Automated Attack Discovery in TCP Congestion Control Using a Model-guided Approach

Samuel Jero[1], Endadul Hoque[2], David Choffnes[3], Alan Mislove[3] , and Cristina Nita-Rotaru[3]
[1]Purdue University, [2]Florida International University, and [3]Northeastern University
Appeared in NDSS 2018

TCP Congestion Control Attacks

Congestion Control
- Prevents Congestion Collapse
- Ensures fairness between flows

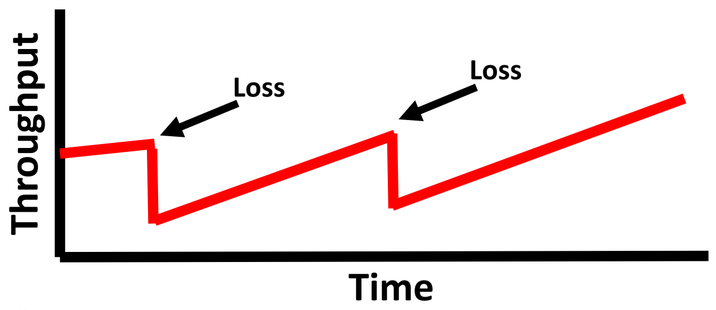

Long history of powerful attacks

Impacts include
- **Decreased Throughput**
- **Increased Throughput, starving other flows**
- **Connection Stalls**

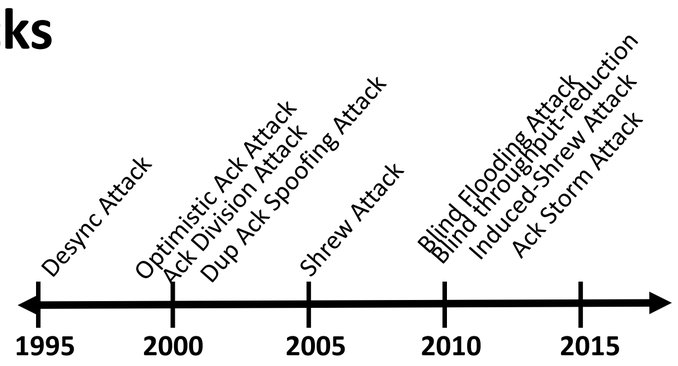

Why so many attacks?
- Hundreds of implementations and variations
- Lack of unified specifications
- Complex, highly dynamic behavior

Can we automatically test implementations for attacks?

Key Challenge: Scalability, attacks are complex, multi-stage and the system is highly dynamic

