

Detailed and understandable network diagnosis


Ratul Mahajan




With Srikanth Kandula, Bongshin Lee,
Zhicheng Liu (*GaTech*), Patrick Verkaik (*UCSD*),
Sharad Agarwal, Jitu Padhye, Victor Bahl

Network diagnosis explains faulty behavior

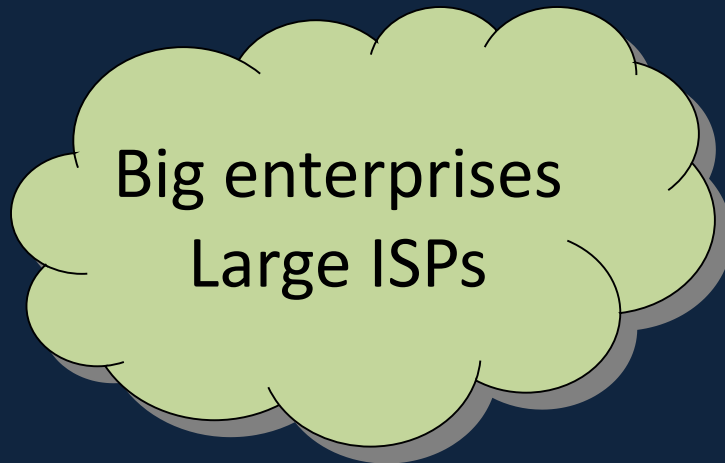
Starts with problem symptoms and ends at likely culprits



Current landscape of network diagnosis systems



Why study small enterprise networks separately?




Small enterprises



Less sophisticated admins

Less rich connectivity

Many shared components

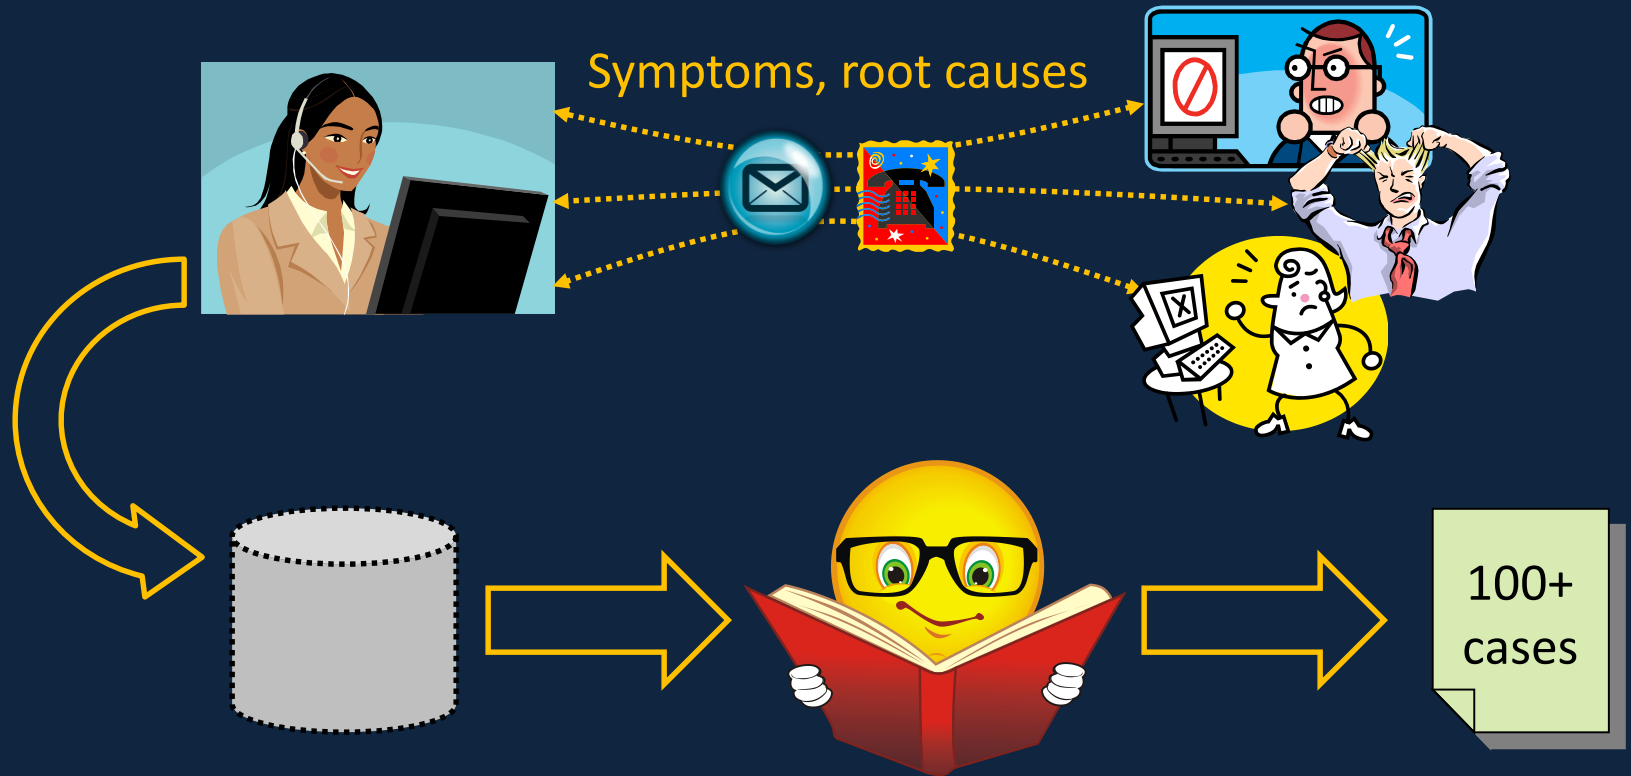


IIS, SQL,
Exchange, ...

Our work

1. Uncovers the need for detailed and understandable diagnosis
2. Develops *NetMedic* for detailed diagnosis
 - Diagnoses application faults without application knowledge
3. Develops *NetClinic* for explaining diagnostic analysis

Understanding problems in small enterprises



And the survey says



Symptom

App-specific	60 %
Failed initialization	13 %
Poor performance	10 %
Hang or crash	10 %
Unreachability	7 %

Handle app-specific
as well as generic faults



Identified cause

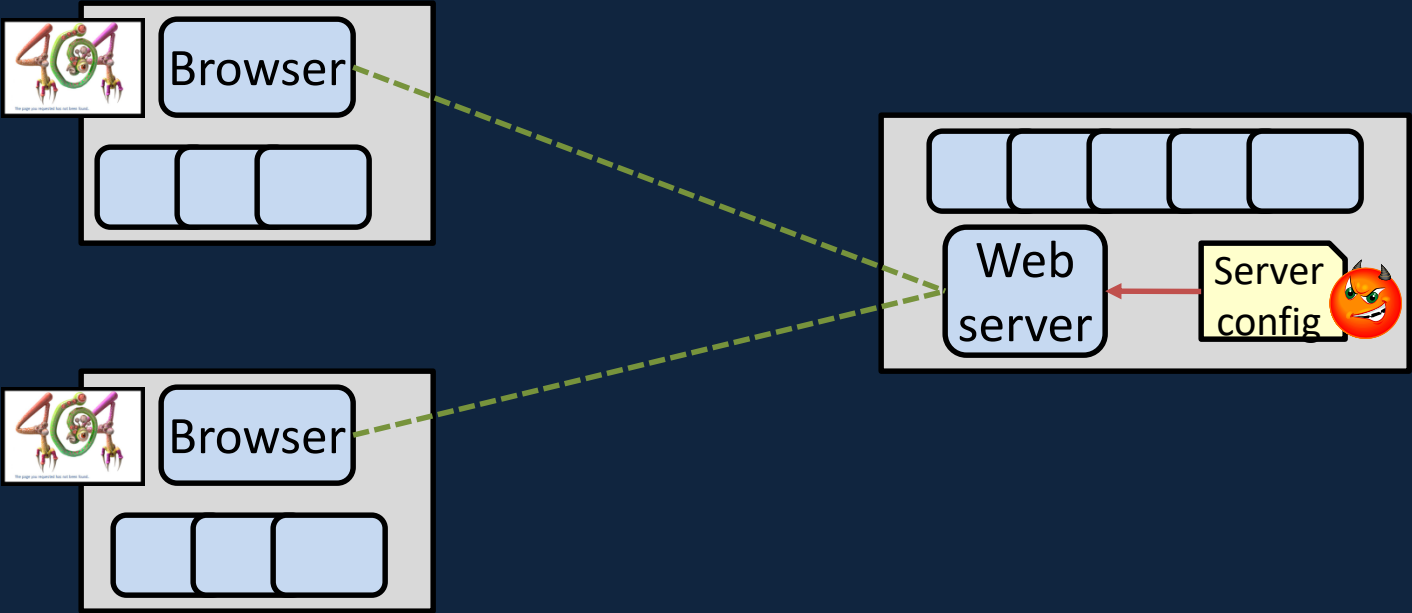
Non-app config (e.g., firewall)	30 %
Software/driver bug	21 %
App config	19 %
Overload	4 %
Hardware fault	2 %

