MuON: Epidemic based Mutual Anonymity in Unstructured P2P Networks

System Model

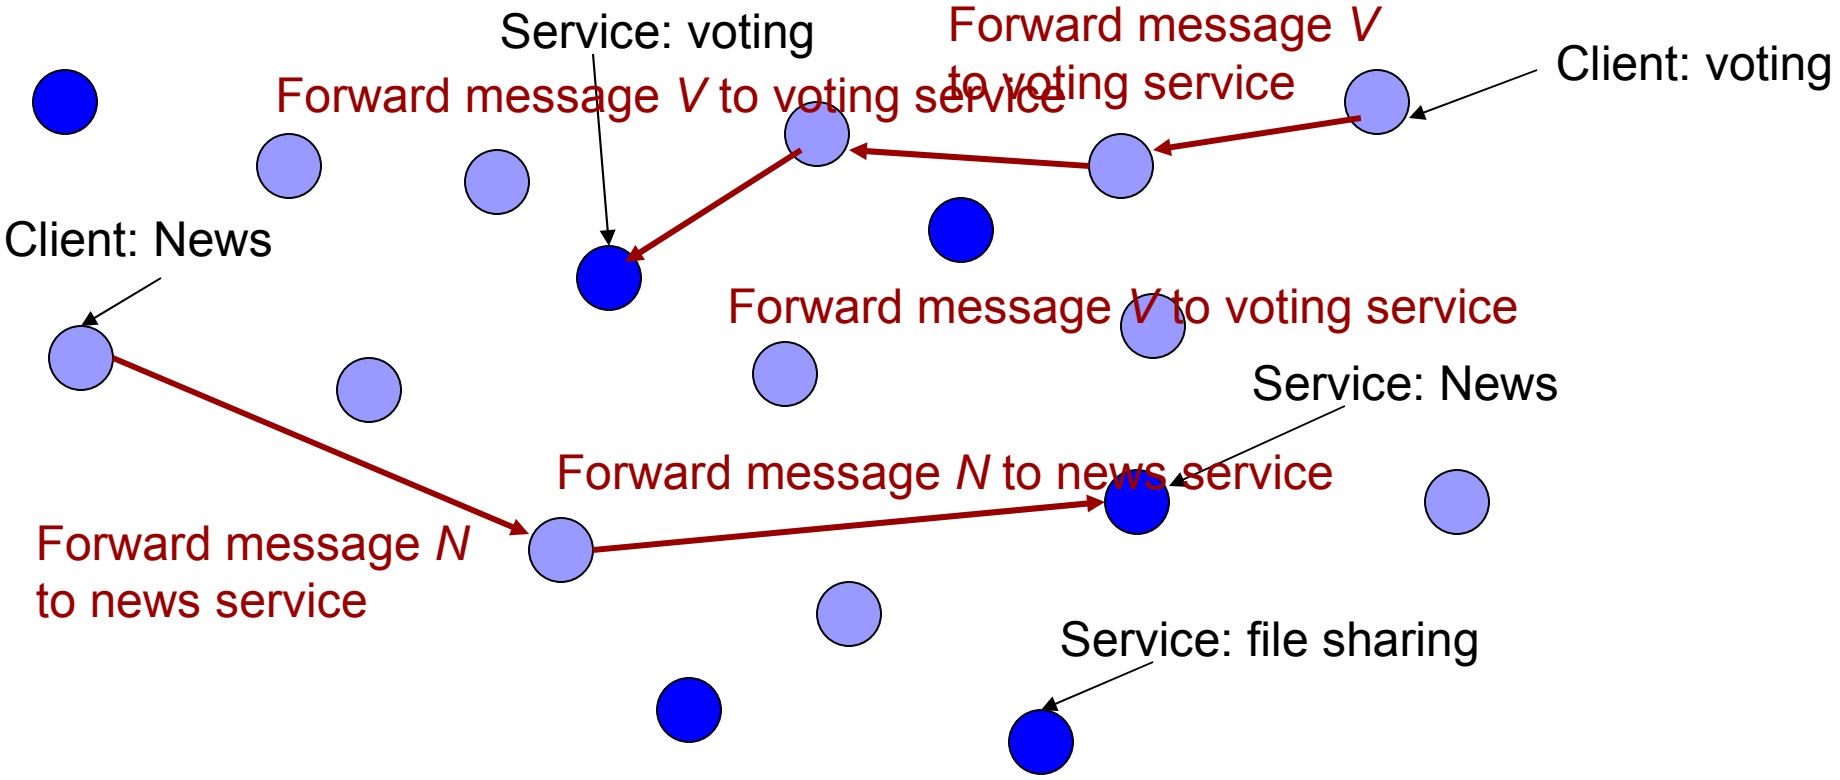

- Mutual anonymity needs to be searchable
  - Initiator knows the service, but not the provider

