Markov modelling for queues of Internet traffic

UK Performance Engineering Workshop (2007)

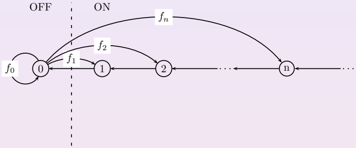

Richard G. Clegg (richard@richardclegg.org)
Dept. of Electronic and Electrical Engineering, UCL
— UK Performance Engineering Workshop, 2007
(Prepared using LaTeX and beamer.)

Talk Overview

Motivation

- Mathematically appealing Markov models of internet data in literature.
- Models capture Long-range dependence of real data (plus other parameters).
- Would like a simple queuing model to do maths with.
- How useful are these models in practice?

Talk Overview

Motivation

- Mathematically appealing Markov models of internet data in literature.
- Models capture Long-range dependence of real data (plus other parameters).
- Would like a simple queuing model to do maths with.
- How useful are these models in practice?

1. Define the problem of simulating “one link” model.

Talk Overview

Motivation

- Mathematically appealing Markov models of internet data in literature.
- Models capture Long-range dependence of real data (plus other parameters).
- Would like a simple queuing model to do maths with.
- How useful are these models in practice?

1. Define the problem of simulating "one link" model.
2. Seven simple ways to model internet traffic (usually with MCs).

Talk Overview

Motivation

- Mathematically appealing Markov models of internet data in literature.
- Models capture Long-range dependence of real data (plus other parameters).
- Would like a simple queuing model to do maths with.
- How useful are these models in practice?

1. Define the problem of simulating "one link" model.
2. Seven simple ways to model internet traffic (usually with MCs).
3. Tests using a very simple infinite buffer queuing model.

Talk Overview

Motivation

- Mathematically appealing Markov models of internet data in literature.
- Models capture Long-range dependence of real data (plus other parameters).
- Would like a simple queuing model to do maths with.
- How useful are these models in practice?

1. Define the problem of simulating "one link" model.
2. Seven simple ways to model internet traffic (usually with MCs).
3. Tests using a very simple infinite buffer queuing model.
4. Compare with freely available real internet data sets.

Irresponsibly hasty guide to Long-Range Dependence

- LRD (also known as long memory) occurs when a data has significant correlations over a number of time scales.
- Imagine that data at a particular time t having some signficant effect on the data at time $t + k$ even if k becomes very large.
- This data might, therefore, have large peaks (or troughs) which cause queuing problems.
- Measured in packets/unit time on internet data [Leland et al '93]. Can cause problems with queuing/delay [Erramilli etc al 96].

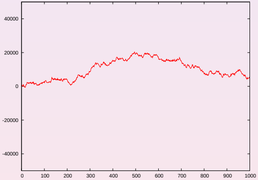

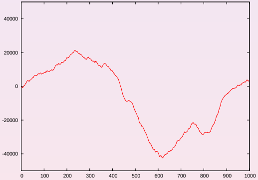

The Markov Model

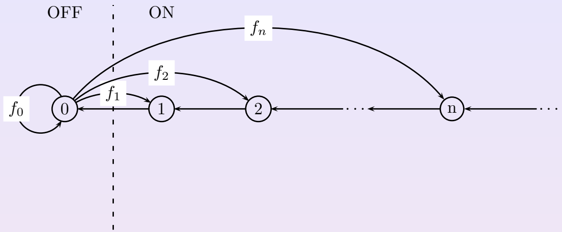

- This is topology of Wang and Clegg/Dodson models.
- If $\{X_t : t \in \mathbb{N}\}$ is generated by chain then generate

$$Y_t = \begin{cases} 0 & X_t = 0 \\ 1 & \text{otherwise.} \end{cases}$$

- Choose f_i so return times have heavy-tails and get binary series with LRD [Heath et al 1998].
- Both models set mean and H parameter.
- Exact solution to discrete queuing model exists.

