Engineering Research Center for
Collaborative Adaptive Sensing of the Atmosphere

Evaluation of P2P Resource Discovery Architectures Using Real-Life Multi-Attribute Resource and Query Characteristics

H. M. N. Dilum Bandara and Anura P. Jayasumana

Electrical and Computer Engineering,

Colorado State University, Fort Collins, CO.

dilumb@engr.colostate.edu


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.

Motivation

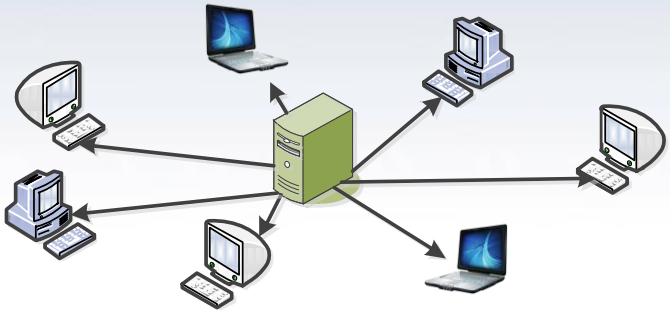

Desktop grids

- Throughput
- Static attributes
  - CPU speed, architecture, GPUs

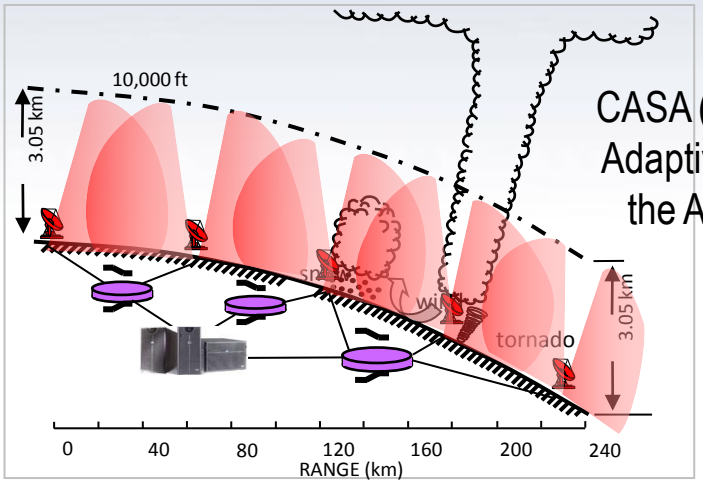

CASA (Collaborative Adaptive Sensing of the Atmosphere)

Community (P2P) clouds

- QoE & QoS
- Dynamic attributes
  - Free CPU, bandwidth

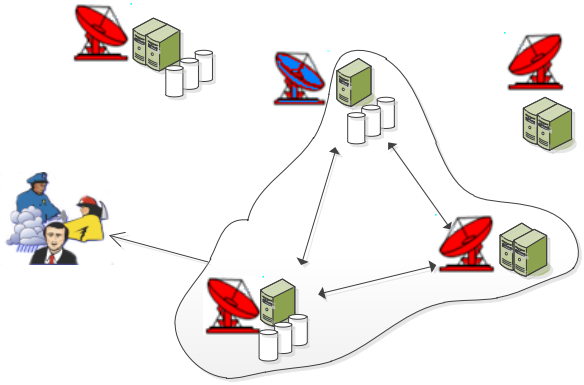

Radar networks

- QoS & latency
- Static & dynamic attributes
  - CPU speed, free CPU, bandwidth

Contribution

Evaluate P2P multi-attribute resource discovery solutions using real-world resources & queries

- Collaborative P2P applications need to discover resources
  - Characteristics & models of resources\queries are essential in design, validation, & performance analysis
  - Existing solutions have relied on many simplifying assumptions
  - Characteristics of real-world resources & queries are different from those assumptions
- Fundamental design choices needs to be reevaluated using behavior learned from actual systems
  - Understand applicability, best practices, what not to do

Objectives

- Understand characteristics & model multi-attribute resources
  - Static attributes [Heien, 2011]
  - Static & dynamic attributes, & queries [Bandara, 2011a]
  - Model resources & synthetic resource generation [Bandara, 2011b]
  - Existing performance studies assume
    - i.i.d attributes, uniform/Zipf's distribution of attributes, queries specify small range of attribute values, ignored dynamic attributes, replication of small datasets, etc.
    - Not valid
- Evaluate fundamental design choices for P2P multi-attribute resource discovery
  - Cost of advertising/querying resources, index size, load balancing

- E. Heien et al., "Correlated resource models of Internet end hosts," ICDCS '11, June 2011.
- H.M.N.D. Bandara & A.P. Jayasumana, "Characteristics of multi-attribute resources/queries and implications on P2P resource discovery," AICCSA '11, Dec. 2011.
- H.M.N.D. Bandara & A.P. Jayasumana, " On Characteristics and Modeling of P2P Resources with Correlated Static and Dynamic Attributes," IEEE GLOBECOM '11, Dec. 2011.

Dataset

- PlanetLab node data
  - Global research network for developing new network services, protocols, & applications
  - Reflects many characteristics of Internet-based distributed systems
    - Heterogeneity, multiple end users, dynamic nodes, & global presence
    - Used to evaluate many preliminary P2P protocols & applications
- Rich dataset
  - Resources from CoMon (comon.cs.princeton.edu)
    - 12 static & 34 dynamic attributes
    - 5 min samples
    - 500-700 active nodes
    - Collected between Nov 1 to 15, 2010
  - Queries from SWORD (sword.cs.williams.edu)
    - 441 queries collected over a year

