Regulating agent interactions in open MAS

Matteo Baldoni e Viviana Patti
Dipartimento di Informatica
University of Torino

Outline

- Part I:
  - Interaction and communication in open MAS: motivation and interoperability issues
  - ACL semantics and standards;
  - Regulating interactions
    - Interaction protocols as patterns of behavior
    - Procedural approach
    - A-priori (design time) verification of interaction properties: interoperability and conformance
- Part II: Commitment Protocols

Interaction in open MAS

- Crucial issue. Why?
- Interaction and communication are fundamental issues of any distributed application, but they becomes even more challenging in open contexts and when cross-business and business-to-business systems are to be developed.

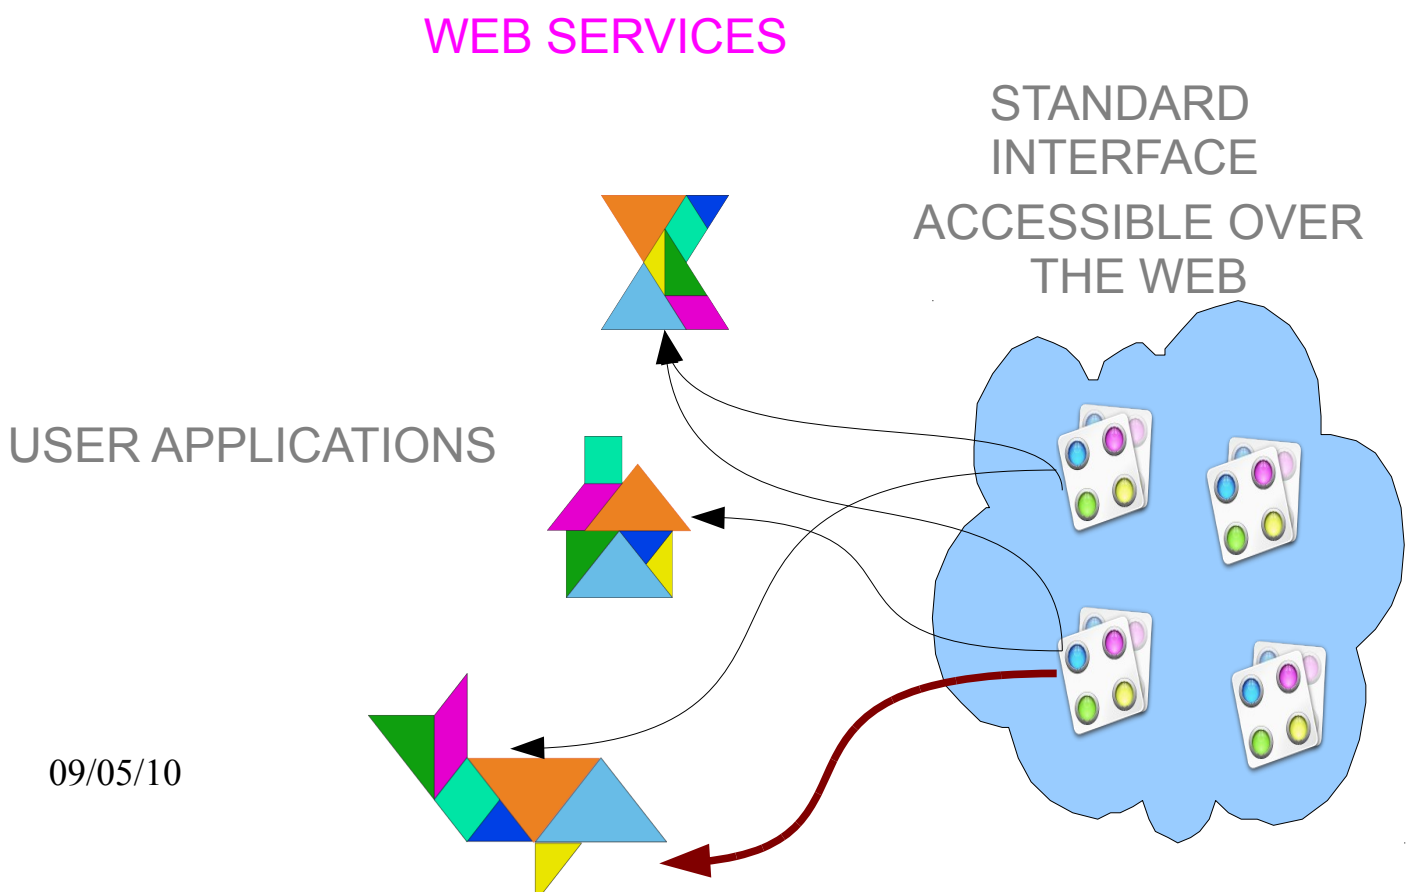

Open distributed systems

- Since the advent of the Internet the landscape of computer science has substantially changed.
- It has become feasible to develop open distributed system
  - components can aggregate dynamically and are free to enter and leave an interaction at their will.
  - the interacting components are typically designed and implemented by different parties
  - may represent conflicting interests, as it happens for example in e-commerce applications.

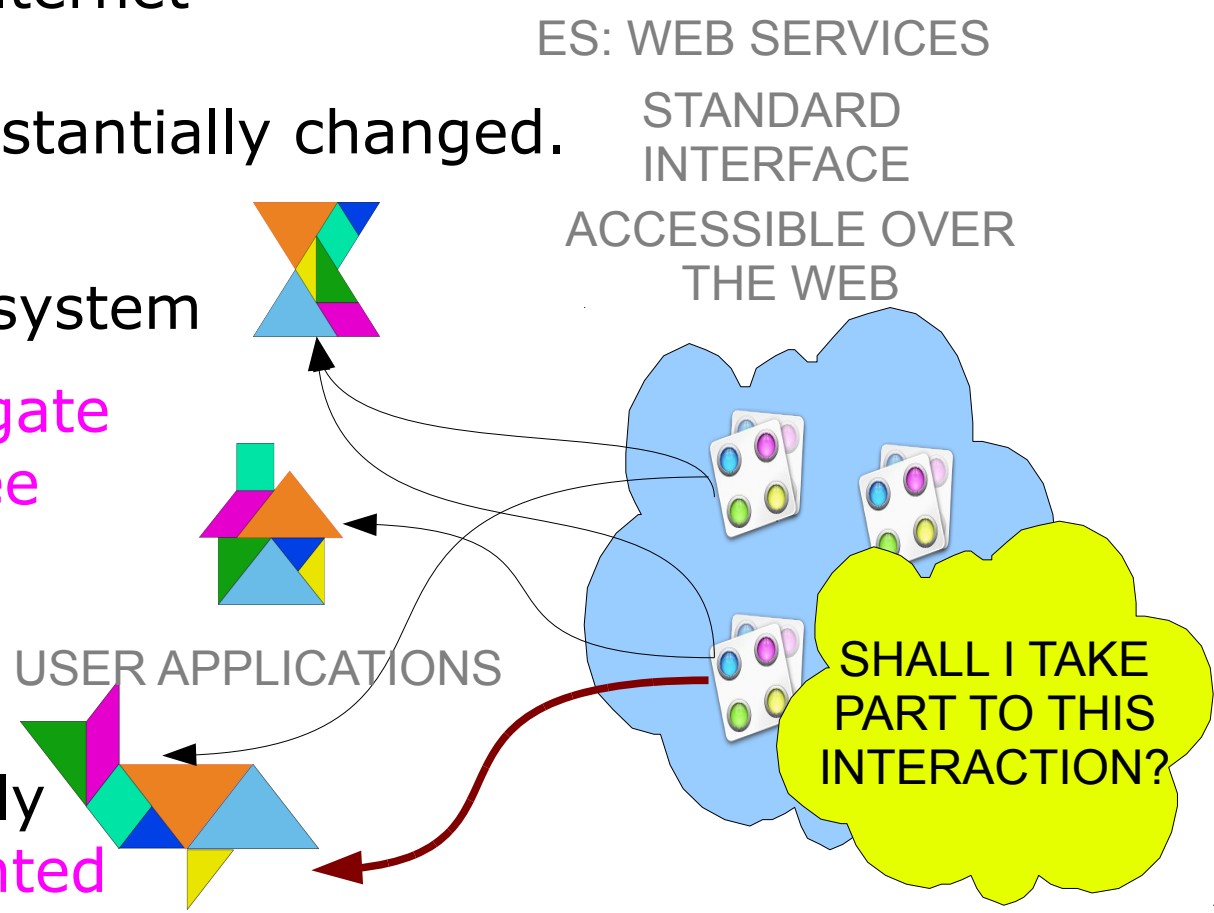

B2B and cross-business solutions

- Interaction and communication are fundamental issues of any distributed application, but they becomes even more challenging in open contexts and when cross-business and business-to-business systems are to be developed.

Support for C-Business collaborative production of goods and services

Facilitate development of new products

Cooperatively exploit resources

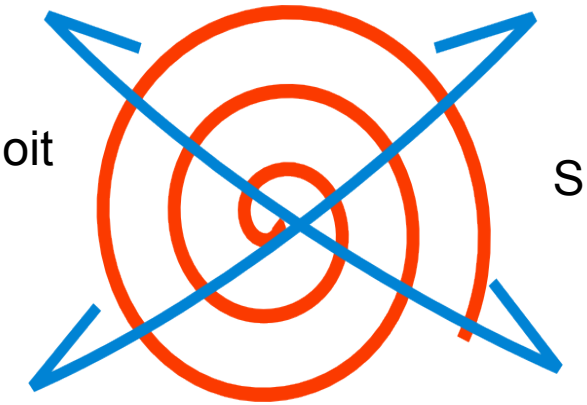

Share best practices

Support for Business integration

Interaction based on agreed contracts

Here a group of heterogeneous and antecedently existing entities need to interact for some time, to share resources, to integrate their capabilities, stick to contracts, etc.

MAS

- Multi-agent systems (MAS) are the tools that currently better meet the needs emerging in this context:
  - Interaction and communication and coordination are being studied since long and are at the core of research on MAS
  - MAS offer proper abstractions

Facilitate development of new products

Cooperatively exploit resources

Share best practices

Integration

Interaction based on agreed contracts

All central issues to the area of MAS!
MAS seem promising for developing B2B/cross-business systems.

Interaction and communication

- Interaction and communication and coordination are being studied since long and are at the core of research on MAS
- The ability to communicate is one of the salient properties of MAS
- "An important part of the agent approach is the principle that agents (like humans) can function more effectively in groups that are characterized by cooperation and division of labor. Agent programs are designed to autonomously collaborate with each other in order to satisfy both their internal goals, and the external demands generated by virtue of their participation in agent societies. The type of collaboration depends on a sophisticated system of inter-agent communication"
- [Dignum and Greaves, 2000].

