Dynamics and Stability of AIMD Algorithms

Fabian Wirth

The Hamilton Institute, NUI Maynooth

Hamilton Institute Workshop on Congestion Control
September 27–28, 2005

based on joint work with:
Abraham Berman, Christopher King, Douglas Leith, Robert Shorten, Rade Stanojević

Many sources, one bottleneck

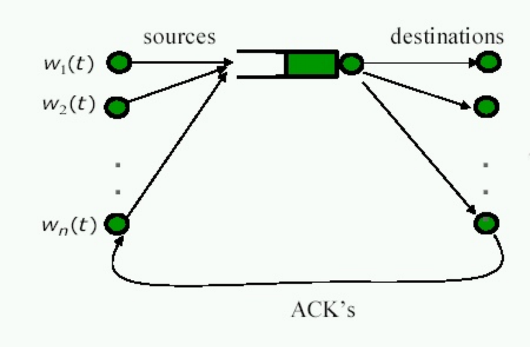

$w_i(k) \triangleq$ the number of packages sent by the i-th source at the k-th congestion event.

Time evolution for one source

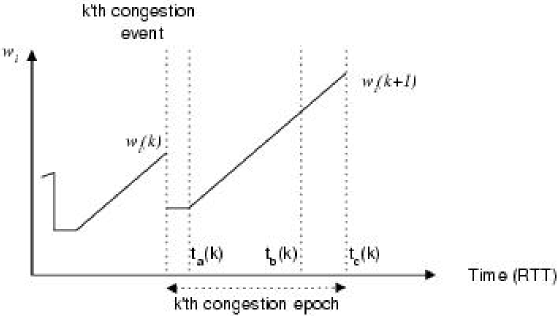

$w_i(k) \hat{=}$ the number of packages sent by the i-th source at the k-th congestion event.

