To Build, Or Not To Build, That Is The Question …

Nitin Vaidya
Georgetown University

Outline

- Brief overview of past work
- Some thoughts on building testbeds
- PG-13

Caveat

- I don’t always follow my own advice

Brief History of Time

ECE @ UMass Ph.D

➔ CS @ Texas A&M

➔ ECE @ UIUC

➔ CS @ Georgetown

Fault-tolerant computing

➔ Wireless networks … systems

➔ Distributed algorithms … theory

Fault-Tolerance

Checkpointing & Rollback Recovery

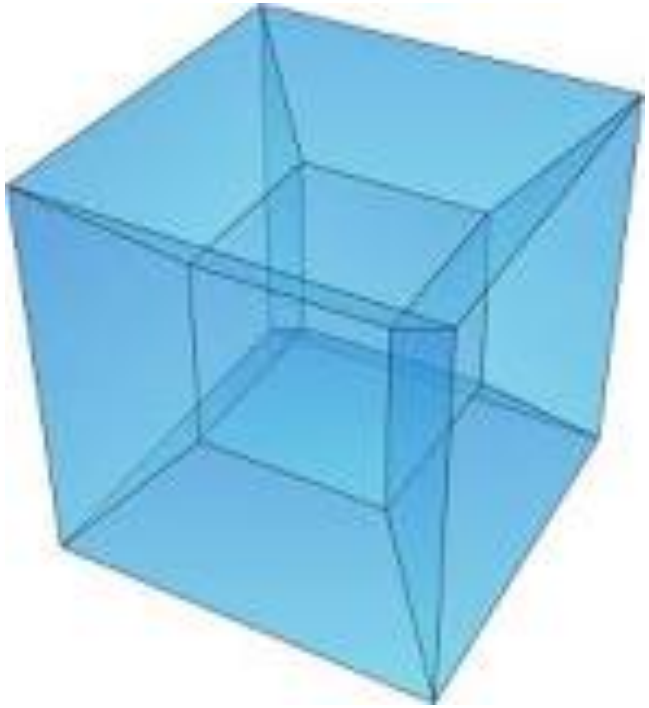

Coordinated Checkpoints

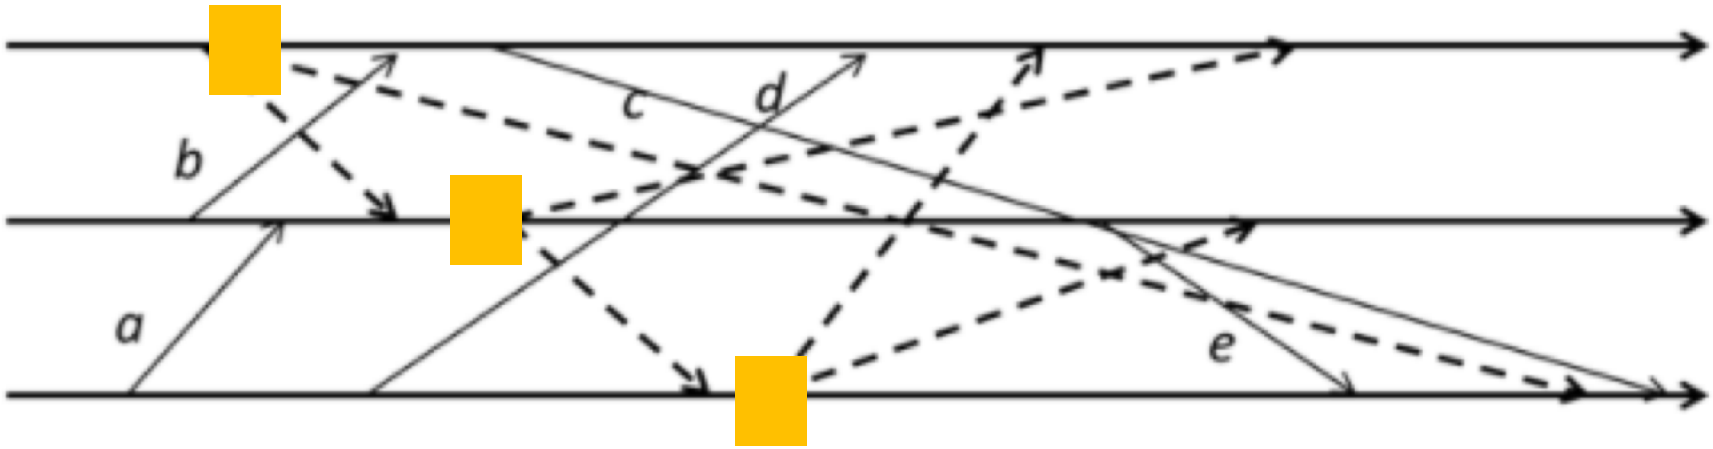

Application messages

