Can Artificial Intelligence meet the Cognitive Networking Challenge?

Dr. Karen Zita Haigh
Raytheon BBN Technologies
July 2011
khaigh@bbn.com
612-308-6726

Raytheon
BBN Technologies

Status Quo in Networking

Raytheon
BBN Technologies

- We have reached the limit of where Human Engineers can predict and design for all contingencies.
- We don't know
  - whose radios (politically) will be in the network,
  - which radios (hardware) will be in the network,
  - where they will be used,
  - by whom, or
  - for what tasks.

Radio setup can not be static or pre-scripted: we need cognitive radios

Requirements for Cognitive Networks

- A Cognitive Radio must have "smart" capabilities:
  - Situation Assessment: identify and forecast network conditions
  - Planning: adapt to constantly changing conditions while balancing the needs of many users
  - Learning: learn from previous experience

- Cognitive networks must operate cooperatively in a distributed, diverse environment.

A cognitive network must adapt to continuous changes rapidly, accurately, and automatically

A solution?

- AI techniques are perfectly positioned to solve the Cognitive Network problem:
  - dynamic,
  - distributed,
  - heterogeneous,
  - low-communication,
  - partially-observable,
  - high-latency,
  - multi-objective optimization

AI enables real-time, context-aware adaptivity

Terminology Caveats

Software-Defined Radio A radio whose hardware can be tweaked to do different things.	**Cognitive Radio** A radio whose software uses smart capabilities to decide what settings to use Needs an SDR as the platform
Adaptive A system whose behaviour changes based on current conditions Action = f (conditions)	**Learning** A system whose prior experience allows it to correct its adaptive models Change the internals of f

Outline

- What is AI?
  - A Sampling of AI Techniques
    - Roles for AI from AI's perspective
    - Techniques already demonstrated in Networking
- Roles for AI from Networking's perspective
  - Cognition Loop Functionality
  - Characteristics of Cognitive Networks
- A Concrete Example
- Architectural Considerations
- Implications

What is AI?

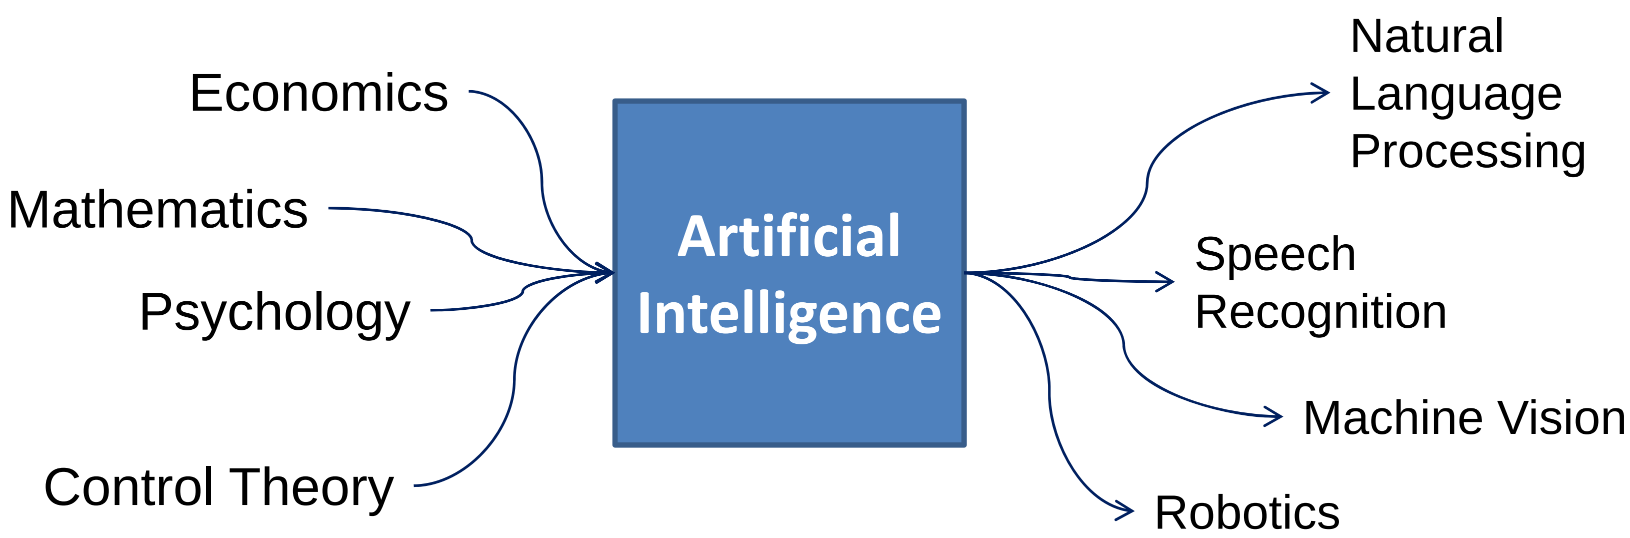

The Odd Paradox

Practical AI successes … were soon assimilated into whatever application domain they were found to be useful in, and became silent partners …, which left AI researchers to deal only with the failures.”

