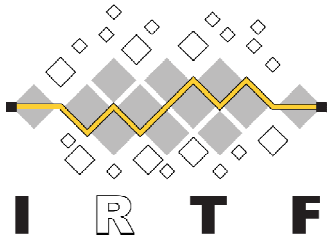

31st NMRG Meeting

Router-based and client-based platforms for performance measurement and censorship detection

Giuseppe Aceto, **Alessio Botta**, Walter de Donato,
Nick Feamster, Antonio Pescapé, Giorgio Ventre

University of Napoli Federico II, Italy
Georgia Institute of Technology, GA, USA

Introduction

- Large scale measurement platforms are necessary for studying residential Internet access

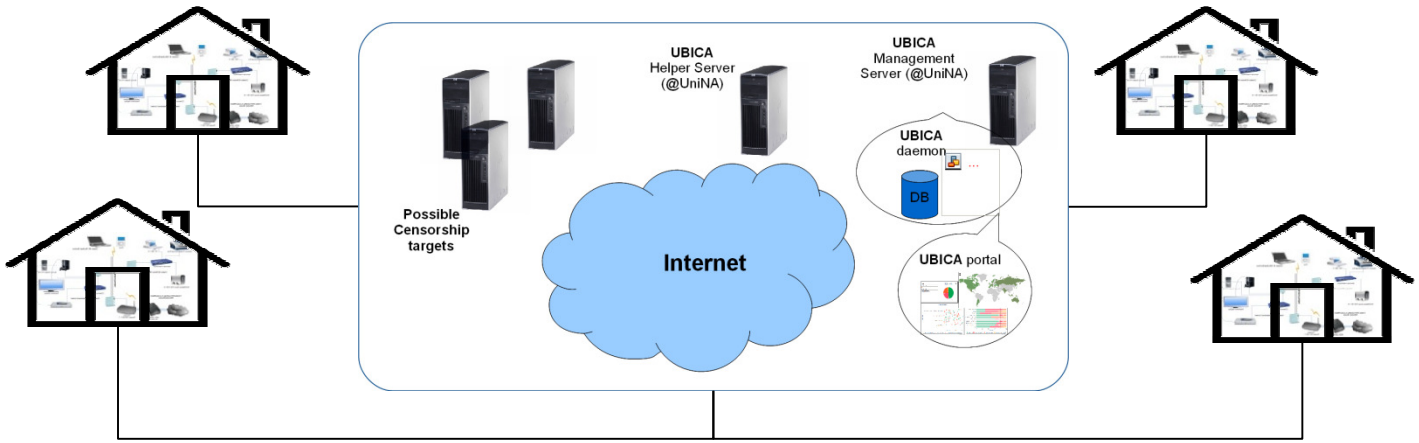

- Several approaches and platforms have been presented
- Standardization effort is ongoing while deployed platforms are not yet interoperable

Studying residential Internet access

Taxonomy of approaches

- Server-based
  - ISP-based
  - Core-based
  - Service-based
- Router-based
- Host-based
  - Client-based
  - Web-based
    - Plugin-based
    - Embedded

