Global Emergent Behaviors in Clouds of Agents

Soummya Kar and José M.F. Moura
Dept. of Electrical and Computer Engineering
Carnegie Mellon University

ICASSP
Special Session: Bio-Inspired Information Processing Networks
May 22-27, 2011
Prague Czech Republic

Acknowledgements: NSF CCF-1011903, NSF CCF-1018509, AFOSR FA95501010291

Electrical & Computer ENGINEERING

Colonies, Herds, Cybernetworks …

- Combined distributed interactions lead to complex behaviors
- Agents perform collective tasks & achieve collective decisions
  - move colony, distribute specialized tasks, defend colony, forage for food …
- Perform these tasks without clear hierarchical structure nor declared leader
- Each agent has narrow spatial sensing, inadequate cognitive ability, access to restricted information, limited ability to motivate or compel others to perform
- Still …. beyond the apparent random behavior of individuals in the colony, coordinated complex behavior arises that exhibits order, possibly inducing the colony to operate in synchrony

Synchrony

- **Synchrony*: coupled dynamical oscillators***
  - **Heart rhythms: Peskin's model for the cardiac pacemaker**
  - **Fireflies flashing in synchrony**
  - **Pulse-coupled biological oscillators**

Kuramoto's model

$$\frac{\partial \theta_i}{\partial t} = \omega_i + \frac{K}{N} \sum_{j=1}^{N} \sin(\theta_j - \theta_i), \qquad i = 1 \ldots N$$

- **Synchrony: stochastic networks + long term limits or averages**
  - **Abstract from individual behaviors**
  - **Ergodic limits (SLLN, WLLN) (sample averages converging to ensemble averages)**
  - **Stoch. (queuing) network models**
  - **Empirical distributions of states of networked systems**
  - **Renormalize, therrmodynamic (mean field) limits**
  - **Under appropriate conditions, limiting behaviors described by equilibria of ODEs**
  - **Synchrony may arise**

Global Behavior Analysis
An Event-based Networked Abstraction

Global Behavior: Mathematical Abstraction

Goal: Model for complex interaction

Mathematical abstraction tools

- Global equilibria- Mean field methods
- Faster scale perturbations – Diffusion (Brownian) Approximations
- Rare Events – Large Deviations