[McCorduck, 2004]

Roles for AI in Networking

Example Functions:

- Performance Analysis
- Traffic Analysis
- Network Configuration (which modules to use)
- Network Control (which parameter settings to use)
- Policy & Constraint Management
- Cyber Security

AI can adapt to:

- Changes in environment
  - Mobility
  - Jamming
- Changes in infrastructure
  - New nodes
  - Interop with legacy systems?
- Changes in user community and their needs

A Sampling of AI Techniques

Knowledge Engineering

- Captures knowledge so that a computer system can solve complex problems, e.g.
  - models of physics and signal propagation, constraints on the system, analysis of interactions, and rules of thumb (e.g., about how to configure the system).
- A formal ontology may help a cognitive system reason about how and when capabilities are interchangeable
  - e.g., QoS and QoI
- Knowledge bases can help optimize the network
  - e.g., By biasing a learning algorithm
  - e.g., By constraining a planner

Planning and Scheduling

- Organizes tasks to meet performance objectives under resource constraints
  - Multi-agent planning, dynamic programming, constraint satisfaction, and distributed or combinatorial optimization algorithms
- Planning and scheduling techniques in networks can decide what content to move, where, when, and how
  - Prefetch / prepush data
  - Power-aware computing
  - Node activity and task scheduling
  - Network management
  - Server placement; when to handle queries
  - Meeting FCC / ROE policy constraints

Raytheon BBN Technologies

Multi-Agent Systems

- Biologically-Inspired approaches are lightweight coordination mechanisms
  - Traditional MAS approaches fail in MANET because they assume that communications are (a) infinite and (b) always available
- Demonstrated Applications:
  - Routing: AntHocNet uses both proactive and reactive schemes to update the routing tables, and outperforms AODV.
  - Network connectivity
  - Dynamic load balancing
  - Service placement

Raytheon
BBN Technologies

Machine Learning

- ML improves the performance of a system by observing the environment and updating models
  - the learner must *generalize* so that the learned model is useful for new (previously unseen) situations.
    - Statistics don't generalize
  - Artificial neural networks, support vector machines, clustering, explanation-based learning, induction, reinforcement learning, genetic algorithms, nearest neighbour methods, and case-based learning.
- Demonstrated Applications
  - Routing
  - Energy management
  - Node mobility
  - Parameter interaction for optimization

