A Conceptual Framework for Self-Organising MAS

Andrea Omicini, Alessandro Ricci, Mirko Viroli

DEIS, Sede di Cesena – Alma Mater Studiorum, Università di Bologna

Cristiano Castelfranchi, Luca Tummolini

ISTC-CNR - Roma, Italy

Outline

- “Cognitive” self-organisation
  - SOS why, where, and how
  - stigmergy, environment & intelligent agents
- BIC & beyond
  - behavioural implicit communication (BIC)
  - generalisation of BIC
- Shared Environment (s-env)
  - observability in s-env & epistemic actions
  - formalisation
- Conclusion

SOS why?

- SOS (Self-Organising Systems) typically exhibit desirable properties like
  - robustness
  - fault-tolerance
  - adaptation to change
- Computational SOS are meant to subsume the same nice features
  - self-healing, self-repairing
  - self-configuring, self-adapting

SOS where?

- Physical systems
  - magnetic materials
- Biological systems
  - cytoskeletal filaments in cytoplasm of eukariotic cells
- Social systems
  - ant nests, swarms
  - human systems
    - “sponteous parking patterns” (Castelfranchi)
- Computational systems
  - peer-to-peer systems

SOS what?

- SOSs are systems that
  - exhibit some forms of *global* order / direction
    - organisation, structure, architecture
  - which emerges from non-ordered, non-directed *local* behaviours / interactions
- As a result, definitory features of SOSs are
  - lack of centralised control
  - locality of interaction between components

Computational SOS bias

- “Intelligent” global behaviour
  - in a very broad sense
- vs. “non-intelligent” individual components
  - where intelligence is not a fundamental feature for individuals
  - e.g.: the behaviour of an ant nest is far more “intelligent” when compared to the single ant’s one
- Seeley 2002
  - “a fundamental flaw in many studies of self-organisation: the assumption that the subunits of a self-organised system are dumb”
- Computational SOS seem biased along this line

The MAS shift

- Local interaction based on agent communication
  - direct interaction
  - non-mediated interaction
- MAS self-organisation based on social interaction
  - communication, negotiation, coordination
- Example: AMAS theory
  - there, self-organisation depends on the ability of the agents to be *locally* “cooperative”
  - based on their ability to subjectively interpret interactions with other agents and the environment
    - cognitive abilities at play
- *However!*
  - it is always EITHER mediated interaction & dumb agents OR direct interaction & cognitive agents

Some points of ours here

- Agents are not ants
  - environment-based coordination and mediated interaction not only for reactive agents, but also for cognitive/intelligent agents
- Generalise stigmergy for MAS
  - mediated interaction
    - like pheromon-based stigmergy
  - with cognitive agents
    - unlike pheromon-based stigmergy
- Understand the role of the (MAS) environment
  - shaping the environment to enable / promote cognitive self-organisation
- Understand the role of the (MAS) infrastructure
  - to shape the environment
  - through suitably expressive abstractions

Interaction *is* mediated

- Social activity is mediated
  - Social-psychological theories
    - Activity Theory, Distributed Cognition
  - AT
    - agent activity in societies / organisations is always mediated by artifacts
      - physical, cognitive, …
- Direct interaction results from an abstraction process
  - abstracting away from the interaction medium
  - sometimes applicable, sometimes not
    - which generates the distinction between direct vs. indirect interaction
- Interaction is always mediated
  - we have just to understand the nature & the role of the medium

What is communication?

- Communication is interaction plus intentionality?
  - Palo Alto: “Any behaviour is communication”
  - An agent selects a behaviour aim at informing another agent
    - typically, a *codified* behaviour
  - *But!*
- Explicit communication is only a part of the story in complex societies
  - Humans and animals usually communicate with no need of codified [= rigid] patterns of action
    - for instance, teaching by example
- Requirements
  - *Observation* / observability *of actions*
  - *Awareness* of observation

Coordination without communication

- More generally, many patterns of interaction / coordination
  - do not require explicit communication
  - do not even require *any* kind of communication
    - e.g., the prey / predator pattern
    - e.g. tracking pattern might be either of the two
      - depending on the intentions of the tracked
- In general
  - observation of actions
  - awareness
    - possibly mutual
  - play a central role in a number of fundamental patterns of coordination within social systems
- Also, they are at the core of several interesting SOS phenomena in social systems