Multi-Attribute Resources & Queries

- Resources
  - *CPUSpeed = 2.4 GHz and CPUFree = 54% and MemorySize = 4096 MB and OS="Linux 2.6"*
  - Specify
    - Set of attributes & their attribute values
    - Static & dynamic attributes
    - Advertised every time an attribute value (significantly) change
- Queries
  - *Find 2 nodes with CPUSpeed ≥ 2.0 GHz and 256 ≤ MemorySize ≤ 512 MB and OS="Linux 2.6"*
  - Specify
    - Required number of resources
    - Multiple attributes
    - Range of attribute values

Resource Characteristics – Distributions

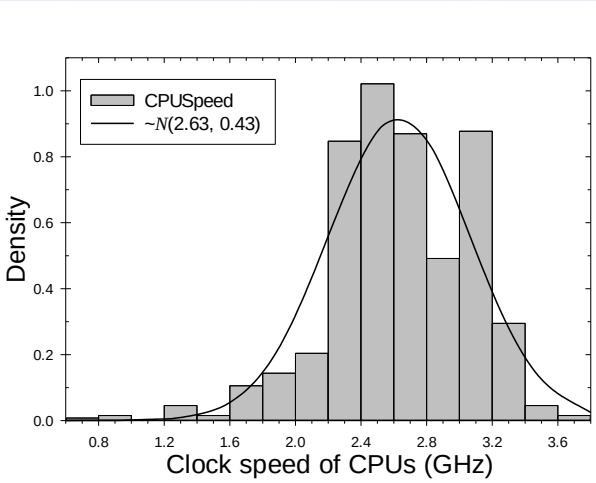

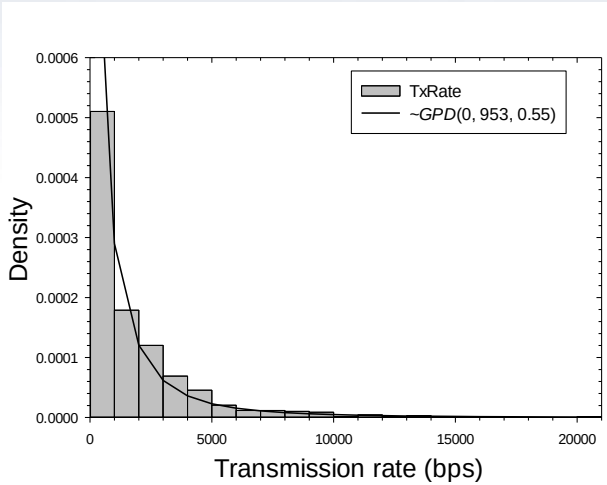

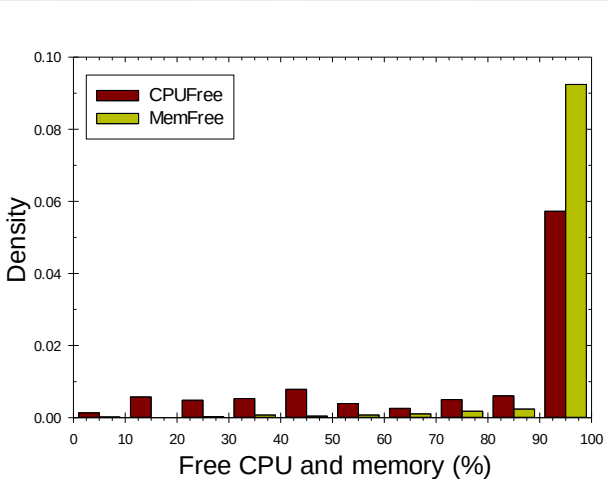

- Resources satisfy a mixture of probability distributions
  - Gaussian – CPUSpeed, MemSize, DiskFree
  - Pareto – TxRate, RxRate
  - Many identical nodes
- Highly skewed distributions
  - CPUFree, MemFree, CPU architecture

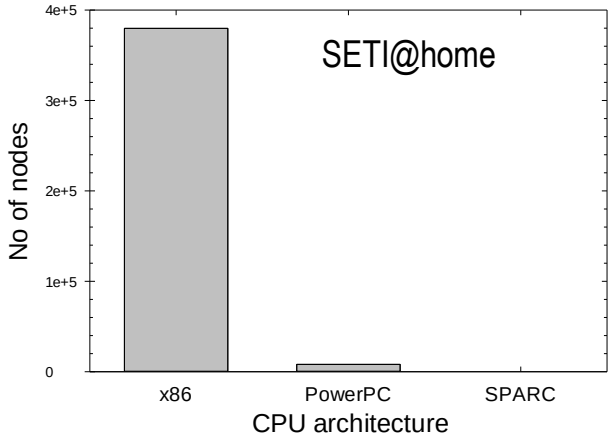

Resource Characteristics – Correlation

Spearman's ranked correlation coefficient ρ

	CPUSpeed	NumCores	CPUFree	1MinLoad	MemSize	MemFree	DiskFree	TxRate
NumCores	0.04							
CPUFree	-0.07	0.67						
1MinLoad	0.10	-0.42	-0.72					
MemSize	0.03	0.37	0.37	-0.33				
MemFree	-0.07	0.37	0.37	-0.38	0.53			
DiskFree	-0.20	0.60	0.52	-0.41	0.44	0.44		
TxRate	0.06	-0.35	-0.39	0.30	-0.07	-0.20	-0.29	
RxRate	0.07	-0.33	-0.42	0.41	-0.11	-0.21	-0.29	0.86

