
# Internet Background Radiation

Ruoming Pang

Princeton University

# ***Cosmic*** Background Radiation

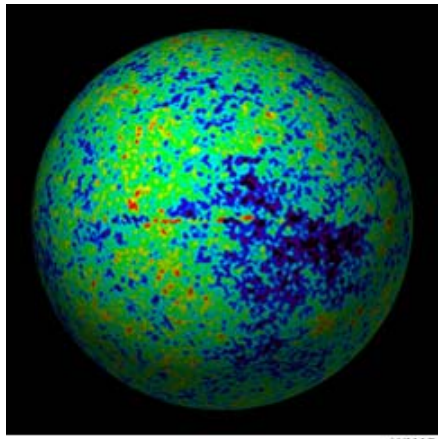


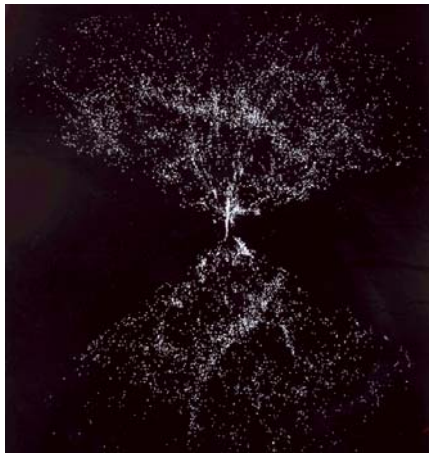

- ***Cosmic Background Radiation (left):*** "in every direction, there is a very low energy and very uniform radiation that we see filling the Universe."
- Afterglow of the "Big Bang": helps us to understand formation and structure of the universe (right)

# Internet Background Radiation

- The Baseline "Noise" of Internet traffic
  - Every IP address---even an unused one---receives packets *constantly*…

This talk:

- What does it look like?
- Why does it exist?
- What can it tell us about the Internet?

# Measurement Apparatus: Network Telescope

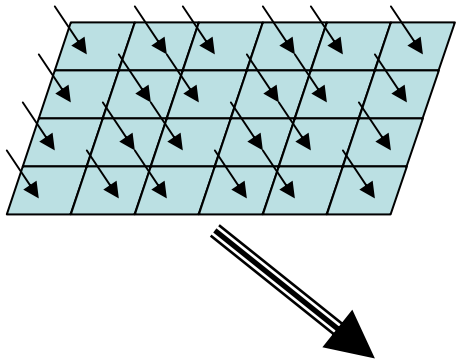

- *Unused but globally reachable IP addresses*

- Our main telescopes:
  - Lawrence Berkeley National Lab
  - Size: 1,280 addresses x 2

# Background

- Why do we study it?
  - To understand Internet malware *in action*
- Related work
  - DoS attack backscatters [Moore]
  - Analysis of specific worms (e.g. Witty)
  - Building honeypots (honeyd, iSink, Potemkin, GQ)
  - Measurements on a set of *distributed* telescopes (Internet Motion Sensors)

# Our Study

The first broad characterization of Internet background radiation

- **Focus: <u>traffic semantics</u>**
  - What is the traffic trying to do at application level?
- Measurement methodology
  - How to extract the meaning of background radiation (on large scales)?

