Integrating Human and Synthetic Reasoning Via Model-Based Analysis

REDJACK

REDJACK

Introduction and Explanation

- This is an experimental idea and very rough
  - Glue together very tame AI and user interface through some fault trees
    - To capture knowledge
    - Improve efficiency
- Overview of work
- (My) Questions!

REDJACK

If you haven't seen this slide, you haven't attended any of my talks

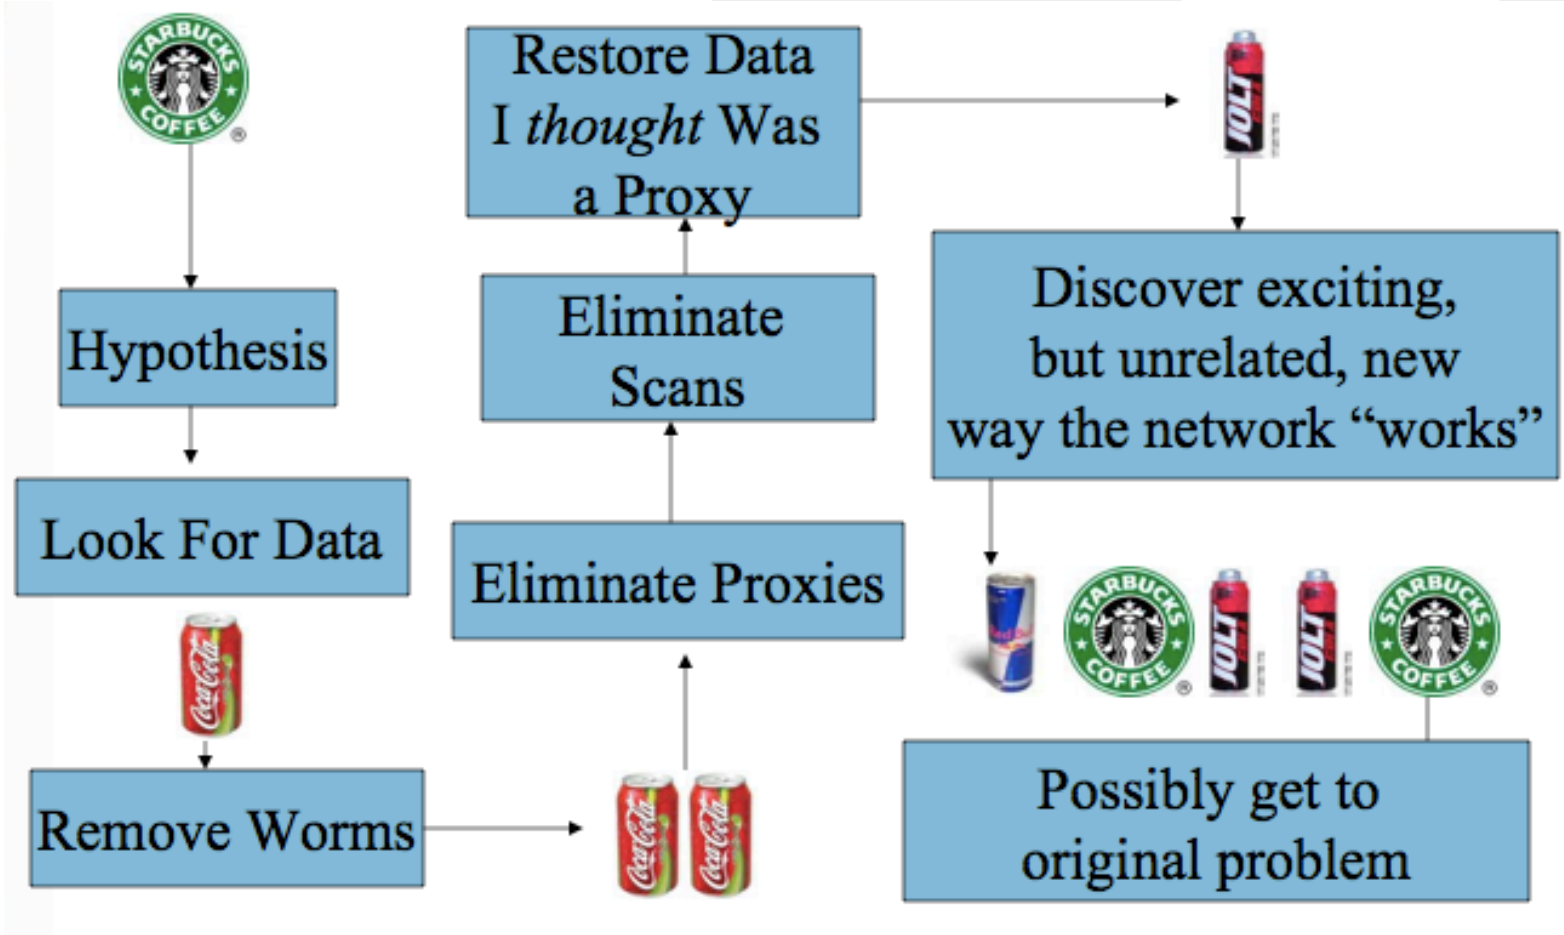

REDJACK

How much do we know about network traffic?

