Resolving the Transport "Tussle"

Recursive InterNetwork Architecture

Ibrahim Matta
Computer Science
Boston U.

The Transport Tussle panel at PFLDNeT 2010

The RINA Team

- Faculty
  - John Day
  - Lou Chitkushev
- Outside Collaborators:
  - Patrick Phelan, TSSG, Ireland
  - Eduard Grasa, i2CAT, Spain
  - Eleni Trouva, i2CAT, Spain
  - Steve Bunch, TRIA Network Systems
  - Peter DeWolf, TRIA Network Systems
  -
- PhD Students:
  - Karim Mattar
  - Vatche Ishakian
  - Gonca Gursun
  - Joseph Akinwumi
  - Flavio Esposito
  - Yuefeng Wang
- Undergraduates:
  - Ethan DeGracia
- Alumni:
  - Gowtham Boddapati, Akamai
  - Peng Ge, Epic Systems
  - Qi Guo, Epic Systems
  - Safae Lahjouji, ENSI-Bourges

What's wrong with today's Internet?

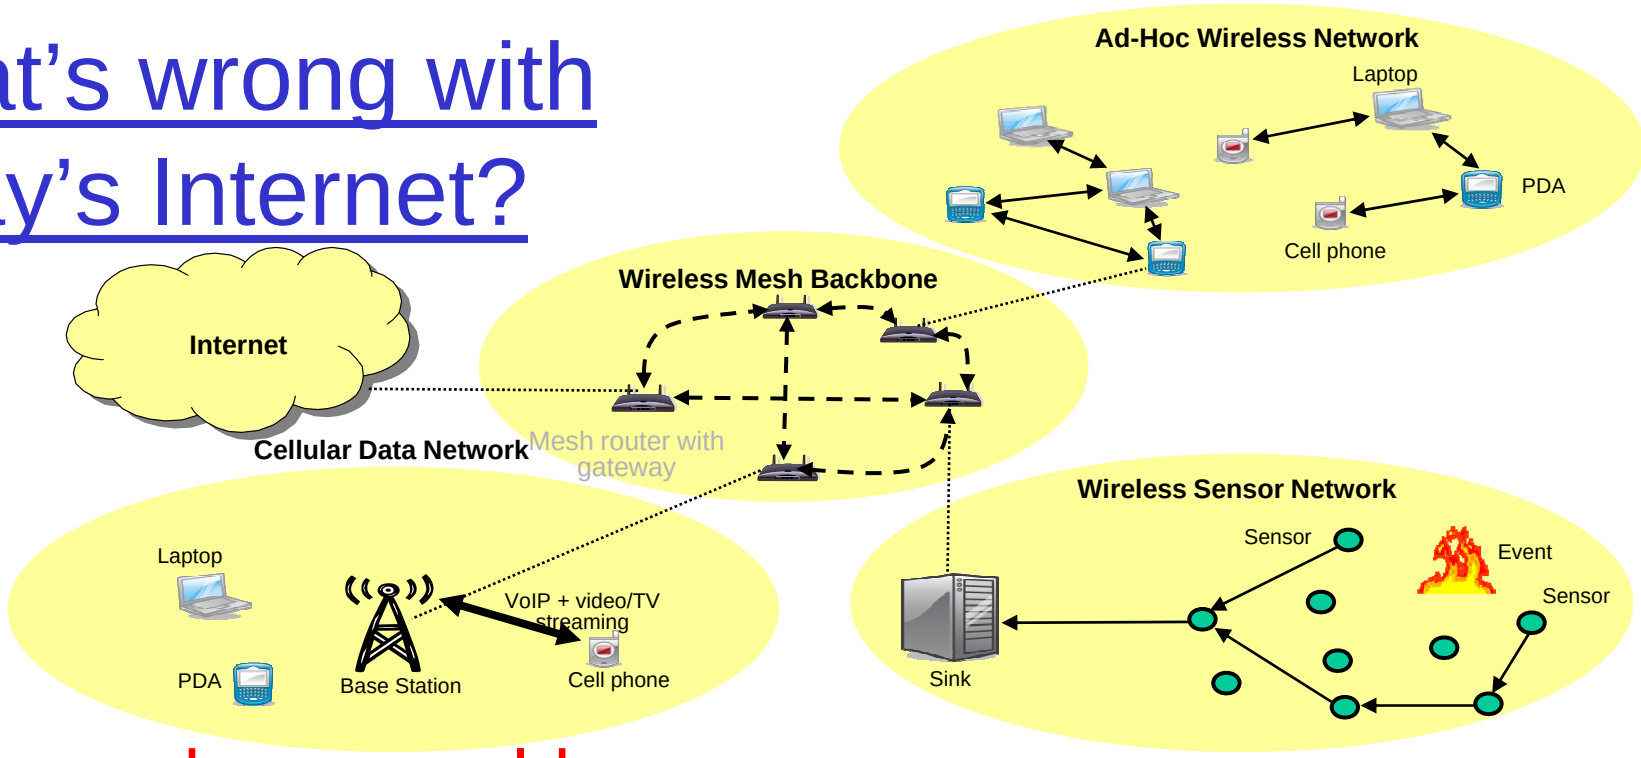

- The new brave world
  - Larger scale, more diverse technologies
  - New services: content-driven, context-aware, mobile, socially-driven, secure, profitable, …
- Custom point-solutions: No or little "science"
- Lots of problems: Denial-of-service attacks, bad performance, hard to manage, …

Questions?

- Is the Internet's architecture fundamentally broken that we need to "clean slate"?
  - Yes
- Can we find a new architecture that is complete, yet minimal? If so, what is it?
  - RINA?

