Self-organisation of Knowledge in Socio-technical Systems: A Coordination Perspective

Stefano Mariani, Andrea Omicini

Università di Bologna

ISTC CNR
Rome, Italy
November 6th, 2015

Outline

Challenges of Socio-technical Systems

- Socio-technical systems (STS) arise when *cognitive* and *social interaction* is mediated by *information technology*, rather than by the natural world (alone) [Whi06]
- STS are heavily *interaction-centred*, thus need to deal with coordination issues at the infrastructural level [MC94]
- Among the many coordination issues in STS are:

  unpredictability — *"humans-in-the-loop"* vs. software programmability and predictability

  ⇒ coordination should account natively for *unpredictability* and *uncertainty*

  scale — *large-scale* distribution, openness, ever-increasing number of users, devices, data

  ⇒ coordination should exploit *decentralised* mechanisms to scale in/out upon need

Challenges of Knowledge-intensive Environments

- Knowledge-intensive Environments (KIE) are workplaces in which *sustainability* of the organisation's long-term goals is *influenced* by the evolution of the body of *knowledge* embodied within the organisation itself [Bha01]
- Usually, KIE are computationally supported by STS, thus they need proper coordination too
- Among the many coordination issues in KIE are:
  - size — *massive amount* of raw data, aggregated information, reification of procedures and best-practices, and the like
    - ⇒ coordination should minimise the *overhead* of information needed for coordination-related (non-) functional requirements
  - pace — *high rate* of information production and consumption, huge frequency of interactions
    - ⇒ coordination mechanisms should be as *simple* and *efficient* as possible

A Path to Follow

Coordination models and technologies have already drawn inspiration from *natural systems*, looking for mechanisms enabling and promoting self-organising and adaptive coordination [VPB12, MZ09, VC09, ZCF+11, MO13]

A novel perspective

Similarly, we focus on the human factor in STS, seeking novel coordination approaches inspired by the latest *cognitive* and *social sciences* theories of *action* and *interaction*

The Approach

1. We *observed* real-world STS/KIE, analysing their (implicit) models of action and interaction
2. We *generalised* such models according to the theoretical framework of *Behavioural Implicit Communication* [CPT10], devising out tacit messages and implicit actions computationally exploited by such STS/KIE
3. We *conceived* the $\mathcal{M}$olecules $\mathcal{o}$f $\mathcal{K}$nowledge model [MO13], promoting *self-organisation of knowledge* in STS/KIE, inspired by the above framework and geared toward the notion of self-organising workspace [Omi11]

Outline

Tacit Messages I

Tacit messages are introduced in [CPT10] to describe the kind of message a *practical action* (and its traces) may *implicitly* send to its observers:

presence — "Agent A is here". Any agent (as well as the environment itself) observing any practical behaviour of A becomes aware of its existence — and, possibly, of contextual information, e.g., its location.

intention — "Agent A plans to do action β". If the agents' workflow determines that action β follows action α, peers (as well as the environment) observing A doing α may assume A next intention to be "do β".

ability — "A is able to do $\phi_{i \in \mathbb{N}}$". Assuming actions $\phi_{i \in \mathbb{N}}$ have similar pre-conditions, agents (and the environment) observing A doing ϕ_i may infer that A is also able to do $\phi_{j \neq i \in \mathbb{N}}$.

Tacit Messages II

opportunity — "$p_{i \in \mathbb{N}}$ is the set of pre-conditions for doing α". Agents observing A doing α may infer that $p_{i \in \mathbb{N}}$ hold, thus, they may take the opportunity to do α as soon as possible.

accomplishment — "A achieved S". If S is the "state of affairs" reachable through action α, agents observing A doing α may infer that A is now in state S.

goal — "A has goal g". By observing A doing action α, peers of A may infer A goal to be g, e.g. because action α is part of a workflow aimed at achieving g.

result — "Result R is available". If peer agents know that action α leads to result R, whenever agent A does α they can expect result R to be soon available.

