Engineering Research Center for
Collaborative Adaptive Sensing of the Atmosphere

Globally Distributed Datacenters: A Collaborative Peer-to-Peer Approach

H. M. N. Dilum Bandara and Anura P. Jayasumana
Electrical and Computer Engineering
Colorado State University
dilumb@engr.colostate.edu

FRHPC 2012

University of Massachusetts Amherst
University of Oklahoma
Colorado State University
University of Puerto Rico Mayaguez

CASA is primarily supported by the Engineering Research Centers Program of the National Science Foundation under NSF award number 0313747.

HPC in the Loop

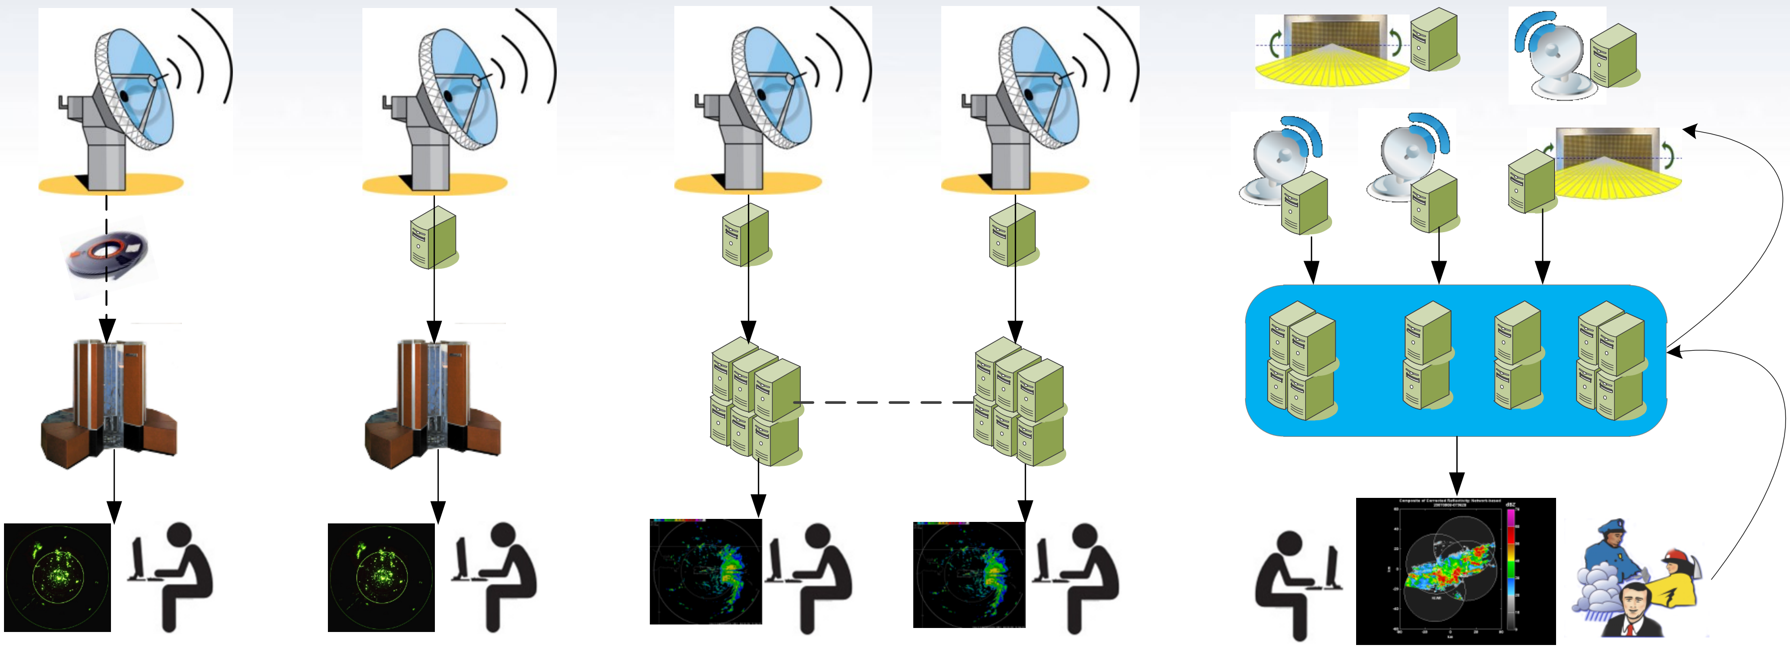

- HPC resources coupled with expensive sensors in real time
- Dynamically grouped sensors & HPC resources adapted in response to changing weather & user needs

