Denial of Service and Anomaly *Detection*

Vasilios A. Siris

Institute of Computer Science (ICS)
FORTH, Crete, Greece
vsiris@ics.forth.gr

SCAMPI BoF, Zagreb, May 21 2002

Overview

- What the problem is and why it is difficult
- Where and why naïve schemes fail
- Consider two algorithms
  - Adaptive Threshold
  - CUSUM (CUmulative SUM)
- Application to SYN attack detection
- Experimental results
- Conclusions and future work

Denial of Service (DoS) attacks

- Aim is to prevent users from receiving service, with some minimum performance
- Achieved by consuming resources
  - Bandwidth
  - Memory
  - Router forwarding capacity
  - Other services: DNS
- Technique: flooding

Importance of DoS attacks

- Recent surveys:
  - 40% of all attacks are DoS (2002 CSI/FBI)
  - 90% of all DoS attacks are TCP attacks (2001 Moore et al)
- Cost of attack = many € or $
  - Several millions to billions $ estimated loss from Feb 2000 attack at Yahoo, CNN, Amazon, etc
- Attacks are increasing
  - DNS route server attack in Oct. 2002
  - DOLnet's attack in Dec. 2002
  - 55% Web attacks are DoS (2002 CSI/FBI)

The DoS problem

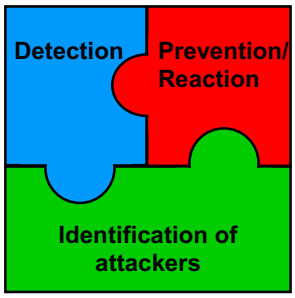

- Our focus on detection of DoS attacks
  - Early and reliable detection of attacks
  - Detection of low intensity attacks

Distributed DoS attack

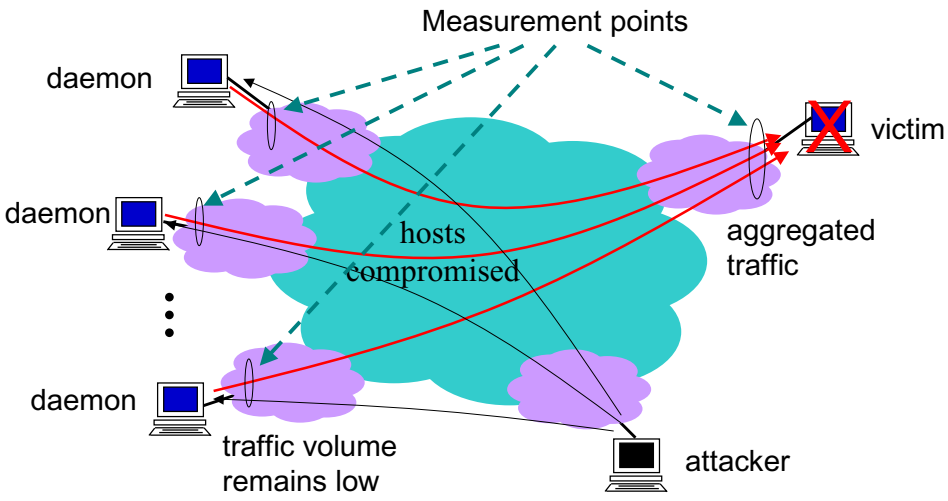

Approaches to anomaly detection

- Alarm when behavior deviates from normal
- Specify normal behavior (operational model)
  - Thresholds: e.g. load < 0.7
- Learn normal behavior
  - Mean and standard deviation statistics
  - Time series analysis: advantage is that they take into account time correlations
    - Change point detection (hypothesis testing)
  - Other approaches: bayesian statistics, neural nets
- DoS attacks one example of anomaly
  - Link/device failures

Non-adaptive approaches not robust

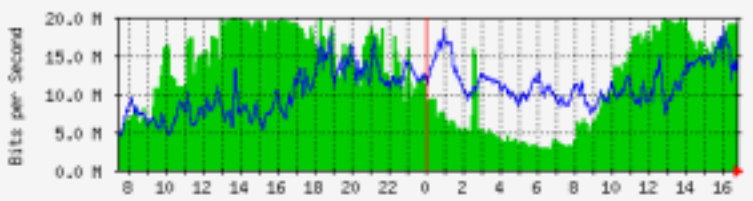

- Fixed threshold tests (e.g. normal < 0.7) will fail due to normal/regular traffic variations
- Why not consider an adaptive threshold ?