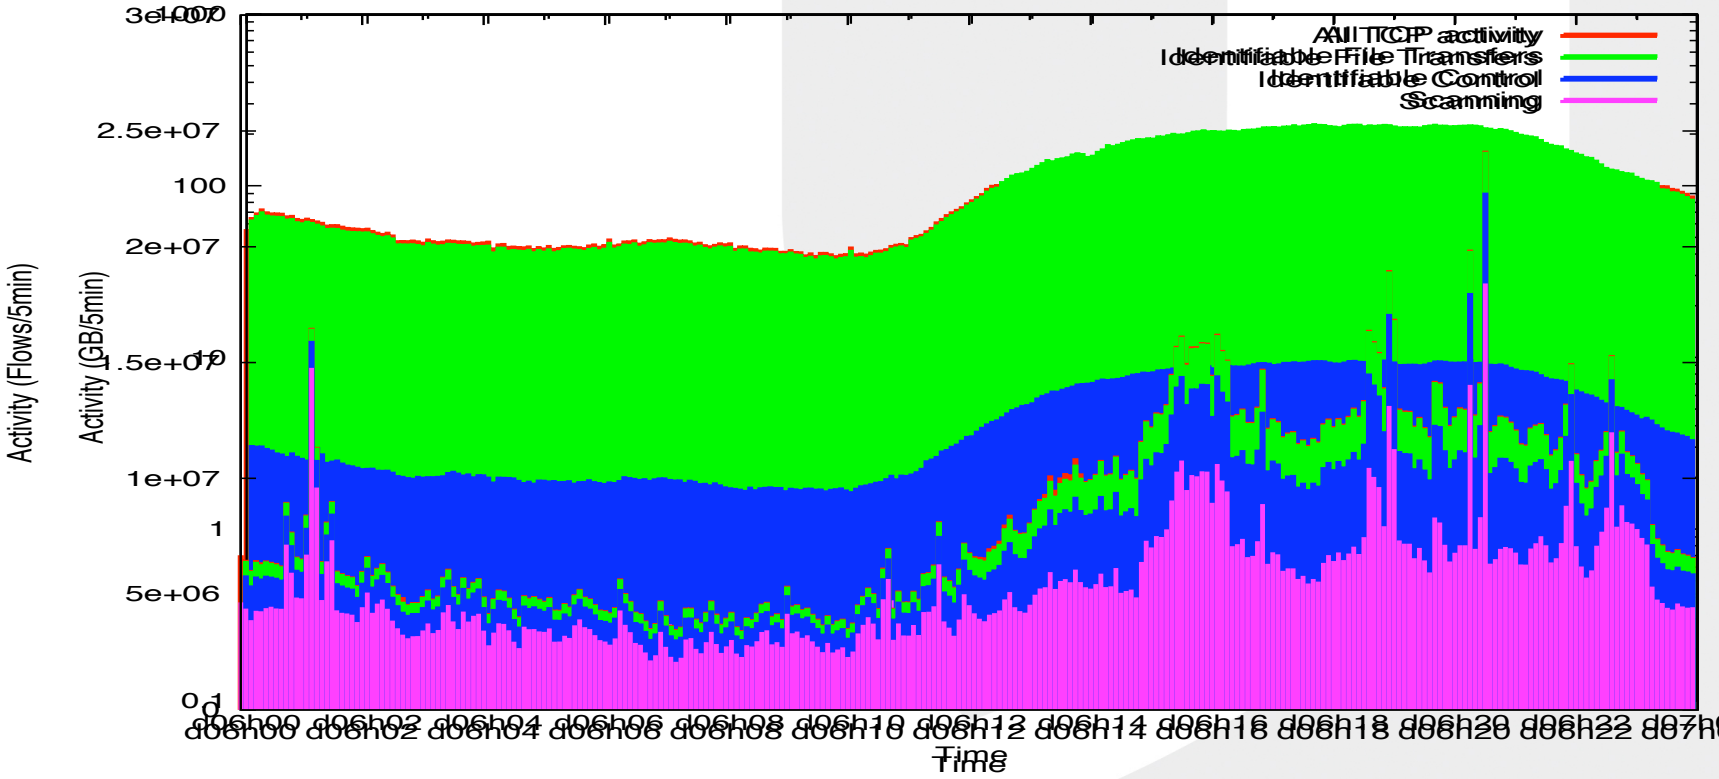

REDJACK

Basic problem

- We don't know what we know and we don't know what we don't know
- Most valuable resource available is analyst head time
- Lots of repetitive mindless attacks
- Lots of low-risk, high-threat attacks
- Have to *automate*
- Also have to ensure *automation isn't self defeating*

REDJACK

A Metric For Knowledge

- Every day we receive k-billion flows
  - We can understand and accurately tag x% of them
  - As x approaches 100%, the better
- We improve x:
  - Hiring more analysts
  - Reducing traffic into the network
  - Automating the process
  - Describing multiple flows at 1 time