# Background Radiation Hit Pattern

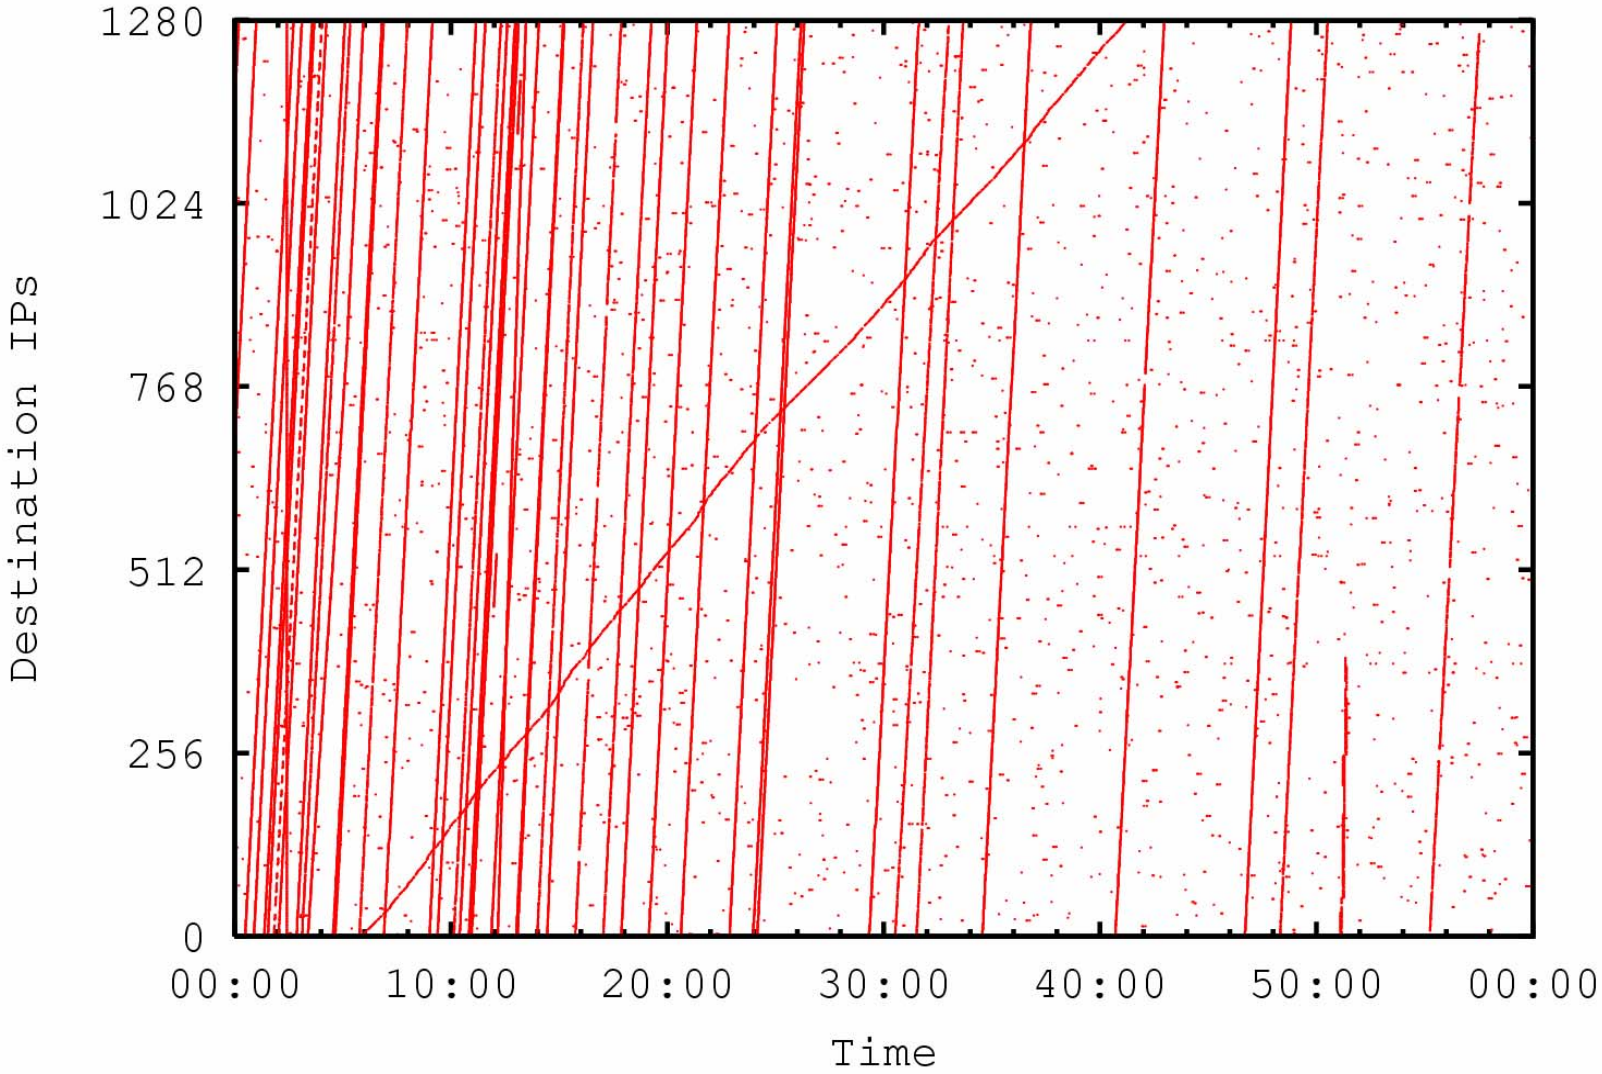

# How Often Do We See a Packet?

- Feb 2006 at Lawrence Berkeley Lab
  (Average on 1,280 IP's over period of a week)
  342 packets / destination IP / day
  $\Rightarrow$ **A packet every 4 minutes on any IP**

- But, how are radiation packets distributed:
  - Among destination IP's? (Hotspot?)
  - Over time? (Burst?)

