The Coherence Phase Transition

Content

- Multiple type particle systems
- Mean-field-type graphs
- Poisson Hypothesis
- Cases of validity: low load (=high T), ...
- Violation of PH: phase transitions at high load (=low T)
- Proofs: Non-linear Markov processes, Fluid networks, stable attractors, convergence, ...
- Joint work with Alexandre Rybko and Alexandre Vladimirov

Multiple type particle systems

We have a network of servers:

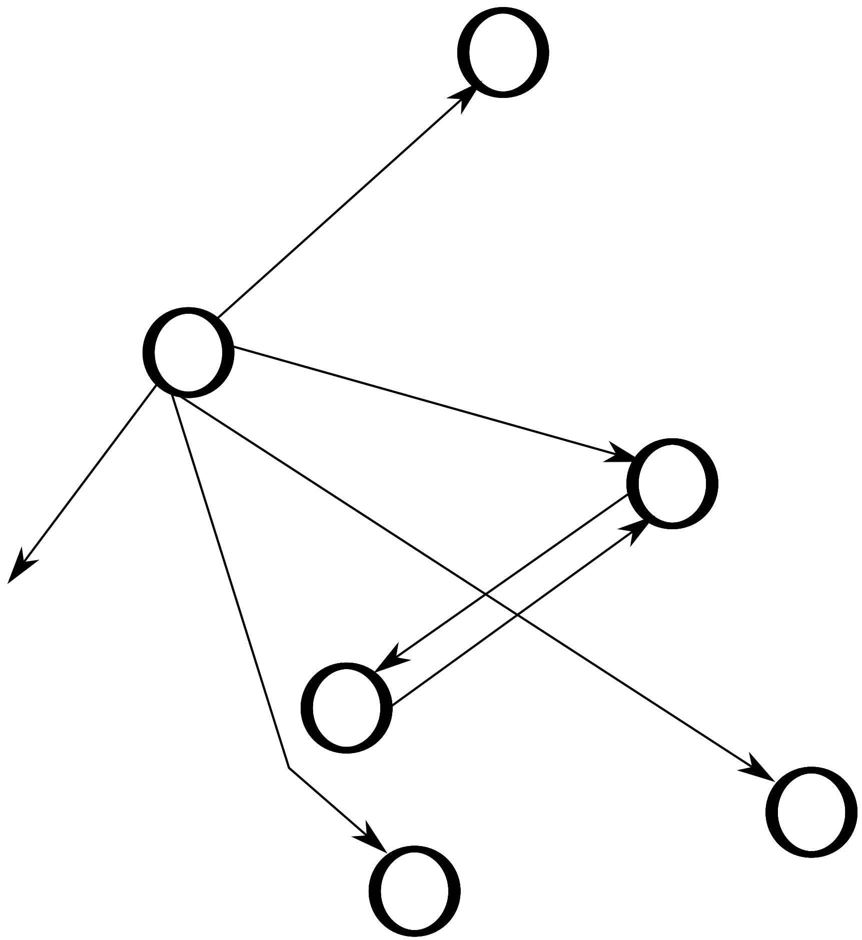

Multiple type particle systems

At every node there is a queue:

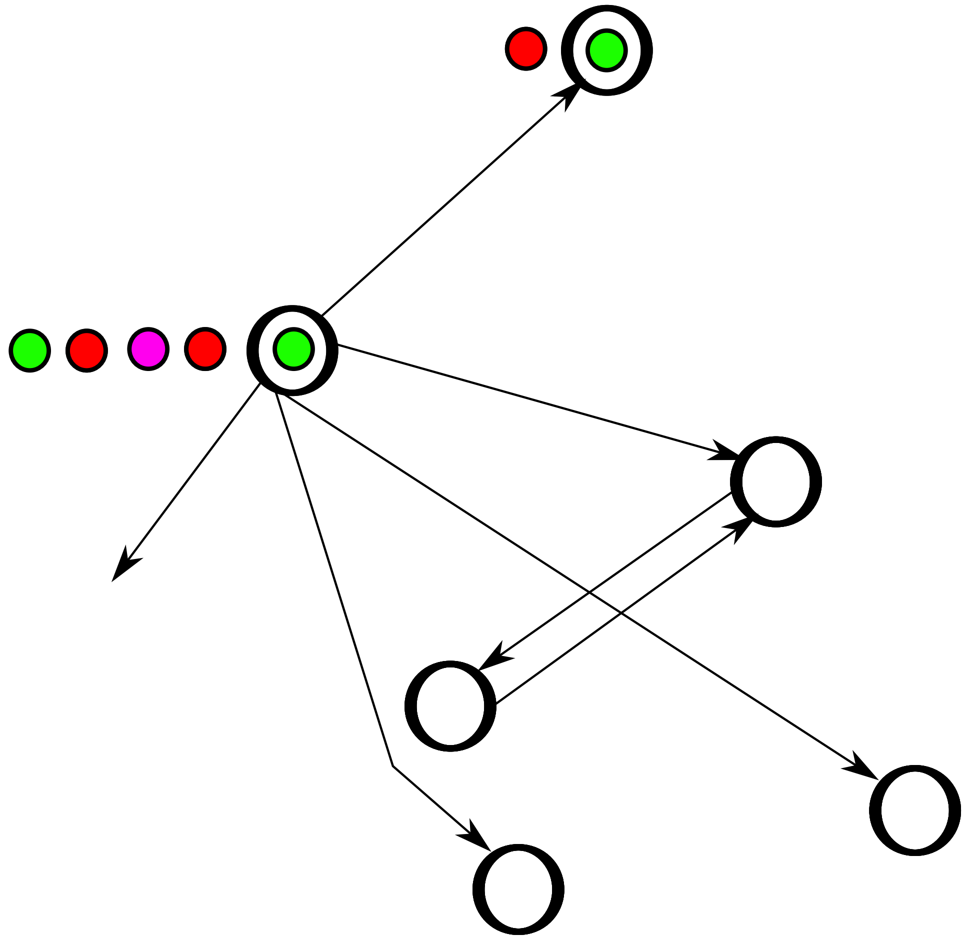

Multiple type particle systems

The service is over:

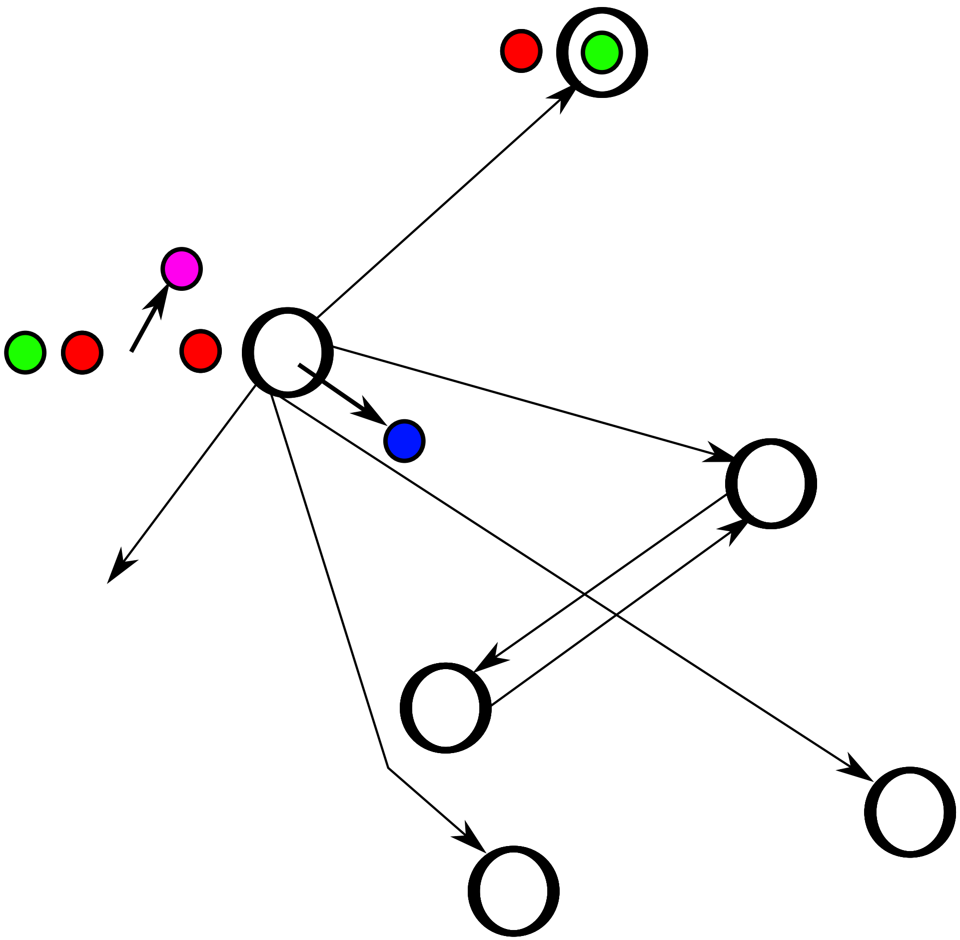

Multiple type particle systems

New clients are being served:

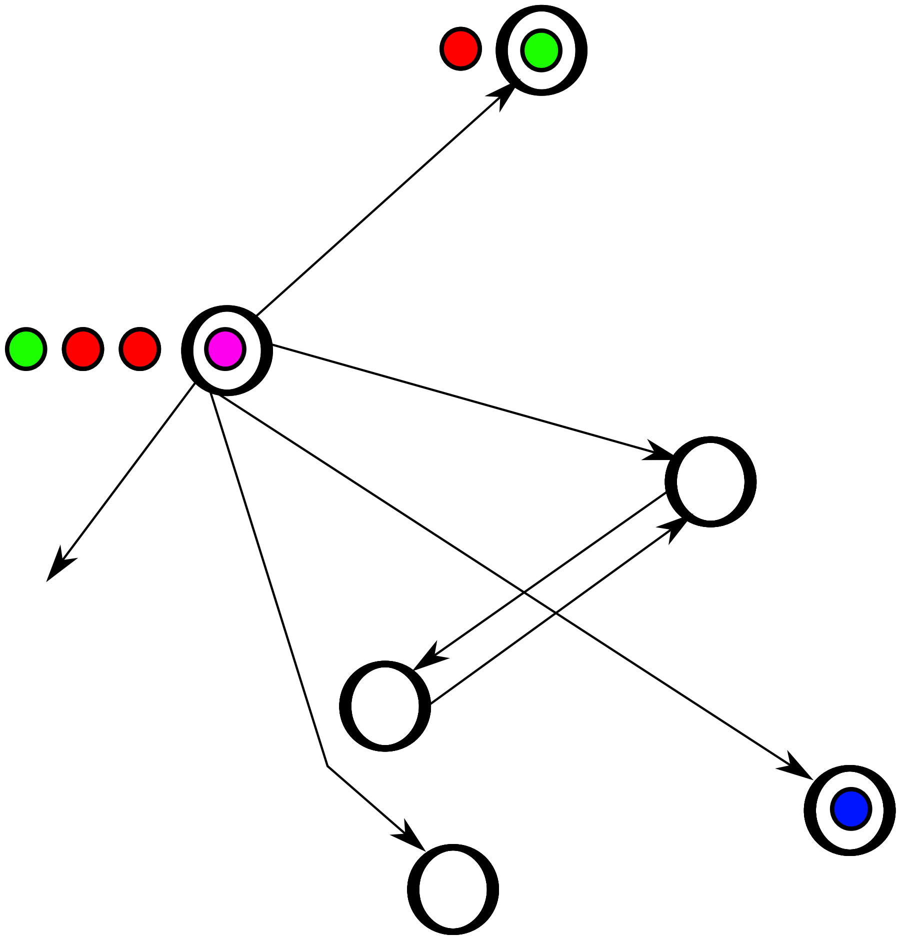

Multiple type particle systems

We study the system in the limit when

- the number of nodes, M, goes to ∞,

Multiple type particle systems

We study the system in the limit when

- the number of nodes, M, goes to ∞,
- the number of connections to (and from) every node goes to ∞ as well,

Multiple type particle systems

We study the system in the limit when

- the number of nodes, M, goes to ∞,
- the number of connections to (and from) every node goes to ∞ as well,
- the number of clients, N, is of the order of M, i.e.

$$N = \rho M.$$

The constant ρ will be called *the load*.

Multiple type particle systems

What to expect?

- (if started from a reasonable initial state):

Multiple type particle systems

What to expect?

- (if started from a reasonable initial state):
- The input flow to every server is Poissonian, with rate $\lambda(t) = \lambda_{node,sort}(t)$.

Multiple type particle systems

What to expect?

- (if started from a reasonable initial state):
- The input flow to every server is Poissonian, with rate $\lambda(t) = \lambda_{node,sort}(t)$.
- All the rates $\lambda(t)$ go to their limiting constant values.

Multiple type particle systems

What to expect?

- (if started from a reasonable initial state):
- The input flow to every server is Poissonian, with rate $\lambda(t) = \lambda_{node,sort}(t)$.
- All the rates $\lambda(t)$ go to their limiting constant values.
- The network equilibrates after time T independent of the size M.

Multiple type particle systems

What to expect?

- (if started from a reasonable initial state):
- The input flow to every server is Poissonian, with rate $\lambda(t) = \lambda_{node,sort}(t)$.
- All the rates $\lambda(t)$ go to their limiting constant values.
- The network equilibrates after time T independent of the size M.

In words, the system looses all the memory about the initial state.

Multiple type particle systems

What to expect?

- (if started from a reasonable initial state):
- The input flow to every server is Poissonian, with rate $\lambda(t) = \lambda_{node,sort}(t)$.
- All the rates $\lambda(t)$ go to their limiting constant values.
- The network equilibrates after time T independent of the size M.

In words, the system looses all the memory about the initial state.
This is called the Poisson Hypothesis behavior.

