

The Limits of Global Scanning Worm Detectors in the Presence of Background Noise

David W. Richardson, Steven D. Gribble,
Edward D. Lazowska

University of Washington

November 11, 2005



Distributed worm detection systems

- Global scanning worms announce themselves
 - infected hosts randomly probe for additional victims
 - as worm grows, probe traffic grows proportionally
 - detectors can monitor unused IP addresses to track probe rate
- Distributed worm detection systems (DWDS) proposed
 - scan packet-based DWDS: count # of probes arriving
 - victims-based DWDS: count # of new hosts sending probes
 - both look for abnormally high rates to spot a growing worm
 - early detection = quick reaction



Background noise

- The Internet has high background “noise”
 - measurement packets, DDoS packets, old worms, port scans, etc.
 - noise rate is ~50 million packets/s Internet-wide, and is trending higher
- Hard to distinguish a noise probe from a worm probe
 - a DWDS sees an aggregate probe arrival rate
 - to spot a worm, it looks for a probe rate that is significantly higher than the typical background noise rate



Our main question

- How does Internet background noise impact the effectiveness of DWDSs?
- Intuition: a DWDS improves as more IPs are monitored
 - spot more worm probes, so able to raise the alarm earlier
 - but, monitoring more IPs also increases arriving noise rate!
- Is the intuition correct?
 - or does the added noise confound worm detection?



Contributions

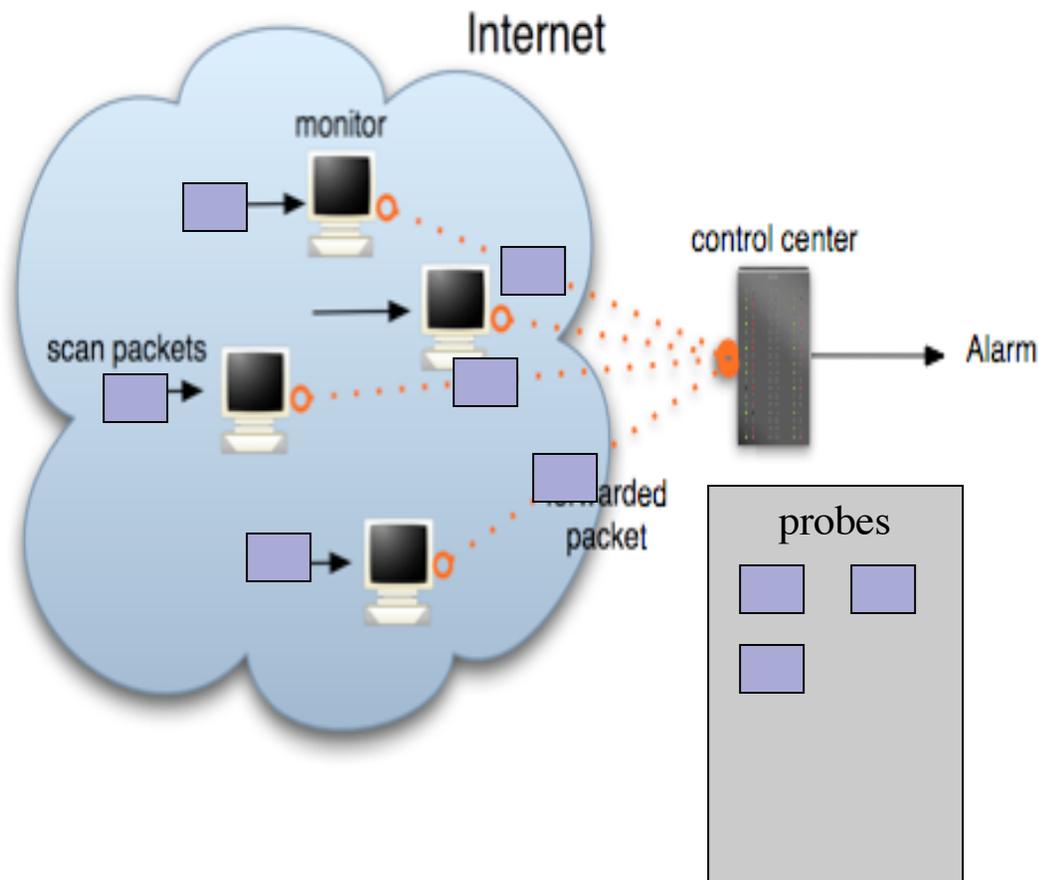
- Demonstrate how and why monitoring more IPs increases DWDS effectiveness, even in the presence of noise
- Using analytical modeling and simulation, quantify:
 1. the tradeoff between DWDS effectiveness and false alarm rate
 2. the impact of noise level on DWDS effectiveness
 3. the impact of # of monitored IPs on DWDS effectiveness
- Explore viability of DWDS if noise levels keep growing