Behavioural Implicit Communication

- The agent (*source*) performs a usual *practical action* (like eating, walking, sitting, cleaning etc.)
- The agent also *knows* (*awareness*) and lets (or makes) the other agent (*addressee*) *observe* and *understand* such a behavior
  - i.e. capture some meaning from that “message”
- because this is part of the agent (motivating or non-motivating) goals in performing that action
- Note
  - stigmergy is a special form of BIC which deals with post-hoc traces of actions
    - no observation of the action in stigmergic coordination

BIC examples

- Starting to cross the road is a tacit message for the car driver to stop
  - then, the car driver might as well ignore your message…
- While moving a table together, the feeling of the other person's movements holding the table enables coordination
  - often, communicatin here mixes codified & non codified messages
- The broom hampering the entrance in the toilet is an easily removable obstacle that is used mainly as a message "do not enter, it is wet"
  - this is a form of stigmergy
- The safe footprints of a scout in a mined field are messages "put your foot here"
  - another form stigmergy

Requirements for BIC

- Enabling conditions for BIC
  - observability of practical actions *and* of their traces
    - the environment might either enable observability or prevent it
  - ability to understand and interpret actions
  - ability to understand the other's perception / understanding of actions
    - the environment might allow an agent to know who is observing, and how it is reacting
- And, of course, in the MAS case
  - an infrastructure might well account for a suitably-shaped environment

Forms of Observation-based Coordination (I)

