HostView: Measuring Internet quality of experience on end-hosts

Renata Teixeira
Muse Team – Inria Paris-Rocquencourt
with
D. Joumblatt (TPT), N. Taft (Google), J. Chandrashekar (Technicolor)
A.-K. Pietilainen, G. Rosca (Inria)
P. Tolmie, T. Rodden, T. Lodge (U. Nottingham)
R. Mortier (U. Cambridge)

Network performance disruptions are frustrating

For users

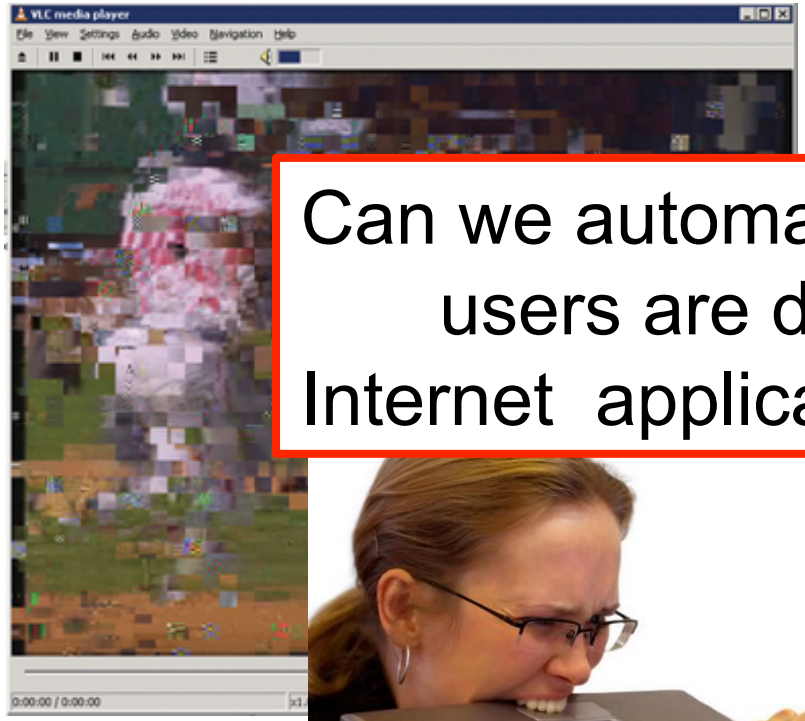

For ISPs

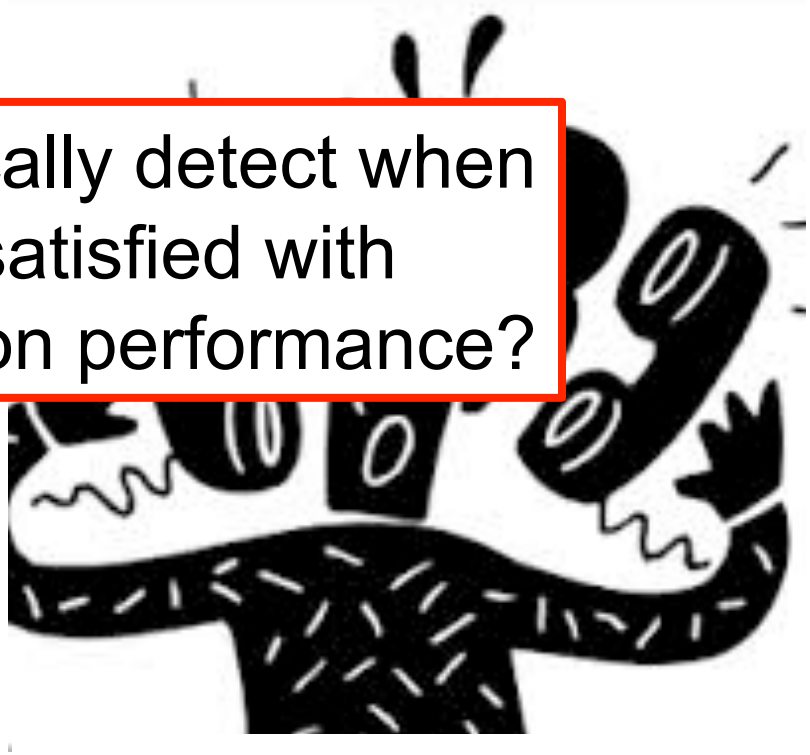

Can we automatically detect when users are dissatisfied with Internet application performance?

