The MobilityFirst Future Internet Architecture and Plans for Experimental Evaluation on GENI

GEC10

March 16, 2011

Ivan Seskar
WINLAB, Rutgers University
Technology Centre of NJ
671 Route 1, North Brunswick,
NJ 08902, USA
seskar (at) winlab (dot) rutgers (dot) edu

MobilityFirst Project: Collaborating Institutions

RUTGERS

(LEAD)

D. Raychaudhuri, M. Gruteser, W. Trappe,
R, Martin, Y. Zhang, I. Seskar,
K. Nagaraja, S. Nelson

A. Venkataramani, J. Kurose, D. Towsley

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

M. Reiter

THE UNIVERSITY of WISCONSIN MADISON

S. Bannerjee

Massachusetts Institute of Technology

W. Lehr

MICHIGAN

Z. Morley Mao

Duke UNIVERSITY

X. Yang, R. RoyChowdhury

UMASS LOWELL

G. Chen

UNIVERSITY OF Nebraska Lincoln

B. Ramamurthy

+ Also industrial R&D collaborations with AT&T Labs,
Bell Labs, NTT DoCoMo,, Toyota ITC, NEC, Ericsson and others

Project Funded by the US National Science Foundation (NSF)

Vision: Mobility as *the* key driver for the future Internet

- Historic shift from PC's to mobile computing and embedded devices…
  - ~4 B cell phones vs. ~1B Internet-connected PC's in 2010
  - Mobile data growing exponentially – Cisco white paper predicts >1exabyte per month (surpassing wired PC traffic) by 2012
  - Sensor deployment just starting, ~5-10B units by 2020

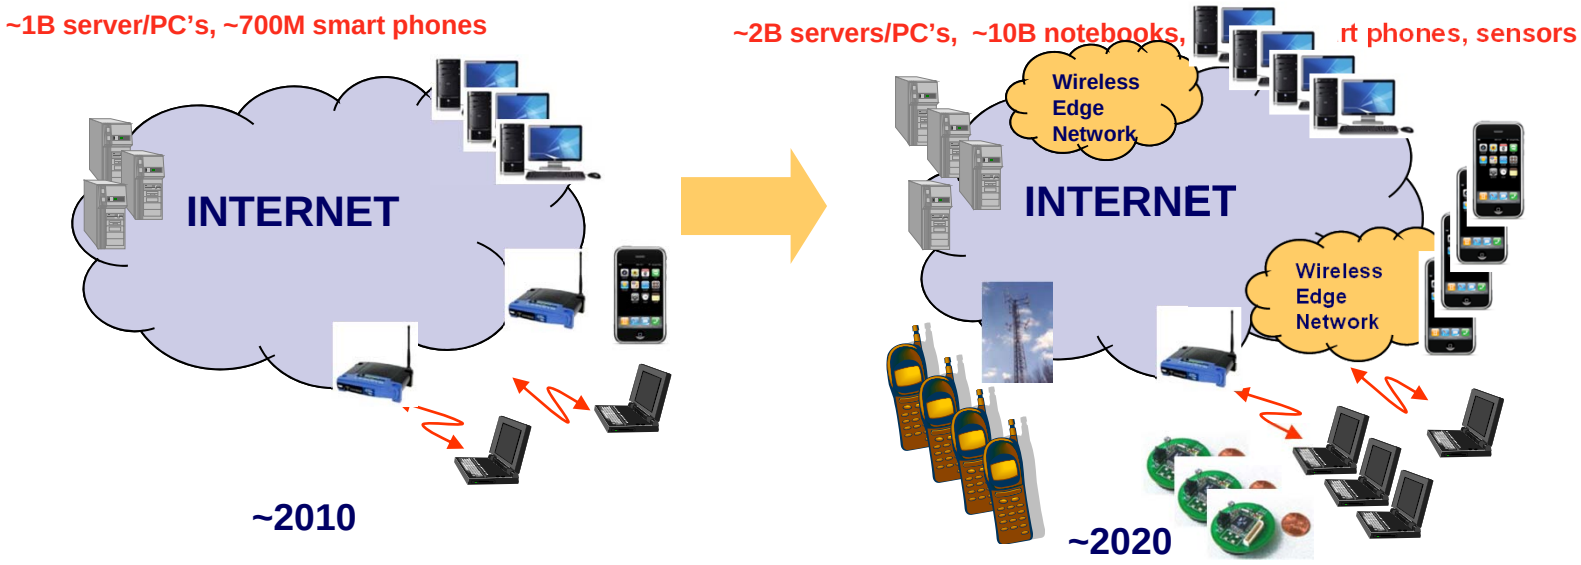

Vision: Protocol Design for the future Mobile/Wireless World

- Fundamental change in design goals and assumptions
  - ~10B+ mobile/wireless end-points as “first-class” Internet devices
  - Mobility as the norm for end-points and access networks
  - Wireless access – varying link BW/quality, multiple radios, disconnections
  - Stronger security/trust requirements due to:
    - open radio medium
    - need for dynamic trust association for mobile devices/users,
    - increased privacy concerns (e.g. location tracking)
    - greater potential for network failure
  - Mobile applications involve location/content/context and energy constraints
- Technology has also changed a lot in the ~40 yrs since IP was designed
  - Moore’s law improvements in computing and storage (~5-6 orders-of-magnitude gain in cost performance since 1970)
  - Edge/core disparity, fast fiber but continuing shortage of radio spectrum

Architecture: *MobilityFirst* Network Overview

- MobilityFirst key protocol features:
  - Separation of naming & addressing
  - Fast global naming service
  - Storage-aware (GDTN) routing
  - Hop-by-hop (segmented) transport
  - Self-certifying public key names
  - Support for content/context/location
  - Programmable computing layer
  - Separate network mgmt plane

