Scalable Distributed Intelligence

Ricardo Sanz

Autonomous Systems Laboratory
Universidad Politécnica de Madrid

IEEE International Symposium on Intelligent Control
Houston, October 6-8, 2003

Motivation

- As plants grow in complexity and control requirements grow in difficulty, there is an increased need for intelligence in control systems.
- Intelligence, traditionally a duty for central brains, has to be available where it is needed and that means, today, everywhere.
- This is achieved putting computing power and software in many places far from central brains. Intelligence is getting ubiquitous in control systems.
- Controllers are becoming societies.

Summary

- In this talk I will address the implementation of complex controllers with distributed intelligence
- Describe some basic technologies and some of the trends
- Propose research topics for future pervasive, highly intelligent, scalable controllers

Contents

- A Fast Review
  - Distributed Control
  - Distributed Artificial Intelligence
  - Distributed Intelligent Control
- Some Trends and Speculations
  - Deeper Intelligence
  - Scalable Intelligence
  - Scalable Distributed Intelligence

Distributed Control

When the control pieces are taken apart

Distributed Plants

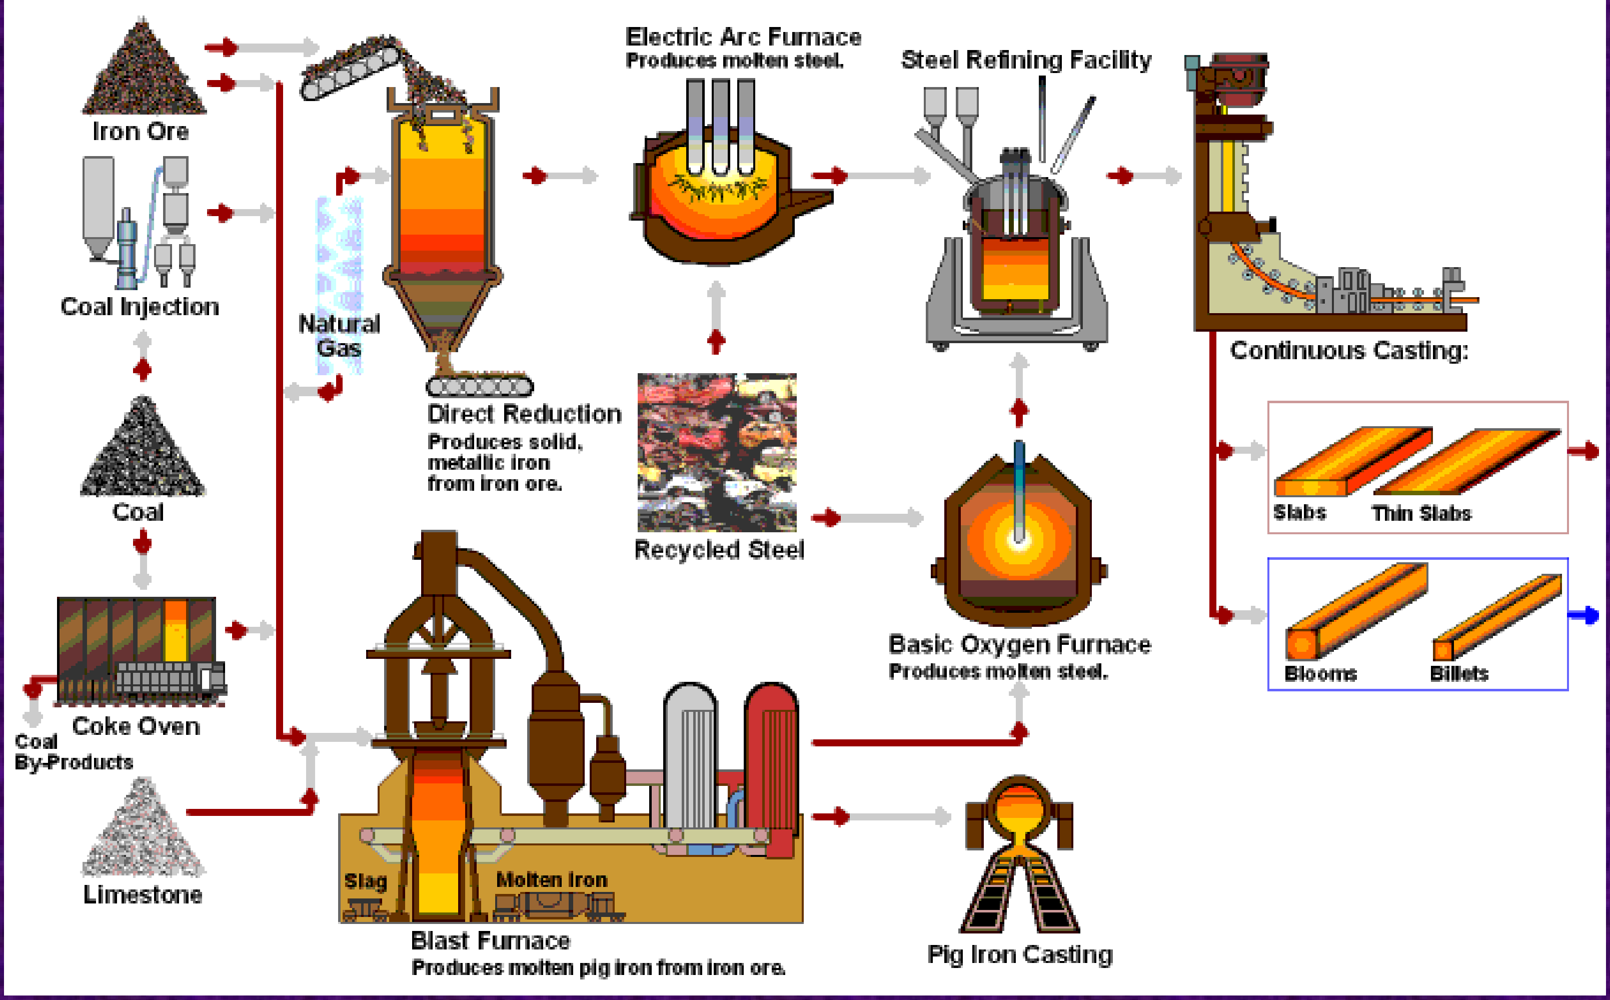

Even in small size

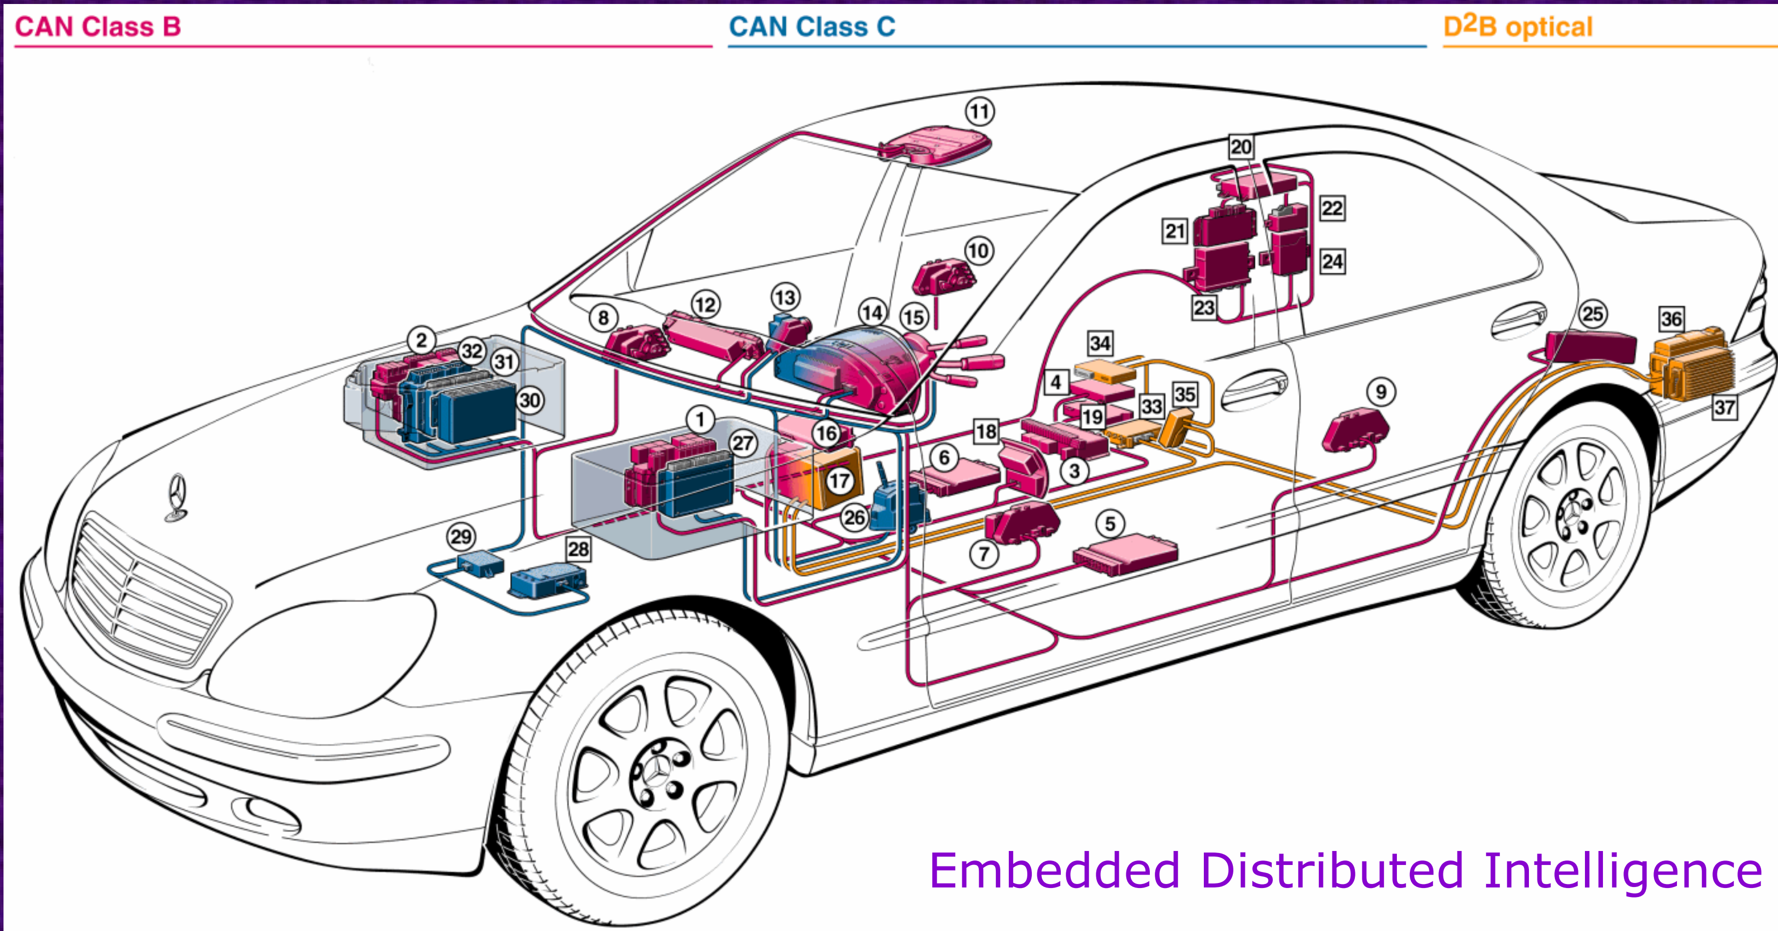

Embedded Distributed Intelligence

Trends

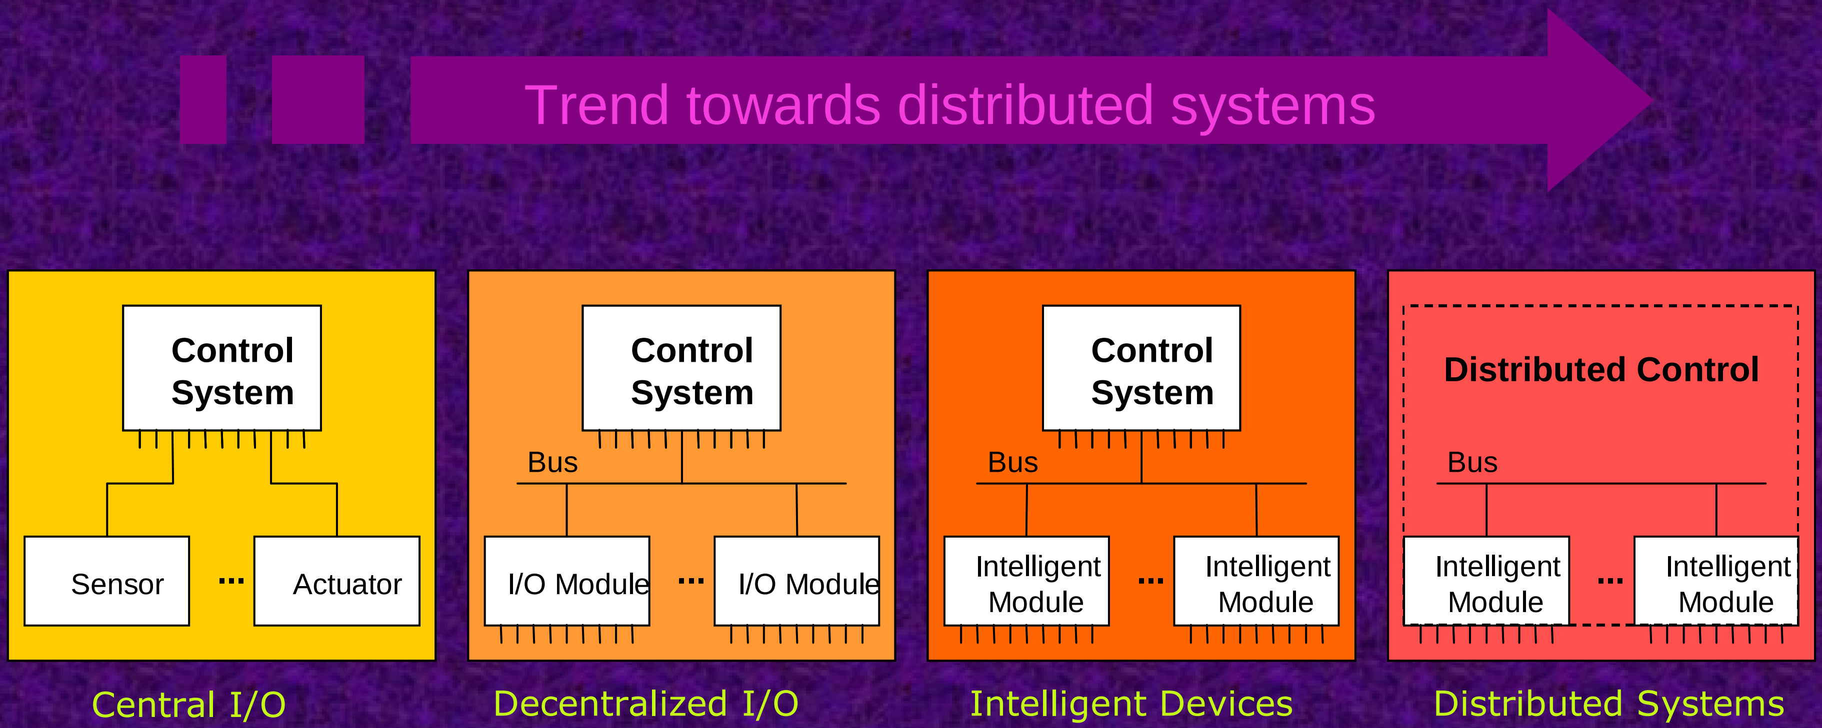
Trend towards distributed systems
Control System
Sensor
...
Actuator
Control System
Bus
I/O Module
...
I/O Module
Control System
Bus
Intelligent Module
...
Intelligent Module
Distributed Control
Bus
Intelligent Module
...
Intelligent Module
Central I/O
Decentralized I/O
Intelligent Devices
Distributed Systems