Detection of some attacks simpler

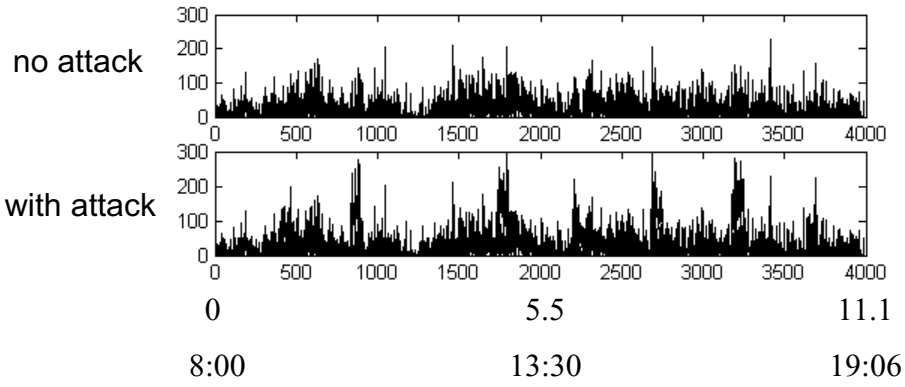

Detection of some attacks simpler

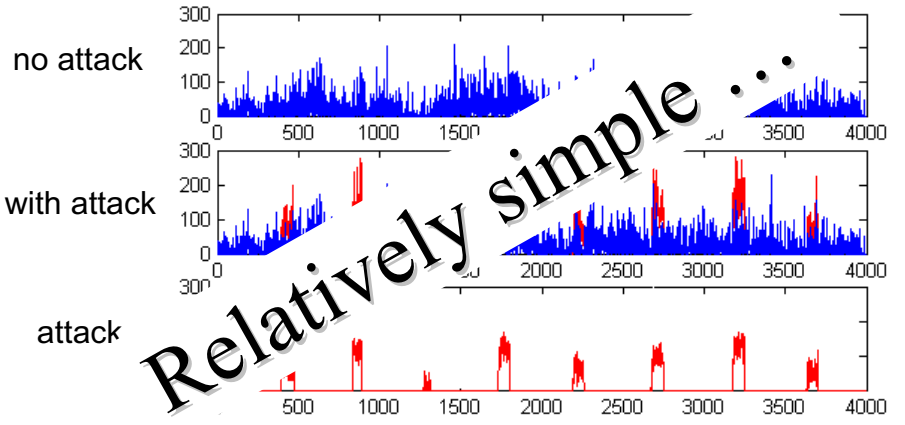

Some attacks are more subtle

no attack

with attack

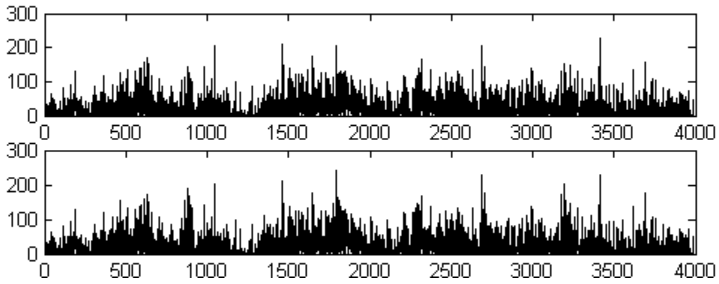

Some attacks are more subtle

no attack

with attack

attack

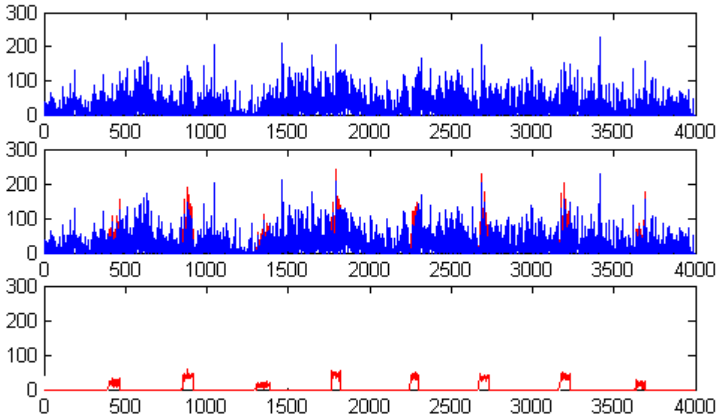

What and when to measure

- Variable measured:
  - Aggregate traffic volume (in fixed time intervals)
  - Traffic volume per flow (in fixed time intervals)
  - # of requests, e.g. TCP, http, …
  - Inter-arrival time of requests
  - Duration of requests (average or bin)
  - Pkt size (average or bin)
- Statistic: Mean, variance, covariance, hurst
- When to measure: order of seconds
  - 10 seconds in our experiments

Algorithms investigated

- Adaptive threshold
  - Adaptively measure mean rate
  - Alarm when rate more than some percentage (e.g. > 150% of mean)
- CUSUM (CUmulative SUM)
  - Adaptively measure mean rate
  - Sum the volume sent above some average factor
  - Alarm when volume more than some threshold

Adaptive Threshold (AT)

- Let y_t be time series of measurements
  - E.g. # of SYN packets in an interval T
- Mean μ_t measured over some past window L
  - By adaptively measuring mean can adjust to periodic (non-stationary) changes