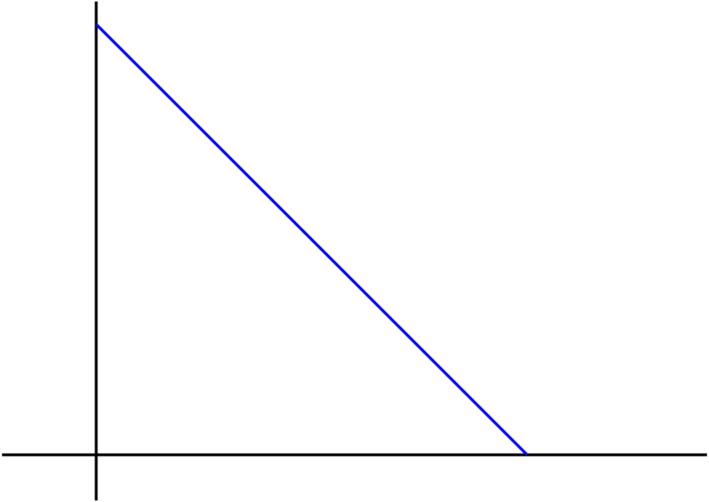

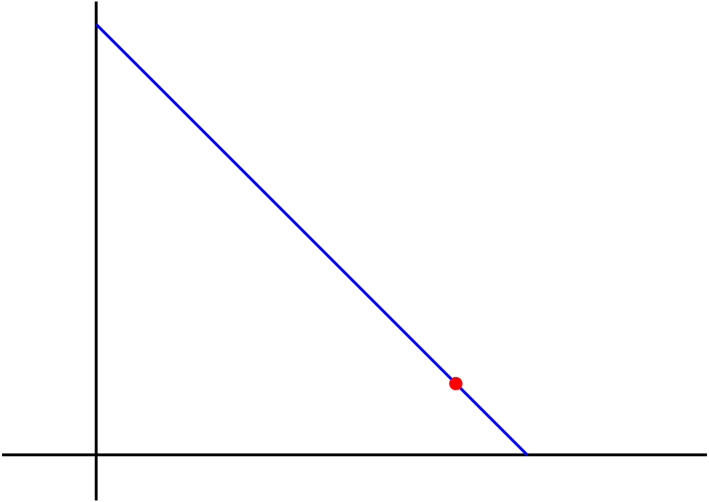

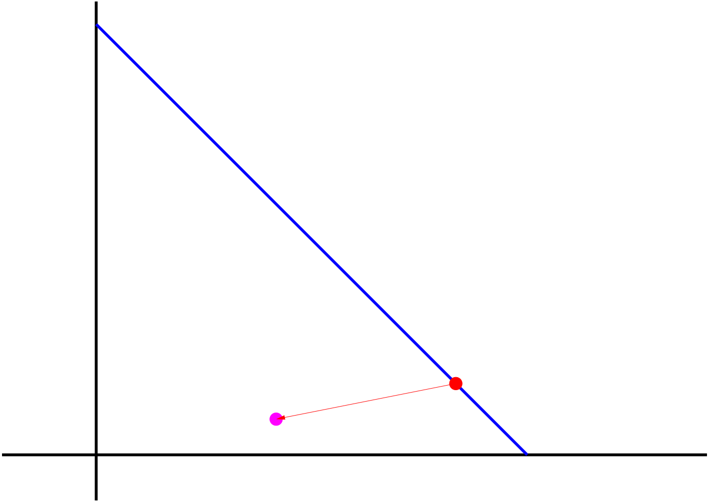

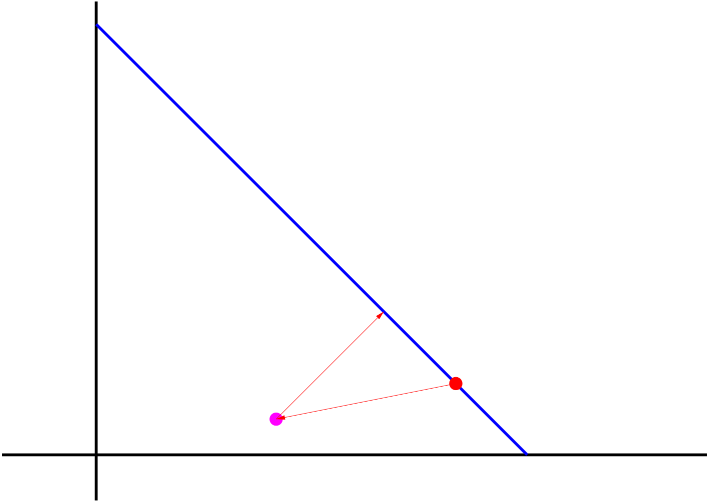

Many sources, one bottleneck

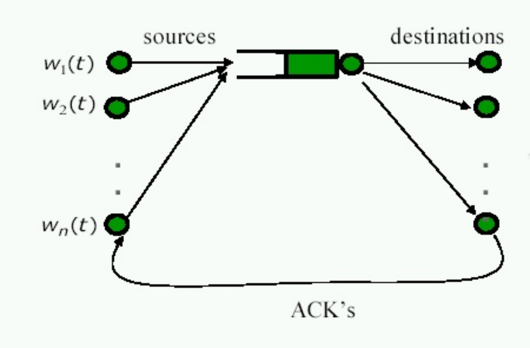

$w_i(k) \triangleq$ the number of packages sent by the i-th source at the k-th congestion event.

Assumptions I

- 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.
- RTT is the same for all sources.
- The network is synchronized.

Many sources, one bottleneck

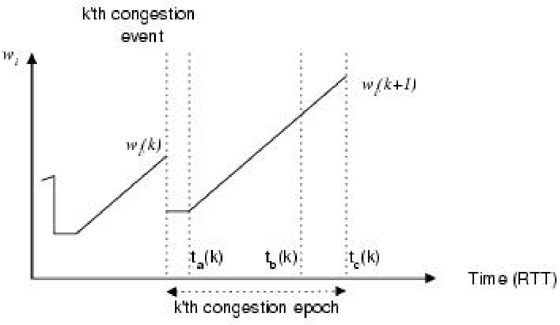

$$w_i(k+1) = \beta_i w_i(k) + \alpha_i(t_c(k) - t_a(k))$$
$$\sum \alpha_i(t_c(k) - t_a(k)) = P - \sum \beta_i w_i(k) + \sum \alpha_i$$
$$\sum w_i(k) = P + \sum \alpha_i$$