Collaborative P2P Systems

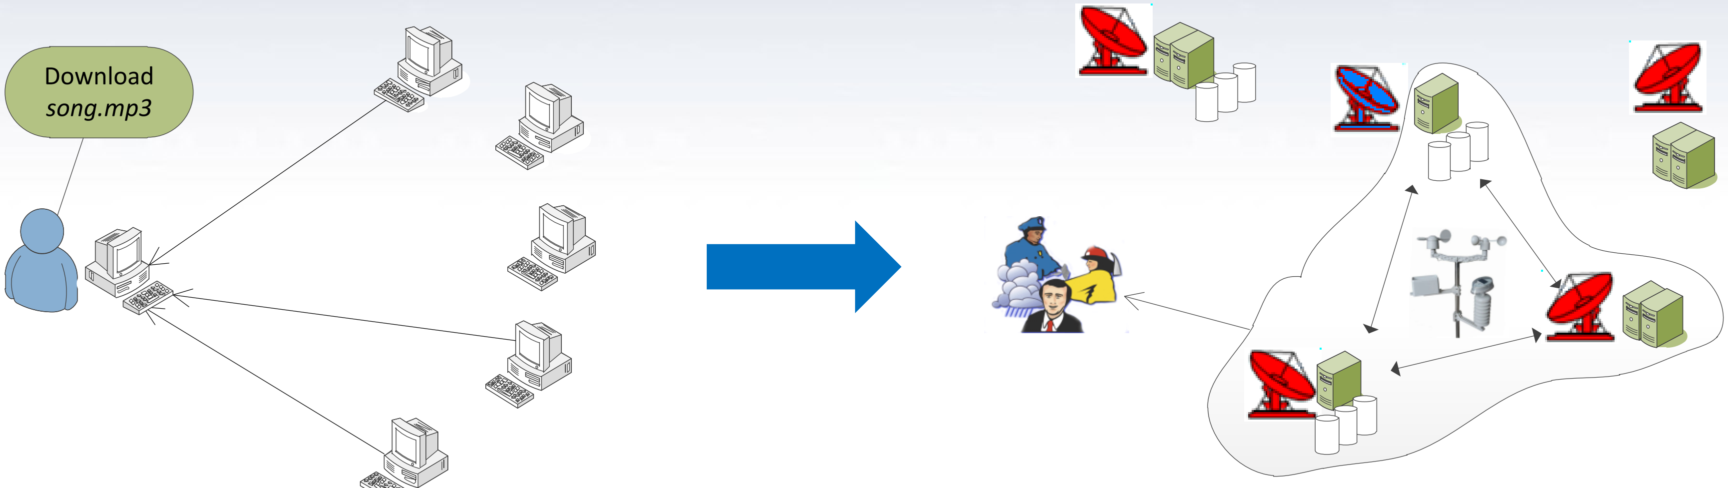

Sharing

Collaboration

- Advances in Web 2.0, ubiquitous high-speed networks, cloud computing, & strong social networks
- P2P systems will play an even greater role in distributed resource collaboration & collaborative applications
- Diverse peers bring in unique resources & capabilities to a virtual community to accomplish something big

Collaborative Adaptive Sensing of the Atmosphere (CASA)

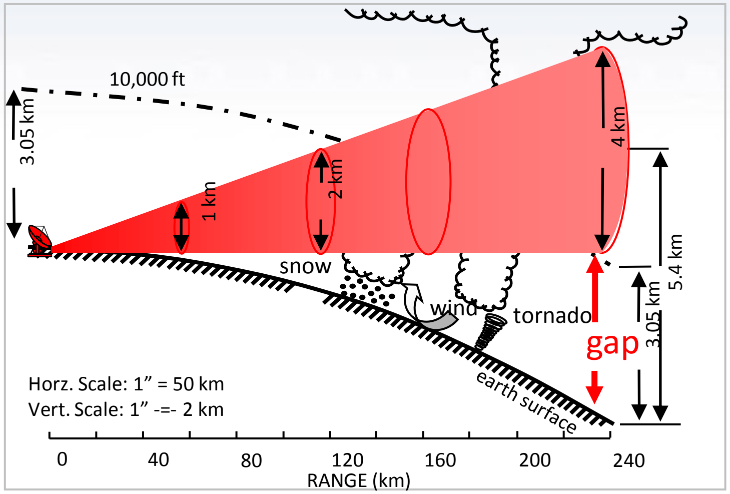

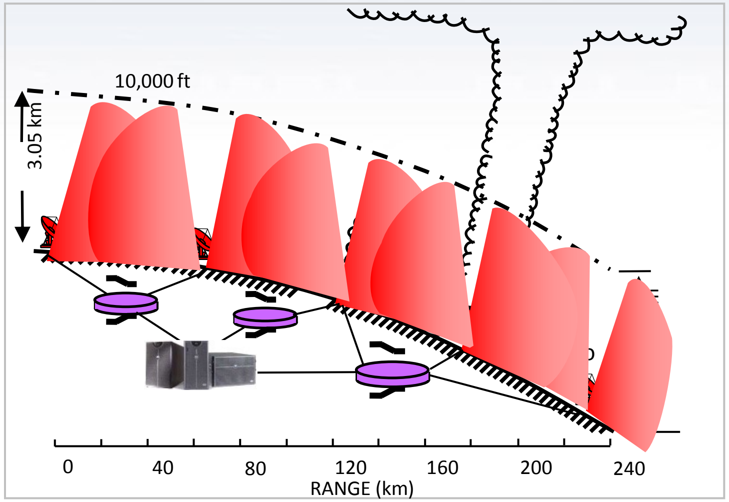

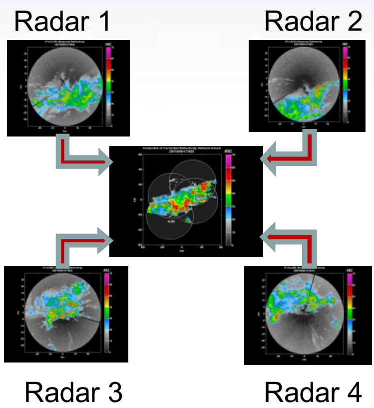

- CASA aggregates groups of resources as and when needed
  - Dedicated & reliable resources
  - Real-time, multi-attribute, heterogeneous, dynamic, & distributed
- Community weather monitoring
  - Voluntary & unreliable resources

Global Environment for Network Innovations (GENI)