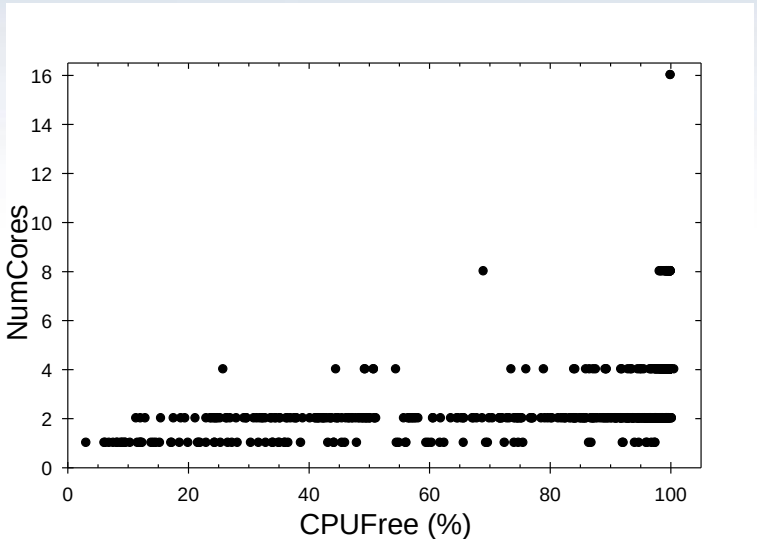

- Complex correlation among attributes
  - Static-dynamic attributes
  - Dynamic-dynamic attributes
- Contemporaneous correlation among time series
- Distinct temporal patterns

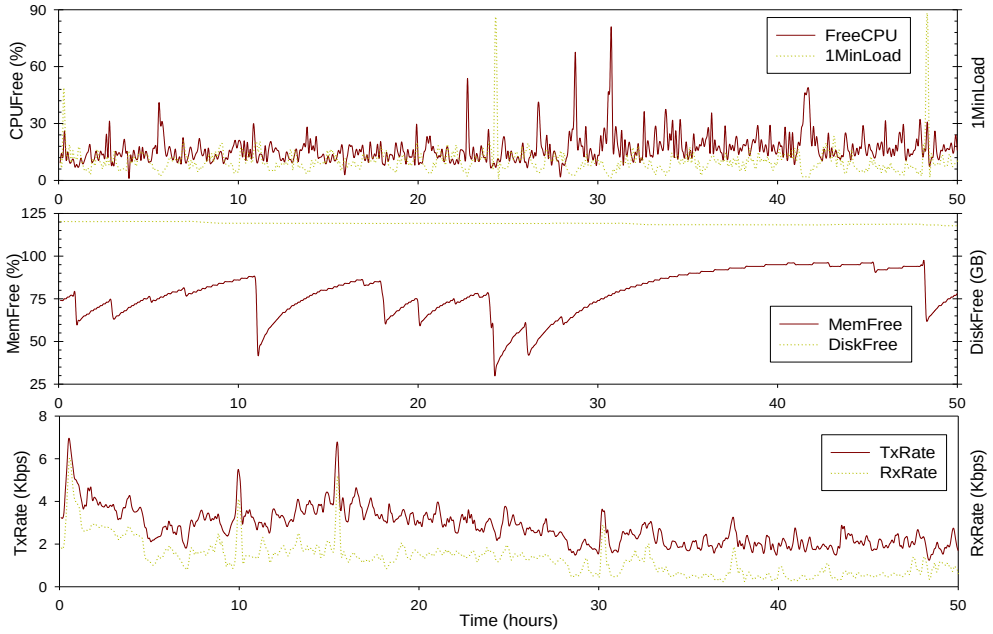

Resource Characteristics – Rate of Change in Dynamic Attributes

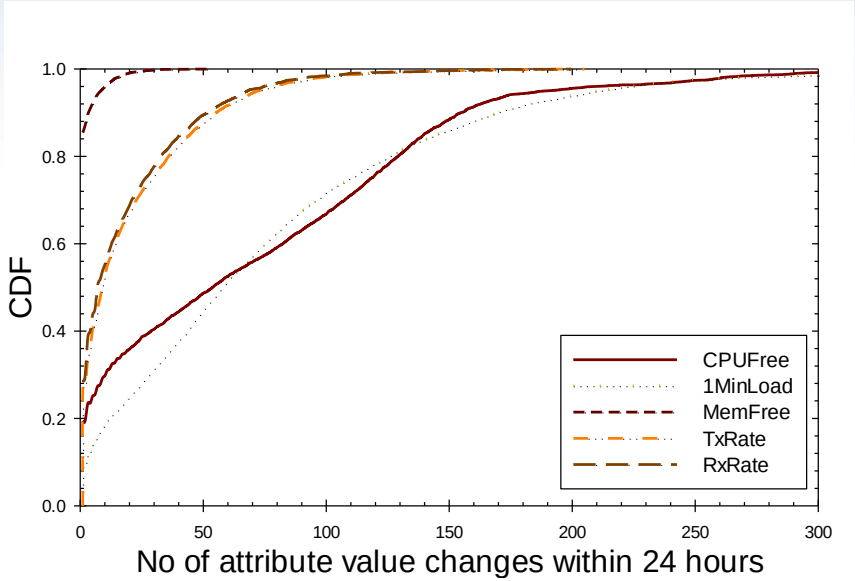

Thresholds: *CPUFree* = *MemFree* = ± 10%,
1MinLoad = ± 2, *TxRate* = *RxRate* = ± 1 Kbps

- Rate of change fits a generalized Pareto distribution
- Wide variation among attributes & their rate of change
  - 32% of nodes significantly change their dynamic attributes 100 or more out of 288 samples
  - Each attribute value change → New advertisement or status update