Mean field graphs

Mean field graphs

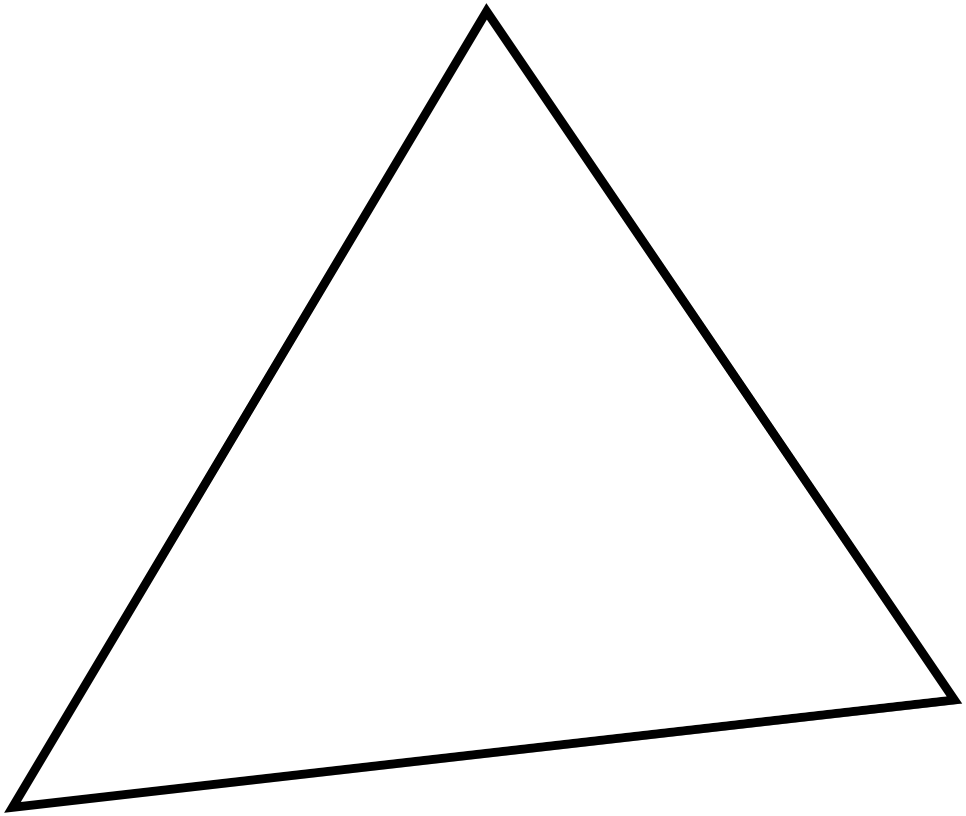

Mean field graphs

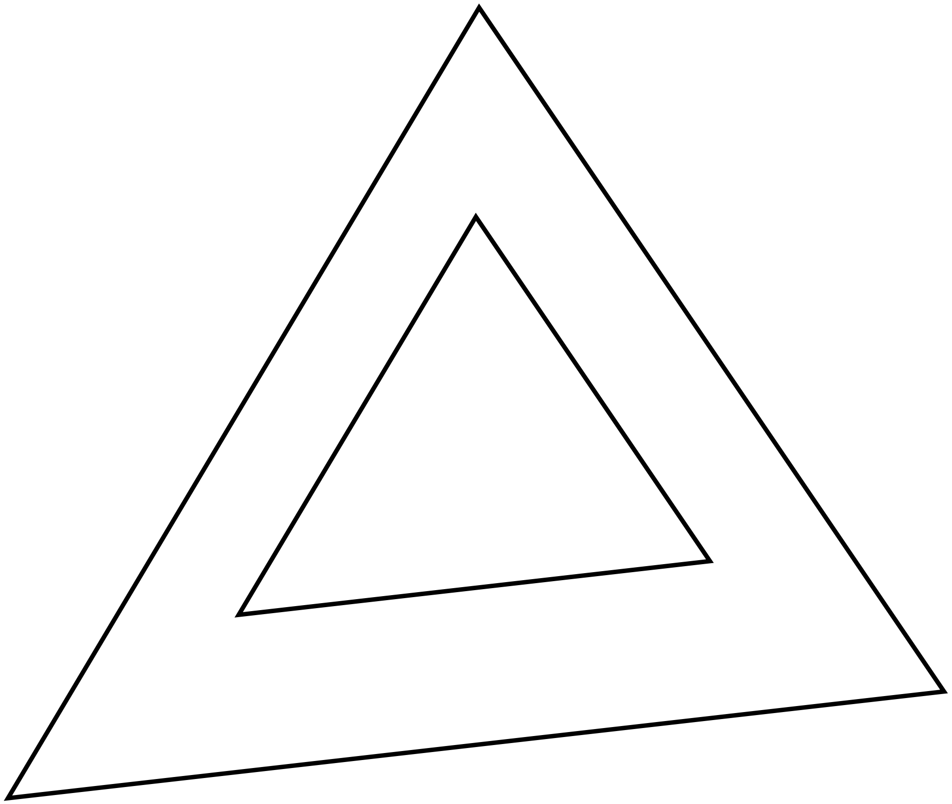

Mean field graphs

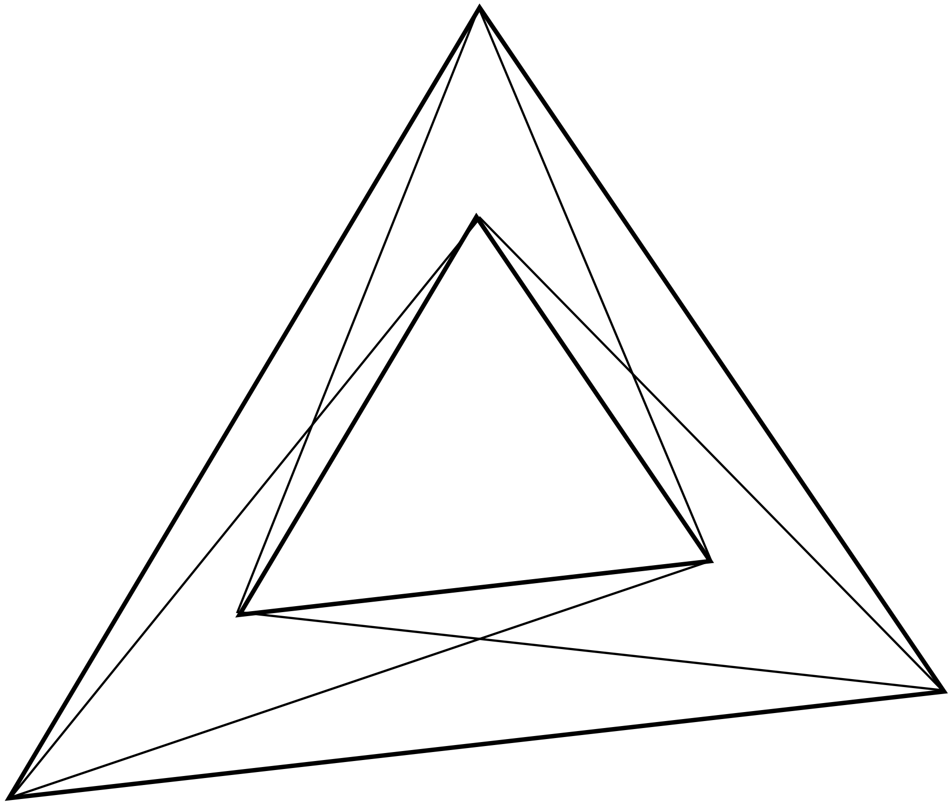

PH: Cases of validity. Single type clients

- Poisson Hypothesis holds "always".

PH: Cases of validity. Single type clients

- Poisson Hypothesis holds "always".
- Due to the self-averaging property:

$$b(t) = \int_{x \geq 0} \lambda(t - x)\, q_{\lambda,t}(x)\ dx$$

Here $\lambda(t)$, $-\infty < t < \infty$ is the rate of the Poisson process of moments of arrivals of customers to our server. Upon completion of the service the customer exits the system. The service time is not necessarily exponential, it can have power law decay.
$b(t)$ is the rate of the exit flow. The kernel $q_{\lambda,t}(x)$ is stochastic.

PH: Cases of validity. Low load, multiple type

Poisson Hypothesis holds since the clients almost never meet each other.

Violation of PH: phase transitions

The elementary network.