Control messages

Coordinated Checkpoints

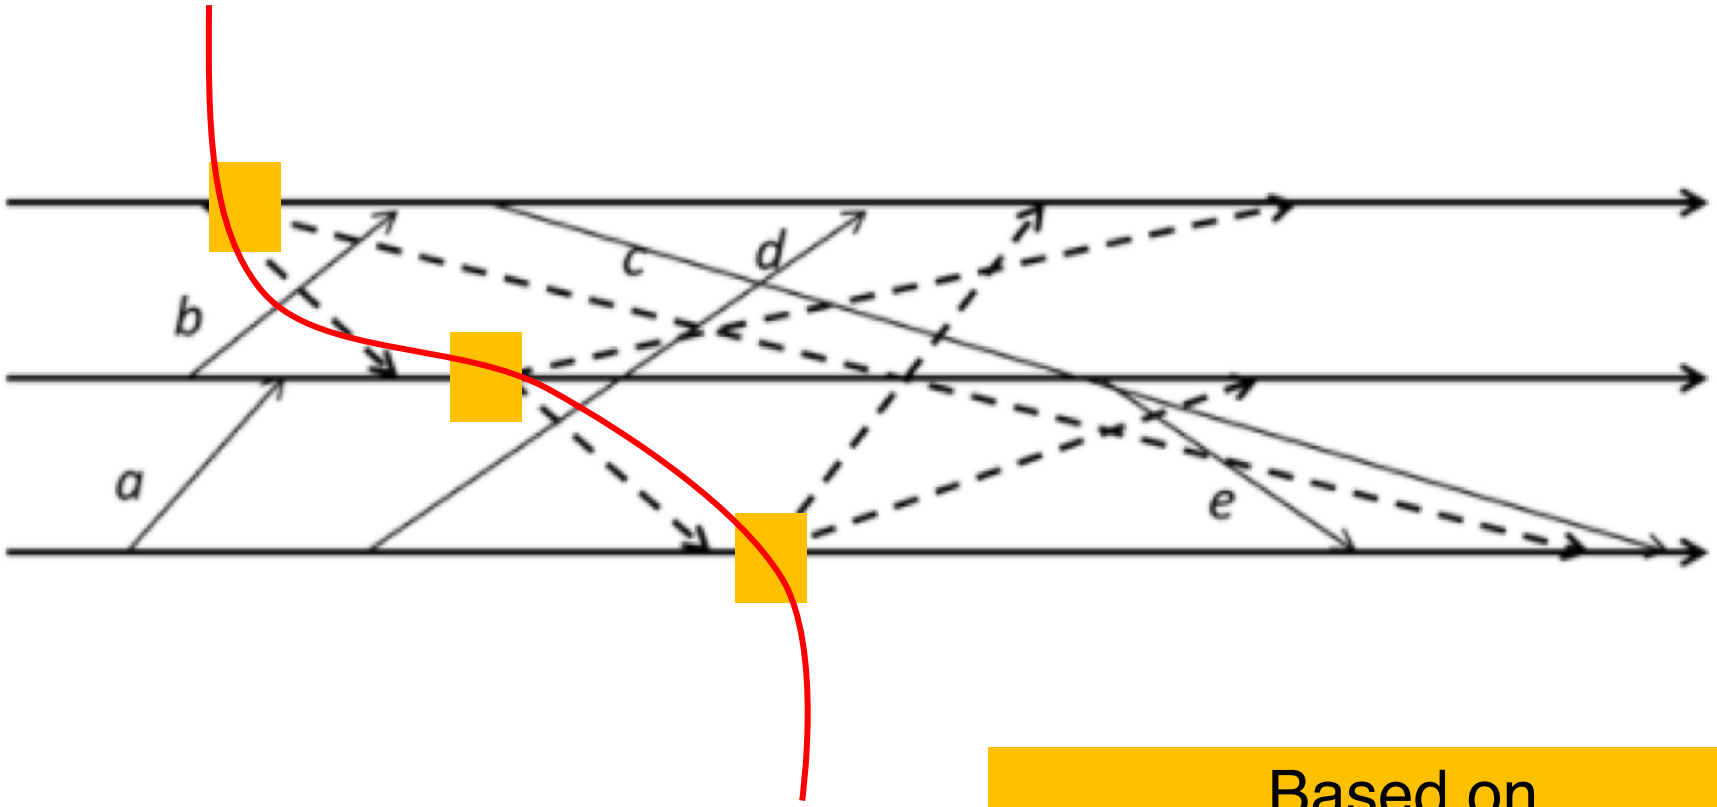

Based on
Chandy-Lamport Snapshot

Coordinated Checkpoints

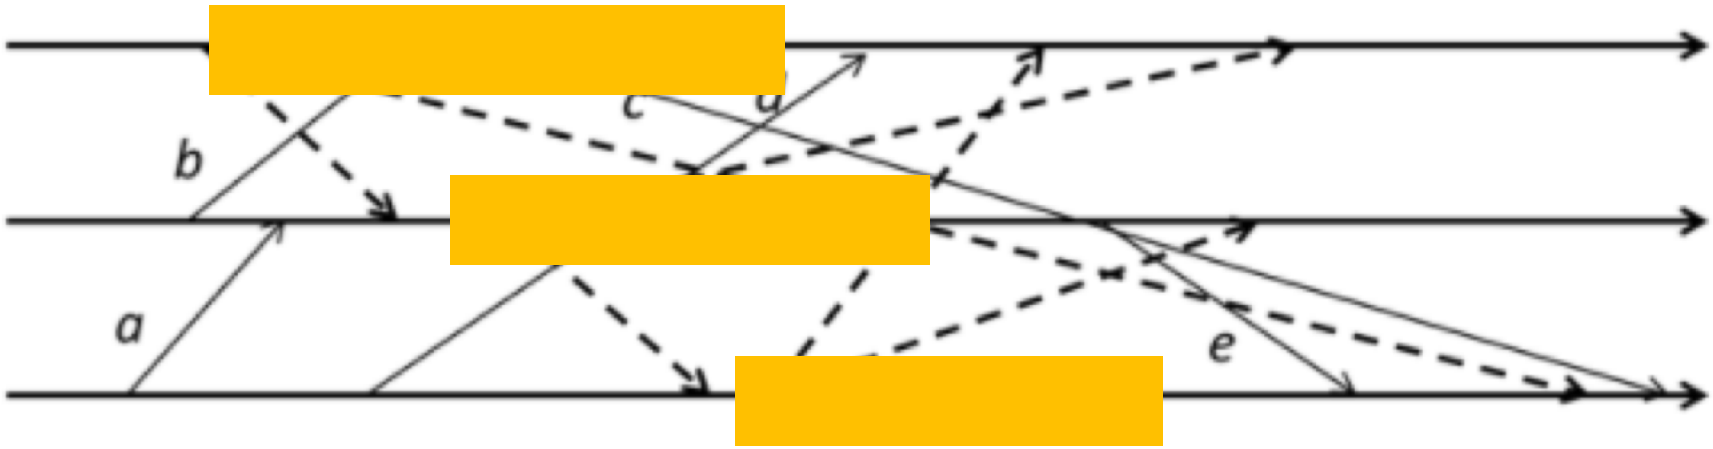

Consistent Logical Checkpoints
Staggered Checkpointing

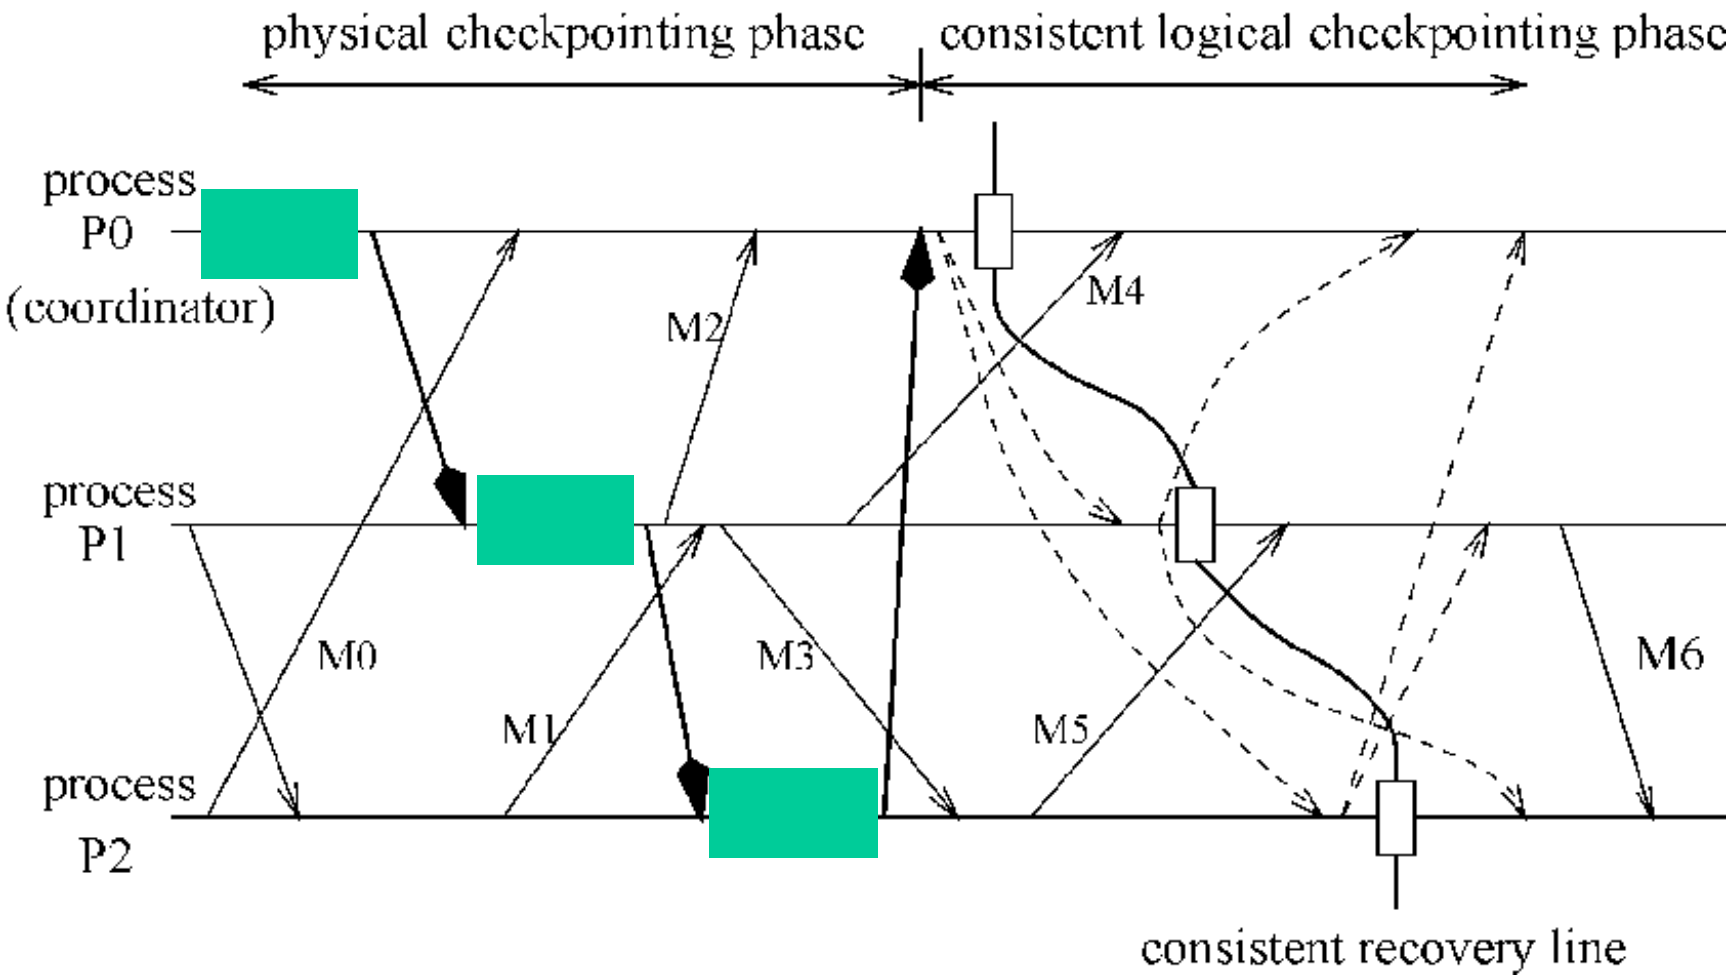

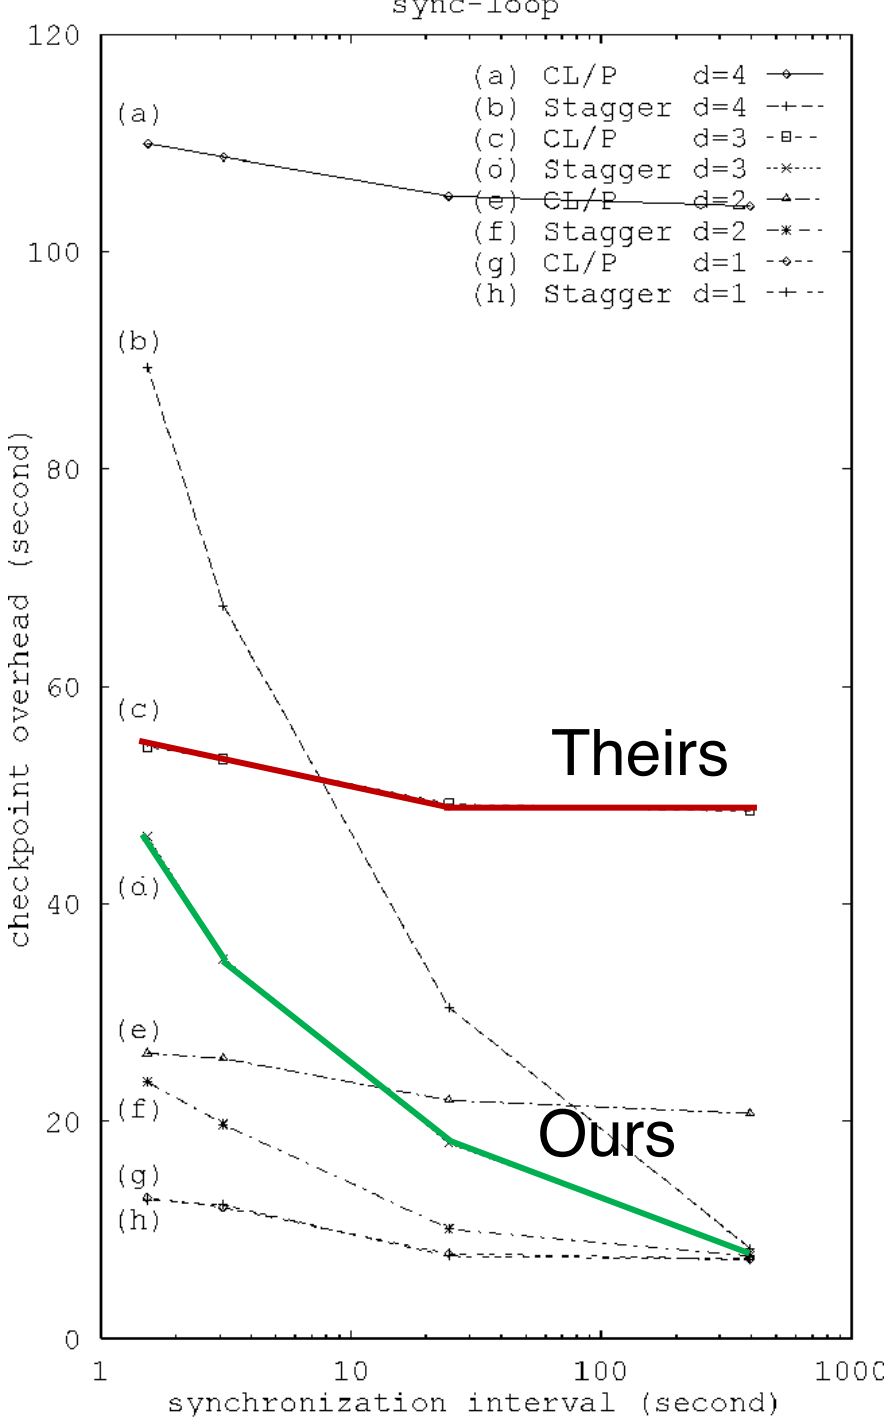
sync-loop
(a) CL/P d=4
(b) Stagger d=4
(c) CL/P d=3
(d) Stagger d=3
(e) CL/P d=2
(f) Stagger d=2
(g) CL/P d=1
(h) Stagger d=1
(a)
(b)
(c)
(d)
(e)
(f)
(g)
(h)
Theirs
Ours
checkpoint overhead (second)
synchronization interval (second)
0
20
40
60
80
100
120
1
10
100
1000

