Multi-Agent Systems

Sept. 2000, Bahia Blanca

University Nacional del Sur

- **Last two weeks in September**.
- **Tentative Dates:** Tuesday, Sept. 19th, Thursday, Sept. 21st, Friday, Sept. 22nd, Tuesday, Sept. 26th, Thursday, Sept. 28th, Friday, Sept. 29th.
- **Time:** From 4–6 pm, unless otherwise indicated.
- Lecture Course is on theoretical issues, emphasis on mathematical-logical foundations.

Overview

1. Introduction, Terminology
2. Three Basic Architectures
3. Logic Based Architectures
4. Distributed Decision Making
5. Contract Nets, Coalition Formation

Chapter 2. Three basic Architectures

2.1 Reactive Agents

2.2 BDI Agents

2.3 Layered Architectures

2 Three Basic Architectures

2.1 Reactive Agents

Intelligent behaviour is *Interaction of the agents with their environment*. It emerges through splitting in simpler interactions.

Subsumption-Architectures:

- Decision making is realized through **goal-directed behaviours**: each behaviour is an individual action.
  nonsymbolic implementation.
- Many behaviours can be applied **concurrently**. How to select between them?
  Implementation through Subsumption-Hierarchies, Layers.
  Upper layers represent abstract behaviour.

Formal Model

- **see**: as up to now, but close relation between observation and action: no transformation of the input .
- **action**: Set of behaviours and inhibition relation.

$$Beh := \{\langle \mathbf{c}, \mathbf{a} \rangle : \mathbf{c} \subseteq \mathbf{P}, \mathbf{a} \in \mathbf{A}\}.$$

$\langle \mathbf{c}, \mathbf{a} \rangle$ “fires” if $\mathbf{see}(\mathbf{s}) \in \mathbf{c}$ ($\mathbf{c}$ stands for “condition”).

$$\prec \subseteq Ag_{rules} \times Ag_{rules}$$

is called inhibition-relation, $Ag_{rules} \subseteq Beh$. We require $\prec$ to be a total ordering. $\mathbf{b_1} \prec \mathbf{b_2}$ means: b_1 inhibits b_2, **$\mathbf{b_1}$ has priority over $\mathbf{b_2}$**.

function $action(p : P) : A$
var $fired : \wp(R)$
var $selected : A$
begin
    $fired := \{(c,a) \mid (c,a) \in R$ and $p \in c\}$
    for each $(c,a) \in fired$ do
        if $\neg(\exists(c',a') \in fired$ such that $(c',a') \prec (c,a))$ then
            return a
        end-if
    end-for
    return $null$
end function $action$

Example 2.1 (Exploring a Planet)
A distant planet (asteroid) is assumed to contain gold. Samples should be brought to a spaceship landed on the planet. It is not known where the gold is. Several autonomous vehicles are available. Due to the topography of the planet there is no connection between the vehicles.

The spaceship sends off radio signals: **gradient field**.

Low Level Behaviour:

(1) **If** detect an obstacle **then** change direction.

2. Layer:

(2) **If** Samples on board **and** at base **then** drop off.

(3) **If** Samples on board **and** not at base **then** follow gradient field.

3. Layer:

(4) **If** Samples found **then** pick them up.

4. Layer:

(5) **If** true **then** take a random walk.

With the following ordering

$$(1) \prec (2) \prec (3) \prec (4) \prec (5).$$

Under which asumptions (on the distribution of the gold) does this work perfectly? What if the distribution is more realistic?

- Vehicles can **communicate indirectly** with each other:
  - they put off, and
  - pick up

  radiactive samples that can be sensed.

Low Level Behaviour:

(1) **If** detect an obstacle **then** change direction.

2. Layer:

(2) **If** Samples on board **and** at base **then** drop off.

(3) **If** Samples on board **and** not at base **then** drop off two radioactive crumbs and follow gradient field.

3. Layer:

(4) **If** Samples found **then** pick them up.

(5) **If** radiactive crumbs found **then** take one and follow the gradient field (away from the spaceship).

4. Layer:

(6) **If** true **then** take a random walk.

With the following ordering $(1) \prec (2) \prec (3) \prec (4) \prec (5) \prec (6)$.

Pro: Simple, economic, efficient, robust, elegant.

