Law of the Jungle in a Linear Network

Mathieu Feuillet

RAP, INRIA Rocquencourt, France

November 20, 2009

Joint work with Thomas Bonald (Orange) and Alexandre Proutière (MSR)

Introduction

- Congestion control
  - Avoidance of packet loss
  - Adaptation of the throughput according to packet loss

Introduction

- Congestion control
  - Avoidance of packet loss
  - Adaptation of the throughput according to packet loss
- No Congestion control
  - No avoidance of packet loss
  - Sources send at their maximum rate: the access rate
  - Recovery mechanism from packet loss (source coding)

Is that sustainable? Is there any congestion collapse?

Law of the Jungle

Stability of a linear network

Theorem of anarchy

Flow Model

Congestion control model introduced by **Massoulié and Roberts** in 2000.

- Network modeled at a flow level, not a packet level
- Flows are going in the network like a fluid
- Users divided in classes to model heterogeneity of the traffic
- Dynamic traffic
- Resource allocation determined by congestion policy

Linear Network

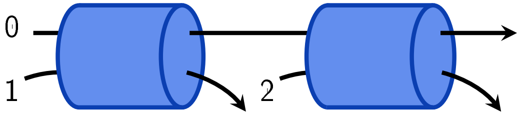

- Two links of capacity 1
- 3 classes of flows
  - Class 0 going through both links with access rate 1
  - Class 1 going through link 1 with access rate 1
  - Class 2 going through link 2 with access rate *a*
- Flow generation rates: λ_i (flows/sec)
- Flow size rates: μ_i (bits^{-1})
- Traffic intensities: $\rho_i = \lambda_i / \mu_i$ (bits/sec)

Resource allocation

Congestion control

- Bandwidth allocation determined by Congestion control algorithm
- Input rate is equal to output rate
- Well studied examples: α-fair allocations (proportional fairness, max-min fairness) (**Mo and Walrand** 2000)

Resource allocation

No congestion control

- Sources send at their maximum rate in the network
- Bandwidth allocation determined by buffer policy in the routers
- Output rate is different of input rate
- Fair policy: Fair Dropping (**Shenker et al** 98)
- Simplest policy: Tail Dropping (Law of the Jungle)

Resource allocation

Tail dropping:

- Sources send at their maximum rate
- At each link, output rates are proportional to input rates

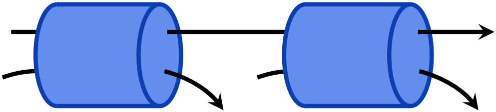

Resource allocation

Tail dropping:

- Sources send at their maximum rate
- At each link, output rates are proportional to input rates

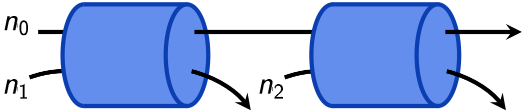

Resource allocation

Tail dropping:

- Sources send at their maximum rate
- At each link, output rates are proportional to input rates

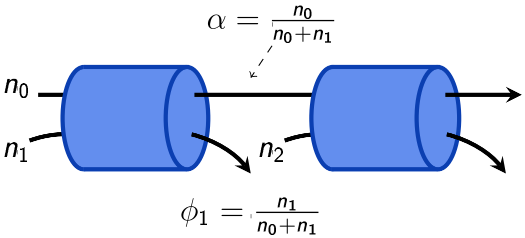

Resource allocation

Tail dropping:

- Sources send at their maximum rate
- At each link, output rates are proportional to input rates

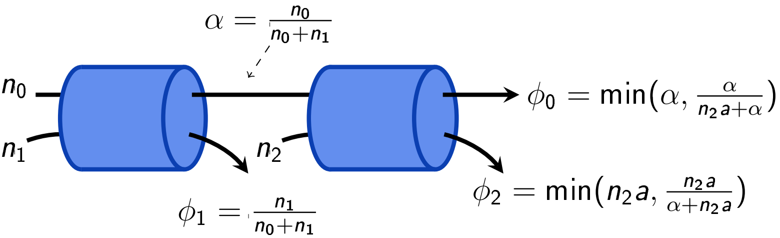

Markov Process (Summary)

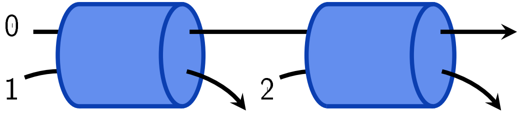

We study the Markov process describing the number of flows in each class:

$$(N_0(t), N_1(t), N_2(t))$$

with transition rates :

$$n \mapsto n + e_i : \lambda_i$$
$$n \mapsto n - e_i : \mu_i \phi_i(n)$$

Performance metrics

- Mean flow response time
  - Ergodicity condition
- Optimal performance
  - The network can be stabilized if and only if

$$\forall l \sum_{i:l\in r_i} \rho_i < C_l$$

  - α-fair policies are optimal (**Bonald and Massoulié** 2001, **de Veciana et al** 2001)
  - Fair dropping is optimal (**Bonald et al** 2009)
  - Tail Dropping is not optimal