Query Characteristics

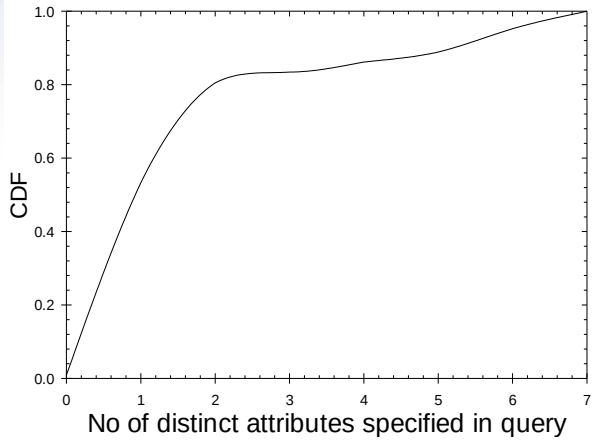

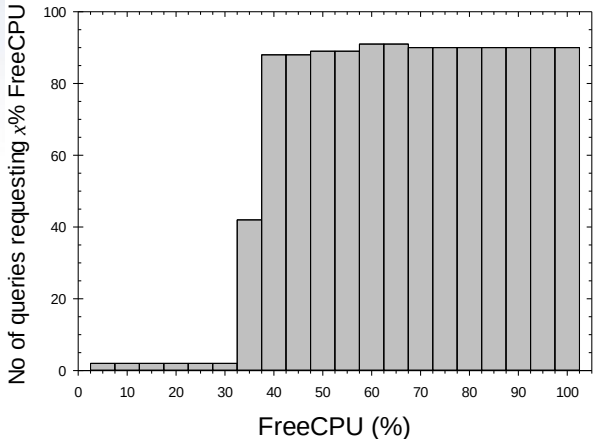

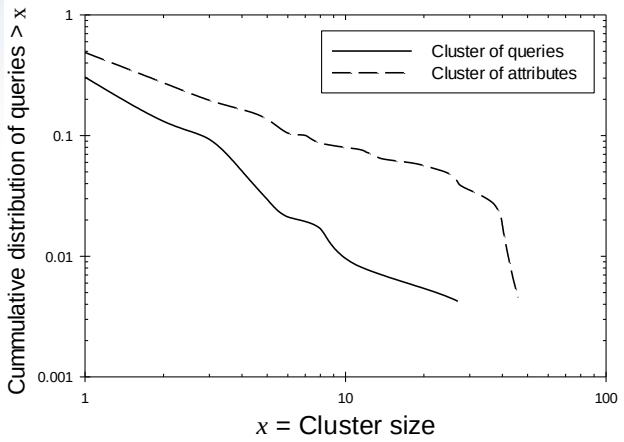

- Few attributes in a query
  - 80% queries specify 1 or 2 attributes
- Less specific attribute ranges
  - 89% queries request CPUFree of 40-100%
  - 70% queries request DiskFree of 5-1000GB
- Skewed but not Zipf's
- Dynamic attributes are popular
- Large number of resources
  - 73 resources per query

