

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Multiagent Planning as an Emerging Behavior in Agent Societies

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/151901> since 2016-06-28T08:39:33Z

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Multiagent Planning as an Emerging Behavior in Agent Societies

Roberto Micalizio

Cristina Baroglio Matteo Baldoni

Dipartimento di Informatica
Università di Torino

NorMAS 2014

Outline

- MOTIVATIONS AND IDEA
 - Multiagent planning as Social Computing
- BACKGROUND
 - Classical Planning
 - Social Commitments & Goals
- SOCIAL CONTINUAL PLANNING by examples

Motivations

- Multiagent planning: synthesis of plans for a number of agents in a given *team*
 - each agent reaches its own goals
 - the agent plans are altogether consistent (i.e., no deadlock, no open preconditions, correct usage of resources)
- Multiagent planning as *distributed problem solving*:
 - agents are homogeneous
 - agents can trust each other
 - agents can inspect each other their beliefs
 - agents do not change over time (the team is fixed at the beginning)

⇒ *agents are not really autonomous*

These assumptions are unpractical when agents constitute a *society* rather than a team

Motivations

- Multiagent planning: synthesis of plans for a number of agents in a given *team*
 - each agent reaches its own goals
 - the agent plans are altogether consistent (i.e., no deadlock, no open preconditions, correct usage of resources)
- Multiagent planning as *distributed problem solving*:
 - agents are homogeneous
 - agents can trust each other
 - agents can inspect each other their beliefs
 - agents do not change over time (the team is fixed at the beginning)

⇒ *agents are not really autonomous*

These assumptions are unpractical when agents constitute a *society* rather than a team

Multiagent Planning as Social Computing

IDEA:

- Enrich the (classical) BDI planning agent with *social capabilities*
- The planning system is thought of as a *normative system*:
 - social norms define the constraints within which agents can operate
 - an agent's plan must be "socially acceptable"

HOW TO GET THERE:

- use of *social commitments* for modeling agent interactions

WHY?

- *commitments have a normative power*
 - ⇒ an agent can create expectations on the behaviors of others just relying on the active commitments
- *commitments are tightly related to goals* [Telang et al. 2011]
 - ⇒ a planning agent can be driven by the commitments it is responsible for
- *commitments enable practical reasoning, that can be seen as a form of planning*

Multiagent Planning as Social Computing

IDEA:

- Enrich the (classical) BDI planning agent with *social capabilities*
- The planning system is thought of as a *normative system*:
 - social norms define the constraints within which agents can operate
 - an agent's plan must be "socially acceptable"

HOW TO GET THERE:

- use of *social commitments* for modeling agent interactions

WHY?

- commitments have a normative power
 - ⇒ an agent can create expectations on the behaviors of others just relying on the active commitments
- commitments are tightly related to goals [Telang et al. 2011]
 - ⇒ a planning agent can be driven by the commitments it is responsible for
- commitments enable *practical reasoning*, that can be seen as a form of planning

Multiagent Planning as Social Computing

IDEA:

- Enrich the (classical) BDI planning agent with *social capabilities*
- The planning system is thought of as a *normative system*:
 - social norms define the constraints within which agents can operate
 - an agent's plan must be "socially acceptable"

HOW TO GET THERE:

- use of *social commitments* for modeling agent interactions

WHY?

- commitments have a normative power
 - ⇒ an agent can create expectations on the behaviors of others just relying on the active commitments
- commitments are tightly related to goals [Telang et al. 2011]
 - ⇒ a planning agent can be driven by the commitments it is responsible for
- commitments enable *practical reasoning*, that can be seen as a form of planning

Multiagent Planning as Social Computing

IDEA:

- Enrich the (classical) BDI planning agent with *social capabilities*
- The planning system is thought of as a *normative system*:
 - social norms define the constraints within which agents can operate
 - an agent's plan must be "socially acceptable"

HOW TO GET THERE:

- use of *social commitments* for modeling agent interactions

WHY?

- *commitments have a normative power*
 - ⇒ an agent can create expectations on the behaviors of others just relying on the active commitments
- *commitments are tightly related to goals* [Telang et al. 2011]
 - ⇒ a planning agent can be driven by the commitments it is responsible for
- *commitments enable practical reasoning, that can be seen as a form of planning*

Multiagent Planning as Social Computing

IDEA:

- Enrich the (classical) BDI planning agent with *social capabilities*
- The planning system is thought of as a *normative system*:
 - social norms define the constraints within which agents can operate
 - an agent's plan must be "socially acceptable"

HOW TO GET THERE:

- use of *social commitments* for modeling agent interactions

WHY?