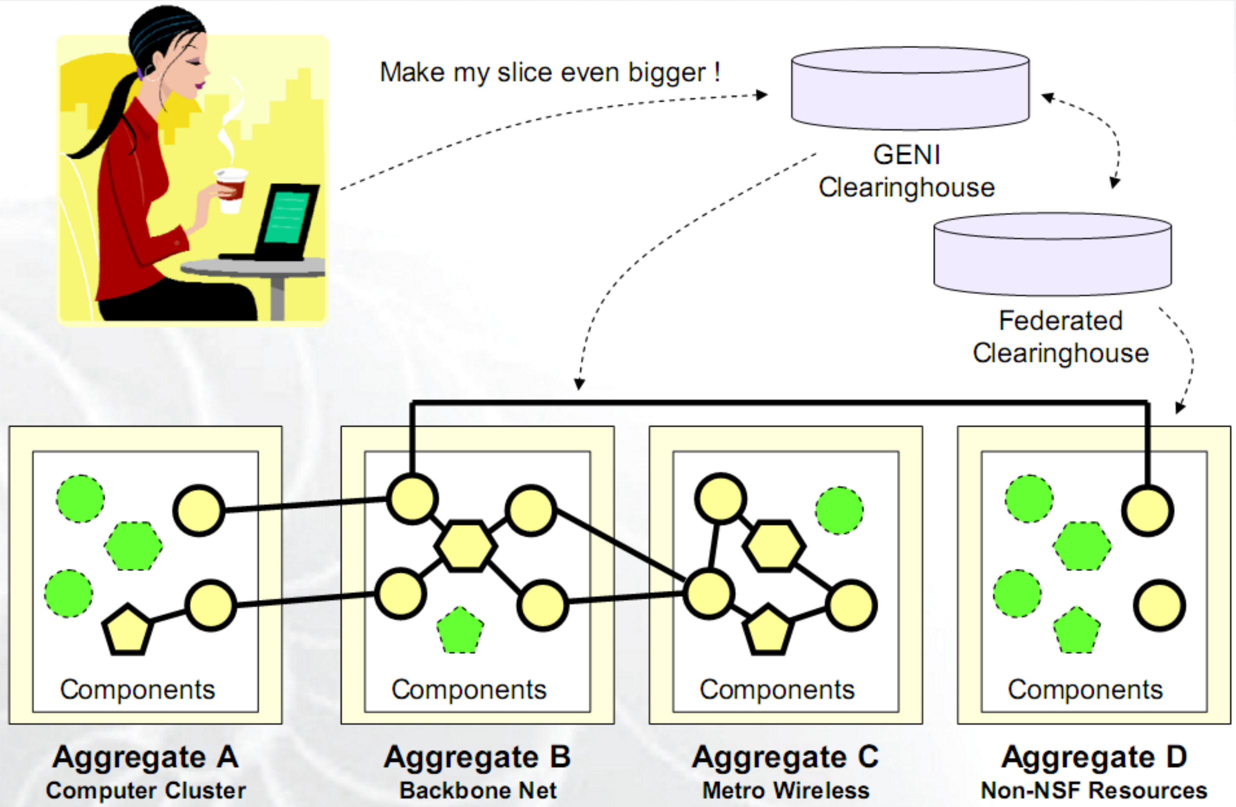

- Sensors
  - Cameras
  - Sensors mounted on busses
  - Micro weather stations
  - Radars
- Processing & storage
  - Amazon EC2
  - Amazon S3
- Networks
  - Internet2
  - Emulab
  - BEN dark fibers

- Collaborative & exploratory platform for innovation
- Aggregating groups of resources across multiple administrative domains

casa

Community (P2P) Cloud Computing

- Resource aggregation within datacenters
  - Data intensive cloud computing
  - Storage, GPUs, FPGAs
  - Encryption, business logic, scientific algorithms
  - Virtual networks in/out & within cloud
    - Sensors can't be inside a datacenter
- Community as a datacenter
  - Resourceful peers, home servers
  - Users govern themselves & hold data
  - Aggregation of bandwidth at edge
  - Ability to scale in/out
  - Monetary & non-monetary benefits

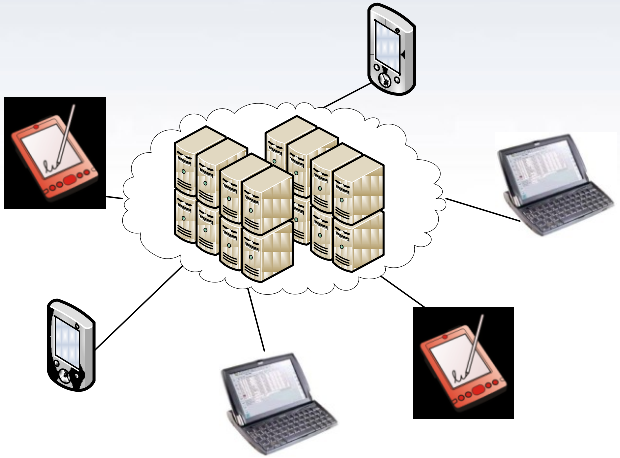

Community (P2P) Cloud Computing

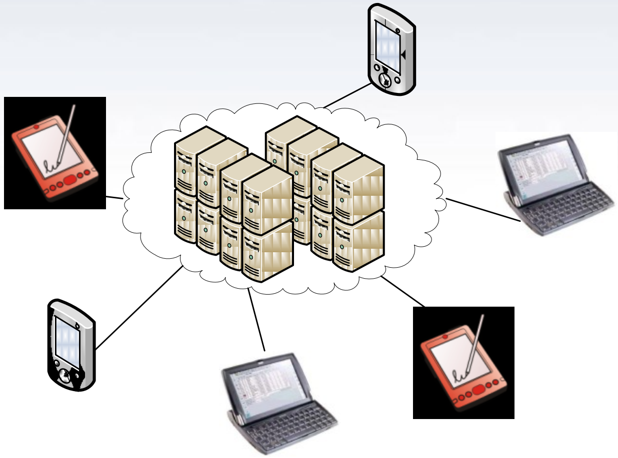

- Resource aggregation within datacenters
  – Data intensive cloud computing
  – Storage, GPUs, FPGAs
  – Encryption, business logic, scientific algorithms
  – Virtual networks in/out & within cloud
    - Sensors can't be inside a datacenter

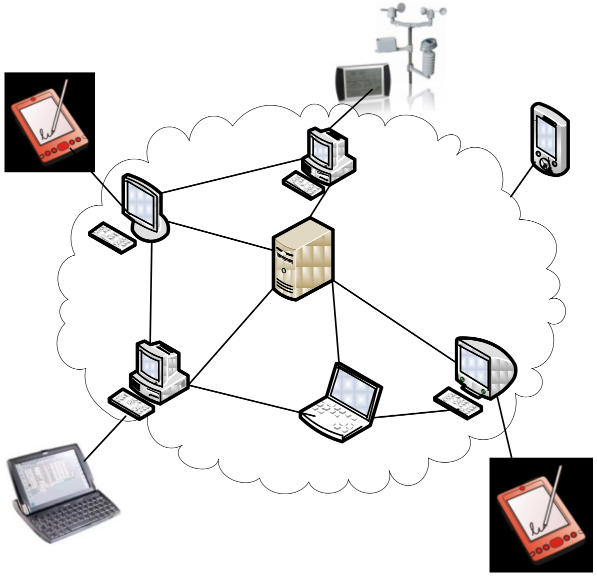

- Community as a datacenter
  – Resourceful peers, home servers
  – Users govern themselves & hold data
  – Aggregation of bandwidth at edge
  – Ability to scale in/out
  – Monetary & non-monetary benefits