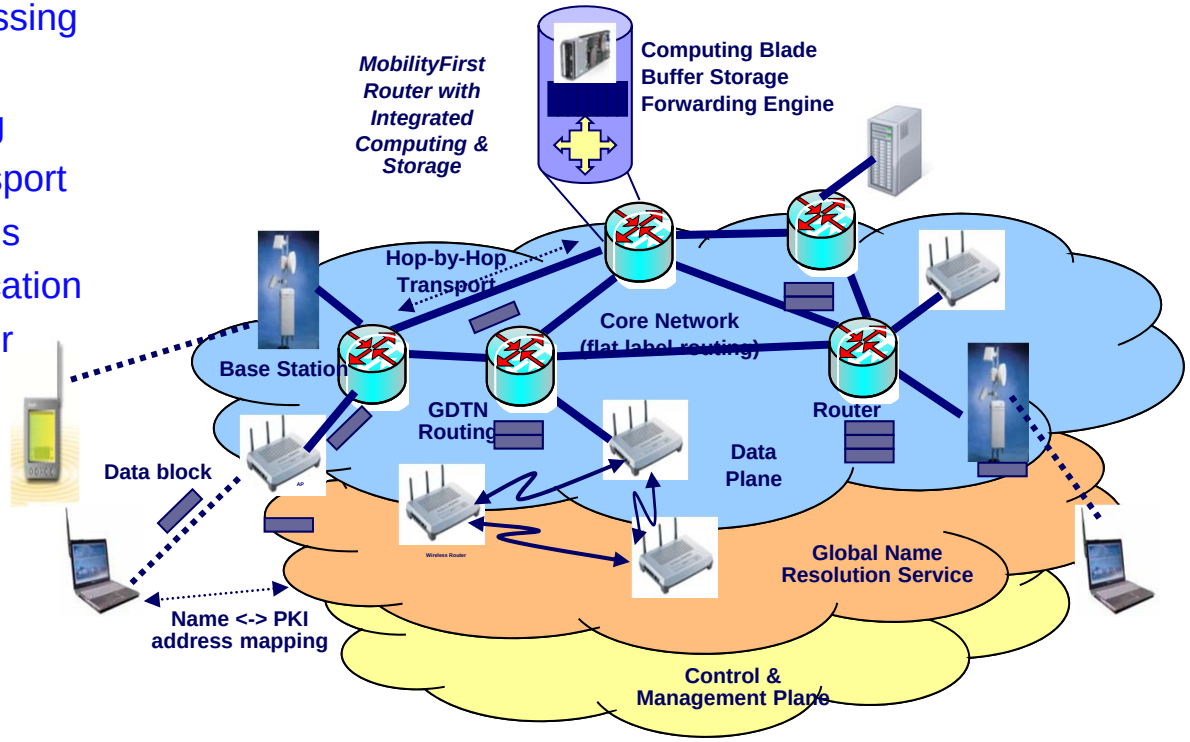

- New components, very distinct from IP, intended to achieve key mobile Internet design goals

Architecture: Name-Address Separation

- Separation of names (ID) from network addresses (NA)
- Globally unique name (GUID) for network attached objects
  - User name, device ID, content, context, AS name, and so on
  - Multiple domain-specific naming services
- Global Name Resolution Service for GUID → NA mappings
- Hybrid GUID/NA approach
  - Both name/address headers in PDU
  - "Fast path" when NA is available
  - GUID resolution, late binding option

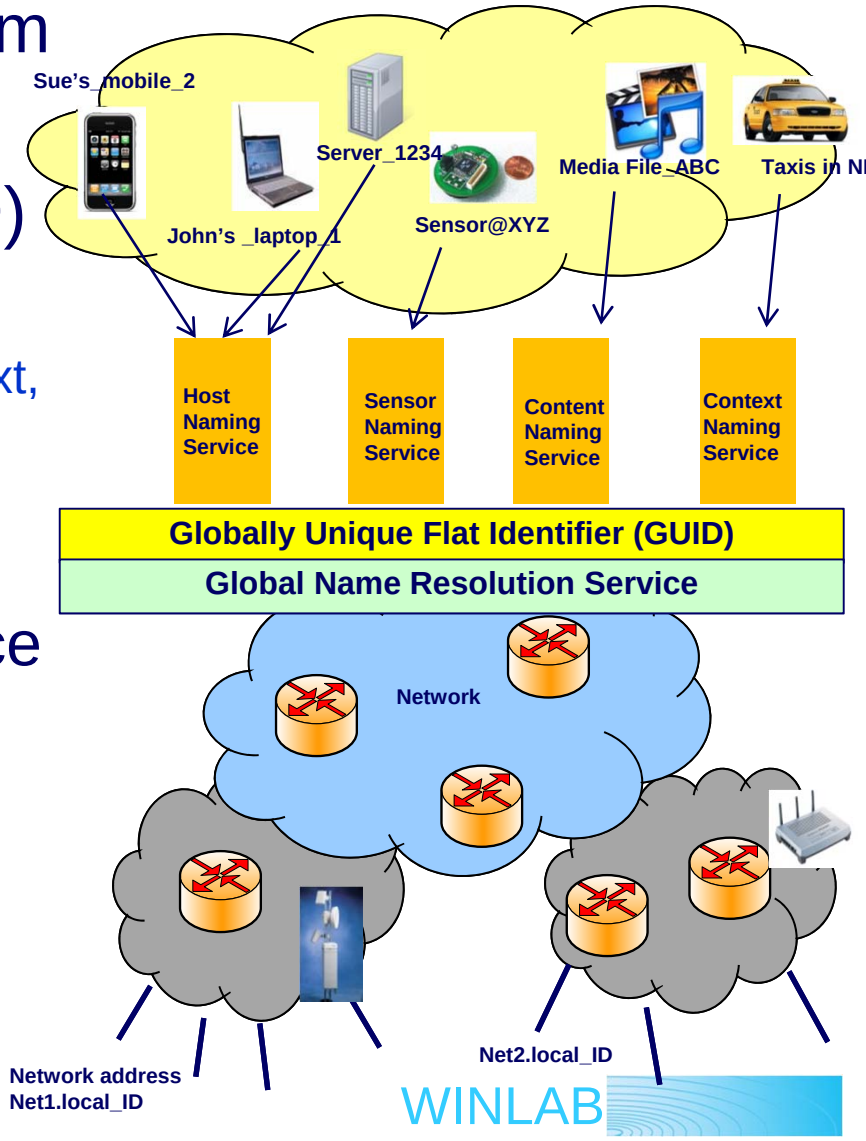

Protocol Design: Storage-Aware Routing

- Storage aware (CNF, generalized DTN) routing exploits in-network storage to deal with varying link quality and disconnection
- Routing algorithm adapts seamlessly adapts from switching (good path) to store-and-forward (poor link BW/disconnected)
- Storage has benefits for wired networks as well..

