Distributed Algorithms for Wireless Networks

Peter W. Glynn

Stanford University

Joint work with Tim Holliday and Andrea Goldsmith

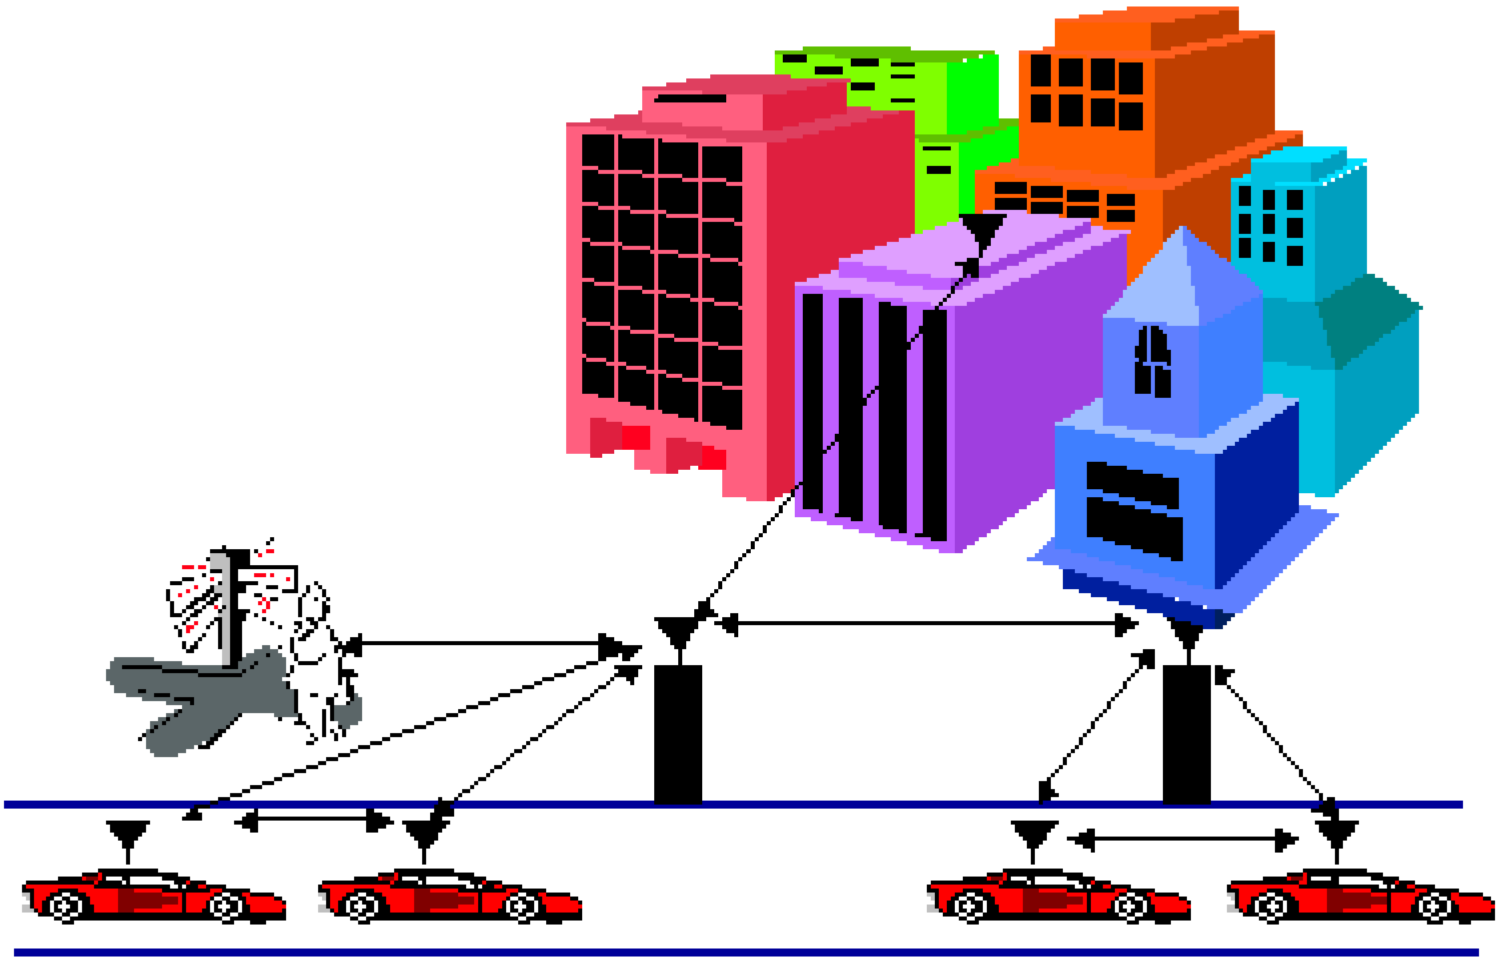

Our Setting:

- m users
- each user wishes to efficiently (e.g. low power) transmit packets in a timely way (e.g. possible deadlines)
- each user has a set of actions available to accomplish these objectives (e.g. power control, coding schemes, etc.)
- a given user's choice of action affects the other users (via interference)
- each user is statistically identical (and follows a common control law)

State of the System:

$X_j(n) =$ state of user j at time n
$\in S$ (finite)

$A_j(n) =$ action chosen by user j at time n
$\in \mathcal{A}$ (finite)

Possible state descriptors:

$$(X_1(n), \ldots, X_m(n)) \in S^m$$

$$\mu_n(\cdot) = \frac{1}{m} \sum_{j=1}^{m} \delta_{X_j(n)}(\cdot) \quad \in \mathcal{P}(S)$$

$$\mathcal{P}(S) = \{(p_1, \ldots, p_d) : \sum_{j=1}^{d} p_j = 1,\ p_j \geq 0,\ 1 \leq j \leq d\}$$

Reward Function

Reward earned at time n is

$$\frac{1}{m}\sum_{j=1}^{m} r\left(X_j(n), A_j(n)\right)$$

Transition Dynamics

$$\begin{aligned}
\mathsf{P}\left(X_1(n+1) = y_1, \ldots, X_m(n+1) = y_m \middle|\, \mathcal{F}_n\right) \\
&= \prod_{j=1}^{m} \mathsf{P}_{A_j(n)}\left(\mu_n\,;\, X_j(n),\, y_j\right) \\