Model-based Attack Discovery

1) Model Congestion Control as a State Machine

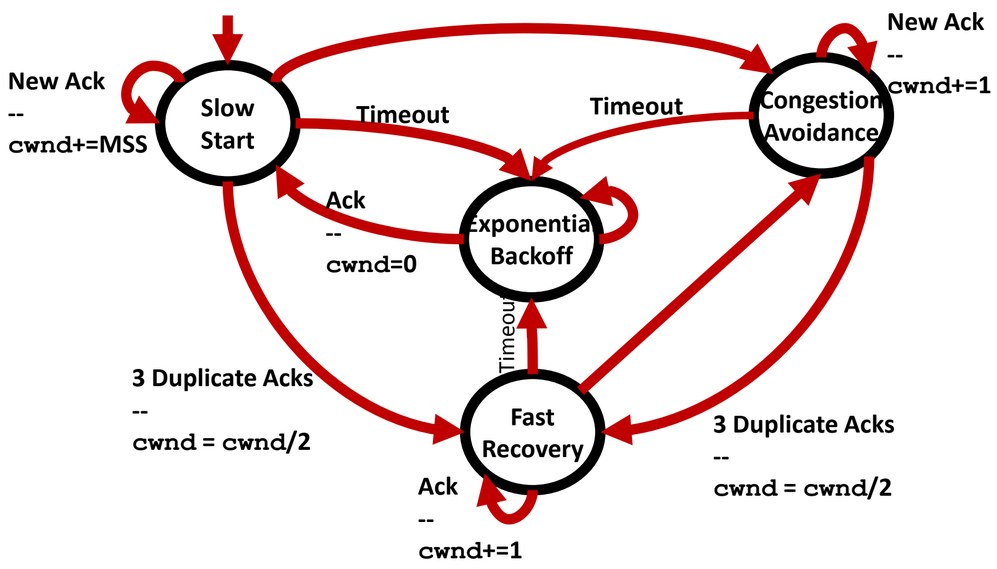

2) Create Abstract Strategies from State Machine

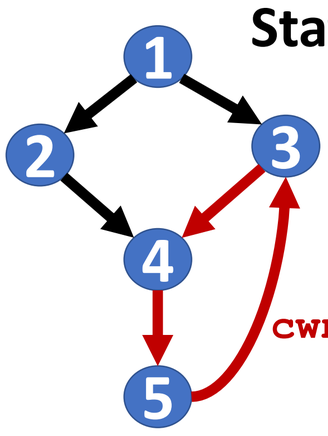

Any Attack MUST:
- **Change `cwnd`**
- **Cause a Cycle**

Enumerate all paths that contain cycles and change `cwnd`

3) Create Concrete Strategies from Abstract Strategies

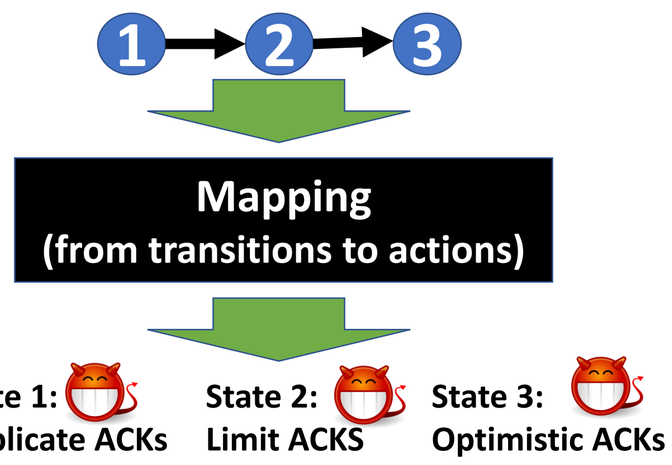

4)Apply Concrete Strategies to Real Implementations

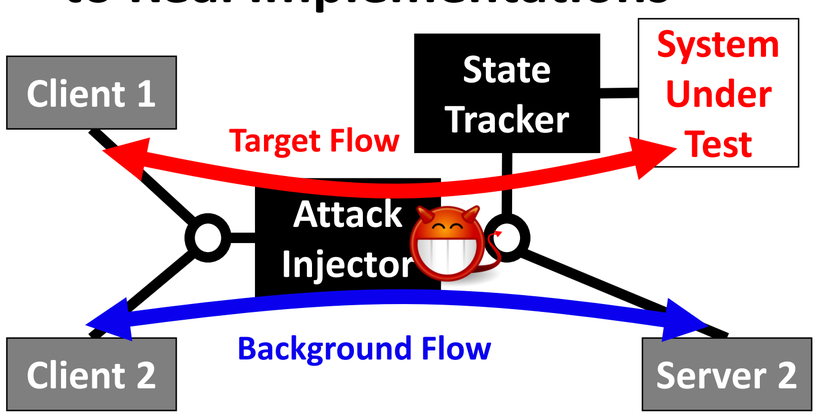

Measure throughput and fairness to identify attacks

Evaluation

Evaluated 5 TCP implementations

Implementation	Date
Ubuntu 16.10 (Linux 4.8)	2016
Ubuntu 14.04 (Linux 3.13)	2014
Ubuntu 11.10 (Linux 3.0)	2011
Debian 2 (Linux 2.0)	1998
Windows 8.1	2014

New Attacks

Attack Class	Impact
On-path Repeated Slow Start	Increased Throughput
Amplified Bursts	Increased Throughput
Ack Lost Data	Connection Stall
Slow Injected Acks	Decreased Throughput
Sawtooth Ack	Decreased Throughput
Dup Ack Injection	Decreased Throughput
Ack Amplification	Increased Throughput
Off-path Repeated Slow Start	Increased Throughput

Found 11 classes of attacks, 8 of which are new

Acknowledgements and Contact Info
For more information about this project, contact: Samuel Jero <sjero@sjero.net>. Or see our paper in NDSS 2018.
This material is also based upon work partially supported by the National Science Foundation under Grant Numbers CNS-1600266, CNS-1617728, and CNS- 1409191. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.