Mini Scuola WOA 2010

Regulating Agent Interactions in Open Multi-Agent Systems

Part II: Commitment protocols

Matteo Baldoni and Viviana Patti

Dipartimento di Informatica — Università degli Studi di Torino

Rimini, September 5, 2010

Overview

1. What is a protocol
2. Are BDI enough to represent interactions in open MAS?
3. Castenfranchi and Singh's commitments
4. Are FSMs enough to represent interactions in open MAS?
5. Singh's commitment protocols
6. On the constitutive and regulative specification of a protocol
7. Verifying a protocol interaction
8. Need for a more powerful regulation
9. 2CL: constraints among commitments
10. Conclusions

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Protocols are public artifacts

- Protocols can be seen as *public artifacts* [Singh and Chopra, 2009], ruling the interaction of agents playing the various roles
- Protocols allow for the *verification* of many properties and "guarantee" them before any interaction takes place
- Patterns of interactions that allow agents to become MAS
- Protocols as *models of the desired interaction*
- Protocols allow to broadly *predict behavior*

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Mentalistic semantics for ACL and verification

Respecting a protocol

A protocol is respected *if and only if* only desired/specified interaction will occur during the execution

- Mentalistic semantics for ACL is well suited for "statically" reasoning about protocols and agent's policies
- However, from the point of view of an agent that partecipates to the interaction, how to be sure that other agents comply with the specified semantics?
- and, from the point of view of a third party that should monitor the interaction, how to be sure that agents comply with the specified semantics?
- Introspection of agents is not possible, agents' mental states are not usually accessible

Mentalistic semantics for ACL and verification

Respecting a protocol

A protocol is respected *if and only if* only desired/specified interaction will occur during the execution

- Mentalistic semantics for ACL is well suited for "statically" reasoning about protocols and agent's policies
- However, from the point of view of an agent that partecipates to the interaction, how to be sure that other agents comply with the specified semantics?
- and, from the point of view of a third party that should monitor the interaction, how to be sure that agents comply with the specified semantics?
- Introspection of agents is not possible, agents' mental states are not usually accessible

Mentalistic semantics for ACL and verification

Respecting a protocol

A protocol is respected *if and only if* only desired/specified interaction will occur during the execution

- Mentalistic semantics for ACL is well suited for "statically" reasoning about protocols and agent's policies
- However, from the point of view of an agent that partecipates to the interaction, how to be sure that other agents comply with the specified semantics?
- and, from the point of view of a third party that should monitor the interaction, how to be sure that agents comply with the specified semantics?
- Introspection of agents is not possible, agents' mental states are not usually accessible

Mentalistic semantics for ACL and verification

Respecting a protocol

A protocol is respected *if and only if* only desired/specified interaction will occur during the execution

- Mentalistic semantics for ACL is well suited for "statically" reasoning about protocols and agent's policies
- However, from the point of view of an agent that partecipates to the interaction, how to be sure that other agents comply with the specified semantics?
- and, from the point of view of a third party that should monitor the interaction, how to be sure that agents comply with the specified semantics?
- Introspection of agents is not possible, agents' mental states are not usually accessible

Mentalistic semantics for ACL and verification

Verifying a protocol

Need for an "*observable*" semantics of ACL

- A semantics that is *independent* from the mental states of the participants
- A semantics that is verifiable both from a participant and a third party that should monitor the interaction
- A semantics that is suitable for open MAS
- A semantics that makes the *violation* of the specification by the participant to the interaction observable, without taking into account the nature of agents or performing an introspection of the agent's mental states

Mentalistic semantics for ACL and verification

Verifying a protocol

Need for an "*observable*" semantics of ACL

- A semantics that is *independent* from the mental states of the participants
- A semantics that is verifiable both from a participant and a third party that should monitor the interaction
- A semantics that is suitable for open MAS
- A semantics that makes the *violation* of the specification by the participant to the interaction observable, without taking into account the nature of agents or performing an introspection of the agent's mental states

Mentalistic semantics for ACL and verification

Verifying a protocol

Need for an "*observable*" semantics of ACL

- A semantics that is *independent* from the mental states of the participants
- A semantics that is verifiable both from a participant and a third party that should monitor the interaction
- A semantics that is suitable for open MAS
- A semantics that makes the *violation* of the specification by the participant to the interaction observable, without taking into account the nature of agents or performing an introspection of the agent's mental states

Mentalistic semantics for ACL and verification

Verifying a protocol

Need for an "*observable*" semantics of ACL

- A semantics that is *independent* from the mental states of the participants
- A semantics that is verifiable both from a participant and a third party that should monitor the interaction
- A semantics that is suitable for open MAS
- A semantics that makes the *violation* of the specification by the participant to the interaction observable, without taking into account the nature of agents or performing an introspection of the agent's mental states

Mentalistic semantics for ACL and verification

Verifying a protocol

Need for an "*observable*" semantics of ACL

- A semantics that is *independent* from the mental states of the participants
- A semantics that is verifiable both from a participant and a third party that should monitor the interaction
- A semantics that is suitable for open MAS
- A semantics that makes the *violation* of the specification by the participant to the interaction observable, without taking into account the nature of agents or performing an introspection of the agent's mental states

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Commitment, intention and interaction

- Choen and Levesque's view: "*Intention is choice with commitment*" [Cohen and Levesque, 1990], intentions are defined as persistent goals, intention is a basis for a commitment
- Cohen and levesque's extension: Joint intention is a basis for a social commitment [Cohen et al., 1997]
- Castelfranchi's view: He agrees with Cohen and Levesque's individual commitment but he *strongly disagrees* on social commitment [Castelfranchi, 1997, Castelfranchi, 1998]

Castelfranchi

in defining social commitments emphasizes their *distributed* and *normative* nature

Commitment, intention and interaction

- Choen and Levesque's view: "*Intention is choice with commitment*" [Cohen and Levesque, 1990], intentions are defined as persistent goals, intention is a basis for a commitment
- Cohen and levesque's extension: Joint intention is a basis for a social commitment [Cohen et al., 1997]
- Castelfranchi's view: He agrees with Cohen and Levesque's individual commitment but he *strongly disagrees* on social commitment [Castelfranchi, 1997, Castelfranchi, 1998]

Castelfranchi

in defining social commitments emphasizes their *distributed* and *normative* nature

Commitment, intention and interaction

- Choen and Levesque's view: "*Intention is choice with commitment*" [Cohen and Levesque, 1990], intentions are defined as persistent goals, intention is a basis for a commitment
- Cohen and levesque's extension: Joint intention is a basis for a social commitment [Cohen et al., 1997]
- Castelfranchi's view: He agrees with Cohen and Levesque's individual commitment but he *strongly disagrees* on social commitment [Castelfranchi, 1997, Castelfranchi, 1998]

Castelfranchi

in defining social commitments emphasizes their *distributed* and *normative* nature

Castelfranchi's view (see also [Norman et al., 2004])

- Social commitment is normative in that it engenders *rights*
- As a difference with Cohen and Levesques's view, the consequences of withdrawing form a social commitment are not simply that the other party(ies) must be informed, but in doing so ...

but in doing so ...

... one *violates obligations*, *frustrates expectations and rights* that the agent created

Castelfranchi's view (see also [Norman et al., 2004])

- Social commitment is normative in that it engenders *rights*
- As a difference with Cohen and Levesques's view, the consequences of withdrawing form a social commitment are not simply that the other party(ies) must be informed, but in doing so ...

but in doing so ...

