

HANDS-ON SDN

Introduction to Software-defined Networking
Block Course – 13-17 March 2015

David Koll

Where we are now

You have now learned about:

- Python programming
- SDN basic principles
 - Basic concepts (CP/DP separation etc.)
 - De-facto standard interfaces (OpenFlow)
 - Controllers (NOX, POX, ...)
 - Virtualization (FlowVisor)

Where we want to go

You have now learned about:

- Python programming
- SDN basic principles
 - Basic concepts (CP/DP separation etc.)
 - De-facto standard interfaces (OpenFlow)
 - Controllers (NOX, POX, ...)
 - Virtualization (FlowVisor)

• Put the stuff learned into practice:

- Implement OpenFlow?
- Implement controllers?
- Implement FlowVisor?

- Rather: *learn how to use and program them!*
 - Hands-on work on state-of-the-art tools

How can we get there?

- Luckily, implementations are available.
 - Switches implementing OF
 - Controllers implementing OF
- So, how do we run them?
 - We don't have a hardware testbed at hand
 - We don't have access to a production network
 - We may want to test different things on different network topologies
 - Simulation?

Emulation of Networks

- Network emulation means to run unmodified code interactively on virtual hardware
- Huge benefit:
 - Can actually port our applications seamlessly to hardware
- Challenges:
 - Scalability: need to model hosts, switches, links, controllers, ...
 - Ease-of-Use: easily allow to create different topologies with varying parameters
 - Accuracy: results have to match results obtained from running same experiment on hardware

Enter Mininet


“Mininet creates a **realistic virtual network**, running **real kernel, switch and application code**, on a single machine (VM, cloud or native), in seconds, with a single command” [1]



[1] mininet.org

Enter Mininet

“Mininet creates a **realistic virtual network**, running **real kernel, switch and application code**, on a single machine (VM, cloud or native), in seconds, with a single command” [1]



[1] mininet.org

Enter Mininet

Mininet offers CLI & API to interact with the network

(see demo)

Sudo mn -> Pingall -> H1 ping h2 -> Iperf -> Nodes -> Xterm h1 -> Ifconfig -a
Complicated MAC addresses -> Sudo mn -c -> Sudo mn -mac ->xterm h1 -> ifconfig -a

Customize Topologies

Mininet is not limited to the very basic setup

(see demo)

`Sudo mn - -topo linear,4 -> dump`

Virtual network -> different namespaces

Everything else: not virtualized -> `h1 ps -a == s1 ps -a`

Can also change link capacities/delays :

`iperf`

`Sudo mn -c`

`sudo mn -link tc,bw=1`

`iperf`

Customize Topologies

```
from mininet.topo import Topo
class MyTopo( Topo ):
    "Simple topology example."

    def __init__( self ):
        "Create custom topo."

        # Initialize topology
        Topo.__init__( self )

        # Add hosts and switches
        leftHost = self.addHost( 'h1' )
        rightHost = self.addHost( 'h2' )
        leftSwitch = self.addSwitch( 's3' )
        rightSwitch = self.addSwitch( 's4' )

        # Add links
        self.addLink( leftHost, leftSwitch )
        self.addLink( leftSwitch, rightSwitch )
        self.addLink( rightSwitch, rightHost )

topos = { 'mytopo': ( lambda: MyTopo() ) }
```

Sudo mn -custom custom_topo.py -topo mytopo -test pingall

Customize Switches and Controllers

You can connect different switches and controllers

(see demo)

Sudo mn –switch ovsk –controller remote
New terminal -> cd pox -> ./pox.py forwarding.hub

Bring Links Up/Down

Change the topology at runtime

(see demo)

Pingall -> link h1 s1 down -> pingall -> link h1 s1 up -> pingall

Python Interpreter

Can access each host, switch and controller with Python

(see demo)

Py „hello“ + „world“
Py h1.IP()

Use of Wireshark

We can use Wireshark to debug our network

(see demo)

Exit -> `sudo mn -c` -> `sudo mn -controller remote`
New terminal -> `sudo wireshark &` -> filter of
New terminal -> `./pox.py forwarding.hub`
See of_hello of_features_request etc in wireshark

Limitations?