# Distribution over Destination IP's

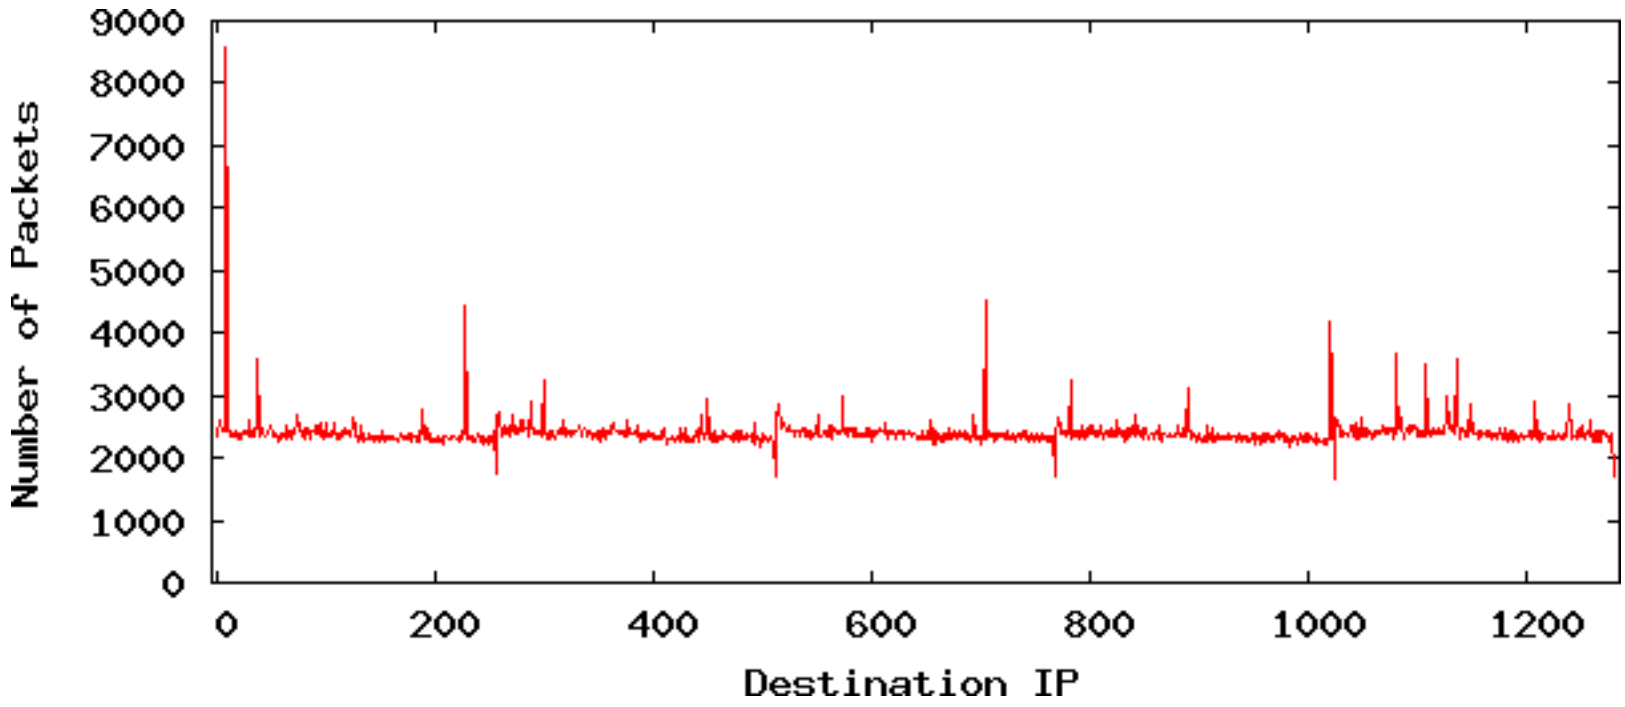


- Number of packets per destination IP received over a week

# Distribution over Destination