NSF's Research Infrastructure Activities

Deep Medhi

National Science Foundation

March 2021

NSF Computer & Information Science & Engineering (CISE) Infrastructure Programs

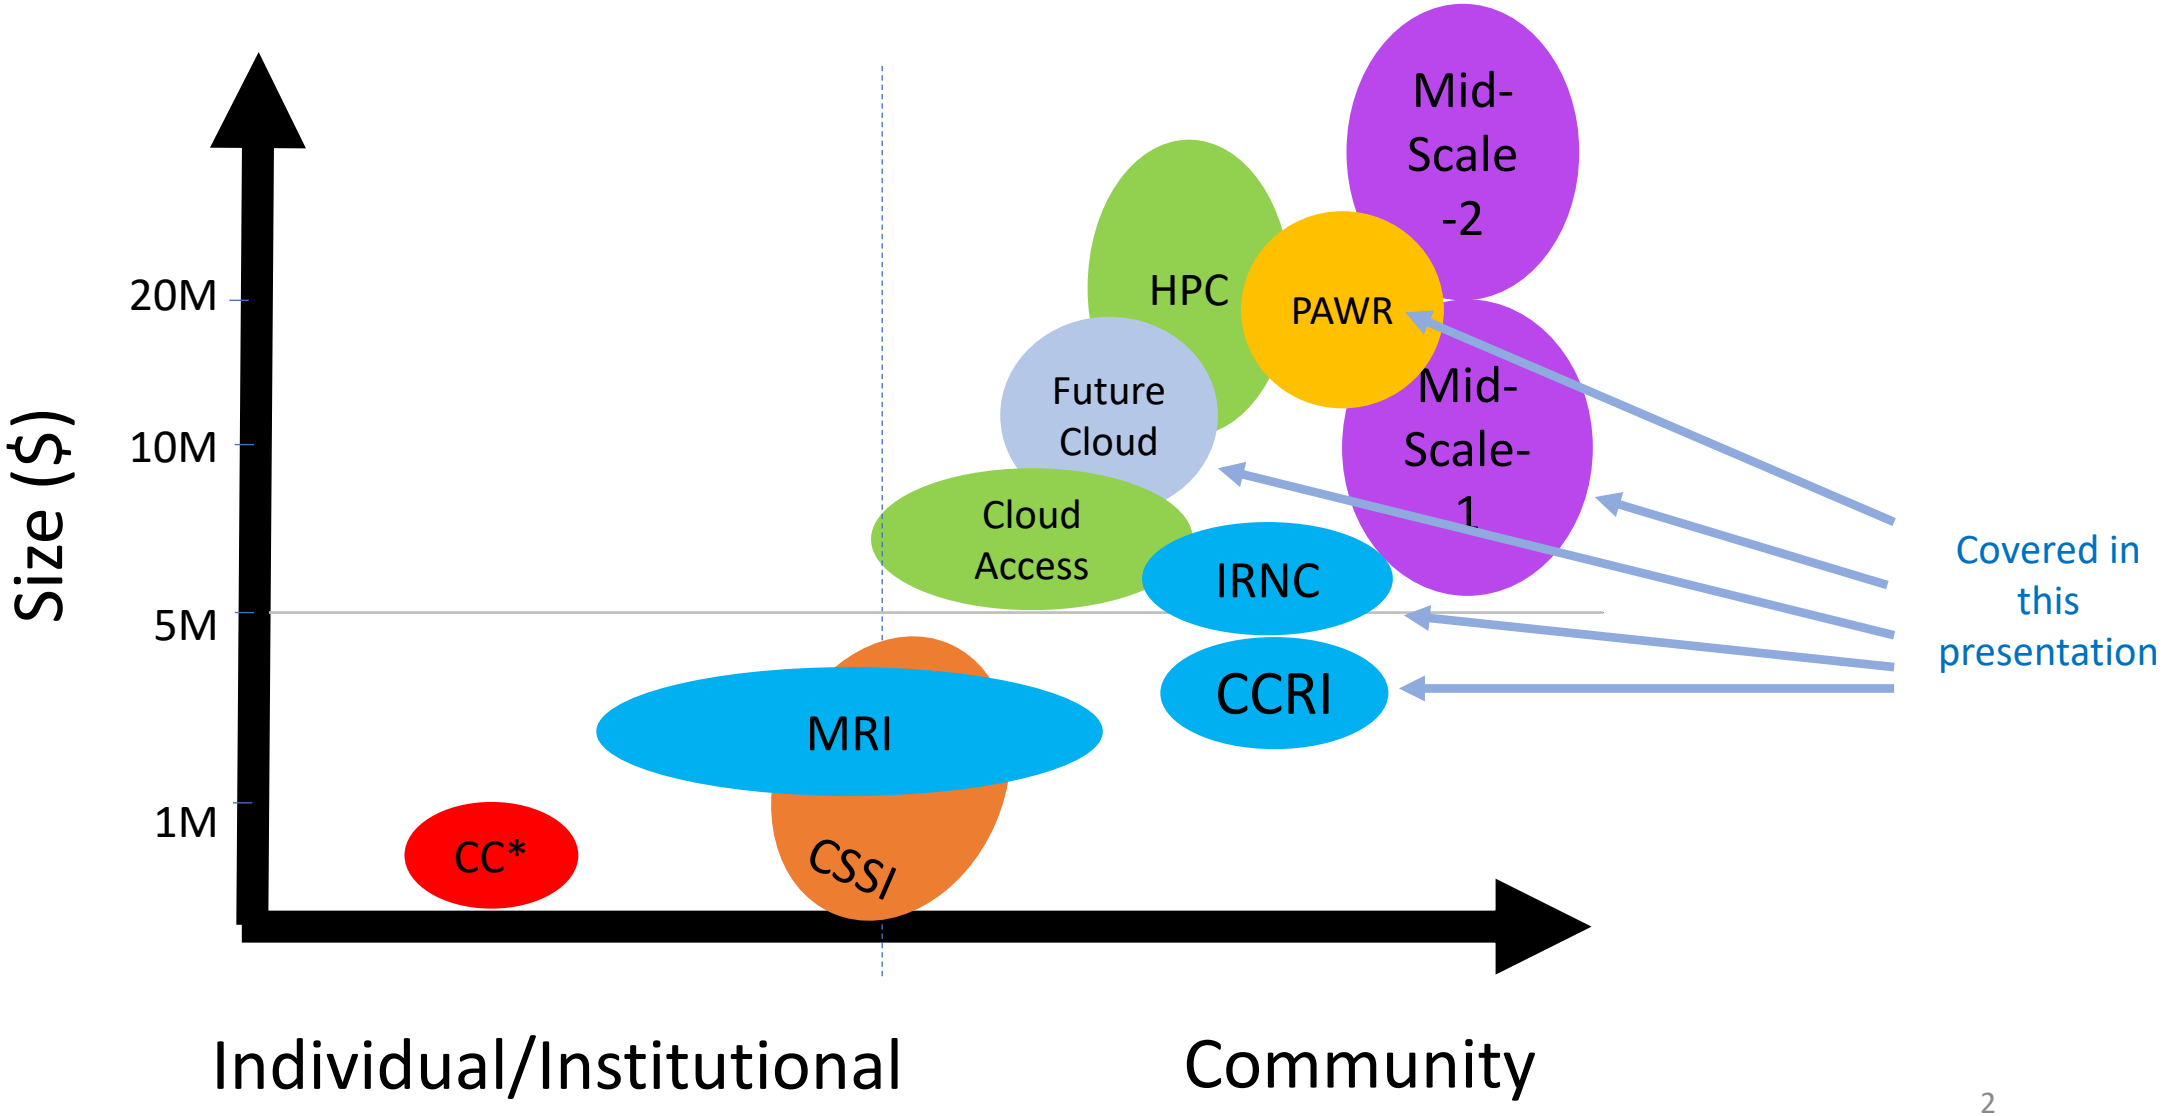

FABRIC: NSF Mid-Scale Research Infrastructure-1 Project https://whatisfabric.net/

A testbed that enables a *new paradigm for distributed applications and Internet protocols:*

- A nation-wide programmable network with in-network compute and storage. Run computationally intensive programs & maintain information in the network.
- GPUs, FPGAs, and network processors (NICs) inside the network
- Quality of service (QoS) - dedicated optical 100Gb
- Interconnects national facilities: HPC, cloud & wireless testbeds, commercial clouds, Internet, and edge
- Design and test applications, protocols and services that run at any node in the network
- Science cases: IoT sensors, Cybersecurity, AI/ML, SDN/P4, Science apps

