

Analyzing Cooperative Containment Of Fast Scanning Worms



Jayanthkumar Kannan

UC Berkeley

Joint work with


Lakshminarayanan Subramanian, Ion Stoica, Randy Katz



Motivation


- Automatic containment of worms required
 - Faster: Slammer infected over 95% of vulnerable population in 10 mins (MPSSSW 03)
 - Easier to write: Worm = "Propagation" toolkit + new exploit

Worm containment strategies



- End-host instrumentation (eg: NS 05)
- Core-router augmentation (eg: WWSGB 04)
- Specialized end-points (eg: honeyfarms - P 04)
- Firewall-level containment (eg: WSP 04)

Decentralized Cooperation



- Internet firewalls exchange information with each other to contain the worm
 - Suggested recently: WSP 04, NRL 03, AGIKL 03
- Pros of decentralization:
 - Scales with the system size
 - No single point of failure / administrative control



Questions we seek to answer


- Cost of decentralization
 - Modes of information exchange
 - Effect of finite communication rate between firewalls on containment
- Effect of malice
 - Trust Model: Only “few” malicious participants
 - How does one deal with malicious firewalls?
- Performance under partial deployment



Roadmap

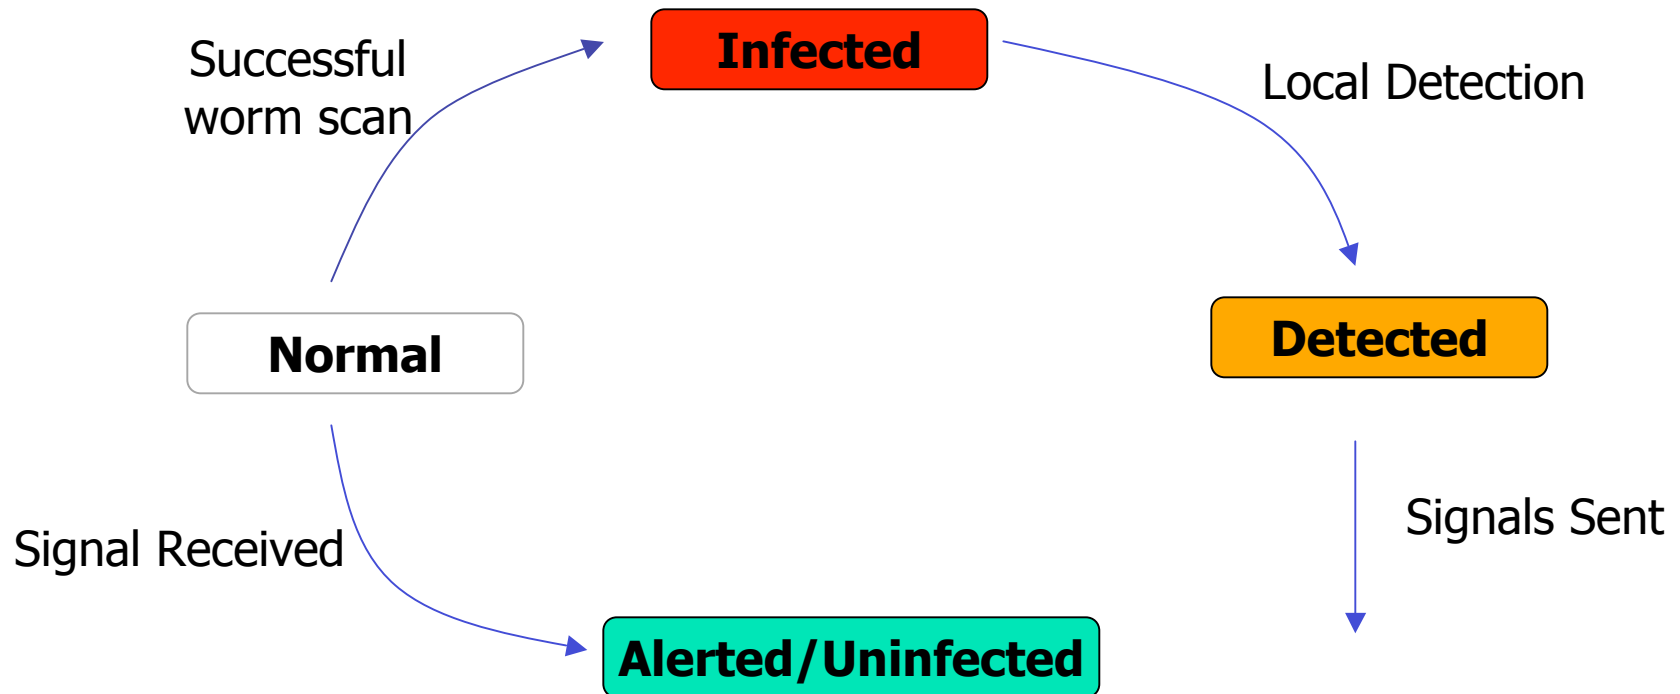
- **Abstract model of cooperation**
- Analysis of cooperation model
- Numerical Results
 - Analytical, Simulation
- Conclusion