IP's

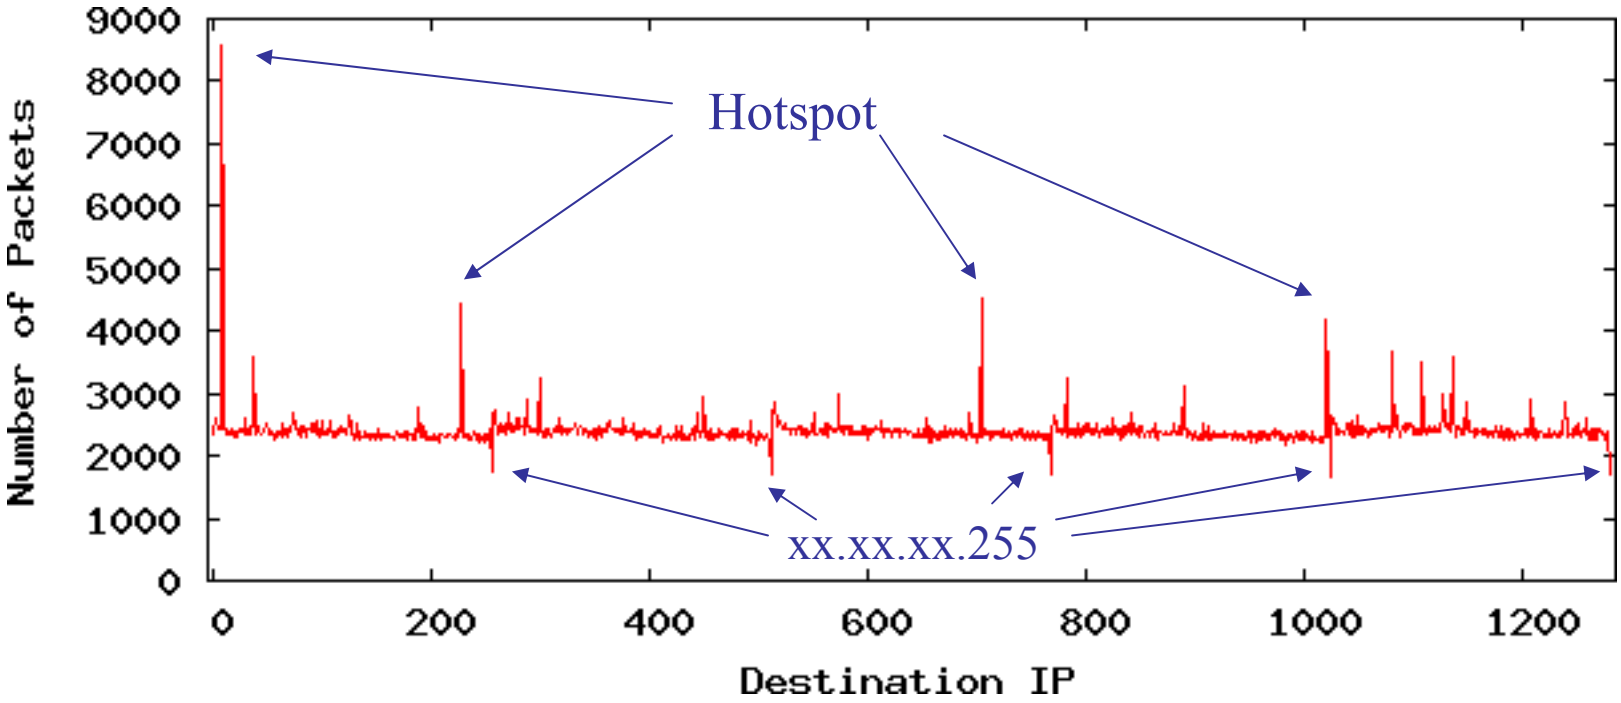


- Packets are in general **evenly distributed** among destinations
- The biggest hotspot receives < 1% of packets

