

Wireless Research with GENI: Machine Learning for Wireless Sensor Networks Track

14 GENI Engineering Conference
Boston, MA
11 July 2012

Fraida Fund, Thanasis Korakis (Polytechnic Institute of NYU)
Abhimanyu Gosain (GENI Project Office, BBN)
Ivan Seskar (Rutgers WINLAB)



Scenario

- Dr. Watson is a researcher in the field of artificial intelligence and machine learning.
- He's developed a novel approach to physical security systems.
- It collects ambient light measurements from distributed sensors and processes them with a machine learning algorithm.
- End result: estimates the probability of a certain security event (e.g., intruder in room).

Scenario (continued)

- Dr. Watson's student wrote a Ruby application to implement the scheme that he developed.

What physical and computing resources does he need to evaluate this implementation?

Required Resources

- Distributed ambient light sensors that communicate with a measurement sink
- A measurement sink that he can configure to run the machine learning application


Optional Resources

- A way to set up and orchestrate complicated experiments with many parts
- Full control over all parts of the experiment, including measurement sources/sinks and all communication links
- Instrumentation and measurement tools for systematically collecting and comparing results

Experiment Topology

- Light sensors at one site communicate measurements over WiMAX network
 - NYU-Poly has thirteen physically distributed nodes equipped with temperature, light, humidity sensors
- Light readings travel over I2 link to a measurement sink at another site
 - NYU-Poly, BBN, WINLAB, all connected over I2
 - We'll configure WiMAX clients to use the I2 link
- Measurement sink receives measurements and processes using student's Ruby application

Experiment Topology, Illustrated



NYU·poly

POLYTECHNIC INSTITUTE OF NEW YORK UNIVERSITY

Fully programmable

geni
Exploring Networks of the Future

Experiment

Does this experiment topology meet Dr. Watson's requirements?

- ✓ Distributed ambient light sensors that communicate with a measurement sink
- ✓ A measurement sink that he can configure to run the machine learning application.

and also

- ✓ Full control over all parts of the experiment, including measurement sources/sinks and all communication links

Accessing a GENI WiMAX Testbed

- Register for an account and make a reservation to use a testbed ahead of time

Pro tip: Two WiMAX testbeds are currently open to experimenters.

- To register at NYU-Poly, visit <http://witestlab.poly.edu> and click “Register” in the top left corner – you’ll get an email later than day with further instructions.
- To register at WINLAB, visit <http://www.orbit-lab.org/userManagement/register>

- At the designated time, use an SSH client to log in to the console of a WiMAX testbed, e.g.
 - `ssh ffund@omfserver-witest.poly.edu`

Setting up Experiment

(Pre-configured BBN – demo this process at NYU-Poly)

- Reset base station to default settings
 - `wget -q0- http://wimaxrf:5052/wimaxrf/bs/default`
 - `wget -q0- http://wimaxrf:5052/wimaxrf/bs/restart`
- Configure datapath (connectivity endpoints) for WiMAX clients
 - `wget -q0- "http://wimaxrf:5052/wimaxrf/datapath/config/load?name=bbn"`

Pro tip: See more information about the pre-configured datapath options available at NYU-Poly at <http://witestlab.poly.edu/index.php/instructions.html>

Setting up Experiment (cont.)

- Install disk images

- `omf-5.3 load -i <image name> -t <list of nodes>`
- The student started with a baseline image that already has WiMAX and sensor functionality set up (baseline-witest.ndz)
- (Images are provided by testbed operators at each testbed)
- He installed the machine learning application he built and its dependencies, some packaged artificial intelligence libraries (ai4r). Then he saved this disk image using `omf-5.3 save -n <node name>`
- From now on, he can just load this saved image onto his nodes every session – no need to set up his environment every time

Pro tip: See more information about the prepared baseline images for NYU-Poly at <http://witestlab.poly.edu/index.php/instructions.html>

Experiment: Sequence of Events

- Dr. Watson's student logs in to each node to begin the experiment
 - `ssh root@omf.witest.node1`
 - `wimaxcu connect network 51`
 - `ifconfig wmx0 10.37.35.151 netmask 255.255.255.0`
 - `sensor -d 10.37.35.175`
- Repeat for each of 13 nodes at NYU-Poly, then for measurement sink at BBN



Experiment: Sequence of Events

- To save time, he learns to use OMF to configure the resources...

```
defProperty('hrnPrefix', "omf.witest.node", "Prefix to use for the HRN of resources")
defProperty('sender', "[13,12,11,10,9,8,7,6,5,4,3,2,1]", "List of IDs for the resources to use as senders")
defProperty('groupSize', 1, "Number of resources to put in each group of senders")
```

```
groupList = []
res = eval(property.sender.value)
groupNumber = res.size >= property.groupSize ? (res.size.to_f / property.groupSize.value.to_f).ceil : 1
(1..groupNumber).each do |i|
  list = []
  (1..property.groupSize).each do |j| popped = res.pop ; list << popped if !popped.nil? end
  senderNames = list.collect do |id| "#{property.hrnPrefix}#{id}" end
  senders = senderNames.join(',')

