On real-time pricing for strategic agents

Cédric Langbort

Department of Aerospace Engineering & Coordinated Science Laboratory
University of Illinois at Urbana-Champaign

Workshop on "Multi-Agent Coordination and Estimation"

Lund Center for Control of Complex Engineering Systems
February 2nd, 2010

What is real-time pricing?

Real-time pricing was proposed for power networks by F. Schweppe et al. in the 80's.

- Goal: induce desired inputs from power plants indirectly.
- Different from spot-pricing: accounts for dynamics of subsystems.

What is real-time pricing?

Real-time pricing was proposed for power networks by F. Schweppe et al. in the 80's.

- Goal: induce desired inputs from power plants indirectly.
- Different from spot-pricing: accounts for dynamics of subsystems.

But, really:

- A general method for controlling dynamical systems using incentives.
- Can account for issues of *segmentation of model information*, *distributed control authority*, *possible selfishness*...

What is real-time pricing?

Much progress since then, e.g.,

- Incorporating dynamics is now common: [Jokic, Lazar & Vanden Hof, '07] (through linear complementarity controllers), [Ilic]...
- Well-developed theory of incentive control in dynamical systems: [Ho et. al], [Cruz], [Başar & Zheng]...
- Application of Mechanism Design to power systems: [Silva et al. '01], [R. Wilson]...

So, why revisit now?

“A case of rock-paper-scissors”...

- Mechanism design-based approach have mostly neglected dynamics so far.
- Optimal control-based dynamic pricing schemes often do not account for possible strategic behavior of subsystems (cf. Schweppe’s own work).
- Incentive control approaches may assume too much common knowledge...

This talk

...makes a simple proposal to try and build on strengths of all three, namely

- pick a simple enough model of individual rationality to enable implementation with no unnecessary shared knowledge assumption,
- use Mechanism Design ideas to systematically construct incentives, and
- take dynamics and causality aspects seriously (the real contribution of control?...)

Original example

(Linearized) load frequency control [Berger & Schweppe, '82]

- Utility wants to find control laws $\{u_1(t)\}_{t=0}^{T-1}$, $\{u_2(t)\}_{t=0}^{T-1}$ to solve problem ($\mathcal{P}$) below:

$$\begin{aligned} \min \ & \frac{1}{2}\sum_{t=0}^{T-1}\left(\sum_{i=1}^{2}\|x_i(t)\|_{Q_i}^2 + \|u_i(t)\|_{R_i}^2 + \|z(t+1)\|_Q^2\right) \\ \text{subject to} \ & x_i(t+1) = A_i x_i(t) + B_i u_i(t) \\ & x_i(0) = \bar{x}_i \ \forall i \\ & z(t+1) = z(t) + M_1 x_1(t) + M_2 x_2(t); \ z(0) = \bar{z} \\ & \|u_i(t)\| \leq 1 \ \forall i, t \end{aligned}$$

- x_i: generator i's internal variable, z: frequency deviation away from 60Hz.

Original example – c'ed

Pictorially:

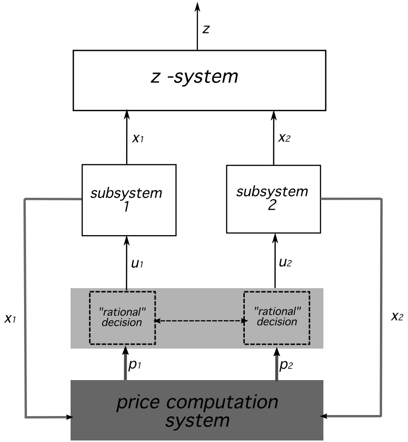

Original example – c'ed

- Subsystems $x_i(t+1) = A_i x_i(t) + B_i u_i(t)$ and cost function parameters Q_i and R_i are known to individual subsystems only.
- They have no *a priori* reason to report them truthfully (cf. Enron, 'ramp constraint gaming': [R. Wilson, Econometrica '02], [Oren]...)
- Power plant i chooses $\{u_i(t)\}_{t=0}^{T-1}$ selfishly, so as to minimize

$$\sum_{t=0}^{T-1} \underbrace{\|x_i(t)\|_{Q_i}^2 + \|u_i(t)\|_{R_i}^2}_{\text{individual cost}} + \underbrace{\pi_i(t)}_{\text{utility payments}}$$

Central question

Can the utility compute payments $\{\pi_i(t)\}_{t=0}^{T-1}$ – with the information available to it – which induce plants to use welfare-optimizing control strategies?

