Modeling end-to-end internet delays using mixtures of Weibull distributions

Iain W. Phillips and José A. Hernández

Computer Science, Loughborough University

July 2004

Introduction

- History of work at Loughborough
- Other Measurement projects
- Visualisation of Measurements
- Mathematical Modelling
- Applications

Ancient History

- In 1994 JANET $\longrightarrow$ SuperJanet, contract won by BT
- Built over SMDS—Switched Multi-megabit Data Service, and ATM networks
- University Research Initiative—Managing Multiservice Networks

MMN—Loughborough

- Performance Monitoring and Measurement
- Researched and built a delay measurement tool
- Active Sender
- Used GPS for synchronisation
- Accurate to about $10\mu s$

Performance Monitoring

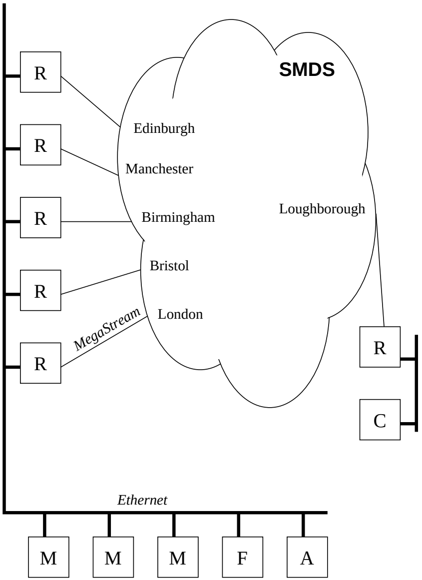

What causes performance problems?

- Routing misconfiguration
- Link or Node failure
- Aggressive Applications
    - Peer-to-peer, video streaming, online gaming etc
- Denial of Service attacks

Visualisation

- Tools to reduce working load of network operators
- FDV—Figurable Deformity Visualisation
- TMT—Trunk Monitoring Tool

FDV

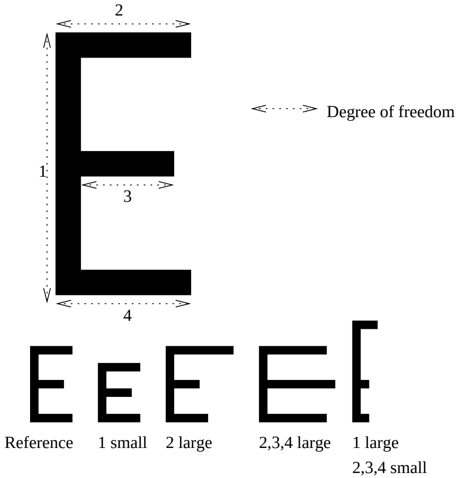
2
1
3
4
Degree of freedom
Reference
1 small
2 large
2,3,4 large
1 large
2,3,4 small

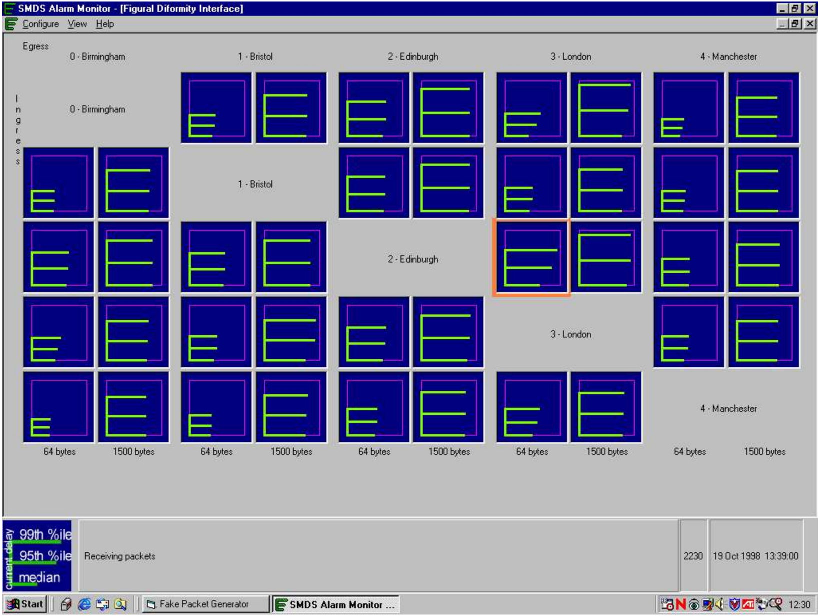
SMDS Alarm Monitor - [Figural Diformity Interface]
Configure View Help
Egress
Ingress
0 - Birmingham
1 - Bristol
2 - Edinburgh
3 - London
4 - Manchester
64 bytes
1500 bytes
current delay
99th %ile
95th %ile
median
Receiving packets
2230
19 Oct 1998 13:39:00
Start
Fake Packet Generator
SMDS Alarm Monitor ...
12:30