Identify culprits
at a fine granularity

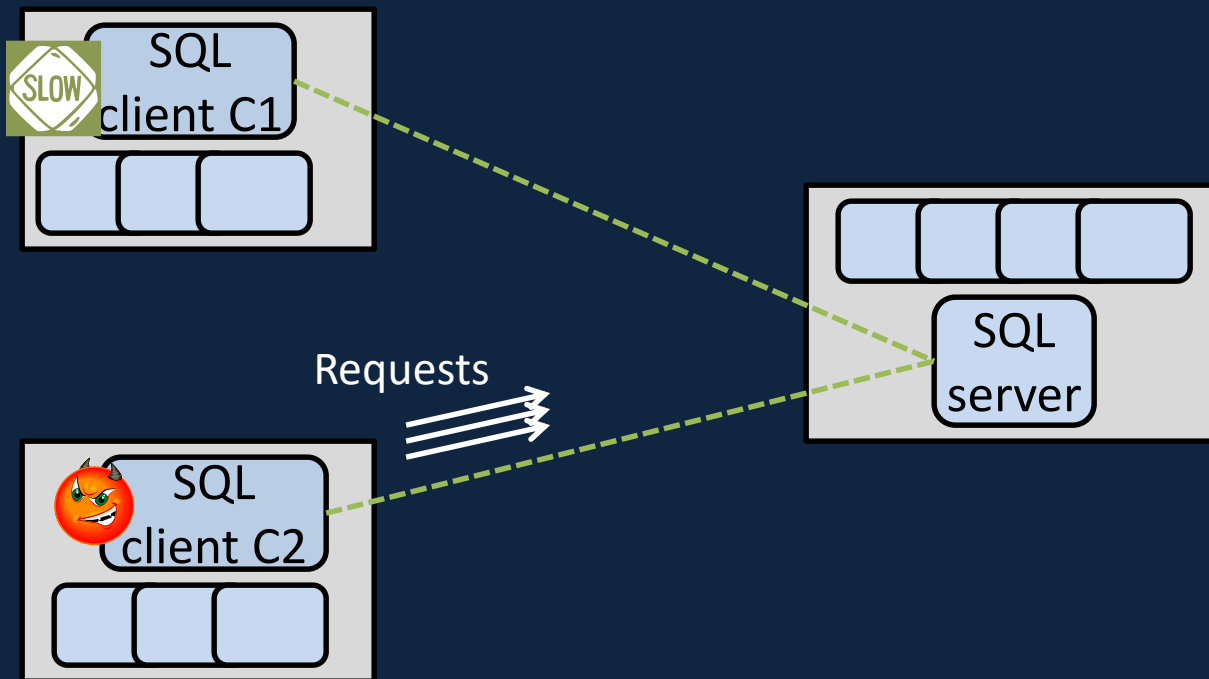


Detailed diagnosis


Example problem 1: Server misconfig



Example problem 2: Buggy client



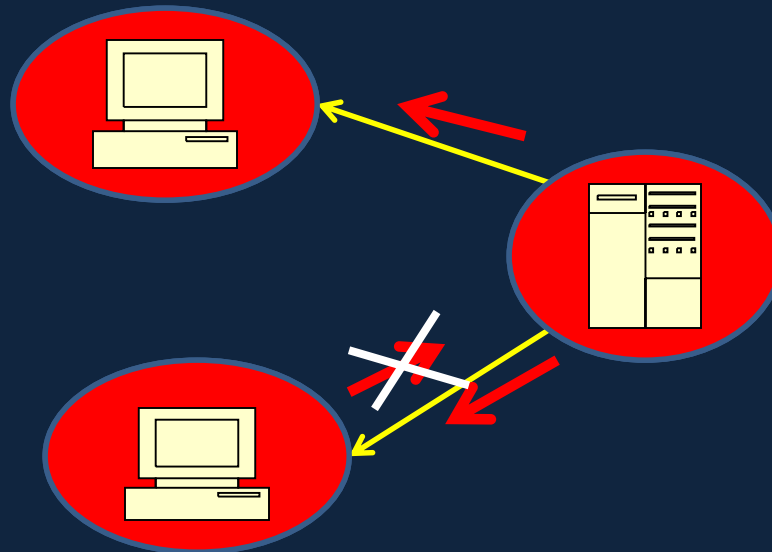
Example problem 3: Client misconfig



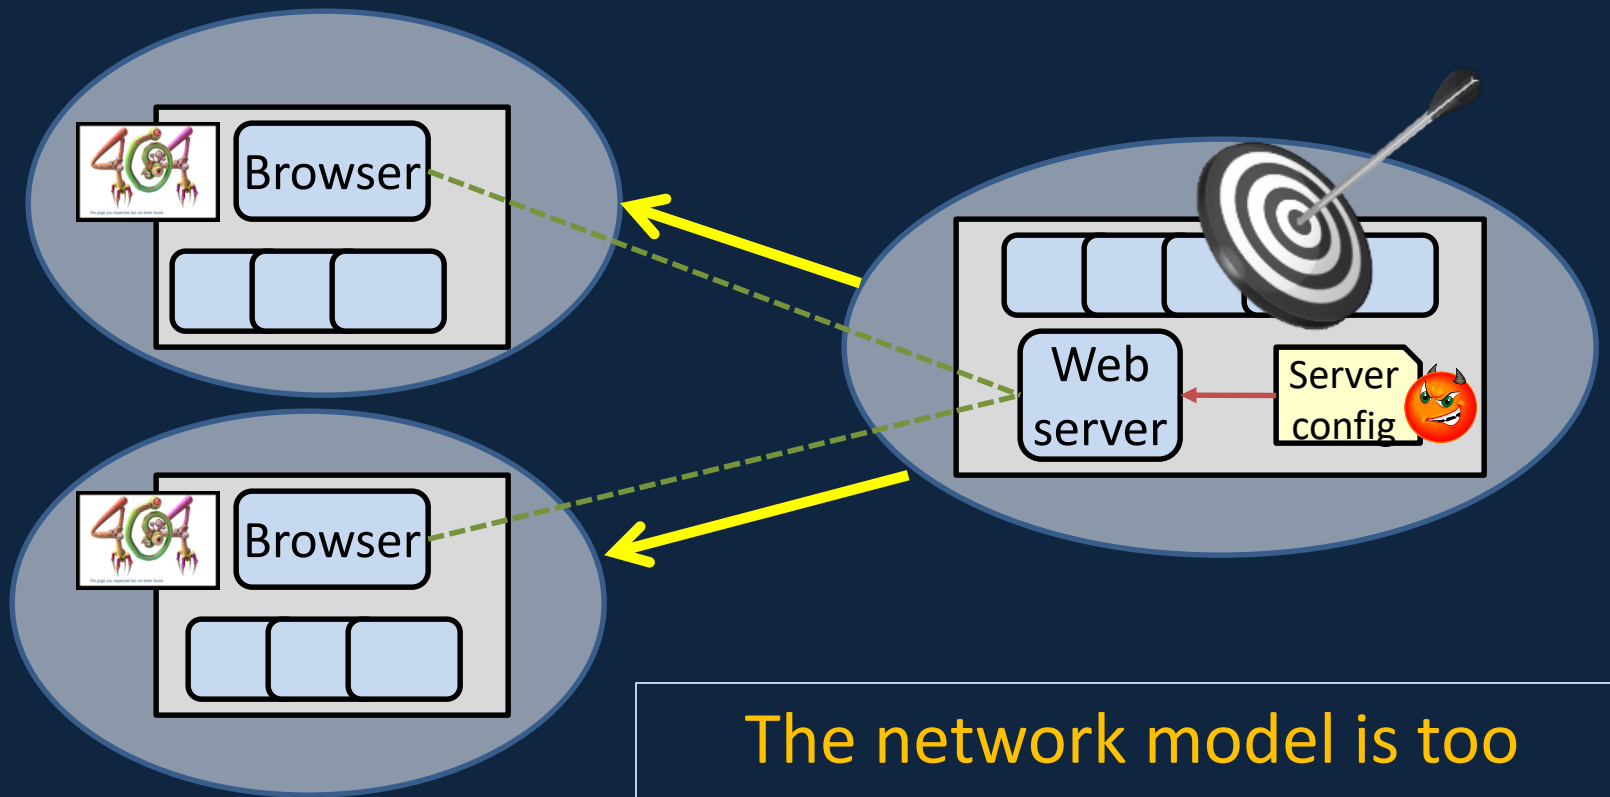
Current formulations sacrifice detail (to scale)

