# Hybrid Deliberative/Reactive Systems

March 27, 2007

Georgia Tech Robots used in Hybrid Systems Research


Robby robot at JPL

# Today's Objectives

- To understand the limitations of purely reactive and purely deliberative methods when each is considered in isolation
- To recognize the issues in establishing interfaces between reactive control and deliberative planners and to express several models for these interfaces
- To study several representative hybrid architectures

# “Philosophy” Behind Hybridization

“In preparing for battle I have always found that plans are useless, but planning is indispensable.” Dwight Eisenhower

“It is a bad plan that admits of no modification.” Publilius Syrus

“Everybody’s got plans … until they get hit.” Mike Tyson

“Few people think more than two or three times a year; I have made an international reputation for myself by thinking once or twice a week.”

George Bernard Shaw

# Why Hybridize?

- Behavior-Based/Reactive systems:
  - Pro:
    - Shown to effectively produce robust performance in complex and dynamic domains
  - Con: But, strong assumptions of reactive systems may not always be true:
    - The environment lacks temporal consistency and stability
    - The robot's immediate sensing is adequate for the task at hand
    - It is difficult to localize a robot relative to a world model
    - Symbolic representational world knowledge is of little or no value

# Without Representation, Reactive Systems are Limited

- Purely reactive robot can't:
  - Plan optimal trajectories
  - Make maps
  - Monitor its own performance
  - Select best behaviors to accomplish a task
- Also:
  - Design of "emergent" behavior is more of an art than a science
- But, consensus that behavior-based/robotic control is best for low-level control because of:
  - Pragmatic success
  - Elegance as a computational theory for both biological and machine intelligence

# Deliberative Systems Sometimes Preferred …

… when:

- World can be accurately modeled
- Uncertainty is restricted
- Some guarantee exists of virtually no change in the world during execution

- But, real world of biological agents isn't usually described in this way

# Introduction of Knowledge
# Makes for More Flexible/General Approach

- Recall Representational Issues:
  - Behavioral and perceptual strategies can be represented as modules and configured to match various missions and environments, adding versatility.
  - A priori world knowledge, when available and stable, can be used to configure or reconfigure these behaviors efficiently.
  - Dynamically acquired world models can be used to prevent certain pitfalls to which non-representational methods are subject.

# Hybrid Deliberative/Reactive Architectures

- Combine aspects of traditional AI symbolic methods and their use of abstract representational knowledge
- But, maintain goal of providing responsiveness, robustness, and flexibility of purely reactive systems
- Permit reconfiguration of reactive control systems based on available world knowledge through their ability to reason over the underlying behavioral components.
- Dynamic control system reconfiguration based on deliberation is an important addition to the overall competence of general purpose robots.

# Current Research Consensus about Hybrid Architectures

- Best general architecture solution because:
  - Use of asynchronous processing techniques (multi-tasking, threads, etc) allow deliberative functions to execute independently of reactive behaviors
  - Good software modularity allows subsystems or objects in Hybrid architectures to be mixed and matched for specific applications

# Recall Hybrid Architecture Paradigm

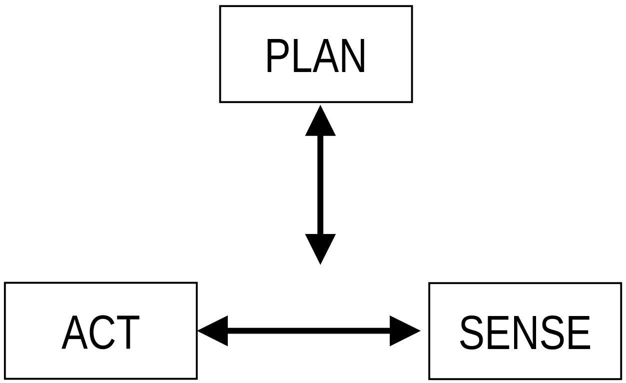


- But note:
  - Building a hybrid system requires compromise from both ends of the robotic spectrum (some say "the worst of both worlds"!)
  - Nature of boundary between deliberation and reactive execution is not well understood at this time, leading to somewhat arbitrary architectural decisions
- Central issue: how to develop a unifying architectural methodology that will ensure a system capable of robust robotic plan execution while taking into account a high-level understanding of the nature of the world and a model of user intent.

# Hybrid Architecture, con't.

- Planning:
  - Covers long time horizon and requires global knowledge
  - Computationally intensive; therefore, should be decoupled from real-time execution
  - Good for setting objectives and selecting methods
  - Not good for making finely grained decisions

- Sensing:
  - Complex organization
  - Hybrid, in that it is available for both reactive and planning levels
  - Perceptual schemas used for behaviors:  can be shared with world-modeling processes

# A Note on Biological Basis of Hybrid Systems

- Experimental evidence for hybrid systems is compelling
- Two distinct modes of behavior:
  - Willed: interface between deliberate conscious control and automatic system
  - Automatic: no awareness, starts without attention, consists of multiple independent parallel activity threads (schemas)
- Normal-Shallice model (1986) for human behavior notes connections between deliberate and automatic control:
  - Automatic schemas modulated by attention arising from deliberate control
  - Schemas (behavioral tasks) compete with each other
  - Schema selection is deliberate control's principle function
  - Neuropsychological experiments are consistent with this model

