

Network Models: Do They Matter?

Matthias Grossglauser

School of Computer and Communication Sciences
EPFL



A Viable Field of Science & Engineering has Accepted Models

- **Examples:**

- Signal processing
 - Fourier transforms, wavelets
- Databases
 - Relational and object-oriented DBs
- Algorithms
 - Turing
- Artificial intelligence & machine learning
 - Bayes' nets, neural nets
- Communications
 - Channels

- **Importance:**

- Indicates stable research agenda
- Didactic tool

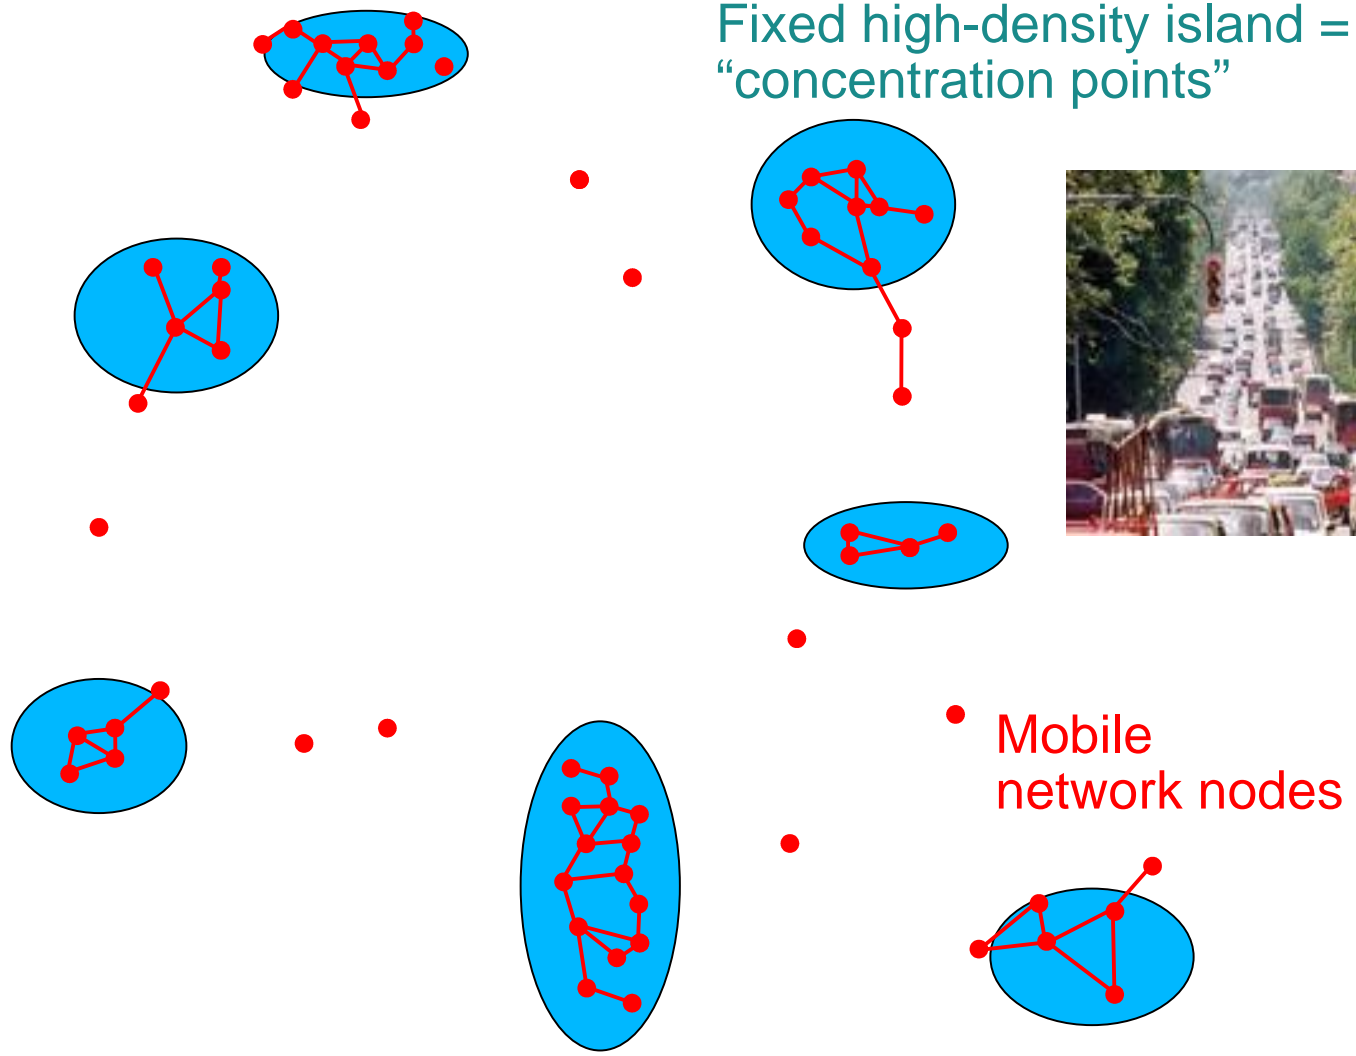
What Models in Networking?

