



# Systems and Internet Infrastructure Security

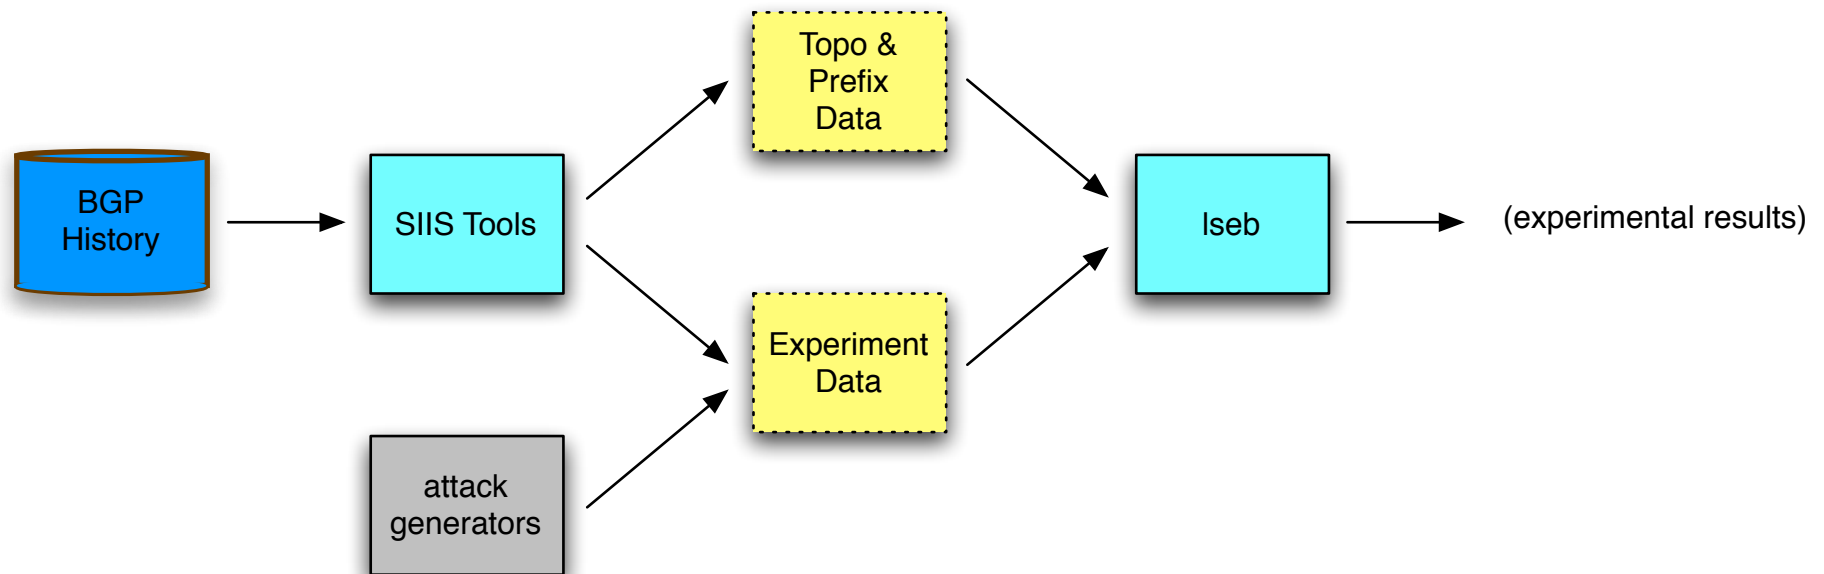
Network and Security Research Center  
Department of Computer Science and Engineering  
Pennsylvania State University, University Park PA

## **Iseb:** Testing Large Scale BGP Security in Replayable Network Environments

DETER/EMIST Workshop  
June 15th, 2006 - Arlington, VA  
Patrick McDaniel and Kevin Butler