- Alarm condition

$$\text{If } \; y_t > \beta\mu_t \; \text{ Alarm at } t$$

- Parameters:
  - T (measurement interval), L (averaging interval), β>1 (threshold)

Adaptive Threshold k (AT-k)

- More robust if alarm set when threshold exceeded for # k of consecutive intervals
- Alarm condition

$$\text{If } \; \sum_{i=t-k}^{t} 1_{\{y_i > \beta\mu_i\}} > k \; \text{ then } \; \text{ALARM at } t$$

- Parameters:
  - T (measurement interval), L (averaging interval), β (threshold), k (# of intervals threshold exceeded)

Adaptive Threshold: intuition

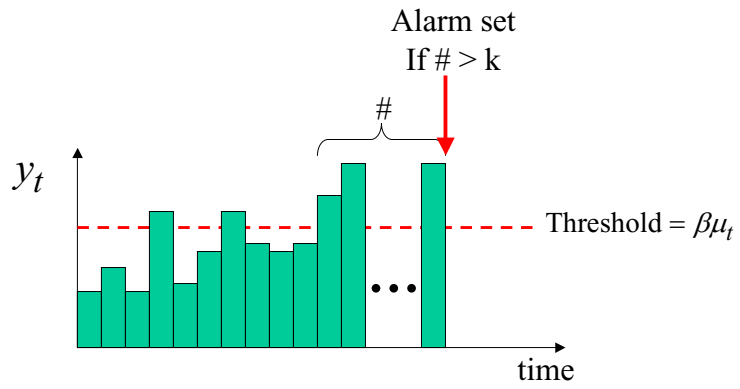

- Assuming fixed mean μ_t

CUSUM algorithm

- Based on hypothesis testing
- Current hypothesis (no attack): θ_0
- Alternative hypothesis θ_1 : $\mu_1 = \beta\mu_0 \quad \sigma_1 = \sigma_0$

$$s_i = \ln \frac{p_{\theta_1}(y_i)}{p_{\theta_0}(y_i)}$$

$$S_t = \sum_{i=0}^{t} s_i \qquad S_{\min} = \min_{0<k\le t} S_k$$

- Alarm condition

$$\text{If} \quad S_t - S_{\min} > h \quad \text{then} \quad \text{ALARM at } t$$

- Parameters: β (surplus), h (alarm threshold)

CUSUM algorithm: another view

- Mean μ estimated using EWMA
- Surplus: $\mu_1 = \mu_1' + \mu = \beta\mu$ (e.g. $\mu_1 = 1.5 \times \mu$)

$$g_t = \left[g_{t-1} + \frac{\mu_1'}{\sigma^2}\left(y_t - \frac{\mu + \mu_1}{2} \right) \right]^+$$