Phases of Collaborative P2P Systems

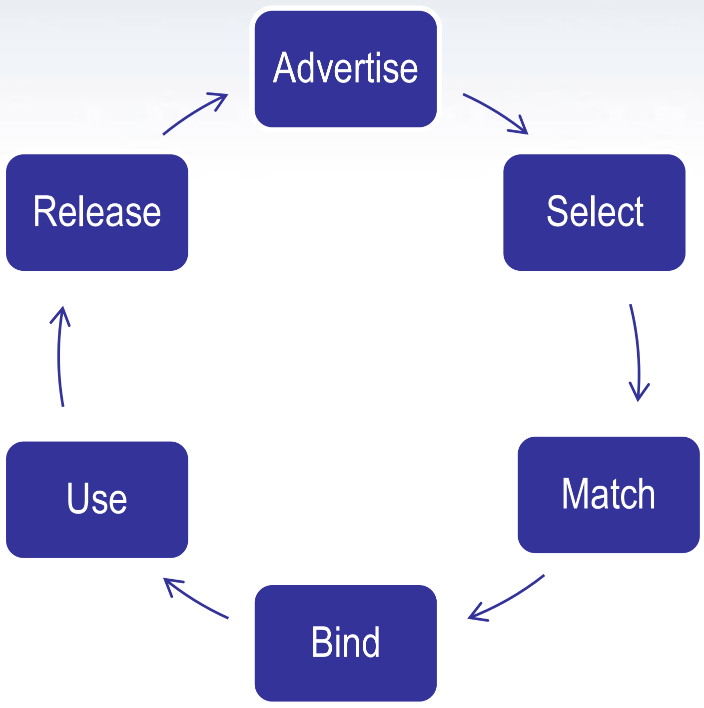

$$
\begin{aligned}
CE = \{&CPUSpeed = 2.4\text{ GHz}, CPUFree = 69\%,\\
&Memory = 4\text{ GB}, Archi = \times 86, OS = \text{“}Linux\_2.6\text{”},\\
&Available = [10\text{pm}, 5\text{am}], Useby = \text{“}Friends\text{”}\}
\end{aligned}
$$

$$
\begin{aligned}
q = \{&\\
&\text{Group } A\\
&\quad CE = \{6, CPUSpeed \in [2.0\text{ GHz}, \text{MAX}],\\
&\qquad DiskFree \in [20\ GB, \text{MAX}]\\
&\qquad Latency \in [0, 50\text{ ms}]\}\\
&\text{Group } B\\
&\quad radars = \{3, Type = \text{“}DualDoppler\text{”},\\
&\qquad Range \in [20\text{ km}, 50\text{ km}]\\
&\qquad Location \in [40^{o}, 102^{o}, 42.5^{o}, 103.7^{o}]\\
&\qquad Bandwidth \in [2\text{ Mbps}, \text{MAX}]\\
&\text{Intergroup}\\
&\quad Latency(A, B) \in [0, 100\ ms]\\
&\quad Bandwidth(A, B) \in [2.5\ Mbps, 10\ Mbps]\\
\}&
\end{aligned}
$$

- Need to support by resource discovery systems, job schedulers, etc.
- Essential for high performance, low latency, & QoS
- Some phases may be combined or skipped

Current Solution Space

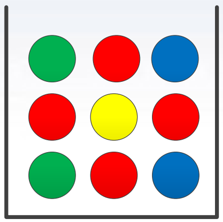

Centralized

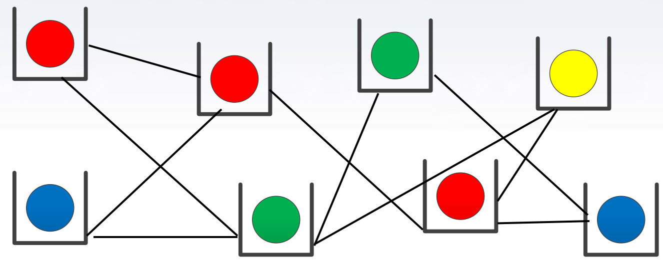

Unstructured P2P

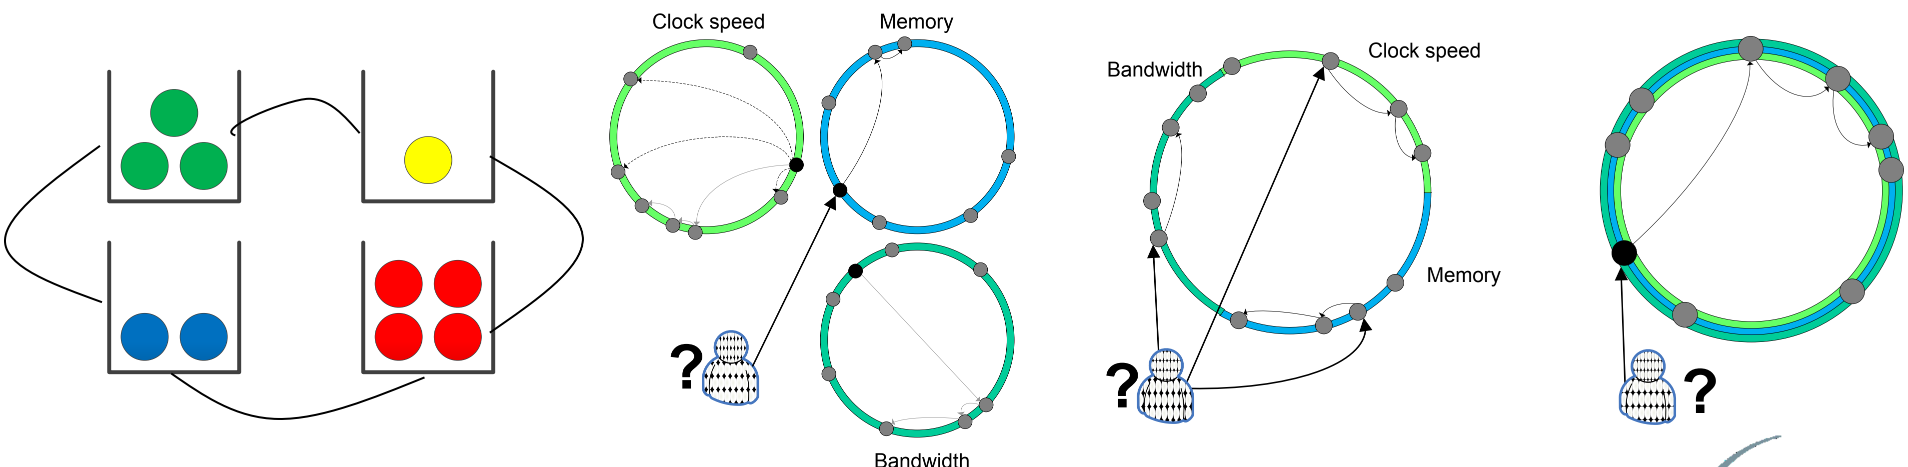

Structured P2P – Distributed Hash Table (DHT)