- Can we transition to it without requiring everyone to adopt it?
  - Yes

Internet's view: one big, flat, open net

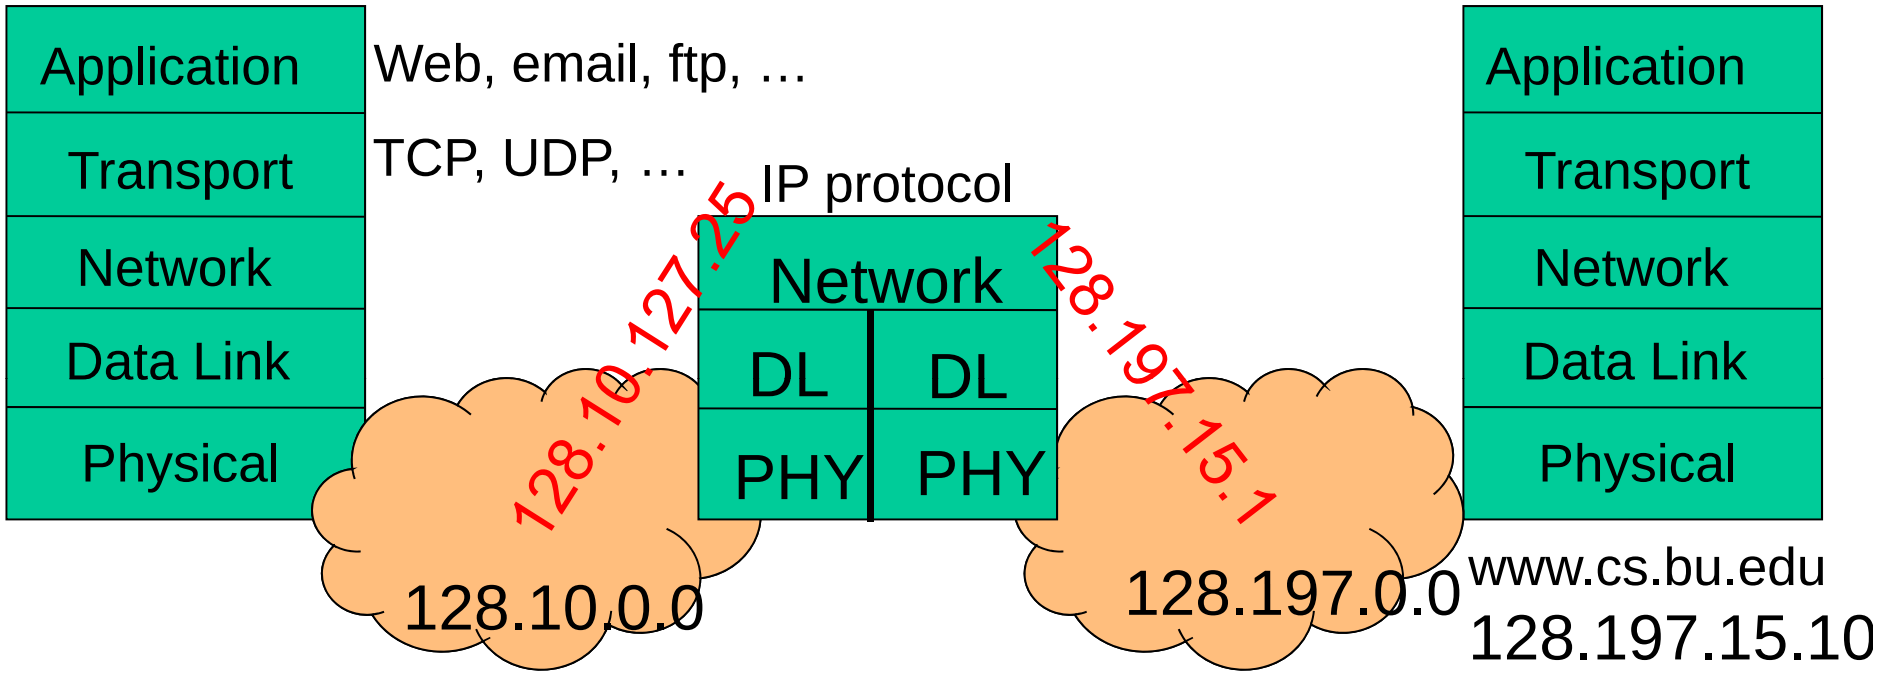

- There's no building block
- The "hour-glass" model imposed a least common denominator
- We named and addressed the wrong things (i.e., interfaces)
- We exposed addresses to applications
- We hacked in "middleboxes"

Our Solution: divide-and-conquer

- Application processes communicate over (distributed) IPC facility
- How IPC managed is hidden ➔ better security

- IPC processes are application processes of lower IPC facilities
- Recurse as needed