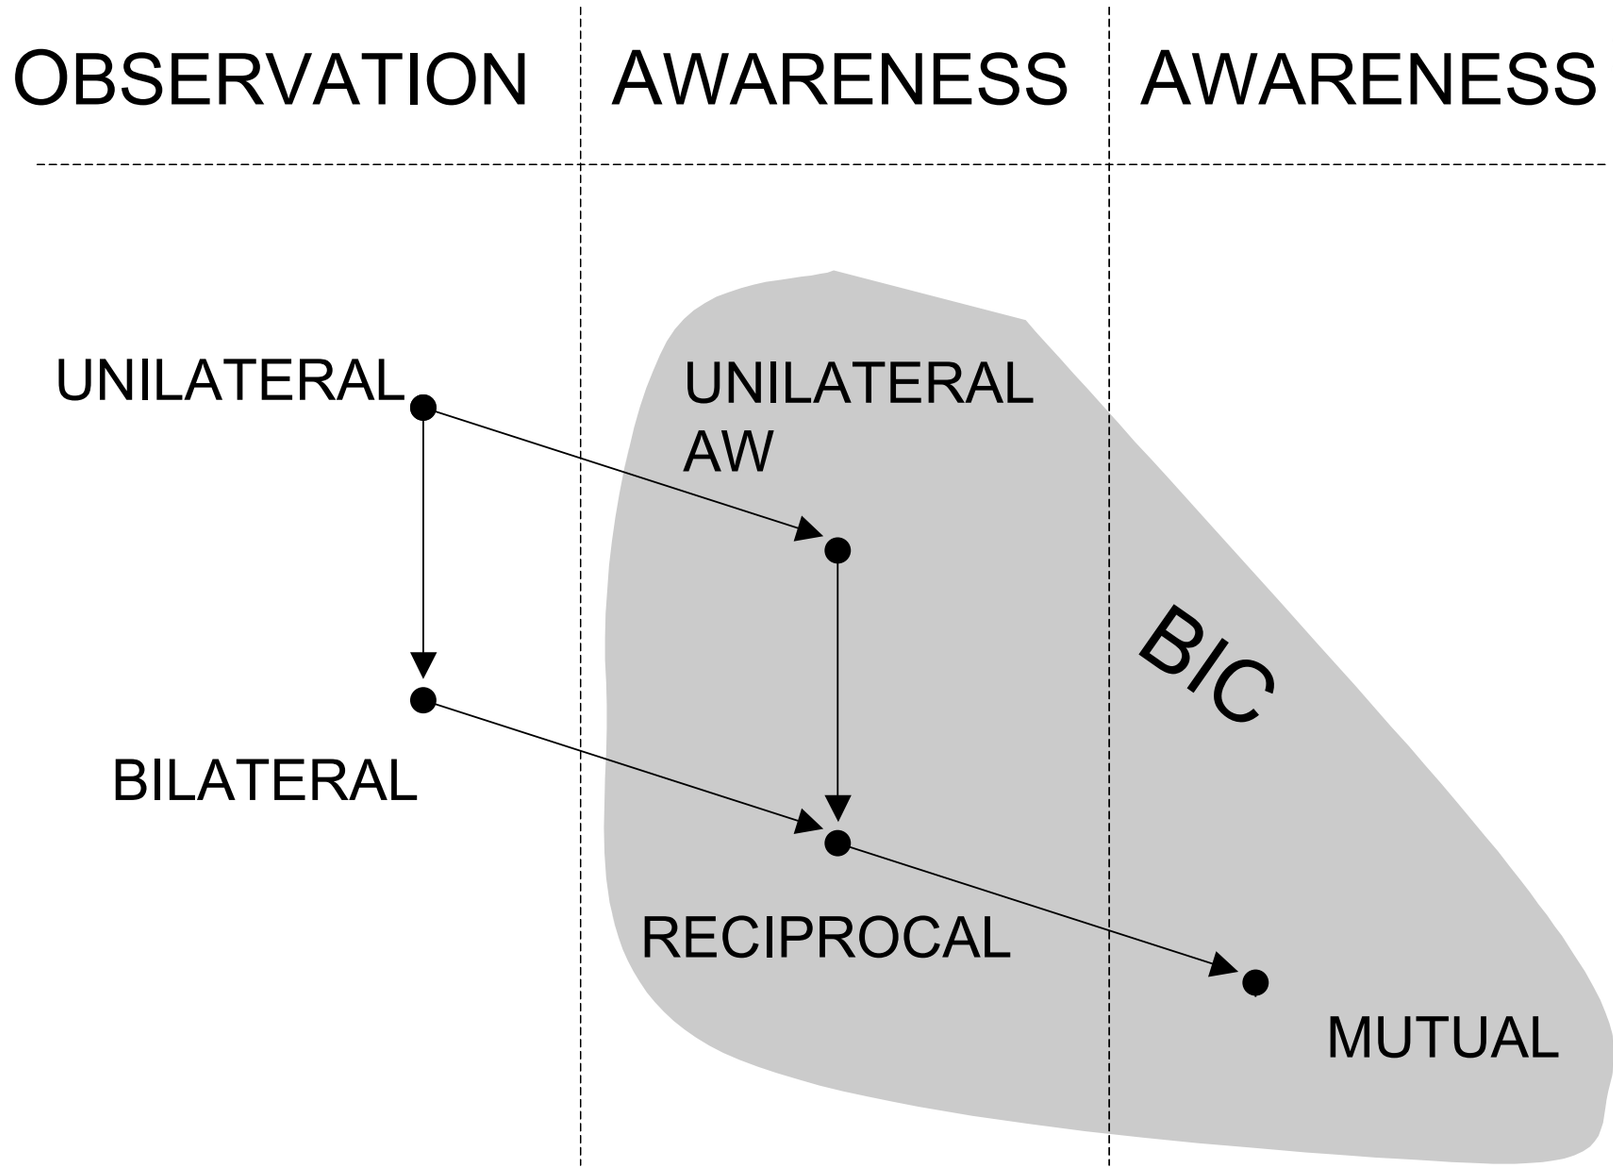

Forms of Observation-based Coordination (II)

$$Uni(x, y, \alpha, S) \triangleq$$
$$Obs(x, y, \alpha) \in S \ \wedge \ I_x coord(x, y, \alpha)$$

$$UniAW(x, y, \alpha, S) \triangleq$$
$$Uni(x, y, \alpha, S) \ \wedge \ B_y obs(x, y, \alpha) \in S$$

$$Bi(x, y, \alpha, S) \triangleq Uni(x, y, \alpha, S) \ \wedge \ Uni(y, x, \alpha, S)$$

$$Rec(x, y, \alpha, S) \triangleq$$
$$UniAW(x, y, \alpha, S) \ \wedge \ UniAW(y, x, \alpha, S)$$

$$Mut(x, y, \alpha, S) \triangleq Rec(x, y, \alpha, S)$$
$$\wedge \ B_x I_y coord(y, x, \alpha) \ \wedge \ B_y I_x coord(x, y, \alpha)$$

