NIMD: A Hybrid Form

Agent Based Modeling and Simulation on Rocks/CUDA Clusters

Kenneth A. Lloyd
Director – Systems Science

WATT SYSTEMS TECHNOLOGIES INC.

Viewpoints – Then and Now

This presentation is based upon a hybrid of Naturally Inspired Computation, of which Particle Swarm, Ant Colony Optimization, Clonal Selection Algorithms, and Mind Evolution Computing are strategies.

The concept of a “supercomputer” has evolved from the mainframe, through homogeneous compute clusters (Beowulf), through hybrid compute clusters (Rocks/CUDA).

“It would appear that we have reached the limits of what it is possible to achieve with computer technology, although one should be careful with such statements, as they tend to sound pretty silly in 5 years.“

-- John Von Neumann, circa 1949

Now using “intelligent compute agents” we can attack not only traditional problems with amazing speed, but also non-analytical, non-deterministic problems by using evolutionary neurocomputation – a genetically inspired form of Mind Evolution Computing.

Presentation Outline

- ABM&S on hybrid, heterogeneous clusters.
- Derivation of agents from UML / SysML.
- Description contexts, applied to ExecutionEnv.
- Evolution of behaviors using MEC TWEANNs.
- Creating the model and the experiment.
- Running the simulation.
- Evaluation. What are solutions?
- An application and implementation of the application and implementation.

Hybrid, Heterogeneous Clusters

- Hybrid
  - The computation is run not only on the CPU, but uses one or more GPUs per CPU as MPP math processors. (up to 1 GPU / CPU Core*)
- Heterogeneous
  - The computation is distributed among the cluster nodes in a manner that allows the most efficient and effective computation on that fabric, like a Jacquard pattern of threaded warps and wefts.

Hybrid, Heterogeneous Clusters

- The coupling between hardware, software and data structure is increasingly close and complex.
- The resulting coupled HDW SFW is referred to as a “compute fabric”.
- Massively parallel computation is non-deterministic (asynchronous) and may be recursive or recurrent.

- Need to resolve the Separation of Concerns between tasks and data structure, given the compute fabric configuration, schedule, priorities (We argue this should be intelligently automated)