  ➔ better management & scalability

- Well-defined interfaces ➔ predictable service

Recursive Architecture based on IPC

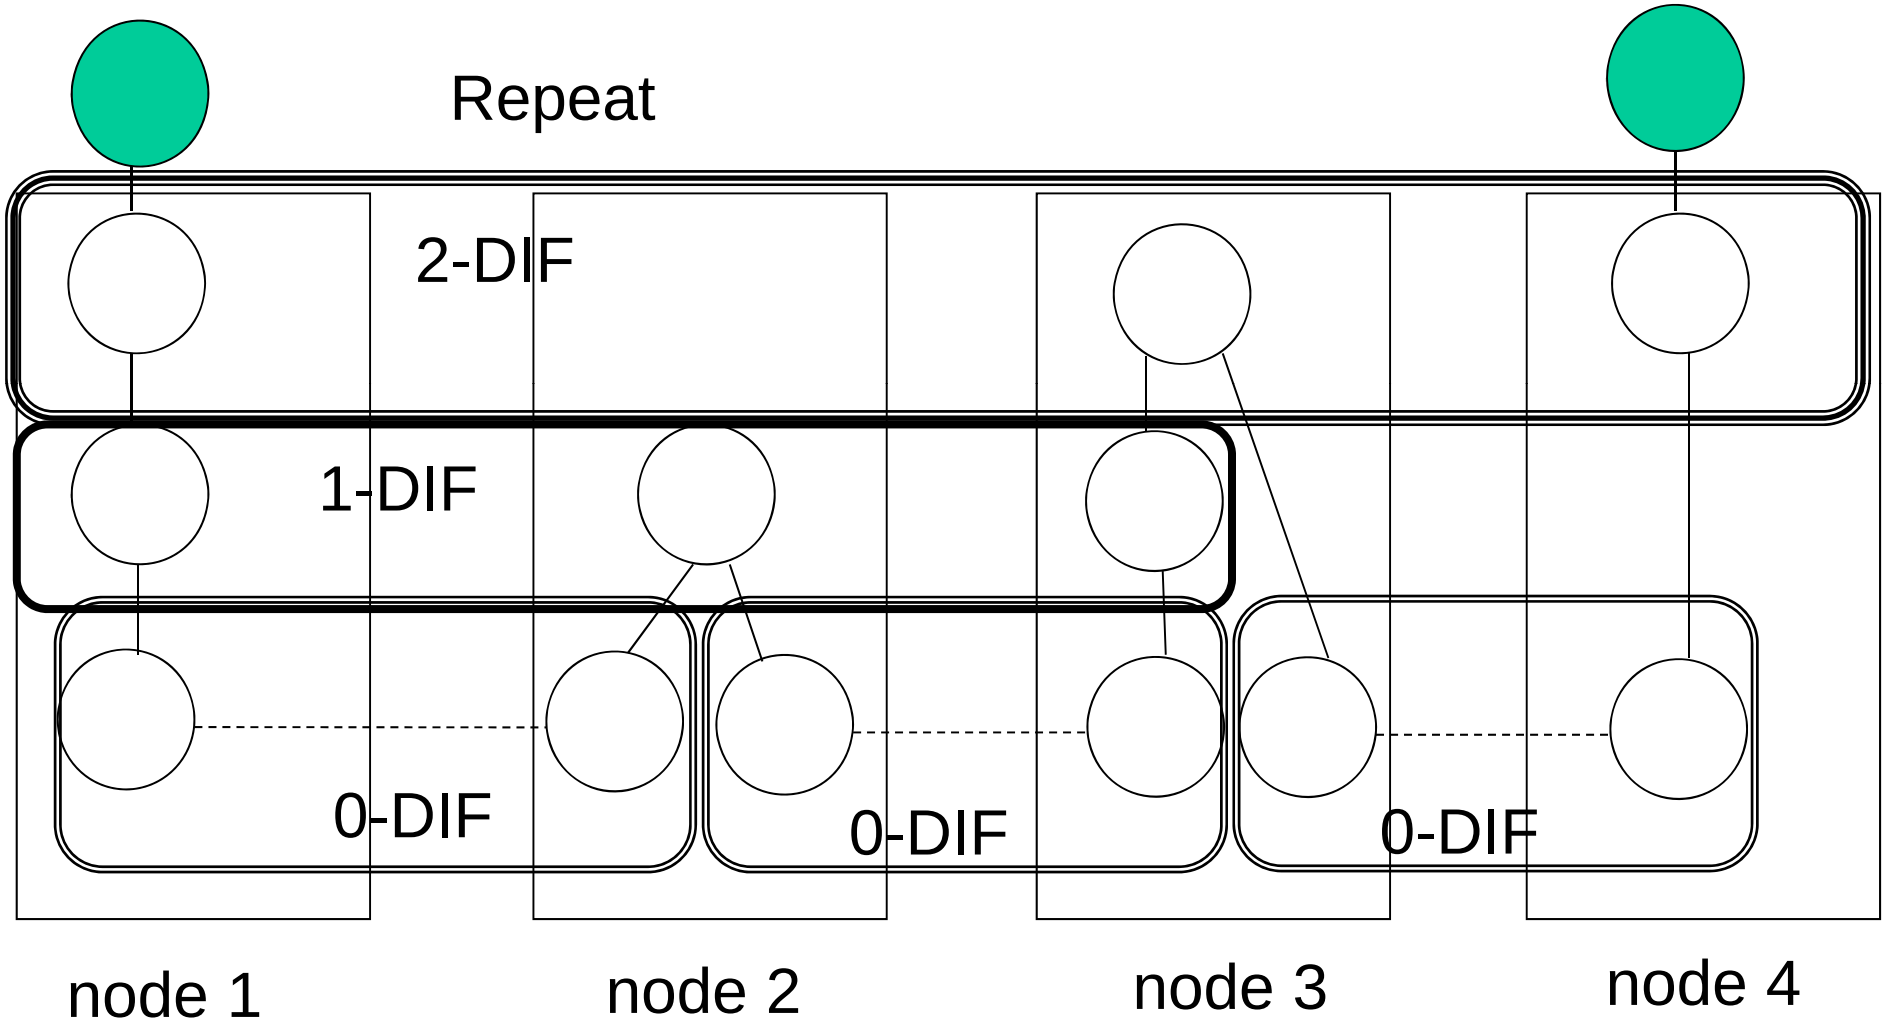

DIF = Distributed IPC Facility (locus of shared state=scope)