Plant-wide Control

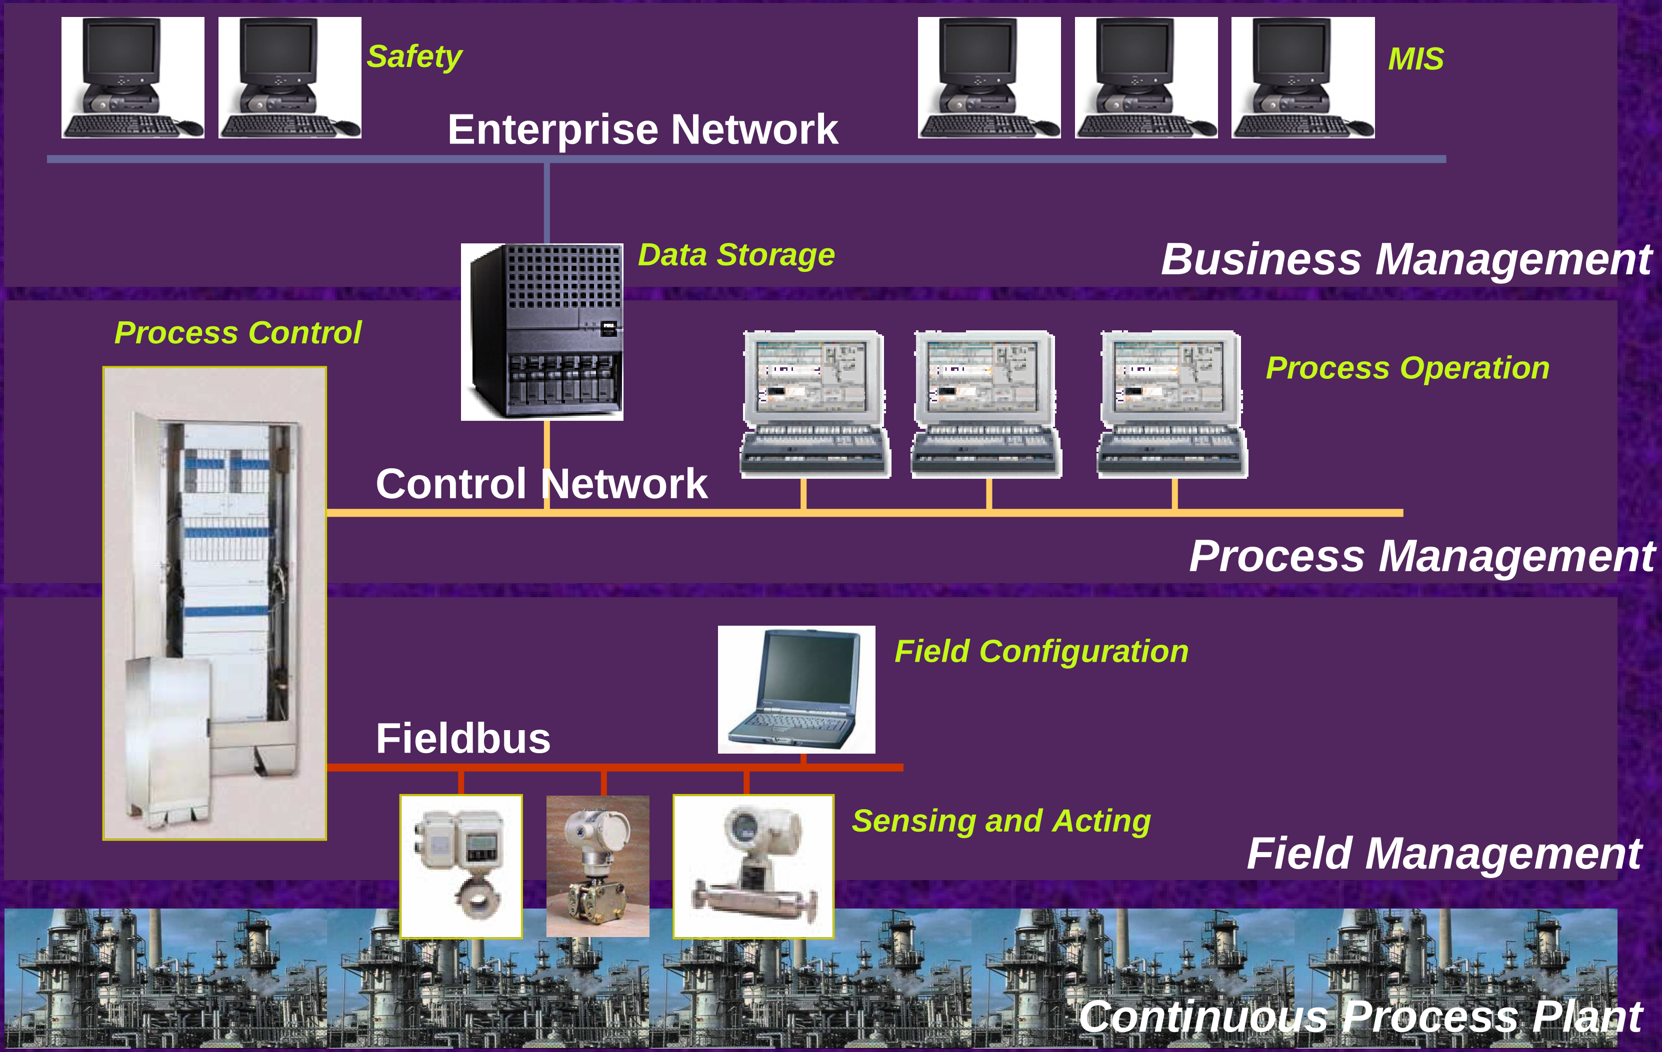

Perspectives

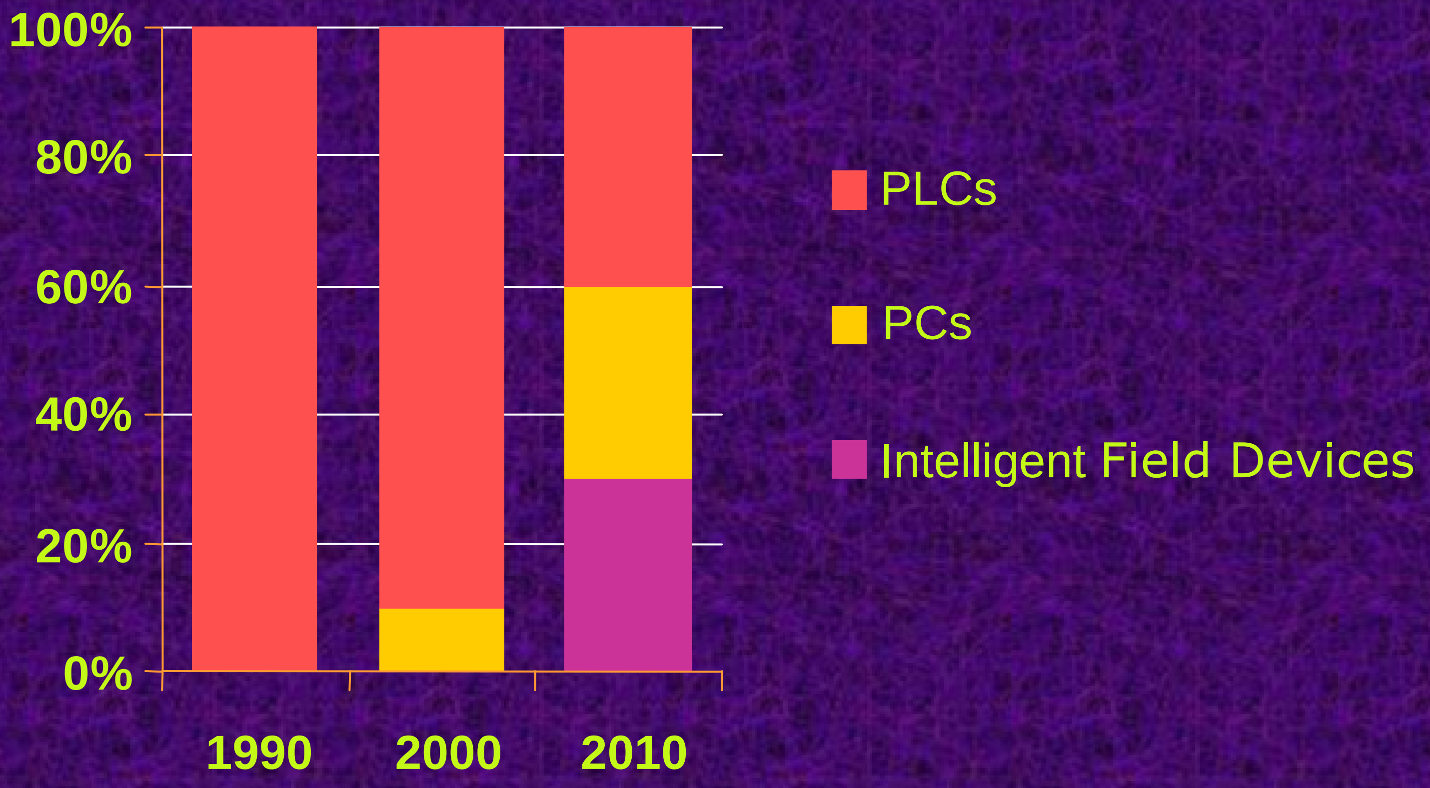

[PNO 2001]

Closing loops trough the net

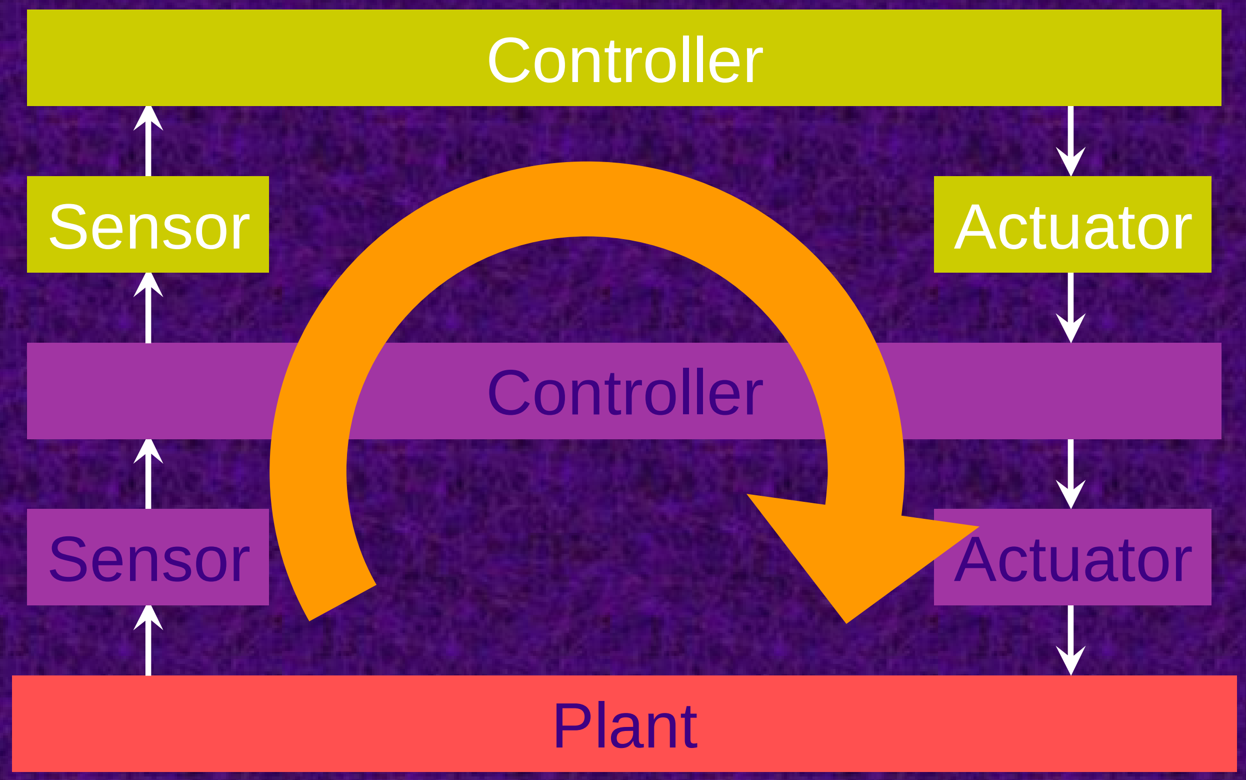

Network Control

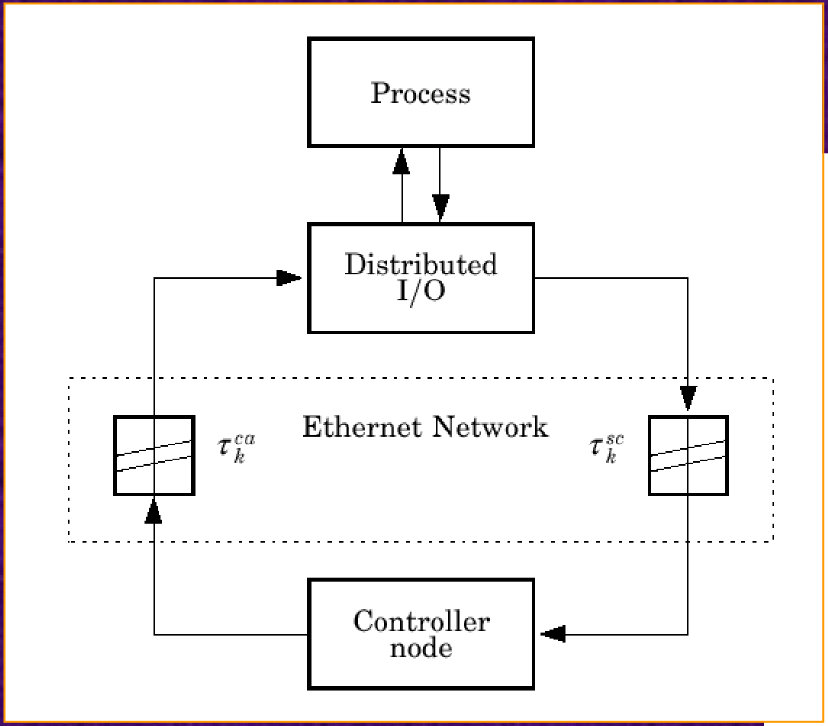

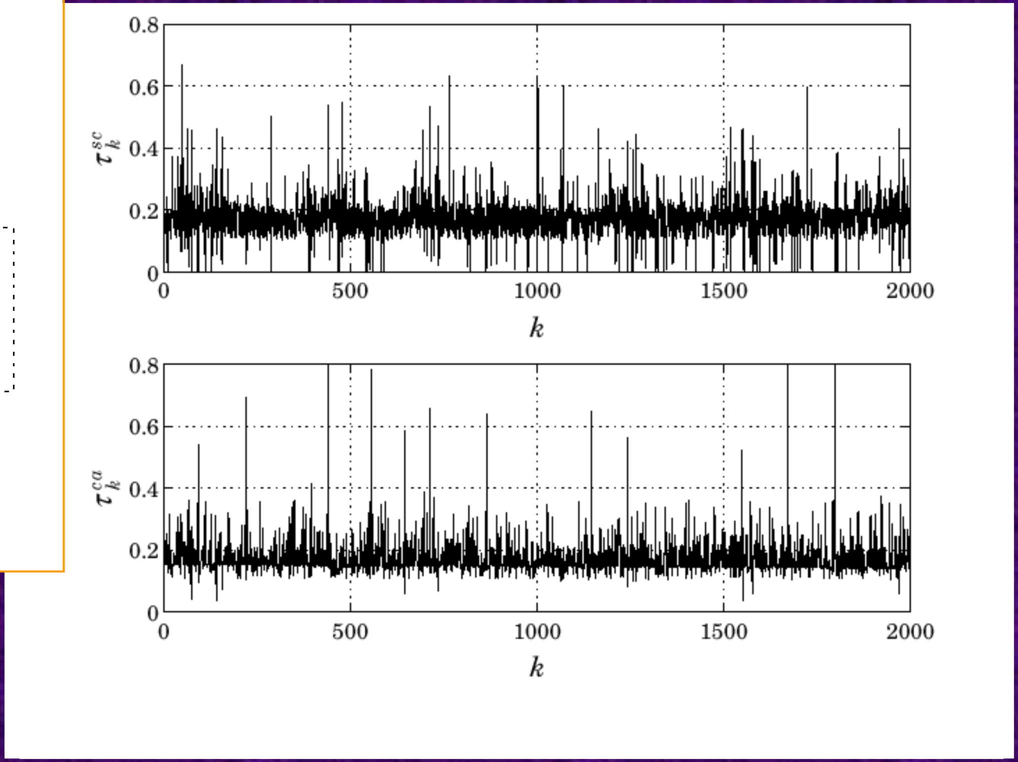

Design Challenges

- Predictability
- Performance
- Scalability
- Openness
- Coordination
- Transparency
- Naming
- Cost
- Consistency
- Failure handling
- Security
- Heterogeneity
- Mobility
- Load sharing
- Footprint
- *etc.*

Where's the infrastructure ?

