Networking with Revd Bayes

Derek McAuley
Senior Vice President
Director, Marconi Labs Cambridge

Introduction

- P(H|D) = P(H)P(D|H) / P(D)
- H the Hypothesis
- P(H) – the “Prior” probability
- Observe data D

Hypothesis “Bayes is dead”

- P(H) .9 from this picture

Marconi

Make an observation

Conclusion

- P(D|H), say .5
- P(D|H'), say .01
- P(D) hence .451

Bayes Clear	Probability		Odds	
P(H)=	.9	Ω_0=	9.000000000	
P(D|H)=	.50	LR=	50	
P(D|H')=	.01			
Compute P(H|D)	0.997782705	Ω_1 =	450.0000000	

- Posterior P(H|D) is .99778..
- Conclusion “Bayes is almost certainly dead”

Marconi

Hardware fails

…a philosophical problem since 1328

Buridan, Jean (1300-1358), French Scholastic philosopher, who held a theory of determinism, contending that the will must choose the greater good. Born in Bethune, Buridan was educated at the University of Paris, where he studied with the English Scholastic philosopher William of Ockham. After his studies were completed, he was appointed professor of philosophy, and later rector, at the same university. Buridan is traditionally, but probably incorrectly, associated with a philosophical dilemma of moral choice called "Buridan's ass." In the problem an ass starves to death between two equally alluring and equidistant bundles of hay because it has no rational basis for preferring one bundle over the other.

