CS 294-28: Network Security

Prof. Vern Paxson

http://inst.eecs.berkeley.edu/~cs294-28/
http://www.icir.org/vern/cs294-28/

vern@cs

January 21, 2009

What Is This Class?

- ~New graduate course on network security
  - ~New = it may be bumpy at times
  - Graduate = focus on reading papers, participatory discussion, major project
  - Network security = how do we keep our computer networks functioning as intended & free of abuse
  - Network = heavy emphasis on global Internet
    - Little emphasis on host-side issues

Target Audience

- Course intended to:
  - Provide grounding necessary for pursuing PhD research in network security
  - Provide breadth for those undertaking research in other areas of security or networking
  - Evolve into regular grad offering complementing CS 261
- Not intended to:
  - Summarize Internet security issues / technology / practices

Prerequisites

- EE 122 (undergrad networking) or equivalent
- Basic network security notions
  - Firewalls, public-key crypto, spoofing, buffer overflow attacks
- Basic probability/statistics
- A willingness to thoughtfully read a lot of technical papers & tackle a hefty/meaningful project

Who Am I?

- Recent professor in EECS (2007)
  - Recent = "it will be bumpy at times" :-)
  - Also affiliated with *International Computer Science Institute* and the *Lawrence Berkeley National Lab*
- Contact:
  - vern@cs, http://www.icir.org/vern/
  - Office hours M 1:30-2:30PM in 737 Soda
    - And by appointment, sometimes at ICSI
      - http://www.icsi.berkeley.edu/where.html
  - Phone: 643-4209, 666-2882
    - Email works *much* better!
  - Hearing impaired: please be ready to repeat questions & comments!

Who Am I?, con't

- Research focuses on network security & network measurement
- Been around the block
  - 10+ years on both topics
  - PC chair/co-chair of SIGCOMM, USESEC, IEEE S&P ("Oakland"), HotNets
- CCIED = NSF Cybertrust *Center for Internet Epidemiology & Defenses*
  - Large-scale compromise, i.e., worms & now botnets
  - 5 year effort joint w/ UCSD (through 2009)
- "Bro" *network intrusion detection system* (NIDS) running 24x7 at LBNL (since 1996!)

My Perspectives/Biases

- I am an empiricist
  - It can be amazing how different a very large system behaves in practice vs. how you would expect it to …
    - … if you only measure in a confined laboratory environment
- A vital, easily overlooked facet of security is *policy* (and accompanying it: operating within *constraints*)
- Much of network security is necessarily reactive, unprincipled, incomplete

Perspectives/Biases, con't

- The goal is risk management, not bulletproof protection.
  - Much of the effort concerns "raising the bar" and *trading off resources*
  - This applies to research as well as practice
- Key notion of threat model: what you are defending against
  - This can differ from what you'd expect
  - Consider the Department of Energy …

General Research Themes

- All papers have shortcomings
  - Doesn't mean you can't extract value
- For your own work:
  - Frame limitations
  - Be thorough & generous towards prior work
  - Provide insight into tradeoffs
- Methodological issues
  - Gauging data quality
  - Bootstrapping (perhaps) ground truth
  - Partition development vs. assessment data

General Research Themes

- Replication/criticism of prior work is unfortunately very rare
  - Corollary: little research upside to publishing data
- Research does not proceed as presented in a well-written paper
- Topics can heat up excessively
  - Multicast, QoS; Traceback, worm models
  - Crucial task for successful research is problem selection

Network Security Research Themes

- Evasion-proof is not a realistic goal
  - Research progresses in often-pretty-modest steps (*building blocks*)
  - "Raising the bar" has definite utility
  - Today's evasion problem looks different tomorrow
  - But: *do* frame evasion picture
- Field changes very fast
  - Including serendipity
  - You need to figure out how to be nimble

Research Themes, con't

- Beware the problem of Crud
  - Surprising diversity of benign activity
  - Great utility in obtaining real data
- We're constantly trading off
  - Especially false positives vs. negatives
- Beware funding ecosystems (and popular press)
  - E.g., DARPA's need for metrics
- Historically, publishing attacks has been worthwhile
  - But not guaranteed

What's Expected of You?

- Read 2 (sometimes 3) papers/week
  - There is an art here regarding figuring out which facets to spend time on and which not
- Write mini-reviews of each paper
  - Mini-review = a few sentences for each of
    - What are the paper's main contributions?
    - What parts of the paper do you find unclear?
    - What parts of the paper are questionable?
      - E.g., methodology, omissions, relevance
    - Given the contributions, what issues remain? What related ideas does it bring to mind?
  - Email me your reviews prior to corresponding lecture (**Tue 9AM** for Weds; **Fri 1PM** for Mon)
    - Late = 50% penalty (no credit if after lecture summary)

What's Expected of You?, con't

- Participate in lecture discussion of the paper & the topic
- "Scribe" a couple of lectures/semester
  - Scribe = write up summary of lecture suitable for posting on course web site
  - Due 1 week after lecture
    - Send me LaTeX, HTML or Word (editable)
  - *Inspect syllabus* and tell me which lecture(s) you'd like to scribe (FCFS)
  - # of lectures to scribe depends on final class size