Multi-Level Checkpointing

- Different (cost) checkpoints for different faults

Mesh Networks

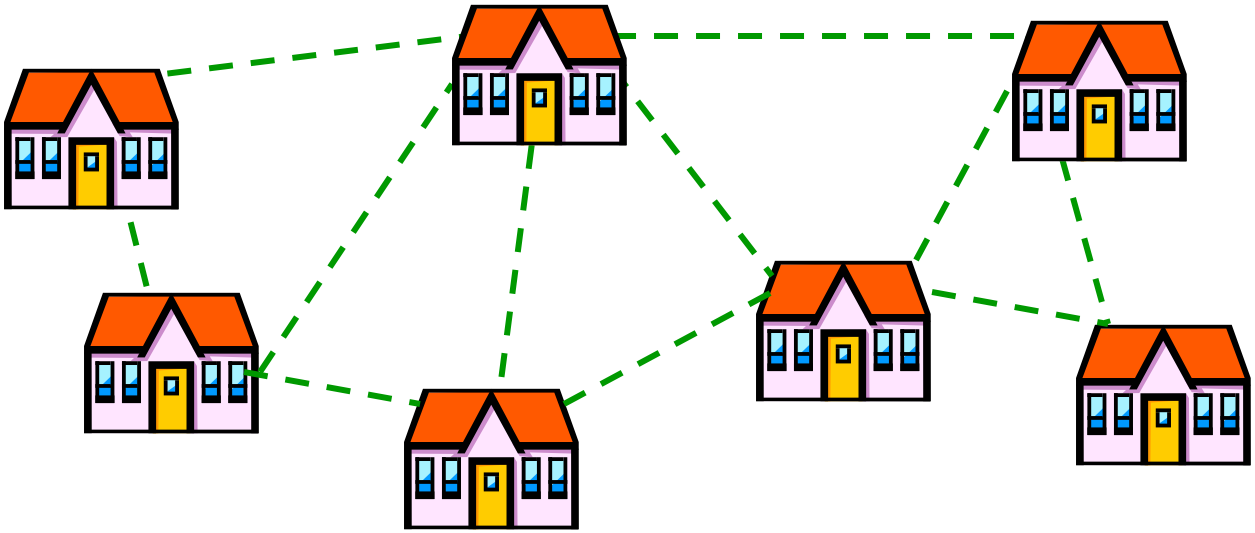

Multi-Channel Systems

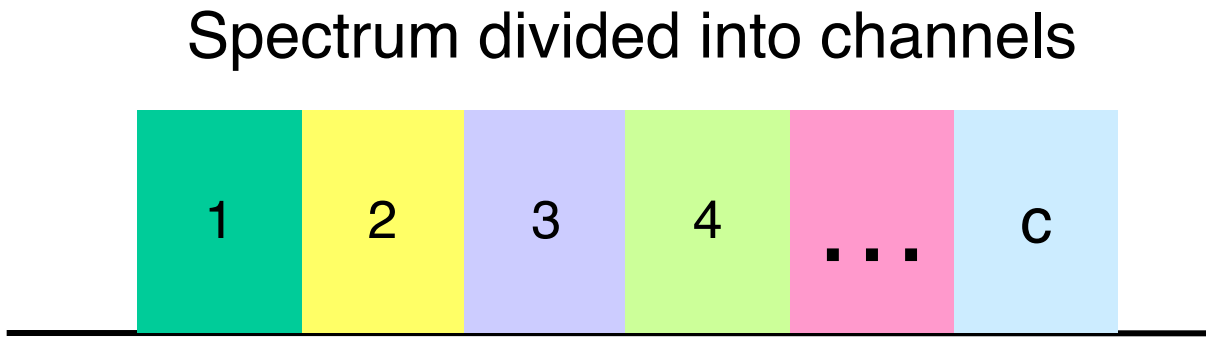

Practical Scenario

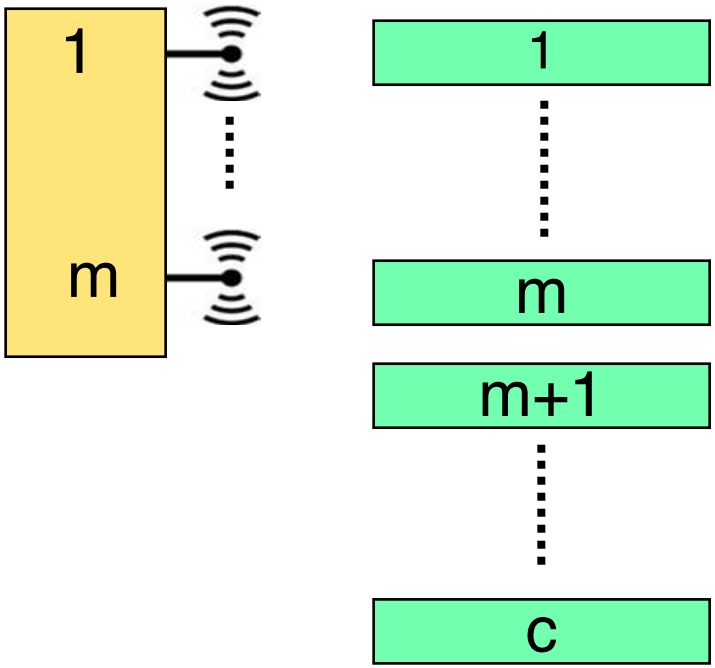

c–m unused channels at each node

Practical Scenario

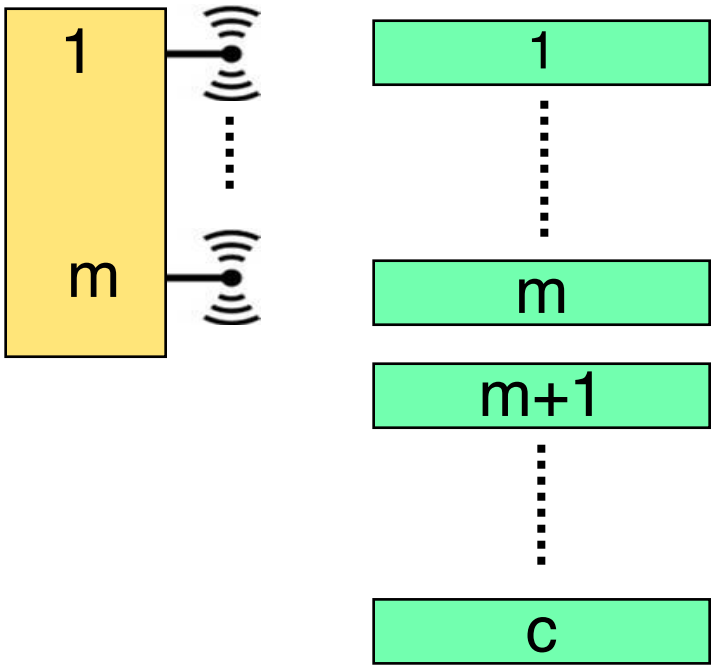

c–m unused channels at each node

How does mesh performance scale with c and m ?