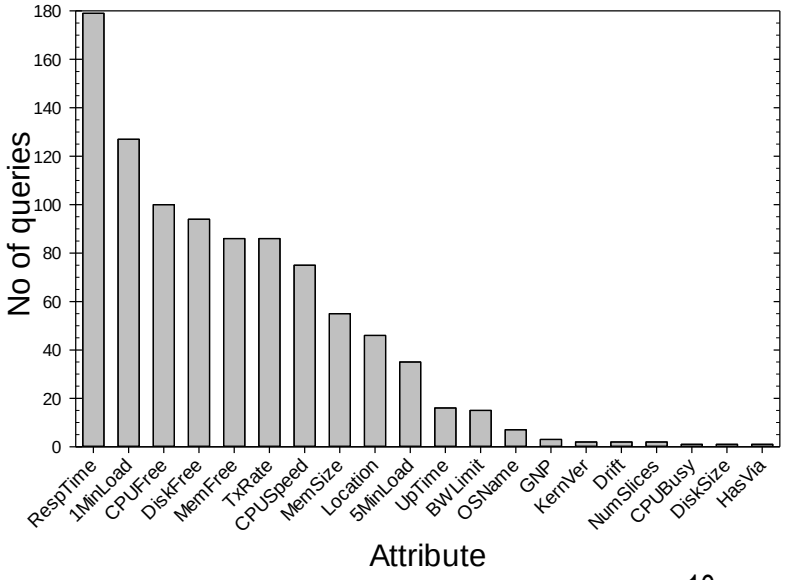

Design Choices – Unstructured P2P

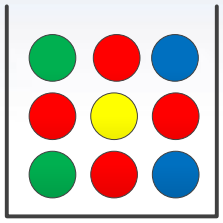

Centralized

$O(1)$

Single point of failure

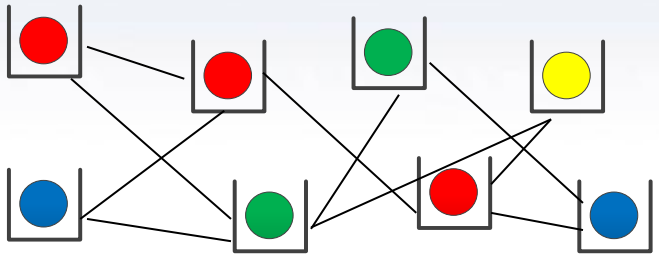

Unstructured

$O(hops_{max})$

Not guaranteed to find resources

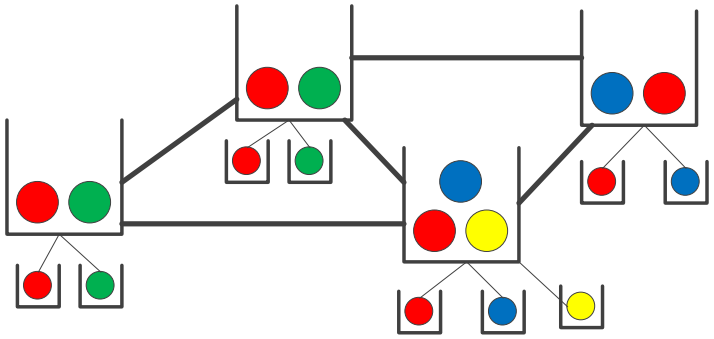

Superpeer

$O(hops_{max})$

Not guaranteed to find resources

Design Choices – Structured P2P

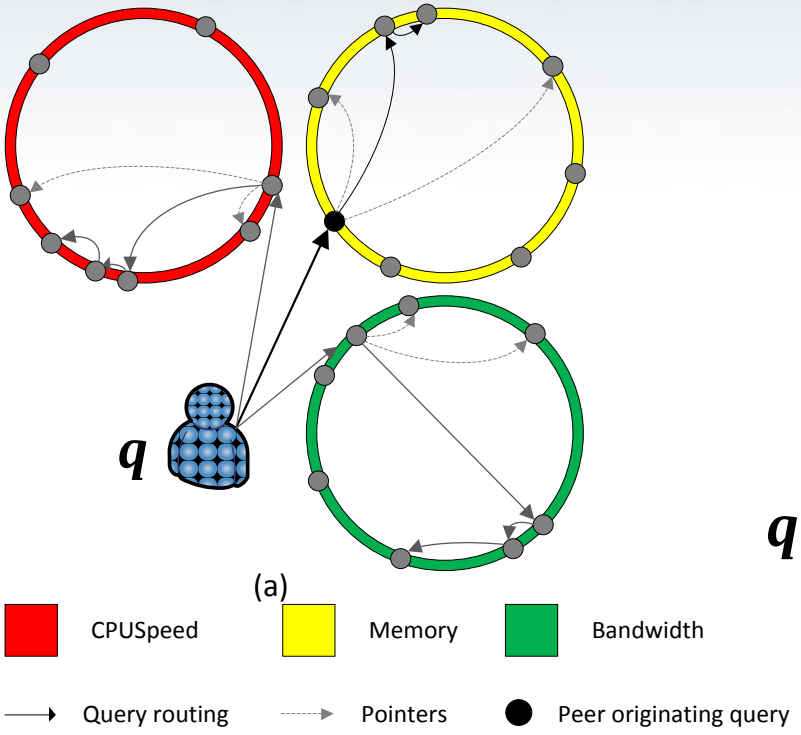

Multiple Rings

Mercury [Bharambe, 2004]

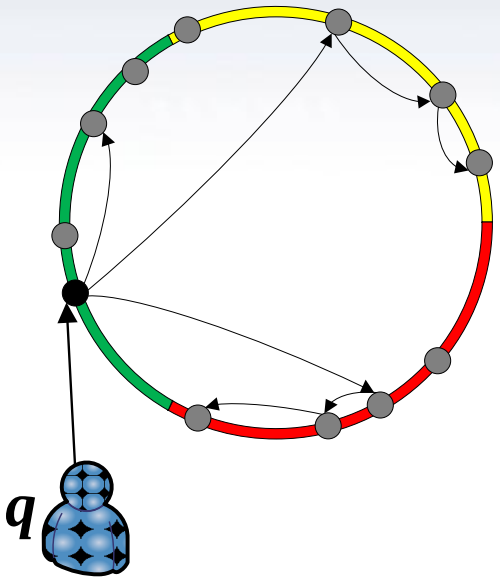

Partitioned Ring

LORM [Shen, 2007]

SWORD [Albrecht, 2008]

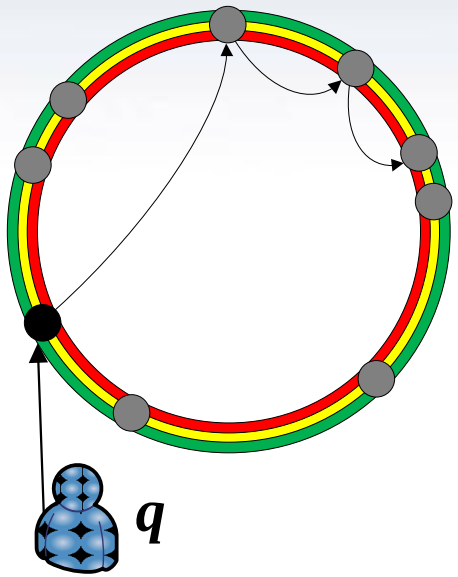

Overlapped Ring

MAAN [Cai, 2004]

$$C_{QUE}^{q} = \sum_{i \in q_A} \left(h_i + \left\lceil \frac{r_q^i}{r_{\max}^i} \times \frac{N}{A} \right\rceil - 1 \right)$$

Query modes

1. Sub–queries
2. Single Attribute Dominated Queries (SADQ)

Design Choices – Structured P2P (Cont.)

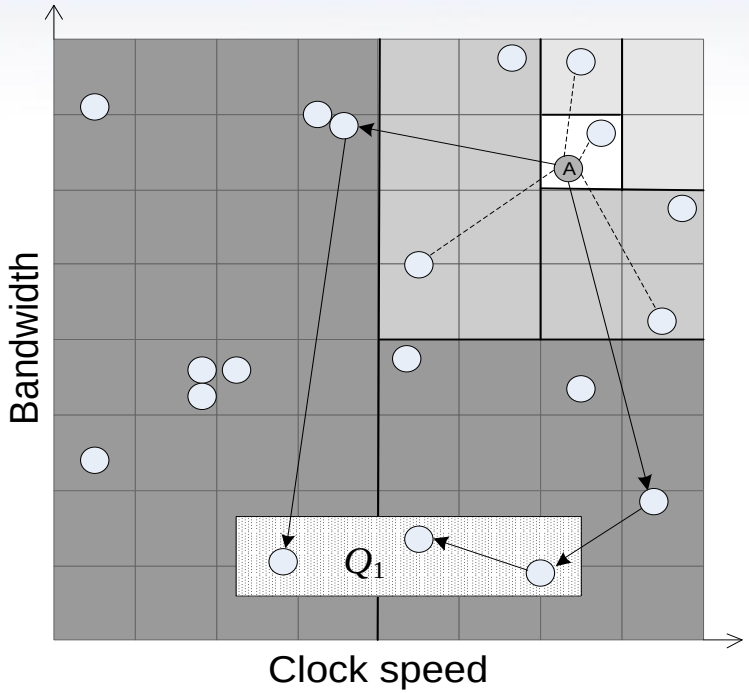

d-Torus [Costa, 2009]