FABRIC Topology

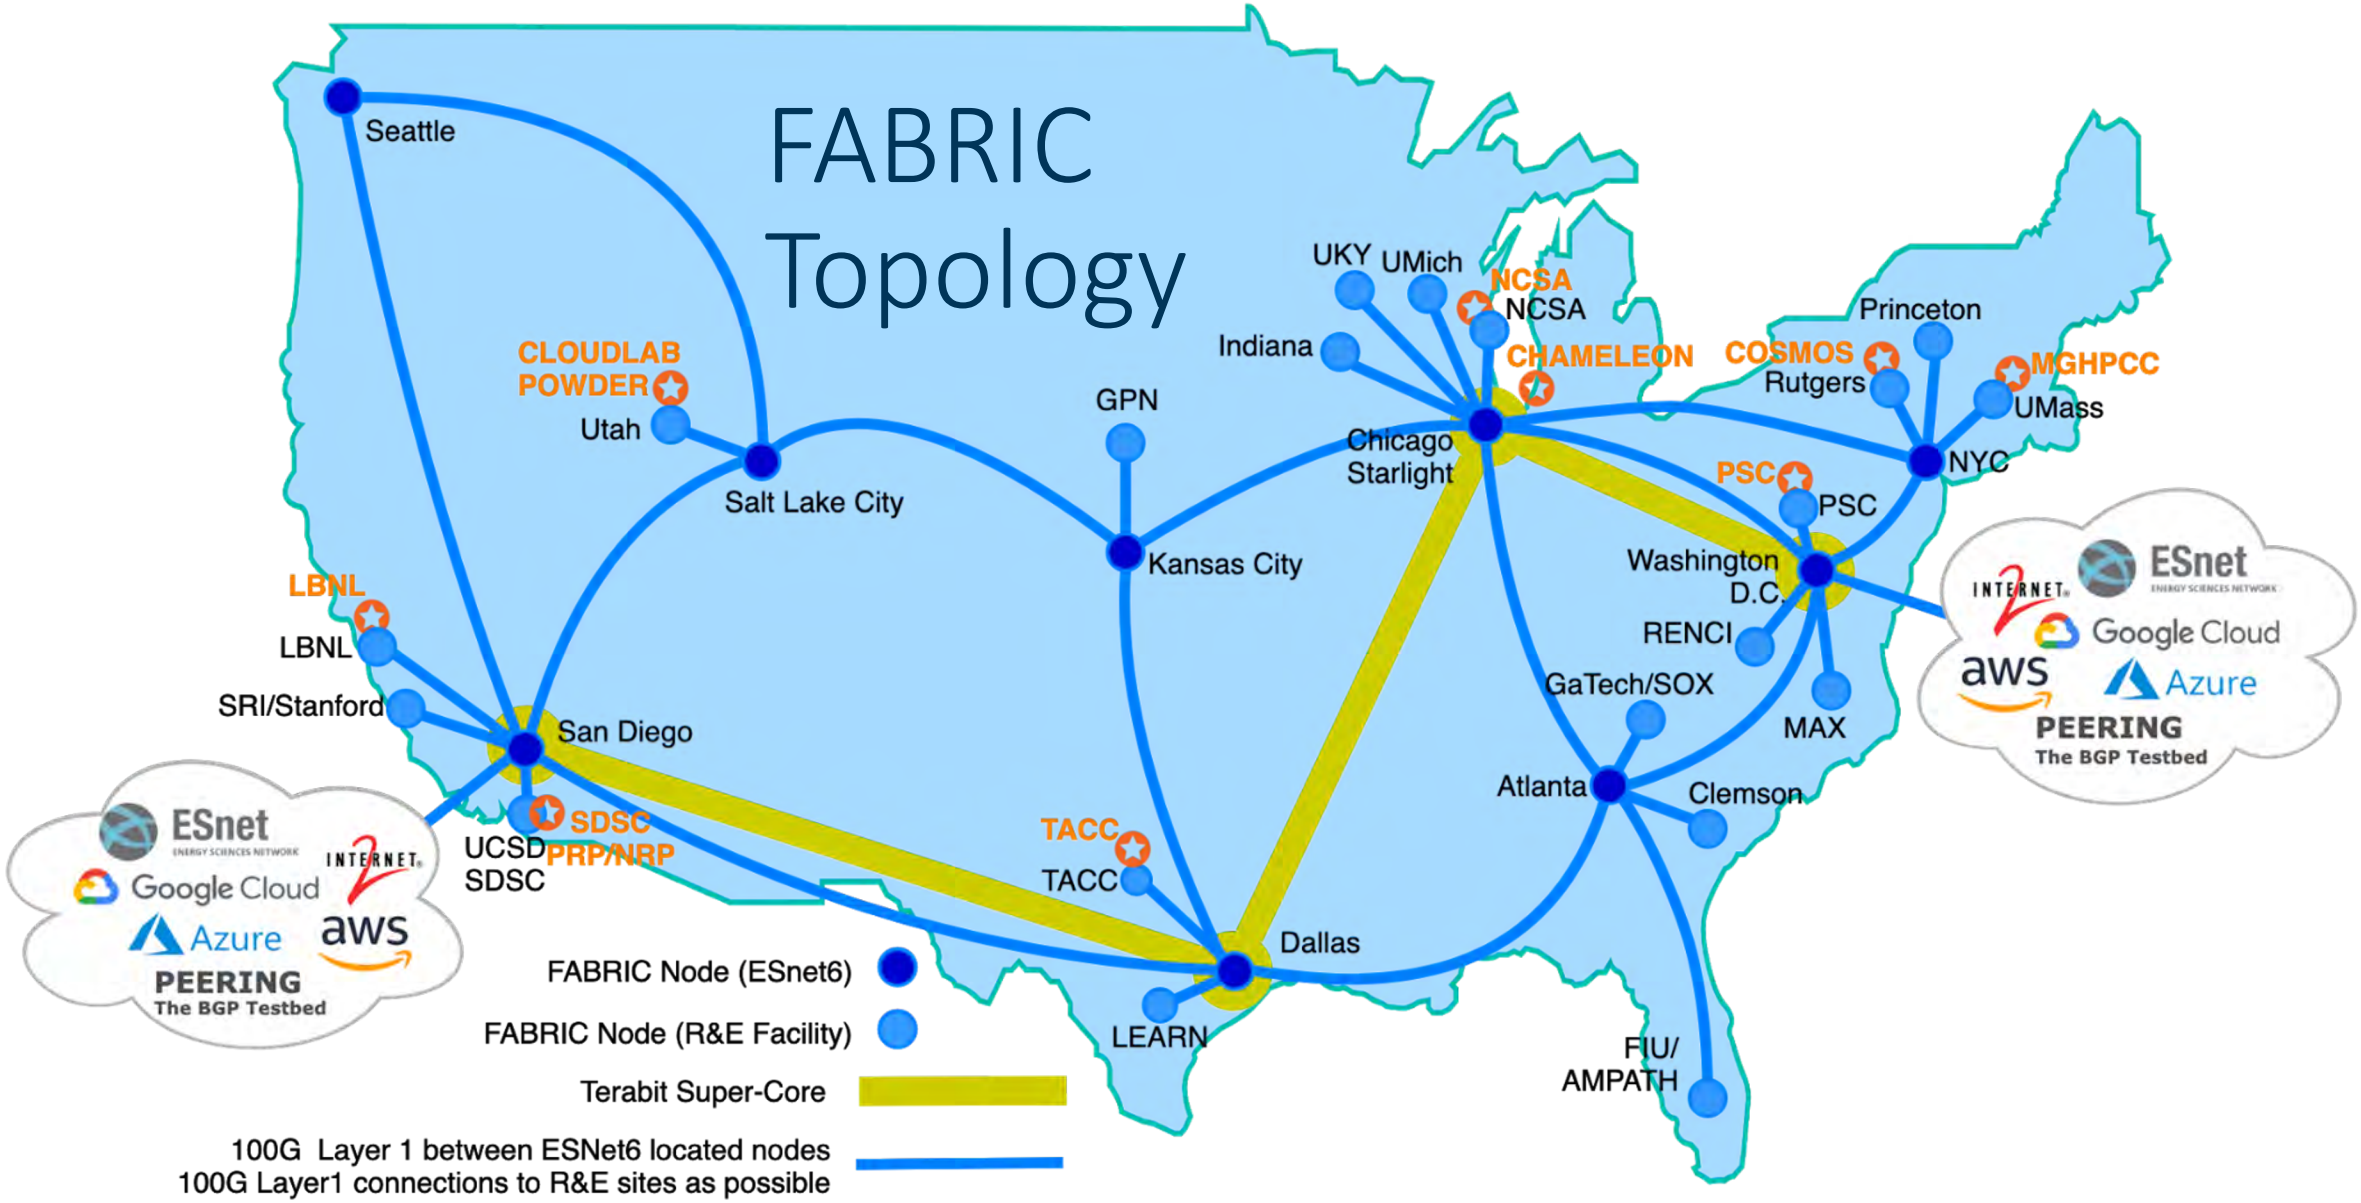

FAB (FABRIC Across Borders):

NSF International Research and education Network Connection (IRNC)Project

A testbed that enables a *new paradigm for distributed applications and Internet protocols, at International Scale*

- **International** programmable network with in-network compute and storage. Run computationally intensive programs & maintain information in the network.
- GPUs, FPGAs, and network processors (NICs) inside the network
- Quality of service (QoS) - optical 100Gb & **use of transoceanic links**
- Interconnects **International** facilities: HPC, clouds, **5G** & wireless testbeds, Internet, edge, **Smart Cities testbeds and science instruments**
- Design and test apps, protocols and services that run at any node in network
- **Partners with domain scientists to optimize the global scientific workflow**
- **Develop software building blocks and demonstrate their use toward a high-speed, programmable, multipoint, customizable data transfer service**