Net-X: Multi-Channel Mesh

Theory to Practice

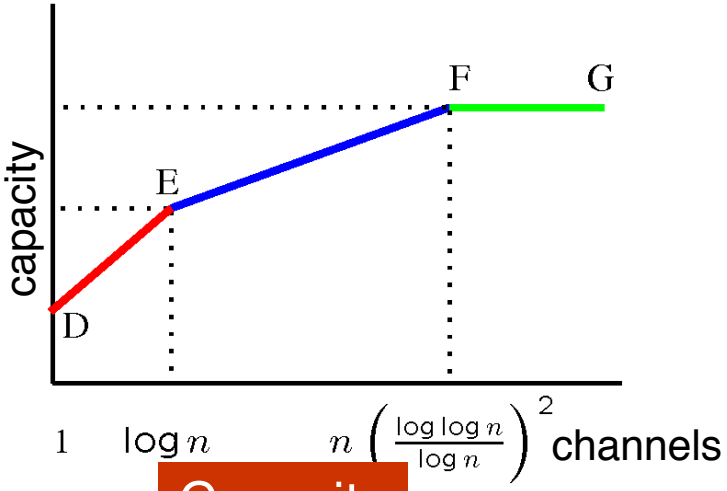

Capacity bounds

Net-X: Multi-Channel Mesh

Theory to Practice

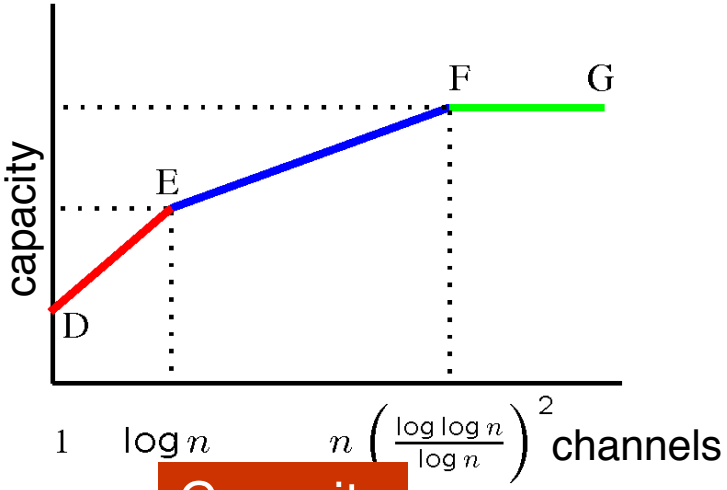

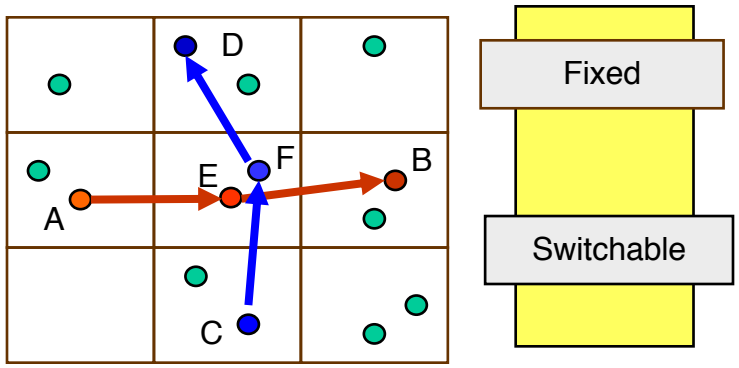

Capacity bounds → Insights on protocol design

Net-X: Multi-Channel Mesh

Theory to Practice

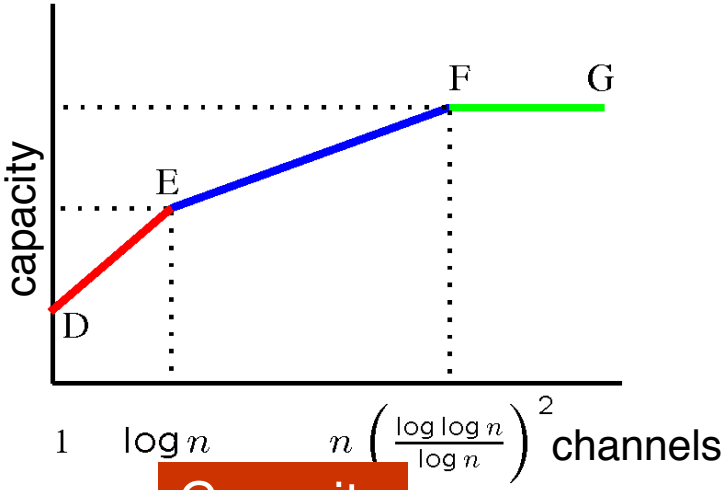

Capacity bounds

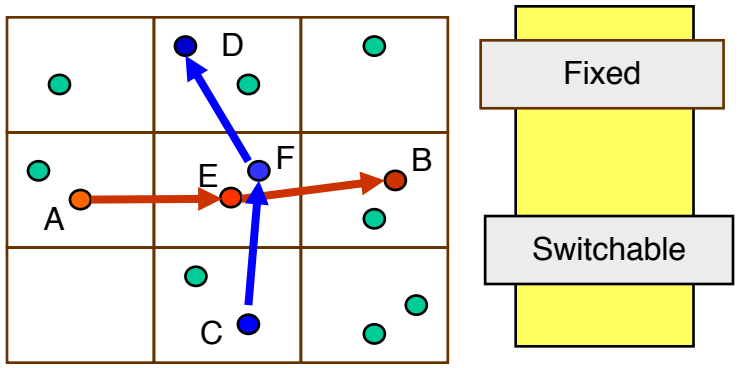

Insights on protocol design

OS improvements
Software architecture

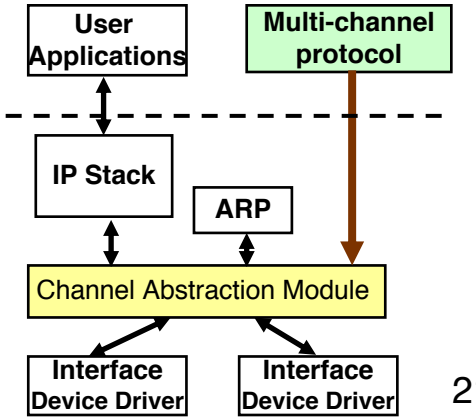

Net-X: Multi-Channel Mesh

Theory to Practice

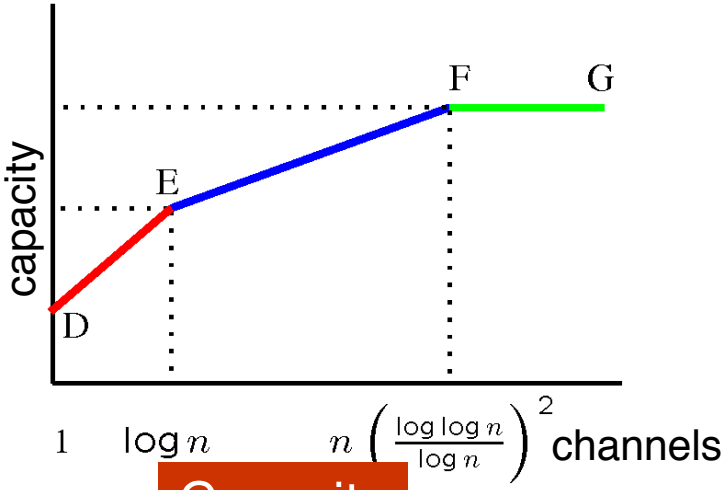

Capacity bounds

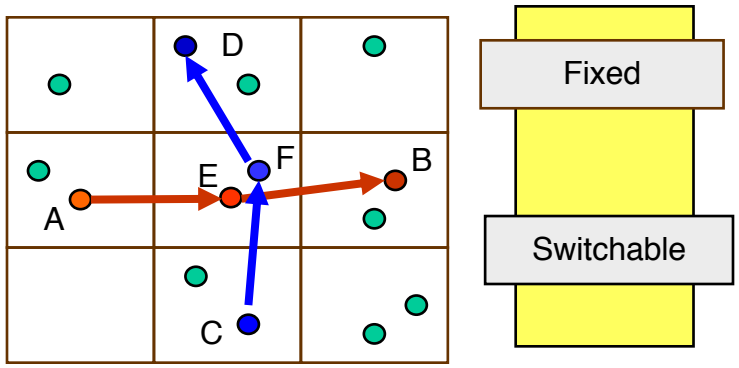

Insights on protocol design

OS improvements
Software architecture

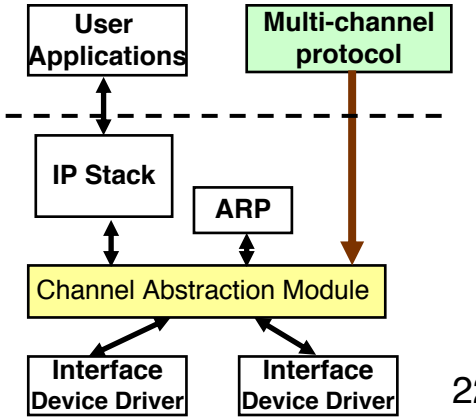

Net-X testbed

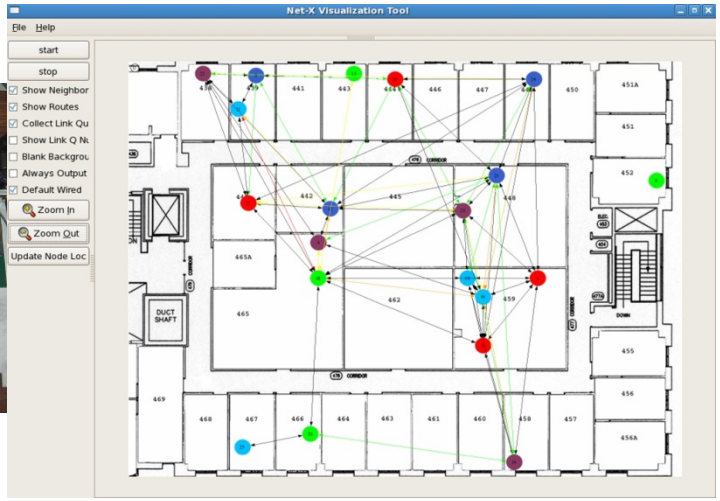

Linux box

Distributed Algorithms

"Local Computations"

