PSU worm modeling and emulation project

George Kesidis
CSE and EE Depts
kesidis@engr.psu.edu

School of CS, Carleton University, Ottawa, Canada
Tues. Oct. 18, 2005

Outline

- Scanning worms
- Slammer worm's spread in the Internet
- Homogeneous network model
- Model extensions

- Elements of this work in collaboration with:
  - V. Paxson and N. Weaver, ICSI, Berkeley
  - P. Liu, School of IST, Penn State
  - M. Vojnovic, Microsoft, Cambridge, UK
  - PSU students: I. Hamadeh, S. Jiwasurat, L. Li, Y. Jin

Scanning worm defenses

- End-systems infected with scanning worms automatically search the IPv4 address space using one of several different strategies that have already been observed.
- They automatically scan (attempt session initiation) with potential victim end-systems.
- Defense/containment devices assumed deployed in peripheral enterprise networks
  - End-hosts and/or network nodes, e.g., access router
  - Stand alone or collaborative
- Zero-day defenses detect anomalously
  - large destination IP addresses contacted per unit time
  - large freq of failed scans, scans to dark addresses in particular
  - large number of packets with certain src/dst ports
  - few DNS precursors (may require DPI, i.e., payload info)

Enterprise network defense DUT

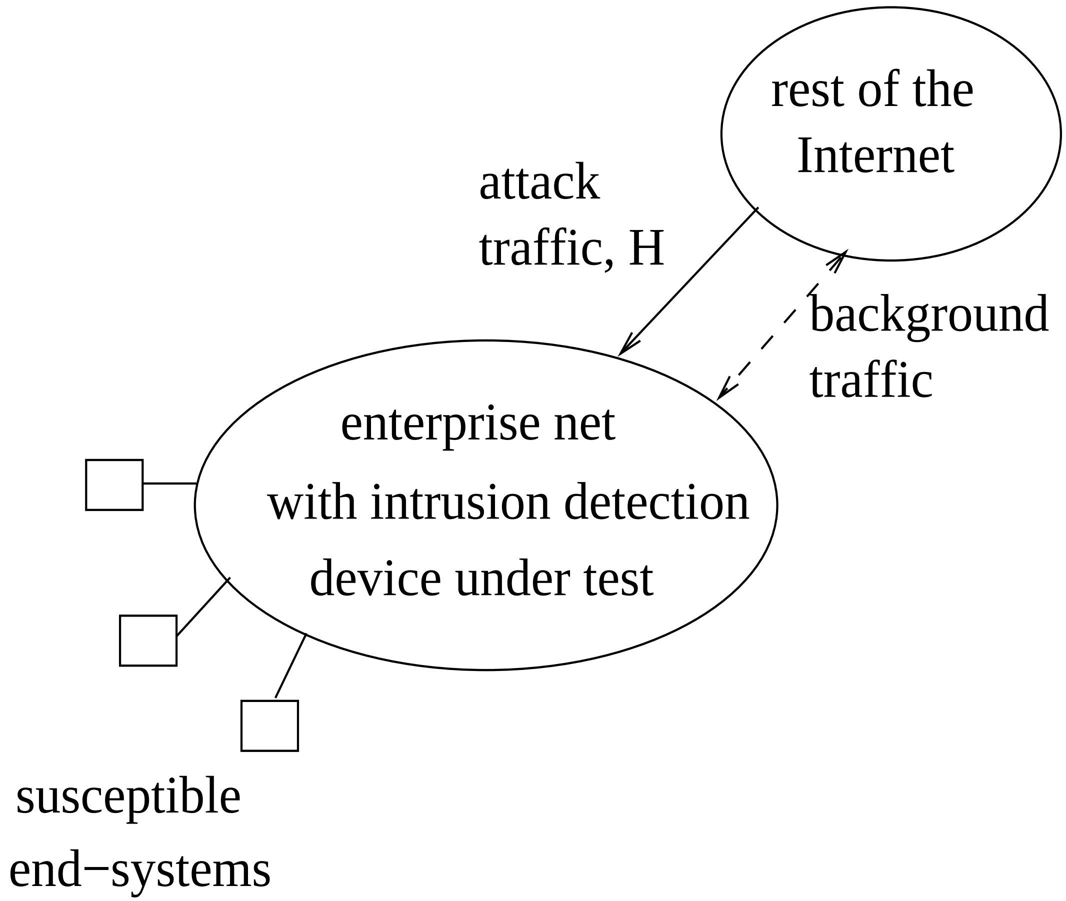

Evaluation of Scanning worm defenses

- Need background traffic for evaluation of false-positives.
- Need attack traffic for evaluation of false-negatives.
- In practice, most defenses are evaluated using
  - worst-case traffic scenarios (→over-engineering), and
  - limited deployments in operational networks (representative?).
- So, need to realistically model the worm probing (scanning) activity *from* the Internet *to* the enterprise network under test.

Scanning worm attack recreation

- We assume that the scans generated from a given enterprise to the rest of the (much larger) Internet and the scanning activity directed at the enterprise from without are negligibly dependent.
- The scan-rate directed at the enterprise under simulation could be approximated as $H(t) = S(t) \cdot A/2^{32}$, where
  - $S(t)$ is the total (Internet-wide) instantaneous scan-rate of the worm at time t, and
  - A is the size of its address space
- Alternatively, a random thinning of *S* could be used to determine H.

Enterprise network defense DUT

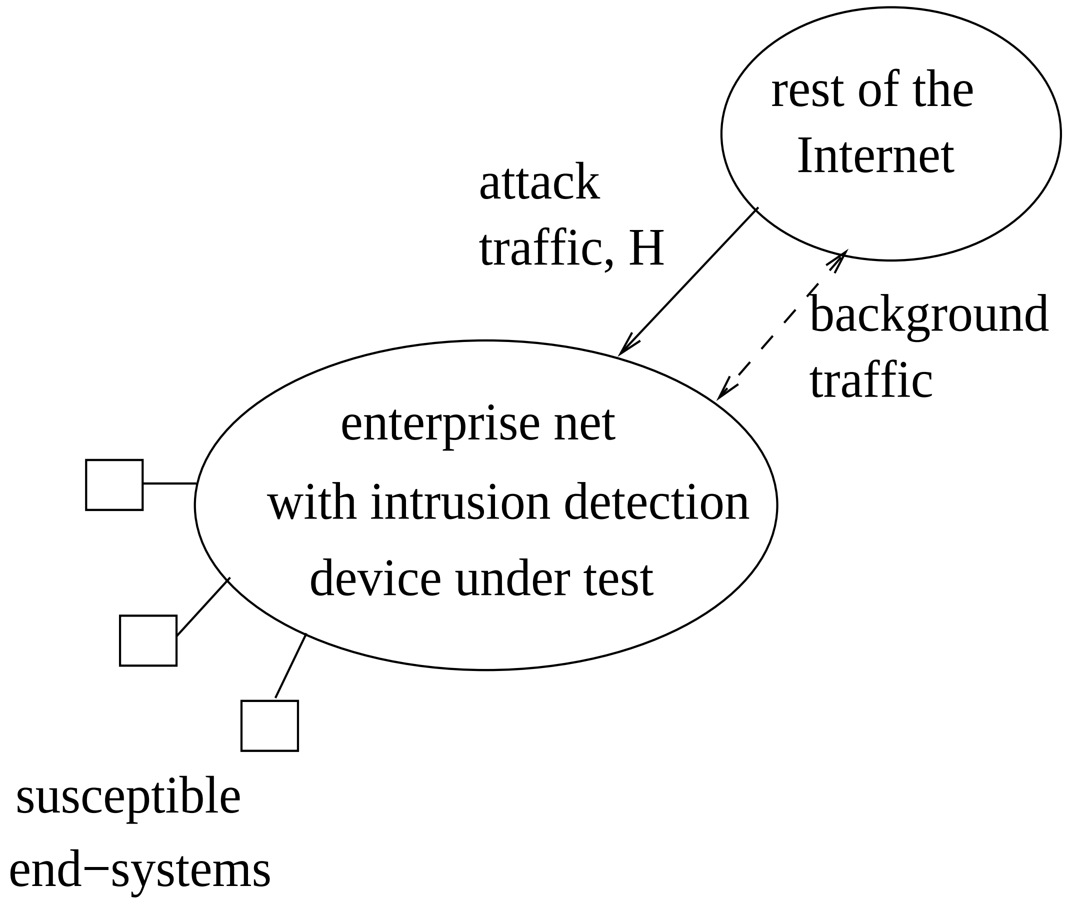

Scanning worm attack recreation (cont)

