Diurnal Availability in un-structured P2P Systems

Nashid Shahriar[1], Mahfuza Sharmin[1], Reaz Ahmed[1],
Md. Mustafizur Rahman[1], Raouf Boutaba[2], Bertrand Mathieu[3]

[1]Department of Computer Science and Engineering
Bangladesh University of Engineering and Technology
[2]School of Computer Science, University of Waterloo
[3]Orange Labs

Overview

- **Introduction**
- State of the Art
- Our Contributions
  - Mathematical Model for Availability
  - A protocol: DATA
- Performance Comparison
- Future Directions

Introduction

- A persistent P2P system leveraging 24x7 availability through virtual assemblage of peers having diurnal uptime
- Possible applications
  - Content sharing
  - Distributed web hosting through peers
  - Distributed social networks
- The challenges are to ensure:
  - Consistent availability of data
  - Minimum overhead
  - Fairness in load distribution

University of Waterloo

orange™

Overview

- Introduction
- **State of the Art**
- Our Contributions
  - Mathematical Model for Availability
  - A protocol: DATA
- Performance Comparison
- Future Directions

Existing Strategies

- High availability is achieved mostly by Replication
- Replica management policies
  - Reactive - System reacts to a change in the adopted criteria
  - Proactive - System predicts the future behavior through monitoring
- Uncontrolled replication impose additional burden on the
  - Storage requirement
  - Bandwidth consumption
  - Update propagation
  - Routing process
- Major objective is to keep the number of replicas small

Shortcomings of Existing Strategies

- Consider only the current but single score of availability
- Do not utilize continuous connectivity information across time

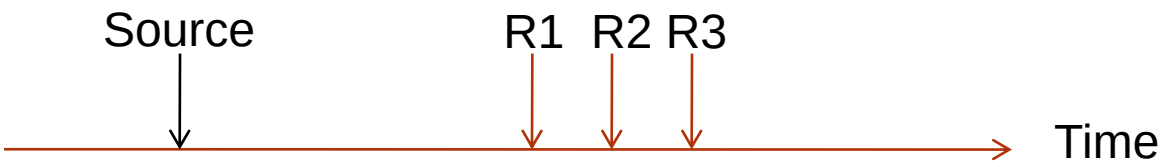

- Reactives approaches are bandwidth hungry
- Proactive ones require complex predictive knowledge
- Burden highly available peers with a skewed load distribution
- Make no distinctions between transient and permanent failure

Overview

- Introduction
- State of the Art
- **Our Contributions**
  - **Mathematical Model for Availability**
  - A protocol: DATA
- Performance Comparison
- Future Directions

Our Contributions

- A representation for the availability patterns
- Measurement techniques for comparing patterns
- A gossip based peer discovery mechanism
- A distributed algorithm to form small groups of peers
- Altogether a protocol named *DATA (Diurnal Availability by Temporal Assemblage of peers)*
- Performance evaluation of the protocol

Global Availability Pattern

- Peers in Skype, Gnutella, Overnet, eDonkey, etc follow diurnal (daily cyclic) pattern governed by time-of-day effect
- The cyclic behavior of the peers situated on different time zones can be found overlapping or complementary
- Combined availability can be improved through virtual assemblage of out of phase peers

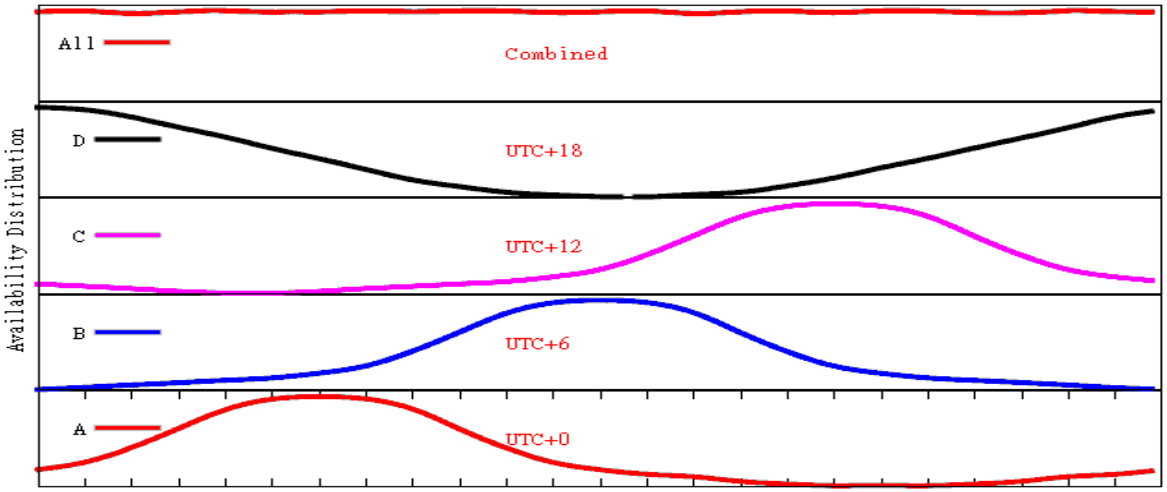

Discrete Model of Availability

- We propose a discrete model to represent availability across time
- Divide a day into *n* slots of equal length *l*, *n=24/l* slots per day
- The availability of a peer in a slot is the probability of its being available at that slot
- Computed from previous history and updated by periodic scan
- Availability in all slots is represented by an *Availability Vector* (A_i)