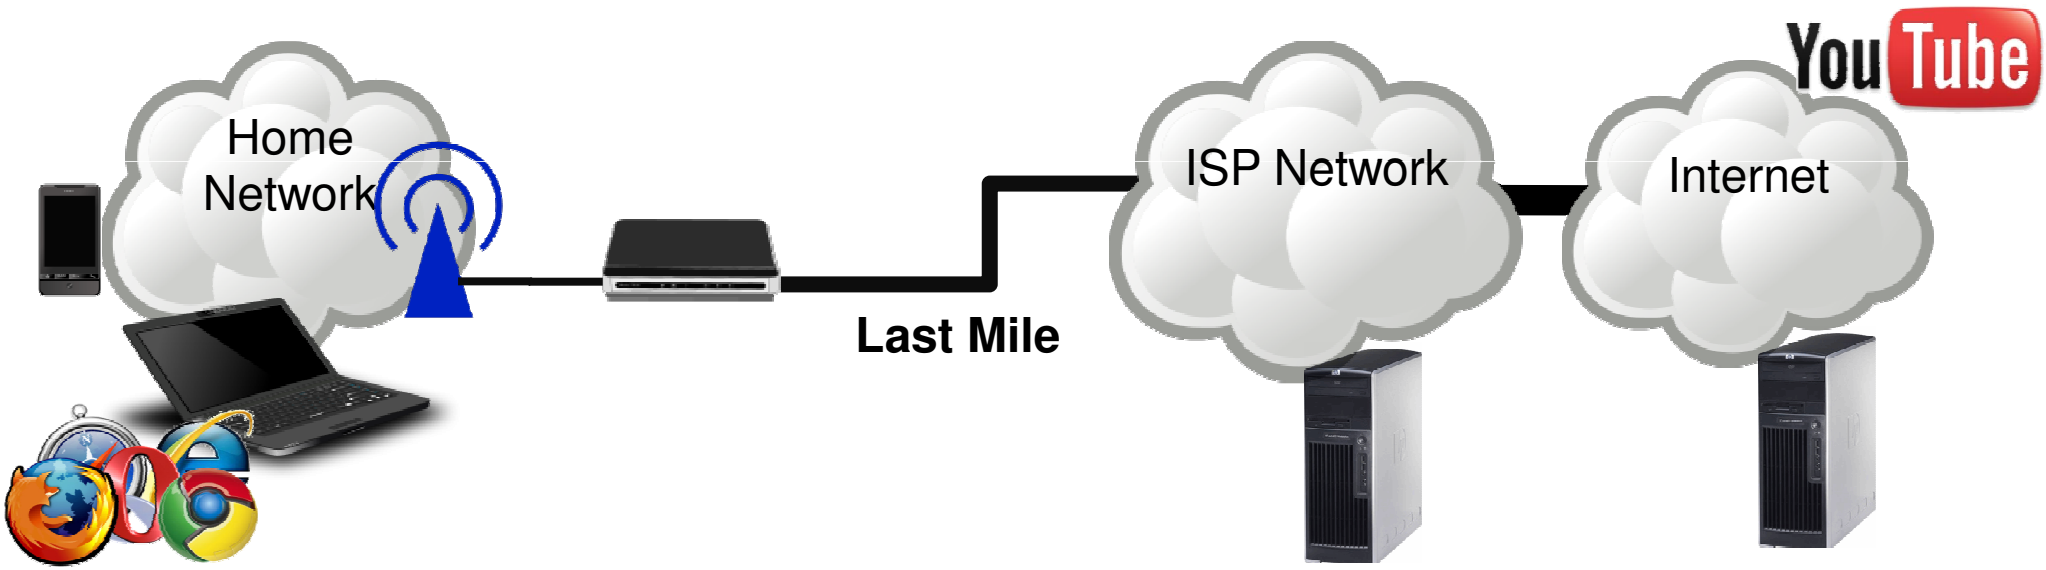

Different projects with the same architecture

- Router-based
  - BISmark → performance and censorship
- Host-based
  - HoBBIT → performance
  - Ubica → censorship