Outline

- Introduction
- **Scan Packet-Based DWDS analysis**
 - Intuition
 - Formal analysis
 - Results
- Victims-Based DWDS analysis
- Conclusions

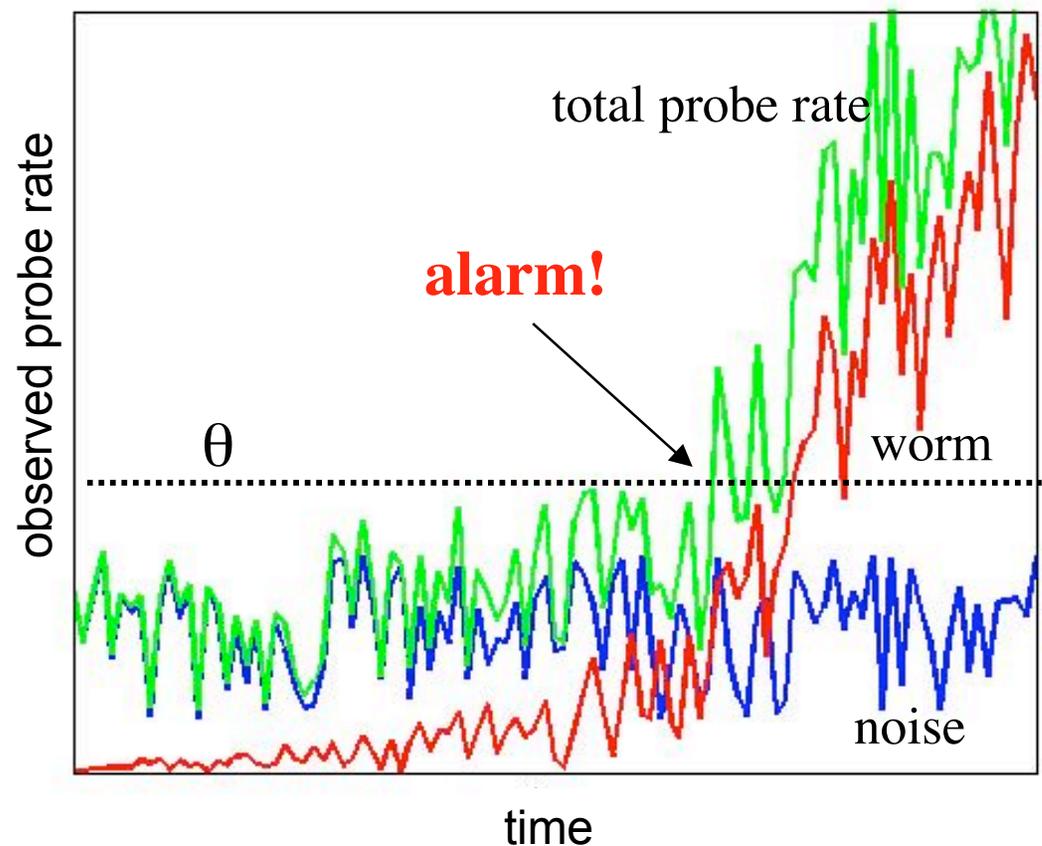
Scan packet-based DWDSs



- Monitor: 
 - IP address
- Probe: 
 - Noise or worm probe
- Control Center:
 - Stores/analyzes aggregate scan packets.
 - Raises alarm if total probe traffic \gg alarm threshold θ

The big picture

- DWDS observes normal range of fluctuations in noise
- Set threshold $\theta >$ likely noise fluctuations
- A fast-growing worm causes abnormal increases in *total* observed probe rate
- Raise alarm when:
total probe rate $> \theta$