Open MAS

- When using MAS for specifying open systems we must take into account the characteristics of open contexts
  - In open MAS agents are contributed by several sources and serve different interests
- MAS offer the proper abstractions
  - Crucial notions: autonomy, heterogeneity;

    Agents must be treated as:
  - autonomous: few constraints on behavior reflecting the independence of their users
  - heterogeneous: few constraints on construction, reflecting the independence of the designers

Interaction and communication in open MAS

- MAS offer proper abstractions for dealing with challenging issues in open distributed systems which concern interoperability
- We will focus on:

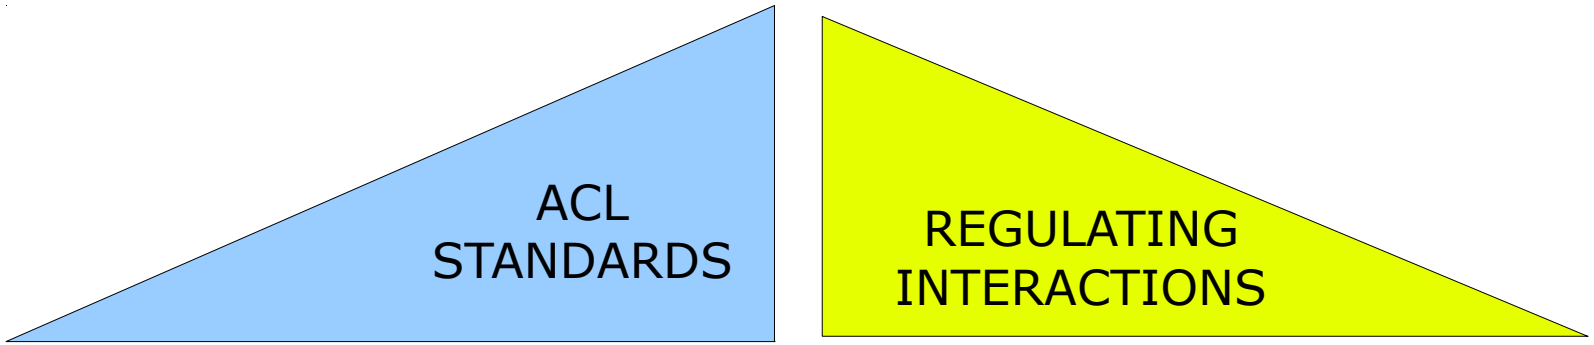

Open MAS, ACL and standards

- Software agents' ability to interact within different open systems, designed by different groups, presupposes an agreement on an unambiguous definition of a set of concepts, used to describe the context of the interaction and the communication language the agents can use.
- Interaction would not even be possible unless agents are designed to comply with well-defined standards.

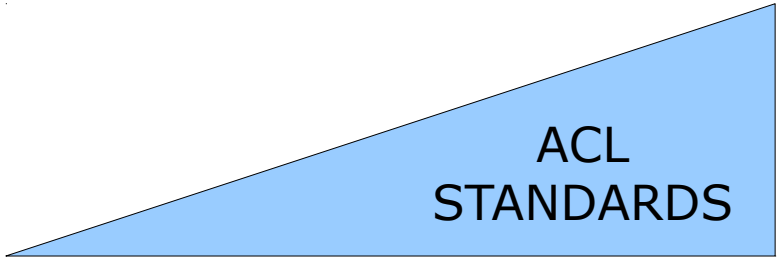

Open MAS, ACL and standards

- The diffusion of open multi-agent system has led the agent community to focus on the creation of standardized communication languages (ACL), that, having an explicit, general and well-defined semantics, could be used by heterogeneous agent programs
- It gives an answer to the interoperability issue at the semantic level: the ability of two or more software systems of exchanging information and of automatically interpreting its meaning.

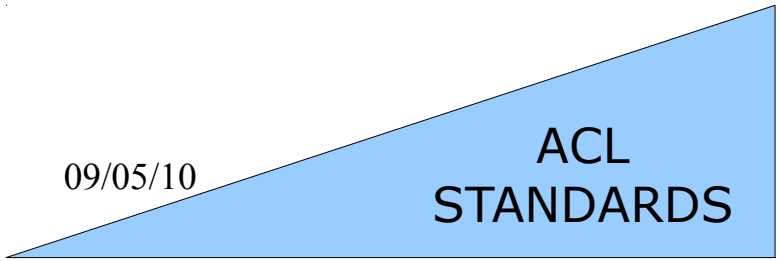

Open MAS, ACL and standards

- The common target indeed was making possible for "agents built by different organizations, using different hardware and software platforms, to safely communicate with one-another via a common language with a universally agreed semantics" [Wooldridge, 2000].
- One example is the FIPA Agent Communication Language, that specifies the semantics of a set of utterances that agents can use to exchange messages. [FIPA 1997]
- KQML [Labrou, 1997; Finin et al., 1995]

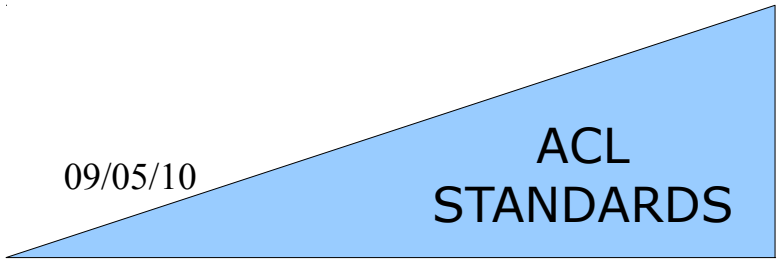

Open MAS, ACL and standards

- Openness asks for a clear semantics for ACL showing the following characteristics:
  - formal/declarative/verifiable (Wooldridge, Guerin & Pitt, Singh)
- Two main approaches
  - Mentalistic semantics (FIPA 1997-2002, Labrou, 1997; Finin et al., 1995)
  - Social semantics (Singh 1999, Verdicchio & Colombetti, 2009)

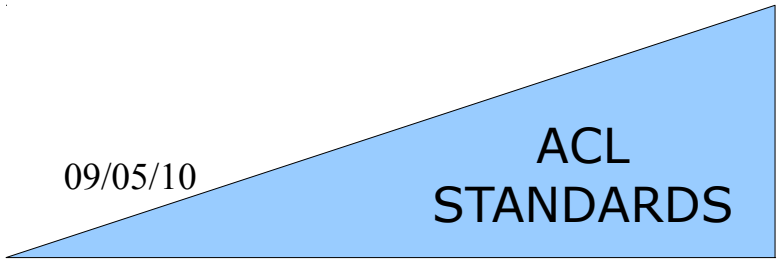

Regulating interaction in Open MAS

- Standard ACLs do not solve all the problems concerning interoperability in open MAS
- A key issue in designing open MAS is to regulate interactions between the heterogeneous and autonomous parties
- In open environments:
  - interacting agents are developed independently;
  - e-commerce: vendors might not desire to publicize the internal code of their agents
  - coalitions are formed dynamically;
  - the partners can be replaced or upgraded at any moment;
- We need means for guaranteeing that agents of the system are actually being able to interoperate by message exchange (interoperability at the communication level) and that interactions produce the desirable outcome for all agents;
- Interactions ought to allow for reliable expectations on the possible evolution of the system.

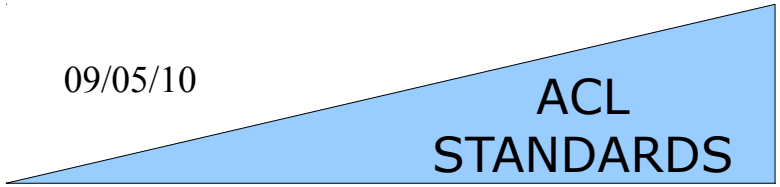

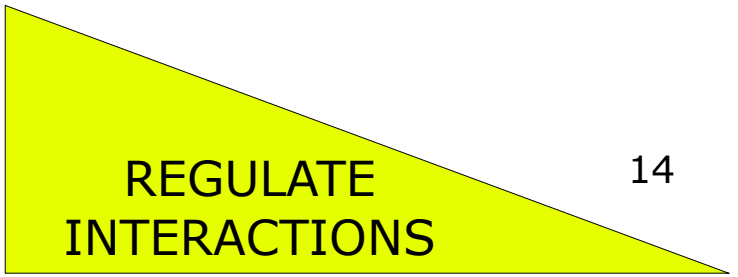

Regulating interaction in Open MAS

- A key issue in designing open MAS is to regulate interactions between the heterogeneous and autonomous parties
- Possible solutions:
  - Interaction protocols: interoperability is guaranteed by adopting an interaction protocol/choreography (pattern of behavior), whose roles are by design interoperable
  - Norms: to define interaction environments including a normative component, with suitable rules to regulate the behaviour of agents
    - institutions/organizations

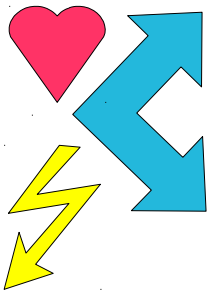

REGULATE INTERACTIONS

Interaction protocols as patterns of behavior

- Traditionally, heterogeneity is accommodated by specifying interaction protocols
  - Interaction protocols: shared specifications of behavioral patterns which allow a set of agents to engage cooperations with one another when playing their respective roles.

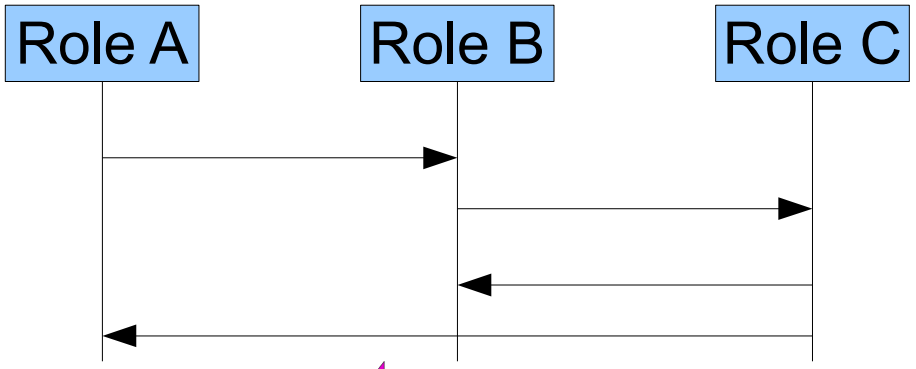

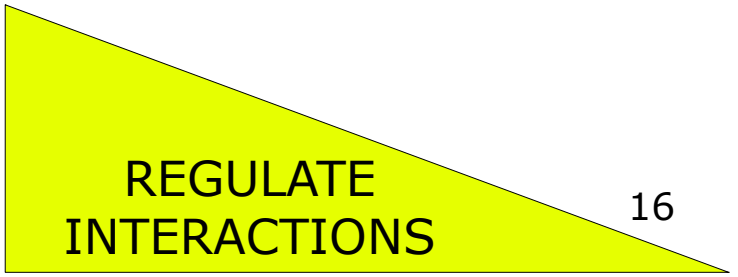

Protocols and verification

- Regulation and expectations: protocols should allow participants to broadly predict each other's behavior
- Besides simplifying the coordination problems, protocols, when a formal semantics is given, introduce the possibility of performing verification tasks
- Considering protocols as models of the desired interaction:
  - A-priori (design time) verification of many properties (interoperability, conformance,...); we guarantee them before any interaction takes place.
  - Run-time compliance testing: checking the behavior of the system at run time to determine if it behave as desired
- What kind of verification and testing may possible depends on information available: agent internals, observable behavior, procedural specification of protocols, normative specifications.