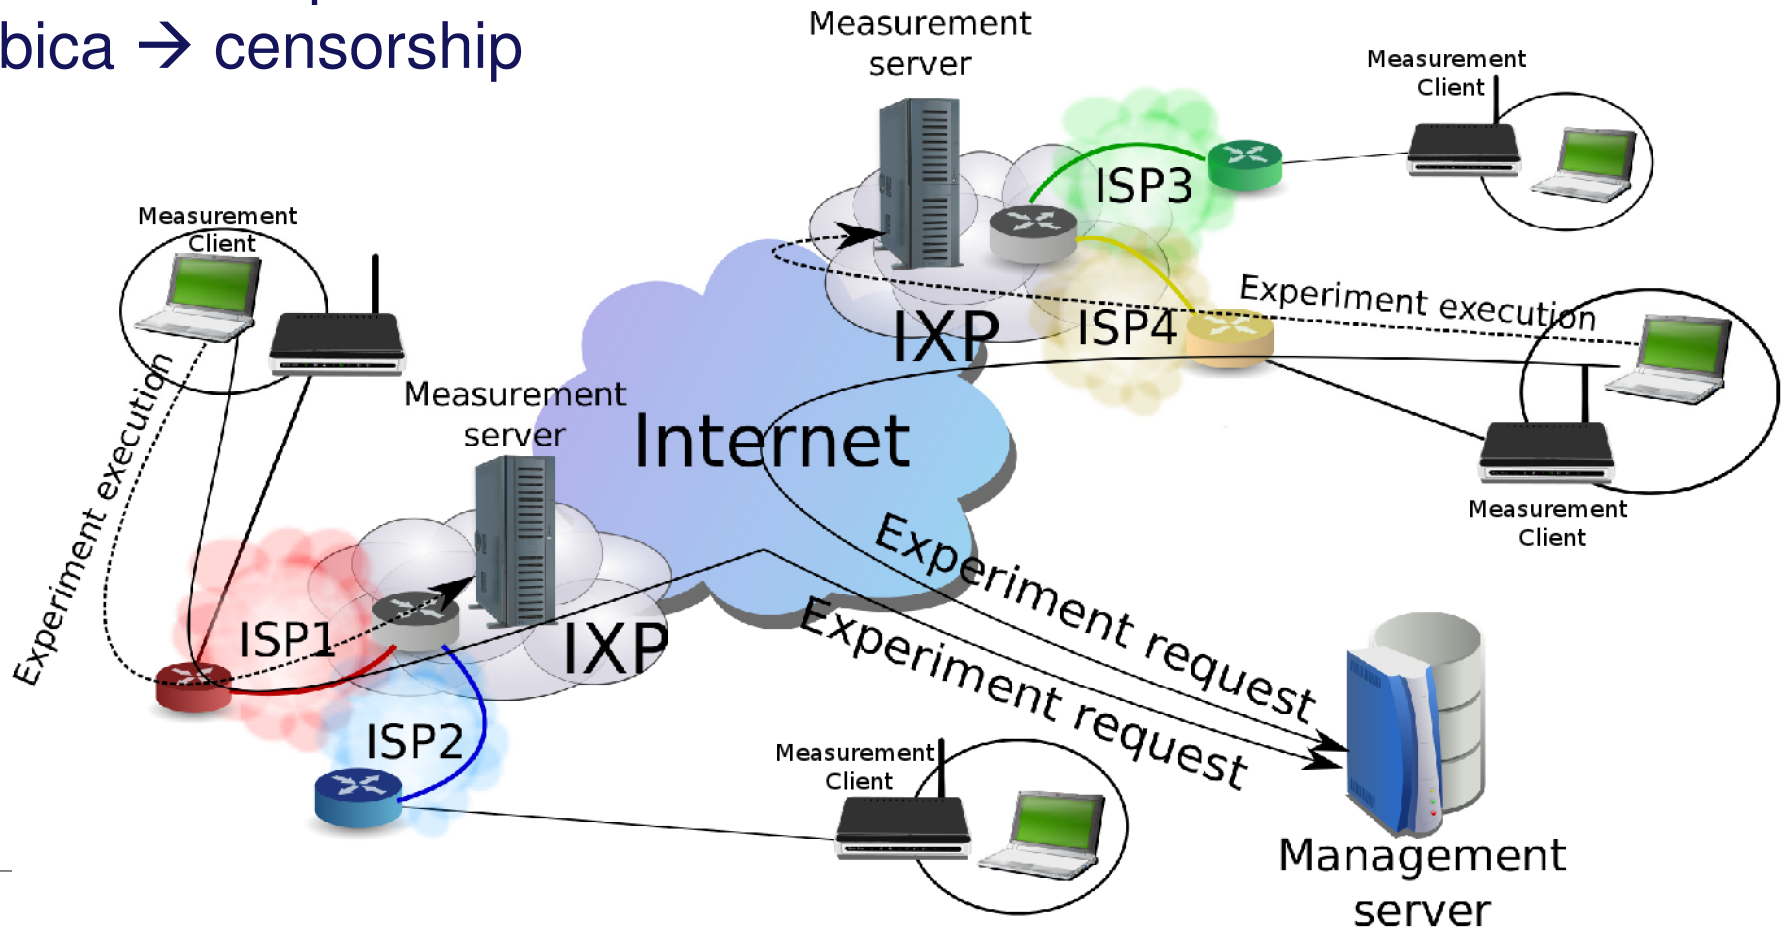

Common features

- Support for pre-existing measurement tools
  - Well tested tools are more accurate
- Implicit management of measurement resources
  - Measurement of existing services
- Flexibility in the configuration of the tools
  - Selection of the best technique depending on the context
- Detection of ISP and position
  - Proper assignment of obtained results
- Accessible real-time reporting
  - Transparency on obtained results

Different features: Router- vs Host-based

- Router-based
  - Modified OS (OpenWRT)
  - Passive analysis, also to understand interferences
  - Remote access to router console for debugging
- Client-based
  - Multiplatform based on Qt libraries and bash/awk ports
  - Detection of connection used
  - Real-time information provided to user
  - Possibility to suspend the measurements

