

Convex Optimization and Congestion Control

Fabian Wirth

University of Würzburg

23.07. – 03.08.2012

- Part I: Convexity and Convex Functions Lectures 1, 2, 3
- Part II: Convex Optimization Lectures 4 and 5
- Part III: Numerical Methods Lectures 6 and 7
- Part IV: Congestion Control Lecture 8
- Part V: Utility Based Congestion Control Lecture 9
- Part VI: Miscellaneous Problems in Networks Lecture 10 (we shall see)

- III.1: Unconstrained Problems
- III.2: Constrained Problems
- III.3: Interior Point Methods

- 1 Unconstrained Problems
- 2 Constrained Problems
- 3 Interior Point Methods

- 1 Unconstrained Problems
- 2 **Constrained Problems**
- 3 Interior Point Methods

- 1 Unconstrained Problems
- 2 Constrained Problems
- 3 Interior Point Methods**

The Convex Optimization Problem with Inequality Constraint

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && g_j(x) \leq 0 && j = 1, \dots, m \\ & && Ax = b. \end{aligned} \tag{1}$$

Assumptions

- (i) f, g_j are convex and twice continuously differentiable,
- (ii) $\mathcal{D} \subset \mathbb{R}^n$ is open,
- (iii) $A \in \mathbb{R}^{p \times n}$ with $\text{rank } A = p < n$
- (iv) Slater condition is satisfied, i.e. there exists a strictly feasible point.

Replace Inequality Constraints by Indicator Function

Define l_- as the indicator function of the nonpositive reals by

$$l_-(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ \infty & \text{if } x > 0. \end{cases}$$

The problem can be reformulated as

$$\begin{aligned} & \text{minimize} && f(x) + \sum_{j=1}^m l_-(g_j(x)) \\ & \text{subject to} && Ax = b, \end{aligned}$$

This new formulation has no inequality constraints, but an unpleasant discontinuity at the boundary of the feasibility set.

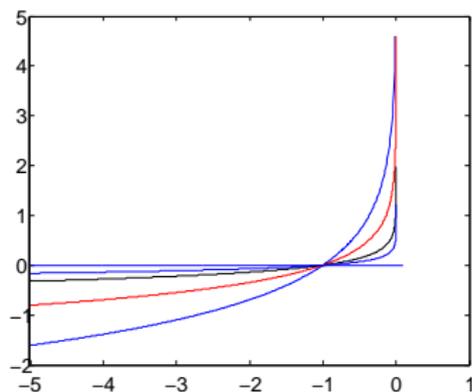
Approximating the Indicator Function

The Approximation

Define I_- as the indicator function of the nonpositive reals by

$$I_-(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ \infty & \text{if } x > 0. \end{cases}$$

$$\hat{l}_t(x) = -\frac{1}{t} \log(-x), \quad x \in (-\infty, 0).$$



A Family of Approximating Optimization Problems

$$\text{minimize } f(x) + \sum_{j=1}^m \hat{l}_t(g_j(x))$$

$$\text{subject to } Ax = b,$$

or equivalently

$$\text{minimize } f(x) + \sum_{j=1}^m -\frac{1}{t} \log(-g_j(x))$$

$$\text{subject to } Ax = b.$$

The Logarithmic Barrier Function

$$\phi(x) := -\sum_{j=1}^m \log(-g_j(x))$$

Central Path

We consider the equivalent problem obtained by multiplying the objective function by t . We thus consider

$$\begin{array}{ll} \text{minimize} & tf(x) + \phi(x) \\ \text{subject to} & Ax = b. \end{array}$$

Definition

Consider the family of optimization problems for $t > 0$. Assume that for each $t > 0$ there is a unique optimal solution $x^*(t)$. The *central path* is defined as the set of optimal points for the problem with parameter t , i.e.

$$\{x^*(t) \mid t > 0\}.$$

The points $x^*(t)$ are called *central points*.

KKT Conditions for the Central Path

the central points are characterized by the necessary and sufficient conditions that they are strictly feasible, i.e.

$$Ax^*(t) = b, \quad g_j(x^*(t)) < 0, \quad j = 1, \dots, m$$

and there exists a $\nu^* \in \mathbb{R}^p$ such that

$$\begin{aligned} 0 &= t \nabla f(x^*(t)) + \nabla \phi(x^*(t)) + A^\top \nu^* \\ &= t \nabla f(x^*(t)) + \sum_{j=1}^m \frac{1}{-g_j(x^*(t))} \nabla g_j(x^*(t)) + A^\top \nu^* . \end{aligned}$$

Lagrange Multipliers

Using strict feasibility of x^* we can define

$$\lambda_j^*(t) = -\frac{1}{tg_j(x^*(t))} > 0,$$

Hence $(\lambda^*(t), \nu^*/t)$ is feasible for the dual problem.

We now divide this equation from the KKT conditions by t

$$0 = t \nabla f(x^*(t)) + \sum_{j=1}^m \frac{1}{-g_j(x^*(t))} \nabla g_j(x^*(t)) + A^\top \nu^* .$$

and obtain

$$\nabla f(x^*(t)) + \sum_{j=1}^m \lambda_j^* \nabla g_j(x^*(t)) + \frac{1}{t} A^\top \nu^* = 0 .$$

Optimality

$x^*(t)$ is a minimizer for the Lagrangian function $L(x, \lambda^*, \nu^*/t)$.

Bounds

As $x^*(t)$ is a minimizer for the Lagrangian function $L(x, \lambda^*, \nu^*/t)$ we have

$$\begin{aligned} g_L(\lambda^*, \nu^*/t) &= f(x^*(t)) + \sum_{j=1}^m \lambda_j^* g_j(x^*(t)) + \frac{1}{t} (\nu^*)^\top (Ax^*(t) - b) \\ &= f(x^*(t)) - \frac{m}{t}, \end{aligned}$$

Bounds

As $x^*(t)$ is a minimizer for the Lagrangian function $L(x, \lambda^*, \nu^*/t)$ we have

$$\begin{aligned} g_L(\lambda^*, \nu^*/t) &= f(x^*(t)) + \sum_{j=1}^m \lambda_j^* g_j(x^*(t)) + \frac{1}{t} (\nu^*)^\top (Ax^*(t) - b) \\ &= f(x^*(t)) - \frac{m}{t}, \end{aligned}$$