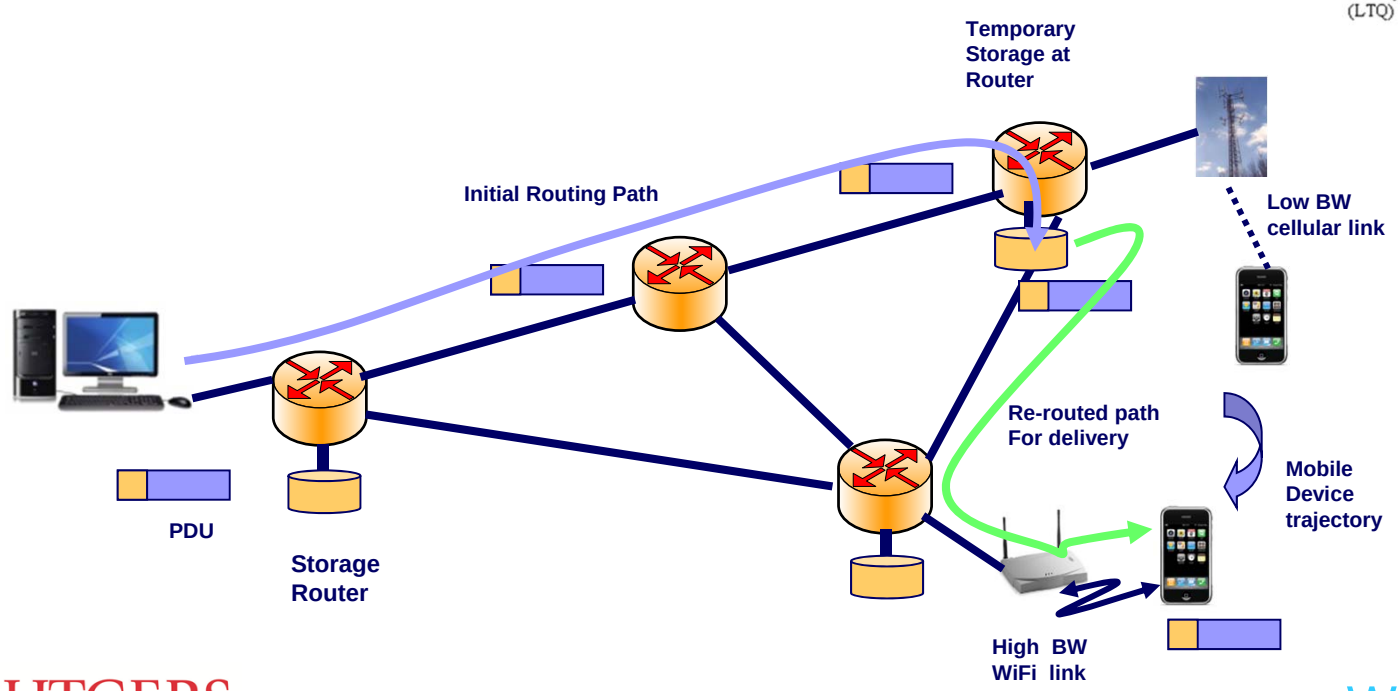

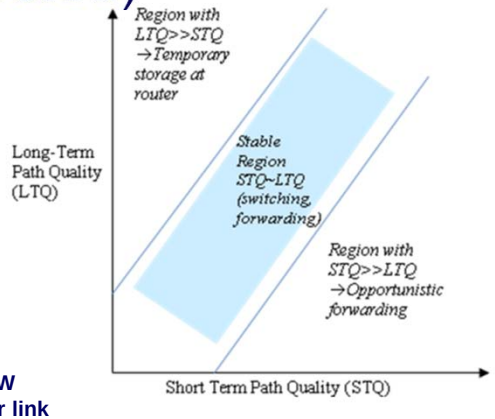

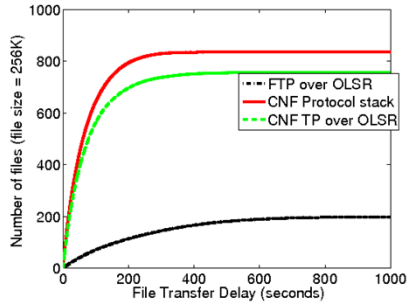

Sample CNF routing result

Protocol Design: Segmented Transport

- Segment-by-segment transport between routers with storage, in contrast to end-to-end TCP used today
- Unit of transport (PDU) is a content file or max size fragment
- Hop TP provides improved throughput for time-varying wireless links, and also helps deal with disconnections
- Also supports content caching, location services, etc.