The details

- Model of noise behavior

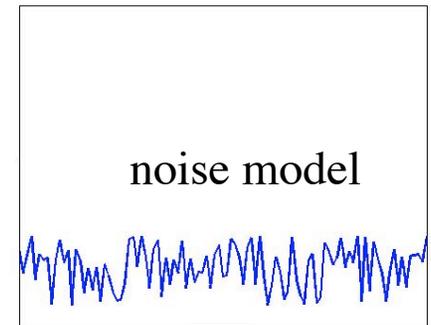
- describes mathematically the range of noise fluctuations
- allows us to derive appropriate θ

- Model of worm growth behavior

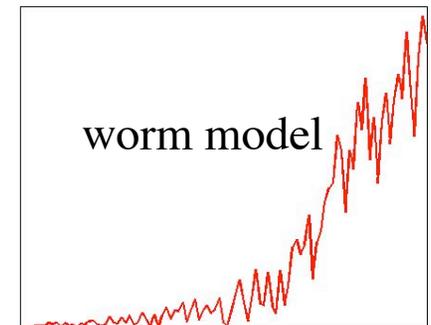
- Describes mathematically how worms grow

- Combine the 2 models:

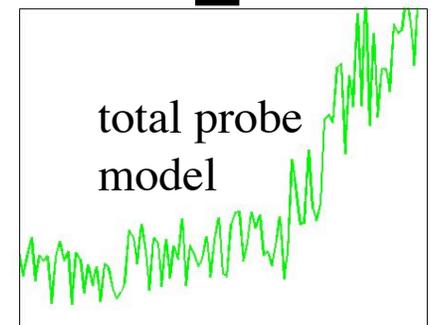
- describes total probes (noise + worm) seen during a worm outbreak
- allows us to calculate probability this model exceeds θ
 - i.e., probability an alarm is raised by DWDS