Dependency graph based formulations (e.g., Sherlock [SIGCOMM2007])

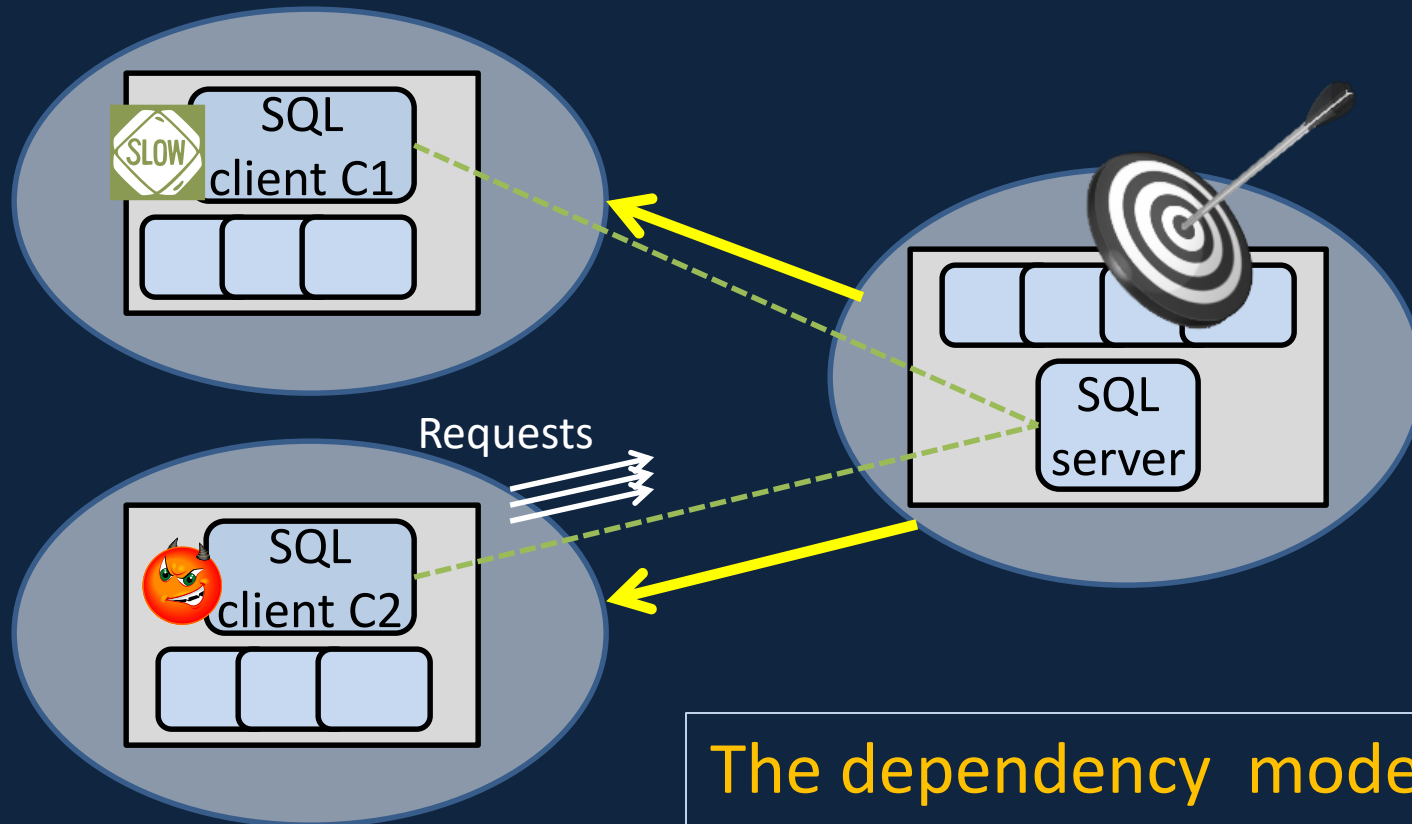
- Model the network as a dependency graph at a coarse level
- Simple dependency model



