GlobeTraff
A traffic workload generator for the performance evaluation of ICN architectures

K.V. Katsaros, G. Xylomenos, G.C. Polyzos
A.U.E.B.

(presented by K. Pentikousis on behalf of the authors)

Work supported by the FP7 PURSUIT Project

ICNRG Meeting @ IETF-88, Vancouver, Canada

Outline

- Traffic characteristics
- Models Implemented
  - Web, P2P, Video, Other
- The GlobeTraff tool
- Conclusions

Traffic characteristics

- Traffic patterns play a vital role
  - Past research has led to many traffic models
    - Web, P2P, Video
  - Application specific traces are not enough
  - We need to consider global traffic effects
- GlobeTraff is a synthetic traffic workload generator
  - Based on current Internet traffic models
  - Allows the creation of various traffic mixes
    - What applications? How much do they contribute?
  - Allows changing the characteristics of each model
    - Of course we cannot predict future applications

Traffic characteristics

- Popularity characteristics
  - Popularity distribution
    - Number/fraction of requests for object
  - Temporal locality
    - How are requests distributed in time?
  - Spatial locality (not modeled)
    - Distribution of requests across the network
- Object sizes
  - Distribution of sizes for items
    - Direct effect on transport
    - Indirect effect on caching

Models Implemented

- GlobeTraff supports several traffic types
  - Web, P2P, Video, Other
  - Models from recent literature
  - Fully parameterized via GlobeTraff's GUI
- Traffic mixture
  - Based on measurements (with DPI techniques) [1][2]

Traffic Type	**Percent of Total Traffic**
Web	31.2-39%
P2P	17-19%
Video	13- 20.8%
Other	29-31%

[1] C. Labovitz et al., "Internet inter-domain traffic," in ACM SIGCOMM. ACM, 2010, pp. 75–86.
[2] G. Maier et al., "On dominant characteristics of residential broadband internet traffic," in ACM IMC, 2009, pp. 90–102.

Models Implemented: Web

- Popularity distribution: Zipf-like
  - $p(i) = K/i^a$
    - i: popularity rank, N: total items
    - $K=1/Sum(1/i^a)$
    - a: slope of distribution, values 0.64-0.84
- Temporal Locality
  - Ordering via LRU stack model
  - Exact timing via exponential distribution
- Object Sizes
  - Concatenation of Lognormal (body) and Pareto (tail)

Models Implemented: P2P

- Popularity distribution: Mandelbrot-Zipf
  - $p(i) = K/((i+q)^a)$
    - q: plateau factor, 5 to 100
    - Flatter head than in Zipf-like distribution (where q=0)
- Temporal Locality: based on BitTorrent
  - Average arrival rate of 0.9454 torrents per hour
  - Peers in a swarm arrive as $\lambda(t)=\lambda_0 e^{-t/\tau}$
    - λ_0: initial arrival rate (87.74 average)
    - τ: object popularity (1.16 average)

Models Implemented: P2P

- Temporal Locality: based on BitTorrent
  - Random ordering of swarm births (first request)
  - For each swarm we calculate a different τ
    - Based on average τ and object popularity
  - Exponential decay rule for subsequent requests
- Object Sizes
  - Wide variation on torrent sizes
  - No analytical model exists
  - Either sampling of real BitTorrent traces
  - Or use of a fixed value

Models Implemented: Video

- Popularity distribution: based on YouTube
  - Weibull distribution ($k=0.513$, $\lambda=6010$)
  - Gamma distribution ($k=0.372$, $\theta=23910$)
- Temporal Locality
  - No analytical models available
  - Random distribution across total duration
    - Total duration is determined by the P2P trace
- Object Sizes: based on YouTube
  - Concatenated normal distribution for duration
  - Same for size since most videos are 330 Kbps

Models Implemented: Other

- Popularity distribution
  - Zipf-like distribution as it is the most common
- Temporal Locality
  - Same approach as for video traffic
  - Possibility of using the web traffic model
    - Exponential distribution of inter-arrival times
- Object Sizes
  - GlobeTraff allows the user to set the size
  - We expect huge amounts of *small* items
    - Internet of Things, machine-to-machine communication

The GlobeTraff tool

- Implementation
  - Based on the ProWGen tool
    - Used for the Web traffic model
    - Extended with the other models
  - Command line tool written in C++
  - Java GUI to drive the tool
- Usage
  - Composition of the generated traffic mix
  - Total size for the trace
  - Parameters for each model
    - Also distributions where multiple options exist