Bounding the optimal value

The duality gap between $x^*(t)$ and $g_L(\lambda^*, \nu^*/t)$ is m/t , so

$$f(x^*(t)) - p^* \leq \frac{m}{t}.$$

This shows

$$f(x^*(t)) \rightarrow p^* \quad \text{as} \quad t \rightarrow \infty.$$

KKT

$x^*(t)$ is characterized by the conditions that there exist (λ^*, ν^*) such that

$$g_j(x^*(t)) \leq 0, \quad j = 1, \dots, m$$

$$Ax^*(t) - b = 0$$

$$\lambda_j^* \geq 0 \quad j = 1, \dots, m$$

$$-\lambda_j^* g_j(x^*) = \frac{1}{t} \quad j = 1, \dots, m$$

$$\nabla f(x^*(t)) + \sum_{j=1}^m \lambda_j^* \nabla g_j(x^*(t)) + A^\top \nu^* = 0.$$

KKT

$x^*(t)$ is characterized by the conditions that there exist (λ^*, ν^*) such that

$$g_j(x^*(t)) \leq 0, \quad j = 1, \dots, m$$

$$Ax^*(t) - b = 0$$

$$\lambda_j^* \geq 0 \quad j = 1, \dots, m$$

$$-\lambda_j^* g_j(x^*) = \frac{1}{t} \quad j = 1, \dots, m$$

$$\nabla f(x^*(t)) + \sum_{j=1}^m \lambda_j^* \nabla g_j(x^*(t)) + A^\top \nu^* = 0.$$

The central points are approximately solutions of the KKT conditions.
This approximation improves as $t \rightarrow \infty$

Algorithm

Input Strictly feasible point $x \in \text{dom } f$, $t^0 > 0$, $\mu > 1$ and error bound $\varepsilon > 0$.

Repeat

- 1 Centering step:
Compute $x^*(t)$ by minimizing $tf(x) + \phi(x)$ subject to $Ax = b$
Use Newton's method.
- 2 Update: $x := x^*(t)$
- 3 Stopping criterion: **Quit** if $m/t \leq \varepsilon$.
- 4 Increase t : $t := \mu t$.

Choosing μ : A tradeoff

- μ is close to 1: t^k, t^{k+1} are close to one another. $x^*(t^k)$ may be expected to be a good initial guess for the next problem. The inner Newton algorithm may be expected to converge quickly. But: many outer approximations as t increases rather slowly.
- μ large: Few outer iterations. But: inner iterations may need a large number of steps.

Choosing μ : A tradeoff

- μ is close to 1: t^k, t^{k+1} are close to one another. $x^*(t^k)$ may be expected to be a good initial guess for the next problem. The inner Newton algorithm may be expected to converge quickly. But: many outer approximations as t increases rather slowly.
- μ large: Few outer iterations. But: inner iterations may need a large number of steps.

Choosing t^0

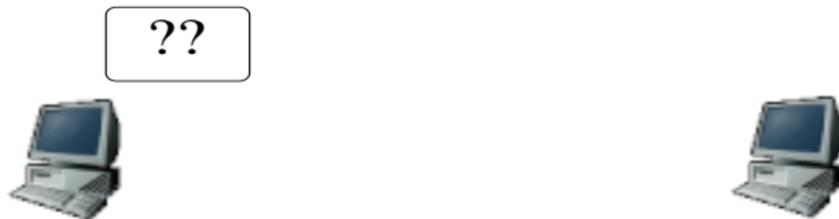
- t^0 large: very likely the resulting optimization problem is numerically unpleasant and the initial Newton algorithm may take considerable time.
- t^0 small: many outer steps again reducing performance.

$$\inf_{\nu} \left\| t \nabla f(x^0) + \nabla \phi(x^0) + A^T \nu \right\|_2$$

- Part I: Convexity and Convex Functions Lectures 1, 2, 3
- Part II: Convex Optimization Lectures 4 and 5
- Part III: Numerical Methods Lectures 6 and 7
- Part IV: Congestion Control Lecture 8
- Part V: Utility Based Congestion Control Lecture 9
- Part VI: Miscellaneous Problems in Networks Lecture 10 (we shall see)

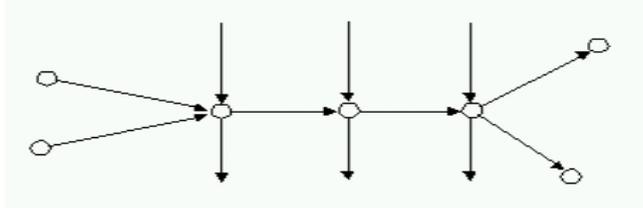
- IV.1: Basics of TCP
- IV.2: Dynamics of Deterministic AIMD
- IV.3: Utility Based Congestion Control

- 4 Modeling of TCP flows
 - The AIMD algorithm
 - A linear model I
 - A linear model II
 - Model Validation
- 5 Analysis of TCP



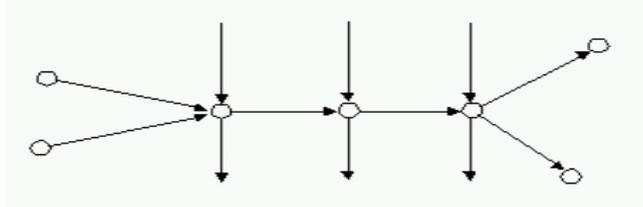
1	2	3	4	5	6
---	---	---	---	---	---	-----	-----

The network consists of a number of sources and sinks that communicate



via network links ('wires') and routers ('queues').

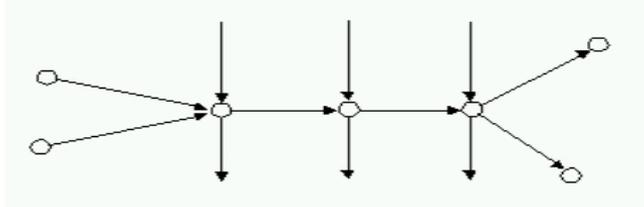
The network consists of a number of sources and sinks that communicate



via network links ('wires') and routers ('queues').

Packets are sent from sources to sinks and 'in flight' packets are on the wire or in the network queues.

The network consists of a number of sources and sinks that communicate



via network links ('wires') and routers ('queues').

Packets are sent from sources to sinks and 'in flight' packets are on the wire or in the network queues.

If packets arrive at a router at a greater rate than the service rate, the queue starts to fill and eventually packets are lost.