... one *violates obligations*, *frustrates expectations and rights* that the agent created

Singh's view [Singh, 1999]

- Commitments can even be used to handle the normative concepts (like pledges, claims, rights, privileges, powers) in a unified framework
- The behavior of agents is affected by the commitments because the underlieing assuption is that agents respect the commitments they have taken

Therefore

Since protocols are contracts and all their normative aspects can be represented by means of commitments and of operations on them, we can use commitments to respresents protocols

Singh's view [Singh, 1999]

- Commitments can even be used to handle the normative concepts (like pledges, claims, rights, privileges, powers) in a unified framework
- The behavior of agents is affected by the commitments because the underlieing assuption is that agents respect the commitments they have taken

Therefore

Since protocols are contracts and all their normative aspects can be represented by means of commitments and of operations on them, we can use commitments to respresents protocols

Commitments

$$C(debtor, creditor, antecedent, consequence)$$

debtor is socially bound to *creditor* to bring about the *consequent* condition if the *antecedent* condition holds

Commitments: Examples [Chopra, 2009]

$$C(debtor, creditor, antecedent, consequence)$$

Commitment

$C(Bookie, Alice, \$12, BeatingtheOdds)$ means that Bookie commits to Alice that if she pays \$12, then Bookie will sends her the book Beating the Odds

Detach

If Alice makes the payment, that is, if \$12 holds, then $C(Bookie, Alice, \$12, BeatingtheOdds)$ is detached. In other words, $C(Bookie, Alice, \$12, BeatingtheOdds) \wedge \$12 \rightarrow C(Bookie, Alice, \top, BeatingtheOdds)$

Discharge

If Bookie sends the book, that is, if BeatingtheOdds holds, then both $C(Bookie, Alice, \$12, BeatingtheOdds)$ and $C(Bookie, Alice, \top, BeatingtheOdds)$ are discharged. In other words, $BeatingtheOdds \rightarrow \neg C(Bookie, Alice, \$12, BeatingtheOdds) \wedge \neg C(Bookie, Alice, \top, BeatingtheOdds)$

Commitments: Examples [Chopra, 2009]

$C(debtor, creditor, antecedent, consequence)$

Commitment

$C(Bookie, Alice, \$12, BeatingtheOdds)$ means that Bookie commits to Alice that if she pays \$12, then Bookie will sends her the book Beating the Odds

Detach

If Alice makes the payment, that is, if \$12 holds, then $C(Bookie, Alice, \$12, BeatingtheOdds)$ is detached. In other words, $C(Bookie, Alice, \$12, BeatingtheOdds) \wedge \$12 \rightarrow C(Bookie, Alice, \top, BeatingtheOdds)$

Discharge

If Bookie sends the book, that is, if BeatingtheOdds holds, then both $C(Bookie, Alice, \$12, BeatingtheOdds)$ and $C(Bookie, Alice, \top, BeatingtheOdds)$ are discharged. In other words, $BeatingtheOdds \rightarrow \neg C(Bookie, Alice, \$12, BeatingtheOdds) \wedge \neg C(Bookie, Alice, \top, BeatingtheOdds)$

Commitments: Examples [Chopra, 2009]

$C(debtor, creditor, antecedent, consequence)$

Commitment

$C(Bookie, Alice, \$12, BeatingtheOdds)$ means that Bookie commits to Alice that if she pays \$12, then Bookie will sends her the book Beating the Odds

Detach

If Alice makes the payment, that is, if \$12 holds, then $C(Bookie, Alice, \$12, BeatingtheOdds)$ is detached. In other words, $C(Bookie, Alice, \$12, BeatingtheOdds) \wedge \$12 \rightarrow C(Bookie, Alice, \top, BeatingtheOdds)$

Discharge

If Bookie sends the book, that is, if BeatingtheOdds holds, then both $C(Bookie, Alice, \$12, BeatingtheOdds)$ and $C(Bookie, Alice, \top, BeatingtheOdds)$ are discharged. In other words, $BeatingtheOdds \rightarrow \neg C(Bookie, Alice, \$12, BeatingtheOdds) \wedge \neg C(Bookie, Alice, \top, BeatingtheOdds)$

Commitments' postulates and operations

Postulates [Singh, 2008]

1. *discharge*: $u \rightarrow \neg C(r, u)$
2. *detach*: $C(r \wedge s, u) \wedge r \rightarrow C(s, u)$
3. *augment*: From $C(r, u) \wedge s \vdash r$ infer $C(s, u)$
4. *l-disjoin*: $C(r, u) \wedge C(s, u) \rightarrow C(r \vee s, u)$
5. *r-conjoin*: $C(r, u) \wedge C(r, v) \rightarrow C(r, u \wedge v)$
6. *consistency*: $\neg C(r, \perp)$
7. *nonvacuity*: From $r \vdash u$, infer $\neg C(r, u)$
8. *weaken*: $C(r, u \wedge v) \wedge \neg u \rightarrow C(r, u)$

Operations [Singh, 1999]

1. $\textsc{create}(x, y, r, u)$ is performed by x, and it causes $C(x, y, r, u)$ to hold
2. $\textsc{cancel}(x, y, r, u)$ is performed by x, and it causes $C(x, y, r, u)$ to not hold
3. $\textsc{release}(x, y, r, u)$ is performed by y, and it causes $C(x, y, r, u)$ to not hold
4. $\textsc{delegate}(x, y, z, r, u)$ is performed by x, and it causes $C(z, y, r, u)$ to hold
5. $\textsc{assign}(x, y, z, r, u)$ is performed by y, and it causes $C(x, z, r, u)$ to hold

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

The example of NetBill protocol

The NetBill protocol has been developed for buying and selling goods on the Internet

- It can be executed in many different ways, eg. the merchant may send a quote before the customer asks for it, to advertise his goods, **or** the merchant wants to send the goods without a prior acceptance of the customer
- Due to this flexibility, the standard approaches to protocol representation are inadequate

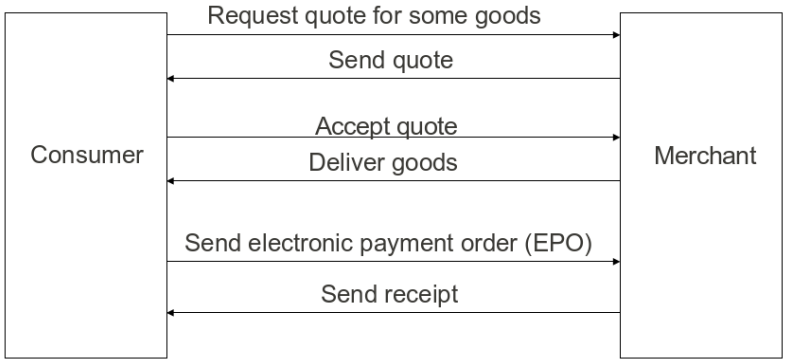

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Commitment-based protocols

act_1 **means** *this*

$\vdots$

act_n **means** *that*

- A commitment-based protocol is a *set of actions* whose meaning in terms of effects on the social state is agreed upon by all the interacting agents. In particular, the effects can be *operations* on commitments (create, delete, release, delegate, assign, discharge)
- The idea behind is to capture the *counts as relationships* that describe the institutional meanings of actions/messages

[Singh, 1999, Yolum and Singh, 2001, Chopra, 2009]

Commitment-based protocols: an example

- A protocol between a *merchant* and a *customer* (taken from [Chopra, 2009])
- It simply states what the meanings of the messages/actions are in terms of commitments between the merchant and customer.