REDJACK

Prototype System Diagram

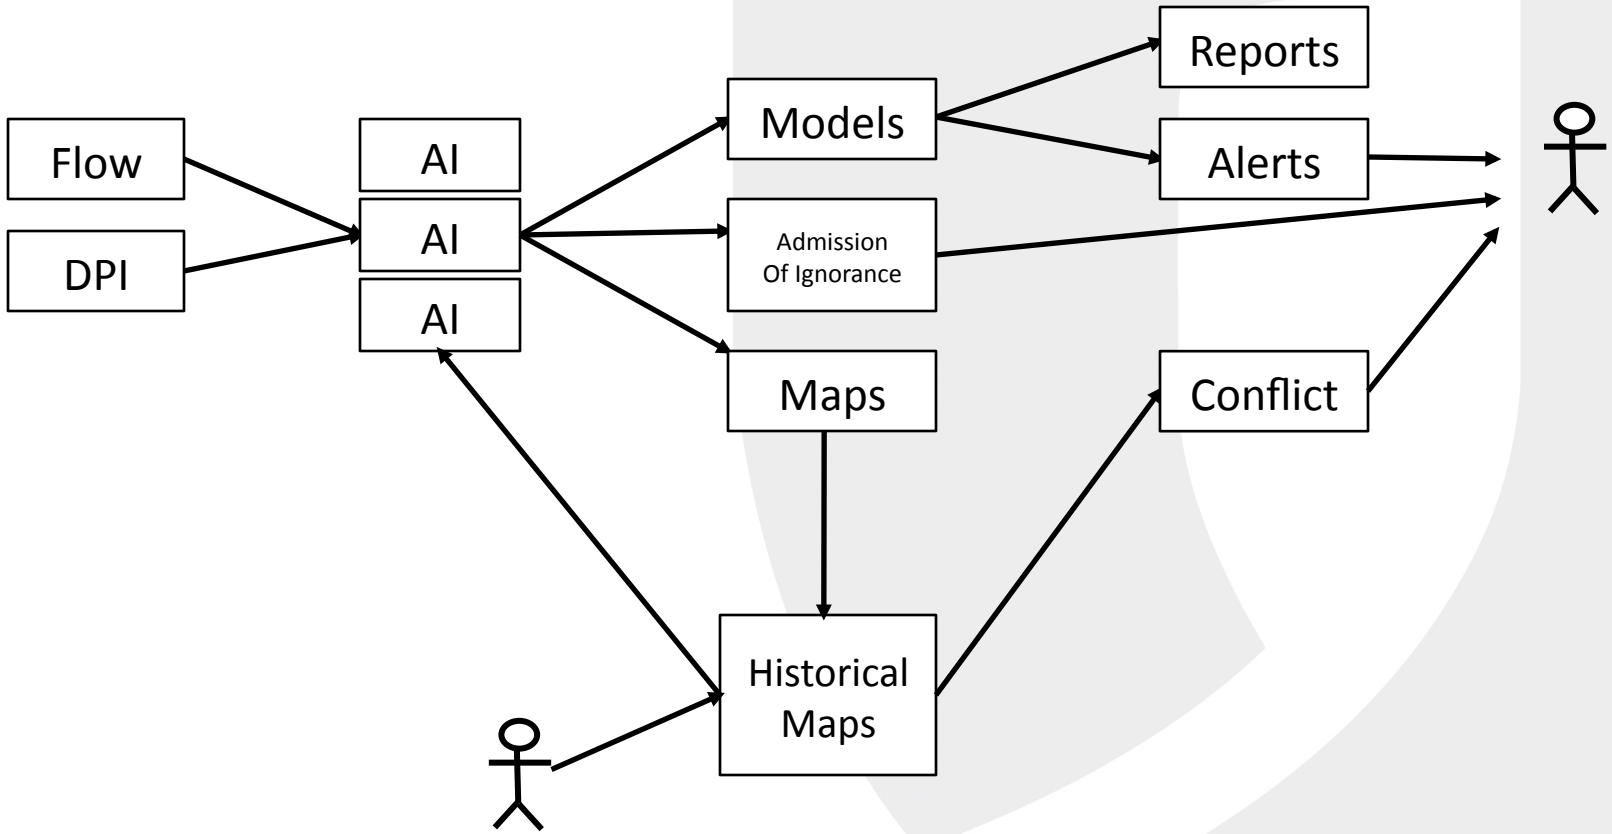

REDJACK

What is an AI?