Hardware Fabric Component

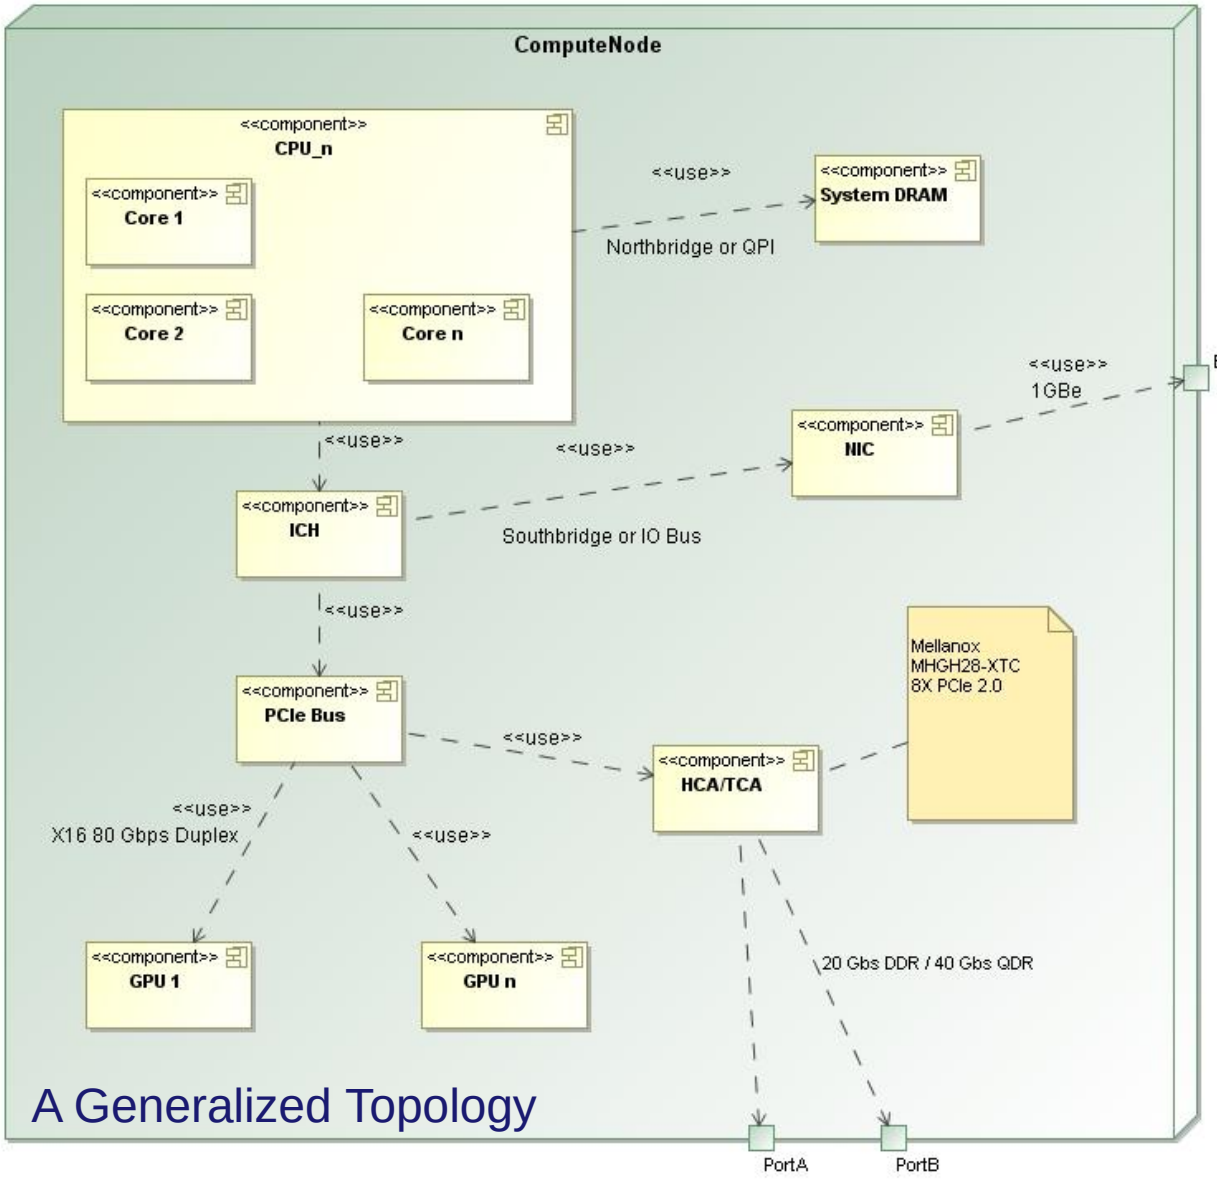

A Generalized Topology

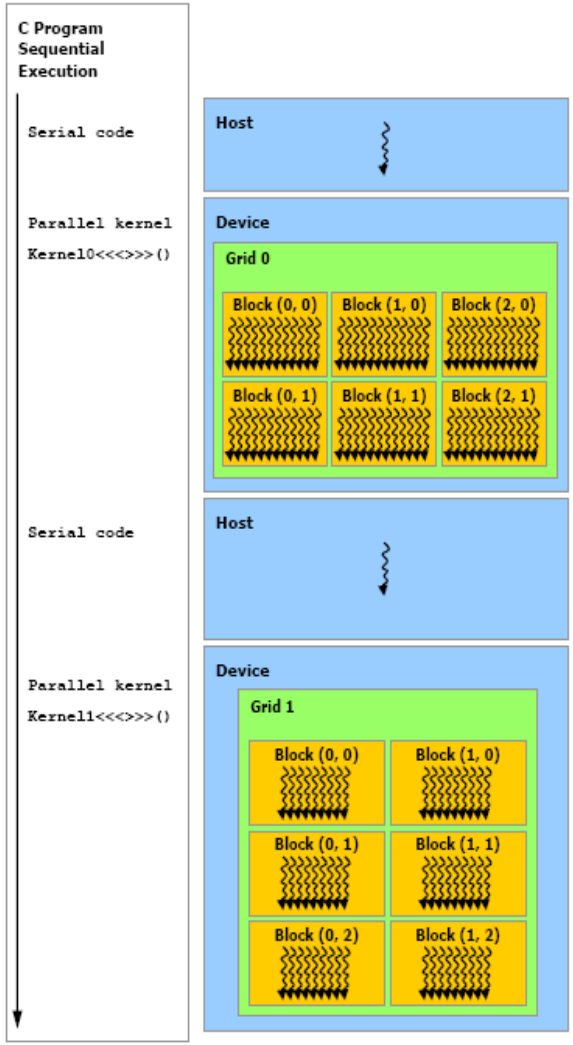

From CUDA Programming Guide 2.0
June 2008 – NVidia Corporation

CPU / GPU Pipeline

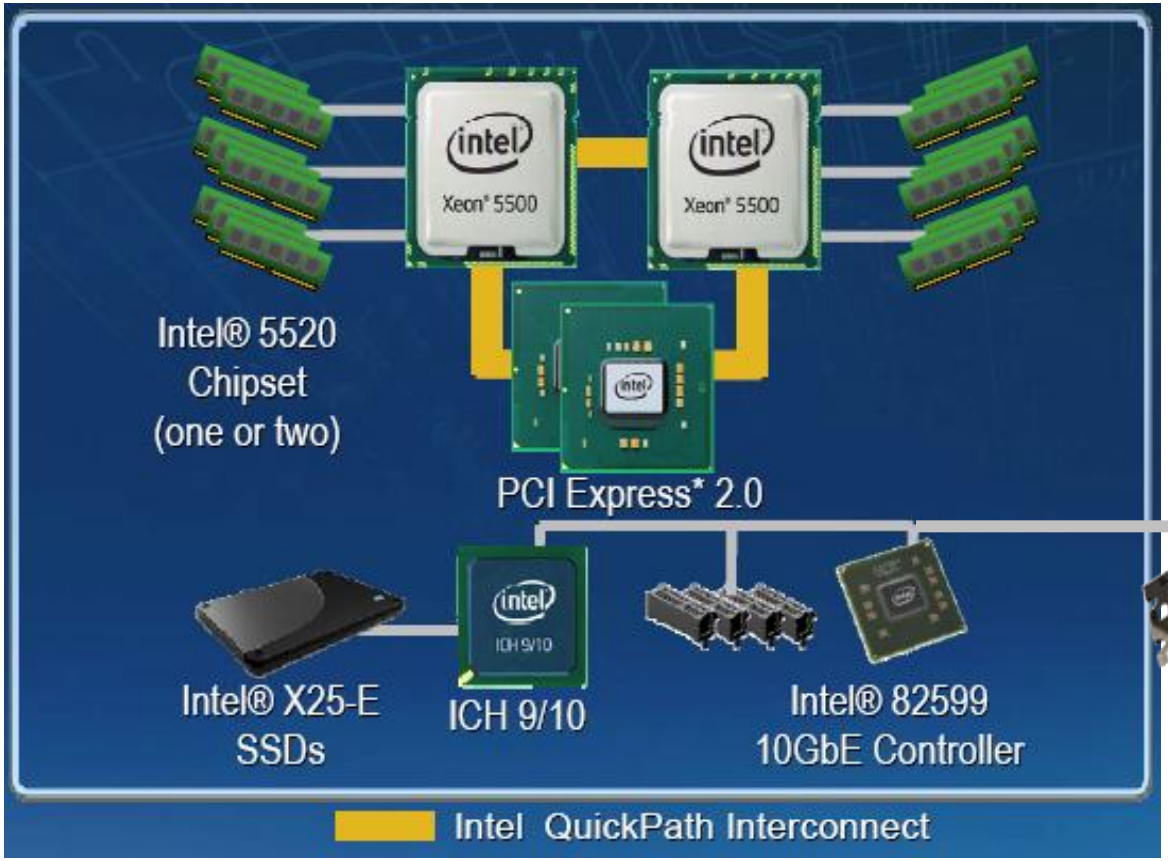

NVidia's GeForce GTX 295
480 Process Cores
~1788 GFLOPS*
(1.7 TFLOPS MAX EACH)

PCI Express 2.0 offers 4GB/s of peak bandwidth per direction for a x16 link and 8 GB/s concurrent bandwidth (duplex). ICH = I/O Control Hub

Xeon 5500 Graphic: Intel Corporation
GTX Graphic :NVidia Corporation

PCI Express Connection of Devices

There is communication between Host (CPU core) and Device(s) through The PCIe bus.

There is also communication between Devices.

The HCA or TCA is a Device on the PCIe Bus. This is where things Get interesting!

InfiniBand™ packets carry messages between Nodes and route them on the PCIe bus through the bus controller, NOT necessarily the CPU.