Law of the Jungle

Stability of a linear network

Theorem of anarchy

Stability of class 2

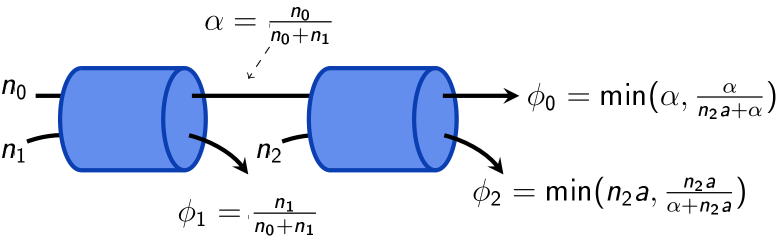

We have $0 \leq \alpha \leq 1$. Thus, class 2 is stable if $\rho_2 < 1$ and transient if $\rho_2 > 1$ whatever the conditions on 0 and 1.
Now, we suppose that

$$\rho_2 < 1$$

Stability of class 2

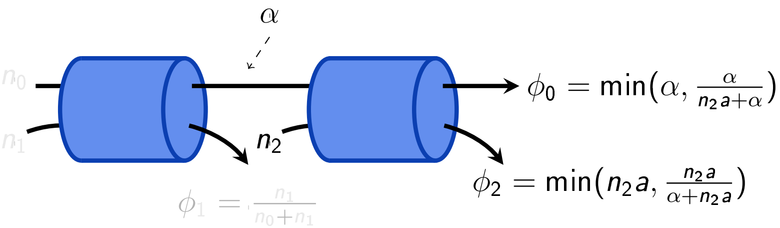

We freeze the number of flows in classes 0 and 1 and then α.
There is a stationary distribution π^α for class 2.
We then define the averaged throughput of class 0:

$$\bar{\phi}_0(\alpha) = \sum_{n_2 \in \mathbb{N}} \pi^\alpha(n_2) \min\left(\alpha, \frac{\alpha}{n_2 a + \alpha}\right)$$

Fluid limit of the system

Class 2 is stable, so there is no need to scale it! We perform the scaling only on classes 0 and 1:

$$(N_0(0), N_1(0)) = n_k, \quad \lim_{k\to\infty} \|n_k\| = \infty,$$

$$\lim_{k\to\infty} \frac{1}{\|n_k\|}(N_0(\|n_k\|t), N_1(\|n_k\|t)) \stackrel{d}{=} (z_0(t), z_1(t))$$

and $(z_0(t), z_1(t))$ satisfies:

$$\dot{z}_0(t) = \lambda_0 - \mu_0 \bar{\phi}_0 \left(\frac{z_0(t)}{z_0(t) + z_1(t)} \right)$$

$$\dot{z}_1(t) = \lambda_1 - \mu_1 \frac{z_1(t)}{z_0(t) + z_1(t)}$$

The proof is similar to the one of **Hunt and Kurtz** in 1994 about Loss Networks.

Fluid Limit of the system

$$\dot{z}_0(t) = \lambda_0 - \mu_0 \bar{\phi}_0 \left(\frac{z_0(t)}{z_0(t) + z_1(t)} \right)$$

$$\dot{z}_1(t) = \lambda_1 - \mu_1 \frac{z_1(t)}{z_0(t) + z_1(t)}$$

There is a separation of time scale between the fluid limit (z_0, z_1) and class 2: class 2 is always at equilibrium and there is an averaging on class 2 for fluid limit z_0.

Fluid Limit dynamics

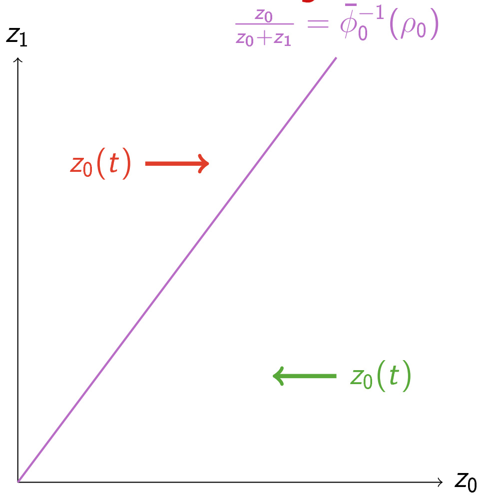

$$\dot{z}_0 = \lambda_0 - \mu_0 \bar{\phi}_0 \left(\frac{z_0}{z_0 + z_1} \right)$$