- Congestion control
 - Borrowed from control theory
- Routing
 - Borrowed from graph theory + algorithms
- Design
 - Borrowed from optimization
- Peer-to-peer and mobile wireless
 - Borrowed from stochastic geometry, random graphs
- QoS
 - Borrowed from queuing theory, scheduling
- Networking
 - Not really a proper field?
 - Will it dissolve into other disciplines as the focus changes?

Engineering Principles vs Models

- We have design and operating principles proper to our field
 - Soft state, end-to-end, stateless, in-band vs out-of-band control, data/control plane, scaling, self-organization, autonomous, adaptive,...
 - The “-ilities”
- But:
 - Just engineering “best current practices”?
 - How do you teach these? How do we reason about them?
 - Qualitatively different from e.g. software design?
- Need more domain-specific models

New Models 1: Explicit Heterogeneity



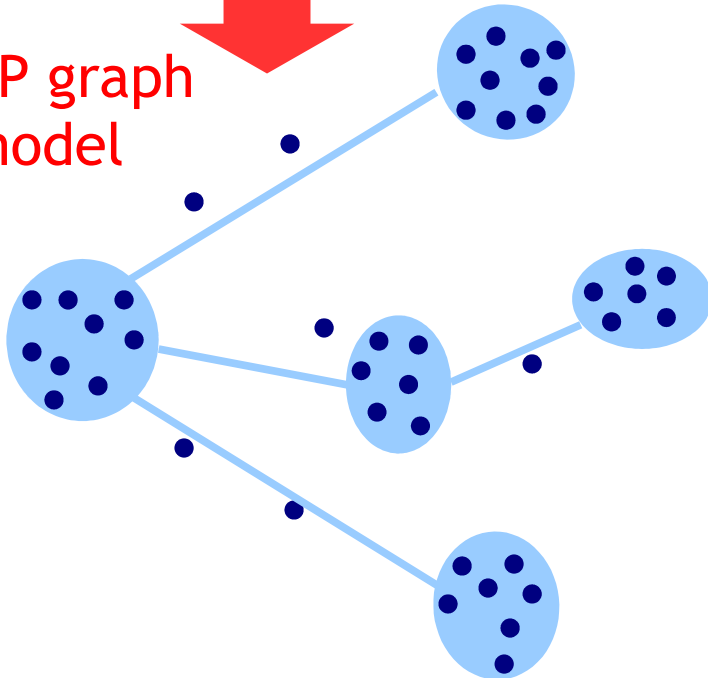
- Real networks vs statistical physics

Mobility Model with Stable Concentration Points

Taxis in Warsaw



CP graph model



- CP graph:
 - Vertices = CPs
 - Edges = flows of nodes between CPs
- Worst-case connectivity & mobility model:
 - Only nodes at the same vertex can communicate
 - A node in transit can not communicate with anyone
 - Nodes perform independent random walks on CP graph

