MoK: Stigmergy Meets Chemistry to Exploit Social Actions for Coordination Purposes

Stefano Mariani, Andrea Omicini
{s.mariani, andrea.omicini}@unibo.it

Dipartimento di Informatica: Scienza e Ingegneria (DISI)
Alma Mater Studiorum—Università di Bologna

SOCIAL.PATH 2013
Social Coordination: Principles, Artifacts and Theories
Exeter, UK – 3rd of April 2013

Outline

The Challenge

Socio-technical systems are becoming increasingly complex and thus difficult to design, mostly due to the unpredictability of human interactions [Omicini, 2012].

KIE

Furthermore, such systems often represent Knowledge Intensive Environments (KIE) [Bhatt, 2001], that is, they are meant to manage a huge amount of (possibly heterogeneous) information [Ossowski and Omicini, 2002].

A Path To Follow I

Coordination models have been developed to cope with software systems' increasing complexity, whose main source is the interaction space such systems have to manage [Gelernter and Carriero, 1992, Papadopoulos and Arbab, 1998, Omicini and Viroli, 2011].

Non-determinism

Such interactions have been for long recognised as an undesirable source of (uncontrollable) non-determinism, harnessing correctness, reliability, and predictability of systems.

A Path To Follow II

Then, probability entered the picture as a means to model, govern, and predict non-determinism, making it become a source of solutions rather than problems.

Nature-inspired Systems

Natural systems from chemistry, biology, physics, sociology, and the like are widely recognised for their capability to "reach order out of chaos" through adaptiveness and self-organisation, hence they have been proficiently used as metaphors for the engineering of coordination models & systems [Omicini, 2013].

Outline

(Cognitive) Stigmergy & BIC I

In complex social systems, such as human organisations, animal societies, and artificial multi-agent systems as well, interactions between individuals are *mediated by the environment*, which "records" all the *traces* left by their actions [Weyns et al., 2007].

Stigmergy

Trace-based communication is related to the notion of stigmergy, firstly introduced in the biological study of social insects [Grassé, 1959].

(Cognitive) Stigmergy & BIC II

Furthermore:

- modifications to the environment are often amenable of a *symbolic interpretation*
- interacting agents feature *cognitive abilities* that can be proficiently exploited in stigmergy-based coordination

Cognitive Stigmergy

When traces becomes signs, stigmergy becomes cognitive stigmergy [Omicini, 2012].

(Cognitive) Stigmergy & BIC III

Another step beyond (cognitive) stigmergy is taken by Behavioral Implicit Communication (BIC), in which coordination among agents can be based on the observation and interpretation of *actions as wholes*, rather than solely of their effects on the environment.

BIC

In BIC, actions themselves become the "message", often intentionally sent through the environment in order to obtain collaboration [Castelfranchi et al., 2010].

Computational requirements I

Moving from stigmergy to BIC, a list of *desiderata* emerge to be supported by the coordination abstraction:

- reification of agent's (inter)actions
- recording of their *contextual properties*
- recording of their "traces"
- capability of reaction to traces emission
- topology-related aspects management
- ontology management for symbolic interpretation

Computational requirements II

Among the many sorts of coordination models [Gelernter and Carriero, 1992], tuple-based ones [Gelernter, 1985] have been already taken as a reference for stigmergic coordination—including cognitive and BIC [Omicini, 2012].

Tuple-based Coordination

(Logic) Tuple-based infrastructures such as TuCSoN [Omicini and Zambonelli, 1999], feature space-time situatedness, probabilistic primitives and coordination programmability[a], providing us with all the necessary tools to fully support stigmergy, cognitive stigmergy, and BIC.

[a] Thanks to the ReSpecT language [Omicini, 2007].

(Bio)Chemistry & Biochemical Tuple Spaces

Biochemical Tuple Spaces

Biochemical tuple spaces are a stochastic extension of the LINDA framework [Gelernter, 1985] inspired both by chemistry and biology proposed in [Viroli and Casadei, 2009].