Graphic courtesy:
http://computer.howstuffworks.com/pci-express2.htm

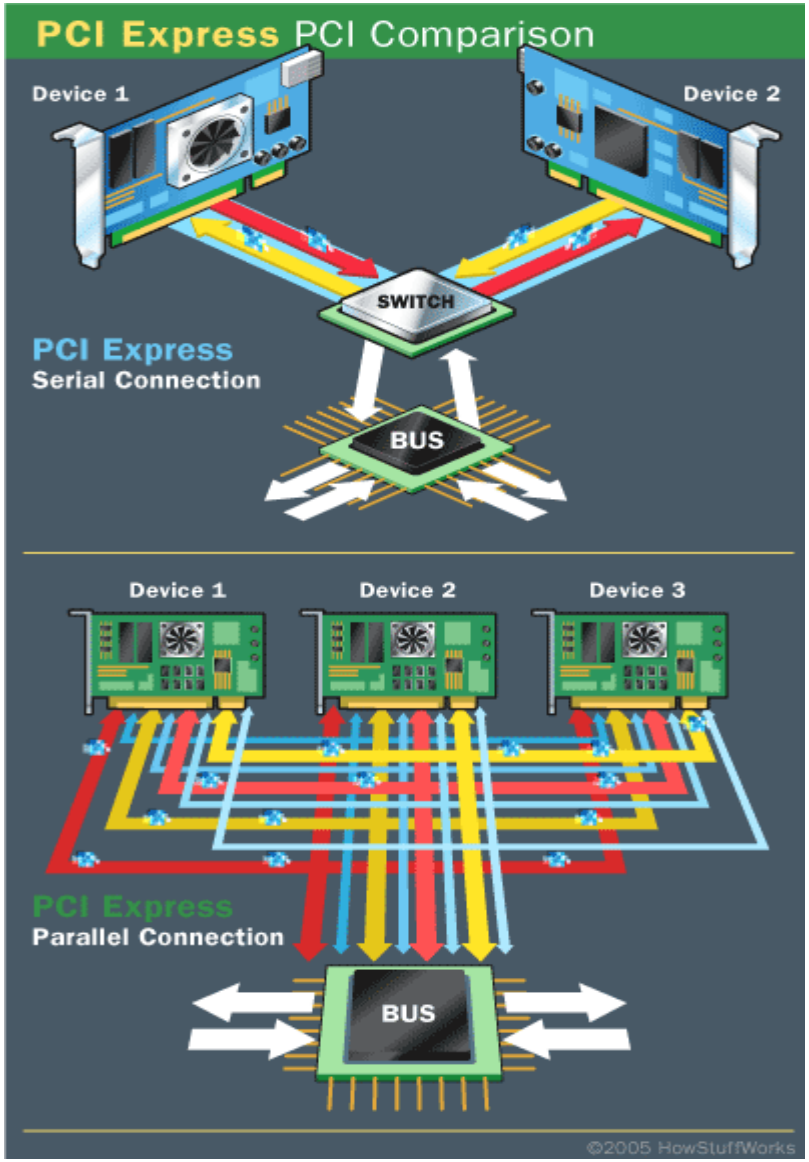

Hybrid, Heterogeneous Cluster HDW

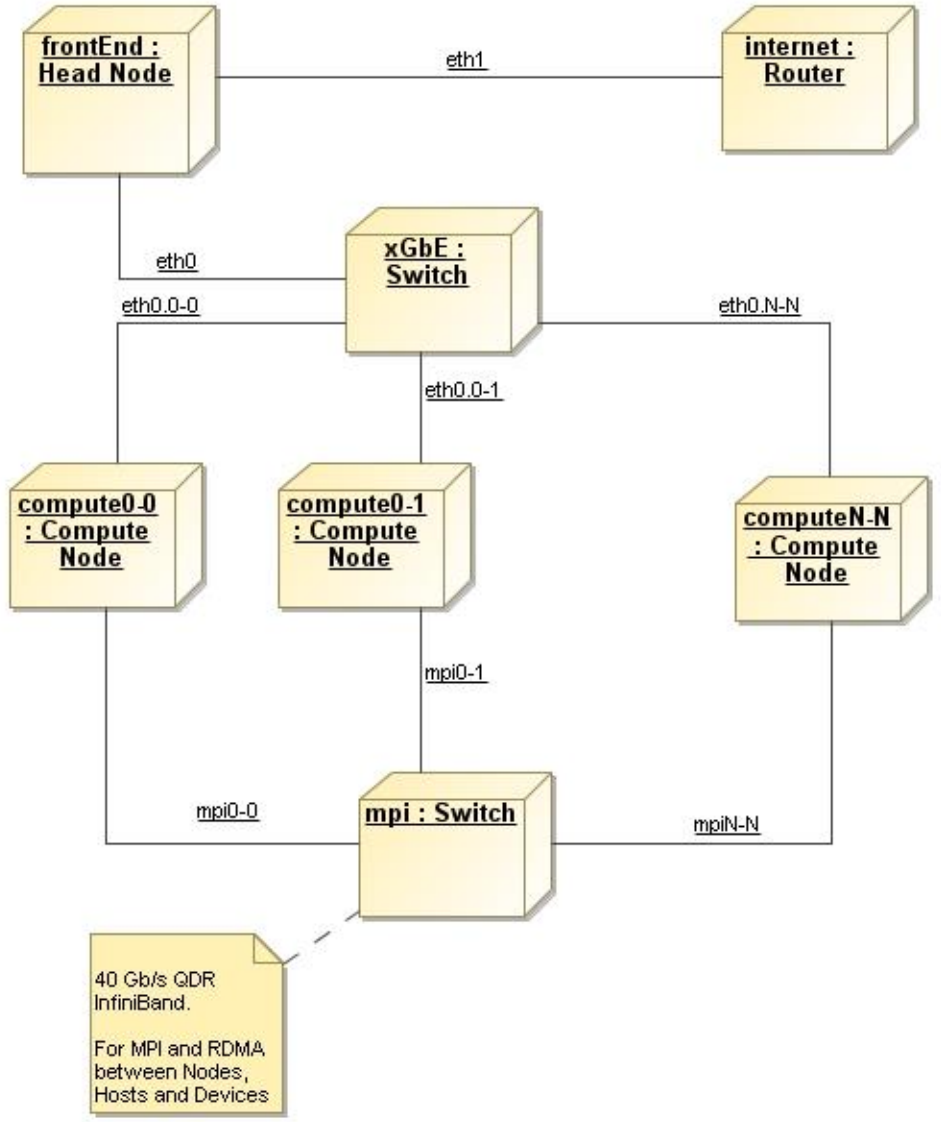

The Client passes simulation contexts, data and the models to be simulated to the cluster.

The actual simulation app is distributed between the compute nodes, then further distributed between the CPU cores and the GPU process cores (the kernel) to be executed on that compute node.

Data is partitioned (and passed) between host and device Memory (traditionally).

A Familiar Methodology

- 1. Set up the physical system to be studied.
- 2. Set up and calibrate the instrumentation
- 3. Run the experimental system and record the outputs of the instrumentation.
- 4. Analyze results.
- 5. Repeat 1 - 4 as necessary.
- 6. Publish, Tenure, Fame, Fortune …

From: Documentation Set for Swarm 2.2

Methodology Implementation

- Models will be developed from a minimal system and a series of contexts and objectives in which the system exists. This static, state-based approach is familiar to most ABM-ers.
- Models will be developed in UML and/or SysML and generated by Factory Classes to hybrid CUDA from XMI.
- “Therefore, a very important aspect of setting up an experiment in a computer is how one weaves the multiple threads of [in] time that must be woven together coherently in order to produce reliable, repeatable results. Much of our work on Swarm has been devoted to not only making the task of managing concurrency manageable, but towards mechanisms to make people aware that they are always making implicit assumptions about how multiple threads of time are interacting with one another …” [from Swarm Development 2.2]
- Time permitting, we will answer questions re: WattSysML, an extension of SysML for Agents in both Temporal and Environmental Contexts.