Questions:

How much information should be provided to the individual source, so that it can send data in an intelligent way ?

Questions:

How much information should be provided to the individual source, so that it can send data in an intelligent way ?

Desired network properties:

Network fairness and efficient use of network resources.

Questions:

How much information should be provided to the individual source, so that it can send data in an intelligent way ?

Desired network properties:

Network fairness and efficient use of network resources.

- (One) definition of network fairness is that n sources competing for a bandwidth B obtain a share of B/n each.

Questions:

How much information should be provided to the individual source, so that it can send data in an intelligent way ?

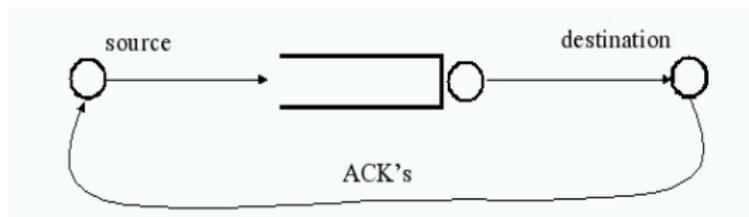
Desired network properties:

Network fairness and efficient use of network resources.

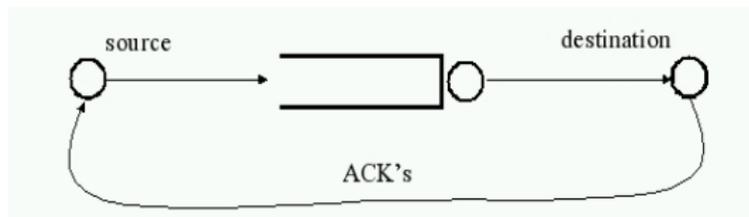
- (One) definition of network fairness is that n sources competing for a bandwidth B obtain a share of B/n each.
- Network resources such as available bandwidth should be fully utilised by the network sources as they becomes available without causing congestion collapse.

TCP is an acknowledgement based protocol that ensures reliable packet delivery and uses “minimal” information.

TCP is an acknowledgement based protocol that ensures reliable packet delivery and uses “minimal” information.



TCP is an acknowledgement based protocol that ensures reliable packet delivery and uses “minimal” information.



Missing acknowledgements are used to infer lost packets and to initiate a resend - many versions of TCP.

Senders receive information about lost packets only after a transport delay. This delay, the RTT, is time-varying due to the nature of the queues.

The AIMD algorithm

AIMD = additive increase, multiplicative decrease.



1	2	3	4	5	6
---	---	---	---	---	---	-----	-----

The AIMD algorithm

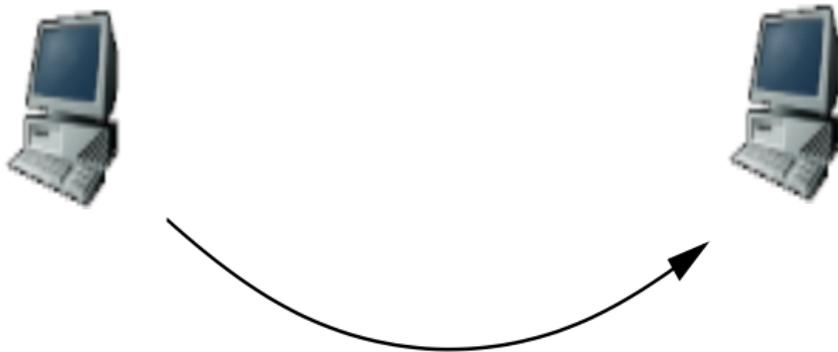
AIMD = additive increase, multiplicative decrease.



1	2
---	---

The AIMD algorithm

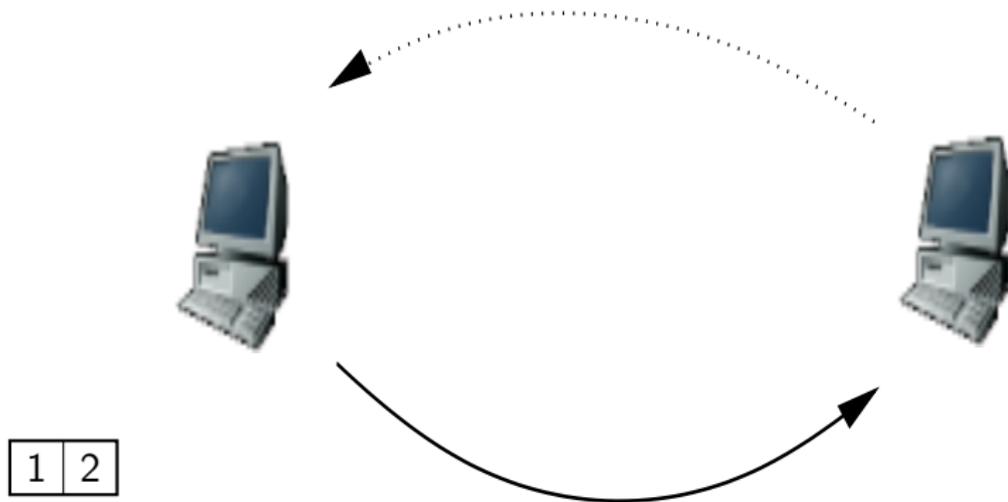
AIMD = additive increase, multiplicative decrease.



1	2
---	---

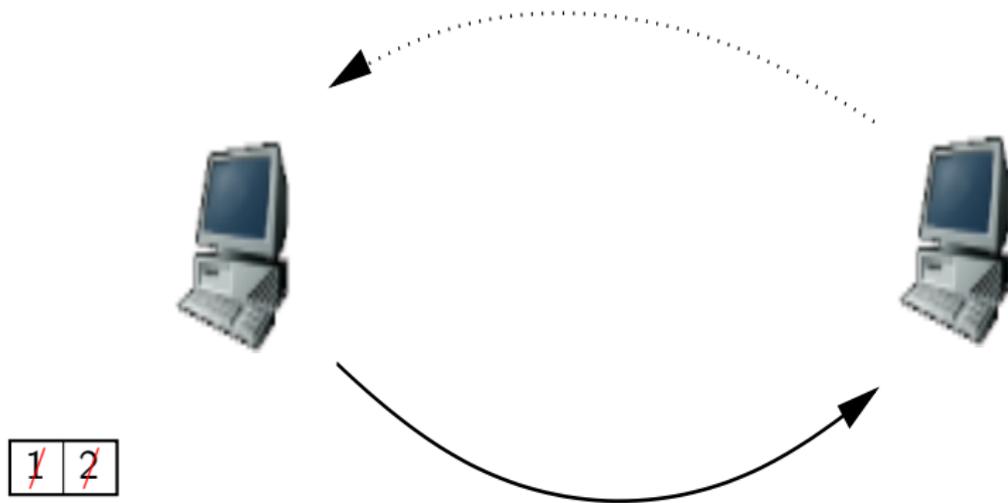
The AIMD algorithm

AIMD = additive increase, multiplicative decrease.



