

Lecture 3

Lecturer: Giovanni Neglia

Scribe: Mohamad El laz

NOTE: The content of these notes has not been formally reviewed by the lecturer. It is recommended that they are read critically.

1 Introduction

In the previous lectures we considered examples where users wanted to transmit at a fixed rate, and were able to choose their transmission paths, their routes.

In this lecture, we study the case where routes are fixed, but users can change their transmission rate. This is the usual situation for TCP flows over the Internet. For this reason, we could title our lecture “Bandwidth sharing over the Internet.”

2 Problem definition

Let us consider the case where a user wants to download some content from the Internet (Figure 1).




Figure 1: server-client connection

In order to download some contents from the Internet, a connection might be established between the user and the server, this process is controlled by TCP (*Transmission Connection Protocol*). TCP is a connection-oriented protocol, it determines how to break application data into packets that networks can deliver, sends packets to and accepts packets from the network layer, manages flow control, and handles retransmission of dropped or corrupted packets as well as acknowledgment of all packets that arrive.

In this example the route is fixed by the Internet provider. For the user, the faster its downloading rate, the better, but link capacities are limited and moreover they need to be shared with other users. We consider a first client (c_1) downloading a file from the server on route A-B-D-E, a first constraint is that the download speed cannot exceed 1 Mbps, the capacity of link B-D. The second constraint would be that another client (c_2) is trying to download from the same source, e.g. through route A-C-D-E or in any case he is sharing a common link (e.g. in Route C-B-D-E). To see how the Internet splits the downloading rate we will consider a general model.

3 A general model for the problem

We can consider that the “happiness” of a user is a function of his downloading rate. A very common way to model such happiness is by an increasing concave function (Figure 2):

- Increasing as the higher the rate, the happier a user is.
- Concave because of the law of diminishing returns.

Remark 1. *The law of diminishing return is a commonly used concept in the field of economy. For example, let us consider the diminishing marginal utility (happiness) of income, which means that as income increases, individuals gain correspondingly a smaller increase in happiness. To explain it better let us consider the example below.*

If you have zero income and gain 10\$ per day, this will improve your living standards significantly as you will be able to pay for basic necessity of life (i.e. food).

However, if you already gain 500\$ a day, an extra 10\$ leads to a smaller increase in utility (it does not significantly affect your standard of living and happiness).

Now, if you are earning 15,000\$ a day, you would hardly notice an extra 10\$. Therefore, then the marginal utility of an extra 10\$ at this income level is very limited.

The same idea is applied to the marginal happiness of users downloading from Internet, where the user is unhappy when there is no downloading rate and happier if the rate is higher. When having a twice faster connection, the user will be happier but not twice as much.

Thus, as income increases, the extra marginal benefit to individuals declines.




Figure 2: An increasing concave function

3.1 Solving the problem by the Lagrange multipliers method

We will just identify the users as the routes (one user corresponds to one route):

- R is the set of all routes/users.
- E is the set of edges.
- $l \in r$ where link l is a part of route r .
- C_l is the capacity of link l (the maximum traffic rate that l can support).
- $U_r(x_r)$ is the utility of user r for downloading at rate x_r .

Thus our objective is to maximize the sum of users happiness.

The optimization problem can be modeled as follows:

$$\begin{aligned}
& \text{maximize} && \sum_{r \in R} U_r(x_r) \\
& \text{subject to} && \sum_{r|l \in r} x_r \leq C_l, \forall l \in E \\
& && x_r \geq 0, \forall r \in R
\end{aligned} \tag{1}$$

We also denote $y_l = \sum_{r|l \in r} x_r$.

For convenience, we also assume that $U_r(x_r)$ is strictly concave, and continuously differentiable on $[0, \infty)$.