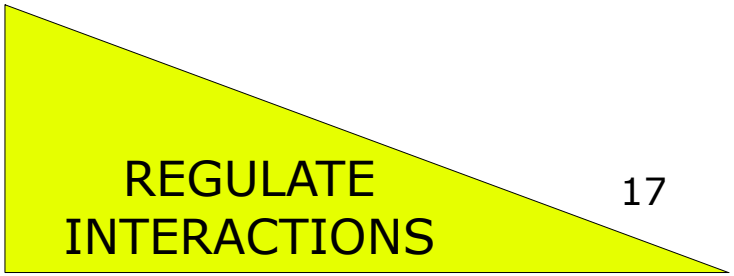

ACL: philosophical ancestors

- Most treatments of communication in (multi-)agent systems borrow their inspiration from speech act theory.
- Speech act theories are pragmatic theories of language, i.e., theories of language use: they attempt to account for how language is used by people every day to achieve their goals and intentions.
- The origin of speech act theories are usually traced to Austin's 1962 book, How to Do Things with Words.
- Main idea: to treat communications as actions
- Further contributions: Searle and Grice
  - "Speech Acts: An Essay in the Philosophy of Language" (1969)

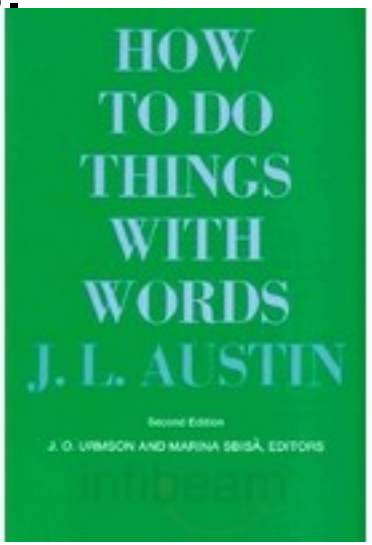

ACL: philosophical ancestors

- Searle and Grice: focus on the speaker intent
  - the illocution of a speech act is what the speaker believed or intended to be.
  - ACL mentalistic semantics focus on this aspect of speech act theory;
- Searle: recognizes the institutional nature of language
  - "The Construction of social reality" (1995)
  - Distinction among constitutive and regulative norms
  - ACL social semantics put in evidence this aspect of speech act theory

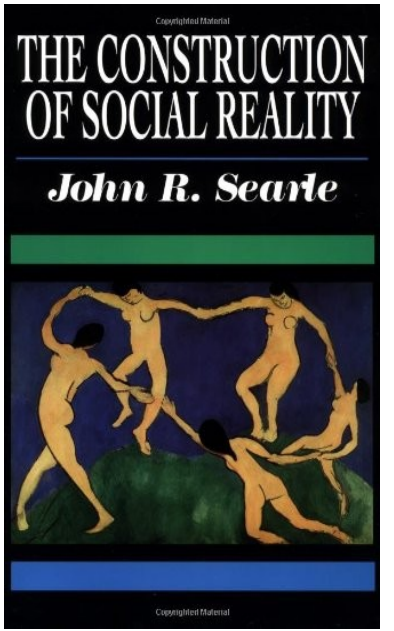

AI: Plan-based semantics for speech acts

- Plan-based semantics: one of the most popular approaches takes as key concept rationality: speech acts are interpreted as rational actions, i.e. they are planned and interpreted under the assumption that agents act rationally.
- The conceptual tools of this research line have been developed (among the others) by Cohen and Perrault [Cohen and Perrault, 1979], Allen [Allen, 1983] and Cohen and Levesque [Cohen and Levesque, 1990b]
- Cohen & Perrault (1979) defined semantics of speech acts using the precondition-delete-add list formalism of planning research.

Mentalistic semantics

- The semantics of the individual speech acts is specified in terms of their preconditions and effects on mental states
- Formal theories of speech acts based on the rationality assumption model individual communicative acts in terms of
  - feasability preconditions and effects, as agent's other actions, with the difference that communicative actions affect the agent's mental state -i.e. beliefs and intentions of participating agents - instead that the physical world.
- Cognitive primitives represented in the various theories differ according to the rationality model they refer to (BDI models)

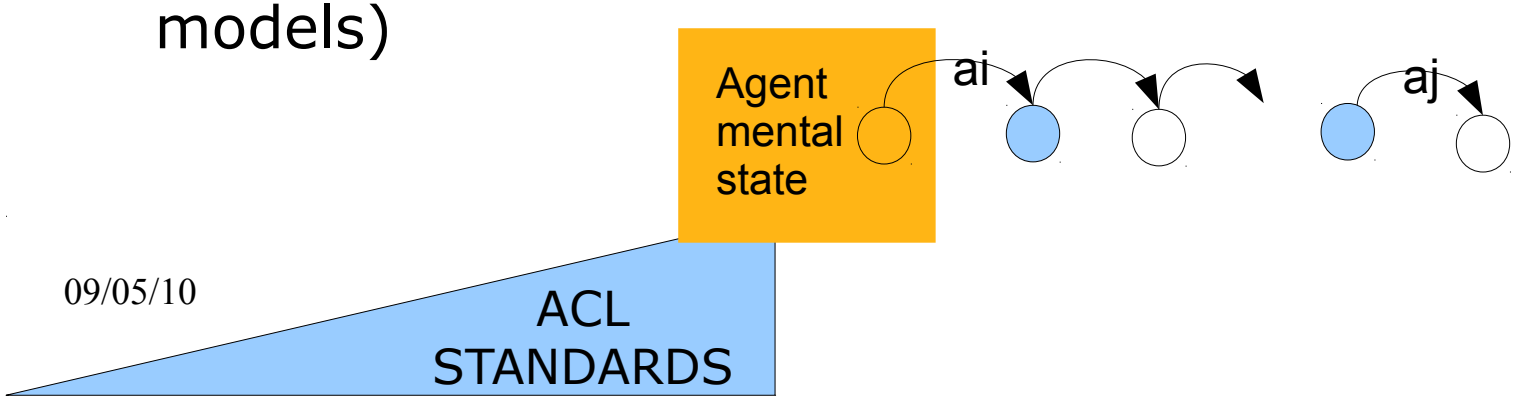

FIPA ACL

- This research line underlies the definition of action communication language as KQML [Labrou, 1997; Finin et al., 1995] and FIPA ACL [FIPA, 1997].
- Two parts in a FIPA ACL message:
  - Administrative part (sender, receiver, protocol, ontology, etc.)
  - Message part (communicative act, content)
- An example:

```
(request
    :sender (agent-identifier :name matteo)
    :receiver (set (agent-identifier :name claudio))
    :content "(shut the door)"
    :language Natural Language English
)
```

FIPA ACL

- The meaning of a communicative act is specified as sets of formulas that describe the act's *feasibility preconditions* and *rational effects*.
  - feasability pre-conditions – what must be true in order for the speech act to succeed.
  - “rational effect” – what the sender of the message hopes to bring about.
- “Inform” and “Request” are the two basic performatives in FIPA. Others are macro de nitions, de ned in terms of these.

FIPA-ACL and autonomy of agents

```
(request
      :sender (agent-identifier :name matteo)
      :receiver (set (agent-dentifier :name claudio))
      :content "(shut the door)"
      :language Natuaral Language English
)
```

- Note that a speaker cannot (generally) force a hearer to accept some desired mental state
- In heterogeneous distributed system of entities (agents), the sender has ***no control*** over the *state* and *behavior* of the receiver of its message (as in object oriented languages)
- Autonomy of the entities (agents)

Mentalistic semantics

- Mentalistic semantics of FIPA-ACL: enables reasoning on the activity of a single agent from its subjective point of view
- Can be the right approach in many situations:
  - Intelligent human-machine dialogue [Bretier and Sadek, 1997; Herzig and Longin, 1999]; target: developing techniques for representing and reasoning about other's mental state in order to recognize his(her) intentions in communication and, then to produce a suitable reply
  - Programming rational agents [DyLOG, CONGOLOG,...]: mentalistic semantics is useful to an agent designer who wants to decide when an agent will plan to send a message and also what way the agent will update its internal state upon receiving a message

Social semantics

- Mentalist semantics do not provide a normative specification for communication in the system
  - this is necessary to ensure that the multi-agent system does not become dysfunctional, for example to ensure that agents respond when spoken to and that they honour their commitments.
- Social commitment-based semantics: a new recent trend focusses on the so-called normative side of communication and interprets communicative acts as actions creating obligations, permissions and other kind of deontic states [Singh 1999, Castelfranchi 1995, Fornara & Colombetti, 2000].
- Commitments provide meaning to agent interaction. During the execution of a protocol, agents enter into, modify, and satisfy commitments [Singh, 2006]

Conversation policies and protocols

- Programming rational agents
- The mentalistic semantics of communicative acts is too complex to determine the possible answer to a message by just reasoning on mental states.
- An agent must implement tractable decision procedures that allow it to select and produce ACL messages that are appropriate to its intentions: agent conversation policies
- FIPA-ACL protocols [FIPA, 2002].

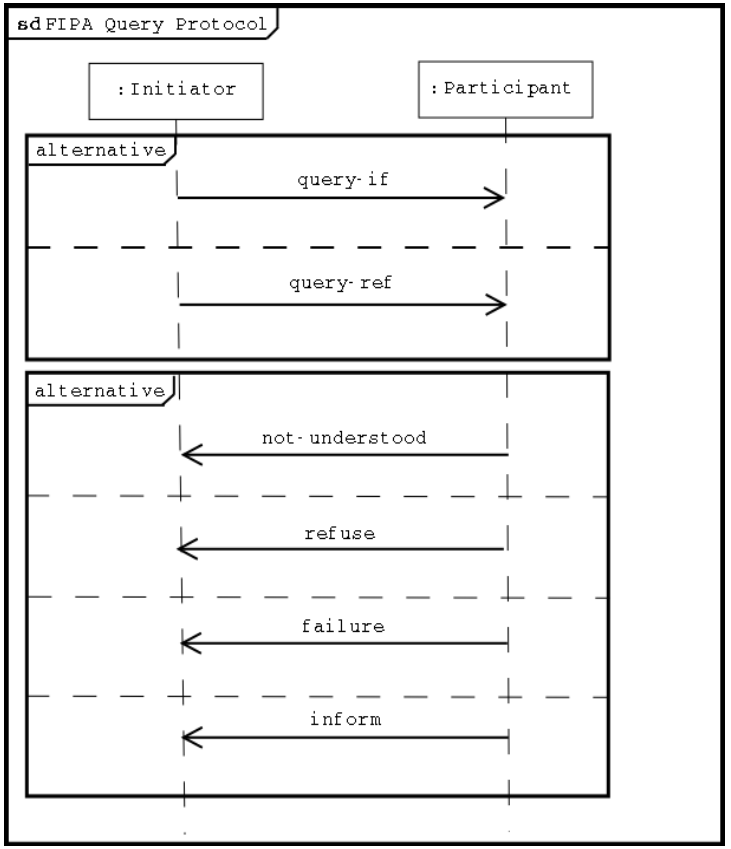

ACL STANDARDS

Regulating interaction in Open MAS

- A key issue in designing open MAS is to regulate interactions between the heterogeneous and autonomous parties
- Possible solutions:
  - Interaction protocols: interoperability is guaranteed by adopting an interaction protocol/choreography (pattern of behavior), whose roles are by design interoperable
  - Norms: to define interaction environments including a normative component, with suitable rules to regulate the behaviour of agents
    - institutions/organizations

REGULATE INTERACTIONS

Formal specification: different approaches

- Interaction protocols can be formally specified in different ways.
- Procedural approach:
  - Traditionally interaction protocol specifications have a procedural nature that captures the allowed interaction flows.
  - Finite State Automata, Petri Nets, Process Algebra, (A)UML
- Commitment-based approach: Singh and colleagues criticize the use of procedural specifications as being too rigid and propose the more flexible commitment-based approach to protocol specification.