$Offer(mer, cus, price, item)$ **means** $\textsc{create}(mer, cus, price, item)$
$Accept(cus, mer, price, item)$ **means** $\textsc{create}(cus, mer, item, price)$
$Reject(cus, mer, price, item)$ **means** $\textsc{release}(mer, cus, price, item)$
$Deliver(mer, cus, item)$ **means** $\textsc{inform}(mer, cus, item)$
$Pay(cus, mer, price)$ **means** $\textsc{inform}(cus, mer, price)$

Commitment-based protocols: some characteristics

The *only constraint* that commitment-based protocols include, to specify that an interaction is successful, is that *all commitments are discharged*

- A *shared* meaning of the *shared* actions
- *Observational semantics* rather than mentalistic semantics: Commitments have a normative nature
- They *do not over-constrain* the behaviour of the agent by imposing an ordering on the execution of the shared actions (*vs* algorithmic or procedural approaches): Agents are free to choose their actions
- Agents are liable for the violation of the commitments they have taken
- Agent can always decide to break a commitment if an *opportunity*, it can take advantage of, arises (of course, at its own risk!)

Commitment-based protocols: some characteristics

The *only constraint* that commitment-based protocols include, to specify that an interaction is successful, is that *all commitments are discharged*

- A *shared* meaning of the *shared* actions
- *Observational semantics* rather than mentalistic semantics: Commitments have a normative nature
- They *do not over-constrain* the behaviour of the agent by imposing an ordering on the execution of the shared actions (*vs* algorithmic or procedural approaches): Agents are free to choose their actions
- Agents are liable for the violation of the commitments they have taken
- Agent can always decide to break a commitment if an *opportunity*, it can take advantage of, arises (of course, at its own risk!)

Commitment-based protocols: some characteristics

The *only constraint* that commitment-based protocols include, to specify that an interaction is successful, is that *all commitments are discharged*

- A *shared* meaning of the *shared* actions
- *Observational semantics* rather than mentalistic semantics: Commitments have a normative nature
- They *do not over-constrain* the behaviour of the agent by imposing an ordering on the execution of the shared actions (*vs* algorithmic or procedural approaches): Agents are free to choose their actions
- Agents are liable for the violation of the commitments they have taken
- Agent can always decide to break a commitment if an *opportunity*, it can take advantage of, arises (of course, at its own risk!)

Commitment-based protocols: some characteristics

The *only constraint* that commitment-based protocols include, to specify that an interaction is successful, is that *all commitments are discharged*

- A *shared* meaning of the *shared* actions
- *Observational semantics* rather than mentalistic semantics: Commitments have a normative nature
- They *do not over-constrain* the behaviour of the agent by imposing an ordering on the execution of the shared actions (*vs* algorithmic or procedural approaches): Agents are free to choose their actions
- Agents are liable for the violation of the commitments they have taken
- Agent can always decide to break a commitment if an *opportunity*, it can take advantage of, arises (of course, at its own risk!)

Commitment-based protocols: some characteristics

The *only constraint* that commitment-based protocols include, to specify that an interaction is successful, is that *all commitments are discharged*

- A *shared* meaning of the *shared* actions
- *Observational semantics* rather than mentalistic semantics: Commitments have a normative nature
- They *do not over-constrain* the behaviour of the agent by imposing an ordering on the execution of the shared actions (*vs* algorithmic or procedural approaches): Agents are free to choose their actions
- Agents are liable for the violation of the commitments they have taken
- Agent can always decide to break a commitment if an *opportunity*, it can take advantage of, arises (of course, at its own risk!)

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Searle's view [Searle, 1969]

He introduces a distiction between the *constitutive* and *regulative* specifications of interaction

Constitutive and regulative rules

- *Constitutive rules*, by identifying certain behaviors as foundational of a certain type of activity, create that activity
- *Regulative rules*, in contrast, contingently constrain a previously constituted activity

Searle's view [Searle, 1969]

He introduces a distiction between the *constitutive* and *regulative* specifications of interaction

Constitutive and regulative rules

- *Constitutive rules*, by identifying certain behaviors as foundational of a certain type of activity, create that activity
- *Regulative rules*, in contrast, contingently constrain a previously constituted activity

Searle's view [Searle, 1969]

He introduces a distiction between the *constitutive* and *regulative* specifications of interaction

Constitutive and regulative rules

- *Constitutive rules*, by identifying certain behaviors as foundational of a certain type of activity, create that activity
- *Regulative rules*, in contrast, contingently constrain a previously constituted activity

Searle's view [Searle, 1969]

He introduces a distiction between the *constitutive* and *regulative* specifications of interaction

Constitutive and regulative rules

- *Constitutive rules*, by identifying certain behaviors as foundational of a certain type of activity, create that activity
- *Regulative rules*, in contrast, contingently constrain a previously constituted activity

Constitutive and regulative specification in a commitment-based protocol

- A commitment protocol is made of a set of actions, most of works adopt a *precondition-effect* view of actions (preconditions to the action execution)
- The agents *respect* the commitments they have taken

In other words

1. Constitutive rules specify the [illegible] while regulative rules specify the [illegible]
2. Commitments have a regulative nature in the sense that people that take a commitment are bound to make the [illegible] true, [illegible] commitment becomes true

Constitutive and regulative specification in a commitment-based protocol

- A commitment protocol is made of a set of actions, most of works adopt a *precondition-effect* view of actions (preconditions to the action execution)
- The agents *respect* the commitments they have taken

In other words

- Constitutive rules specify the *semantics of actions* while regulative rules specify the *flow of execution*
- Commitments have a *regulative* nature in the sense that agents that take a commitment are bound to make the condition in the commitment become true

Constitutive and regulative specification in a commitment-based protocol

- A commitment protocol is made of a set of actions, most of works adopt a *precondition-effect* view of actions (preconditions to the action execution)
- The agents *respect* the commitments they have taken

In other words

- Constitutive rules specify the *semantics of actions* while regulative rules specify the *flow of execution*
- Commitments have a *regulative* nature in the sense that agents that take a commitment are bound to make the condition in the commitment become true

Constitutive and regulative specification in a commitment-based protocol

- A commitment protocol is made of a set of actions, most of works adopt a *precondition-effect* view of actions (preconditions to the action execution)
- The agents *respect* the commitments they have taken

In other words

- Constitutive rules specify the *semantics of actions* while regulative rules specify the *flow of execution*
- Commitments have a *regulative* nature in the sense that agents that take a commitment are bound to make the condition in the commitment become true

Constitutive and regulative specification enables verifications and predictability of the agents' behaviour

- Regulative specifications have a normative nature and create free social obligations, expectations, rights that anyone can observe *without the need of any introspection* of the agents' mental states
- By *autonomously deciding* whether entering the protocol an agent 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 participant

Constitutive and regulative specification enables verifications and predictability of the agents' behaviour

- Regulative specifications have a normative nature and create free social obligations, expectations, rights that anyone can observe *without the need of any introspection* of the agents' mental states
- By *autonomously deciding* whether entering the protocol an agent 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 participant

Constitutive and regulative specification enables verifications and predictability of the agents' behaviour

- Regulative specifications have a normative nature and create free social obligations, expectations, rights that anyone can observe *without the need of any introspection* of the agents' mental states
- By *autonomously deciding* whether entering the protocol an agent 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 participant

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

E-Institutions and Organizations

- Commitment-based protocols enable the verification of compliance of an open MAS, but how actually to perform this task?
- E-institutions [Arcos et al., 2007] and Artificial institutions [Fornara et al., 2008] solve exactly this task
  - They represent the social state of interaction and monitor the execution of the institutional actions
  - They assure the enforcement of institutional rules and the institutional policy
- Organizations [Zambonelli et al., 2003, Hübner et al., 2007, Hübner et al., 2009] may help also to regiment behavior of agents and to structure the system