Heterogeneous Traffic in Space

- Location-driven services, sensor-actuator networks, smart spaces, vehicular applications, ...
 - No uniform traffic from anywhere to anywhere -> locality
 - Predictable patterns?
 - Impact on routing, dissemination, caching, etc.
- Analogy: Liben-Nowell et al.'s study of geography of social networks
 - Robust law describing $P(u,v \text{ connected})$ as $f(d(u,v))$



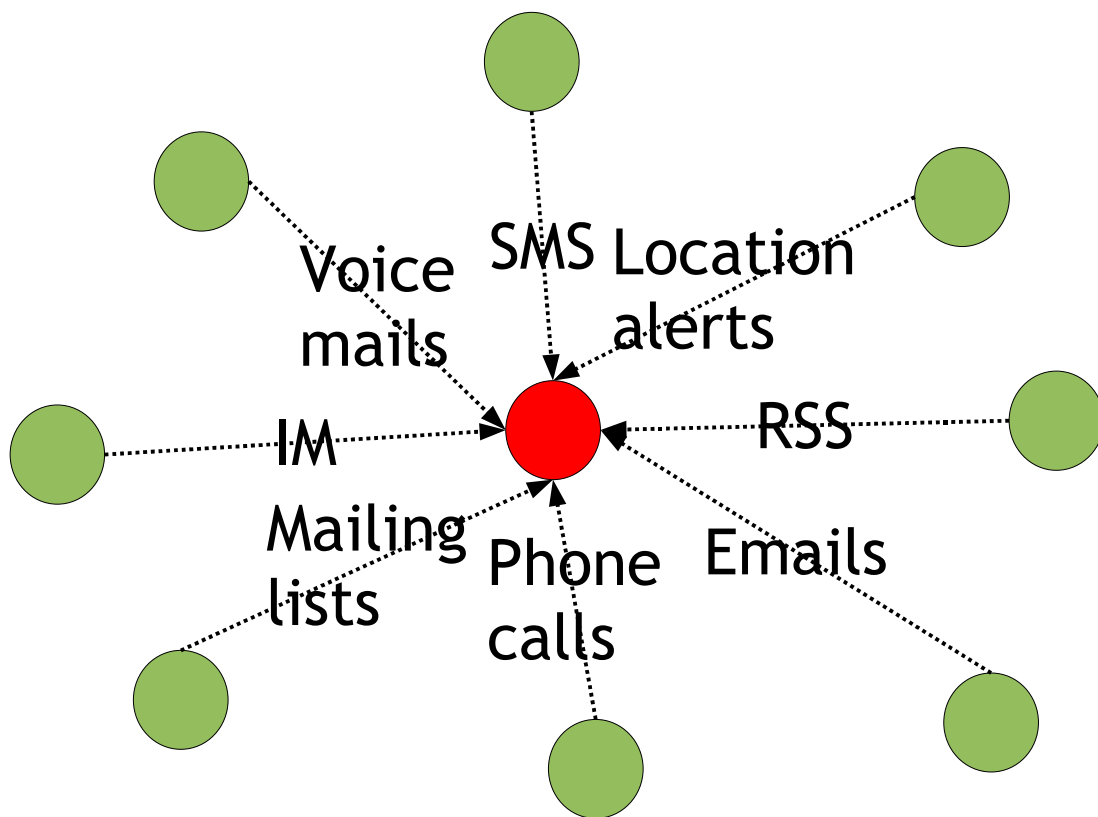
New Models 2: Models for “Flows of Attention”

- Where are the bottlenecks?