Contra:

- Without knowledge about the environment agents need to know about the own local environment.
- Decisions only based on local information.
- How about bringing in **learning**?
- Relation between agents, environment and behaviours is not clear.
- Agents with ≤ 10 behaviours are doable. But the more layers the more complicated to understand what is going on.

2.2 BDI-Architecture

Belief, **D**esire, **I**ntention.

From time to time intentions need to be re-examined. But they also should persist, normally. (**Pro-active vs. reactive**).

Extreme: *stubborn agents*, *unsure agents*.

What is better? Depends on the environment.
Let γ the **rate of world change**.

1. γ small: stubbornness pays off.
2. γ big: unsureness pays off.

Belief 1:	Going to lectures is worth doing to learn something.
Belief 2:	Dix is a decent lecturer.
Desire 1:	Visit Dix-Lecture, in addition read books.
Intention:	Getting knowledge about Distributed Systems.
New Belief:	Alejandro makes it much better. Therefore revise your Desire.
Desire 2:	Visit Garcia-Lecture, in addition read books.

Of course, Alejandro may turn out to be the worst lecturer from all . . .

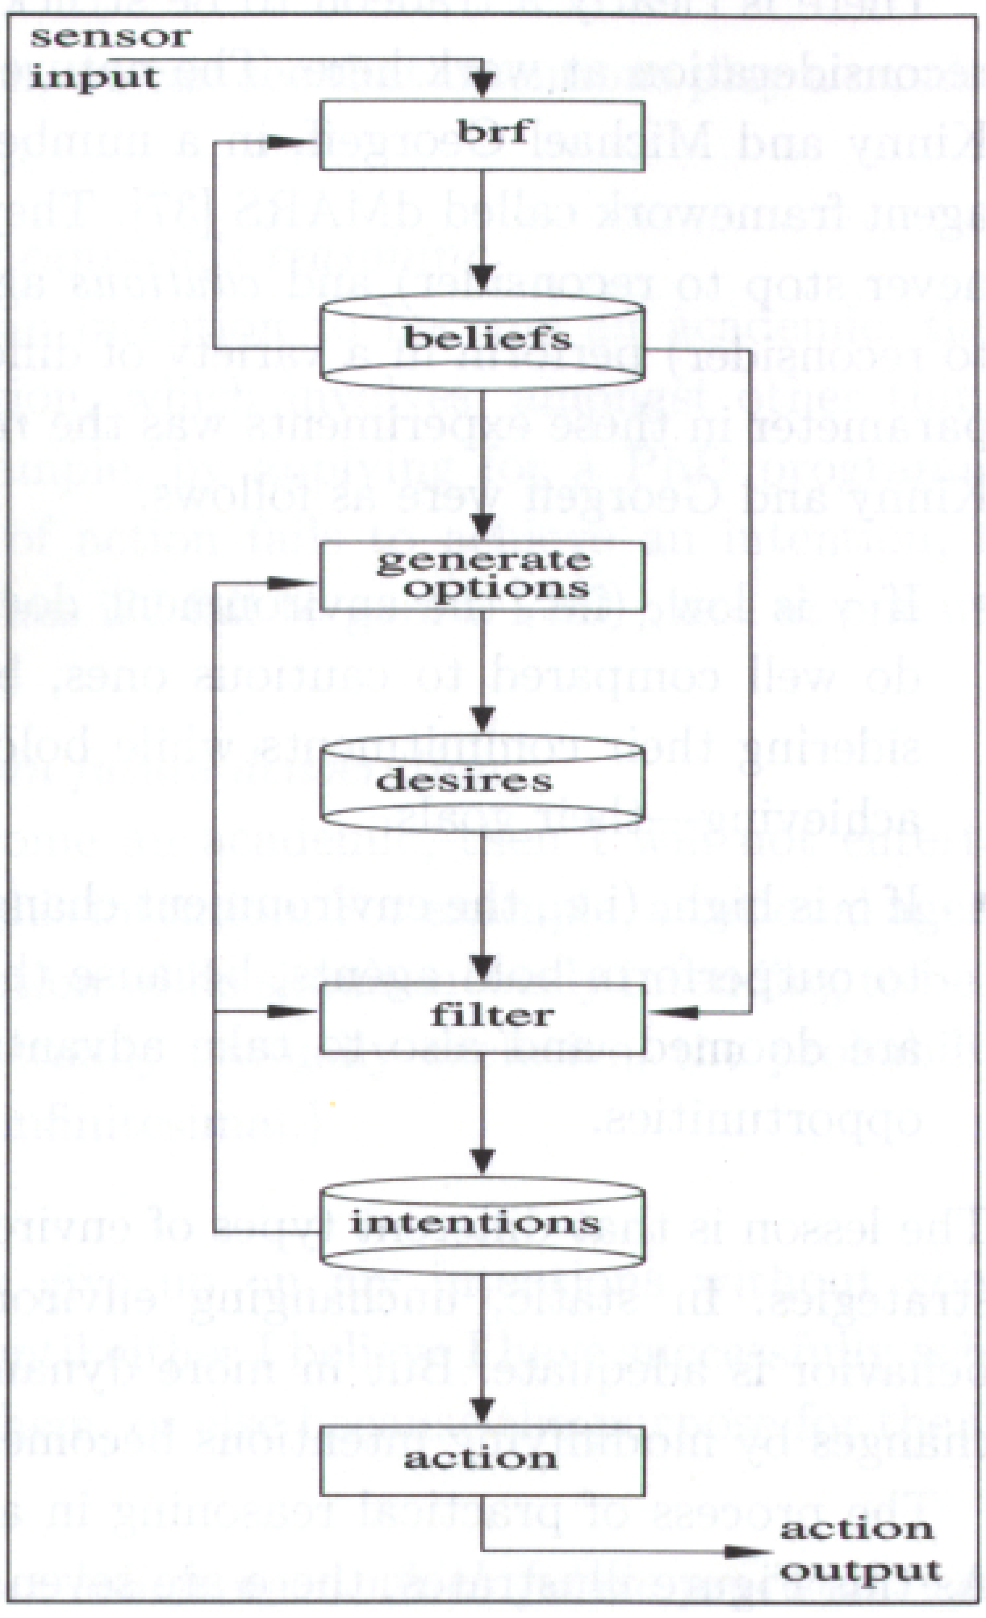
sensor input
brf
beliefs
generate options
desires
filter
intentions
action
action output

```
function action(p : P) : A
begin
        B := brf(B, p)
        D := options(D, I)
        I := filter(B, D, I)
        return execute(I)
end function action
```