It's not just bad engineering

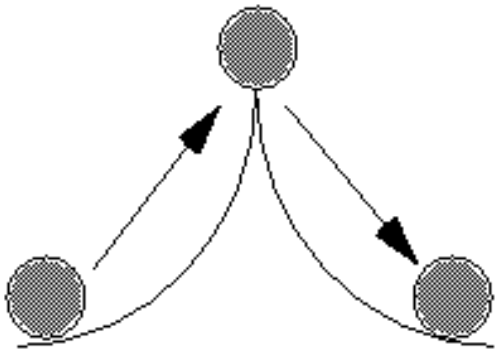

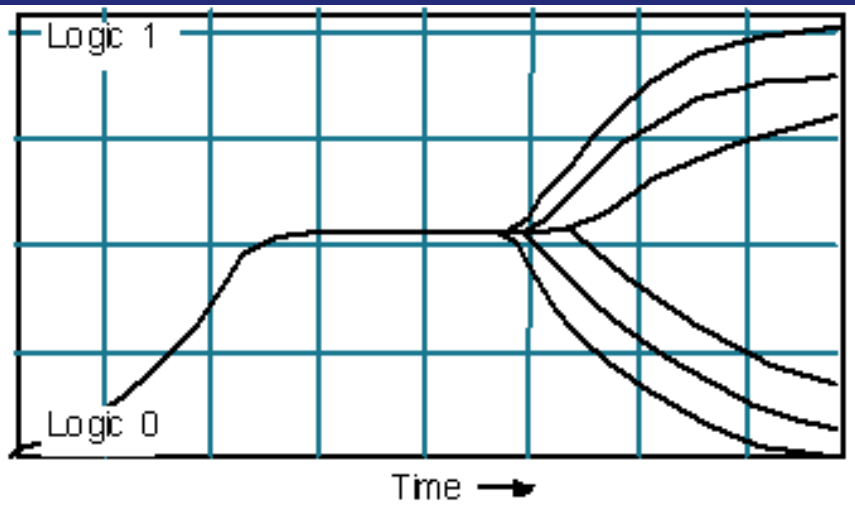

Marconi

Networks drop packets

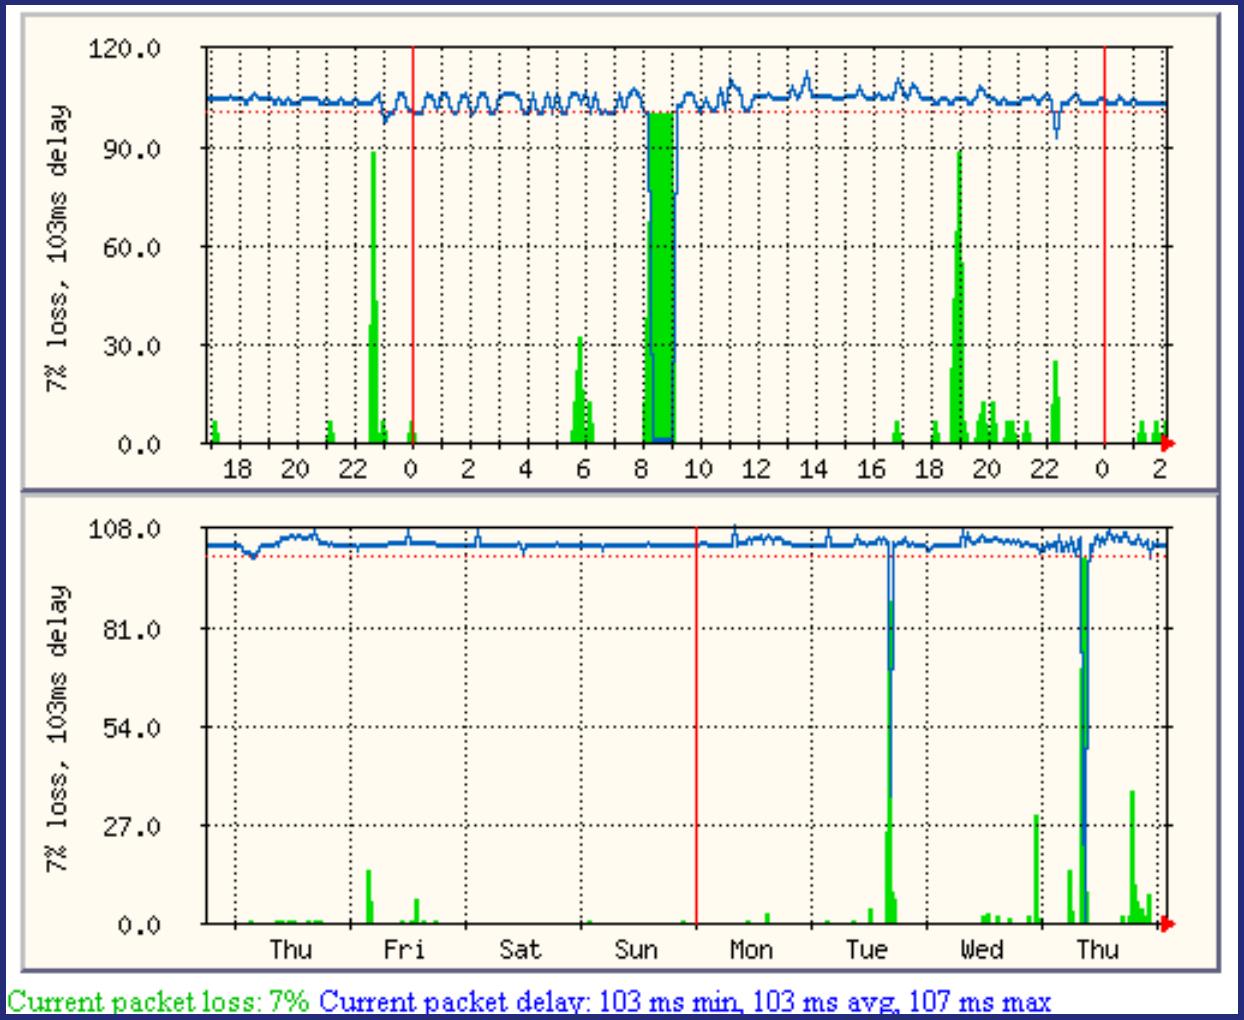

Marconi

Consider a well known game…

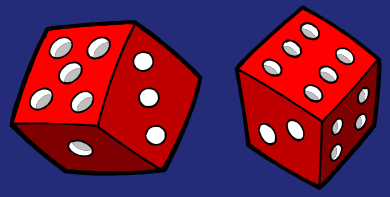

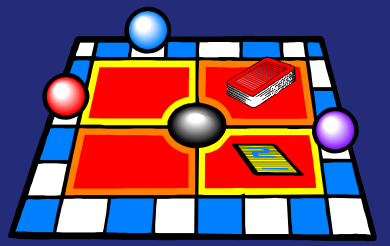

A board of 41 squares
Moves based on the roll of two dice
Some rules…

Marconi

Our Markov Matrix

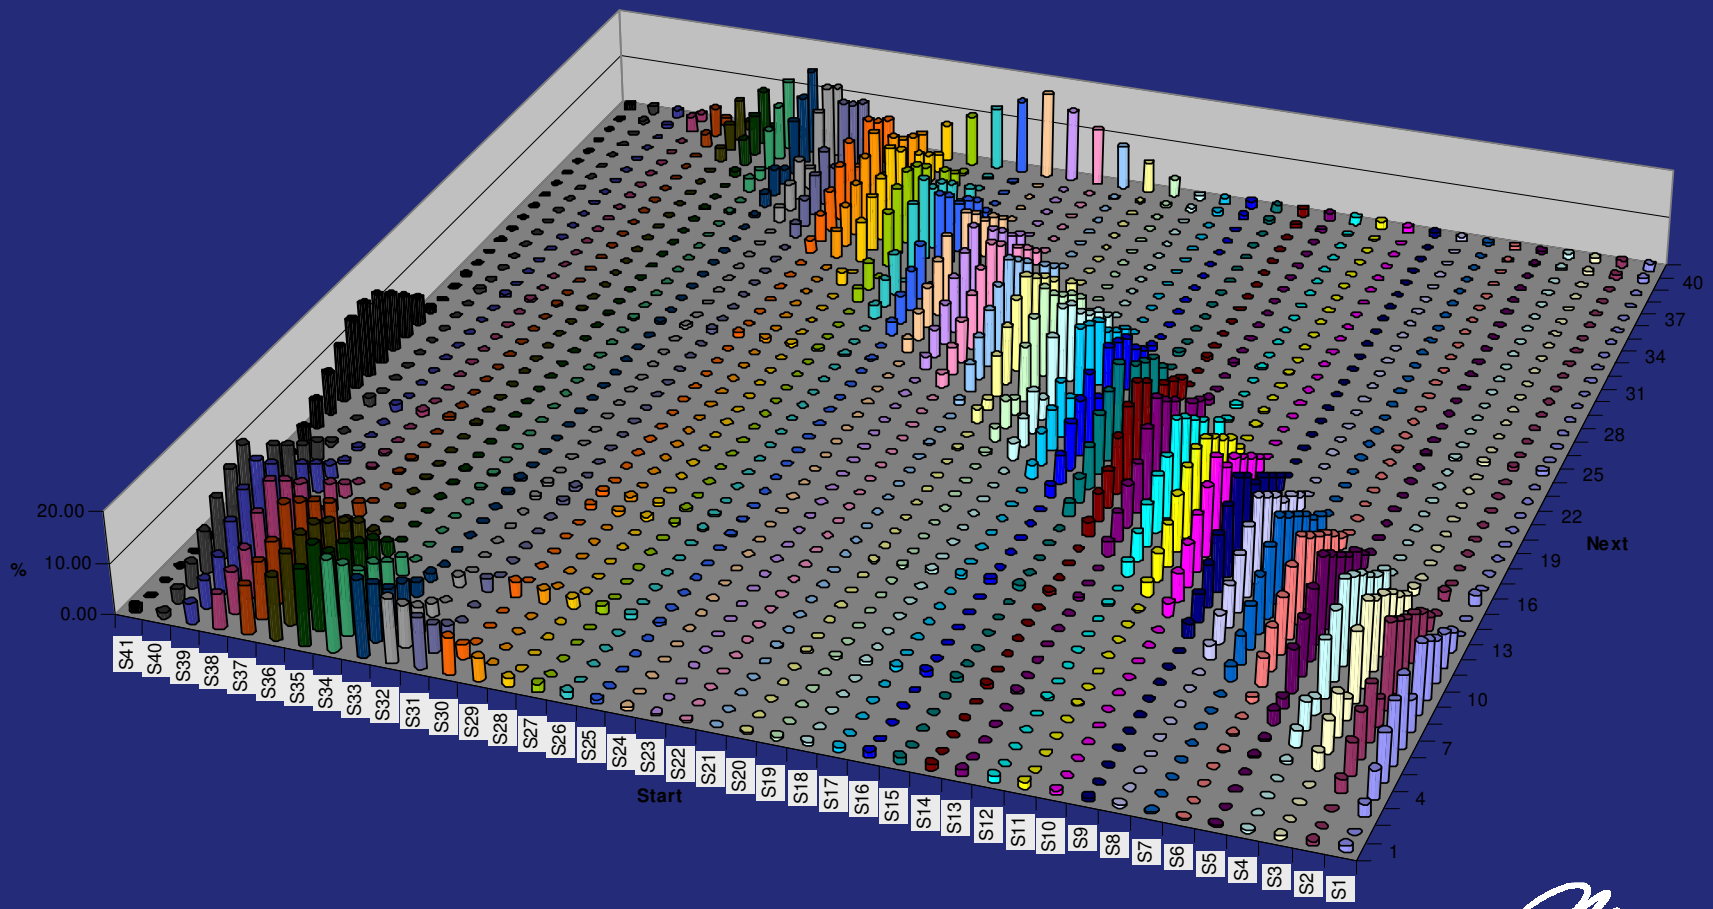

Marconi

The Autonet Skeptic

Roddeheffer & Schroder, DEC SRC, 1991:

"The skeptic limits the failure rate of a link by delaying its recovery if it has a bad history."

Requirements:

- A link with a good history must be allowed to fail and recover several times without significant penalty.
- In the worst case, a link's average long-term failure rate must not be allowed to exceed some low rate.
- Common behaviors shown by bad links should result in exceedingly low average long-term failure rates.
- A link that stops being bad must eventually be forgiven its bad history.

Marconi

Why learning

Our example had rules, but maybe:

- We don't know the rules
- We can't work out the rules
- There are too many rules
- The rules change (argh!!)

Examples:

- Internet Router configurations
- Link failures
- System log files
- Napster

Marconi

Argh!!! The rules might change

Machine learning is applied to lots of non-stationary data:

Ælfred kyning hāteð grētan Wærferð biscep his wordum luflīce ond frēondlīce; ond ðē cȳðan hāte ðæt mē cōm swīðe oft on gemynd, hwelce wiotan iū wæron giond Angelcynn, ægðer ge godcundra hāda ge woruldcundra;