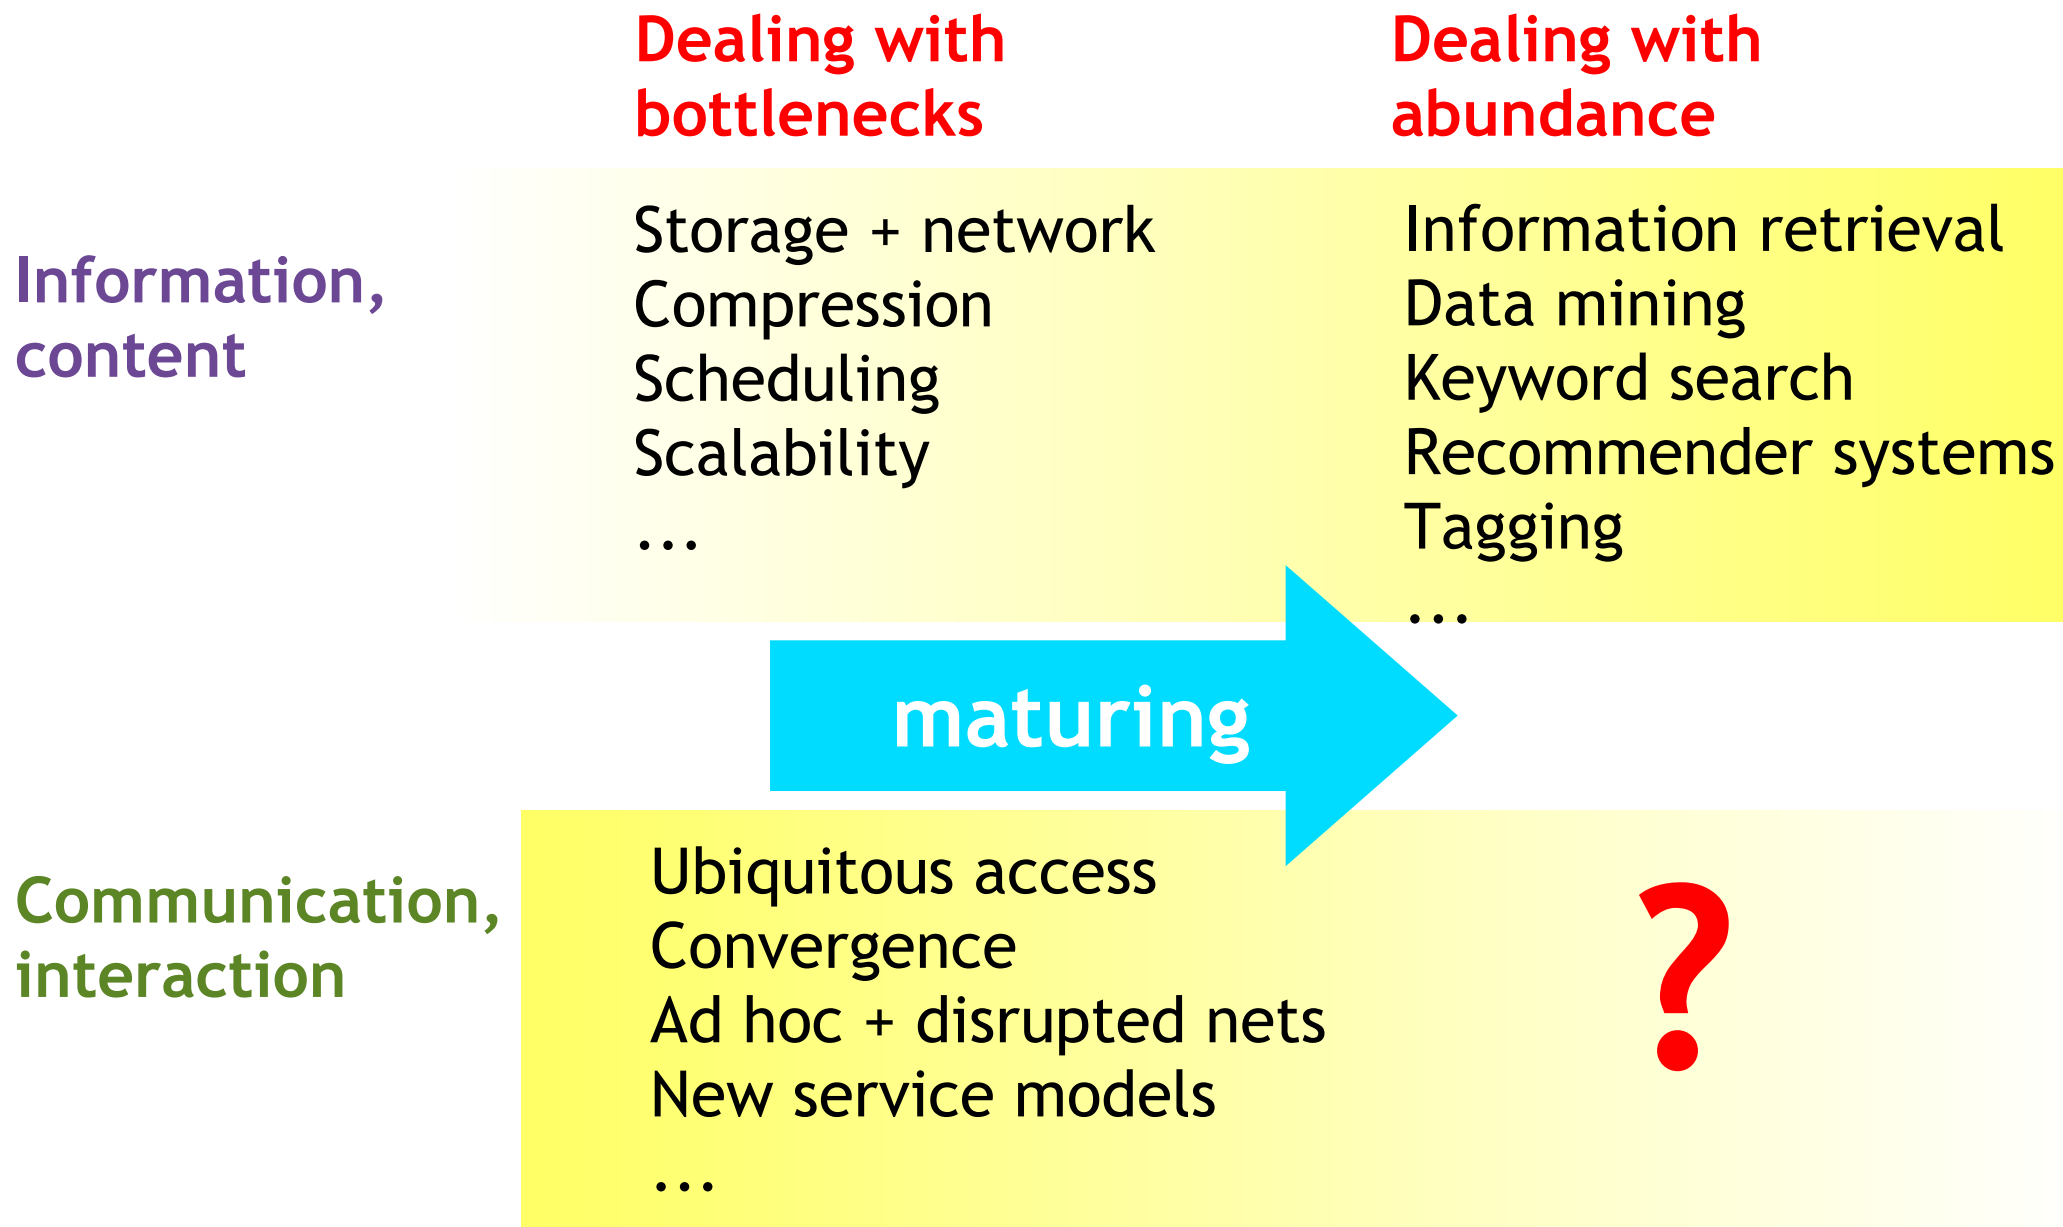
- Human attention is increasingly the most critical resource
- Success-disaster!