Derivation of Agents in SysML

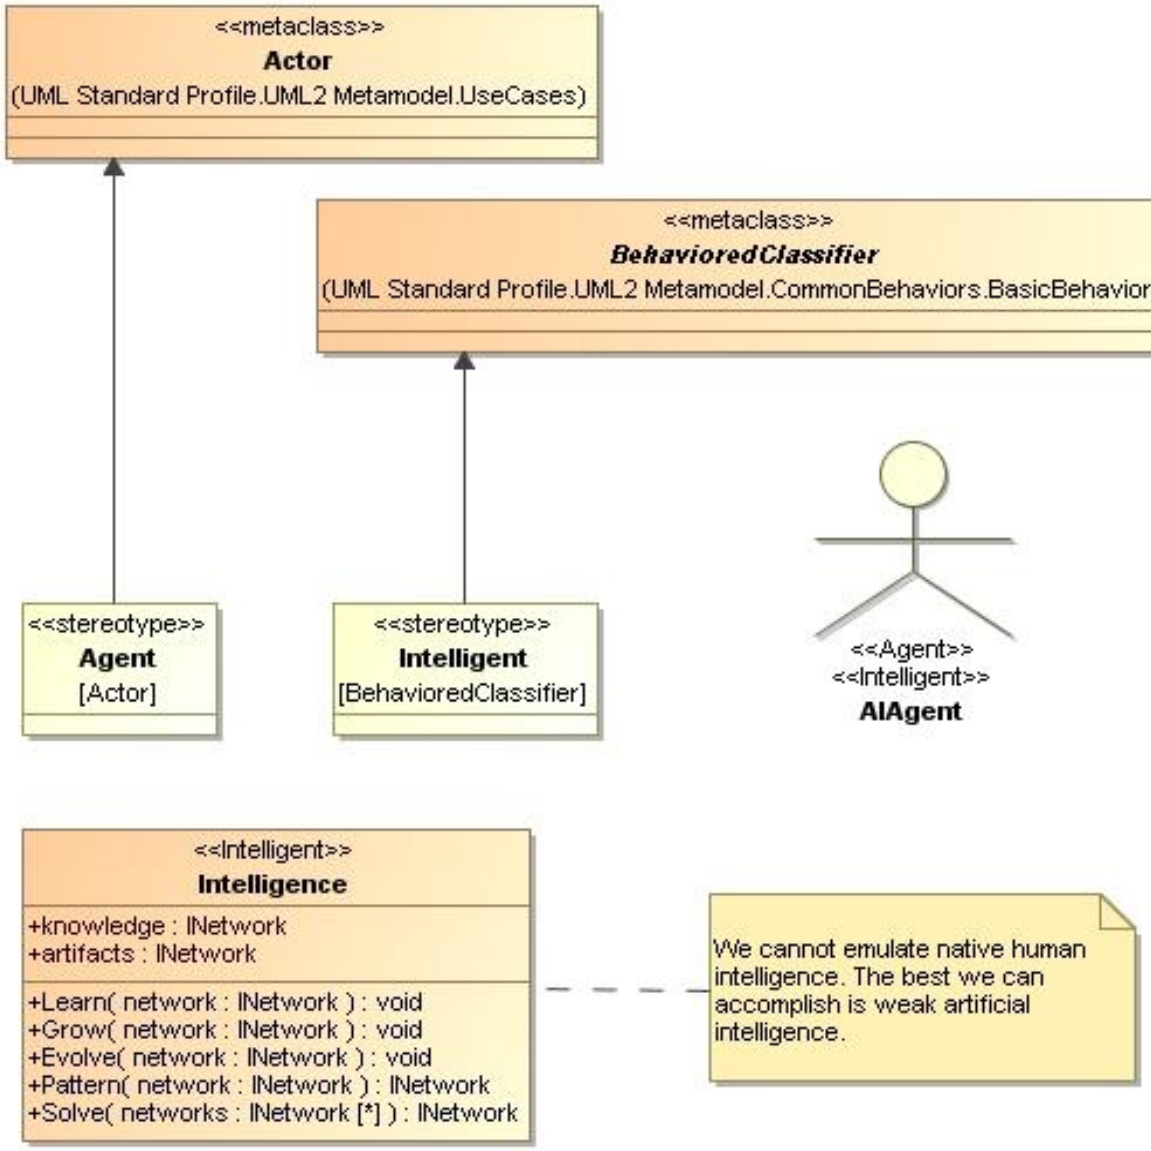

Derivation of Agents

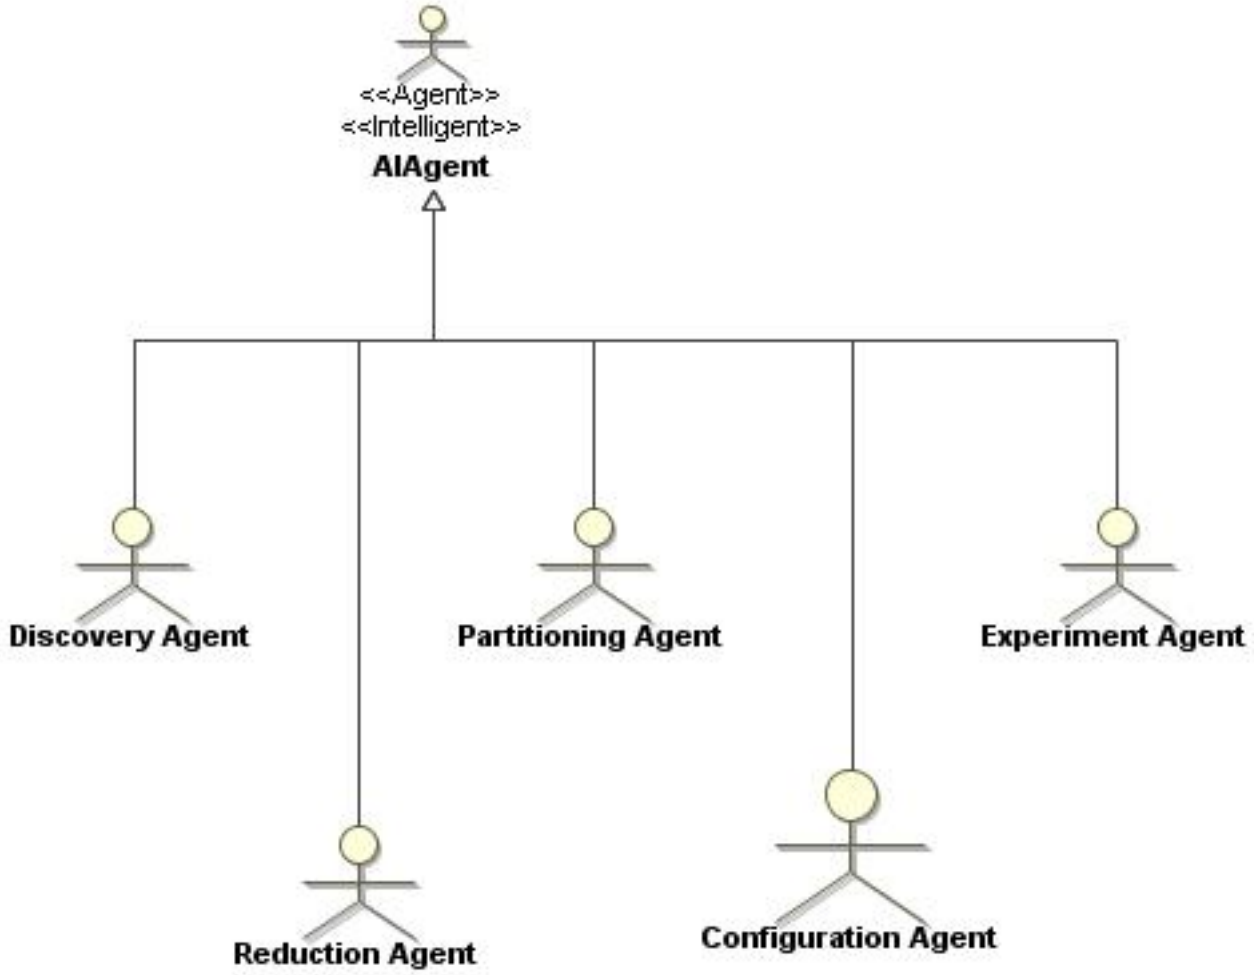

A Model World

“In a computer, you don't just drag the pieces of your experiment in from the outside world and hook them up. You have to create a world with space and time, a bunch of objects in that world …, schedules of events over those objects, all sorts of computer widgetry to interact with that artificial world and to manage multiple experimental runs and the data that they generate, and so forth. In other words, in a computer, one usually has to first *create* from scratch all of the bits and pieces of the experimental setup … “

-- Swarm Development Guide 2.2

Our Model Based Approach …
But, that is exactly what we are going to do … Drag things into a context, being a hybrid model execution environment and a SysML context .

We create a minimal system and drop it into some context, then evolve agents that learn how best to accomplish the objectives (the requirements).