- Connect resources based on their static attributes

Simulation Setup

- 7 representative architectures
- PlanetLab data to preserve mixture of distributions, complex correlation, temporal patters, etc.
  - Nodes active over 3 days→ 527 active nodes
  - Use 441 queries to generate large query set
    - Empirical distributions derived using no of attributes per query, popularity of attributes, range of attribute values, no of resources
    - Conditional probabilities to preserver correlation
- 750 & 1,000 node networks generated using our synthetic data generation tool [Bandara, 2011b]
- Unstructured & superpeer topologies
  - B-A scale-free network generator (available at mathworks.com)
  - hop_{max} = 100 (unstructured) & 10 (superpeer)
- Structured overlays based on Chord

Performance of Design Choices

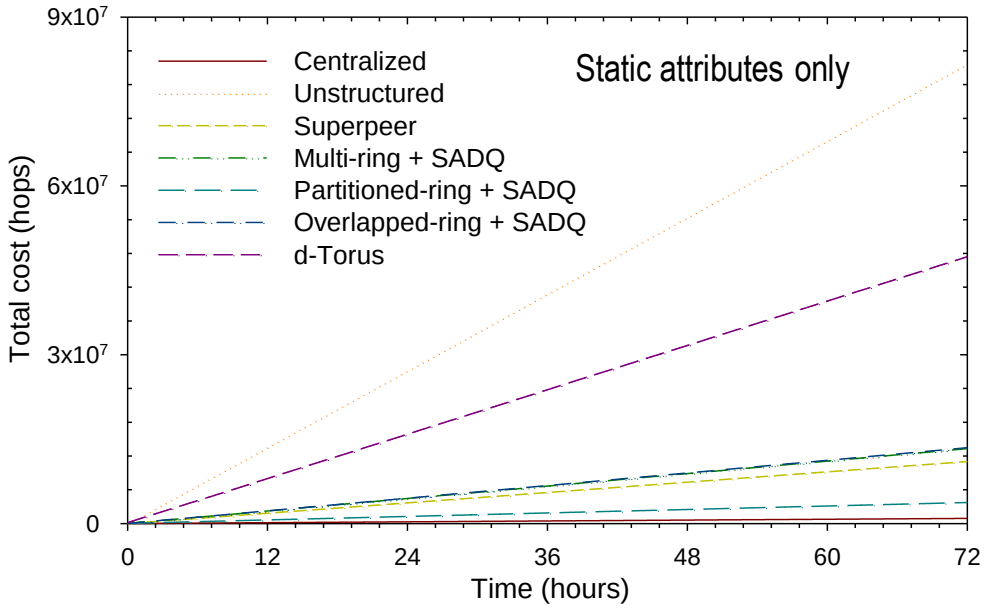

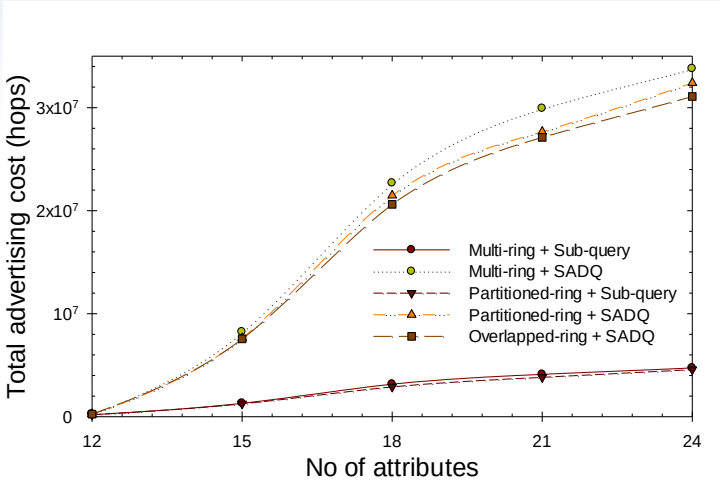

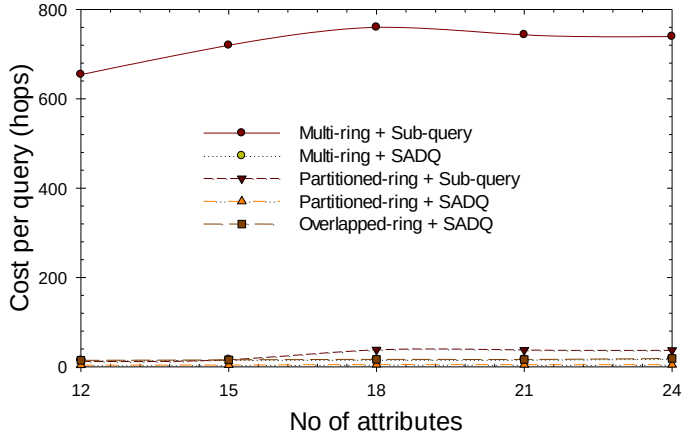

- Low cost
  - Centralized, superpeer, & partitioned ring
- Sub-queries
  - Low advertising cost & high query cost
- SADQ
  - Low query cost & high advertising cost

Simulation of Different Solutions (cont.)