The role of BIC in dynamic social order

- Global social order cannot be mainly created and maintained by explicit and formal norms, supported only by a centralised control, formal monitoring, reporting and surveillance protocols
- *Social order* needs to be *self-organising*, spontaneous and informal, with spontaneous and decentralised forms of control and of sanction
- Examples
  - imitation for rule propagation
  - fulfilment of social commitments
  - local reissuing of norms

Shaping the environment

- Advanced forms of *cognitive self-organisation* require a suitably shaped environment
  - what does “suitably shaped” means?
- Common environment (c-env)
  - allowing agents to to keep track of its state and evolution, and possibly affect it
- Shared environment (s-env)
  - a c-env, that also enables
    - different forms of observability of other agents’ actions
    - awareness of such observability
      - that is, an agent is allowed to know is someone is observing its actions
- Note
  - a s-env inherently supports unilateral, bilateral, reciprocal and mutual coordination

Defining s-envs

- The level of observability of a s-env defines the s-env itself
- Observability is expressed in terms of
  - *Pow(x, y, α)*
    - the power of agent *x* to observe action α executed by agent *y*
    - Power relation
      - describes the set of opportunities and constraints for agent observation in a s-env
  - *Obs(x, y, α)*
    - the fact that the environment is making *x* observe actions α executed by agent *y*
    - Observability relation
      - describes the state of observation in a s-env
  - *Pow* vs. *Obs.* is potential vs. actual
- To be fully understood, *Pow* and *Obs* requires the agent viewpoint over observation to be accounted for

Epistemic State

- Epistemic State (ES)
  - the beliefs the agent has due to its observation role
- The ES of an agent includes its *environmental knowledge*
  - knowledge about the agents it is observing
  - knowledge about the agents that are observing it
  - knowledge about the action execution it is observing
- modelled as agent beliefs
  - *e.g. Bz obs(x,y,*α)
- ES evolves through Epistemic Actions
  - actions aimed at acquiring knowledge

Motivational State

- Motivational State (MS)
  - includes all intentions of an agent at a given time
- MS and Epistemic Actions
  - intentions to acquire knowledge
    - to observe another agent
    - to check whether the agent is observed
    - to stop observing another agent

MAS configuration & evolution

- A MAS configuration is a composition of both agent and environmental properties
  - environment configuration
    - a composition of *Pow* and *Obs* terms
  - agent configuration
    - a composition of mental properties
      - beliefs *B* and intentions *I*
- MAS evolution
  - the environment reacts to the MS updating the ES according to the rewrite rules specified in the operational semantics

Syntax of MAS configurations

S	$::=$	$0 \mid A \mid E \mid S \,\|\|\, S$	MAS configuration
E	$::=$	0	environment configuration
	$\mid$	$Pow(x, y, \alpha)$	x has the power to observe y's α
	$\mid$	$Obs(x, y, \alpha)$	x is observing y's α
	$\mid$	$E \,\|\|\, E$	composition
A	$::=$	0	agent configuration
	$\mid$	$B_x \phi$	belief of x
	$\mid$	$I_x \phi$	intention of x
	$\mid$	$A \,\|\|\, A$	composition
ϕ	$::=$		formulas
		$obs(x, y, \alpha)$	x is observing y's α
	$\mid$	$coord(x, y, \alpha)$	x coordinates with y through α
	$\mid$	$check(x, y, \alpha)$	check whether x is observing y's α
	$\mid$	$drop(x, y, \alpha)$	prevent x from observing y's α
	$\mid$	$done(x, \alpha)$	x executes actions α
	$\mid$	$\neg\phi \mid I_x \phi \mid B_x \phi$	structured formulas

