



Pollution Attacks and Defenses for Internet Caching Systems

Leiwon Deng, Yan Gao, Aleksandar Kuzmanovic, and Yan Chen (Northwestern Univ.)



<http://networks.cs.northwestern.edu/AE/>

NSF Grant CNS-0627715

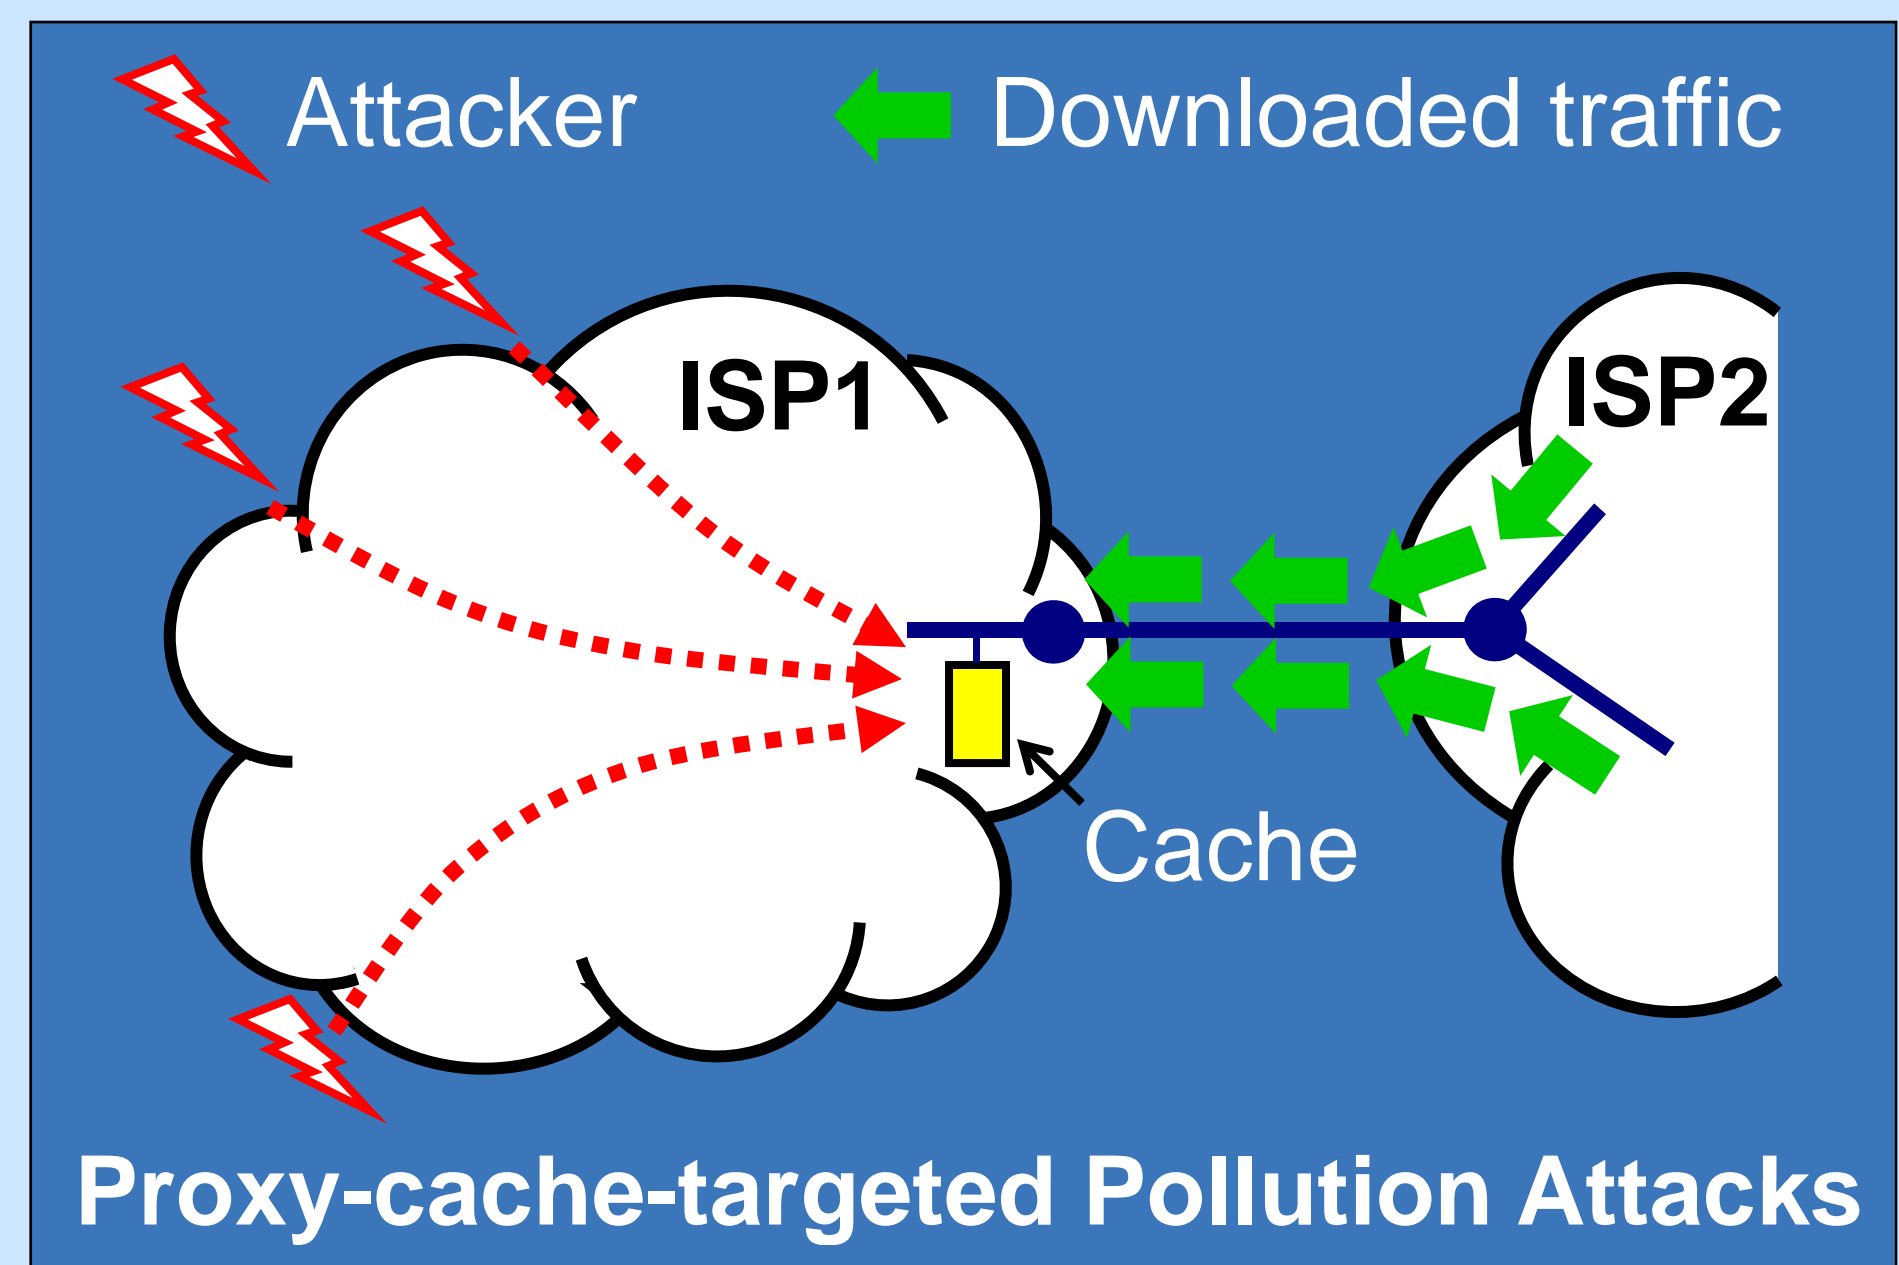
Problem

Proxy caching servers are widely deployed in today's Internet. While cooperation among proxy caches can significantly improve a network's resilience to denial-of-service (DoS) attacks, lack of cooperation can transform such servers into viable DoS targets.