System Model

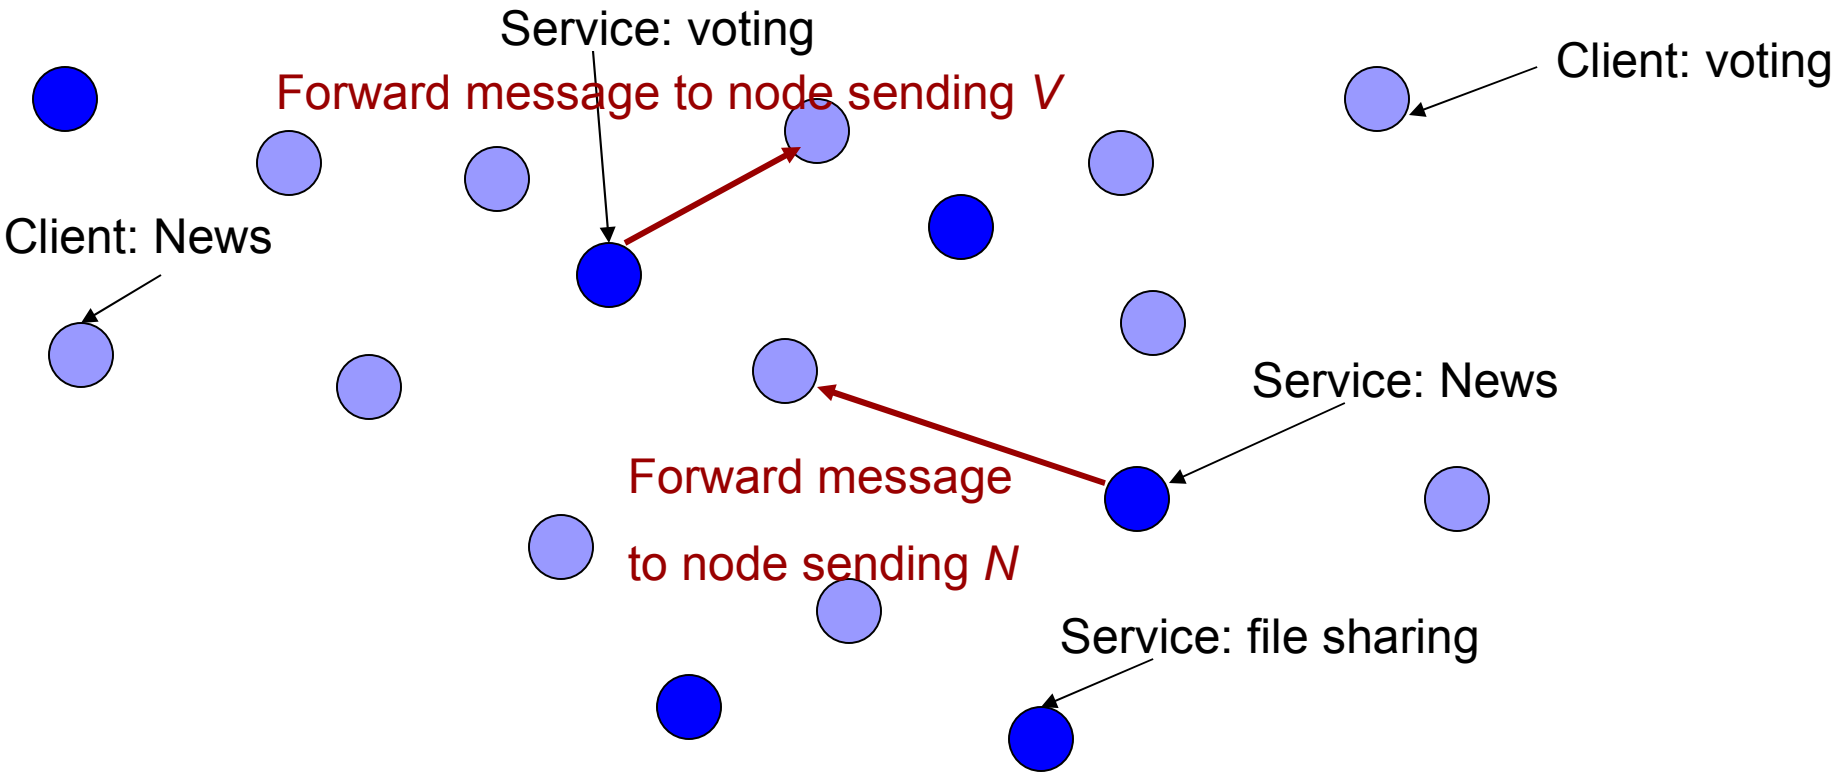

- Mutual anonymity needs to be searchable
  - Initiator knows the service, but not the provider

Goals of the Protocol (MuON)

- Mutual anonymity and unlinkability
- Group communication-based
- Single group
- Peer-to-Peer voluntary networks
- Low communication overhead
- Scale with increasing overlay size and churn
- High reliability
- Bounded latency

Design Decision

- Use of Multicasting
  - Provides mutual anonymity
  - Tolerates peer dynamics
  - Short latencies
  - Since many nodes get the message, any one of them could be the initiator/responder

- Overheads in multicast
  - High because so many nodes get the message
  - Consumes bandwidth and computational resources

Basic Principle

- A header is generated for each message