Tacit Messages in Real-world STS I

- We identified a set of (virtual) practical actions, fairly common in real-world STS despite the diversity in scope of each specific STS — e.g. Facebook vs. Mendeley[1] vs. Storify[2]
- For each, we point to a few *tacit messages* they may convey:

quote/share — re-publishing or mentioning someone else's information can convey, e.g., tacit messages *presence*, *ability*, *accomplishment*. If X shares Y's information through action a, every other agent observing a becomes aware of existence and location of both X and Y (*presence*). The fact that X is sharing information I from source S lets X's peers infer X can manipulate S (*ability*). If X shared I with Z, Z may infer that X expects Z to somehow use it (*accomplishment*).

Tacit Messages in Real-world STS II

like/favourite — marking as relevant a piece of information can convey tacit messages *presence*, *opportunity*. If the socio-technical platform lets X be aware of Y marking information I as relevant, X may infer that Y exists (*presence*). If Y marks as relevant I belonging to X, X may infer that Y is interested in her work, perhaps seeking for collaborations (*opportunity*).

follow — subscribing for updates regarding a piece of information or a user can convey tacit messages *intention*, *opportunity*. Since X manifested interest in Y's work through subscription, Y may infer X intention to use it somehow (*intention*). Accordingly, Y may infer the opportunity for collaboration (*opportunity*).

Tacit Messages in Real-world STS III

search — performing a search query to retrieve information can convey, e.g., tacit messages *presence*, *intention*, *opportunity*. If X search query is observable by peer agents, they can infer X existence and location (*presence*). Also, they can infer X goal to acquire knowledge related to its search query (*intention*). Finally, along the same line, they can take the chance to provide matching information (*opportunity*).

Now the question is

How to computationally exploit the envisioned *mind-reading* and *signification* abilities from a coordination perspective?

[1] https://www.mendeley.com

[2] https://storify.com

Perturbation Actions I

Perturbation actions are *computational functions* changing the state of a STS, in *response* to users' *interactions*, but *transparently* to them [MO15]

A possible answer is

Perturbation actions may then exploit the implicit information conveyed by tacit messages to leverage *mind-reading* and *signification* for coordination purposes

Perturbation Actions II

- Accordingly, perturbation actions may:
  - spread discovery messages informing agents about the presence and location of another (tacit message *presence*)
  - establish privileged communication channels between frequently interacting agents (*opportunity*)
  - undertake coordination actions enabling/hindering some desirable/dangerous interaction protocol (*intention*, *ability*, *goal*)
  - autonomously notify users about availability of novel, potentially interesting information (*accomplishment*, *result*)

Perturbation Actions in Real-world STS I

The virtual practical actions already identified are likely to (transparently) cause perturbation actions *under-the-hood*:

quote/share — provided by Facebook, Twitter (retweet), G+, LinkedIN, Mendeley (post), Academia.edu (publish), ResearchGate (publish), Storify, etc. It is likely to help the STS platform, underlying the social network application, in:

- suggesting novel connections
- ranking feeds in the newsfeed timeline

like/favourite — provided by Facebook, Twitter, G+ (+1), LinkedIN (suggest), Mendeley, Academia.edu (bookmark), ResearchGate (follow/download), Storify, etc. It is likely to influence the STS as above.

Perturbation Actions in Real-world STS II

follow — provided by Facebook (add friend), Twitter, G+ (add), LinkedIN (connect), Mendeley, Academia.edu, ResearchGate, etc. It is likely to help the STS platform by:

- suggesting further connections
- activating/ranking feeds in the newsfeed timeline

search — provided by almost every social network, it is the epistemic action by its very definition, thus may be exploited by the STS platform in a number of ways:

- re-organising the knowledge graph internally used by the STS
- tune the algorithm providing suggestions
- improve personalised advertising policies
- and many more. . .

A Systematic Analysis

. . . can we frame the above observations

within a

coherent computational framework

?

Of course we can.