Operational semantics of agent configurations

$$\frac{Obs(x,y,\alpha) \in S}{I_z check(x,y,\alpha) \,||\, S \rightarrow B_z obs(x,y,\alpha) \,||\, S} \quad \text{[CHECK]}$$

$$\frac{Obs(x,y,\alpha) \notin S}{I_z check(x,y,\alpha) \,||\, S \rightarrow B_z \neg obs(x,y,\alpha) \,||\, S} \quad \text{[N-CHECK]}$$

$$\frac{-}{I_z drop(x,y,\alpha) \,||\, B_z obs(x,y,\alpha) \,||\, Obs(x,y,\alpha) \,||\, S \rightarrow B_z \neg obs(x,y,\alpha) \,||\, S} \quad \text{[DROP-Y]}$$

$$\frac{Obs(x,y,\alpha) \notin S}{I_z drop(x,y,\alpha) \,||\, B_z obs(x,y,\alpha) \,||\, S \rightarrow B_z \neg obs(x,y,\alpha) \,||\, S} \quad \text{[DROP-N]}$$

$$\frac{-}{I_z obs(x,y,\alpha) \,||\, Pow(x,y,\alpha) \,||\, S \rightarrow B_z obs(x,y,\alpha) \,||\, Pow(x,y,\alpha) \,||\, Obs(x,y,\alpha) \,||\, S} \quad \text{[ASK]}$$

$$\frac{I_x done(x,\alpha) \,||\, S \rightarrow I_x done(x,\alpha) \,||\, S'}{I_x done(x,\alpha) \,||\, Obs(y,x,\alpha) \,||\, S \rightarrow I_x done(x,\alpha) \,||\, Obs(y,x,\alpha) \,||\, B_y done(x,\alpha) \,||\, S'} \quad \text{[OBS-R]}$$

$$\frac{Obs(y,x,\alpha) \notin S}{I_x done(x,\alpha) \,||\, S \rightarrow B_x done(x,\alpha) \,||\, S} \quad \text{[OBS-F]}$$

$$\frac{-}{A \,||\, S \rightarrow A' \,||\, S} \quad \text{[AGENT]}$$

Notes upon the formal model

- Why a formal model if (for now) it is not used to prove properties?
  - typical remark / criticism of awful reviewers…
- The formal system forced us to *obliterate ambiguities* obvious emerging from the cooperation between research groups with different experiences, competences and skills
- The formal system is simple enough to work as a element of *clarity* and understanding for the reader, rather than to introduce further complexity to the global picture
- The formal system works as a *specification* for our infrastructures and systems

S-env in TuCSoN

- The required features of the shared environment translate to the requirements for the MAS infrastructure
  - believe it or not, we are using TuCSoN :)
- Each agent is assigned its own ACC (agent coordination context) that records the allowed agent's actions
  - *Pow(x,y,α)*
    - The ACC labels the actions that are observable
      - publish-subscribe like mechanism
  - *Obs(x,y,α)*
    - The intention *Ix obs(x,y,α)* of agent *x* to observe agent *y*'s action makes the infrastructure enable the observation
    - The ACCs record any observation action of their agents
  - Awareness
    - ACCs provide agents with a service to detect if they are observed, by connecting epistemic knowledge from different agents / ACCs