Different features: Performance vs Censorship

- Performance
  - Multiple parameters emulating different applications
  - Centralized scheduling of "heavy" measurements
  - Measurements at different hours of the day
  - Detect potential internal interference
- Censorship
  - Enhanced user privacy
  - Coarse-grain localization
  - More incentives (e.g. submit your url)
  - Controlled deployment

Performance tests

- Hobbit
  - First campaign to understand basic performance (i.e. to infer the contract with the operator)
  - Then period campaigns to infer the maximum throughput, the latency jitter, and losses with TCP and UDP and on different ports

CPE campaign experiments

ID	Measured metrics	Required bandwidth (kbps)	Duration (sec)	Measurement class
#1	Latency, Jitter & Packet loss	512	30	LIGHT
#2	Upstream throughput	2000	15	INVASIVE
#3	Downstream throughput	20000	15	INVASIVE

BPE & BTPE campaigns experiments

ID	Measured metrics	Required bandwidth (kbps)	Duration (sec)	Measurement class
#1	Latency, Jitter & Packet loss	512	30	LIGHT
#2	Upstream UDP throughput	2000	15	INVASIVE
#3	Downstream UDP throughput	20000	15	INVASIVE
#4	Upstream TCP throughput	2000	15	INVASIVE
#5	Downstream TCP throughput	20000	15	INVASIVE

- BISmark
  - Also passive tools (TIE)

Censorship tests

- Tests
  - DNS resolution
  - HTTP reachability
  - Content size
  - IP/Port filtering
- Destinations
  - Lists from
    - http://propakistani.pk/wp-content/uploads/2010/05/blocked.html
    - http://www.herdict.org/explore/data/view?ft=inaccessible
- Origins
  - Large campaign from several sources
  - Tailored tests from China and Pakistan

Current deployments

BISmark

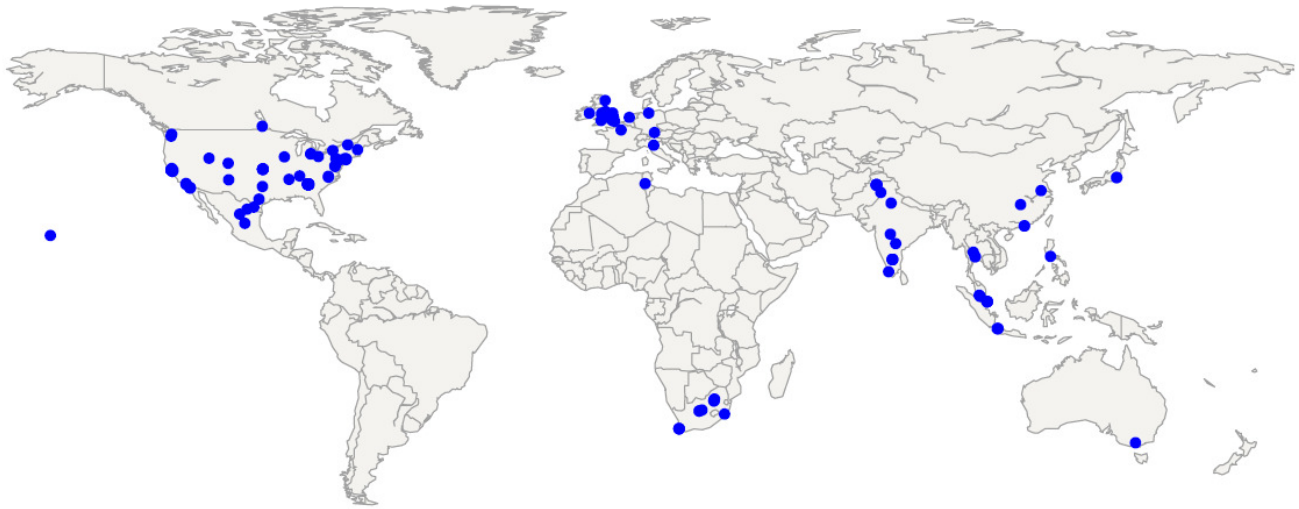

HoBBIT

We are running some tests also from selected PlanetLab nodes using the same framework