The AIMD algorithm

AIMD = additive increase, multiplicative decrease.



The AIMD algorithm

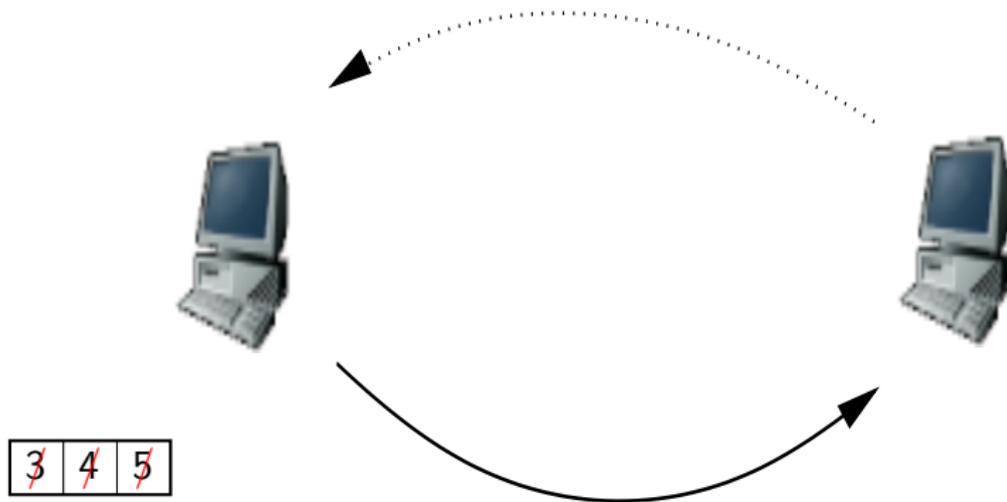
AIMD = additive increase, multiplicative decrease.



3	4	5
---	---	---

The AIMD algorithm

AIMD = additive increase, multiplicative decrease.



The AIMD algorithm

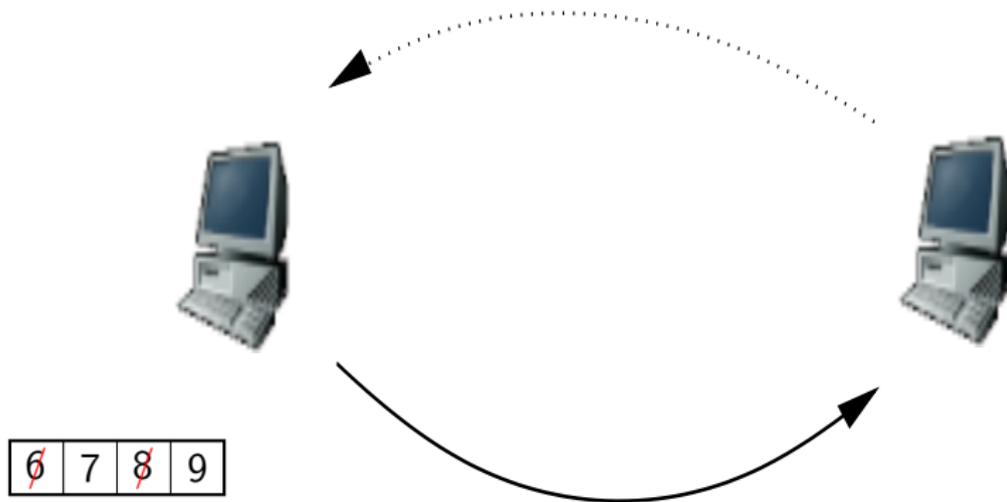
AIMD = additive increase, multiplicative decrease.



6	7	8	9
---	---	---	---

The AIMD algorithm

AIMD = additive increase, multiplicative decrease.



The AIMD algorithm

AIMD = additive increase, multiplicative decrease.



7	9
---	---

Current TCP implementations are characterized by two constants:

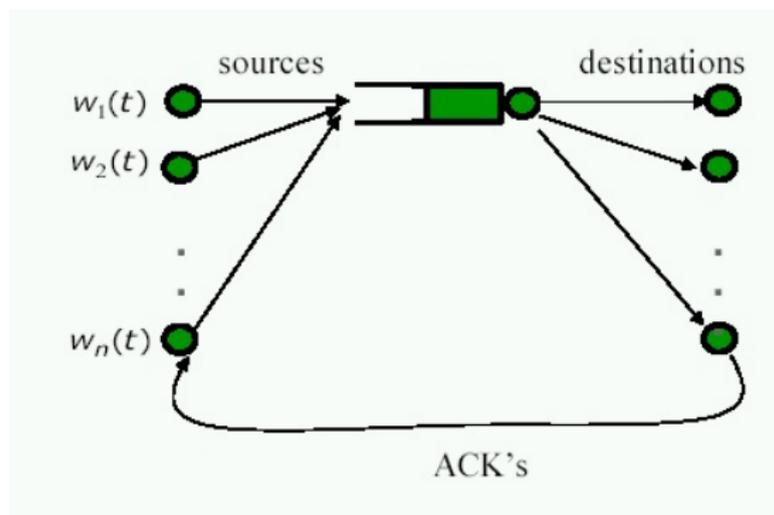
Current TCP implementations are characterized by two constants:

- the additive increase constant $\alpha \geq 1$.
(In your computer $\alpha = 1$.)