  - size of header much less than size of message
  - Headers contain signed checksums for integrity
  - Messages are encrypted for confidentiality
- Dissemination of data
  - Headers are disseminated to all N nodes
  - Messages are disseminated to $p*N$ nodes ($0 < p < 1$)
  - p is a protocol parameter, tradeoff between anonymity and performance
- Based on a class of protocols called *epidemic protocols*

Data Dissemination: Epidemics

- Mimic spread of a contagious disease
- Each node exchanges state with randomly chosen peers.
- Ensures each message copy is sent to ~ *log(n)* members

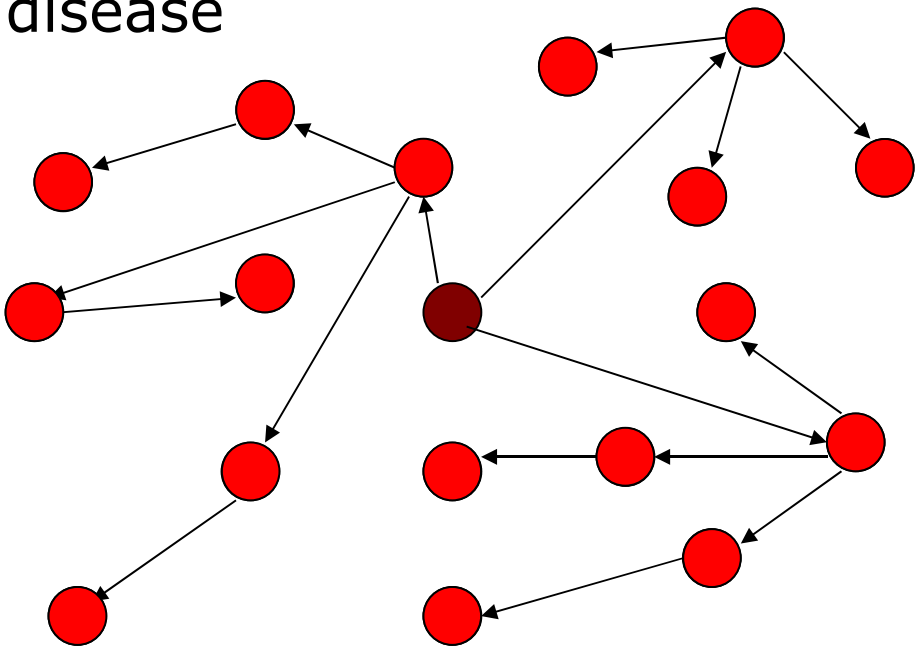

- Easy to deploy, robust, highly resilient to failures
- Shown to be more scalable than SRM and more efficient than flooding