Resource Characteristics

- Data from 4 systems
  - PlanetLab
  - SETI@home
  - EGI grid
  - CSU

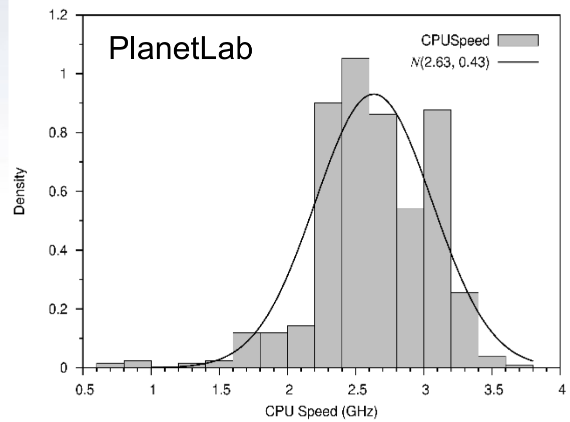

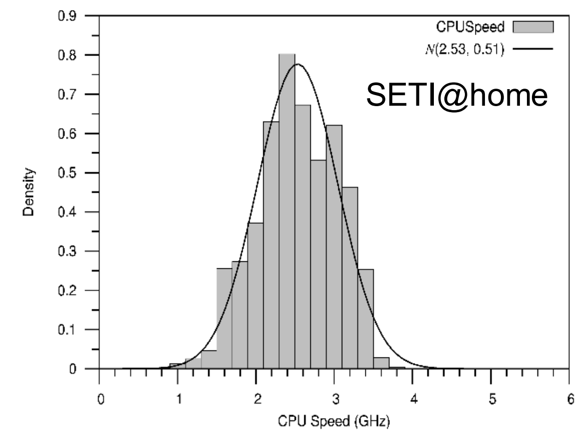

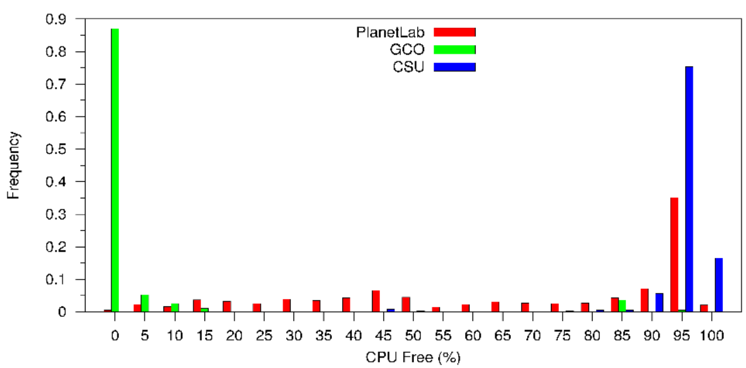

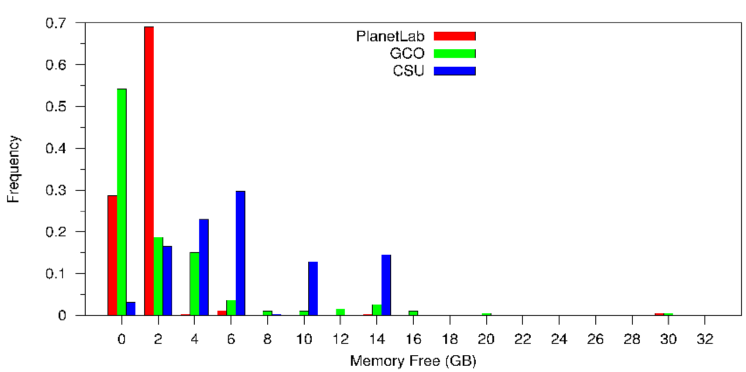

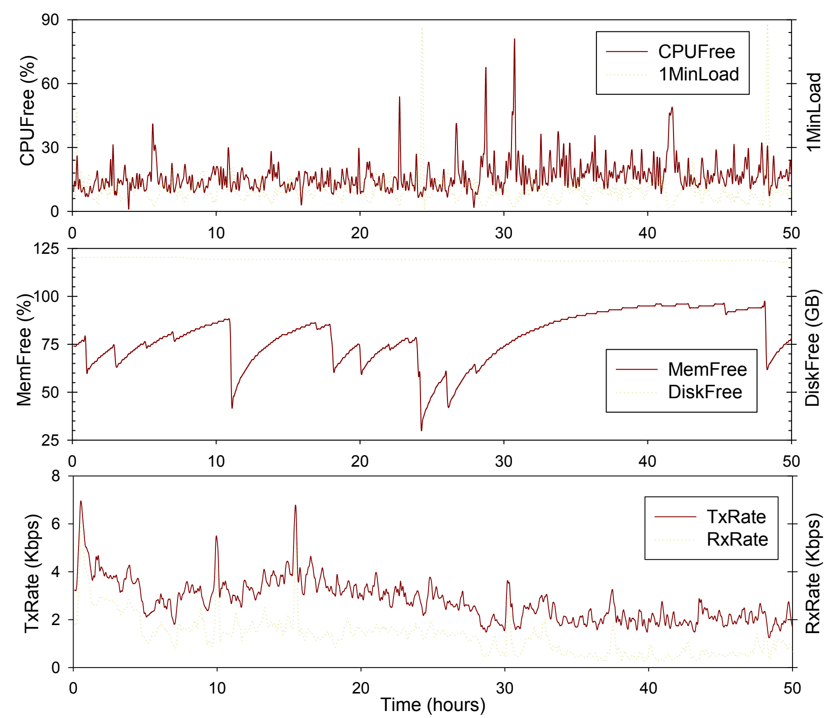

Resource Characteristics (cont.)