This is a problem for which we know that the Lagrangian is a necessary and sufficient condition (because the sum of concave functions is a concave function and the constraints are linear), so we can apply *Theorem 2* of lecture 2a for concave maximization. We can also apply *Theorem 1* of lecture 2a for convex minimization to the following problem:

$$\begin{aligned}
& \text{minimize} && \sum_{r \in R} -U_r(x_r) \\
& \text{subject to} && \sum_{r|l \in r} x_r \leq C_l, \forall l \in E \\
& && x_r \geq 0, \forall r \in R
\end{aligned} \tag{2}$$

The resulting Lagrangian function:

$$L(\mathbf{x}, \boldsymbol{\mu}) = \sum_{r \in R} -U_r(x_r) + \sum_{l \in E} \mu_l \left(\sum_{r: l \in r} x_r - C_l \right) \tag{3}$$

When differentiated with respect to the rate on link \bar{r} :

$$\frac{\partial L}{\partial x_{\bar{r}}} = -U'_{\bar{r}}(x_{\bar{r}}) + \sum_{l \in \bar{r}} \mu_l \tag{4}$$

\mathbf{x}^* is a global minimizer iff:

1. x^* is feasible.
2. There exists $\mu_l^* \geq 0$, such that $\nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\mu}^*)^T (\mathbf{x} - \mathbf{x}^*) \geq 0, \forall \mathbf{x} \in X$, where X is the set of constraints we left out of the Lagrangian, i.e. $X = \{\mathbf{x}, x_r \geq 0 \forall r\}$.

The slack conditions are

$$\mu_l^* \left(\sum_{r: l \in r} x_r^* - C_l \right) = 0 \tag{5}$$

One of the two factors is equal to zero, then

$$\mu_l^* \begin{cases} \geq 0 & \text{if the link is fully used, } (\sum_{r: l \in r} x_r^* - C_l) = 0, \\ = 0 & \text{if the link is under utilized, } (\sum_{r: l \in r} x_r^* - C_l) < 0, \end{cases} \tag{6}$$

3.1.1 The optimum value

Let $x_{\bar{r}}^*$ be the optimum value of $x_{\bar{r}}$.

Observe that if $x_{\bar{r}}^* > 0$ we can consider feasible solutions $\mathbf{x} \in X$, whose \bar{r} -th component is larger or smaller than $x_{\bar{r}}^*$. We can find $\Delta x > 0$ such that both $\mathbf{x}^* + \Delta x \mathbf{e}_{\bar{r}}$ and $\mathbf{x}^* - \Delta x \mathbf{e}_{\bar{r}}$ belong to X . Then, from the condition on the gradient of the Lagrangian, it follows:

$$\left. \frac{\partial L}{\partial x_{\bar{r}}} \right|_{\mathbf{x}^*, \boldsymbol{\mu}^*} \Delta x \geq 0,$$

$$\left. \frac{\partial L}{\partial x_{\bar{r}}} \right|_{\mathbf{x}^*, \mu^*} (-\Delta x) \geq 0,$$

From these two results we can see that

$$\frac{\partial L}{\partial x_{\bar{r}}} = 0, \text{ if } x_{\bar{r}}^* > 0. \quad (7)$$

Instead if $x_{\bar{r}}^* = 0$, it is only possible to increase it and we get:

$$\left. \frac{\partial L}{\partial x_{\bar{r}}} \right|_{\mathbf{x}^*, \mu^*} \geq 0 \text{ if } x_{\bar{r}}^* = 0 \quad (8)$$

From (7) we get that:

$$U'_{\bar{r}}(x_{\bar{r}}^*) = \sum_{l \in \bar{r}} \mu_l^*, \text{ if } x_{\bar{r}}^* > 0 \quad (9)$$

The structure is similar to the one we studied in the other examples. Again, we have that at the optimum, the marginal utility $U'_{\bar{r}}(x_{\bar{r}}^*)$ should be equal to the sum over all the links in the route of the Lagrange multipliers associated to the link constraints. Moreover, the multipliers are *equal to zero* over all the unused links and can be *greater than zero* only over the utilized links. It is then natural to consider a Lagrange multiplier as the marginal congestion cost of the link, that expresses the damage that a user causes to other users transmitting over the same link. $\sum_{l \in \bar{r}} \mu_l^*$ is the marginal congestion cost of the route \bar{r} .