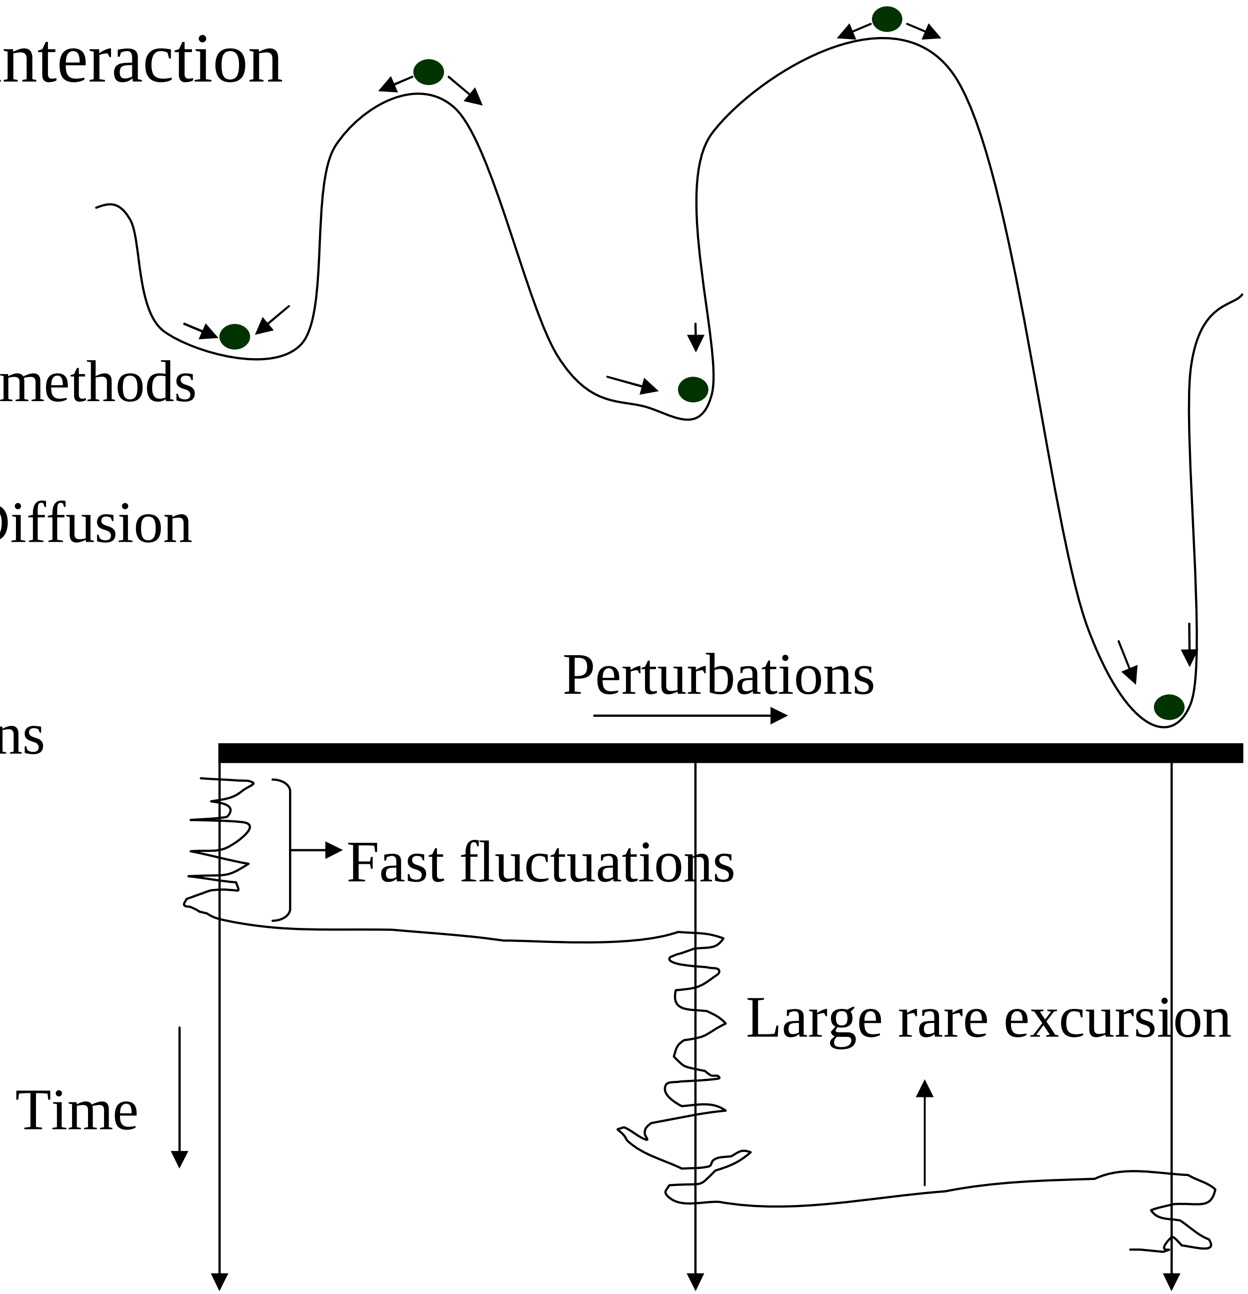

Event-based Network Model

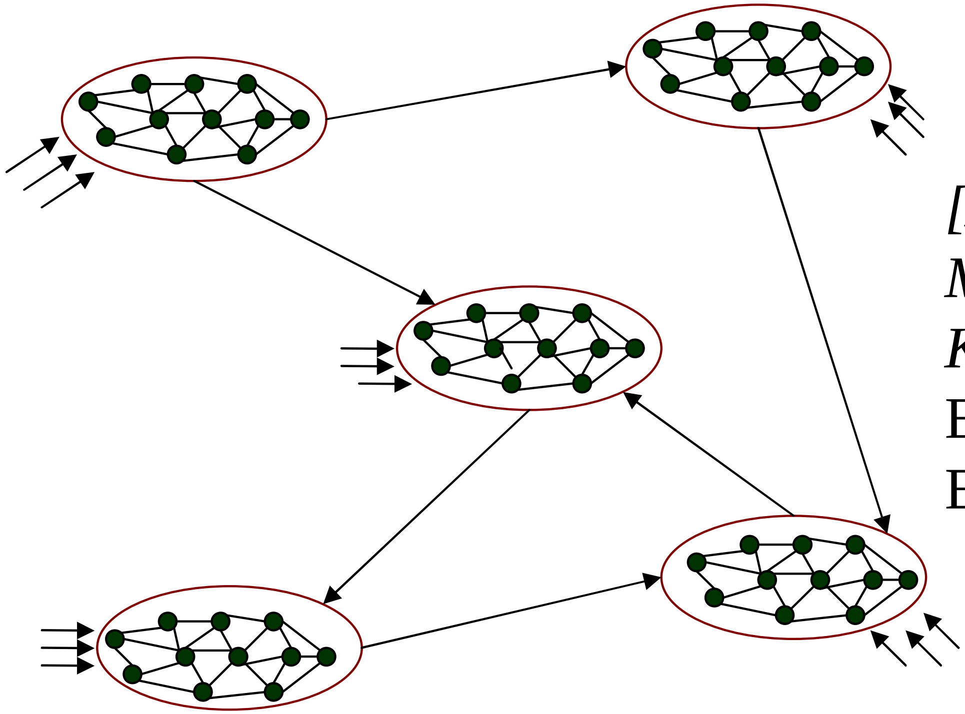

[Antunes, 2006]
M supernodes
K classes local events
Each node *C* units capacity
Events occur, interact, random times

For each class *k* event at each agent in each supernode *i*:

- External occurrences: Poisson, rate λ_{ik}
- Event influence time: exponential, mean $1/\mu_{ik}$
- Intra-(super)node interaction: Poisson, rate γ_{ik}
- Inter-(super)node interaction: Poisson, rate γ^{ij}
- Node with zero capacity rejects other events

Local Configuration Space

- ***Local state*** of agent n at supernode i: # of ≠ event types affecting n

$$X^N_{in}(t) = \mathbf{n} = (n_1, \cdots, n_K)$$

- Space of local interactions:

$$\mathcal{X} = \left\{ \mathbf{n} = (n_1, \cdots, n_K) \,\middle|\, \sum_{k=1}^{K} A_k n_k \leq C_k \right\}$$

- Vector of local node states

$$\mathbf{X}^N(t) = \{X^N_{in}(t),\ i \leq M,\ \mathbf{n} \in \mathcal{X}\}$$

- Detailed network dynamics given by $\mathbf{X}^N(t)$, embeds *global* behvior, intractable analysis for large N

Global Configuration Space

- Let $\mathbf{Y}_{i\mathbf{n}}^{N}(t)$ % nodes in supernode i with configuration $\mathbf{n}$ at time t
- ***Global network state*** set of *empirical distribution* processes

$$\mathbf{Y}^{N}(t) = \left\{\mathbf{Y}_{i\mathbf{n}}^{N}(t), i \leq M, \mathbf{n} \in \mathcal{X}\right\} \in \mathbb{R}^{M\mathcal{X}+1}$$

- $\mathbf{Y}^{N}(t)$ represents emerging global behaviors in terms of empiricalk averages
- Global concerns: what % nodes failed, not which nodes failed
- Dimension $\mathbf{Y}^{N}(t)$ (probability simplex-valued process) independent of N
- $\mathbf{Y}^{N}(t)$ is jump Markov process

