

# Panel

# Shaping the Future of Distributed Computing: Bottom up, or Top down

---

June 2, 2022 Hamburg, Germany

**ExaComm 2022**

Seventh International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale

Barney Maccabe

Executive Director, Institute for the Future of Data and Computation

Professor, School of Information

University of Arizona



# Institute for the Future of Data and Computation

Cultivating the information capabilities and trust needed to address grand challenges

## New UArizona institute focuses on making sense of data

Led by a University of Arizona alumnus, the Institute for the Future of Data and Computing also will explore the societal questions driving how and why data is used.

Research, Innovation & Impact

May 4, 2022

[Share](#) [Tweet](#) [in Share](#) [Email](#) [Print](#)



Maccabe@Arizona.edu



**Infrastructure:** Deploying the information resources needed for research



**Workforce:** Growing peer communities and industrial partnerships



**Sustainability:** Developing business models needed to sustain infrastructure and artifacts



**Informed:** Integrating cognitive and social sciences to build ethically informed systems



**Foundations:** Creating virtuous cycles supporting use-inspired foundational research

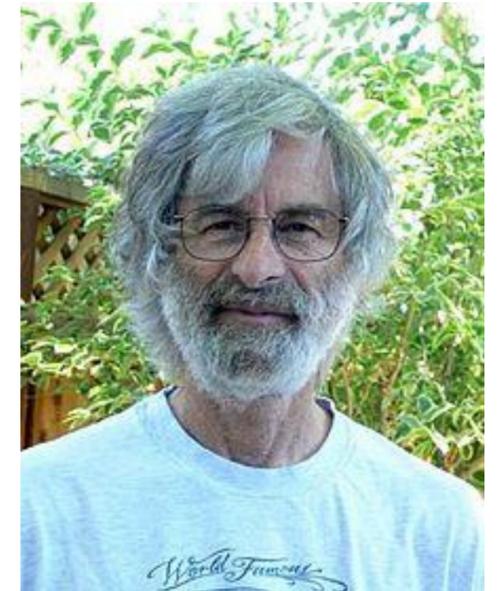


THE UNIVERSITY OF ARIZONA  
Research, Innovation  
& Impact

# Distributed Systems

---

“A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable.”



Leslie Lamport, 1986

# What is a “Distributed” Data Processing System?

Philip H Enslow, Jr., *IEEE Computer* (Volume: 11, Issue: 1, January 1978)

## Definition (essential characteristics)

- A *multiplicity* of general-purpose resource components, including both physical and logical resources, that can be assigned to specific tasks on a dynamic basis. Homogeneity of physical resources is not essential.
- A *physical distribution* of these physical and logical components of the system interacting through a common network.
- A *high-level operating system* that unifies and integrates the control of the distributed components. Individual processors each have their own local operating system, and these may be unique.
- *System transparency*, permitting services to be requested by name only. The server does not have to be identified.
- *Cooperative autonomy*, characterizing the operation and interaction of both physical and logical resources.

## Benefits

- High system performance, fast response, high throughput
- High availability
- Reduced network costs
- Graceful degradation (fail-soft capability)
- Ease of modular, incremental growth and configuration flexibility
- Resource sharing
- Automatic load sharing
- High adaptability to changes in world load
- Incremental replacement and/or upgrading of components (both hardware and software)
- Easy expansion in both capacity and function
- Easy adaptation to new functions
- Good response to temporary overloads

Critical issue: the need to span *multiple administrative domains*

# Classic Bottom-up Approach

- Build the system by assembling from available components
- Potential for heterogeneity in hardware and local operating systems
- Incremental growth, incremental upgrades

“Federated Instruments”

# An Alternate Conceptualization (Top-down)

- Single systems are becoming too large to fit into a single facility
  - Can't continue growth in power, space, and cooling
  - Instruments generate large amounts of data, requiring local computation
- Potential for campus-wide interconnects
- HPC gets pushed to scientific instruments or instruments/data pulls HPC

Nodes/Cabinets of an HPC system are pushed and pulled to the “edge”

# Déjà vu, All Over Again

- Parallel programming was generally expected to be a small number of tasks
  - Each task had a specific purpose
  - Coordination (synchronization) of tasks was critical
- Massively parallel processing, and the emphasis on **scale**, changed our perspective
  - SPMD: Single Program Multiple Data
  - BSP: Block Synchronous Programming (“naturally synchronized” by boundary exchanges)
- Aside: will Extreme Heterogeneity lead to a “reversion to the mean”?

# The Panel

---

# Charge questions for the panel

- Which approach, bottom-up or top-down, is most likely to shape the future of distributed computing?
- To what extent will (should) economic considerations drive the future of distributed computing?
- To what extent does the notion of technology refresh rate factor into how distributed systems should be built?
- To what degree will developments in interconnect technologies, possibly including wireless technologies, be important in shaping this future?
- To what degree will the required software infrastructure be critical in shaping this future?

# Panelists



- Sadaf Alam, CSCS (Swiss National Supercomputing Center)
- Debbie Bard, Lawrence Berkeley National Laboratory
- Norbert Eicker, Jülich Supercomputer Center
- Rich Graham, Nvidia/Mellanox