CUSUM algorithm: another view

- Mean μ estimated using EWMA
- Surplus: $\mu_1 = \mu_1' + \mu = \beta\mu$ (e.g. $\mu_1 = 1.5 \times \mu$)

$$g_t = \left[g_{t-1} + \frac{\mu_1'}{\sigma^2}\left(y_t - \frac{\mu + \mu_1}{2} \right) \right]^+$$

- Alarm condition

$$\text{If} \quad g_t > h \quad \text{then} \quad \text{ALARM at } t$$

- Parameters:
  - $\beta > 1$ (surplus), h (alarm threshold)

CUSUM algorithm: intuition

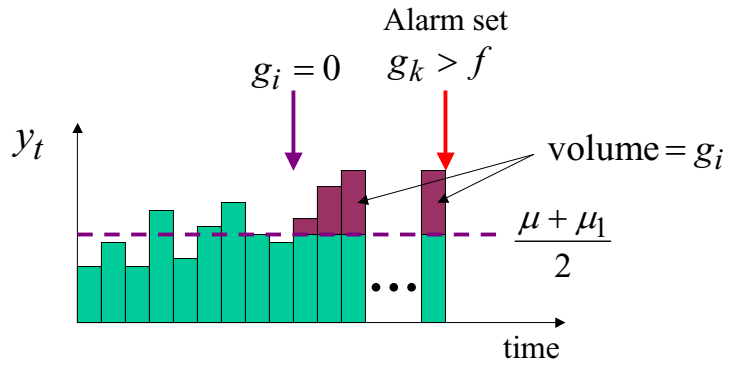

- Assuming $\frac{\mu + \mu_1}{2}$ constant
- Accumulates excess traffic (memory)

Types of DoS attacks

- TCP SYN flooding
- ICMP flooding
- UDP flooding
- SMURF attack

Application to SYN attack detection

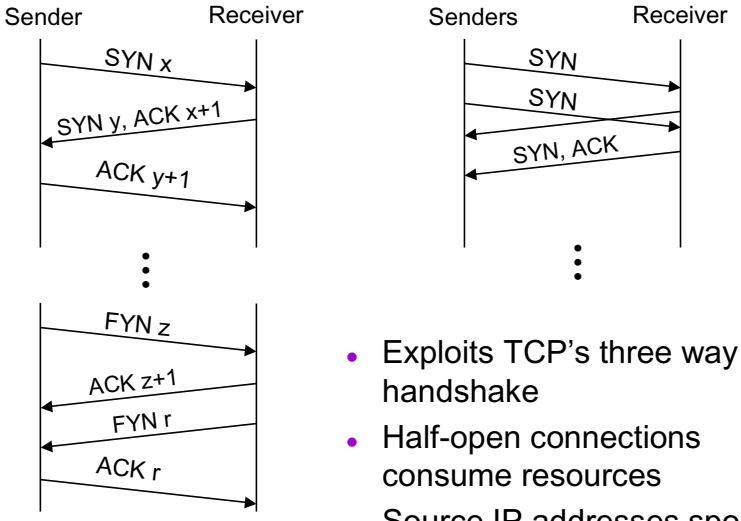

- Exploits TCP's three way handshake
- Half-open connections consume resources
- Source IP addresses spoofed

Performance measures

- Attack detection ratio
- False alarm ratio (false positives)
- Detection delay
- Robustness
- How tunable the algorithm is
  - Tradeoff between detection ratio, false alarm ratio and detection delay
- Evaluate above for different attack types
  - Intensity of attack (amplitude)
  - How fast it reaches peak amplitude

Experiments

- Considered real trace without attacks ~ 11 hours
  - # of SYN pkts in 10 second intervals
- 50 runs, 95% confidence interval
- Synthetic attacks
  - Intensity of attack (peak)
  - Time to reach peak
  - Inter-arrival: exponential, 400 sec