# Model for Integrated Automatic and Willed Behavior

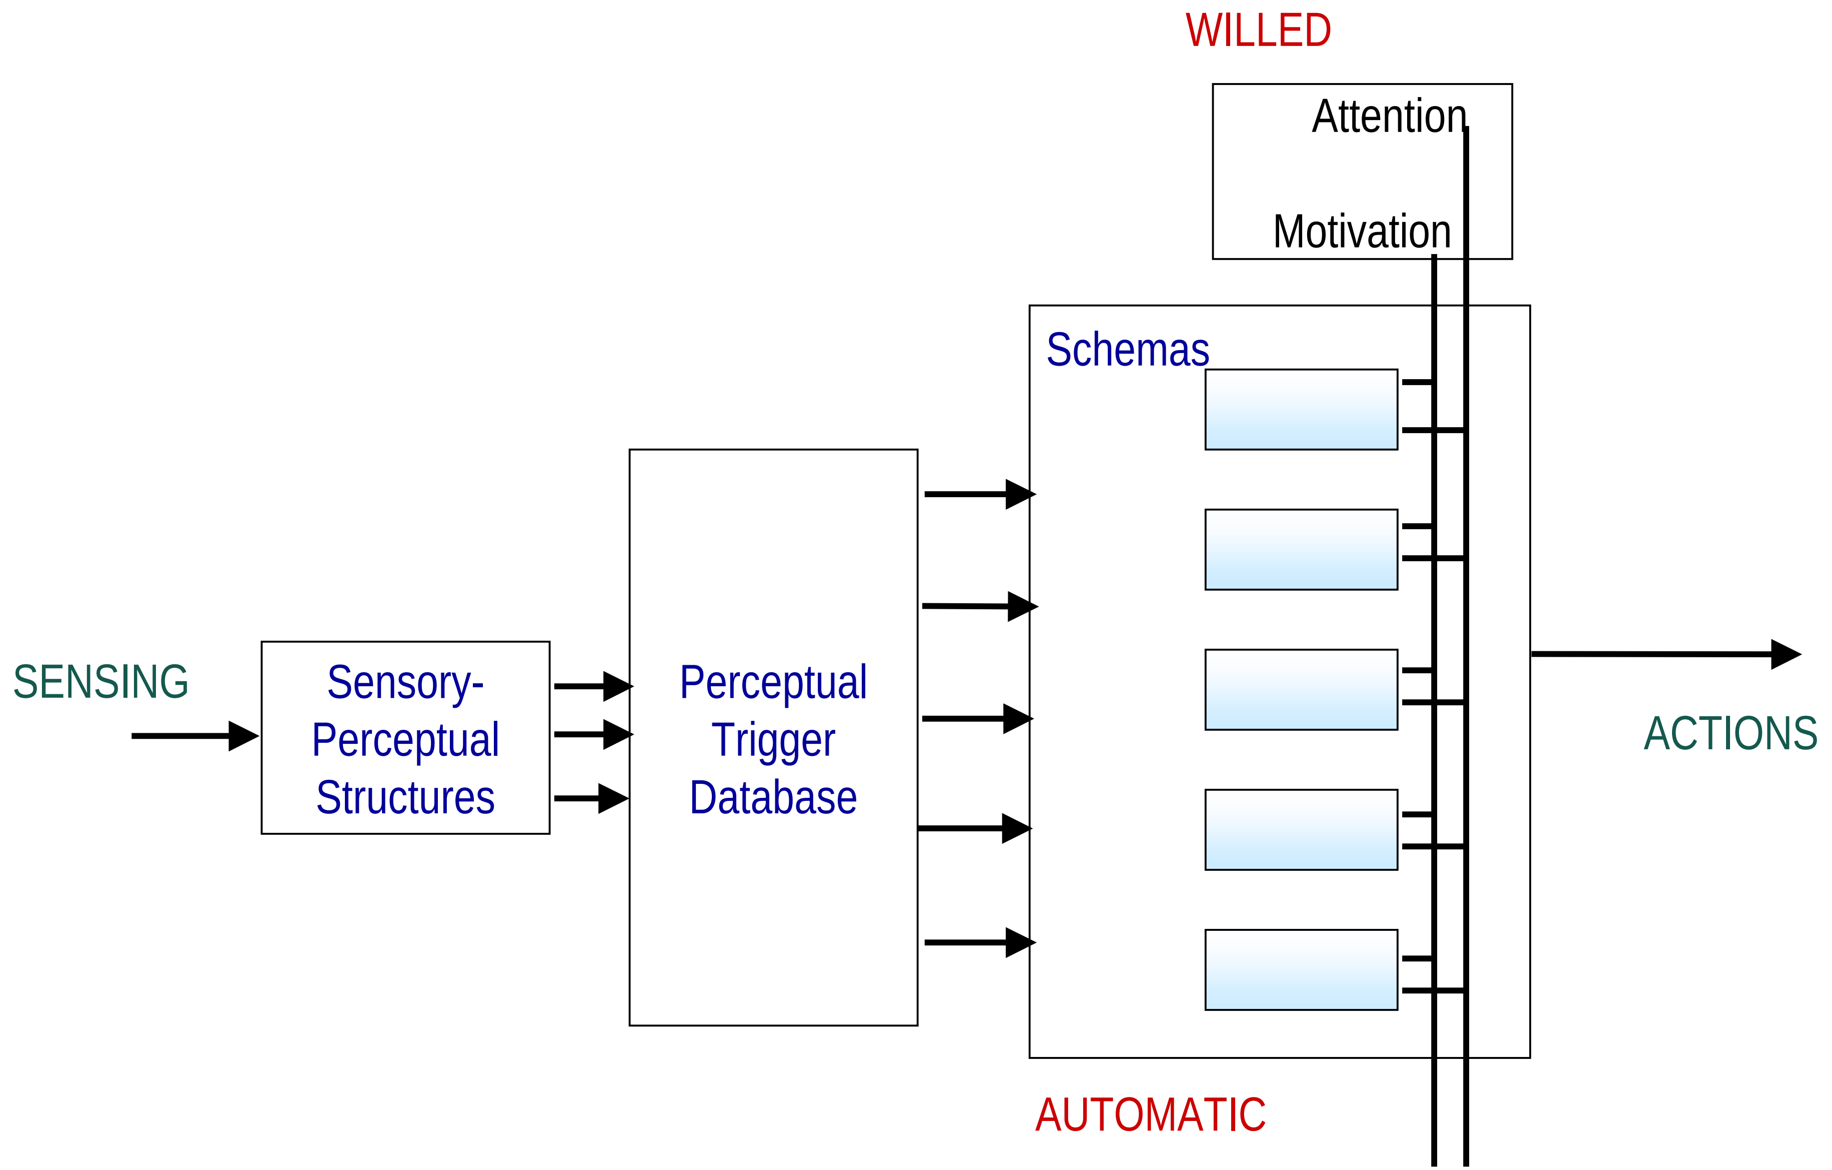

# Two Parts to Hybrid Architecture

- Two parts are named differently by different researchers:
  - Reactive and Deliberative (our choice)
  - Reactor and Deliberator
  - Reactor and Planner
  - Subcognitive and Cognitive

