


NetPrints: Diagnosing Home Network Misconfigurations using Shared Knowledge

Venkat Padmanabhan

*Joint work with: Bhavish Aggarwal, Ranjita Bhagwan,
Tathagata Das, Siddharth Eswaran (intern, IIT Delhi),
Geoff Voelker (visiting researcher, UC San Diego)*

(To appear at NSDI 2009)

Typical Home Network



User Applications:

- IM, VoIP
- Email
- File Transfer (MSN Messenger/ FTP)
- Remote Desktop Connection (RDC)
- Web Hosting
- File and Printer Sharing
- SSH
- VPN

Network Components:

- User Desktops/Laptops
- Host-based Firewall
- NAT router/WiFi Access point
- Modem

No network admin!

Examples of Problems

Problem	Solution
VPN client does not connect from home	Turn on PPTP passthrough on router, use a subnet that is either 192.168.0.x or 192.168.1.x
XBOX doesn't connect to the Live service	Turn up your MTU above 1365, change NAT settings to full-cone, turn on UPnP
My IM client doesn't work from home	Turn off the DNS proxy on the router
File sharing doesn't seem to work at home	Make sure you and the file server are on the same domain/workgroup
Printing doesn't work from my laptop	Turn on correct firewall rules on the server machine
Cannot send large emails	Turn down MTU on your router