```

```
info "Group Sender #{i}: '#{senders}'"
groupList << "Sender#{i}"
defGroup("Sender#{i}", senders) do |node|
  node.net.x0.profile = '51'
  node.net.x0.ip = "10.37.35.#{i+150}"
  node.net.x0.netmask = "255.255.255.0"
  node.addApplication("test:app:sensor") do |app|
    app.setProperty('interval',5)
    app.setProperty('destination', '10.37.35.175')
    app.measure('stats')
  end
end
end
end
```

Set up WiMAX connectivity

Configure sensor application

Experiment: Sequence of Events

- .. and to define the sequence of events

```
onEvent(:ALL_UP_AND_INSTALLED) do |event|
  wait 10
  allGroups.startApplications
  wait 120
  allGroups.stopApplications
  Experiment.done
end
```

- He writes a similar script for BBN
- Now, to run the experiment, he only needs to do

```
omf-5.3 exec sensorapp-source.rb (at NYU-Poly)
omf-5.3 exec sensorapp-sink.rb (at BBN)
```

Experiment: I&M

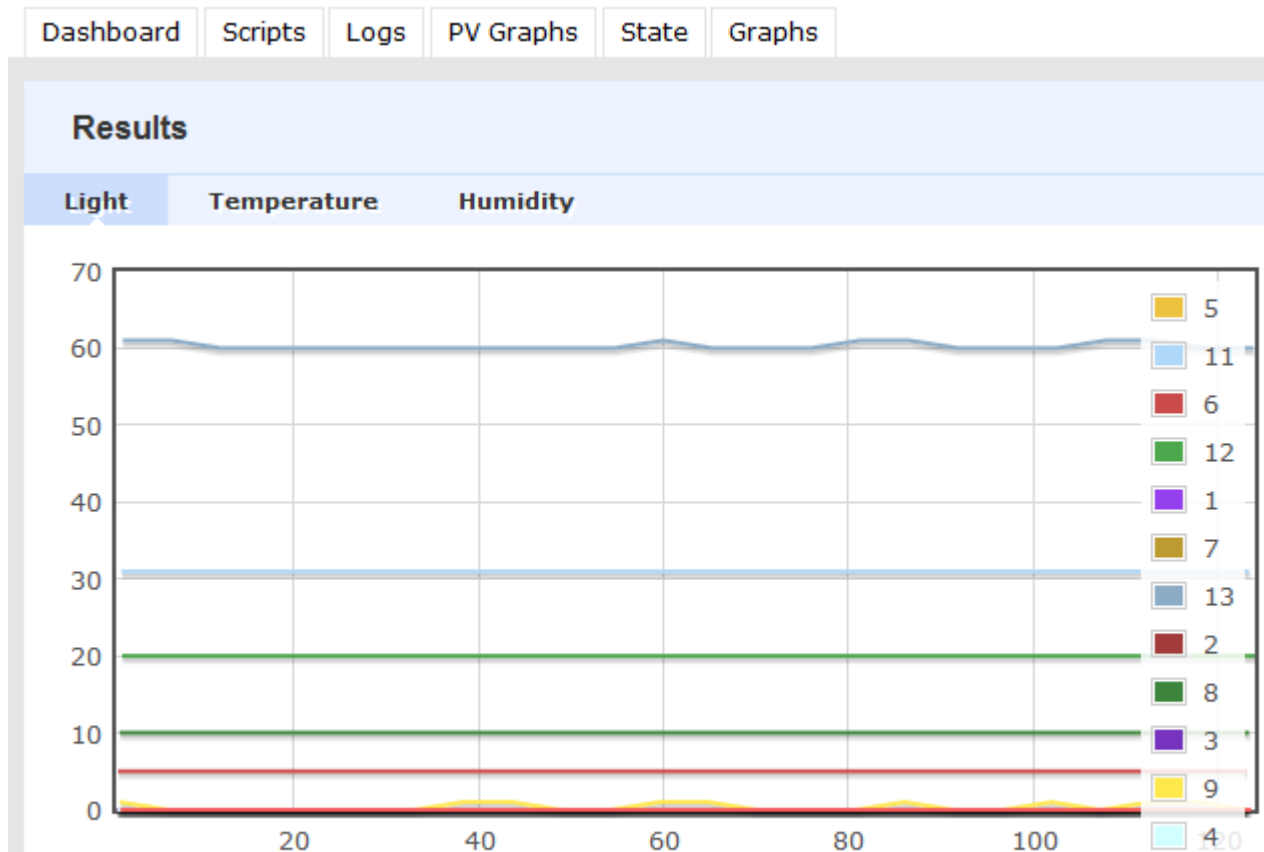
- The *sensor* applications saves measurements for him in a database, so he doesn't need to worry about capturing the results.

```
CREATE TABLE _experiment_metadata (key TEXT PRIMARY KEY, value TEXT);
CREATE TABLE _senders (name TEXT PRIMARY KEY, id INTEGER UNIQUE);
CREATE TABLE "sensor_predicted" (oml_sender_id INTEGER, oml_seq INTEGER, oml_ts_client REAL,
oml_ts_server REAL, "probability" REAL, "ts" INTEGER);
CREATE TABLE "sensor_received" (oml_sender_id INTEGER, oml_seq INTEGER, oml_ts_client REAL,
oml_ts_server REAL, "ts_server" INTEGER, "ts_client" INTEGER, "sender" TEXT, "source" TEXT,
"temperature" REAL, "humidity" INTEGER, "light" INTEGER);
CREATE TABLE "sensor_stats" (oml_sender_id INTEGER, oml_seq INTEGER, oml_ts_client REAL,
oml_ts_server REAL, "ts" INTEGER, "source" TEXT, "temperature" REAL, "humidity" INTEGER,
"light" INTEGER);
```

