

Understanding End-User Perception of Network Problems

J. Scott Miller Amit Mondal Rahul Potharaju
Peter A. Dinda Aleksander Kuzmanovic

Empathic Systems Project
empathicsystems.org

Northwestern University

Overview

- **Overall Question:** How does **end-user perception** of network performance correlate with low-level network measurements?
- **Approach:** User study that collects network, host, and user measurements on end-user machines, *labeled by end-user with “irritation events”*; comparison of labeled/unlabeled measurements
- **Context:** How to incorporate end-user satisfaction and guidance in computer systems, computer architecture (and network?) design?


SoylentLogger

- Windows service that monitors network, host, and user context and uploads to our server
 - Negligible network, cpu, and memory overheads on laptops
 - Packet-level inspection, connections tagged with applications
 - Periodic measurement, and irritation-driven measurement
- Measurements of
 - **User:** application focus, user activity, web traffic (URLs),
 - **Host:** CPU utilization, process statistics
 - **Network from the perspective of the host:** offered throughput, application RTT, receiver signaling duplicate packets, link properties, wireless interface properties, ping/traceroute probes

User Irritation

Press F8 when irritated
by the network ↓

Press F8 when irritated
by the network ↓



- User prep document focuses on network performance and states “We ask that you press (the irritation button) when you are uncomfortable or dissatisfied with the network service being provided to the applications you are using.”


User Study

- 32 users recruited using broad, IRB-approved advertising at Northwestern
 - Almost all ***non-technical users***
 - \$25 for participation
- SoyLentLogger installed and tested by us on each user's personal machine
- Controlled interaction with users
 - Users read standard preparatory document
 - Users told to use their machines normally
- Operation over same three week period

Data Set

- Immediately consecutive irritation events filtered
 - User mashing F8 repeatedly counts once
- Only data within 60 s of user activity considered
 - User must have opportunity to express irritation
- 20 GB of raw data
- 899 irritation events
 - ~1.2 events/user/day (varies across users, does not vary much across time)
 - Apparent power law interarrival times per user
 - 50% of irritation events occur within 17 minutes of a previous event

What does an irritation event label?



ω : window of experience leading to event

τ : delay from experience to keypress

We evaluate sensitivity of results to ω and τ

Most sensitivity is to ω

Hypotheses Evaluated

- **Supported**

- Users can distinguish between local and network sources of irritation
- User irritation is dependent on the applications and services with which the user interacts
- User irritation is stateful
- User irritation is affected by user location (wireless access point)

- **Supported with other observations**

- Most irritation is associated with small flows

- **Not supported**

- Users are more sensitive to the network when using streaming applications
- RSSI and link quality indicators predict user irritation on wireless networks

Hypotheses Evaluated

- **Supported**

- Users can distinguish between local and network sources of irritation
- User irritation is dependent on the applications and services with which the user interacts
- User irritation is stateful
- User irritation is affected by user location (wireless access point)


- **Supported with other observations**

- Most irritation is associated with small flows

- **Not supported**

- Users are more sensitive to the network when using streaming applications
- RSSI and link quality indicators predict user irritation on wireless networks


Do Applications Matter?




Do WiFi Signal Quality Metrics Help?

- Not supported by our evidence
- NIC RSSI metric seems completely uncorrelated with user irritation events
- Windows signal quality metric is slightly correlated with user irritation events

But Windows Signal Quality Is A Poor Predictor...



Yet Location (which WAP is used) Matters



Conclusions

- Attempt to correlate user irritation due to perceived network problems with low level network measurement
 - Using a feedback mechanism that could be continuously employed
 - Contra QoE, OneClick, EmNet, Vienna Surfing, HostView, LRD,...
- More detailed technical report and study materials available on line
- We are working on making data available

For More Information

- pdinda@northwestern.edu
- empathicsystems.org
- presciencelab.org

- Sponsored by NSF CNS-070691

Top Five ASes By Traffic Volume

Host	No Irritation		Irritation		Total Traffic (MB)	%Bytes in Irritation
	Traffic (MB)	Flows	Traffic (MB)	Flows		
Google	8402	85133	295	1376	8698	3.4
Comcast	6475	7084	<1	3	6475	0.01
NU	4242	88970	66	908	4308	1.53
Level 3	3988	18024	234	582	4222	5.54
Limelight	3155	14608	3	110	3157	0.08

Top Three ASes By Irritation

Host	No Irritation		Irritation		Total Traffic (MB)	%Bytes in Irritation
	Traffic (MB)	Flows	Traffic (MB)	Flows		
Advanced Video Commun.	767	3032	452	10	1219	37
Global Crossing	480	1325	240	19	721	33
NTT America	560	5379	246	45	805	31

5.1% of observed traffic, but 48.9% of all bytes associated with irritation

Are Users Particularly Sensitive During Streaming?

- Not supported by our evidence
- Irritation events during times when at least one streaming flow exists: 0.41/hour
- Irritation events when no streaming flow exists: 0.81/hour
- We tag flows as “streaming” based on size, port, “Googling the Internet” technique, and destination ASN
 - There are caveats

Is User Irritation Stateful?