Pseudo-Self-Similar Traffic Model (PSST)

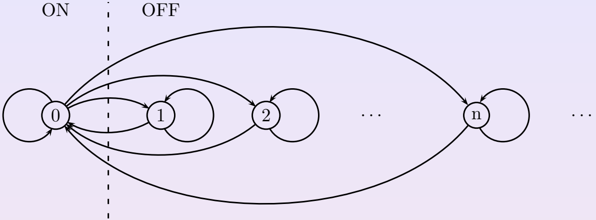

- Introduced in [Robert et al '97] no proof of LRD.
- Parameters: q relates to mean a has no obvious interpretation.

$$\mathbf{P} = \begin{bmatrix} \Sigma_0 & \frac{1}{a} & \frac{1}{a^2} & \cdots \\ \frac{q}{a} & \Sigma_1 & 0 & \cdots \\ \left(\frac{q}{a}\right)^2 & 0 & \Sigma_2 & \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

Arrowsmith/Barenco Model

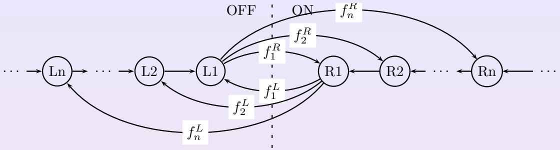

- General class of models described in [Barenco & Arrowsmith '04] proof of strong result giving LRD.
- Think of as double-sided version of Wang topology.
- Could set model to use LRD with Wang or Clegg/Dodson probabilities but theoretical issues cause problem with mean and stability.
- Instead use on/off length distributions for real data.
- Results here not be a criticism of this family of models.

Models Used

- Simple and tractable packet generation models.
- Models are "clocked" and "binary". Fixed width packets generated at times $n\Delta t : n \in \mathbb{N}$.
- Generating Models (listed in chronological order):
  1. Poisson process (strictly speaking Bernoulli process) (mean only).
  2. Fractional Brownian Motion model (mean and Hurst parameter).
  3. Wang model [Wang '89] — Markov Modulated process (mean and H).
  4. Pseudo Self-Similar Traffic (PSST) [Robert et al '97] — MMP (mean and ?).
  5. Arrowsmith/Barenco [Barenco & Arrowsmith '04] — MMP (mean and on/off dist).
  6. Clegg/Dodson [Clegg & Dodson '05] — MMP (mean and H).
  7. UH model (Bernouilli–Zeta) [Conversation in pub '07] — MMP (mean and H).

Queuing Model

- Assume a single FIFO server with an infinite buffer and output bandwidth b.
- Takes time l/b to process a packet of length l.
- If l is constant then this is a G/D/1 type queue.
- Measure $\mathrm{E}[q]$ the expected queue length (in packets or in bits) as function of b.
- Input to the queue maybe from "real" traffic traces or from models.

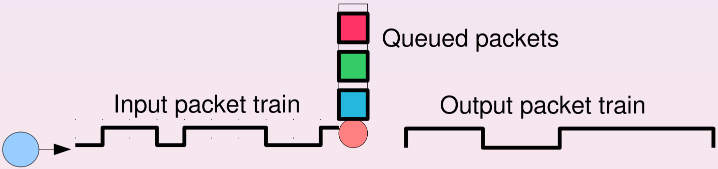

Real Traffic Traces

- 100,000 packets from real life traffic sources which give times and packet lengths.
- Establish base case — use arrivals times and lengths as input to queue. Pick b to get approx 10% occupancy.
- Get "digitised" version of real data by only allowing output of fixed l bit packets at times $n\Delta t$.
- CAIDA OC48 data, two sets $(H = 0.6)$ $(H = 0.65)$ from spring 2003. High speed link (2.45 Gb/s). Available from CAIDA website.
- Bellcore data $(H = 0.8)$ $(H = 0.9?)$ much beloved historic data from autumn 1989. Available from Internet Traffic Archive.