Average Consensus

Average Consensus

Initially, state = input

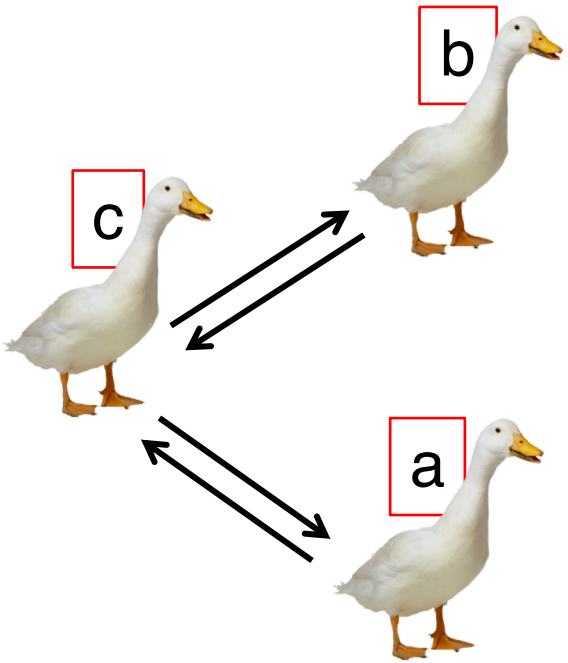

b = 3b/4+ c/4

c = a/4+b/4+c/2

a = 3a/4+ c/4

Average Consensus

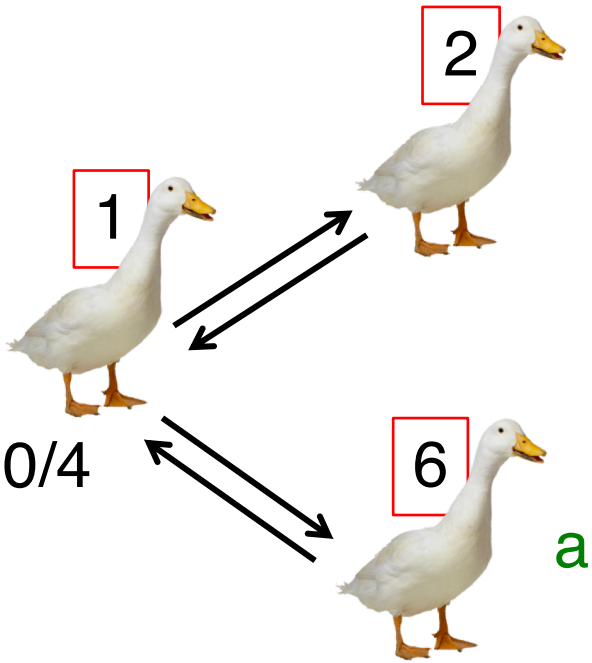

b = (6/4+1/4)=7/4

c = (2/4+6/4+1/2) = 10/4

a = (18/4+1/4) = 19/4

Average Consensus

Values converge to *average* of inputs

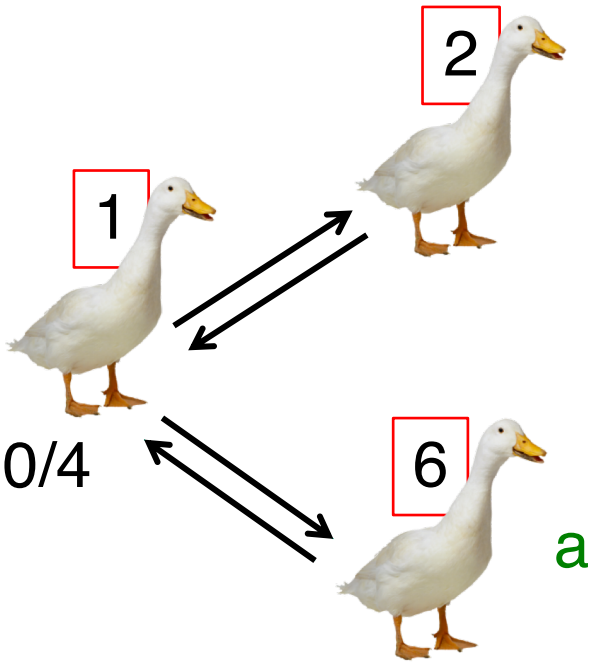

b = (6/4+1/4)=7/4

c = (2/4+6/4+1/2) = 10/4

a = (18/4+1/4) = 19/4

Implementing Local Algorithms

- Too much work to implement (in wireless networks)
- Software environment to make life easier
- Programmer provides pseudo-code
- Rest automated