Protocols and verification

- Protocols introduce the possibility of performing verification tasks
- Procedural approach and focus on:
  - A-priori (design time) verification of property of the interaction like interoperability and conformance

REGULATE INTERACTIONS

The summer school example

- In open enviroments agents are identified and composed on demand:
  - retrieval done component by component
  - components behaviour could be private, not accessible for inspection

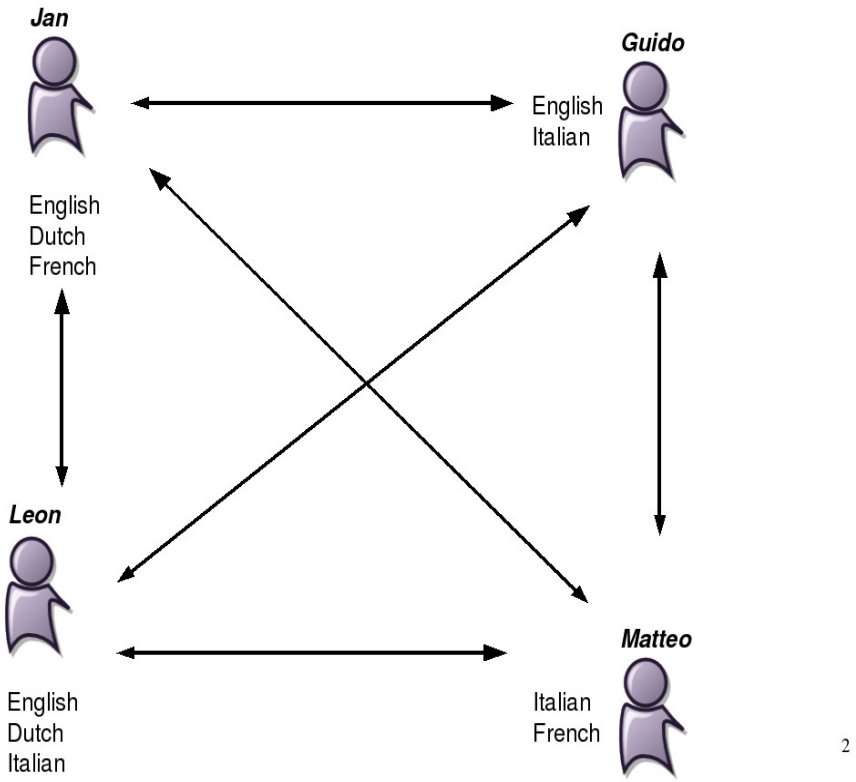

Interoperability

- Agents must coordinate their executions and must be able to interact with each other.
- Interoperability = stuck-free: a set of agents is interoperable when it is stuck-free: i.e. whatever point of the interaction maybe reached the system will not be blocked

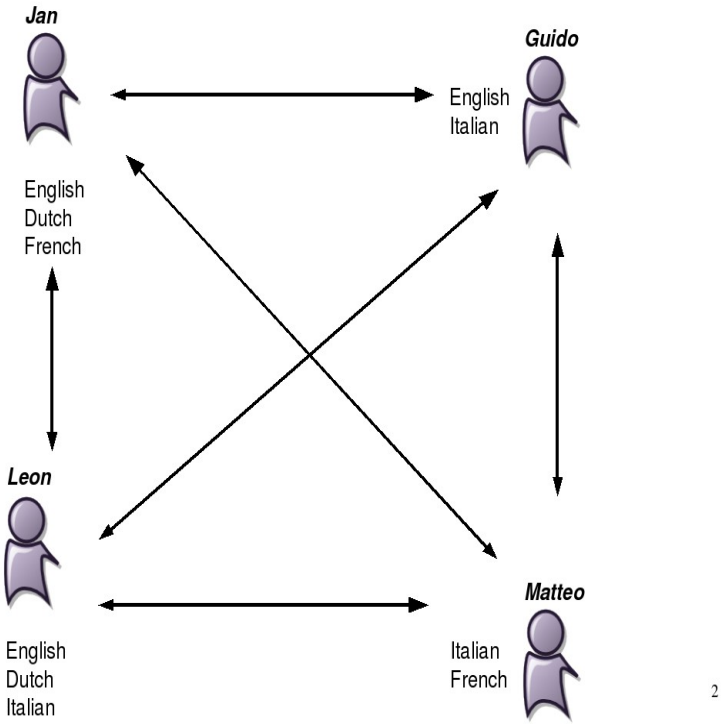

REGULATE INTERACTIONS

Interoperability

- Need for a "distributed verification" of interoperability
- It should be possible to perform each verification without knowledge of the actual partners
- The use of protocols / choreographies enables this distribution
- Interoperability according to the protocol

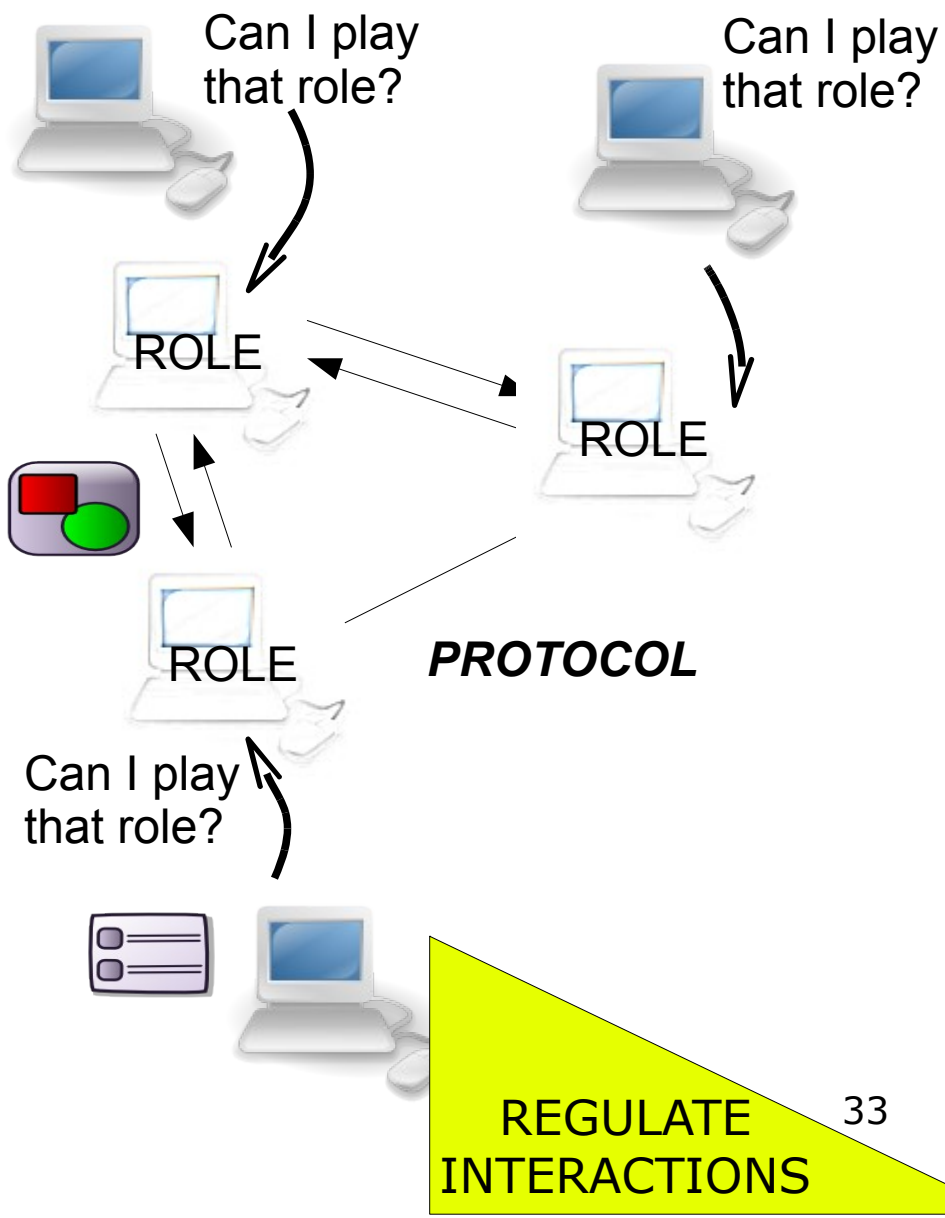

Distributed Verification

- **Step 1**:
  Adoption of a protocol / choreography that **already guarantees** the interoperability
- A choreography abstractly specifies an interaction in terms of its *roles*, the messages exchanged among the roles, and regulate the message exchanges.