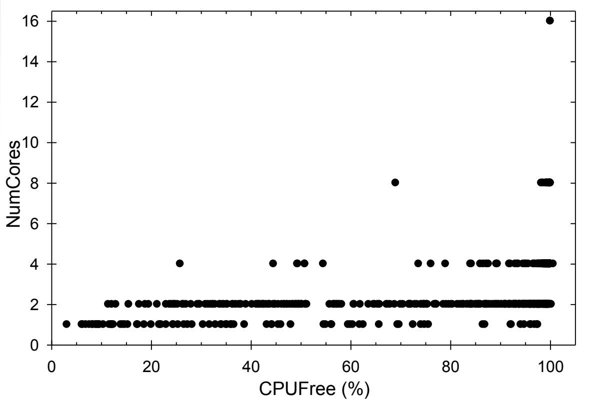

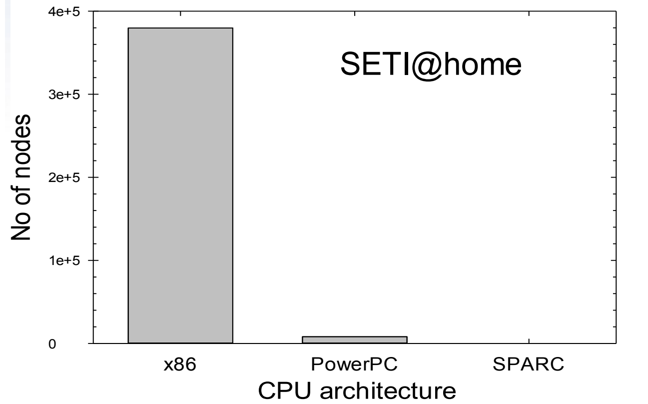

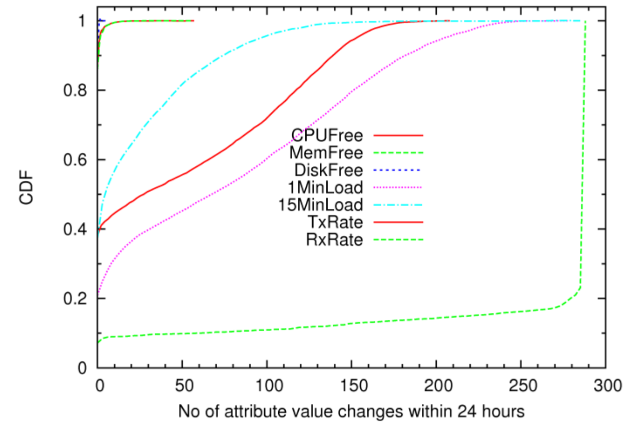

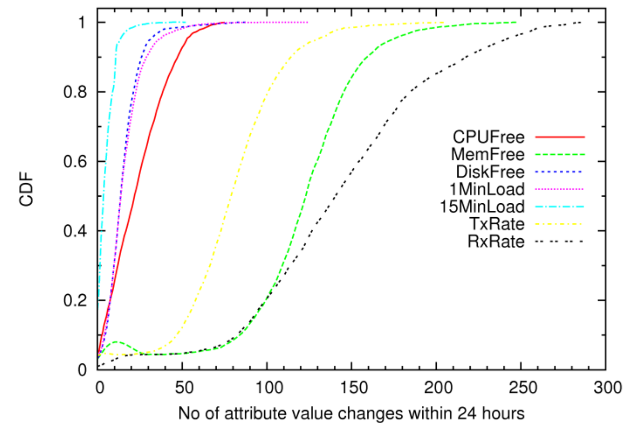

- Attribute values are skewed & have different marginal distributions
  - Most of them don't fit a known distribution
- Complex correlation patterns
- Few attribute values
- Dynamic attributes change at different rates

Query Characteristics

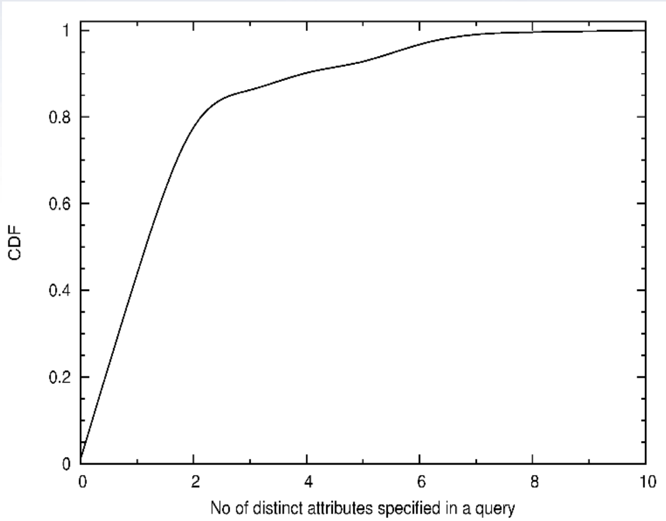

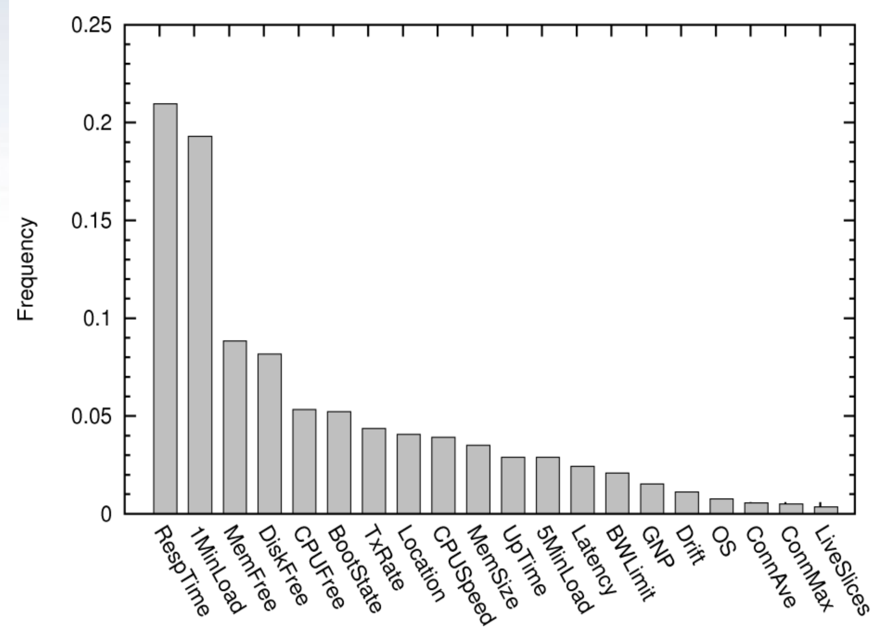

- Less specific queries
  - Few attributes
  - Large ranges of attribute values
- Dynamic attributes are more popular
- Skewed attributes