+

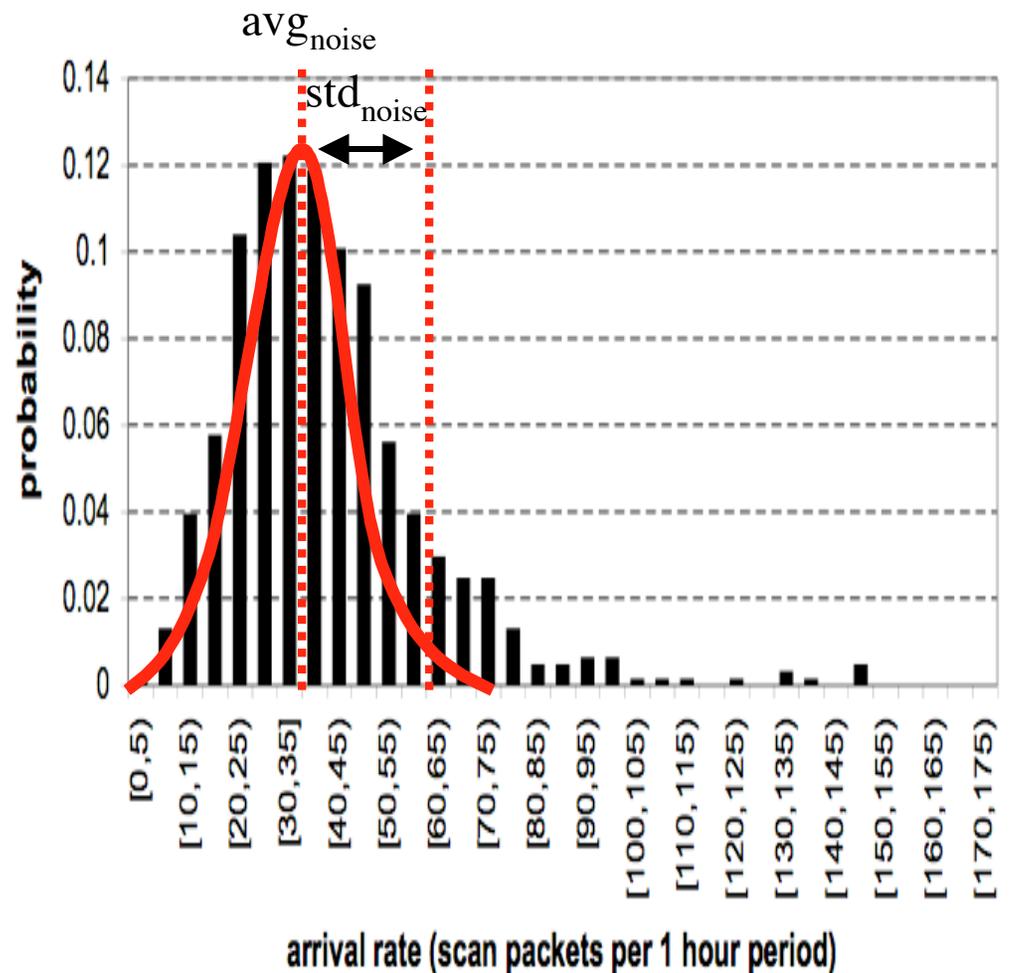


=



Noise behavior \Rightarrow model

- Measured noise behavior:
 - A distribution with an average and standard deviation
- Derived noise model to fit measurements
 - Insight: noise probes = Bernoulli trials
 - Binomial distribution of noise probes at any monitor
 - Analysis verified binomial is a “good fit”



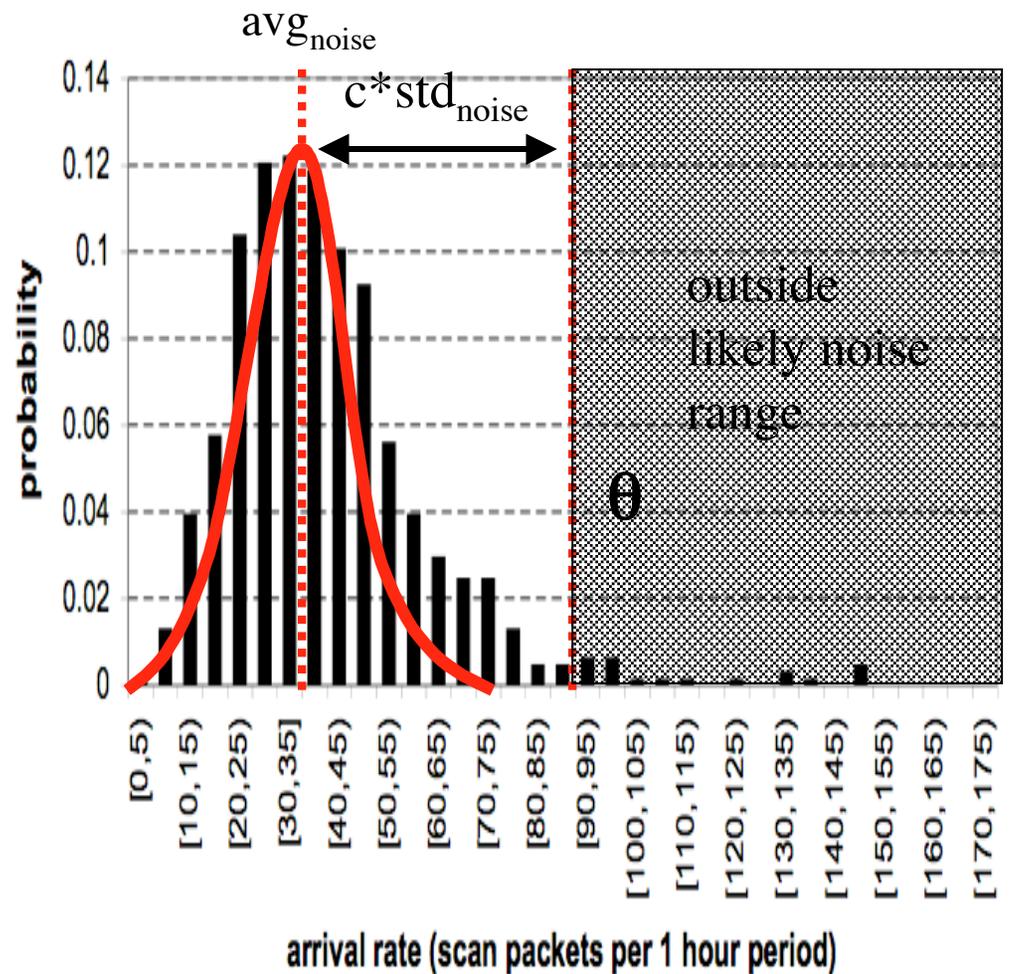
Deriving θ from noise model

- Alarm threshold $>$ likely noise variations
 - “likely” range is below c standard deviations above the average

- Thus, alarm threshold:

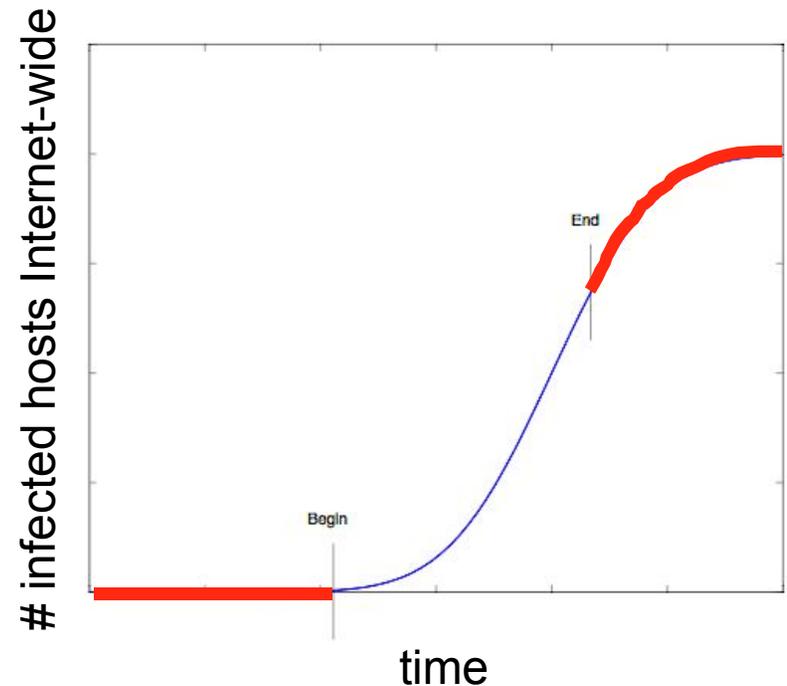
$$\theta = \text{average}_{\text{noise}} + c * \text{std}_{\text{noise}}$$

- What value of c do we choose?



Worm growth model

- We use a pre-existing model that's known to model random-scanning worms accurately (e.g. CodeRedII)
- Random constant spread (RCS) model
 - Exponential growth phase in victims and probe rate until population saturates
 - Signifies “abnormal” probe trend in Internet



Tail details don't matter --
useful detection must occur
during exponential growth
phase



Combined worm + noise model

- Combined model lets us quantify DWDS *fidelity*
 - = probability alarm raised *by* time t in a worm's growth
 - calculated from combined distribution and θ
 - can vary model parameters (noise level, # monitors) to quantify their effect on fidelity



Evaluation

- Parameterize DWDS fidelity model with CodeRedII worm data
 - Vary number of monitors and noise to see how fidelity changes

- Ask and answer 3 questions:
 - 1) What is the tradeoff between detection fidelity and false alarm rate?
 - 2) How is detection fidelity affected by the amount of noise in the Internet?
 - 3) How is detection fidelity affected by the number of monitors in the system?

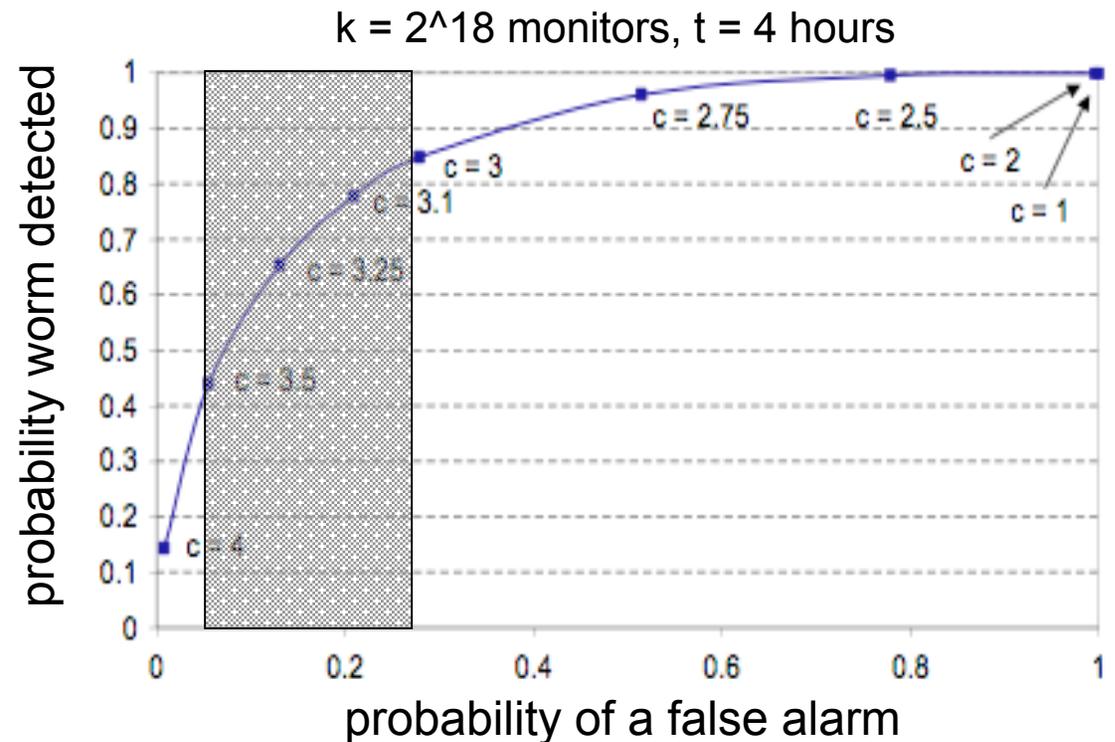
Q1: Detection fidelity vs. false alarm rate?