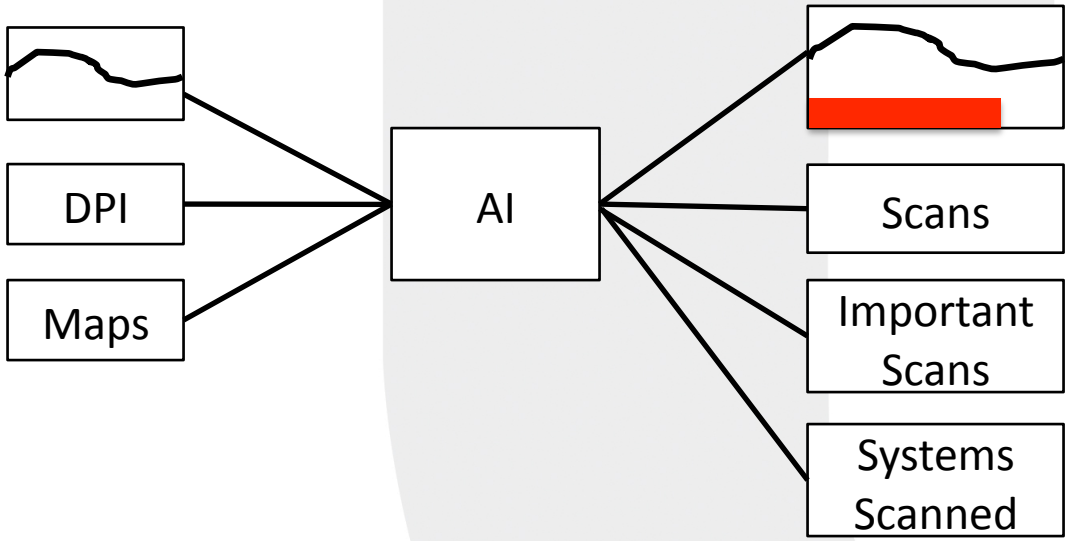

- An AI is a system that reads in network data and outputs:
  - A *domain*
  - Some *models*
  - *Alerts*
  - Inventory data

REDJACK

Complementary AIs

- Accurate
- Predictable
- Unambiguous human/ machine communication
- Humans serve as the final judge
- Don’t overwhelm with trivia

REDJACK

Accuracy

- Control ambiguity
  - ROC curves provide us with a measure of accuracy
  - But we've generally been unsure about what TP to use
- AIs will not guess in order to avoid a bad guess

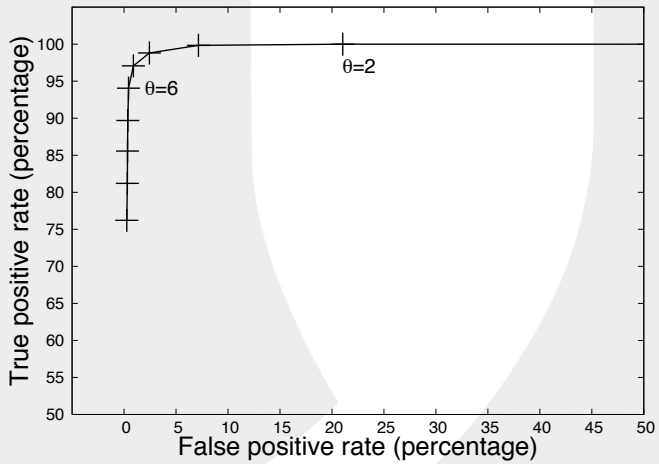

REDJACK

Predictable

- There isn't much we can do…
  - Reports: periodic and predictable information on the state of the system (e.g., scanning)
  - Alerts: When an *actionable* event occurs, a notice of the event and a recommended strategy (alter fw rules, take down machine, send people with guns)
  - Internal Intelligence: maps of the inside of the network
  - External Intelligence: maps of the outside world
- Inventory is central

REDJACK

Conflict Resolution

- We know that something is something
  - By fiat ("It's my webserver")
  - By published reference (port 80 is http)
  - Deep packet inspection (HTTP/1.0…)
  - Behaviorally (short requests, big transfers)
- Hierarchy of certainty
  - DPI >> Fiat >> Behavioral >> Published reference

REDJACK

Managing Conflict

DPI	Fiat	Behavioral	Published	Result
A	-	-	-	Map as A, alert on lack of published info
A	!A	X	X	Map as A, alert on conflict
A	A	X	X	Map as A
A	-	-	!A	Map as A, alert on masquerade
-	A	-	A	Map as A
-	A	!A	A	Report anomaly, Map as A
-	A	-	-	Map as A