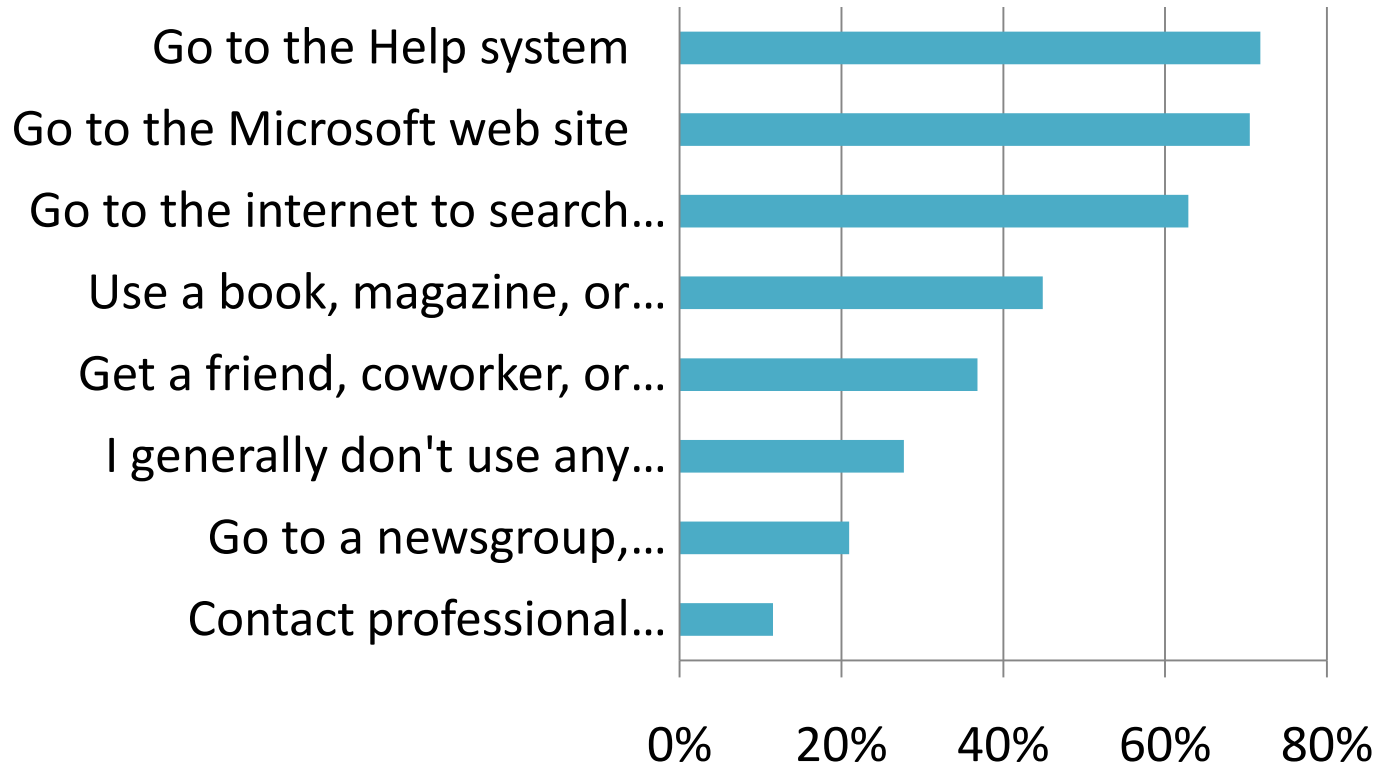
Router changes

End-host changes

Remote problem, local changes

Diversity \Rightarrow home network troubleshooting is hard

What Do Users Do Today?



Reference: Learning Orientation Survey, October 2004)

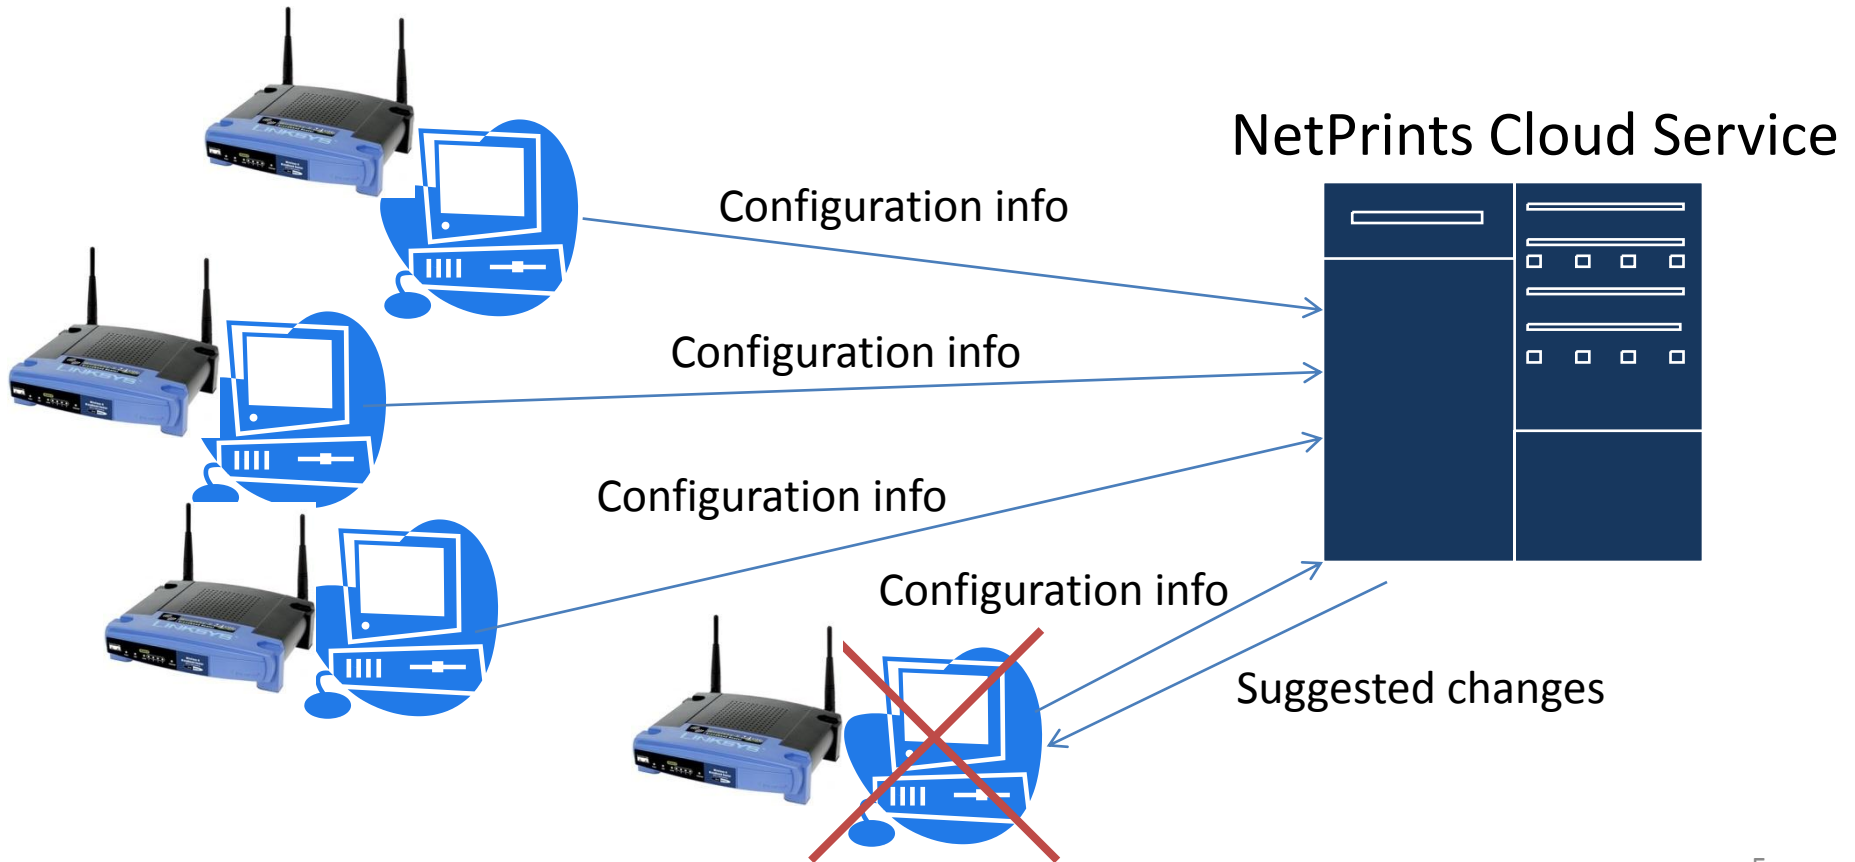
Only 55% people had their problems resolved!