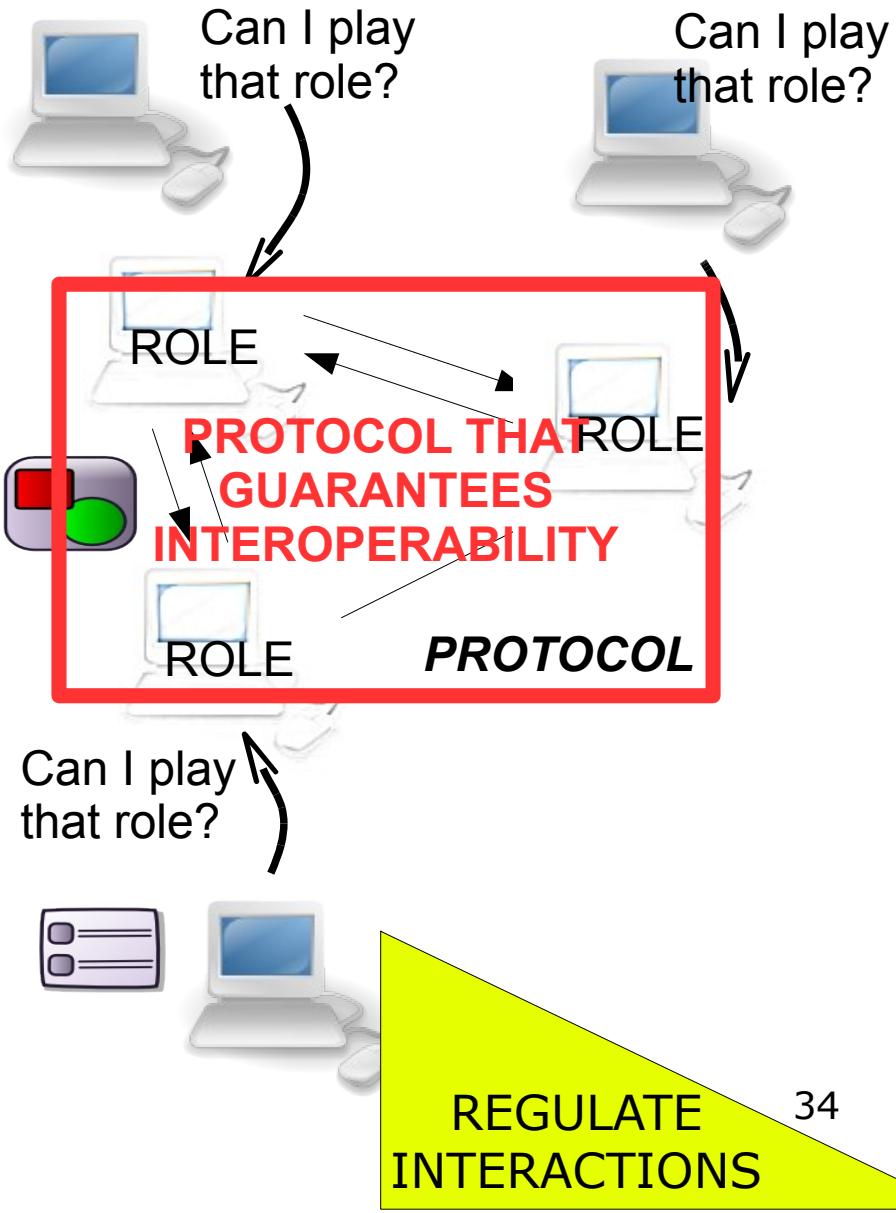

Distributed Verification

- Step 2:
  To perform a conformance test against the role specification: “a-priori” / static interoperability checking
- The agents commit to not execute actions that are not foreseen by the protocols
- *Substituting a role by a conformant player it is preserved the interoperability of the agent/coalition service composition.*

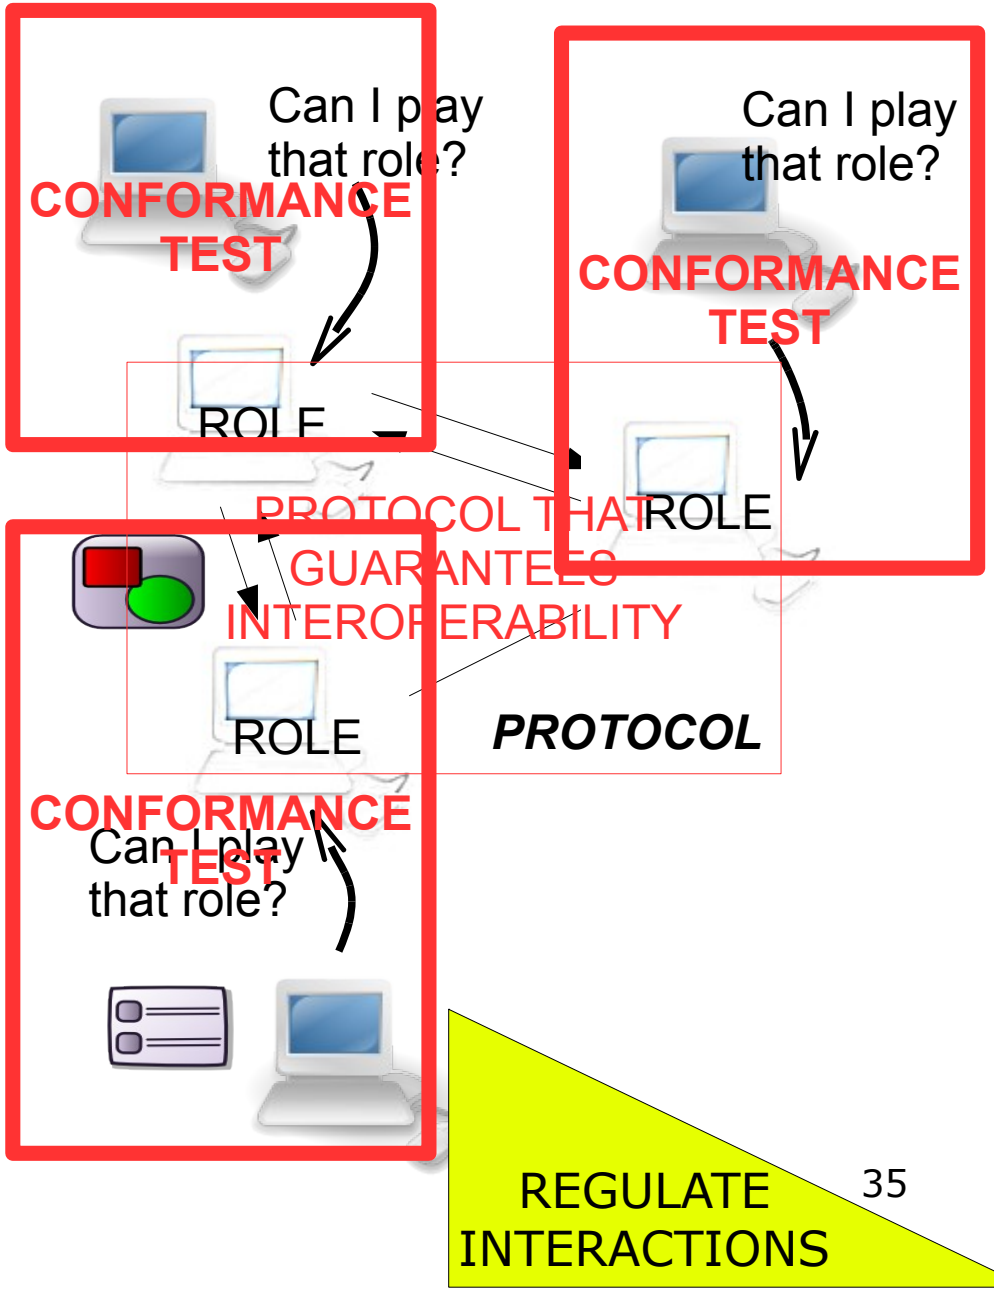

Distributed Verification

- Step 1:
  Adoption of a protocol / choreography that alread guarantees the interoperability

- Step 2:
  To perform a conformanc test againt the role specification: “a-priori” / static interoperability checking

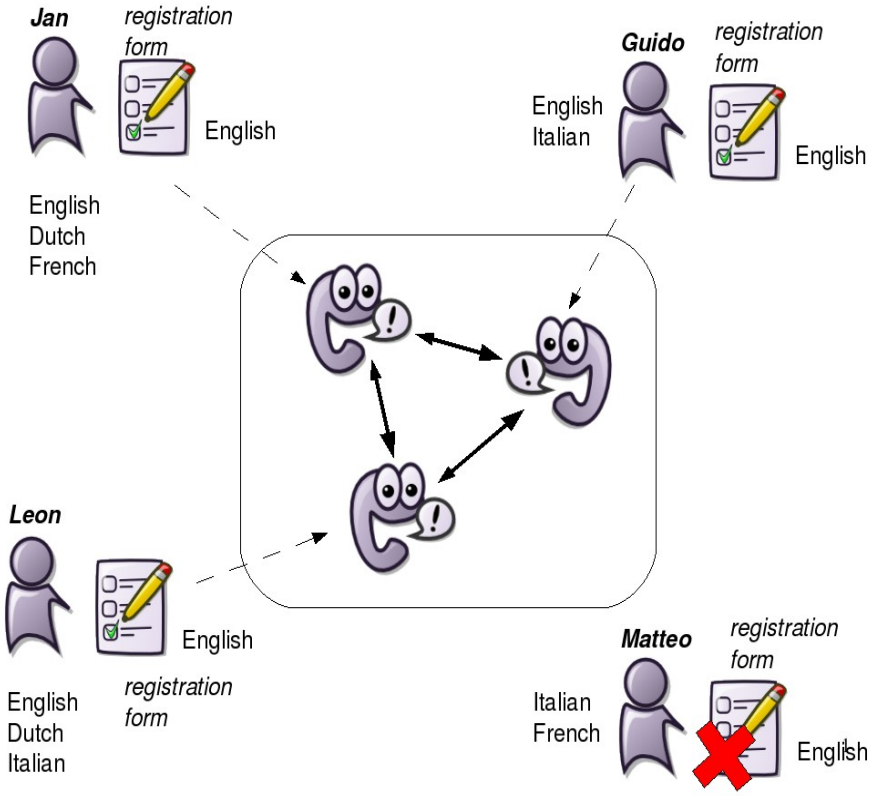

A-priori interoperability verification

- To perform the ***conformance test*** against the role specification
  - ***"a-priori" interoperability test***
  - ***static interoperability cheking***

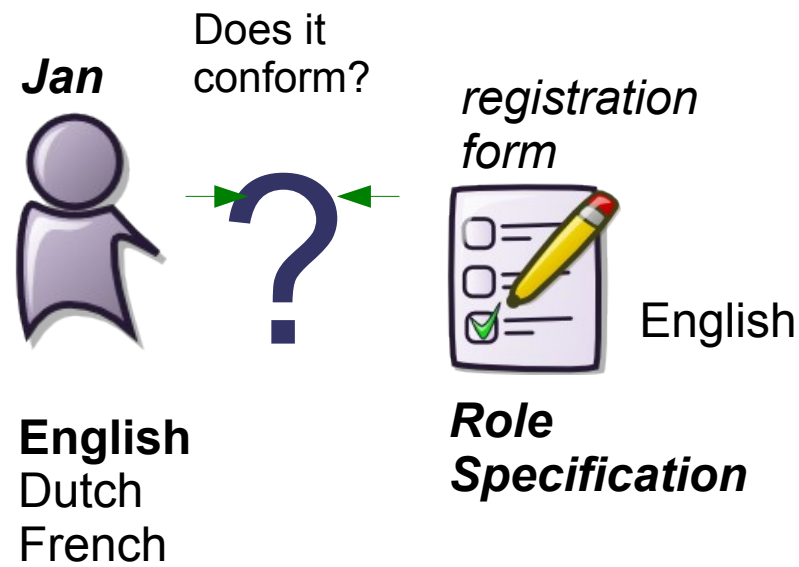

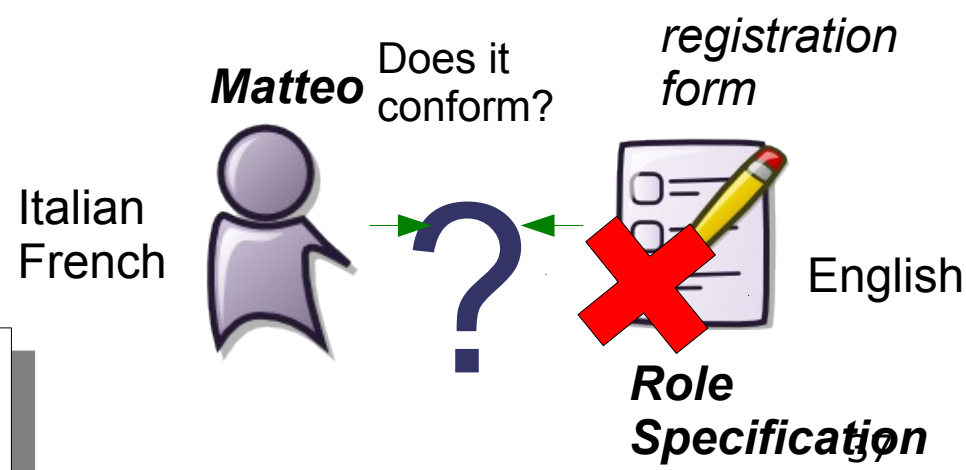

Verify-once run-always

Verification techniques

- Conformance should guarantee substitutability
  - reachability algorithms,
  - techniques for the verification of deadlock-freeness,
  - alternated simulation
  - (bi)simulation techniques

  have been proposed for verifying the conformance of agents to roles and for guaranteeing the substitutability of implementations (agents) to specifications (roles) in a way that preserves the interoperability of the involved parties.