Performance results: Hobbit

Average performance over different regions/municipalities

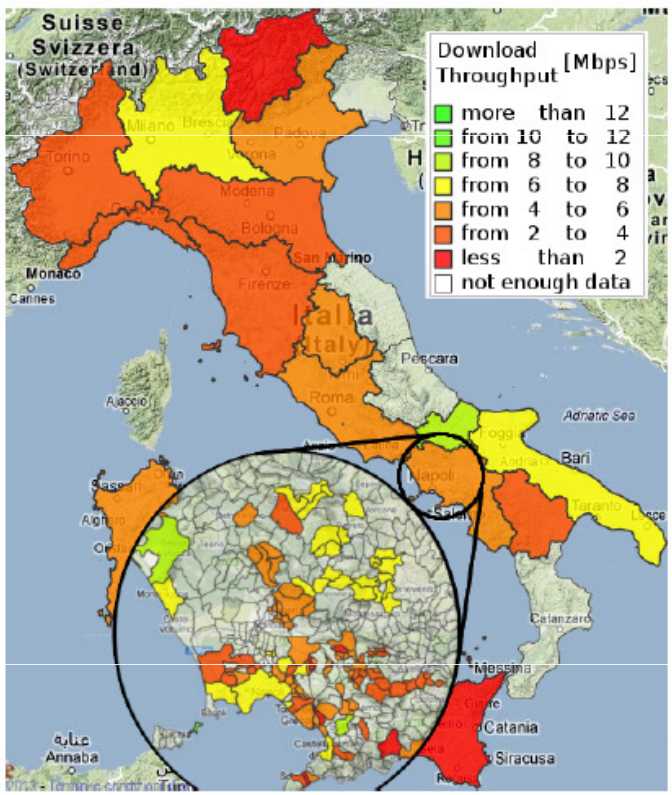

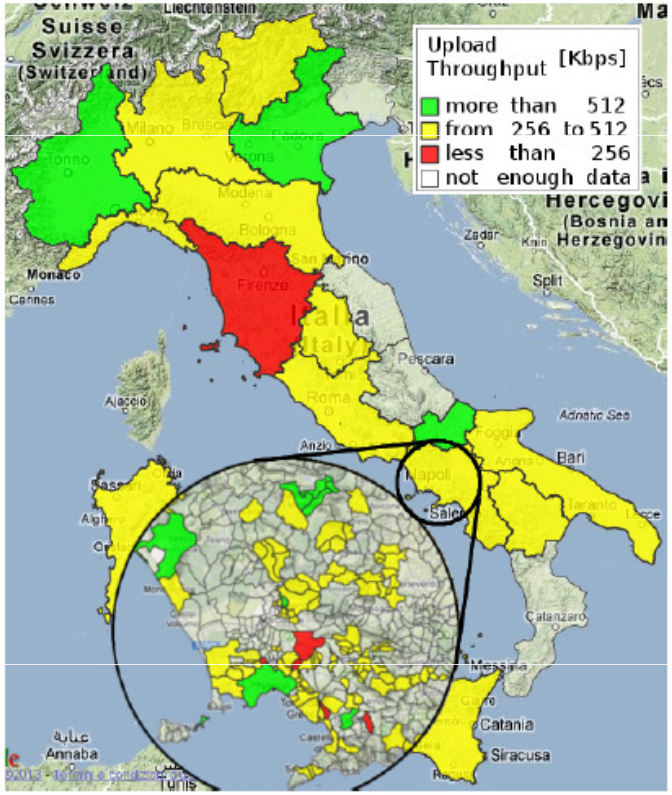

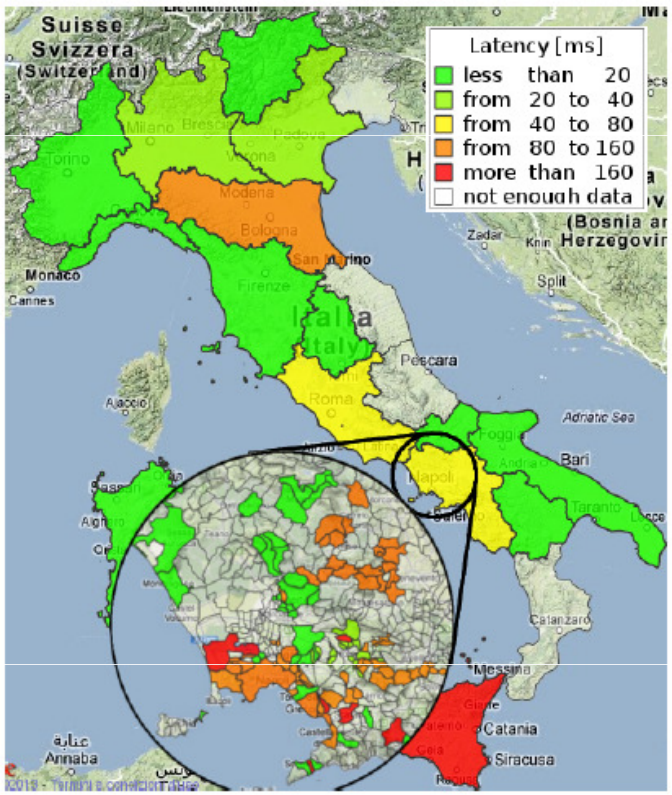

- Maps allow to have a quick sketch of average performance over the geographical areas
- Downlink throughput is mostly between 4-6Mbps, and sometimes between 2-4Mbps
- Uplink throughput is mostly between 256-512Kbps, and sometimes lower than 256Kps