# Number of Source IP's Seen Per Destination

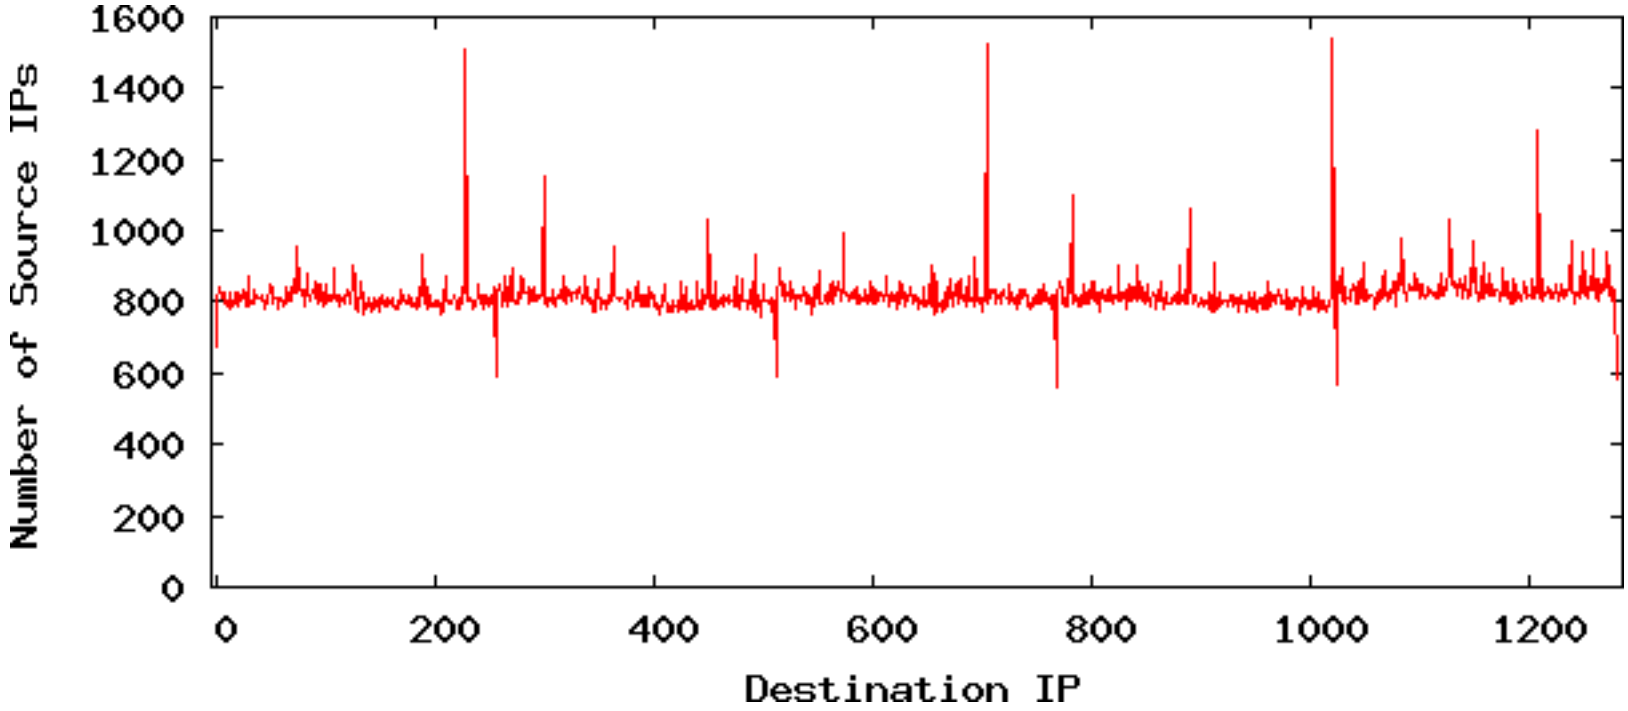


- Number of source IP's also quite evenly distributed

# Number of Packets Per Hour

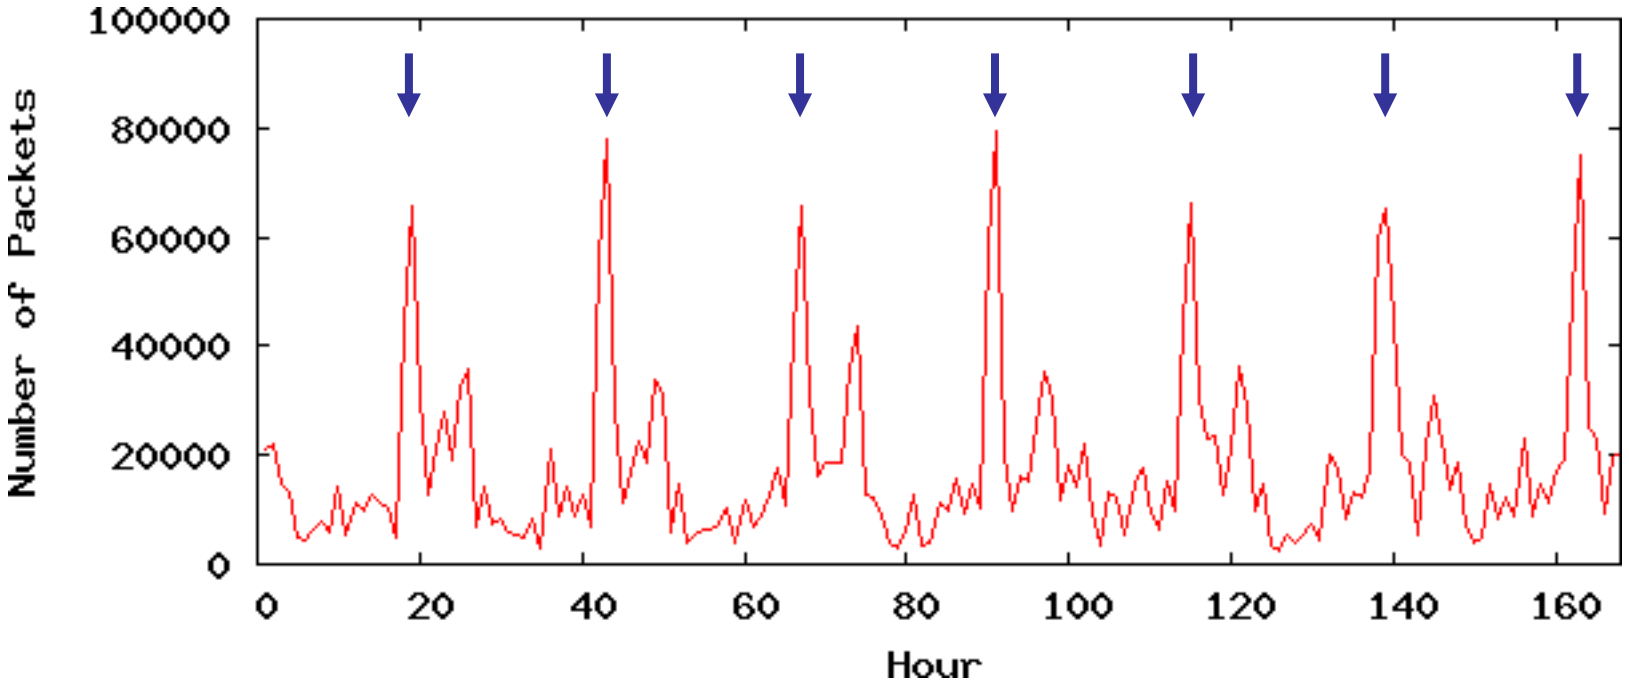


- Daily patterns (and lunch time spikes)