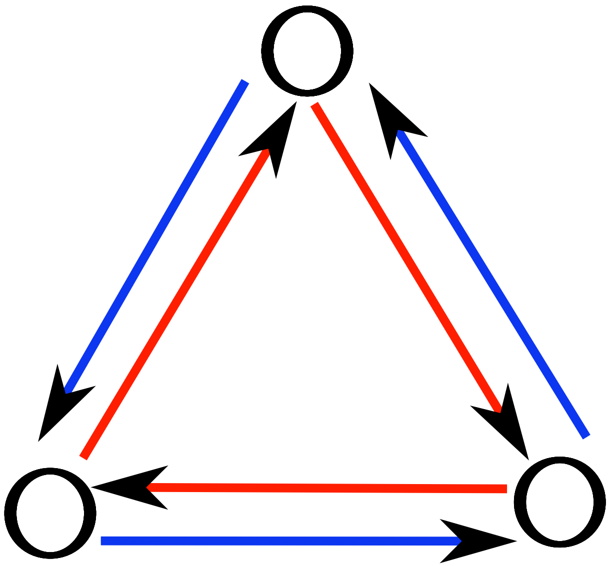

Violation of PH: phase transitions

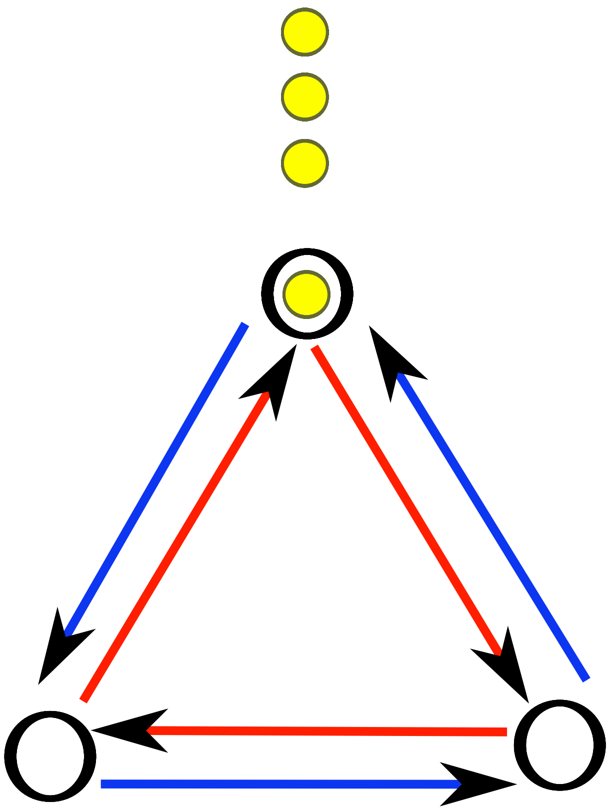

Violation of PH: phase transitions

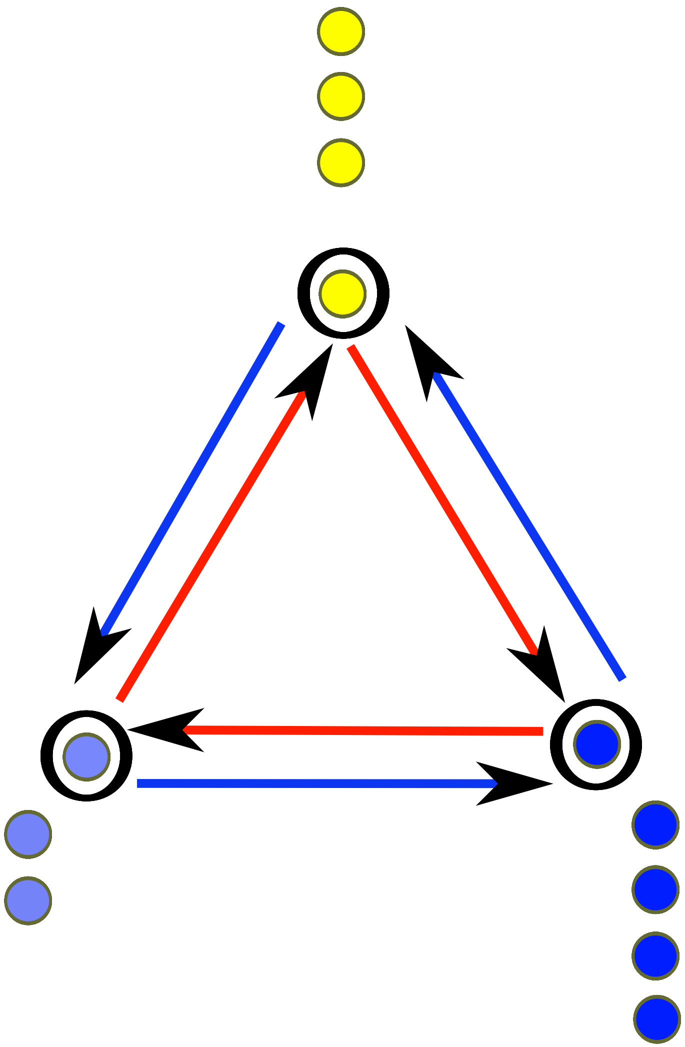

Violation of PH: phase transitions

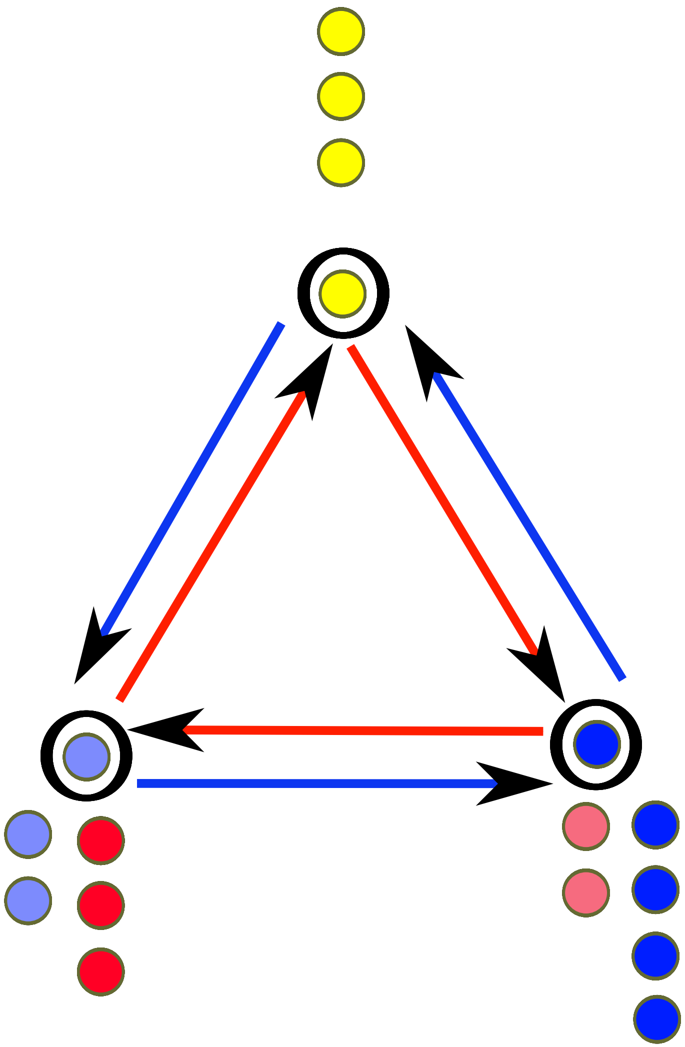

Violation of PH: phase transitions

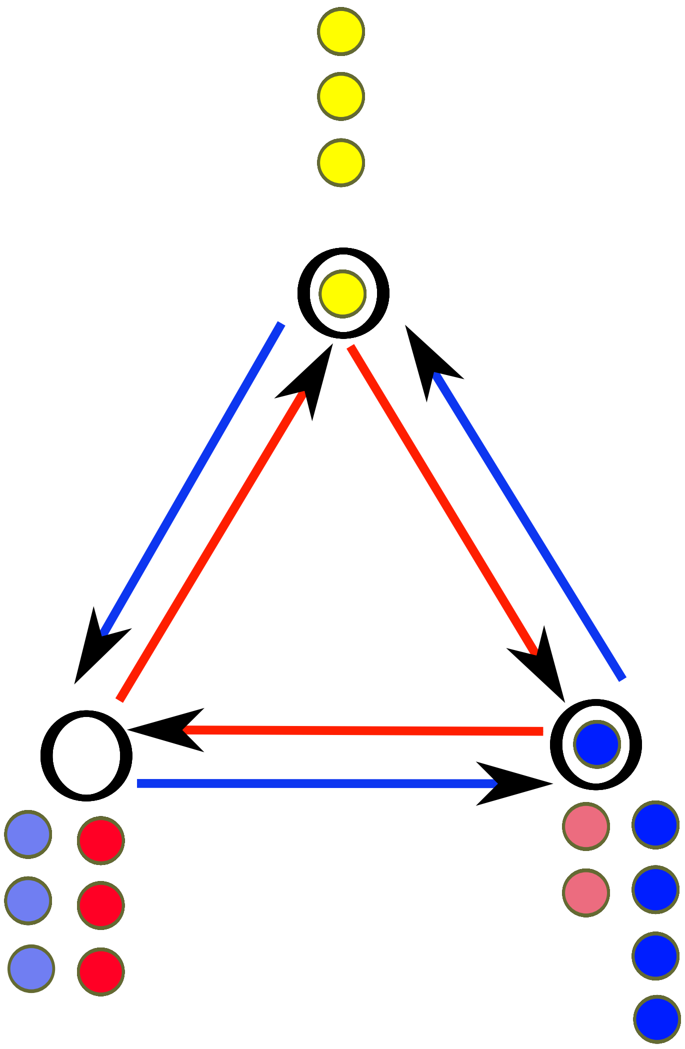