The GlobeTraff tool: Web traffic

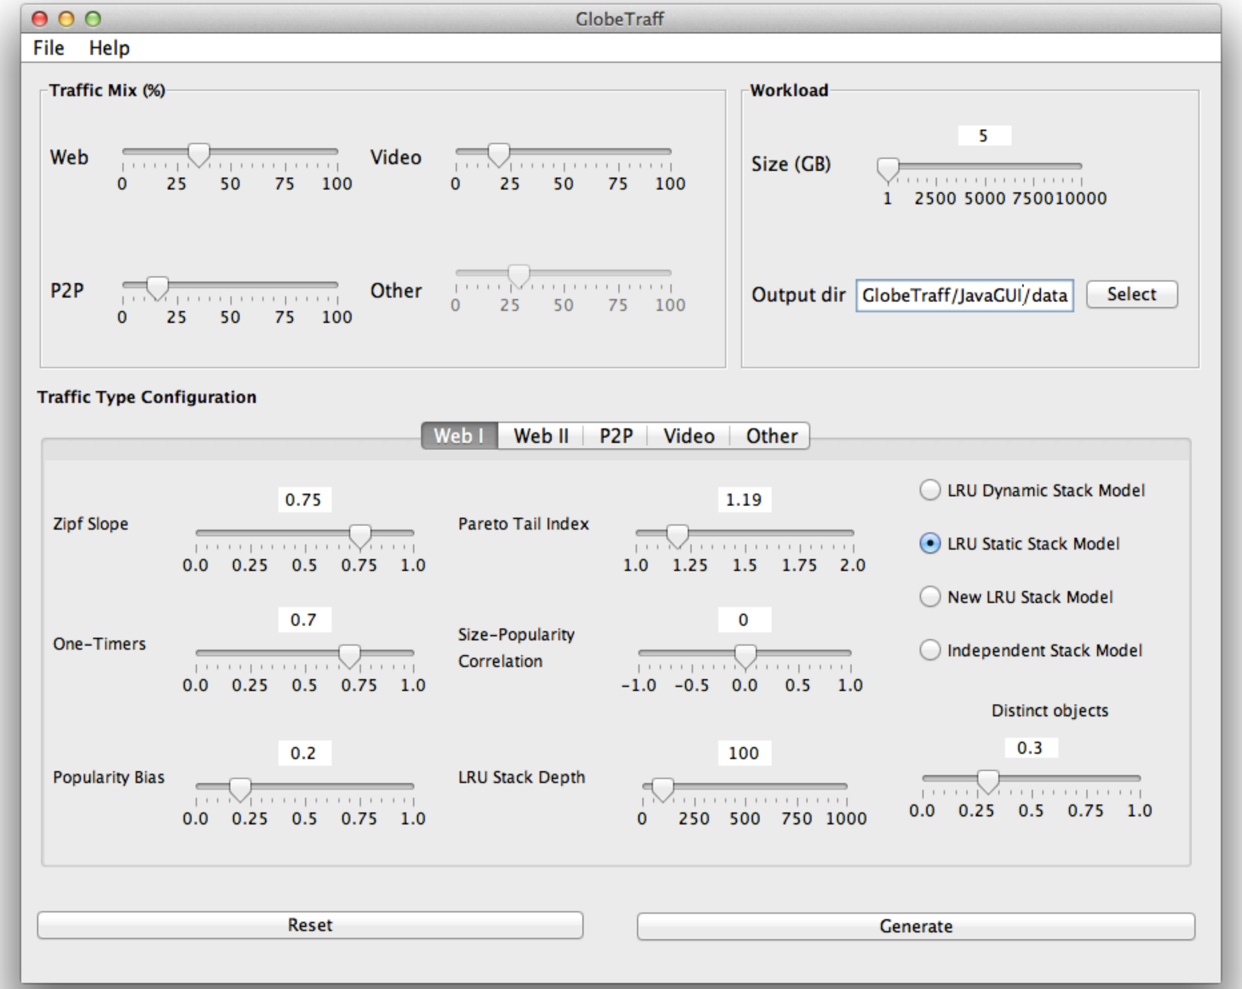

The GlobeTraff tool: P2P

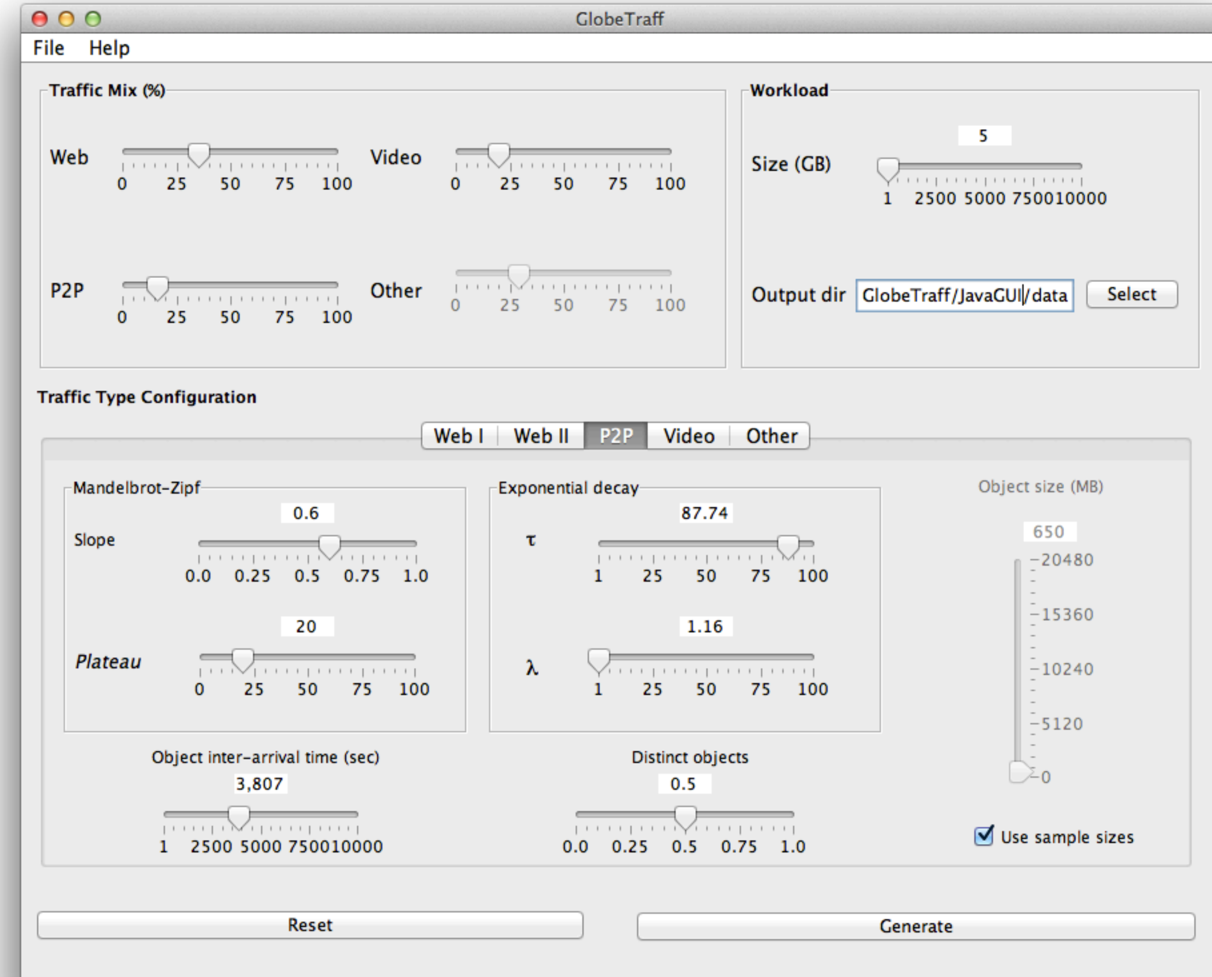

The GlobeTraff tool: Video

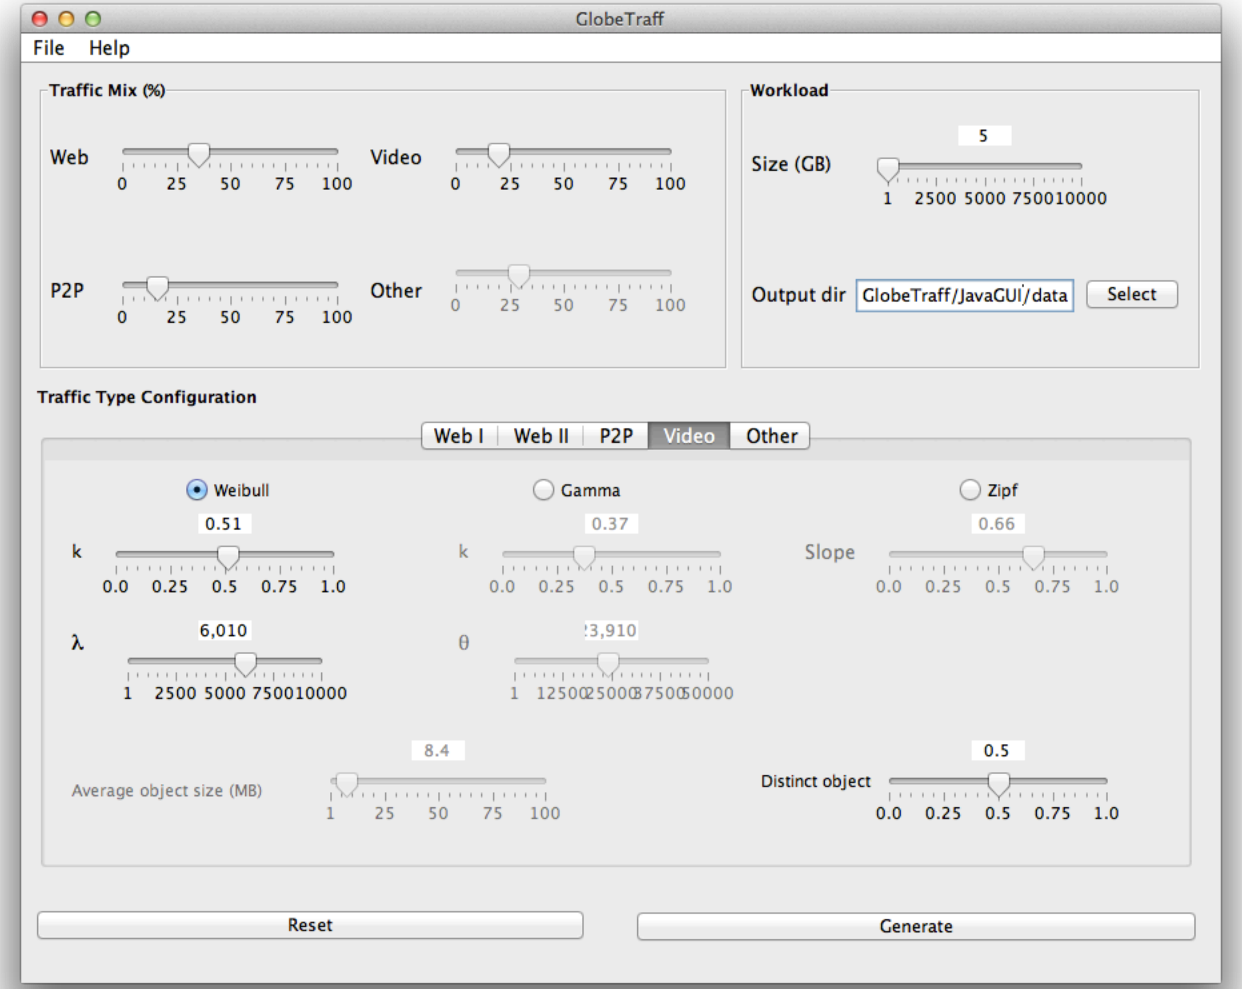

The GlobeTraff tool: *Other* traffic

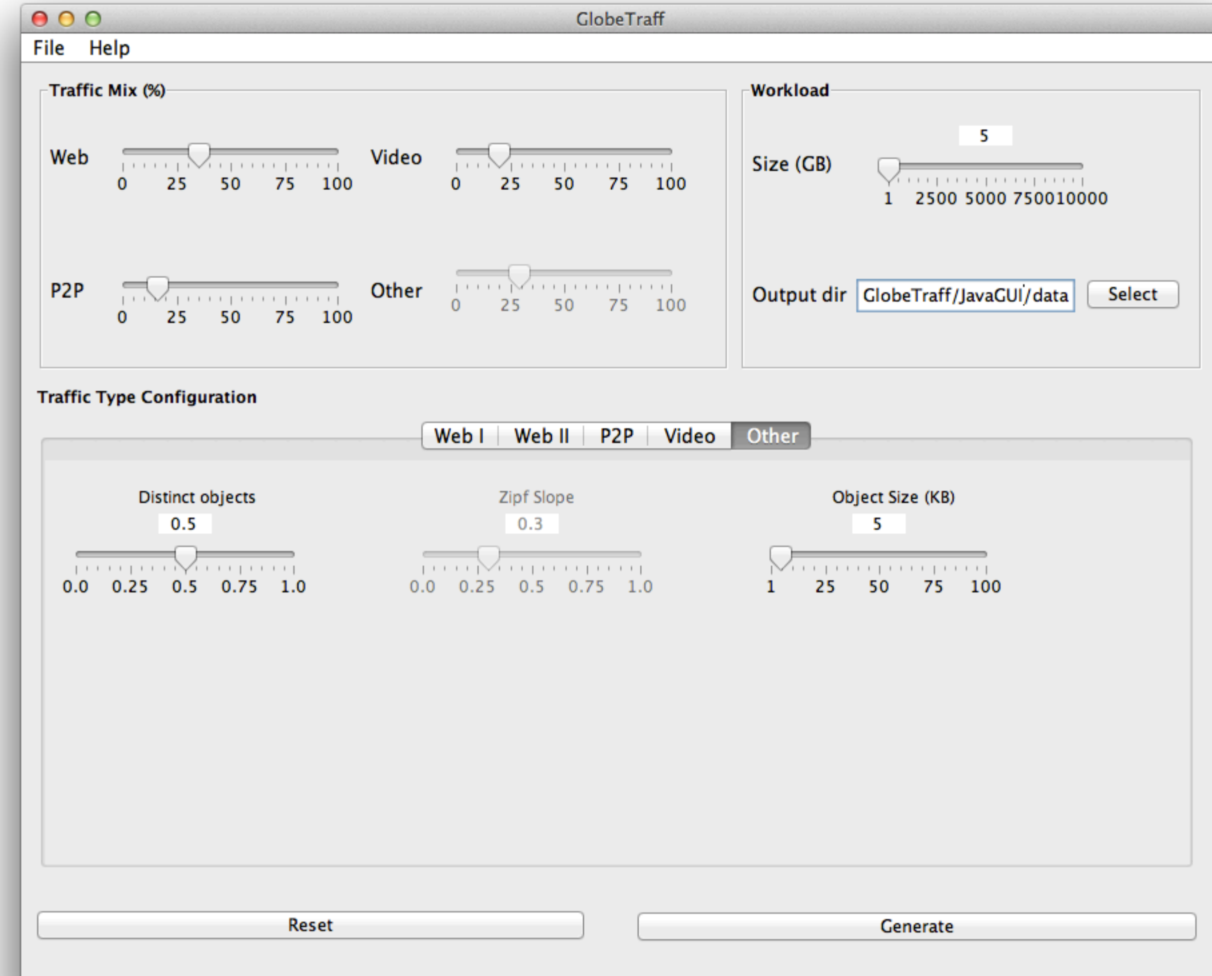

The GlobeTraff tool

- Output
  - In two files as in ProWGen
    - Per traffic type and for the entire mix
  - Table 1: Objects in the workload
    - Unique ID for the object
    - Popularity expressed as total number of requests
    - Size in bytes
