Software That Can Think and Do
Capacity on Demand Across Platforms

Dr. Bernie Domanski
The City University of New York
The College of Staten Island

Motivation

- Rapid advances in research and technology now allow data analysis and modeling of extremely complex systems.
- Methods from artificial intelligence (AI) such as *Neural Networks* have been explored to determine how they apply to systems management.

Software That Can Think & *Do*

- "*Go the next step*" - from analysis to action
- Software systems can now be automated to
  - *analyze* measurements being generated by a complex system in *real time*,
  - automatically *model* its different facets, and then
  - *initiate actions automatically* to maintain effective service time and resource utilization.

Automation is Within Our Grasp

- By combining emerging technologies from
  - AI
  - Performance Measurement
  - Workload Modeling
- … data centers can begin to enjoy a new level of availability and guaranteed service-levels.

Automation … Within Our Grasp?

- Implies the *automatic configuration and deployment* of new servers in real time.
- With such functionality, can spare capacity for individual systems be minimized and can *Capacity on Demand* become a reality?

Background - Neural Nets

- Neural Nets …
  - Recognize patterns and relationships
  - Learn from experience
  - Adapt dynamically
  - Predict outcomes
- How Old Are They?
  - Original ideas developed in the 40's
  - Research moved from labs to commercial world in the 80's

Applications of the Technology

- Performance & Capacity Planning
- Customer profiling/segmentation
- Market penetration analysis
- Fraud detection
- Financial forecasting
- Help desk
- Process optimization

A Simple Neural Net Example

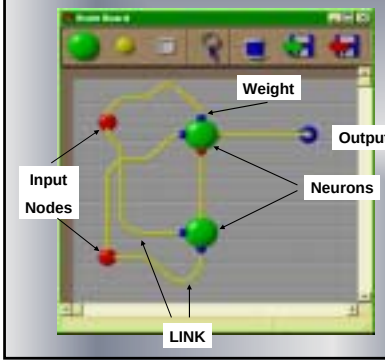

How Does the Network Work?

- Assign weights to each input-link
- Multiply each weight by the input value (0 or 1)
- Sum all the weight-firing input combinations
- If Sum > Threshold for the Neuron
  - THEN Output = +1
  - ELSE Output = -1

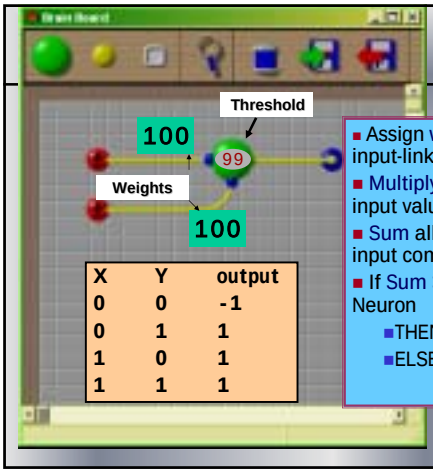

OR

- Assign weights to each input-link
- Multiply each weight by the input value (0 or 1)
- Sum all the weight-firing input combinations
- If Sum > Threshold for the Neuron
  - THEN Output = +1
  - ELSE Output = -1

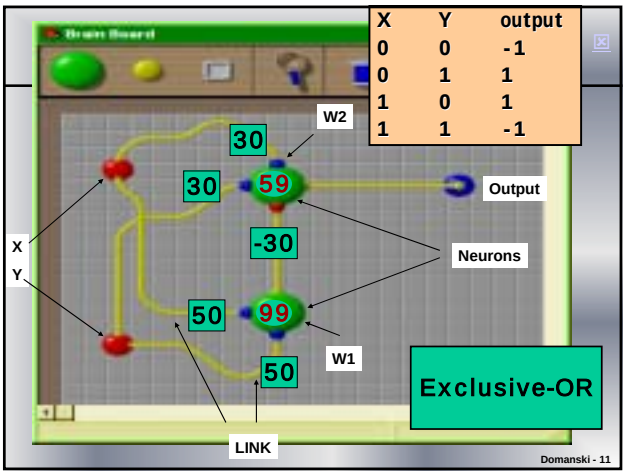

The Training Rule Explained

- Modify the *weights* (w_i) according to the *Training Rule:*

$$w_i = w_i + \triangle w_i \text{ where}$$
$$\triangle w_i = r * (t - a) * x_i$$