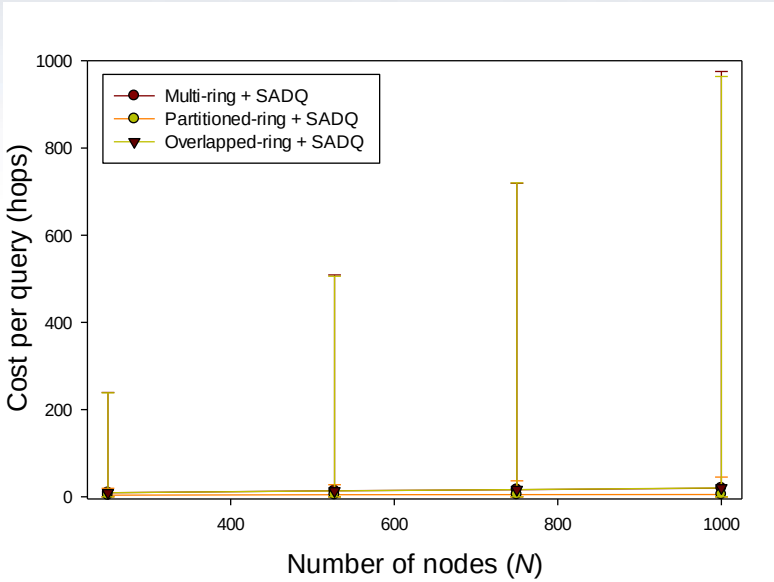

N	Multi-ring + SADQ			Partitioned-ring + SADQ			Overlapped-ring + SADQ		
	Min	Ave	Max	Min	Ave	Max	Min	Ave	Max
250	0	9.2	239.1	0	3.7	19.4	0	9.1	238.4
527	0	13.7	509.0	0	4.6	27.6	0	13.5	506.0
750	0	16.2	719.1	0	4.9	36.6	0	16.5	719.9
1000	0	19.8	975.5	0	5.3	45.3	0	20.4	963.8

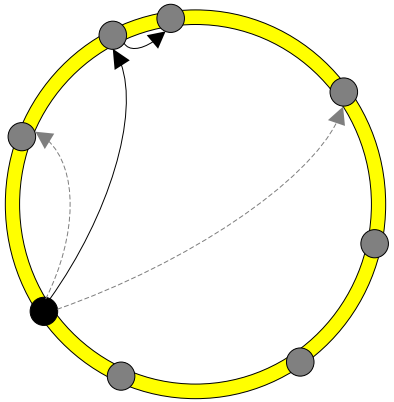

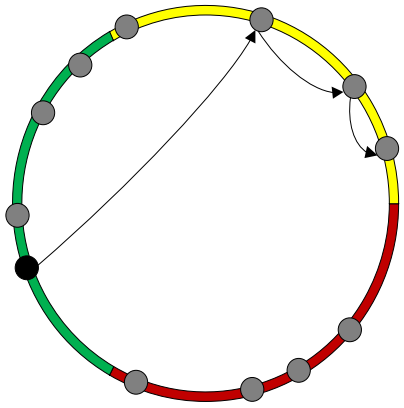

- Large range of attribute values → cost of ring-based designs $O(N)$

Load Distribution

Architecture	Total Cost per Query		Query Load				Index Size	
			Min		Max		Min	Max
	SWORD	Uniform	SWORD	Uniform	SWORD	Uniform		
Centralized	2.03	2.03	950,000	950,000	950,000	950,000	527	527
Unstructured	69.5	94.8	4,859	1,272	268,497	37,824	1	1
Superpeer	6.5	9.5	81,021	22,390	289,626	87,209	17	36
Multi-ring + SADQ	48.3	69.0	0	0	178,492	22,943	0	527
Multi-ring + Sub-queries	398.8	120.8	0	0	624,837	57,518	0	230
Partitioned-ring + SADQ	36.6	37.0	0	0	185,972	15,840	0	527
Partitioned-ring + Sub-queries	40.7	16.4	0	0	432,859	46,946	0	527
Overlapped-ring + SADQ	46.0	67.2	0	0	391,738	57,524	0	527

N = 527, *Attributes* = 24

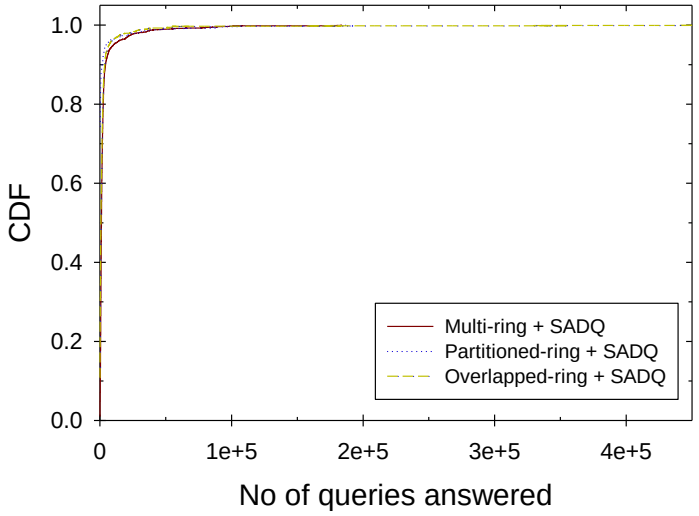

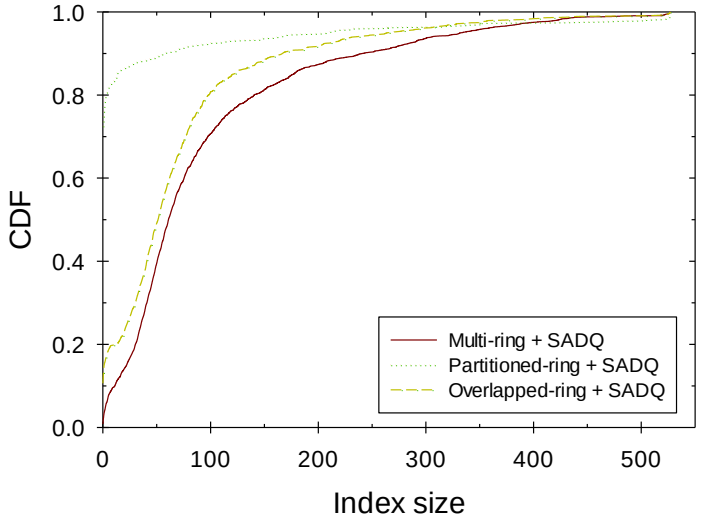

- Unbalanced index size & query load

Summary of Findings

- Rapidly changing dynamic attributes
  - High advertising cost
  - Their impact should not be ignored
- Less specific range queries
  - Resources are relatively easier to find using unstructured & superpeer
  - Partitioned-ring have lower cost
  - Sub-queries are expensive
  - Cost of ring-based designs is effectively $O(N)$
- Skewed distribution of attribute values & popularity
  - Few nodes index most resources
  - Few nodes answer most queries
  - In partitioned & multiple rings – some nodes don't index or answer queries at all