E-Institutions and Organizations

- Commitment-based protocols enable the verification of compliance of an open MAS, but how actually to perform this task?
- E-institutions [Arcos et al., 2007] and Artificial institutions [Fornara et al., 2008] solve exactly this task
  - They represent the social state of interaction and monitor the execution of the institutional actions
  - They detect the violation of regulative rules and decide the sanction to apply
- Organizations [Zambonelli et al., 2003, Hübner et al., 2007, Hübner et al., 2009] may help also to regiment behavior of agents and to structure the system

E-Institutions and Organizations

- Commitment-based protocols enable the verification of compliance of an open MAS, but how actually to perform this task?
- E-institutions [Arcos et al., 2007] and Artificial institutions [Fornara et al., 2008] solve exactly this task
  - They represent the social state of interaction and monitor the execution of the institutional actions
  - They detect the violation of regulative rules and decide the sanction to apply
- Organizations [Zambonelli et al., 2003, Hübner et al., 2007, Hübner et al., 2009] may help also to regiment behavior of agents and to structure the system

E-Institutions and Organizations

- Commitment-based protocols enable the verification of compliance of an open MAS, but how actually to perform this task?
- E-institutions [Arcos et al., 2007] and Artificial institutions [Fornara et al., 2008] solve exactly this task
  - They represent the social state of interaction and monitor the execution of the institutional actions
  - They detect the violation of regulative rules and decide the sanction to apply
- Organizations [Zambonelli et al., 2003, Hübner et al., 2007, Hübner et al., 2009] may help also to regiment behavior of agents and to structure the system

E-Institutions and Organizations

- Commitment-based protocols enable the verification of compliance of an open MAS, but how actually to perform this task?
- E-institutions [Arcos et al., 2007] and Artificial institutions [Fornara et al., 2008] solve exactly this task
  - They represent the social state of interaction and monitor the execution of the institutional actions
  - They detect the violation of regulative rules and decide the sanction to apply
- Organizations [Zambonelli et al., 2003, Hübner et al., 2007, Hübner et al., 2009] may help also to regiment behavior of agents and to structure the system

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

What is missing in this picture?

Many practical contexts require to express rules that regulate the behavior of agents, the way in which things are done

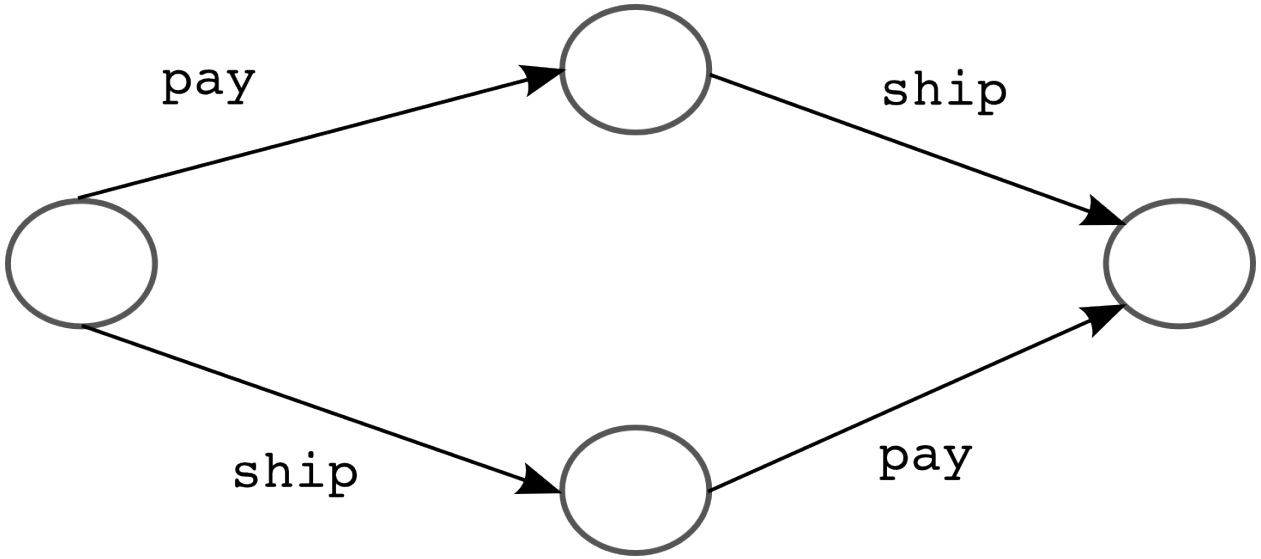

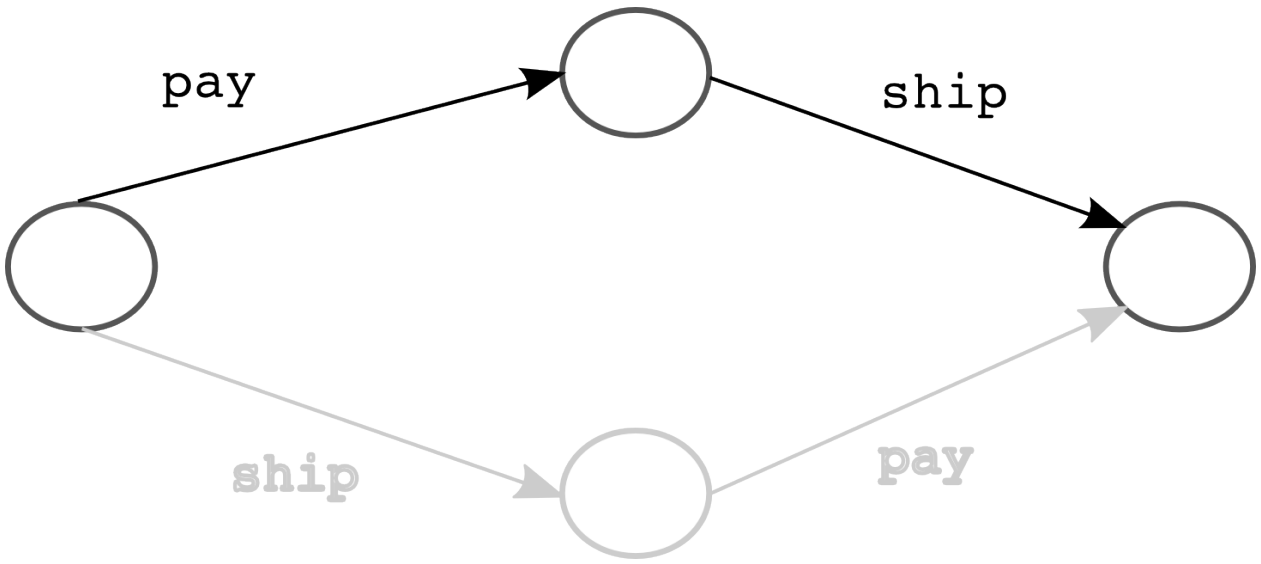

They bring to the same state but ... one may wish to say that only payment followed by shipment is acceptable

Working at the level of actions

Solution 1: use preconditions to the executability [Winikoff et al., 2004, Chopra and Singh, 2006, Giordano et al., 2007, Chopra, 2009]

shipping **possibleIf** *paid*

Conceptually not adequate

The act of shipping must follow payment by convention, because the shipper does not trust the buyer. It is a *regulation* to be imposed on the existing behaviors of paying and of shipping

Practically not adequate

Little reusability of actions (they are contextualized) and of regulations (hidden inside actions)

Working at the level of actions

Solution 1: use preconditions to the executability [Winikoff et al., 2004, Chopra and Singh, 2006, Giordano et al., 2007, Chopra, 2009]

shipping **possibleIf** *paid*

Conceptually not adequate

The act of shipping must follow payment by convention, because the shipper does not trust the buyer. It is a *regulation* to be imposed on the existing behaviors of paying and of shipping

Practically not adequate

Little reusability of actions (they are contextualized) and of regulations (hidden inside actions)

Working at the level of actions

Solution 1: use preconditions to the executability [Winikoff et al., 2004, Chopra and Singh, 2006, Giordano et al., 2007, Chopra, 2009]

shipping **possibleIf** *paid*

Conceptually not adequate

The act of shipping must follow payment by convention, because the shipper does not trust the buyer. It is a *regulation* to be imposed on the existing behaviors of paying and of shipping

Practically not adequate

Little reusability of actions (they are contextualized) and of regulations (hidden inside actions)

Commitment-based protocols lack of means for:

- *Prevision*: allowing agents to have expectations on each others' behavior. Commitments create expectations on what will be achieved but not on how;
- *Ordering*: expressing patterns of interaction, keeping the same flexibility shown by commitments.