- The total scan-traffic generation S can be estimated from extrapolations of measured data for a particular worm *when this is available*, e.g., from the University of Wisconsin's, Michigan's or CAIDA's (UCSD's) tarpit.
- Alternatively, one could use a mathematical model whose parameters can be
  - fit to the salient data of a given worm (again, if that data is available) or
  - varied in an attempt to capture the behavior of actual worms for which measured Internet data is unavailable or set for hypothetical worms.
- A mathematical model also:
  - Has insight and computational advantages over the potentially more accurate approach based on scale-down techniques and parallel simulation.
  - Allows for convenient study of hypothetical worms that are necessary to consider when evaluating defenses to be deployed.
  - Does not have privacy issues associated with dissemination of tarpit data.

Bandwidth Limited Scanners

- Propagation of Blaster, Slammer and Witty worms:
  - congested network links thereby creating a temporary denial-of-access to the Internet for large population of end-hosts.
  - resulted in a significant direct expenditure for patching and very significant aggregate loss of productivity.
- We focus on bandwidth-limited, random UDP-scanning worms like Slammer and Witty that spread extremely rapidly in the wild.
- Slammer infected about 75 thousand SQL servers (nearly the entire population of susceptibles) in less than 10 minutes and caused significant congestion in the stub-links connecting peripheral enterprise networks to the Internet core.

Slammer's spread

- The success of the simple Kermack-McKendrick (SIR) model for the Code Red worm has been demonstrated.
- Modeling Slammer and Witty is substantially more complex because network bandwidth limitations mitigated the spread of the worm.
- Beyond just spreading very quickly, Slammer was the first significant worm without a constant scanning rate:

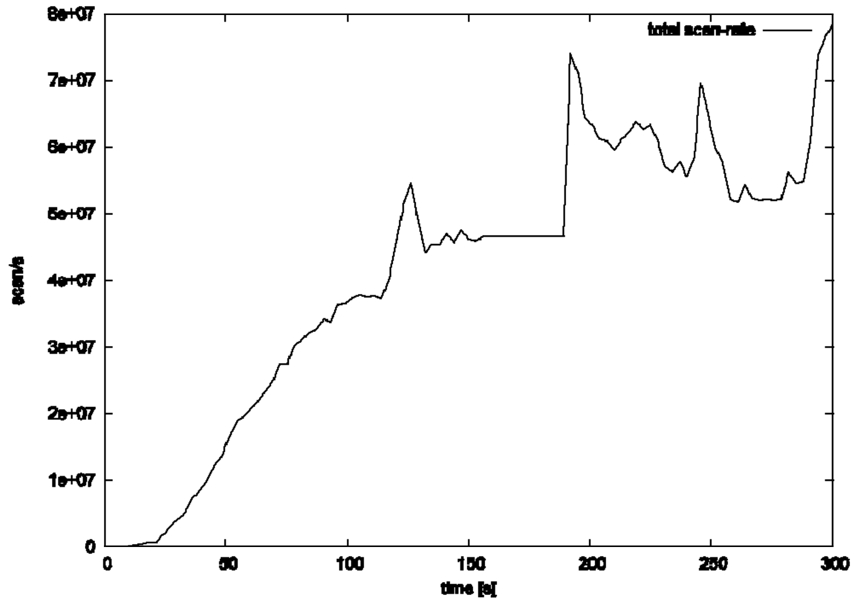

Slammer's scan-rate per worm (infective)

