

Lec17

Sunday, February 23, 2020 10:02 AM

HW3 is on the web.

Need for Cross-layer Design - 10min

Saturday, January 31, 2009

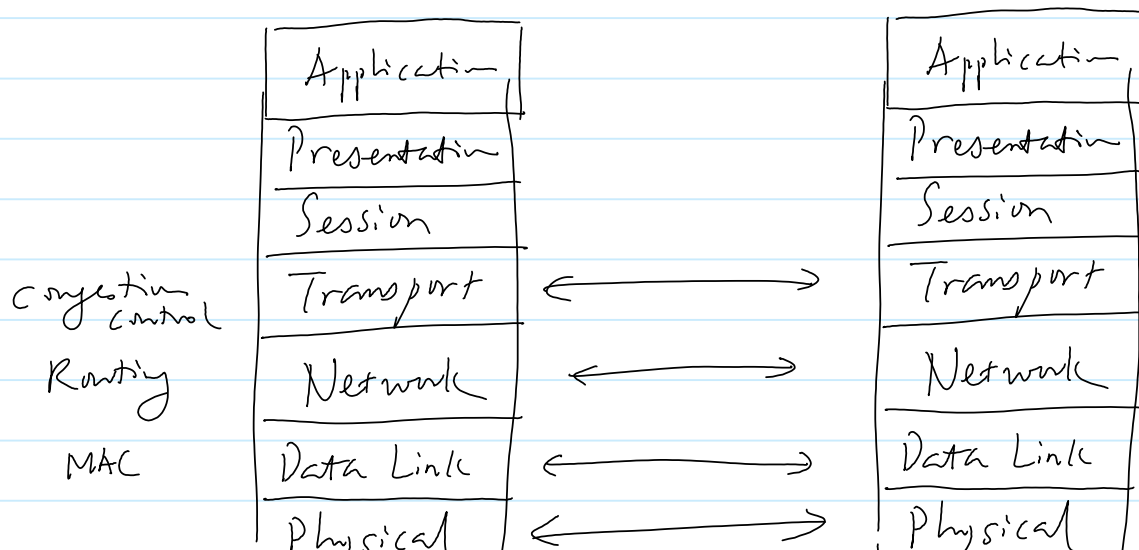
5:37 PM

We have seen an example of joint congestion control & multipath routing.


This is just one example of "cross-layer control"

① Why do we want to consider multiple layers together?

- In Wireline networks, often the protocols are classified into layers.
- Layering is a form of hierarchical modularity.
- The higher layer uses the service provided by the lower layer. But it does not need to know the inner working of the lower layer



MAC




- Benefits of Modularity.

- easy to understand
- easy to change.

- However, for wireless networks, examples have been found where such a layering architecture can limit performance.

Example.

- Typically, routing is designed to minimize the # of hops



- Tend to use "long" links.

- In wireless networks, long transmission links can suffer from a low SNR

⇒ poor end-to-end performance.


- It would be better if the routing protocol takes into account the physical-layer characteristics.
- Pitfalls of Cross-layer Design
 - loss of modularity
 - fragile solution that is hard to change.
- Nonetheless, the Lyapunov-based approach allows us to easily integrate different functionality together
 - Many of the solutions consist of a MAC component of the form
$$\max_{\vec{p} \in \mathcal{H}} \sum_l g^l(t) r_l$$
$$\vec{r} = g(\vec{p}, K(t))$$
- The queue-length becomes the "glue" across layers.
- We will not study all of these problems. Instead, focus on routing.

The project¹ deals with one flavor of
rate-control & fairness.

Joint congestion control - 5min

Sunday, February 11, 2018 9:48 AM

- Suppose that we want to also choose the arrival rate x_s each flow s to maximize some system objective
- Define a utility function $U(x_s)$
 - increasing, concave
 - e.g. $U(x_s) = \log x_s$
 - no flow's rate can be too large or too small (fair)
 - logarithmic utility corresponds to "proportional fairness"
- Let us consider the problem of
$$\begin{aligned} \max \quad & \sum_s U(x_s) \\ \text{sub to} \quad & \left(\sum_s H_{s,k} x_s \right) \in \Lambda \triangleq \sum_k \lambda_k \text{Conv-hull} \{ g(\vec{p}, k) \mid \vec{p} \in \mathcal{P} \} \end{aligned}$$
- Let $[x_s^*]$ be this optimal vector.