Outline

BIC in a Nutshell

Implicit interaction

Behavioural implicit communication (BIC) is a form of *implicit interaction* where no specialised signal conveys the message, since the message is the *practical behaviour* itself — and possibly, its post hoc *traces* [CPT10]

- BIC presupposes advanced observation capabilities: agents should be able to observe others' actions (and traces), as well as to *mind-read* the intentions behind them, so as to leverage *signification*
- BIC applies to human beings, to both cognitive and non-cognitive agents, and to *computational environments* as well [WOO07]
- Through BIC, such environments can become smart environments, namely pro-active, intelligent workplaces able to *autonomously adapt* their configuration and behaviour according to users' interactions [CPT10]

Cognitive Stigmergy in a Nutshell

Trace-based Communication

The notion of stigmergy has been introduced in the biological study of social insects [Gra59], to characterise how termites (*unintentionally*) coordinate themselves during nest construction, with no need of exchanging direct messages, but relying solely on *local* interactions instead

- Stigmergy is a special form of BIC, where the addressee does not directly perceive the behaviour, but just other *post-hoc traces* — in the form of environment modifications
- Such modifications are amenable of a symbolic interpretation, thus exploitable by agents featuring *cognitive* abilities — either humans or software
- When traces become signs, stigmergy becomes cognitive stigmergy, which involves agents able to correctly understand traces as signs *intentionally* left in the environment [Omi12]

Computational Smart Environments in a Nutshell

- In [TCR+05], an abstract model for computational *smart environments* is proposed, which defines two types of environment

  c-env — common environment, where agents can observe only the state of the environment (including actions' traces), *not the actions* of their peers

  s-env — shared environment, enabling different forms of *observability of actions*, and *awareness* of this observability

- Then, three requirements enabling them are devised
  1. observability of agents' actions and traces should be enabled by default
  2. the environment should be able to *understand* actions and their traces, possibly inferring intentions and goals motivating them
  3. agents should be able to understand the effects of their activity on the environment as well as on the other agents, so as to *opportunistically* obtain a reaction

Outline

Putting All Together

- Quoting from [Omi11], a STS for working in knowledge-intensive environments requires

  > *"that all the relevant information sources are made available to the user in a complete yet usable format, [...] that the working environment autonomously evolves and adapts to the individual uses and work habits"*

- Furthermore,

  > *"the main point here is the explicit representation, memorisation and exploitation of user actions in the workspace"*

Our master equations

BIC + (cognitive) stigmergy = Smart Environments

Smart Environments + Self-organisation = Self-organising Workspaces

Outline

Outline

MoK in a Nutshell I

MoK

Molecules of Knowledge (MoK) is a coordination model promoting *self-organisation* of knowledge [MO13]

- Inspired to *biochemical tuple spaces* [VC09], *stigmergic coordination* [Par06], and *BIC* [CPT10]
- Main goals
    1. self-aggregation of information into more complex heaps, possibly reifying useful knowledge previously hidden
    2. autonomous diffusion of information toward the interested agents, that is, those needing it to achieve their goals

MoK in a Nutshell II

- A *MoK*-coordinated system is
  - a network of *MoK* compartments (*tuple-space* like information repositories)...
  - ...in which *MoK* seeds (*sources* of information) autonomously inject *MoK* atoms (*information* pieces)
  - atoms undergo *autonomous* and *decentralised* reactions
    - aggregate into molecules (*composite* information chunks)
    - diffuse to neighbourhoods
    - get reinforced and perturbed by users
    - decay as time flows
  - reactions are influenced by enzymes (reification of users' *epistemic actions*) and traces (their *(side) effects*)...
  - ...and scheduled according to Gillespie's chemical dynamics simulation algorithm [Gil77]

$\mathcal{M}o\mathcal{K}$ Enzymes and Traces as BIC Enablers I

Model

$$\texttt{enzyme}(species,\ s,\ mol)_c$$

$$\texttt{enzyme}(species,\ s,\ mol') + mol_c \xrightarrow{r_{reinf}} \texttt{enzyme}(species,\ s,\ mol') + mol_{c+s}$$

$$\texttt{trace}(enzyme,\ p,\ mol)_c$$

$$\texttt{trace}(enzyme,\ p,\ mol') + mol_c \xrightarrow{r_{pert}} .exec(p,\ trace,\ mol)$$

$$enzyme \xrightarrow{r_{dep}} enzyme + \texttt{trace}(enzyme,\ p[species],\ mol)$$

species defines the epistemic nature of the action

s *strength* of reinforcement

p the *perturbation* the trace wants to perform

.exec starts execution of perturbation *p* [3]