Top 20 OEMs: \$350m/year on warrantys!

NetPrints

NetPrints = Network Problem Fingerprinting

Automate problem diagnosis using “*shared knowledge*”



Why Use Shared Knowledge?

Rule-based techniques Learning-based techniques

Windows Diagnostics Framework ~~Sk~~ rider+PeerPressure

Network Magic Indexing/retrieving System History

Apple's Diagnostics

Autobash


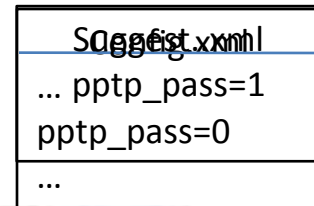

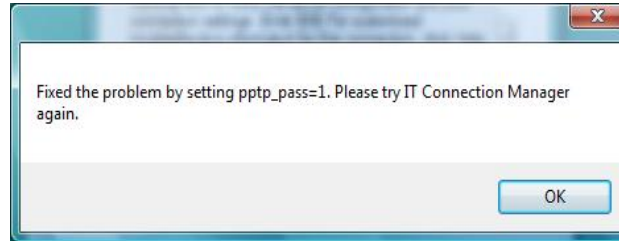
*Rule-based diagnosis Statistical/history-based diagnosis
(for basic connectivity issues)*

NetPrints

- Unstructured, heterogeneous environment
- Distributed configuration information
- Problems due to interaction of multiple configuration parameters