PROGRAM	TOPIC	STD CANDIDATES
ANTS	Negotiation-based resource management	Architectural framework Negotiation protocols
NEST	Robust, compact coordination services for large-scale distributed embedded systems	Coordination service API-s Composable, lean micro-middleware
PCES	Language and compiler technology for DRE-s	Aspect/weaving API-s Componentized/aspectified middleware
ARMS	Adaptive,reflective middleware for simultaneous control of multiple QoS	QoS API-s ORB services to support QoS
SEC	High-performance control algorithms and systems	Controls API (facility) Design patterns for assurance
MoBIES	Model-based Integration of Embedded Systems via domain-specific tools	Domain specific modeling languages
NEPHEST	National Experimental Platform for Hybrid and Embedded Systems	A relatively new, long-term NSF program

[Watson]

Distributed Artificial Intelligence

When the network thinks

What’s intelligence?

- The capability of exploit knowledge

- What’s knowledge then?
  - Machine and task specific meaningful information

Why distribute intelligence?

- Intelligence is distributed across the plants for many reasons (both needs and enablers):
  - availability of suitable embeddable processors,
  - timing requirements that forbid communication due to latencies,
  - need of increased levels of performance that is achieved through parallelism,
  - simplification of construction and maintenance tasks through modularity,
  - reduction of cost and time-to-market by means of component-based reuse,
  - integration of legacy systems
  - availability of software platforms for AI
  - …

Distributed Artificial Intelligence

- Distributed Problem Solving (DPS)
  - Subalgoritmhs that collaborate
  - E.g. blackboard systems
- Parallel AI (PAI)
  - Splitting an algorithm for performance
  - Typically on multiprocessors
- Multi-agent systems (MAS)
  - When each component thinks

[Nwana]

A Definition of Agent

- An agent is an autonomous computer-based system capable of independent action in some environment with control over its internal state.

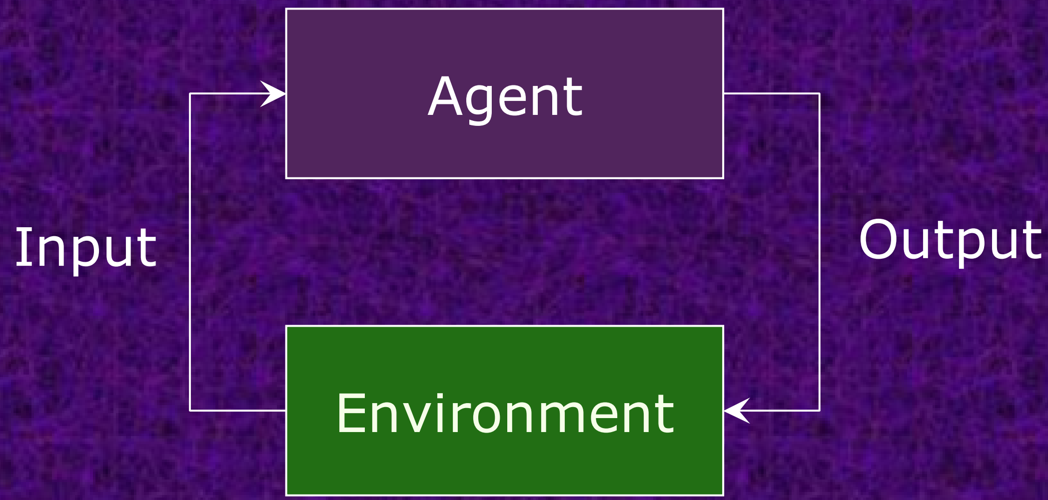

Agent Typology

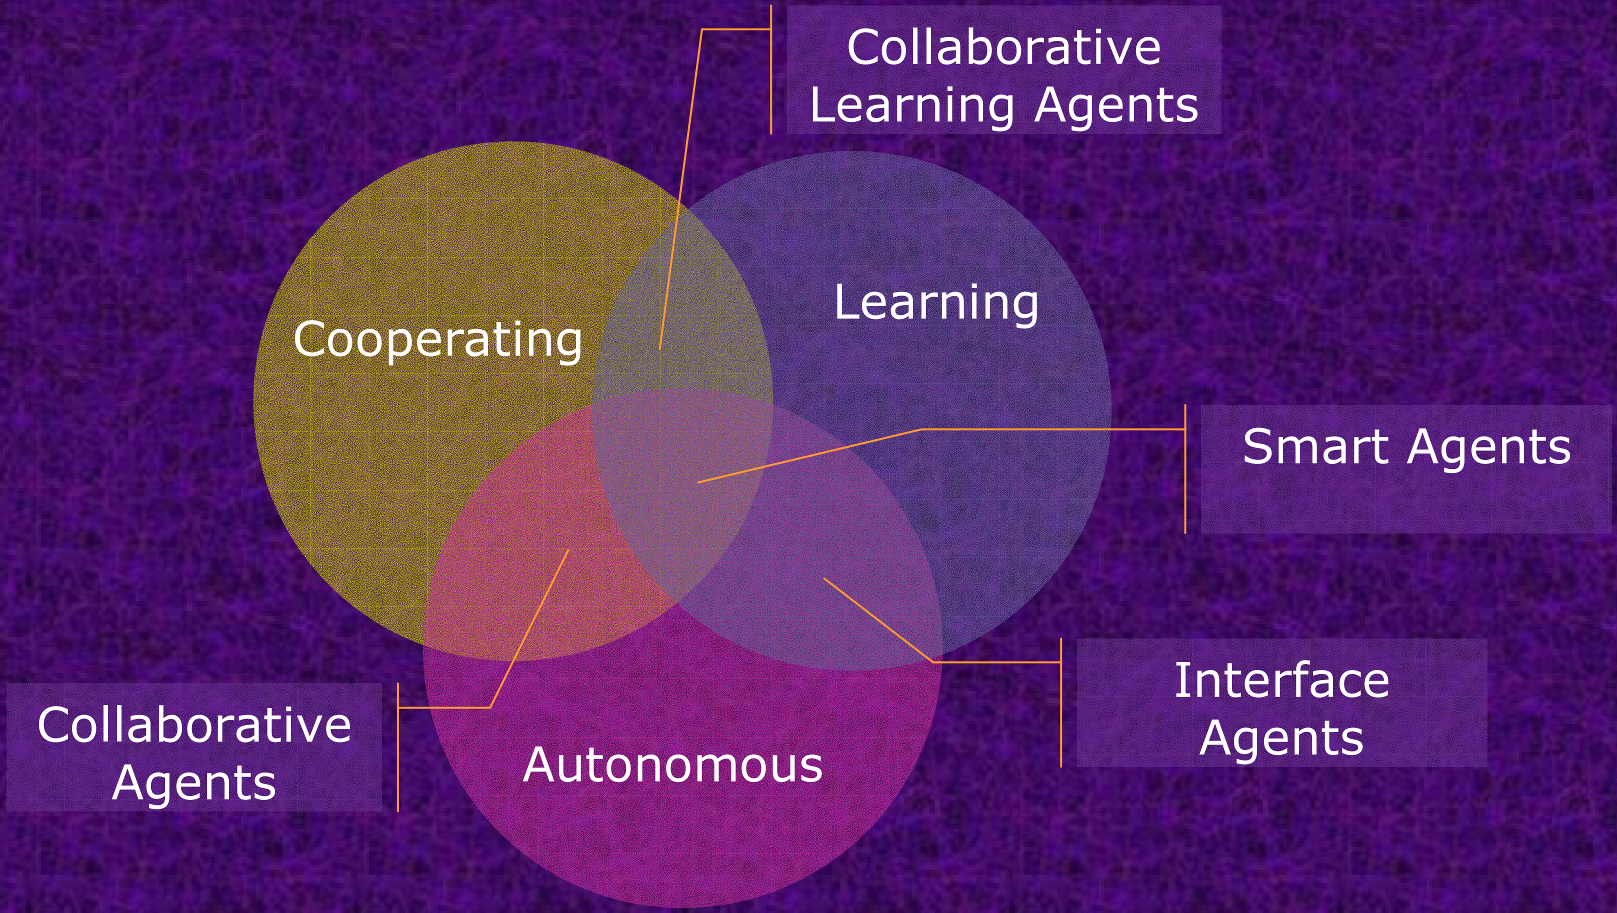

[Nwana]

Agent Typology

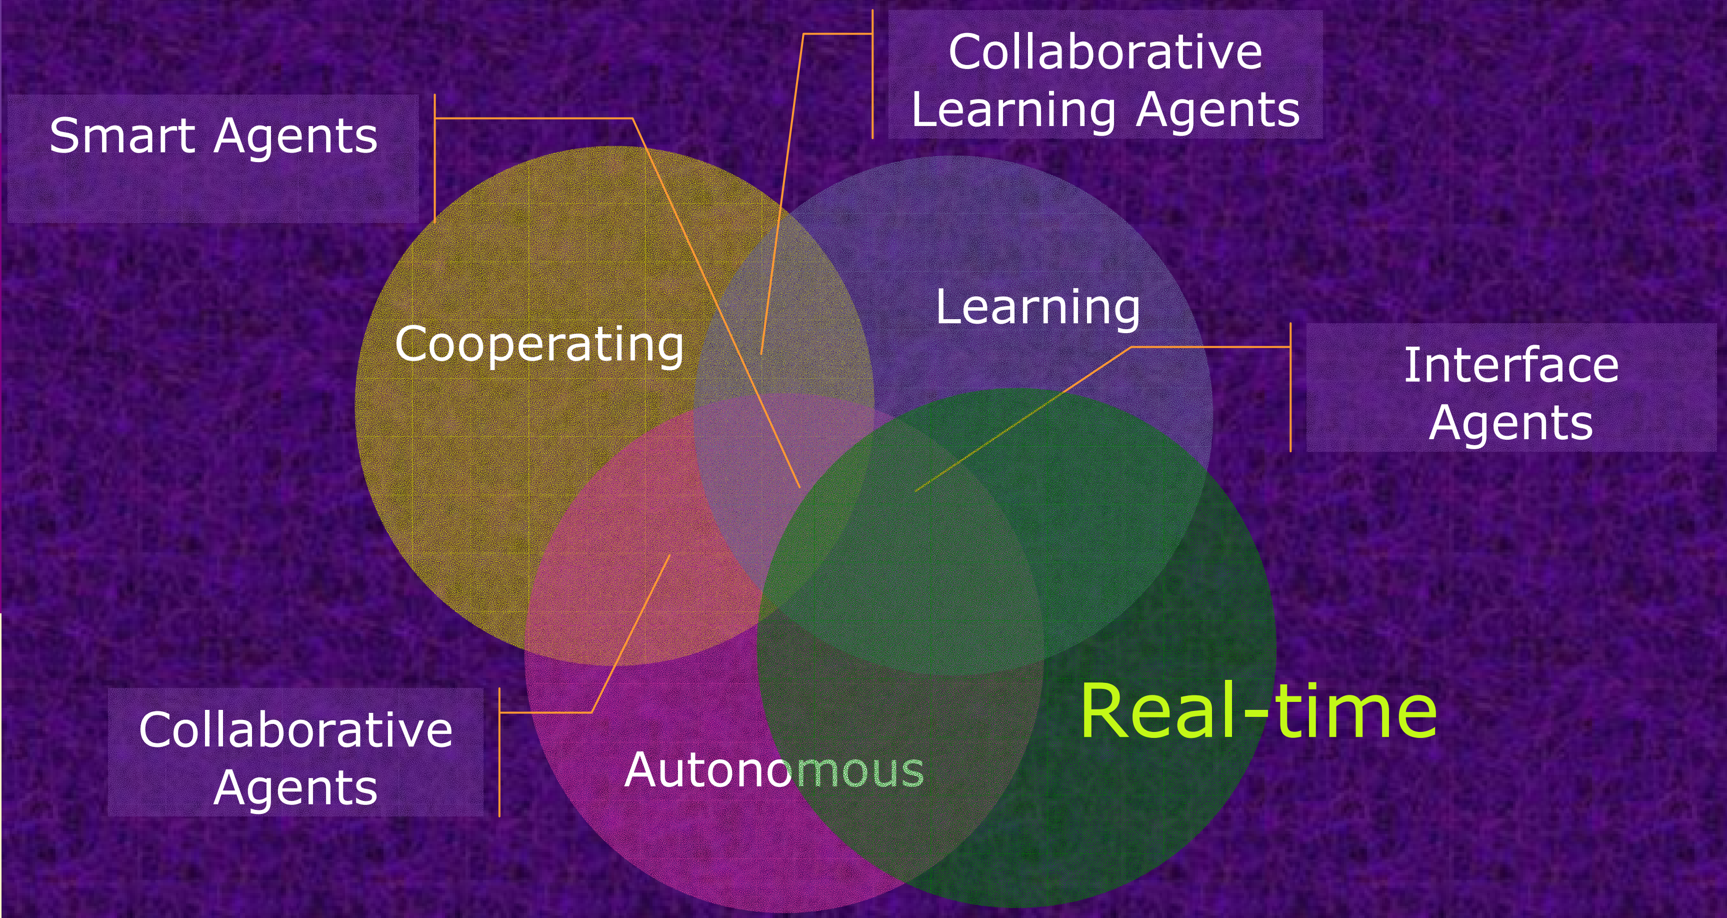

[Nwana]

Multi-agent Systems (MAS)

A number of agents which interact with one another through communication. Each agent will act upon or influence different parts of the environment.

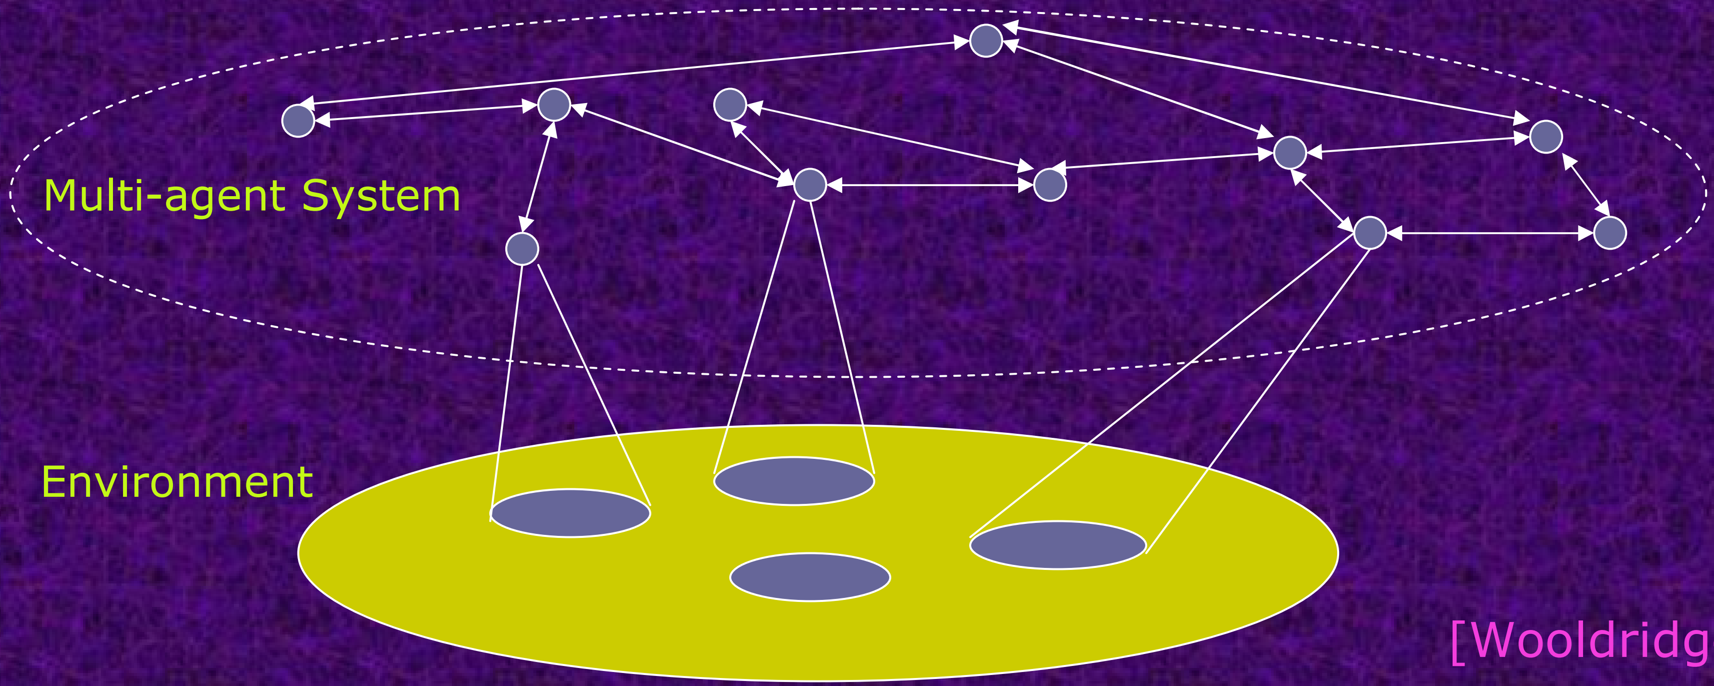

[Wooldridge]