A linear model

A little manipulation shows that

$$w(k+1) = Aw(k)\,, \quad k \geq 1$$

where A is given by

$$\begin{bmatrix} \beta_1 & 0 & \dots & 0 \\ 0 & \beta_2 & & \vdots \\ \vdots & & \ddots & \\ 0 \dots & 0 & & \beta_n \end{bmatrix} + \frac{1}{\sum_i \alpha_i} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \dots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1-\beta_1 & \cdots & 1-\beta_n \end{bmatrix}.$$

Here α_i is the additive increase parameter of the i-th source and β_i is the multiplicative decrease parameter.

Digression: Different RTT's

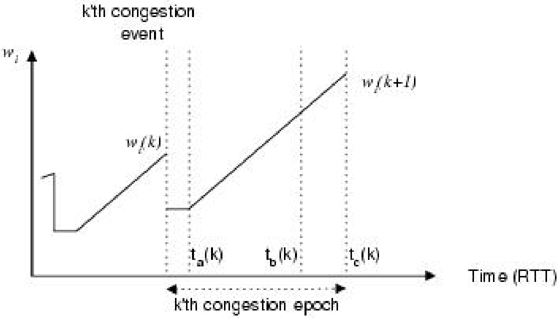

$$w_i(k+1) = \beta_i w_i(k) + \frac{\alpha_i}{\mathrm{RTT}_i}(t_c(k) - t_a(k))$$
$$\sum \frac{\alpha_i}{\mathrm{RTT}_i}(t_c(k) - t_a(k)) = P - \sum \beta_i w_i(k) + \sum \frac{\alpha_i}{\mathrm{RTT}_i}$$
$$\sum w_i(k) = P + \sum \frac{\alpha_i}{\mathrm{RTT}_i}$$

Digression: Different RTT's

A little manipulation leads to a linear equation

$$w(k+1) = \tilde{A} w(k) \,.$$

The similarity transformation

$$\hat{w} := \begin{bmatrix} \mathrm{RTT}_1^{-1} & 0 & \dots & 0 \\ 0 & \mathrm{RTT}_2^{-1} & \dots & 0 \\ & & \ddots & \\ 0 & \dots & & \mathrm{RTT}_n^{-1} \end{bmatrix} w$$

leads to an evolution equation for $\hat{w}$ of the same form as in the case for equal RTT's.

A linear model

$$w(k+1) = Aw(k)\,, \quad k \geq 1$$

where A is given by

$$\begin{bmatrix} \beta_1 & 0 & \dots & 0 \\ 0 & \beta_2 & & \vdots \\ \vdots & & \ddots & \\ 0\dots & 0 & & \beta_n \end{bmatrix} + \frac{1}{\sum_i \alpha_i} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \dots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1-\beta_1 & \cdots & 1-\beta_n \end{bmatrix}.$$

Here α_i is the additive increase parameter of the i-th source and β_i is the multiplicative decrease parameter.

The Synchronized Case

Assumption: Every source experiences all congestions

$$A = \begin{bmatrix} \beta_1 & 0 & \dots & 0 \\ 0 & \beta_2 & & \vdots \\ \vdots & & \ddots & \\ 0 \dots & 0 & & \beta_n \end{bmatrix} + \frac{1}{\sum_i \alpha_i} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \dots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1 - \beta_1 & \cdots & 1 - \beta_n \end{bmatrix}.$$

The matrix A is positive and column stochastic. Thus by the Perron-Frobenius Theorem the evolution of A^k is very well understood.

The Synchronized Case

$$A = \begin{bmatrix} \beta_1 & 0 & \dots & 0 \\ 0 & \beta_2 & & \vdots \\ \vdots & & \ddots & \\ 0\dots & 0 & & \beta_n \end{bmatrix} + \frac{1}{\sum_i \alpha_i} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \dots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1-\beta_1 & \cdots & 1-\beta_n \end{bmatrix}.$$

Theorem (Berman, Shorten, Leith)

1. A has an eigenvalue one with eigenvector
$$v = \begin{bmatrix} \frac{\alpha_1}{1-\beta_1} & \cdots & \frac{\alpha_n}{1-\beta_n} \end{bmatrix}.$$
2. For all initial conditions w_0 we have
$$A^k w_0 \to \theta v .$$
3. the rate of convergence is exponential, bounds for this rate can be given in terms of the β_i.