…it's all about timescales

Marconi

QoS the easy way

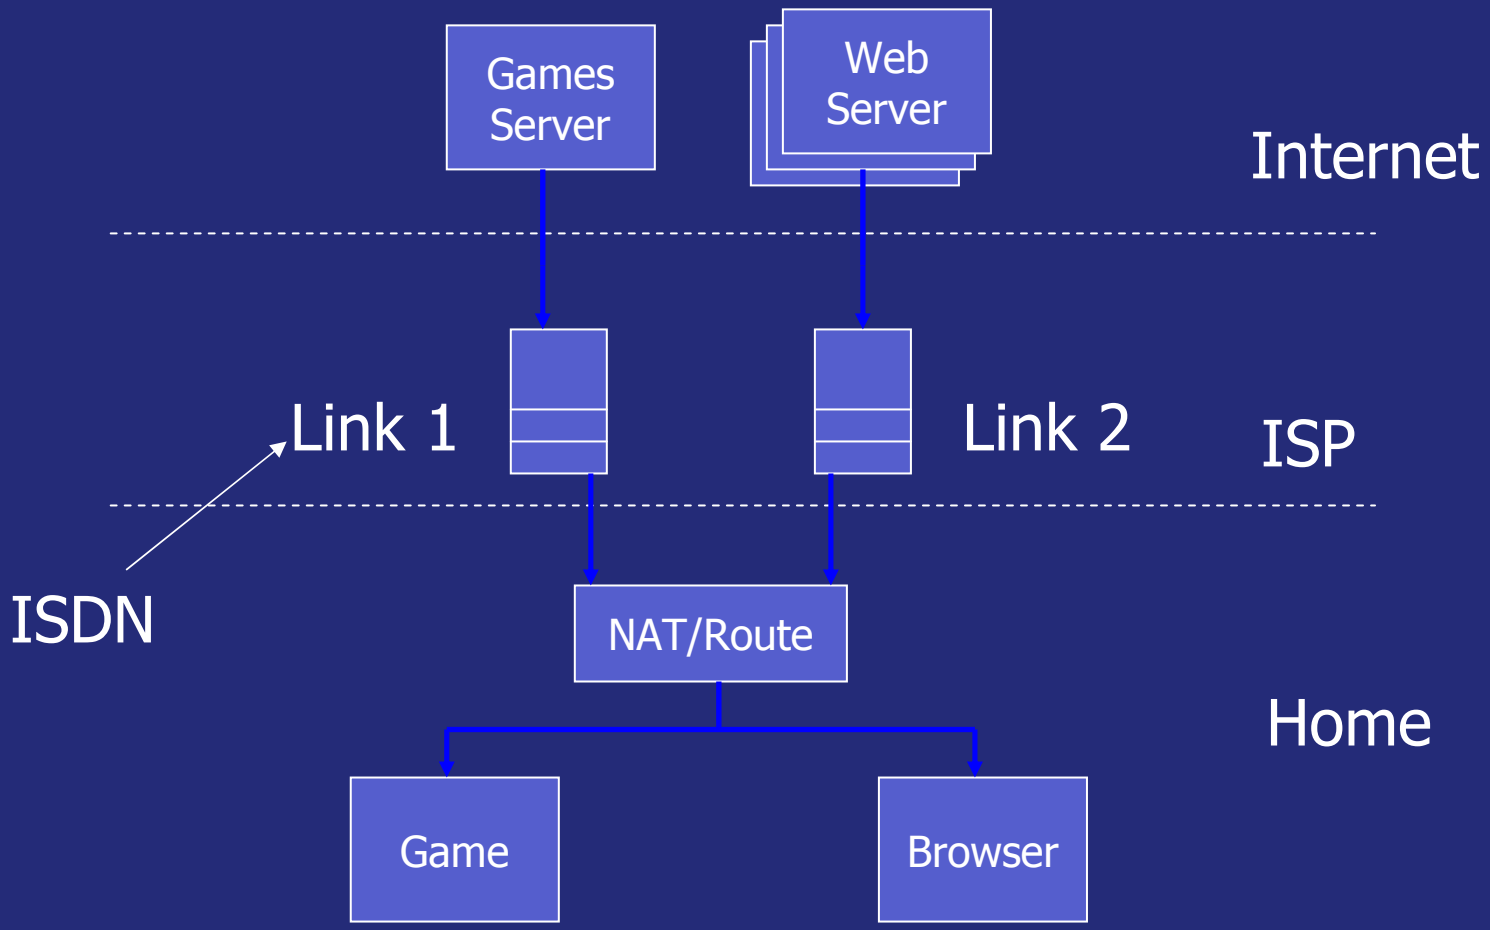

Marconi

RTTs and history

e.g.