FAB International Connectivity

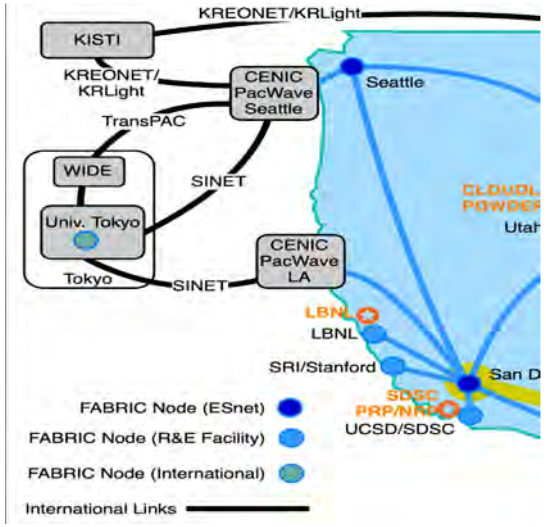

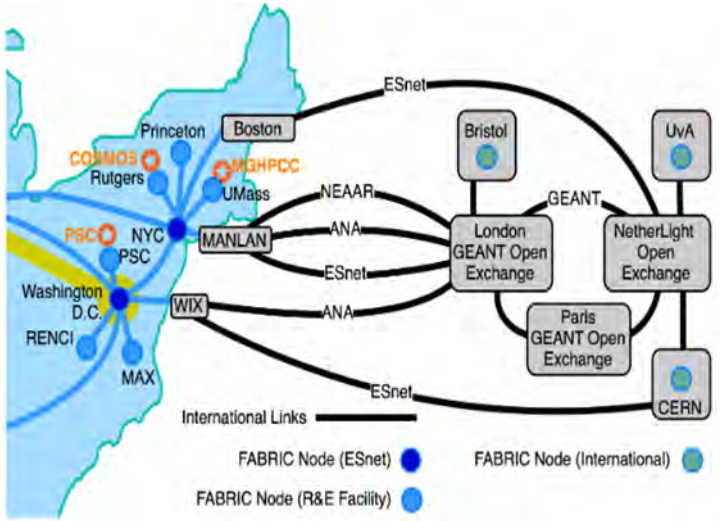

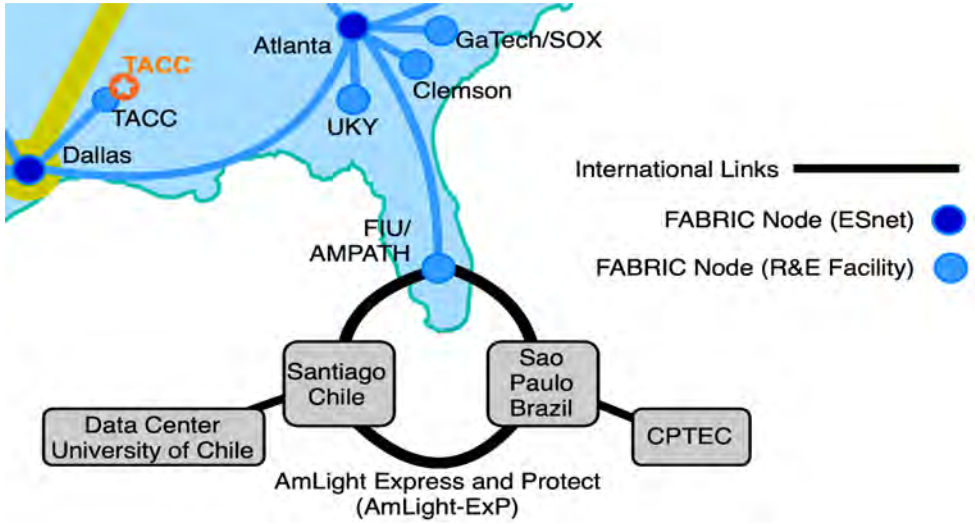

FABRIC/FAB Initial Science Drivers & Primary Partners

- Astronomy/Cosmology
  - Cosmic Microwave Background - CMB-S4, Chile & S Pole
  - Vera Rubin Observatory (LSST), Chile
- Weather Science (Univ. of Miami)
- High-Energy Physics (Co-PI Rob Gardner)
  - LHC ATLAS (CERN)
- Urban Sensing/IoT (Univ. of Bristol)
  - Smart City Bristol
- AI at Edge (SAGE)
  - SAGE
- Computer Science
  - 5G across borders (Univ. of Tokyo)
  - SDX Economics (Univ. of Kentucky)
  - Censorship-Resistant Protocols (Clemson)

SAGE: NSF Mid-Scale-1 Project

http://sagecontinuum.org/

SAGE *A Software-Defined Sensor Network - Cyberinfrastructure for* ***AI at the Edge***

Science Justification

- Understanding resilience —why do ecosystems, cities, neighborhoods respond differently to phenomena such as gradual changes in ecological diversity, increased severe weather, or catastrophic events? What are the mechanisms of resilience, or indices of fragility?

Research Community Benefits

- New cyberinfrastructure for exploring machine learning in a national-scale instrument network
- Flexible national-scale measurement infrastructure
- New measurements via software-defined sensors

Cyberinfrastructure Use Modalities

- Open Data, Open Calibration, Education
- Triggered Computation
- Scheduled Edge Computation
- Machine Learning for the Edge (AI@Edge)