Let us begin with the same Lyapunov function as before:

$$V(\vec{q}) = \frac{1}{2} \sum_i (q^i)^2$$

Since

$$\begin{aligned} \Rightarrow (q^l(t+1))^2 &\leq (q^l(t) + \sum_j H_j^l x_j - r^l(t))^2 \\ &= (q^l(t))^2 + 2q^l(t) \left[\sum_j H_j^l x_j - r^l(t) \right] \\ &\quad + \left[\sum_j H_j^l x_j - r^l(t) \right]^2 \end{aligned}$$

Since x_j is finite, and $r^l(t)$ is bounded, there exists a constant M_l such that

$$(q^l(t+1))^2 \leq (q^l(t))^2 + 2q^l(t) \left[\sum_j H_j^l x_j - r^l(t) \right] + M_l$$

$$\begin{aligned} \Rightarrow V(\vec{q}(t+1)) &\leq V(\vec{q}(t)) + \sum_l q^l(t) \left[\sum_j H_j^l x_j - r^l(t) \right] \\ &\quad + \frac{1}{2} \sum_l M_l^2 \end{aligned}$$

$$\begin{aligned} \Rightarrow E[V(\vec{q}(t+1)) - V(\vec{q}(t)) | \vec{q}(t)] &\leq \sum_l q^l(t) \sum_j H_j^l x_j - \sum_l q^l(t) E[r^l(t) | q^l(t)] \\ &\quad + \frac{1}{2} \sum_l M_l^2 \end{aligned}$$

However, maximizing $\sum_l q^l(t) E[r^l(t) | q^l(t)]$ does not help us to choose x_j .

Drift + Penalty

- Add the penalty term

$$\Delta(t) = \sum_j u(x_j^*) - \sum_j u(x_j)$$

- How suboptimal $[x_s]$ is.

- The sum is

$$V(\vec{p}(t+1)) - V(\vec{p}(t)) + \alpha \Delta(t)$$

$$\leq \left[\sum_l p'_l(t) \sum_s H'_s x_s - \alpha \sum_s u(x_s) \right] \xrightarrow{A}$$

$$= \sum_l p'_l(t) \mathbb{E}[r'_l(t) | p'_l(t)] + \frac{1}{2} \sum_l M_l^2 + \alpha \sum_{i,j} u(\lambda_{ij}^{c,*})$$

\xrightarrow{B}

- We again try to minimize the drift plus-penalty

- Minimizing A is the same as

$$\max_{[x_s]} \alpha \sum_s u(x_s) - \sum_l p'_l(t) \sum_s H'_s x_s$$

$$= \max_{[x_s]} \sum_s \left[\alpha u(x_s) - x_s \cdot \sum_l H'_s p'_l(t) \right]$$

$$= \sum_s \max_{x_s} \left[\alpha u(x_s) - x_s \cdot \sum_l H'_s p'_l(t) \right]$$

- Hence, each flow chooses x_s

$$x_s(t) = \operatorname{argmax}_{x_s} \alpha u(x_s) - x_s \cdot \sum_l H'_s p'_l(t)$$

- We may interpret $\frac{p'_l(t)}{\alpha}$ as a "price" of link l

- $\sum_l H'_s \frac{p'_l(t)}{\alpha}$ is the total "price" along the path of flow s .

- $x_s(t)$ is thus chosen to maximize the net utility.

- For $u(x_s) = \log x_s$

$$\Rightarrow \frac{\alpha}{x_s} = \sum_l H'_s p'_l(t) \Rightarrow$$

$$x_s = \frac{\alpha}{\sum_l H'_s p'_l(t)}$$

- maximizing B leads to the same max-weight scheduling decision.

Why does this work?