Model Of Cooperation



- **Local Detection:** Identify when its network is infected by analyzing outgoing traffic
- **Signaling:** Informs other firewalls of its own infection along with filters
- **Filtering:** An informed firewall drops incoming packets

Firewall states






Model of Signaling

- Two kinds of signaling:
 - Implicit: Piggyback signals on outgoing packets
 - Explicit: Signals addressed to other firewalls
- How to do robust signaling in face of malicious firewalls?

Robust Signaling



■ Security Parameter: T

- Attack: Firewalls suppress signaling keys infected
- Defense: Challenge response verification state signals
- Even if about 25% firewalls behave this way, good containment is possible



Roadmap

- Abstract model of cooperation
- **Analysis of cooperation model**
- Numerical Results
 - Analytical, Simulation
- Conclusion



Analytical results

- Main focus: Containment metric C :
 - C = fraction of networks that escape infection
- Cost of Decentralization
 - Effect of type of signaling:
 - Dependence of containment on signaling rate
 - Is Signaling Necessary?
- Effect of malice:
 - Dependence of containment on Threshold T



Parameters used in analysis

- Worm model:
 - Scanning: Topological scanning (zero time) followed by global uniform scanning
 - Scanning rate = \mathbf{s}
 - Probability of successful probe = \mathbf{p}
 - Vulnerable hosts uniformly distributed behind these firewalls, initial number of seeds small
- Local detection model:
 - After infection, the time required for the infection to be detected is an exponential variable with mean \mathbf{t}_d
- Signaling model:
 - Explicit signals sent at rate \mathbf{E}



No Signaling

- Worm probes only in interval between “infection” and “detection”
- λ is the expected number of successful infections made by a infected network before detection
 - $\lambda = p s t_d$
- Result: If $\lambda < 1$, $C = 1$ for large N (WSP 04)
 - Analogy to birth-death process
- Implications
 - Earlier worms like Blaster satisfied this constraint



No Signaling (2)

- Surprisingly, even if $\beta > 1$, containment possible without signaling for **random** scanning worm
- Intuition:
 - As the infection proceeds, harder to find new victims
 - β ($= p s t_d$) effectively decreases over time
- For $\beta = 1.5$, about 40% containment
- For $\beta = 2.0$, about 20% containment
 - $\beta = O(2)$ for a Slammer-like worm



Need for Signaling

- Signaling required if $\beta > 1$
- Differential equation model
- For $\beta > 1$ and $\beta = (\beta - 1)/t_d$, the containment metric C is lower-bounded by

$$1 - \frac{(\log(N) + (T-1)\log(\log(N)))t_d\sigma^2}{(s+E)} \left(\frac{1}{t_d\sigma} + 1 \right)$$



Need for Signaling (2)

- **Implicit Signaling:**
 - Spread rate of worm (ps) outpaced by signaling rate (s)
 - Implicit signaling relies on ($p \ll 1$)
 - Linear drop with time to detection (t_d)
 - Linear drop with threshold (T)
- **Explicit Signaling:**
 - Explicit signals essential for high p
 - Linear drop with $1/E$
 - Tunable parameter



Summary

- $\beta < 1$: no signaling required for good containment
- $\beta \geq 1$: without signaling, only moderate containment
- $\beta \geq 1$, low p : implicit signaling works
- $\beta \geq 1$, high p : explicit signaling required