Assumptions II

- 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.
- RTT is the same for all sources.
- ~~The network is synchronized.~~

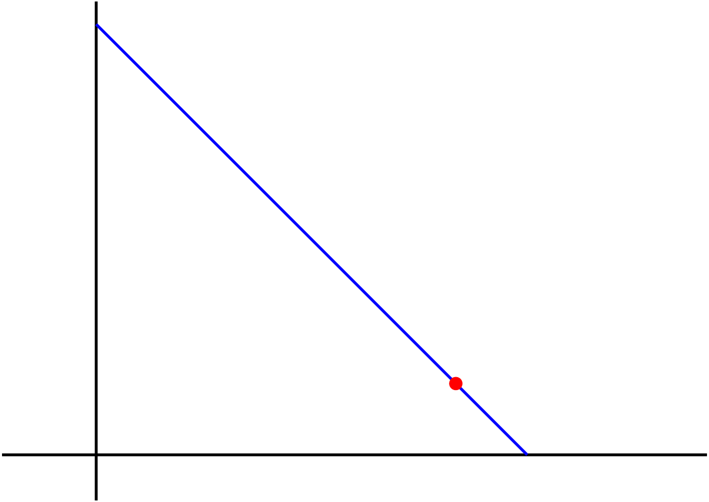

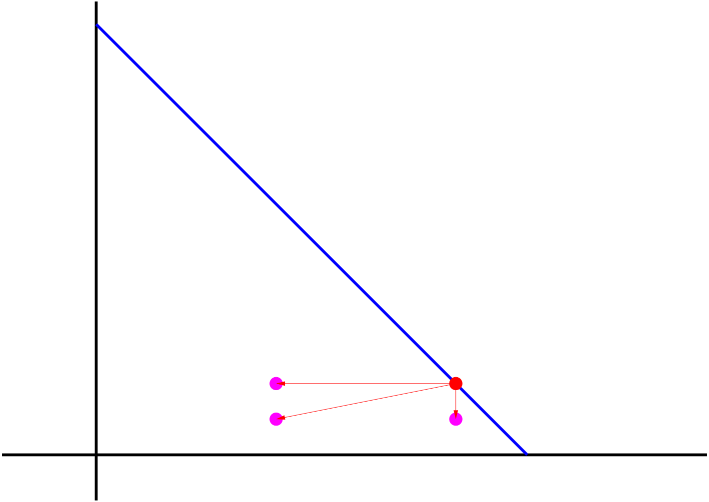

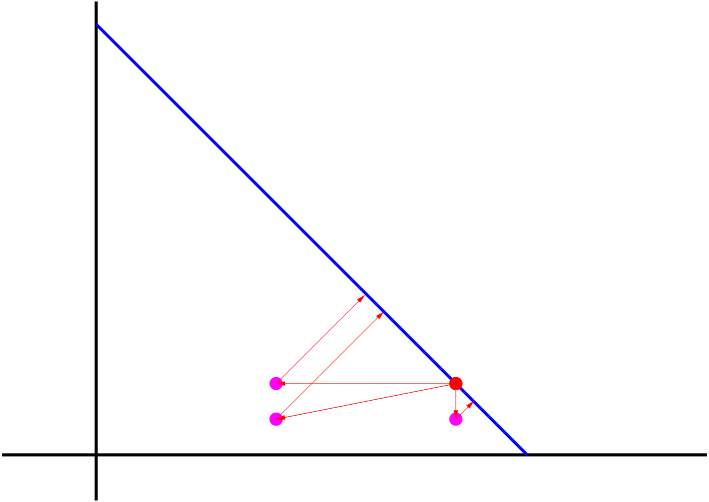

The unsynchronized case

$$w(k+1) = A(k)w(k)\,, \quad k \geq 1$$

where $A(k)$ is given by

$$\begin{bmatrix} \beta_1(k) & 0 & \dots & 0 \\ 0 & \beta_2(k) & & \vdots \\ \vdots & & \ddots & \\ 0\dots & 0 & & \beta_n(k) \end{bmatrix} + \frac{1}{\sum_i \alpha_i} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \dots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1-\beta_1(k) & \cdots & 1-\beta_n(k) \end{bmatrix}$$