OOPDAL: The Cognition Loop

The OODA Loop

AI Techniques
Cognition Loop
Characteristics

Raytheon BBN Technologies

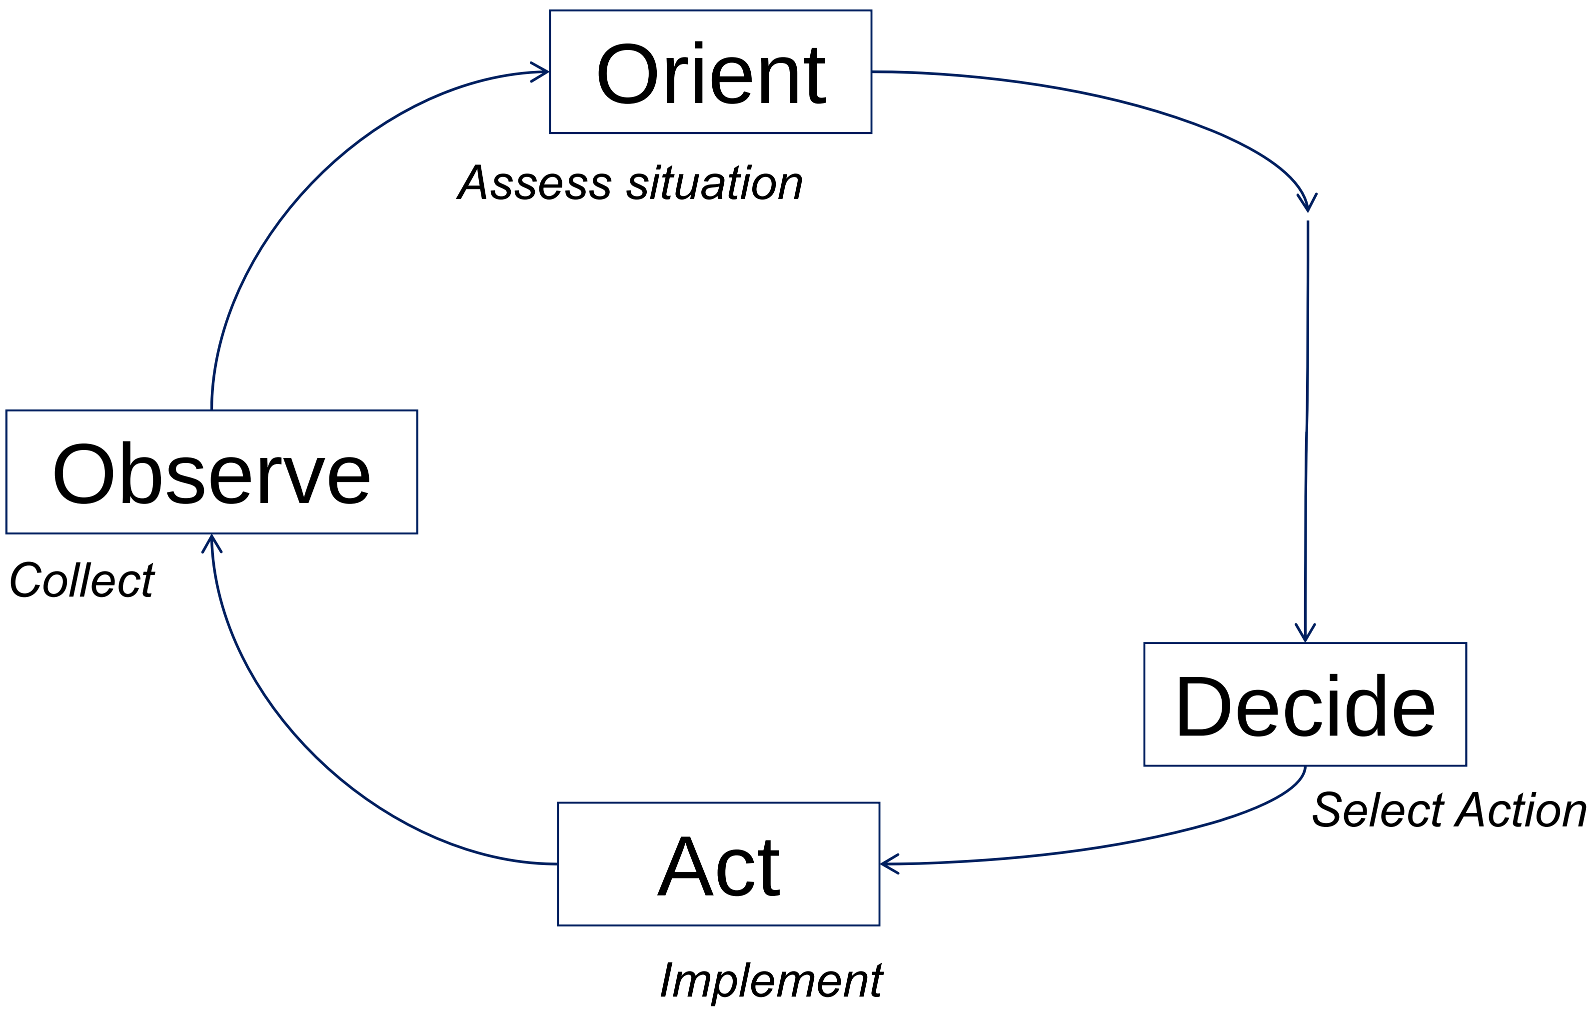

Raytheon BBN Technologies

The Cognition Loop

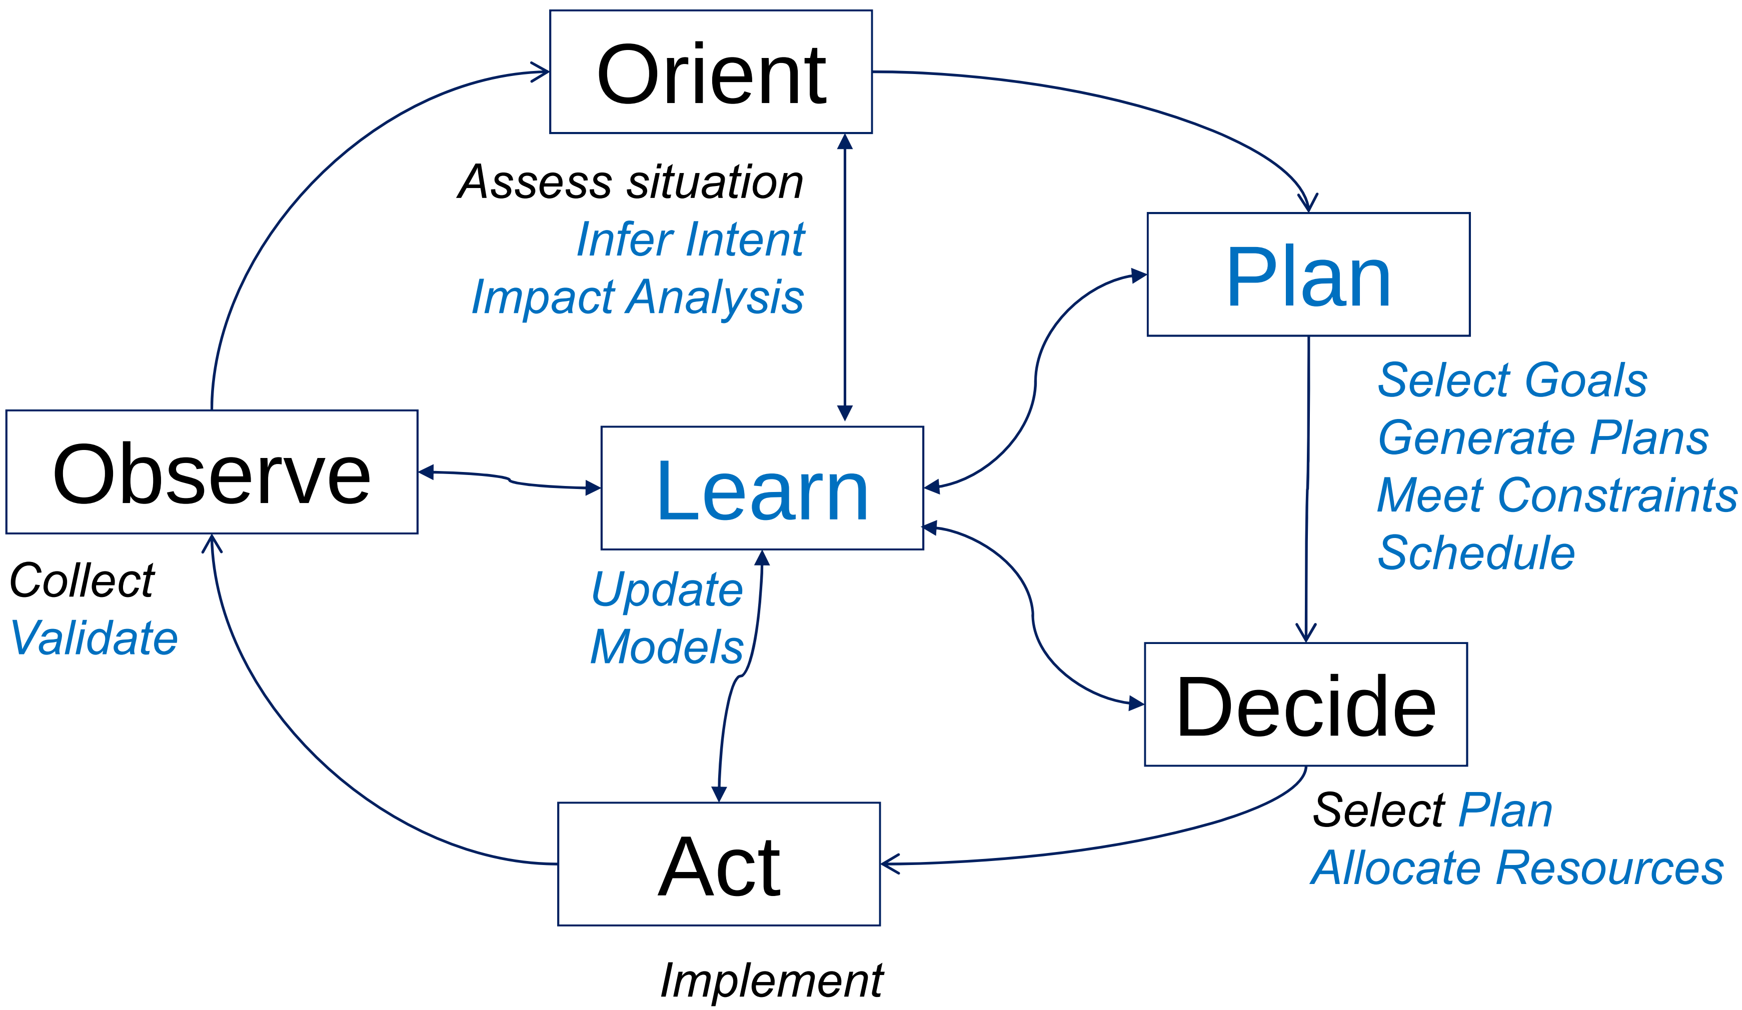

Raytheon
BBN Technologies

The relationship between MANET Characteristics & AI

MANET Characteristics

AI Techniques
Cognition Loop
Characteristics

Raytheon BBN Technologies

What AI is good at
(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI
(and Networking techniques don't handle)

Complex Temporal Feedback Loops
Complex Interactions
Heterogeneous Intercommunication

Raytheon BBN Technologies

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

Everything Changes in a MANET

- Situation Assessment quickly recognizes changes and their impact
- ML can update models and generalize from prior experiences
- Planning Under Uncertainty can select actions that leave flexibility