Performance results: Hobbit

Paid vs achieved performance

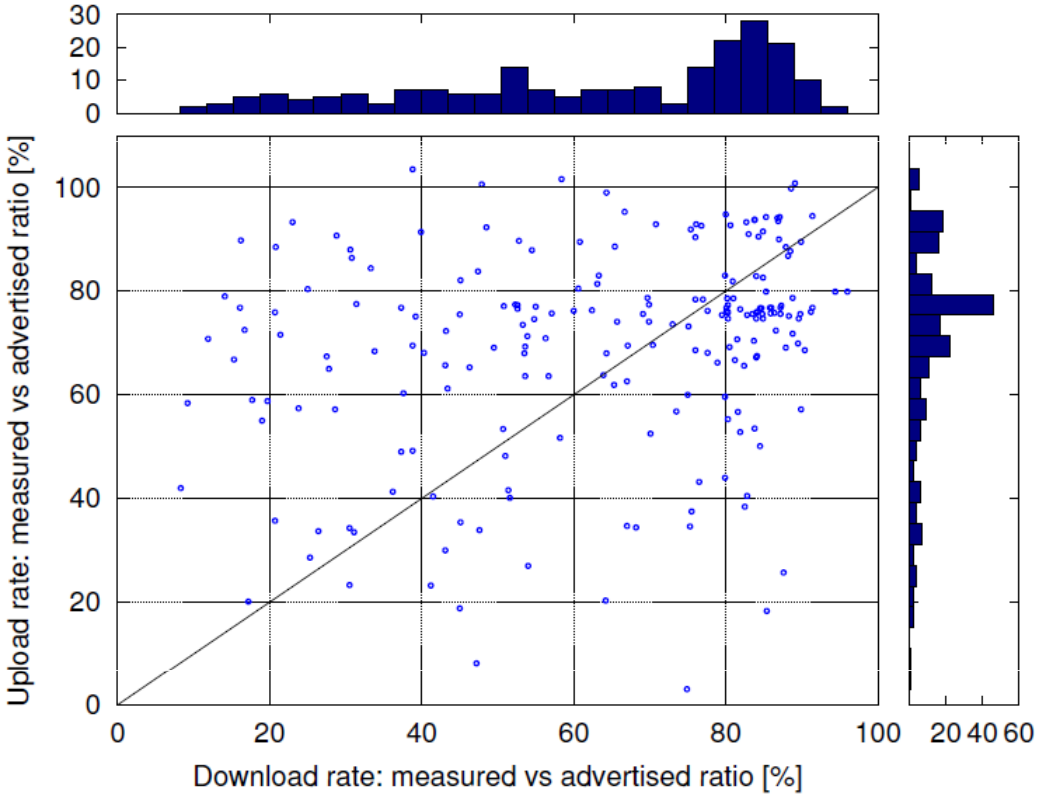

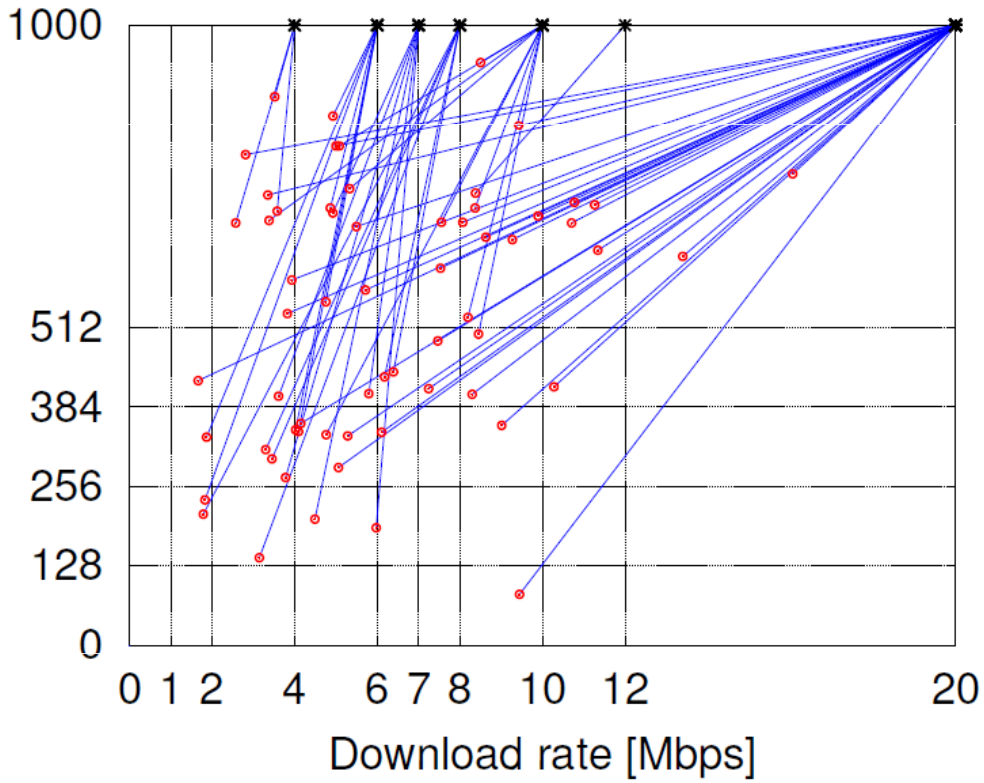

- Most of the users have much less than what they pay for, with no symmetry between uplink and downlink
- Especially when they pay more!

Censorship results 1/2

cc	avgreachperc	avgunreachperc	span(m)
US	76.81	23.19	213.9
CN	35.23	64.77	212.8
JP	75.11	24.89	207.5
NZ	60.49	39.51	207.2
UY	71.16	28.84	209.3
BR	80.97	19.03	208
CA	83.98	16.02	207.1
RU	70.3	29.7	207.6
XX	77.96	22.04	207.3
JO	0	100	201.4
GB	71.85	28.15	206.7
HK	85.25	14.75	209.2
00	19.3	80.7	207.6
BD	35.69	64.31	206.9
IN	85.31	14.69	207.1
KR	79.21	20.79	205.3
AU	79.83	20.17	206.4
AR	76.04	23.96	205.9
CZ	85.09	14.91	209.1
TR	85.54	14.46	207.2
EC	85.2	14.8	208
SE	81.94	18.06	209.6

- No response to HTTP request or response empty, grouped by country
- "Expected" (CN, BD) and unexpected (NZ) countries
- Currently under investigation!!!

- Different html size retrieved for the same site from different countries
- Block page is typically small
- The size ratio may be a good censorship detector

URL	$< size_{PK} >$	$< size_{US} >$	Ratio
https://barenakedislam.wordpress.com	453.0	49095.63	0.01
http://www.internationalfreepresssociety.org	443.5	38085.32	0.01
http://ninjaproxy.com	342.45	14085.42	0.02
NinjaProxy.com	342.39	13154.06	0.03
http://www.similarsites.com	375.33	13701.44	0.03
http://www.youtube.com	4183.91	144177.2	0.03
http://www.freefacebookproxies.com/	9041.17	241485.33	0.04
http://friendlyatheist.com	7881.34	205294.23	0.04
http://www.loonwatch.com	2661.73	65075.19	0.04
http://www.sodahead.com	3575.67	73969.7	0.05
http://www.hotspotshield.com/	731.8	10789.91	0.07
http://face-of-muhammed.blogspot.com/	6208.7	85342.93	0.07
http://www.foxnews.com	4705.53	63425.26	0.07
http://www.buzzfeed.com	22097.93	287001.77	0.08