# Three ways for Tying Planning and Reaction

1. Hierarchical integration:
   - Deliberative planning and reactive execution are involved with different activities, time scales, and spatial scope
   - Planning vs. reacting depends on situation at hand
   - Closely related to “hierarchical paradigm”, EXCEPT nature and type of knowledge and reasoning is distinctly different in the layers of a deliberative architecture

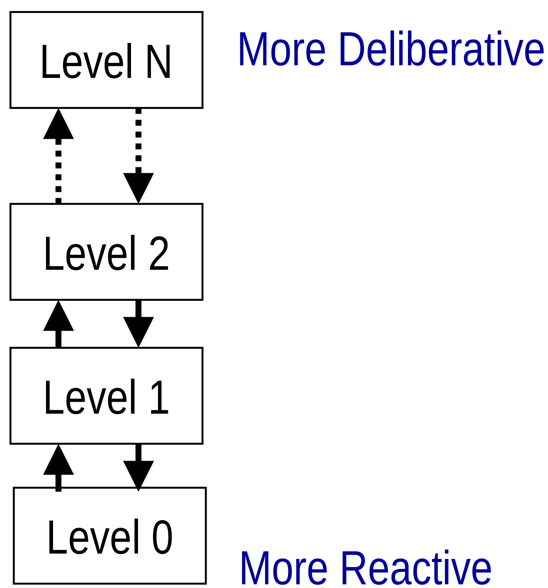

# Three ways for Tying Planning and Reaction (con't.)

2. Planning to guide reaction:
   - Planning configures and sets parameters for the reactive control system
   - Execution occurs only due to reactive system
   - Planning occurs prior to and concurrent to reactive system

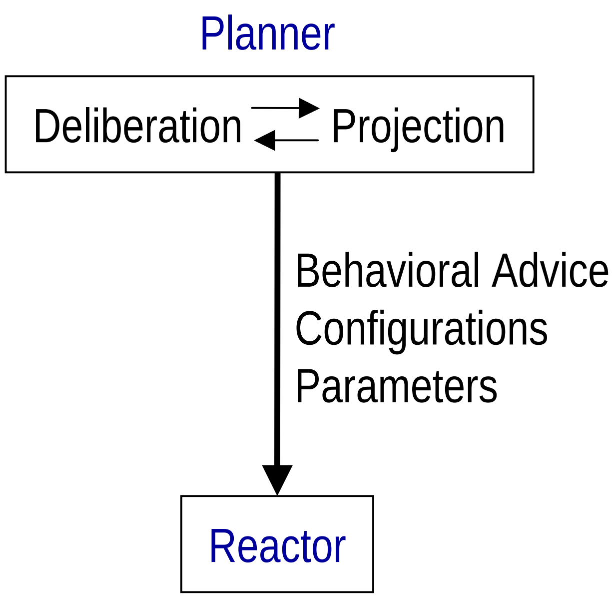

# Three ways for Tying Planning and Reaction (con't.)

3. Coupled planning-reacting:
   – Planning and reacting are concurrent activities, each guiding the other

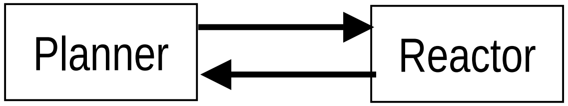

# Behaviors in Hybrid Architectures

- “Behaviors” aren’t identical in reactive/behavior-based systems and hybrid systems

- In hybrid systems:
  - Behaviors include reflexive, innate, and learned behaviors (or skills)
  - Tend to have use of assemblages of behaviors sequenced over time
  - Typically have more diversity in methods for combining output from concurrent behaviors

# Several Hybrid Approaches Have Been Developed

- AuRA (Arkin 1986)
- Atlantis (Gat 1991)
- Sensor-Fusion Effects (SFX) (Murphy 1996)
- 3-Tiered (3T) (JPL1990s)
- Saphira (Konolige 1998)
- Tack Control Architecture (Simmons 1997)
- Planner-Reactor (Lyons and Hendriks 1992)
- Procedural Reasoning System (PRS) (Georgeff and Lansky 1987)
- SSS (Connell 1992)
- Multi-Valued Logic (Saffiotti 1995)
- SOMASS Hybrid Assembly System (Malcom and Smithers 1990)
- Agent Architecture (Hayes-Roth 1993)
- Etc., Etc.

# How to Compare/Contrast Hybrid Architectures?

- Differences in:
  - How architecture distinguishes between reaction and deliberation
  - How architecture organizes responsibilities in the deliberative portion
  - How overall behavior emerges

# Common Components

- Sequencer: agent that generates set of behaviors to use in order to accomplish the subtask, and determines any sequences and activation conditions
- Resource manager: allocates resources to behaviors, including selecting from libraries of schemas (e.g., determining whether to use sonar or laser in a perceptual schema, depending upon the task)
- Cartographer: responsible for creating, storing, and maintaining a map or spatial information, plus methods for accessing data. Contains global world model and knowledge representation.
- Mission Planner: interacts with human, operationalizes commands into robot terms, and constructs mission plan.
- Performance monitoring and problem solving agent: allows robot to notice if it is making progress or not.