Peer	Slot 1 00:00-01:00	Slot 2 01:00-02:00	....	Slot 23 22:00-23:00	Slot 24 23:00-00:00
P_1	a_{11}	A_{12}	…	a_{123}	a_{124}
P_2	0.80	0.75	…	0.05	0.10
P_3	0.55	0.35	…	0.85	0.70

Availability Measurement

- $C_{i,j}$ is a measure of the improvement on availability after P_i and P_j form a group
- Our protocol uses this metric to find the best matching peer
- More complementary availability vectors should result in a higher $C_{i,j}$
- Availability of a group is measured by the probability of at least one member of the group being present

Availability Measurement

- A "*conservative*" equation

$$C_{i,j} = \sum_{k=1}^{K} \frac{C_{ijk}}{|G_i \cup G_j|}$$

$$C_{ijk} = \begin{cases} J_{ijk}^{(\frac{a_{ik}}{a_{jk}})} - J_{ijk} & \text{if } a_{ik} \leq a_{jk}; \\ J_{ijk}^{(\frac{a_{jk}}{a_{ik}})} - J_{ijk} & \text{if } a_{ik} > a_{jk}. \end{cases}$$

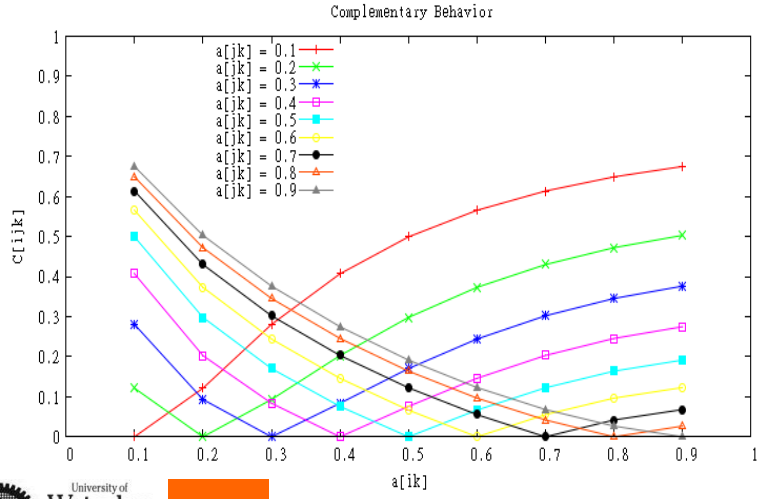

- A "*general*" equation

$$C_{i,j} = \frac{(U_{ig} + U_{jg})}{|G_i \cup G_j|}$$

$$U_{ig} = \sum^{K} (a_{gk} - a_{ik})$$

$$U_{jg} = \sum_{k=1}^{K} (a_{gk} - a_{jk})$$

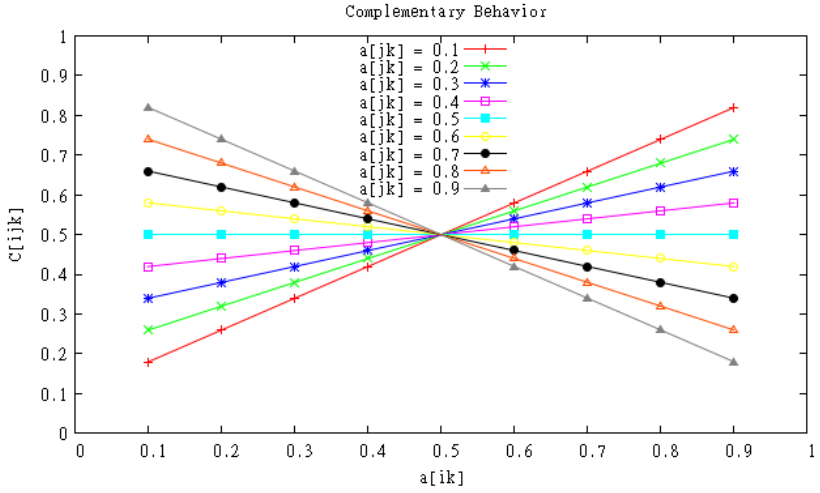

University of Waterloo

orange

Overview

- Introduction
- State of the Art
- **Our Contributions**
  - Mathematical Model for Availability
  - **A protocol: DATA**
- Performance Comparison
- Future Directions