Searle's Proposal [Searle, 1969]:

- *Constitutive Specification*: It identifies certain actions as foundational of a certain type of activity and specifies their semantics;
- *Regulative Specification*: contingently constrains a previously constituted activity. It captures how things should be carried on.

Commitment-based protocols lack of means for:

- *Prevision*: allowing agents to have expectations on each others' behavior. Commitments create expectations on what will be achieved but not on how;
- *Ordering*: expressing patterns of interaction, keeping the same flexibility shown by commitments.

Searle's Proposal [Searle, 1969]:

- *Constitutive Specification*: It identifies certain actions as foundational of a certain type of activity and specifies their semantics;
- *Regulative Specification*: contingently constrains a previously constituted activity. It captures how things should be carried on.

Which is the best way to specify patterns of interaction?

In a way that:

- does not compromise the *flexibility* of commitment-based protocols;
- fosters *openness* (agents can easily enter/leave a system);
- introduces modularity:
  - easier *re-use* of actions in different contexts
  - easier *customization* of protocols
  - easier *composition* of protocols
  - allows the *verification* of properties like interoperability and conformance

Interaction diagrams

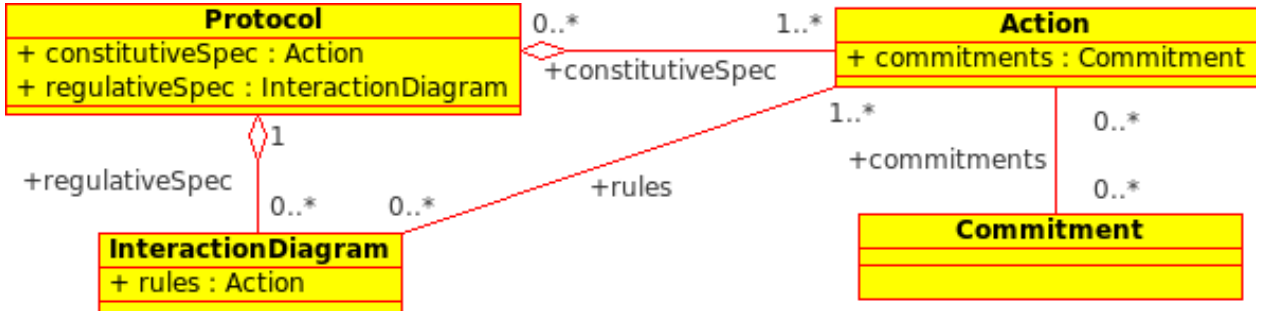

- Specify the allowed sequences of actions (interaction diagrams [Fornara and Colombetti, 2004])
  - *intuitive* to write and to read;
  - *strict* because prescriptive: agents' behavior limited to the specified sequences (e.g. agents are not free to take advantage of opportunities);
  - *little openness*: adding a new action means updating the set of allowed sequences;
  - *little reusability*.

Dependencies

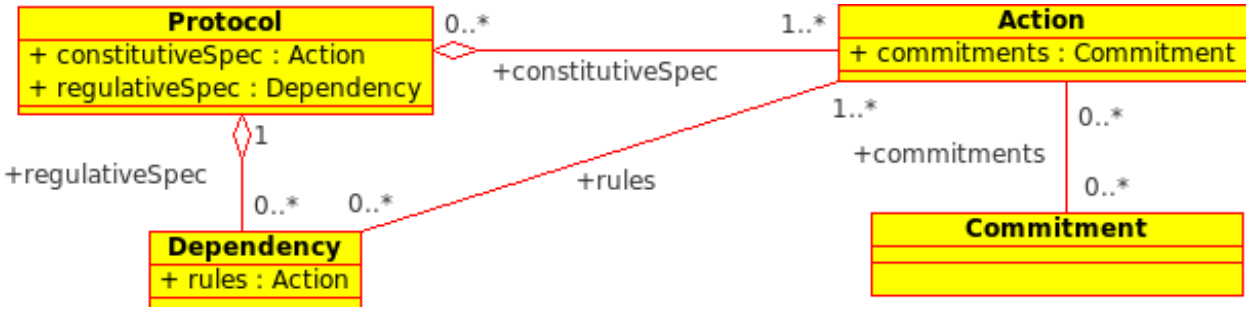

- Ordering the possible executions according to a preference criterion given in terms of actions ([Mallya and Singh, 2006]) or defining a *before* relation applied to events ([Singh, 2003])
  - *little openness*: order is given in terms of actions;
  - *little reusability*.

Solutions coming from the Business Processes' literature

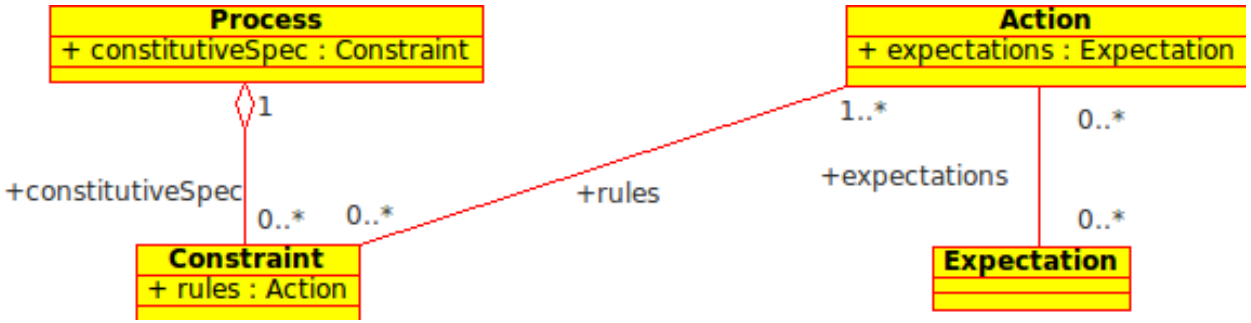

- [Pesic and van der Aalst, 2006] uses a *declarative* language (ConDec) for pecifying business processes.
- [Montali, 2009, M. et al., 2009] integrates ConDec with SCIFF thus giving a semantics to actions based on *expectations*
- ConDec specifications have a *constitutive nature*.