# Four Principle Interface Strategies in Hybrid Systems

- Selection: Planning is viewed as configuration
  - Example architecture: AuRA
- Advising: Planning is viewed as advice giving
  - Example architecture: Atlantis
- Adaptation: Planning is viewed as adaptation
  - Example architecture: Planner-Reactor
- Postponing: Planning is viewed as a least commitment process.
  - Example architecture: Procedural Reasoning System (PRS)

# Case Study of Hybrid Architecture: AuRA (Autonomous Robot Architecture)

- Arkin (1986): One of first to advocate use of hybrid deliberative and reactive control systems, through use of AuRA:
  - Deliberative: traditional AI-based planner
  - Reactive: based on schema theory
- AuRA: first robot navigational system to be presented in this integrated manner
- Major planning and execution components:
  - Hierarchical system consisting of mission planner, spatial reasoner, and plan sequencer coupled with reactive system (schema controller)

# Original AuRA Architecture

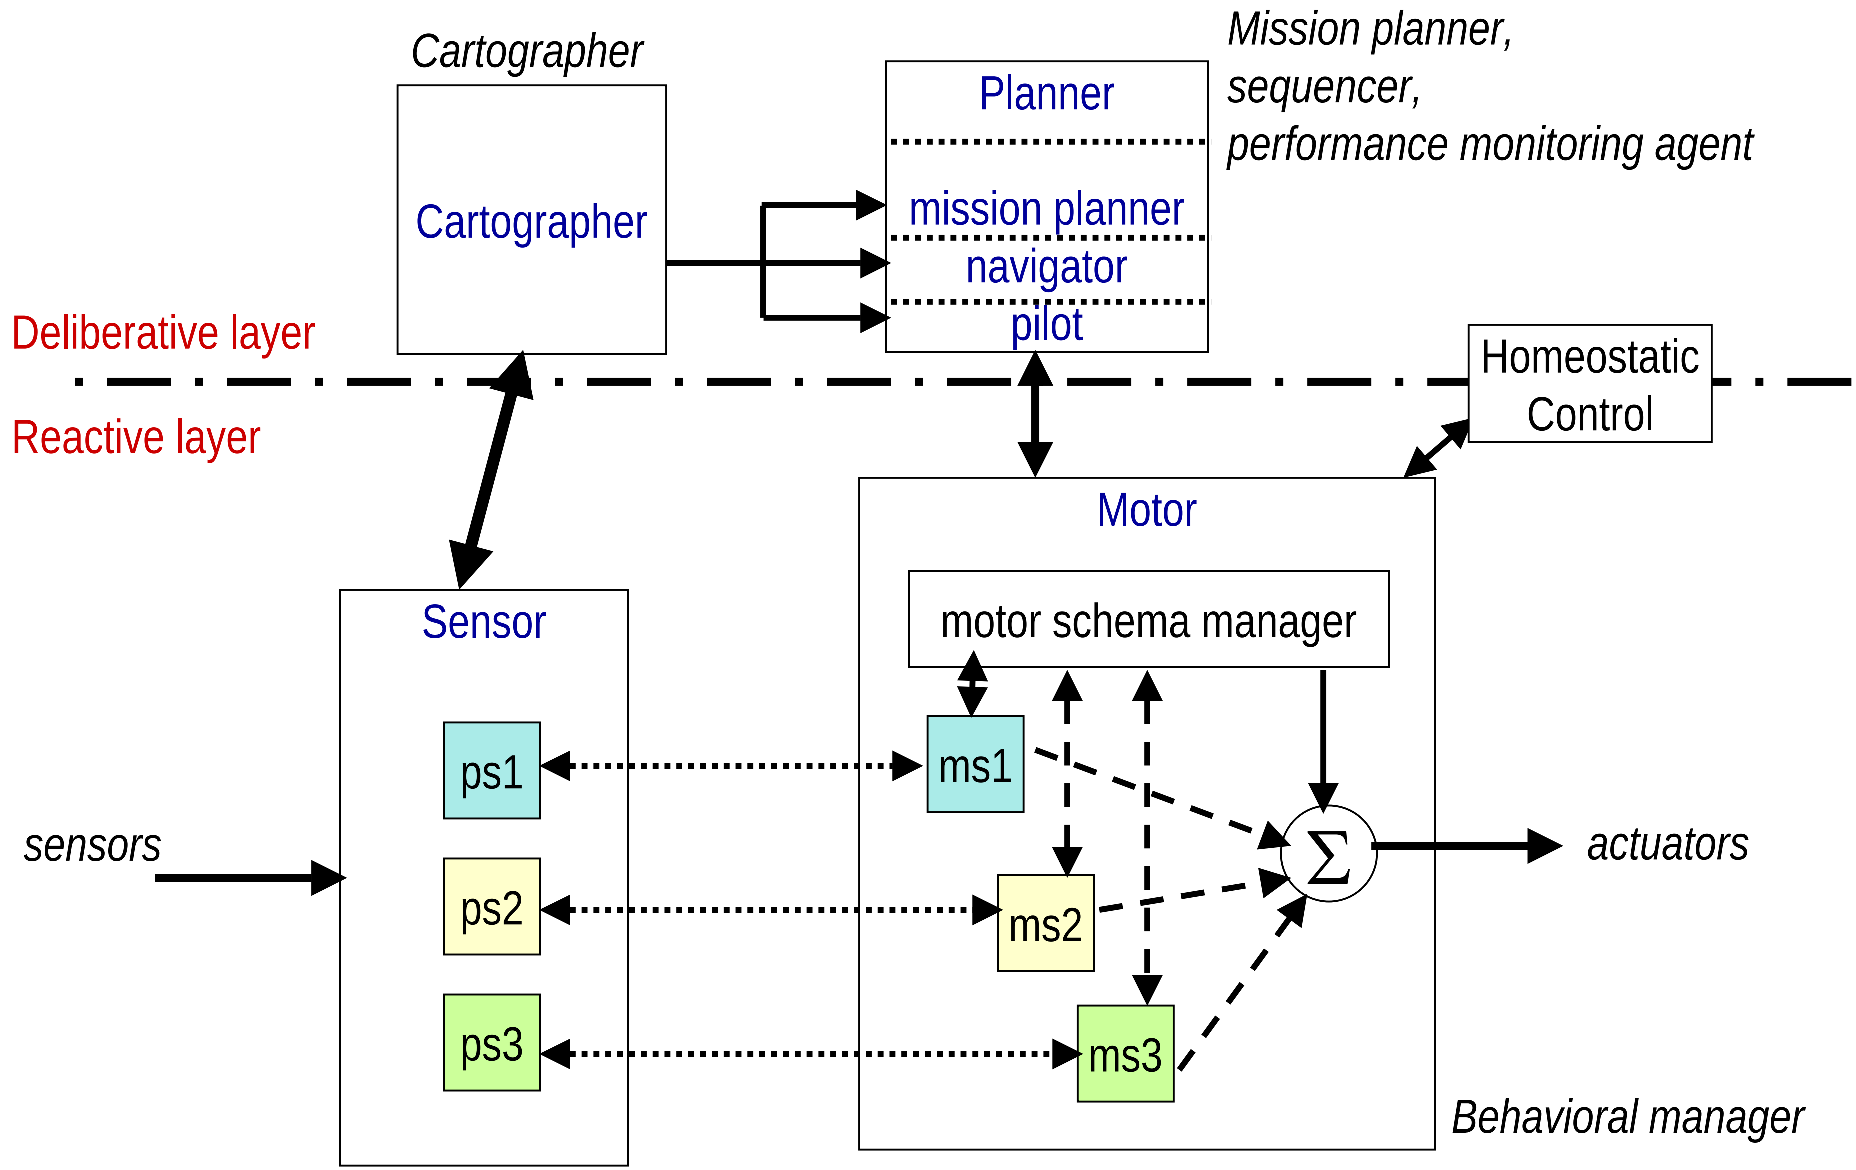

# More Recent AuRA Architecture

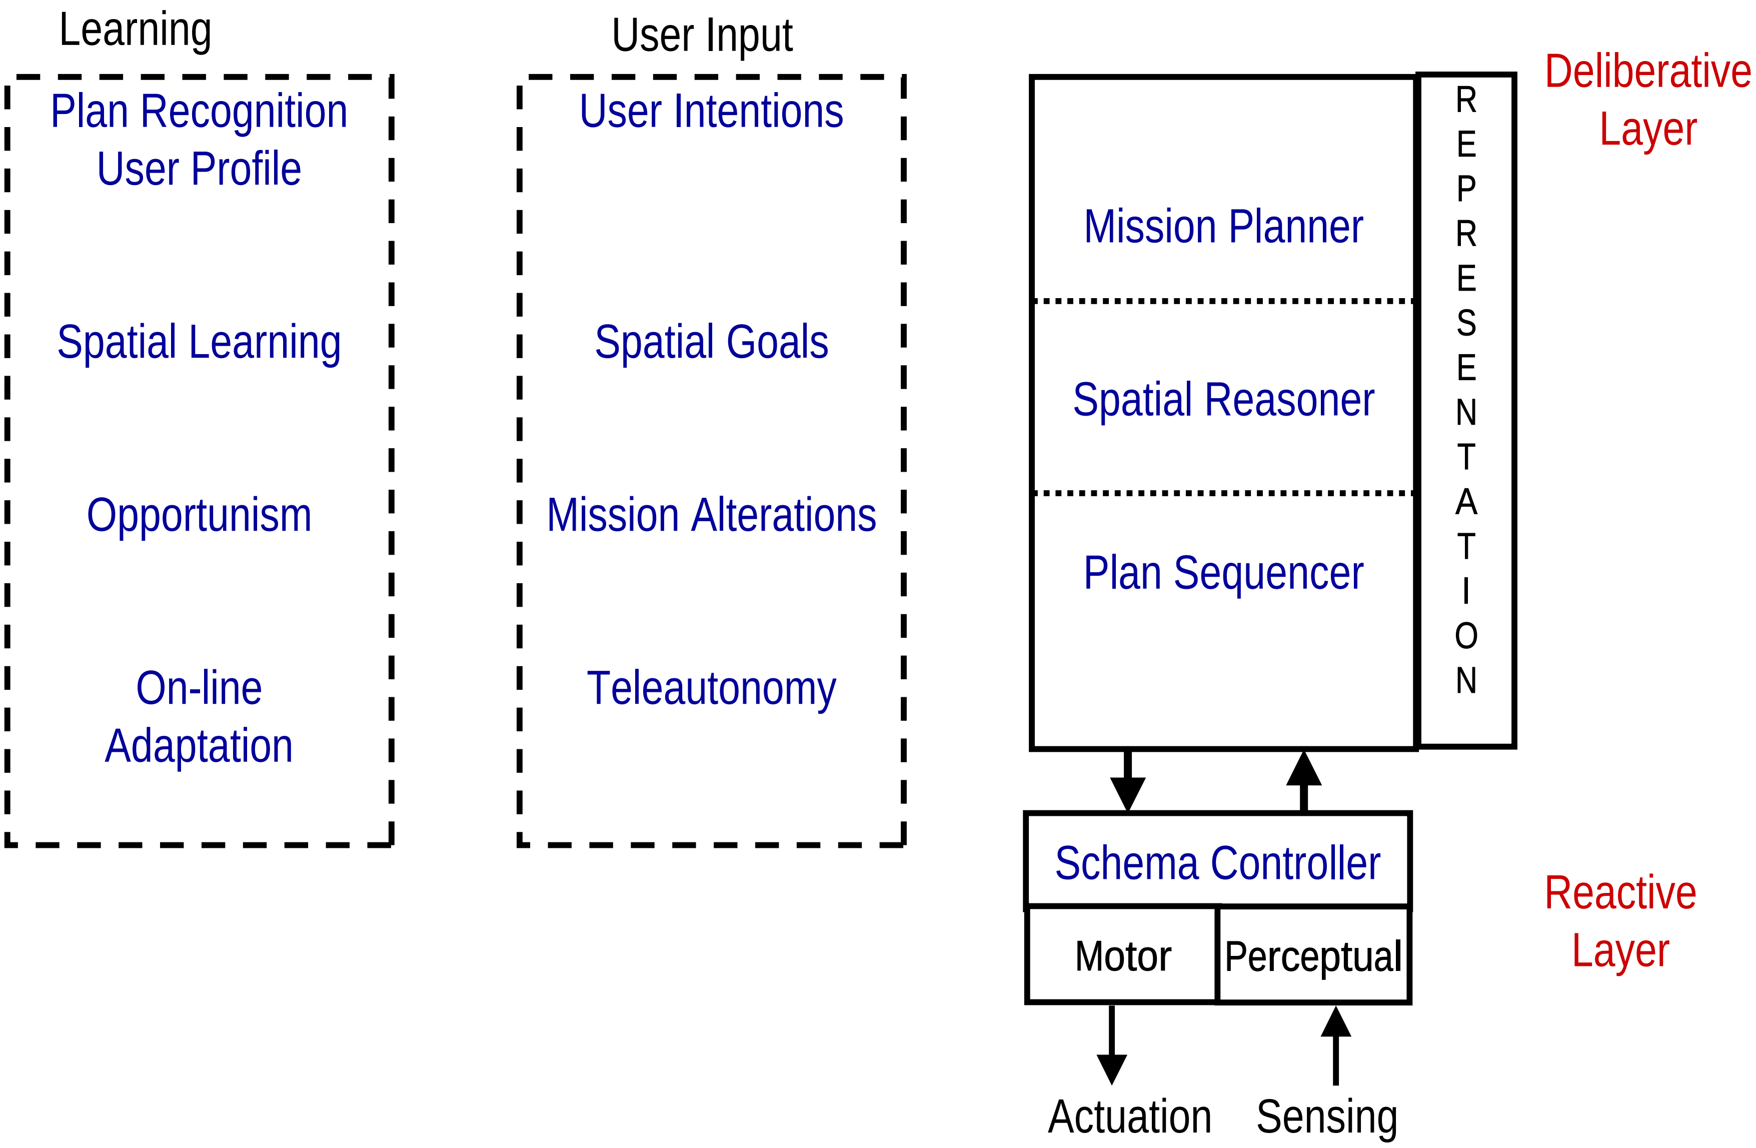

# Advantage of AuRA: Replacing Components

- In research, helpful to test out various approaches to sub-components
- Examples of actual component testing:
  - Specialized mission planner developed for assembly
  - Spatial reasoner can vary (A*, Router, etc.)
  - Perceptual schemas can include sensor fusion
  - Plan sequencer: can be rule-based plan or FSA temporal sequencer

# Implementations of AuRA

- Manufacturing environments
- 3D navigation (e.g., aerial or underwater)
- Indoor and outdoor navigation
- Robot competitions
- Vacuuming
- Military scenarios
- Mobile manipulation
- Multi-robot teams

Georgia Tech Robots used in Hybrid Systems Research

# Another Case Study: Atlantis (Gat 1991)

- Three-level hybrid system
- Incorporates:
  - Deliberator that handles planning and world modeling
  - Sequencer that handles initiation and termination of low-level activities, addressing reactive-system failures to complete the task
  - Reactive controller that manages collections of primitive activities
- Architecture is both asynchronous and heterogeneous
  - No layers are in charge of others
  - Activity is spread throughout architecture
- Planning viewed as advice-giving