- My prediction:

- Dealing with the abundance of communication will be the most pressing problem over the next decade
- Can we model these flows? Are there interesting/useful invariants? Importance of context?
- At the boundary between networking and data mining/knowledge discovery



From Too Little to Too Much



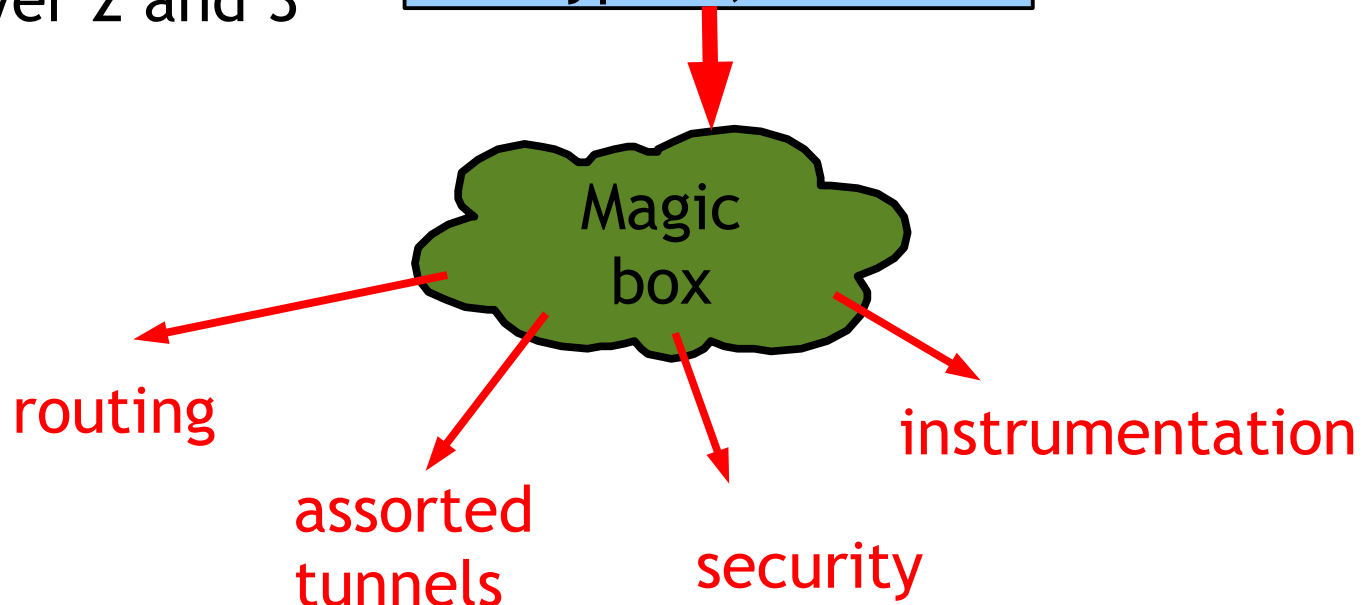
New Models 3: Still no VHDL for Networks?

- Network engineers today:
 - 10% TCP performance et al.
 - 90% “logical” problems
- Teaching: how to go beyond mountains of examples? How to make Cisco certification unnecessary?

- Cross-product of:

- MPLS
- routing, incl. BGP
- security
- virtualization (layer 2 and 3 and hosts)
- firewalls
- pricing
- ...

I need a VPN with failover protection, carrying low-delay VoIP + data, class x encrypted,...



Conclusion

- Models are the hallmark of a viable field
- We need more domain-specific models
- Examples:
 - Explicit heterogeneity: node and traffic distribution
 - The human user in the loop: flows of attention
 - Higher-level service abstractions: VHDL for networks