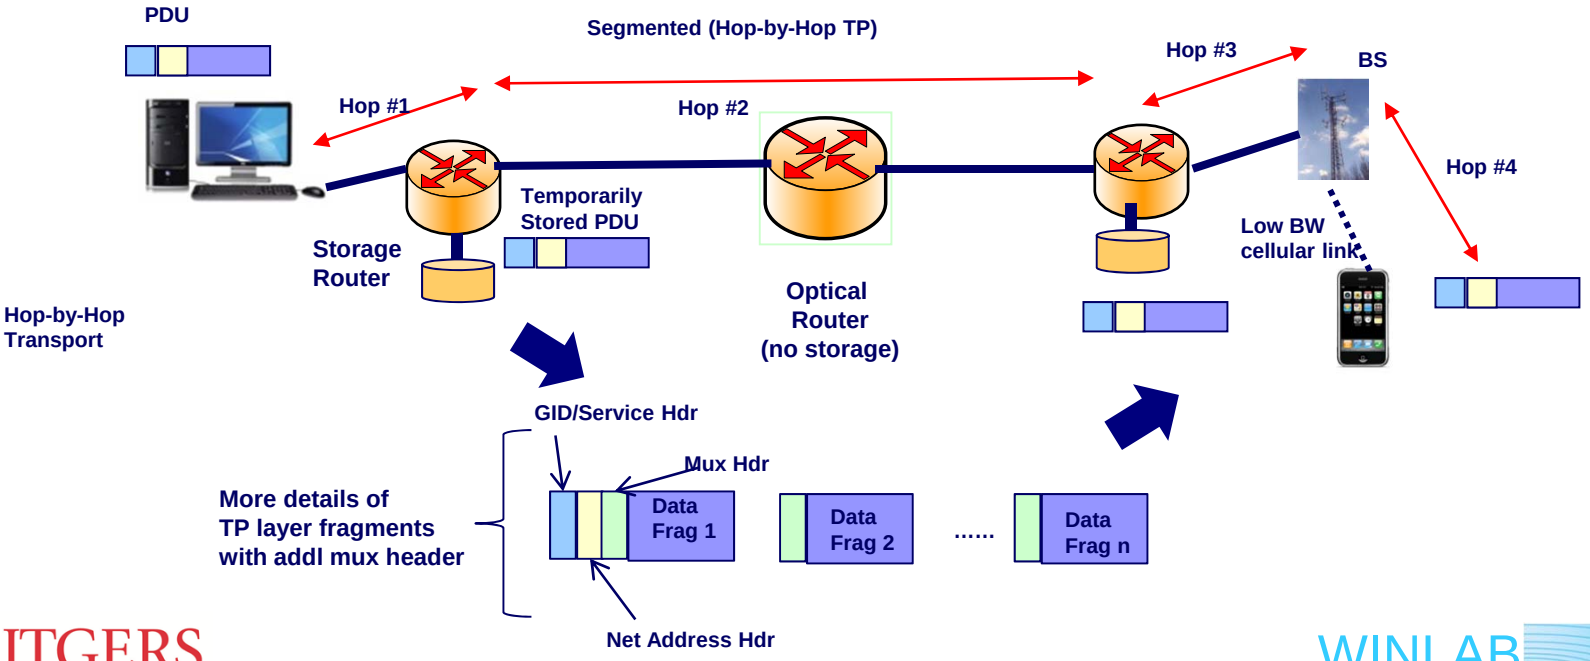

Protocol Design: Context Aware Scenario

- Context-aware network services supported at two layers by MF architecture
  - Dynamic mapping of arbitrary context or content label by global name service
  - Per-packet multicast capabilities based on name resolution at each router
  - Same mechanism used to handle named content
  - Optional software services implemented at the computing layer

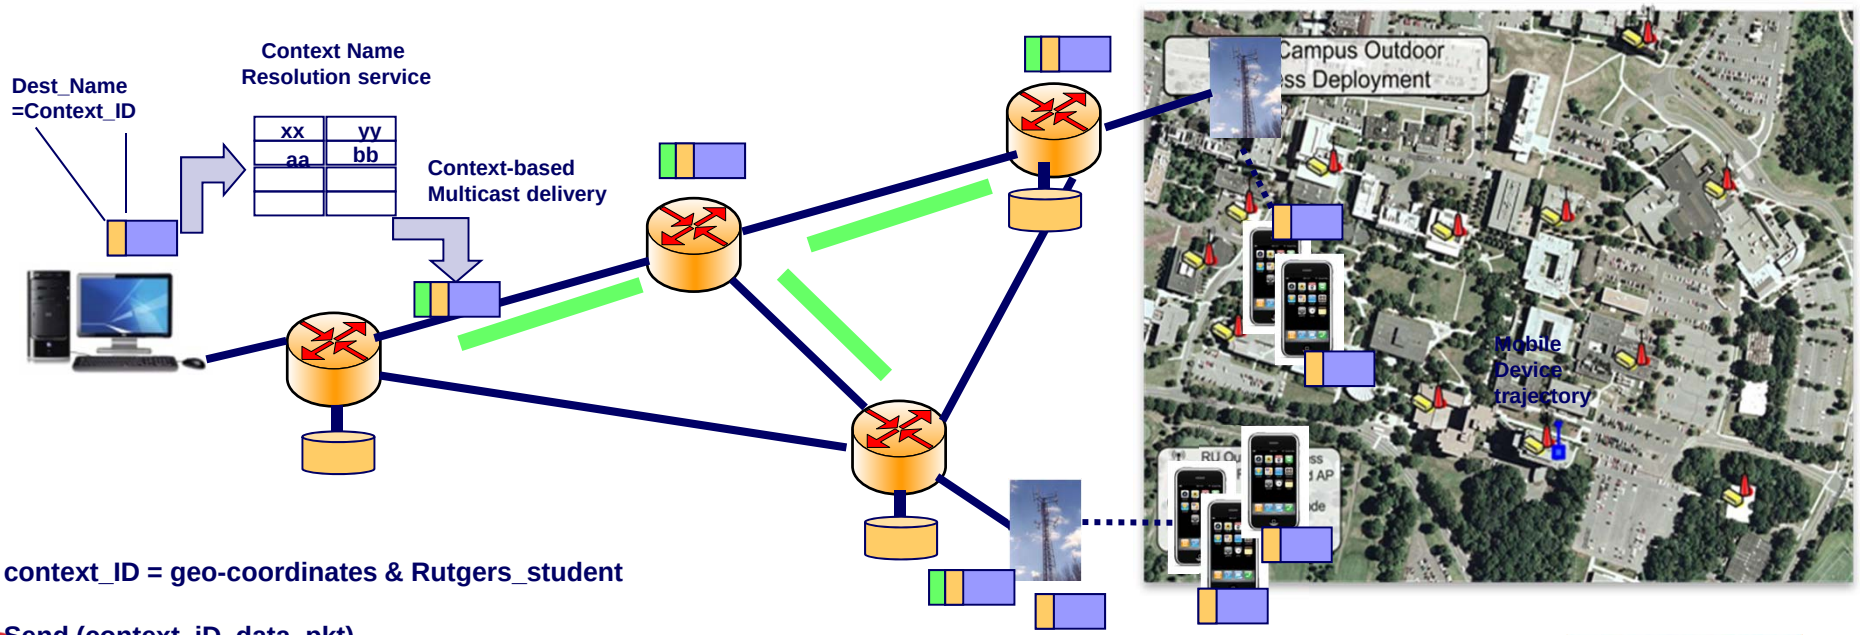

Protocol Design: Security Aspects

1. Public keys addresses for hosts & networks; forms basis for
   - Ensuring accountability of traffic
   - Ubiquitous access-control infrastructure
   - Robust routing protocols
   - Preventing address hijacking
2. Support deployment of policies that constrain the traffic that a network or node receives
   - In the limit, a “default-disconnected” posture
3. No globally trusted root for naming or addressing
   - Opens naming to innovation to combat naming-related abuses
   - Removes obstacles to adoption of secure routing protocols
4. Systematically consider Trusted Computing Base of designs
   - Promote TCB reduction technologies (e.g., Byzantine fault tolerance)

Research Agenda: More Specifics

- Fast global fast name resolution service for ~10-100B end-points
- Unified support for mobile devices, context, location, content,… as an integral capability of the network
- Scalable routing protocol for flat name identifiers, taking into account unique mobility requirements (multi-homing, anycast, multicast, disconnection, location-awareness, …)
  - Design of storage-aware routing algorithms which are robust to disconnection and bandwidth variation across wired & wireless nets
  - Techniques for interdomain network aggregation applicable to name-based and storage-aware routing under consideration
- Integrating security and privacy features into network services – self-certification, multiple roots of trust, Byzantine fault-tolerance, …
- Design of network management plane features to achieve a high level of transparency and accountability
- Network neutrality, economic feasibility, …
- Evaluation & validation at scale!