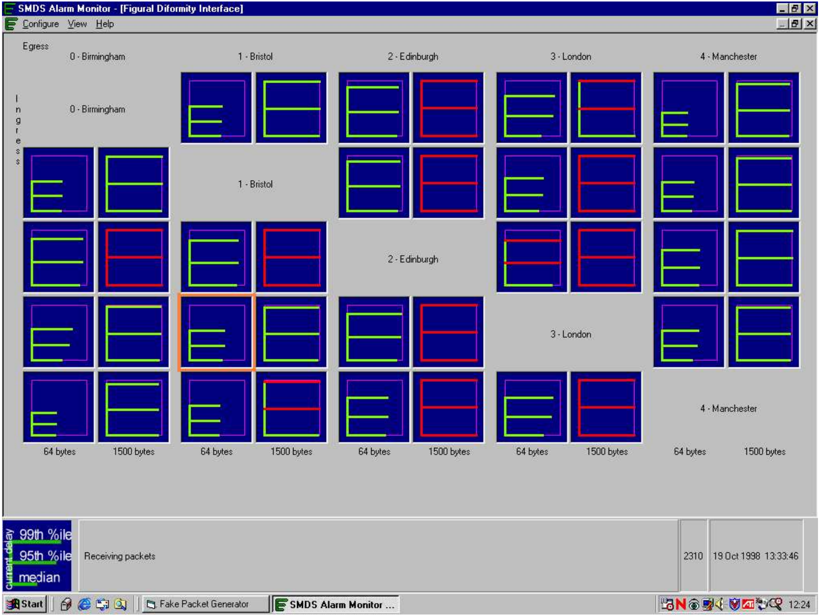
SMDS Alarm Monitor - [Figural Diformity Interface]
Configure View Help
Egress
0 - Birmingham
1 - Bristol
2 - Edinburgh
3 - London
4 - Manchester
Ingress
64 bytes
1500 bytes
current delay
99th %ile
95th %ile
median
Receiving packets
2310
19 Oct 1998 13:33:46
Start
Fake Packet Generator
SMDS Alarm Monitor ...
12:24

TMT

Trunk Monitoring Tool

- Uses SNMP to query trunk information from SMDS switches
- Presents this in a “single-look” view to operators.
- Deployed April 200