Our approach

- Learn user satisfaction models empirically
- HostView: end-host data collection tool
  - Measure network activity and performance annotated with user feedback

Challenges in measuring user perception

- User perception varies
  - Per user, per environment, per application
  - For a given user according to external factors
- Imbalance in number of samples
  - Can't collect frequent user feedback (~10 per day)
  - Orders of magnitude more network measurements (~10^3/10^4/10^5/…)
- End-host data collection raises issues
  - Privacy
  - Machine overhead

HostView data collection

- Network performance
  - Packet headers (anonymized IP source)
- User environment and system performance
  - CPU load, OS, time zone, country
- Application-level context
  - Content-type, referer for HTTP responses
  - Application (process names)
- User feedback
  - System-triggered questionnaires (3 times a day)
  - I'm annoyed button

User feedback mechanisms

- System Trigged feedback
  - Experience sampling methodology (ESM)
  - Triggered based on state of machine
  - 5 short questions about network performance
  - At most 3 times a day
- User Triggered feedback ☹
  - “I’m annoyed” button
  - Same questions as in ESM
  - Can trigger as often as user wants

Deployment

- Recruiting volunteers
  - Leaflets at IMC 2010 and CS Mailing lists
  - 50 USD Amazon gift cards
  - Real-time feedback about network connection
- Data: 40 users
  - Nov 2010 – Feb 2011
  - 26 Mac OS and 14 Linux
  - 14 countries
  - Most users ran tool for one month

Can we predict user dissatisfaction with network performance?

- Train predictors using HostView data
  - For some applications we achieve high accuracy
    - Non-linear SVM
    - Train predictors per application
- Challenges
  - User feedback is scarce
  - Each user feedback is limited to few apps
  - User feedback is unbalanced
  - Predictor may work for one app, but not others
  - Hard to identify user activity

New HostView

- Design informed by ethnographic study
  - 12 participants in France
  - Heavy-duty system/network monitoring
  - Participants filled diaries
  - Interviews using data visualization tool
- Improvements
  - Windows version (maybe Android)
  - Better capture user activity and context
  - Simpler questionnaires
  - Better incentives: Visualization of data