Raytheon BBN Technologies

MANET Characteristics

What AI is good at
(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI
(and Networking techniques don't handle)

Complex Temporal

Many factors cannot be measured or have multiple possible causes

- Situation Assessment can infer missing data or recognize ambiguous situations
- ML can generalize from related concepts
- Planning can choose to gather more information, or choose actions that cover multiple possibilities

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

Bandwidth, energy, and compute power are always restricted

- Situation Assessment can compute the impact of current conditions and the risk of taking (or not taking) corrective actions
- ML are very good at incremental approaches
- Planning can allocate tasks to available resources

(Most techniques have "anytime" or scalable versions)

In AI, Diversity is considered a benefit

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

Complex Temporal

Most decisions need to be made locally

- Situation Assessment can infer missing data or recognize ambiguous situations
- ML can estimate impact of local decisions on global network
- Planning can choose what information to share with the broader network

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

As many as 600 observable and 400 controllable parameters *per node per timestep*

- Situation Assessment can identify critical parameters
- ML has worked on massive datasets, and incremental approaches are very effective
- Planning can decompose actions hierarchically

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

Complex Temporal Feedback Loops

Diversity of data, nodes, and users may restrict communication patterns

- Situation Assessment can identify unexpected interactions
- ML can predict traffic needs
- Planning can manage and allocate tasks according to constraints

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective
  Performance Requirements

Main challenges for AI

The network should fairly support needs of multiple users (QoS, power, security, policies, connectivity, priority…) *even though they may be far away*

- Situation Assessment can match current conditions to mission requirements
- ML can estimate impact of local actions on global performance
- Planning can manage and allocate tasks according to constraints

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

At least 3 temporal loops
(1) MAC/PHY
(2) Routing to App layer
(3) Across nodes
Makes it harder to correlate cause & effect

AI has good techniques for correlating cause & effect, but not typically as rich as this

Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

- Complex Temporal Feedback Loops
- Complex Interactions
- Heterogeneous Intercommunication

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Pair-wise parameter interactions are extremely complex

Makes it harder to generalize patterns

AI will need to get better about hybrid learning approaches to bootstrap empirical approaches

Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

Complex Temporal Feedback Loops

Complex Interactions

Heterogeneous Intercommunication

MANET Characteristics

Raytheon BBN Technologies

What AI is good at

(and Networking techniques only partially handle)

Strong tradition that networks have static, homogeneous configurations (to reduce complexity)

Heterogeneity may cause non-interoperable conditions

AI assumes "safe" communications, and will need mechanisms for fallback

Main challenges for AI

(and Networking techniques don't handle)

- Complex Temporal Feedback Loops
- Complex Interactions
- Heterogeneous Intercommunication

Raytheon BBN Technologies

MANET Characteristics

What AI is good at

(and Networking techniques only partially handle)

Dynamic
Partially-Observable
Ambiguous Observations
Resource Constrained
Diverse
Discrete
Massive Scale
Complex Access Policies
Complex Multi-objective Performance Requirements

Main challenges for AI

(and Networking techniques don't handle)

Complex Temporal Feedback Loops
Complex Interactions
Heterogeneous Intercommunication

Union of Networking & AI may solve previously unsolvable problems

Machine Learning in ADROIT

CONCRETE EXAMPLE

Concrete Example: ML in ADROIT

- Adaptive Dynamic Radio Open-source Intelligent Team (ADROIT)
- Goal: Create cognitive radio teams that
  - Recognize that the situation has changed
  - Anticipates changes in networking needs
  - Learns about the operating environment
  - Adapts the network, in real-time, for improved performance
    - Real-time composability of the stack
    - Real-time control of parameters
    - On one node and across the network