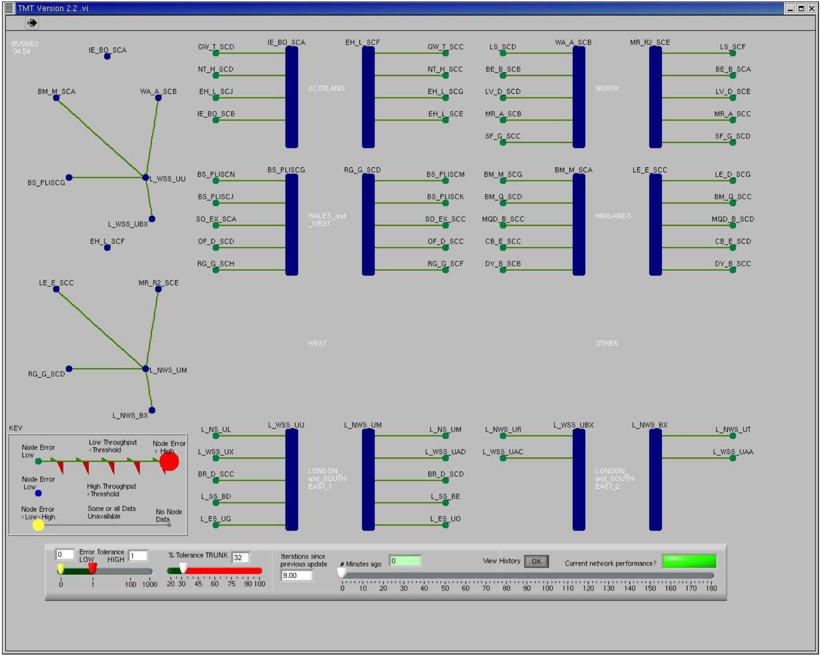
TMT Version 2.2 .vi
SCOTLAND
NORTH
WALES_and_WEST
MIDLANDS
WEST
OTHER
LONDON_and_SOUTH-EAST_1
LONDON_and_SOUTH-EAST_2
KEY
Node Error Low
Low Throughput <Threshold
Node Error > High
High Throughput >Threshold
Node Error >Low<High
Some or all Data Unavailable
No Node Data
Error Tolerance LOW HIGH
% Tolerance TRUNK
Iterations since previous update
Minutes ago
View History
Current network performance?

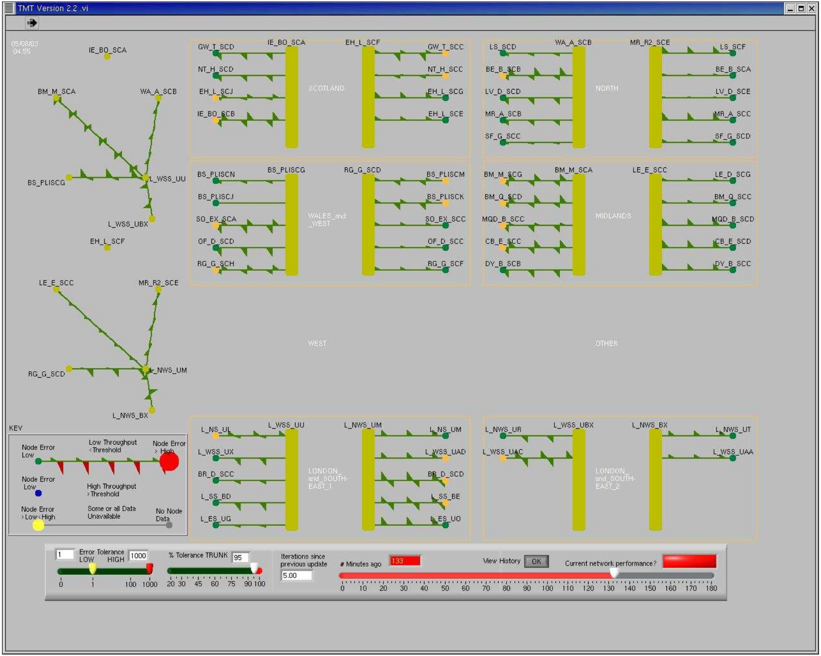
TMT Version 2.2 .vi
05/08/03
04:55
IE_BO_SCA
BM_M_SCA
WA_A_SCB
BS_PLISCG
L_WSS_UU
L_WSS_UBX
EH_L_SCF
LE_E_SCC
MR_R2_SCE
RG_G_SCD
L_NWS_UM
L_NWS_BX
SCOTLAND
NORTH
WALES_and_WEST
MIDLANDS
WEST
OTHER
LONDON_and_SOUTH-EAST_1
LONDON_and_SOUTH-EAST_2
GW_T_SCD
NT_H_SCD
EH_L_SCJ
IE_BO_SCB
GW_T_SCC
NT_H_SCC
EH_L_SCG
EH_L_SCE
LS_SCD
BE_B_SCB
LV_D_SCD
MR_A_SCB
SF_G_SCC
LS_SCF
BE_B_SCA
LV_D_SCE
MR_A_SCC
SF_G_SCD
BS_PLISCN
BS_PLISCJ
SO_EX_SCA
OF_D_SCD
RG_G_SCH
BS_PLISCM
BS_PLISCK
SO_EX_SCC
OF_D_SCC
RG_G_SCF
BM_M_SCG
BM_Q_SCD
MQD_B_SCC
CB_E_SCC
DY_B_SCB
LE_D_SCG
BM_Q_SCC
MQD_B_SCD
CB_E_SCD
DY_B_SCC
L_NS_UL
L_WSS_UX
BR_D_SCC
L_SS_BD
L_ES_UG
L_NS_UM
L_WSS_UAD
BR_D_SCD
L_SS_BE
L_ES_UO
L_NWS_UR
L_WSS_UAC
L_WSS_UBX
L_NWS_BX
L_NWS_UT
L_WSS_UAA
KEY
Node Error Low
Low Throughput <Threshold
Node Error >High
Node Error Low
High Throughput >Threshold
Node Error >Low<High
Some or all Data Unavailable
No Node Data
Error Tolerance LOW HIGH
1
1000
0 1 100 1000
% Tolerance TRUNK
95
20 30 45 60 75 90 100
Iterations since previous update
5.00
Minutes ago
133
View History
OK
Current network performance?
0 10 20 30 40 50 60 70 80 90 100 110 120 130 140 150 160 170 180