Limited by single system resources
Limited to Linux kernel (e.g., portability to Windows?)
Limited to real-time

Exercise!

Time for Exercise 7 and 8

Custom Topologies with Mininet Python API

Mininet offers some topologies!

Eg: single switch, linear, tree

What if you want to replicate your very own production network?

Create a custom topology!

Low-level API: Nodes and Links

```
h1 = Host( 'h1' )
h2 = Host( 'h2' )
s1 = OVSSwitch( 's1', inNamespace=False )
c0 = Controller( 'c0', inNamespace=False )
Link( h1, s1 )
Link( h2, s1 )
h1.setIP( '10.1/8' )
h2.setIP( '10.2/8' )
c0.start()
s1.start( [ c0 ] )
print h1.cmd( 'ping -c1', h2.IP() )
s1.stop()
c0.stop()
```

Mid-level API: Network Object


```
net = Mininet()  
h1 = net.addHost( 'h1' )  
h2 = net.addHost( 'h2' )  
s1 = net.addSwitch( 's1' )  
c0 = net.addController( 'c0' )  
net.addLink( h1, s1 )  
net.addLink( h2, s1 )  
net.start()  
print h1.cmd( 'ping -c1', h2.IP() )  
CLI( net )  
net.stop()
```

High-level API: Topology templates

```
class SingleSwitchTopo( Topo ):
    "Single Switch Topology"
    def __init__( self, count=1):
        Topo.__init__(self)
        hosts = [ self.addHost( 'h%d' % i )
                  for i in range( 1, count + 1 ) ]
        s1 = self.addSwitch( 's1' )
        for h in hosts:
            self.addLink( h, s1 )

topos = {'single' : (lambda: SingleSwitchTopo())}
```

Example Topology – Research Lab



Example Topology – Research Lab

```
1  #!/usr/bin/python
2  from mininet.topo import Topo
3
4  class ResearchLab(Topo):
5      """Research Lab Topology"""
6      def __init__(self):
7
8          Topo.__init__(self)
9          testbedhosts = [self.addHost('h%d' % i) for i in range(1, 9)]
10         simservers = [self.addHost('sim%d' % i) for i in range(1, 4)]
11         s1 = self.addSwitch('s1') # TOR switch
12         s2 = self.addSwitch('s2') # Testbed switch
13         s3 = self.addSwitch('s3') # Server switch
14
15         for h in testbedhosts:
16             self.addLink(h, s2, bw=1, delay='15ms')
17
18         for srv in simservers:
19             self.addLink(srv, s3, bw=10, delay='1ms')
20
21         self.addLink(s2, s1, bw=100)
22         self.addLink(s3, s1, bw=100)
23
24     topos = {'rlab': (lambda: ResearchLab())}
```

```
sudo mn
--custom rlab.py
--topo rlab
--link=tc
```

The POX Controller

- Invoke with: `./pox.py [options] <component>`
- <options> can be:
 - `--verbose` : display debugging info
 - `--no-openflow`: do not automatically listen for OpenFlow connections
- <components> are the real meat!
 - There are some basic components we will use for this class
 - Intention: developers will build their own components

The POX Controller - Components

- Some stock components:
 - py
 - forwarding.hub
 - forwarding.l2_learning
 - forwarding.l2_pairs
 - forwarding.....
- openflow.webservice
 - Creates a webinterface to interact with OpenFlow
- openflow.of_01
 - Communicates with OpenFlow 1.0 switches