A general formulation

Welfare problem:

$$\begin{aligned}(u_1^\star, u_2^\star) = \arg\min \;& J_1(u_1) + J_2(u_2) + J(z) \\ \text{subject to } \;& z = \underbrace{H_1(u_1) + H_2(u_2)}_{H(u_1,u_2)},\; u_i \in \mathcal{U}_i, \forall i\end{aligned}$$

where $u_i := \{u_i(t)\}_{t=0}^{T-1}$, $\mathcal{U}_i$ compact convex

Subsystems:

- Pick $\bar{u}_i$ which minimizes net cost $J_i(u_i) + \pi_i(u_i)$,*where*π_i is the payment made by the utility,
- know J_i privately.

Problem:

Determine payment strategies $\{\pi_i\}_{i=1}^2$ such that $\bar{u}_i = u_i^\star$ *regardless of* $\{J_i, \mathcal{U}_i\}_{i=1}^2$.

Two traditional approaches... and their inadequacy

From KKT conditions:

$$\pi_i^\star(u_i) = [\nabla_{u_i}(J \circ H)(u_1^\star, u_2^\star)]\, u_i \text{ for all } u_i \in \mathcal{U}_i$$

induces $\bar{u}_i = u_i^\star$.

From dual decomposition:

$$\pi_i^{dual}(u_i) = {p_i^\star}^T H_i(u_i) \text{ for all } u_i \in \mathcal{U}_i,$$

where $p_i^\star$ is the Lagrange multiplier of coupling constraint at optimality, also induces $\bar{u}_i = u_i^\star$.

Two traditional approaches... and their inadequacy

Both price functions require that full information about $\{J_i, \mathcal{U}_i\}_{i=1}^2$ *be revealed to utility* since $u^\star$ must be computed.

- The first approach coincides with payments derived using incentive control/ Stackelberg games techniques.
- The second payment is the one originally proposed by Schweppe et al.

Mechanism design approach

Typical setup [Vickrey-Clarke-Groves]:

$$\min_d \sum_{i=1}^{n} v_i(\theta_i, d) \tag{1}$$

- θ_i is privately known true type
- social decision d depends on reported types: $d = d(\bar{\theta}_i, \bar{\theta}_{-i})$
- agent i reports type $\bar{\theta}_i$ such that

$$v_i(\theta_i, d(\bar{\theta}_i, \theta_{-i})) + t(\bar{\theta}_i, \theta_{-i}) < v_i(\theta_i, d(\theta_i', \theta_{-i})) + t(\theta_i', \theta_{-i})$$

  for all $\theta_i', \ \theta_{-i}$ (dominant strategy)
- VCG mechanism constructs payments t such that $\bar{\theta}_i = \theta_i$.

Mechanism design approach –c'ed

- VCG payment is of the form:

$$t(\theta_i, \theta_{-i}) = \sum_{-i} v_{-i}(\theta_{-i}, d_{opt}(\theta)) + F_i(\theta_{-i})$$

  for some F_i. A good choice of F_i leads to reinterpretation as player's marginal contribution to the optimal welfare...

- It uses optimal decision map $d_{opt}(\{\bar{\theta}_i\}_{i=1}^n)$ which solves (1) for given $\{\bar{\theta}_i\}_{i=1}^n$ BUT

- It incentivizes truth-telling without a priori knowledge of what the truth is (as opposed to "incentive control")!