Conformance should guarantee substitutability

- ***Substitutability***: *If "I conforms to R_i" and $P=R_1|...|R_i|...|R_n$ is any protocol such that $R_1|...|R_i|...|R_n$ is interoperable then $R_1|...|I|...|R_n$ is interoperable*

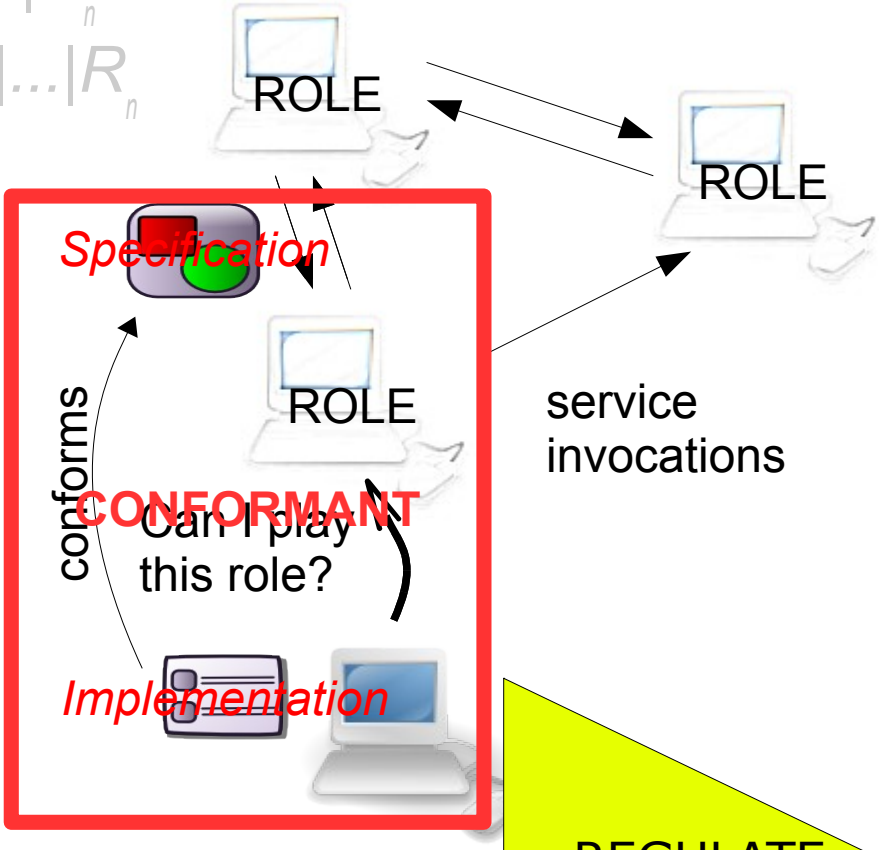

REGULATE INTERACTIONS

Conformance should guarantee substitutability

- ***Substitutability***: *If "I conforms to R_i" and $P=R_1|...|R_i|...|R_n$ is any protocol such that $R_1|...|R_i|...|R_n$ is interoperable then $R_1|...|I|...|R_n$ is interoperable*

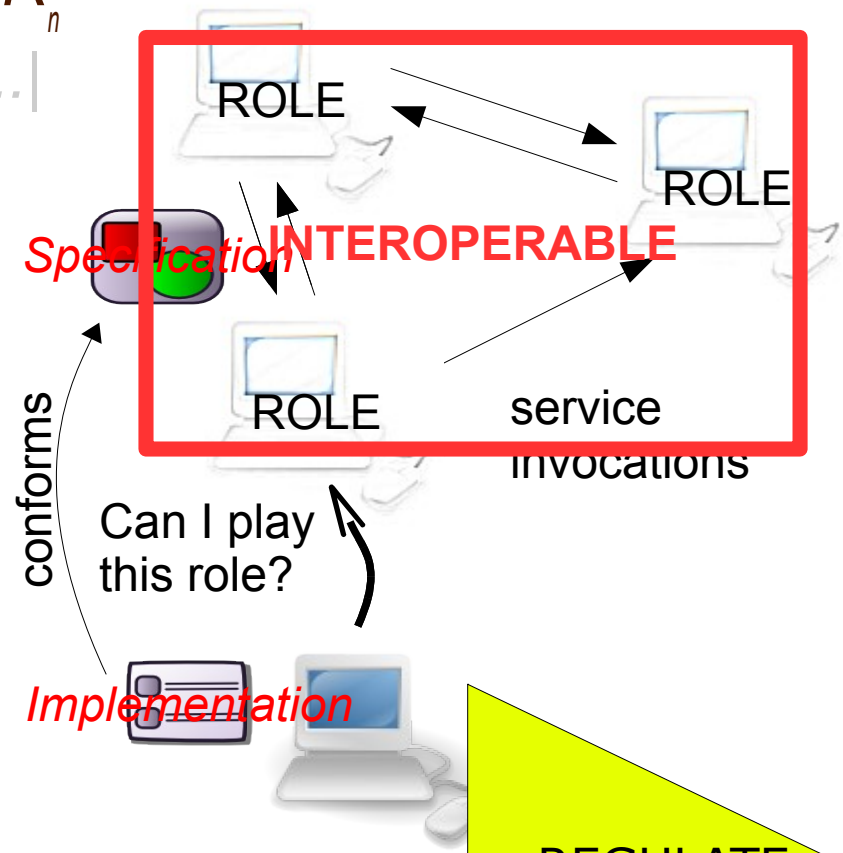

REGULATE INTERACTIONS

Conformance should guarantee substitutability

- ***Substitutability***: *If "I conforms to R_i" and $P=R_1|...|R_i|...|R_n$ is any protocol such that $R_1|...|R_i|...|R_n$ is interoperable then $R_1|...|I|...|R_n$ is interoperable*
- Conformance does not require any knowledge about the other involved parties
- *The only assumption is that all of them respect the agreed upon choreography*

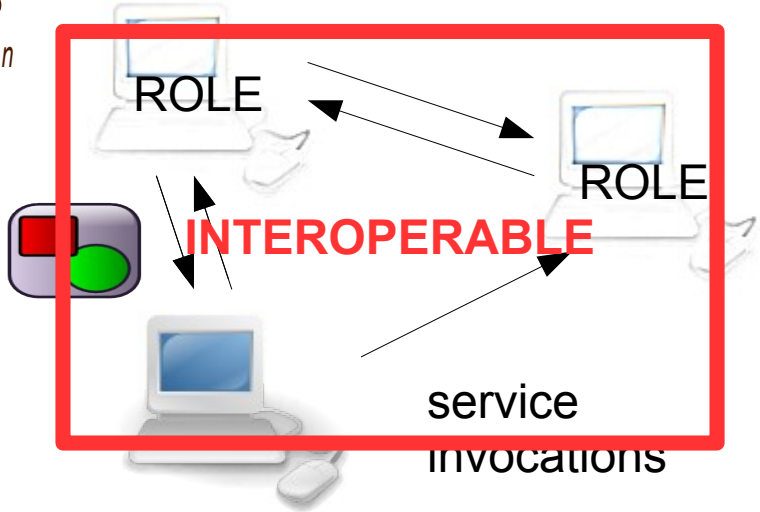

REGULATE INTERACTIONS

Conformance is a subtype relation

- An agent/service playing a role may be slightly different than the role specification
- A sort of subtype relation
- Software re-use

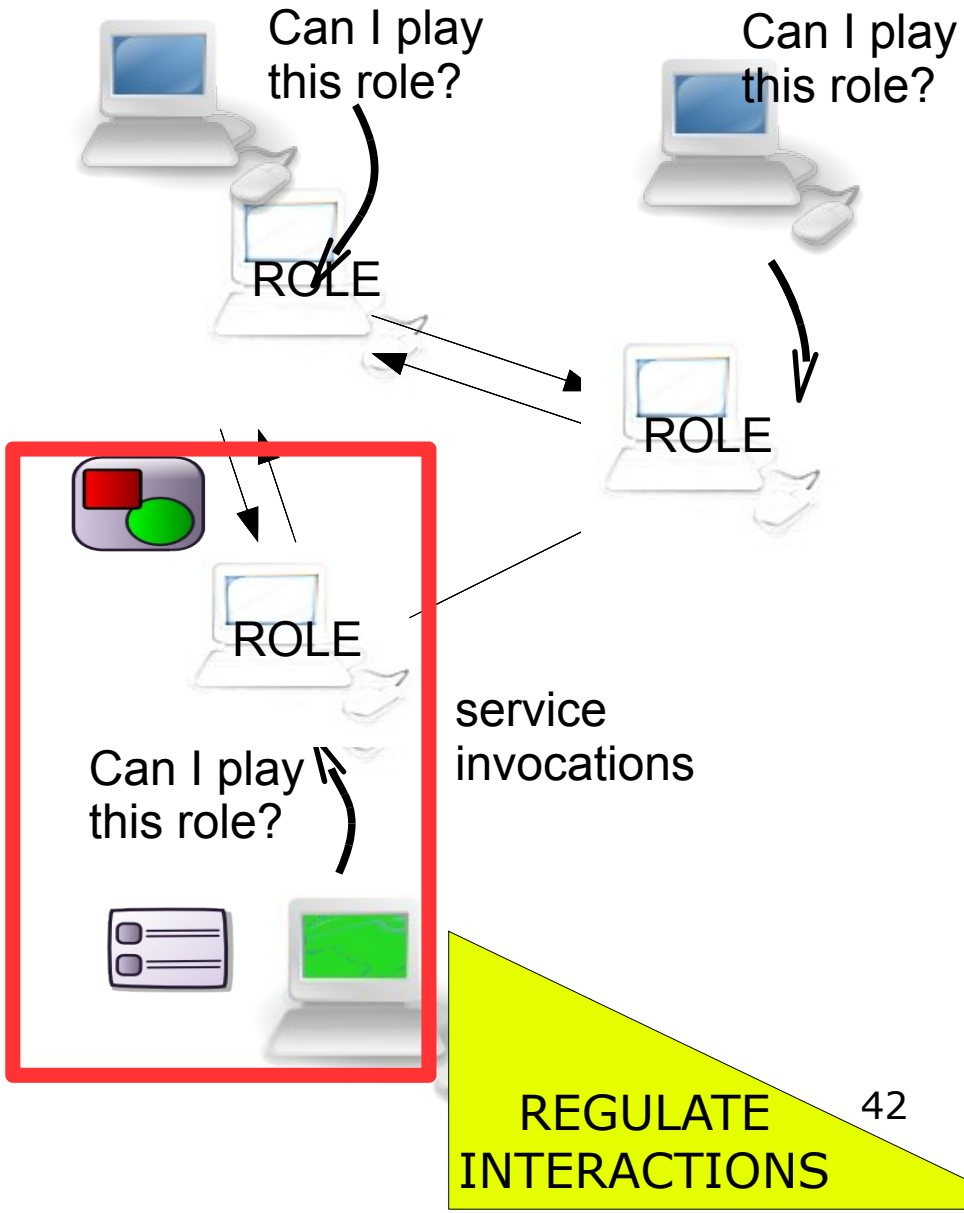

Conformance as a tool for the designer

- How can a behavior be modified ***without compromising*** interoperability? (and ***without re-checking*** the conformance)
- What is the ***degree of freedom of a designer***?
- Upgrades