*Focus on application-specific
network connectivity problems*

NetPrints in Action

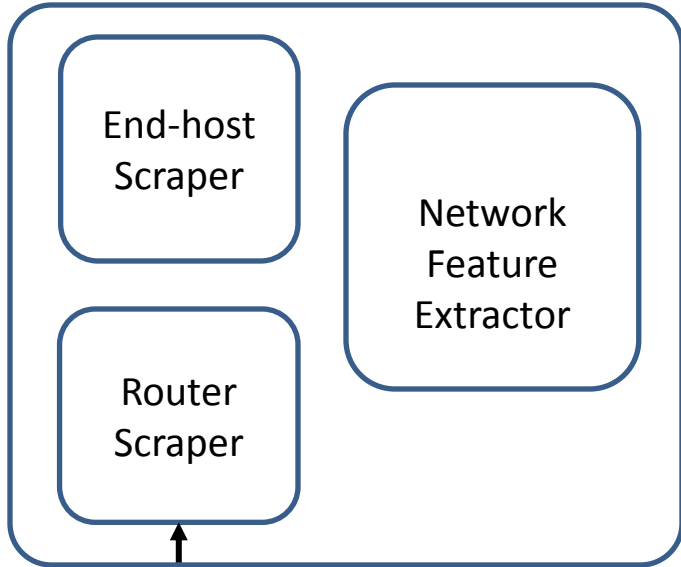


Training data for
Linksys WRT54G
router



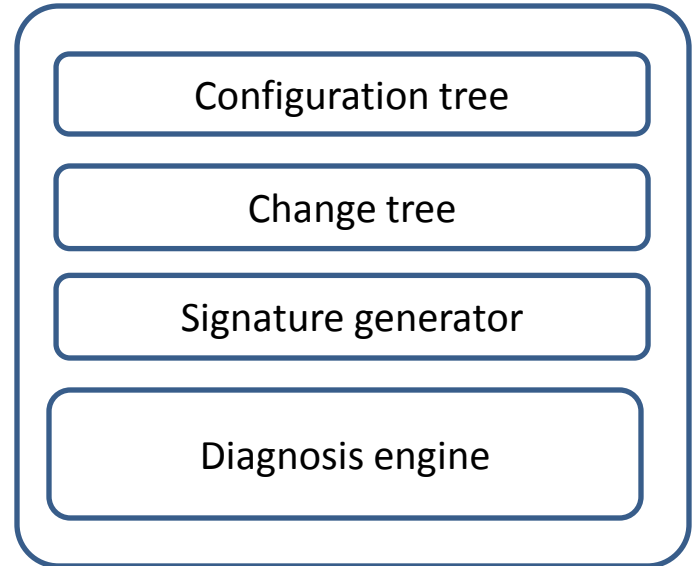


System Design

NetPrints Client



NetPrints Server



Normal Mode

Periodically, *scrape configuration* and network features.

Occasionally, prompt user for success/failure info.

Send labeled configuration to the server.

Server trains its *knowledgebase* with this data.

Diagnosis Mode

User identifies the problematic application.





Scrape configuration and network features for the problematic application. Send to server.



Server uses the knowledgebase to suggest changes using ***configuration mutation*** algorithm.

#1: Configuration Scraper

Problem	Solution
VPN client does not connect from home	Turn on PPTP passthrough on router, use a subnet that is either 192.168.0.x or 192.168.1.x
XBOX doesn't connect to the Live service	Turn up your MTU above 1365, change NAT settings to full-cone, turn on UPnP
My IM client doesn't work from home	Turn off the DNS proxy on the router
File sharing doesn't seem to work at home	Make sure you and the file server are on the same domain/workgroup
Printing doesn't work from my laptop	Turn on correct firewall rules on the server machine
Cannot send large emails	Turn down MTU on your router



Configuration Scraper

- Router scraper
 - UPnP
 - Web Interface (HTTP Request Hijacking)
- End-host scraper: local host & remote host
 - Interface-specific parameters
 - Patches and software versions
 - Firewall rules
- Composition of local and remote configs