&= f_n\left(\mu_n, M_{a,x}(n) : x \in S,\, a \in \mathcal{A}\right)
\end{aligned}$$

We will focus here on the case where we have
a "scalar interaction" effect (e.g. interference)

i.e. there exists $h : S \longrightarrow \mathbb{R}$ such that

$$\mathsf{P}_a\,(\mu\,;\, x,\, y) = \mathsf{P}_a\,(\mu h\,;\, x,\, y)$$

where

$$\mathsf{P}_a(\cdot) = (\mathsf{P}_a(\cdot\,;\, x,\, y) : x, y \in S)$$

is assumed continuous on $[\,\underline{h}\,, \overline{h}\,]$

$(\underline{h} = \min(h(x) : x \in S), \overline{h} = \max(h(x) : x \in S))$

Some of what we have done generalizes to more general interaction effects.

Centralized Controller

$$A_j(n) \in \mathcal{F}_n = \sigma\left(X_j(i), \mu_i, U_j(i) : 0 \leq i \leq n\right)$$

Decentralized Controller

$$A_j(n) \in \mathcal{G}_n = \sigma\left(X_j(i), U_j(i) : 0 \leq i \leq n\right)$$

Analysis of Distributed Controls

Conditional on μ_n,

$$m\mu_{n+1} \stackrel{\mathcal{D}}{=} \sum_{x \in S} \text{multinomial}\left(m\mu_n(x)\,;\ (\mathsf{P}\left(\mu_n h\,;\ x,\ y\right) y \in S)\right)$$

where

$$\mathsf{P}\left(c\,;\ x,\ y\right) = \sum_{a \in \mathcal{A}} \pi(a|x)\mathsf{P}_a\left(c\,;\ x,\ y\right)$$

i.e. each user in state x independently chooses action a with probability $\pi(a|x)$

One could also deterministically associate action a with a fraction $\pi(a|x)$ of the $m\mu_n(x)$ users in state x.