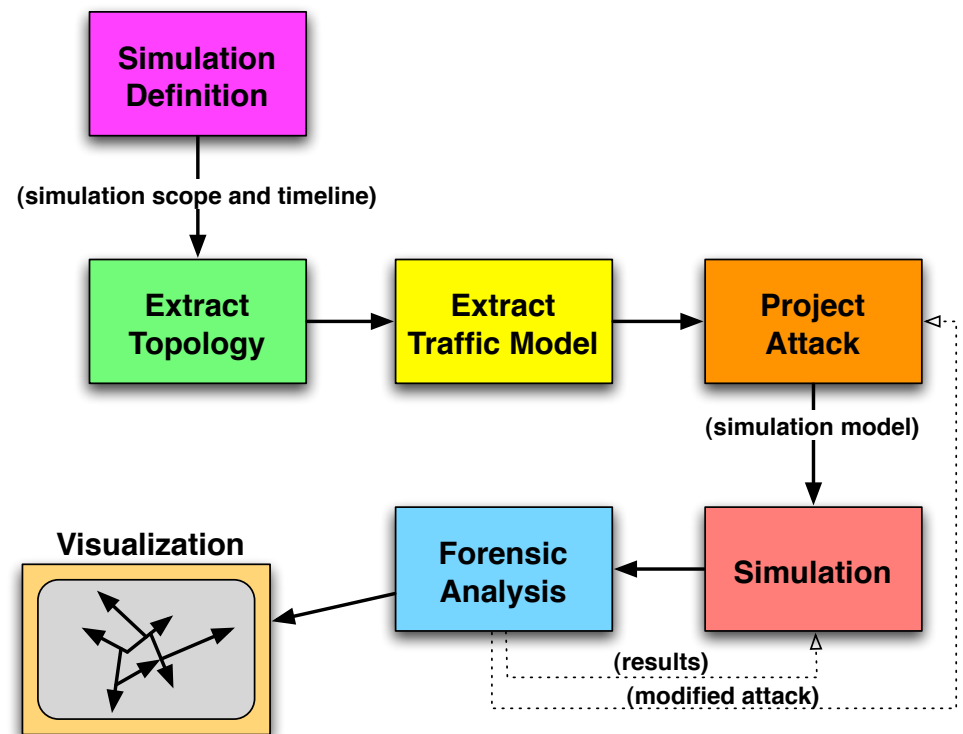
# Iseb: a BGP “way-back” machine?

- *Large-Scale eBGP Simulation (Iseb)*
  - ▶ **Goal:** an experimental apparatus to *insert any attack* in *any place* at *any time* in the recorded history of BGP.
  - ▶ Massive-scale distributed simulation