Roadmap

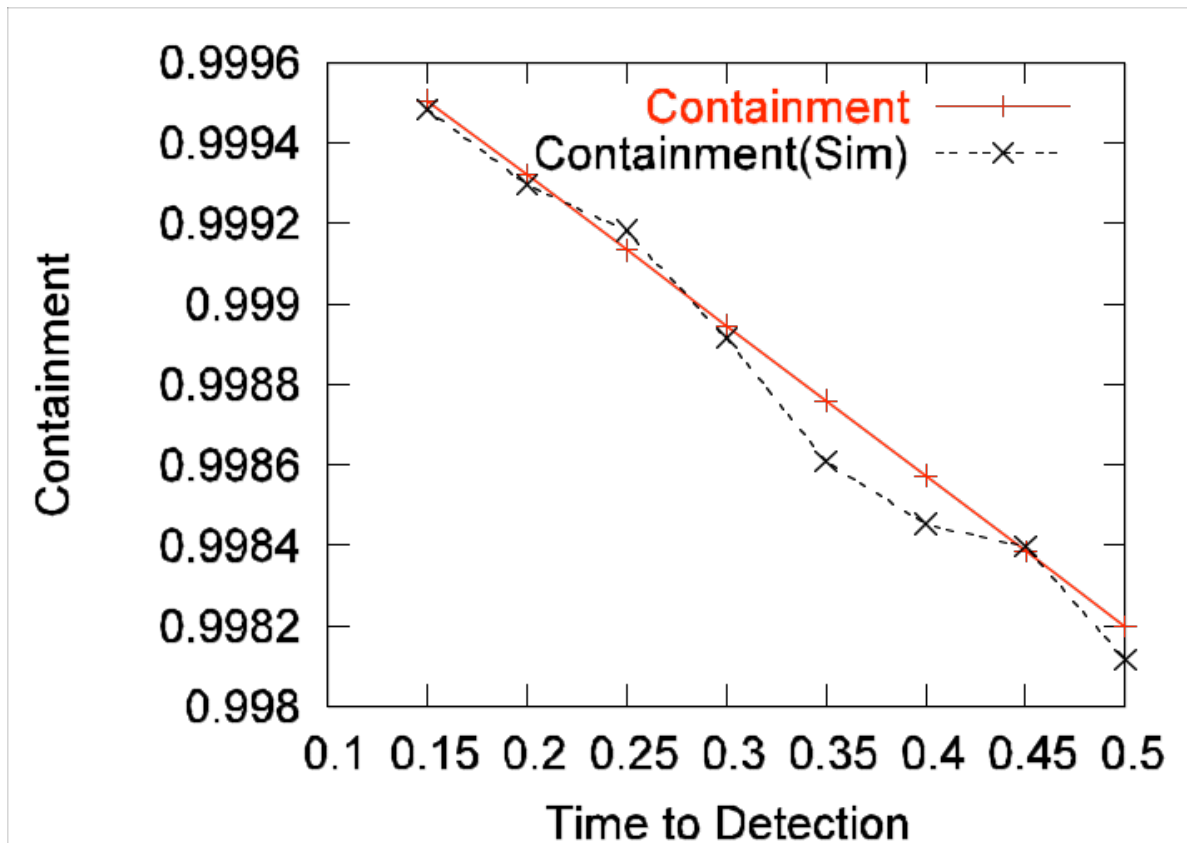
- Abstract model of cooperation
- Analysis of cooperation model
- Numerical Results
 - Analytical, Simulation
- Conclusion