The POX Controller - Components

- Developing your own components:
 - <https://openflow.stanford.edu/display/ONL/POX+Wiki#POXWiki-DevelopingyourOwnComponents>
- In general: POX wiki a good place to look for help
 - <https://openflow.stanford.edu/display/ONL/POX+Wiki>

POX APIs

- When writing or modifying components (you will do the latter in this course), POX offers some helpful API.
 - E.g.: API for packet handling: `pox.lib.packet`

Example: Get L2 source and destination from a packet

```
def _handle_PacketIn(self, event):  
    packet = event.parsed # POX is based on events!  
    src_of_packet = packet.src #returns an EthAddr  
    dst_of_packet = packet.dst #also returns an EthAddr
```

POX APIs

- When writing or modifying components (you will do the latter in this course), POX offers some helpful API.
 - E.g.: API for packet handling: `pox.lib.packet`

Example: Get source IP from a packet

```
def _handle_PacketIn(self, event):  
    "check if packet is an IP packet"  
    packet = event.parsed  
    ip = packet.find('ipv4') #check if packet is IP  
    if ip is None: #packet is not IP  
        return  
    print "Source IP: ", ip.srcip
```

POX and Openflow

- Up front: Best to read POX wiki:
 - <https://openflow.stanford.edu/display/ONL/POX+Wiki#POXWiki-OpenFlowinPOX>
- Usually, switches connect to POX automatically via OpenFlow
 - Exception: no-openflow option (see previous slides)
- So – how do we communicate with them?

Coding in POX – Connection Elements

- Upon connecting to POX, a switch is associated with a `Connection` object
- Use that object's `send()` method to send messages to the switch
- `Connection` object will raise events on the corresponding switch
 - Create **event handlers** for events you are interested in

In Practice

- Launch our component.
- Add one event listener for `PacketIn`

```
from pox.core import core
import pox.openflow.libopenflow_01 as of

log = core.getLogger()

def launch ():
    "Starts the Component"
    core.openflow.addListenerByName("PacketIn",
        _handle_packetin)

    log.info("Switch running.")
```

In Practice

- Write packet handler (here: flood packet)

```
def _handle_packetin (event):  
    "Handle PacketIn"  
    packet = event.parsed  
    send_packet(event, of.OFPP_ALL) #broadcast  
  
    log.debug("Broadcasting %s.%i -> %s.%i" %  
              (packet.src, event.ofp.in_port,  
               packet.dst, of.OFPP_ALL))
```

In Practice

- Write `send_packet` method (simplified)

```
def send_packet (event, dst_port):  
    "Instructs switch to send packet via dst_port"  
    msg = of.ofp_packet_out(in_port=event.ofp.in_port)  
    msg.data = event.ofp.data  
    msg.actions.append(of.ofp_action_output(port = dst_port))  
  
    event.connection.send(msg)
```


In Practice


- Code on previous slides implemented a hub behaviour
- Exercise: modify hub behaviour to learning switch behaviour

Exercise!

Time for Exercise 9

FlowVisor

- Exercise 5: You have already installed FlowVisor
- Recall: FlowVisor is an extra layer between controllers and switches



FlowVisor

- Basic procedure:
 - Create and start your network topology with Mininet
 - Connect Flowvisor to switches on standard port
 - Slice network with Flowvisor
 - Connect Controllers to Flowvisor slices

FlowVisor

- Basic procedure:
 - Create and start your network topology with Mininet
 - Connect Flowvisor to switches on standard port
 - Slice network with Flowvisor
 - Connect Controllers to Flowvisor slices

Connecting FlowVisor

- FlowVisor operates outside of Mininet!

```
$ sudo /etc/init.d/flowvisor start
```

(see demo)

- Afterwards: use flowvisor control (command: `fvctl`) to slice

```
$ sudo /etc/init.d/flowvisor start
```

Slicing the Network with FlowVisor

- First: enable topology controller


```
$ fvctl -f /dev/null set-config --enable-topo-ctrl  
$ sudo /etc/init.d/flowvisor restart
```