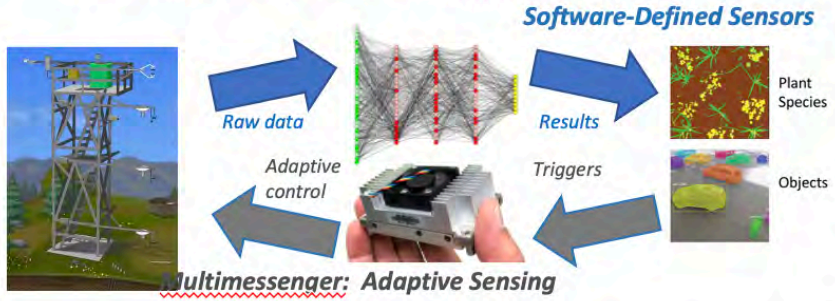

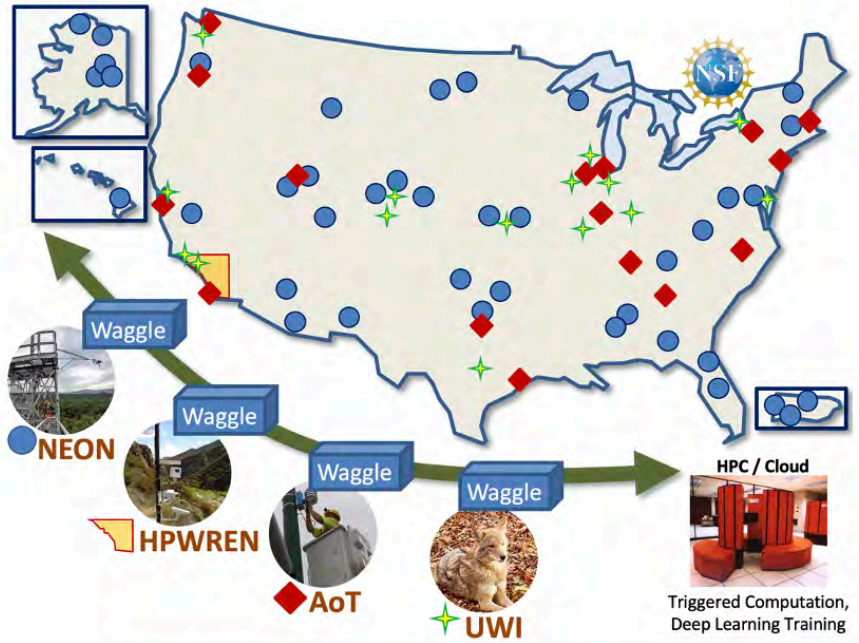

SAGE: A cyberinfrastructure to bring diverse disciplinary expertise to bear on complex questions, *revealing the hidden layer of mechanisms and relationships to understand the causality* inherent in these systems and at their intersections. SAGE will enable convergent science teams to integrate concepts and methods from different disciplines and conducting research on shared infrastructure that is generative and diverse.

NSF Future Cloud (research cloud program started in 2014, now in Phase-III)

Chameleon

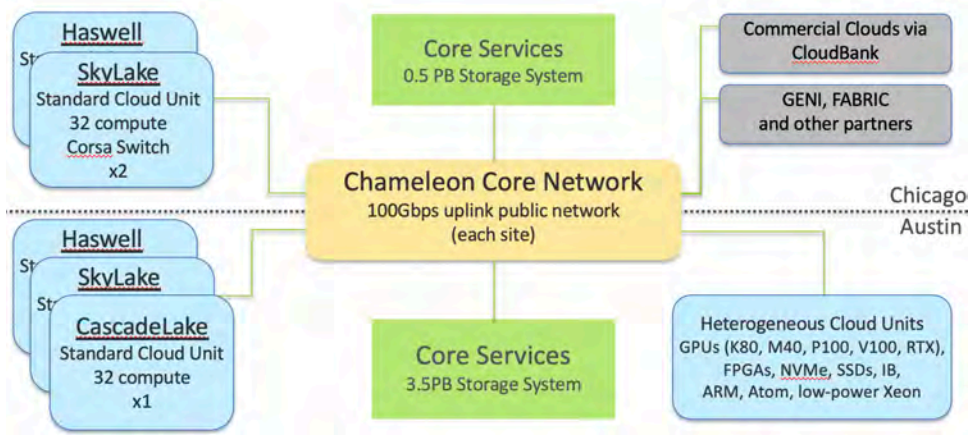

- "Start with large-scale homogenous partition"
  - 12 Haswell Standard Cloud Units (48 node racks), each with 42 Dell R630 compute servers with dual-socket Intel Haswell processors (24 cores) and 128GB RAM and 4 Dell FX2 storage servers with 16 2TB drives each; Force10 s6000 OpenFlow-enabled switches 10Gb to hosts, 40Gb uplinks to Chameleon core network
  - 3 SkyLake Standard Cloud Units (32 node racks); Corsa (DP2400 & DP2200) switches, 100Gb ulpinks to Chameleon core network
  - CascadeLake Standard Cloud Units (32 node rack) , 100Gb ulpinks to Chameleon core network
  - Allocations can be an entire rack, multiple racks, nodes within a single rack or across racks (e.g., storage servers across racks forming a Hadoop cluster)
- Shared infrastructure
  - 3.6 + 0.5 PB global storage, 100Gb Internet connection between sites
- "Graft on heterogeneous features"
  - Infiniband with SR-IOV support, High-mem, NVMe, SSDs, P100 GPUs (total of 22 nodes), RTX GPUs (40 nodes), FPGAs (4 nodes), V100 GPUs (3 nodes)
  - ARM microservers (24) and Atom microservers (8), low-power Xeons (8)

https://www.chameleoncloud.org/

CloudLab

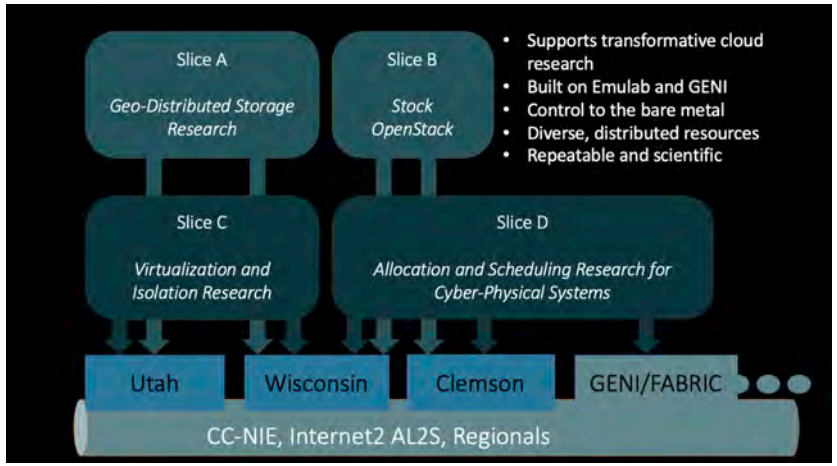

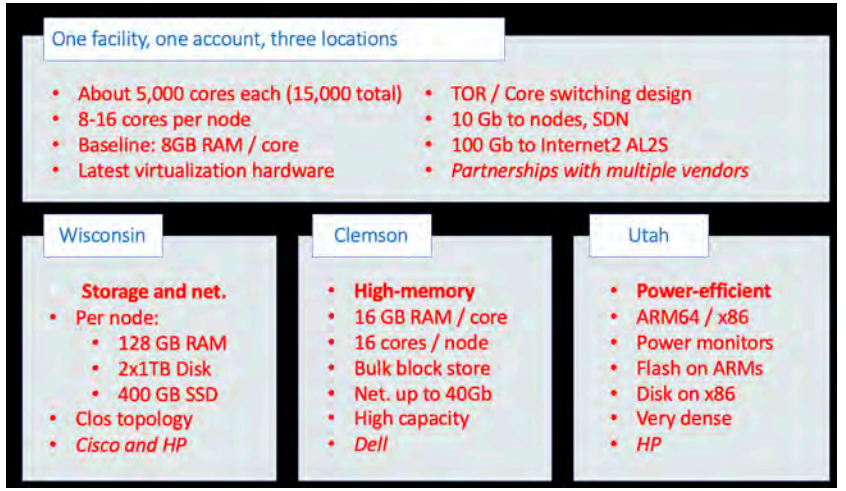

https://cloudlab.us/

PAWR: Platforms for Advanced Wireless Research

Created by the National Science Foundation

PAWR platforms will allow investigation into some of the most important areas of wireless innovation and development (Beyond 5G)

mmWave/THz: enabling R&D and systems testing with millimeter-wave bands including 28GHz and 60GHz with a target of delivering 100 Gbps across multiple city blocks.