Here α_i is the additive increase parameter of the i-th source and $\beta_i(k)$ is equal to the multiplicative decrease parameter or 1, depending on whether the $i-th$ source experiences congestion at the k-th congestion event or not.

The unsynchronized case

In the analysis of the dynamics of TCP flows we are led to the consideration of a linear inclusion of the form

$$w(k+1) \in \{Aw(k) \mid A \in \mathcal{M}\}\,,$$

where $\mathcal{M} \subset \mathbb{R}^{n\times n}$ is the set of 2^{n-1} matrices obtained by setting $\beta_i, i = 1, \dots, n$ either to 1 or to a constant in $(0,1)$.
(Note that the identity is omitted.)

Model Validation I

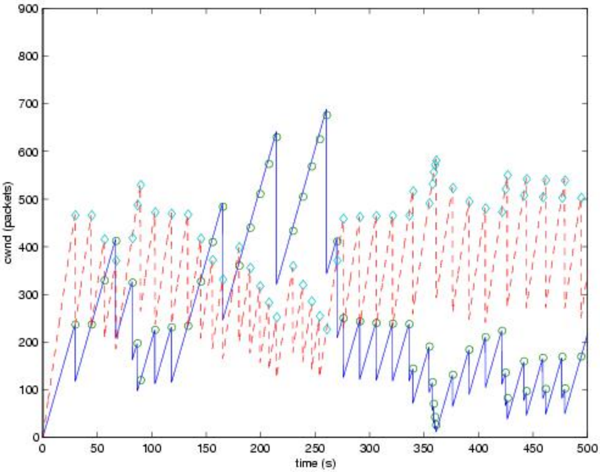

Application to the TCP model

The matrices modeling TCP flows were of the form

$$\begin{bmatrix} \beta_1(k) & 0 & \dots & 0 \\ 0 & \beta_2(k) & & \vdots \\ \vdots & & \ddots & \\ 0\dots & 0 & & \beta_n(k) \end{bmatrix} + \frac{1}{\sum_i \alpha_i} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \dots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1-\beta_1(k) & \cdots & 1-\beta_n(k) \end{bmatrix}$$

The matrices are all column stochastic.
The set $\mathcal{M}$ is irreducible on the invariant subspace

$$S := \begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}^{\perp}.$$

Weak ergodicity

Theorem Let $\{A(k)\}_{k\in\mathbb{N}} \subset \mathcal{M}^{\mathbb{N}}$ be a sequence with the property that at most one source does not see a drop infinitely often, then

$$\lim_{k\to\infty} A(k)A(k-1)\dots A(0)_{|S} = 0\,.$$

In other words, the sequence $\{A(k)\}_{k\in\mathbb{N}}$ is weakly ergodic.
This result extends results of Leizarowitz (LAA, 2000) in this special case.
Proof builds on properties of extremal norms

A fairness result

Assume now that the matrices $A(k)$ are i.i.d. random variables. These random variables induce for the i-th source a probability

$\lambda_i := P($ source i experiences congestion at the k-th congestion event$)$.

Theorem
If for each $i = 1, \dots, n$ we have

$$\lambda_i > 0\,,$$

then, almost surely, (for the right γ)

$$\lim_{k\to\infty} \frac{1}{k+1} \sum_{l=0}^{k} w(l) = \gamma \begin{bmatrix} \frac{\alpha_1}{\lambda_1(1-\beta_1)} & \cdots & \frac{\alpha_n}{\lambda_n(1-\beta_n)} \end{bmatrix}^T .$$