# Evaluation Testbed: an opportunity?

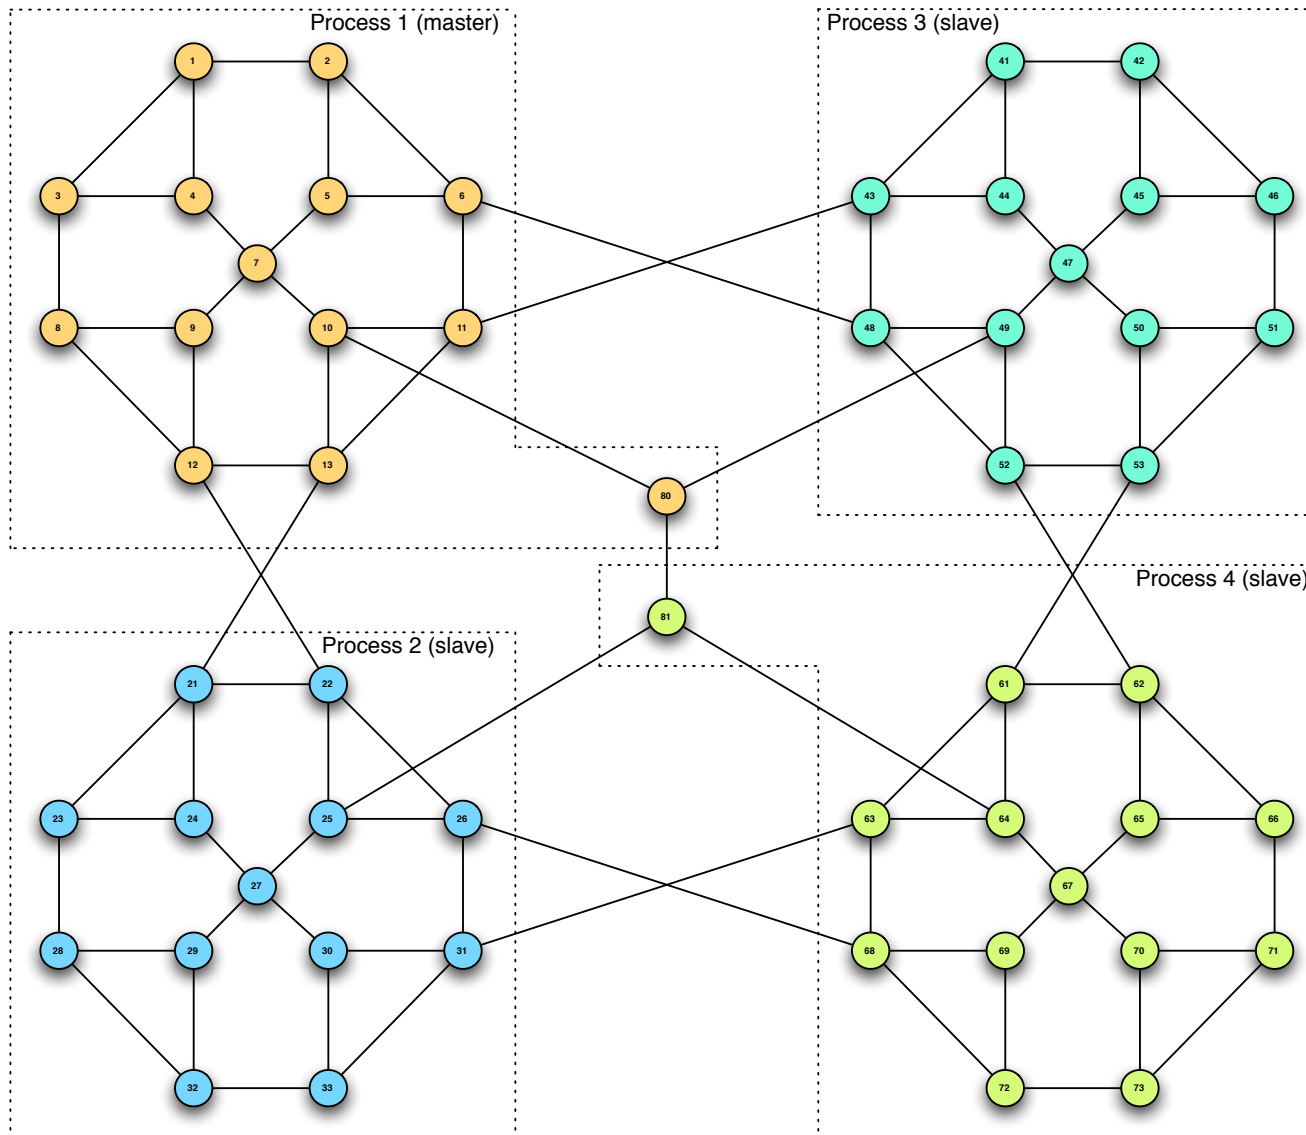
- We have a *global, diverse, detailed* history of BGP behavior from many vantage points over extended period (years ...)
  - PREDICT, PCH, RouteViews, RIPE, ...
- Q: How can we exploit these unique repositories?
  - Analysis
  - Simulation



# Simulating the AS

- *Domain*: understanding global effects of BGP attacks and countermeasures
  - ▶ protocol simplifications (e.g., iBGP, simplified policy)
- Hypothesis: real traffic traces will preserve realism of simulation BGP behavior
  - ▶ Extract meaningful BGP structure and events from history
  - ▶ *Advantage*: verifiable via traces
- Simulation of *entire* Internet eBGP protocol
  - ▶ e.g., 20,000+ simulated autonomous systems
  - ▶ Each AS executes eBGP and maintains BGP table (**166Mb**)
  - ▶ Model **every** BGP message on Internet
- Scalability lies in the clustering of simulated ASes