Numerical Results

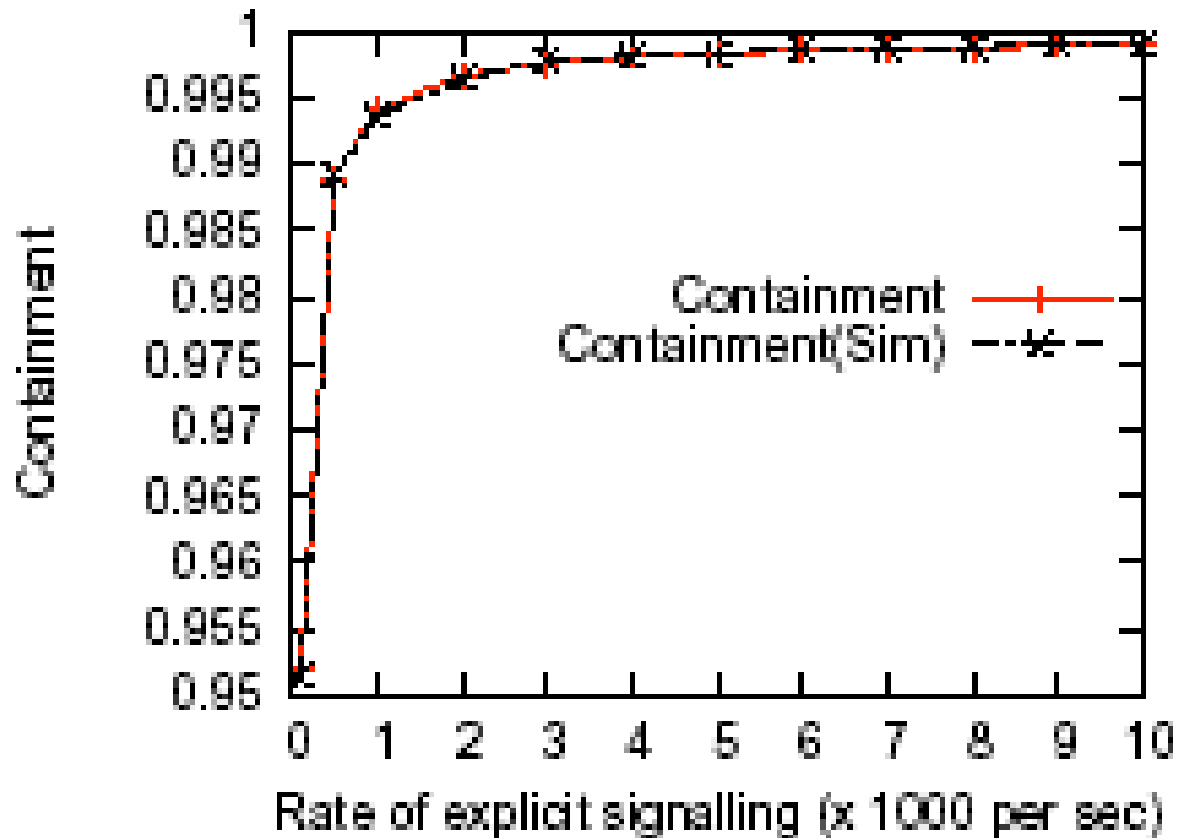
- Parameter Settings:
 - Scan rate set to that of Slammer
 - Size of vulnerable population = 2 x Blaster
 - 100,000 networks: 20 vulnerable hosts per network
 - Start out with 10 infected networks and track worm propagation
 - Time to infect is about 2 secs