A Possible Implementation in TuCSoN

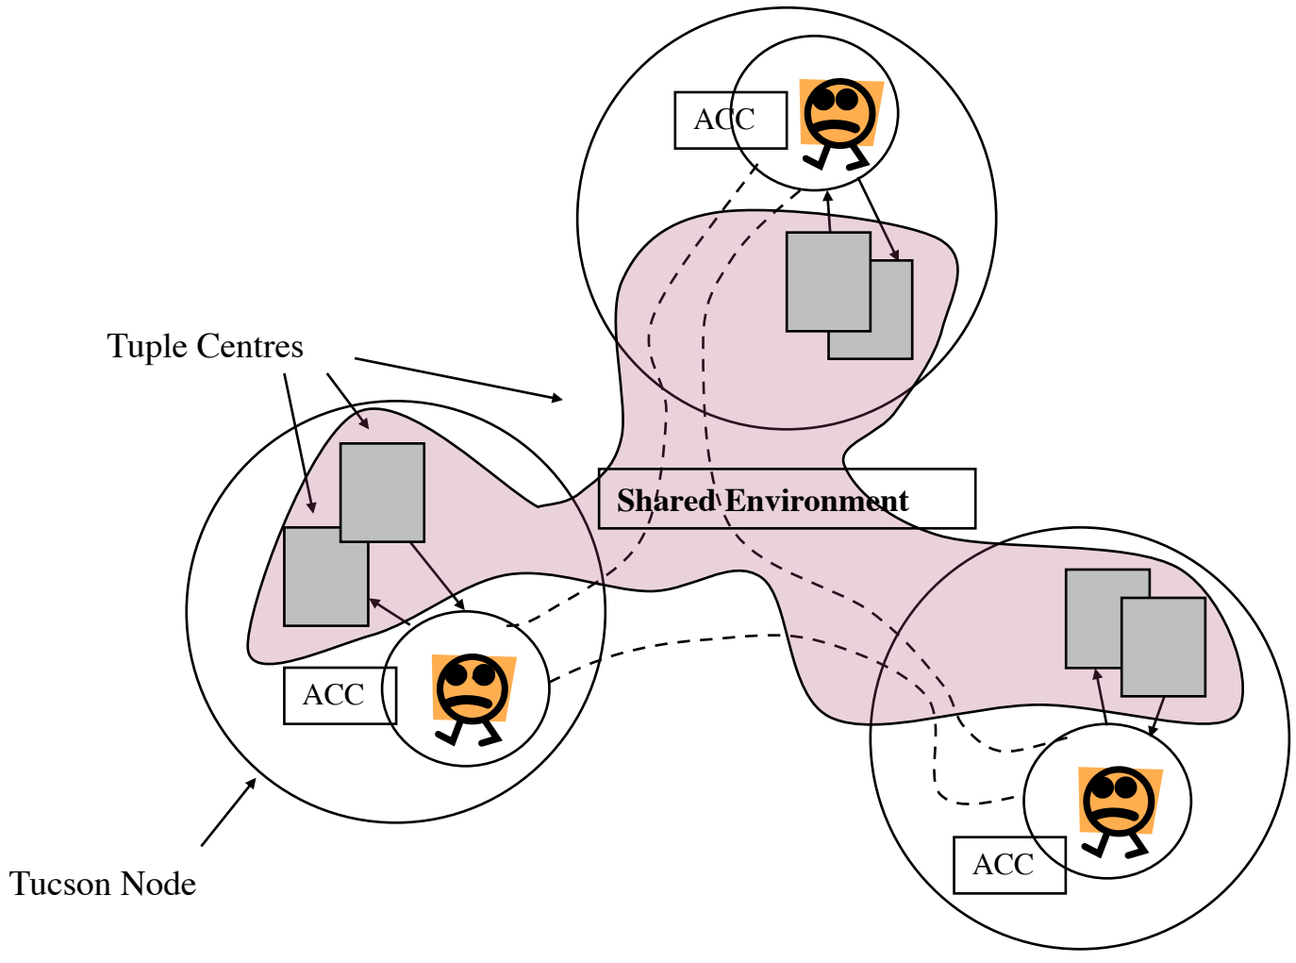

Summing up

- MASs built on top of a BIC-oriented infrastructure exhibit the basic enabling principles which typically characterise self-organisation
  - local interaction
  - decentralised control
  - emergent patterns
- Besides, other interesting principles of SOS can be re-casted in our framework
  - individual-based models
  - (thermodynamic) opennes
  - non-linearity & (positive) feedback
  - dissipative structures
- Along this line, complex systems engineers might find a solution to the “global vs. local control” dichotomy

A Conceptual Framework for Self-Organising MAS

Andrea Omicini, Alessandro Ricci, Mirko Viroli

DEIS, Sede di Cesena – Alma Mater Studiorum, Università di Bologna

Cristiano Castelfranchi, Luca Tummolini

ISTC-CNR - Roma, Italy