Mean-Field Limit

Theorem 1 *As $N \to \infty$, the processes $\{\mathbf{Y}^N(t)\}$ converge weakly (in distribution w.r.t. the Skorokhod topology) to a deterministic trajectory $\{\mathbf{Y}(t)\}$. The limiting (mean-field) trajectory $\mathbf{Y}(t)$ is absolutely continuous and satisfies the ODE,*

$$\acute{Y}_{i,\mathbf{n}}(t) = f_{i,\mathbf{n}}(\mathbf{Y}(t)),\ i \leq M, \mathbf{n} \in \mathcal{X}$$

where the vector fields $f_{i,\mathbf{n}}(\cdot)$ are given by

$$f_{i,\mathbf{n}}(\mathbf{y}) = \sum_{k=1}^{K}\Bigg[\lambda_{i,k} + \gamma_{i,k}\sum_{\mathbf{m}\in\mathcal{X}} m_k y_{i,\mathbf{m}} + \sum_{j\neq i}\gamma_k^{j,i}\sum_{\mathbf{m}\in\mathcal{X}} m_k y_{j,\mathbf{m}}\Bigg]\left[y_{i,\mathbf{n}-\mathbf{e}_k}\mathbb{I}_{(n_k\geq 1)} - y_{i,\mathbf{n}}\mathbb{I}_{(\mathbf{n}+\mathbf{e}_k\in\mathcal{X})}\right]$$
$$+\sum_{k=1}^{K}\Bigg[\gamma_{i,k} + \mu_{i,k} + \sum_{j\neq i}\gamma_k^{i,j}\Bigg]\left[(n_k+1)y_{i,\mathbf{n}+\mathbf{e}_k}\mathbb{I}_{(\mathbf{n}+\mathbf{e}_k\in\mathcal{X})} - n_k y_{i,\mathbf{n}}\right]$$

Theorem 1: Discussion

- The proof is based on a martingale representation of the jump Markov processes $\mathbf{Y}^N(t)$.
- The trajectory $\mathbf{Y}(t)$ essentially is a probability measure valued process, taking values in the $M\mathcal{X}$ dimensional probability simplex.
- The dynamical equations (fluid limit) may be viewed intuitively as a coupled collection of the limiting differential equations obtained in [Antunes-2006], the coupling being attributed to the presence of multiple supernodes interacting with each other.
- The fluid limit provides qualitative characterization of the network model, for example, the stability, attractors of the asymptotic dynamical system.
- The fixed points of the ODE, i.e., $\mathbf{f}_{i,\mathbf{n}}(\mathbf{y}^*) = 0$, correspond to the ***global equilibria***.
- Such an equilibrium always exists.

Global Equilibria: Existence

Some Notation

For a vector $\rho = (\rho_1, \cdots, \rho_K) \in \mathbb{R}^K$, consider the function $v_\rho : \mathcal{X} \longmapsto \mathbb{R}_+$,

$$v_\rho(\mathbf{n}) = \frac{1}{Z(\rho)} \prod_{l=1}^{K} \frac{(\rho_l)^{n_l}}{n_l!}, \quad \mathbf{n} \in \mathcal{X}$$

where $Z(\cdot)$ is the partition function given by

$$Z(\rho) = \sum_{\mathbf{n}\in\mathcal{X}} \prod_{l=1}^{K} \frac{(\rho_l)^{n_l}}{n_l!}$$

Thus v_ρ represents a family of probability measures on the configuration space $\mathcal{X}$, being parameterized by the K dimensional vector ρ.

Global Equilibria: Existence (*Contd.*)

Theorem 2:

- *Let* $\mathbf{y}^* \in \mathbb{R}^{M\mathcal{X}}$ *be an equilibrium of the mean field dynamical system. Then, there exists vectors* $\rho^1, \cdots, \rho^M \in \mathbb{R}^K$, *such that,* $y^*_{i,\mathbf{n}} = v_{\rho^i}(\mathbf{n})$, *for all* $i \leq M$ *and* $\mathbf{n} \in \mathcal{X}$ *and the parameterizations* $\rho^i = (\rho^1, \cdots, \rho^K)$ *satisfy the fixed point equations,*

$$\rho^i_k = \frac{\frac{\lambda_{i,k}}{\gamma_{i,k}} + < I_k, v_{\rho^i} > + \sum_{j \neq i} \frac{\gamma^{j,i}_k}{\gamma_{i,k}} < I_k, v_{\rho^j} >}{1 + \frac{\mu_{i,k}}{\gamma_{i,k}} + \sum_{j \neq i} \frac{\gamma^{i,j}_k}{\gamma_{i,k}}}$$

  and $< I_k, v_{\rho^j} >= \sum_{\mathbf{m} \in \mathcal{X}} m_k v_{\rho^j}(\mathbf{m})$ *for all* k, j.

- *There always exists an equilibrium point of the mean-field ODE.*

- *If, in addition,* $K = 1$, *i.e., there is only one class of events, the equilibrium of the mean-field ODE is unique.*

Global Equilibria: Existence (*Contd.*)

- Theorem 2 presents a parametric representation of the equilibrium points of the fluid dynamical system.
- Provides a characterization of the equilibrium points of the fluid ODE in terms of solutions of the fixed point equations.
- Shows the existence of at least one equilibrium always.
- The equilibria may not be unique, leading to the possibility of complicated system phenomena like metastability.
- However, if the number of classes $K = 1$, then the assertions conclude the existence of a unique equilibrium of the fluid ODE.

Synchronous Global Equilibria

Definition 1 : A fixed point $\mathbf{y}^*$ *of the mean-field dynamical equations, i.e.,* $V(\mathbf{y}^*) = 0$, *is said to be a* ***synchronous*** *equilibrium of the asymptotic system if there exists* $\mathbf{a} \in \mathbb{R}^{\mathcal{X}}$, *such that,*

$$\mathbf{y}^* = [\mathbf{a}^T, \mathbf{a}^T, \cdots, \mathbf{a}^T]^T$$

Equivalently, by the parametric representation of equilibria, this implies the existence of $\rho \in \mathbb{R}^K$, *such that* $(v_\rho, \cdots, v_\rho)$ *is an equilibrium of the asymptotic dynamical system.*

- Synchronous equilibria (if they exist) may lead to eventual (w.r.t. time) synchronous behavior of the supernodes.
- The synchrony between supernodes in terms of the global equilibrium configuration is important in many applications including load balancing etc.
- May also be viewed from a system design perspective, where the goal is to design the inter-supernode coupling parameters $\gamma_k^{i,j}$ under network constraints to lead to synchronous supernode operation.

Synchronous Global Equilibria: Existence Cdtns

Theorem 3: Let the network parameters satisfy the following additional conditions:

(i) *The entry rates of events of different classes are same for all the supernodes, i.e.,* $\lambda_{i,k} = \lambda_k,\ \forall i, k$.

(ii) $\gamma_{i,k} + \sum_{j \neq i} \gamma_k^{j,i} = x_k \geq 0, \quad \forall i, k$

(iii) $\gamma_{i,k} + \mu_{i,k} + \sum_{j \neq i} \gamma_k^{i,j} = y_k, \quad \forall i, k$

Then,

(1) *There exists a* synchronous *fixed point. In other words, there exists* $\rho \in \mathbb{R}_+^K$*, such that* $(v_\rho, \cdots, v_\rho)$ *is an equilibrium distribution for the asymptotic dynamical system.*

(2) *If, in addition,* $K = 1$ *(only one class of events), the above synchronous fixed point is the unique equilibrium of the mean-field limit* $\mathbf{Y}(t)$.

Conclusions

- We introduced a queueing stochastic network model to describe the interactions of clouds of agents and studied its asymptotic behavior after suitable normalization.

- The asymptotic behavior of the network is described by an ordinary differential equation (ODE) whose fixed points correspond to the global equilibria.

- Different global behaviors may emerge as different equilibria of this ODE (metastability). Further study is needed to determine the regions of the parameters spaces that lead to metastability.

- Prelimit behavior (finite N) requires a large deviation analysis of the convergence to the mean-field limit and offers an important direction for future research.