ADROIT's Experimental Testbed

Raytheon
BBN Technologies

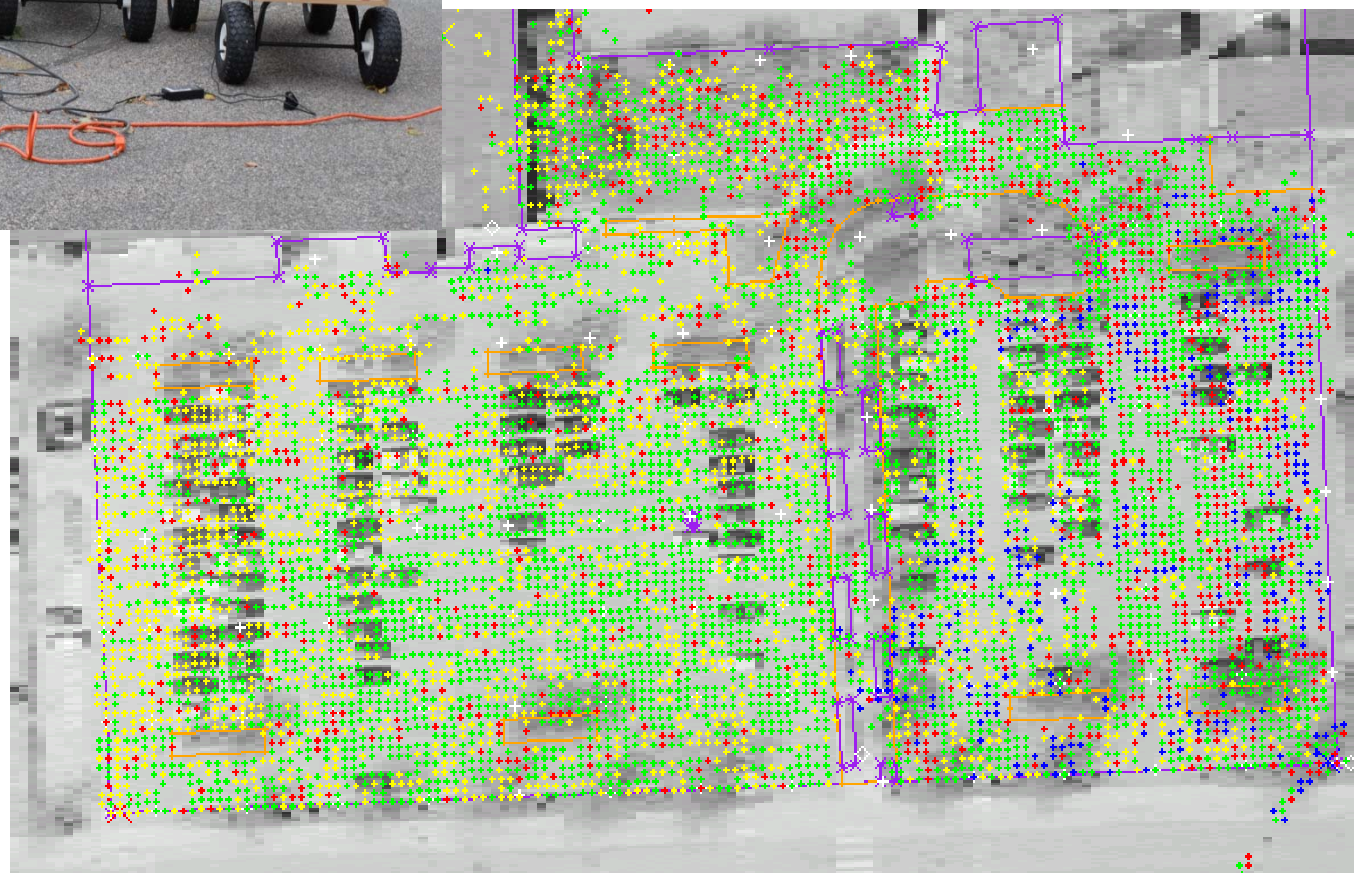

Maximize % of shared map of the environment

Experimental Results

Training Run:

- In first run nodes learn about environment
- Train neural nets with

  (Conditions,Strategy)➔Performance

  tuples
  - Every 5s, measure and record progress, conditions, & strategy
  - Observations are local, so each node learns different model !

Real-time learning run:

- In second run, nodes adapt behaviour to perform better.
- Adapt each minute by changing strategy according to current conditions

Real-time cognitive control of a real-world wireless network

Observations from Learning

Raytheon BBN Technologies

System performed better with learning

- Selected configurations *explainable* but not *predictable*
  - Farthest-refraining was usually better
    - congestion, not loss dominated
  - Unicast/Multicast was far more complex
    - close: unicast wins (high data rates)
    - medium: multicast wins (sharing gain)
    - far: unicast wins (reliability)

[Troxel2007] Troxel, Caro, Castineyra, Goffee, Haigh, Hussain, Kawadia, Rubel and Wiggins, http://www.cs.cmu.edu/~khaigh/papers/troxel07-globecom.pdf

Knowledge Transfer

Raytheon BBN Technologies

- Can we transfer learned knowledge from one environment to a totally new one?