The idea is to attach to each tuple a "concentration", which can be seen as a measure of the *pertinency/activity* of the tuple: the higher it is, the more likely and frequently the tuple will influence system coordination.

Concentration of tuples evolves at a certain rate due to stochastic chemical rules installed into the tuple space, which acts as a chemical simulator [Gillespie, 1977].

Outline

Outline

MoK Overview

The MoK model was introduced in [Mariani and Omicini, 2013] as a framework to conceive, design, and describe knowledge-oriented, self-organising coordination systems.

Main Ideas

- information should autonomously link together
- users should see information spontaneously manifest to and diffuse toward them

Formal MoK I

Atoms

Produced by a source and conveying an atomic piece of information, atoms should also store ontological metadata to ease automatic processing:

$$\texttt{atom}(src, val, attr)^c$$

Molecules

MoK heaps for information aggregation, molecules cluster together semantically related atoms:

$$\texttt{molecule}(Atoms)^c$$

Formal MoK II

Enzymes

Enzymes represent the reification of (epistemic) knowledge-oriented (inter-)actions, and are meant to participate biochemical reactions to properly increase molecules' concentration[a]:

$$\texttt{enzyme}(\mathit{Molecule})^{c}$$

[a] The term "molecule" will be used also for "atom" in the following.

MoK Function

As a knowledge-oriented model, MoK must have a way to determine the semantic correlation between information. Therefore, the MoK function should be defined, taking two molecules and returning a value $m \in [0, 1]$:

$$\mathcal{F}_{mok}\colon \mathit{Molecule} \times \mathit{Molecule} \longmapsto [0, 1]$$

Formal MoK III

Biochemical Reactions

The behaviour of a MoK system is determined by biochemical reactions, which stochastically drive molecules aggregation, as well as reinforcement, decay, and diffusion:

Aggregation — Bounds together semantically related molecules

$$\texttt{molecule}(Atoms_1) \ + \ \texttt{molecule}(Atoms_2) \longmapsto^{r_{agg}} \\ \texttt{molecule}(Atoms_1 \uplus Atoms_2) \ + \ Residual(Atoms_1 \bigcup Atoms_2)$$

Reinforcement — Consumes an enzyme to reinforce the related molecule

$$\texttt{enzyme}(Molecule_1) \ + \ Molecule_1^c \longmapsto^{r_{reinf}} \ Molecule_1^{c+1}$$

Decay — Enforcing time situatedness, molecules should fade away as time passes

$$Molecule^c \longmapsto^{r_{decay}} \ Molecule^{c-1}$$

Diffusion — Analogously, space situatedness is inspired by biology and therefore based upon *diffusion*

$$\{Molecule_1 \bigcup Molecules_1\}_{\sigma_i} + \{Molecules_2\}_{\sigma_{ii}} \longmapsto^{r_{diffusion}} \\ \{Molecules_1\}_{\sigma_i} + \{Molecules_2 \bigcup Molecule_1\}_{\sigma_{ii}}$$

Formal MoK IV

Other aspects like topology, information production and consumption are addressed by:

Compartments — the conceptual *loci* for all other MoK abstractions, providing the notions of locality and neighbourhood

Sources — each one associated to a compartment, MoK sources are the origins of atoms, which are continuously injected at a given rate

Catalysts — the abstraction for knowledge prosumers, who emit enzymes whenever they interact with their compartment

Outline

"Smart migration" I

A "producer" compartment diffuses a collection of different atoms to the neighbouring compartments "economics" and "sports". We expect the system to reach an "equilibrium" in which the two topic-oriented compartments are mainly populated by topic-compliant news molecules.