MoK Enzymes and Traces as BIC Enablers II

- Reinforcement *influences* relevance of information according to the (epistemic) nature and frequency of their *actions* and *interactions*
- Enzymes situate actions, e.g., at a precise *time* as well as in a precise *space*
- *Mind-reading* and *signification* are enabled by assuming that users manipulating a given corpus of information are interested in that information more than other
- Perturbation *influences* location, content, any domain-specific trait of information, according to users' inferred goals, with the goal of easing and optimising their workflows
- Traces enable the environment to exploit users' actions (possibly, inferred) *side-effects* for the profit of the coordination process — promoting the *distributed collective intelligence* leading to anticipatory coordination

[3] Notice, p is implicitly defined by *species*, as highlighted by notation $p[species]$.

Outline

MoK Ecosystem Architecture

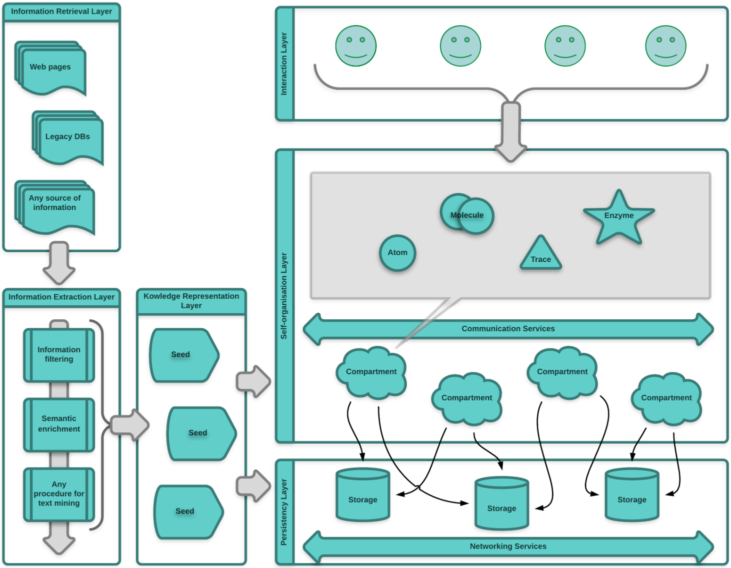

Outline

1. Coordination Issues in STS
2. Tacit Messages and Perturbation Actions in Real-world STS
3. BIC, Stigmergy, and Smart Environments
4. Toward Self-organising Workspaces
5. Molecules of Knowledge
   - Model
   - Architecture
   - **Early Results**
6. Conclusion
7. Outlook

Simulated Scenario

- Citizen journalism scenario:
  - users *share* a MoK-coordinated IT platform for retrieving and publishing news stories
  - they have personal devices (smartphones, tablets, pcs, workstations), running the MoK middleware, which they use to *search* within the IT platform relevant information
  - search actions can spread up to a *neighbourhood* of compartments — e.g., to limit bandwidth consumption, boost security, optimise information location, etc.
  - search actions leave *traces* the MoK middleware exploits to *attract similar information*

Anticipatory coordination

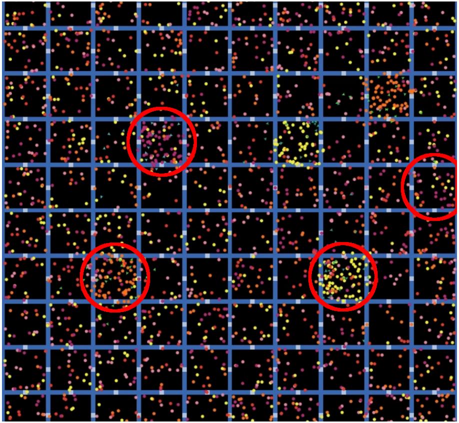

Figure: Whereas data is initially randomly scattered across workspaces, as soon as users interact clusters appear by emergence thanks to *BIC-driven self-organisation*. Whenever new actions are performed by catalysts, the MoK infrastructure *adaptively* re-organises the spatial configuration of molecules so as to better tackle the new coordination needs.