Policies are tailored to scope of DIF

RINA allows scoping of services

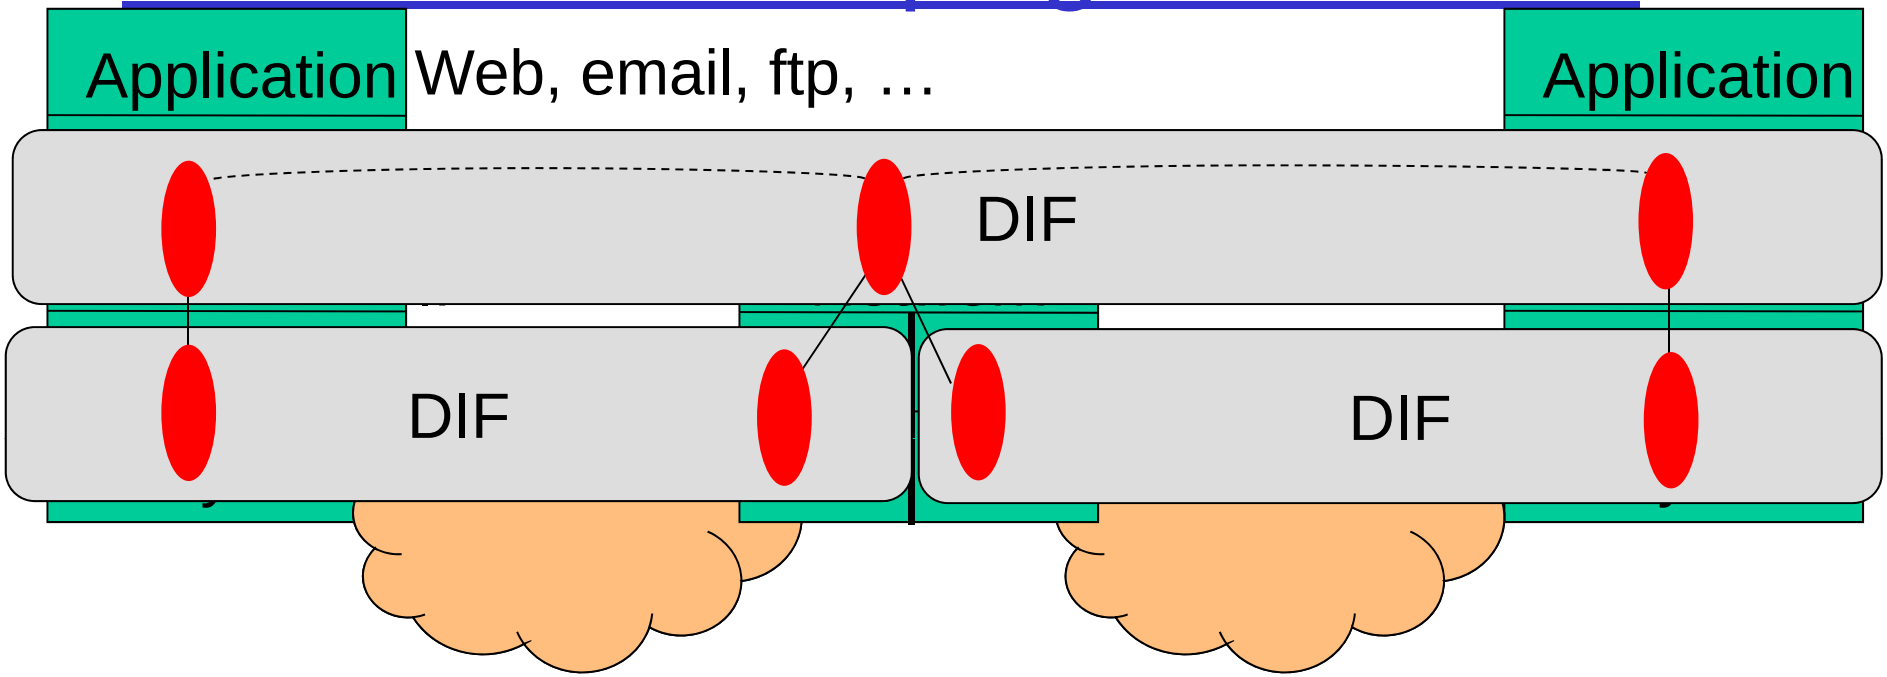

- The DIF is the building block and can be composed
  - A DIF has all what is needed to manage a "private" network, i.e. it integrates routing, transport and management
- E2E (end-to-end principle) is not relevant
  - Each DIF layer provides (transport) service / QoS over its scope
- IPv6 is/was a waste of time!
  - We can have many layers / levels without too many addresses per DIF layer

RINA: some features

- Each DIF is privately managed
  - It assigns private node addresses to IPC processes
  - It internally maps app/service name to node address
  - Addressing is relative: node address is name for lower DIF, and point-of-attachment (PoA) for higher DIF
- Routing is done hop-by-hop over node addresses (names), and next-hop node name is late bound to PoA by lower DIF
- No “middleboxes”
  - The role of a machine is determined by IPC processes running on it and which DIF they are members of
  - Communication is nothing more than explicit negotiation whereby application processes enroll into (join) the same DIF to form a secure container of coordinated IPC processes
- A healthy marketplace
  - Individual services of DIFs can be (recursively) composed to offer new user services

More @
http://csr.bu.edu/rina