# What's inside Lunch Time Spikes?

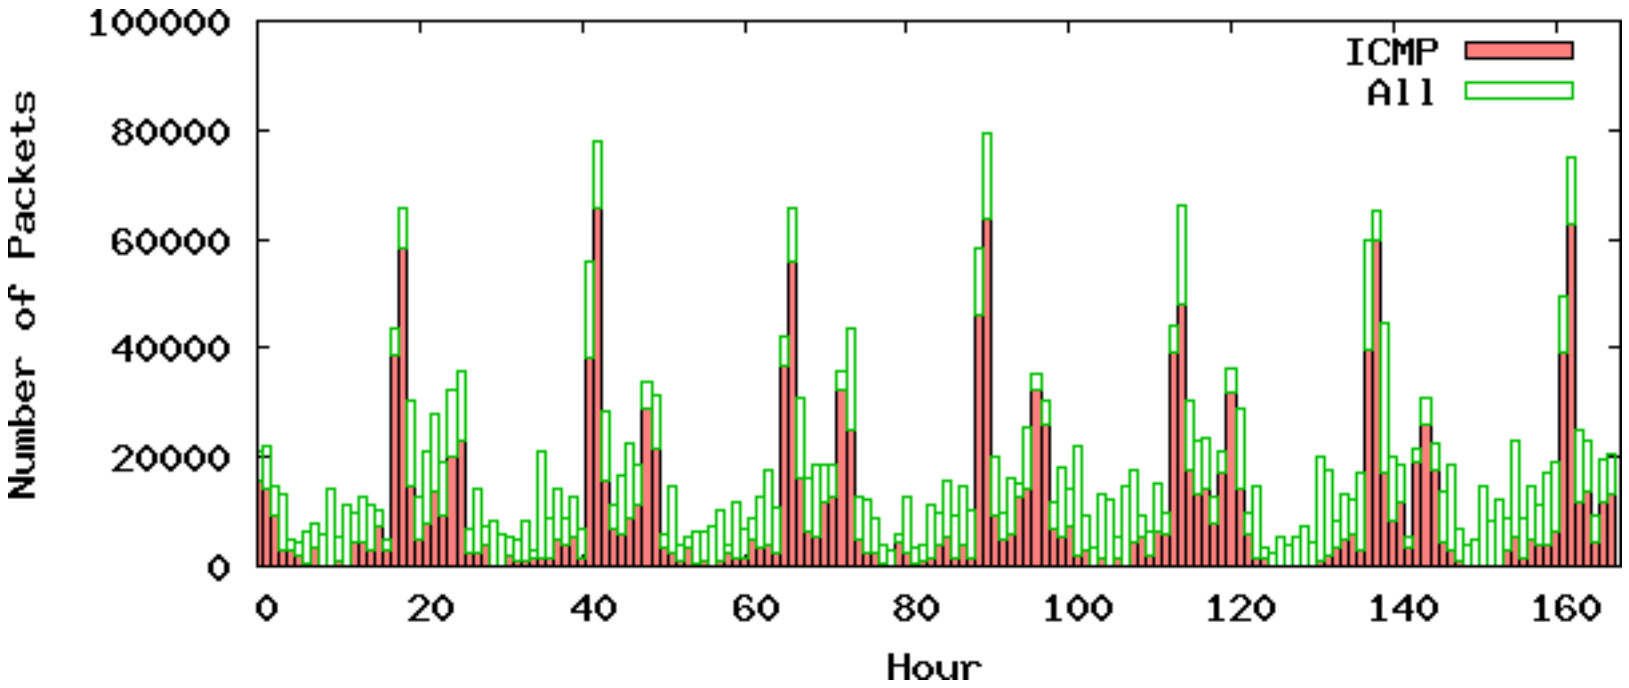


- Lunch time spikes are dominated by ICMP "Ping"