- He can retrieve these measurements as an *sql* or *csv* anytime after the experiment runs, and plot them using standard data analysis tools like *gnuplot* or *R*.

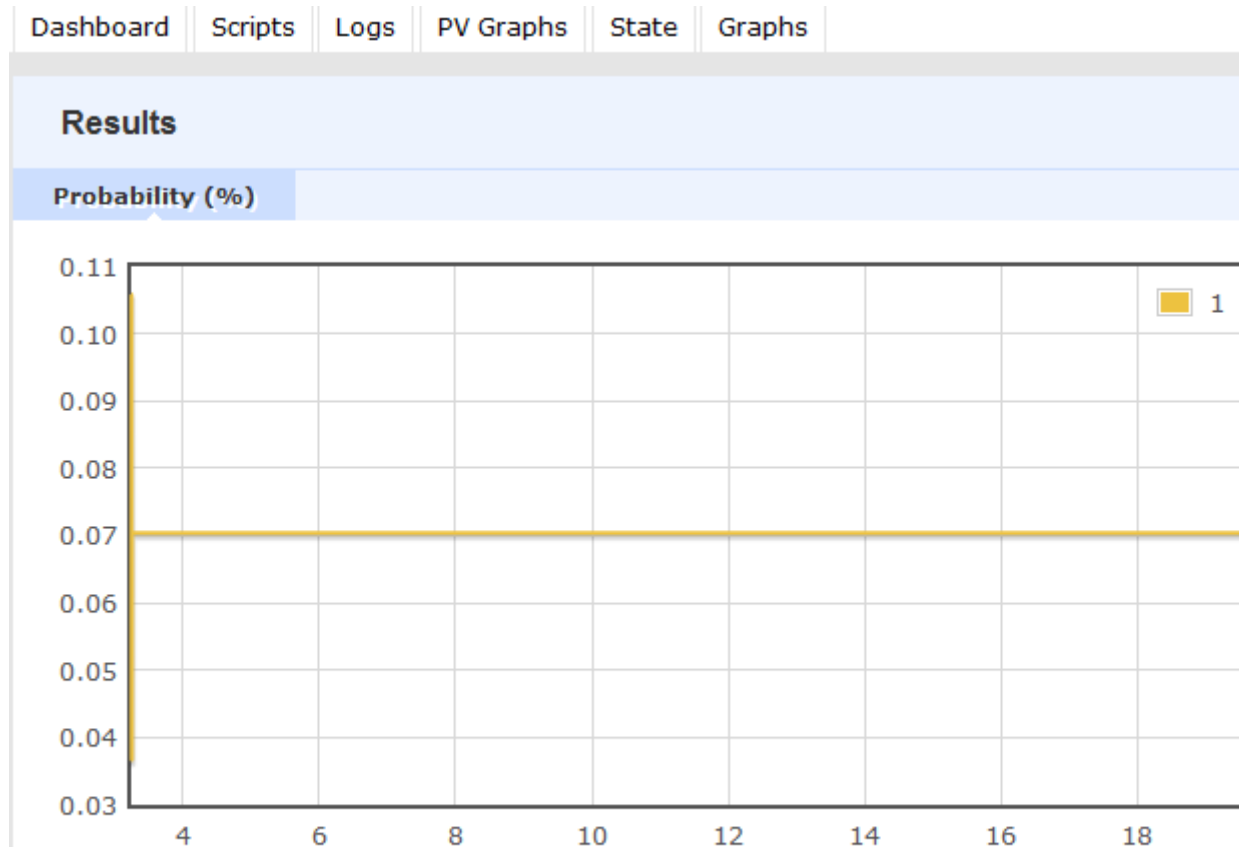
Experiment: I&M

- He even configured graphs in the experiment script, so he can view the visualization live and check on progress during the experiment



Experiment: I&M

Pro tip: Click on “Logs” to see the console output and “Scripts” to see the experiment script.



Experiment – with OMF

Does OMF give Dr. Watson useful functionality?

- ✓ A way to set up and orchestrate complicated experiments with many parts
- ✓ Instrumentation and measurement tools for systematically collecting and comparing results

Discussion

Questions?

(Someone will explain this experiment to the other groups –approximately 2 minutes)