Exploration

- *DATA* use a gossip based network discovery called *Exploration*
- Each peer gathers availability vectors of the peers within its two-hop distance
- Using the collected information it constructs a local list of current best candidates named as *knownlist*
- *Knownlist* is used in the next phase to keep the search space relatively smaller

Grouping Phase

- Peers communicate each other to form groups utilizing the complementary uptime distribution
- A group is constructed incrementally, i.e., forming groups with two single peers initially then growing in size
- Later, two non-Singleton groups merge into a larger one
- Grouping continues if the resultant availability of the new group increases from that of the former two groups

Grouping Phase

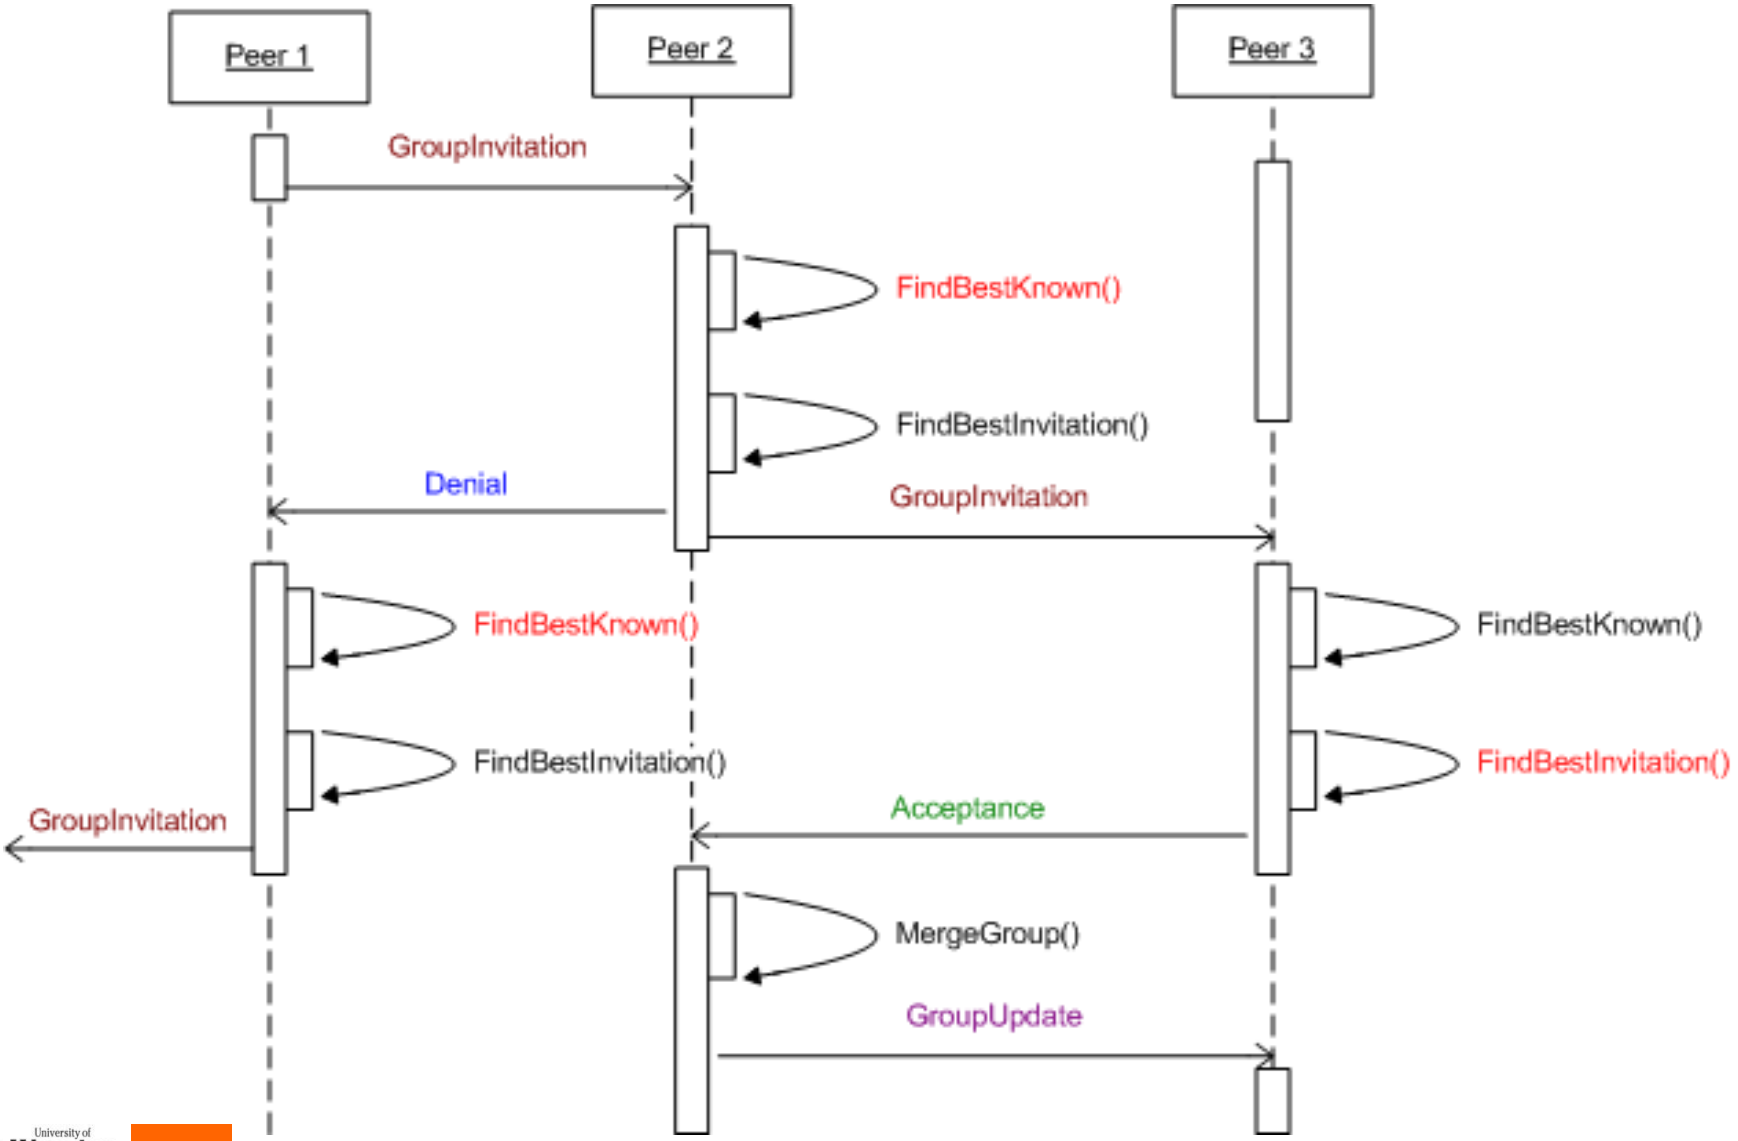

University of Waterloo

orange™

Group Maintenance

- Groups can revoke the membership of a peer
- If it is no more contributing to the availability increase
- The peer that is *Revoked* from a group,
  - executes exploration with its modified vector
  - go through grouping activities again

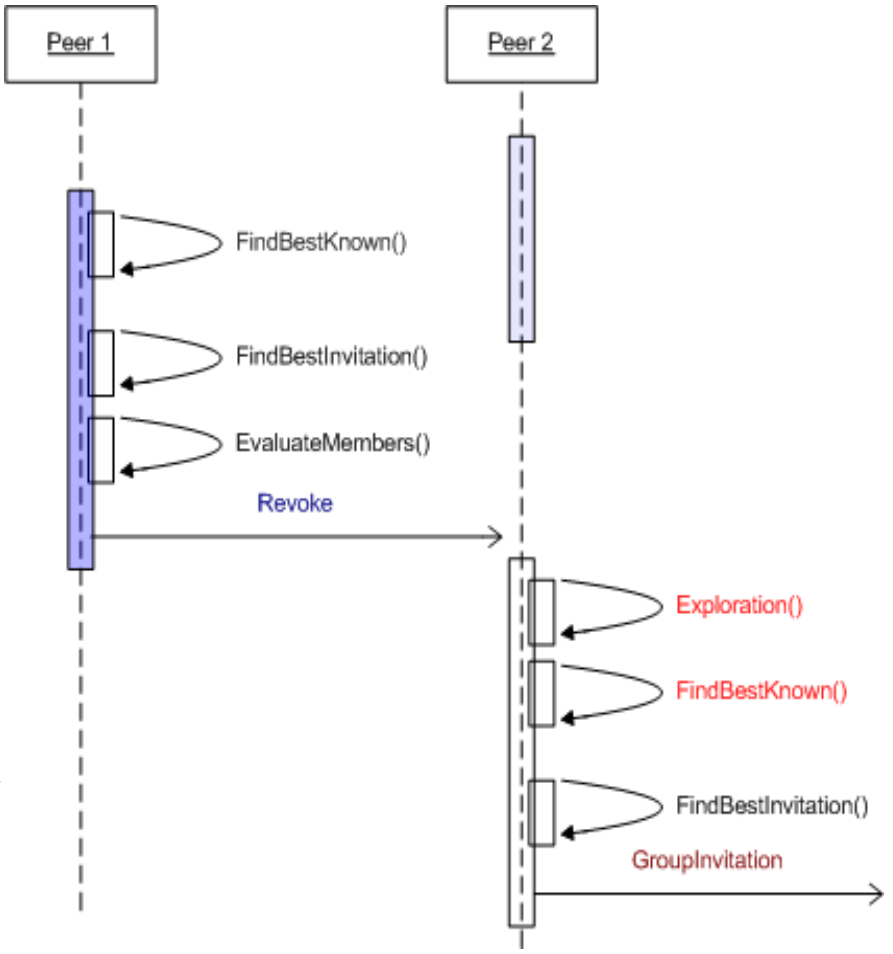

Overview

- Introduction
- State of the Art
- Our Contributions
  - Mathematical Model for Availability
  - A protocol: DATA
- **Performance Comparison**
- Future Directions