Laws of large numbers identical,
but deterministic assignment not easily distributed

Fixed User population Analysis

- Assume $\mathsf{P}\,(c) \gg \mathsf{P}$ for each $c \in [\,\underline{h}\,, \overline{h}\,]$, where P is irreducible and aperiodic

- Then $(\mu_n : n \geq 0)$ is irreducible and aperiodic on

$$\mathcal{P}_m(S) = \left\{ \left(\frac{k_1}{m}, \ldots, \frac{k_d}{m} \right) : \sum_{j=1}^{d} k_j = m,\ k_j \geq 0, 1 \leq j \leq d \right\}$$

so

$$\mu_n \Rightarrow \mu_\infty$$

- Remark. If $\mathsf{P}(c) = \mathsf{P}$,

$$\|\mathsf{P}(\mu_n \in \cdot) - \mathsf{P}(\mu_\infty \in \cdot)\| \to 1 \qquad \text{as } m \to \infty$$

so total variation convergence rate is not uniform in m (no uniform coupling)

Analysis of Distributed Control: Finite time analysis

Recall multinomial representation ...
Conditional on μ_n,

$$\mu^m_{n+1}(y) \Rightarrow \sum_{x \in S} \mu_n(x) \mathsf{P}(\mu_n h\,;\, x,\, y)$$

so if

$$\mu^m_0 \Rightarrow \nu_0$$

then

$$\mu^m_n \Rightarrow \nu_n$$

where $(\nu_n : n \geq 0)$ is the deterministic mean-field limit satisfying the non-linear iteration

$$\nu_{n+1} = \nu_n \mathsf{P}(\nu_n h)$$

Fixed Points for Mean-Field Limit

Brower Fixed Point Theorem establishes existence of at least one fixed point of

$$\nu = \nu \mathsf{P}(\nu h)$$

But mean-field generally has multiple fixed points

Multiple Fixed Points for Mean-Field Limit

For $c \in [\,\underline{h}\,, \overline{h}\,]$, choose a smooth family of positive stochastic vectors $\nu(c)$ with $\nu(c)h = c$ and set

$$\mathsf{P}(c) = \begin{pmatrix} \nu(c) \\ \text{- - - -} \\ \vdots \\ \text{- - - -} \\ \nu(c) \end{pmatrix}$$

Then, for each c,

$$\nu \mathsf{P}(c) = \nu(c)$$

for all $\nu(c) \in \mathcal{P}(S)$, so $\nu(c)$ is a fixed point for each c!

When is Fixed Point Unique?

ν is fixed point of

$$\nu = \nu \mathsf{P}(\nu h)$$

if and only if $(\nu(c), c)$ is a fixed point of

$$\begin{aligned} \nu(c) &= \nu(c)\mathsf{P}(c) \\ c &= \nu(c)h \end{aligned}$$

Note that

$$\nu'(c) = \nu'(c)\mathsf{P}(c) + \nu(c)\mathsf{P}'(c)$$

so

$$\nu'(c)(I - \mathsf{P}(c)) = \nu(c)\mathsf{P}'(c)$$

$$\nu'(c) = \nu(c)\mathsf{P}'(c)\left(I - \mathsf{P}(c) + \Pi(c)\right)^{-1}$$

If

$$\sup_{c \in [\underline{h}, \overline{h}]} \|\mathsf{P}'(c)\| \leq (1 - \varepsilon) \left(\sup_{c \in [\underline{h}, \overline{h}]} \left\| (I - \mathsf{P}(c) + \Pi(c))^{-1} \right\| \right)^{-1}$$

then we have a strict contraction (uniqueness)
Note: counterexample has $\mathsf{P}'(c)h = e$

Convergence to Fixed Point

If

$$\sup_{c \in [\underline{h}, \overline{h}]} \|\mathsf{P}'(c)\|$$

is small, then $\nu = \nu \mathsf{P}(\nu h)$ has a unique fixed point ν^* and

$$\nu_{n+1} = \nu_n \mathsf{P}(\nu_n h) \to \nu^*$$

as $n \to \infty$.