- Suppose that we choose $X_S(t) = X_S^*$.
- Then, since $[X_S^*] \in \mathcal{L}$, we must have

$$\begin{aligned} E[V(\vec{q}(t+1)) - V(\vec{q}(t)) \mid \vec{q}(t)] \\ \leq \sum_i q_i'(t) \sum_j H_j^i(X_S) - \sum_i q_i'(t) E[r_i'(t) \mid q_i'(t)] \\ \quad + \frac{1}{2} \sum_i M_i^2 \\ \leq \frac{1}{2} \sum_i M_i^2 \end{aligned}$$

- Further,

$$\begin{aligned} \Delta(t) &= \sum_S u(X_S^*) - \sum_S u(X_S) \\ &= 0 \end{aligned}$$

- Together, the total drift + penalty would have been

$$\begin{aligned} V(\vec{q}(t+1)) - V(\vec{q}(t)) + \alpha \Delta(t) \\ \leq \frac{1}{2} \sum_i M_i^2 \end{aligned}$$

- Now, since our choice of $X_S(t)$ & $r_i'(t)$ minimize the drift + penalty, the same must be true for our decision (even though we do not know X_S^*).

- Summing over all t , divided by T

$$\frac{\mathbb{E}[V(\vec{q}(T+1)) - \mathbb{E}[V(\vec{q}(0))]]}{T}$$

$$+ \alpha \sum_s U(x_s^*) - \alpha \frac{1}{T} \sum_{t=1}^T \sum_s \mathbb{E}[U(x_s(t))] \\ \leq M \triangleq \frac{1}{2} \sum_s M_s^2$$

$$\Rightarrow \frac{1}{T} \sum_{t=1}^T \sum_s \mathbb{E}[U(x_s(t))]$$

$$\geq \sum_s U(x_s^*) - \underbrace{\frac{M}{\alpha}}_{\rightarrow 0 \text{ as } \alpha \rightarrow +\infty} + \underbrace{\frac{\mathbb{E}[V(\vec{q}(T+1))]}{\alpha T}}_{\geq 0} - \underbrace{\frac{\mathbb{E}[V(\vec{q}(0))]}{\alpha T}}_{\rightarrow 0 \text{ as } T \rightarrow +\infty}$$

- As $\alpha \rightarrow +\infty$, the loss of average utility $\rightarrow 0$.

Reference:

Michael J. Neely, Eytan Modiano, Chih-ping Li, "Fairness and Optimal Stochastic Control for Heterogeneous Networks," *IEEE/ACM Transactions on Networking*, April 2008

Let us begin with the same Lyapunov function as before:

$$V(\vec{q}) = \frac{1}{2} \sum_i (q^i)^2$$

Since

$$q^i(t+1) = [q^i(t) + \sum_j H_j^i x_j - r^i(t)]^+$$

$$\begin{aligned} \Rightarrow (q^i(t+1))^2 &\leq (q^i(t) + \sum_j H_j^i x_j - r^i(t))^2 \\ &= (q^i(t))^2 + 2q^i(t) \left[\sum_j H_j^i x_j - r^i(t) \right] \\ &\quad + \left[\sum_j H_j^i x_j - r^i(t) \right]^2 \end{aligned}$$

Since x_j is finite, bounded, there exists a constant M_j such that

$$(q^i(t+1))^2 \leq (q^i(t))^2 + 2q^i(t) \left[\sum_j H_j^i x_j - r^i(t) \right] + M_i$$

$$\Rightarrow V(\vec{q}(t+1)) \leq V(\vec{q}(t)) + \sum_i q^i(t) \left[\sum_j H_j^i x_j - r^i(t) \right] + \frac{1}{2} \sum_i M_i^2$$

$$\begin{aligned} \Rightarrow E[V(\vec{q}(t+1)) - V(\vec{q}(t)) | \vec{q}(t)] \\ \leq \sum_i q^i(t) \sum_j H_j^i x_j - \sum_i q^i(t) E[r^i(t) | q^i(t)] \\ + \frac{1}{2} \sum_i M_i^2 \end{aligned}$$

However, maximizing $\sum_i q^i(t) E[r^i(t) | q^i(t)]$ does not help us to choose x_j .