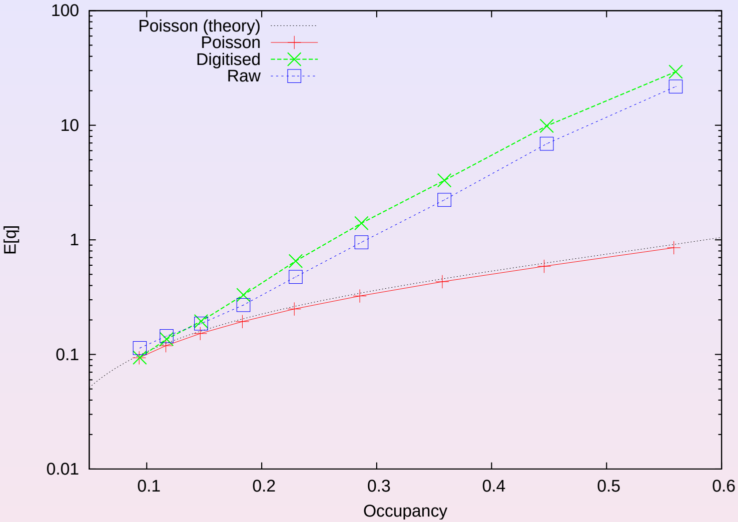

Poisson vs Bellcore (theory line is from P-K theorem).

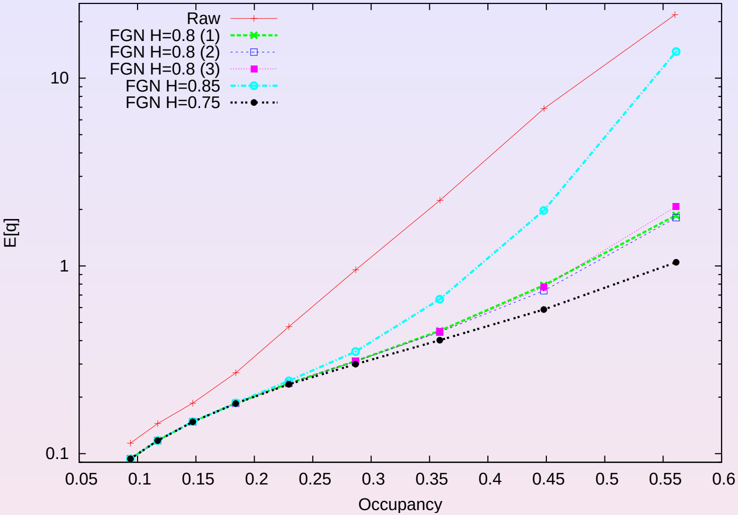

Bellcore vs FGN (several realisations with $H = 0.8$ and one each of $H = 0.75$ and $H = 0.85$).

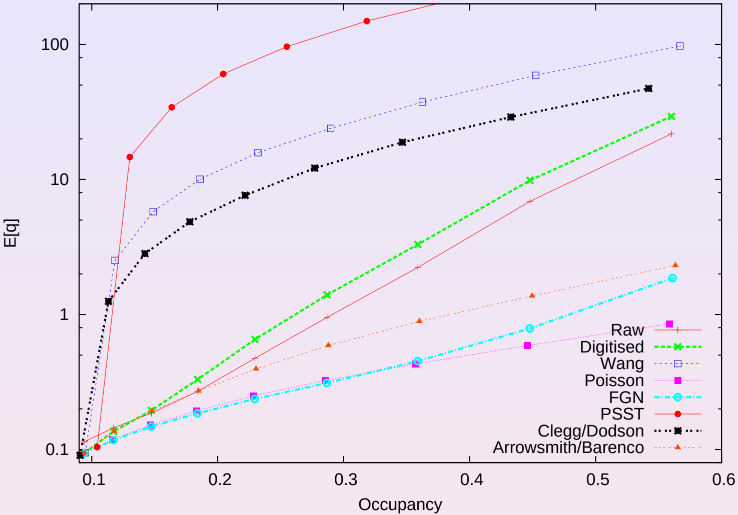

Bellcore – All models compared with real data.