Overview - Component Architecture

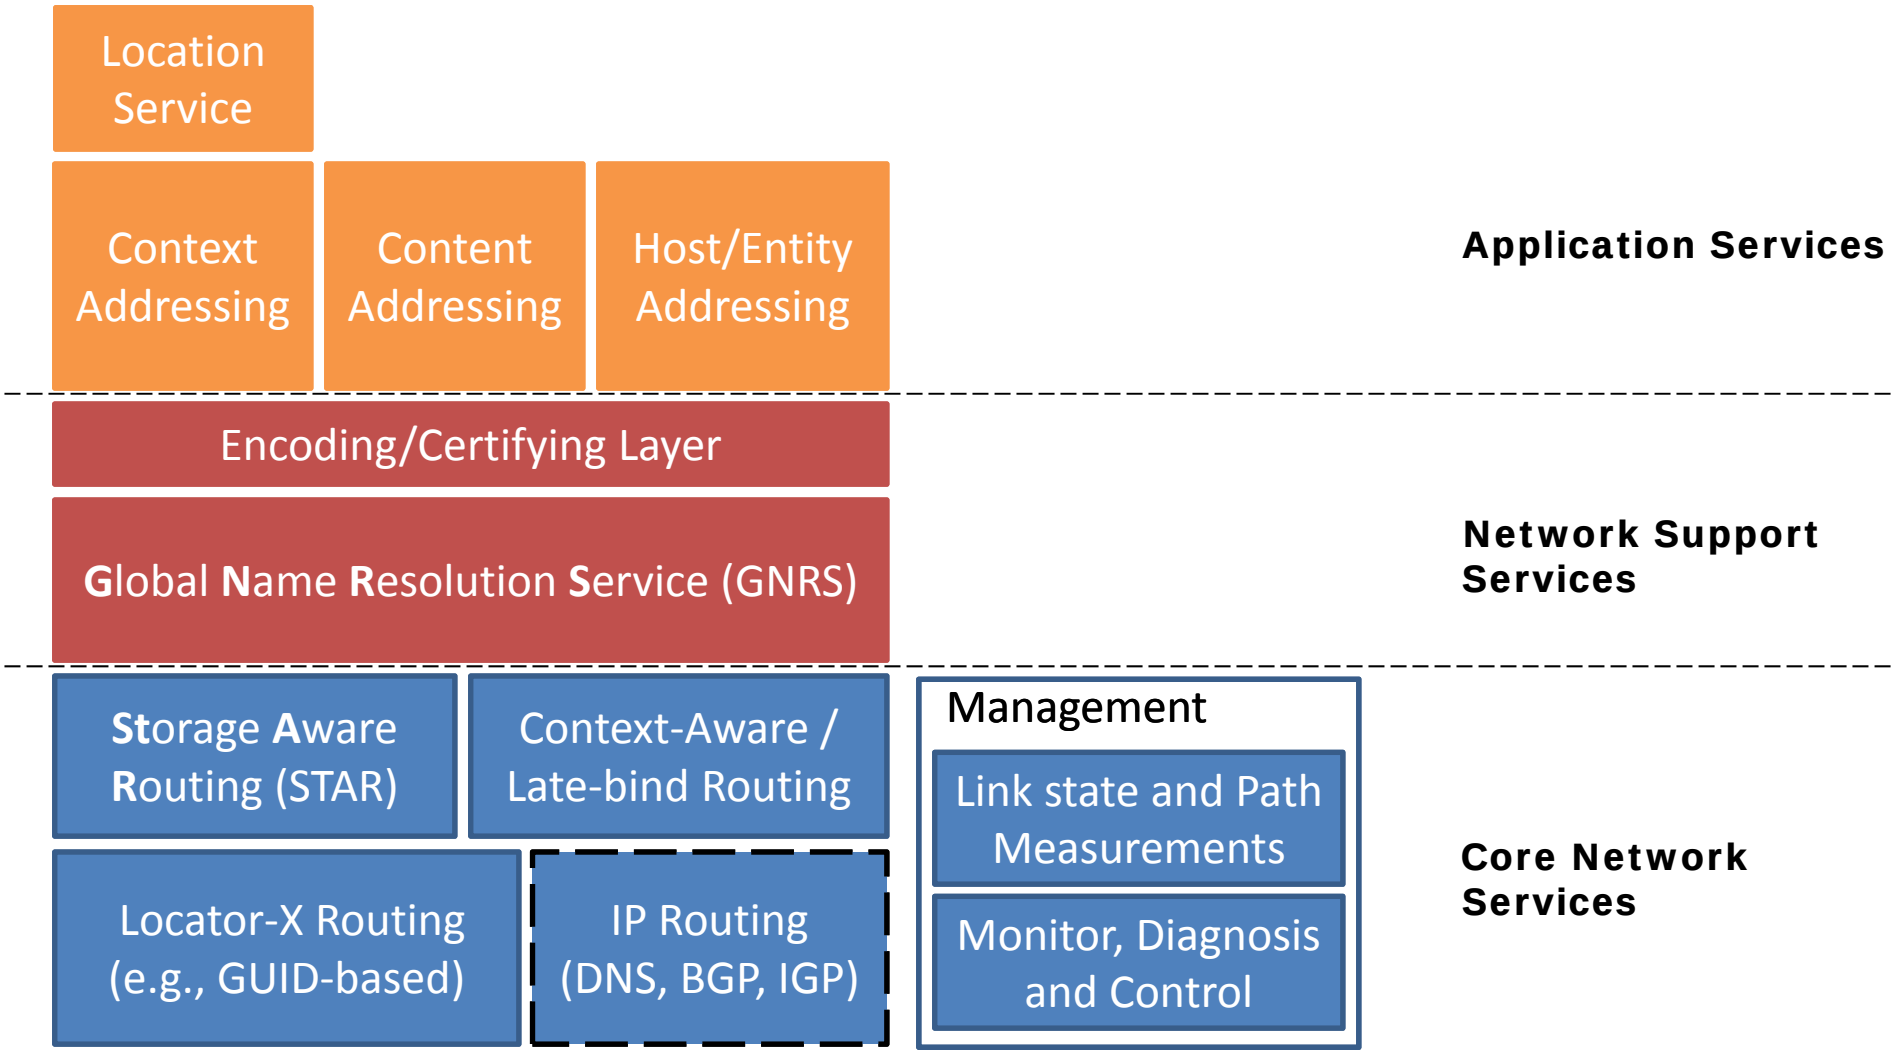

Phased Implementation Approach

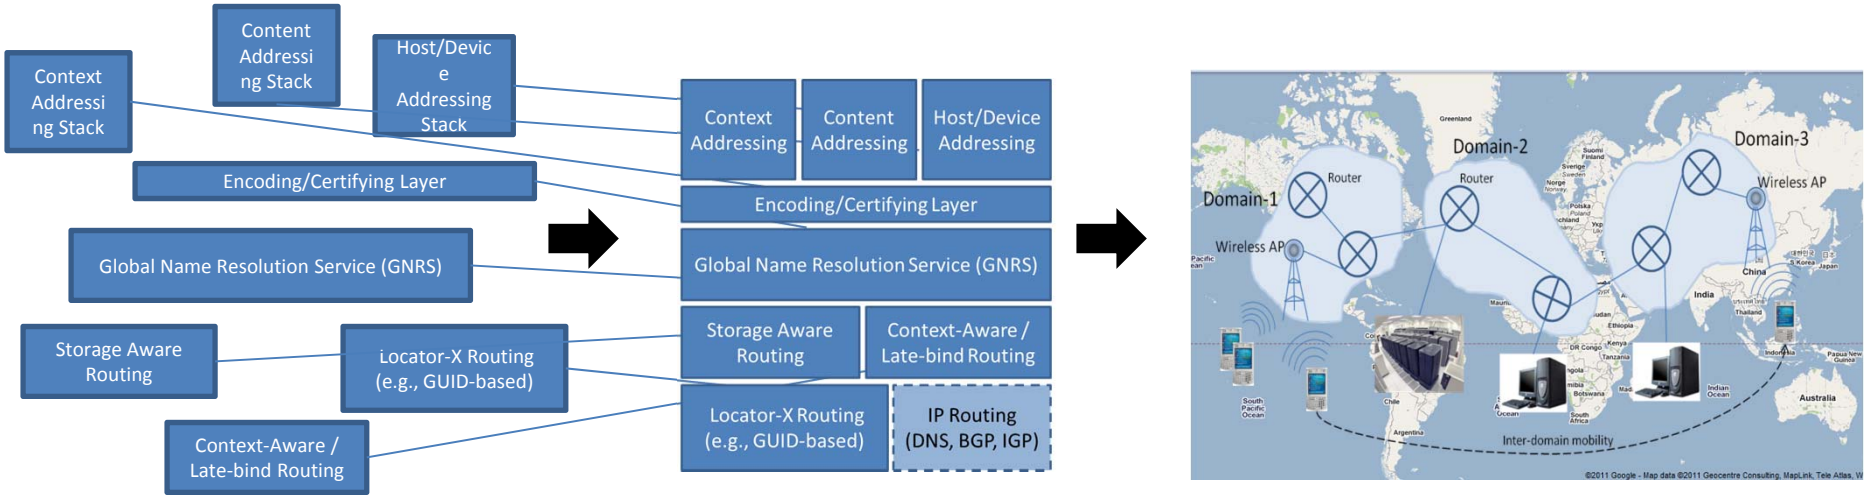

Simulation/Emulation | **Emulation/Limited Testbed** | **Testbed/'Live' Deployment**

Evaluation Platform

Standalone Components | **Cross Layer Integration** | **Production, Hardware Ready**

Prototyping Status

Testbed/Deployment Options

- WINLAB ORBIT – 400-node Grid and Outdoor (GENI)
  - Wireless and mobile focus; WiMAX, WiFi, OpenFlow
- UMass DOME/DieselNet (GENI)
  - Mobile, wireless focus. WiMAX, campus bus network
- Wisconsin WiNGS - 60+ node Indoor/Outdoor (GENI)
  - WiMAX, metro and campus bus networks, WiFi mesh, 3G
- ProtoGENI
  - 10-12 core nodes interconnected by Internet2 backbone
  - PrimoGENI - Large scale network simulator embeddable into GENI
- PlanetLab, VINI
  - Wired network emulation/overlay, most widely distributed
- Emulab
  - Mostly wired with flexible topology and delay/bandwidth control
- GpENI
  - 4+ clusters of nodes, with L1 Midwest optical backbone between clusters