REDJACK

Human/Machine Communication

- AIs don't raise alerts on normal behavior
  - Reports are for that
- AIs raise alerts on *actionable* anomalies
  - Provide diagnostics, inventory and history
- AIs raise alerts on conflicts
  - Rely on the user to resolve the conflict and move on

REDJACK

User Controls

- Everyone controls domains: sip, dip, sport, dport, time and protocol value
  - Domains have wildcards
- Agents *mark* or *subscribe* domain:
  - Mark: this happened in the past and I can infer what happened
    - For AI's, Mark indicates "I recognize this"
  - Subscribe: I will control and worry about this from now onto the future
    - For users, subscription says "This is my territory"

REDJACK

Models

- AIs don't output flow data
  - They mark off some segment of flows and group them together as a separate structure
- For example:
  - A "scan"
  - A "Bittorrent Network"
  - A "Surfing session"
- These models, in turn, have questions and structures that are more relevant to analysis
  - Who did the scan hit?
  - How much traffic was transferred in BT?

REDJACK

A Really Ugly UI

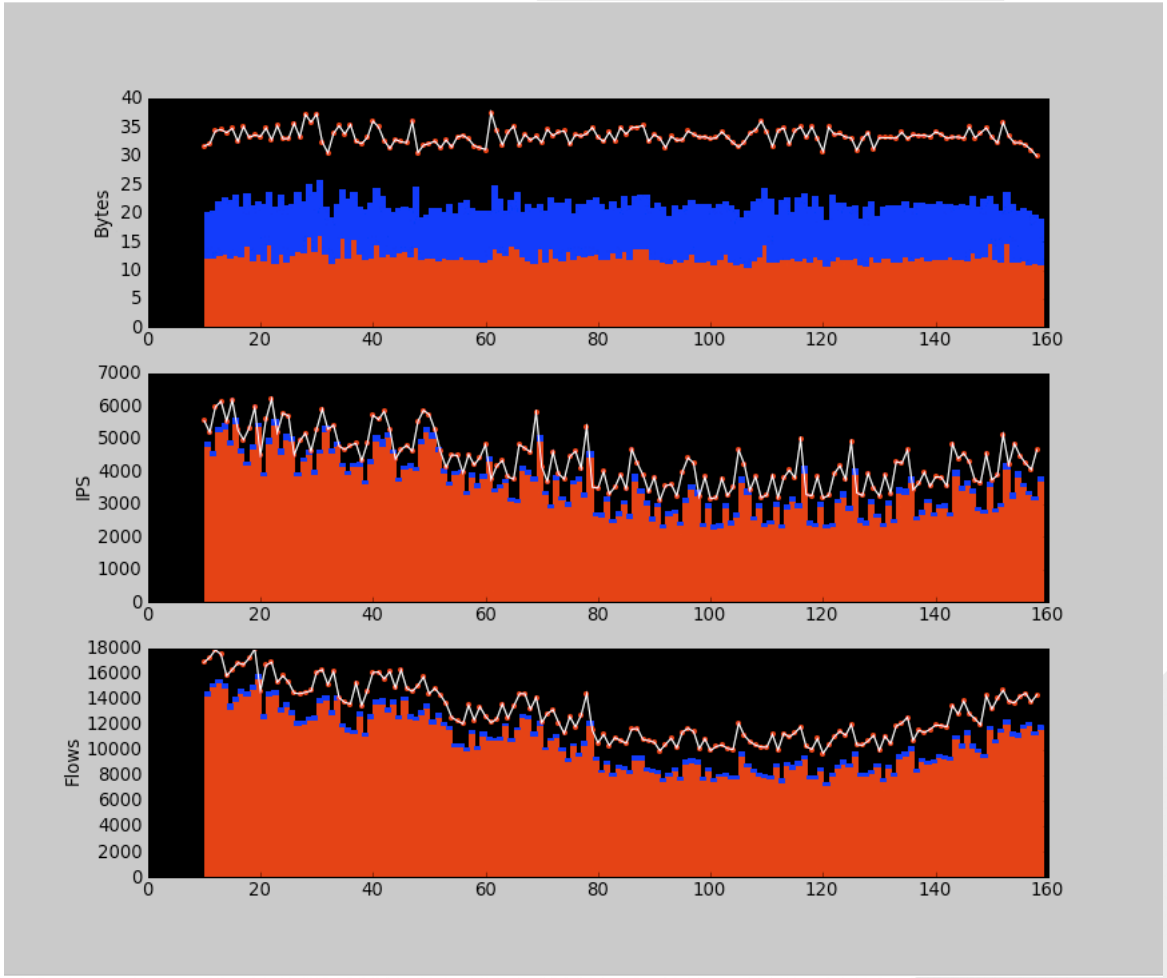

REDJACK