Discussion

- *MoK* anticipates users' needs, not based on behaviour *prediction*, but on present actions and its *mind-reading* and *signification* abilities
  - addressing unpredictability
- *MoK* reactions act only *locally*, thus exploit local information solely
  - addressing scale
- *MoK* decay destroys information as *time* passes — furthermore, the *overhead* brought by *MoK* is minimal, since it exploits solely information already in the system
  - addressing size
- *MoK* reaction execution and BIC-related mechanisms are rather efficient, mostly due to their *local* nature and absence of complex *reasoning*
  - addressing pace

Outline

Summing Up

- Engineering effective coordination for large-scale, knowledge-intensive STS is a difficult task
- Nature-inspired approaches proven successful in mitigating the issue, by leveraging self-organisation and adaptiveness
- We may further improve by shifting attention toward the social side of STS, transparently exploiting the epistemic nature of users' (inter-) actions for coordination purposes

The tools in our hands

BIC, (cognitive) stigmergy, and biochemical coordination give us the right models and approaches to do so

Outline

1. Coordination Issues in STS
2. Tacit Messages and Perturbation Actions in Real-world STS
3. BIC, Stigmergy, and Smart Environments
4. Toward Self-organising Workspaces
5. $\mathcal{M}$olecules *of* $\mathcal{K}$nowledge
6. Conclusion
7. **Outlook**

A Bright and Exciting Future Awaits I

- The world needs *efficient* and *smart* ways of preserving, managing, and analysing the astonishing amount of knowledge it produces and consumes every day
- Big data approaches are more or less the standard now, mostly because they are good in finding *patterns* of knowledge, but:
  - they mostly fail in discovering *anti-patterns*, e.g., detecting outliers
  - they mostly fail in accommodating *ever-changing*, *heterogeneous* knowledge discovery needs
  - they mostly neglect "humans-in-the-loop", relying on algorithms and measures (e.g. of similarity) which are completely *user-neutral* and *goal-independent*
  - they won't scale forever
  - they are not suitable for pervasive and privacy-demanding scenarios

A Bright and Exciting Future Awaits II

- We are in the perfect spot to start a paradigm shift toward self-organising knowledge, where:
  - user-centric adaptiveness of knowledge discovery processes is the foremost goal
  - measures and algorithms exploited for knowledge discovery, inference, management and analysis natively account for users' goals
  - seamlessly scale up/down/out/in naturally, being operating on the assumption that only local-information is available consistently
- As witnessed by the latest H2020 calls, increasingly demanding *user-inclusive* policy making, governance *participation*, *user-centric* knowledge sharing platforms, etc.
  - H2020-SC6-CO-CREATION-2016-2017
  - H2020-EINFRA-2016-2017
  - H2020-FETPROACT-2016-2017

References I

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

Cristiano Castelfranchi, Giovanni Pezzullo, and Luca Tummolini.
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, January–March 2010.

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

Pierre-Paul Grassé.
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, March 1959.

References II

Thomas W. Malone and Kevin Crowston.
The interdisciplinary study of coordination.
ACM Computing Surveys, 26(1):87–119, 1994.

Stefano Mariani and Andrea Omicini.
Molecules of Knowledge: Self-organisation in knowledge-intensive environments.
In Giancarlo Fortino, Costin Bădică, Michele Malgeri, and Rainer Unland, editors, *Intelligent Distributed Computing VI*, volume 446 of *Studies in Computational Intelligence*, pages 17–22. Springer, 2013.