Current TCP implementations are characterized by two constants:

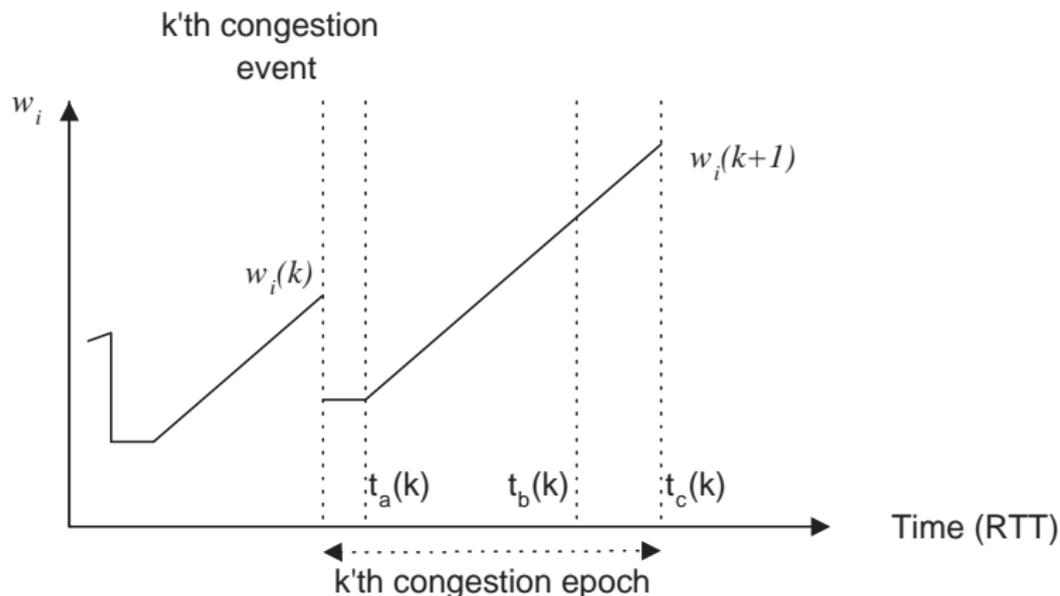
- the additive increase constant $\alpha \geq 1$.
(In your computer $\alpha = 1$.)
- the multiplicative decrease factor $0 < \beta < 1$.
(In your computer $\beta = 1/2$.)