What that is

- Certainly not a testament to my visualization skills
- Prototype using two systems
  - Simple scan detection
  - BitTorrent detection
- The black is what's left

REDJACK

Problems

- As I said, this is all very rough right now
- Problems remaining:
  - Application/Knowledge Layering
  - Model Taxonomy
  - User experience
  - Backtracking
  - Metadata

REDJACK

Application Layering

- Make judgments at different levels of the stack
- Different inferential resolution:
  - Does this IP exist?
  - Does this IP communicate?
  - What does this IP communciate?
  - Is this IP significant in its network?

REDJACK

Model Taxonomy

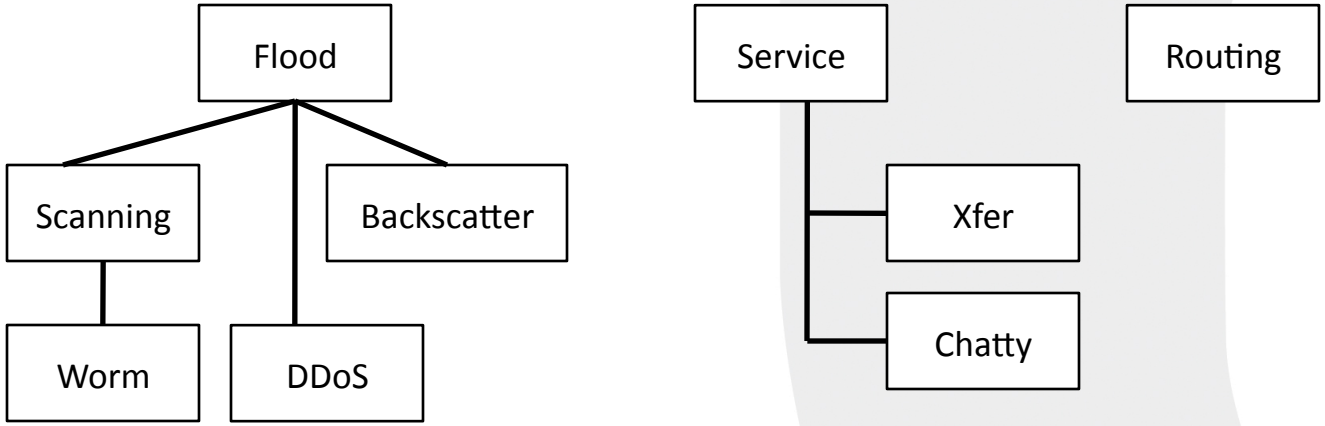

- Models replace flows with more compact descriptions of phenomena
  - E.g., “A Scan” is a list of the scanned IP’s, and anything that responded
- Trying to begin with broad behavioral descriptions and move down from there

REDJACK

Unsolved Problems

- Weirdometer metrics
  - Flows/IPS/bytes/IP pairs?
- Backtracking
  - How much do we want to see flow vs. model vs. map?
- Response Mechanism
  - What can a CSIRT do?
- Meta metrics
  - How much of the traffic do we understand?