GENI glue

- ProtoGENI control framework
- 1-Gbps L2-connection to GENI core from sites

Evaluation Scenario 1 - Core Network Svcs

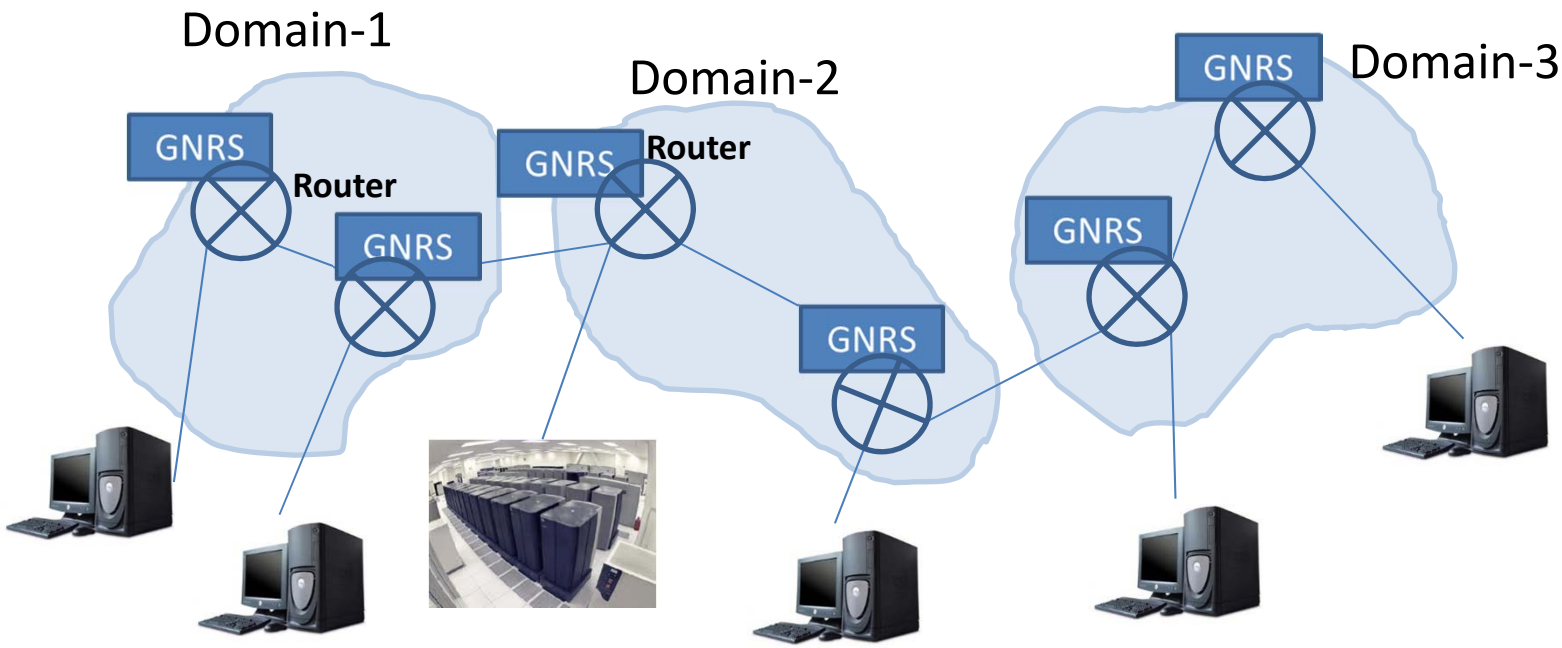

Testbeds:
ProtoGENI
Emulab
PlanetLab
VINI

- Multiple domains, core + edge routers, reliable connectivity with redundant paths
- Explore inter-domain routing, global services – e.g., GNRS

Evaluation Scenario 2 – Edge Only

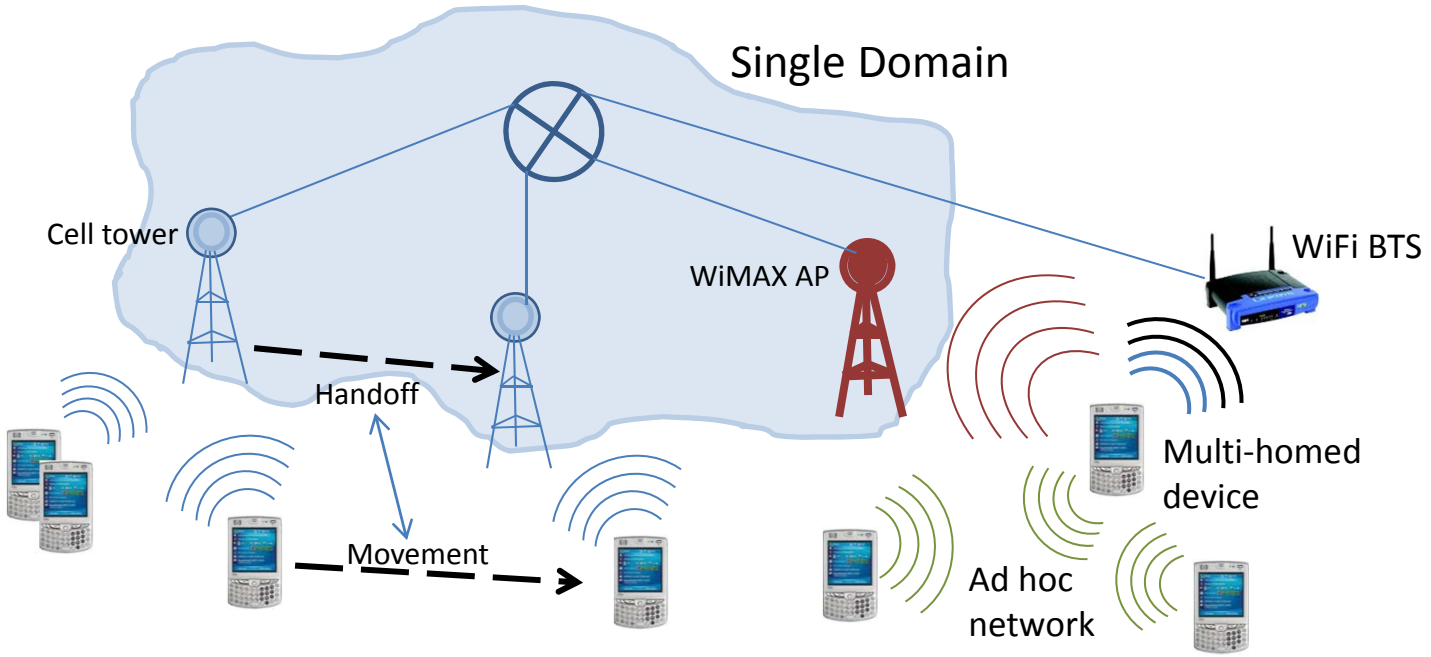

Testbeds:
ORBIT
DieselNET
WiNGS

- Ad hoc, multiple wireless technologies – WiFi, WiMAX
- Explore routing with mobility, handoff, multi-homing within single domain

Scenario 3 – Core + Edge

- Similar to 2, except edge network access services within core.
- Explore:
  - Core-edge routing
  - Cross-layer interaction between global naming and routing services
  - In-core storage resources
- Testbed candidates: ORBIT GENI (outdoor + indoor), ProtoGENI + WiMAX, GENI