■ False alarm

- DWDS alarm fires even though there is no worm
- Caused by abnormally high noise fluctuation

■ “c” trades fidelity for low false alarm rate

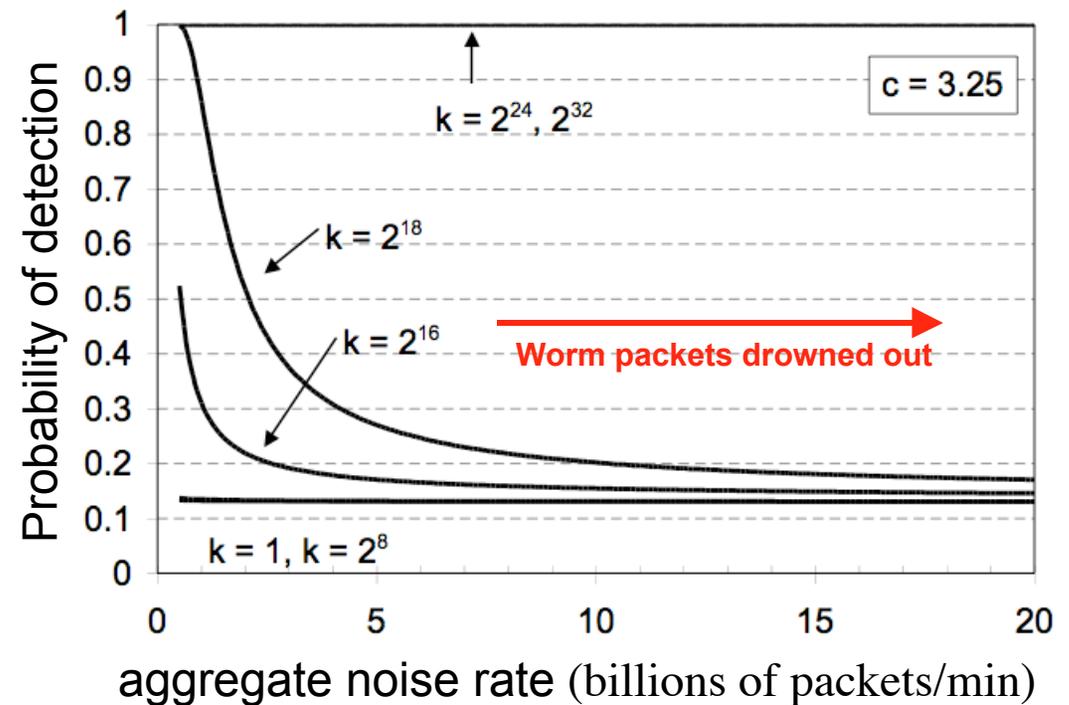
- $\theta = \text{average}_{\text{noise}} + c * \sigma_{\text{noise}}$



Conclusion: $3 < c < 3.5$ lowers false alarms while keeping detection probability high.