$$(B, D, I) \text{ where } B \subseteq \mathbf{Bel}, D \subseteq \mathbf{Des}, I \subseteq \mathbf{Int}$$

I can be represented as a stack (priorities are available)

- BDI dates back to (Bratman, Israel, and Pollack 1988).
- PRS (*procedural reasoning system*, (Georgeff and Lansky 1987)) uses BDI. Applications: Space Shuttle (Diagnosis), Sydney Airport (air traffic control).
- BDI-Logics: (Rao and Georgeff 1991; Rao and Georgeff 1995; Rao 1995).

2.3 Layered Architectures

At least 2 layers: reactive (event-driven), pro-active (goal directed).

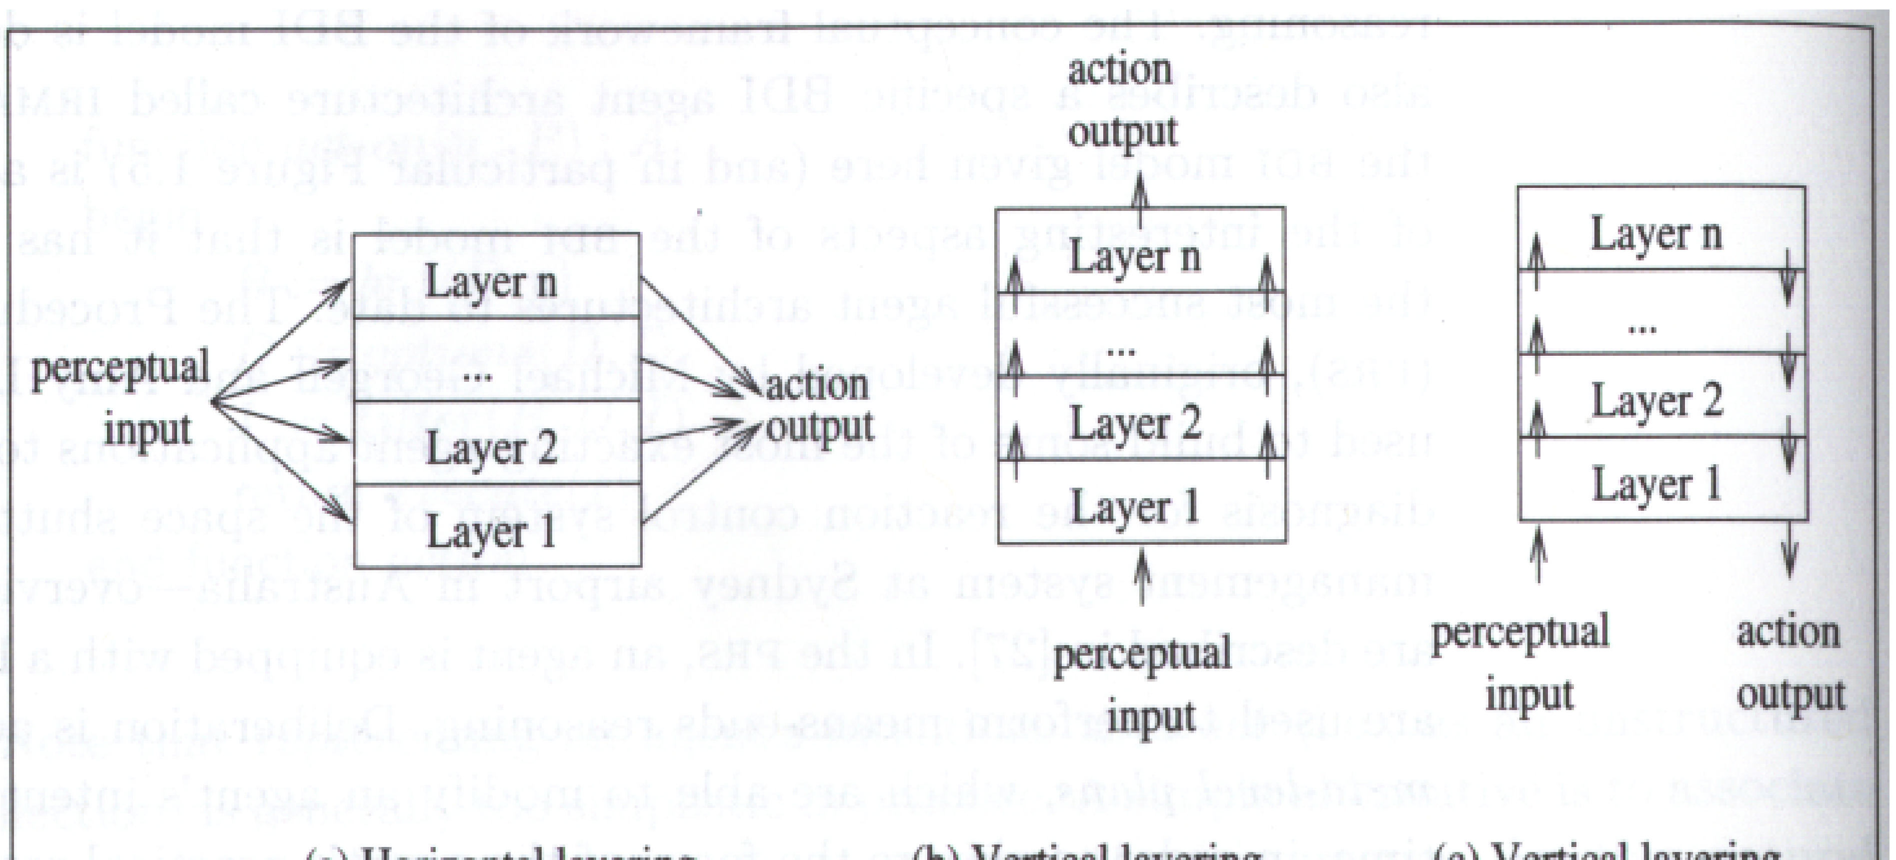

(a) Horizontal layering (b) Vertical layering (One pass control) (c) Vertical layering (Two pass control)

Horizontal:

- simpel (n behaviours, n layers),
- overall behaviour might be inconsistent,
- Interaction between layers: m^n (m = # actions per layer)
- Control-system is needed.