# DETER Simulation Topology



# Simulation Setup

- Single file configuration (process self organization)
  - ▶ identify the ASes, links, and topology
  - ▶ As gleaned from BGP data (e.g., SIIS-topo-extract)

```
% Demo network Setup
...
4 AS 72
4 AS 73
4 AS 81
% Links
4 PEER 61 62
4 PEER 61 63
4 PEER 61 64
4 PEER 62 65
...
% Prefixes
4 PREFIX 61 61.0.0.0/8
4 PREFIX 73 73.0.0.0/8
...
```

# Scripting the experiment ...

- Single file script delivered to all via master
  - ▶ identify the network events, experimental signaling
  - ▶ As gleaned from BGP data (future work)

```
% Demo network Experiment
START
SLEEP 5000
DUMP
DROP 1 1.0.0.0/8
...
ADD 62 62.0.0.0/8
...
FAIL 80 81
...
RECOVER 80 81
ADD 1 1.0.0.0/8
...
STOP
```

# IPFilter (router filters)


- Allows us to create route filtering rules for the control of routing
  - Commonly used in networks avoid obvious problems
  - E.g., implement *bogon* blacklist

- Syntax,

```
IPFILTER <AS #> <PERMIT/DENY> <PRFX #> <NAS #> <IN/OUT>
```

- Example: *avoiding the valley*

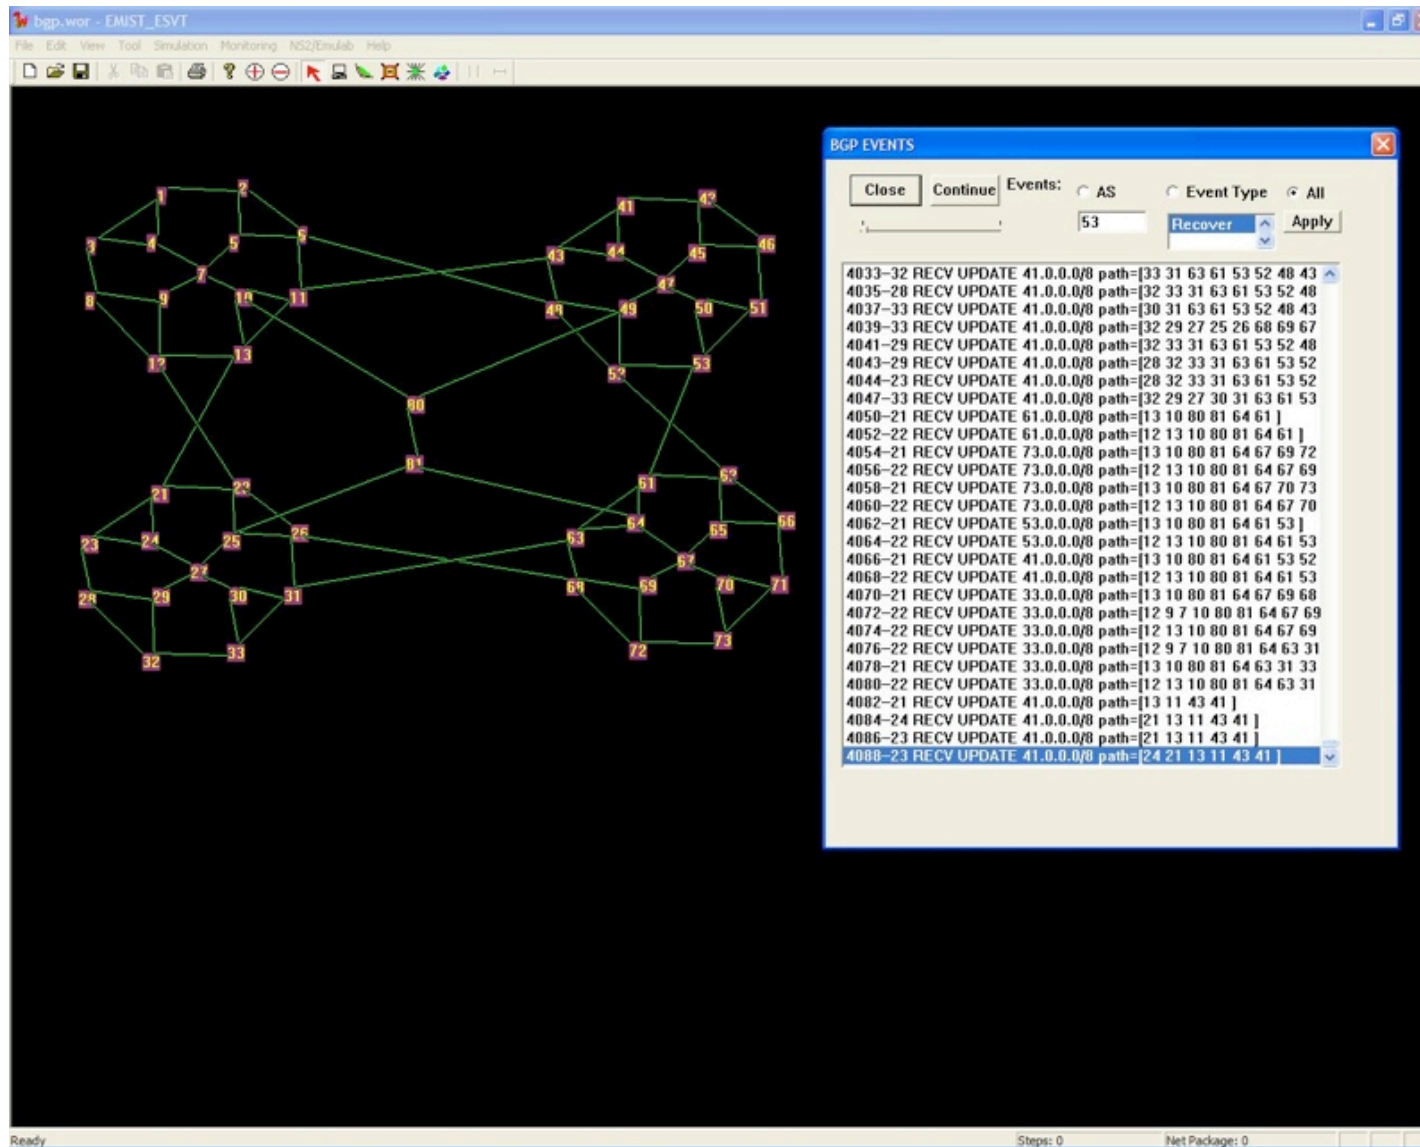
```
IPFILTER 12 DENY 12.0.0.0/8 14 IN
```