- *commitments have a normative power*
 - ⇒ an agent can create expectations on the behaviors of others just relying on the active commitments
- *commitments are tightly related to goals* [Telang et al. 2011]
 - ⇒ a planning agent can be driven by the commitments it is responsible for
- *commitments enable practical reasoning, that can be seen as a form of planning*

Multiagent Planning as Social Computing

IDEA:

- Enrich the (classical) BDI planning agent with *social capabilities*
- The planning system is thought of as a *normative system*:
 - social norms define the constraints within which agents can operate
 - an agent's plan must be "socially acceptable"

HOW TO GET THERE:

- use of *social commitments* for modeling agent interactions

WHY?

- *commitments have a normative power*
 - ⇒ an agent can create expectations on the behaviors of others just relying on the active commitments
- *commitments are tightly related to goals* [Telang et al. 2011]
 - ⇒ a planning agent can be driven by the commitments it is responsible for
- *commitments enable practical reasoning, that can be seen as a form of planning*

Background: Classical Planning

- a single-agent *planning domain* $D : \langle P, S, A, R \rangle$
 - P is the (finite) set of atomic propositions
 - $S \subseteq 2^P$ is the set of possible states
 - A is the (finite) set of actions
 - $R \subseteq S \times A \times S$ is a transition relation
- a single-agent *planning problem* $Pr : \langle D, I, G \rangle$
 - D is the a planning domain
 - $I \subseteq S$ initial state
 - $G \subseteq S$ goal state
- a *solution* π for Pr is a sequence of actions $\langle a_1, \dots, a_n \rangle$ such that:
 - a_1 is applicable to the initial state I
 - a_i is applicable to the state resulting after the application of a_{i-1} (for $i : 2..n$)
 - G holds after the application of a_n

Background: Classical Planning


- a single-agent *planning domain* $D : \langle P, S, A, R \rangle$
 - P is the (finite) set of atomic propositions
 - $S \subseteq 2^P$ is the set of possible states
 - A is the (finite) set of actions
 - $R \subseteq S \times A \times S$ is a transition relation
- a single-agent *planning problem* $Pr : \langle D, I, G \rangle$
 - D is the a planning domain
 - $I \subseteq S$ initial state
 - $G \subseteq S$ goal state
- a *solution* π for Pr is a sequence of actions $\langle a_1, \dots, a_n \rangle$ such that:
 - a_1 is applicable to the initial state I
 - a_i is applicable to the state resulting after the application of a_{i-1} (for $i : 2..n$)
 - G holds after the application of a_n

Background: Classical Planning


- a single-agent *planning domain* $D : \langle P, S, A, R \rangle$
 - P is the (finite) set of atomic propositions
 - $S \subseteq 2^P$ is the set of possible states
 - A is the (finite) set of actions
 - $R \subseteq S \times A \times S$ is a transition relation
- a single-agent *planning problem* $Pr : \langle D, I, G \rangle$
 - D is the a planning domain
 - $I \subseteq S$ initial state
 - $G \subseteq S$ goal state
- a *solution* π for Pr is a sequence of actions $\langle a_1, \dots, a_n \rangle$ such that:
 - a_1 is applicable to the initial state I
 - a_i is applicable to the state resulting after the application of a_{i-1} (for $i : 2..n$)
 - G holds after the application of a_n

Background: Commitments and Goals

Life cycle of a commitment



Life cycle of a goal



Background: Commitments and Goals

- the relation between commitments and goals has been captured by a set of rules [Telang et al. 2011]:
 - *structural rules*: complete and deterministic, describe how commitment and goal states evolve
 - *pragmatical rules*: describe patterns of practical reasoning over commitments and goals; these rules are neither complete nor deterministic

Background: Pragmatical Rules

$$\frac{guard}{S_1 \rightarrow S_2}$$

- *guard* is a condition over an agent beliefs and over the active commitments
- $S_1 \rightarrow S_2$ is a state transition defining how goals and commitments change

Background: Pragmatical Rules

$$\frac{\textit{guard}}{S_1 \rightarrow S_2}$$

- *guard* is a condition over an agent beliefs and over the active commitments
- $S_1 \rightarrow S_2$ is a state transition defining how goals and commitments change

- Pragmatical Rules are divided into:

- rules from **goals to commitments**

$$\frac{\langle G^A, C^N \rangle}{\textit{create}(C)} \text{ ENTICE}$$

- rules from **commitments to goals**

$$\frac{\langle G^N, C^D \rangle}{\textit{consider}(G), \textit{activate}(G)} \text{ DELIVERY}$$