Composing Local & Remote Configs

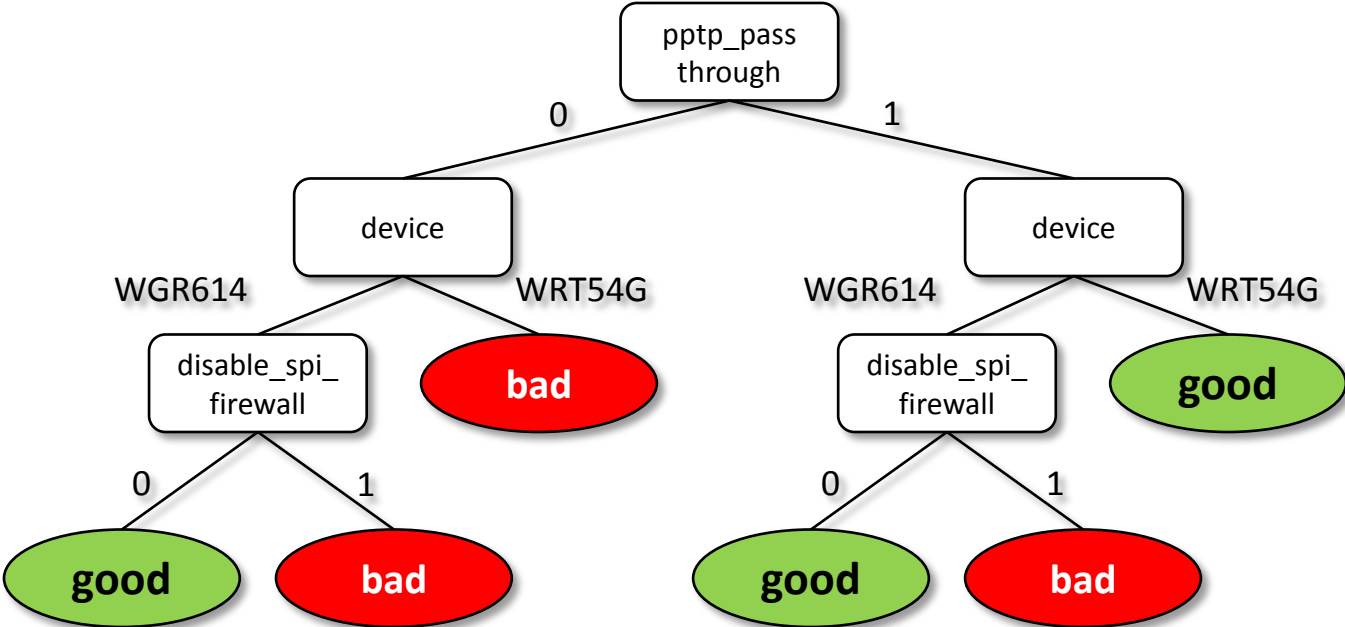
Problem	Solution
VPN client does not connect from home	Turn on PPTP passthrough on router, use a subnet that is either 192.168.0.x or 192.168.1.x
XBOX doesn't connect to the Live service	Turn up your MTU above 1365, change NAT settings to full-cone, turn on UPnP
My IM client doesn't work from home	Turn off the DNS proxy on the router
File sharing doesn't seem to work at home	Make sure you and the file server are on the same domain/workgroup.
Printing doesn't work from my laptop	Turn on correct firewall rules on print server machine
Cannot send large emails	Turn down MTU on your router

Sometimes it is the *combination* of local and remote configs that is the problem