From (8) we get that:

$$U'_{\bar{r}}(x_{\bar{r}}^*) \leq \sum_{l \in \bar{r}} \mu_l^*, \text{ if } x_{\bar{r}}^* = 0 \quad (10)$$

In this case, user \bar{r} does not transmit, because its marginal utility at $x_{\bar{r}} = 0$ is smaller than the congestion cost on the route (and because of concavity it would be even smaller for positive transmission rates). An increase of its rate would cause a significant loss of utility to all the other users sharing one of the links on route \bar{r} . Therefore, user \bar{r} gets nothing, because the increase of his happiness is not worthy the increase of the congestion cost on the route and the decrease of other users' happiness.

3.2 A first attempt to introduce dynamics

Until now, we have characterized the equilibrium. We would like to see also by which dynamics the system can converge to the optimum solution. The analysis above suggests the following possible operation: each link transmits to each user its current marginal congestion cost and users react to the current marginal congestion cost of the route by increasing/decreasing their rates.

In table 1, we consider the case of three users sharing the same congested link (capacity is equal to 10) according to the following simple discrete-time dynamics. Each user increases its rate by one if the marginal congestion cost is null and decreases it if it is positive. Step by step, by increasing time, users increase their rate and the network tells them the cost of the congestion associated to the link. Once the congestion is reached (> 0), the network tells the users to decrease their rate and the process repeats itself, without never reaching the equilibrium. This example suggests that an (ON-OFF) congestion signal, where links only reveal if there is or not congestion, may not be lead the system to converge to the optimum.

In the next section we will slightly modify the problem by looking at a system for which links provide feedback even before being completely utilized. We will see a particular dynamic that actually converges to the solution of the optimum.

4 A relaxed version of the problem

What we have seen until now is the system where the cost of congestion is zero up to where we have no congestion, then positive when we have congestion. But it is not the best way for a dynamic system, it would be better if we have a situation where we already pay something when we approach congestion. We look at a slightly different problem where we can even violate the constraints on the capacity but at the same time we will have an additional cost related to the congestion.

Table 1: Congested link

Time	User r_1	User r_2	User r_3	Congestion
0	0	0	0	0
1	1	1	1	0
2	2	2	2	0
3	3	3	3	0
4	4	4	4	> 0
3	3	3	3	0
4	4	4	4	> 0
...

Let us imagine that for every link there is some cost $M_l(y_l)$ (figure) This function should be increasing (increases with the rate), strictly convex (because it should grow faster when we arrive to C_l) and non negative. Note that we could assume that $M_l(y_l)$ has a vertical asymptote for $y_l = C_l$, but we will consider that $M_l(y_l)$ has finite value also for $y_l \geq C_l$. This will allow us to take into account the possibility that the traffic on a link may temporarily exceed the link capacity (this extra traffic may be accommodated by link buffers).




Figure 3: An increasing convex function [RATES CANNOT BE NEGATIVE]

Another technical condition is that this function should grow faster than the utility. This guarantees that the optimal solution of the problem below corresponds to finite rate values. For this reason, we introduce the following condition:

$$\lim_{x \rightarrow \infty} U_r(x_r) - M_l(x_r) = -\infty, \forall r \in R, \forall l \in E \quad (11)$$

So now, we consider a relaxed version of our initial problem in which we have no more hard constraint but a softer constraint transformed into a cost.

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{y}}{\text{maximize}} && \sum_{r \in R} U_r(x_r) - \sum_{l \in E} M_l(y_l) \\ & \text{subject to} && y_l = \sum_{r|l \in r} x_r \\ & && x_r \geq 0, \forall r \in R \end{aligned} \quad (12)$$

Where $M_l(y_l)$ is the cost of congestion on link l .