What's Expected of You?, con't

- Undertake a significant project
  - Individually or in a team of two (encouraged)
    - Discuss w/ me if you want a larger team
- Can involve:
  - Measurement study characterizing/exploring a network security issue
  - Substantive analysis/assessment of security issues for a given network system
  - Development of a new mechanism or technique
  - Deep, thoughtful literature survey of an area
  - Develop & assess a new threat

Project, con't

- Proposals due within a couple of weeks
  - To be commented upon by your peers
- *Related Work* writeup due before Spring Break
- Short status report due a few weeks later
- Class presentations in early May
- Final project due at end of semester
  - Written as a conference-style paper

Project, con't

- Aim high!
  - End result should be workshop-caliber
  - The best should be within shouting distance of publication-caliber
- Find a topic that grabs you
  - Feel free to run preliminary ideas by me

Grading

Homework + Participation	20% + 15%
Scribing	15%
Project	50%

FAQ: Can I audit the course?

A: Instead, please take it P/F. To pass, you need to then do a solid job on either homework/participation/scribing, or a project. Let me know up front which you're pursuing.

Lecture Format

- Each lecture has at its heart a core paper (sometimes 2)
- For the most part, seminal paper that opened new area or developed key new insight
  - Not "bleeding edge" or comprehensive or perfect
- Lecture will cover main contributions …
- … but then go from there into related considerations (sometimes taken from the optional reading) in an interactive fashion
- What to cover & where to go driven in part by thoughts/considerations from HW writeups

Ethics

- We will be discussing attacks - some quite nasty! - and powerful eavesdropping technology
- None of this is in *any way* an invitation to undertake these in any fashion other than with informed consent of all involved parties
- If in some context there's any question in your mind, come talk with me first
- Oh and: for homeworks, please do your own work

A Look At The (Tentative) Topics

- Denial-of-Service
- Traceback
- Network Capabilities
- DoS Defense

Tentative Topics, con't

- Network intrusion detection
  - Systems
  - Evasion
  - Evaluation
- Worms
  - Threat
  - Distilling signatures
  - Detection mechanisms

Tentative Topics, con't

- Scanning
- Forensics
- Traffic Analysis
- Web Authentication & Attacks
- Anonymity

Tentative Topics, con't

- Botnets
- Architecture
- Legal & Ethical Issues
- Infrastructure Protection
- Wireless

Give Feedback

- Regarding syllabus
  - Topics/subtopics you'd like explored
  - Particular papers
- Post-lecture
  - We can revisit at beginning of next lecture
- Course mechanics
- Anonymous is fine if you want
  - Either using a remailer
  - Or just a note under my door (737 Soda)

Some Project Ideas - SP08

- Dynamic firewalls for data centers
- Security analysis of AirBears
- Distributed detection of spam sources
- Detecting "fast flux" DNS domains in real-time
- Literature survey of forensics
- Survey of SCADA security issues
- Assessment of the relationship between users and overall system security
- Automated vulnerability diagnosis for network services
- Efficacy of heuristics for detecting phishing sites

Some Project Ideas - Elsewhere

- Privacy exposure of social networking sites
- Software updater vulnerabilities
- Fingerprinting spam-generation software
- Detecting phishing sites by logo-matching
- (*) Documenting and explaining the "wholesale" traffic delivery business

Project Ideas, con't

- Reproduce a result from the literature
- Build a detector for traffic injection (DNS, ARP) and run it widely
- Javascript analysis / rewriting (*)
- Constructing services with specific vulnerabilities (*)
- Spam classification & mining (*)
- Counterspam (*)
- XML/AJAX analysis & attacks

Project Ideas, con't

- Attribution architecture (*)
- Finding exploitable flaws in botnet C&C (*)
- Relationship between whois data & malice (*)
- Mining network tools for protocol archaeology (*)
- End-user/middlebox negotiation architecture (*)
- Evidence of spammers hijacking address blocks via BGP
- Longitudinal traffic analysis (scanning, service flux; *)
- Cloud computing security analysis

Project Ideas, con't

- Is on-line poker fair?
- Automated analysis of IRC chat focussed on illegal transactions (*)
- Spoofing GPS via software radio?
- How does web malice change depending on user agent / fingerprint / IP address?
- Assessment of accuracy/efficacy/overlap of blacklist feeds

Project Ideas, con't

- Build software to inject artificial information into keyloggers/spyware/email harvesters, verify that it works
- Can you detect vote fraud in YouTube/EBay/Amazon?
- Analysis/detection of blog spam
- Analysis of DDoS traces
- Fingerprinting malware family trees by the data structures they use

Some Possible Project Resources

- Trace/log analysis *mediation*
- whois data, DNS churn
- Blacklist feeds
- Malware specimens
- Javascript corpus
- NetGear boxes

Next Lecture

- Denial-of-Service
- Homework #1 due Friday 1PM
  - Writeup for “Backscatter” paper
  - Check out the syllabus
  - Background survey
  - Optional: read/write up *TCP DoS* or *Reflector Attacks* papers