- **Hybrid Learning** combines a priori engineering knowledge with emipirical data
  - Analytical Models: a priori models of network behaviour
    - Capture useful general principles
    - But are incomplete, incorrect, and static
  - Traditional Machine Learning: empirical models built from experience
    - Capture actual operating conditions
    - But poorly transfer knowledge to new domains or objective functions

- Simulation Results show learning is more effective than traditional config (static homogeneous) and smart human red team for
  - New Comms Environment
  - New Traffic Patterns

	Throughput	% Improvement
Static Homogeneous	2,449,002,560	-
Human Red Team	2,910,847,040	119%
Basic Learner	3,061,248,416	125%
Hybrid Learner	3,944,319,424	161%

[Haigh2006] Haigh, Varadarajan and Tang, http://www.cs.cmu.edu/~khaigh/papers/haigh06a-configuration.pdf

Building a cognitively-controllable system

SOFTWARE ARCHITECTURE

A Need for Restructuring

- SDR gives opportunity to create highly-adaptable systems, BUT
  - They usually require network experts to exploit the capabilities!
  - They usually rely on module APIs that are carefully designed to expose each parameter separately.
- This approach is not maintainable
  - e.g. as protocols are redesigned or new parameters are exposed.
- This approach is not amenable to real-time cognitive control
  - Hard to upgrade
  - Conflicts between module & AI

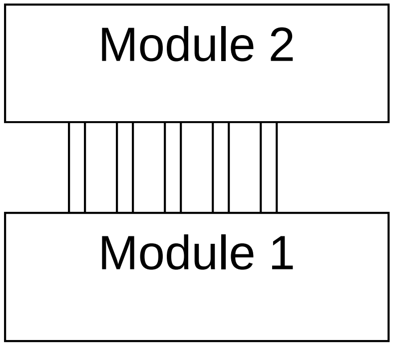

A Need for Restructuring

- We need one consistent, generic, interface for all modules to expose their parameters and dependencies.