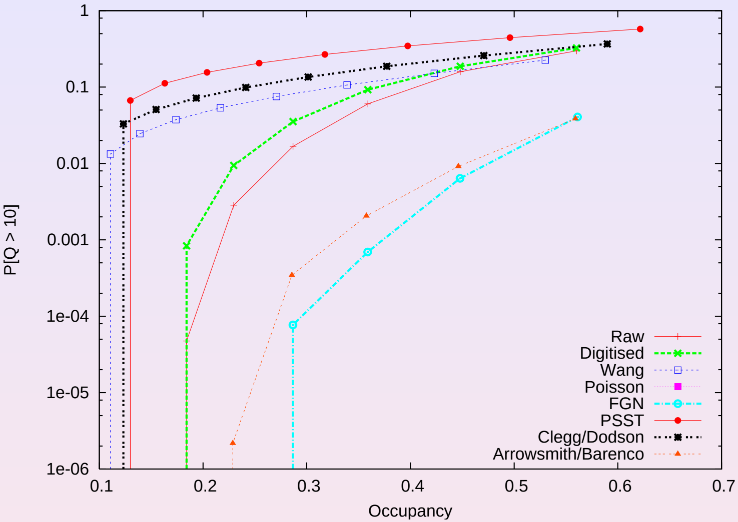

Bellcore – $\mathbb{P}\left[Q > 10\right]$.

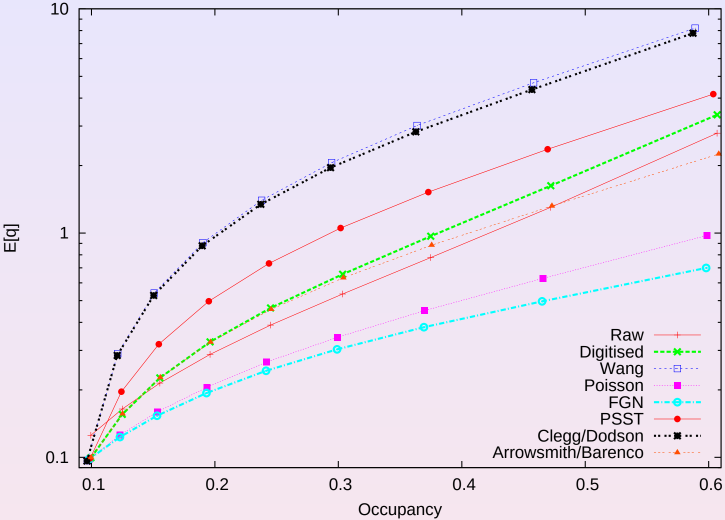

CAIDA – All models compared with real data.

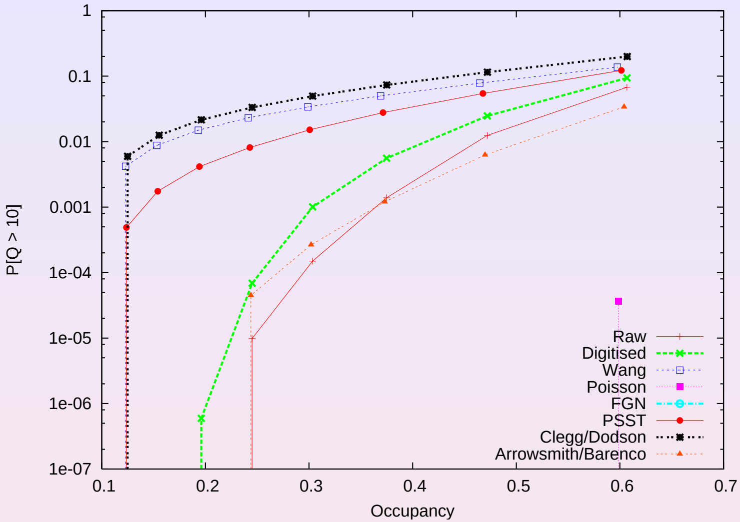

CAIDA – $\mathbb{P}[Q > 10]$.