Violation of PH: phase transitions

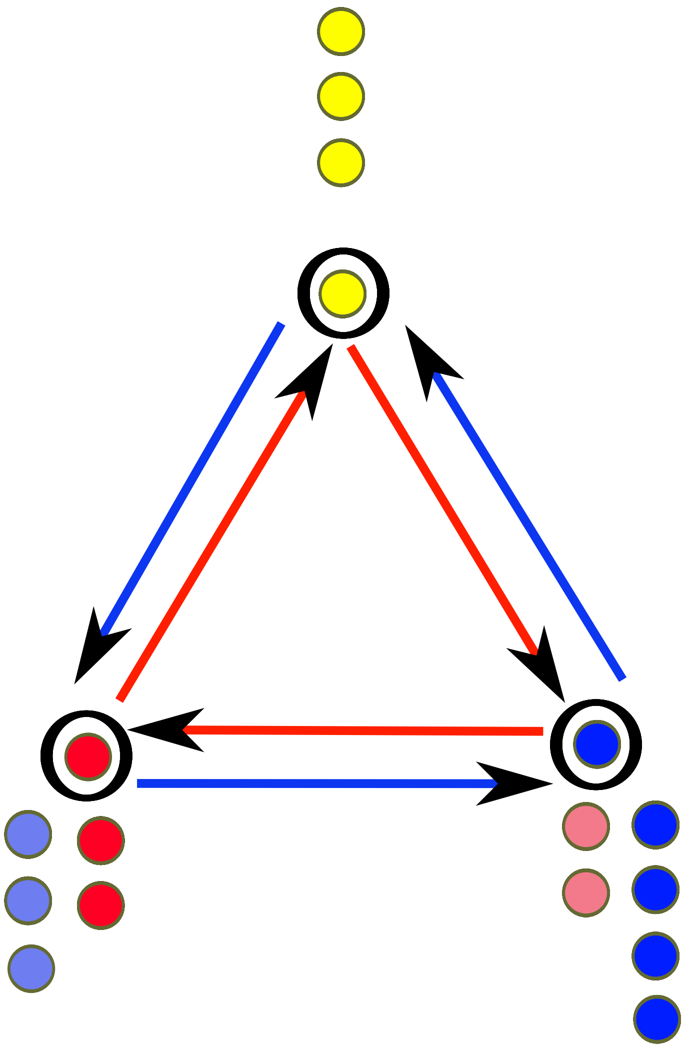

Violation of PH: phase transitions

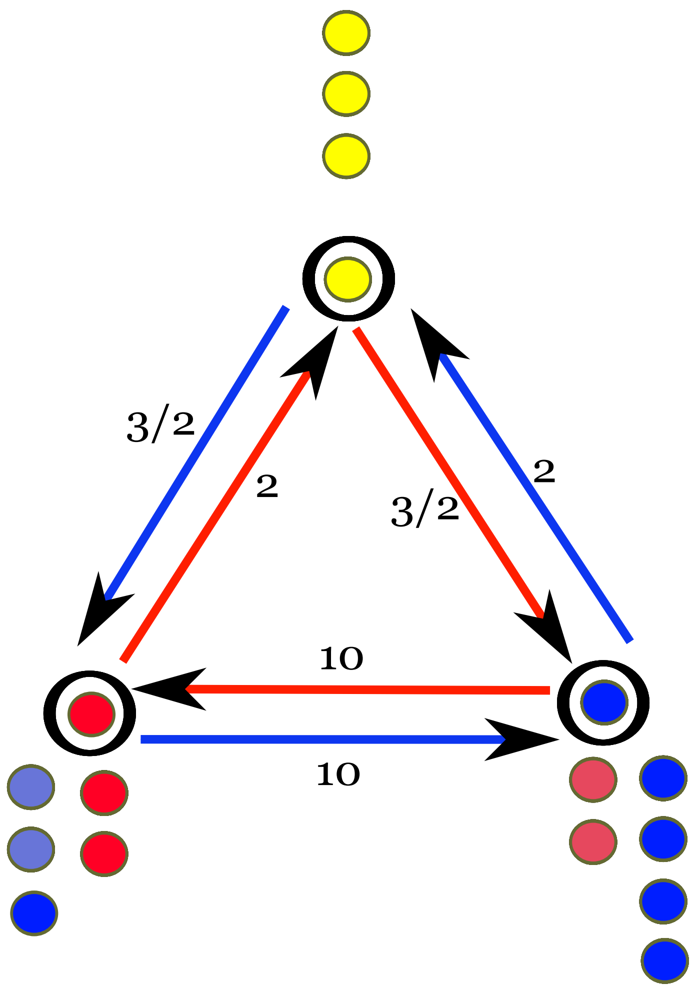

Fluid systems

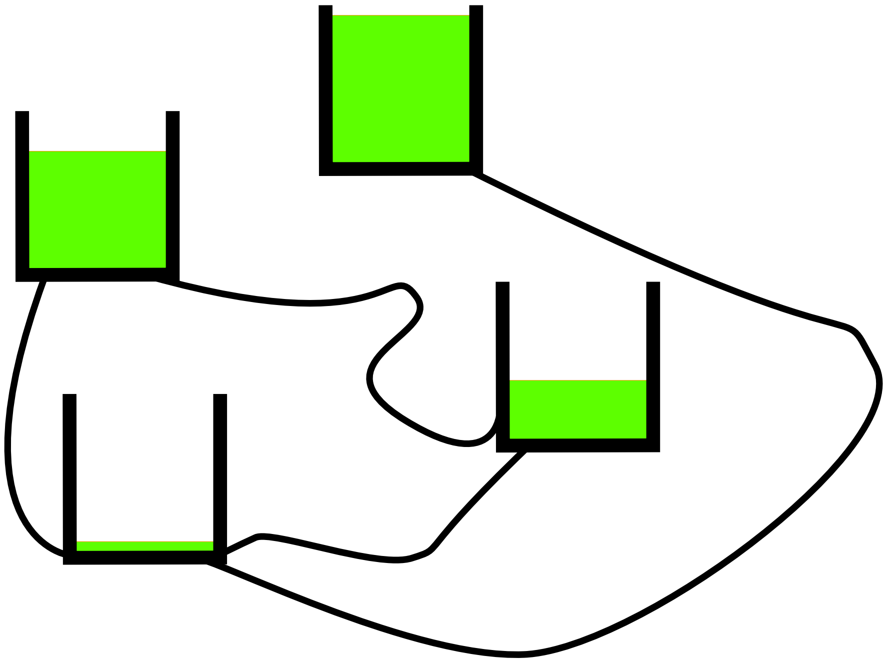

Fluid system with several fluids

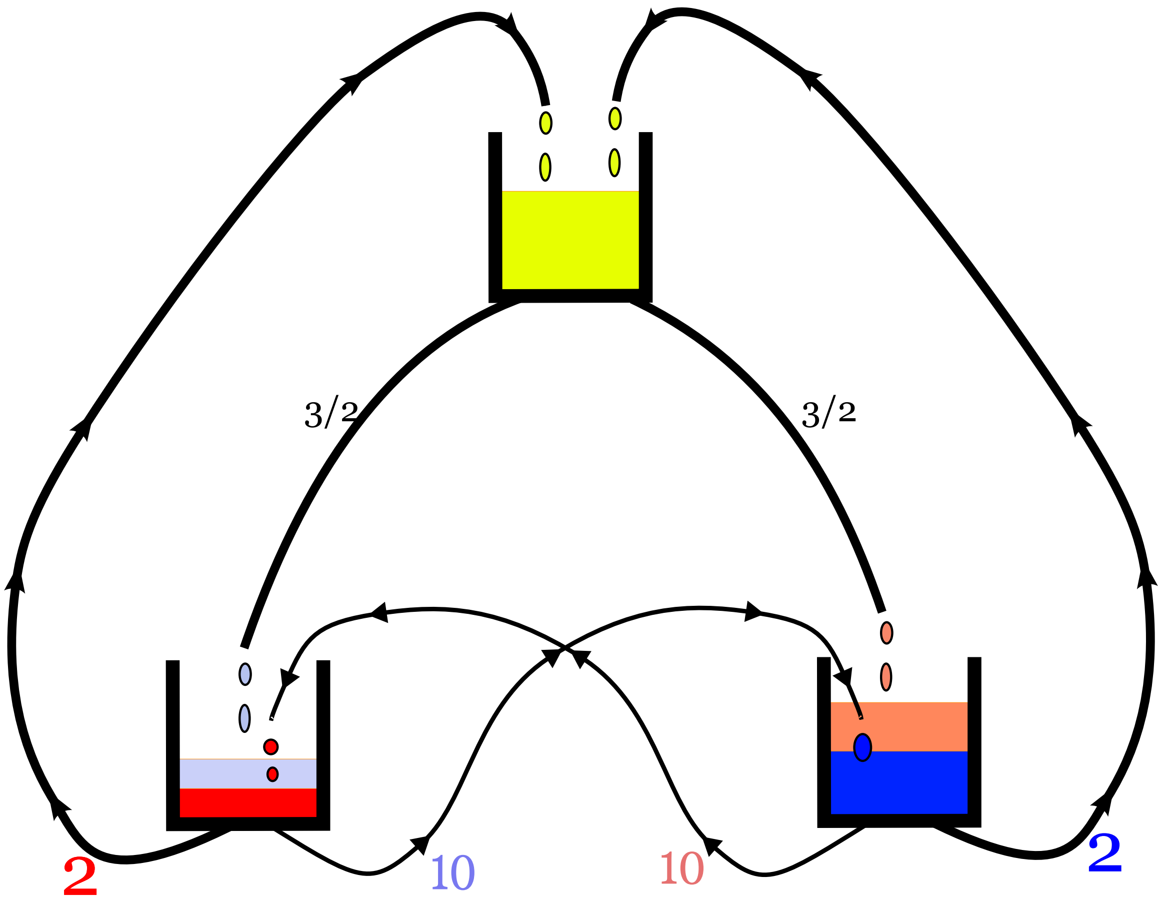

Fluid system with several fluids

Cyclic behavior:

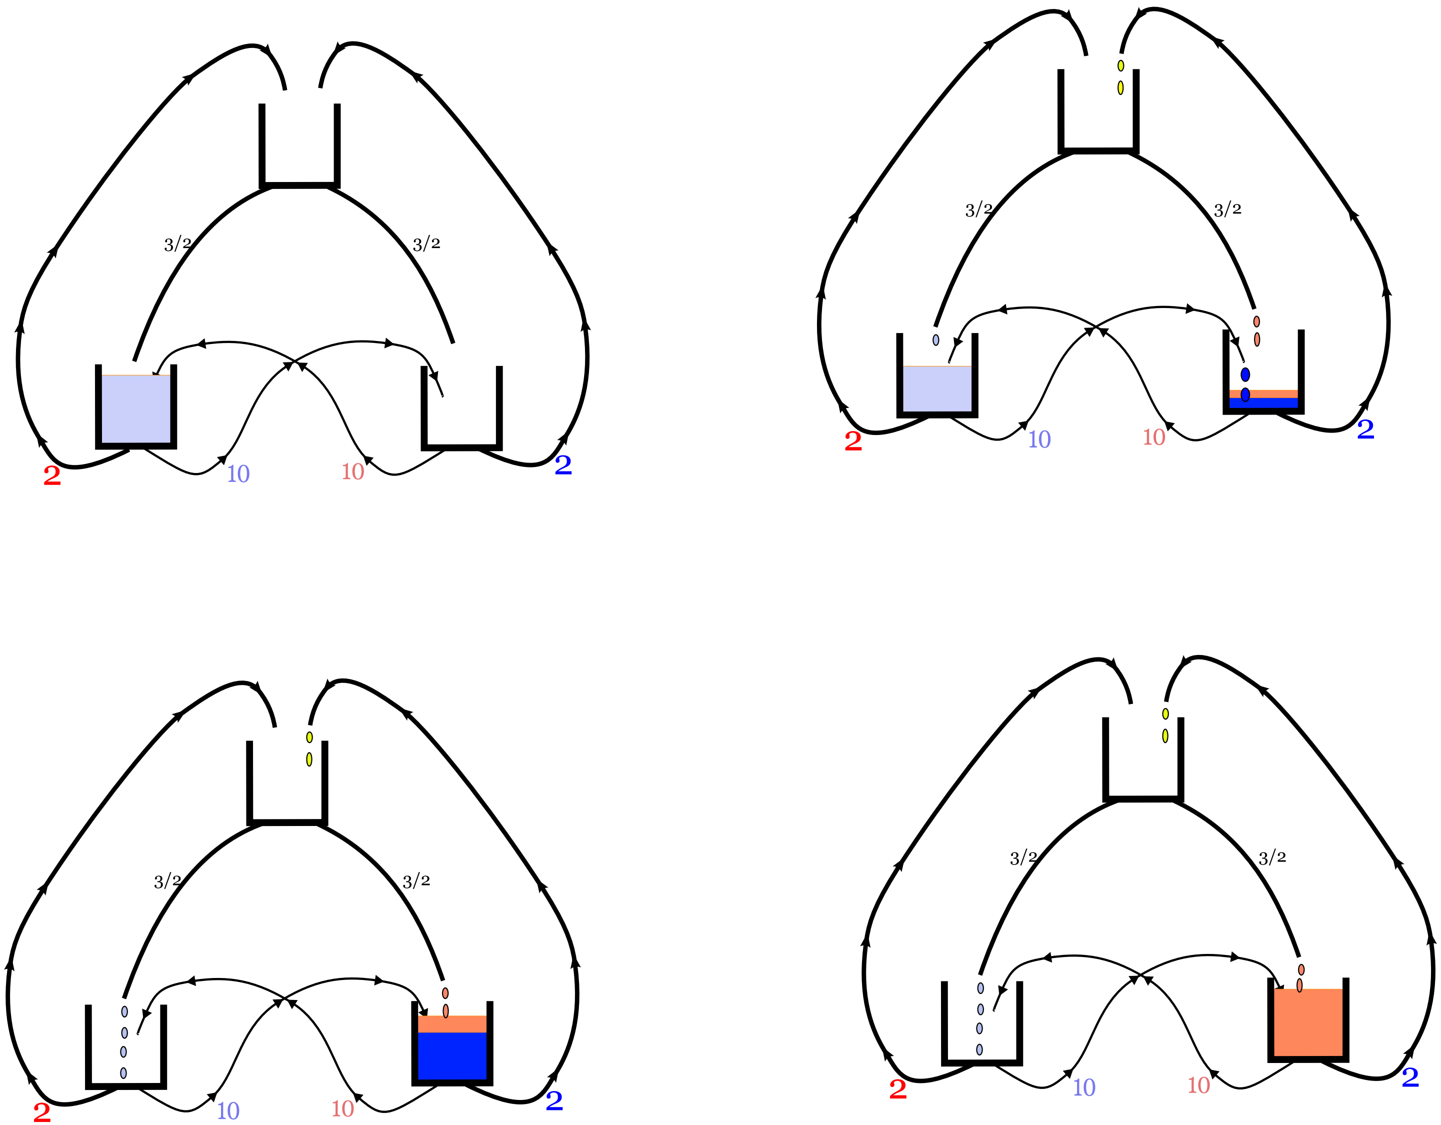

Construction of the network

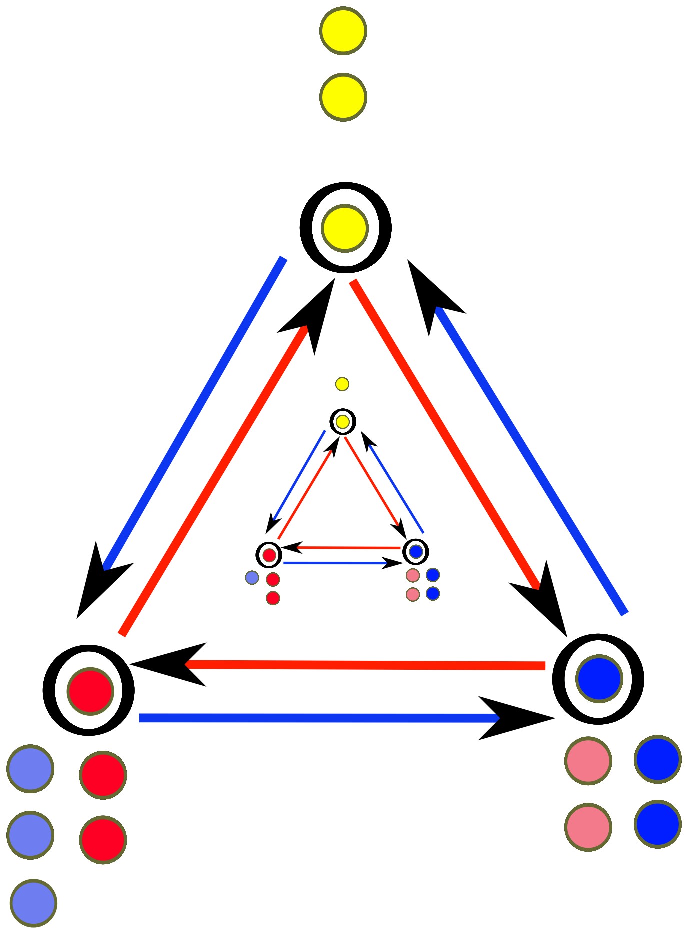

Construction of the network

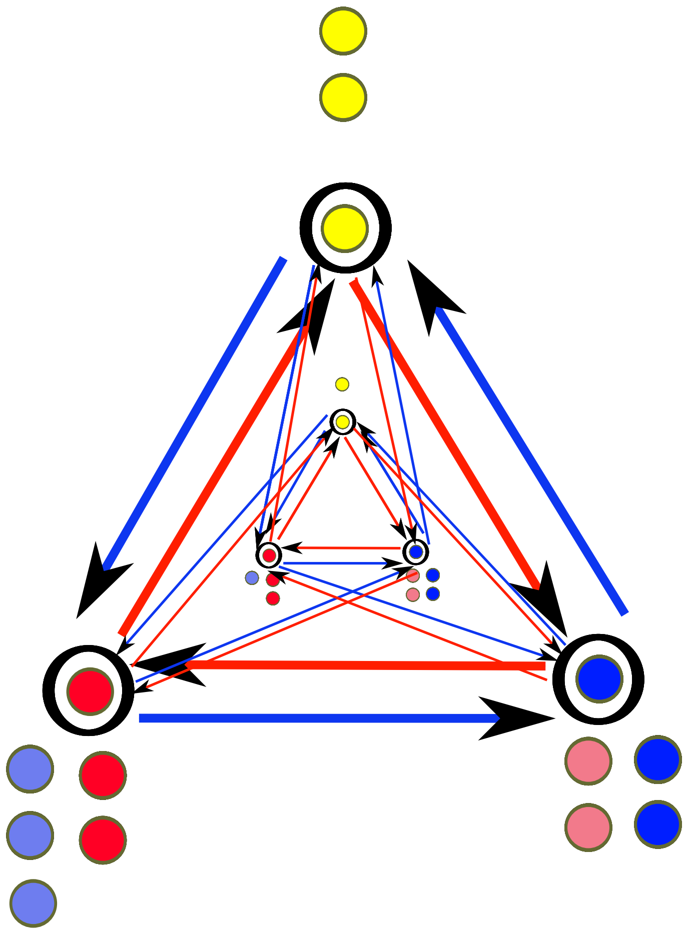

Phase transition

- Let us start the network ∇_M, made from M triangles, in a state with $\leq R$ clients per server.

Phase transition

- Let us start the network ∇_M, made from M triangles, in a state with $\leq R$ clients per server.
- **Theorem**
  - For $\rho \leq \rho_0$, ρ_0 small, the relaxation time $T_r(M, R, \rho)$ is bounded by $C(R, \rho_0)$, uniformly in M.
  - For $\rho \geq \rho_1$, ρ_1 large, there are initial states with $\leq R$ clients per server, such that the relaxation time $T_r(M, R, \rho) \to \infty$ as $M \to \infty$.

Phase transition

Theorem

- The marginals of ∇_M converge, as $M \to \infty$, to a non-linear Markov process, ∇_∞.
- For ρ large the process ∇_∞ is not ergodic.

Non-linear Markov Processes

(Linear) Markov chain $X_n \in \Omega$, $|\Omega| = k < \infty$.
Configurations $=$ points in Ω.
State $=$ probability measure on Ω.
S_k – the simplex of probability measures on Ω.
Transition matrix

$$\mathbf{P} = P\left(s, t\right), \quad \sum_t P\left(s, t\right) = 1$$

State μ is transformed to ν by

$$\nu = \mu \mathbf{P}$$

The map: $\mathbf{P} : S_k \to S_k$ is linear.

Non-linear Markov Processes

Non-linear Markov chain:
Transition probability to go from s to t depends also on the state μ.
The Non-linear Markov chain is defined by the collection of transition matrices

$$\mathbf{P}_\mu = P_\mu(s,t), \quad \sum_t P_\mu(s,t) = 1,$$

and state μ is transformed to ν by

$$\nu = \mu \mathbf{P}_\mu.$$

The map: $\mathbf{P} : S_k \longrightarrow S_k$ is non-linear.

The process ∇_∞ on $\left(\mathbb{Z}^5\right)^+$

Let $(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \in \left(\mathbb{Z}^5\right)^+$ is drawn from the state $\nu = \nu^\rho$. The rates:

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O - 1, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \qquad \gamma_{OA} + \gamma_{OB} = 3$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O, x_{OA}, x_{OB}, x_{AB} - 1, x_{BA}) \qquad \gamma_{BO} = 2$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA} - 1) \qquad \gamma_{AO} = 2$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA} = 0) \longrightarrow (x_O, x_{OA} - 1, x_{OB}, x_{AB}, x_{BA} = 0) \qquad \gamma_{AB} = 10$$

$$(x_O, x_{OA}, x_{OB}, x_{AB} = 0, x_{BA}) \longrightarrow (x_O, x_{OA}, x_{OB} - 1, x_{AB} = 0, x_{BA}) \qquad \gamma_{BA} = 10$$

The process ∇_∞ on $\left(\mathbb{Z}^5\right)^+$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O + 1, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \qquad \gamma_{BO}\nu\,(x_{AB} > 0) + \gamma_{AO}\nu\,(x_{BA} > 0)$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O, x_{OA}, x_{OB}, x_{AB} + 1, x_{BA}) \qquad \gamma_{AB}\nu \begin{pmatrix} x_{OA} > 0, \\ x_{BA} = 0 \end{pmatrix}$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA} + 1) \qquad \gamma_{BA}\nu \begin{pmatrix} x_{OB} > 0, \\ x_{AB} = 0 \end{pmatrix}$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O, x_{OA} + 1, x_{OB}, x_{AB}, x_{BA}) \qquad \gamma_{OA}\nu\,(x_O > 0)$$

$$(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) \longrightarrow (x_O, x_{OA}, x_{OB} + 1, x_{AB}, x_{BA}) \qquad \gamma_{OB}\nu\,(x_O > 0)$$

The process ∇_∞ on $\left(\mathbb{Z}^5\right)^+$

To show that for ρ large the process ∇_∞ is not ergodic, we consider its fluid limit as $\rho \longrightarrow \infty$, the dynamical system Δ_∞, which is obtained by applying the Euler scaling. If ν_t^ρ is our evolution, then we put

$$\mu_t(A) = \lim_{\rho\to\infty} \nu_{\rho t}^\rho (\rho A),$$

$A \subset \left(\mathbb{R}^5\right)^+$.