Example problem 1: Server misconfig

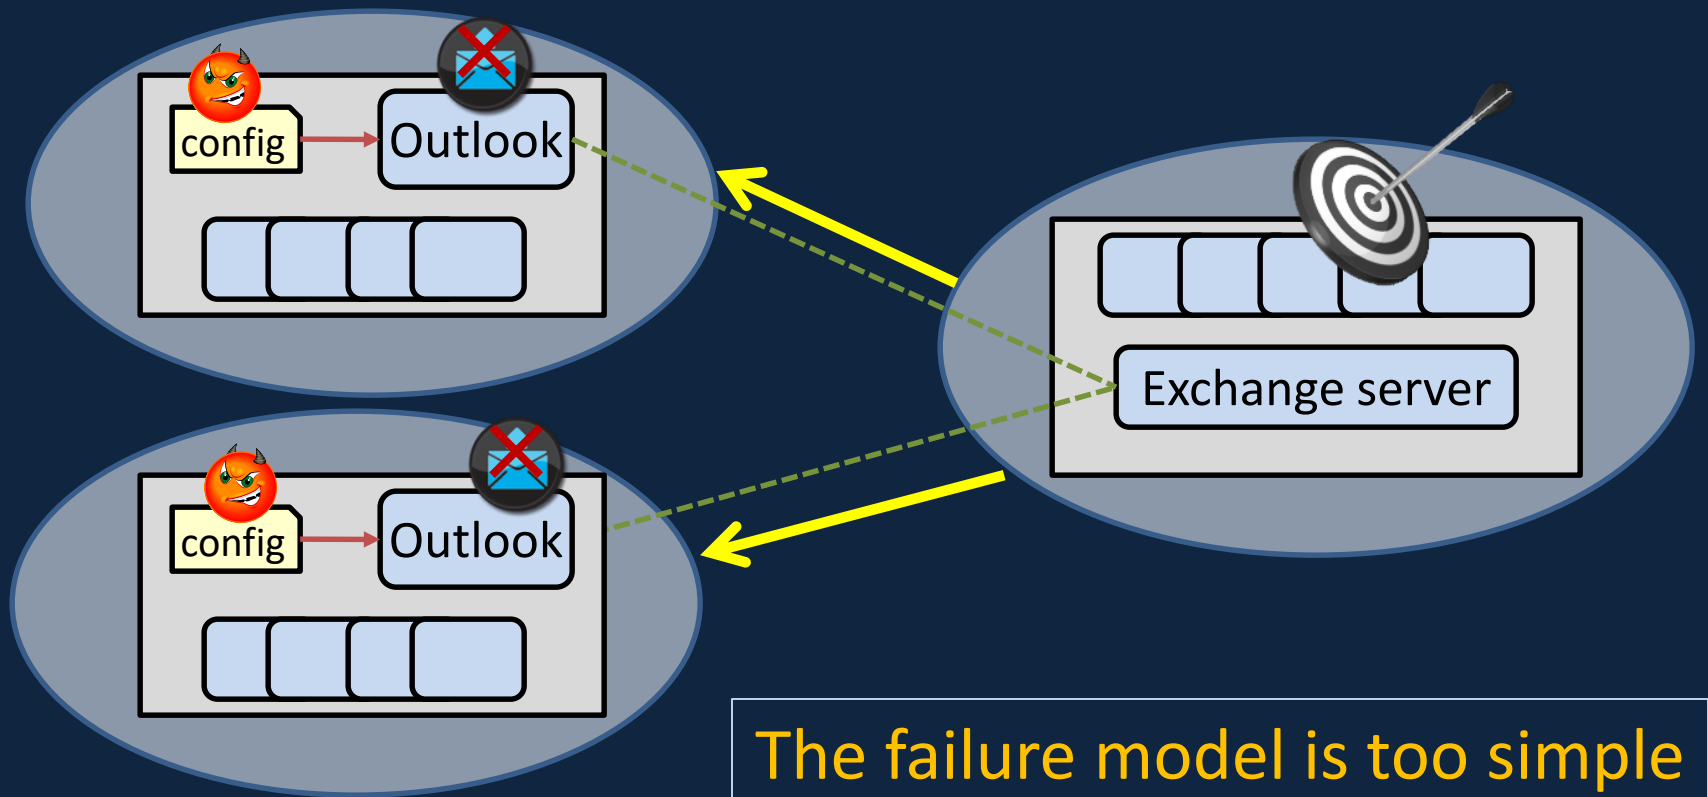


Example problem 2: Buggy client



The dependency model is too simple in current formulations

Example problem 3: Client misconfig

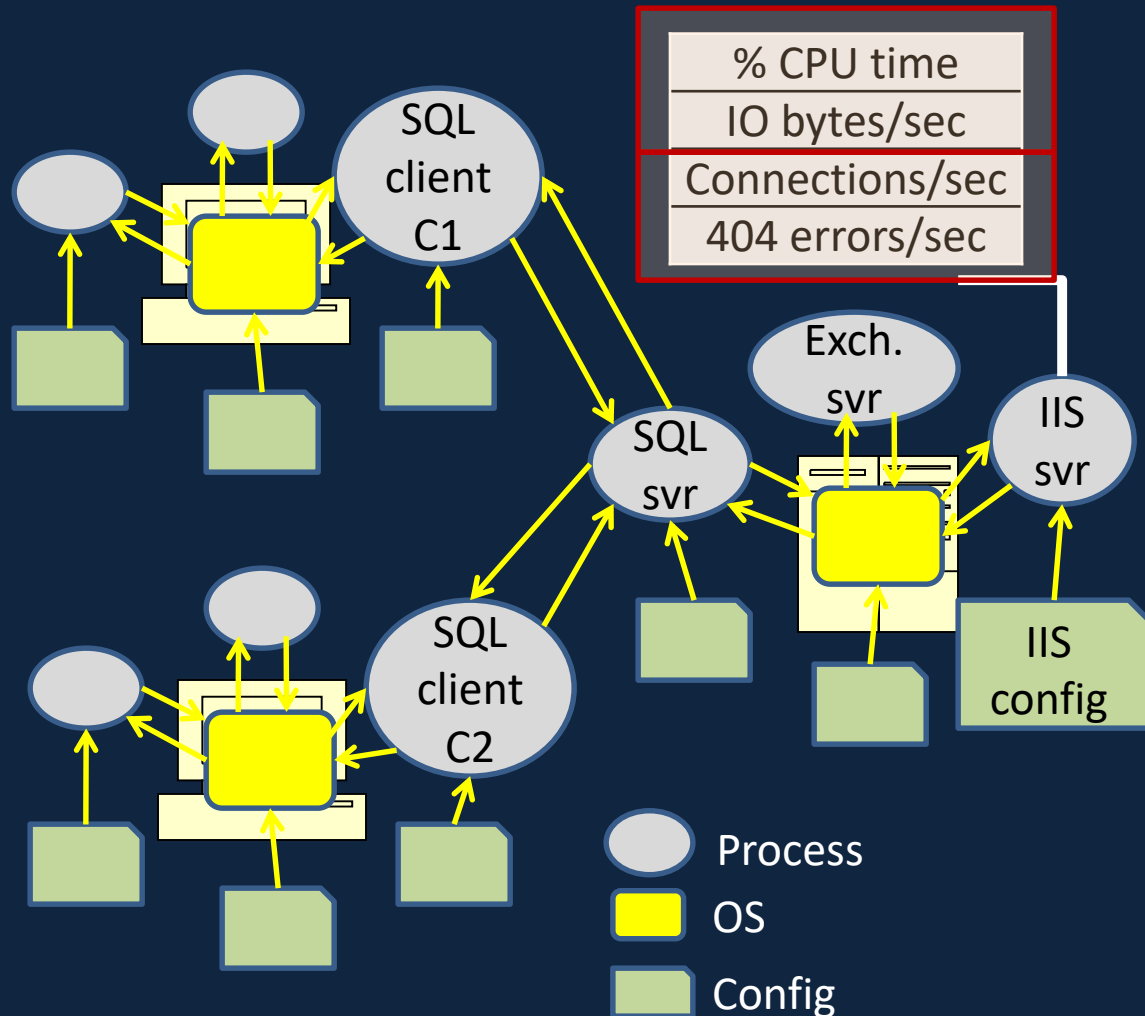


The failure model is too simple in current formulations

A formulation for detailed diagnosis

Dependency graph of fine-grained components

Component state is a multi-dimensional vector

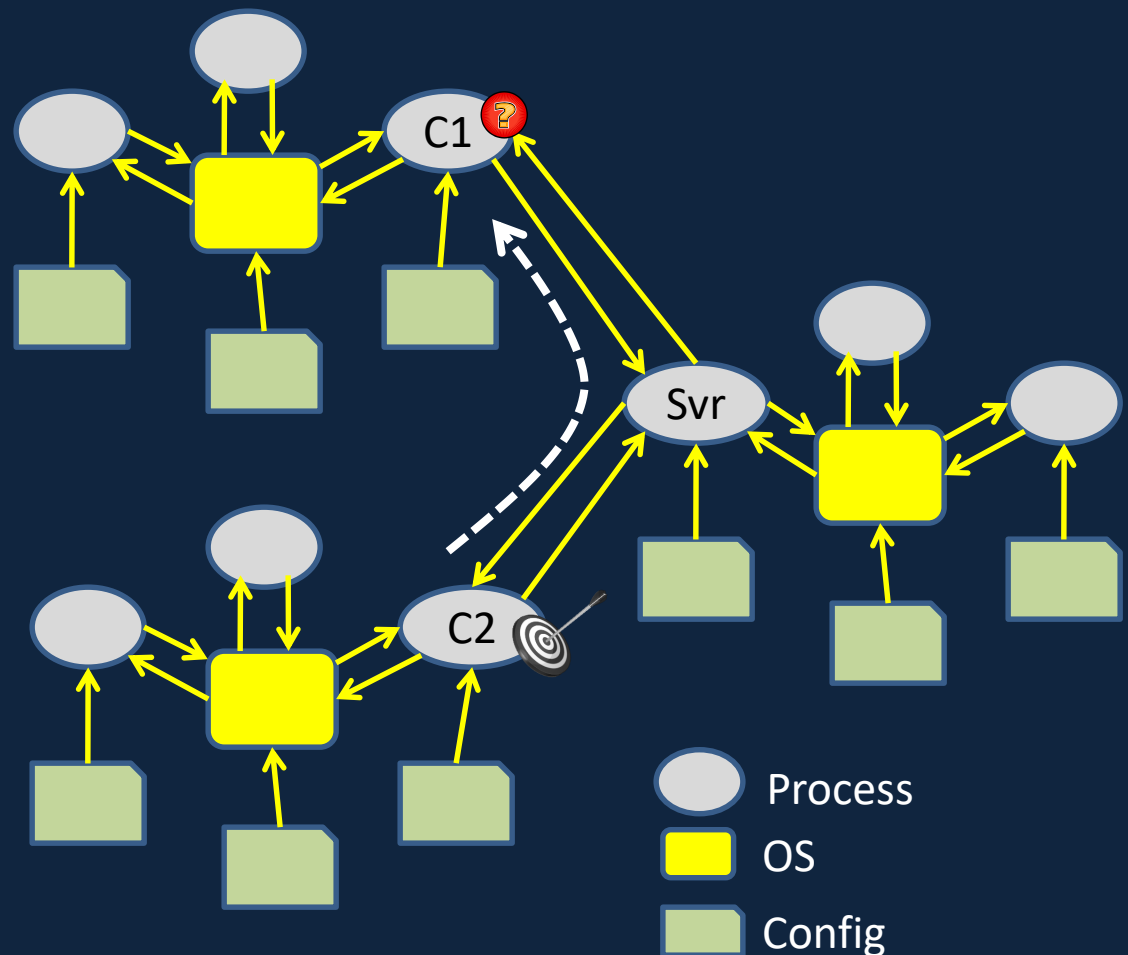


The goal of diagnosis

Identify likely culprits
for components of
interest

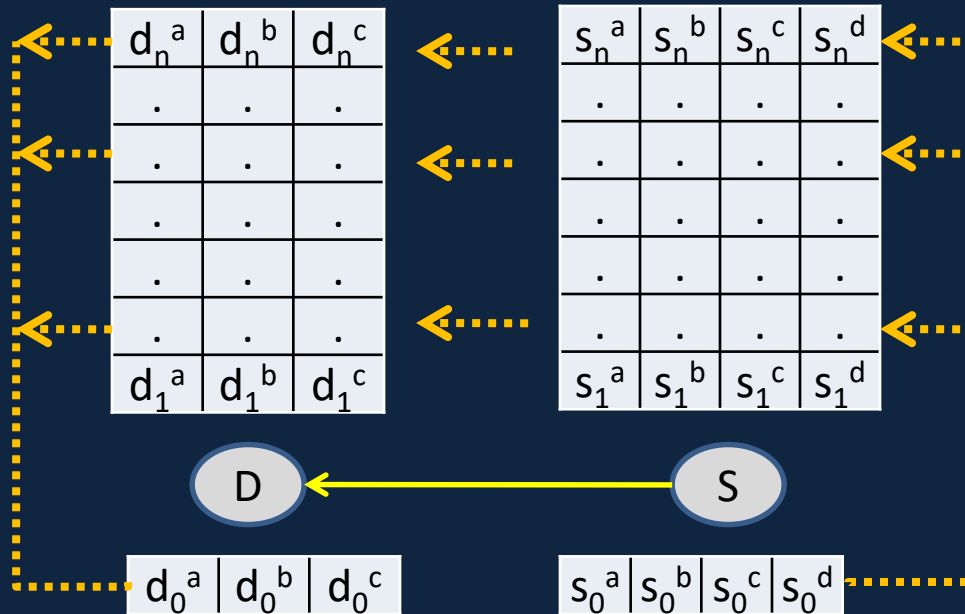
Without using
semantics of state
variables

→ No application
knowledge




Using joint historical behavior to estimate impact

How “similar”
on average
states of D are
at those times



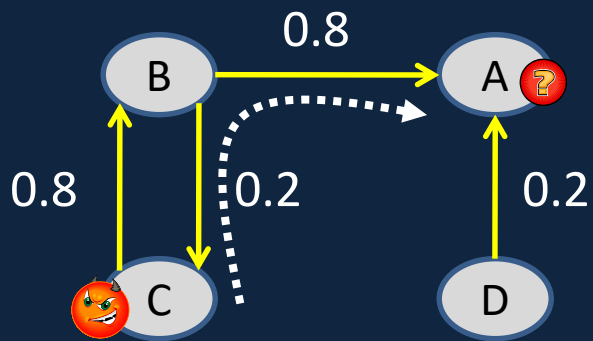
Identify time
periods when
state of S was
“similar”





Robust impact estimation

- Ignore state variables that represent redundant info
- Place higher weight on state variables likely related to fault being diagnosed
- Ignore state variables irrelevant to interaction with neighbor
- Account for aggregate relationships among state variables of neighboring components
- Account for disparate ranges of state variables