Drift + Penalty

- Add the penalty term

$$\Delta(t) = \sum_s u(x_s^*) - \sum_s u(x_s)$$

— How suboptimal $[x_s]$ is.

— The sum is

$$V(\vec{p}(t+1)) - V(\vec{p}(t)) + \alpha \Delta(t)$$

\leq

— We again try to minimize the drift plus-penalty

— Minimizing A is the same as

$$\max_{[x_s]} \alpha \sum_s u(x_s) - \sum_i p'_i(t) \sum_s H_s^i x_s$$

$=$

$=$

— Hence, each flow chooses x_s

$$x_s(t) = \arg \max_{x_s} \alpha u(x_s) - x_s \cdot \sum_i H_s^i p'_i(t)$$

$$\text{— For } u(x_s) = \log x_s$$

$$\Rightarrow \frac{\alpha}{x_s} = \sum_i H_s^i p'_i(t) \Rightarrow$$

$$x_s = \frac{\alpha}{\sum_i H_s^i p'_i(t)}$$

— maximizing B leads to the same max-weights

scheduling decision,

Why does this work?

- Suppose that we choose $X_S(t) = X_S^*$.
- Then, since $[X_S^*] \in \mathcal{A}$, we must have

$$\mathbb{E}[V(\vec{r}(t+1)) - V(\vec{r}(t)) \mid \vec{r}(t)]$$

- Further,

$$\Delta(t)$$

- Together, the total drift + penalty would have been

$$\begin{aligned} & V(\vec{r}(t+1)) - V(\vec{r}(t)) + \alpha \Delta(t) \\ & \leq \frac{1}{2} \mathbb{E} M_t^2 \end{aligned}$$

- Now, since our choice of $X_S(t)$ & $r^1(t)$ minimize the drift + penalty, the same must be true for our decision (even though we do not know X_S^*).

— Summing over all t , divided by T

$$\begin{aligned} \Rightarrow & \frac{1}{T} \sum_{t=1}^T \sum_s \mathbb{E} \left(u(x_s(t)) \right) \\ & \geq \sum_s u(x_s^*) - \frac{M}{\alpha} - \underbrace{\frac{\mathbb{E}(V(\vec{r}(0)))}{\alpha T}}_{\rightarrow 0 \text{ as } T \rightarrow +\infty} + \underbrace{\frac{\mathbb{E}(V(\vec{r}(T+\nu)))}{\alpha T}}_{\rightarrow 0 \text{ as } T \rightarrow +\infty} \end{aligned}$$

— As $\alpha \rightarrow +\infty$, the loss of average utility $\rightarrow 0$.

Queue / Optimizing tradeoff

— As $\alpha \rightarrow +\infty$, we expect that $X_S(\alpha)$ is around X_S^*

— From (A)

$$\max \alpha U(X_S) - X_S \sum_i H_i' p_i'(\alpha)$$

$$\Rightarrow \alpha U'(X_S) = \sum_i H_i' p_i'(\alpha)$$

— As $\alpha \uparrow$, so is $p_i'(\alpha) \uparrow$.

Closer optimality is achieved at the cost of higher queue length.

— Similar approach can be applied to other objectives, such as energy consumption.

- Optimal Joint Routing & Scheduling

Model:

- Link model same as that in the scheduling problem

$$\vec{r} = f(\vec{p}, K(t))$$

$$\vec{p} \in \Theta$$

$K(t)$ i.i.d over time.

- Multiple commodities $c=1, 2, \dots, C$, each with a destination node $d(c)$

- A number of nodes may generate packets of commodity c .

λ_i^c = rate of new packets of commodity c generated by node i

- No routes are specified yet.