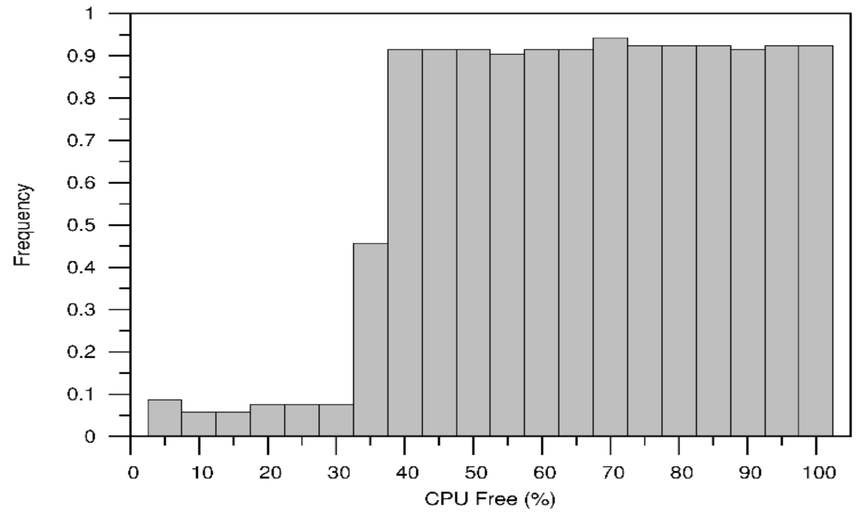

Performance Under Real Workloads

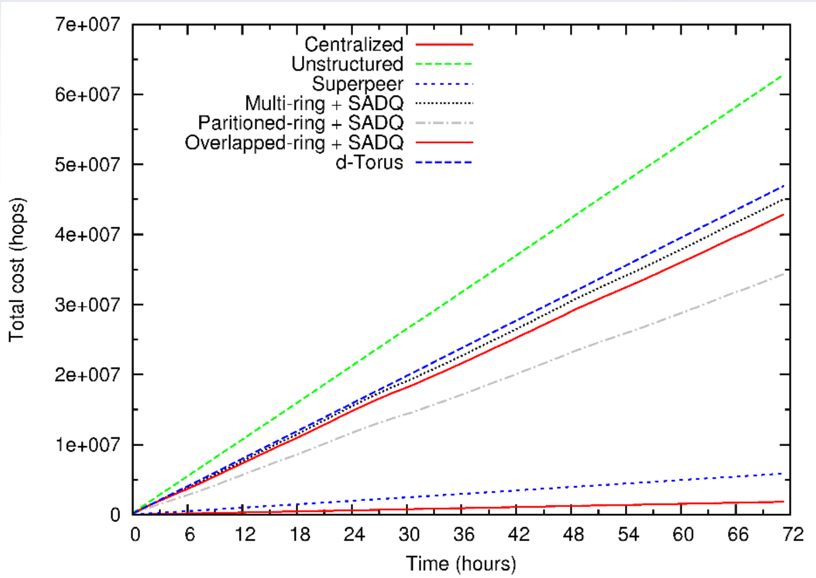

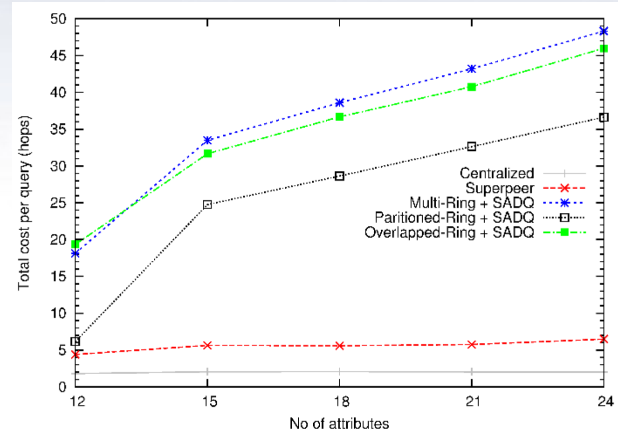

n	Multi-ring + SADQ			Partitioned-ring + SADQ			Overlapped-ring + SADQ		
	Min	Ave	Max	Min	Ave	Max	Min	Ave	Max
250	0	9.2	239.1	0	3.7	19.4	0	9.1	238.4
527	0	13.7	509.0	0	4.6	27.6	0	13.5	506.0
750	0	16.2	719.1	0	4.9	36.6	0	16.5	719.9
1000	0	19.8	975.5	0	5.3	45.3	0	20.4	963.8

- Lowest cost – Centralized
- Highest cost – Unstructured
- Cost increase with no of attributes & nodes
  - Cost of ring-based designs $O(n)$
- Unbalanced query & index load