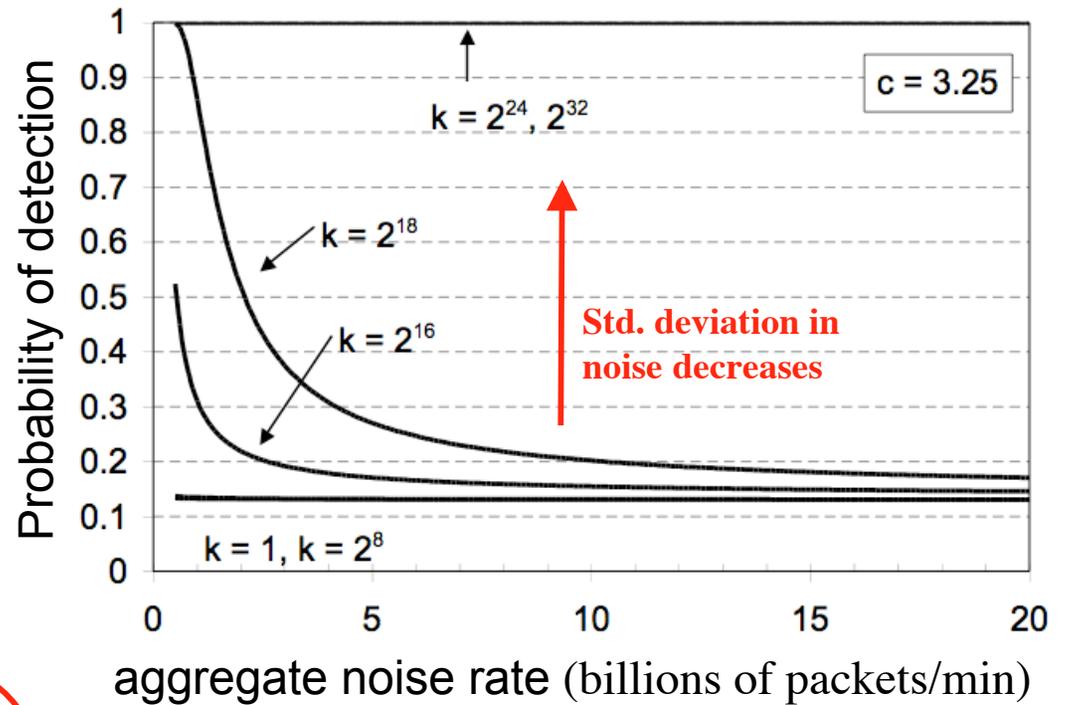
Q2: Detection fidelity vs. average noise rate

- More noise means less fidelity:
 - Value of $\text{std}_{\text{noise}}$ increases
 - Larger noise fluctuations must be tolerated
 - Worm packets are drowned out



Q3: Detection fidelity vs. number of monitors?

- More monitors improves fidelity!
 - Adding monitors decreases variability in observed noise



$$std_{noise} = \sqrt{average_{noise} * \left(1 - \frac{k}{2^{32}}\right)}$$



Summary

- Analytical models for worm growth, noise, fidelity for scan packet-based DWDSs
 - Key insights:
 - Worm and noise probes are Bernoulli trials
 - Model detection fidelity in terms of probability distributions

- Numerical calculations to show:
 - Tradeoff between fidelity and low false alarm rate
 - More noise means less fidelity
 - More monitors increases fidelity



Outline

- Introduction
- Scan Packet-Based DWDS analysis
- **Victims-Based DWDS analysis**
 - Differences compared to scan packet-based analysis (details in paper)
 - Results
- Conclusions



Victims-based DWDS

■ Victim:

- a host that sends probe traffic
- noise or worm

■ Victims-based DWDS similar to scan packet based, but:

- measure rate at which new victims are observed
 - not rate at which probes are sent
- much less susceptible to bursty victims
 - e.g., portscan



Victims-based analysis

- Noise and worm models more complicated
 - Lose independence of Bernoulli trials
 - # new victims depends on what DWDS saw in the past
 - Noise model in terms of birth rate of new noisy victims in Internet
- Calculate detection probabilities via simulation
 - Simulate CodeRedII and parameterized noise model to calculate detection probs
- Details in paper



Evaluation

- Vary number of monitors and noise to see how fidelity changes
 - ◆ Difference: noise metric is now in terms of noise victim birth rate in Internet!
- Ask and answer same 3 questions:
 - 1) What is the tradeoff between detection fidelity and false alarm rate?
 - 2) How is detection fidelity affected by the amount of noise in the Internet?
 - 3) How is detection fidelity affected by the number of monitors in the system?