Approach

We investigate a class of pollution attacks that aim to degrade a proxy's caching capabilities, either by ruining the cache file locality (*locality-disruption* attack), or by inducing false file locality (*false-locality* attack).

We evaluate the effects of pollution attacks both in Web and peer-to-peer (p2p) scenarios. We also develop efficient countermeasures against pollution attacks.



Approach and Impact

New approach

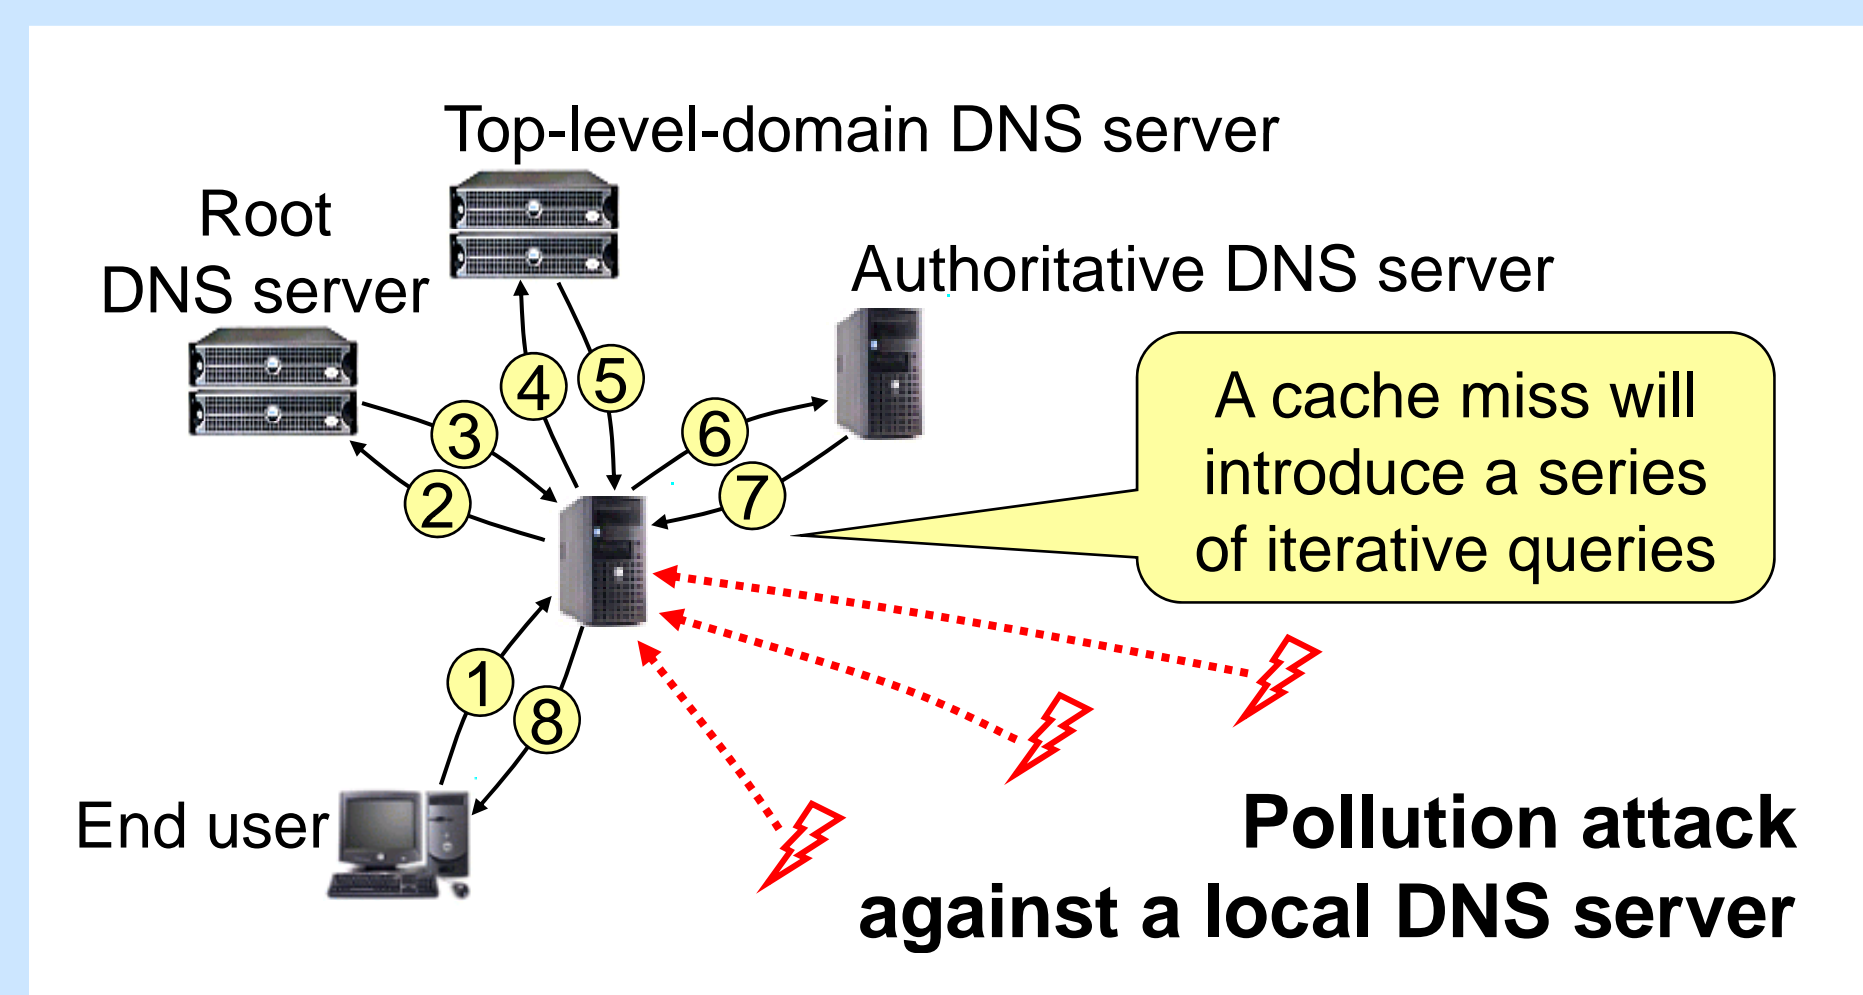
- Investigate a new class of DoS attacks
- Scalable countermeasures
- Anti-pollution engine (a Squid-based implementation)

Research Impact

- Analyze cache's resilience to pollution attacks
- Counter-pollution techniques
- Attacker-based and object-based detections

Technical Description

- *Locality-disruption* attacks continuously generate requests for new unpopular files, thus ruining the cache file locality. *False-locality* attacks repeatedly request the same set of files, thus creating a false file locality at proxy caches.
- Using simulations, we reveal dramatic variability in cache's resilience to pollution attacks among several cache replacement algorithms (GDSF, LRU, LFU).
- The cache pollution attacks are very stealthy because they can be easily mixed with and regarded as normal clients' requests. Thus, no existing schemes are capable of detecting such attacks.
- Local (low-level) DNS servers can be a potential target of pollution attacks. A set of malicious attackers may pollute the local DNS



server's cache with unpopular entries, thus significantly reducing the performance experienced by regular clients.

- We develop efficient methods to detect both false-locality and locality-disruption attacks, as well as a combination of the two. We leverage streaming computation techniques such as bloom filter and probabilistic counting to achieve high scalability for a large number of clients/requests without sacrificing detection accuracy.