Models in Contexts

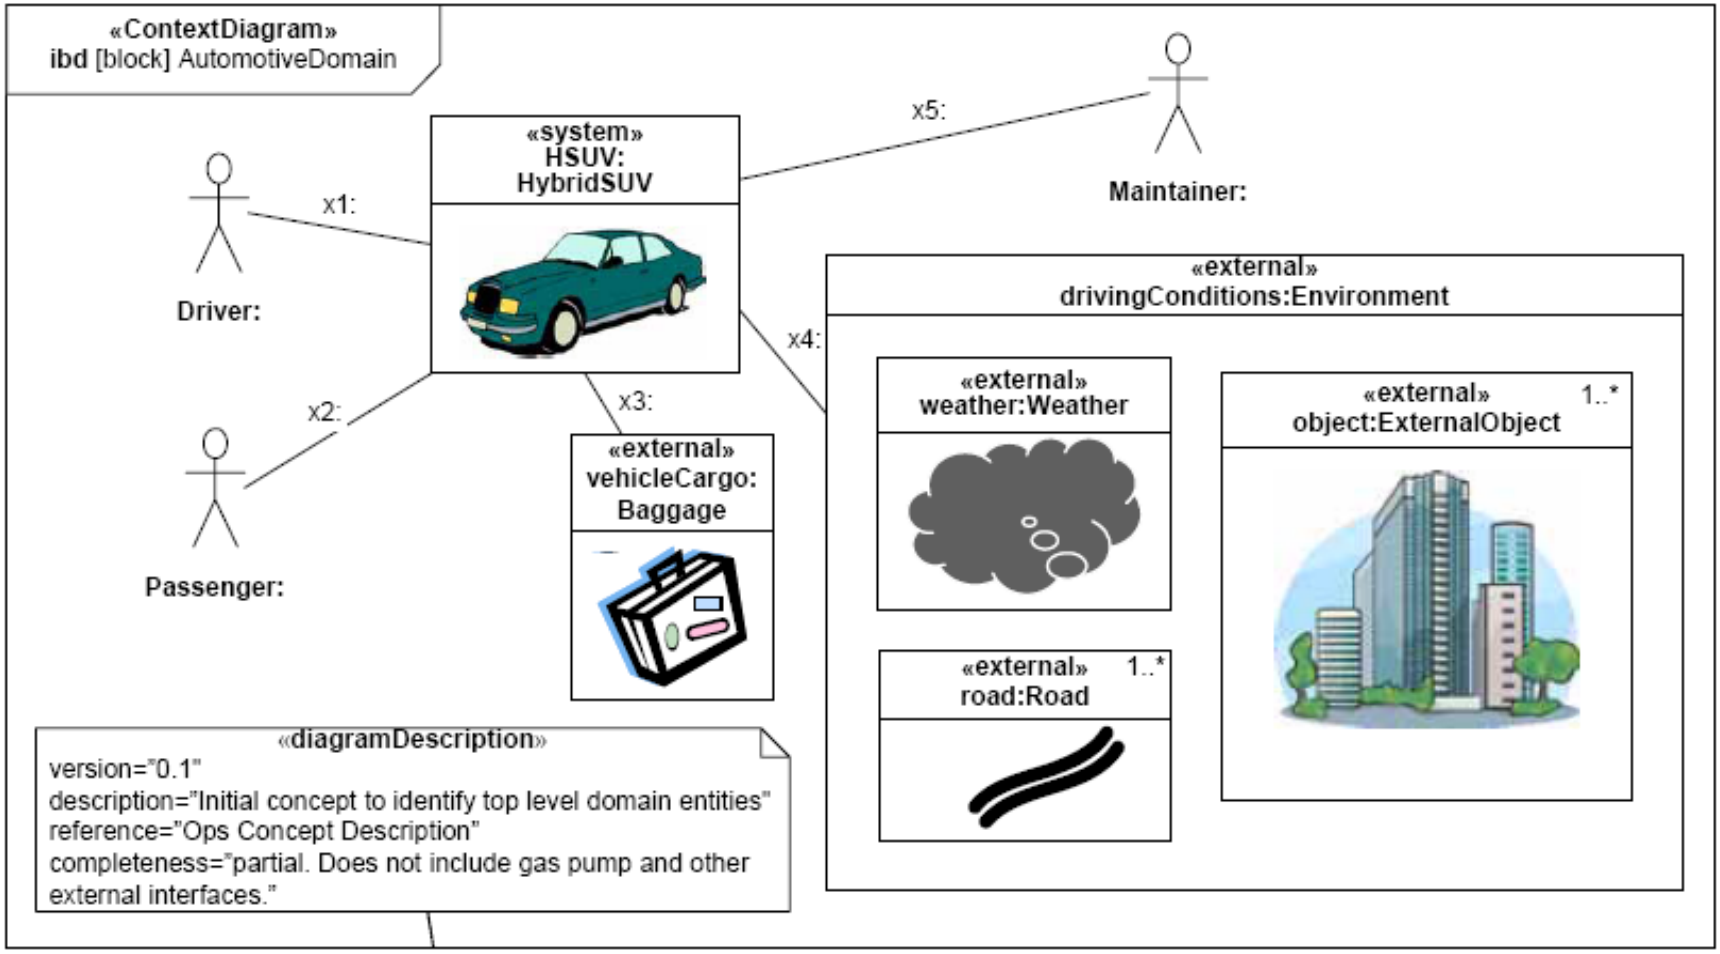

A System Context element is a virtual container that includes the entire system and its actors.

Graphic courtesy of OMG SysML Modeling Language formal/2008-11-01

Executing Models

“ExecutionEnvironment instances are assigned to node instances by using composite associations between nodes and ExecutionEnvironments, where the ExecutionEnvironment plays the role of the part. ExecutionEnvironments can be nested (e.g., a database ExecutionEnvironment may be nested in an operating system ExecutionEnvironment). Components of the appropriate type are then deployed to specific ExecutionEnvironment nodes. Typical examples of standard ExecutionEnvironments that specific profiles might define stereotypes for are «OS», «workflow engine», «database system», and «J2EE container». An ExecutionEnvironment can optionally have an explicit interface of system level services that can be called by the deployed elements, in those cases where the modeler wants to make the ExecutionEnvironment software execution environment services explicit.”

-- UML Superstructure v2.2

A Trivial Model, Trivial Experiment

- Simple, double pole balancing.
- Context is gravity here on earth.
- Send out swarms of agents to evolve solutions from minimal structure.
- Pick a fit solution
  - Is the most fit solution the best?

Models (on a napkin)

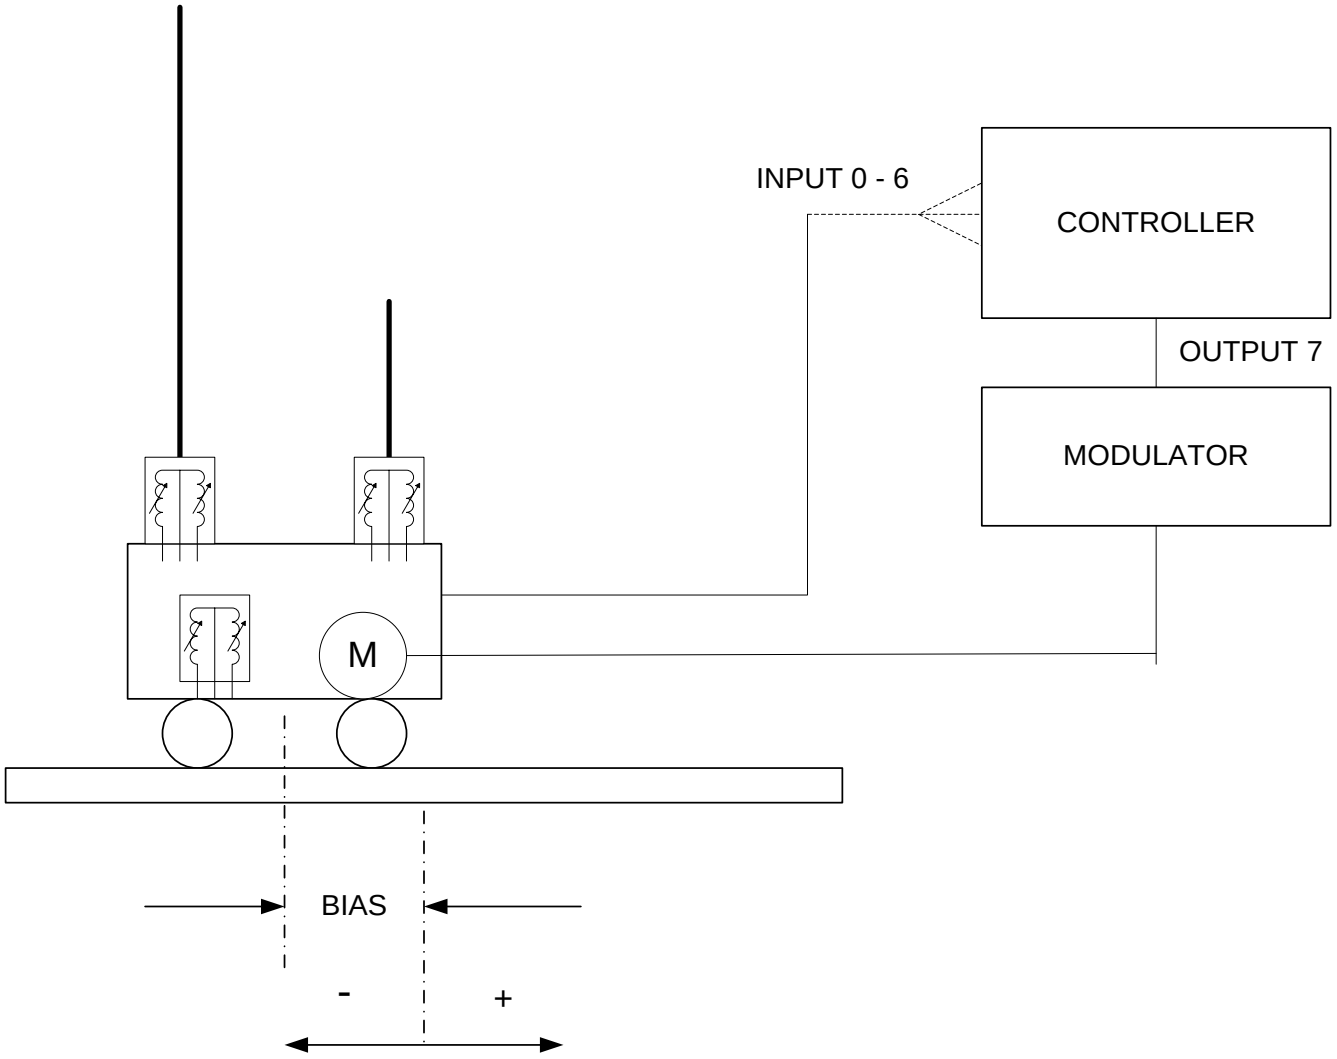

Models in SysML

<<requirement>>
Balance

Id = " Req1"
Text = "Balance unequal poles for 5 minutes "