# Number of Source IP's Per Hour

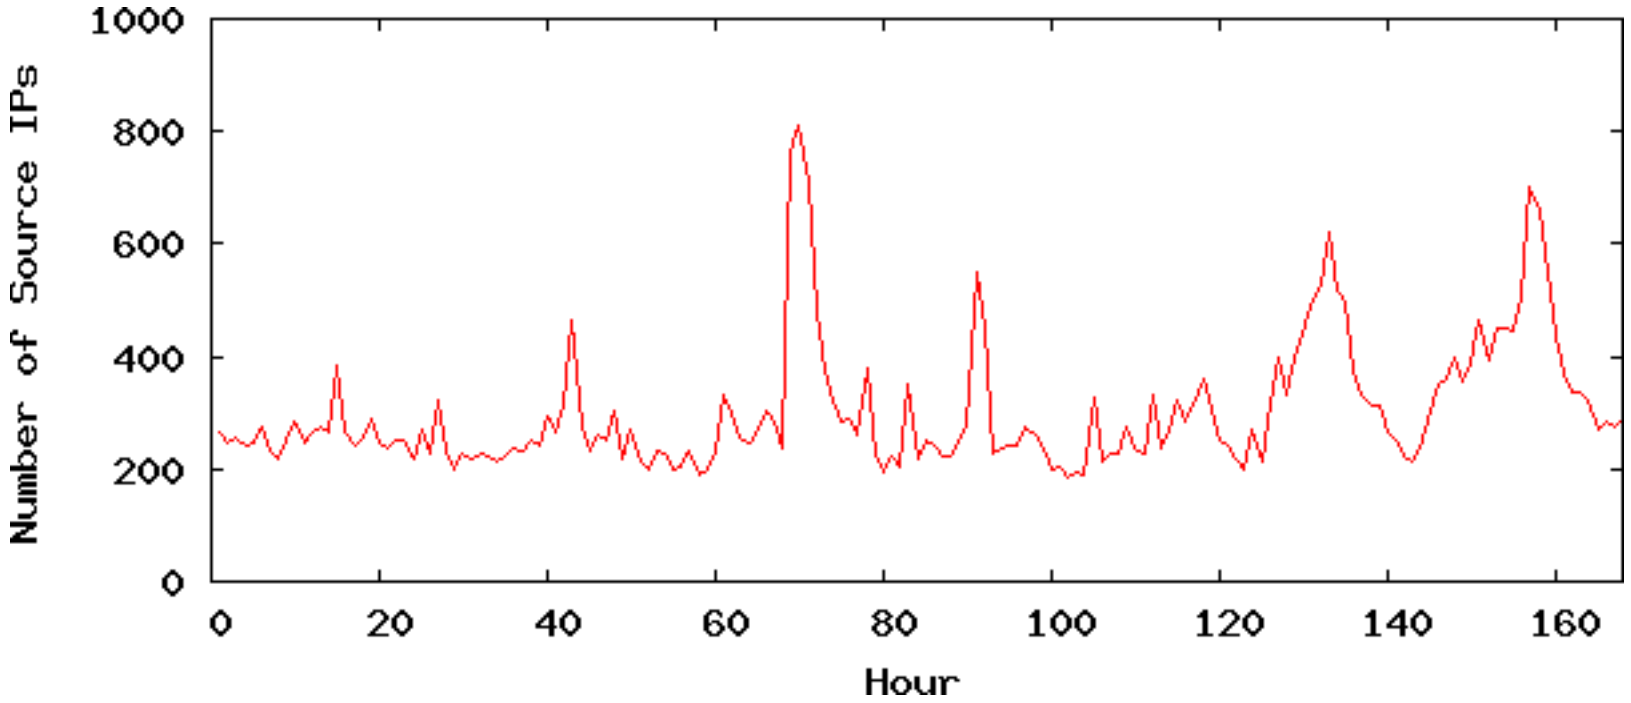


- Number of source IP's also vary over time
  But <u>not correlated</u> with packet volume