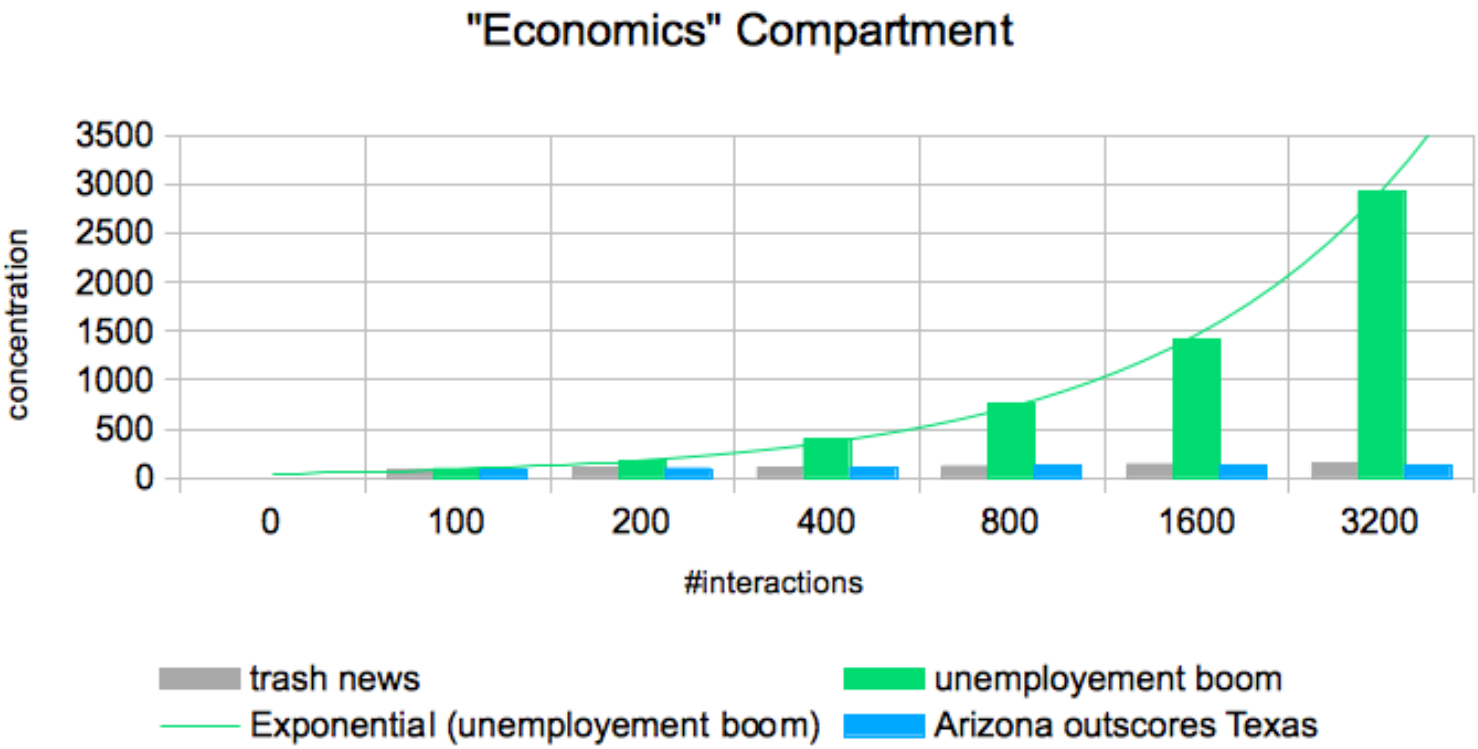

"Smart migration" II

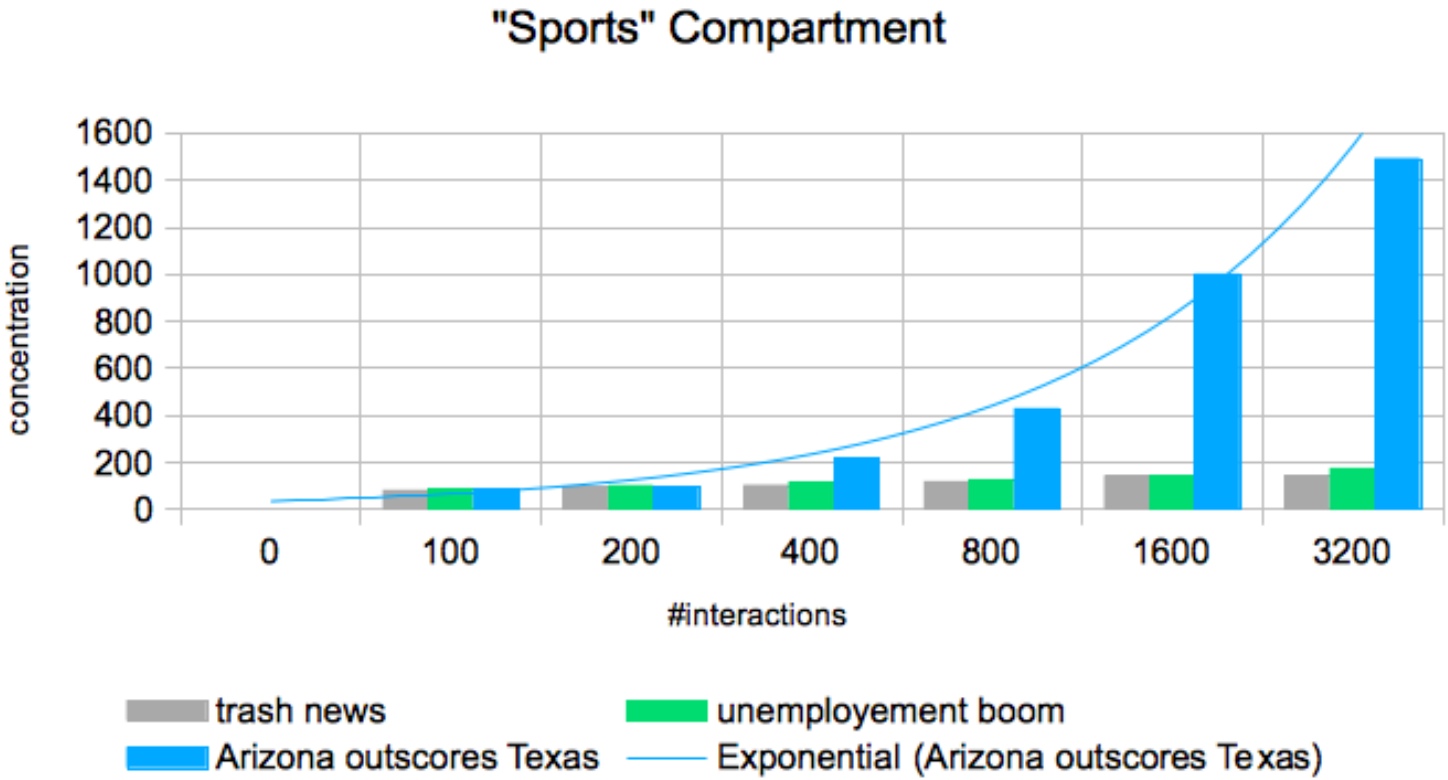

The stochastic equilibrium between diffusion, reinforcement and decay laws, makes a "smart migration" pattern appear by emergence.

Outline

MoK Limitations

MoK currently lacks three of the computational requirements needed to support BIC-based coordination:

- making traces of agents' interactions available for observation to other agents
- making agents' interactions themselves available for other agents inspection
- explicitly record contextual information about such actions

In fact:

- enzymes are consumed and cannot diffuse, hence their observation is restricted to the environment
- contextual information is only implicitly conveyed by enzymes, being them produced at a given time in users' own compartment

MoK Extension I

σ Descriptor

We can keep track of contextual information regarding agents interactions by associating each enzyme to a descriptor σ of the compartment they were released into

$$\texttt{enzyme}(\sigma, \mathit{Molecule})^c$$

Such a descriptor could store any meta-information about the compartment useful to better understand the action: its current time, place, and so on.