So let us write again the Lagrangian as:

$$L(\mathbf{x}, \mathbf{y}, \boldsymbol{\nu}) = - \sum_{r \in R} U_r(x_r) + \sum_{l \in E} M_l(y_l) - \sum_{l \in E} \nu_l \left(y_l - \sum_{r: l \in r} x_r \right) \quad (13)$$

Then the derivatives are:

$$\begin{cases} \frac{\partial L}{\partial y_l} = M'_l(y_l) - \nu_l \\ \frac{\partial L}{\partial x_{\bar{r}}} = -U'(x_{\bar{r}}) + \sum_{l \in \bar{r}} \nu_l \end{cases} \quad (14)$$

Following the same reasoning as in the last lesson we have:

$$\frac{\partial L}{\partial x_{\bar{r}}} \Big|_{\mathbf{x}^*, \mathbf{y}^*, \boldsymbol{\nu}^*} \begin{cases} = 0 & \text{if } x_{\bar{r}}^* > 0, \\ \geq 0 & \text{if } x_{\bar{r}}^* = 0, \end{cases} \quad (15)$$

So from (14) we get:

$$\nu_l^* = M'_l(y_l^*)$$

and

$$U'(x_{\bar{r}}^*) \begin{cases} = \sum_{l \in \bar{r}} \nu_l^* & \text{if } x_{\bar{r}}^* > 0, \\ \leq \sum_{l \in \bar{r}} \nu_l^* & \text{if } x_{\bar{r}}^* = 0. \end{cases} \quad (16)$$

As above, at the equilibrium on each route the marginal utility is equal to the total marginal congestion cost of the route. In this case the marginal congestion cost of a link is expressed by $M'_l(y_l)$. Because of the convexity of $M_l(y_l)$, the marginal congestion cost is increasing in the traffic on the link.

4.1 Convergence

Now we can study a new system dynamics. Let $x_r(t)$ and $y_l(t)$ denote respectively the transmission rate along route r and the traffic on link l at time t . We can think that user r increases/decreases its rate proportionally to the difference between his marginal utility and the marginal congestion cost along the route:

$$\frac{dx_r}{dt} = k_r (U'(x_r(t)) - \sum_{l \in r} M'_l(y_l(t))) \quad (17)$$

Where do we expect this system to converge? Imagine that the system stops at a given configuration. Let \hat{x}_r and \hat{y}_l denote the corresponding quantities. If this is an equilibrium, then (17) should predict that the rates do not change anymore, but from the right hand side of the equation, it follows: $U'(\hat{x}_r) = \sum_{l \in r} M'_l(\hat{y}_l)$

The equilibrium of the dynamics indicated by (17) corresponds then to the optimal solution of our optimization problem. We can conclude that, if the dynamics stop, then they stop at a global optimizer.

So if everyone adapts his rate according to (17) and the system stops, this will be at the optimum. In fact this approach is not rigorous because we actually do not know if it will stop.

Next we show that the congestion avoidance algorithm of TCP, due to Jacobson (1988), can be described in this framework.

5 Modeling TCP

We start describing a simplified operation for TCP.

Let T_r , be the round-trip time along route r , i.e. the time between a source sending a packet and the source receiving an acknowledgement. The source maintains a window (of size cwnd) of packets that have been sent but not yet acknowledged. The transmission rate x can then be calculated as cwnd/T_r . During one round-trip time T_r , the source should receive cwnd acknowledgements (then one every T_r/cwnd). If a given acknowledgement is received, then cwnd is increased by $1/\text{cwnd}$. If the acknowledgement is not received, cwnd is halved.

Let p_r be the packet loss probability over route r . During the time T_r/cwnd , cwnd increases of $1/\text{cwnd}$ with probability $1 - p_r$ and decreases of $\text{cwnd}/2$ with probability p_r . Similarly, the instantaneous rate will change as follows:

- if there is no packet loss (and then with probability $1 - p_r$), it will increase by $\frac{1/cwnd}{T_r}$
- if packet loss (and then with probability p_r), it will decrease by $-\frac{cwnd}{2T_r}$.