Vertical:

- Only $m^2(n-1)$ interactions between layers.
- Not fault tolerant: If one layer fails, everything brakes down.

Touring Machine

Autonomous Vehicle.

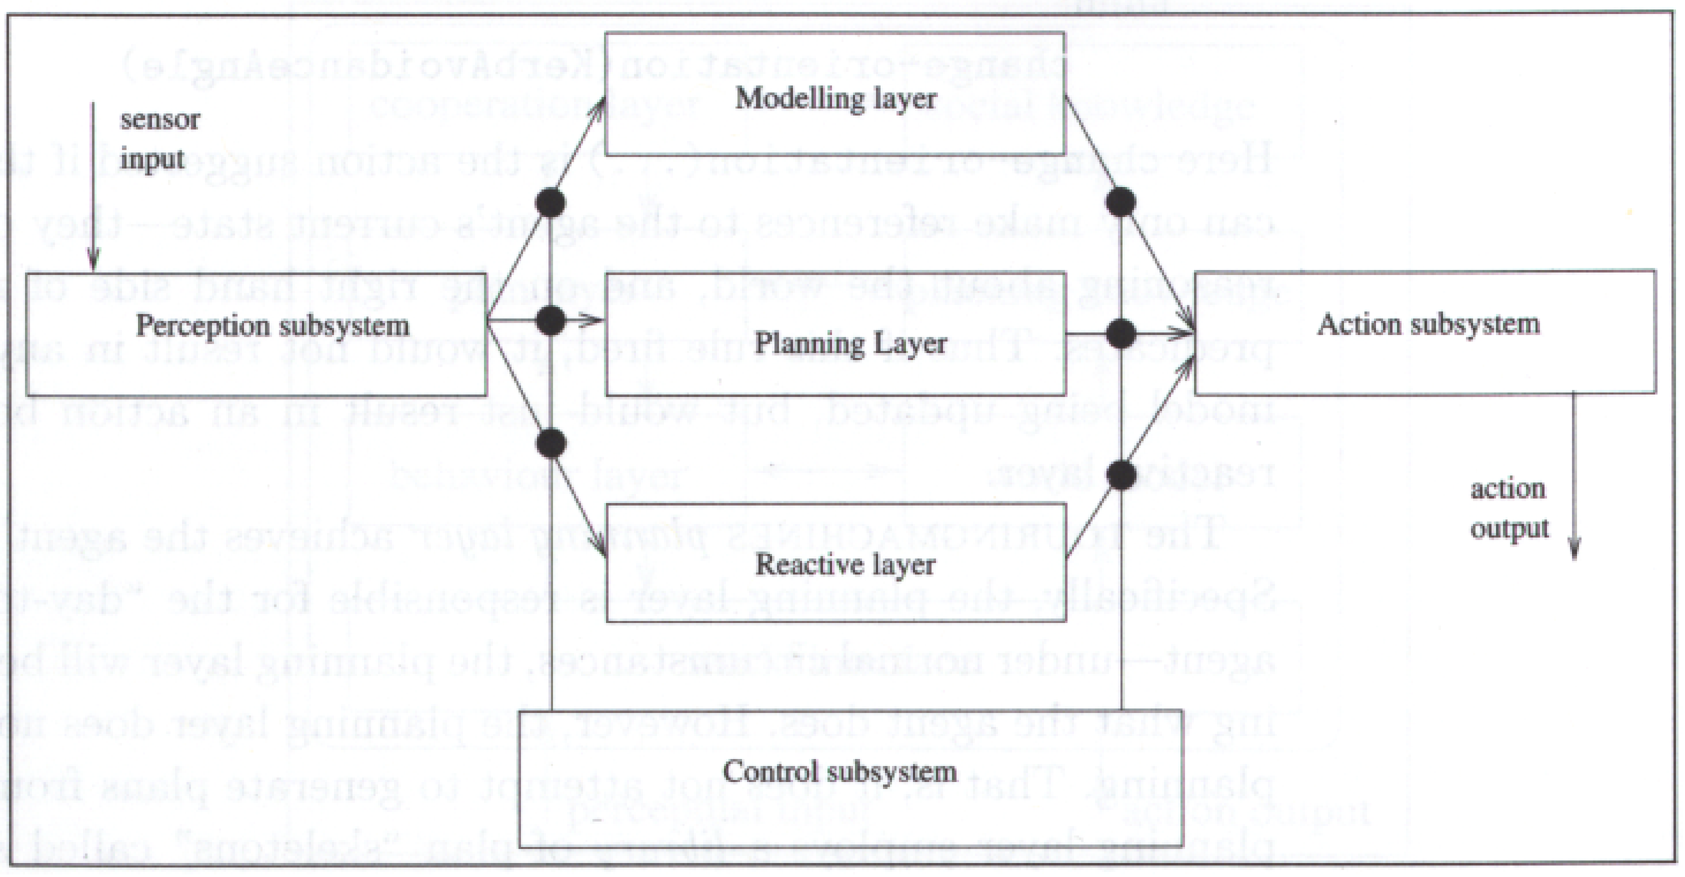

Rule 1: *Avoid curb*

$$
\begin{array}{ll}
\underline{\text{if}} & is\_in\_front(curb, observer) \text{ and} \\
 & speed(observer) > 0 \text{ and} \\
 & seperation(curb, observer) < curb\_threshold \\
\underline{\text{then}} & change\_orientation(curb\_avoidance\_angle)
\end{array}
$$

Planning-Layer: Pro-active behaviour

Modeling Layer: updating of the world, beliefs, predicts conflicts between agents, changes planning-goals

Control-subsystem: Decides about who is active. Certain observations should never reach certain layers.

Layered architectures do not have a clear semantics and the horizontal interaction is diffcult.

References

Arisha, K., F. Ozcan, R. Ross, V. S. Subrahmanian, T. Eiter, and S. Kraus (1999, March/April). IMPACT: A Platform for Collaborating Agents. *IEEE Intelligent Systems 14*, 64–72.

Bratman, M., D. Israel, and M. Pollack (1988). Plans and Resource-Bounded Practical Reasoning. *Computational Intelligence 4*(4), 349–355.

Dix, J., S. Kraus, and V. Subrahmanian (2001). Temporal agent reasoning. *Artificial Intelligence to appear*.

Dix, J., M. Nanni, and V. S. Subrahmanian (2000). Probabilistic agent reasoning. *Transactions of Computational Logic 1*(2).

Dix, J., V. S. Subrahmanian, and G. Pick (2000). Meta Agent Programs. *Journal of Logic Programming 46*(1-2), 1–60.