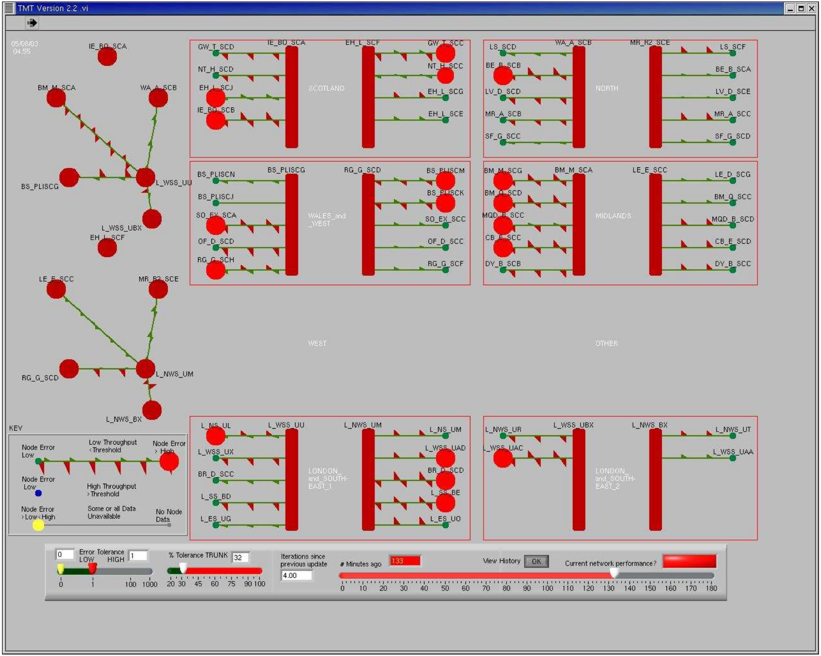
TMT Version 2.2 .vi
05/08/03
04:55
SCOTLAND
NORTH
WALES_and _WEST
MIDLANDS
WEST
OTHER
LONDON_ and_SOUTH-EAST_1
LONDON_ and_SOUTH-EAST_2
L_WSS_UU
L_WSS_UBX
L_NWS_UM
L_NWS_BX
KEY
Node Error Low
Low Throughput <Threshold
Node Error > High
High Throughput >Threshold
Node Error >Low<High
Some or all Data Unavailable
No Node Data
Error Tolerance
LOW
HIGH
% Tolerance TRUNK
Iterations since previous update
Minutes ago
View History
OK
Current network performance?

Exceptions

Interesting Network Events, detected by:

- Manual
- Rule-based
- Neural networks

All based on simple statistics, max in day, min in day, mean, max - min, variance etc

Other Monitoring Projects

- RIPE-NCC—Monitoring (mostly) European Delays
- SPRINT (US)—Monitoring for Traffic Engineering
- NLANR—Traceroute/ping delays
- Waikato (NZ) DAG hardware traffic capture
- Cambridge/Loughborough (EE) passive monitoring
- *new* UKLIGHTmas*(t)*

What to do next ...