The expected change of the rate is then $\Delta x_r = \frac{1/cwnd}{T_r}(1 - p_r) - \frac{cwnd}{2T_r}p_r$. The evolution of the rate can then be approximated as follows:

$$\begin{aligned}
\frac{dx_r}{dt} &= \frac{\Delta x_r}{T_r/cwnd} \\
&= \frac{cwnd}{T_r} \left(\frac{1/cwnd}{T_r}(1 - p_r) - \frac{cwnd}{2T_r}p_r \right) \\
&= \frac{1}{T_r^2}(1 - p_r) - \frac{cwnd^2}{2T_r^2}p_r \\
&= \frac{1}{T_r^2} - p_r \left(\frac{1}{T_r^2} + \frac{x_r^2}{2} \right) \\
&= \frac{1}{T_r^2} \left(1 - p_r \left(1 + \frac{x_r^2 T_r^2}{2} \right) \right) \\
&= \frac{1 + \frac{x_r^2 T_r^2}{2}}{T_r} \left(\frac{1}{1 + \frac{x_r^2 T_r^2}{2}} - p_r \right)
\end{aligned} \tag{18}$$

By a comparison with (17), the first term in (18) plays the same role of $k_r > 0$, that is a parameter affecting the speed of convergence, but not its direction.

The term $1/\left(1 + \frac{x_r^2 T_r^2}{2}\right)$ corresponds to the marginal utility $U'_r(x_r)$ that a user gets from transmitting at rate x_r . By integrating, we obtain then that:

$$U_r(x_r) = \frac{\sqrt{2}}{T_r} \arctan \frac{x_r T_r}{\sqrt{2}} \tag{19}$$

The term p_r should correspond to the marginal cost of congestion on route r . p_r is the packet loss probability, that is indeed an indication of congestion. We would like the marginal congestion cost on a route to be expressed as sum of link costs, but the probability to lose a packet on a route is not the sum of probability to lose the packet over all links. For example consider a route r with 3 links. The probability of losing a packet is:

$$p_r = 1 - (1 - p_1)(1 - p_2)(1 - p_3) \tag{20}$$

But if the probability to lose a packet is not high then (20) would be equal to:

$$p_r \approx 1 - 1 + p_1 + p_2 + p_3 = p_1 + p_2 + p_3 \tag{21}$$

If we assume that the network is working with low packet loss probabilities, then we can identify:

$$M'_l(y_l) = p_l(y_l), \tag{22}$$

from which we obtain that

$$M_l(y_l) = \int_0^{y_l} p_l(t) dt \tag{23}$$

We conclude that all TCP sources together are solving the following problem:

$$\begin{aligned}
&\text{maximize} && \sum_{r \in R} \frac{\sqrt{2}}{T} \arctan \frac{(x_r)(T)}{\sqrt{2}} - \sum_{l \in E} \int_0^{y_l} p_l(t) dt \\
&\text{subject to} && y_l = \sum_{r|l \in r} x_r, \forall l \in E \\
&&& x_r \geq 0, \forall r \in R
\end{aligned} \tag{24}$$

This is the global optimization problem that TCP is solving in a distributed way. We see that TCP (or at least our differential equation model of it) is behaving as if it is maximizing the sum of user utilities, subject to a cost function penalizing proximity to the capacity constraints. We saw how every TCP source cares only of packet losses along its path. The specific optimization problem solved by TCP sources is difficult to justify, but TCP congestion control was originally proposed to avoid congestion collapses and not to solve a meaningful optimization problem. Later, once this optimization framework had been recognized, other congestion control mechanisms have been engineered. In theory these other mechanisms are better than TCP, but they were not widely adopted. The problem is that we cannot start from scratch. In general if we could restart the Internet, we know how to do it better. But we cannot switch off all the transmitting computers and start them with the new congestion protocols. All the new protocols suffer for the fact that they are less aggressive than TCP, and then they perform worse in a heterogeneous scenario where other users still adopt TCP. Some of these protocols are now used for inter-machine communication inside data centers.

References

- [1] Frank Kelly and Elena Yudovina, Stochastic Networks. *Cambridge Press*, 2014.