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Our proposal [Baldoni et al., 2010a, Baldoni et al., 2010b]

A decoupled representation of the constitutive and of the regulative specifications of the protocol, which are both based on commitments

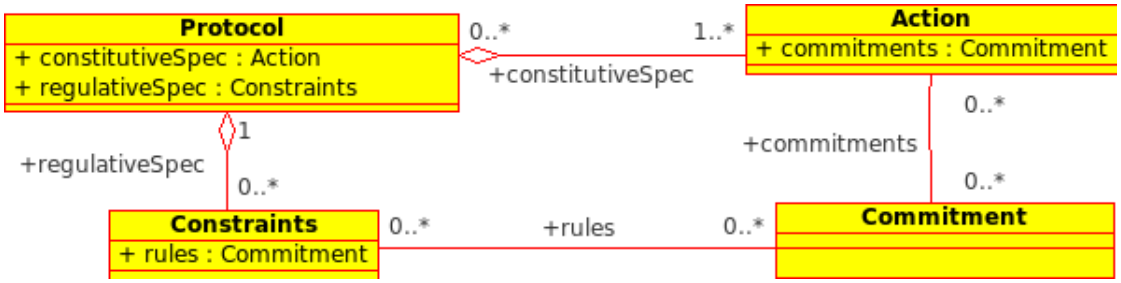

- the *constitutive* specification defines the meaning of actions based on their effects on the social state
- the *regulative* specification is a set of behavioral rules, given in terms of *constraints among commitments*, which regulate the evolution of the social state independently from the executed actions

Our proposal [Baldoni et al., 2010a, Baldoni et al., 2010b]

A decoupled representation of the constitutive and of the regulative specifications of the protocol, which are both based on commitments

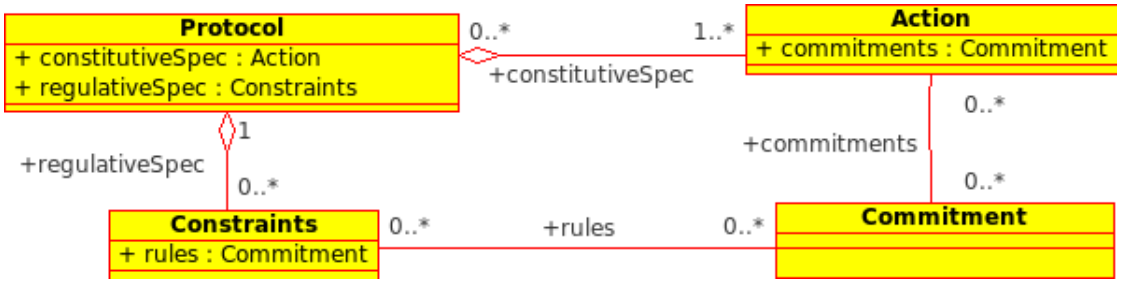

- the *constitutive* specification defines the meaning of actions based on their effects on the social state
- the *regulative* specification is a set of behavioral rules, given in terms of *constraints among commitments*, which regulate the evolution of the social state independently from the executed actions

Our proposal [Baldoni et al., 2010a, Baldoni et al., 2010b]

A decoupled representation of the constitutive and of the regulative specifications of the protocol, which are both based on commitments

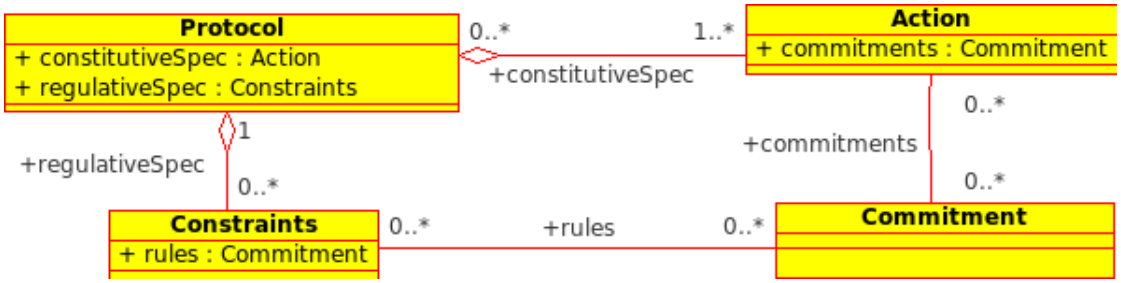

- the *constitutive* specification defines the meaning of actions based on their effects on the social state
- the *regulative* specification is a set of behavioral rules, given in terms of *constraints among commitments*, which regulate the evolution of the social state independently from the executed actions

Protocols: Constitutive specification

- The *constitutive specification* defines *the meaning of actions*
- We adopt the solution in [Chopra, 2009] (and also [Singh, 1999, Yolum and Singh, 2001])

Constitutive Specification of RONR

(a) $motion(c, m)$ **means** $\forall_{p_i \in P} \text{CREATE}(C(c, p_i, cfv(m)))$
(b) $openDebate(c, p, m)$ **means** $\text{CREATE}(C(c, p, assignFloor(p, m)))$
(c) $refuseFloor(p, c, m)$ **means** $refusedFloor(p, m) \wedge$
$\text{RELEASE}(C(c, p, assignFloor(p, m)))$
(d) $askFloor(p, c, m)$ **means** $\text{CREATE}(C(p, c, discussed(p, m)))$
(e) $recognition(c, p, m)$ **means** $assignFloor(p, m)$
(f) $startTalk(p, m)$ **means** $discussing(p, m)$
(g) $stopTalk(p, m)$ **means** $discussed(p, m) \wedge \neg discussing(p, m) \wedge$
$\neg assignFloor(p, m)$
(h) $timeOut(c, p)$ **means** $discussed(p, m) \wedge \neg discussing(p, m) \wedge$
$\text{RELEASE}(C(p, c, discussed(p, m))) \wedge \neg assignFloor(p, m)$
(i) $cfv(c, m)$ **means** $cfv(m)$
(l) $vote(p, m)$ **means** $voted(p, m)$

2CL: *Constraints among Commitments*

We defined 2CL:

- it is grounded on LTL and allows declarative, constraint-based representation of pattern of interaction;
- graphical language;
- we defined a set of operators and their negations:
  - constraints of type *base* express relations between the fluents, saying what and when should become true in the social state.
  - constraints of type *persistence* expresses the fact that a condition of interest holds in all the states until another condition of interest becomes true.

Relation	*Type*	*Positive*	*LTL meaning*	*Negative*	*LTL meaning*
Correlation	base	$A \bullet\!\!- B$	$\Diamond A \supset \Diamond B$	$A \bullet\!\!\not- B$	$\Diamond A \supset \neg\Diamond B$
	persistence	$A \bullet\!\!= B$	$\Box(A \supset (A \wedge B))$	$A \bullet\!\!\not= B$	$\Box(A \supset \neg(A \wedge B))$
Co-existence	base	$A \bullet\!\!-\!\!\bullet B$	$A \bullet\!\!- B \wedge B \bullet\!\!- A$	$A \bullet\!\!\not-\!\!\bullet B$	$A \bullet\!\!\not- B \wedge B \bullet\!\!\not- A$
	persistence	$A \bullet\!\!=\!\!\bullet B$	$A \bullet\!\!= B \wedge B \bullet\!\!= A$	$A \bullet\!\!\not=\!\!\bullet B$	$A \bullet\!\!\not= B \wedge B \bullet\!\!\not= A$
Response	base	$A \bullet\!\!-\!\!\triangleright B$	$\Box(A \supset \Diamond B)$	$A \bullet\!\!\not-\!\!\triangleright B$	$\Box(A \supset \neg\Diamond B)$
	persistence	$A \bullet\!\!=\!\!\triangleright B$	$\Box(A \supset (\Diamond B \wedge (A \cup B)))$	$A \bullet\!\!\not=\!\!\triangleright B$	$\Box(A \supset \neg(A \wedge B))$
Before	base	$A -\!\!\triangleright\!\!\bullet B$	$\neg B \cup A$	$A \not-\!\!\triangleright\!\!\bullet B$	$\Box(\Diamond B \supset \neg A)$
	persistence	$A =\!\!\triangleright\!\!\bullet B$	$\neg B \cup (A \cup B)$	$A \not=\!\!\triangleright\!\!\bullet B$	$\Box(\Diamond B \supset \neg A)$
Cause	base	$A \bullet\!\!-\!\!\triangleright\!\!\bullet B$	$A \bullet\!\!-\!\!\triangleright B \wedge A -\!\!\triangleright\!\!\bullet B$	$A \bullet\!\!\not-\!\!\triangleright\!\!\bullet B$	$A \bullet\!\!\not-\!\!\triangleright B \wedge A \not-\!\!\triangleright\!\!\bullet B$
	persistence	$A \bullet\!\!=\!\!\triangleright\!\!\bullet B$	$A =\!\!\triangleright\!\!\bullet B \wedge A \bullet\!\!=\!\!\triangleright B$	$A \bullet\!\!\not=\!\!\triangleright\!\!\bullet B$	$A \not=\!\!\triangleright\!\!\bullet B \wedge A \bullet\!\!\not=\!\!\triangleright B$
Premise	base	$A \triangleright\!\!- B$	$\Box(\bigcirc B \supset A)$	$A \not\triangleright\!\!- B$	$\Box(\bigcirc B \supset \neg A)$
Immediate after	base	$A -\!\!\triangleright B$	$\Box(A \supset \bigcirc B)$	$A \not-\!\!\triangleright B$	$\Box(A \supset \bigcirc\neg B)$