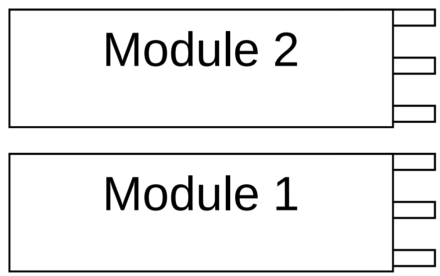

A Generic Network Architecture

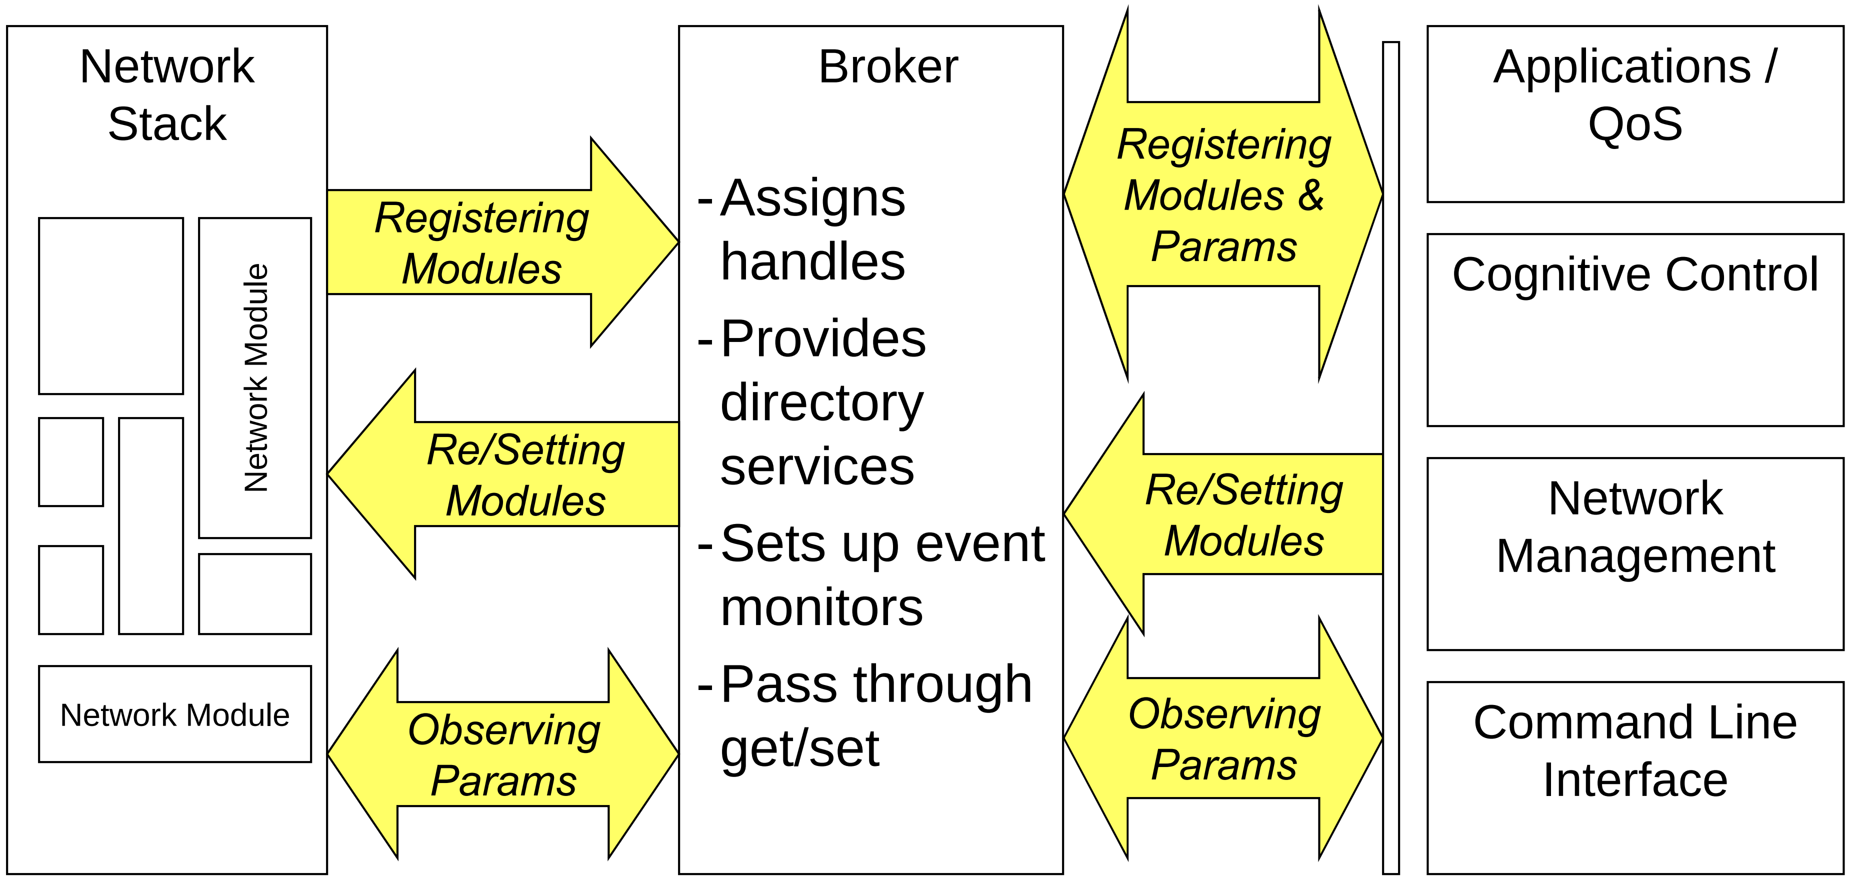

exposeParameter(*parameter_name*, *parameter_properties*)
setValue(*parameter_handle*, *parameter_value*)
getValue(*parameter_handle*)

Benefits of a Generic Architecture

Raytheon BBN Technologies

- It supports network architecture design & maintenance
  - Solves the *nxm* problem (upgrades or replacements of network modules)
  - Solves the problem of coordination of control (multiple controllers)
- It doesn't restrict the form of cognition
  - Open to just about any form of cognition you can imagine
  - Supports *multiple* forms of cognition on each node
  - Supports *different* forms across nodes
- It doesn't *mandate* cognition

IMPLICATIONS

Benefits & Scope of Cross-layer Design

- Much broader scope than previously imagined
  - More than 2 layers!
  - More than 2-3 parameters per layer
- Potentially gives much greater benefit
  - ➢ Drill-down walkthroughs highlighted benefits and explained restrictions
  - ➢ Simulation results for specific scenarios demonstrate the power

- Traditional network design includes adaptation
  - This works against cognition: it is hard to manage *global* scope
  - But network module may be better at doing something focussed

  - ➢ Design must include *constraining* how a protocol adapts

Cross-Layer Design on Steroids

Heterogeneous Nodes

- Local conditions may create a heterogeneous network
  - Different configurations depending on current (local) conditions
  - May result in *non-interoperability* of nodes

  - “Orderwire” bootstrap channel as backup
  - “Coordination Manager” that decides what information to share

- Traditional network design expects static homogeneous configurations
  - Can be written to ensure interoperability
  - But cannot guarantee maximal performance
- Traditional AI always assumes safe communications

Capability Boundaries

- Blurring of “controller” and “controllee”
  - Any client can expose parameters
  - Any client can control parameters
  - e.g., MAC can set video compression rate

- Traditional Networking has very clear boundary between “network” and “application”
  - Often user/kernel boundary

➢ Take care that the additional complexity does not yield an unreliable system
  - Priority control
  - “Failsafe” mode

Conclusion

Raytheon BBN Technologies

- **AI** techniques are ready to be challenged with this complex real-world domain
- **Networking** requirements are reaching the limits of what can be done without AI.

- A union will overcome the most significant challenges that remain for cognitive networking.

- A Cognitive Radio must have "smart" capabilities:
  - Situation Assessment
  - Planning
  - Learning
- AI enables a Cognitive Network to manage the complex mix of
  - whose radios (politically),
  - which radios (hardware),
  - locations,
  - users, and
  - tasks.

A cognitive network trained in a desert can learn how to perform well under water.