Back to general formulation

Our problem differs from this setup in two ways:

- Type is (A_i, B_i, Q_i, R_i) or J_i: complicated and not directly price-able. Must price u_i or $M_i x_i$ instead...
- Optimal decision map is not available: cannot use dominant strategy implementation

Idea:

- Use an *indirect* mechanism with $M_i x_i$ or u_i as "messages"
- Implement in Nash equilibrium, using payments depending on *both* inputs and assuming

$$J_i(\bar{u}_i, \bar{u}_{-i}) \leq J_i(u_i, \bar{u}_{-i}) \;\; \forall u_i \in \mathcal{U}_i.$$

- decompose payments over time

Main result

Theorem:
Price functions $\{\pi_i\}_{i=1}^2$ are smooth and implement the optimal decisions $(u_1^\star, u_2^\star)$ in Nash equilibrium for any convex functions J_1 and J_2 if and only if there exist arbitrary smooth functions $\{F_i\}_{i=1}^2$ such that

$$\pi_i(u_i, u_{-i}) = [J \circ H](u_1, u_2) + F_i(u_{-i}) \tag{2}$$

for all i and all $(u_i, u_{-i}) \in int\mathcal{U}_i \times int\mathcal{U}_{-i}$.

This shows that the “Wonderful Life” utility is the only possible choice [Wolpert, Marden & Shamma...]

Price decomposition

Must rewrite a payment satisfying (2) as a sum of incremental causal payments:

$$\pi_i(u_i, u_{-i}) = \sum_{t=0}^{T-1} \pi_i^t(\{x_i(s)\}_{s=0}^t, \{x_{-i}(s)\}_{s=0}^t).$$

Going back to the original problem...

Possible solutions

Choice #1 (independent from N):

$$\pi_i^0(x_i(0), x_{-i}(0)) = \frac{1}{2}x_i(0)^T M_i^T Q M_i x_i(0) + x_{-i}(0)^T M_{-i}^T Q M_i x_i(0)$$
$$+ \bar{z}^T Q M_i x_i(0)$$

$$\pi_i^{t+1}(\{x_i(s)\}_{s=0}^t, \{x_{-i}(s)\}_{s=0}^t) = \pi_i^t + \bar{z}^T Q M_i x_i(t+1)$$
$$+ \frac{1}{2}x_i(t+1)^T M_i^T Q M_i x_i(t+1) + x_{-i}(t+1)^T M_{-i}^T Q M_i x_i(t+1)$$
$$+ \sum_{s \le t} \Big(x_i(s)^T M_i^T Q M_i x_i(t+1) + x_{-i}(t+1)^T M_{-i}^T Q M_i x_i(s)$$
$$+ x_{-i}(s)^T M_{-i}^T Q M_i x_i(t+1) \Big)$$

for all $0 < t \le T-1$.

Possible solutions

Choice #2 (N-dependent):

$$\tilde{\pi}_i^t(\{x_i(s)\}_{s=0}^t, \{x_{-i}(s)\}_{s=0}^t) =$$

$$(N-t)\left[\frac{1}{2}x_i(t)^T M_i^T Q M_i x_i(t) + x_{-i}(t)^T M_{-i}^T Q M_i x_i(t) + \bar{z}^T Q M_i x_i(t)\right.$$

$$+\sum_{s<t}\Big(x_i(s)^T M_i^T Q M_i x_i(t) + x_{-i}(t)^T M_{-i}^T Q M_i x_i(s)$$

$$\left.+x_{-i}(s)^T M_{-i}^T Q M_i x_i(t)\Big)\right]$$

for all $0 \leq t \leq T-1$

Remaining issues & future work

- Decisions $\bar{u}_i$ do not depend causally on the incremental payments.
- This is due to finite horizon problem formulation and the non-separability of cost-to-go.
- Different from recent work in Dynamic Mechanism Design (e.g., [Cavallo & Parkes '06-'08], [Bergemann & Valimäki, '06]), where *type* is time-varying and there is no coupling between type dynamics ("private dynamic utility") $\Rightarrow$ cost-to-go function is separable...

Remaining issues & future work

What to do with MPC?

- should plant return full $\{u_i(t)\}_{t=0}^{T-1}$ and utility pay at the end of horizon?
- how can utility pay for $u_i(0)$ only?

Beyond Nash implementation??