Table: 2CL constraint relations and their semantics in LTL.

Protocols: Regulative specification

- Constraints have a *normative nature* but they *do not* specify *which* actions should bring conditions about
- *Any* evolution that respects the relations involving the specified literals (and *commitments*) respects the protocol

Regulative Specification of RONR

c1: $C(c, p, cfv(m)) \bullet\!\!\rightarrow\!\!\bullet\ C(c, p, assignFloor(p, m))$

c2: $C(c, p, assignFloor(p, m)) \bullet\!\!\rightarrow\!\!\bullet$
$C(p, c, discuss(p, m))$ XOR $refusedFloor(p, m)$

c3: $C(p, c, discuss(p, m)) \bullet\!\!\rightarrow\!\!\bullet\ assignFloor(p, m)$

c4: $assignFloor(p, m) \bullet\!\!\Rrightarrow\ discussed(p, m)$

c5: $assignFloor(p, m) \rightarrowtail\!\!\!\!-\ discussing(p, m)$

c6: $discussing(p, m) \bullet\!\!\rightarrow\ discussed(p, m)$

c7: $discussed(p, m) \bullet\!\!\rightarrow\!\!\bullet\ \neg assignFloor(p, m)$

c8: $refusedFloor(p, m) \vee discussed(p, m) \bullet\!\!\rightarrow\!\!\bullet\ cfv(m)$

c9: $cfv(m) \rightarrow\!\!\bullet\ voted(p, m)$

c10: $assignFloor(X, _) \bullet\!\!\not\Rrightarrow\!\!\bullet\ assignFloor(Y, _) \wedge X \neq Y$

No-Flow-In-Flow

- they rule the evolution of the social state by expressing only what is mandatory for the protocol and what is forbidden;
- constraints do not specify a procedure;
- *perception of a flow*, normally missing in declarative approaches as criticised by [Miller and McGinnis, 2008].

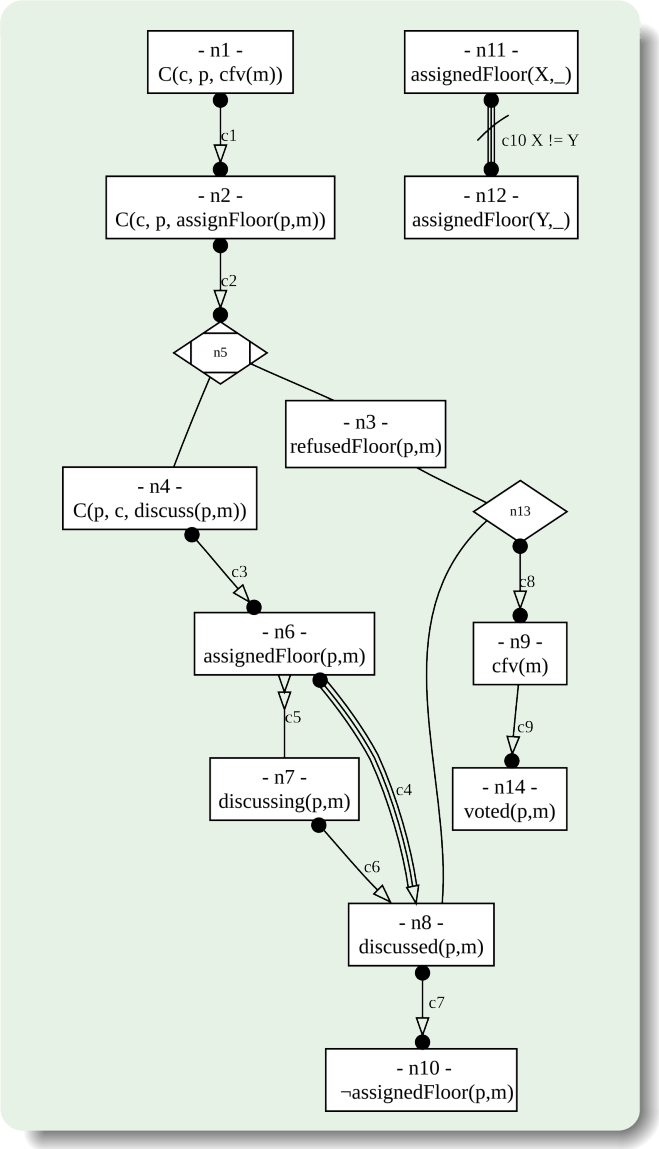

Advantages of this specification of protocols

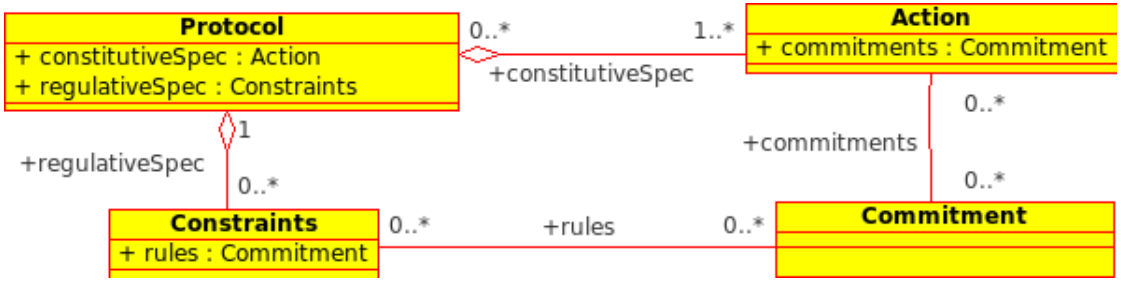

Modularity

- constraints are not defined on actions
- constitutive and regulative specifications share only the *domain model*
- the constitutive part can be modified without modifying the regulative part, and the other way around