MoK Extension II

Then, we should:

- allow enzymes to participate in MoK diffusion reactions, so that both other agents and compartments can observe and use them
- produce a "dead enzyme" whenever enzymes are involved in a reinforcement reaction, subject to decay and diffusion

"Trace" Reinforcement

MoK reinforcement reaction should then be rewritten as follows:

$$\texttt{enzyme}(\sigma, \mathit{Molecule}_1) \; + \; \mathit{Molecule}_1^c \xrightarrow{r_{reinf}} \mathit{Molecule}_1^{c+1} \; + \; \texttt{dead}(\texttt{enzyme}(\sigma, \mathit{Molecule}_1))$$

Furthermore, decay and diffusion reactions should apply respectively to dead enzymes solely (decay) and both enzymes (diffusion).

Outline

Current Prototype

A first prototype of the MoK model is actually implemented upon the TuCSoN coordination infrastructure [Omicini and Zambonelli, 1999] featuring ReSpecT *tuple centres* [Omicini, 2007].

TuCSoN Mapping

- logic tuples are MoK atoms, molecules, enzymes and biochemical reactions declarative specification
- ReSpecT reactions installed in tuple centres implement Gillespie's chemical simulation algorithm [Gillespie, 1977]
- ReSpecT tuple centres represent MoK compartments
- TuCSoN agents reify with MoK sources and users (catalysts)

Next Steps I

Such prototype implementation paves the way toward a much more complex and general idea of self-organising workspaces [Omicini, 2011] a full-fledged MoK system would support.

What Is Done

In this context, MoK already support knowledge aggregation and organisation:

- the former by the molecule abstraction, actually reifying semantic relationships among different information chunks
- the latter by the combined contribution of MoK diffusion, reinforcement, and decay, in which enzymes play a critical role

Next Steps II

Socio-technical system for knowledge-oriented coordination should properly handle pervasive computing scenarios where knowledge is pervasively distributed and is to be accessible ubiquitously.

What Still To Do

The *data in the cloud* paradigm could inspire a novel knowledge in the cloud paradigm, aiming at transforming data clouds into semantic clouds, spontaneously and continuously *self-(re)organising* as a consequence of semantic correlations between information chunks and social actions carried out by knowledge workers.

Outline

Conclusion

- We discussed principles borrowed from cognitive and behavioural (social) sciences...
- ...then showed how to exploit them in computational systems...
- ...thanks to the biochemically-inspired MoK model, aimed to deal with the issue of coordination in knowledge-intensive environments

Future

By sharing some of our ideas regarding the future of socio-technical systems, we hope new collaborations between the research fields of sociology, cognitive sciences, and coordination models and languages will arise.

Thanks to...

...everybody here for listening.

...the SAPERE team for supporting our work[1].

[1]This work has been supported by the EU-FP7-FET Proactive project SAPERE *Self-aware Pervasive Service Ecosystems*, under contract no.256873.

Bibliography I


Bhatt, G. D. (2001).
Knowledge management in organizations: examining the interaction between technologies, techniques, and people.
Journal of Knowledge Management, 5(1):68–75.


Castelfranchi, C., Pezzullo, G., and Tummolini, L. (2010).
Behavioral implicit communication (BIC): Communicating with smart environments via our practical behavior and its traces.
International Journal of Ambient Computing and Intelligence, 2(1):1–12.


Gelernter, D. (1985).
Generative communication in Linda.
ACM Transactions on Programming Languages and Systems, 7(1):80–112.


Gelernter, D. and Carriero, N. (1992).
Coordination languages and their significance.
Communications of the ACM, 35(2):97–107.


Gillespie, D. T. (1977).
Exact stochastic simulation of coupled chemical reactions.
The Journal of Physical Chemistry, 81(25):2340–2361.