# Simulation Walkthrough ...

- A process is started at each host associated with the simulation
  1. The hosts *self-organize* by contacting the master process (this is the control channel)
  2. The master process advertises the mapping of process IDs to IP address/port numbers
  3. The master informs the processes of the setup of the experiment (AS assignments, topology, prefixes)
  4. Each host contacts the peer processes as needed by experimental topology
- The experimental script is executed in order
  - ▶ All data is logged in master log (over control channel)
  - ▶ Models “real time” BGP via messaging
  - ▶ Limited synchronization (*scalability*)

# Experimental Visualization



# Experimental Visualization (cont.)

The screenshot displays the 'bgp-wor - EMIST\_ESVT' application interface. The main window shows a network diagram with nodes and connections. Two dialog boxes are open:

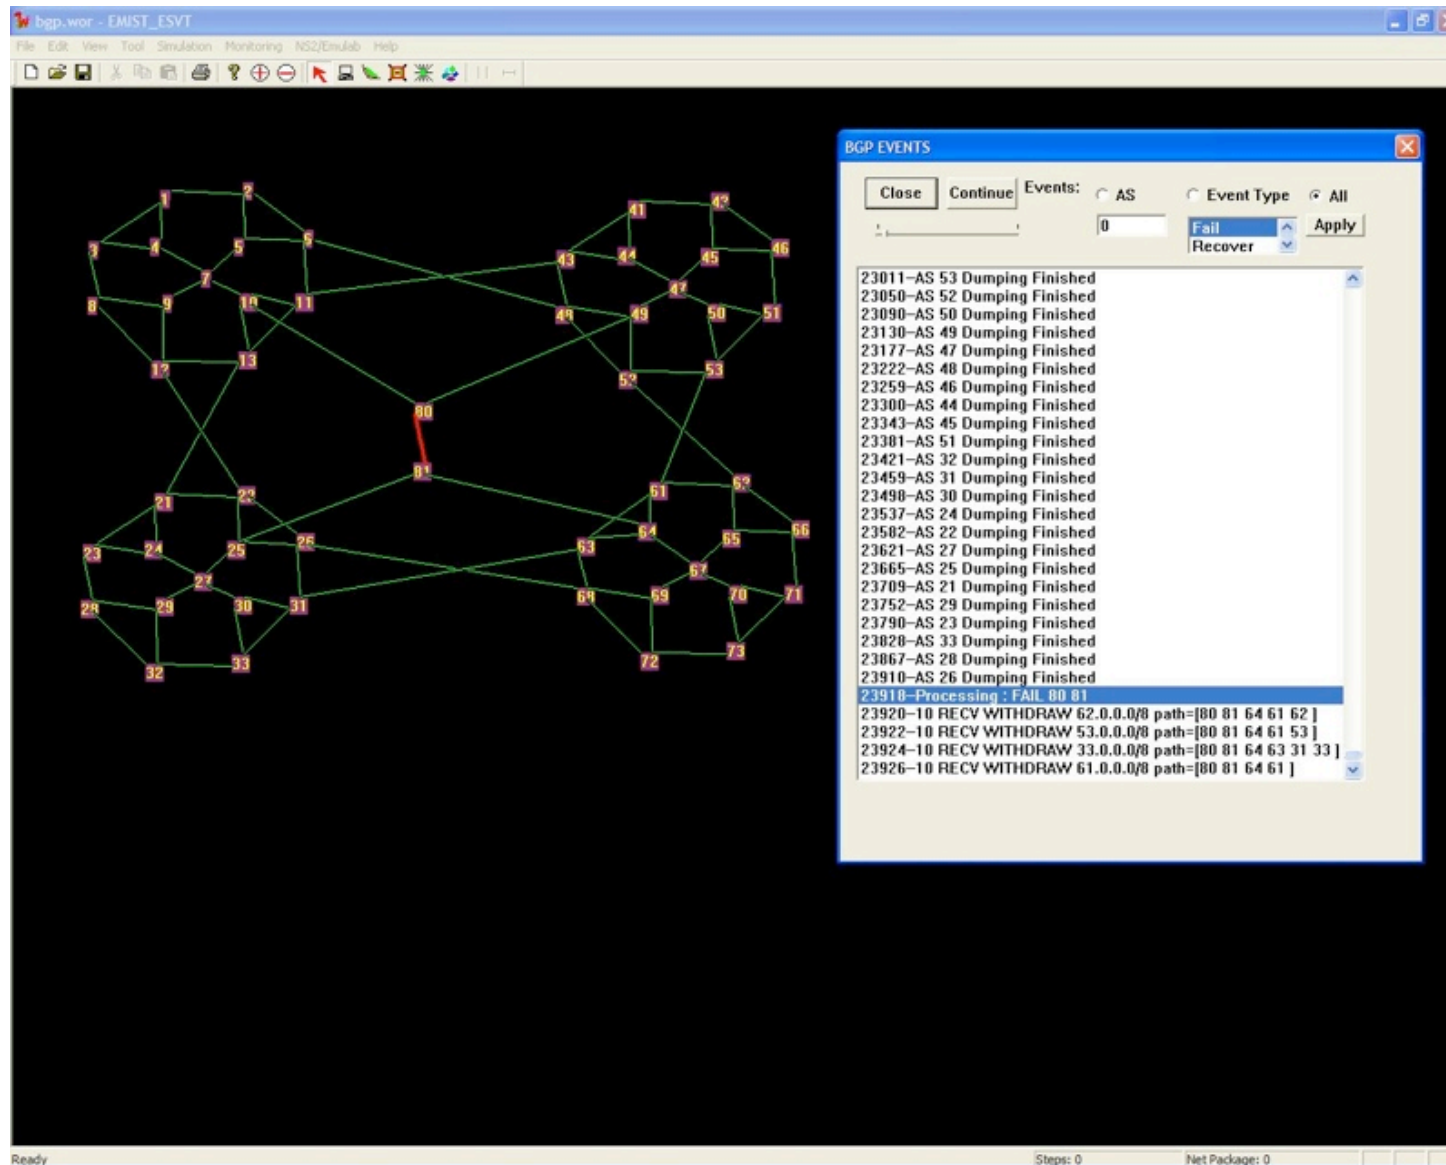
**AS Information Dialog:**

- AS: 53
- AS Prefix List: 53.0.0.0/8
- AS Routing Table:
  - \*\* AS 53
  - PEERS : 52 51 61 50
  - ORIGINS : 53.0.0.0/8
  - ROUTES —
  - Prefix 62.0.0.0/8 \*\*\* PATHS \*\*\*
