


# Future Directions ACM Sigmetrics

Albert Greenberg  
Principal Researcher  
Microsoft Corporation

# Agenda

- An example
- Some learnings
- Hard problems

# Robust IP Network Design, Engineering, and Operations

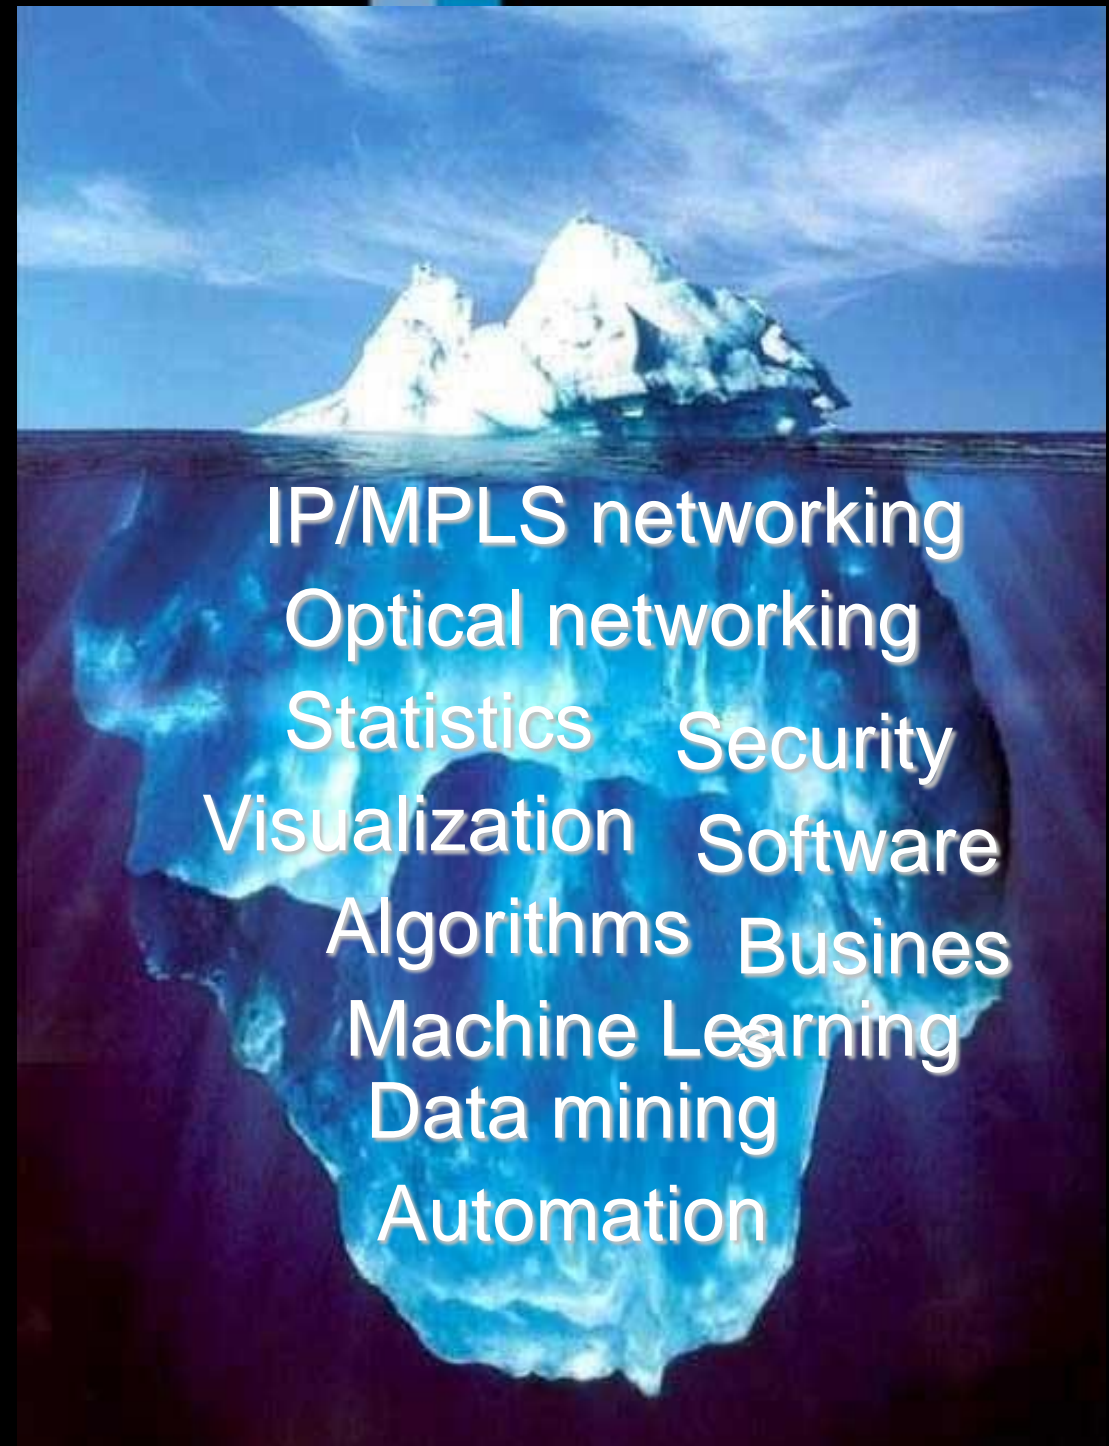


# Oops

- We don't know the topology
- We don't know the risks
- We don't know the traffic matrix
- We don't know how to optimize routing

# Research


- Risk?
- Topology?
- **Traffic?**
- Optimization?



# Traffic Analysis Research

(Long and Windy Road to Effective Traffic Matrix Estimation)