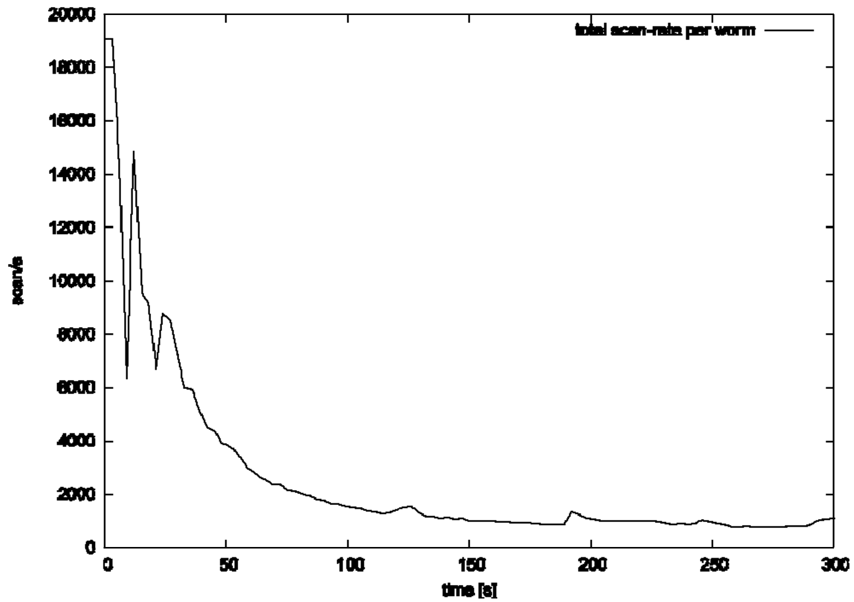

- Note that the oscillations in these curves are largely due to measurement error that is magnified by extrapolation.

Simple SIR epidemiological model

- Consider a population of N hosts which are all susceptible to an infectious disease.
- Suppose that hosts are either in an infected or uninfected state.
- At time t, let $y(t)$ be the number of infected hosts so that $N-y(t)$ are uninfected.
- Suppose each infected host generates “contagion”, each targeted at a specific host, at a constant rate σ.
- Suppose that each contagion will select a single host (infected or otherwise) at random; the probability that a given host is selected is η.

Simple SIR epidemiological model (cont)

- Therefore, over the interval of time $(t,t+dt)$, the expected change in the number of infected hosts, $dy(t)$ will be equal to
  - the total amount of contagion generated in the interval, $\sigma y(t)dt$,
  - times the expected amount of infection caused by a given scan, $\eta(N-y(t))$ (a small fractional quantity),
- I.e., $dy/dt = \beta\, y(t)(N-y(t))$ where $\beta=\sigma\eta$ so that
  $y(t) = Ny(0)/[y(0)+(N-y(0))\exp(\beta Nt)]$ for $t\geq 0$.
- Note that we are not considering countermeasures (inoculations or cures) or deaths: Slammer was a very rapidly spreading worm that was otherwise benign to its host (unlike Witty).

Homogeneous model with instantaneously saturating links

- Assume the Internet core connecting peripheral enterprise networks only negligibly affects any scanning traffic they generate.
- Consider now a population of N enterprise networks.
- For a homogeneous Internet model, assume each enterprise has the same number C of susceptible (SQL server) nodes.
- Each enterprise is in one of C+1 states where state i connotes exactly i worms (infectives) for $0 \leq i \leq C$.
- For the entire network, define the state variables $y_i(t)$ representing the number of enterprises in state i at time t.
- Clearly, for all time $t \geq 0$, $\Sigma^{C}_{i=0}\, y_i(t) = N$.

Homogeneous model (cont)

- Define $Y(t) = \Sigma^{C}_{i=1} y_i(t) = N - y_0(t)$ as the number of enterprises with one or more worms (infectives).
- Assume that each such infected enterprise transmits exactly σ scans/s into the Internet irrespective of the "degree" of its infection, i.e., we assume that a single infective saturates the stub-link bandwidth of the enterprise.
- Finally, an implicit assumption of the following is that "local" infections (between nodes in the same enterprise) are negligible in number.
- Thus, the total rate of scanning (causing infection) into the Internet at time t is $S(t) = \sigma Y(t)$.