Ranking likely culprits



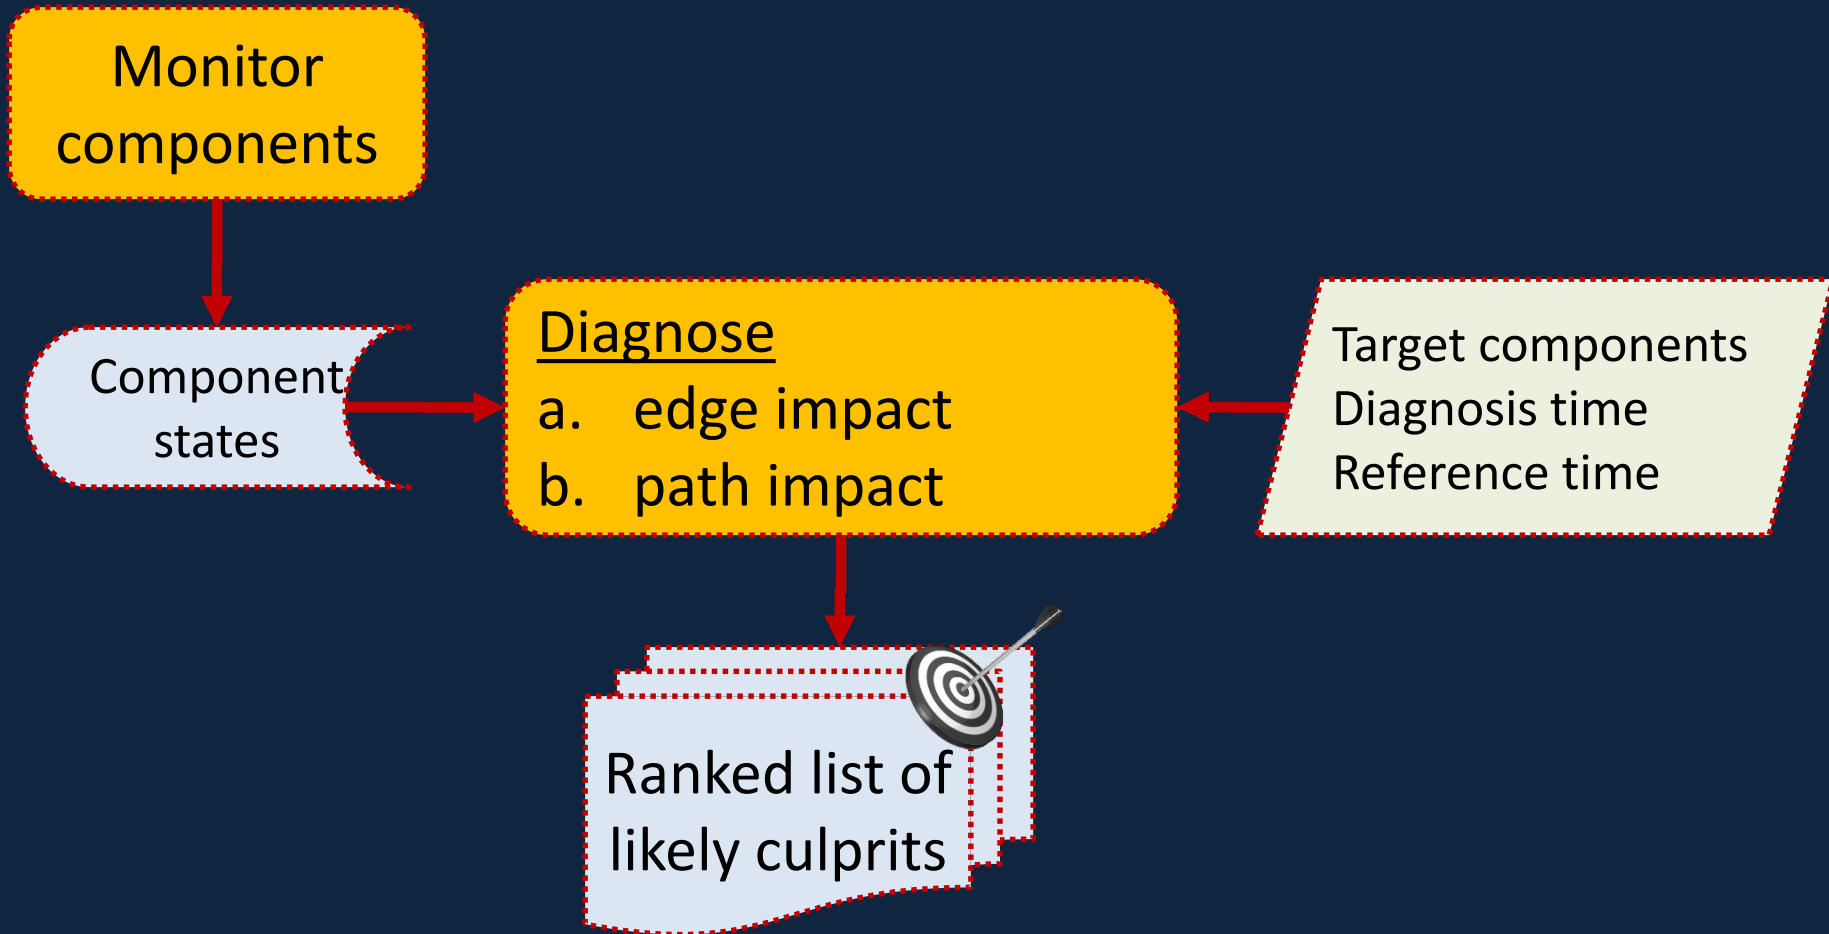
Path weight




Global impact



Implementation of NetMedic



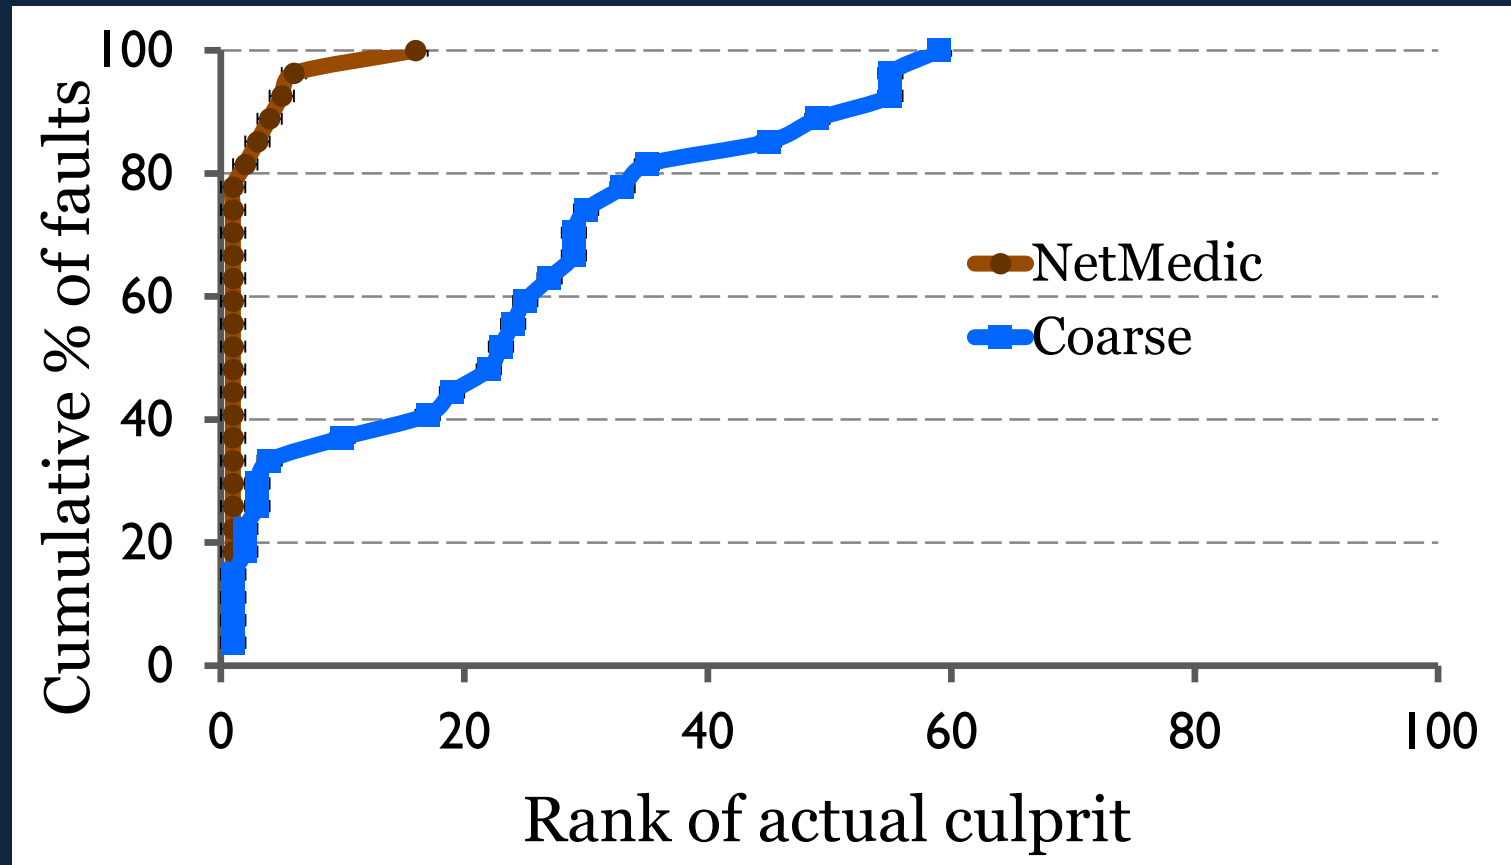
Evaluation setup




#components	~1000
#dimensions per component (avg)	35

Diverse set of faults observed in the logs

NetMedic assigns low ranks to actual culprits



NetMedic handles concurrent faults well



2 simultaneous faults

Other empirical results

Netmedic needs a modest amount (~60 mins) of history

The key to effectiveness is correctly identifying many low impact edges

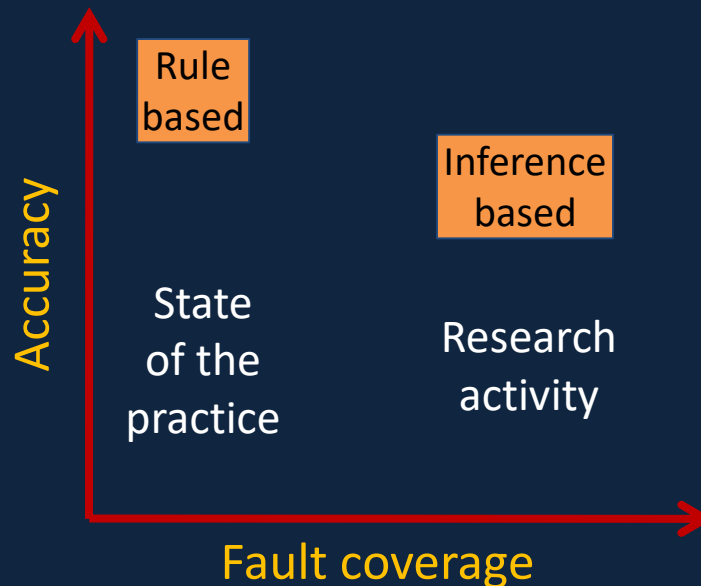
It compares favorably with a method that understands variable semantics

Unleashing (systems like) NetMedic on admins

How to present the analysis results?