To Build, Or Not To Build,
That Is The Question …

Staggered Checkpointing

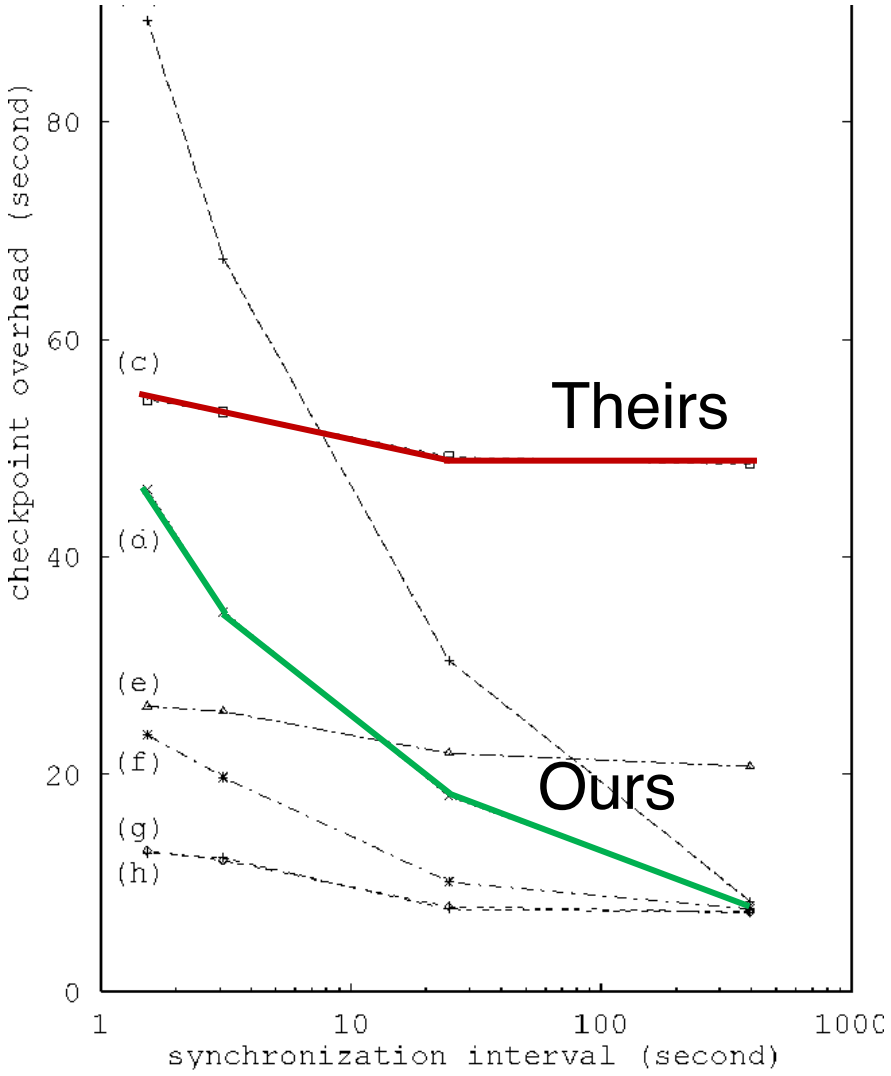

Staggered Checkpointing

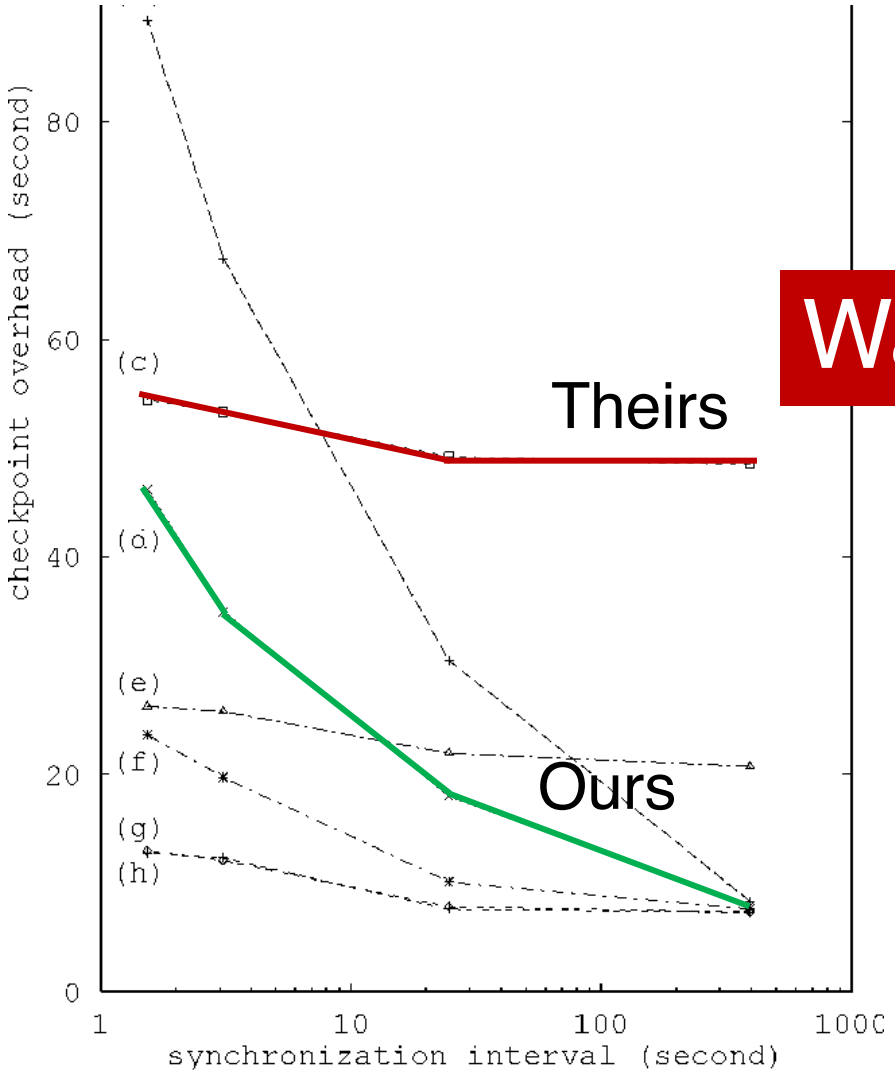

Waste of time

Average Consensus

Software Toolkit for Local Computations

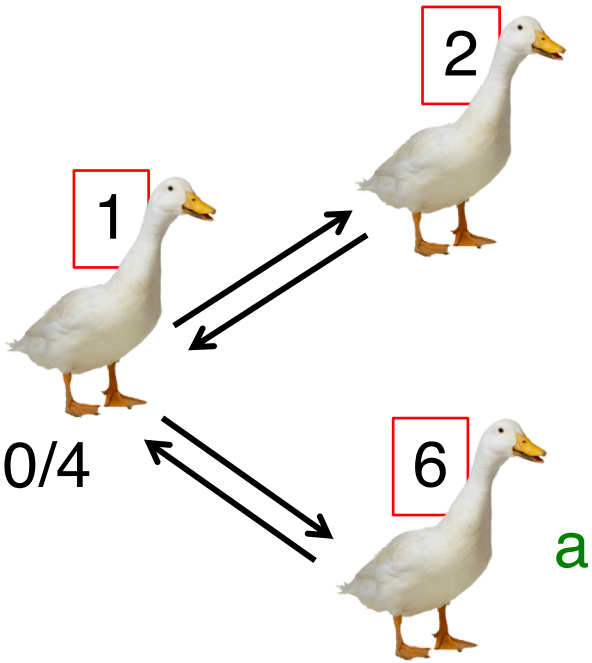

b = (6/4+1/4)=7/4

c = (2/4+6/4+1/2) = 10/4

a = (18/4+1/4) = 19/4

Average Consensus

Software Toolkit for Local Computations

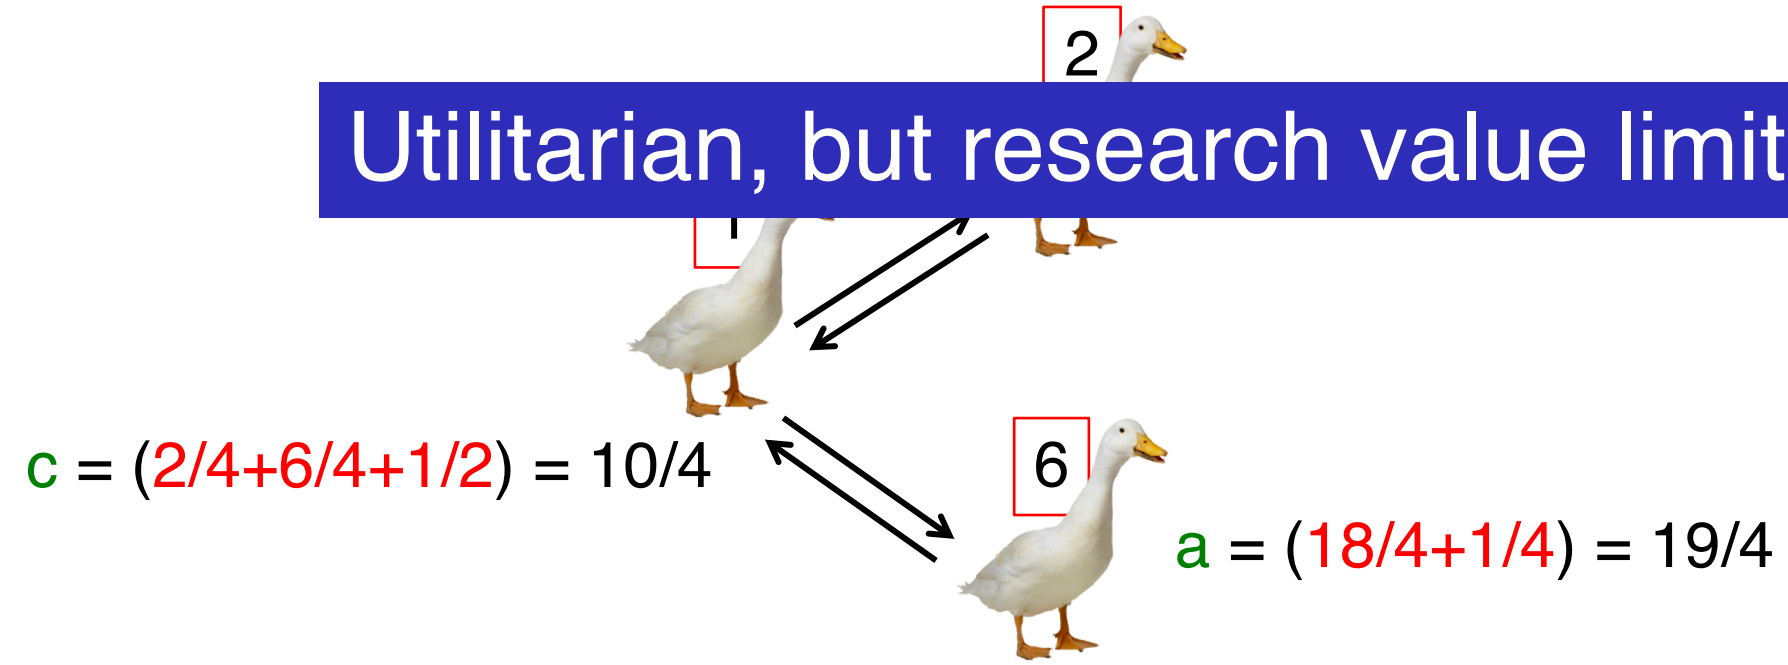

Utilitarian, but research value limited

Net-X: Multi-Channel Mesh

Theory to Practice

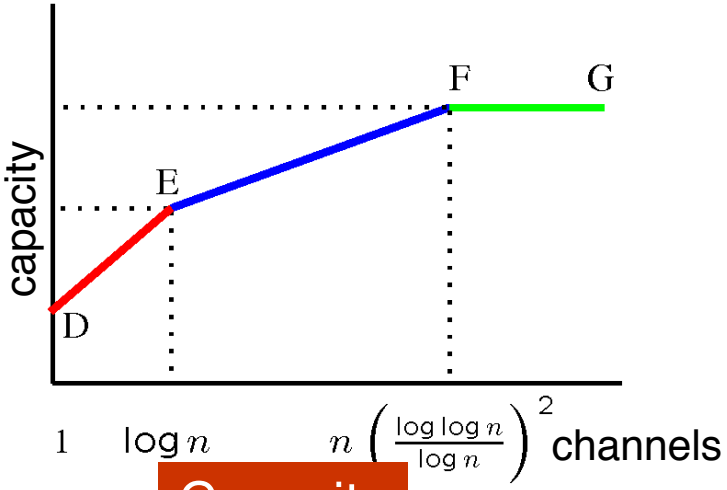

Capacity bounds

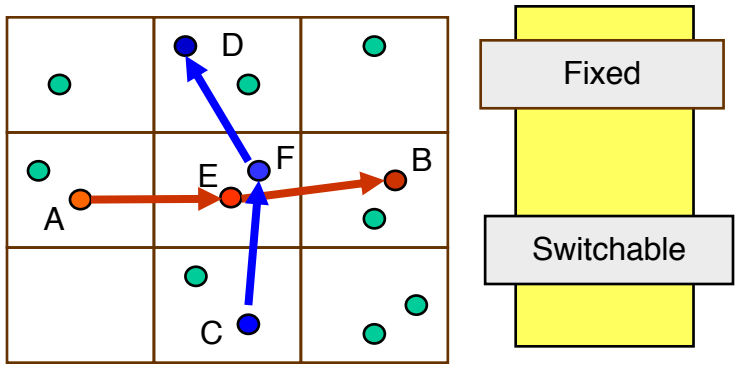