Model Validation II, NS Simulation

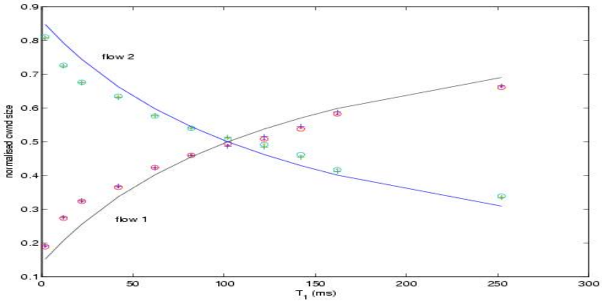

Variation of mean of $w_1(k), w_2(k)$ with propagation delay T_1.
Key: $+NS$ simulation result; $\cdot$ mathematical model; $\circ$ Theorem on long time-averages; solid lines correspond to synchronised case.

Model Validation III, Data from a Real Network

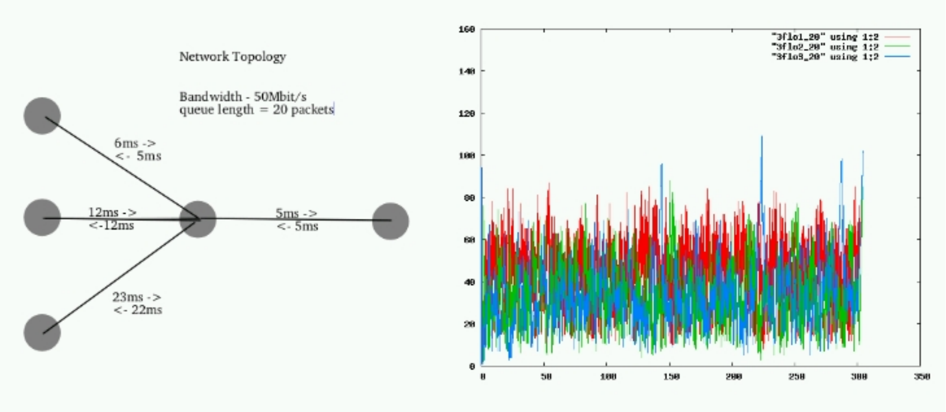

Model Validation III, Data from a Real Network

Queue		rtt=21ms	rtt=34ms	rtt=55ms
20	Measurement	0.4055	0.3014	0.2931
	Theorem 3.3	0.4054	0.3061	0.2886
	% difference	0.0247	1.5594	1.5353
40	Measurement	0.4122	0.2849	0.3029
	Theorem 3.3	0.4121	0.2915	0.2964
	% difference	0.0243	2.3166	2.1459
60	Measurement	0.4024	0.3093	0.2882
	Theorem 3.3	0.4204	0.3087	0.2709
	% difference	4.4732	0.1940	6.0028
80	Measurement	0.416	0.3293	0.2547
	Theorem 3.3	0.3835	0.2891	0.3274
	% difference	7.8125	12.2077	28.5434

A Markov model

We are interested in the system

$$W(k+1) = A(k)W(k)\,, \quad A(k) \in \mathcal{M} := \{M_1, \ldots, M_{2^n-1}\}\,.$$

All matrices A_i are column stochastic, and window sizes are non-negative.
We may thus restrict our attention to the simplex

$$\Sigma := \left\{ x \in \mathbb{R}^n_+ \mid \sum_{i=1}^{n} x_i = 1 \right\} ,$$

We may thus restrict our attention to the simplex

$$\Sigma := \left\{ x \in \mathbb{R}^n_+ \mid \sum_{i=1}^{n} x_i = 1 \right\} ,$$

Assume the random variables $A(k), k = 0, 1, \ldots$ are i.i.d. and

$$P(A(k) = M_i) = \rho_i \, , \quad i = 1, \ldots, \mu \, .$$

Assume also the "everyone drops" condition

$$\lambda_j = \sum_{\beta_j(M_i) < 1} \rho_i > 0 \, ,$$

The assumptions lead to a Markov process on Σ. The transition kernel of that process acts on continuous functions $h : \Sigma \to \mathbb{R}$ through

$$Ph(x) = \int_\Sigma h(y) P(x, dy) = \sum_{i=1}^{\mu} \rho_i h(A_i x) \, .$$

For each continuous h the sequence

$$P^k h \, , \quad k \in \mathbb{N} \, ,$$

is equicontinuous.
Markov chains with this property are called e-chains.