Δ_∞ - the non-linear dynamical system

Linear dynamical systems

A vector field V_x on $\mathbb{R}^n$ produces a flow along it, $F_t : \mathbb{R}^n \longrightarrow \mathbb{R}^n$.
It also defines the flow $\mathcal{F}_t$ on measures $\mathcal{M}(\mathbb{R}^n)$:

$$\mu_t(A) = \mu_0\left(F_t^{-1} A\right).$$

$\mathcal{F}_t$ acts linearly on measures $\mathcal{M}(\mathbb{R}^n)$

Non-Linear dynamical systems

A semigroup $\mathcal{F}_t$ of (not necessarily linear) transformations of the space $\mathcal{M}\left(\mathbb{R}^n\right)$.
Example: let $V_\mu = \int V_x \, d\mu\left(x\right)$, and

$$\mu_{t+\Delta t}\left(A\right) \approx \mu_t\left(A - \Delta t V_{\mu_t}\right)$$

The fluid network

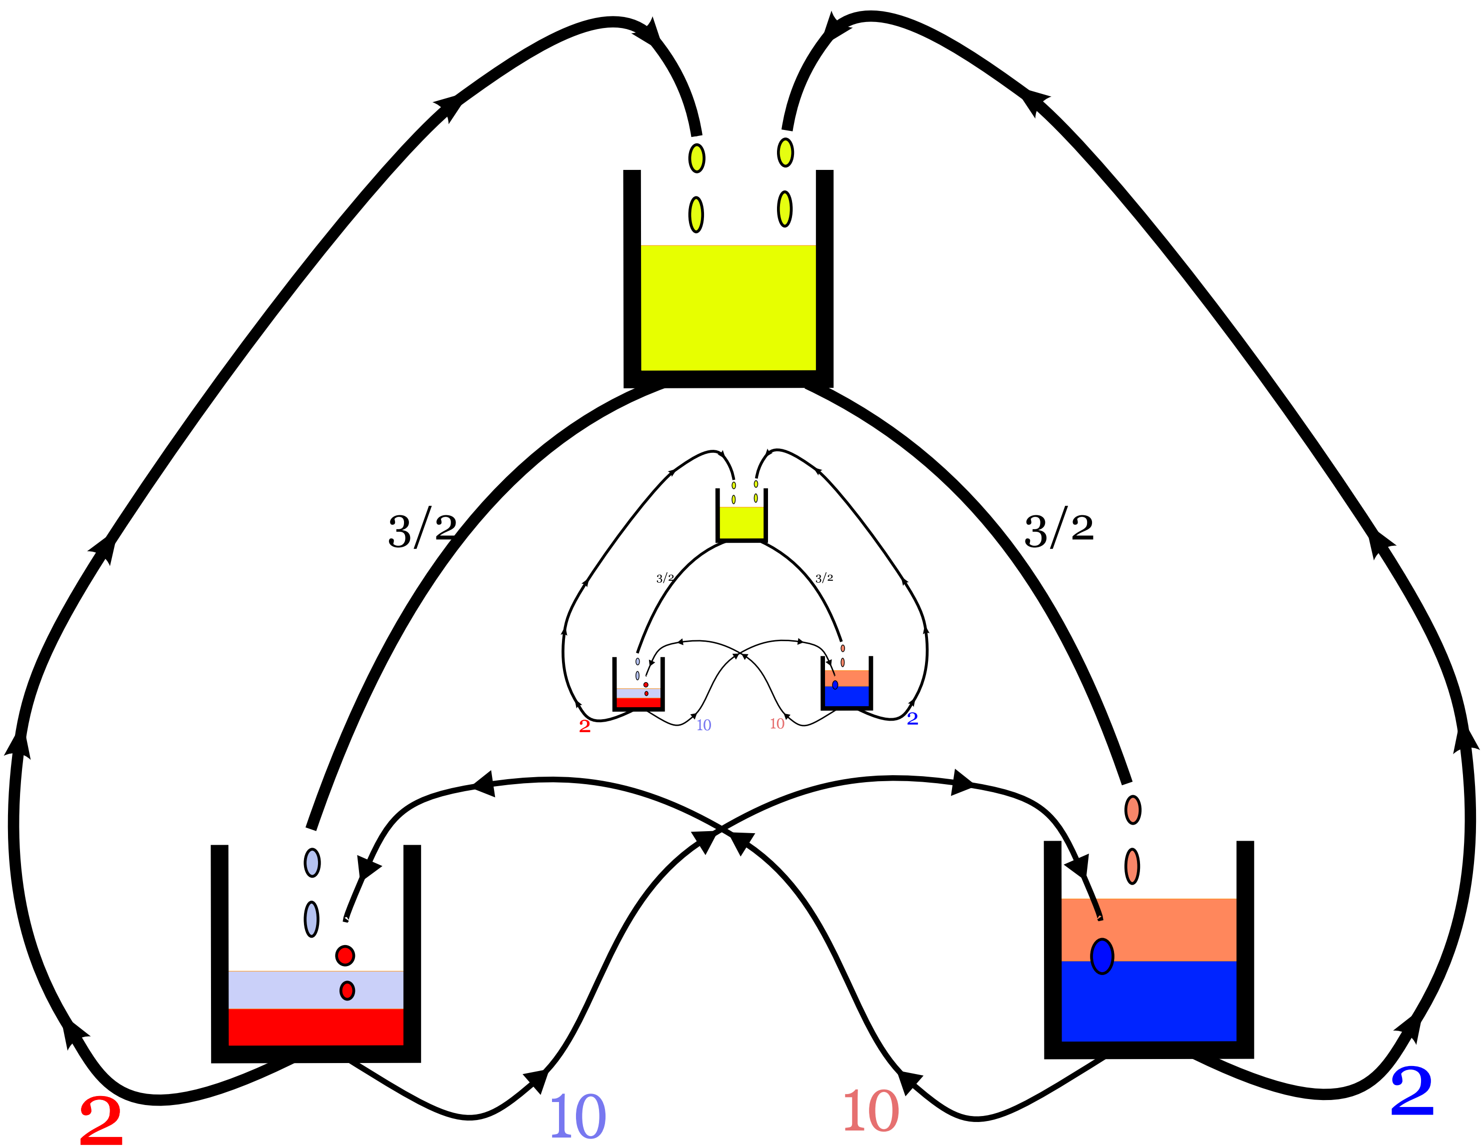

The fluid network

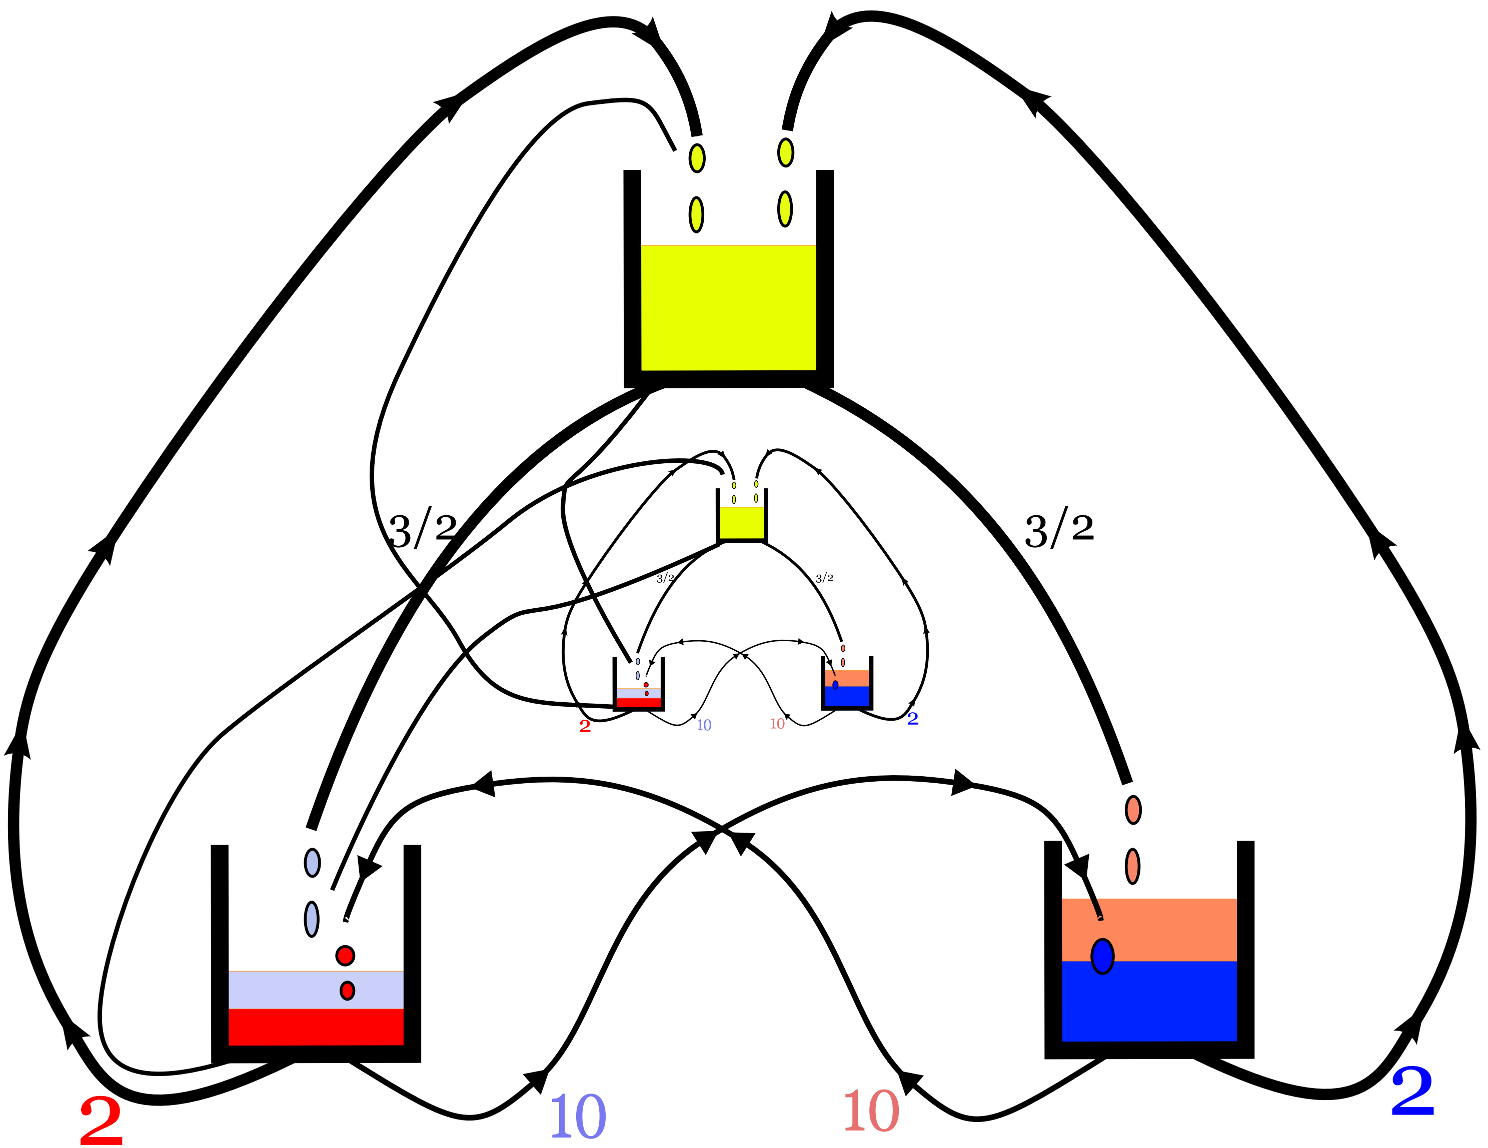

The fluid network

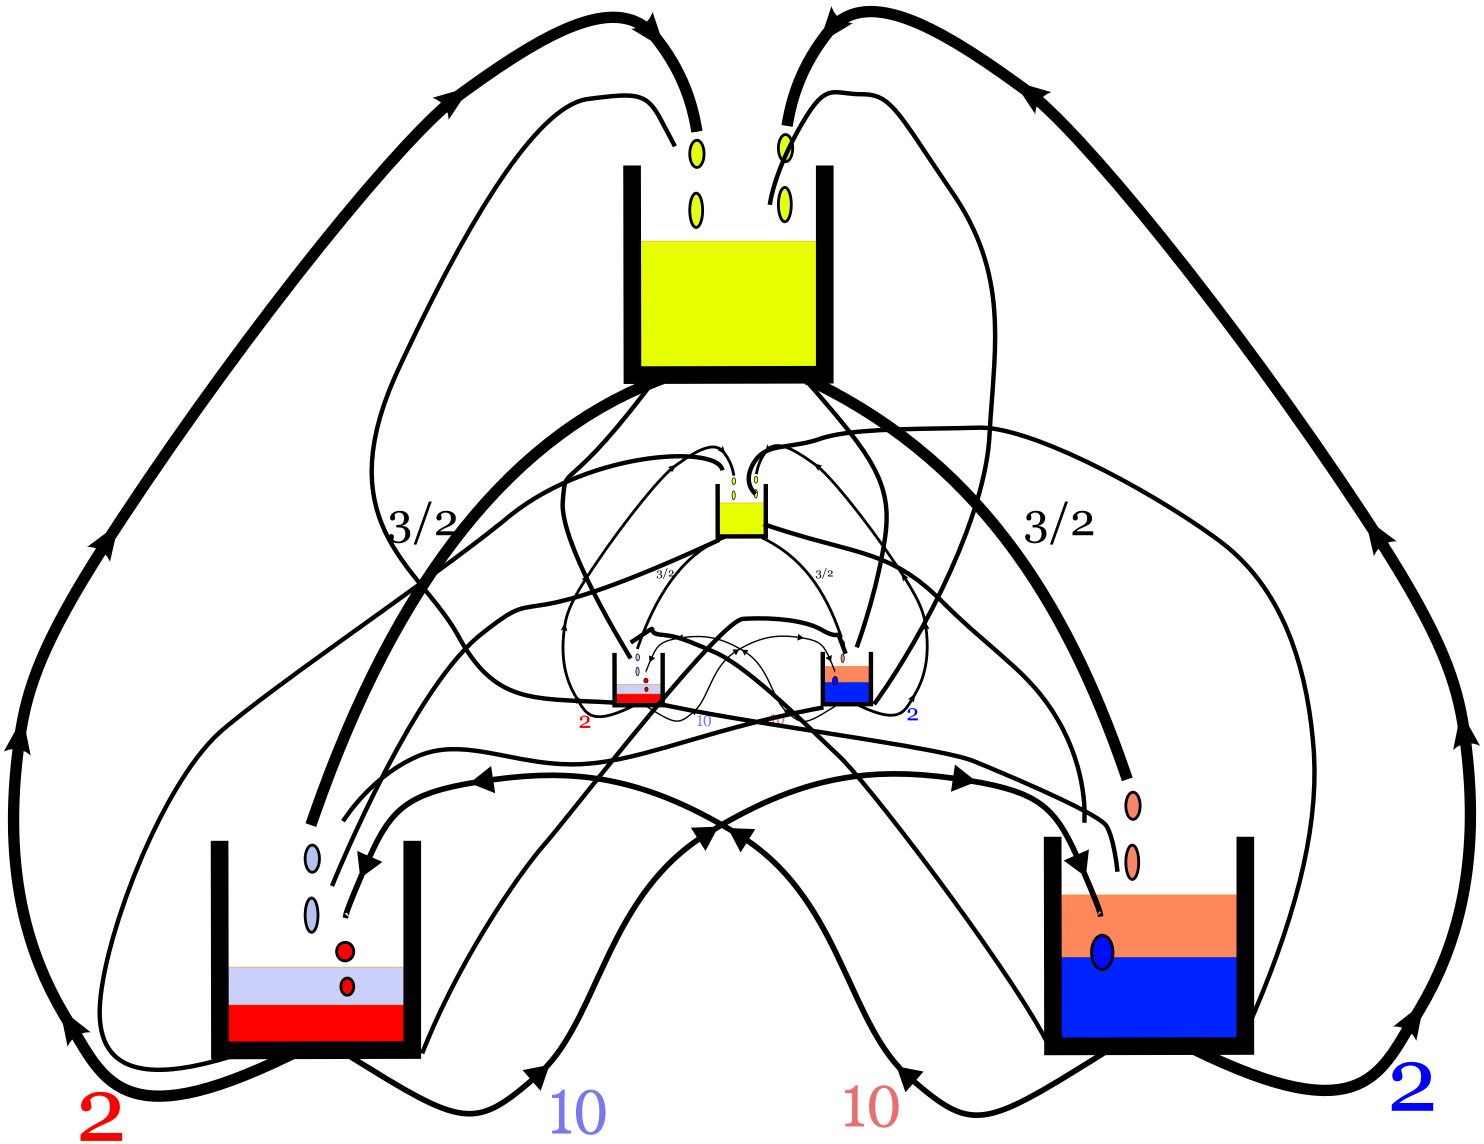

The limiting fluid network Δ_∞

Non-linear dynamical system on $\mathcal{M}\left(\left(\mathbb{R}^5\right)^+\right)$.