Insights on protocol design

OS improvements
Software architecture

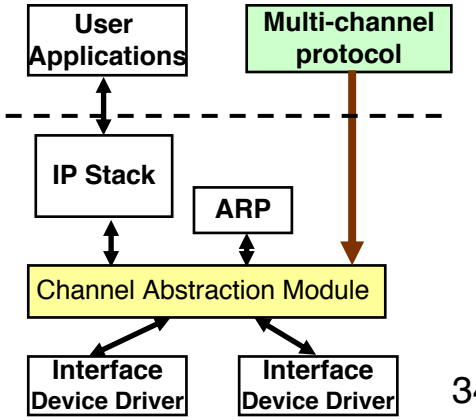

Net-X testbed

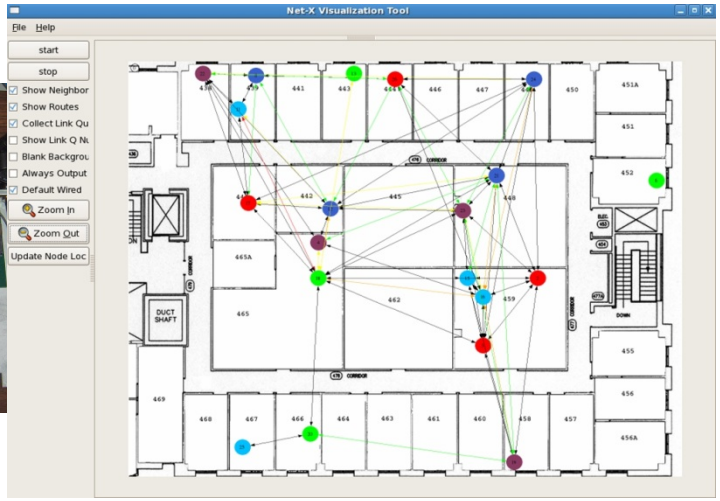

Linux box

Net-X: Multi-Channel Mesh

Theory to Practice

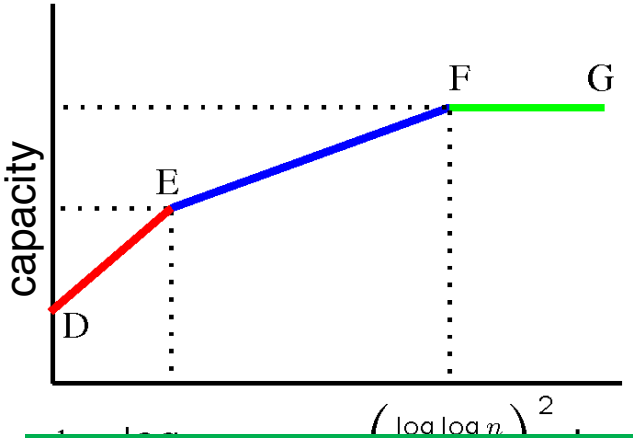

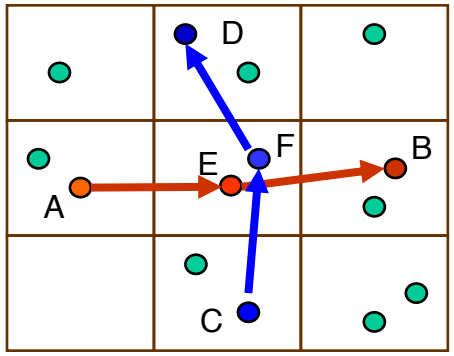

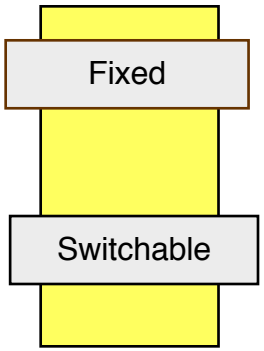

Best Case Scenario

Linux box

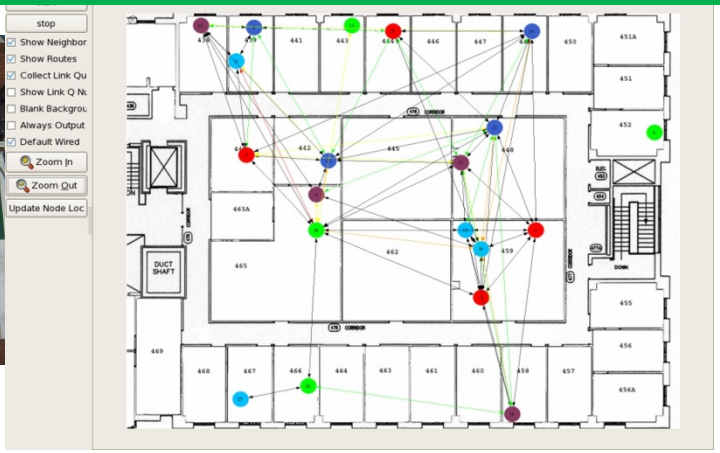

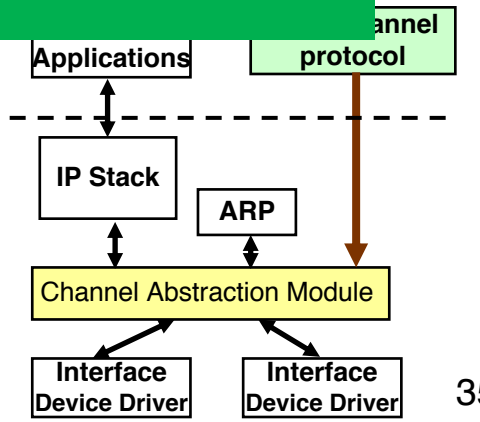

Multi-Channel Systems

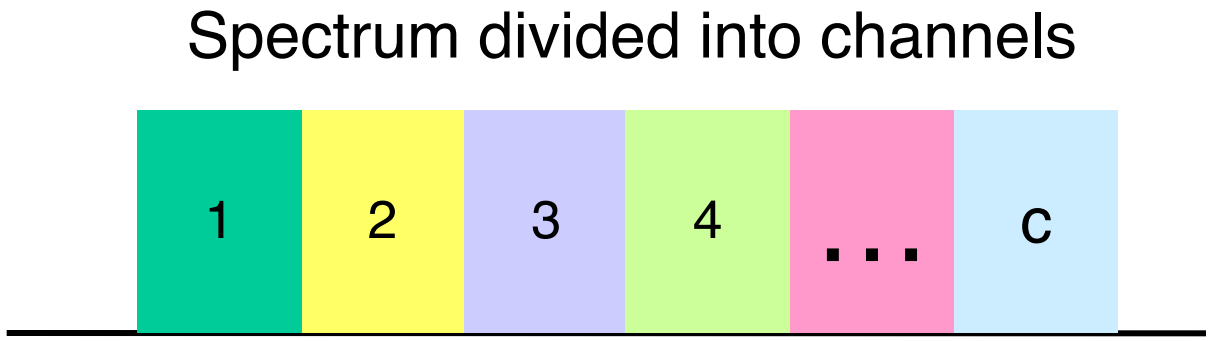

Multi-Channel Systems

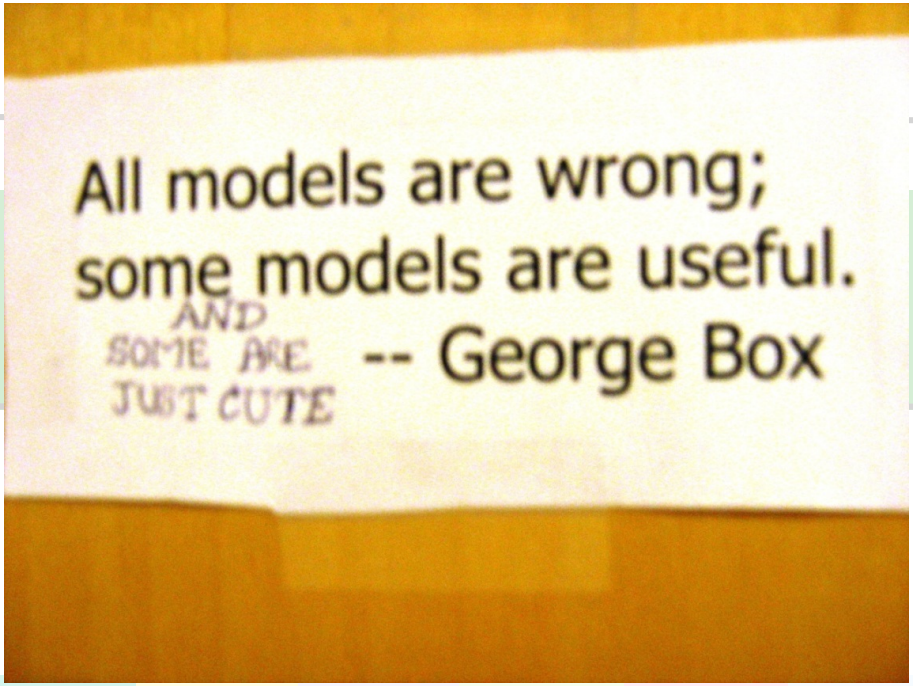

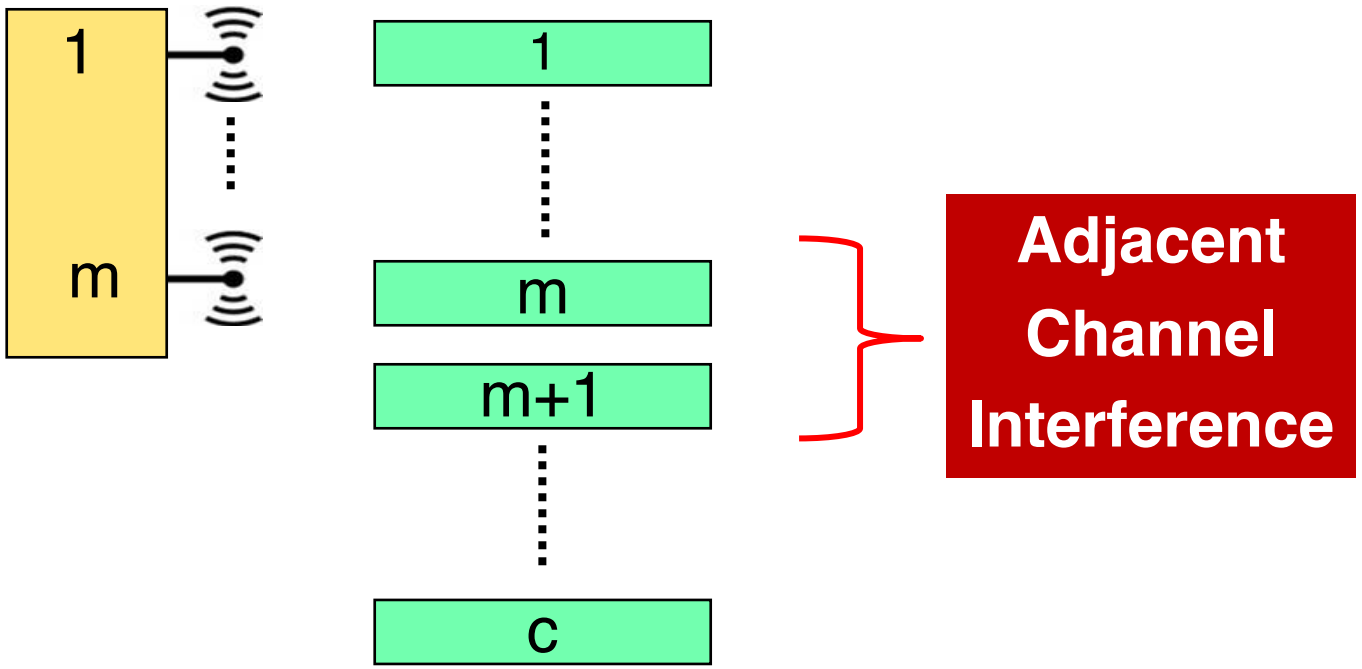
1
m
1
m
m+1
c
Adjacent
Channel
Interference