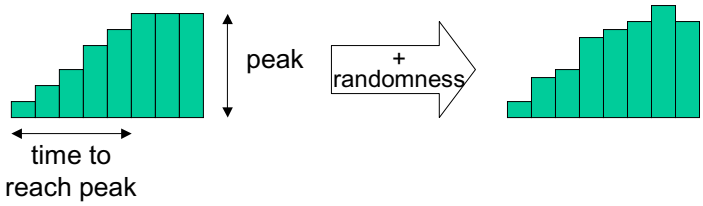

Adaptive Threshold – k

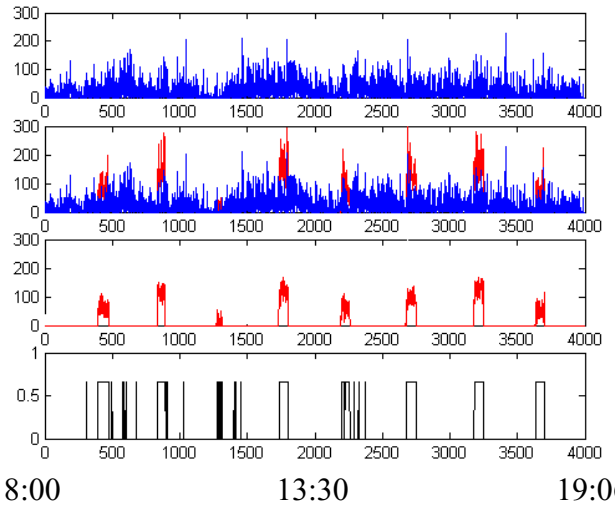

- Intense attack: rate ~ 250% mean

CUSUM

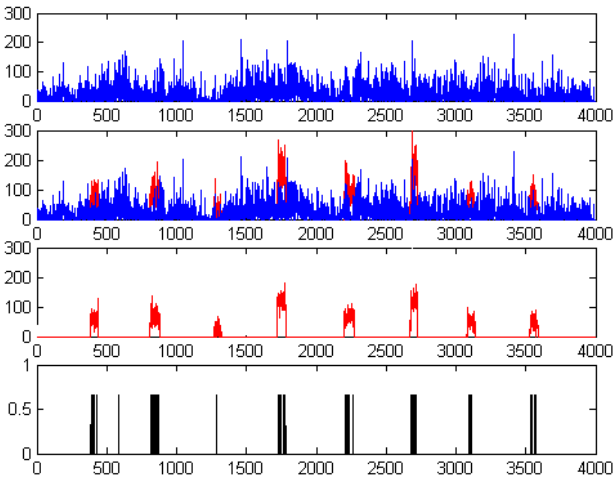

- Intense attack: rate ~ 250% mean

Adaptive Threshold – k

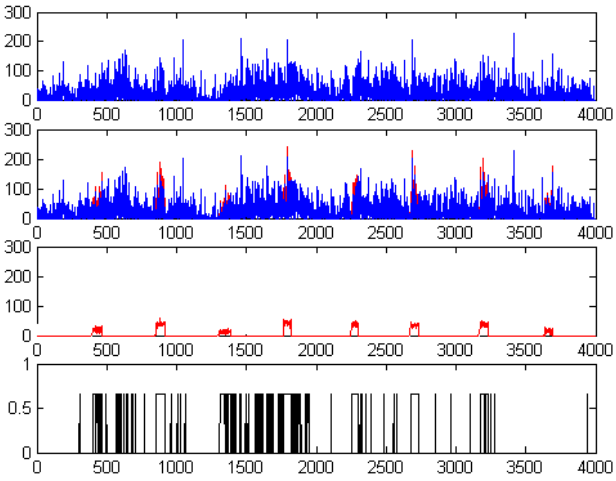

- small attack: rate ~ 10% mean