- Here –
  - r is the *learning rate* (eg. 0.2)
  - t = target output
  - a = actual output
  - x_i = i-th input value

Training for ‘OR’

Training Set:

X_1	X_2	target
0	0	-1
0	1	1
1	0	1
1	1	1

Initial Random Weights

W1 = .3
W2 = .7

Learning Rate

r = .2

Applying the Training Set for OR - 1

0 0 = -1
0 1 = -1 X

$\triangle w_1 = r * (t - a) * x_1$
$= .2 * (1-(-1)) * x_1$
$= .2 * (2) * 0$
$= \mathbf{0}$

$\triangle w_2 = .2 * (1-(-1)) * x_2$
$= .2 * (2) * 1$
$= \mathbf{.4}$

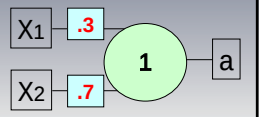

$w_1 = w_1 + \triangle w_1 = .3 + 0 = \mathbf{.3}$

$w_2 = w_2 + \triangle w_2 = .7 + .4 = \mathbf{1.1}$

Applying the Training Set for OR- 2

0 0 = -1
0 1 = +1
1 0 = -1 X

$\triangle w_1 = r * (t - a) * x_1$
$= .2 * (1-(-1)) * x_1$
$= .2 * (2) * 1$
$= \mathbf{.4}$

$\triangle w_2 = .2 * (1-(-1)) * x_2$
$= .2 * (2) * 0$
$= \mathbf{0}$

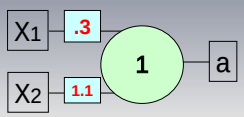

$w_1 = w_1 + \triangle w_1 = .3 + .4 = \mathbf{.7}$

$w_2 = w_2 + \triangle w_2 = 1.1 + 0 = \mathbf{1.1}$

Applying the Training Set for OR- 3

0 0 = -1
0 1 = +1
1 0 = -1 X

$\triangle w_1 = r * (t - a) * x_1$
$= .2 * (1-(-1)) * x_1$
$= .2 * (2) * 1$
$= \mathbf{.4}$

$\triangle w_2 = .2 * (1-(-1)) * x_2$
$= .2 * (2) * 0$
$= \mathbf{0}$

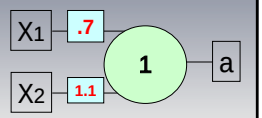

$w_1 = w_1 + \triangle w_1 = .7 + .4 = \mathbf{1.1}$

$w_2 = w_2 + \triangle w_2 = 1.1 + 0 = \mathbf{1.1}$

Applying the Training Set for OR- 4

0 0 = -1
0 1 = +1
1 0 = +1
1 1 = +1 ☑

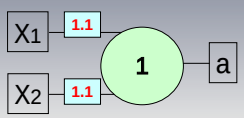

Training for ‘AND’

Training Set:

X_1	X_2	target
0	0	-1
0	1	-1
1	0	-1
1	1	1

Initial Random Weights

W1 = .3
W2 = .7

Learning Rate

r = .2

Applying the Training Set for AND-1

0 0 = -1
0 1 = -1
1 0 = -1
1 1 = -1 **X**

$\triangle w_1 = r * (t - a) * x_1$
$= .2 * (1-(-1)) * 1$
$= .4$

$\triangle w_2 = .2 * (1-(-1)) * 1$
$= .4$

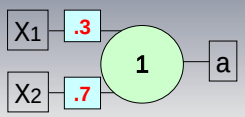

$w_1 = w_1 + \triangle w_1$
$= .3 + .4 = .7$

$w_2 = w_2 + \triangle w_2$
$= .7 + .4 = 1.1$

Applying the Training Set for AND-2

0 0 = -1
0 1 = +1 **X**

$\triangle w_1 = r * (t - a) * x_1$
$= .2 * (-1-(+1)) * 0$
$= 0$

$\triangle w_2 = .2 * (-1-(+1)) * 1$
$= -.4$

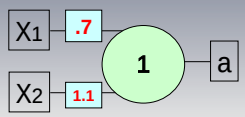

$w_1 = w_1 + \triangle w_1$
$= .7 + 0 = .7$

$w_2 = w_2 + \triangle w_2$
$= 1.1 - .4 = .7$

Applying the Training Set for AND - 3

0 0 = -1
0 1 = -1
1 0 = -1
1 1 = +1 ☑

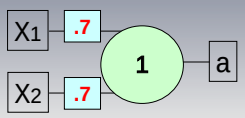

Why Is It Useful?

- NNets are unlike *AI* ***Rule-Based*** software
- A NNet is ***trained to learn relationships*** in the data they have been given.
- A NNet learns by being given a ***training set***.
  - Like a child learns the difference between a chair and a table by being shown examples
- Due to its complex, non-linear structure, the NNet can find relationships in data that humans are unable to do.

Do Neural Nets Work All the Time?

- NNets can only learn if the training set consists of good examples
- *'garbage in - garbage out'* is doubly true
- Great care should be taken to
  - present decorrelated inputs
  - remove outliers in the data
  - use as much prior knowledge to find relevant inputs as possible
- Care must also be taken that the training set is representative
- NNets cannot learn from just a few examples

What To Expect …

- At 11:00AM, we predict a 95% probability that server AB232 will run out of virtual memory in approximately 45 minutes.
- The problem appears in the metrics below.
  - Network Maintenance Message Input (ICMP:Messages Received per second) is higher than expected. It is 1.75 and has never been higher than 1.324.
  - WWW Input (HTTP Service:Bytes Received/sec) is higher than expected. It is 1482 and has never been higher than 1268.
  - Virtual memory read from disk (Memory:Pages Input/sec) is lower than expected. It is 26.8 and is expected to be greater than 43.1.
  - Available RAM (Memory:Available Bytes) is within limits.
  - Virtual memory swapping (Memory:Pages/sec) is within limits.
  - Virtual memory read and write operations (Memory:Page Reads/sec) are within limits.

It Aint A Straight Line !

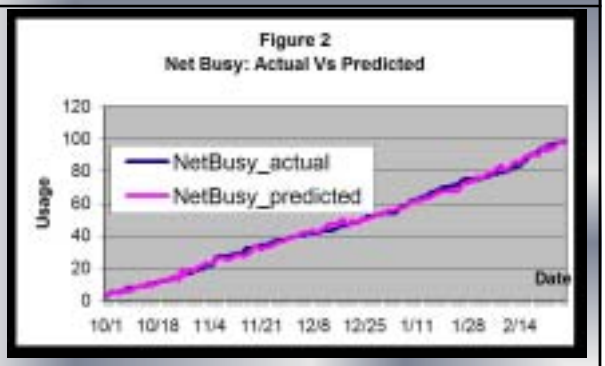

What Can This Do For Me?

Crisis Management - with many systems on the floor, *systems management* becomes complex

- ***Automated processes -***
  - Free your time
  - Are more accurate
  - Don't miss things
  - Increase productivity
- ***Allows you to address more difficult problems***

Tastes Great ... Less Filling

- Customers really want computers that are *less filling*.
- With data centers costing $50 million to build,
  - companies need devices with smaller footprints
  - they can cram more high-octane computing "*oomph*" into less space.
  - They want machines that use less power, both to save on energy costs and so that machines can be packed close together without frying each other.
  - Companies want gear that can take better care of itself.
- The sorry truth is that the expense of configuring and fixing computers is higher than the boxes themselves.