# Atlantis Architecture

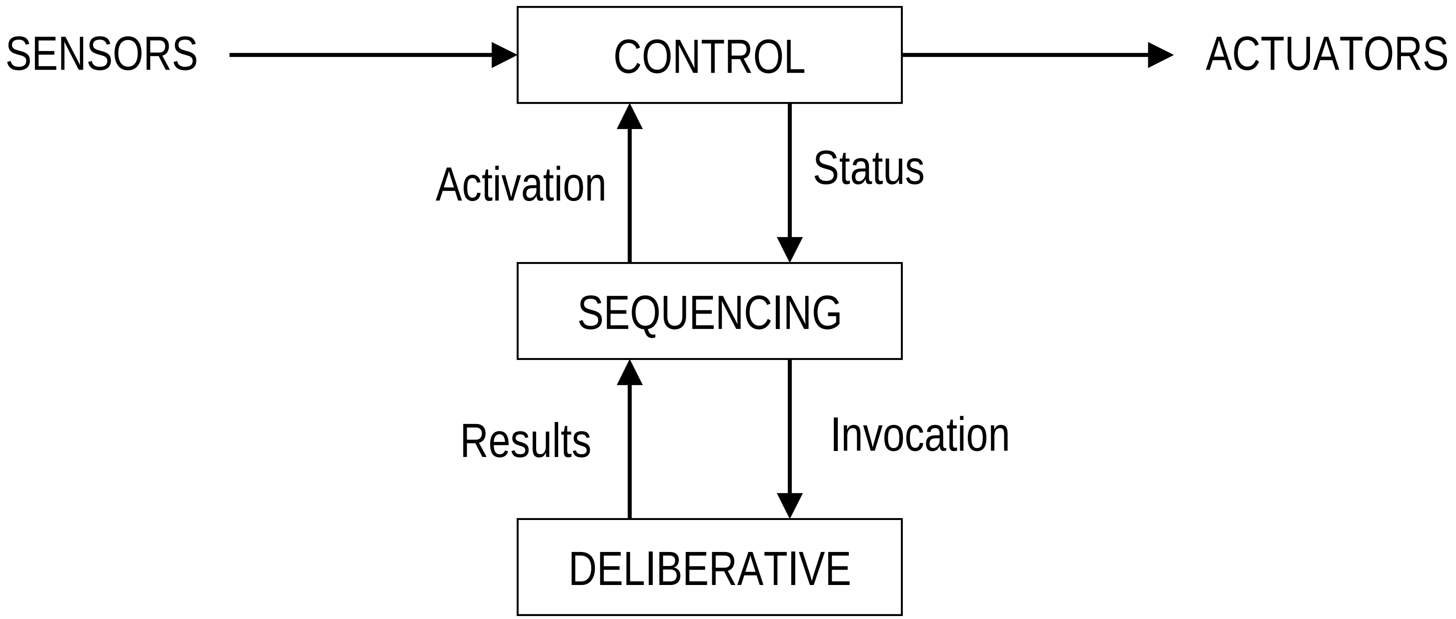

# Atlantis Architecture (con't.)

- Control layer implemented in ALFA – LISP-based programming language used to program reactive modules configured in networks connected via communication channels
- Sequencing layer:
  - Conditional sequencing occurs upon completion of subtasks or detection of failure
  - Cognizant failure: robot's ability to recognize on its own when it has not or cannot complete its tasks
  - Monitor routines added to make this determination
- Deliberation:
  - Occurs at sequencing layer's request
  - Consists of traditional LISP-based AI planning algorithms specific to problem at hand
  - Planner's output is treated as advice

# Design in Atlantis

- Design proceeds from bottom up:
  - Low-level abilities capable of being executed in reactive controller level are built first
  - Sequences of primitive behaviors are then constructed for use within sequencing level
  - Deliberative methods to assist in the sequencer’s decisions are then designed
- Experiments performed on large outdoor JPL MARS rover
  - Complex navigational tasks in rough outdoor terrain

Robby robot at JPL

# Summary of Atlantis' Important Features

- Three-layered architecture consisting of controller, sequencer, and deliberator
- Asynchronous, heterogeneous reactivity and deliberation are used
- The results of deliberation are viewed as advice, not decree
- Classical AI is merged effectively with behavior-based reactive control methods
- Cognizant failures provide an opportunity for plan restructuring
- The system has been exercised successfully on both indoor and outdoor robotic systems

# Hybrid Architectures vs. Hierarchical Architectures: What is the Difference?

- Hybrids:
  - Explicitly reflect more principles of software engineering (modularity, coherence, design for reuse, etc.)
  - Global models used only for symbolic functions
  - Frame problem doesn't exist because:
    - Execution is reactive and therefore well-suited for unstructured environments
    - Or, software agents can use agent-specific abstractions to exploit structure of environment in order to fulfill their particular role in deliberation
  - Robot "thinks" in terms of a closed world, while it acts in an open world
  - Planning not performed for every move
  - Hybrid paradigm has roots in ethology/biology, and provides a framework for exploring cognitive science

# Planner-Reactor Architecture

- Developed by Lyons and Hendricks (1992)
- Philosophy: use of planner as mechanism to continuously modify an executing reactive control system
- Planner is in essence an execution monitor that adapts the underlying behavioral control system in light of the changing environment and the agent's underlying goals
- Planner's goal is to improve the performance of reactor at all times
- This is a form of anytime planning:
  - Providing approximate answers in a time-critical manner such that:
    - At any point a plan is available for execution
    - The quality of the plan as a whole increases over time

# Planner-Reactor Architecture

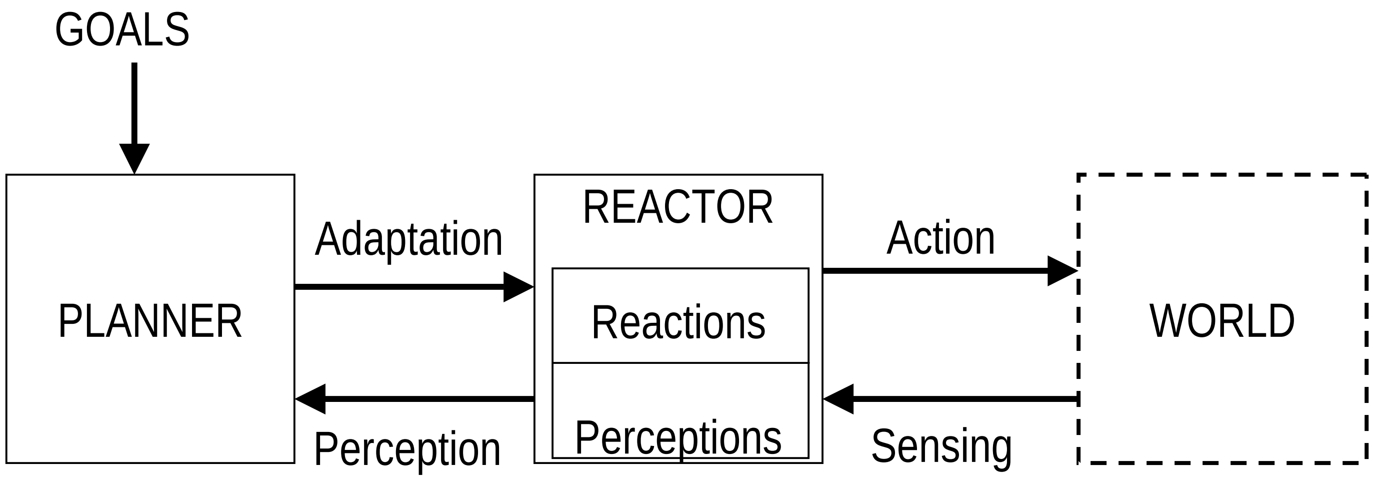

# Key Features of Planner-Reactor Methodology

- Deliberation and reactivity are integrated through asynchronous interaction of a planner and a concurrent reactive control system
- Planning is viewed as a form of reactor adaptation
- Adaptation is an on-line process rather than an off-line deliberation
- Planning is used to remove errors in performance when they occur
- The reactor undergoes situationally dependent on-line performance improvement
- The basic techniques, tested in assembly work cell tasks and grasp planning for a robotic hand, are believed to be applicable for a wide range of applications