Conclusions

- Complex real-world resources & queries
  - Mixture of distributions, correlated, skewed, dynamic attributes, less specific queries
- Solutions lack scalability
  - Each solution is applicable under very specific conditions
    - Real-world queries are relatively easier to resolve using unstructured & superpeer
    - Effective cost of ring-based designs $O(N)$
  - Load balancing issues in all solutions (except unstructured P2P)
    - P2P load balancing solutions also fail under real-world characteristic
- Better solutions are need
  - More scalable, efficient, & robust
  - Consider real-world resource & query characteristic at design time
  - Hybrid designs

RESQUE – RESource & QUEry Generator

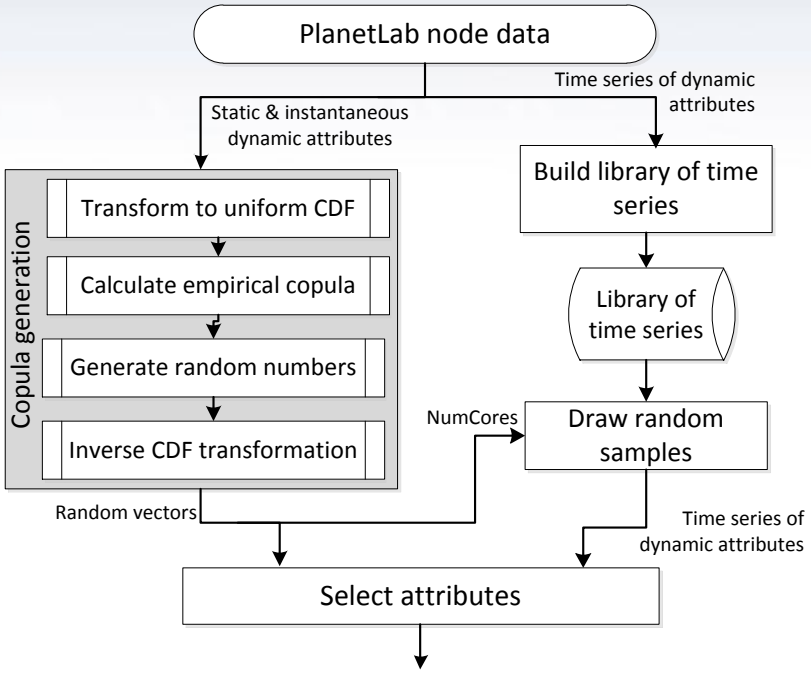

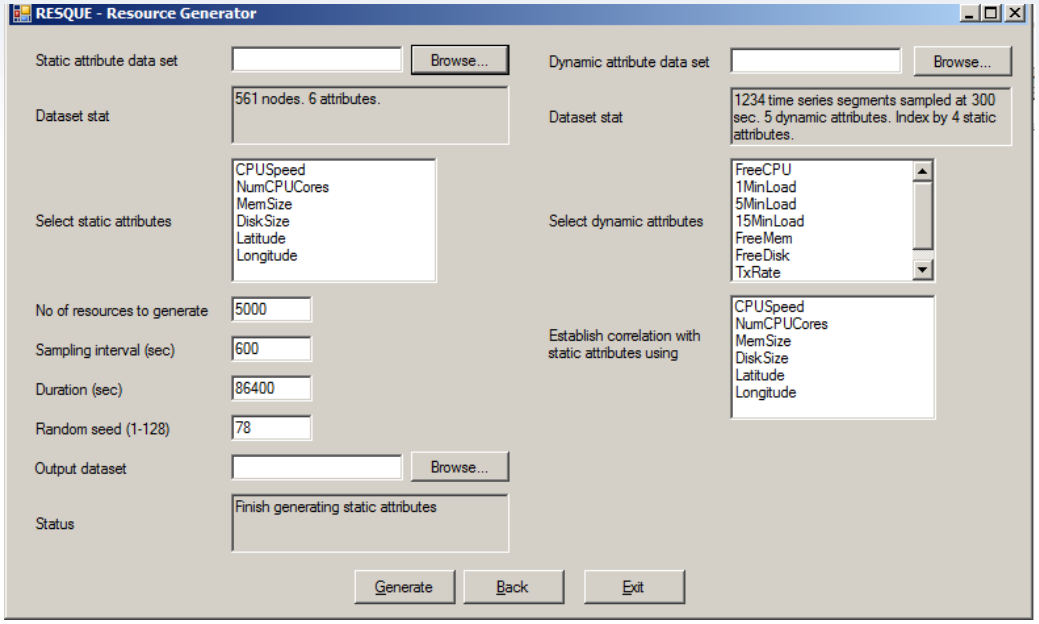

- Generate large synthetic traces of multi-attribute resources & queries while preserving correlation & temporal patterns
- Useful for large-scale performance studies in P2P, grid, & cloud computing
- Beta version available - www.cnrl.colostate.edu/Projects/CP2P/

Engineering Research Center for
Collaborative Adaptive Sensing of the Atmosphere

Questions/Comments

Anura.Jayasumana@ColoState.edu

www.cnrl.colostate.edu/Projects/CP2P/


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.