Homogeneous model (cont)

- The likelihood that a particular susceptible is infected by a scan is $\eta=2^{32}$ (purely random scanning in the 32-bit IPv4 address space).
- Thus, the likelihood that a scan causes an enterprise in state i at time t to transition to state i+1 is (C-i)η because there are C-i susceptible but not infected nodes in the enterprise at time t.
- Thus, define $\beta_i = \sigma\eta(C-i)$.
- The y_i are governed by the following *coupled* Kermack-McKendrick equations: For times $t \geq 0$,
- $dy_C(t)/dt = \beta_{C-1} y_{C-1}(t) Y(t)$,
- $dy_i(t)/dt = (\beta_{i-1} y_{i-1}(t) - \beta_i y_i(t))\, Y(t)$ for $1 \leq i \leq C-1$
- $dy_0(t)/dt = -\beta_0 y_0(t) Y(t)$

DETER figure of our WORM'04 scaledown simulation

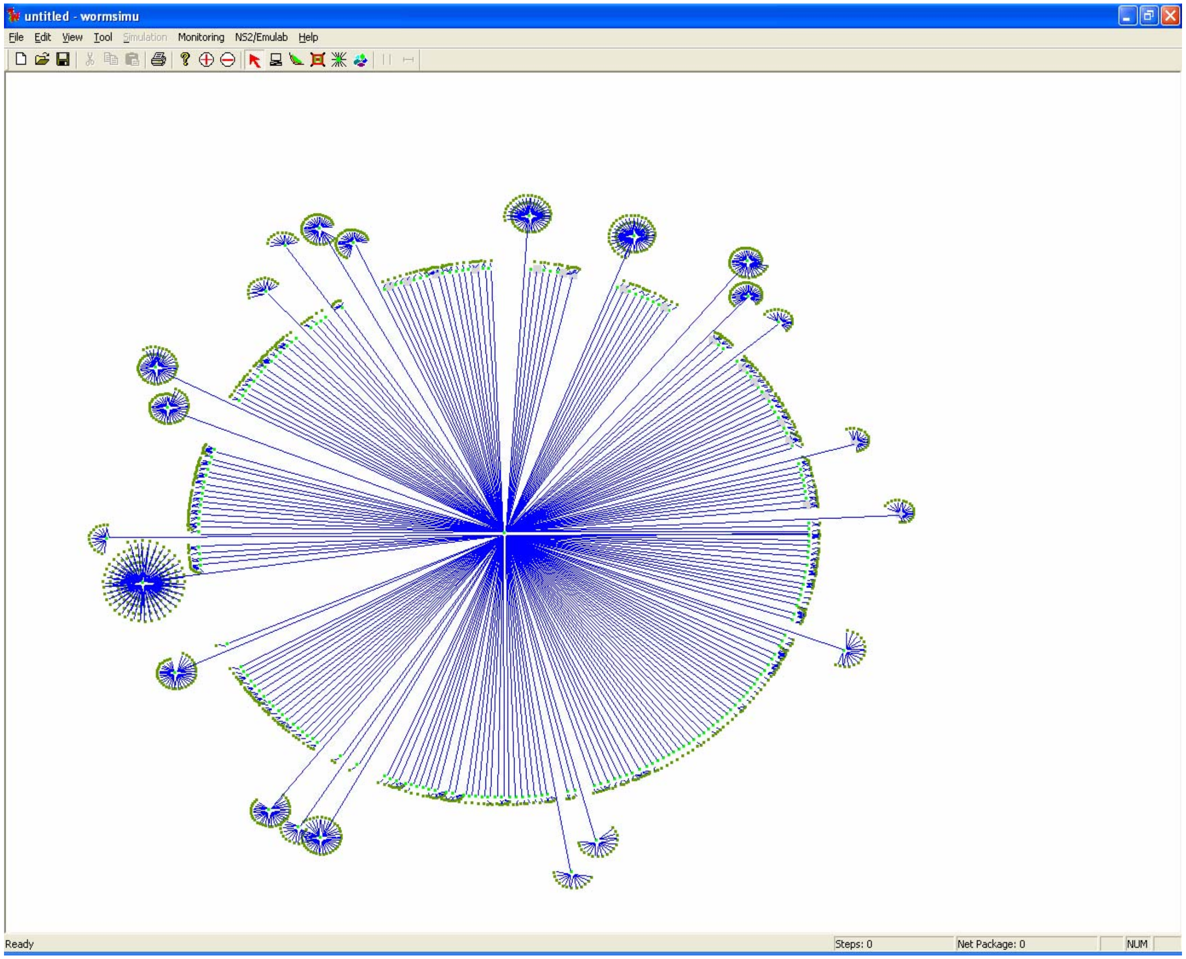

Homogeneous model (cont)