Want Fewer People

- ***"If you look at the total cost of ownership of any hardware, the biggest cost is people"***
- ***"Equipment needs to be able to manage itself much more than it does today."***
  - Jerry A. Scaggs, Technology VP, UPS

Pizza Box Servers

- Pizza box servers
  - stacked in racks and
  - take up far less space than traditional machines
  - as big as a college-dorm refrigerator.
- IDC: 1 million of these servers were sold, vs. 3.6 million bigger ones

Blade Servers

- Book-size "*blade*" computers are available
  - a processor on a circuit card
  - 20 can fit in the space of 1 pizza box
  - use chips that require less power.
  - come preconfigured to handle just one task, eg. serving up Web pages
  - they can be set up in minutes.
  - *Thinking software* links the blades together so that if one unit fails, it can be replaced quickly.

Lose The General?

- "*We want to offer the best foot soldiers, not the best General.*
- *If you lose the General, you're screwed.*
- *If you lose the foot soldier, you've got 10,000 behind him*"
  - Giovanni Coglitore, CEO of Rackable Systems

Startups Are Into This …

- *Thinking software* that makes the growing potpourri of servers operate in concert.
- Like a conductor, *automatically* direct servers to the greatest enterprise need at any moment –
  - say, speeding up a company's e-mail system as workers sit down at their desks in the morning ...
  - then shifting to printing end-of-day reports in the afternoon.

Startups Are Into This …

Examples of *software that can think* -

- its about finding what the most pressing needs are at any moment, and
- then make changes (*do*) to maintain efficient service levels.
  - *Think Dynamics,*
  - *Ejasent,*
  - *NetScaler,* and
  - *Peakstone*

… And So Is IBM …

Are Self-managing Computers Still A Fantasy?

- IBM has pulled a number of efforts together into one grand, multibillion-dollar research and development project dubbed *eLiza*.
- It uses the z-Series *Intelligent Resource Director*
  - software that anticipates glitches within a server and
  - shifts the workload to other parts of the system.
- *Heterogeneous Workload Manager* (HWM)
  - will handle as many as 64 classes of work,
  - spanning the entire network rather than just a single server.
  - will track each job throughout the network, and
  - automatically tune both the network and operating system to reduce delays and
  - assist various classes of work in achieving performance goals.

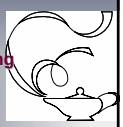

Océano = Scalable Cluster Mgmt

- A highly integrated, parallel system that monitors everything from the computing and network resource usage to application workload and database performance
- Automatically allocates computing resources to various workloads.
- Océano is a scalable infrastructure that allows for multi-customer hosting on a collection of *virtualized* hardware.

Océano – Scalable Cluster Management

- In English: a group of servers can be automated to handle the IT needs of many users, including on-the-fly changes in the workload requirements.
- Océano will introduce high levels of automation to dynamically adjust web sites to actual traffic demands over a massively parallel array of shared and distributed LINUX servers.
- This is similar to the *Thinking Software* of the startups.

How Close Is This To Becoming Reality?

Paul Horn, the Director of IBM Research –

- "We actually have a prototype of this [Océano] operating. I'll give you an example: Today, if we were doing a Victoria's Secret fashion show on a Web server, we would gather a bunch of servers, plug them in, and load them up with the [Web server] application.
- You would get the ability to deal with a lot of customers through brute force by just plugging in more servers.
- What Océano can do is pull that server farm together automatically by sensing what the demand is going to be.

How Close Is This To Becoming Reality?

- It goes out and automatically brings down servers not being heavily used for, say, a sports Web site.
- It will clean off that site and bring those servers back up but running the Victoria's Secret application.
- It can do this without any human intervention.
- It can be a great productivity boost for IT. We think this is going to be pretty exciting technology." *

*http://iwsun4.infoworld.com/articles/hn/xml/01/03/13/010313hnhorn.xml

Scalable Cluster Mgmt

- **Isolation For Security –**
  - so that a breech of one application environment doesn't jeopardize another;
- **Fault Isolation –**
  - so faults in one application environment can't effect another; and
- **Performance Isolation –**
  - high demand for one application can't effect the performance of another.