Reasoning about Goal and Commitments via **Social Continual Planning**

MAIN IDEA:

- interleave *planning* phases with *execution* and *negotiation* phases
- the **planning phase** involves both:
 - “physical” actions: directly change the world
 - *pragmatical actions*: (indirectly) change the social state
- during the **execution phase**:
 - a physical action is directly performed by an agent
 - a *pragmatical action* triggers a negotiation with others
- **negotiation** involves operations on commitments and it is driven by *pragmatical rules*

Reasoning about Goal and Commitments via **Social Continual Planning**

MAIN IDEA:

- interleave *planning* phases with *execution* and *negotiation* phases
- the **planning phase** involves both:
 - “physical” actions: directly change the world
 - **pragmatical actions**: (indirectly) change the social state
- during the **execution phase**:
 - a physical action is directly performed by an agent
 - a pragmatical action triggers a negotiation with others
- **negotiation** involves operations on commitments and it is driven by pragmatical rules

Reasoning about Goal and Commitments via **Social Continual Planning**

MAIN IDEA:

- interleave *planning* phases with *execution* and *negotiation* phases
- the **planning phase** involves both:
 - “physical” actions: directly change the world
 - *pragmatical actions*: (indirectly) change the social state
- during the **execution phase**:
 - a physical action is directly performed by an agent
 - a pragmatical action triggers a negotiation with others
- **negotiation** involves operations on commitments and it is driven by pragmatical rules

Reasoning about Goal and Commitments via **Social Continual Planning**

MAIN IDEA:

- interleave *planning* phases with *execution* and *negotiation* phases
- the **planning phase** involves both:
 - “physical” actions: directly change the world
 - *pragmatical actions*: (indirectly) change the social state
- during the **execution phase**:
 - a physical action is directly performed by an agent
 - a pragmatical action triggers a negotiation with others
- **negotiation** involves operations on commitments and it is driven by pragmatical rules

Pragmatical Rules to Define Agent's Strategy

- pragmatical rules **from commitments to goals** define the *strategy* of an agent (i.e., *when to trigger a planning phase*)

Pragmatical Rules to Define Agent's Strategy

- pragmatical rules **from commitments to goals** define the *strategy* of an agent (i.e., *when to trigger a planning phase*)
- e.g.

$$\frac{\langle G^N, C^D \rangle}{\text{consider}(G), \text{activate}(G)} \text{ DELIVERY}$$

“an honest agent activates a goal G when G appears as a consequent of a detached commitments it responsible for”


(but are all agents honest?)

Pragmatical Rules as Pragmatical Actions

- pragmatical rules **from goals to commitments** are thought of as pragmatical actions

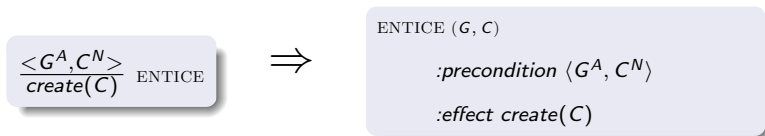
Pragmatic Rules as Pragmatic Actions

- pragmatic rules **from goals to commitments** are thought of as pragmatic actions



Pragmatical Rules as Pragmatical Actions

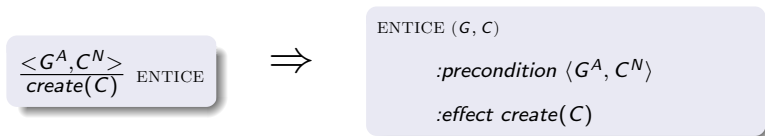
- pragmatical rules **from goals to commitments** are thought of as pragmatical actions



- ISSUE
 - how to determine over which goals and commitments these actions are defined?

Pragmatical Rules as Pragmatical Actions

- pragmatical rules **from goals to commitments** are thought of as pragmatical actions



- ISSUE
 - how to determine over which goals and commitments these actions are defined?
- SOLUTION
 - **blackboard of services**

Example: World-Wide Delivery Service

Problem: sending a parcel from Oklahoma City (Oklahoma) to Bertinoro (Italy)



four shipping agencies:

- *AmericanTrucks*: operates only in north America
- *EuropeanTrucks*: operates only in Europe
- *BlueVector* (flight company): blue connections
- *RedVector* (flight company): red connection

Conclusions

Social Continual Planning:

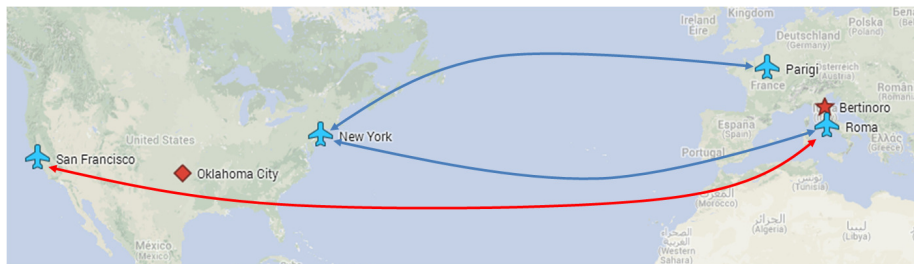
- practical reasoning as a form of planning
- agent's autonomy is preserved
 - an agent can adopt local optimization strategies
 - each agent can use the planner that suits it most
- commitments support flexible planning solutions
 - help agents take advantage of the opportunities available in a given time
 - help agents find alternative solutions when something wrong happens

multiagent planning = local agents' planning + social state

Thank you!

Example: World-Wide Delivery Service

Problem: sending a parcel from Oklahoma City (Oklahoma) to Bertinoro (Italy)



four shipping agencies:

- *AmericanTrucks*: operates only in north America
- *EuropeanTrucks*: operates only in Europe
- *BlueVector* (flight company): blue connections
- *RedVector* (flight company): red connection

Physical Actions

A subset of physical actions for the truck agencies

load(?t - truck ?p - parcel ?l - location)

:precondition $at(?t, ?l) \wedge at(?p, ?l)$

:effect $\neg at(?p, ?l) \wedge loaded(?p, ?t)$

drive(?t - truck ?l1, ?l2 - location)

:precondition $at(?t, ?l1)$

:effect $\neg at(?t, ?l1) \wedge at(?t, ?l2)$

deliver(?t - truck ?p - parcel ?l - location)

:precondition $at(?t, ?l) \wedge loaded(?p, ?t) \wedge dest(?p, ?l)$

:effect $\neg loaded(?p, ?t) \wedge at(?p, ?l) \wedge delivered(?p)$

Blackboard of Services

agent	service	price
AmericanTrucks	$at(?p, Oklahoma) \wedge delivered(?p)$	$\$?x$
	$at(?p, New York) \wedge delivered(?p)$	$\$?x$
	$at(?p, San Francisco) \wedge delivered(?p)$	$\$?x$

EuropeanTrucks	$at(?p, Rome) \wedge delivered(?p)$	$\$?x$
	$at(?p, Paris) \wedge delivered(?p)$	$\$?x$
	$at(?p, Bertinoro) \wedge delivered(?p)$	$\$?x$

BlueVector	$at(?p, Rome)$	$\$?x$
	$at(?p, Paris)$	$\$?x$
	$at(?p, New York)$	$\$?x$

RedVector	$at(?p, Rome)$	$\$?x$
	$at(?p, San Fransisco)$	$\$?x$

Pragmatical Actions

From the point of view of AmericanTrucks (AmT):

entice_delivery(?a - agent ?p - parcel ?l - location)

:precondition

$G^A(at(?p, ?l) \wedge delivery(?p)), C^N(AmT, ?a, at(?p, ?l) \wedge delivery(?p), \$?x)$

:effect create(C)

entice_at(?a - agent ?p - parcel ?l - location)

:precondition $G^A(at(?p, ?l), C^N(AmT, ?a, at(?p, ?l), \$?x)$

:effect create(C)

These new actions are made available to an off-the-shelf planner

Solving the Problem

- AmericanTrucks has to deliver parcel $p1$, initially located in Oklahoma City, to Bertinoro

`entice_delivery(AmT, EuT, {at(p1, Bertinoro), delivery(p1)}, $?x)`

- The planner finds a trivial plan: “ask EuropeanTrucks to deliver $p1$ ”
- The execution of such a pragmatic action triggers a negotiation phase between AmericanTrucks and EuropeanTrucks

Solving the Problem

As an effect of the negotiation...