Many sources, one bottleneck

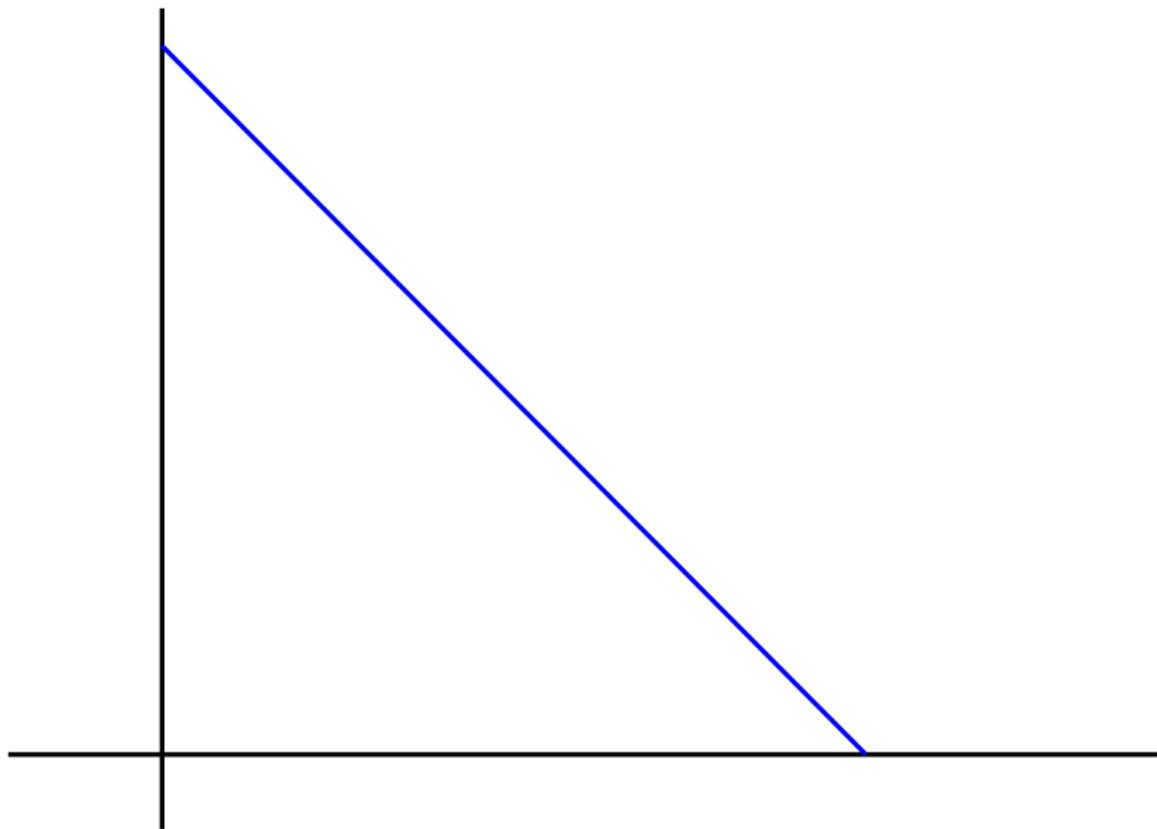


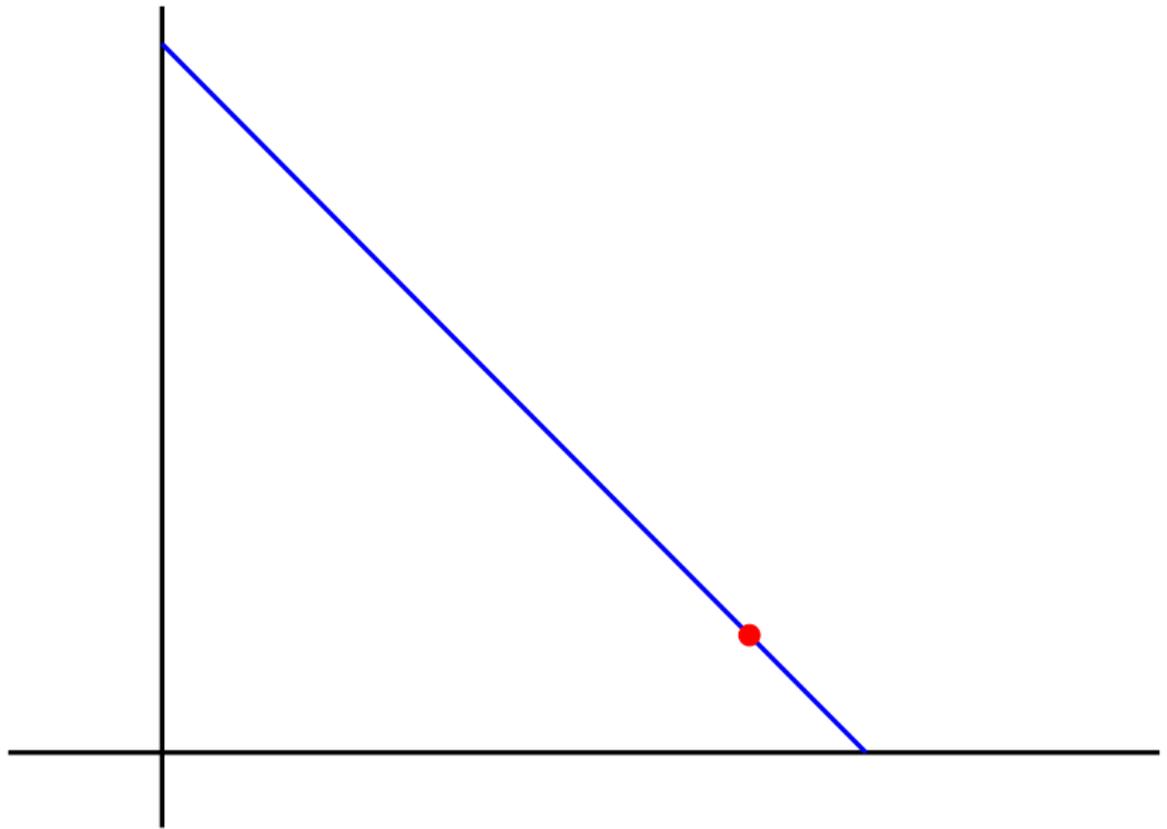
$w_i(k) \hat{=}$ the number of packages sent by the i -th source at the k -th congestion event.