- The total number of worms at time t is clearly $\Sigma^{C}_{i=1}\, iy_i(t)$.
- Thus, the scan-rate per worm (per infective) is $Y(t) / \Sigma^{C}_{i=1}\, iy_i(t) = \Sigma^{C}_{i=1}\, y_i(t) / \Sigma^{C}_{i=1}\, iy_i(t)$.
- Note that summing the coupled equations indexed i=1 to C yields the "standard" Kermack-McKendrick equation:

  $dY/dt = \beta_0 y_0(t)Y(t) = \beta_0(N\text{-}Y(t))Y(t)$

  whose solution is

  $Y(t) = Y(0)/(1+\exp(-\beta_0 Nt)$.
- If we change time to $du(t) = (N - y_0(t))dt$, then a system of linear ODEs results and this leads to a closed form solution.

Fitting to measured data

- We fitted just three parameters to measured data:
  - The initial value of scan-rate per worm: σ = 15000
  - The ratio of initial to final value of scan-rate per worm: C = 18
  - The final value of total instantaneous scan-rate, $NC\sigma$ (or simply from the total number of initially susceptible (and ultimately infected) end-systems, NC=73784) giving N=4099.
- Numerically solving starting from i.c.'s $y_0(0)=N-1$ and $y_1(0)=1$ (i.e., one initially infected server) yielded the following “homog-i” curves.
- This simple deterministic mathematical model of a homogeneous network with instantaneous link saturation yielded numerical results similar to those obtained by simulation of the “homogeneous clusters” model in our WORM'04 paper.

Total instantaneous scan-rate

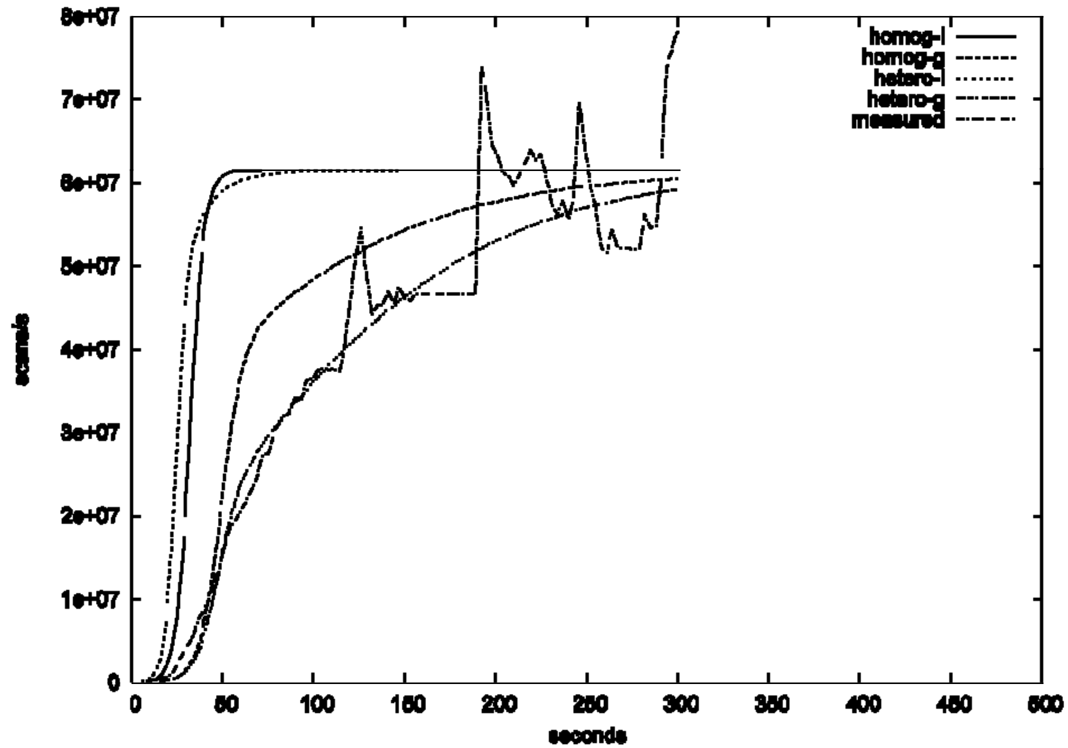

Scan-rate per worm (infective)