Virtualization and Network Slicing: disaggregating network functions and enabling software control and/or dedicated delivery through network slicing.

Microservices Architecture: implementing systems for assembling, controlling, and composing network services.

Artificial Intelligence: using AI-driven systems to automate network optimization.

Spectrum Optimization: experimenting with signal propagation, spectrum interference, and shared frequencies.

Massive MIMO: using fully programmable massive antenna arrays in the 2.5-2.7GHz and 3.5-3.7GHz frequency bands to enable FDD, full duplex research.

Virtualized and Open RAN: examining the RAN CU-DU split and opportunities for open source RAN management.

Resilient and Secure 5G and Beyond Networks: protecting against attack, intrusion, and degradation.

Applications/Services: scaled environments for testing advanced applications for Smart and Connected Community networks that involve Cyber-Physical Systems, Cyber-Security, Internet of Things, Robotics, Smart and Connected Health, and Big Data.

Researchers will use PAWR testbeds to prove out concepts in an array of disciplines

Cybersecurity Testing

Programmable Massive MIMO Arrays

Millimeter Wave Performance Optimization

Open Source Hardware & Software Development

AI-Enabled Network Functions

Accelerated Virtualization of Network Architectures

Dynamic Spectrum Management

Communications Impacts of Unmanned Aerial Vehicles

PAWR platforms were chosen to be geographically diverse and research focus independent

POWDER

Salt Lake City, UT

Software defined networks and massive MIMO

COSMOS

West Harlem, NY

Millimeter wave and backhaul research

AERPAW

Raleigh, NC

Unmanned aerial vehicles and mobility

Rural Broadband Platform

TBD

Coming early.mid 2021

FCC Innovation Zone

Colosseum – *World's largest RF emulator, located at Northeastern University in Boston*

POWDER (https://powderwireless.net)

Partners

A collaboration of the University of Utah, Rice University, the Utah Education and Telehealth Network, and Salt Lake City

Location

Designed to cover 2.3 square miles of the Utah campus, 1.2 square miles of downtown Salt Lake City, and a 2-mile corridor in between, reaching a population of 40,000

Platform Resources

Includes software defined rooftop and base station radios; access to compute, storage, and cloud resources; a control framework for remote access and control

Research Focus

End-to-end programmable networking, spectrum management, massive MIMO

COSMOS (https://www.cosmos-lab.org)

Partners

Partnership of Rutgers University, Columbia University, NYU, and New York City, with support from The City College of New York, the University of Arizona, Silicon Harlem, and IBM

Location

1 square mile in West Harlem

Platform Resources

Includes large and medium-sized software defined base station radios; fiber connections to and from radio sites and to local data centers; interaction with other smart community and innovation initiatives

Research Focus

Millimeter wave radio communications and dynamic optical switching technologies

AERPAW (https://aerpaw.org)

Partners

North Carolina State University and the city of Raleigh and town of Cary, with support from Wireless Research Center of North Carolina, Mississippi State University, and Renaissance Computing Institute (RENCI) at the University of North Carolina at Chapel Hill

Location

Town of Cary, Lake Wheeler Field Laboratory, Centennial Campus in Raleigh

Platform Resources

Fixed software defined radio nodes; unmanned aerial and ground vehicles carrying AERPAW nodes; rural and urban environments; remote access control framework

Research Focus

Advanced wireless services to support unmanned aerial systems and advanced UAS to enable new wireless services

Rural Broadband Platform

The final PAWR RFP focuses on applying advanced wireless technologies in rural and other low-density geographic areas in an efficient and affordable way. Proposers are asked to create a testbed for experimenting with advanced wireless technologies and network architectures – combined with existing technologies – that may transform the existing broadband deployment cost curve through innovations in technologies and engineering processes.

Colosseum

Colosseum is the world's largest wireless network emulator with granularity at the RF signal level

- 256 x 256 100 MHz RF channel emulation
- 128 Programmable Radio Nodes
- Computing resources (CPU, GPU, FPGA)
- Access control and scheduling infrastructure
- Supports remote shared access
- Colosseum supports a containerized software environment with full-stack, end-to-end experiments