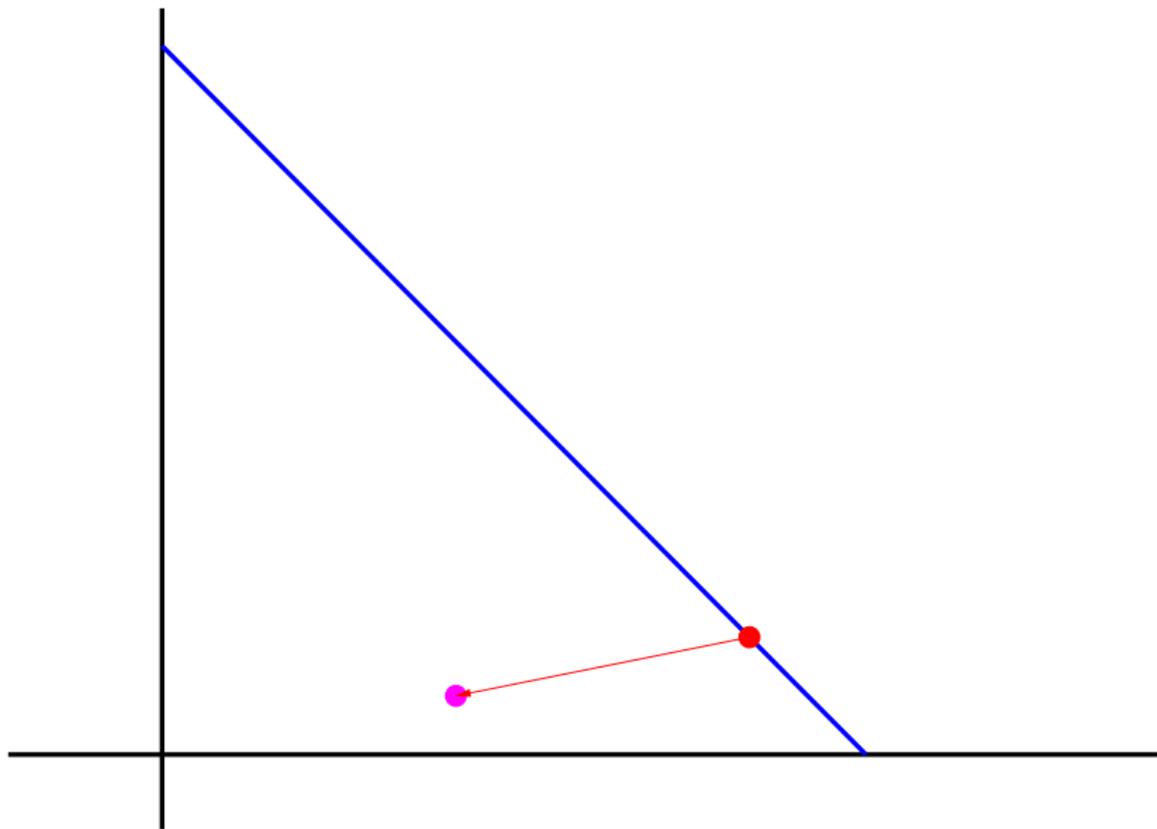
Time evolution for one source

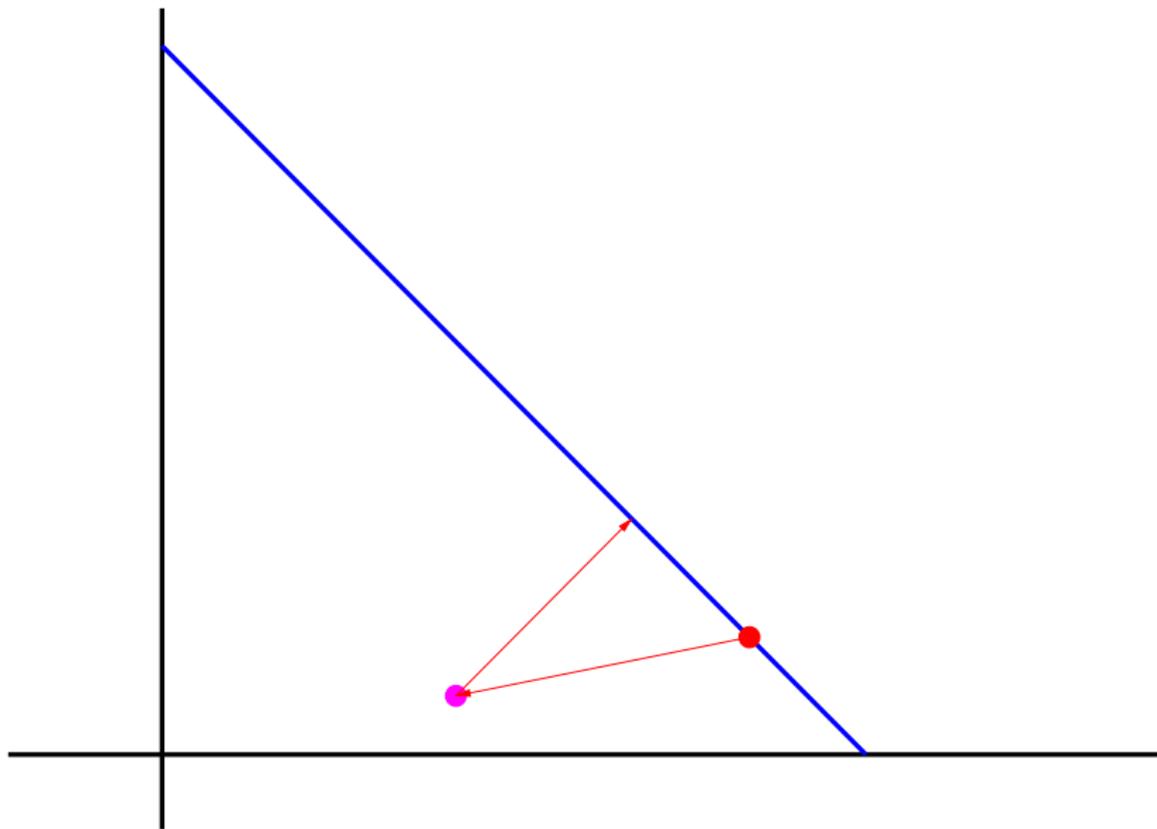


$w_i(k) \hat{=}$ the number of packages sent by the i -th source at the k -th congestion event.

