

# Beyond Simulation: Large-Scale Distributed Emulation of P2P Protocols

Nathan S. Evans    Christian Grothoff

Technische Universität München




August 8, 2011



**Presented by: Bartłomiej Polot and Matthias Wachs**


## Systems Research — Simulation



Ideal use of simulation for security testing.




## Systems Research — Simulation



Ideal use of simulation for security testing.




## Systems Research — Simulation




## Systems Research — Emulation



DETER




# Our Emulation Approach




# GNUnet Architecture

- P2P framework
- Focus on security
- Written in C
- Portable & extendable
- Multi-process architecture & IPC
- Extensive utility library



# The Transport Service

- Low-level P2P connectivity
- Transport plugins:  
provide many connection options
- Unix domain sockets
- Blacklisting & whitelisting




## P2P Emulation Steps


- Design P2P application
- Implement as GUNet service
- Use built-in statistics or design logging facility
- Create test control application
  - Links against emulation library
  - Peer group startup/shutdown
  - Utilizes API to access service



# Our Emulation Approach



# Single Peer Startup Sequence



# Peer Group

- “Peer group” is the handle to running peers
- Layering — peer group reuses single peer startup code
- Peer group features
  - Configuration mangling
  - Resource allocation, throttling
  - Connects peers in desired topology
  - Capture running topology/statistics
  - Start/stop/reconfigure peers
  - Induce churn
  - Provide handles to specific peers



## Peer Group Startup, Code Example

```

1 struct GNUNET_TESTING_PeerGroup *
2 GNUNET_TESTING_peergroup_start(const struct GNUNET_CONFIGURATION_Handle *cfg ,
3                               unsigned int total ,
4                               struct GNUNET_TIME_Relative timeout ,
5                               GNUNET_TESTING_NotifyConnection connect_cb ,
6                               GNUNET_TESTING_NotifyCompletion peergroup_cb ,
7                               void *peergroup_cls ,
8                               const struct GNUNET_TESTING_Host *hosts);

```

```


1 GNUNET_CONFIGURATION_load (testing_cfg , "~/test.conf");
2 struct MyClosure *data; /* your data here */
3 struct GNUNET_TESTING_Host *hosts = GNUNET_TESTING_HOSTS_load("~/hosts.conf");
4 pg = GNUNET_TESTING_peergroup_start (testing_cfg , 20000, TIMEOUT, &connect_cb ,
5                                     &peergroup_cb , data , hosts);
6 /* peergroup_cb must eventually call: */
7 GNUNET_TESTING_daemons_stop (pg, TIMEOUT, &shutdown_cb , data);

```




# Network Topologies


*Simple*  
topology  
creation/im-  
port/export




(a) 2d-grid



(b) Small-World




(c) Erdos-Renyi




(d) InterNAT




# Topology Generation and Evolution




(a) Initial




(b) 2 Minutes




(c) 5 Minutes



(d) 10 minutes



(e) 15 minutes



(f) 30 minutes



# Limitations of Emulation

- Timing accuracy
  - Network latency
  - Throughput
- Underlying OS interference
  - CPU scheduling
  - Disk access
  - Memory usage
- Speed
- Shared IP/hostnames
- Peer diversity
- GUNet





# Overcoming Limitations

- Single OS per peer
  - ⇒ Testing framework can be used on lower level emulators focused on timing accurate results
    - PlanetLab, Emulab, DETER, etc.
    - Sacrifice scalability
- Shared IP/hostnames — Virtual addresses, VMs
- Peer diversity — Configure per-peer bandwidth, VMs
- GUNet — Benefit and limitation



# Important Lessons Learned

- Cryptography
- Start-up time
- Periodic tasks
- Sockets
- Memory



## Peer and Emulation Performance

## Memory consumption

<b>Service</b>	<b>Non-shared</b>	<b>Heap</b>	<b>Shared</b>
supervisor	228 KB	32 KB	2,364 KB
transport	359 KB	99 KB	2,888 KB
core	300 KB	84 KB	2,428 KB
dht	536 KB	240 KB	3,684 KB
<b>total</b>	<b>1,424 KB</b>	<b>456 KB</b>	<b>11,364 KB</b>



## Peer and Emulation Performance

Architecture	Hosts	Cores (Total)	Memory (Total)	Peers	Connections per second	Time to start peer
Cortex-A8	1	1	512 MB	100	~ 1	~ 206 ms
Xeon W3505	1	2	12 GB	2,025	~ 60	~ 12 ms
Xeon W3520	1	8	12 GB	2,025	~ 188	~ 5 ms
Opteron 8222	1	16	64 GB	10,000	~ 327	~ 27 ms
Opteron 850	31	124	217 GB	80,000	~ 559	~ 1 ms



## Example: Comparison of DHT Performance

- Performance comparison of different DHT implementations
- 60,000 peers
- Specific peers were changed into malicious sybil nodes
- Success rate of requests measured



## Example: NSE Implementation

- Network Size Estimation algorithm
- 2 days to implement
- 2 weeks from idea to paper
- Single host: 4,000 peers



# Conclusion

- Framework available at <https://gnunet.org>
- We encourage people to use our framework
- 80,000 peers on cluster:  
what happens on supercomputer?
- at least consider:  
emulation vs. simulation even at large scale



# Questions?



The End.

N. Evans, C. Grothoff (TUM)

August 8, 2011

24