northeastern.edu/colosseum

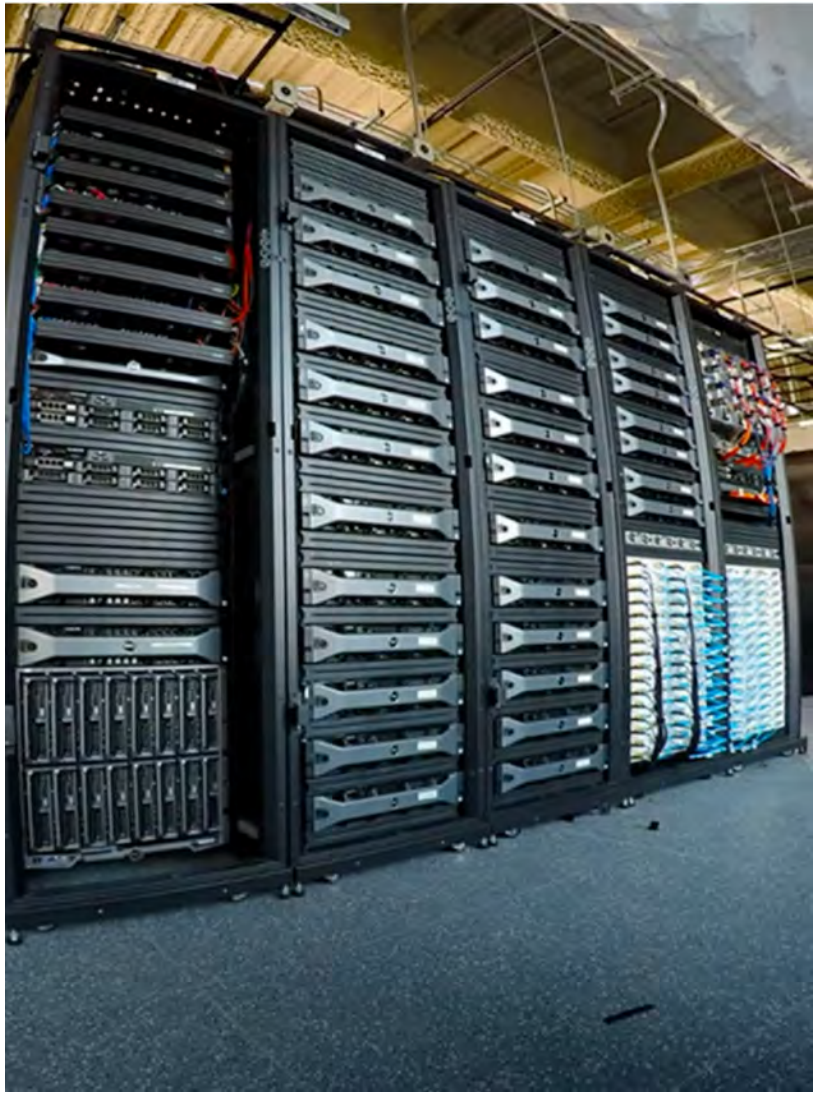

CCRI Example: Infrastructure for Mobile Underwater Wireless Networking Research

Community-shared, open-source, open-architecture infrastructure, with a goal is to establish an ecosystem that can support research/education agendas across CISE communities including:

- Miniaturized aquatic robots for lab tests==➔ indoor testbed
- Field-deployable autonomous surface vehicles (ASVs) and autonomous underwater vehicles (AUVs) for field deployments==➔ lake testbed
- Re-configurable software suites, µNet software
- Various resources & services to support the usage from the community users