<<requirement>>
Pole limits

Id = "Req2"
Text = "Keep poles within +/- 36 degrees from vertical"

<<requirement>>
Track limits

Id = "Req3"
Text = "Keep cart on track"

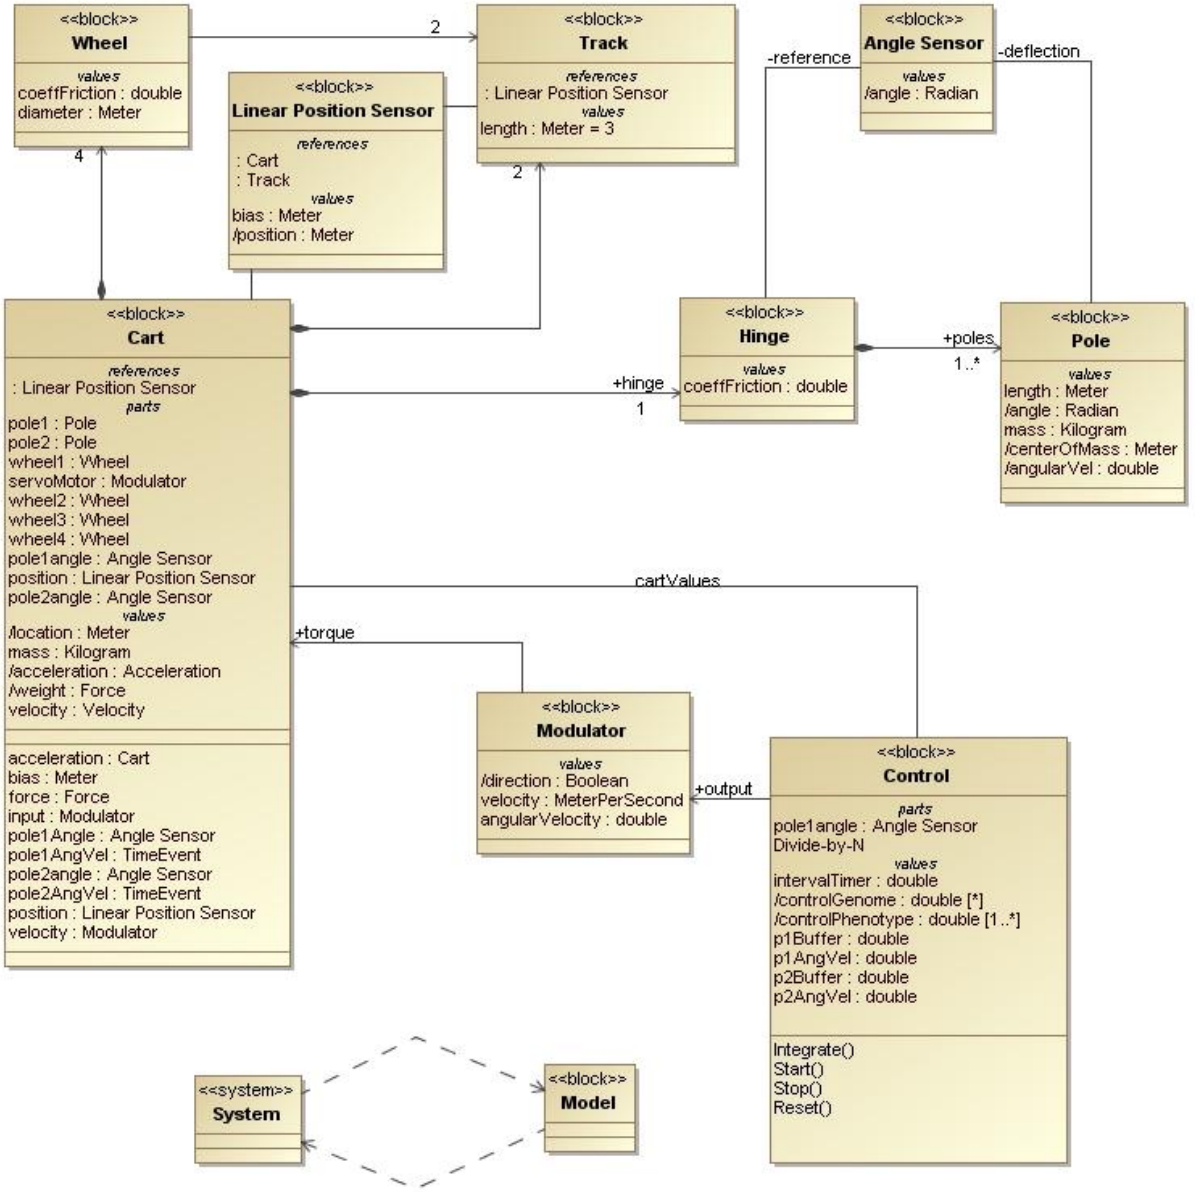

States and Parameters

```
/// <summary>
/// [0] - Cart Offset from Track C/L (bias)
/// [1] - Cart Diff Position (meters).
/// [2] - Cart Velocity (m/s) - derived.
/// [3] - Pole 1 Angle (radians)
/// [4] - Pole 1 Angular velocity (radians/sec) - derived.
/// [5] - Pole 2 Angle (radians)
/// [6] - Pole 2 Angular velocity (radians/sec) - derived.
/// </summary>
```

Sym.	Description	Value
x	Position of cart on track	[-2.4,2.4] m
θ	Angle of pole from vertical	[-36,36] deg.
F	Force applied to cart	[-10,10] N
l_i	Half length of i^{th} pole	l_1 = 0.5m l_2 = 0.05m
M	Mass of cart	1.0 kg
m_i	Mass of i^{th} pole	m_1 = 0.1 kg & m_2 = 0.01 kg
μ_c	Coefficient of friction of cart on track	0.0005
μ_p	Coefficient of friction if i^{th} pole's hinge	0.000002

```
protected float[] state = new float[7];
```