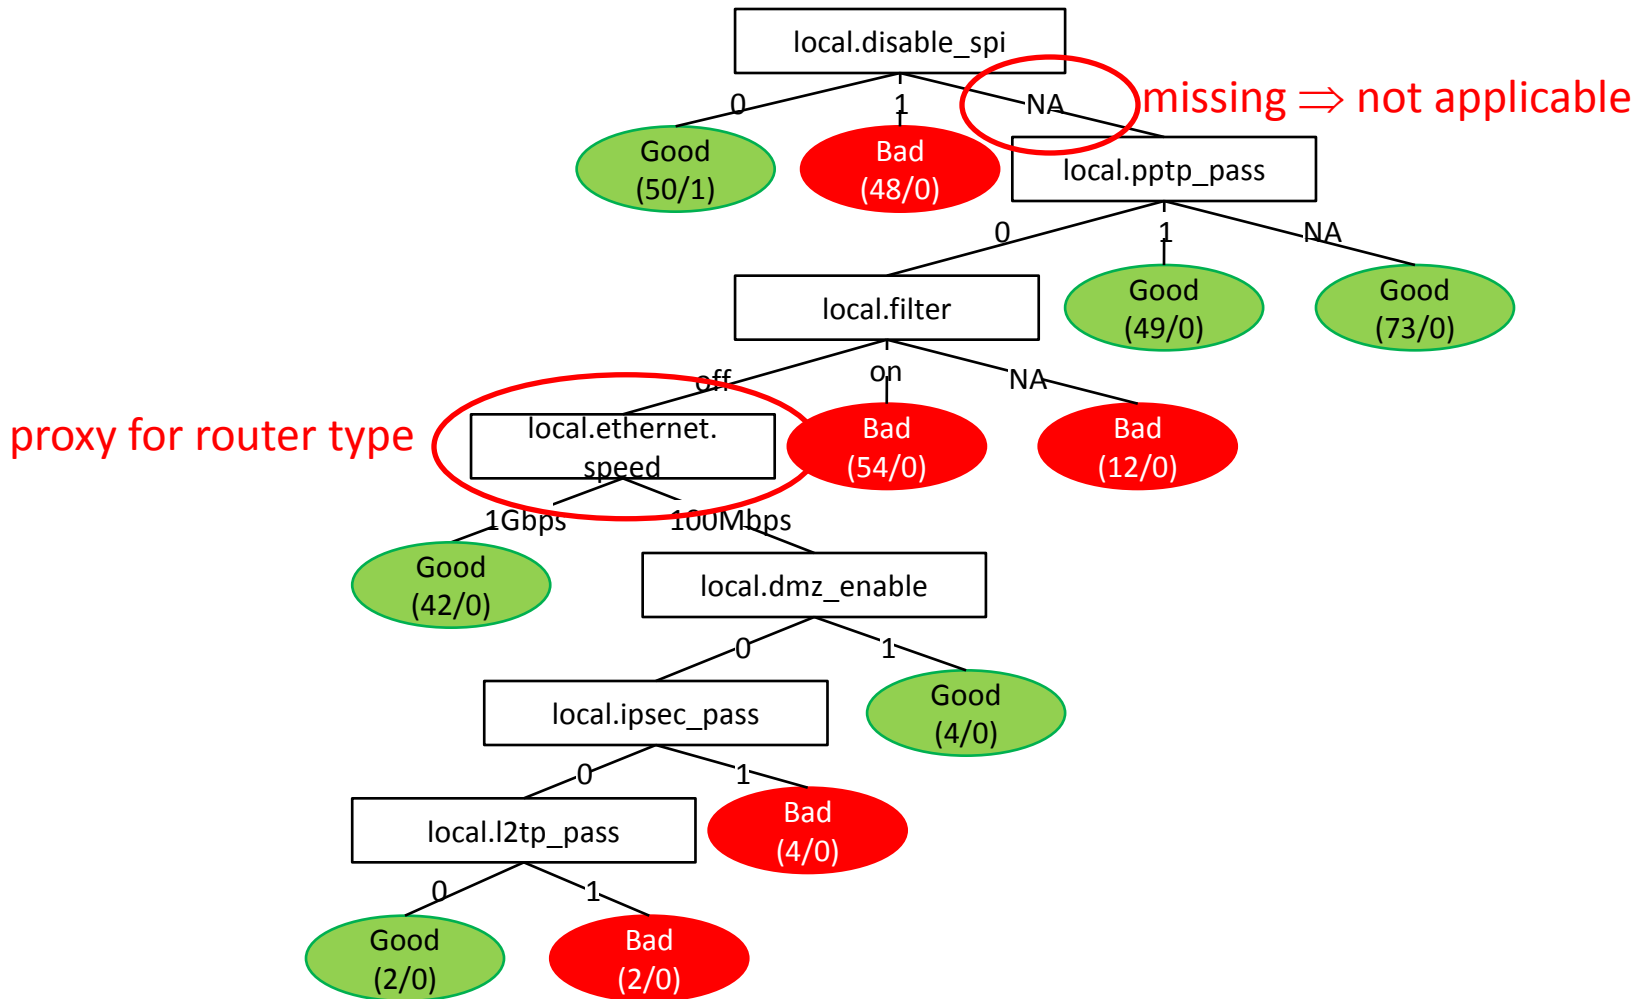
#2: Server Knowledgebase

- Per-application decision trees constructed using labeled configuration snapshots
 - decision trees aid interpretability
 - C4.5 decision tree learning algorithm
- Configuration tree
 - based on static snapshots of “good” and “bad” configs
- Change trees
 - based on change from “bad” config to “good” config