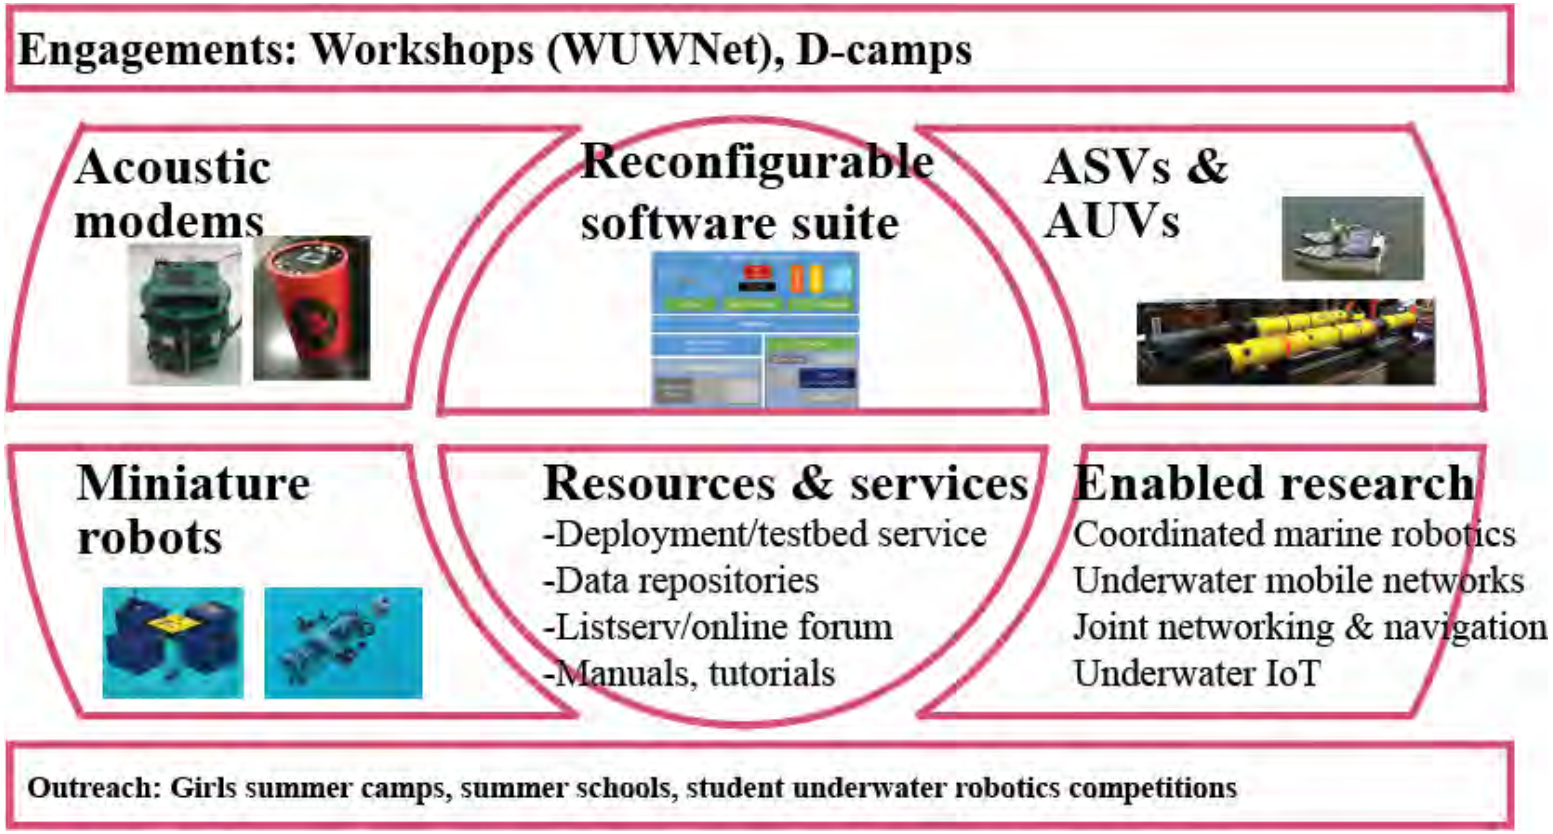

http://gtsr.gatech.edu/munet.html

Indoor testbed for swarming/networking

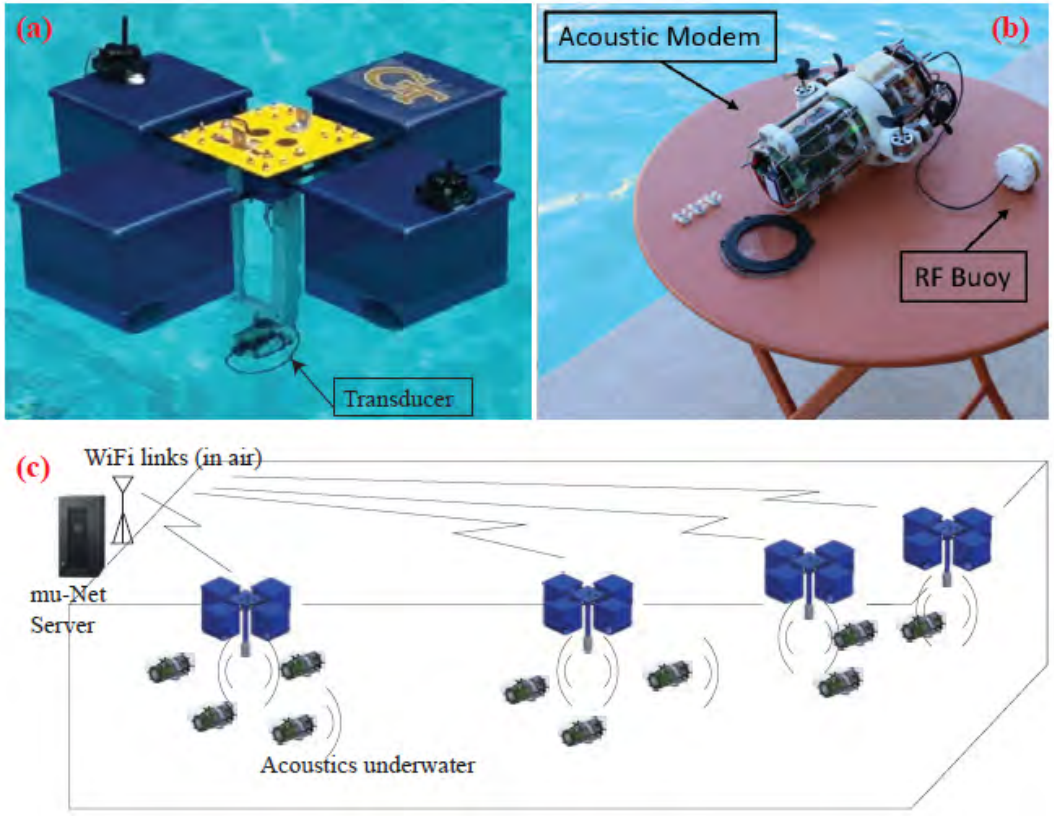

- Main assets (miniaturized robots): GT-OSV (a) + GT-MUR (b) + acoustic modems
- Deployment setup in a swimming pool (c)
- Goal: Remote automatic access, aqua-robotarium.

Lake testbed for underwater communications/networking