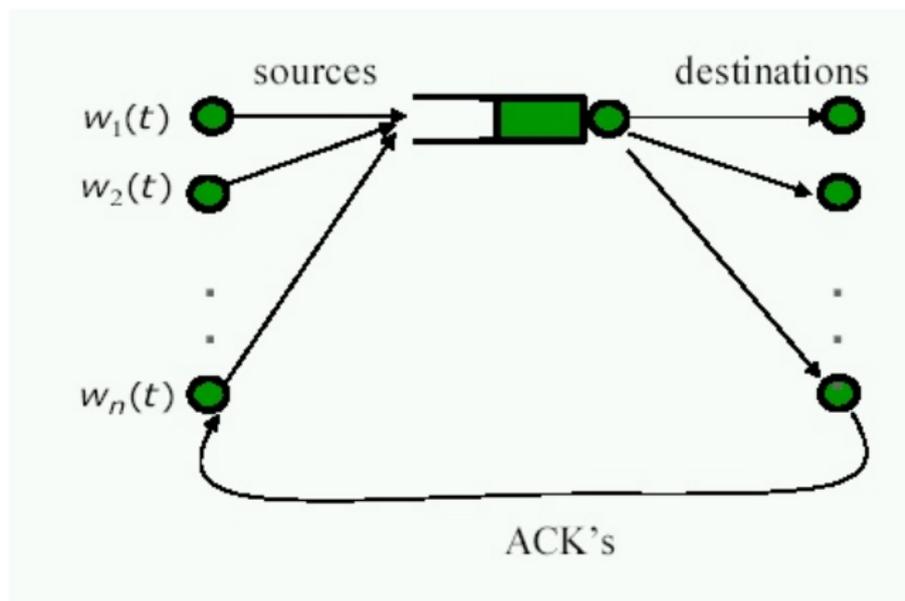








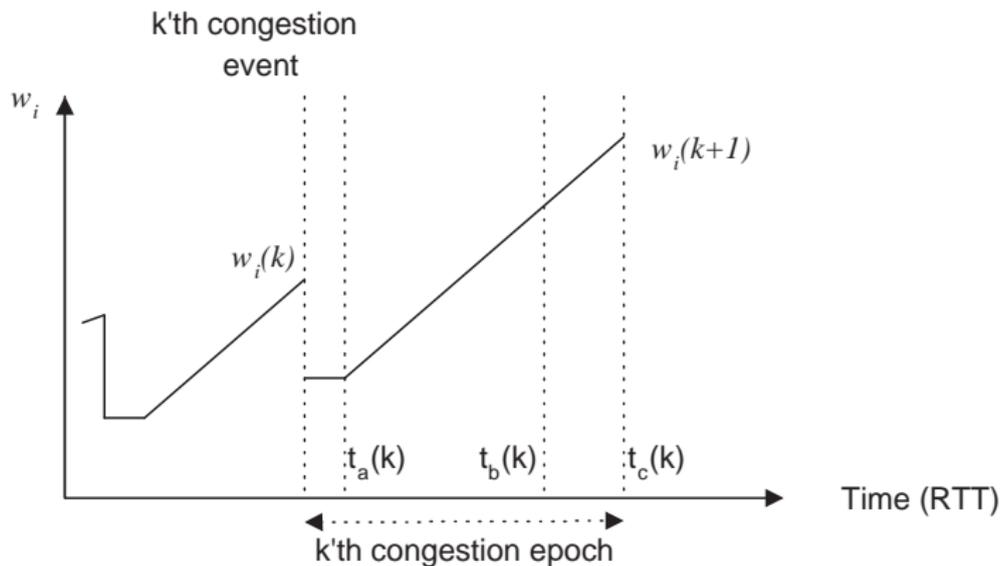
Many sources, one bottleneck



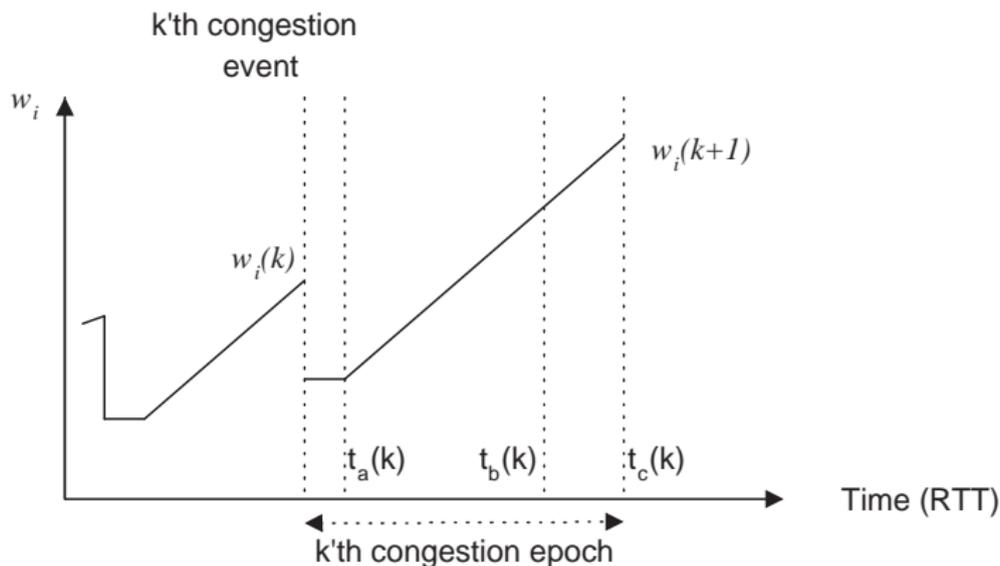
$w_i(k) \hat{=}$ the number of packages sent by the i -th source at the k -th congestion event.

- Congestion occurs in a single bottleneck.
- Congestion is noticed one RTT after it happens.
- Buffer size of bottleneck is small, i.e. RTT can be approximated by a constant.
- ACKs are not lost.
- RTT is the same for all sources.
- The network is synchronized.

Many sources, one bottleneck

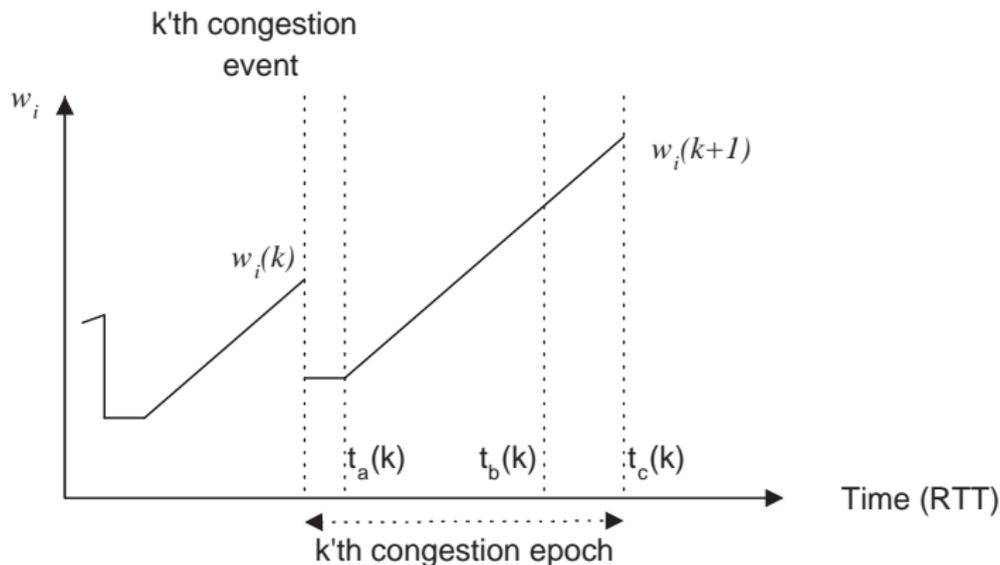


Many sources, one bottleneck



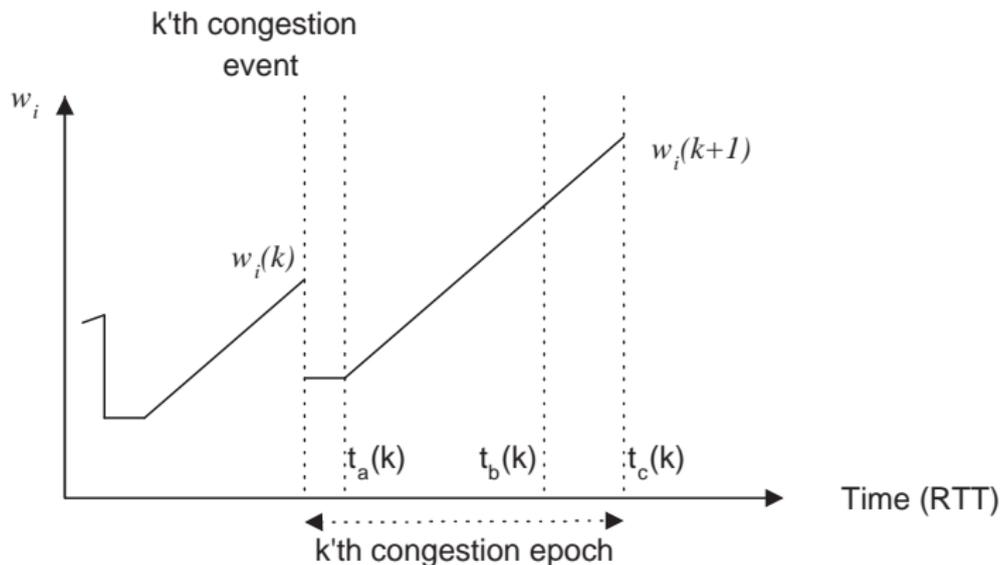
$$w_i(k+1) = \beta_i w_i(k) + \alpha_i (t_c(k) - t_a(k))$$

Many sources, one bottleneck



$$w_i(k+1) = \beta_i w_i(k) + \alpha_i (t_c(k) - t_a(k))$$
$$\sum \alpha_i (t_c(k) - t_a(k)) = P - \sum \beta_i w_i(k) + \sum \alpha_i$$

Many sources, one bottleneck



$$w_i(k+1) = \beta_i w_i(k) + \alpha_i (t_c(k) - t_a(k))$$
$$\sum \alpha_i (t_c(k) - t_a(k)) = P - \sum \beta_i w_i(k) + \sum \alpha_i$$
$$\sum w_i(k) = P + \sum \alpha_i$$