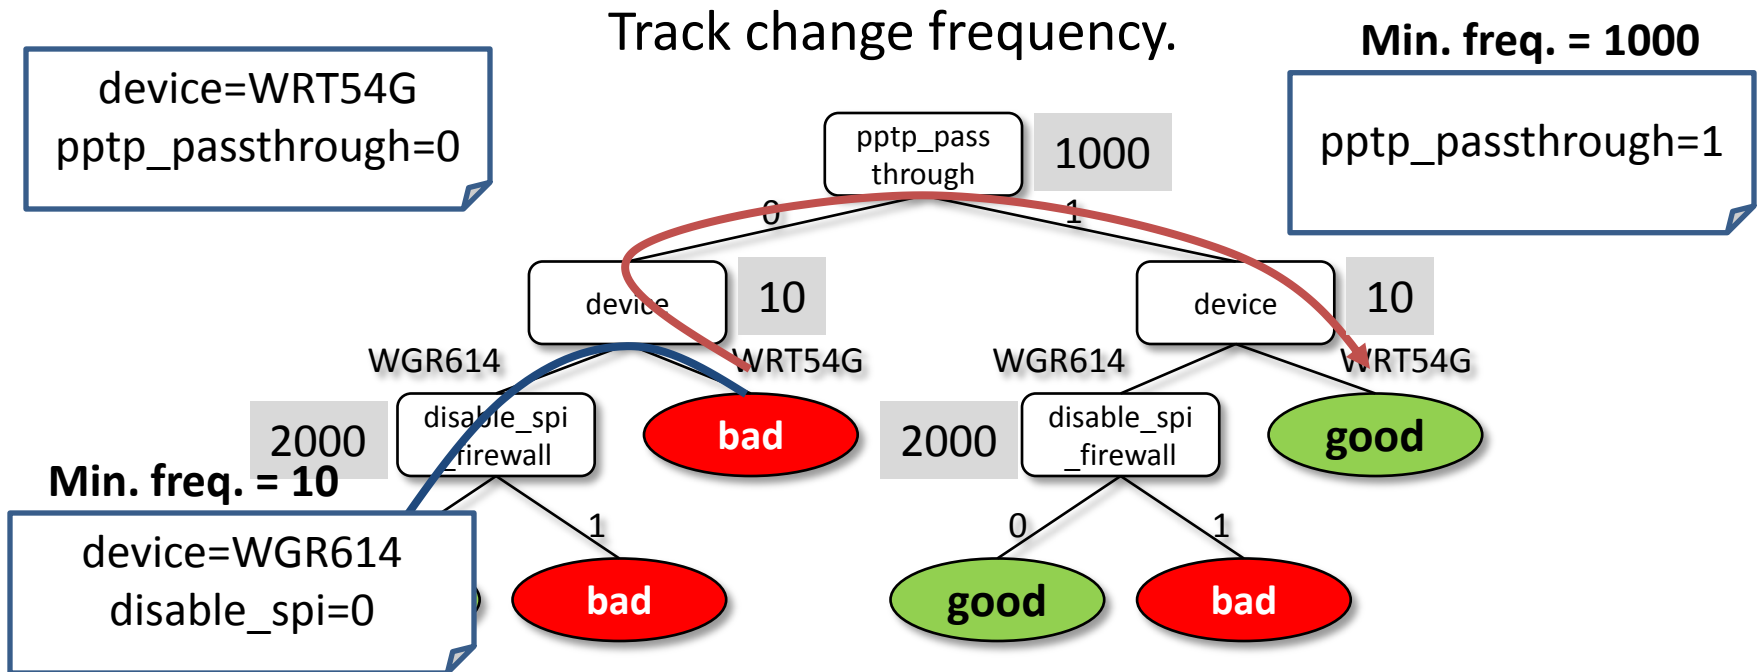
Example of Configuration Tree



Automatically Generated Tree



#3: Configuration mutation



- Preference for mutations involving frequently changing parameters
- Assumption: higher the frequency, less disruptive the change
- E.g., easier to change parameter setting than to change router