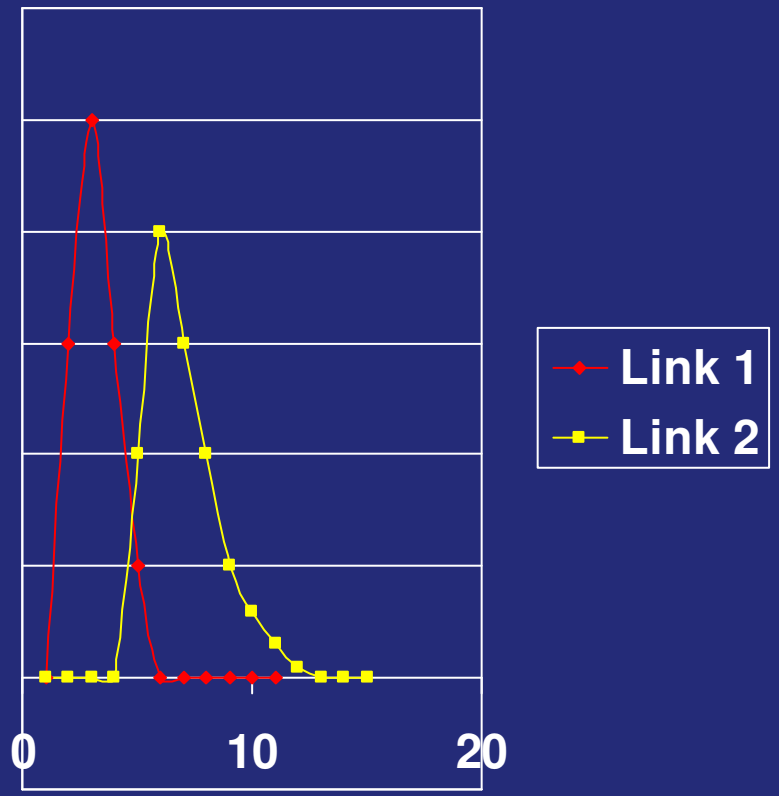

Keeping history (like CM) needs some thought

E.g. Dynamic routing

- Might need clustering of RTT history before use
- Change estimation algorithm to depend on RTT classification

Marconi

Link reliability

Say we set out to measure link reliability

- Track up/down events and link BER
- Might try to allocate to “classes” based on “type”
- Or … might try to learn “classes”
  - Might find a T1 link from vendor A and vendor B differ radically
  - After all, OS TCP implementations do!

Marconi

Silly routing example

- Now presume for a class of links we have a measured reliability of “4 nines” – in any given 30 second interval there is a 99.99% chance the link works OK.
  - Prior P(H) = 0.0001
- When we receive a LSA down what should we do?
  - What is P(H|D)?
- Understanding that “pathologies” exist in some BGP implementations:
  - Erroneous LSA down/up pair due to transient or load
  - Bad software (never surely!)
  - Operational issues
  - Investigate influence of P(D|H’)….

Marconi

Two (or three) eyes are better than one!

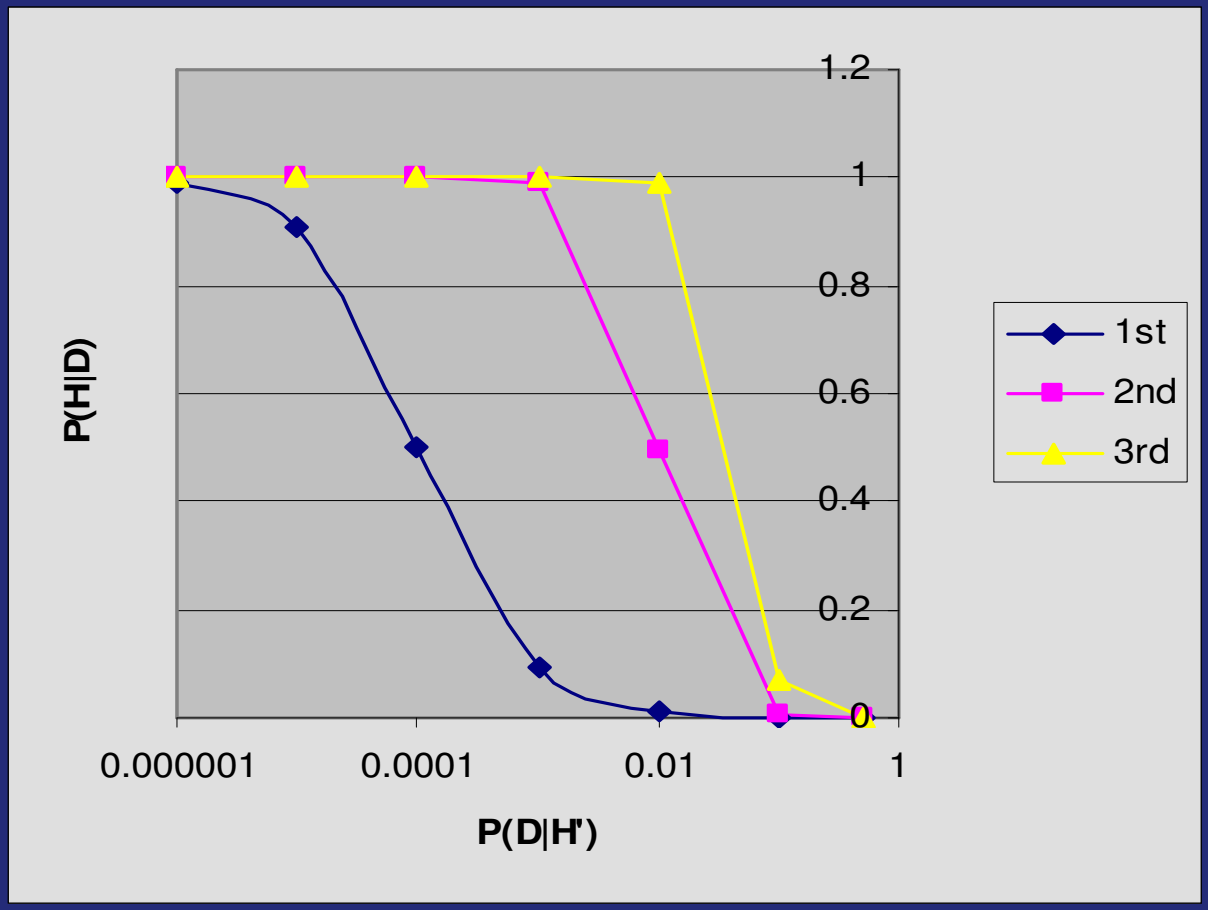

Confidence after 1, 2 and 3 LSA (independent) down messages.