Cost of Decentralization



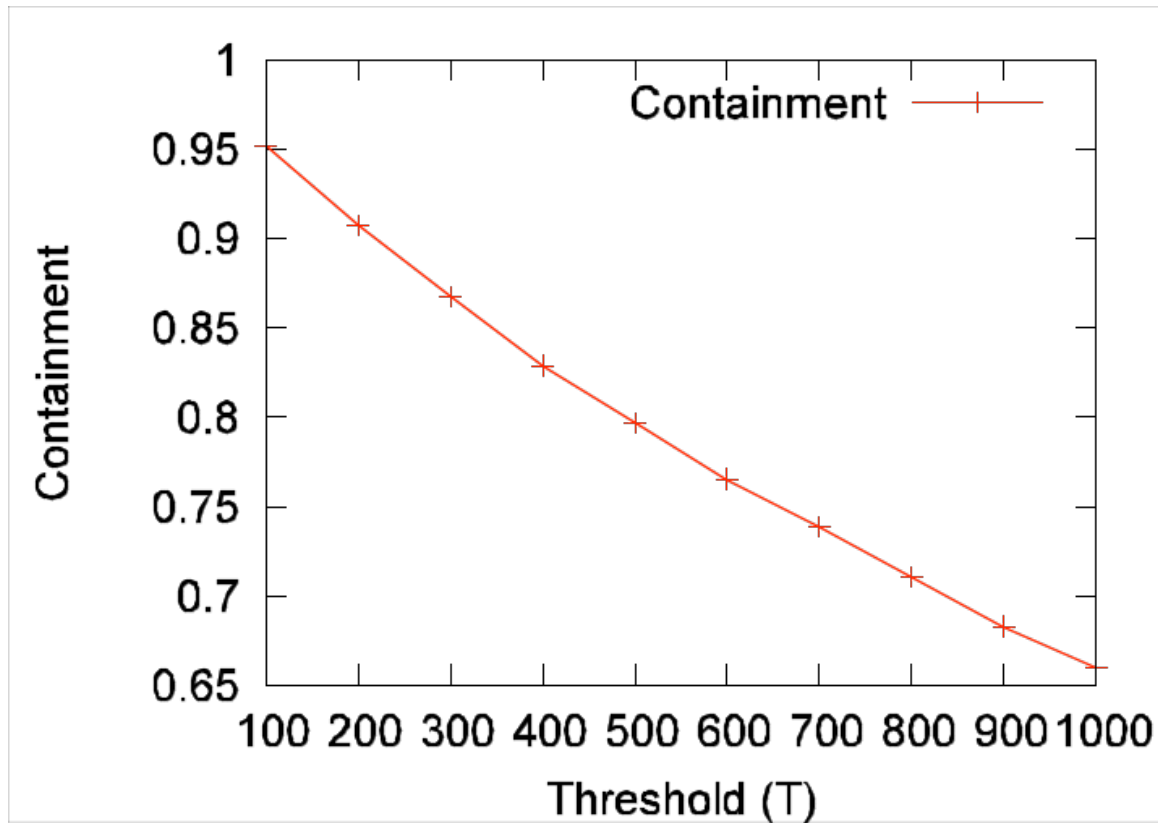
Higher the detection time, lower the containment

Cost Of Decentralization (2)



Even for low explicit signaling rate, good containment

Effect of Malice



Defends against a few hundred malicious firewalls



Conclusion

- Contribution: Characterize necessity, efficacy, and limitations of cooperative worm containment
- Cost of Decentralization:
 - With moderate overhead, good containment can be achieved
- Effect of Malice:
 - Can handle a few hundred malicious firewalls in the cooperative
- Cost of Deployment:
 - Even with deployment levels as low as 10%, good containment can be achieved