Distributed Intelligent Control

Taking apart control thought processes

Elementary case

- Remote Expert Control
- Expert system running on special processors controlling remote system
- Typically multi-layered control
  - Reactive real-time controller
  - Intelligent layers

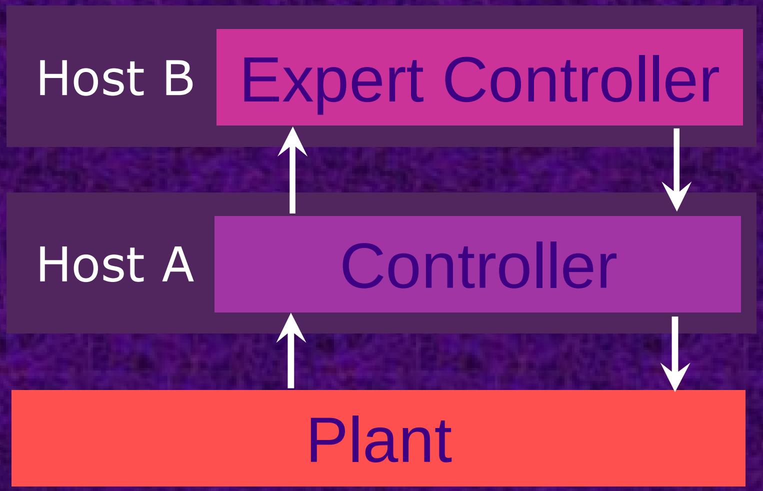

Strategic Process Control

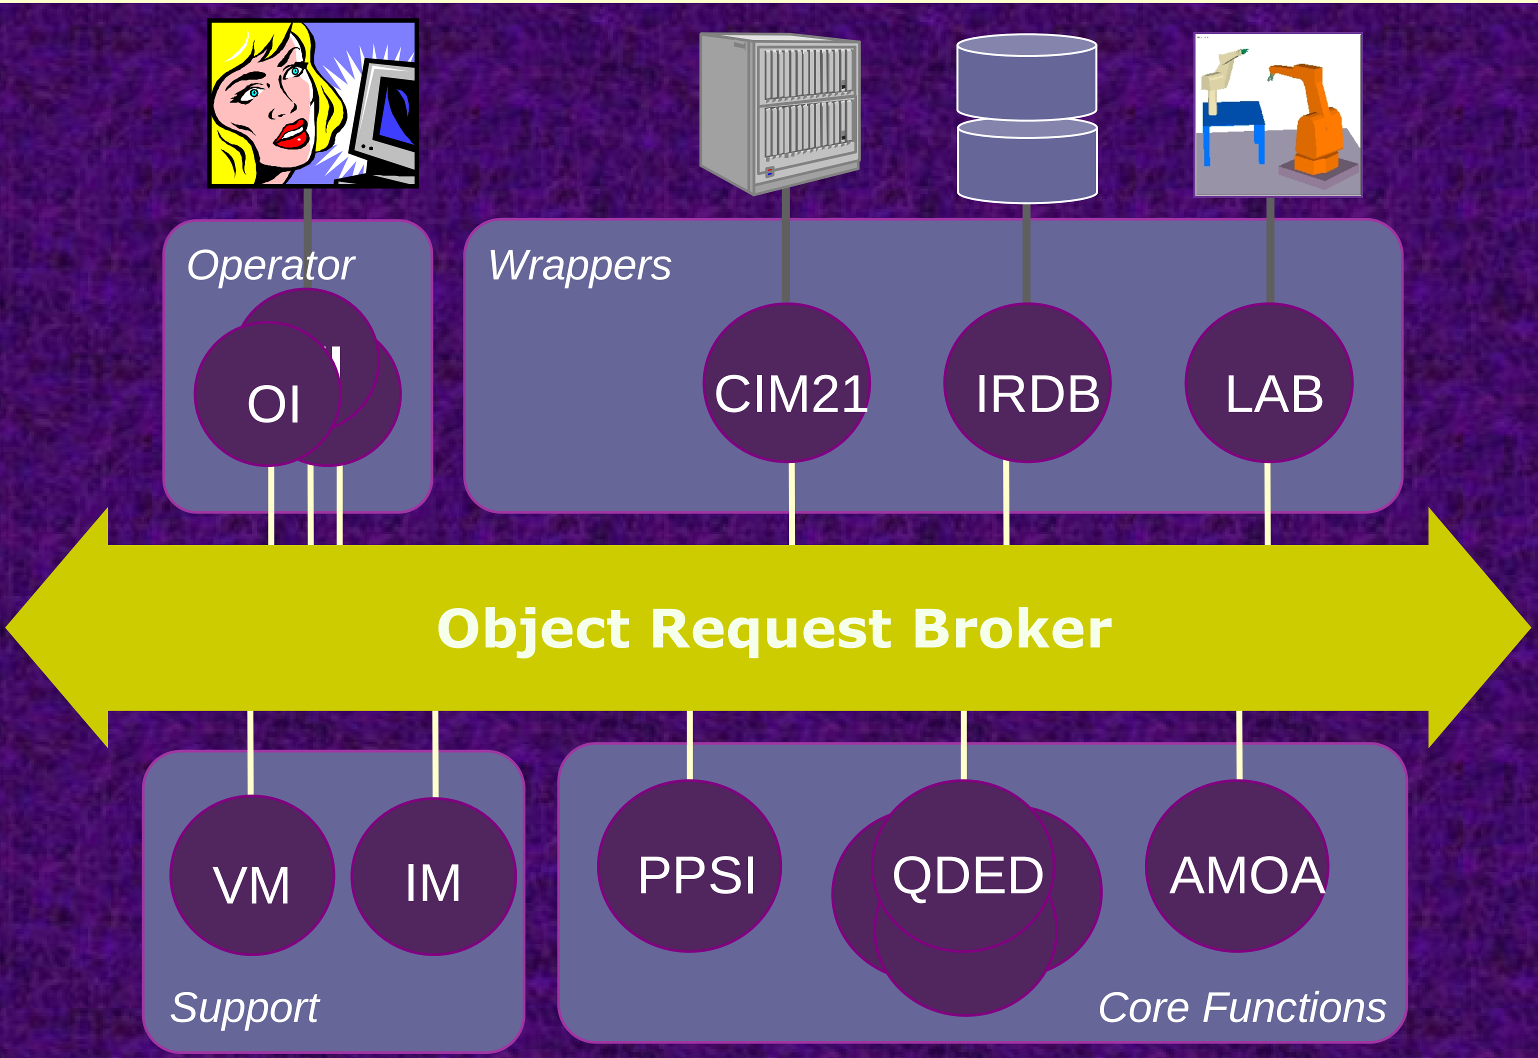

Total-Ship Systems

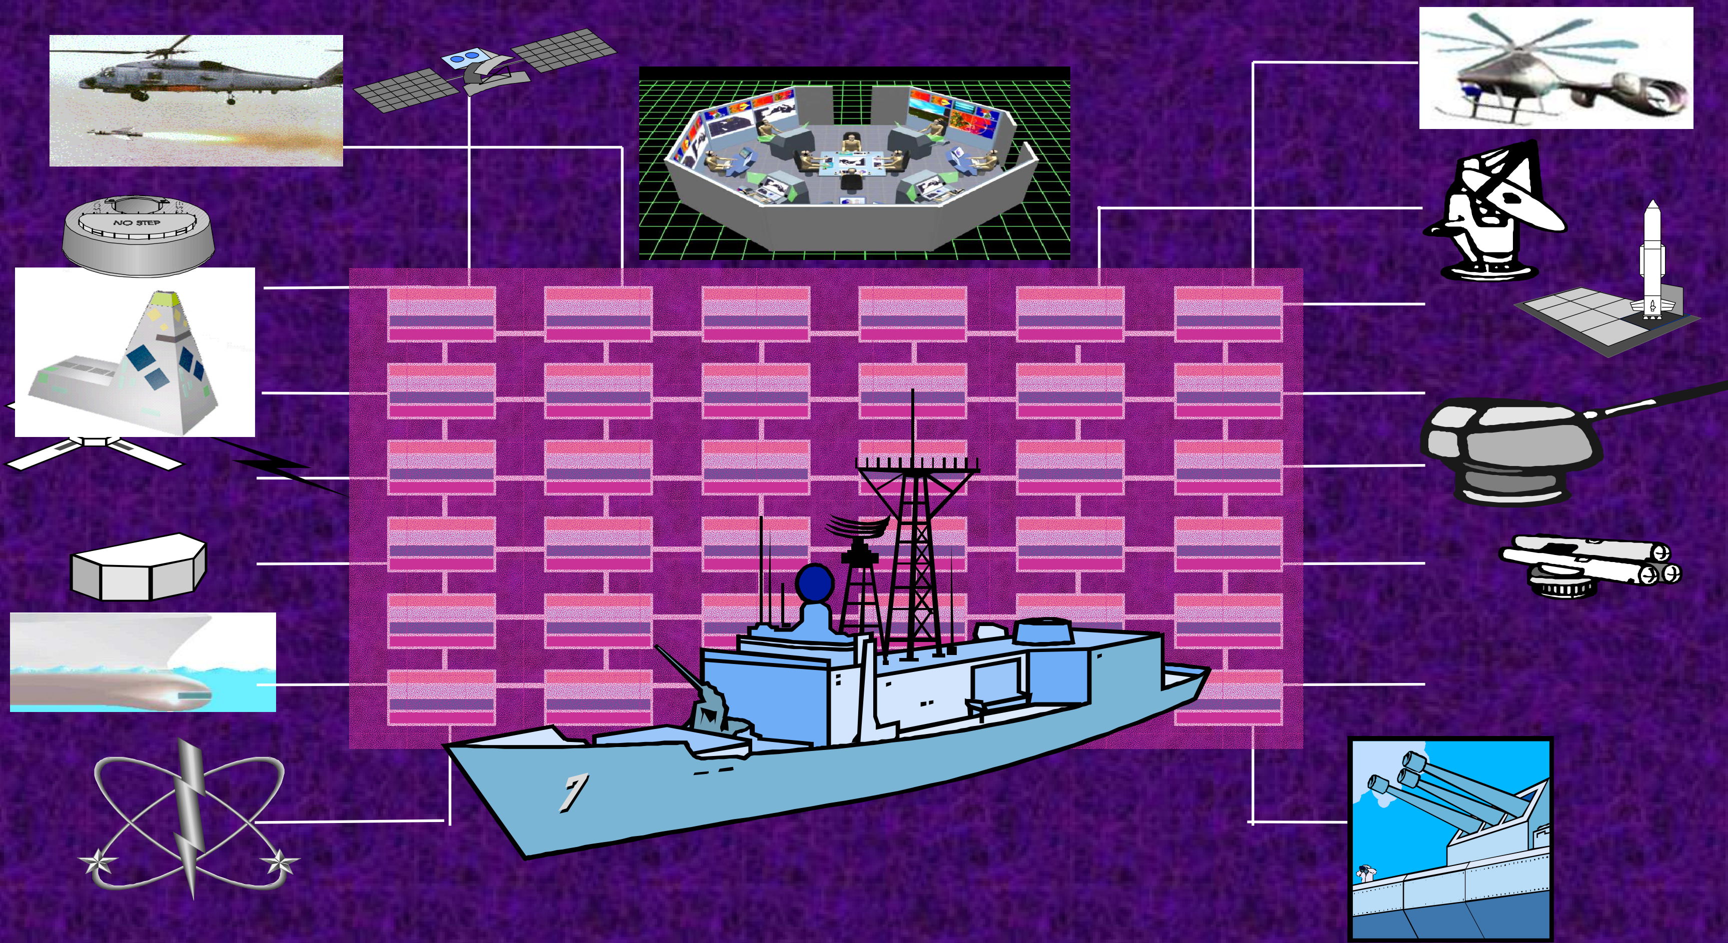

Plant-wide Risk Management

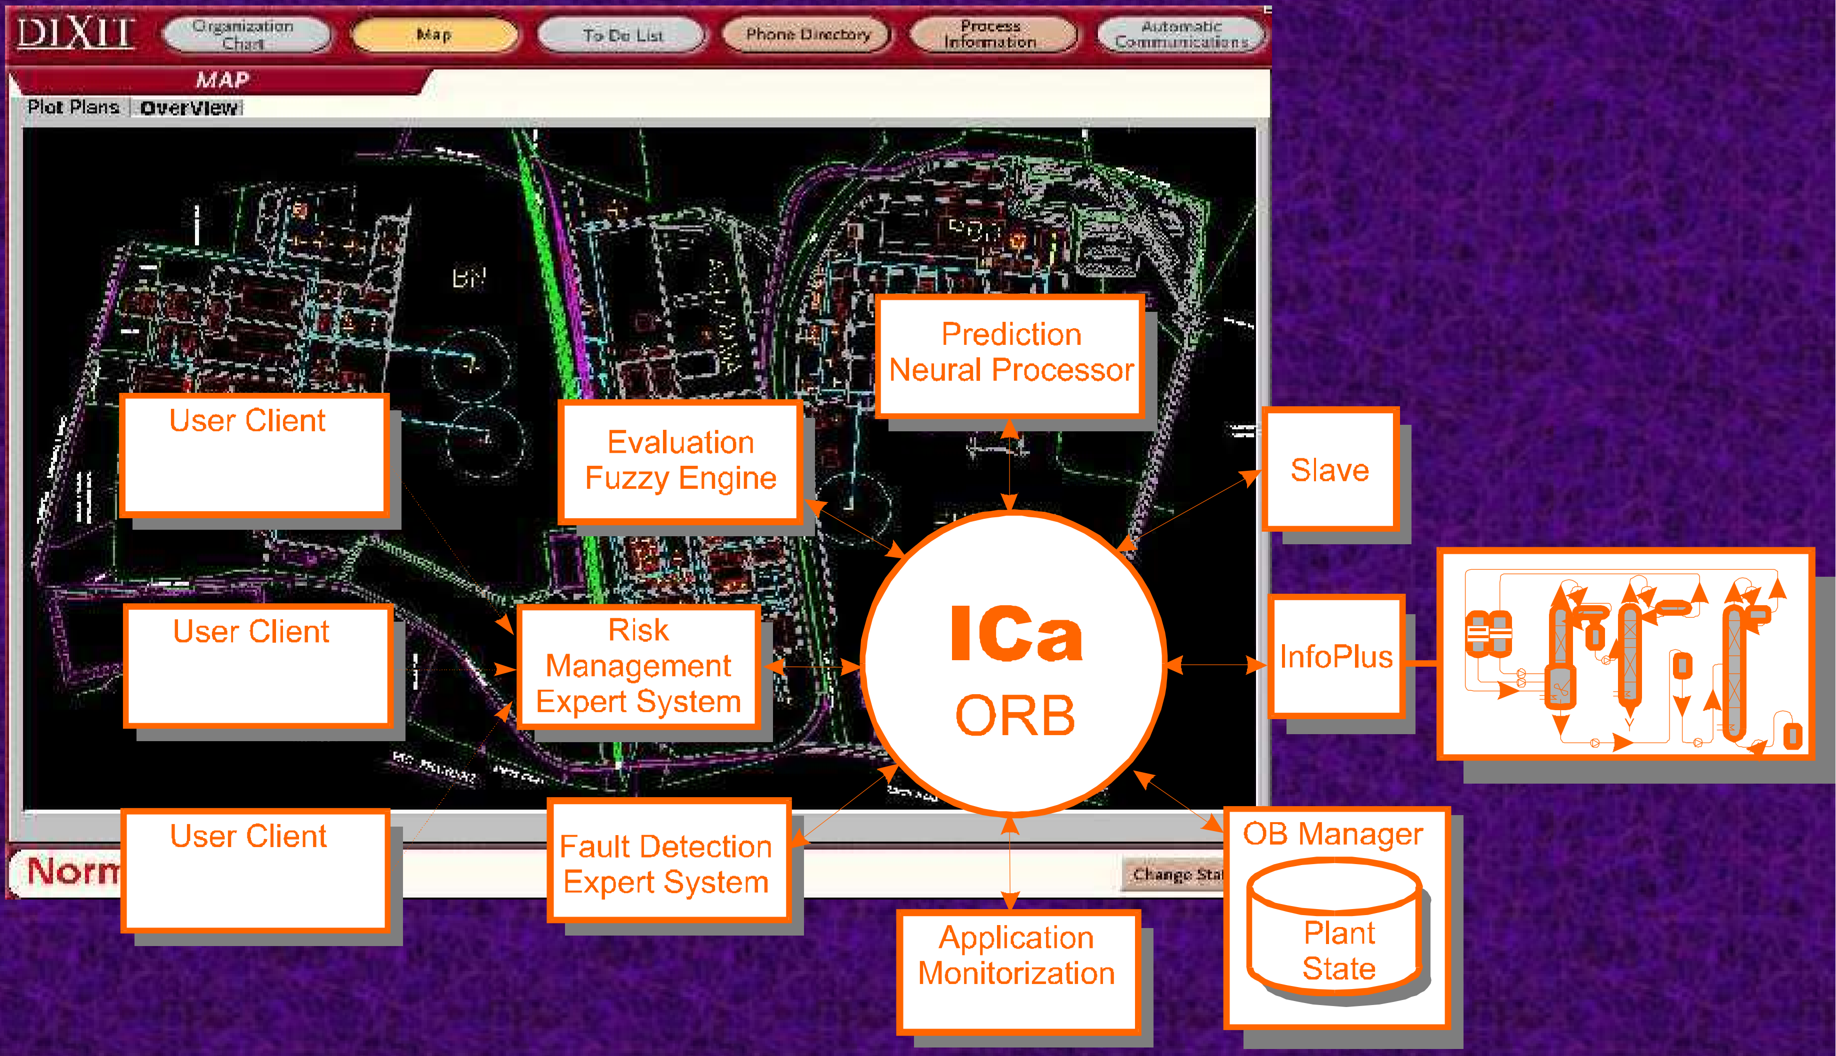

Deeper Intelligence