Theorem(Meyn & Tweedie)
Consider an e-chain on a compact space Σ. If there is a positive and aperiodic state $z \in \Sigma$, then

(i) there exists a unique invariant probability π,

(ii) for every $x \in \Sigma$ and every continuous function $h : \Sigma \to \mathbb{R}$ we have that if $W(0) = x$, then

$$\lim_{k \to \infty} \frac{1}{k} \sum_{j=0}^{k-1} h(W(j)) = \int_{\Sigma} h(y) d\pi(y) , \qquad \text{in probability,}$$

(iii) for every $x \in \Sigma$ and every continuous function $h : \Sigma \to \mathbb{R}$ we have

$$\int_{\Sigma} h(y) P^k(x, dy) \to \int_{\Sigma} h(y) d\pi(y) , \qquad \text{as } k \to \infty .$$

Positive and Aperiodic States

$$\mathcal{L} := \{\{A_k\}_{k\in\mathbb{N}} \in \mathcal{M}^{\mathbb{N}} \mid \{A_k\}_{k\in\mathbb{N}} \text{ each source sees infinitely many drops}\}$$

$$\mathcal{R}_{\mathcal{L}} := \{R \mid \operatorname{rank} R = 1,\ \exists \{A_k\}_{k\in\mathbb{N}} \in \mathcal{L}, k_l \to \infty : \lim_{l\to\infty} \Pi(k_l) = R\}\,.$$

$$\mathcal{C} := \{z \in \Sigma \mid zy^T \in \mathcal{R}_{\mathcal{L}}\}\,.$$

Positive and Aperiodic States

Proposition

(i) $\mathcal{C}$ is compact and forward invariant,

(ii) for any $z \in \mathcal{C}$ and any open neighborhood $U \subset \Sigma$ of z there is a $k_0 > 0$ such that $P^k(x, U) > \delta > 0$, for all $k \geq k_0$ and all $x \in \Sigma$,

(iii) For any initial condition $W_0 \in \Sigma$ we have almost surely

$$\lim_{k \to \infty} \operatorname{dist}(W(k), \mathcal{C}) = 0 \,.$$

By (ii)& (iii): $\mathcal{C}$ is the set of positive, aperiodic states.
$\mathcal{C}$ is the support of the unique invariant probability measure π.

$$W(k+1) = A(k)W(k)\,.$$

Theorem

Consider the Markov chain W on Σ. For this chain there is a positive and aperiodic state $z \in \Sigma$, and so

(i) there exists a unique invariant probability π,

(ii) for every $x \in \Sigma$ and every continuous function $h : \Sigma \to \mathbb{R}$ we have that if $W(0) = x$, then

$$\lim_{k\to\infty} \frac{1}{k} \sum_{j=0}^{k-1} h(W(j)) = \int_{\Sigma} h(y)d\pi(y)\,, \qquad \text{in probability,}$$

What does the support look like ?

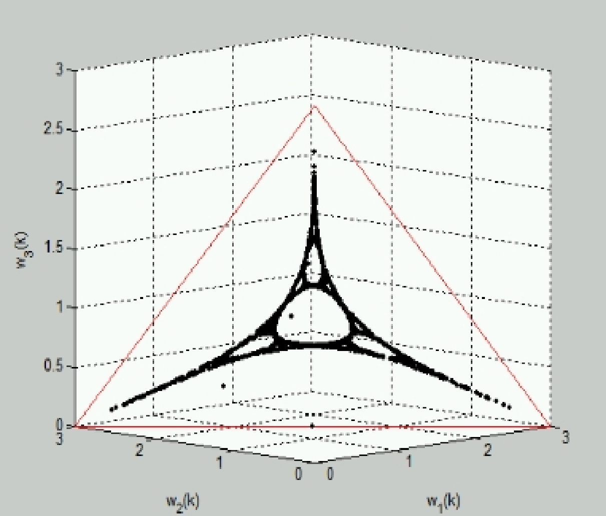

Thank you !