  - \*[52 62 ], length 2
  - [61 62 ], length 2
  - Prefix 53.0.0.0/8 \*\*\* PATHS \*\*\*
  - \*[], length 1
  - Prefix 33.0.0.0/8 \*\*\* PATHS \*\*\*
  - \*[61 63 31 33 ], length 4
  - Prefix 1.0.0.0/8 \*\*\* PATHS \*\*\*
  - [50 47 49 48 6 11 10 7 4 1 ], length 10
  - \*[52 49 48 6 11 10 7 4 1 ], length 9
  - [51 50 47 49 48 6 11 10 7 4 1 ], length 11
  - [61 62 52 49 48 6 11 10 7 4 1 ], length 11
  - Prefix 13.0.0.0/8 \*\*\* PATHS \*\*\*
  - [61 63 68 26 22 21 13 ], length 7
  - \*[52 48 6 11 13 ], length 5
  - [50 47 44 43 11 13 ], length 6
  - [51 46 42 41 43 11 13 ], length 7

**BGP EVENTS Log:**

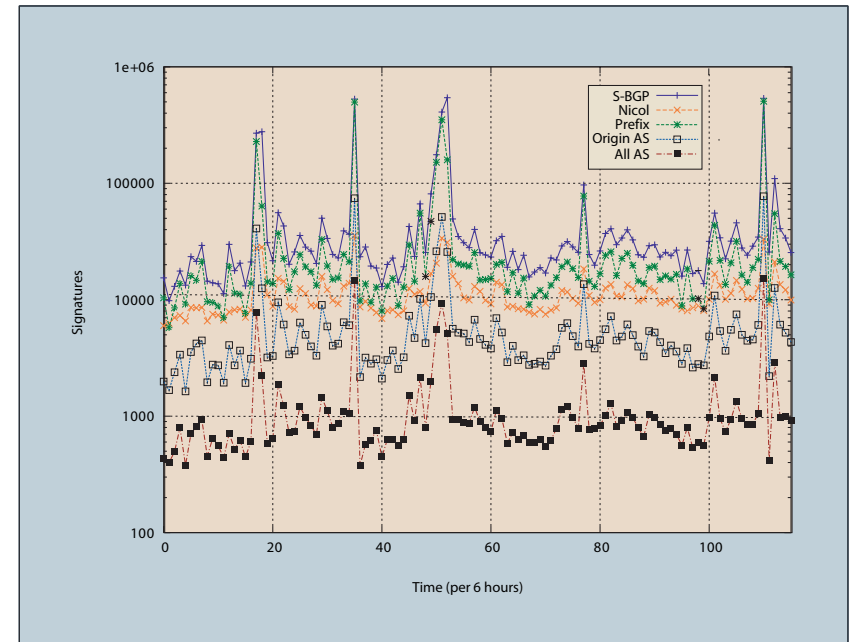
```
12943-53 RECV UPDATE 1.0.0.0/8 path=[61 63 64 81 80 10 11 6 2 ^
12945-53 RECV UPDATE 1.0.0.0/8 path=[61 64 81 80 10 11 6 5 2 ^
12947-53 RECV UPDATE 1.0.0.0/8 path=[61 64 81 80 10 7 9 8 3 1 ^
12949-53 RECV WITHDRAW 1.0.0.0/8 path=[61 64 81 80 10 7 9 8 ^
12951-53 RECV UPDATE 1.0.0.0/8 path=[61 62 52 49 48 6 11 10 7 ^
12953-53 RECV UPDATE 1.0.0.0/8 path=[61 64 81 80 10 11 6 5 2 ^
12955-53 RECV WITHDRAW 1.0.0.0/8 path=[61 64 81 80 10 11 6 !
12957-53 RECV UPDATE 1.0.0.0/8 path=[61 62 52 49 48 6 11 10 7 ^