Censorship results 2/2

Evidences and visualization of DNS hijacking in Italy

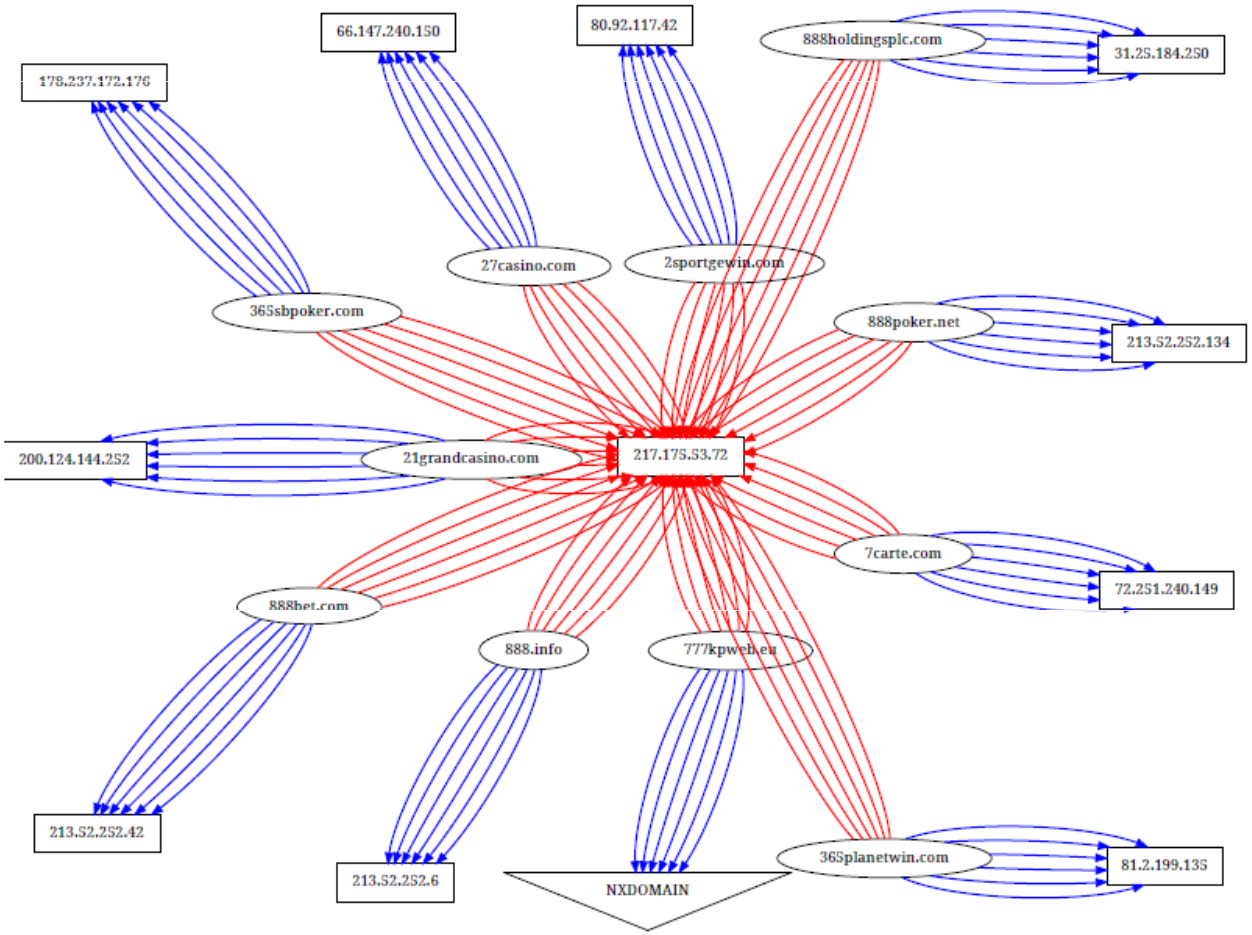

UBICA: IT, home (red) BISMark: US, university (blue)

- Non-authorized online betting services are hijacked to warning page
- The graph visualization and graph properties allow to detect this behavior

Conclusions

- Large scale platforms are necessary for understanding both residential broadband performance and country-wide censorship
- We designed, developed and deployed different platforms both router- and host-based and for both performance and censorship analysis
- Common platforms can be used, but specific requirements have to be considered for router/host-based approaches and performance/censorship applications
- For performance analysis there is necessity of collaboration among running platforms
- For censorship detection there is still necessity of novel tests and metrics

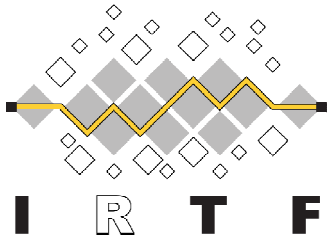

Thanks!!!

More info at:

http://traffic.comics.unina.it

a.botta@unina.it