Social State

CC(AmT, EuT, {at(p1, Bertinoro), delivery(p1)}, \$100)

CONDITIONAL

CC(EuT, AmT, at(p1, Rome),{at(p1, Bertinoro), delivery(p1)})

CONDITIONAL

- AmericanTrucks has now a new goal: $at(p1, Rome)$
- A new planning phase is activated

Solving the Problem

A new trivial plan is found:

entice_at(AmT, BlueV, at(p1, Rome), \$?x)

which triggers a new negotiation phase:

Social State

CC(AmT, EuT, {at(p1, Bertinoro), delivery(p1)}, \$100)

CONDITIONAL

CC(EuT, AmT, at(p1, Rome), {at(p1, Bertinoro), delivery(p1)})

CONDITIONAL

CC(AmT, BlueV, at(p1, Rome), \$500)

CONDITIONAL

CC(BlueV, AmT, at(p1, New York), at(p1, Rome))

CONDITIONAL

Solving the Problem

AmericanTrucks

load(AmTruck27, p1, OC)

drive(AmTruck27, OC, NY)

unload(AmTruck27, p1, OC)

CC(BlueV, AmT, at(p1, New York), at(p1, Rome))

CONDITIONAL

Solving the Problem

AmericanTrucks

load(AmTruck27, p1, OC)

drive(AmTruck27, OC, NY)

unload(AmTruck27, p1, OC)

CC(BlueV, AmT, T, at(p1, Rome))

DETACHED

Solving the Problem

AmericanTrucks

load(AmTruck27, p1, OC)

drive(AmTruck27, OC, NY)

unload(AmTruck27, p1, OC)

CC(BlueV, AmT, T, at(p1, Rome))

DETACHED

BlueVector

embark(BV5, p1, NY)

fly(BV5, NY, RM)

disembark(BV5, p1, RM)

CC(EuT, AmT, at(p1, Rome), {at(p1, Bertinoro), delivery(p1)})

CONDITIONAL

CC(AmT, BlueV, at(p1, Rome), \$500)

CONDITIONAL

Solving the Problem

AmericanTrucks

load(AmTruck27, p1, OC)

drive(AmTruck27, OC, NY)

unload(AmTruck27, p1, OC)

CC(BlueV, AmT, T, at(p1, Rome))
SATISFIED

BlueVector

embark(BV5, p1, NY)

fly(BV5, NY, RM)

disembark(BV5, p1, RM)

CC(EuT, AmT, T, {at(p1, Bertinoro), delivery(p1)})
DETACHED

CC(AmT, BlueV, T, \$500)
DETACHED

Solving the Problem

AmericanTrucks

load(AmTruck27, p1, OC)

drive(AmTruck27, OC, NY)

unload(AmTruck27, p1, OC)

CC(BlueV, AmT, T, at(p1, Rome))

SATISFIED

pay(BlueV, \$500)

BlueVector

embark(BV5, p1, NY)

fly(BV5, NY, RM)

disembark(BV5, p1, RM)

CC(EuT, AmT, T, {at(p1, Bertinoro), delivery(p1)})

SATISFIED

CC(AmT, BlueV, T, \$500)

SATISFIED

EuropeanTrucks

load(EuTruck13, p1, RM)

drive(EuTruck13, RM, BR)

deliver(EuTruck13, p1, BR)

CC(AmT, EuT, T, \$100)

DETACHED

Solving the Problem

AmericanTrucks

load(AmTruck27, p1, OC)

drive(AmTruck27, OC, NY)

unload(AmTruck27, p1, OC)

CC(BlueV, AmT, T, at(p1, Rome))

SATISFIED

pay(BlueV, \$500)

pay(EuT, \$100)

BlueVector

embark(BV5, p1, NY)

fly(BV5, NY, RM)

disembark(BV5, p1, RM)

CC(EuT, AmT, T, {at(p1, Bertinoro), delivery(p1)})

SATISFIED

CC(AmT, BlueV, T, \$500)

SATISFIED

EuropeanTrucks

load(EuTruck13, p1, RM)

drive(EuTruck13, RM, BR)

deliver(EuTruck13, p1, BR)

CC(AmT, EuT, T, \$100)

SATISFIED

BACKUP

Reasoning about Goal and Commitments via Continual Planning

- Given an agent x , its configuration is $S_x : \langle B, C, G \rangle$ [Telang]:
 - B : set of beliefs about the world state
(including beliefs about itself and others)
 - C : set of commitments of the form $C(x, y, s, u)$ (public)
 - G : set of goals of the form $G(x, p, r, q, s, f)$ (private)
- Extended agent configuration $S_x : \langle B, C, G, A_x, A_x^{gc}, R_x^{cg} \rangle$:
 - A_x : set of primitive actions for agent x (change a portion of the world)
 - A_x^{gc} : set of actions corresponding to pragmatical rules from goals to commitments (change the social state)
 - R_x^{cg} : set of reactive rules corresponding to pragmatical rules from commitments to goals (trigger planning phases)