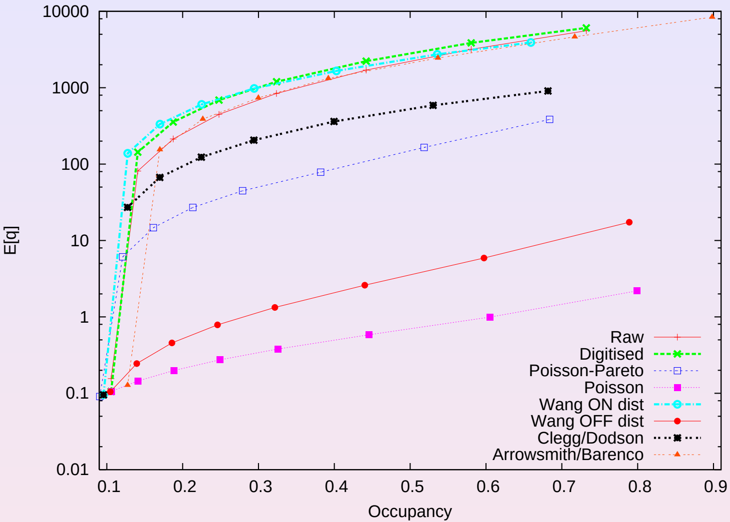

Bellcore trace 2 – Selected models compared with real data.

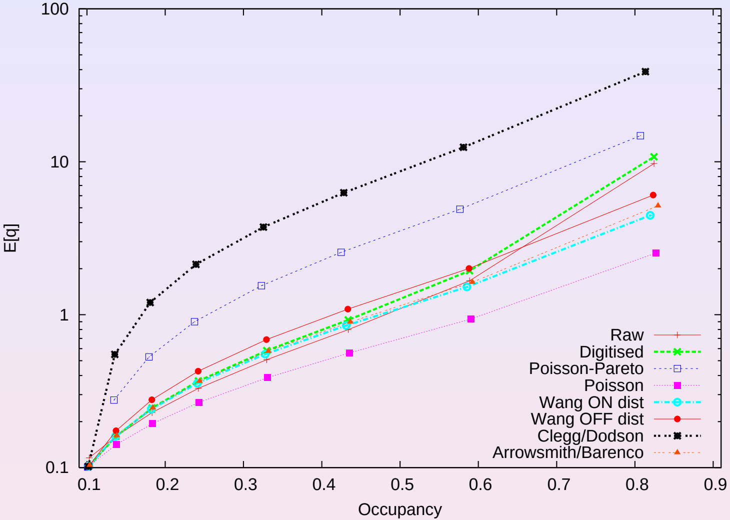

CAIDA trace 2 – Selected models compared with real data.

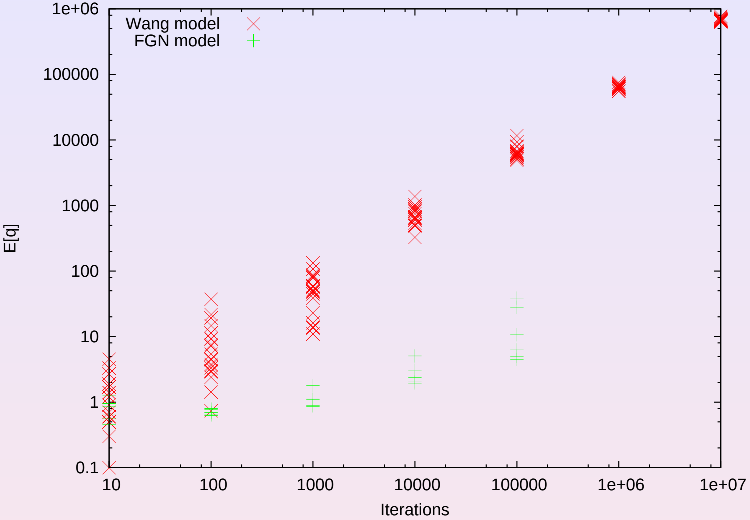

The effect of increasing the number of iterations on $\mathrm{E}[q]$ in two LRD models.