Contexts – Using A New Profile

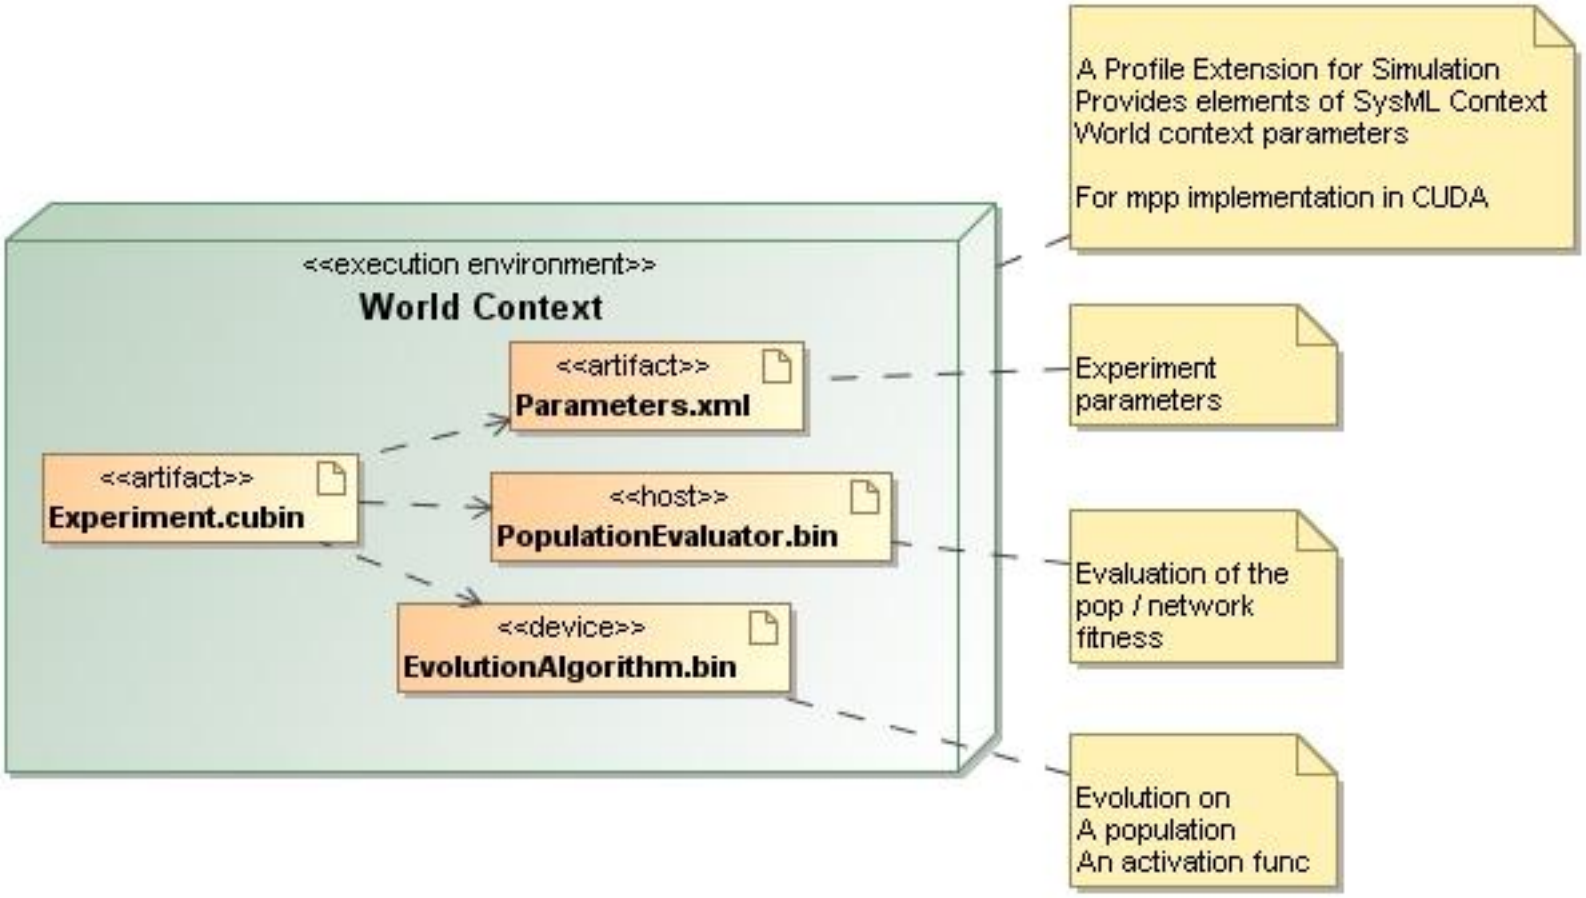

The Experiment

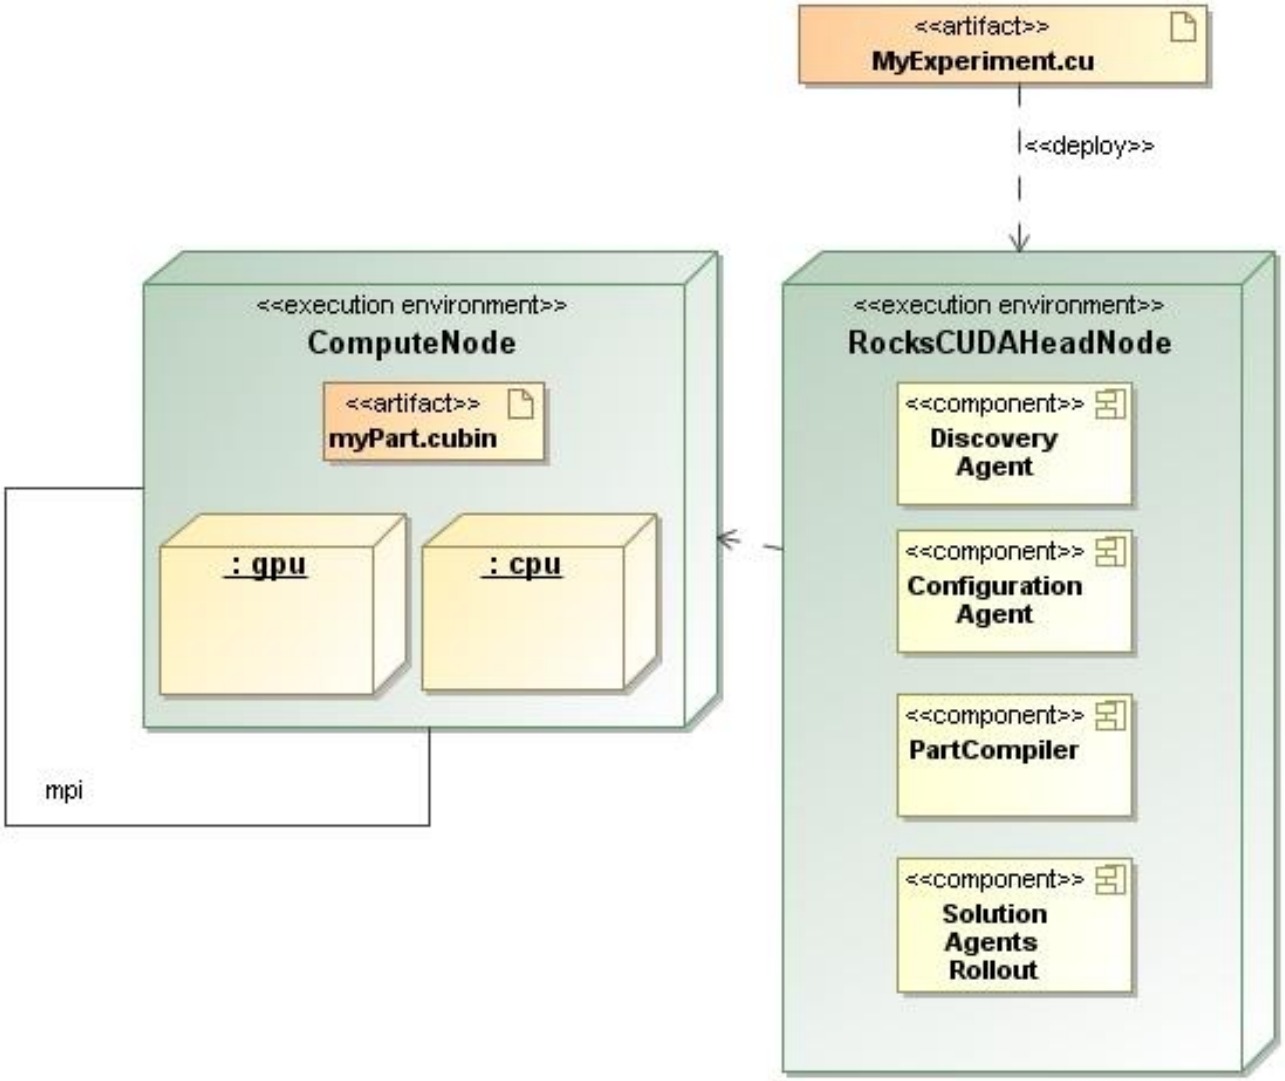

MPP Models in Superposition

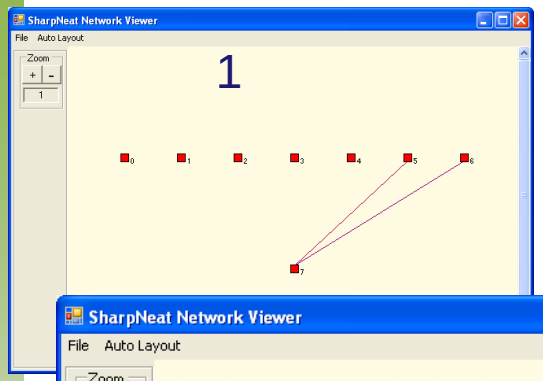

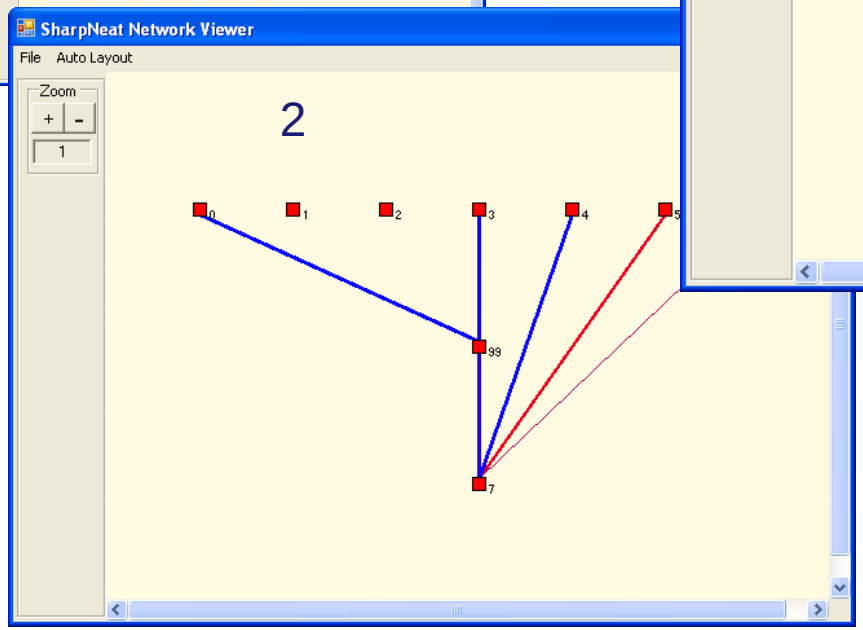

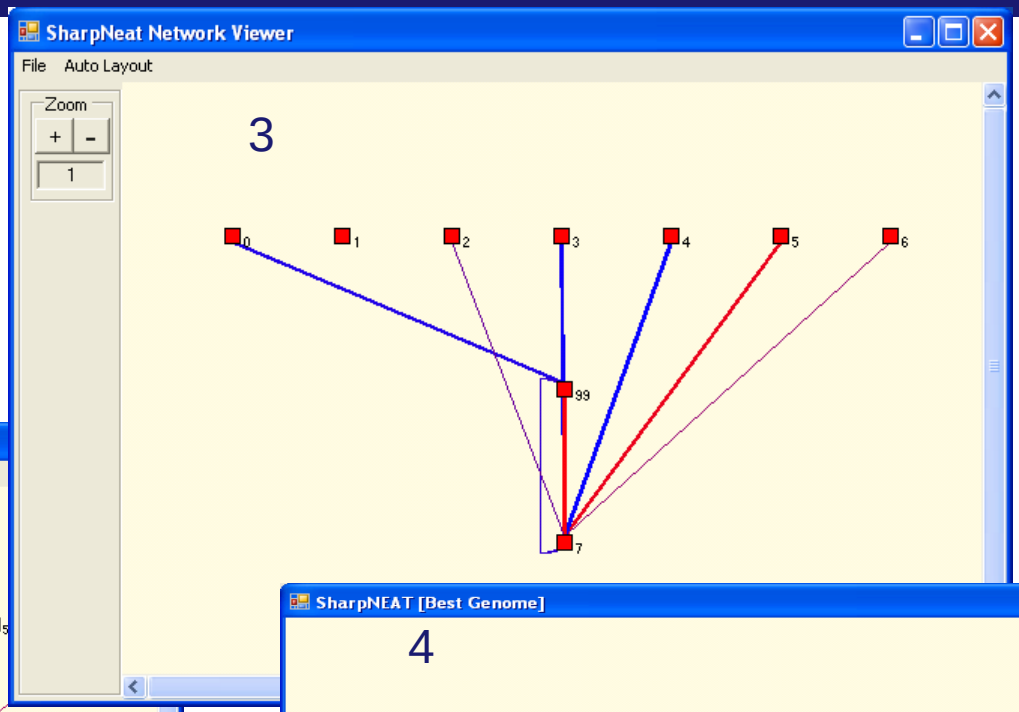

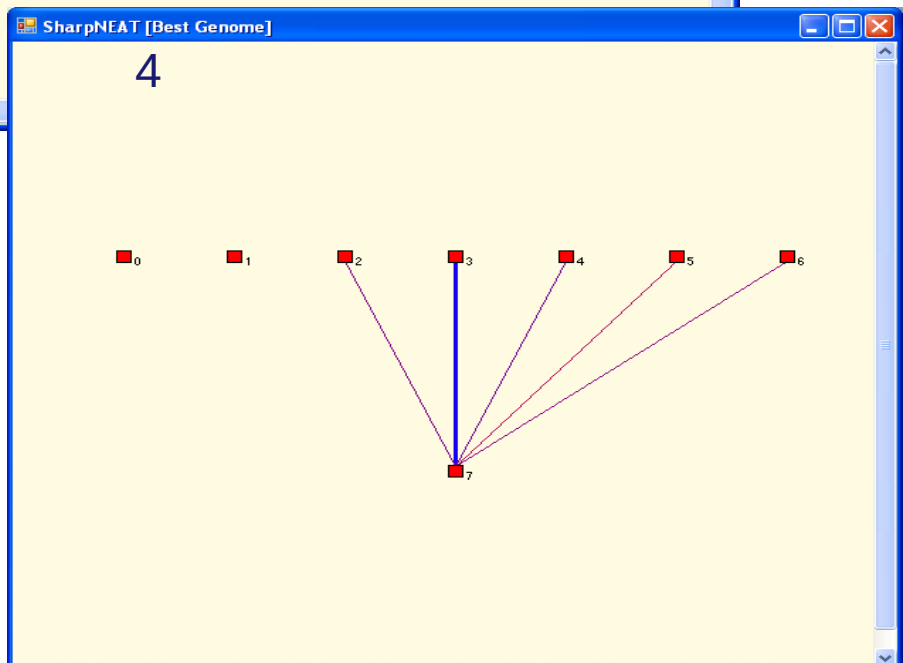

Conclusions

- The use of a form of weak AI called a Topology and Weight Evolving ANN may overcome limited human knowledge and human artifacts in solving extremely complex problems.
- The configuration and simulation of Swarm-like experiments on MPP hybrid, heterogeneous compute fabric is precisely such a problem.

Resulting Work

- The partitioning of experimental tasks for massively parallel computation is dependent upon the reduction and partitioning of large data structures.
- We conclude this must be evolved by an agent in a hybrid, heterogeneous cluster and is dependent upon contextual factors, such as existing schedule, priorities, and the current state of the cluster's configuration.
- We offer a demonstration of the power of this technique using a "toy" Rocks/CUDA cluster.

References

- Burkhart, R. – ppt. *Modeling System Structure and Behavior with SysML Blocks*, March, 2006, Frontiers in Design and Simulation Research, Georgia Tech. http://www.pslm.gatech.edu/events/frontiers2006/proceedings/2006-03-16-Frontiers2006-Burkhart-2.pdf
- Wang, Xiaolei – "*Hybrid Nature-Inspired Computation Methods for Optimization*", PhD Thesis, Helsinki University of Technology, Espoo, Finland, 2009
- Lloyd, K., - *WattSys: A Methodology and Formalism for Modeling, Simulating, Visualizing and Analyzing Complex Systems,* 2008
- Lloyd, K., - US Patent App 20080033897 "*An Object Oriented System and Method for Graphically Displaying and Analyzing Complex Systems"*

Contact Information

Kenneth A. Lloyd, Jr.
CEO – Director, Systems Science
Watt Systems Technologies Inc.
9270 Eagle Ranch Road NW #1818
Albuquerque, NM 87114 USA
(505) 918-6988

kenneth.lloyd@wattsys.com (preferred)
All non-attributed graphics by author