## • Netflow?



- Invented to support flow-switching. Oops

- Partial support. Oops

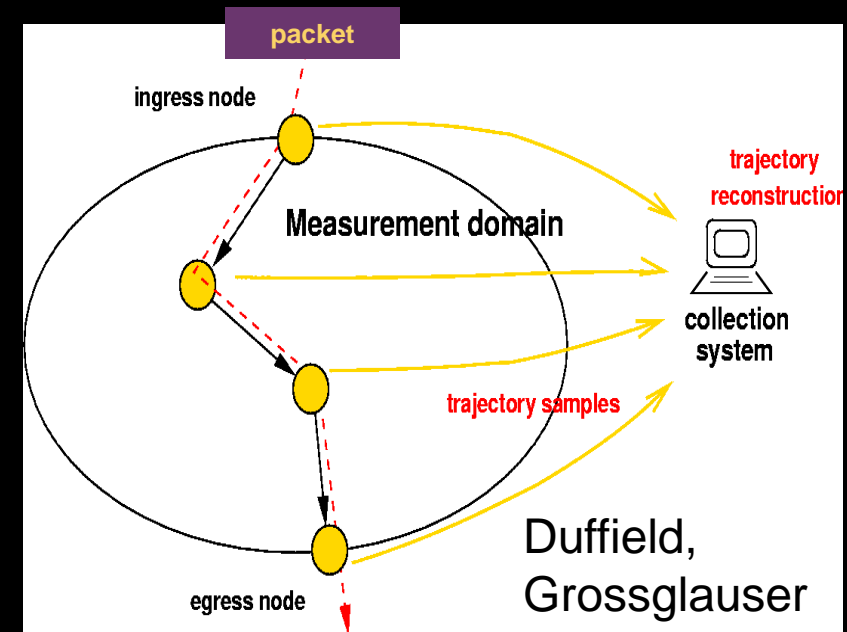
## • Something Better?

- Measurement chip?

- Market realities. Oops

- Trajectory sampling!

- Implement on every linecard? Oops





# A Killer App Appears – Survivability

- All services/networks converge to IP

- SLAs and expectations escalate


- Survivability

- Failure in LA. Maintenance scheduled at 2 AM in NY.

- Impact of the maintenance?
    - Go or no go?