Example question

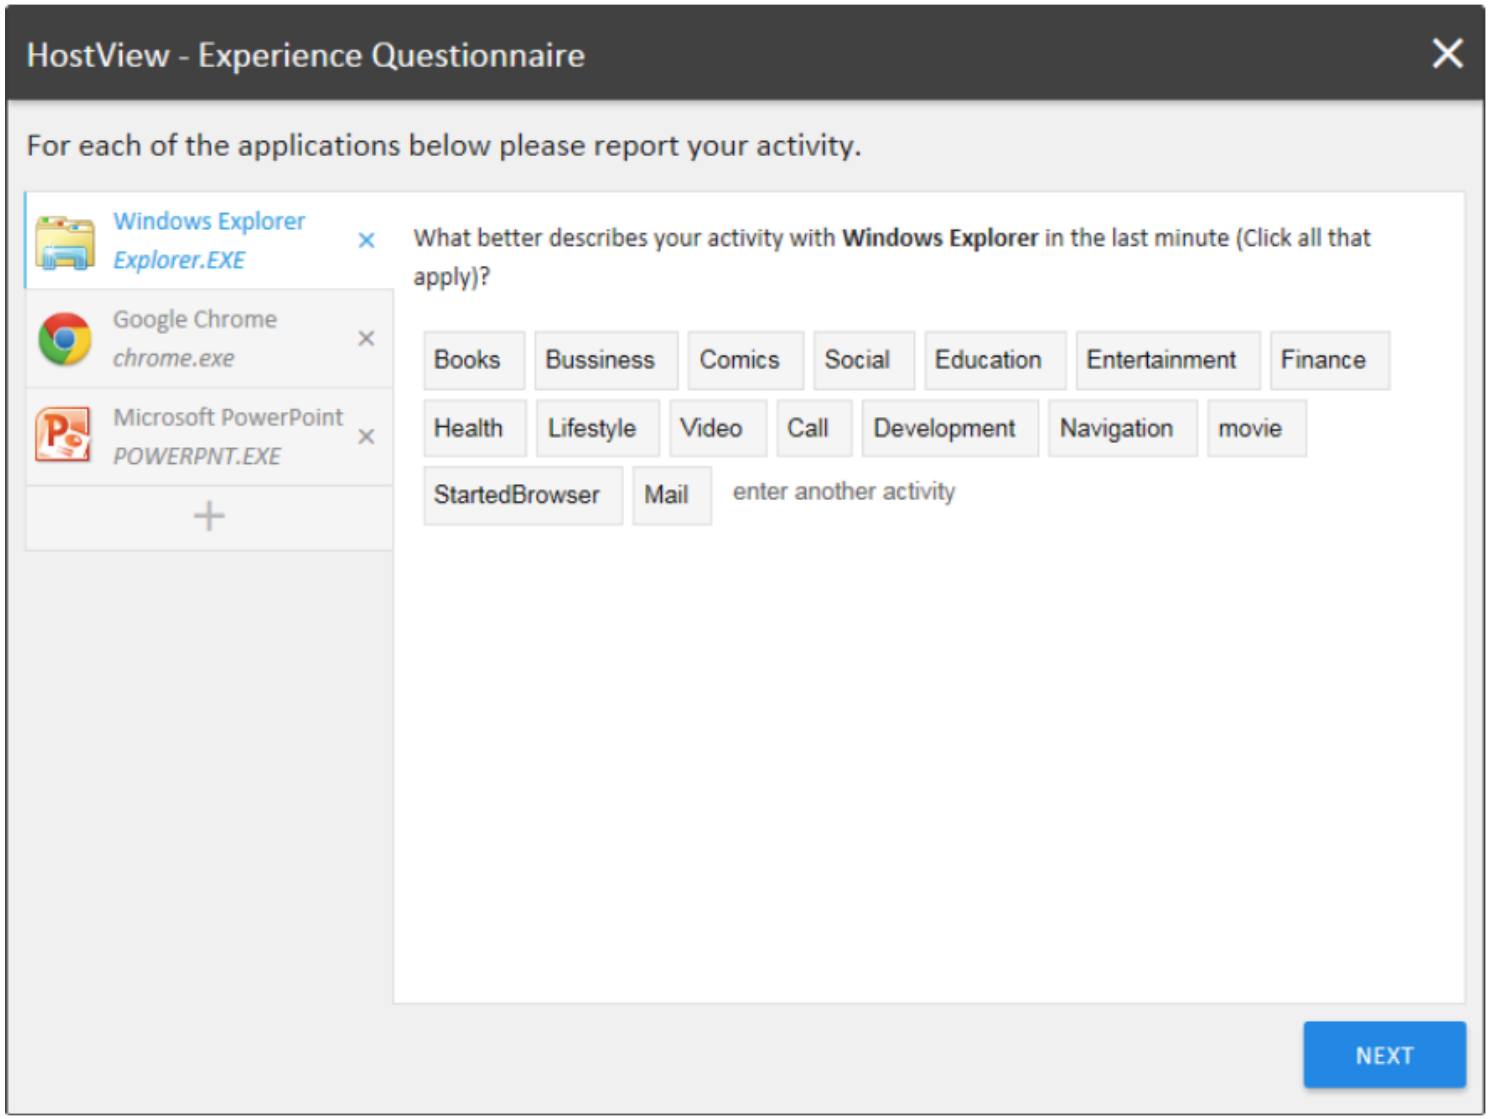

Questions?