- Need human verification

(Fundamental?) trade-off between coverage and accuracy



The understandability challenge

Admins should be able to verify the correctness of the analysis

- Identify culprits themselves if analysis is incorrect

Two sub-problems at the intersection with HCI

- Visualizing complex analysis (NetClinic)
- Intuitiveness of analysis (ongoing work)

NetClinic: Visualizing diagnostic analysis

Underlying assumption: Admins can verify analysis if information is presented appropriately

- They have expert, out-of-band information

Views diagnosis as multi-level analysis

Makes results at all levels accessible on top of a semantic graph layout


Allows top-down and bottom-up navigation across levels while retaining context

Load Data

Network View

Show Machine Labels Show Outgoing Edges on Mouseover Edge Visibility Sort Processes by Abnormality Find: Go

- Selected
- Neighbors of Selected
- In Path
- Focus in Path



Diagnoses

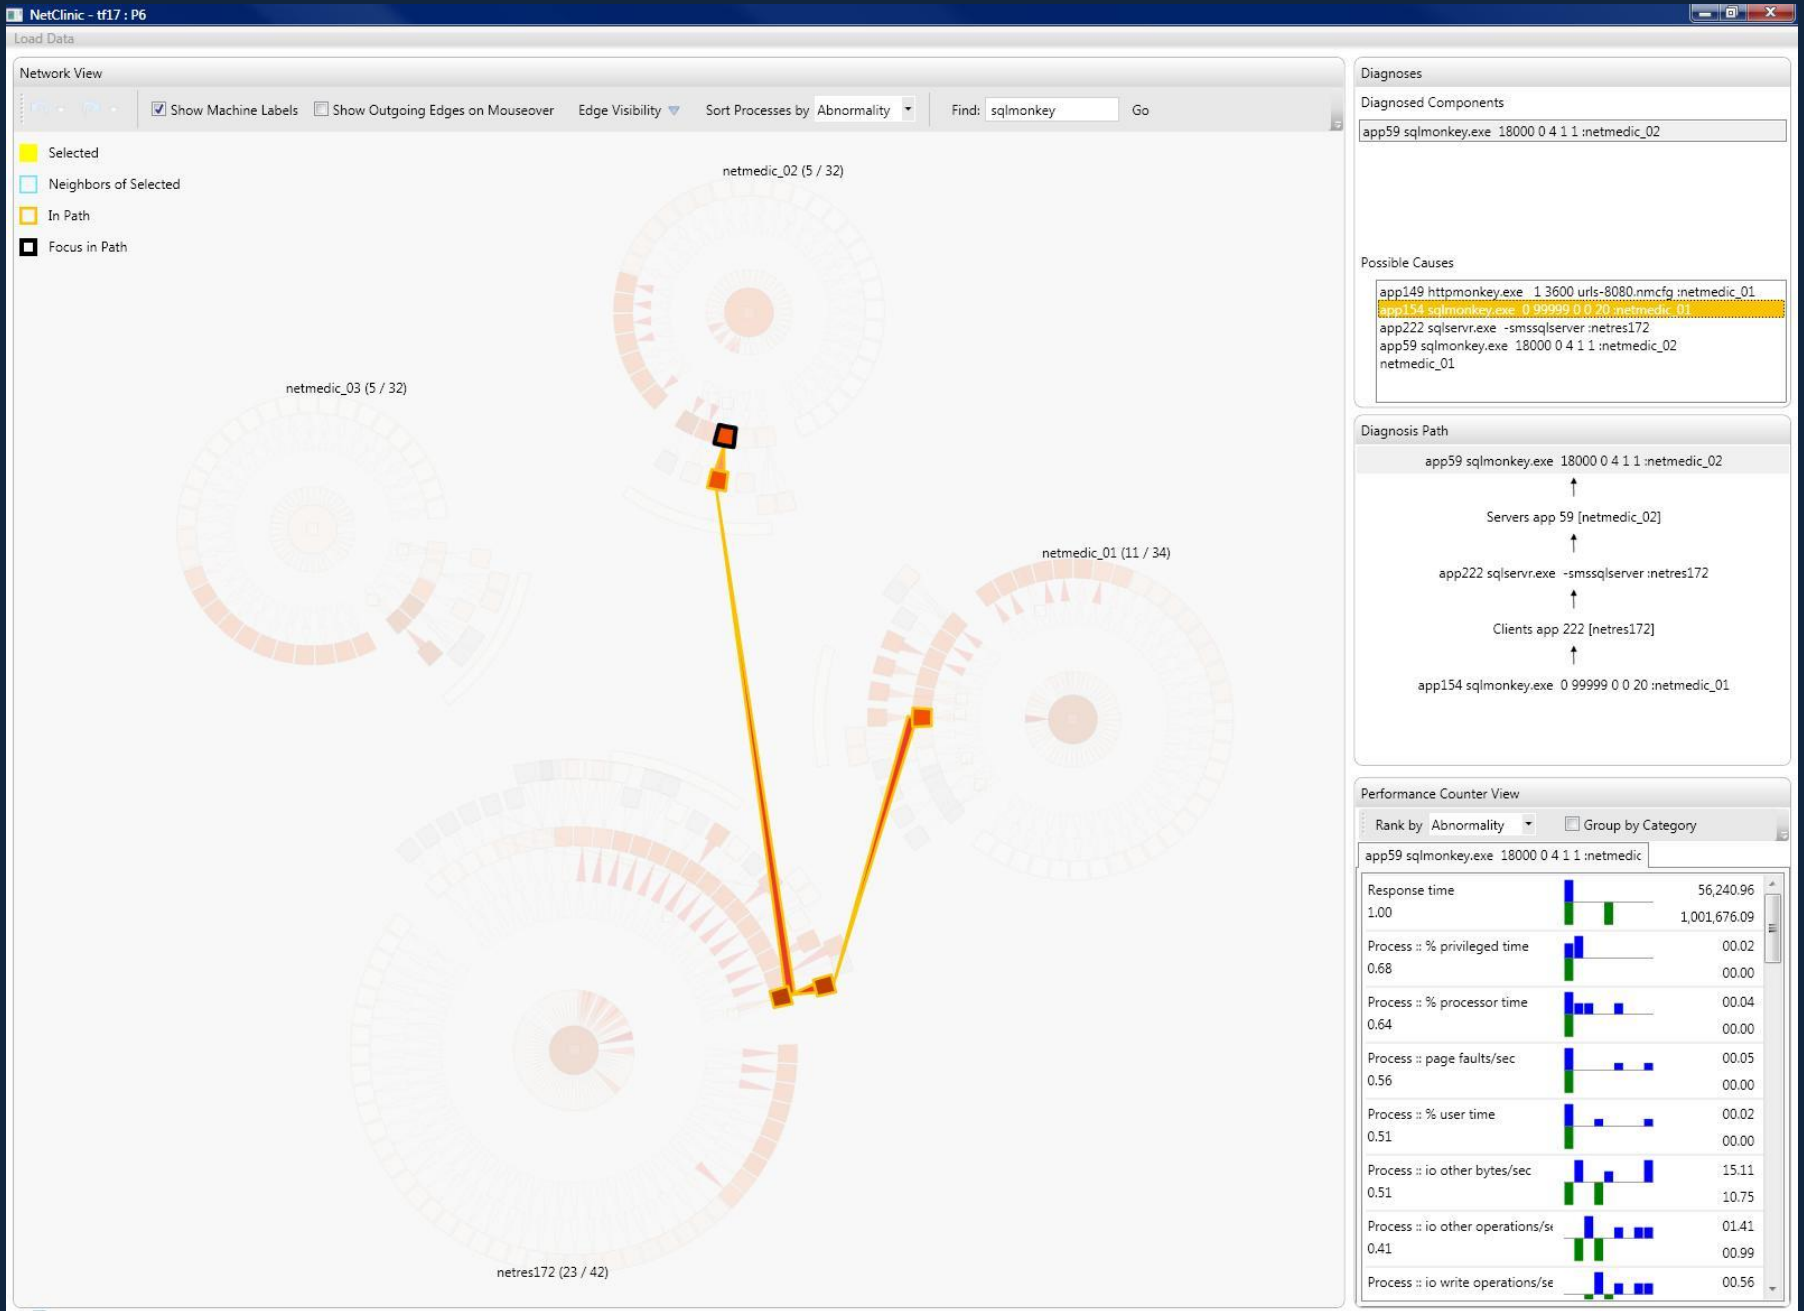
Diagnosed Components

Possible Causes

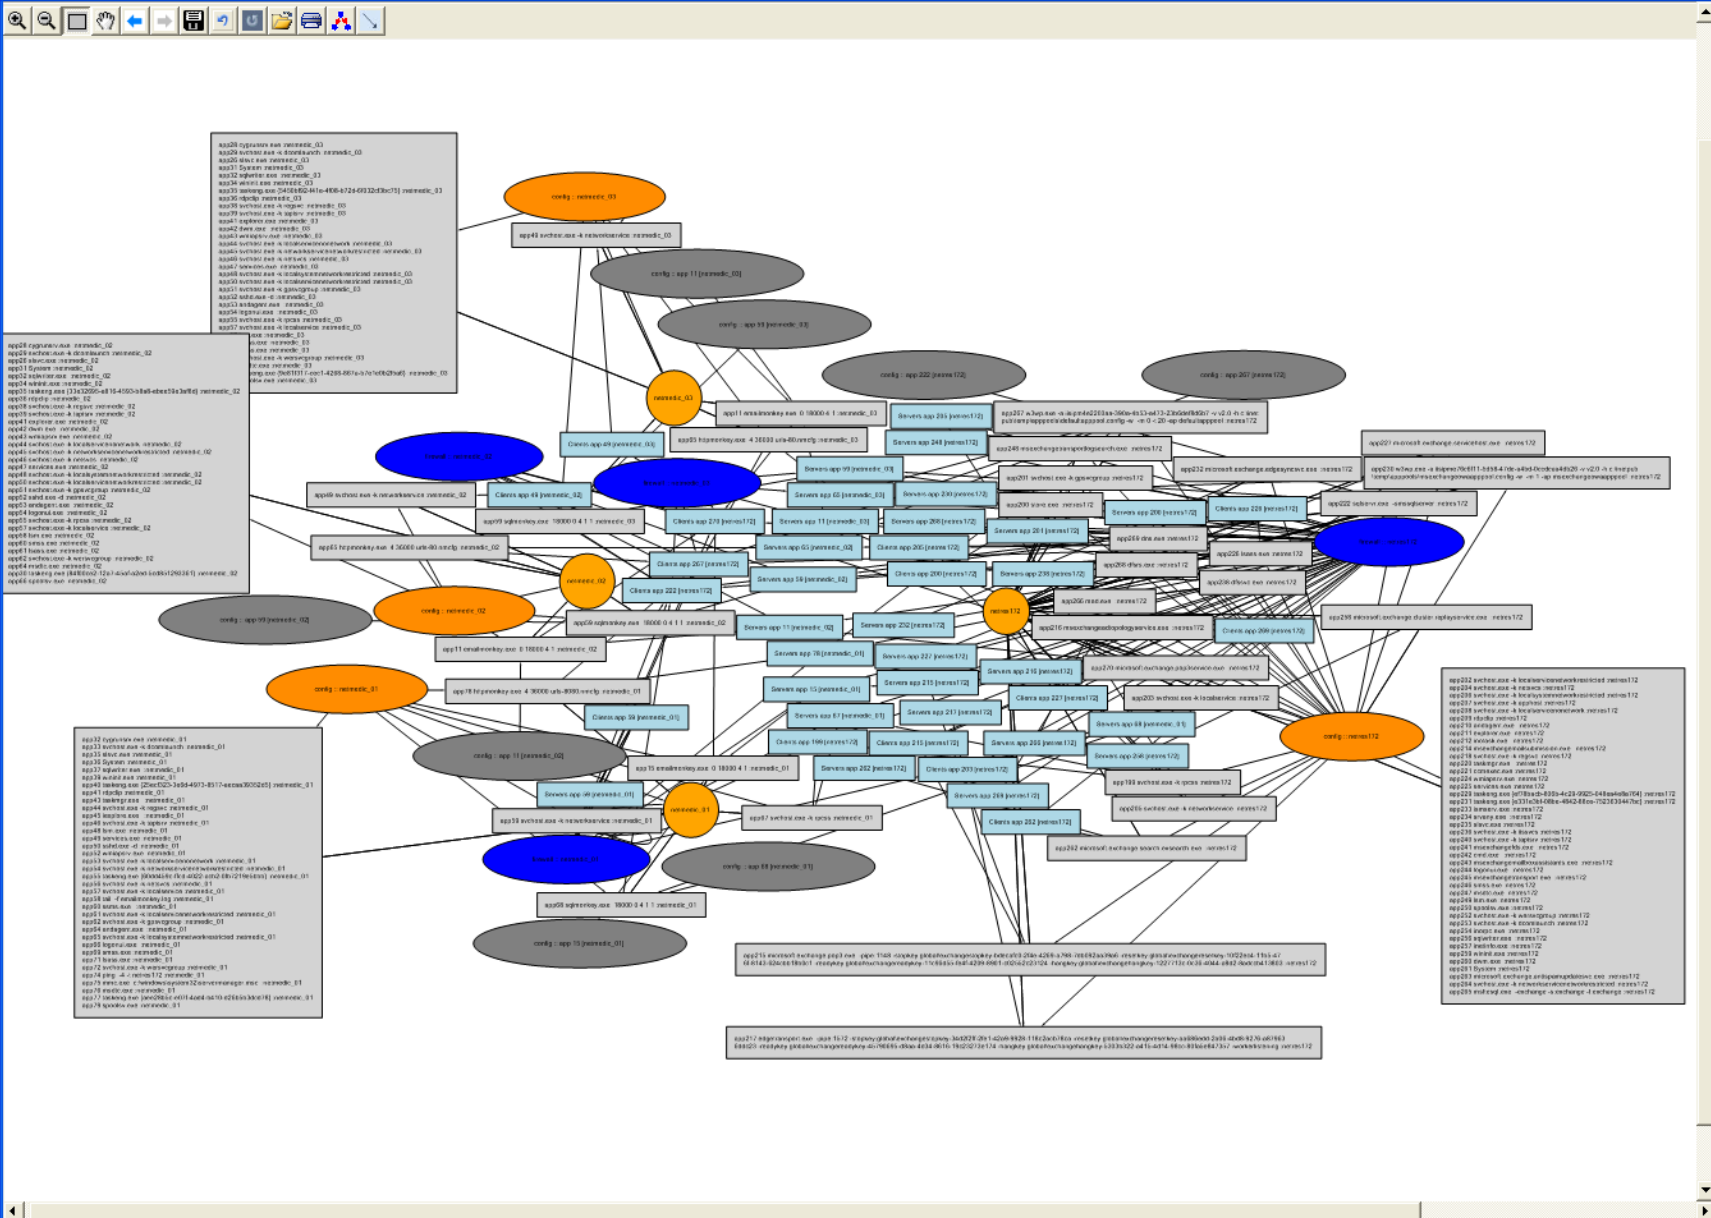