Goal: Live Edge-Core-Edge Slice

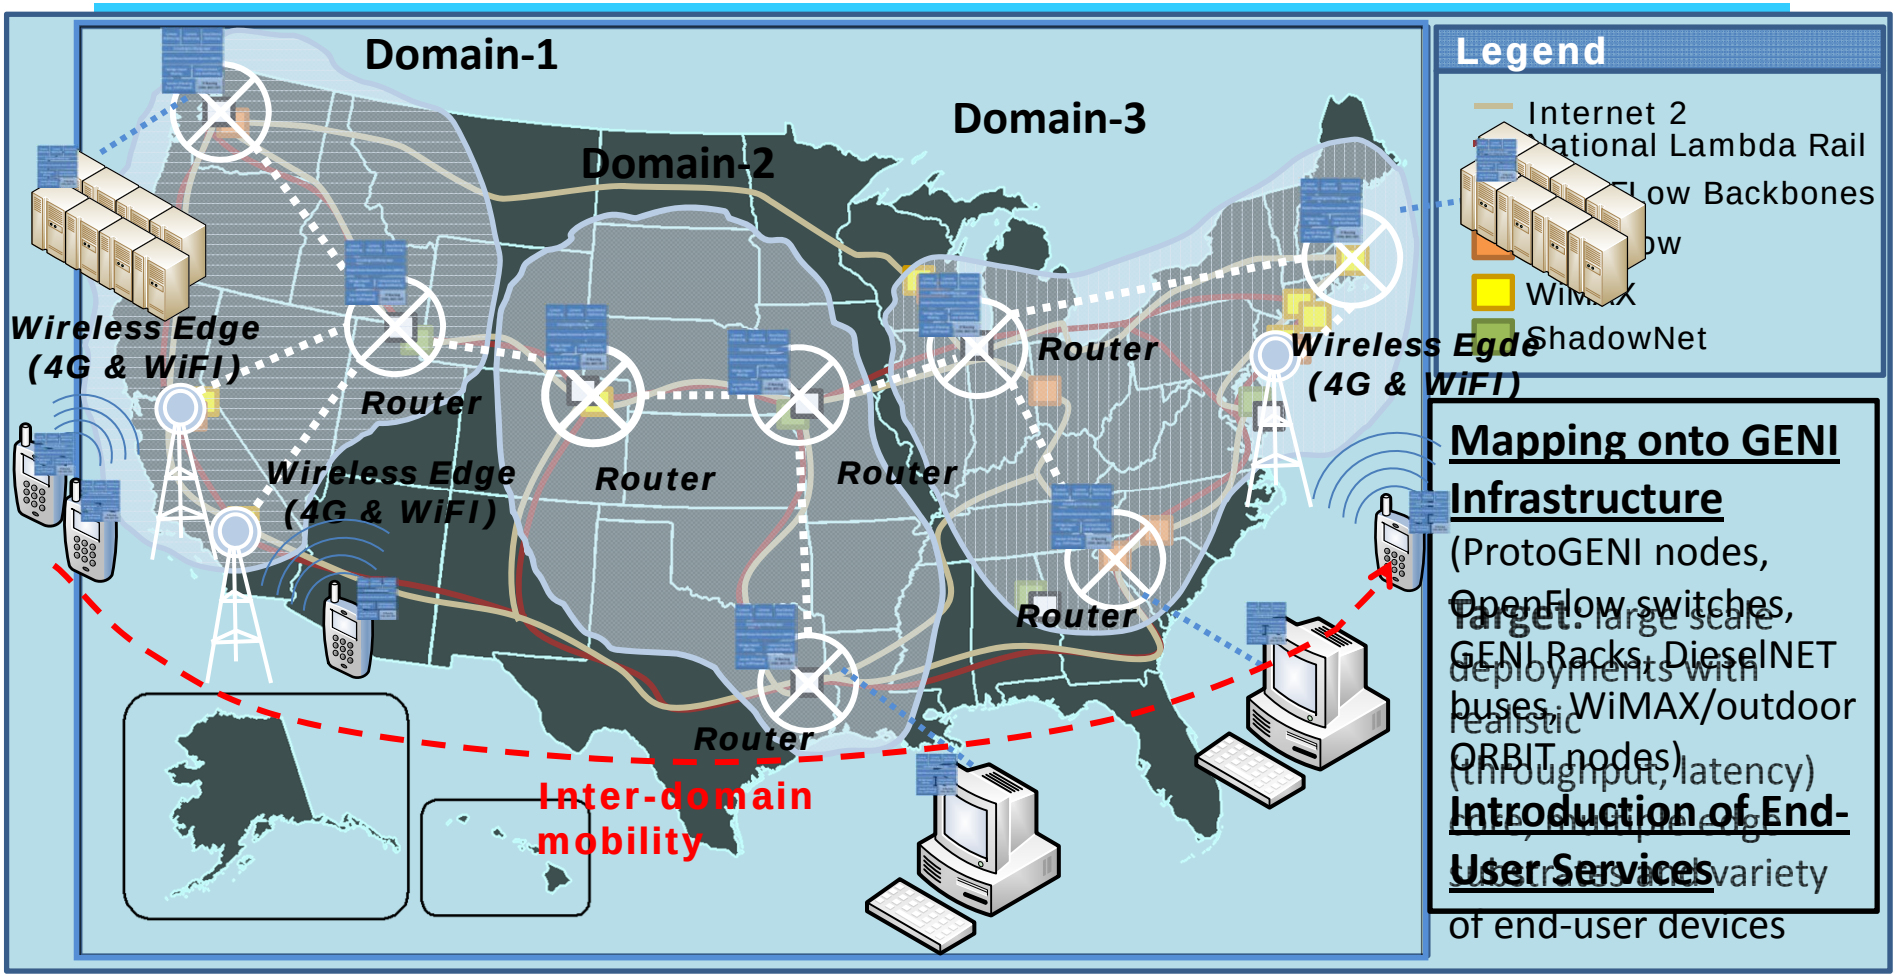

- Entire MobilityFirst stack on network devices
- Explore inter-domain mobility, e.g., emulate as process migration
- **Real traffic** through applications: media, social, location, etc.
  Live slice deployed in multiple sites/campuses with opt-in users

Prototyping - Work in Progress

- Evaluation of emulation and testbed options
  - XORP/Click, OpenFlow/NetFPGA, Emulab/ProtoGENI
- Extraction and mapping of GNRS, Routing algorithms/code from simulation to prototype framework
  - Targeting limited GENI (ORBIT, ProtoGENI,?) deployment/demo of GNRS/routing in Nov '11 – GEC12
- More details:

mobilityfirst.winlab.rutgers.edu