Tailoring

- Protocols can easily be adapted to the needs of different contexts

Advantages of this specification of protocols

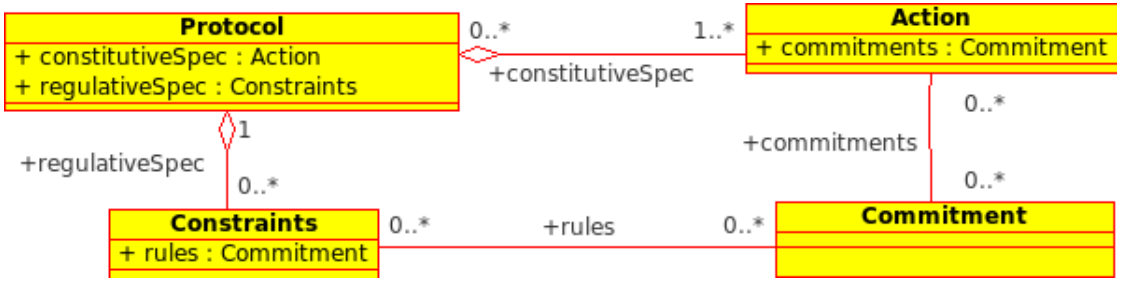

Modularity

- constraints are not defined on actions
- constitutive and regulative specifications share only the *domain model*
- the constitutive part can be modified without modifying the regulative part, and the other way around

Tailoring

- Protocols can easily be adapted to the needs of different contexts

1 What is a protocol

2 Are BDI enough to represent interactions in open MAS?

3 Castenfranchi and Singh's commitments

4 Are FSMs enough to represent interactions in open MAS?

5 Singh's commitment protocols

6 On the constitutive and regulative specification of a protocol

7 Verifying a protocol interaction

8 Need for a more powerful regulation

9 2CL: constraints among commitments

10 Conclusions

Conclusions

- Verification of properties such as interoperability, operability, substitutability, conformance, compliance is still an open problem for commitment-based approaches
- Direct communication and communication mediated by the environment: The MERCURIO proposal
- Efficent techniques to model check specifications
- Exploring graphical approaches for designing protocols: composition of protocols, bottom-up development, ...

Thanks to ...

Viviana and I ...

...would like to thank the members of our research group in Torino and in particular, *Cristina Baroglio*, *Elisa Marengo*, and *Alberto Martelli*. Most of the works that we presented have been done with their important support.

Arcos, J. L., Noriega, P., Rodríguez-Aguilar, J. A., and Sierra, C. (2007).
E4MAS Through Electronic Institutions.
In Weyns, D., Van Dyke Parunak, H., and Michel, F., editors, *Environments for Multi-Agent Systems III, Third International Workshop, E4MAS 2006*, volume 4389 of *Lecture Notes in Computer Science*, pages 184–202. Springer.

Baldoni, M., Baroglio, C., and Marengo, E. (2010a).
Behavior-oriented Commitment-based Protocols.
In Coelho, H., Studer, R., and Wooldridge, M., editors, *Proc. of 19th European Conference on Artificial Intelligence, ECAI 2010*, pages 137–142, Lisbon, Portugal. IOS Press.

Baldoni, M., Baroglio, C., and Marengo, E. (2010b).
Constraints among Commitments: Regulative Specification of Interaction Protocols.
In Artikis, A., Bentahar, J., Artikis, A., and Dignum, F., editors, *Proc. of International Workshop on Agent Communication, AC 2010, held in conjuction with AAMAS 2010*, pages 2–18, Toronto, Canada.

Castelfranchi, C. (1997).
Principles of individual social action.
In Tuomela R., H. G., editor, *Contemporary Action Theory*, pages 24–29. Kluwer, Dordrecht.

Castelfranchi, C. (1998).
Modelling social action for ai agents.
Artificial Intelligence, N° 103:pp.157–182.

Chopra, A. K. (2009).
Commitment Alignment: Semantics, Patterns, and Decision Procedures for Distributed Computing.
PhD thesis, North Carolina State University, Raleigh, NC.

Chopra, A. K. and Singh, M. P. (2006).
Contextualizing commitment protocol.
In *Proc. of AAMAS'06*, pages 1345–1352. ACM.

Cohen, P., Levesque, H., and Smith, I. (1997).
On team formation.

In *Contemporary Action Theory. Synthese*, pages 87–114. Kluwer Academic Publishers.

Cohen, P. R. and Levesque, H. J. (1990).
Intention is choice with commitment.
Artif. Intell., 42(2-3):213–261.

Fornara, N. and Colombetti, M. (2004).
A Commitment-Based Approach To Agent Communication.
Applied Artificial Intelligence, 18(9-10):853–866.

Fornara, N., Viganò, F., Verdicchio, M., and Colombetti, M. (2008).
Artificial institutions: a model of institutional reality for open multiagent systems.
Artif. Intell. Law, 16(1):89–105.

Giordano, L., Martelli, A., and Schwind, C. (2007).
Specifying and verifying interaction protocols in a temporal action logic.
J. Applied Logic, 5(2):214–234.

Hübner, J. F., Boissier, O., Kitio, R., and Ricci, A. (2009).

Instrumenting multi-agent organisations with organisational artifacts and agents: "Giving the organisational power back to the agents". *Autonomous Agents and Multi-Agent Systems*, 20.

Hübner, J. F., Sichman, J. S., and Boissier, O. (2007).
Developing organised multiagent systems using the MOISE.
IJAOSE, 1(3/4):370–395.

M., M., Pesic, M., van der Aalst, W. P., Chesani, F., Mello, P., and Storari, S. (2009).
Declarative specification and verification of service choreographies.
ACM Transactions on the Web.

Mallya, A. U. and Singh, M. P. (2006).
Introducing preferences into commitment protocols.
In *AC*, volume 3859 of *LNCS*, pages 136–149. Springer.

Miller, T. and McGinnis, J. (2008).
Amongst first-class protocols.
In *Proc. of Eng. Societies in the Agents World VIII*, volume 4995 of *LNCS*, pages 208–223. Springer.

Montali, M. (2009).
Specification and Verification of Declarative Open Interaction Models - A Logic-based framework.
PhD thesis, University of Bologna.

Norman, T. J., Carbogim, D. V., Krabbe, E. C. W., and Walton, D. N. (2004).
Argument and multi-agent systems.
In *Argumentation Machines: New Frontiers in Argument and Computation*, volume 9 of *Argumentation Library*, pages 15–54. Kluwer Academic Publishers, Dordrecht.

Pesic, M. and van der Aalst, W. M. P. (2006).
A Declarative Approach for Flexible Business Processes Management.
In *Proc. of. Business Process Management Workshops*, volume 4103 of *LNCS*, pages 169–180. Springer.

Searle, J. (1969).
Speech Acts.
Cambridge University Press.

Singh, M. P. (1999).
An ontology for commitments in multiagent systems.
Artif. Intell. Law, 7(1):97–113.

Singh, M. P. (2003).
Distributed enactment of multiagent workflows: temporal logic for web service composition.
In *AAMAS*, pages 907–914. ACM.

Singh, M. P. (2008).
Semantical Considerations on Dialectical and Practical Commitments.
In Fox, D. and Gomes, C. P., editors, *Proc. of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008*, pages 176–181, Chicago, Illinois, USA. AAAI Press.

Singh, M. P. and Chopra, A. K. (2009).
Correctness properties for multiagent systems.
In *DALT*, volume 5948 of *Lecture Notes in Computer Science*, pages 192–207. Springer.

Winikoff, M., Liu, W., and Harland, J. (2004).

Enhancing commitment machines.
In *Proc. of DALT*, volume 3476 of *LNCS*, pages 198–220.

Yolum, P. and Singh, M. P. (2001).
Designing and executing protocols using the event calculus.
In *Agents*, pages 27–28.

Zambonelli, F., Jennings, N. R., and Wooldridge, M. (2003).
Developing multiagent systems: The Gaia methodology.
ACM Trans. Softw. Eng. Methodol., 12(3):317–370.