Conclusions (general conclusions)

- No models were always close to matching queuing behaviour.
- Getting a simple model to match queuing performance is very difficult.
- The "digitisation" in these models is not the reason for the difference.
- Real traffic is variable in ways which simple models cannot be.
- Modelling the distribution of on and off periods is not sufficient in the high Hurst case but may be for low Hurst.
- Models which took the distribution of ON burst lengths were sometimes "good enough".
- I need more data and fewer parameters (good models have many parms).

Conclusions (LRD modelling)

- LRD is a nuisance to work with (poor convergence of mean, hard to measure H) is it fundamental anyway?
- All LRD models matched mean (sort of) and Hurst once aggregated (except PSST) but got different wrong answers.
- The PSST model is very peculiar. (Non-Hurst LRD?)
- Hurst parameter can be "naughty" or "nice" [Neidthardt '98].
- Different models which give the same mean and H give very different queuing performance.
- With an infinite buffer these models are predicting infinite queue and delay.
- The very idea of LRD modelling may be fundamentally broken.

Where to now?

- Multi-parameter models? (Multi-fractal wavelet model? Variants of Arrowsmith/Barenco model? Capture ACF?)
  - Pro: Captures more parameters of traffic.
  - Pro: Mathematics is interesting.
  - Anti: Mathematics is much more difficult (accuracy versus understanding).
- Closed loop models?
  - Pro: Captures importance of TCP feedback mechanism.
  - Anti: Likely to be mathematically intractable.
  - Anti: Does complex simulation gain us understanding?
- What am I missing? (User behaviour? Network behaviour? Misunderstanding theory?)
- Definitely more research required.

References

1. Barenco and Arrowsmith, Dynamical Systems, vol 19(1) p61-74 (2004)
2. Clegg and Dodson, Phys. Rev. E, vol 72 (026118) (2005)
3. Clegg, Int. Journ. Simul.: Sys., Sci. & Tech. vol 7(1) p3-14 (2006)
4. Erramilli, Narayan & Willinger, IEE/ACM Trans. Net vol 4(2) p209-223 (1996)
5. Heath, Resnick & Samorodnitsky, Math. of Oper. Res., vol 23(1) p 145-165 (1998)
6. Leland, Taqqu, Willinger, Wilson, Proc ACM SIGCOMM, p183-193 (1993)
7. Neidhardt & Wang, Proc. ACM SIGMETRICS, p222-232 (1998)
8. Robert & Le Boudec, Perf. Eval., vol 30 p57-68 (1997)
9. Wang, Phys. Rev. A, vol 40(11) p6647-6661 (1989)

This talk, the author's papers referred to above and the software used are all available online at:

`www.richardclegg.org/`.

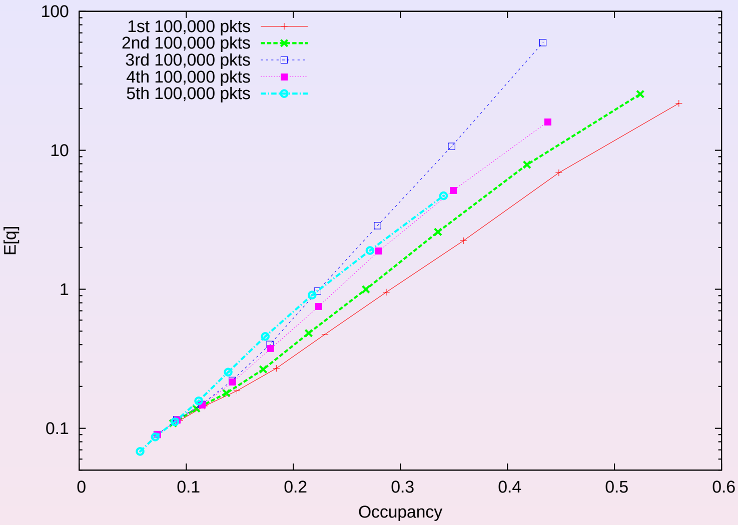

Bellcore – next four blocks of 100,000 packets.