Detection and Filtering

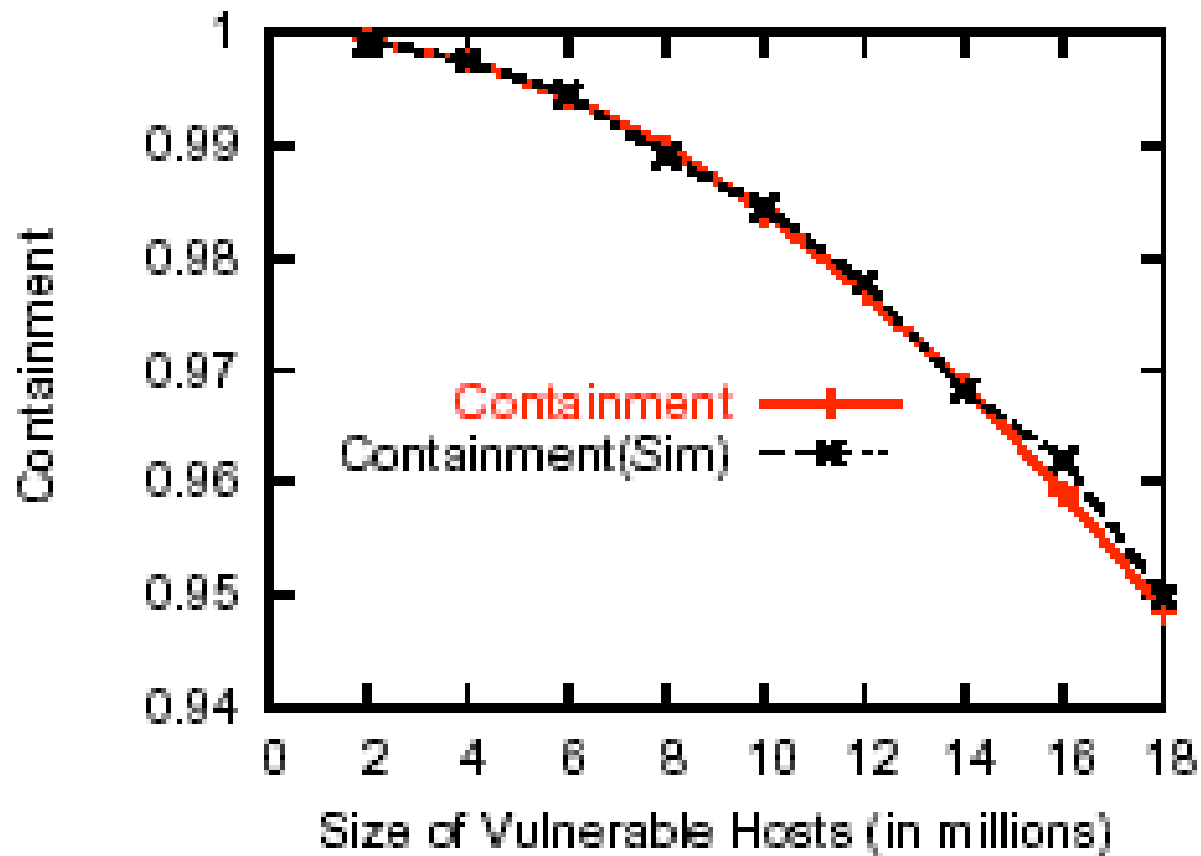
Lemma 2: If $(\lambda > 1)$, assuming $I_0 \ll N$, $C \geq 1 - \min_{k:k>1} \left(\frac{(k\lambda-1)(k+1)}{k\lambda(k-1)} - \frac{2*\log(k\lambda)}{(k-1)\lambda} \right)$ against a random scanning worm (k is a variational parameter used in minimization).