Fluid Limit dynamics

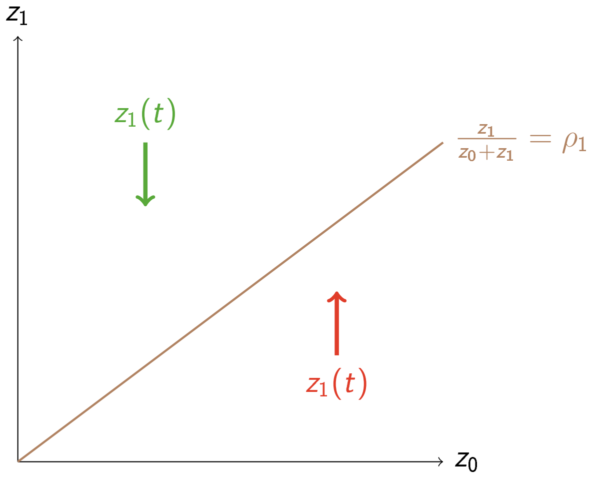

$$\dot{z}_1 = \lambda_1 - \mu_1 \frac{z_1}{z_0 + z_1}$$

Fluid Limit dynamics

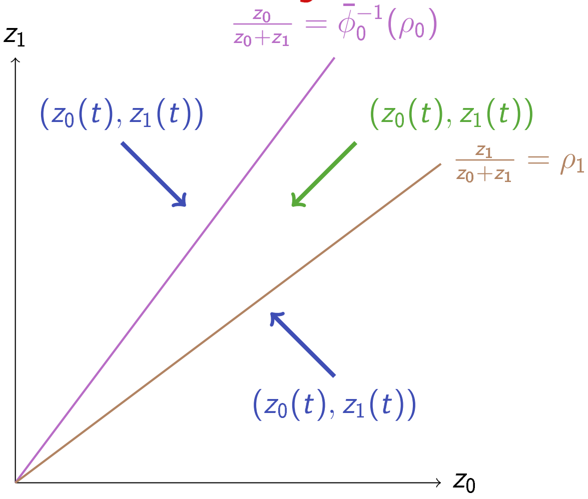

$$\rho_0 < \bar{\phi}_0(1 - \rho_1)$$

Fluid Limit dynamics

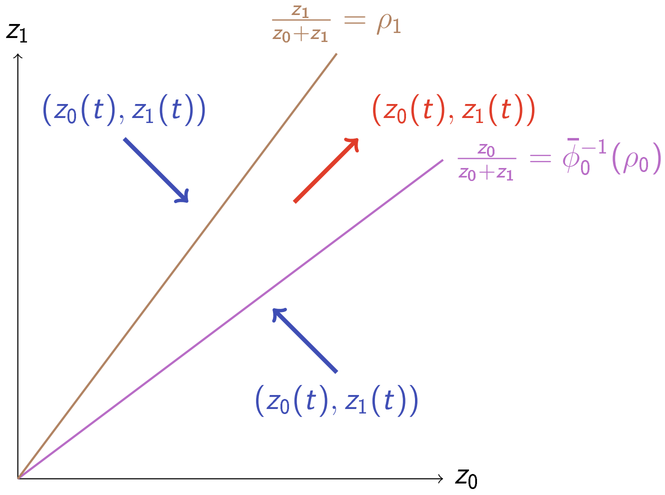

$$\rho_0 > \bar{\phi}_0(1 - \rho_1)$$

Stability Conditions

The exact conditions for stability under the Law of the Jungle are:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \bar{\phi}_0(1 - \rho_1)$$

The optimal stability conditions are:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \min(1 - \rho_1, 1 - \rho_2)$$

But:

$$\bar{\phi}_0(1 - \rho_1) < \min(1 - \rho_2, 1 - \rho_1)$$

The stability conditions are not optimal!

Law of the Jungle

Stability of a linear network

Theorem of anarchy

What is the anarchy?

The price of anarchy quantifies how far the stability conditions are from optimal ones.
In our case, the definition is simply:

$$P(a) = \max_{\rho_1, \rho_2} \left(\min(1 - \rho_1, 1 - \rho_2) - \bar{\phi}_0(1 - \rho_1) \right)$$

The price of anarchy depends on a, the access rate of class 2.

Theorem of anarchy

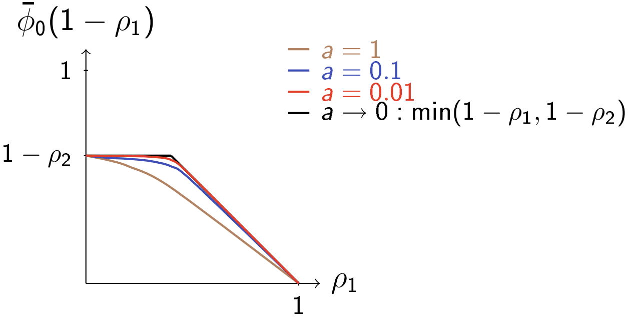

Theorem:

$$\lim_{a \to 0} P(a) = 0$$

For any (ρ_0, ρ_1, ρ_2) satisfying the optimal stability conditions, there exists a small enough such that the network is stable.

Conclusion

- Do we need congestion control in the Internet?
- Evaluation of the impact of big clients on the network (optical networks).

- A fluid limit with an interesting averaging phenomenon
- Can be extended to linear networks with more than two links
- In other contexts?

- Conjecture: The theorem of anarchy is true in acyclic networks.