- In order for the arrival rate vector $[\lambda_i^c]$ to be supported, there must exist $[\bar{r}_{ij}^c]$ such that

- \bar{r}_{ij}^c is the rate that node i forwards commodity c to node j

$$\lambda_i^c + \sum_{j \neq i} \bar{r}_{ji}^c \leq \sum_{j \neq i} \bar{r}_{ij}^c$$

for all node $i \neq d(c)$

$$\left(\sum_c \bar{r}_{ij}^c \right) \in \mathcal{L} \stackrel{\text{def}}{=} \sum_k \lambda_k \text{Conv-Hull} \{ f(\vec{p}, k) \mid \vec{p} \in \Theta \}$$

Queue-length Based Policy

- Each node maintains multiple queues, one for each destination c .

$$q_i^c$$

Let r_{ij}^c be the amount of packets of commodity c that are forwarded from node i to node j .

$$q_i^c(t+1) = \begin{cases} \left[q_i^c(t) + \sum_{j \neq i} r_{ji}^c + \lambda_i^c - \sum_{j \neq i} r_{ij}^c \right]^+ & \text{if } i \neq d(c) \\ 0 & \text{if } i = d(c). \end{cases}$$

- We now see why working with the Lyapunov function also produces a joint routing & scheduling policy that is throughput-optimal.

- Use the Lyapunov function

$$V(\vec{q}) = \frac{1}{2} \sum_{i,c} (q_i^c)^2$$

Note

$$\begin{aligned} & (q_i^c(t+1))^2 - (q_i^c(t))^2 \\ & \leq 2 q_i^c(t) \left[\sum_{j \neq i} r_{ji}^c + \lambda_i^c - \sum_{j \neq i} r_{ij}^c \right] + \text{constant} \end{aligned}$$

$$\Rightarrow V(\vec{q}(t+1)) - V(\vec{q}(t))$$

$$\leq \sum_{i,c} q_i^c(t) \left[\sum_{j \neq i} r_{ji}^c + \lambda_i^c - \sum_{j \neq i} r_{ij}^c \right] + \text{constant}$$

$$= \sum_{i,c} q_i^c(t) \cdot \lambda_i^c - \underbrace{\sum_{(i,j) \in E, c} r_{ij}^c [q_i^c(t) - q_j^c(t)]}_{\text{should maximize this to minimize the drift.}}$$

- Give $\vec{r} = f(\vec{p}, K(t))$, we have $\sum_c r_{ij}^c = r_{ij}$

\Rightarrow maximizing the last term implies that we should only use r_{ij}^c with the largest value of $(r_i^c(t) - r_j^c(t))$

- The last-term becomes $\sum_{ij} r_{ij} \max_c (r_i^c(t) - r_j^c(t))$
 \Rightarrow should then choose $\vec{p}(t)$ to maximize this weighted sum.

Joint Routing & Scheduling Algv.

① Maximum differential backlog.

For each link (i, j) , select the commodity c such that the value

$$r_i^c(t) - r_j^c(t)$$

is the largest.

$$\text{Let } c_{ij}^*(t) = \arg \max_c r_i^c(t) - r_j^c(t)$$

Let $w_{ij}(t) = r_i^{c_{ij}^*(t)} - r_j^{c_{ij}^*(t)}$ denote the maximum differential backlog.

② Scheduling.

Choose $\vec{p}(t)$ such that

$$\vec{p}(t) = \arg \max_{\vec{p} \in \Theta} \sum_{(i,j)} w_{ij}(t) \cdot r_{ij}$$

$$\vec{r} = f(\vec{p}, K(t))$$

$$\text{Let } \vec{r}(t) = f(\vec{p}(t), K(t))$$

③ Routing:

On each link (i, j) , route the commodity $c_{ij}^*(t)$ using the rate r_{ij} i.e.

$$r_{ij}^c(t) = \begin{cases} r_{ij}(t) & \text{if } c = c_{ij}^*(t) \\ 0 & \text{, otherwise} \end{cases}$$

Intuition: As packets are queued, the queue difference forms a "gradient", which points to the optimal direction to forward packets

Can be shown to achieve the largest set of offered loads $\{\lambda_i^*\}$.

Reference: Neely & Modiano.

[Dynamic Power Allocation and Routing for Time Varying Wireless Networks](#), by M. J. Neely, E. Modiano and C. E. Rohrs, in IEEE INFOCOM, April 2003.

(35)