Bibliography II


Grassé, P.-P. (1959).
La reconstruction du nid et les coordinations interindividuelles chez Bellicositermes natalensis et Cubitermes sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs.
Insectes Sociaux, 6(1):41–80.


Mariani, S. and Omicini, A. (2013).
Molecules of Knowledge: Self-organisation in knowledge-intensive environments.
In Fortino, G., Bădică, C., Malgeri, M., and Unland, R., editors, *Intelligent Distributed Computing VI*, volume 446 of *Studies in Computational Intelligence*, pages 17–22. Springer.
6th International Symposium on Intelligent Distributed Computing (IDC 2012), Calabria, Italy, 24-26 September 2012. Proceedings.


Omicini, A. (2007).
Formal ReSpecT in the A&A perspective.
Electronic Notes in Theoretical Computer Science, 175(2):97–117.
5th International Workshop on Foundations of Coordination Languages and Software Architectures (FOCLASA'06), CONCUR'06, Bonn, Germany, 31 August 2006. Post-proceedings.

Bibliography III

Omicini, A. (2011).
Self-organising knowledge-intensive workspaces.
In Ferscha, A., editor, *Pervasive Adaptation. The Next Generation Pervasive Computing Research Agenda*, chapter VII: Human-Centric Adaptation, pages 71–72. Institute for Pervasive Computing, Johannes Kepler University Linz, Austria.

Omicini, A. (2012).
Agents writing on walls: Cognitive stigmergy and beyond.
In Paglieri, F., Tummolini, L., Falcone, R., and Miceli, M., editors, *The Goals of Cognition. Essays in Honor of Cristiano Castelfranchi*, volume 20 of *Tributes*, chapter 29, pages 543–556. College Publications, London.

Omicini, A. (2013).
Nature-inspired coordination models: Current status, future trends.
ISRN Software Engineering, 2013.
Article ID 384903, Review Article.

Omicini, A. and Viroli, M. (2011).
Coordination models and languages: From parallel computing to self-organisation.
The Knowledge Engineering Review, 26(1):53–59.
Special Issue 01 (25th Anniversary Issue).

Bibliography IV

Omicini, A. and Zambonelli, F. (1999).
Coordination for Internet application development.
Autonomous Agents and Multi-Agent Systems, 2(3):251–269.
Special Issue: Coordination Mechanisms for Web Agents.

Ossowski, S. and Omicini, A. (2002).
Coordination knowledge engineering.
The Knowledge Engineering Review, 17(4):309–316.
Special Issue: Coordination and Knowledge Engineering.

Papadopoulos, G. A. and Arbab, F. (1998).
Coordination models and languages.
In Zelkowitz, M. V., editor, *The Engineering of Large Systems*, volume 46 of *Advances in Computers*, pages 329–400. Academic Press.

Viroli, M. and Casadei, M. (2009).
Biochemical tuple spaces for self-organising coordination.
In Field, J. and Vasconcelos, V. T., editors, *Coordination Languages and Models*, volume 5521 of *LNCS*, pages 143–162. Springer, Lisbon, Portugal.
11th International Conference (COORDINATION 2009), Lisbon, Portugal, June 2009. Proceedings.

Bibliography V

Weyns, D., Omicini, A., and Odell, J. J. (2007).
Environment as a first-class abstraction in multi-agent systems.
Autonomous Agents and Multi-Agent Systems, 14(1):5–30.
Special Issue on Environments for Multi-agent Systems.

MoK: Stigmergy Meets Chemistry to Exploit Social Actions for Coordination Purposes

Stefano Mariani, Andrea Omicini
{s.mariani, andrea.omicini}@unibo.it

Dipartimento di Informatica: Scienza e Ingegneria (DISI)
Alma Mater Studiorum—Università di Bologna

SOCIAL.PATH 2013
Social Coordination: Principles, Artifacts and Theories
Exeter, UK – 3rd of April 2013