Grouping Strategies

- Using local knowledge two methods on the proposed metrics-
  - *Conservative*
  - *General*
- Random grouping among the currently alive two hop neighbors
- Centralized Oracle based scheme using global knowledge
- Even using, Oracle optimal grouping is an NP-complete problem
- Need a greedy strategy for near optimal solution

Experiment Overview

- We simulate an unstructured P2P network using Peersim varying time-zone disparity
- Performance metrics:
  - Group availability- Measured in units of nines defined as -log10(1-T) , where T is the fraction of the total observed time when groups are available
  - Group count- Total number of groups created by each strategy.
  - Message overhead- Total number of messages exchanged by the protocol normalized by the group count

Performance Comparisons

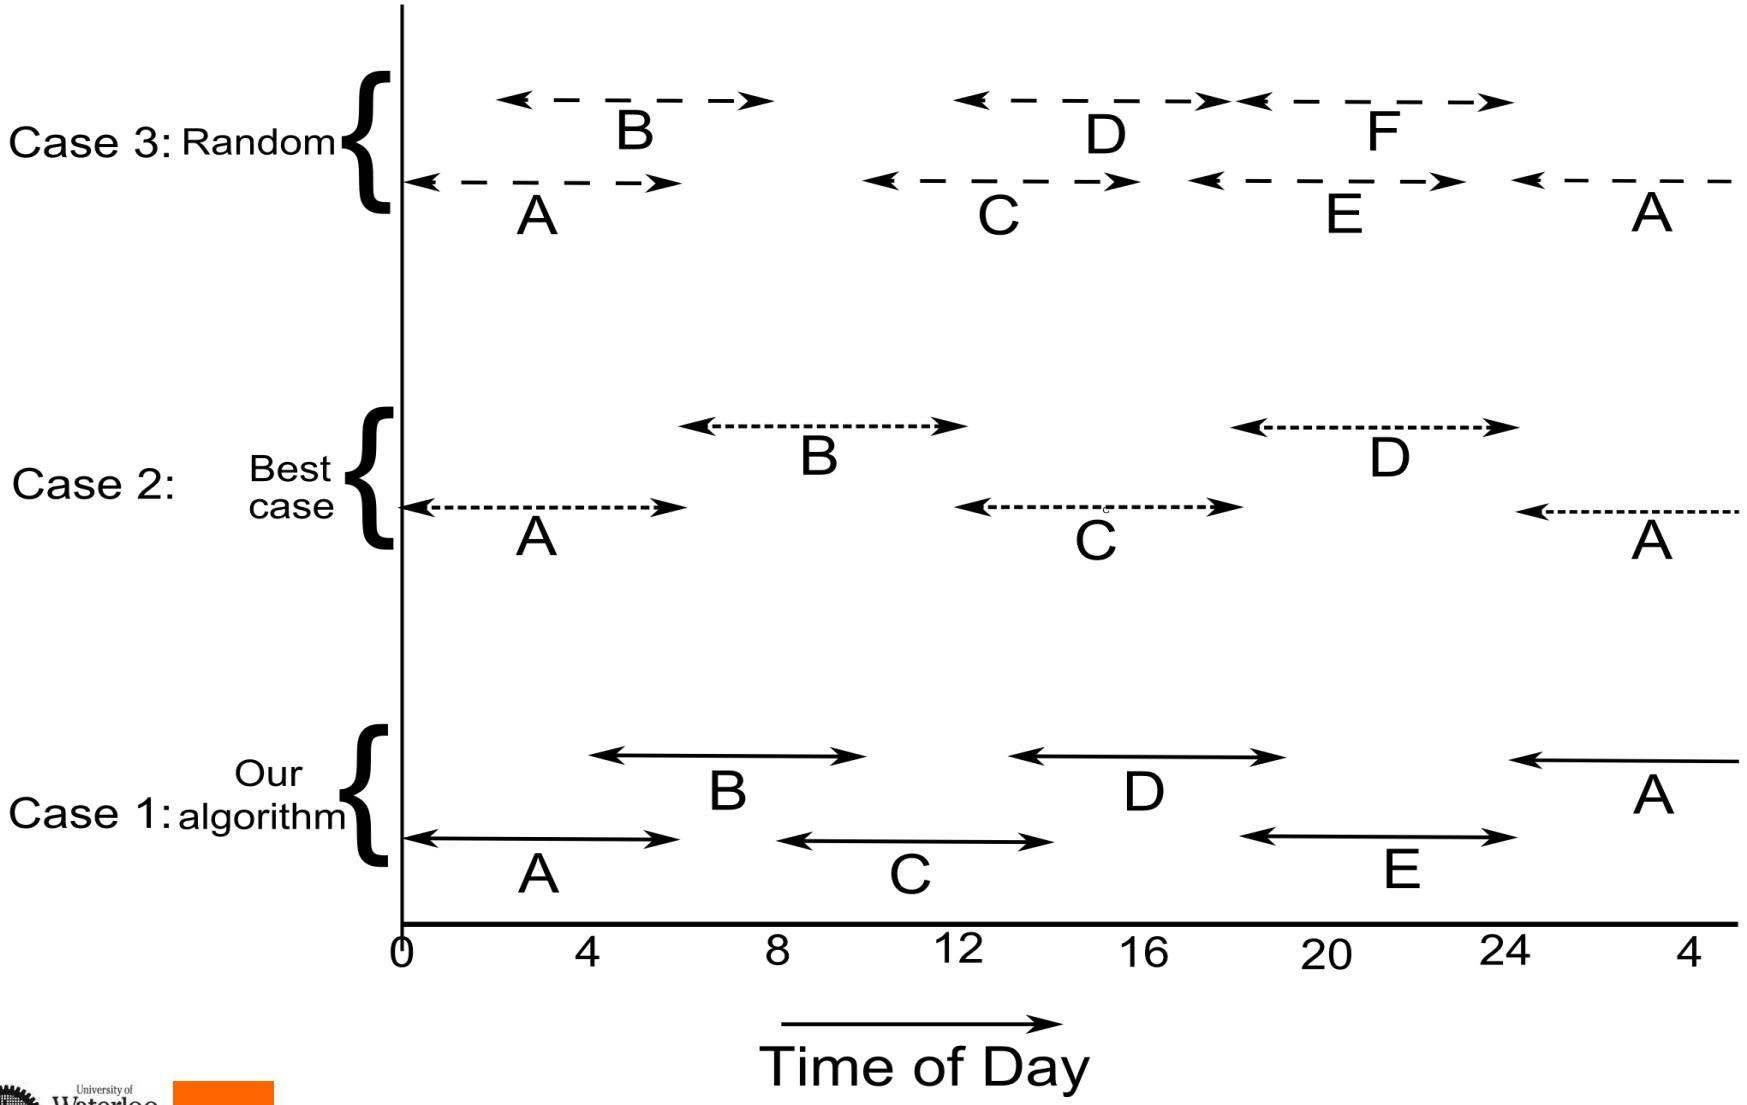

Performance Comparison

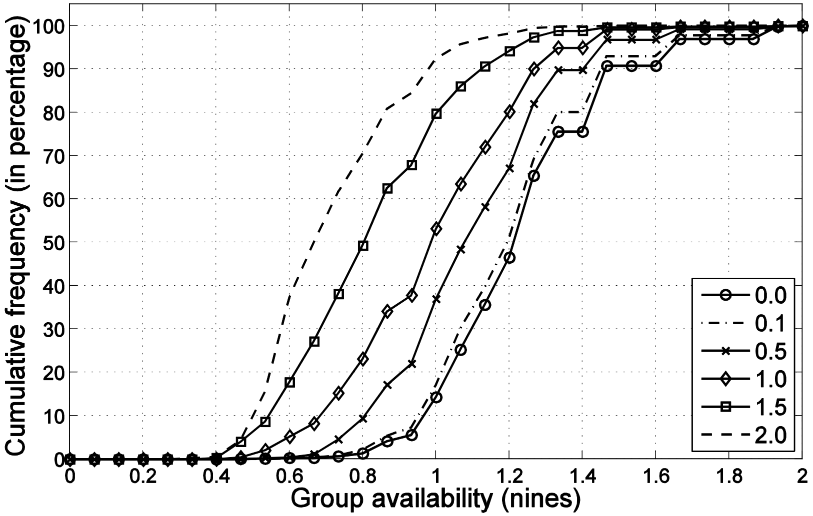

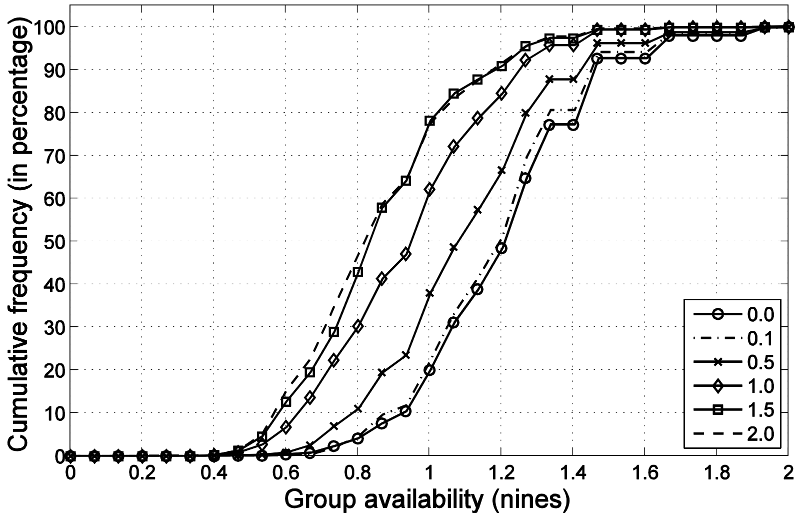

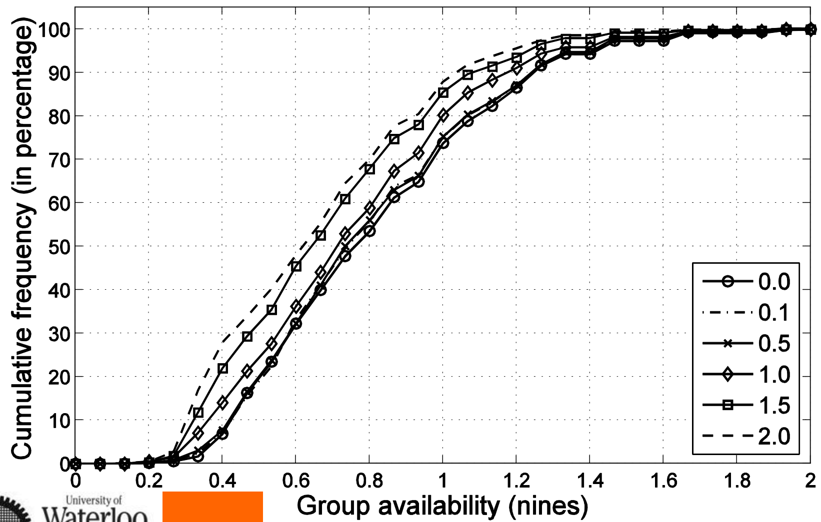

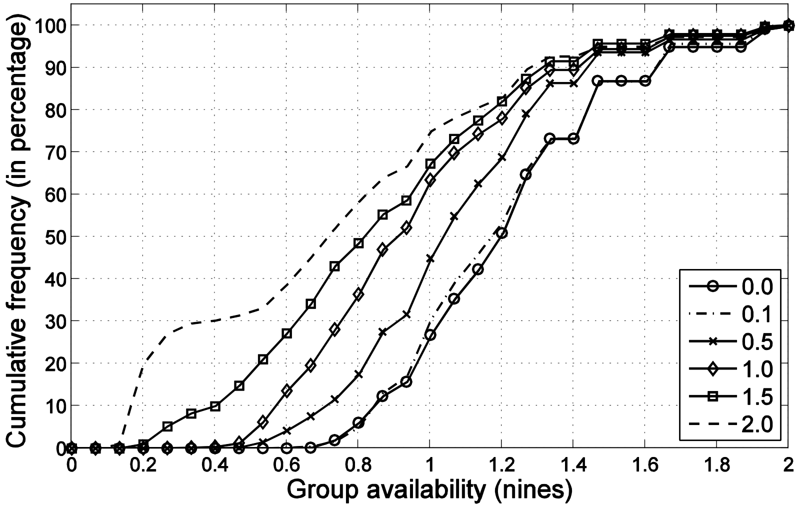

University of Waterloo

orange

Performance Comparison

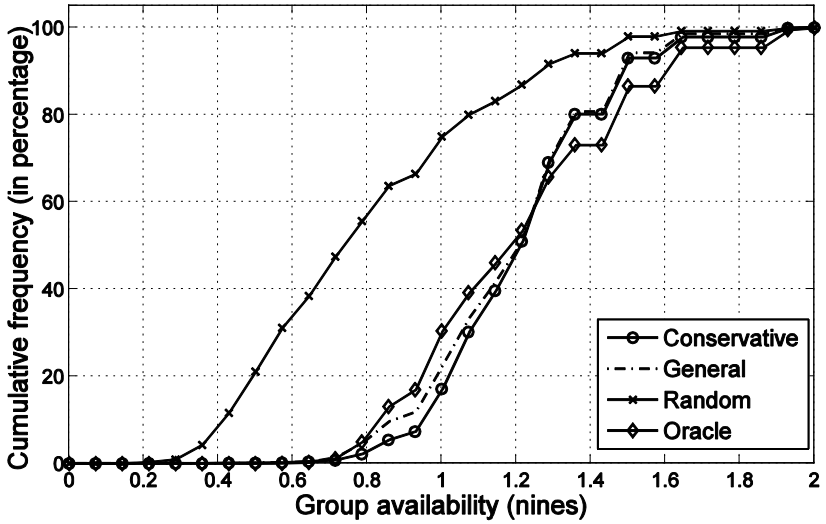

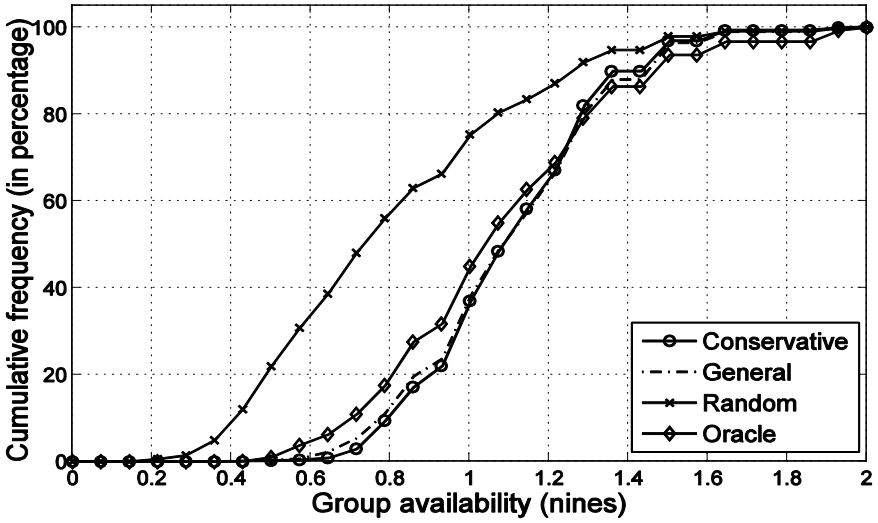