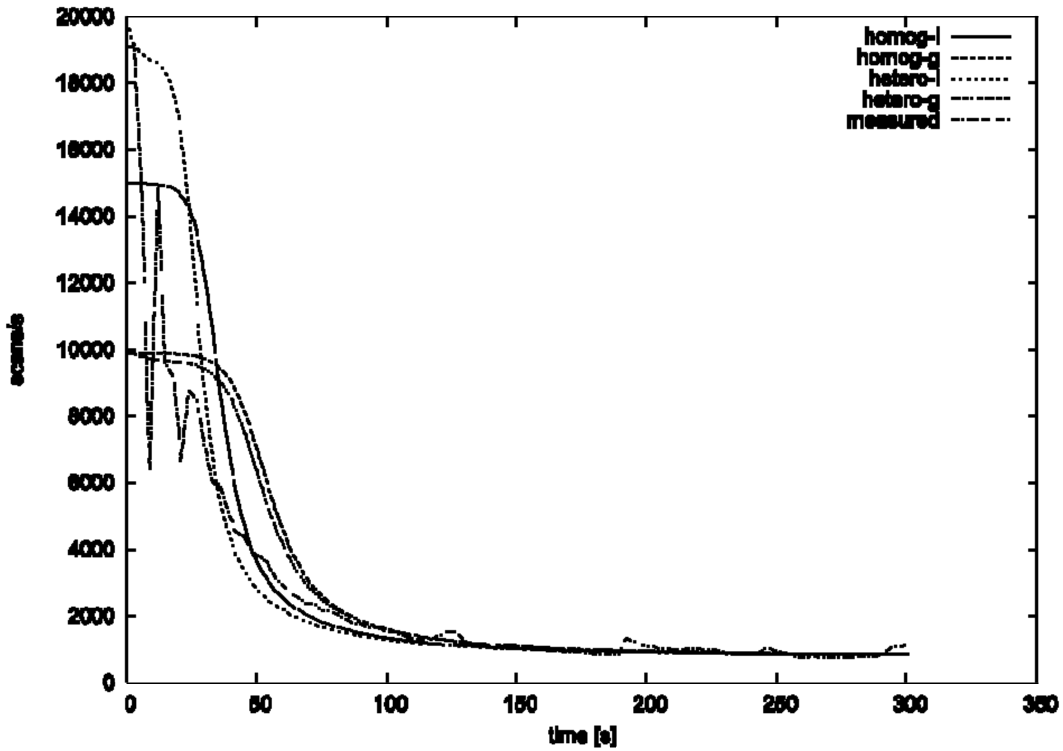

Model Extensions

- Straightforward to extend our model to accommodate:
  - access links that gradually saturate as the number of infectives grow.
  - more heterogeneous enterprise networks with different access link capacities and/or different numbers of susceptible end-systems per enterprise.
  - removals of infectives (patch/crash) or susceptibles (patch).
- See Penn State's KMSim and packet injector tools (open sourced) at http://www.isi.edu/deter

Model Extensions

- In particular, Slammer's **routeview** data can be used to define a number of classes of enterprises with different numbers of susceptibles and all classes having instantly saturating access links with the same bandwidth (as in our WORM’04 paper):
- We have shown that that both the total scan-rate and the scan-rate per infective curves are accurately approximated by this model.
- Furthermore, we have recently shown that this model with countermeasures accurately represents the Witty worm, its non-uniform scanning strategy notwithstanding.

Other projects of G. Kesidis

- NSF cybertrust project on congestion control in non-cooperative networks (with C. Das + Purdue)
- NSF NeTS NoSS: Sensor MANETs (with G. Cao, T. La Porta and C. Das)
- NSF ITR on networking visual sensors (with O. Camps and M. Sznaier)
- NSF ITR on incentive engineering (with R. Acharya and N. Gautam)
- DARPA/ARO Emerging Surveillance Plexsus MURI (ARL)
- Cisco collaborations: attack attribution, reputation systems
- Leadership role in NSF/DHS project that is the sister project of DETER under which our worm research is funded…

Evaluation Methods for Internet Security Technology (EMIST)