Stefano Mariani and Andrea Omicini.
Anticipatory coordination in socio-technical knowledge-intensive environments: Behavioural implicit communication in MoK.
In Marco Gavanelli, Evelina Lamma, and Fabrizio Riguzzi, editors, *AI*IA 2015, Advances in Artificial Intelligence*, volume 9336 of *Lecture Notes in Computer Science*, chapter 8, pages 102–115. Springer International Publishing, 23–25 September 2015.
XIVth International Conference of the Italian Association for Artificial Intelligence, Ferrara, Italy, September 23–25, 2015, Proceedings.

References III

Marco Mamei and Franco Zambonelli.
Programming pervasive and mobile computing applications: The TOTA approach.
ACM Transactions on Software Engineering and Methodology (TOSEM), 18(4), July 2009.

Andrea Omicini.
Self-organising knowledge-intensive workspaces.
In Alois Ferscha, 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, May 2011.

Andrea Omicini.
Agents writing on walls: Cognitive stigmergy and beyond.
In Fabio Paglieri, Luca Tummolini, Rino Falcone, and Maria Miceli, editors, *The Goals of Cognition. Essays in Honor of Cristiano Castelfranchi*, volume 20 of *Tributes*, chapter 29, pages 543–556. College Publications, London, December 2012.

H. Van Dyke Parunak.
A survey of environments and mechanisms for human-human stigmergy.
In Danny Weyns, H. Van Dyke Parunak, and Fabien Michel, editors, *Environments for Multi-Agent Systems II*, volume 3830 of *Lecture Notes in Computer Science*, pages 163–186. Springer, 2006.

References IV

Luca Tummolini, Cristiano Castelfranchi, Alessandro Ricci, Mirko Viroli, and Andrea Omicini.
"Exhibitionists" and "voyeurs" do it better: A shared environment approach for flexible coordination with tacit messages.
In Danny Weyns, H. Van Dyke Parunak, and Fabien Michel, editors, *Environments for Multi-Agent Systems*, volume 3374 of *Lecture Notes in Artificial Intelligence*, pages 215–231. Springer, February 2005.

Mirko Viroli and Matteo Casadei.
Biochemical tuple spaces for self-organising coordination.
In John Field and Vasco T. Vasconcelos, editors, *Coordination Languages and Models*, volume 5521 of *Lecture Notes in Computer Science*, pages 143–162. Springer, Lisbon, Portugal, June 2009.

Mirko Viroli, Danilo Pianini, and Jacob Beal.
Linda in space-time: An adaptive coordination model for mobile ad-hoc environments.
In Marjan Sirjani, editor, *Coordination Models and Languages*, number 7274 in Lecture Notes in Computer Science, pages 212–229. Springer, 2012.

References V

Brian Whitworth.
Socio-technical systems.
Encyclopedia of human computer interaction, pages 533–541, 2006.

Danny Weyns, Andrea Omicini, and James J. Odell.
Environment as a first-class abstraction in multi-agent systems.
Autonomous Agents and Multi-Agent Systems, 14(1):5–30, February 2007.

Franco Zambonelli, Gabriella Castelli, Laura Ferrari, Marco Mamei, Alberto Rosi, Giovanna Di Marzo, Matteo Risoldi, Akla-Esso Tchao, Simon Dobson, Graeme Stevenson, Yuan Ye, Elena Nardini, Andrea Omicini, Sara Montagna, Mirko Viroli, Alois Ferscha, Sascha Maschek, and Bernhard Wally.
Self-aware pervasive service ecosystems.
Procedia Computer Science, 7:197–199, December 2011.

URLs

Slides

- On APICe

  http://apice.unibo.it/xwiki/bin/view/Talks/MokCnr2015

- On SlideShare

  http://www.slideshare.net/andreaomicini/selforganisation-of-knowledge-in-sociotechnical-systems-a-coordination-perspective

MoK

- On APICe

  http://mok.apice.unibo.it/

Self-organisation of Knowledge in Socio-technical Systems: A Coordination Perspective

Stefano Mariani, Andrea Omicini

Università di Bologna

ISTC CNR
Rome, Italy
November 6th, 2015