Eiter, T., V. Subrahmanian, and G. Pick (1999). Heterogeneous Active Agents, I: Semantics. *Artificial Intelligence 108*(1-2), 179–255.

Eiter, T., V. Subrahmanian, and T. J. Rogers (2000). Heterogeneous Active Agents, III: Polynomially Implementable Agents. *Artificial Intelligence 117*(1), 107–167.

Eiter, T. and V. S. Subrahmanian (1999). Heterogeneous Active Agents, II: Algorithms and Complexity. *Artificial Intelligence 108*(1-2), 257–307.

Georgeff, M. and A. Lansky (1987). Reactive Reasoning and Planning. In *Proceedings of the Conference of the American Association of Artificial Intelligence*, Seattle, WA, pp. 677–682.

Rao, A. S. (1995). Decision Procedures for Propositional Linear-Time Belief-Desire-Intention Logics. In M. Wooldridge, J. Müller, and M. Tambe (Eds.), *Intelligent Agents II – Proceedings of the 1995 Workshop on Agent Theories, Architectures and Languages (ATAL-95)*, Volume 890 of *LNAI*, pp. 1–39. Berlin, Germany: Springer-Verlag.

Rao, A. S. and M. Georgeff (1991). Modeling Rational Agents within a BDI-Architecture. In J. F. Allen, R. Fikes, and E. Sandewall (Eds.), *Proceedings of the International Conference on Knowledge Representation and Reasoning*, Cambridge, MA, pp. 473–484. Morgan Kaufmann.

Rao, A. S. and M. Georgeff (1995, June). Formal models and decision procedures for multi-agent systems. Technical Report 61, Australian Artificial Intelligence Institute, Melbourne.

Subrahmanian, V., P. Bonatti, J. Dix, T. Eiter, S. Kraus, F. Özcan, and R. Ross (2000). *Heterogenous Active Agents*. MIT-Press.

Weiss, G. (Ed.) (1999). *Multiagent Systems*. MIT-Press.