Results summary

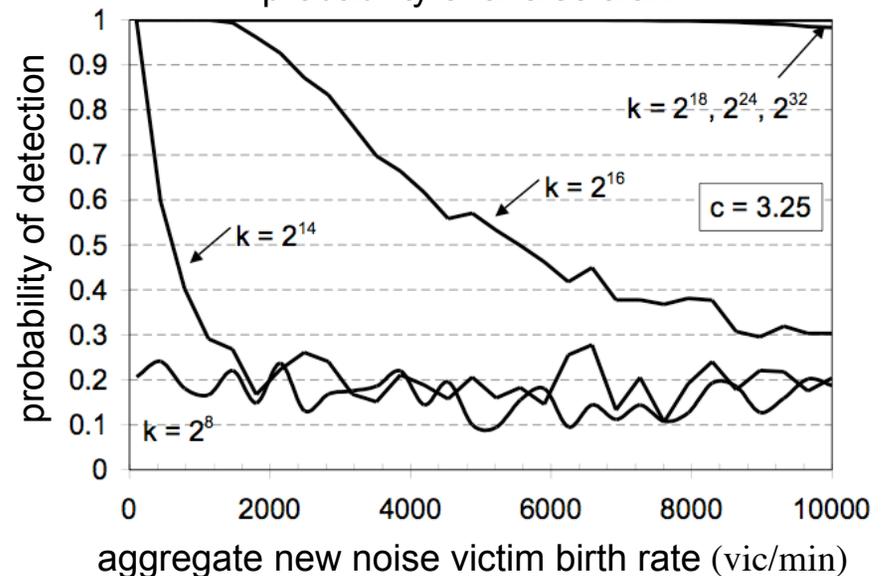
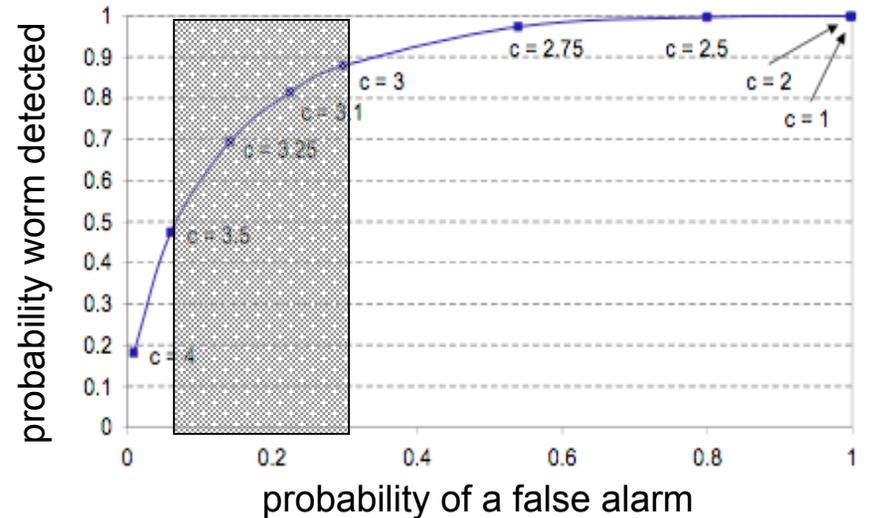
- Detection fidelity vs. false alarm rate?

- “c” again trades fidelity for low false alarm rate

$$\theta = \text{average}_{\text{noise}} + c * \sigma_{\text{noise}}$$

- Detection fidelity vs. average noise rate and number of monitors?

- More noise = less fidelity
- More monitors = better fidelity





Outline

- Introduction
- Scan Packet-Based DWDS analysis
- Victims-Based DWDS analysis
 - Differences compared to scan packet-based analysis (details in paper)
 - Results
- **Conclusions**



Conclusions

- For both victims-based and scan packet-based DWDS's:
 - Tug-of-war between minimizing false alarm rate while maximizing detection fidelity
 - Noise level greatly reduces detection fidelity
 - More monitors increases detection fidelity

- Victims-based DWDS is better?
 - Slightly less affected by noise
 - Less susceptible to bursty victims (portscans)



Bottom line

- For both types of DWDS, need many monitors to be effective
- As noise level in Internet continue to increase, these systems will become less and less effective
- Viable long-term strategies *must* include better signature-detection schemes