# Procedural Reasoning System

- Reactivity refers to postponement of elaboration of plans until necessary (least commitment)
- Least commitment strategy:
  - Defers making a decision until it is absolutely necessary to do so. The information necessary to make a correct decision is assumed to become available late in the process, thus reducing need for backtracking.
- Plans are primary mode of expressing action
- Plans are continuously determined in reaction to current situation
- Previous plans can be abandoned at any time
- Representations of robot's beliefs, desires, and intentions are all used to formulate a plan
- Plan represents robot's desired behavior (as opposed to traditional AI planner's output of goal states)

# Procedural Reasoning System

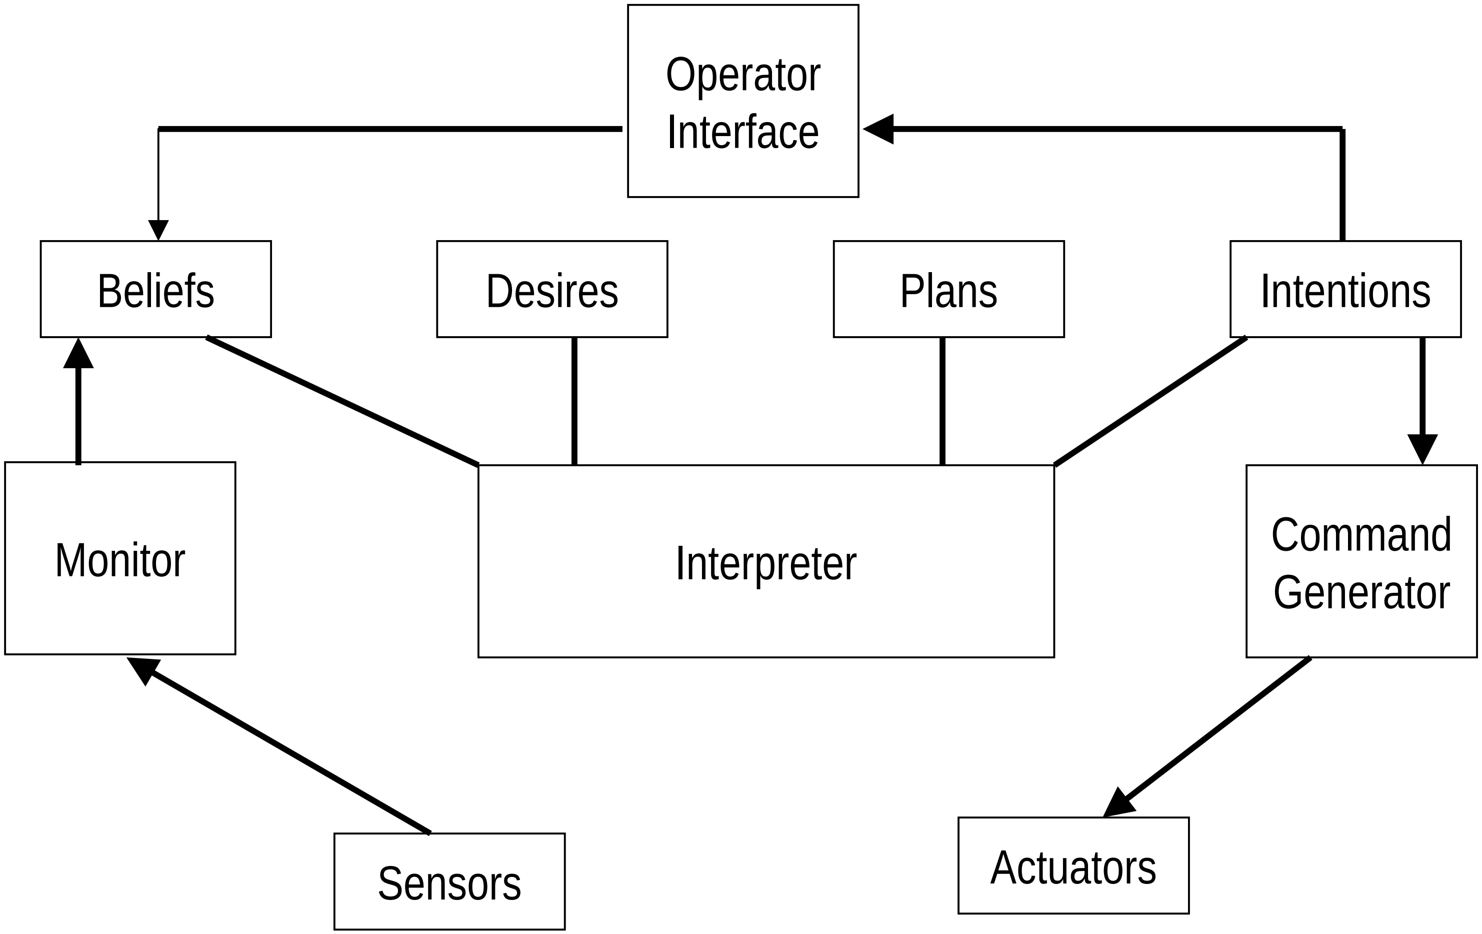

# Summary of Hybrid Systems

- Both deliberative planning systems and purely reactive control systems have limitations when each is considered in isolation
- Deliberative planning systems provide an entry point for the use of traditional AI methods and symbolic representational knowledge in a reactive robotic architecture
- The interface between deliberation and reaction is poorly understood and serves as the focus of research in this area
- Strong evidence exists that hybrid deliberative and behavior-based systems are found in biology
- Hybrid models include hierarchical integration, planning to guide reaction, and coupled planning and reacting

# Summary of Hybrid Systems (con't.)

- Common components of hybrid systems include a sequencer, a resource manager, a cartographer, a mission planner, and a performance monitoring and problem solving agent
- Four principle interface strategies include selection, advising, adaptation, and postponing
- AuRA is an early hybrid deliberative/reactive system using motor schemas and a traditional AI planner
- Atlantis is a three-layer hybrid architecture, introducing notion of cognizant failure. Plans are advice, not decree.
- Planner-Reactor architecture views planning as continuous adaptation
- PRS uses a least-commitment strategy to delay elaboration of plans for execution until necessary