

Step 1, do no harm.... Thoughts on intrusiveness in network measurement

Grenville Armitage

garmitage@swin.edu.au

Centre for Advanced Internet Architectures (CAIA)
Swinburne University of Technology



Outline



Do no harm

Observing & Sharing

Influencing while observing

Social expectations

Final thoughts

To measure is to meddle



“In the land of the blind, the one-eyed man is king.”
(Desiderius Erasmus)

To measure is to meddle



“In the land of the blind, the one-eyed man is king.”
(Desiderius Erasmus)

- We do evidence-based based research, yes?
 - Hypothesise →
 - Measure→
 - Know

- “belief without evidence” is unpalatable, so we *measure*



“In the land of the blind, the one-eyed man is king.”
(Desiderius Erasmus)

- We do evidence-based based research, yes?
 - Hypothesise →
 - Measure→
 - Know

- “belief without evidence” is unpalatable, so we *measure*

- *Someone* will find your measurements to be
 - intrusive...
 - unwelcome...
 - meddlesome...

The problem with “other people”



- Other people make the most interesting test subjects

(and yet)

- Capturing, recording and analysing what other people do (or own) is fraught with pitfalls



“Curiosity killed the cat”

- Taking measurements
 - ...may alter observed system or end-user behaviours
 - ...may violate social expectations



“Curiosity killed the cat”

- Taking measurements
 - ...may alter observed system or end-user behaviours
 - ...may violate social expectations

- Using measurements
 - ...may reveal unpopular or unauthorised insights



“Curiosity killed the cat”

- Taking measurements
 - ...may alter observed system or end-user behaviours
 - ...may violate social expectations
- Using measurements
 - ...may reveal unpopular or unauthorised insights
- Sharing measurements
 - ...wider analysis → more scientific credibility
 - ...increased risk of measurement data leakage

Step 1, do no harm



- Categories of consequences:
 - Operational (acquisition of accurate observations)
 - Social (clashing with other people's expectations)
 - Political (getting permission 1st, 2nd, Nth time....)
 - Legal (your plans might simply be illegal)



- Categories of consequences:
 - Operational (acquisition of accurate observations)
 - Social (clashing with other people's expectations)
 - Political (getting permission 1st, 2nd, Nth time....)
 - Legal (your plans might simply be illegal)

As evidence-based researchers we should carefully consider how to minimise these consequences

Outline



Do no harm

Observing & Sharing

Influencing while observing

Social expectations

Final thoughts



Measurement impacts the system under observation

- “physical” examples include:
 - Electrical measurements draw finite power
 - Network measurements require CPU time slices
 - Data logging requires I/O bus activity
 -

- “social” examples include:
 - Diluting the observed’s sense of privacy
 - The observed alter their behaviours
 -

(cf. Uncertainty principle: Observer effect at quantum levels)

Sharing: Benefits vs Chilling effects



- Observe then share → Virtual Observers

- Benefits
 - Observer effect is reduced (amortised over recipients)
 - “Many eyes” may discover new phenomena



- Observe then share → Virtual Observers
- Benefits
 - Observer effect is reduced (amortised over recipients)
 - “Many eyes” may discover new phenomena
- Risks
 - Potential for leaks is multiplied
 - Anonymisation not tested against attacks
- Chilling effects
 - Casual/careless de-anonymisation attempts
→ poison the future good-will of data owners

Outline



Do no harm

Observing & Sharing

Influencing while observing

Social expectations

Final thoughts



- Physical intrusion
 - passive fibre tap (miniscule energy draw)
 - active tap (regenerates tapped signal)
- Network port/interface monitoring
 - Mirror ports (an active tap at a higher layer)
 - Listening to WiFi transmissions (broadcast or otherwise)
- Instrumentation, logging system state
 - Packet timestamps, payloads, aggregate rates, instantaneous CPU loads, queue depths...

Requires local access to observation point, and we only learn what existing traffic patterns reveal

Active observations



- Direct engagement with observed system or network
- Probing
 - Scattered (background) probes – connectivity mapping
 - Sustained (bursty) probes – service characterisation
 - ...
- User-triggered testing (“drive-by testing”)
 - scripts embedded in web page, tickling wireless devices
 - ...

(Remote) traffic injection so we can observe system response to externally imposed conditions or stimuli



- Do your measurement activities look malicious?
 - False-positives in IDS or anti-DoS systems
 - Prelude to identity theft
 - Appearance of doing a vulnerability analysis
- Is your anonymisation secure?

Is it okay to *briefly* disrupt?



- “Characterizing Residential Broadband Networks”
<http://conferences.sigcomm.org/imc/2007/papers/imc137.pdf>

“At a high level, our technique is simple – we probe the broadband link with packet trains of different rates, using packets of various types and sizes. We use the responses received to infer a broad range of characteristics,”



- Carna Botnet, “Internet Census 2012”
<http://internetcensus2012.bitbucket.org/paper.html>
- Inject test code into vulnerable home gateways
- Run probe tests from 420K locations
- Restore home gateways to previous state

Stealth measurement – client browsers



- “Mitigating sampling error when measuring internet client IPv6 capabilities” (IMC 2012)
<http://www-net.cs.umass.edu/imc2012/papers/p87.pdf>
- Javascript/Flash embedded in web pages
- Web sites cooperate (Google ads)
- Browsers run tests from client locations



Do no harm

Observing & Sharing

Influencing while observing

Social expectations

Final thoughts

People are strange creatures



- People often desire
 - Agency
 - Capacity to act in the world around us
 - Ability to make choices
 - Awareness
 - Knowing what others know
 - Knowing what others are trying to know about us

(And often perception is reality)



- Measurements create new knowledge / insights
 - Yet to have agency / awareness suggests a need to control data about oneself

- The observed might
 - not want you to know more
 - want some control over what you learn
 - want some control over what you do with the knowledge

- Laws: a social response to express this control

Stealthy is fine, right?



- Is measurement okay if you don't disrupt:
 - the technical system/network?

 - the network/system operator's peace?

 - the observed's (perception of) agency or awareness?



*“Don’t poke a sleeping lion”
(a very wise person)*

- Laws about technical observations pre-date The Internet
- Rule #1 – don’t assume
- Rule #2 – get legal advice
- Different jurisdictions are.... different

Data sharing



- “You” need the help of two (overlapping?) communities
 - People whose behaviour is observed
 - People whose systems enable your observations
- Laws protect the former, the latter don’t want to get sued
- Avoid poisoning relationships
 - Data minimisation – how little do you need?
 - Data storage – how secure?
 - Data obfuscation – how secure?
 - Policies, procedures & agreements– transparent?



- Step 1: Run around the world taking photos
- Step 2: Decide public WLAN packets are fair game
- Step 3: (accidentally) capture & store (partial) payloads
- Step 4: Spend years trying to shake off mud
 - “As of 2012, investigations have gone forward in at least 12 countries, and at least 9 countries have found Google guilty of violating their laws.”
<http://epic.org/privacy/streetview/>
 - US wiretap case (Joffe v. Google) to proceed (Sep 2013)

Outline



Do no harm

Observing & Sharing

Influencing while observing

Social expectations

Final thoughts



- Respect the potential for legal complications
- Observe only what you need
- Secure your captured & shared data from leakage
- Nurture & protect relationships with network operators
- Avoid looking like an attacker or thief
- Minimise conflict with social expectations