    - Application type (1: Web, 2: P2P, 3: Video, 4: Other)

Item ID	Popularity	Size (Bytes)	Application Type
0	34	8145	1
1	18	9033	1
2	12	8751	1
3	12	8145	1

  - Table 2: Workload in time
    - Time the request is submitted
    - ID of the Object referred to
    - Size of the object (same as Table 1)

Time	Item ID	Size (Bytes)
0.018314	0	8145
0.171501	2	8751
1.374289	1	9033
4.240712	1	9033

Future ICN traffic

- How would traffic load look like in an ICN environment?
  - Models based on existing applications and network architecture

- Efficient content delivery (e.g., caching) ➔ P2P traffic could diminish (?)
- New (ICN) applications?
  - E.g., Internet of Things (IoT) traffic
- Increased signaling traffic
  - E.g., CCN Interest packets, IoT name resolution (?)
- …

- Adapting GlobeTraff
  - Models only content requests, not the actual delivery mechanism
  - Easily extensible

Conclusions

- Realistic traffic models are very important for ICN
  - Need to evaluate an entirely new concept
  - Cannot rely on individual traffic models
  - All traffic types end up in the same caches!
- GlobeTraff provides global traces
  - Many traffic models based on literature
  - Fully parameterized via GUI
  - Allows projections on mix and individual types

GlobeTraff

- **Available at:** http://goo.gl/QkBxVf
- **Contact:** Konstantinos V. Katsaros (ntinos@aueb.gr)
- **Also see:**

K. V. Katsaros, G. Xylomenos, and G. C. Polyzos, “GlobeTraff: a traffic workload generator for the performance evaluation of future Internet architectures,” IEEE/IFIP International Conference on New Technologies, Mobility and Security (NTMS), May 2012