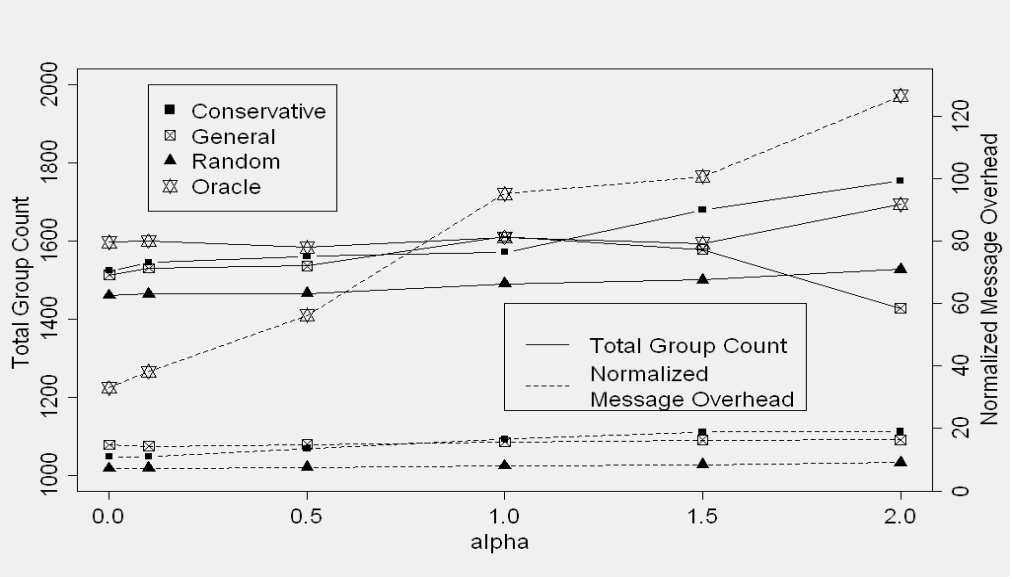

University of Waterloo

orange

Overview

- Introduction
- State of the Art
- Our Contributions
  - Mathematical Model for Availability
  - A protocol: DATA
- Performance Comparison
- **Conclusion**

Conclusion

- Our experimentation reveals following results
  - *DATA* ensures high availability around the clock for transient P2P environments
  - Keeps the number of replicas small
  - Requires no centralized mechanism or global knowledge
  - All the participating peers contribute resources for a portion of day
  - Ensures fairness in load distribution
  - The proposed model becomes more precise when the dimension of the vector $\mathcal{A}_i$ becomes larger

Future Directions

- Investigate performance of the protocol in a real system
- Tackle the malicious peers
- Security issues related to group formation
- Apply a continuous model to estimate availability
  - $\int \mathcal{A}(t)\, dt$
- Refine the technique for a structured network using a DHT
- Globally optimizing the grouping strategy using distributed pattern matching (e.g., Plexus)

Thank You

For any question please email to:
nshahriar@csebuet.org