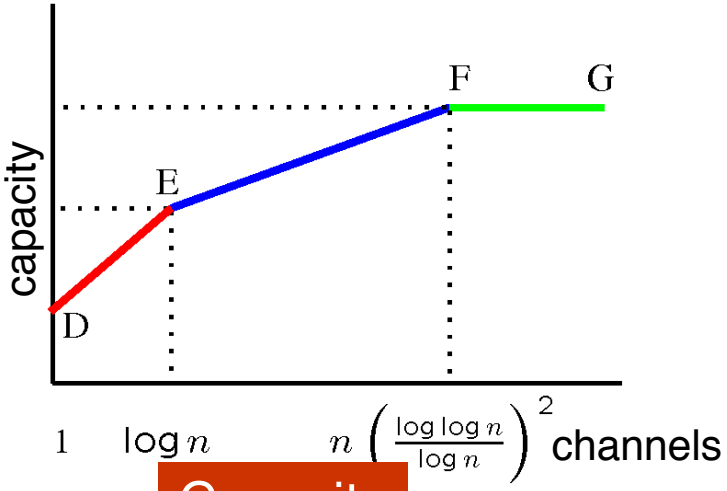

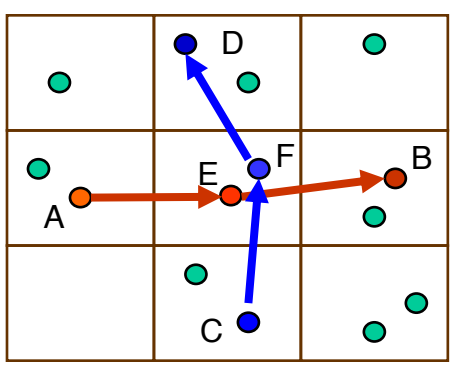

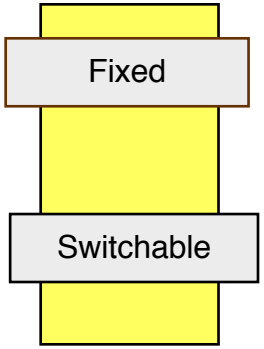

Capacity bounds

Insights on protocol design

OS improvements
Software architecture

Net-X testbed

Linux box

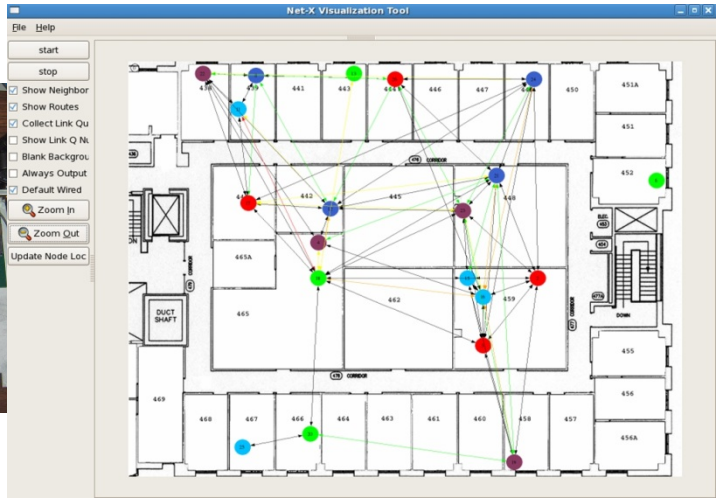

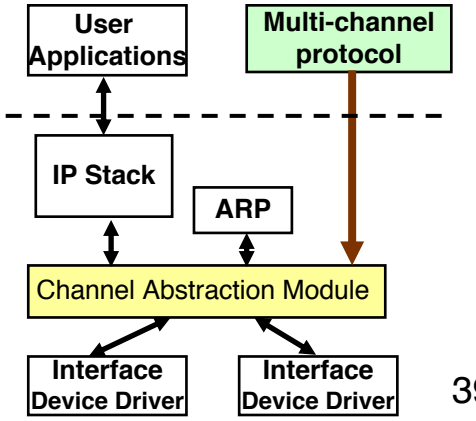

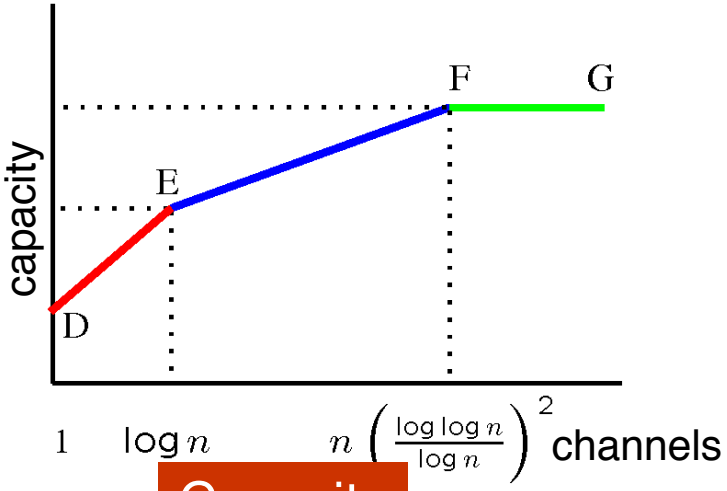

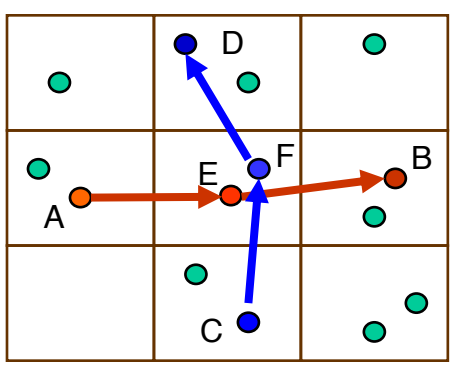

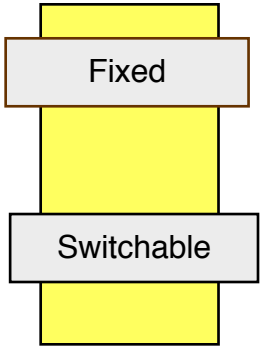

Capacity bounds

Insights on protocol design

OS improvements
Software architecture

Net-X testbed

Linux box

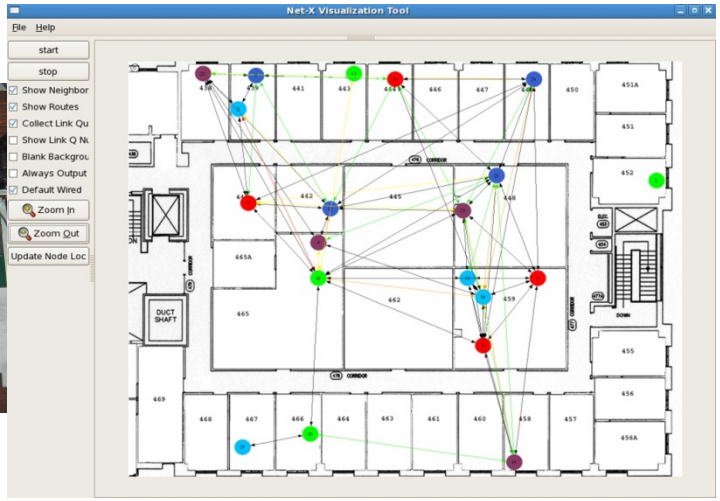

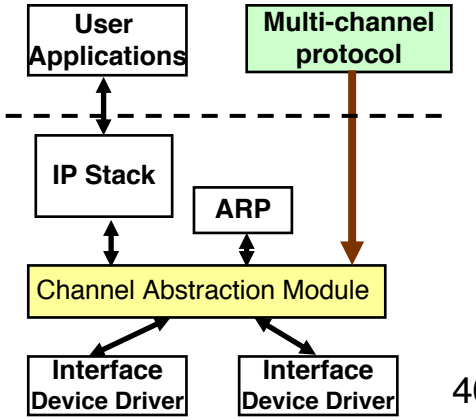

When To Build

When To Build

Theory

Systems

Picture from Wikipedia

When To Build

- When results are not predictable from theory
- For theory & simulations to suffice,
  need accurate system & workload models

When To Build

- When results are not predictable from theory
- For theory & simulations to suffice,
  need accurate system & workload models

➔ Academic architects rarely build physical systems anymore

Bad Reasons To Build

Bad Reasons To Build

- So we can publish the paper

Bad Reasons To Build

- So we can publish the paper

Much of the "systems" literature

Bad Reasons To Build

- So we can publish the paper

Much of the "systems" literature

- Make simple ideas appear "substantive"

Bad Reasons To Build

- So we can publish the paper

Much of the "systems" literature

- Make simple ideas appear "substantive"

- Everybody is doing it
  … so what's wrong with you?

Bad Reasons To Build

- So we can publish the paper

Much of the "systems" literature

- Make simple ideas appear "substantive"

- Everybody is doing it
  … so what's wrong with you?

Funding agencies prone to this

“The MSR Effect” *

“The MSR Effect” *

* Disclaimers:

Replace your favorite lab here

Some of my best friends are at MSR

“The MSR Effect”

- In the *good old days*, industry research labs aspired to do relevant but academic quality research
  - Fundamental research
  - Long timescales
  - “Independence” from products

"The MSR Effect"

- In the *good old days*, industry research labs aspired to do relevant but academic quality research
  - Fundamental research
  - Long timescales
  - "Independence" from products

Today …

- Research labs dominate many conferences

- Academics aspire to emulate industry labs

  ➔ "Systems" communities have succumbed to this

How to unwind this clock?

Break Artificial Boundaries

Theory

Systems

Minimalism

- Often less is more
- Don’t build just because you can
- There may be better things to do with your time and resources

Litmus Test

- Would you be willing to publicly post the *exact problem statement* ?

… before developing the solution

- If not, find something better to do

Thanks!

disc.georgetown.domains