Marconi

Enough Bayes already

Marconi Labs, Cambridge; objective:

"Enhance technological competence and competitiveness of Marconi"

- through world-class research in communications and internet technology
- sponsorship of appropriate research within Univ of Cambridge and partner universities,
- vigorous technology *transfer*.

Marconi

Hot Topics

Optical switching and routing

- Electronic/Optical trade-offs

Network Management

- What does it take to manage a lump of glass

Network Modelling

- Lots of measurements

Machine Learning applied to networks

- Factor uncertainty into decision processes

Marconi

University Projects

In place

- 3 PhD students Security architecture
- Optical topology design
- BGP / ECN / Congestion pricing
- GaN properties

Coming Soon

- 4 more PhDs
- Async clocking
- Optical packet node
- …and a professor ☺ or ☺ ☺

Marconi

Cambridge from the air

Marconi

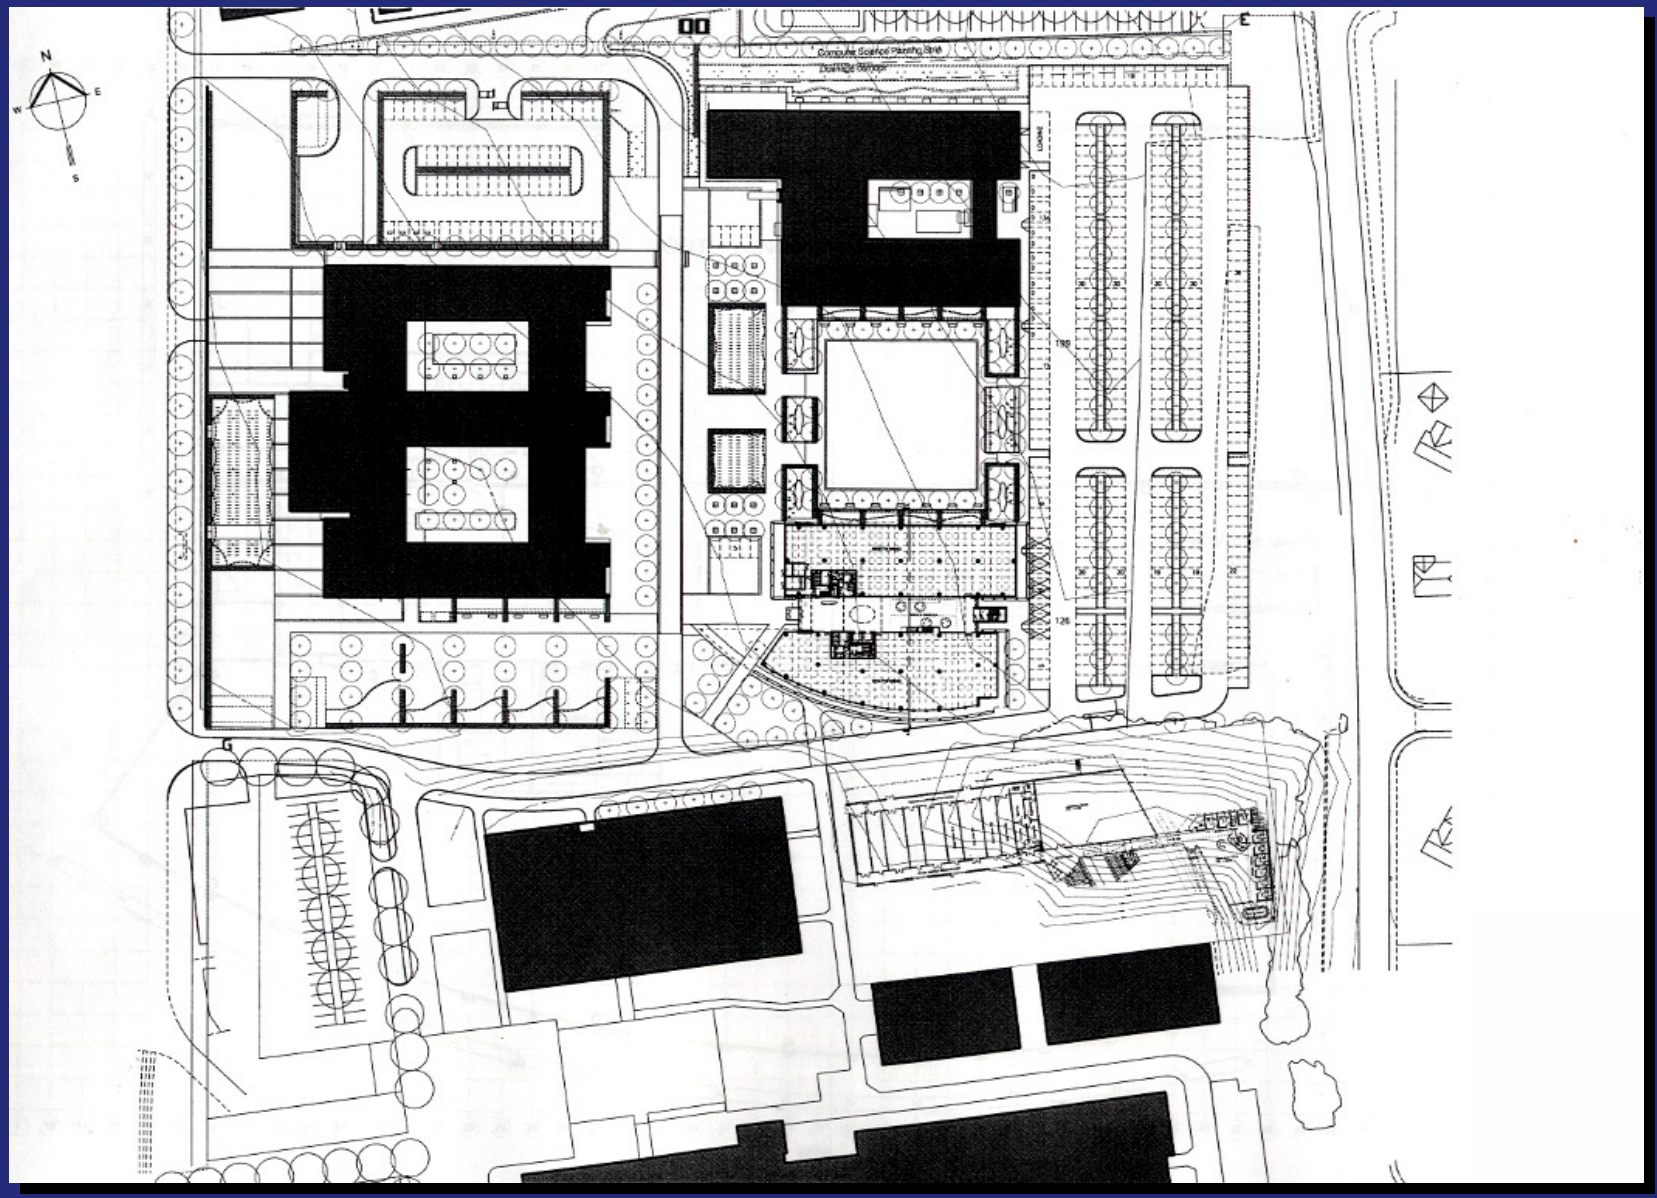

Marconi

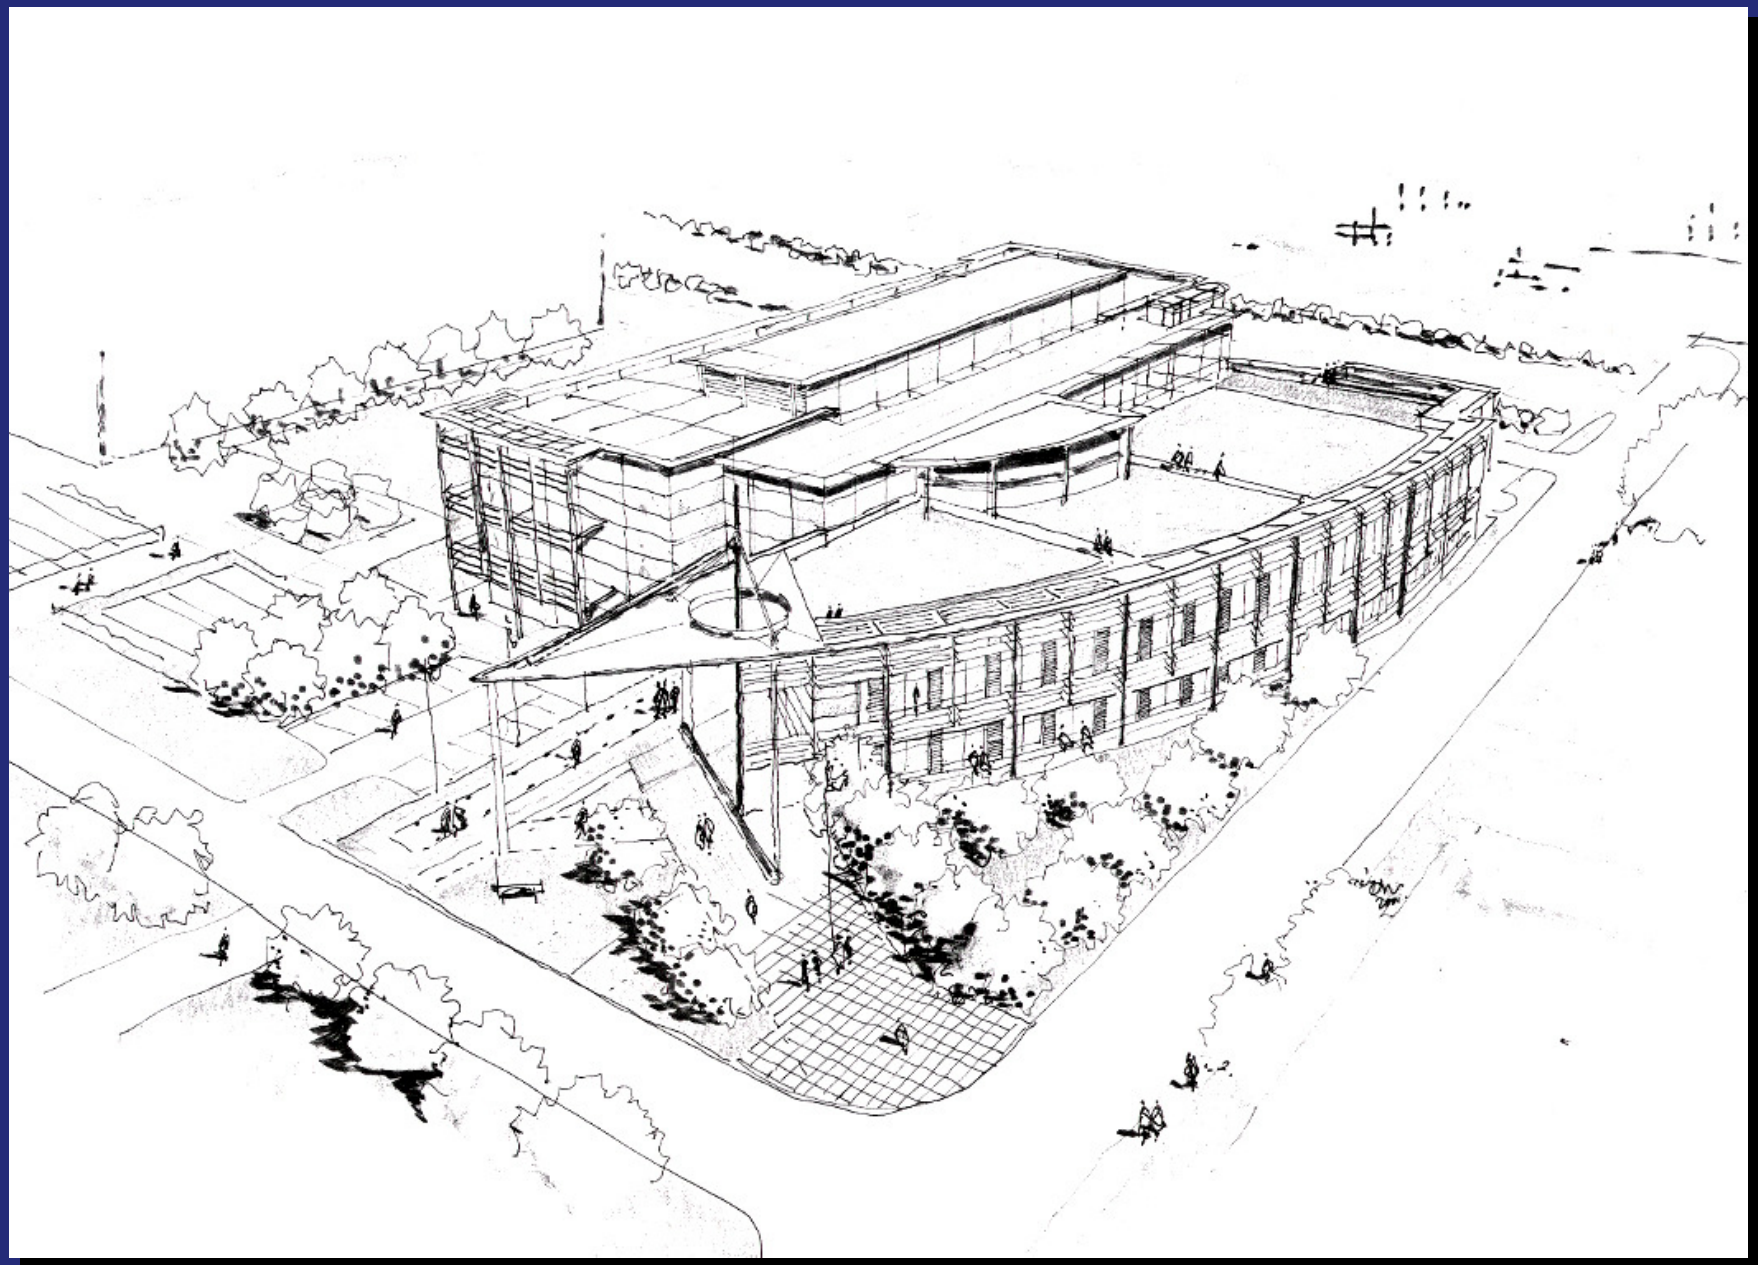

Marconi