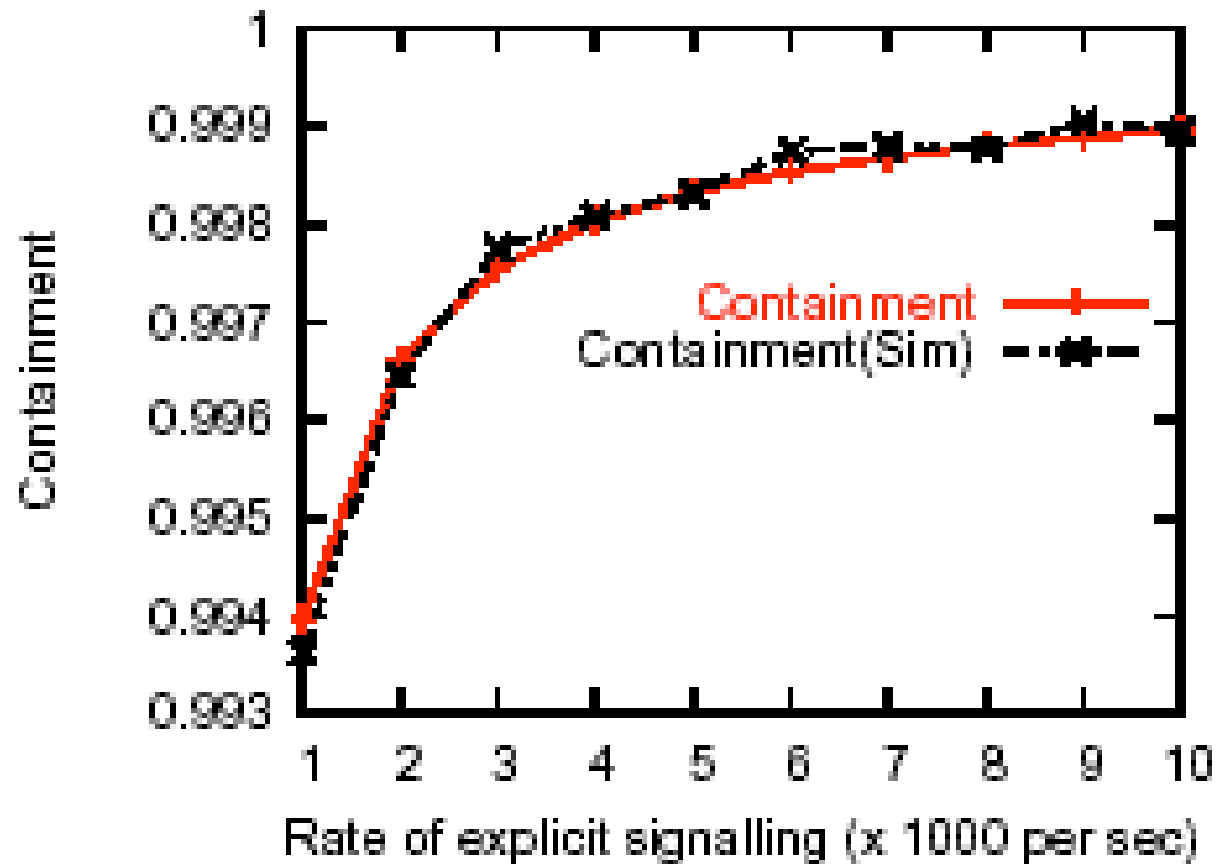
Signaling

Lemma 4: For $\lambda > 1$ and $I_0 \ll N$, the containment metric C obtained by implicit signaling is at least $1 - \frac{(\log(N) + (T-1)\log(\log(N)))t_d\sigma^2}{(s+E)} \left(\frac{1}{t_d\sigma} + 1 \right)$ where $\sigma = \frac{\lambda-1}{t_d}$.

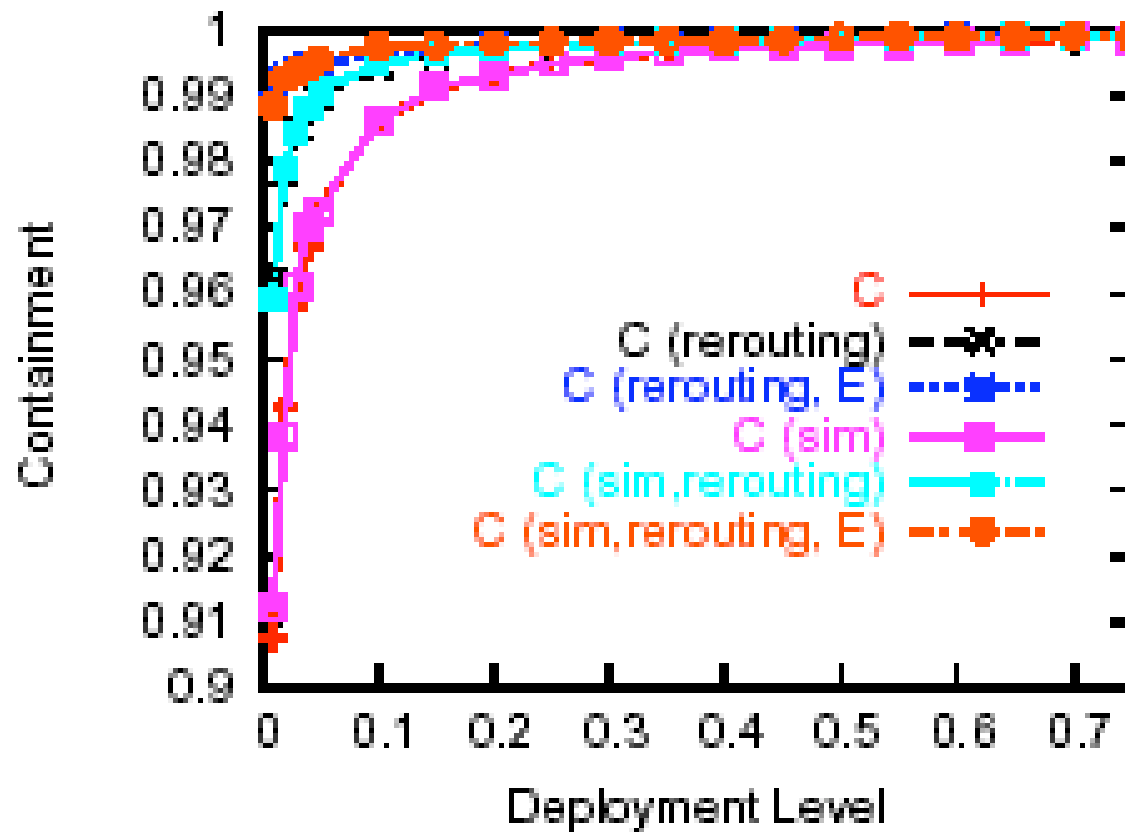
Containment vs Vulnerable population size