Diagnosis Path

Performance Counter View

Rank by Abnormality Group by Category



Dependency graph



NetClinic user study

11 participants with knowledge of computer networks but not of *NetMedic*

Given 3 diagnostic tasks each after training

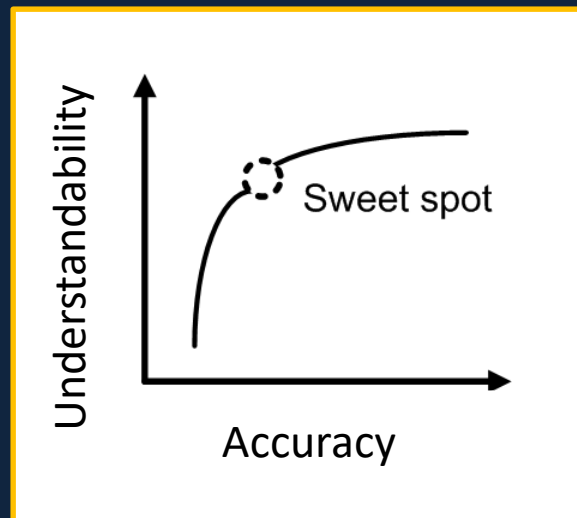
- 88% task completion rate

Uncovered a rich mix of user strategies that the visualization must support

Intuitiveness of analysis

What if you could modify the analysis itself to make it more accessible to humans?

- Counters the tendency to “optimize” for incremental gains in accuracy



Intuitiveness of analysis (2)

Goal: Go from mechanical measures to more human centric measures


- Example: MoS measure for VoIP

Factors to consider

- What information is used? E.g., Local vs. global
- What operations are used? E.g., Arithmetic vs. geometric means

Conclusions

Thinking small (networks) can provide new perspectives



NetMedic enables detailed diagnosis in enterprise networks w/o application knowledge

NetClinic enables admins to understand and verify complex diagnostic analyses