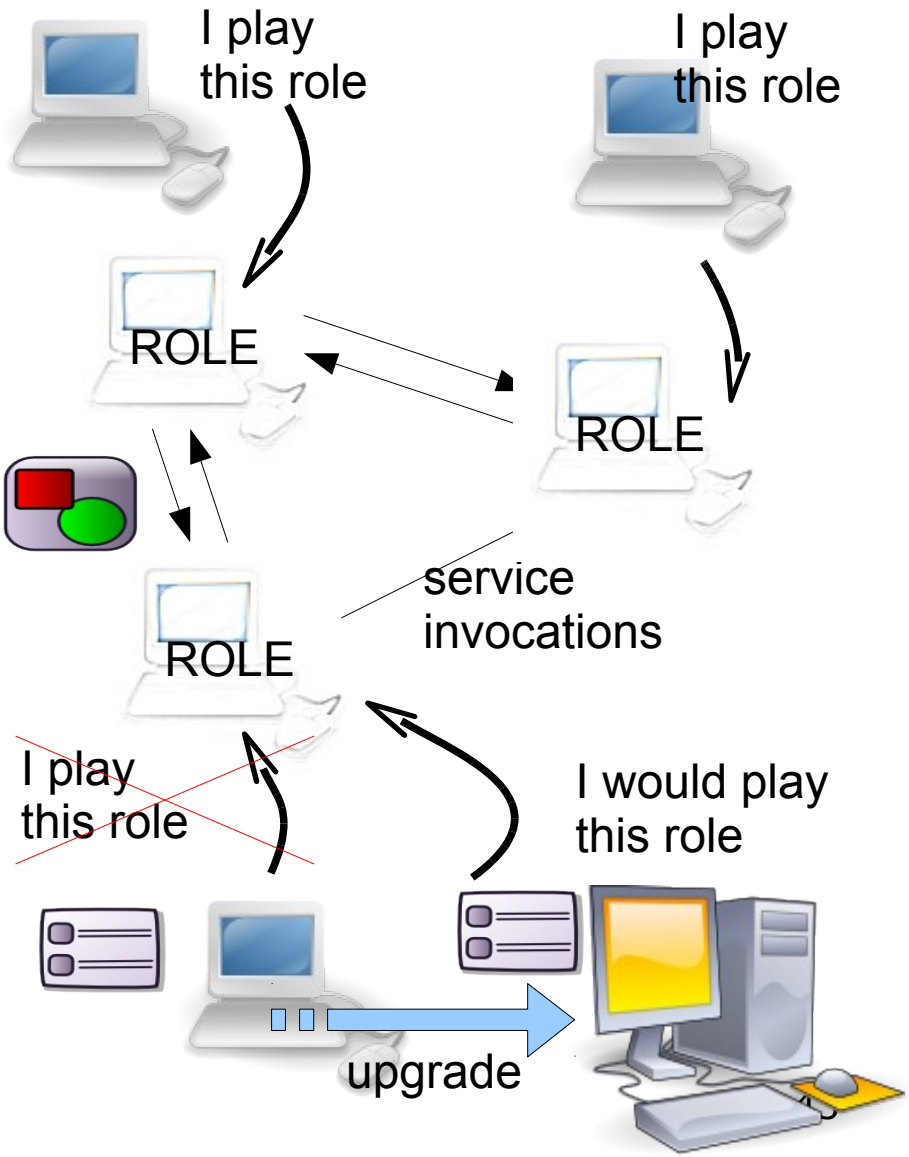

Conformance is a subtype relation

- On this line, some of the most recent approaches to the conformance problem [FHRR04,BSBM05,BBMP06] come from the area of service-oriented computing.
- The exchanged messages are analyzed and it is allowed to substitute a specification (role) with a service (agent) even though it does not exactly match the specification but produces narrower sets of emissions and tackles the same or a broader set of receptions, without compromising interoperability.
  - Bordeaux et al. [BSBM05] codify this intuition as the motto "less emissions, more receptions!".
- A new framework for conformance and interoperability has been proposed, where interactive parties in a system can also lead choices about which message to receive and follow choices of which messages to send, respecting the motto: "lead less, follow more!"[BBCDPS09].
  - The framework allows to tackle also situations where interaction is based on a producer-consumer and not only on the usual sender-receiver pattern. The framework relies on the possibility of verifying conformance and interoperability based on an explicit representation of decision points and can deal with the multi-party case.

REGULATE INTERACTIONS

Choreographies and protocols: public specifications

- Each peer knows:
  - its policy
  - the role specification it would like to play
  - Maybe, the rest of the protocol
- Regulation and a-priori conformance verification enables:
  - dynamic selection of the partners
  - privacy of the used policies

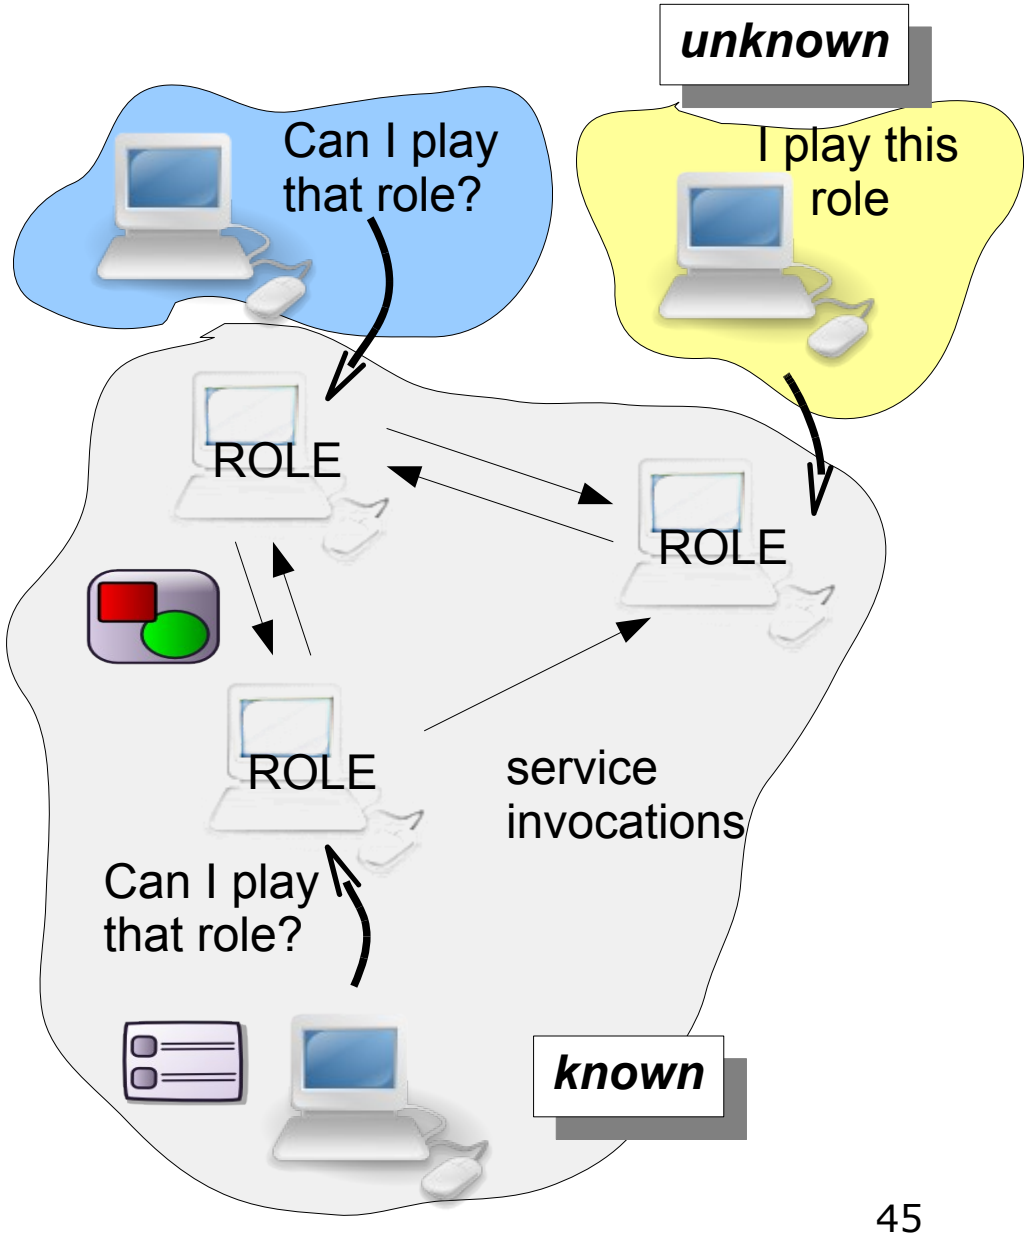

The summer school example

- Desire: to check interoperability without any knowledge of the services that will play other roles
- Solution: a-priori verification of interoperability in terms of conformance relation w.r.t a protocol