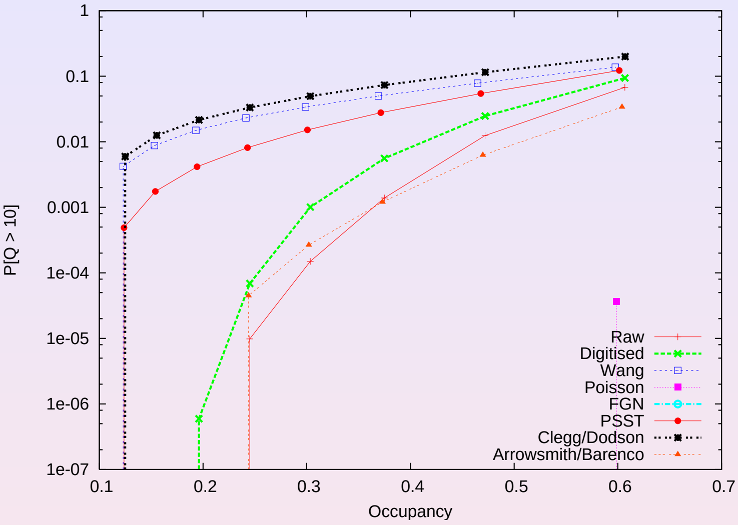

CAIDA – $\mathbb{P}[Q > 10]$.

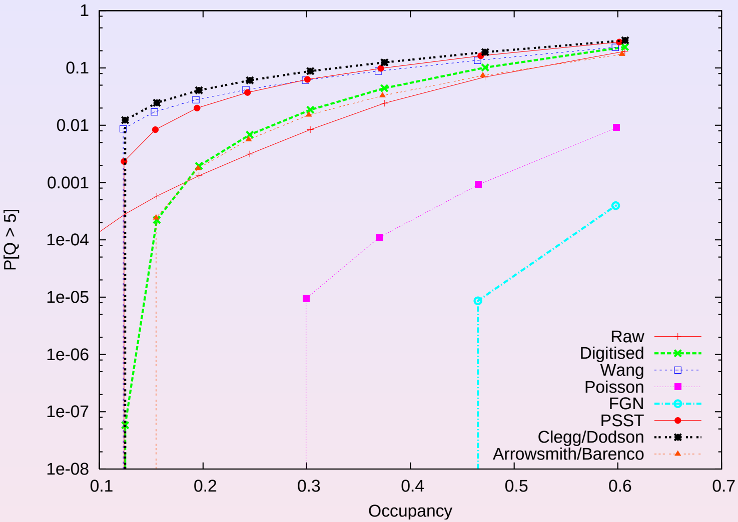

CAIDA – $\mathbb{P}[Q > 5]$.

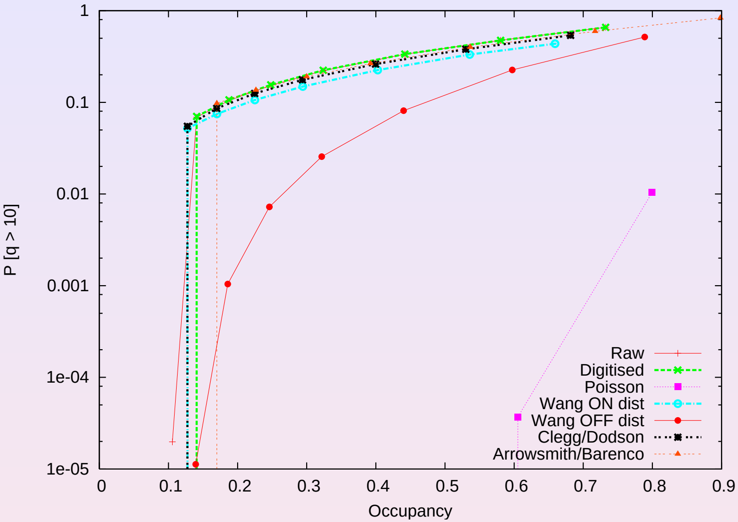

Bellcore 2 – $\mathbb{P}[Q > 10]$.