12959-53 RECV UPDATE 1.0.0.0/8 path=[61 63 64 81 80 10 11 6 5 ^
12961-53 RECV UPDATE 1.0.0.0/8 path=[61 63 64 81 80 10 7 9 8 :
12963-53 RECV WITHDRAW 1.0.0.0/8 path=[61 63 64 81 80 10 7 !
12965-53 RECV UPDATE 1.0.0.0/8 path=[61 62 52 49 48 6 11 10 7 ^
12967-53 RECV UPDATE 1.0.0.0/8 path=[61 64 81 80 10 13 11 6 5 ^
12969-53 RECV UPDATE 1.0.0.0/8 path=[61 64 81 80 10 13 12 9 8 ^
12971-53 RECV WITHDRAW 1.0.0.0/8 path=[61 64 81 80 10 13 12 ^
12973-53 RECV UPDATE 1.0.0.0/8 path=[61 62 52 49 48 6 11 10 7 ^
12975-53 RECV UPDATE 1.0.0.0/8 path=[61 64 81 80 10 13 11 6 5 ^
12977-53 RECV WITHDRAW 1.0.0.0/8 path=[61 64 81 80 10 13 11 ^
12979-53 RECV UPDATE 1.0.0.0/8 path=[61 62 52 49 48 6 11 10 7 ^
20522-AS 53 Dumping Finished
21487-53 RECV UPDATE 62.0.0.0/8 path=[52 62 ]
21506-53 RECV UPDATE 62.0.0.0/8 path=[61 62 ]
23011-AS 53 Dumping Finished
24097-53 RECV WITHDRAW 13.0.0.0/8 path=[61 64 81 80 10 13 ]
24210-53 RECV UPDATE 13.0.0.0/8 path=[61 63 64 81 80 10 13 ]
24212-53 RECV WITHDRAW 13.0.0.0/8 path=[61 63 64 81 80 10 1 ^
24214-53 RECV UPDATE 13.0.0.0/8 path=[61 63 68 26 22 21 13 ]
25865-AS 53 Dumping Finished
```