CUSUM

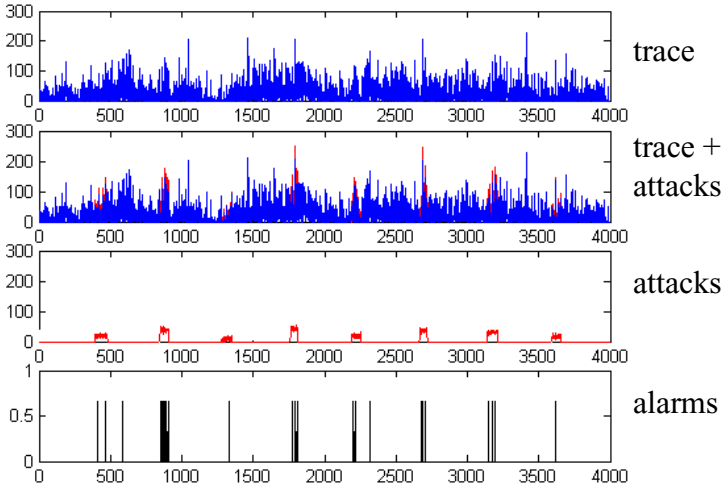

- small attack: rate ~ 10% mean

CUSUM

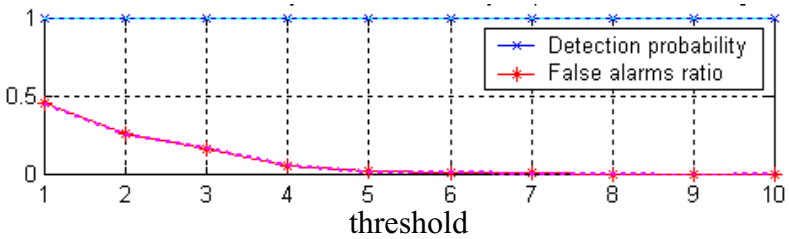

- Attack amplitude: 150% mean
- Time to reach peak: 90 sec

Adaptive Threshold - k

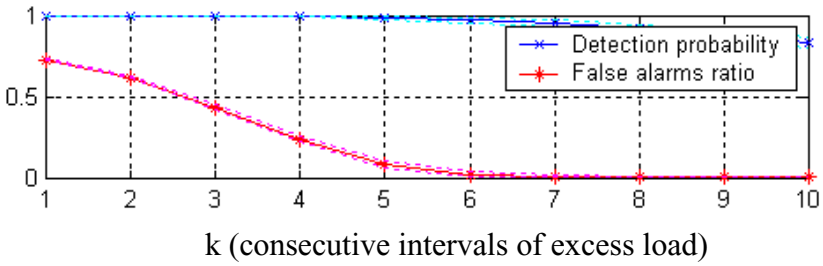

k (consecutive intervals of excess load)