Increased intelligence levels in complex controllers

Pervasive Intelligence

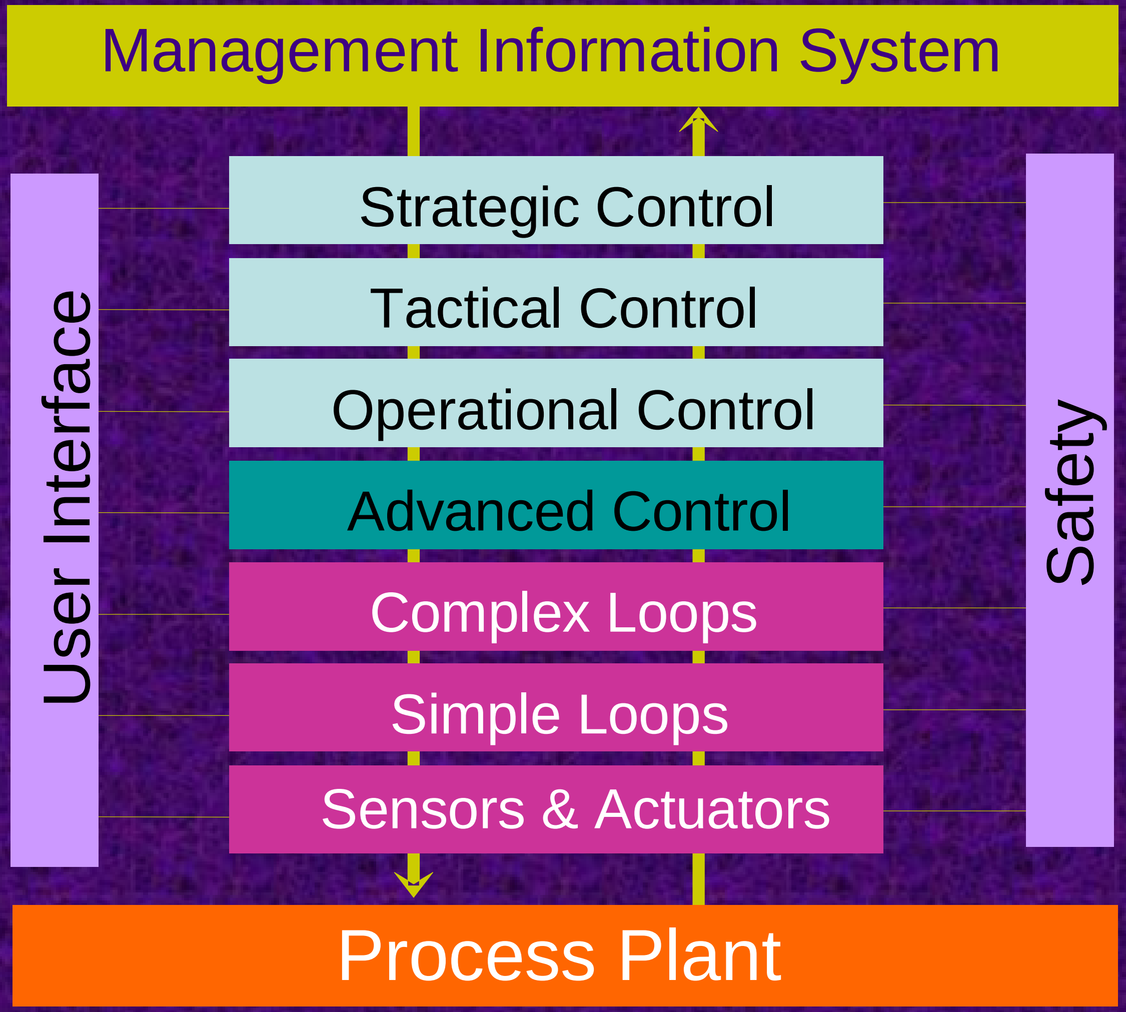

Deeply Embedded Controllers

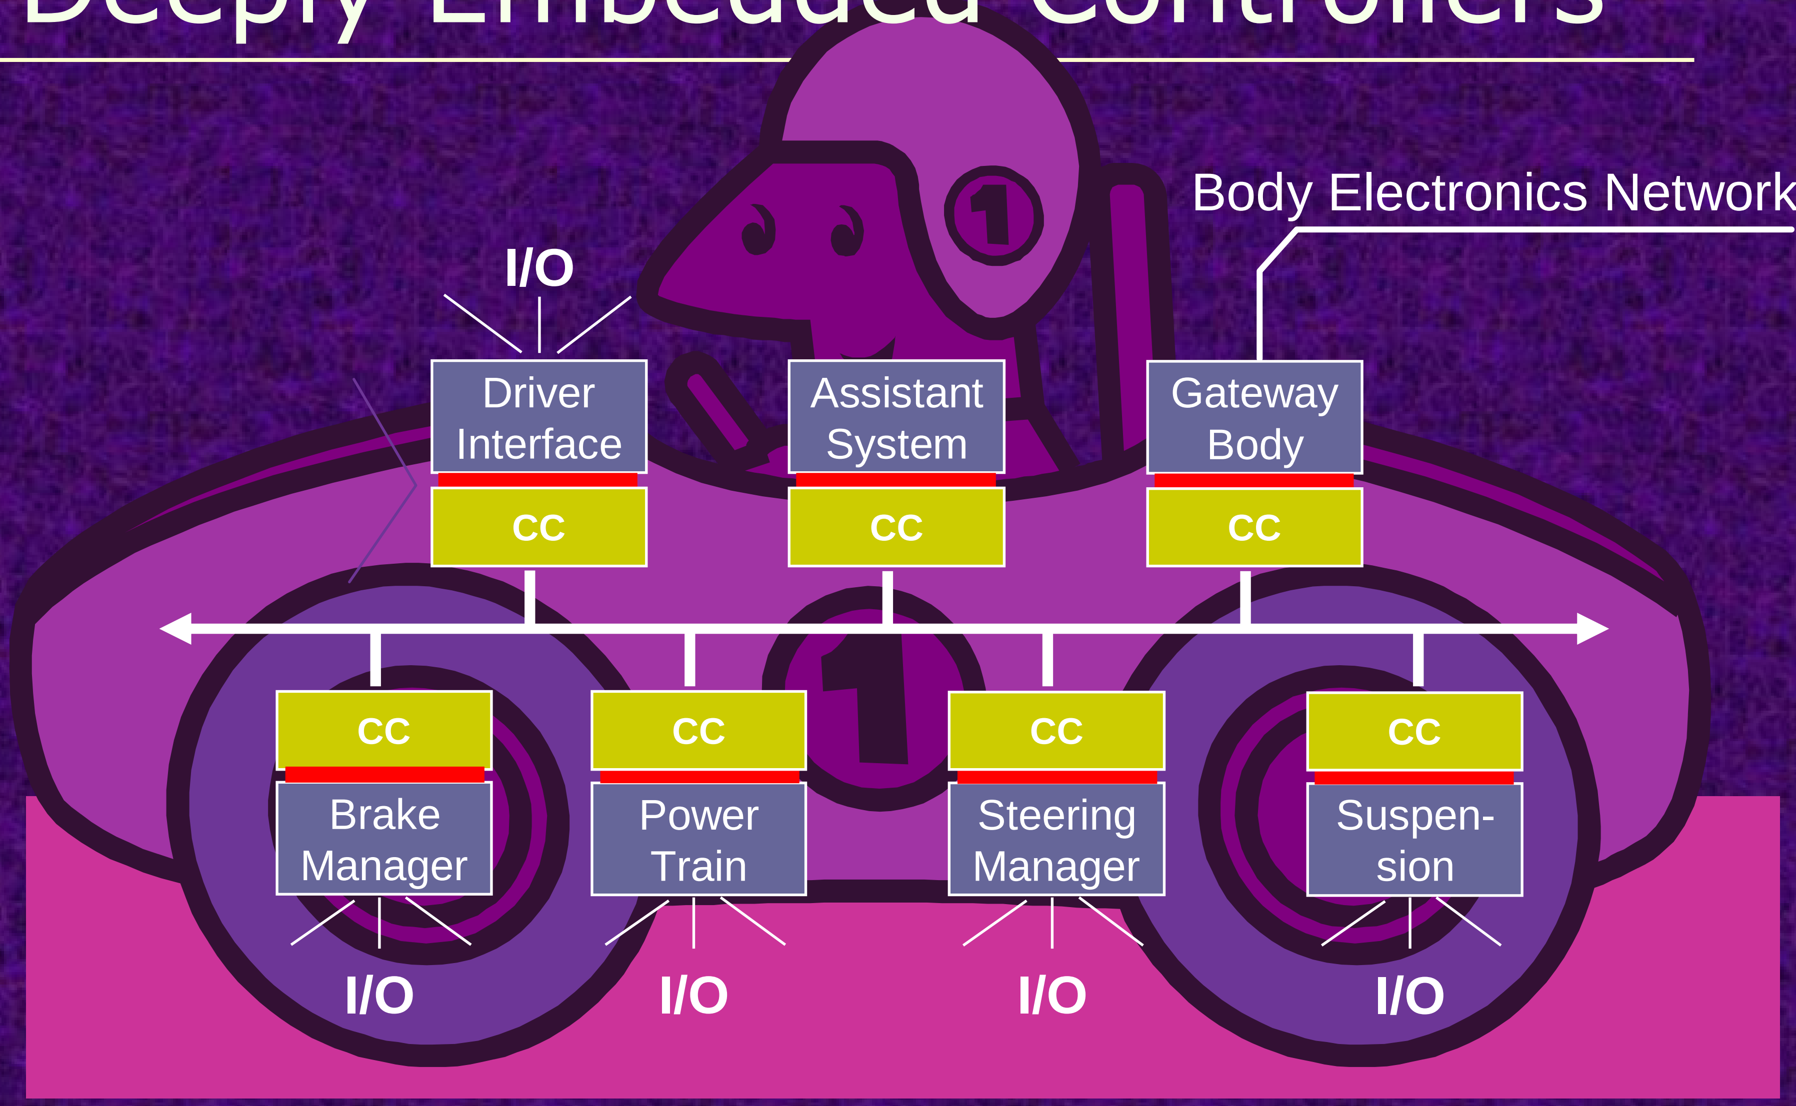

Organic Systems of Objects

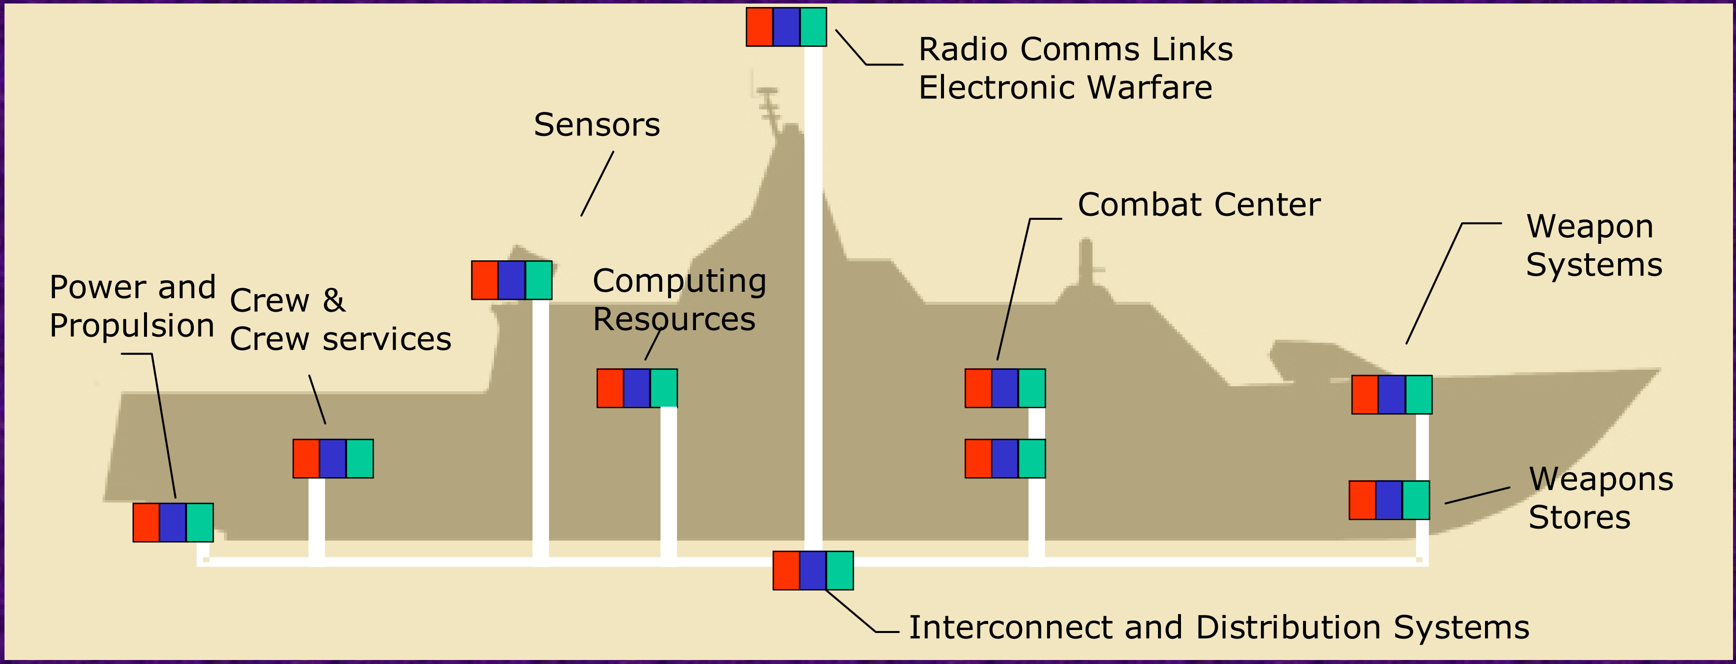

Physical Object | Logical Object (In Sys Mgmt Sense) | Network Connection

Multiple agent architectures

- Coarse grained architectures
  - BDI (Beliefs, desires, intentions)
  - GOMS (Goals, Operators, Methods, Selections)
  - Soar
- Fine grained architectures
  - Custom
  - FIPA.org
  - SIM_Agent
  - OAA