Proof establishes contractivity of original mapping (and uses Birkhoff contraction coefficient)

Connection to Finite-User Systems

Note that μ_n depends on initial distribution μ,

i.e. $$\mu_n = \mu_n(\mu)$$

$$E\left[\|\,\mu_{n+1}(\mu) - \mu_{n+1}(\nu)\|\,|\,\mathcal{F}_n\right]$$

$$\leq c\|\mu_n(\mu) - \mu_n(\nu)\| \qquad (c < 1)$$

uniformly in $m \geq m_0$ and μ, ν

Proof couples the multinomials

So, $c^{-n}\|\mu_{n+1}(\mu) - \mu_{n+1}(\nu)\|$ is a non-negative supermartingale

Convergence to stationarity is uniform in m:

$$E\|\mu_n^m - \mu_\infty^m\| \to 0 \qquad \text{as } n \to \infty$$

uniformly in $m \geq m_0$, and

$$\sup_{n \geq 1} E\|\mu_n^m - \nu_n\| \to 0 \qquad \text{as } m \to \infty$$

Selecting the Optimal Distributed Control in the Mean-Field Limit

$$\max_{\pi} \sum_{x,a} \nu(x)\pi(a|x)r(x,a)$$

s/t

$$\nu(y) = \sum_{x,a} \nu(x)\pi(a|x)\mathsf{P}_a(\nu h\,;\, x,\, y)$$

$$\nu(y) \geq 0\,, \quad \sum_{y\in S} \nu(y) = 1$$

non-linear program

Solving via LP's

$$\max_{c} \gamma(c)$$

where $\gamma(c)$ is the maximum of the LP

$$\max \sum_{x,a} \nu(x)\pi(a|x)r(x,a)$$

s/t

$$\nu(y) = \sum_{x,a} \nu(x)\pi(a|x)\mathsf{P}_a(c\,;\,x,\,y)$$

$$\nu(y) \geq 0\,, \quad \sum_{y\in S} \nu(y) = 1$$

$$\nu h = c$$

Remark: Incorporation of expectation side-constraints is easy.

Back to the Finite User Population Setting

Apply the mean-field optimal control to the finite population models:

$$\underline{\lim}_{m\to\infty} r_m^* \;\geq\; r_\infty^*$$

Apply $\pi_m^*(a|x)$, the optimal distributed control for the m'th system. Choose convergent subsequence:

$$\left(\mu_n^{*\,m_k},\ \pi_{m_k}^*\right) \Rightarrow (\nu_n^*,\ \pi^*)$$

$$r_{m_k}^* \;\to r_\infty \;\leq\; r_\infty^*$$

Optimal control problems converge as $m \to \infty$.

How much efficiency have we lost by restricting ourselves to distributed controls?

Need to study "centralized control problem"

i.e. control problem for measure-valued sequence $(\mu_n : n \geq 0)$

Optimal Control for Mean-Field Limit

Discounted Control

$$V_\alpha(\nu) = \sup_{\overrightarrow{\pi}} \sum_{n=0}^{\infty} e^{-\alpha n} \sum_{x,a} \nu_n(x)\pi_n(a|x)r(x,a)$$