- Can statistics/mathematics improve such displays?
- Can we predict Internet performance like the weather?
- How do we model?

The rest of this talk

- Motivation
- Traffic modelling review
- Mixing Weibull distributions
- Expectation Maximisation algorithm
- Experiments and results
- Applications and discussion

Motivation

The need to model network performance:

- Metrics to define network performance
- Low-level quantities: delay and loss
- End-to-end network performance status
- Packet probes such as ping or one-way delay UDP packets

Previous work:

Traffic modelling and delay distributions:

- Network traffic shows self-similarity and long-range dependency.
- Current traffic strategies search for models compliant to these empirical properties: fBm, fARIMA, FSD, etc.
- When inputting such traffics into routers, the queue distribution exhibit heavy-tail distributions. Such distribution can be approximated to Weibull for the particular case of fBm.
- Such result has been previously validated in a single hop scenario.

Previous work:

Traffic modelling and delay distributions:

- Our aim is to model multiple-hop (or end-to-end) delays with a combination of several Weibull distributions.

Mixing Weibull distributions:

The Weibull distribution $p(x|r,s) = \frac{sx^{s-1}}{r^s}\exp(-\left(\frac{x}{r}\right)^s)$

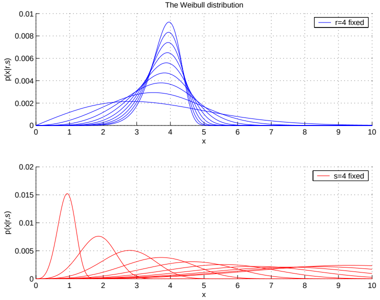

- r is concerned with the mode location.
- s is related to tail behaviour.

Mixing Weibull distributions:

Problem statement:

- Let us assume we are given a sample of N delay measurements $\mathbf{x} = [x_1, .., x_N]$, which are supposed to be drawn from M Weibull distributions: $[p(x|\theta_1), .., p(x|\theta_M)]$
- The result is: $p(x|model) = \sum_{j=1}^{M} \alpha_j p(x|\theta_j)$
- $\alpha_j =$ weight of the j-th component of the mixture. Obviously, $\sum_j \alpha_j = 1$
- $\theta_j = [r_j, s_j]$ shape and scale parameters of the j-th Weibull distribution
- Finding α and θ appropriate to best fit delay histograms represented by the measurements sample $\mathbf{x}$

Expectation Maximisation

- To proceed, second random variable y, referred to as labels, is necessary to complete the problem formulation.
- $p(y_i = j|x_i, \Theta)$ = the probability of data x_i being drawn from the j-th component of the mixture. Obviously,
  - $p(x_i|y_i = j, \Theta) = p(x_i|\theta_j)$, and
  - $p(y_i = j|\Theta) = \alpha_j$
- With this formulation EM defines an iterative procedure to obtain the maximum likelihood estimates, based on two steps:
  - E-step: $Q(\Theta, \Theta^{(t)}) = E[\log L(\Theta|x, y)|\mathbf{x}, \Theta^{(t)}]$
  - M-step: $\Theta^{(t+1)} = \arg\max_\Theta Q(\Theta, \Theta^{(t)})$

Mixing Weibull distributions:

Computing EM

- Expanding E-step:

$$
\begin{aligned}
Q(\Theta, \Theta^{(t)}) = \; & \textstyle\sum_{j=1}^{M} \sum_{i=1}^{N} \big(\log p(x_i|\theta_j) \big) p(y_i = j|x_i, \Theta^{(t)}) \\
& + \textstyle\sum_{j=1}^{M} \sum_{i=1}^{N} \big(\log \alpha_j \big) p(y_i = j|x_i, \Theta^{(t)})
\end{aligned}
$$

- Maximising:

$$
\frac{\partial Q(\Theta, \Theta^{(t)})}{\partial \alpha_j} = 0
$$

$$
\frac{\partial Q(\Theta, \Theta^{(t)})}{\partial \theta_j} = 0
$$

Mixing Weibull distributions:

EM applied to mixtures of Weibull distributions

1. Computing parameters:
$$\alpha_j = \frac{1}{N}\sum_{i=1}^{N} p(y_i = j|x_i, \Theta)$$
$$r_j = \left(\frac{\sum_{i=1}^{N} x_i^{s_j} p(y_i{=}j|x_i,\Theta)}{\sum_{i=1}^{N} p(y_i{=}j|x_i,\Theta)}\right)^{1/s_j}$$
$$s_j = \frac{\sum_{i=1}^{N} p(y_i{=}j|x_i,\Theta)}{\sum_{i=1}^{N}\left(\frac{x_i^{s_j}}{r_j^{s_j}} - 1\right)\log\left(\frac{x_i}{r_j}\right) p(y_i{=}j|x_i,\Theta)}$$
2. Updating hidden probs:
$$p(y_i = j|x_i, \Theta) = \frac{\alpha_j p(x|\theta_j)}{\sum_{k=1}^{M} \alpha_k p(x_i|\theta_k)}$$

Convergence speed - Initialisation

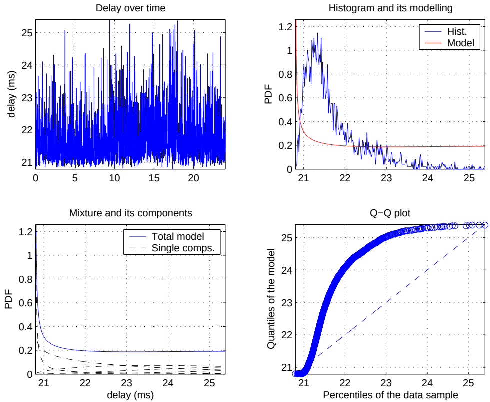

Convergence speed - After 1 iteration

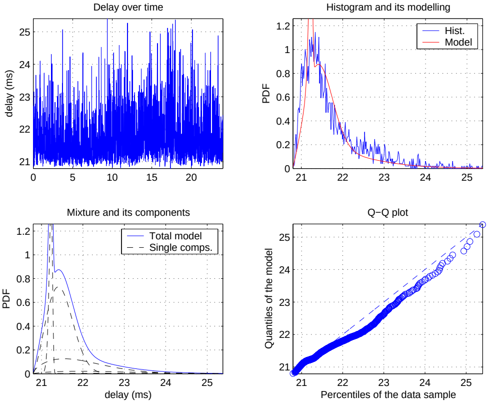

Convergence speed - After 2 iterations

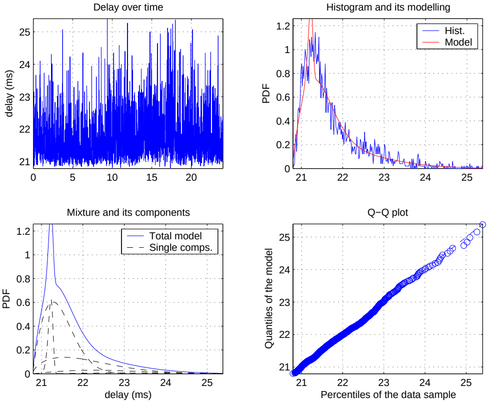

Mixing Weibull distributions

Convergence speed - After 3 iterations

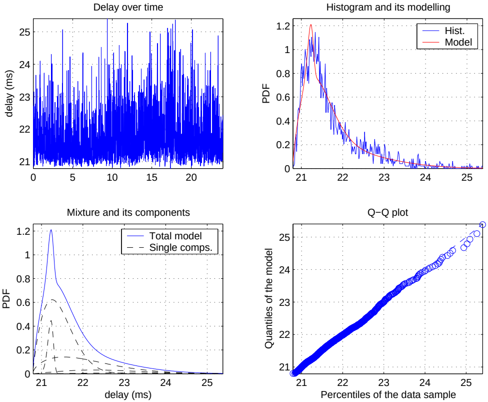

Convergence speed - After 4 iterations

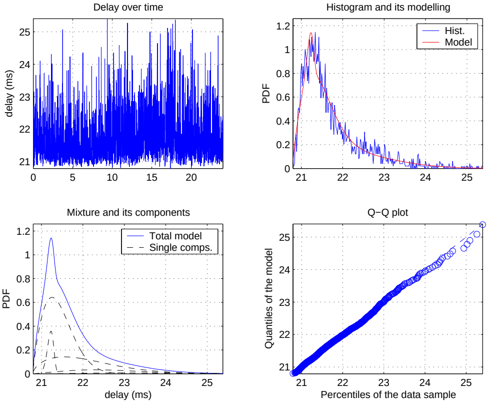

Convergence speed - After 5 iterations