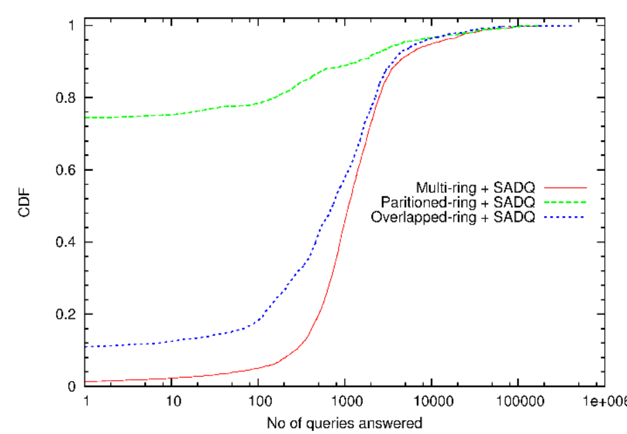

Resource & Query Aware Resource Discovery

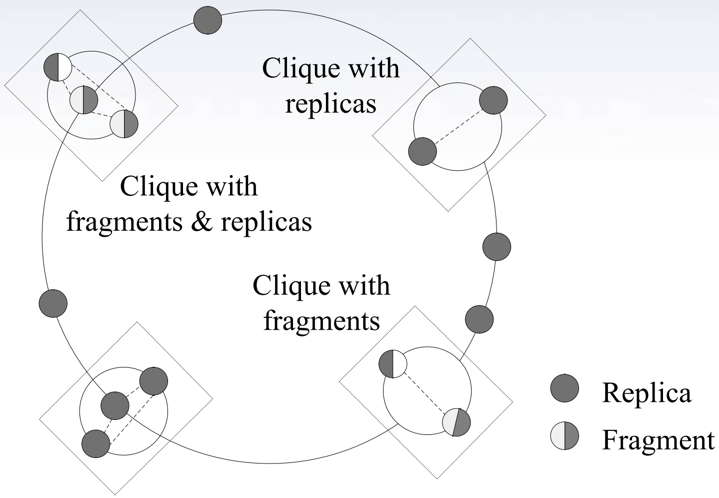

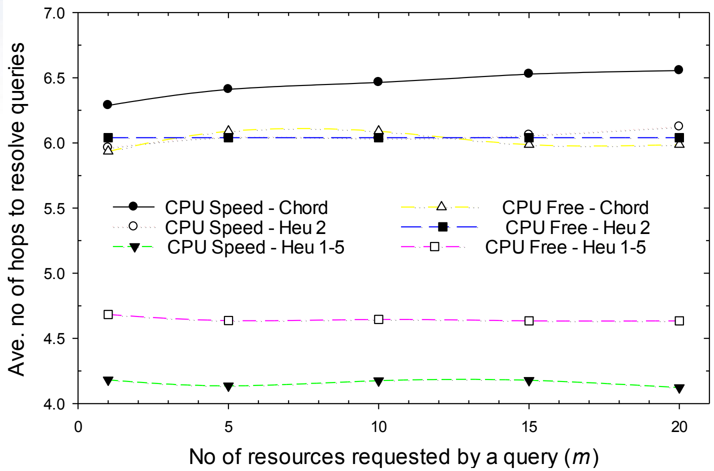

- No of nodes along ring ≤ no of distinct attribute values
- Place fragments & replicas orthogonal to ring
  - Replicas - skewed query load
  - Fragments – skewed/identical resources

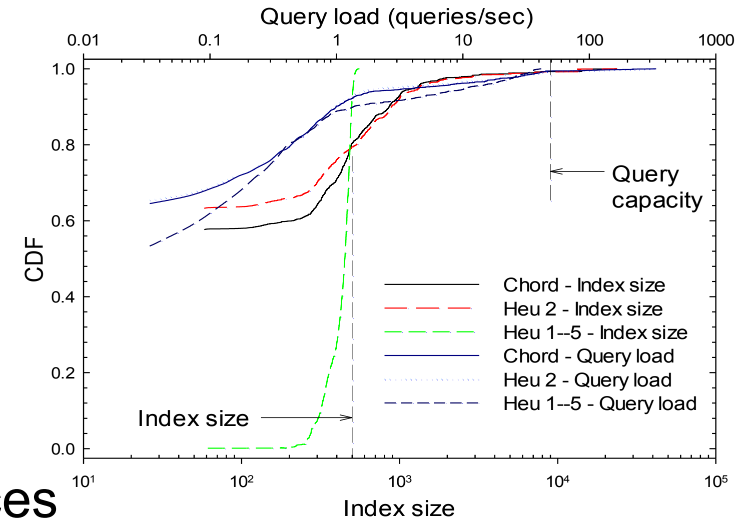

Conclusions

- Emerging collaborative P2P systems need to integrate sensing and computing resources in real time
- Novel solutions are needed to aggregate group(s) of resources as and when needed
  - Multi-attribute, heterogeneous, dynamic, and distributed resources
  - Massive no of possible resources & ways to group them
  - Support all key phases of collaborative P2P systems
  - Develop incentives, trust, privacy, & security solutions
- Collective power of P2P communities & their resources → Globally distributed virtual datacenters for many collaborative applications

casa

Related Publications

- H. M. N. D. Bandara and A. P. Jayasumana, “Collaborative applications over peer-to-peer systems – Challenges and solutions,” Peer-to-Peer Networking and Applications, Springer, 2012.
- H. M. N. D. Bandara and A. P. Jayasumana, “Resource and query aware, peer-to-peer-based multi-attribute resource discovery,” In Proc. 37th IEEE Conf. on Local Computer Networks (LCN ‘12), Oct. 2012, To appear.
- P. Lee, A. P. Jayasumana, H. M. N. D. Bandara, S. Lim, and V. Chandrasekar, “A peer-to-peer collaboration framework for multi-sensor data fusion,” Journal of Network and Computer Applications, vol. 35, no. 2, May 2012, pp. 1052-1066.
- H. M. N. D. Bandara and A. P. Jayasumana, “Evaluation of P2P resource discovery architectures using real-life multi-attribute resource and query characteristics,” In Proc. IEEE Consumer Communications and Networking Conf. (CCNC ‘12), Jan. 2012.
- H. M. N. D. Bandara and A. P. Jayasumana, “On characteristics and modeling of P2P resources with correlated static and dynamic Attributes,” In Proc. IEEE GLOBECOM ‘11, Dec. 2011.
- H. M. N. D. Bandara and A. P. Jayasumana, “Characteristics of multi-attribute resources/queries and implications on P2P resource discovery,” In Proc. 9th ACS/IEEE Int. Conf. On Computer Systems And Applications (AICCSA ‘11), Dec. 2011.
- H. M. N. D. Bandara, A. P. Jayasumana, and M. Zink, “Radar networking in collaborative adaptive sensing of atmosphere: State of the art and research challenges,” Under review.

casa

Questions/Comments

dilumb@engr.colostate.edu
www.cnrl.colostate.edu

casa