During the short time interval the measure μ evolves along the vector field $V_\mu(\bar{x})$;
at a point $\bar{x} = (x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA})$ with all coordinates positive it is

$$V_\mu(x_O, x_{OA}, x_{OB}, x_{AB}, x_{BA}) = \begin{pmatrix} -3 \\ 0 \\ 0 \\ -2 \\ -2 \end{pmatrix} + \begin{pmatrix} \gamma_{BO}\mu\{y_{AB} > 0\} + \gamma_{AO}\mu\{y_{BA} > 0\} \\ \gamma_{OA}\mu\{y_O > 0\} \\ \gamma_{OB}\mu\{y_O > 0\} \\ \gamma_{AB}\mu\{y_A > 0, y_{BA} = 0\} \\ \gamma_{BA}\mu\{y_B > 0, y_{AB} = 0\} \end{pmatrix}.$$

The limiting fluid network Δ_∞

The cycle $\mathcal{C} \subset \mathcal{M}\left(\left(\mathbb{R}^5\right)^+\right)$ of Δ is also a cycle of Δ_∞.

There are other attractors as well. But if

$$\rho_{KROV}\left(\mu_0, \mathcal{C}\right) < \delta,$$

then

$$\rho_{KROV}\left(\mu_t, \mathcal{C}\right) \to 0.$$

Therefore we can prove our theorem by induction in time t.

The End