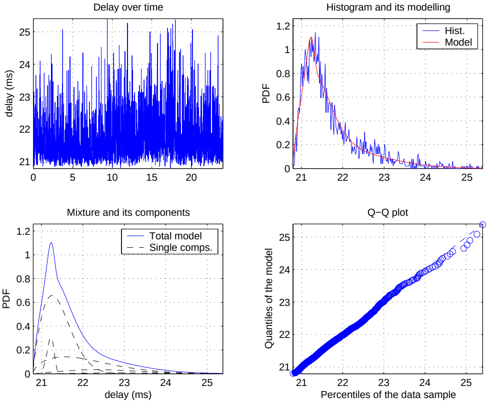

Mixing Weibull distributions

Convergence speed - After 10 iterations

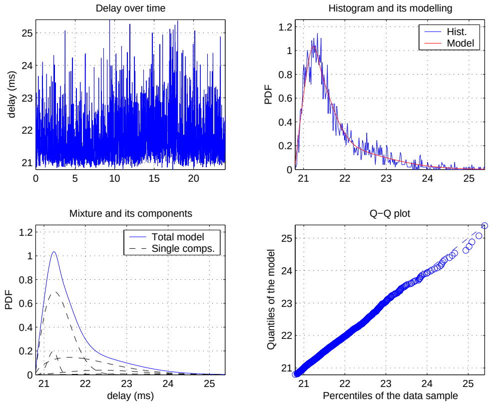

Convergence speed - After 15 iterations

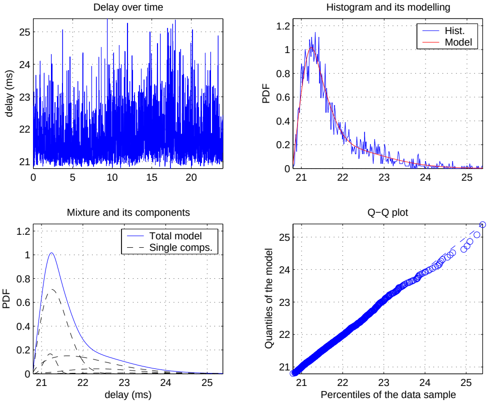

Convergence speed - After 25 iterations

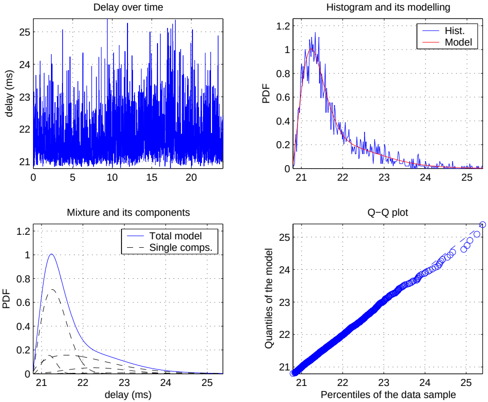

Convergence speed - After 50 iterations

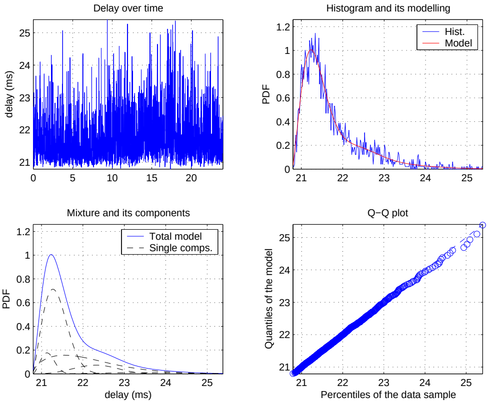

Convergence speed - After 100 iterations

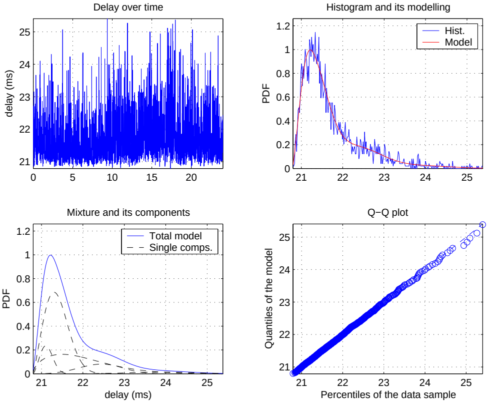

Measurement testbed

- The following delay measurements, provided by RIPE NCC[1], have been utilised for this experiments.

$$\left.\begin{array}{l}\text{245 24-hour exps.}\\ \approx 3000 \text{ meas. per exp.}\end{array}\right\} \text{Total:} \quad \approx 700,000 \text{ measurements}$$

- GPS accuracy $\approx$ few hundred of nanoseconds error.
- Matching error $= \dfrac{\sqrt{\sum(\text{hist-model})^2}}{\sum \text{hist}} \times 100\%$

[1] http://www.ripe.net

Full experiments model validation

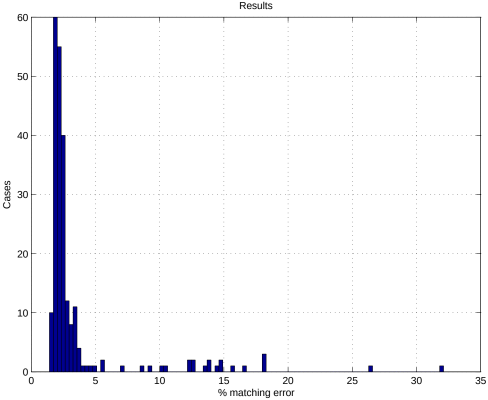