A Cognitive System Anatomy

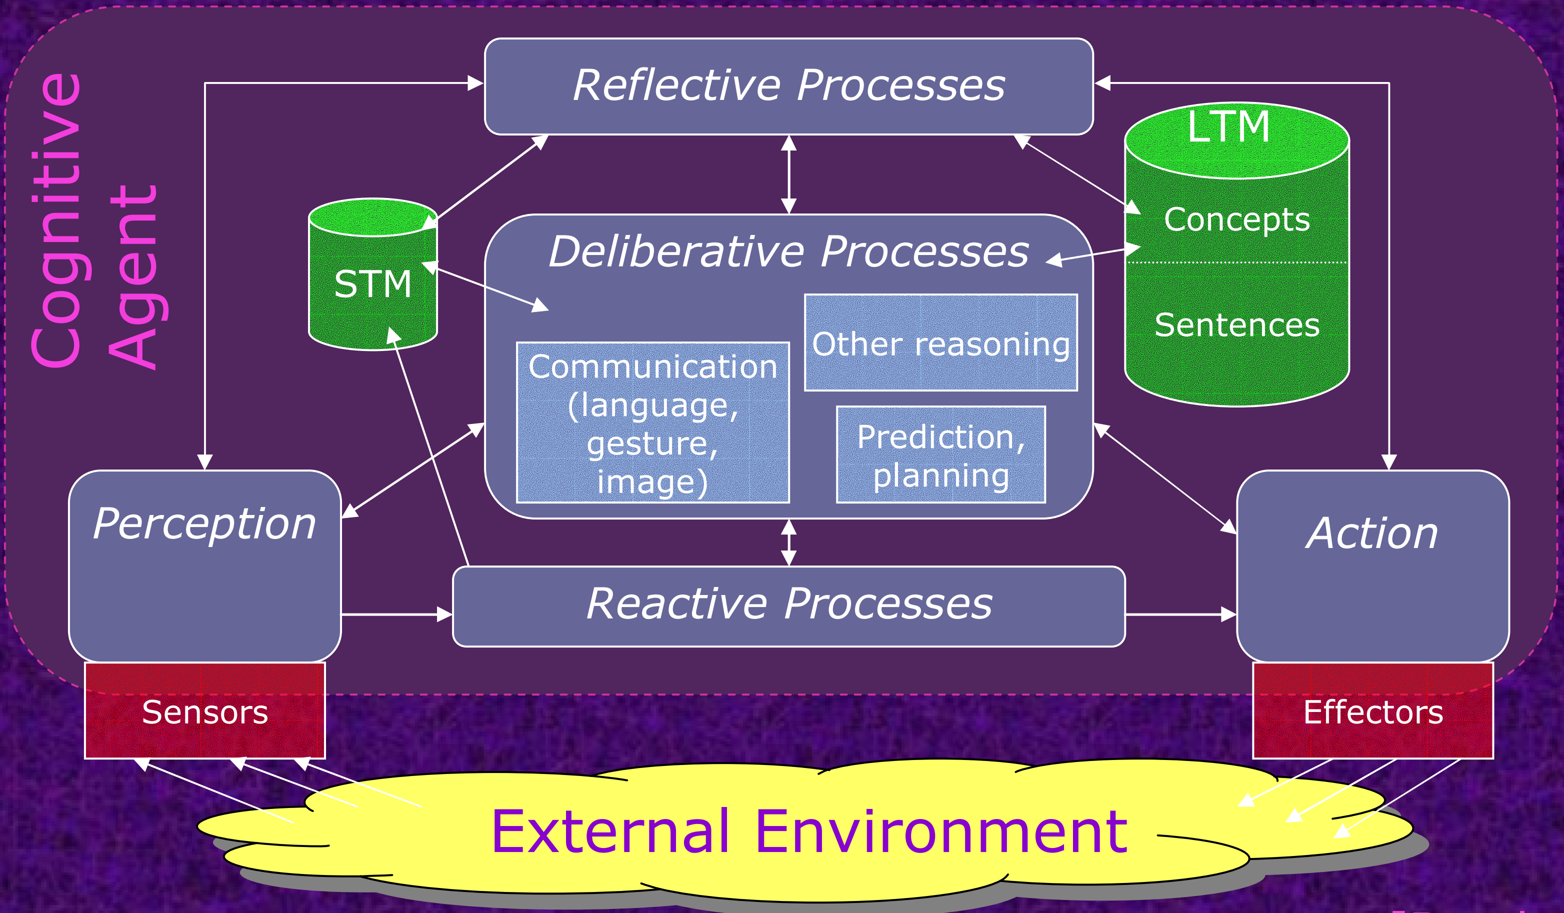

[Brachman]

Thrust in Cognitive systems

- A cognitive system is one that
  - can reason, using substantial amounts of knowledge
  - can learn from its experience so that it performs better tomorrow than it did today
  - can explain itself and be told what to do
  - can be aware of its own capabilities and reflect on itself
  - can respond robustly to surprise

[DARPA]

- A cognitive system is one that
  - is physically instantiated or embodied
  - can perceive and understand (the semantics of information conveyed through their perceptual input)
  - Can interact with their environment
  - can evolve in order to achieve human-like performance in activities requiring context-(situation and task) specific knowledge

[IST]

Beyond "Normal" Agents

- Dependable control agents do have requirements well beyond what is considered "normal" intelligent function:
  - Real-time behavior
  - Embeddability
  - Evolvability
  - Ugradeability
  - Robustness

Mechanisms for robustness

- H_2, H_∞ (robust control)
  - The system tolerates small displacement from design conditions
- Redundancy
  - Increase robustness up to a limit where the increase in dependability is less than the new induced risks
- Fault-tolerance
  - Copes with plant changes due to faults
- Reflection

Robust design

- “The suggestion that robust design is the primary source of complexity is motivated by the observation that for most biological and technological systems, the vast majority of components are present for robustness rather than for basic functionality of the organism or machine.”

[Reynolds]

Two Notions of Complexity

- "Complexity emerges in systems that are otherwise internally homogeneous and simple".
  - Self-organized criticality (SOC) and the edge of chaos suggests that large-scale structure arises naturally and at no apparent cost through collective fluctuations in systems with generic interactions between individual agents.
- "Complexity is associated with intricately designed or highly evolved systems".
  - Highly optimized tolerance (HOT) emphasizes the role of robustness to uncertainties in the environment as a driving force towards increasing complexity in biological evolution and engineering design.

Scalable Intelligence

One-size fits-all intelligences

Levels of Machines

- The machine as a tool: The user can execute an action with the machine
- The machine as an automaton: The user (or condition) can initiate action sequences
- The machine as an agent: The user can request a context dependent function
- The machine as an autonomous agent: The machine executes context dependent actions as is deemed necessary by (heuristic) set of rules
- The machine as a cognitive agent: The machine understands and is aware; is able to execute tasks requiring real intelligence and thought

[Haikonen]

Problems

- Traditional control systems construction processes are based on proven methods that can lead to well-performing systems.
- Conventional control systems modules for distributed controllers can be built and are built this way.

- But there is a major problem when we address complex systems:
  - traditional engineering methods for high-assurance, distributed or intelligent control systems do not scale well.

Scalable Control

- **Scalable**
  \Scal"a*ble\, a. Capable of being scaled.
- Scaled to what ?
- Dimensions
  - Space: complex wide-area plants
  - Time: multiple time-scale loops
  - Rationality: levels of thought
  - Size: down to embeddability

Societies of Intelligences

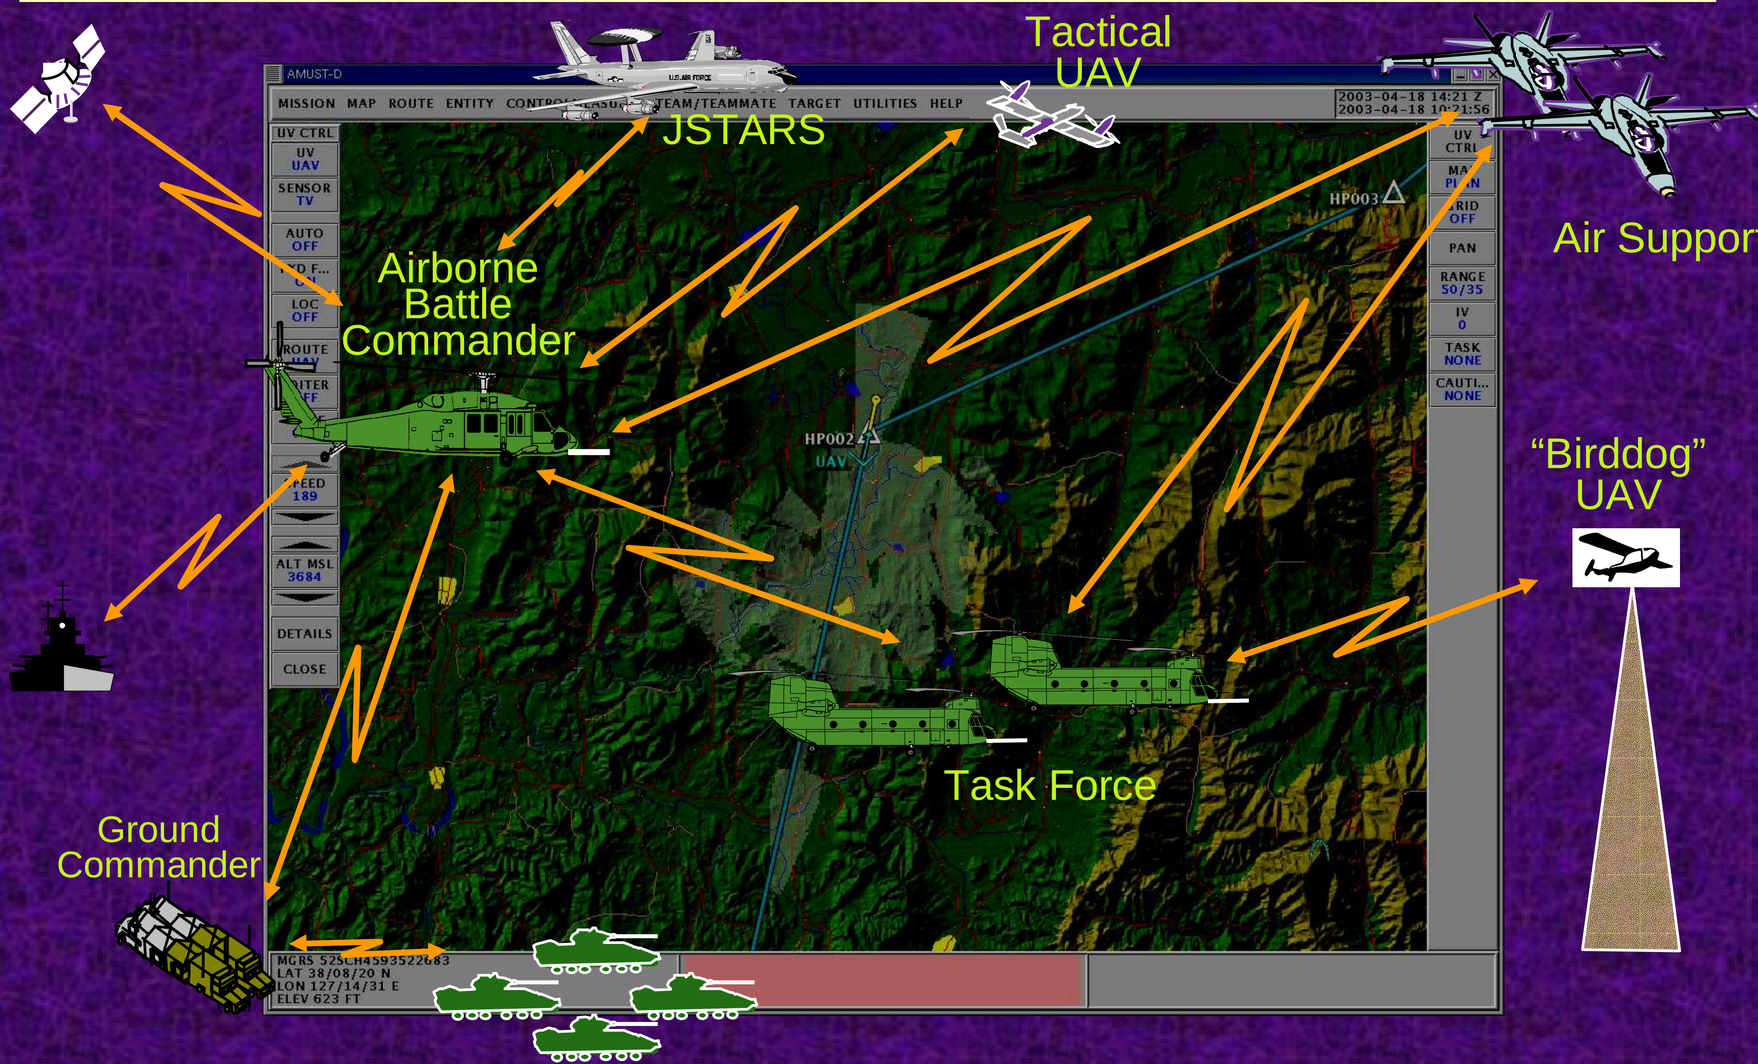

The Layered Agent Pattern

Mobility	When in different locations, this layer is required for message transmission
Translation	Agent formulates a message for another agent, translating it into other semantics
Collaboration	Agent determines how to collaborate with another agent
Actions	Actions are carried out. Pending actions are queued
Reasoning	Agent determines what to do
Beliefs	Agent's models of itself and its environment
Sensory	Agent gathers sensory input from its environment

[Kendall]

Scaling Intelligence

- Increasing intelligence levels to meet requirements
- Use of heterogeneous technologies
- Use of control encapsulation
  - PID
  - Fuzzy
  - Rule-based Learning

- Now touching the frontiers of being-human

Integrated Control Architecture

- Real-time Middleware
- Generic agents
  - Architectural
    - Core, RT, FT, BDI
  - Task
    - PlantState, Simulator, etc.
  - Interface
    - Wrappers, Probes, etc.
- Design Patterns

Adding Reflection

- Incorporates structures representing (aspects of) itself
- Go beyond what is called fault-tolerant control
- The self representation is causally connected to the aspects of the system it represents:
  - The system has always an accurate representation of itself
  - The status and computation of the system are always in compliance with this representation