# Experimental Visualization (cont.)



# DETER Simulated Attacks

- Prefix/AS Hijacking
  - ▶ outright stealing prefixes
  - ▶ whack-a-mole ASes
- Worm holing
  - ▶ Illusory link creation
- Link cutting
  - ▶ Observing the affects of link cuts on selected paths
  - ▶ Simple replication/confirmation of Bellovin result
- Countermeasure testing
  - ▶ Filtering rule effectiveness on route manipulation
  - ▶ Working : policy driven path selection, e.g., attributes, static routes



# Iseb Status

- Working simulator
  - ▶ Large Java-based implementation
  - ▶ Multiprocessor
  - ▶ Self organization
  - ▶ Complete recreation of BGP messaging
  - ▶ Forensic visualization (thanks Peng/Lunquan)
  - ▶ Running on DETER (thanks Ihab)
- Working experiments
  - ▶ 54 AS topology executing on single processor and 4 processors
  - ▶ Executed experiments in real time (scaling?)
  - ▶ Submission impending for large-scale experiment

# Future work/next steps

- *Realism*
  - ▶ Extension of path selection algorithm (policy extraction)
  - ▶ Delay modeling (trace analysis, PREDICT)
- *Attack Modeling*
  - ▶ Extending experimental lexicon to model real attacks
  - ▶ Tools integration
- *Optimization*
  - ▶ Shared tables, message aggregation, ...
- *Analysis*
  - ▶ Extending measurement apparatus
  - ▶ Extending modeled countermeasures
  - ▶ DETER: validation

# Availability

- Release 1.1 is will be made available within the next few months.
  - ▶ With bgpevent, bgptopo, and other tools
  - ▶ If you would like the pre-release, please contact

[siis@cse.psu.edu](mailto:siis@cse.psu.edu)