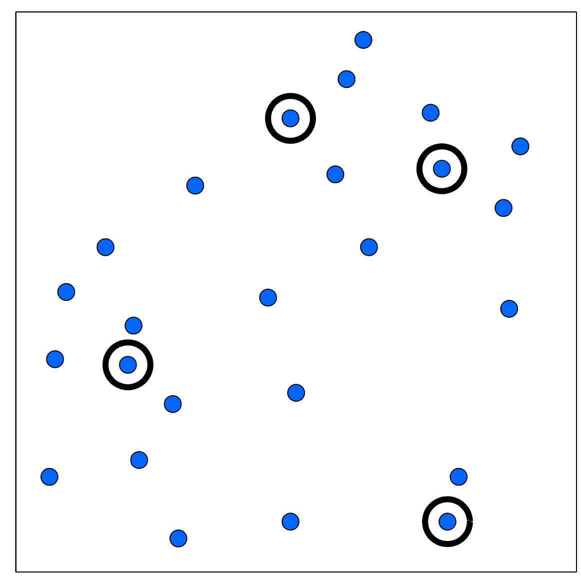

Action $u(t) = i$ leads to

- Server $i$ moves to $r(t)$
- cost incurred is $\|x_i(t) - r(t)\|$

# K-Server Problem

We construct a *decentralized policy* $u(t) = \mu(x(t), r)$ that achieves

$$J(\mu) \leq 2J_{\text{opt}}$$

- Policy is heuristic, since even centralized optimal control is intractable.
- $J_{\text{opt}} = \inf_{\nu} J(\nu)$ is optimal cost achievable by *any* policy
- Decentralized policy performs well compared to centralized policies
- DTRP of Bertsimas and van Ryzin, 1991, Frazzoli and Bullo, 2004

# K-Medians Problem

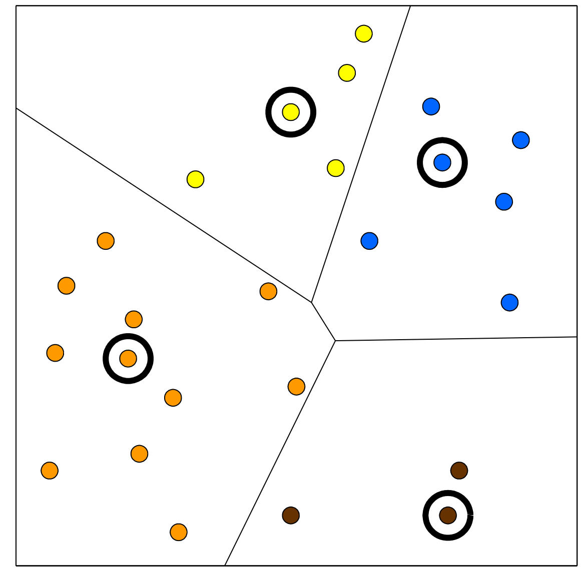

Pick $y_1, \ldots, y_k \in S$ to minimize

$$m(y) = \sum_{x \in S} p(x) \min_i \|y_i - x\|$$

- Let $m_{\text{opt}}$ be the minimum.
- NP-hard, even for metric case.
- This partitions the domain $S$ into $S_1, \ldots, S_k$.
- Simple decentralized policy: server $i$ responds if request falls in $S_i$

# Bounds

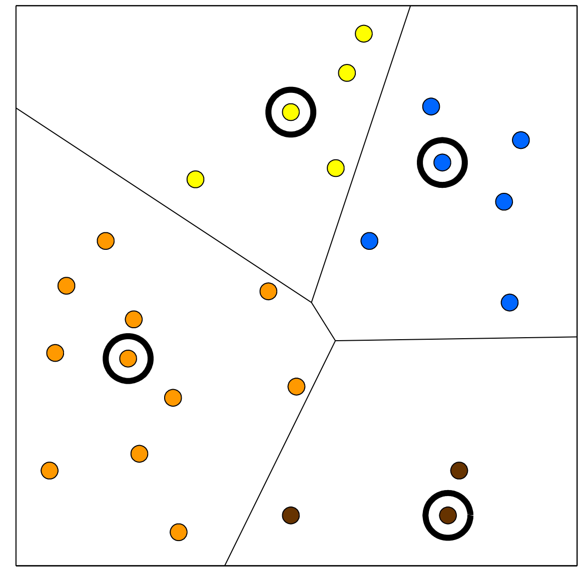

- *Lower bound* on achievable cost

$$J_{\text{opt}} \geq m_{\text{opt}}$$

Certified by $h(x,r) = \min_i \|x_i - r\|$

- *Upper bound:* the above decentralized policy $\mu$ achieves

$$J(\mu) \leq 2m_{\text{opt}}$$

Certified by function

$$h(x,r) = \sum_i \|x_i - y_i\| + 2\min_j \|y_j - r\|$$

# Summary: Approximation Bounds

- *Approximation bounds* for general Markov decision processes.
  - Intractability leads us to look at heuristic policies.
  - Upper bounds on performance of a given policy
  - Lower bounds on performance achievable by any policy
- Applications
  - Variants of $k$-*server problem*
  - Average backlog for *CIOQ switches*
  - Model reduction of *particle dynamics*

# Approximate Dynamic Programming for Decentralized Control of MDPs

with Mike Rotkowitz, Randy Cogill, Ben Van Roy

# Decentralized Control of MDPs

We have a Markov decision process, with states $x_0, x_1, \ldots$ and actions $u_0, u_1, \ldots$, and cost

$$\lim_{t\to\infty} \mathrm{E}\left( \sum_{t=0}^{\infty} \alpha^t g(x_t, u_t) \,\middle|\, x_0 = z \right)$$

- Approximate dynamic programming approach
- Linear programming to compute approximate $Q$ function
- Does not require knowing the optimal $Q$ or value functions

# Decentralized Control of MDPs

This *cost-to-go* function satisfies the *Bellman equation*

$$J^*(x) = \min_u \left( c(x,u) + \alpha \sum_{y \in \mathcal{X}} p(y \,|\, x,u) J^*(y) \right)$$

Optimal controller chooses actions to minimize the $Q-$*function*

$$Q^*(x,u) = c(x,u) + \alpha \sum_{y \in \mathcal{X}} p(y \,|\, x,u) J^*(y)$$

Decentralized controller results from minimizing a *structured* $Q$ function

$$Q(x,u) = \sum_{i=1}^{k} Q_i(x_i, u_i)$$

# Optimal Value Functions

We expect $Q$ to approximate $Q^*$, for example

$$\begin{aligned} \text{minimize} \quad & \|Q - Q^*\|_1 \\ \text{subject to} \quad & Q(x,u) \leq c(x,u) + \alpha \sum_y p(y \,|\, x,u) \Big( \min_w Q(y,w) \Big) \\ & Q \in S \end{aligned}$$

- We can find the optimal $Q$ *without* knowing $Q^*$
- Approximation ratio

$$\|J^\mu - J^*\|_\nu \leq \frac{1}{1-\alpha} \|Q^*_\mu - Q_\mu\|_\omega$$

# Summary: Approximate DP

- Approximate DP gives a *heuristic* for decentralized control

- Applies to finite state MDPs

- Cheaper to compute than centralized optimal controller

- Simple a-posteriori error bounds

# Conclusions

- Linear dynamical systems

  *Quadratic invariance* leads to convex optimal control problems

- Team decision problems

  They are *approximable*, and we have new approximation algorithms

- Markov decision processes

  A new approach for computing *approximation bounds*

- Approximate dynamic programming

  An efficient approach for *synthesis* for decentralized control of MDPs