Shortcoming of Configuration Trees

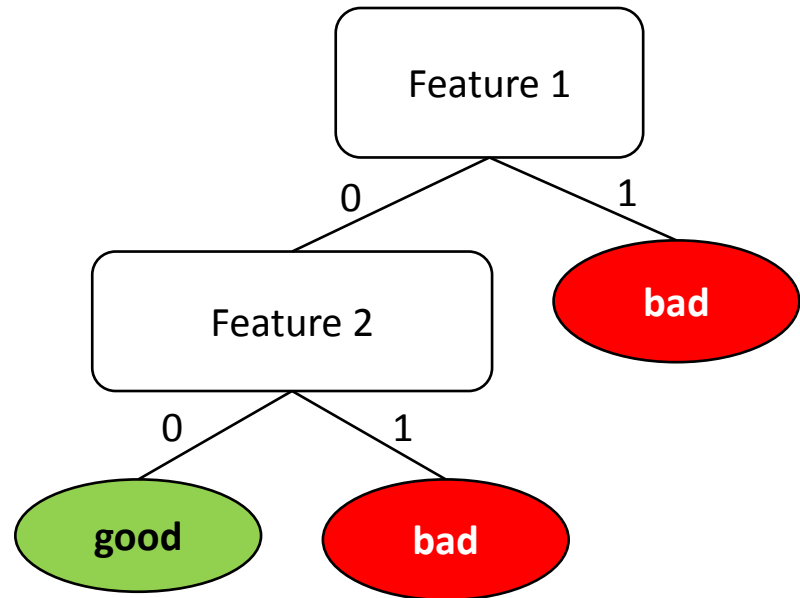
- Some config info may not be learned
- So traversal of config tree may end in a “good” leaf even if config is problematic
- Reasons:
 - **Insufficient data**
 - e.g., a new router enters the market
 - **Hidden configurations**
 - e.g., remote FTP server only performs active FTP (fix: turn on firewall rule on FTP client)

Change Trees

- Only consider configs for which the config tree traversal ends in a “good” leaf
- If config change results in a resolution of the problem, use the before config (“bad”) and the after config (“good”) to learn change tree
- Change trees
 - attached to “good” leaves of config tree
 - indexed by network traffic signature

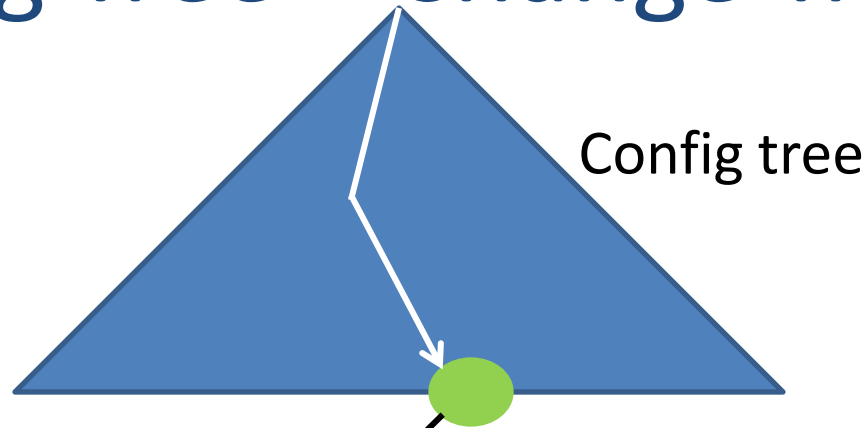
Network Problem Signature

- Apps can fail in various ways
- Signatures help tell different failure modes apart
- Constructed based on features extracted from app's network traffic



	Feature Description	Evaluation Type
1	TCP: Three SYN no response	Per-connection
2	TCP:RST after SYN	Per-connection
3	TCP:RST after no activity for 2 minutes	Per-connection
4	TCP:RST after some data exchanged	Per-connection
5	UDP: Data sent but not received	Per-four-tuple
6	Other: Data sent but not received	Per-IP address pair
7	All: No data sent or received	Overall

Config Tree + Change Trees



Problem Signature	Change trees
1XXXXXX	
0XXX X1X	

Summary of Diagnosis Procedure

- Look for solution in configuration tree for the application
- If traversal ends in “good” leaf, look up change tree for the given network signature
- If solution found, report to the user, and automatically fix problem wherever applicable
- If no solution found, give up.

Experimental Evaluation

- Testbed comprising 7 different routers
 - various makes: Netgear, Linksys, D-Link, Belkin
- Experiments with 4 applications
 - FTP client, VPN client, file sharing, Xbox wireless
- Robustness to mislabeling
 - 13-17% mislabeling \Rightarrow 1% error in diagnosis
 - tolerant to skew in configuration diversity

Summary

- Home network diagnostics is challenging
 - diversity of apps and configs
 - absence of an admin
- NetPrints leverages info community info to perform *automated* diagnosis
 - decision tree based learning
- More info: NSDI 2009 paper to be available at www.research.microsoft.com/research/mns/