IDA’s Architecture

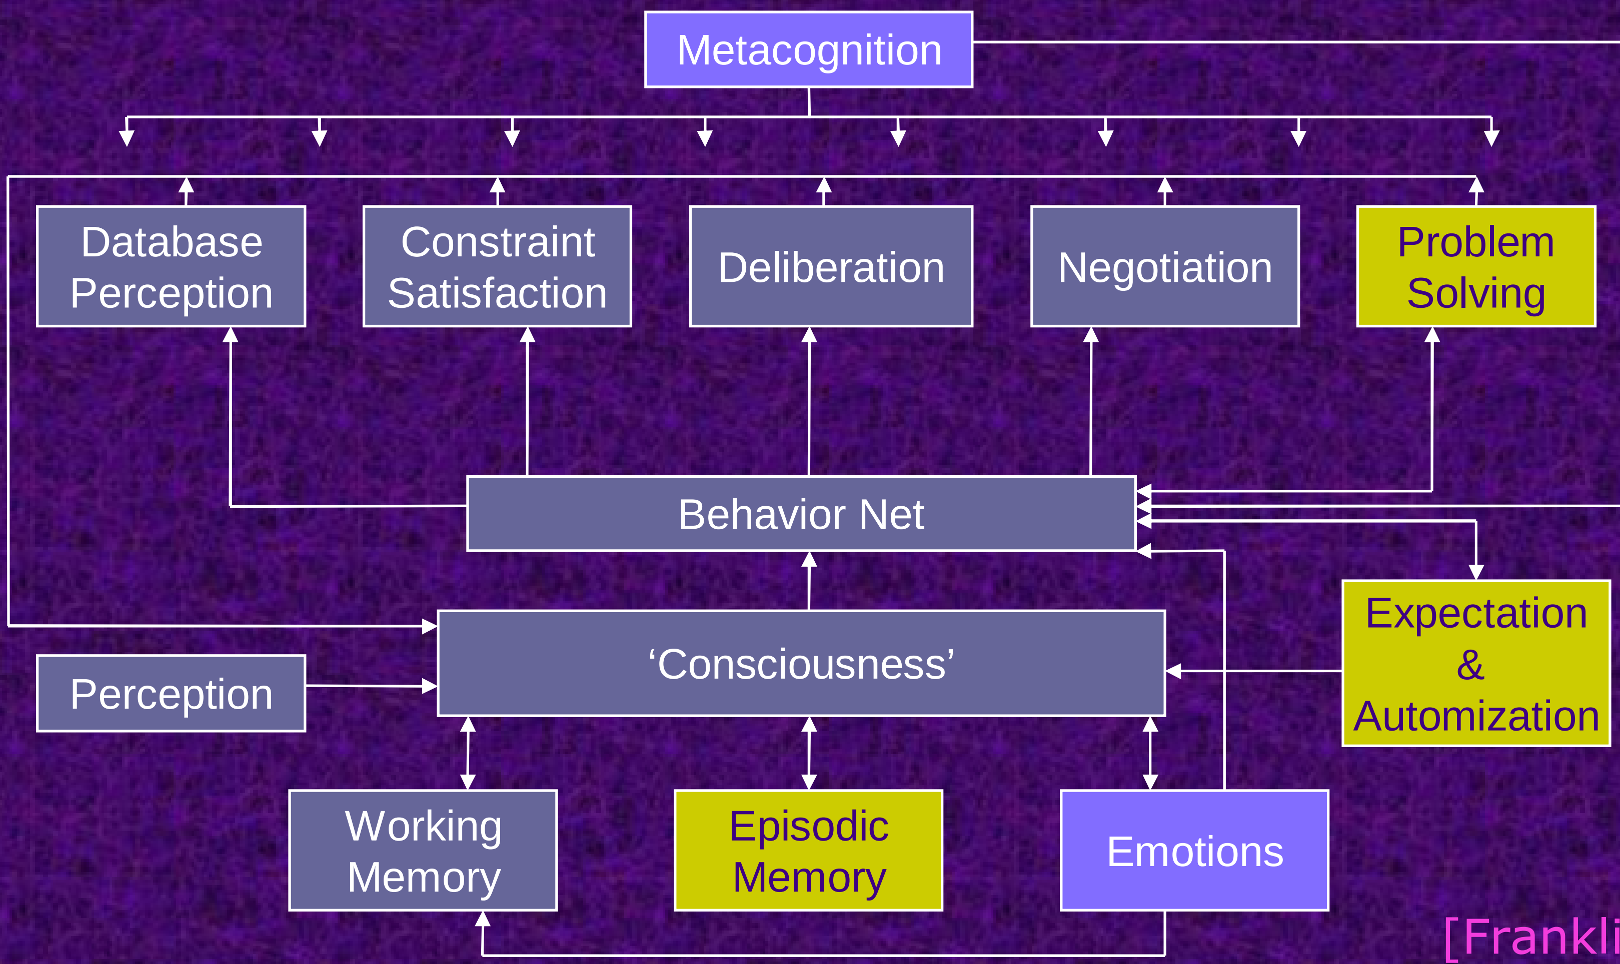

[Franklin]

Autonomic Systems

(as IBM sees them)

- Adapts to changes in its environment
- Strives to improve its performance
- Heals when it is damaged
- Defends itself against attackers
- Exchanges resources with unfamiliar systems
- Communicates through open standards
- Anticipates users' actions

- Possesses a sense of self

[SciAm May 06, 2002]

Hot words

- Meaning
- Value
- Awareness
- Consciousness
- Self
- Emotion
- Imagination
- Qualia
- Wisdom

The question ?

- Is there any agent architecture that could be scaled across dimensions and still guarantee critical control properties?

Scalable Distributed Intelligence

This is the wild speculation part

House Energy Agency

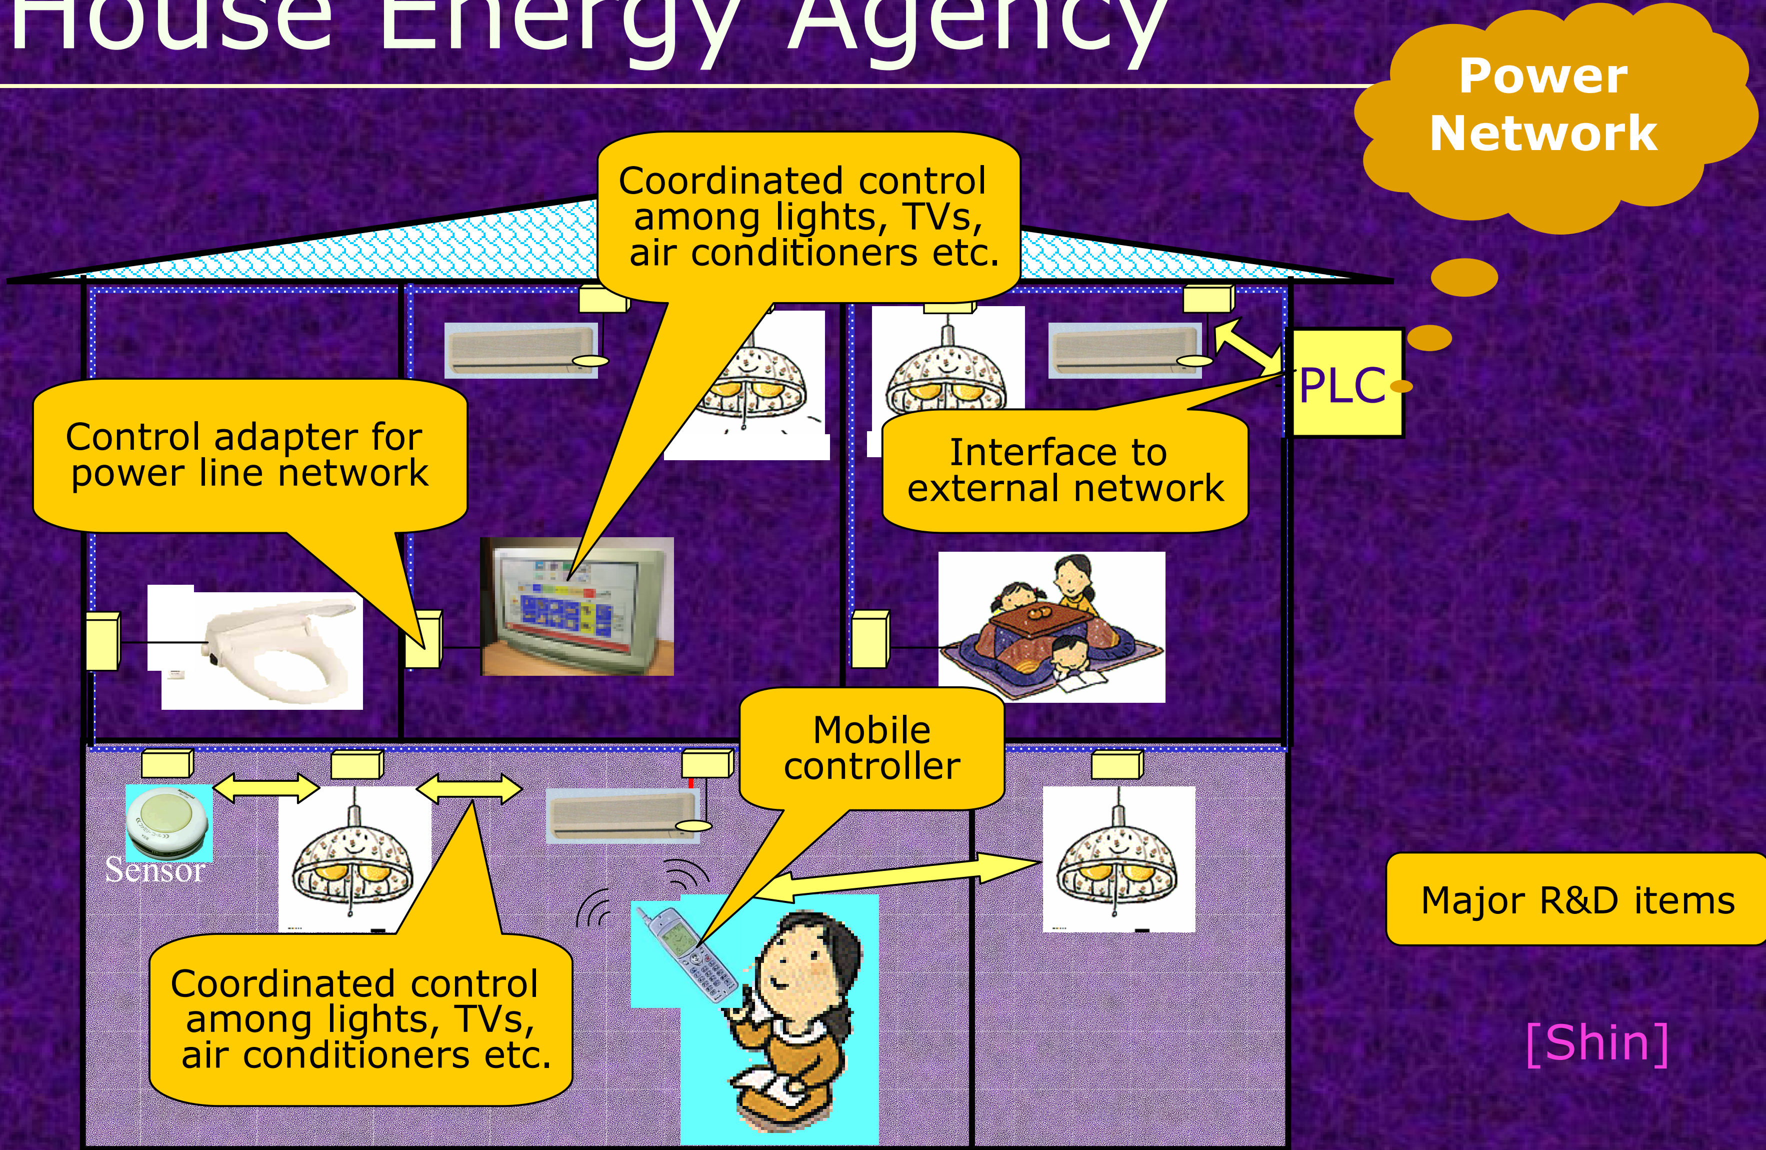

Major R&D items

[Shin]

Agents of Agents

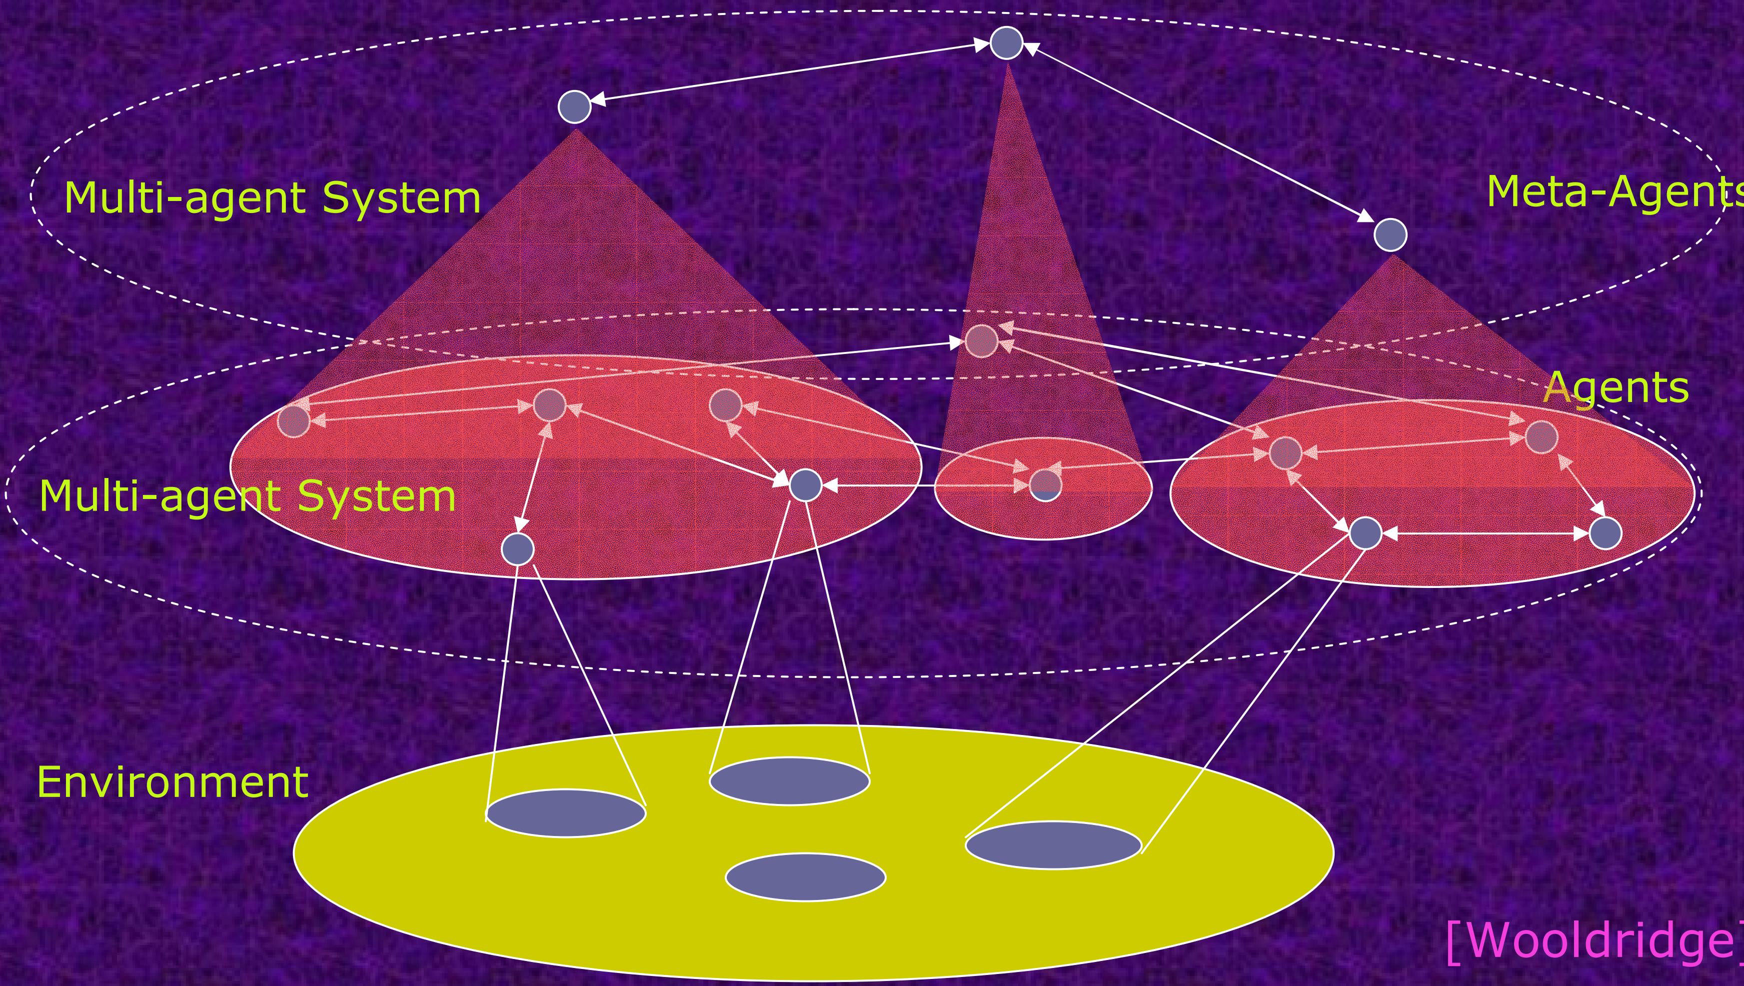

[Wooldridge]

Modules and protocols

- Protocols and modules are complementary (dual) notions
- Primitive technologies
  - Modules are more important than protocols
- Advanced technologies
  - Protocols are at least as important
- Even bacteria are “advanced technology”

[Doyle]

Biochemical Network

- E. Coli Metabolism

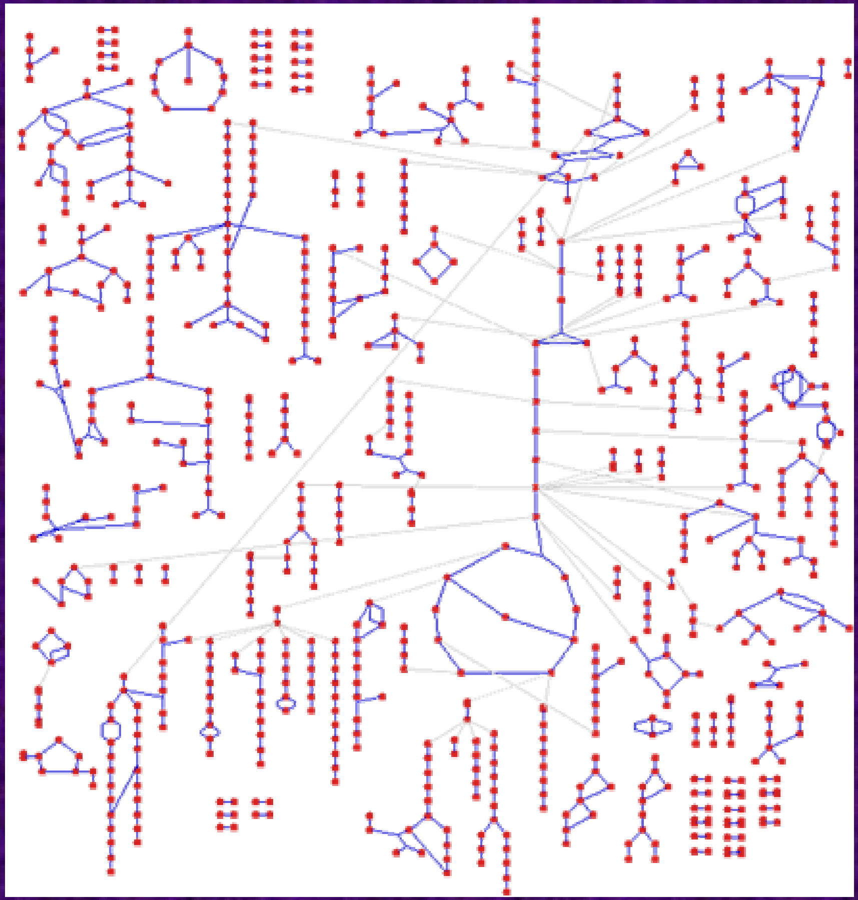

[Arkin]

A "universal" architecture?