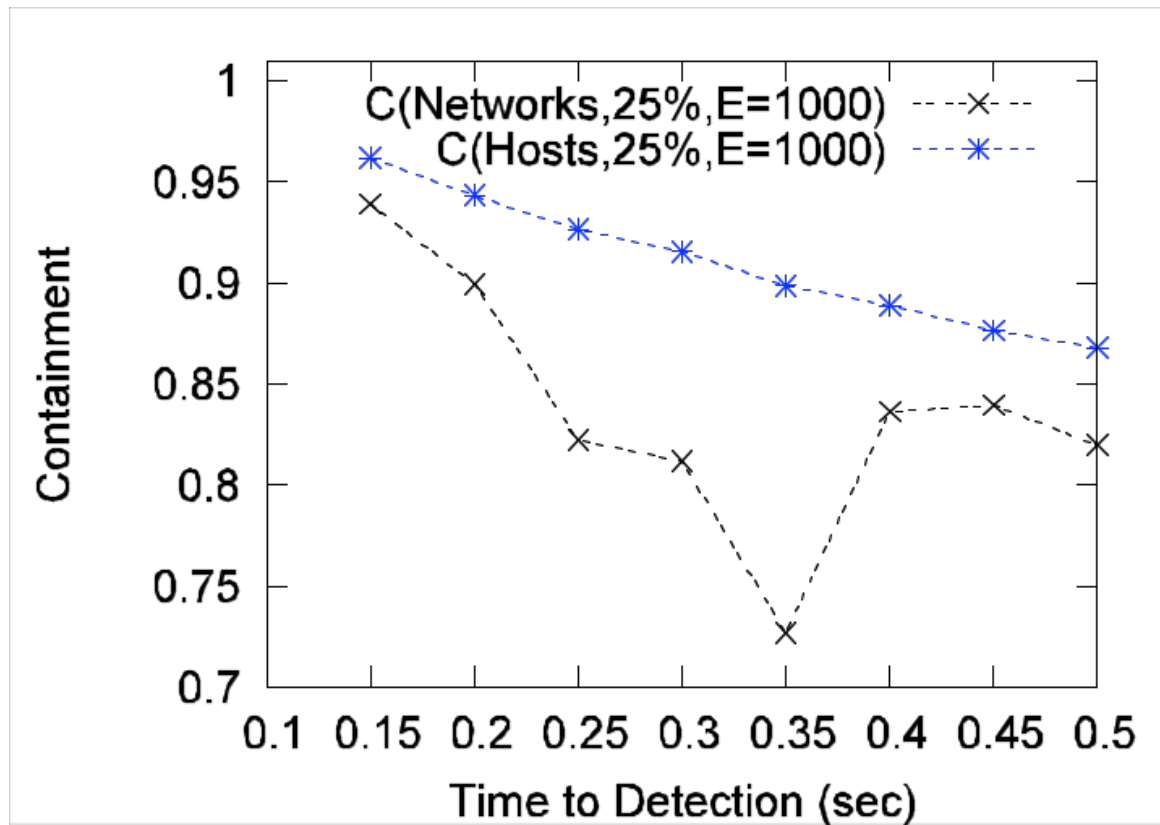
Containment vs Signaling Rate



Containment vs Deployment



Internet-like Scenario



Works well even under non-uniform distributions



Conclusions

- Main result: with moderate overhead, cooperation can provide good containment even under partial deployment
 - For earlier worms, cooperation may have been unnecessary
 - Required for the fast scanning worms of today
- Our results can be used to benchmark local detection schemes in their suitability for cooperation
- Our model and results can be applied to:
 - Internet-level / enterprise-level cooperation
 - More sophisticated worms like hit-list worms
- Room for improvement in terms of robustness
 - Verifiable signals
- Hybrid architecture:
 - Fit in “well-informed” participants in the cooperative