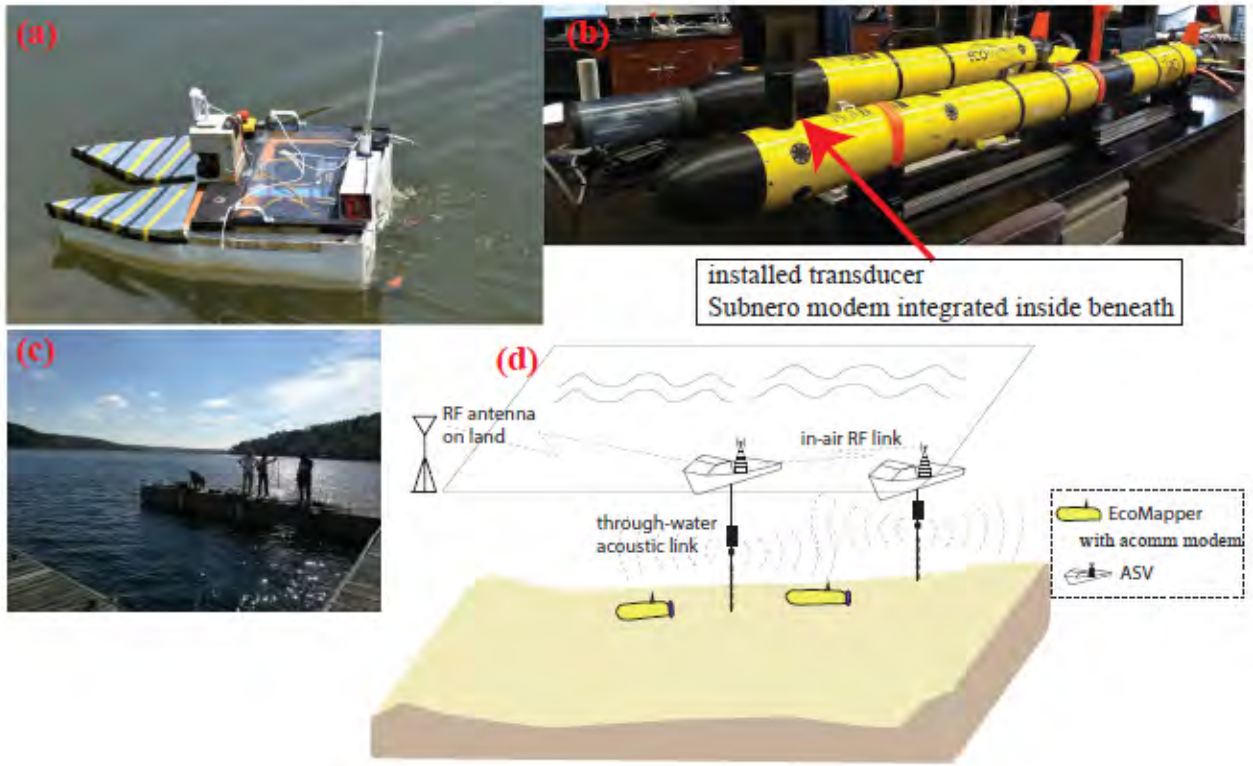

- Main assets: Autonomous surface vehicle (ASV) (a) + EcoMapper-AUVs (b)
- Field deployment site (c): 20-m depth (Lake of Tuscaloosa)
- Deployment setup in the lake (d)
- Goal: Manned service

End-to-End *Federated* Research Opportunities

Thank you!

Declaration: Many individuals contributed to the content of presentation