(see demo)

- -f /dev/null option: -f points to pwd file – in our case: empty pw

```
$ sudo fvctl -f /dev/null set-config --enable-topo-ctrl  
$ sudo /etc/init.d/flowvisor restart
```

Let's slice the research lab



Slicing the Network with FlowVisor

- Want to create slice for servers. Have a look at topology:

```
$ fvctl -f /dev/null list-slices
$ fvctl -f /dev/null list-flowspace
$ fvctl -f /dev/null list-datapaths
$ fvctl -f /dev/null list-links
```

(see demo)

```
$ fvctl -f /dev/null list-slices
$ sudo /etc/init.d/flowvisor restart
```

Slicing the Network with FlowVisor

- Add slices with

```
fvctl add-slice [options] <slicename>  
                <controller-url> <admin-email>
```

```
$ fvctl -f /dev/null add-slice servers  
                tcp:localhost:10001 admin@servers
```

(see demo)

```
$ fvctl -f /dev/null add-slice servers tcp:localhost:10001  
admin@servers
```

Add Flowspaces

- Add flowspaces with

```
fvctl add-flowspace [options] <flowspace-name> <dpid>
                    <priority> <match> <slice-perm>
```

```
$ fvctl -f /dev/null add-flowspace switch1-port2
    1 1 in_port=2 servers=7
```

- Permissions: Bitmask
 - 1=DELEGATE, 2=READ, 4=WRITE

(see demo)

```
$ fvctl -f /dev/null add-flowspace switch1-port2 1 1 in_port=2
servers=7
$ fvctl -f /dev/null add-flowspace switch3-port1 3 1 in_port=1
servers=7
$ fvctl -f /dev/null add-flowspace switch3-port2 3 1 in_port=2
servers=7
$ fvctl -f /dev/null add-flowspace switch3-port3 3 1 in_port=3
servers=7
$ fvctl -f /dev/null add-flowspace switch3-port4 3 1 in_port=4
servers=7

$ fvctl -f /dev/null list-flowspace
```

Connect Controllers

- Start controller and connect to FlowVisor

(see demo)

NEW TERMINAL

```
$ ./pox.py openflow.of_01 --port 10001 forwarding.l2_pairs
```

Comment: need to load openflow with custom parameters here (instead of default that would be loaded when just loading forwarding.l2_pairs)

Test Slicing

- Servers should be able to ping each other, but not any hosts

(see demo)

```
Mininet console> sim1 ping -c1 sim2  
Mininet console> sim1 ping -c1 h1
```

Paper Presentations

- Friday, April 22nd, Room 0.101
- 8.30am - 12am

- Groups of 3: 30 minutes presentation, 5 mins Q&A
- Groups of 2: 20 minutes presentation, 5 mins Q&A
 - Split time equally among speakers

- Single presenters: 10 minutes presentation, 5 mins Q&A
 - Pick one of the two papers two present, write review for the other

Paper Presentations

- General rule of thumb: you will need 1 to 1.5 minutes per slide in your presentation
- Plan number of slides appropriately
- Every member of the group needs to read all papers
- Different ways of approaching a presentation:
 - Present your set of papers at a relatively high level
 - Present one paper briefly and the other one in more detail
 - Up to you... but you have to know all papers in detail for creating the presentation and for Q&A.
- Practice your presentation before – make sure each speaker gets enough time and you don't overuse time
- We will run a clock and stop you
- Imbalanced talk quotas will negatively affect the grade

Paper Reviews

- You will receive a form with guidelines
- Similar to real-world paper review

Exercise Submissions

- Code: submit .py files (executable where required)
- Text exercise: submit .pdf, can include screenshots or copy/pasted console output where required
- Put all items (slides presented, text exercises, code exercises, paper review) into one .zip archive. Send to Sameer (see email in first lecture slides)

Exercise!

Time for Exercise 10