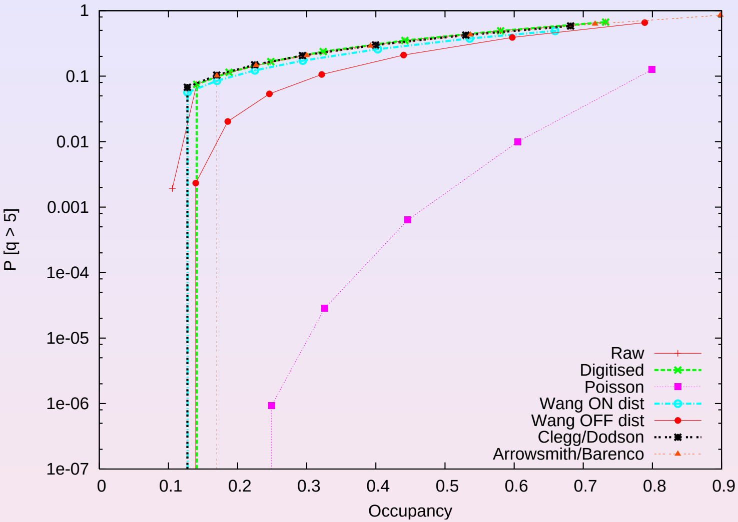

Bellcore 2 – $\mathbb{P}[Q > 5]$.

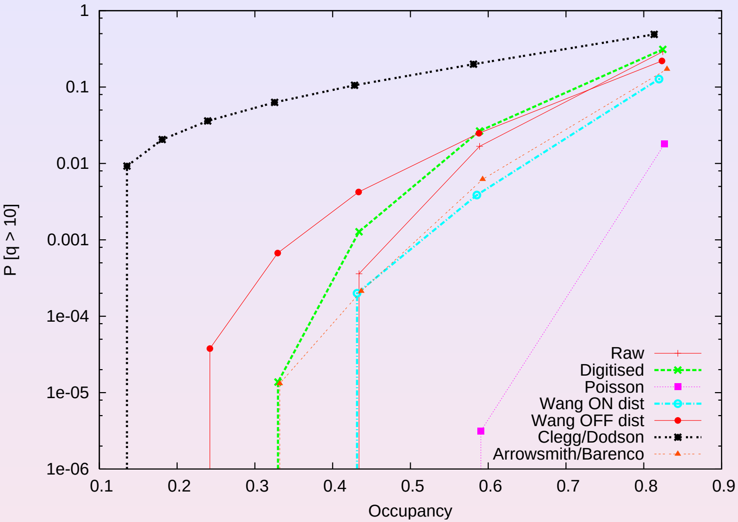

CAIDA 2 – $\mathbb{P}[Q > 10]$.

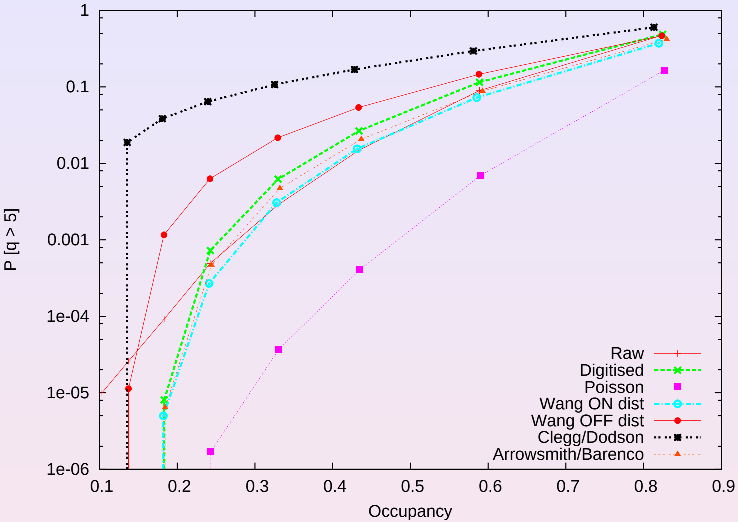

CAIDA 2 – $\mathbb{P}[Q > 5]$.