# Variation of Number of Source IP's

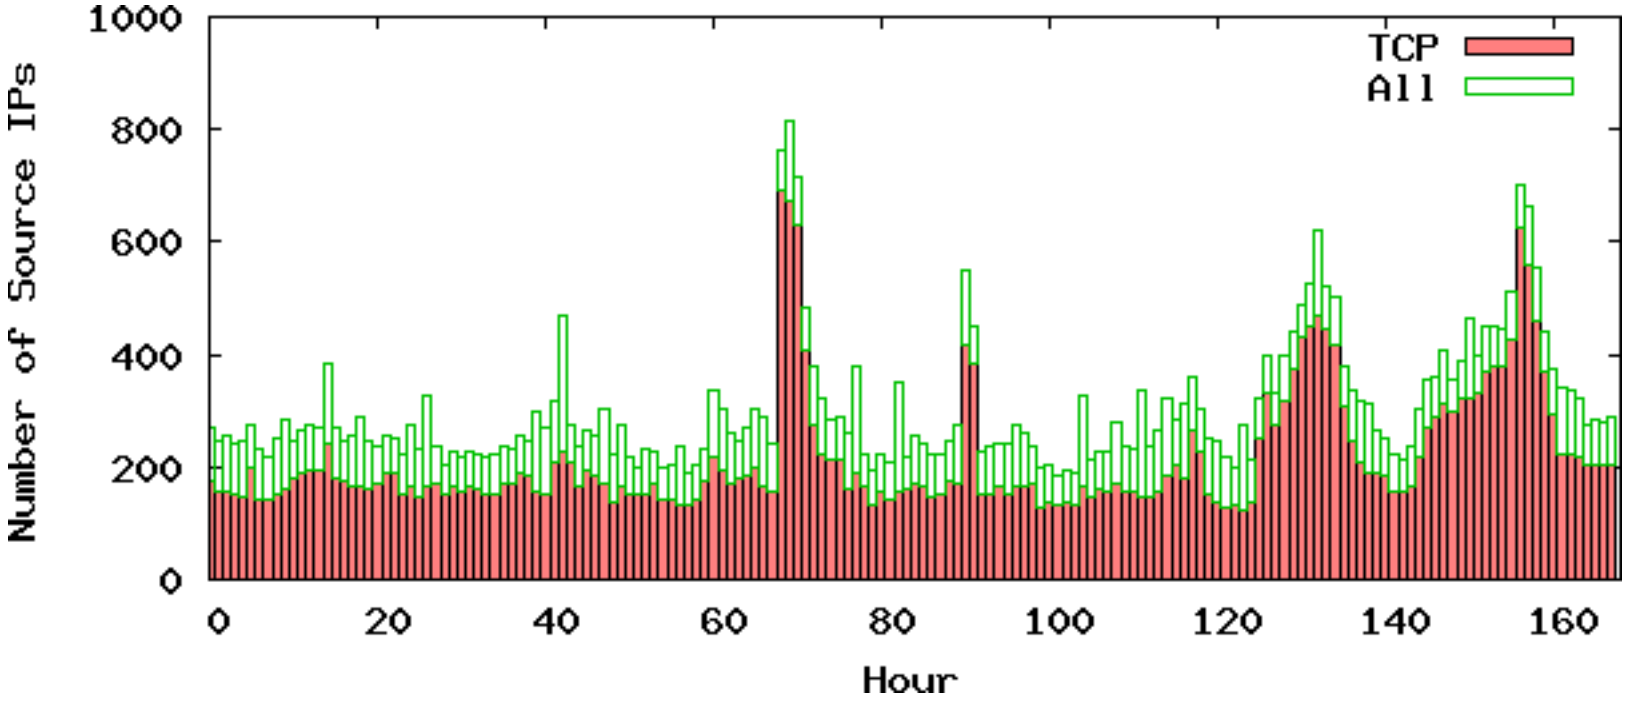


- Most sources are sending TCP/SYN's to initiate connections

# Summary of Passive Observation

- Near uniformity among destinations (for nearby addresses)
  - Hit pattern: sweeping or random
- Variation over time
  - Diurnal pattern and lunch time ICMP spikes
- Dominated by ICMP-ping and TCP/SYN
  - The sources are expecting replies!

- Verdict: highly automated and intentional

  *What do they want?*