where

$$\nu_{n+1}(y) = \sum_{x,a} \nu_n(x)\pi_n(a|x)\mathsf{P}_a(\nu_n h\,;\,x,\,y)$$

Then

$$\begin{aligned} V_\alpha(\nu) \;&=\; \max_\pi \Bigg[\sum_{x,a} \nu(x)\pi(a|x)r(x,a) \\ &\quad + \quad e^{-\alpha} V\Bigg(\Bigg(\sum_{x,a} \nu(x)\pi(a|x)\mathsf{P}_a(\nu h\,;\,x,\,y) : y \in S \Bigg) \Bigg) \Bigg] \end{aligned}$$

numerically easier than centralized DP for finite population system

If

$$\max_{x,a} \sup_{c \in [\underline{h}, \overline{h}]} \|(\mathsf{P}'_a(c\,;\,x,\,y) : y \in S)\ h\|$$

is sufficiently small, then

$$|V_\alpha(\nu) - V_\alpha(\tilde{\nu})| \le \beta \|\nu - \tilde{\nu}\|$$

uniformly in α. (Use coupling based on:

starting at ν: use $\pi^*_\alpha(\nu_0), \pi^*_\alpha(\nu_1), \ldots, \pi^*_\alpha(\nu_n) \quad (\nu_0 = \nu)$

starting at $\tilde{\nu}$: use $\pi^*_\alpha(\nu_0), \pi^*_\alpha(\nu_1), \ldots, \pi^*_\alpha(\nu_n) \quad (\nu_0 = \nu)$

then interchange).

Time Average Control

Let $\alpha \searrow 0$

Then, there exists a Lipschitz solution (V, g) such that

$$V(\nu) + g \;=\; \max_{\pi} \left[\sum_{x,a} \nu(x)\pi(a|x)r(x,a) \right.$$
$$\left. +\quad V\Big(\Big(\sum_{x,a} \nu(x)\pi(a|x)\mathsf{P}_a(\nu h\,;\,x,\,y) : y \in S\Big)\Big) \right]$$

Let $\pi^*(a|\nu, x)$ be the maximizing π;
and set

$$\mathsf{P}^*(\nu\,;\,x,\,y) = \sum_{x,a} \nu(x)\pi^*(a|\nu,x)\mathsf{P}_a(\nu h\,;\,x,\,y)$$

Time Average Optimal Trajectory

- Fix $\nu \in \mathcal{P}(S)$. Put $\nu_0^* = \nu$ and
$$\nu_{n+1}^* = \nu_n^* \mathsf{P}^*(\nu_n^*).$$

- $(\nu_n^* : n \geq 0)$ need not converge

- In fact, no guarantee of fixed points

Suppose there exists a fixed point $\tilde{\nu}$.

Then,

$$g = \sum_{x,a} \tilde{\nu}(x)\pi^*(a|\,\tilde{\nu},\,x)r(x,a)$$

We get optimal reward of g on every transition.

Put

$$\pi(a|x) = \pi^*(a|\tilde{\nu}, x)$$

This is a distributed control! Use it globally

$$\nu_n(\nu) \to \tilde{\nu}$$

(because distributed policies are contractive)

We get same average reward as in centralized case.

When does there exist a fixed point $\tilde{\nu}$?

$$\mathsf{P}_a^\varepsilon(\mu h\,;\,x,\,y) = \mathsf{P}_a(\varepsilon(\mu - \mu_0)h\,;\,x,\,y)$$

Then,

$$\sup_c \left\| \mathsf{P}_a'^{\varepsilon}(c\,;\,x,\,\cdot) \right\| = O(\varepsilon)$$

When $\varepsilon = 0$, there are no interference interactions (m non-interacting users)

maximizer in HJB equation unique
at $\varepsilon = 0$ (use standard DP ideas)

maximizer unique for small ε

continuity of maximizer in small neighborhood of zero

Conclusion

In the setting of "low sensitivity" to interference

optimal distributed control is also optimal for centralized control problem

optimal distributed control computable as solution to NLP

Future Work

- What happens when system is highly sensitive to interference?
- What about more complex interaction effects?

To see a set of worked out examples (with full numerics included) please go to the website of Tim Holliday at:

http://systems.stanford.edu