- Tomo-gravity invented

- Proven in the field
  - Proven on the blackboard
    - Minimizing Mutual Information (Donoho)
  - Used every day



Zhang et al.

# Learnings

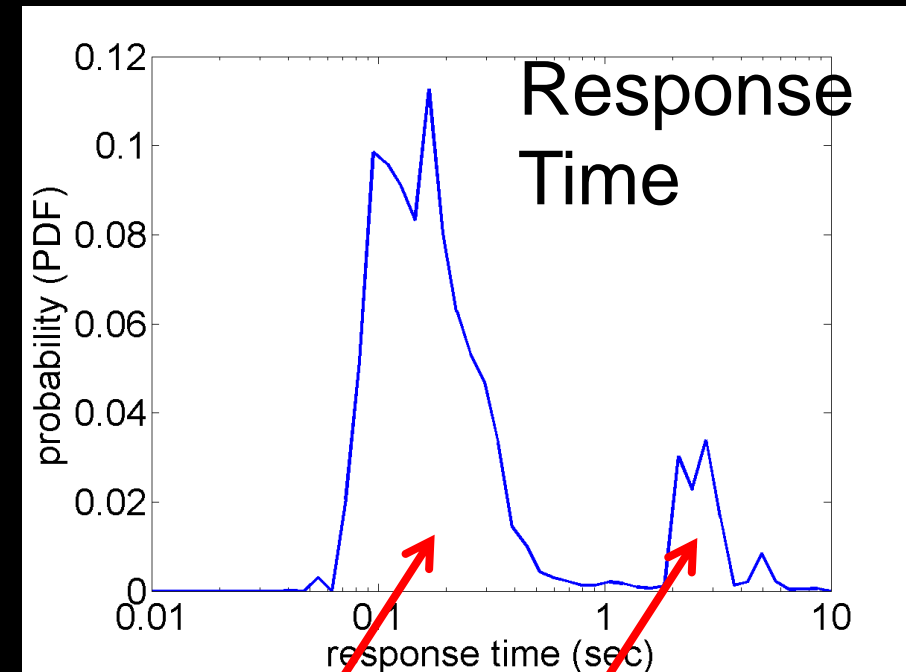
- Big payoffs from abstracting simple, fundamental questions
  - What inputs/precision do we really need?
  - What model and prediction?
  - What coarse grained action?
- Helpful pressure and “clue” from reality were key ingredients
  - At some point, we had to solve it
    - Parallel programming today; can no longer crank up clock speed
- **Importance of DATA, DATA, DATA**
- Computer Science, with help welcomed from any other domain that might provide insight!
  - Algorithms, Statistics, ...
  - A lot of code and creativity devoted to instrumentation and data management

???



# Harder Problems Remain...

- What is Normal? What is correct?
  - Arif's wonderful remarks on HP Delphi & workload characterization...
- What anomalies are user perceptible?
- Where is the problem located?
  - Among the 1000s of components broken or troubled?
  - Hardware and software (correctness) problems
- How can it be mitigated?
- How can it be repaired?



Baseline

Why???

# Hard Problems

- Automation and reliability for **large scale** critical infrastructure
  - Networks, data centers, enterprises, e-commerce ...
    - Hardware and **software** artifacts
    - Unpredictable demands, potential failure modes, ...
- Priorities: Agility, Robustness, Managed cost
  - Oops. All three often not achievable today.
- Oops. Agile and not robust
- Oops. Robust and not agile (cannot innovate)
  - Innovation stopped dead to maintain reliability in a corner of the design space ...
    - Not good for service providers, vendors, users

# Modeling

- Modeling is huge in systems and capacity planning
  - <http://www.microsoft.com/business/dsi/default.mspx>
- Data driven modeling is huge in running systems
  - Self-adaptive, resilient, secure
    - Fine grained control
    - Coarse grained management
- Manage scale and complexity dilemmas
  - Better test before deployment
  - Instrumentation default on & test after deployment
    - Unless tested in the field, you really don't know ...