- Attack amplitude: 150% mean
- Time to reach peak: 90 sec

AT-k versus CUSUM

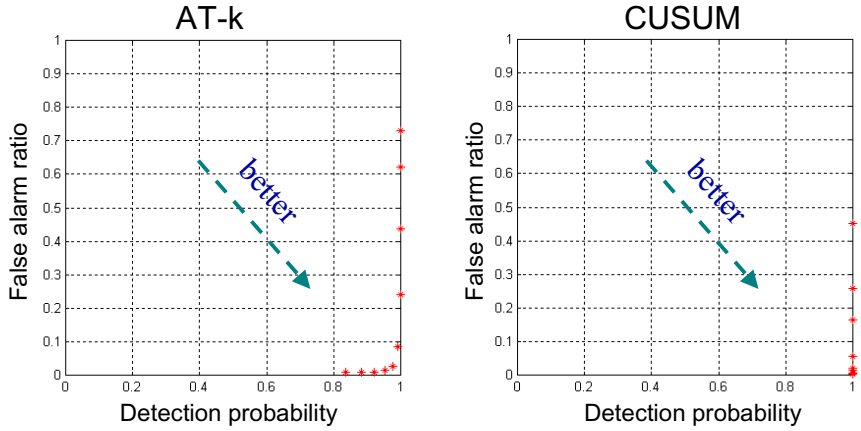

- Attack amplitude: 150% mean
- Time to reach peak: 90 sec

AT-k versus CUSUM

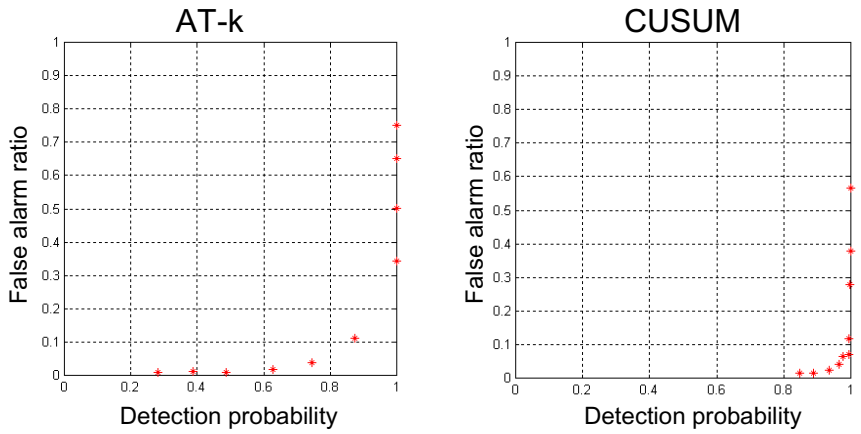

- Attack amplitude: 50% mean
- Time to reach peak: 90 sec

Adaptive Threshold - k

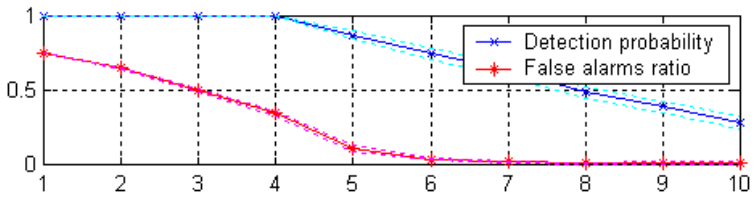

k (consecutive intervals of excess load)

- Attack amplitude: 50% mean
- Time to reach peak: 90 sec

CUSUM

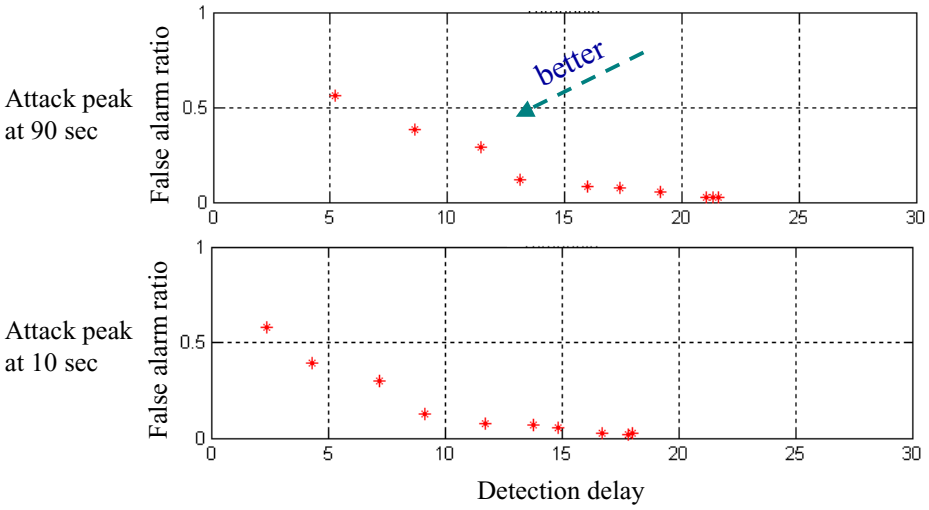

- Attack amplitude: 50% mean

Experiment results

- Performance depends on attack characteristics
- For some (intense) attack types straightforward procedures can be effective
- But simple procedures are not robust for different attacks
- Sound statistical methods are robust and not necessarily complex
- Intuition on how to tune parameters important

Future work

- Application to other measures & statistics
- Combination of alarms
- Application to QoS measurements
  - Measurements: delay, jitter, throughput
  - Up to now: alert when measurements exceed guarantees
  - Idea: apply anomaly detection to measurements => early detection of QoS violations

Denial of Service and Anomaly *Detection*

Vasilios A. Siris

Institute of Computer Science (ICS)
FORTH, Crete, Greece
vsiris@ics.forth.gr

SCAMPI BoF, Zagreb, May 21 2002