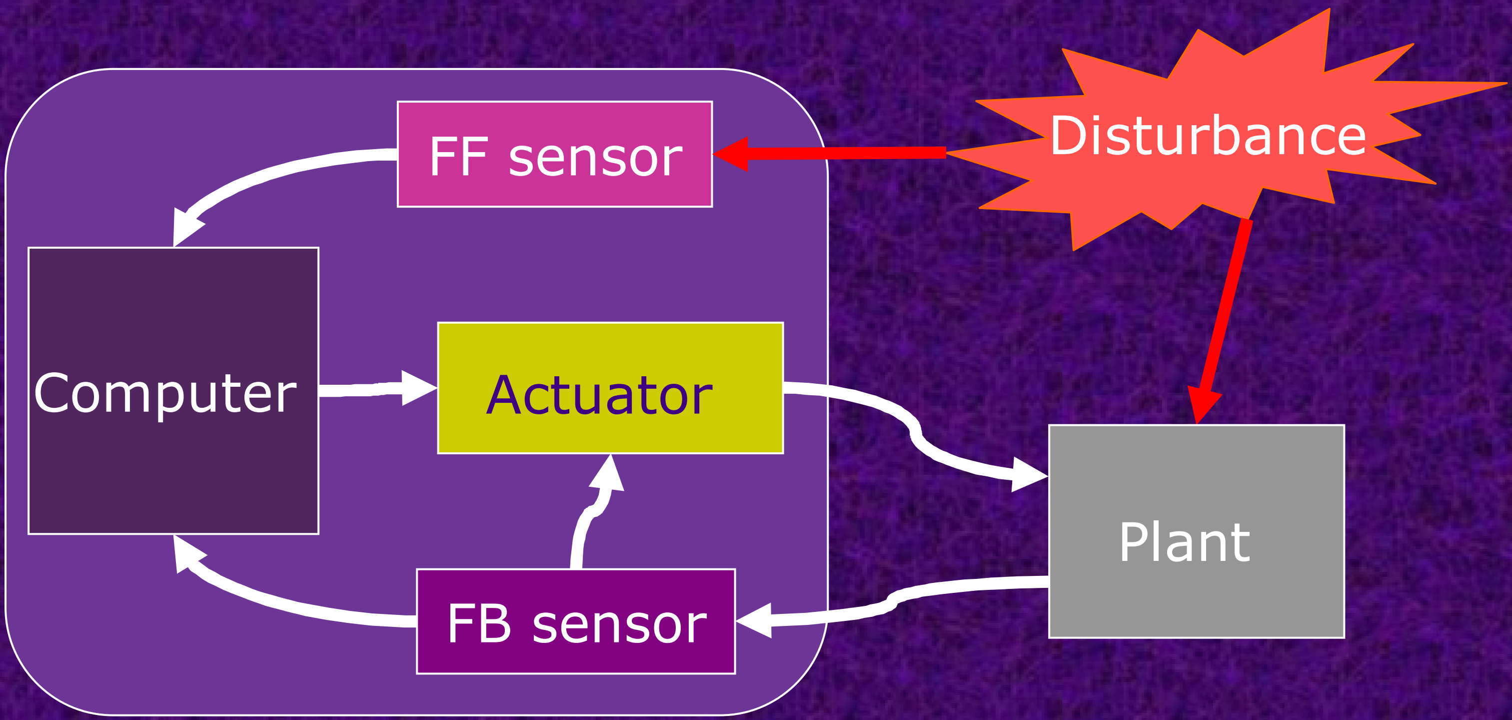

[Doyle]

A "universal" architecture?

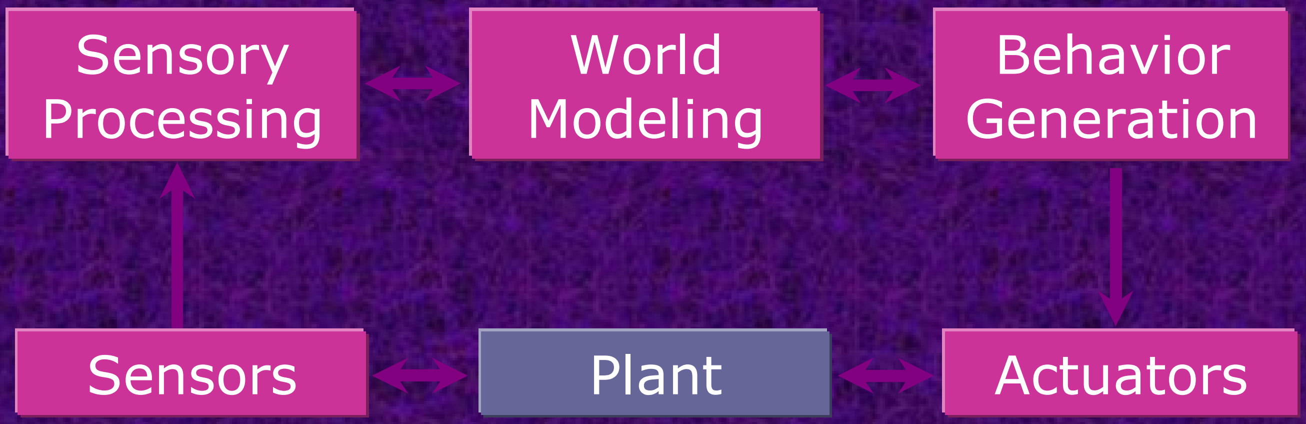

[Meystel]

A “universal” architecture?

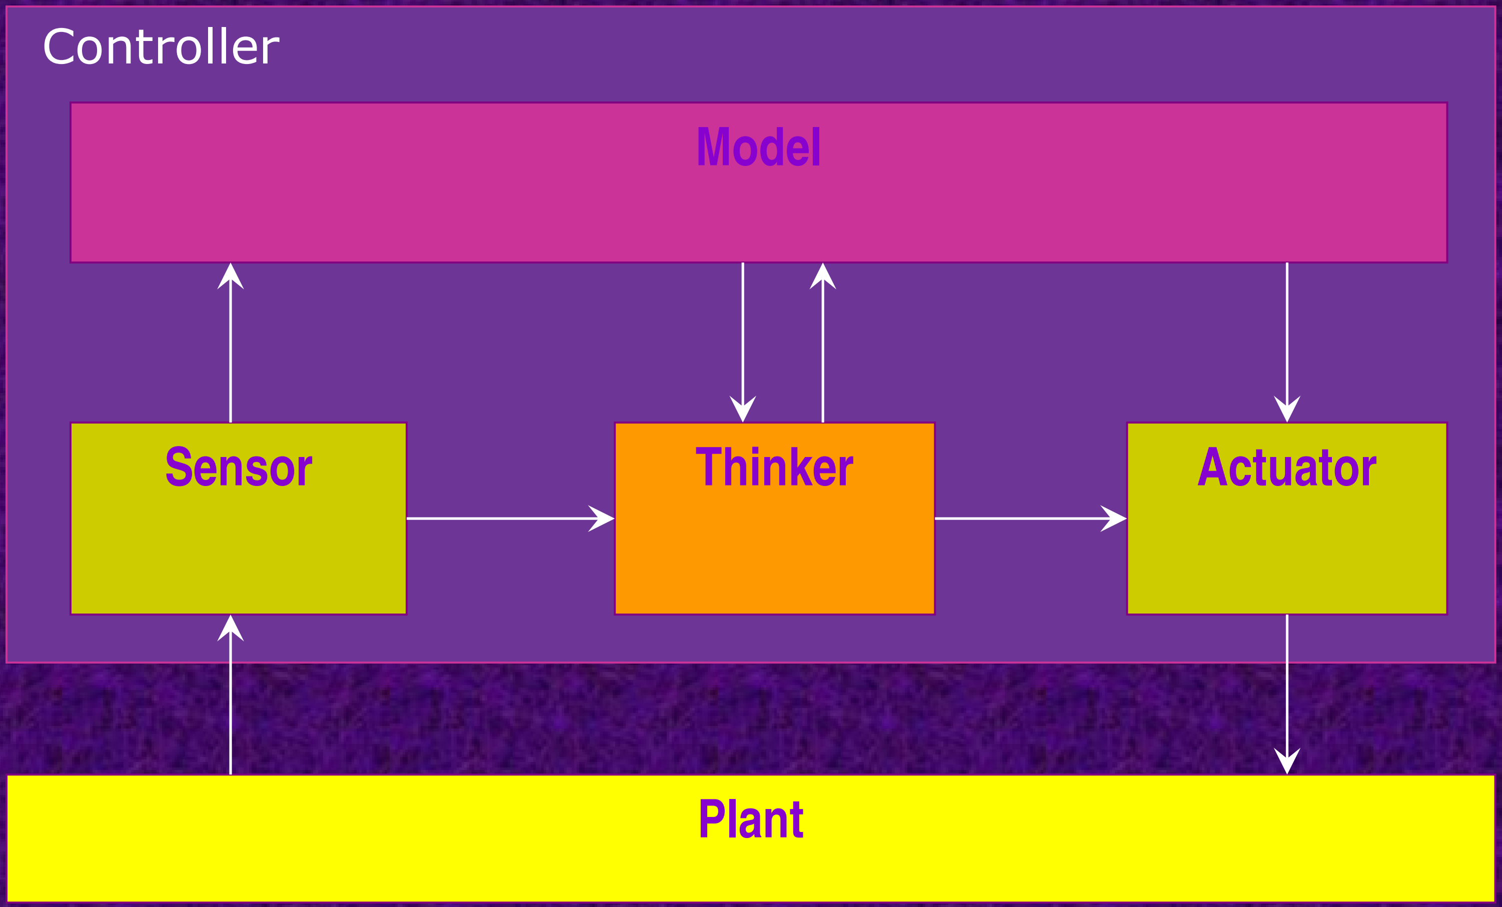

Is there a core architecture ?

- Feedback-feedforward loop?
- Meystel-Albus' elementary loop of functioning?
- Doyle's bowties?
- Autopoietic networks?
- Interactivist-constructivist systems?

- Why is it necessary ?
  - Homogeneity
  - Emergence

The Autopoietic Mind?

- The difference between code and data is contextual not ontological
- Control systems that generate themselves as products

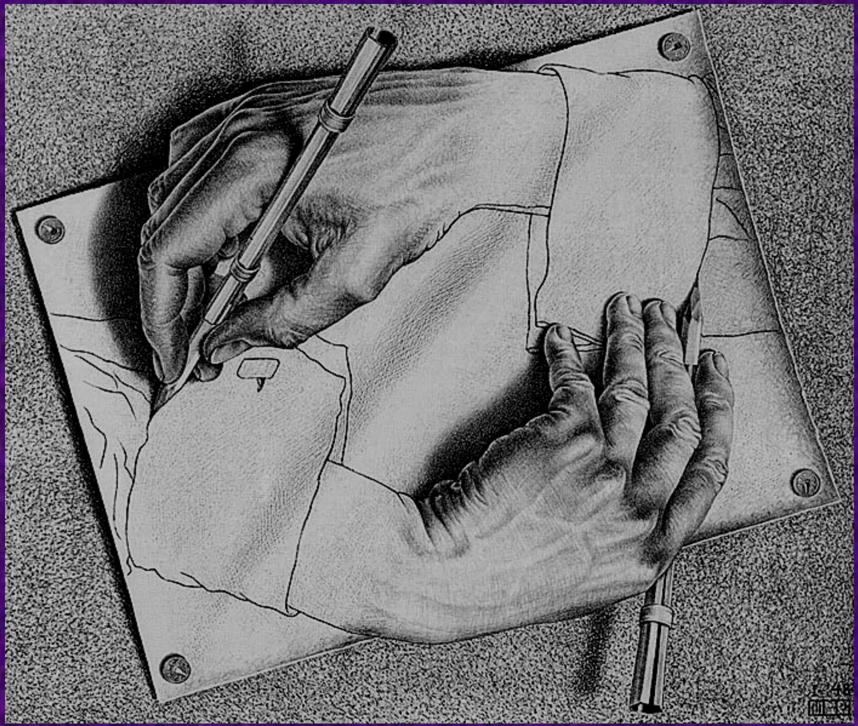

Theoretical foundations

- Control theory
- Dynamical systems
- Computational complexity
- Information theory
- Statistical physics
- Cognitive science
- Neuroscience (ground truth)

Largely fragmented within isolated technical disciplines.

Unified theory would be both intellectually satisfying and of enormous practical value.

[Doyle+]

Summary

- Mind is a multiresolutional phenomenon of model-based adaptive-predictive control
- Minds generate and use dynamic models of reality including their bodies
- Minds can be scaled and aggregated to for new minds
- At any resolution level, meaning generation generates awareness
- At any resolution level, mind reflection generates consciousness
- Scalable controllers should imitate? this basic design

Thanks

Questions ?

Abstract

As plants grow in complexity and control requirements grow in difficulty, there is an increased need for intelligence in control systems. Intelligence, traditionally a duty for central brains, has to be available where it is needed and that means, today, everywhere. This is achieved putting computing power and software in many places far from central brains. Intelligence is getting ubiquitous in control systems. Controllers are becoming societies.

Intelligence is distributed across the plants for several reasons (both needs and enablers): availability of suitable embeddable processors, timing requirements that forbid communication due to latencies, need of increased levels of performance that is achieved through parallelism, simplification of construction and maintenance tasks through modularity, reduction of cost and time-to-market by means of component-based reuse, etc.

Traditional control systems construction processes are based on proven methods that can lead to well-performing systems. Conventional control systems modules for distributed controllers can be built and are built this way. But there is a major problem when we address complex systems: traditional engineering methods for high-assurance, distributed or intelligent control systems do not scale well. High performance plant-wide distributed controllers are getting so complex that conventional engineering methodologies fail to achieve the degree of predictability needed. Obviously one of the reasons is that the tools used by the control engineers are not suitable for this task; but to my understanding the most critical problem is that conventional control architectures do not provide the degree of scalability of intelligence that is needed.

Much work has been done in complex controller understanding (see for example the work in learning control, real-time artificial intelligence, reflective controllers, multiresolutional hierarchical control or emergence in control agent societies) but this work does not properly address the issue of scalability.

In this talk I will address the current state of the art in the implementation of complex controllers with distributed intelligence; will describe the basic technologies, the trends and will propose research topics for future pervasive, highly intelligent, scalable controllers.