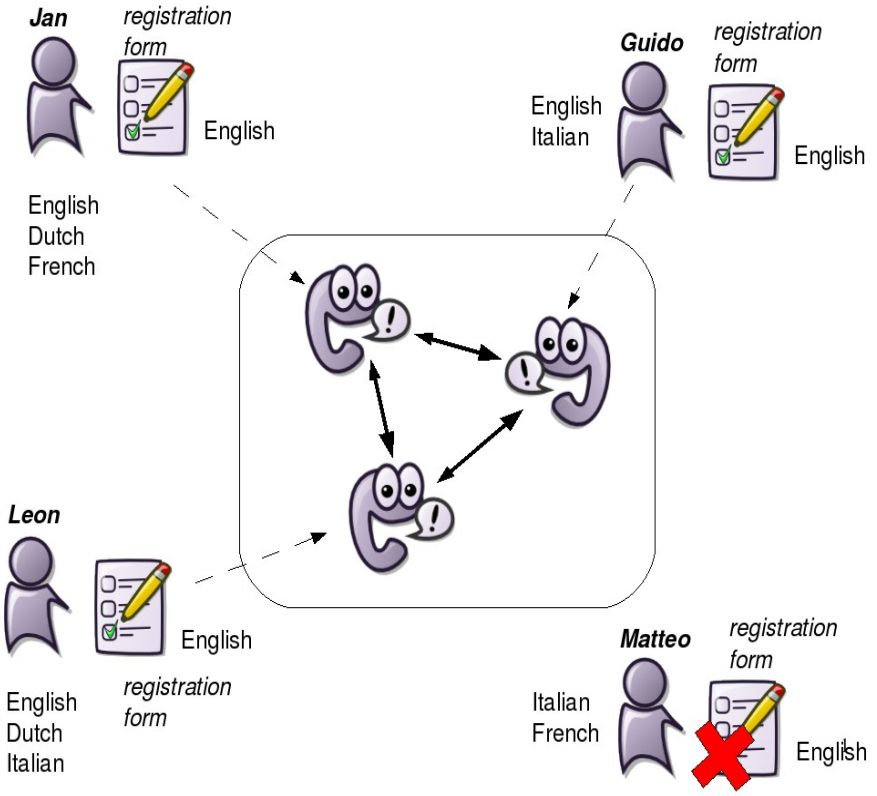

Final remarks

- Protocols and autonomy of agents
- ACL and protocols
- Protocol versus norms

- Part II: commitment protocols

Final remarks: protocols and autonomy

- An agent, willing to play a role in a protocol, has the means for understanding if that pattern of interaction meets its goals.
- By autonomously deciding whether entering the protocol it, however, commits to respect its rules along the whole interaction.
- The fact that all agents accept to respect the rules has the advantage of making the course of interaction predictable and, therefore, of giving guarantees to all of the participants.

Final remarks: protocols and autonomy

- For example: at the school all ***might*** speak another common language (French)

- ... however at the school all ***might not*** speak another common language...

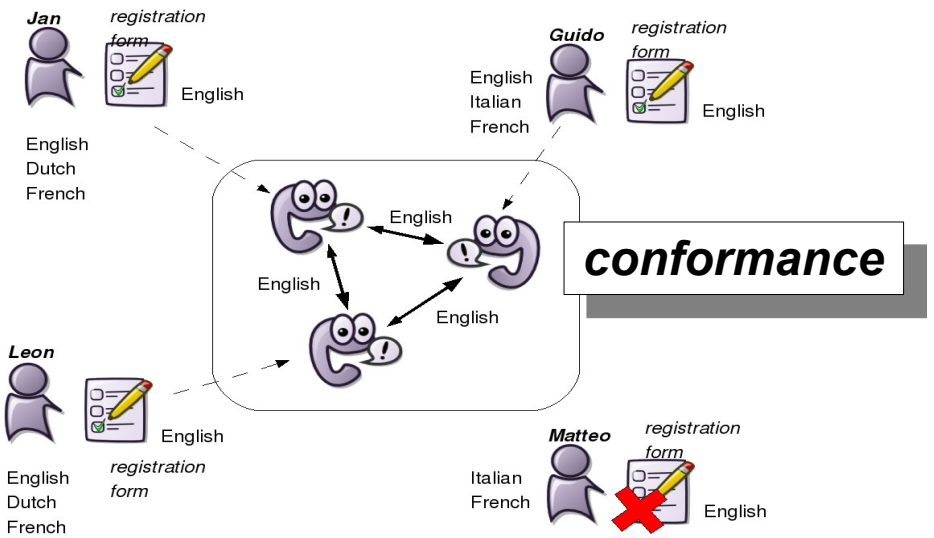

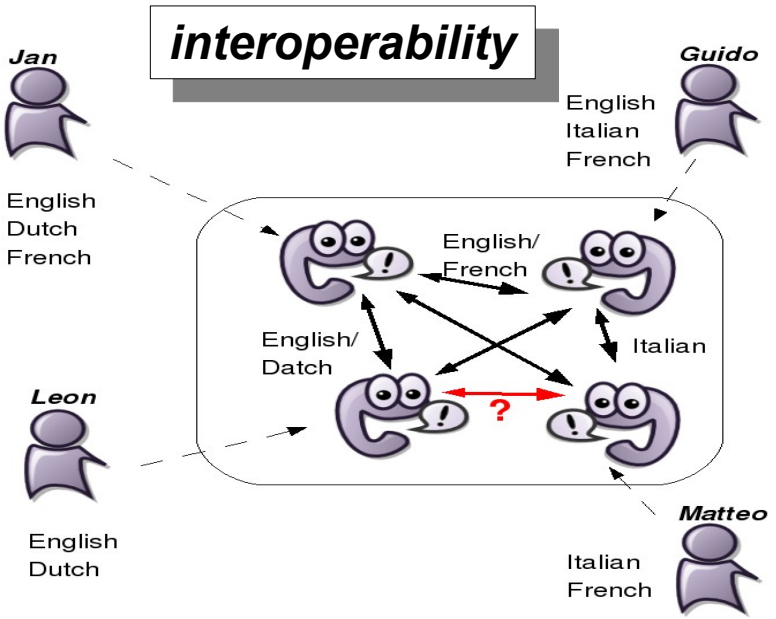

Final remarks: autonomy and protocol violation

- Notice that, even though an agent decided to play a role in a protocol that imposes some rules, it can always decide to break the rules if an opportunity, it can take advantage of, arises.
- Suppose that at a summer school the official language everybody should speak is English. Every student at the school has a badge reporting the name and the nationality of the person.
- If a French attendee meets a colleague whose badge says she is from France, the first student might decide to speak in French with the other even though the official language she should speak is English. There is a clear expected advantage in doing so: a better understanding
- However, the violation of the rules of the protocol introduces a certain amount of risk. For example, if the second student by mistake took the badge of her roommate, she might not understand the former one because the protocol is not attended.
- This would not have happened if the first agent did not violate the rules!

Final remarks: ACL and protocols

- I protocolli hanno come mattoncini base atti comunicativi
- Quale collegamento fra gli approcci in letteratura per la specifica e la verifica di protocolli di interazione e la semantica degli ACL?
- 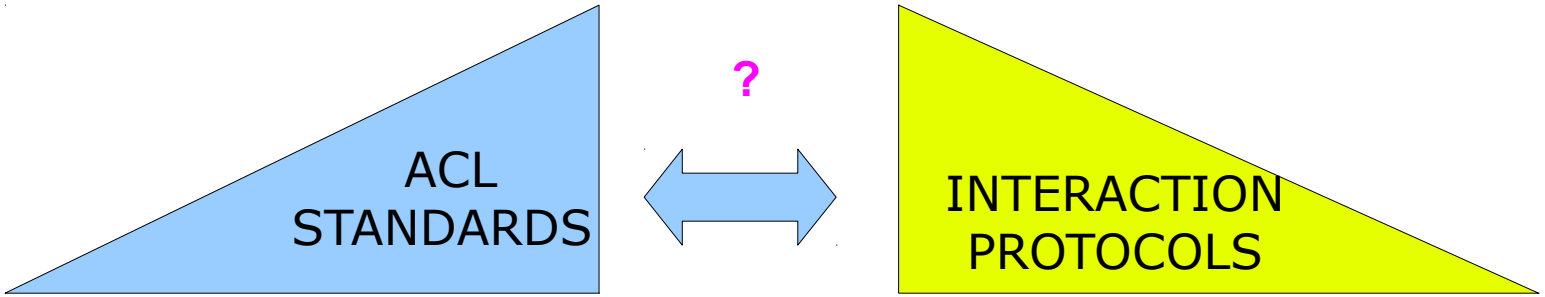

-
- Gli approcci procedurali alla specifica e verifica dei protocolli non tengono conto della semantica dei messaggi che li compongono
- Per Singh e' una prova delle limitazioni della semantica mentalistica degli ACL
  - Social semantics and commitment-based protocols in order to introduce an explicit notion of responsibility and commitment

Final remarks: interaction protocol versus norms

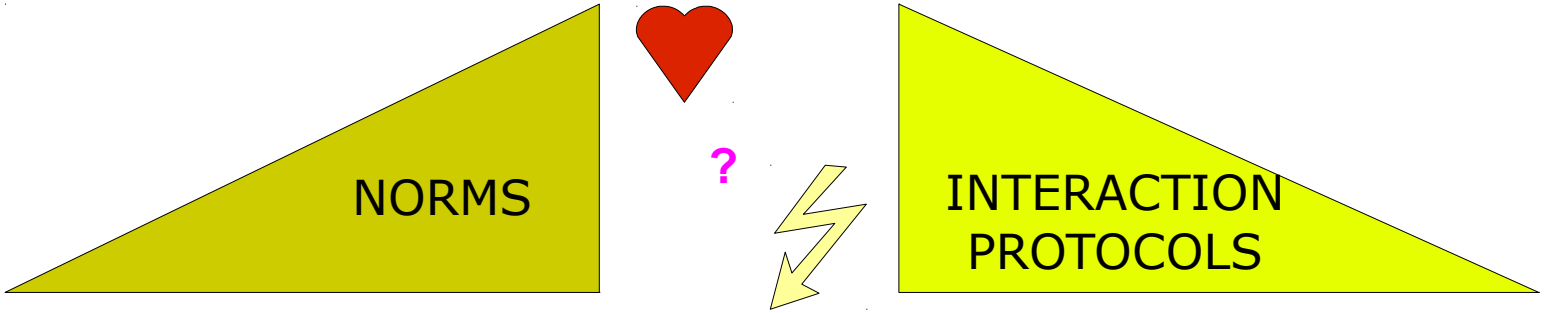

- Interaction protocols themselves can be conceived as sets of norms (about what an agent may do and when).
- Anche i protocolli specificati in maniera procedurale sono normativi per gli agenti che si impegnano a rispettarli: norme che regolano il comportamento degli agenti qualificandolo deonticamente in termini di permesso o vietato
- La regolamentazione e' data in modo molto rigido: e' permesso fare solo quello che e' previsto dal protocollo, quando e' previsto dal protocollo. Tutto il resto e' vietato -> violazione.
- Gli agenti si impegnano a eseguire solo le azioni previste dal protocollo in ogni stato della conversazione normata dal protocollo. In cambio ricevono garanzia di interoperabilita' con gli altri agenti conformi al protocollo, interoperabilita' finalizzata al compimento di un business o allo scambio di informazioni o allo svolgimento in cooperazione di un compito di loro interesse.
- In qualche modo possiamo dire sono garantiti che se seguono gli schemi di comportamento specificati nel protocollo sono sicuri di che le interazioni saranno quelle previste,che onoreranno le reciproche aspettative e che riusciranno quindi a cooperare profiquamente.

MA

- Le nozioni di commitment e responsibility non vengono specificate ad alto livello
- I vincoli sul comportamento dell'agente sono troppo forti

Thanks to

- Members of our research group in Torino and in particular

  Cristina Baroglio,
  Elisa Marengo and
  Alberto Martelli