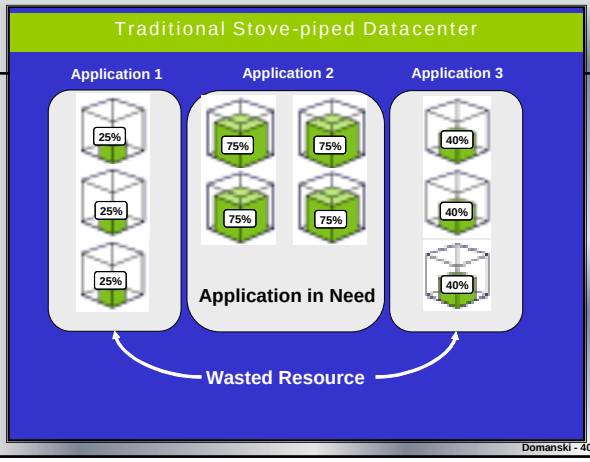

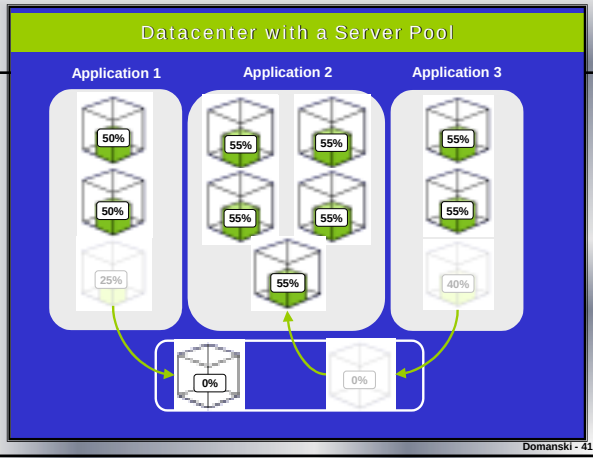

Who Would "*Thinking / Doing*" Software Affect Immediately?

- Delivery of managed or outsourcing services by application infrastructure providers (AIPs)
- Based on managing *dedicated* customer systems from remote "Internet" data centers.
- Is the launch of computing utility services a watershed moment?
  - because the provisioning of computing services as a set of utility offerings could radically alter the cost structure of the services and
  - place tremendous pressure on existing players to alter their service infrastructures.

Conclusions, Futures, Questions

- What is the significance of thinking/doing software?
  - impact long-term trends in providing computing services from remote datacenters as well as traditional outsourcing.
  - Immediate benefit = *defer spending on upgrades*
  - *Datacenters will simply make better use of available resources*.
  - *Automation* minimizes the time required to add additional capacity when needed.
  - At an unexpected workload peak, *capacity-on-demand* will smooth the performance spikes, lessening the impact on service delivery and response time.

Price of Automation

- cap… aut… …t are
- rep… out… …onal
- In … *pe…* …*n of*
- soc…
- Rea… cor… the… …e

So if new services will reduce costs over time and will be adopted as part of the service delivery infrastructure, the long-term results will impact *pricing strategies* and *outsourcing decisions*.

Pricing Strategies

- Will require service providers to develop sophisticated platforms to monitor real-time consumption of com… …l-time data…
- This… …mers as p…
- Cus… envi… …gle bill a… …on envi…
- Note… …cing model is undergoing massive change as well. We can expect a large impact on how we distribute those workloads across a *SYSPLEX*.

From the now fixed monthly fee, computing utility services will create a disruption in this pricing model as the inherent value of a utility is its variable cost/price.

Outsourcing Strategies

- Success in offering these capacity-on-demand services could accelerate the trend toward outsourcing.
- Key components in enabling this acceleration include
  - providing a higher quality of service,
  - significantly reducing costs, and
  - remaining in business "forever."

Final Thought

Paul Horn
IBM's Director of Research:

Building an autonomic computing network is not optional; in fact, the future of the Internet and e-business depends on it.

The End

- Questions, comments, … ??
- Finding me –
  - Dr Bernie Domanski
  - Email: domanski@postbox.csi.cuny.edu
  - Website: http://domanski.cs.csi.cuny.edu
  - Cell: (908) 415-6105
  - Thanks for coming and listening !