Example of a five Weibull matching result

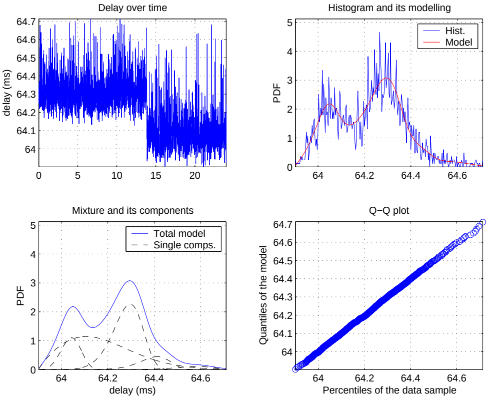

Example of Parameter Evolution

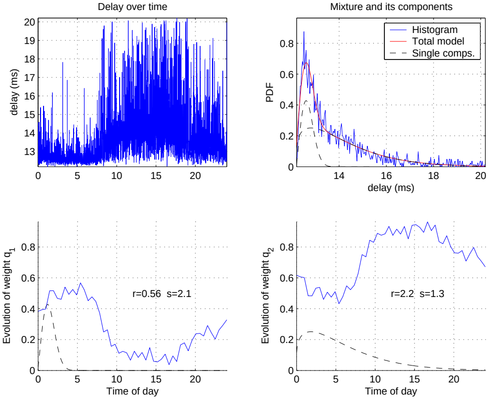

Discussion

Two main conclusions arise from this work:

- A combination of Weibull distributions look very suitable to match end-to-end delay histograms.
- The Weibull parameters impact the appearance of the Weibull distribution.
  - r is related to the location of the mode/maximum/peak for that particular Weibull component.
  - s concerns tail behaviour: the smaller the slower the tail decays.
- Expectation Maximisation is a suitable algorithm to find the parameters defining such model, both easily and optimally.

Applications

Performance related applications:

- Traffic engineering.
- Fault tolerance and troubleshooting.
- Provisioning.
- Admission control.
- ...

Thanks

- Other researchers at Loughborough especially: David Parish, Omar Bashir, Mark Sandford, Antony Pagonis.
- Jose's PhD is a Loughborough CS Scholarship.
- RIPE for 1 year's $(75\mathrm{GB})$ measurements.