We adapt epidemics so that message don't have to be multicast to all nodes

Epidemics/Gossip Protocols

- Form of application level multicast
- Mathematically proven that data is multicast in *log(n)* steps.
- Overlay almost a multicast group.
- Shown to be highly reliable and scalable in dynamic overlays

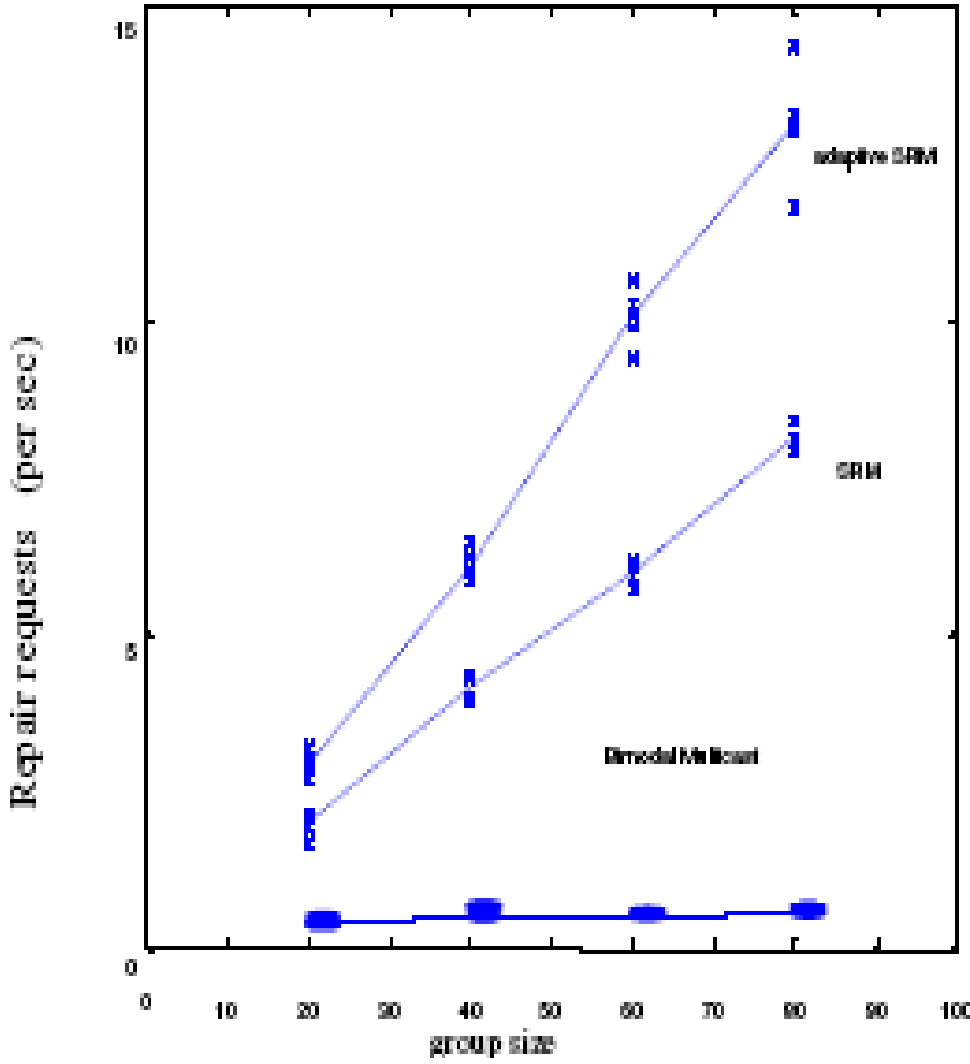

MuON: Basic idea

- Public key is used to identify recipient
- Small header created for each message
  - Header is encrypted using public key of recipient
  - Message and header have suitable encryption, checksums and nonce for confidentiality and integrity.
  - Non-encrypted field in header called *owner*
    - IP of node (or *owner*) with message

MuON: Basic idea

- Header is multicast to all members
  - Use an epidemic protocol for multicasting
    - Reliable delivery to all members
    - Bounded latency in large groups

- Larger message sent to subgroup
  - Dynamically created subgroups
    - For a given header, a peer pulls the corresponding message from the owner with probability P_{inter} intermediate probability
      - Node becomes the new owner
    - If a node can decrypt the header, it pulls the message
      - Reliable delivery at recipient

Anonymous Dissemination

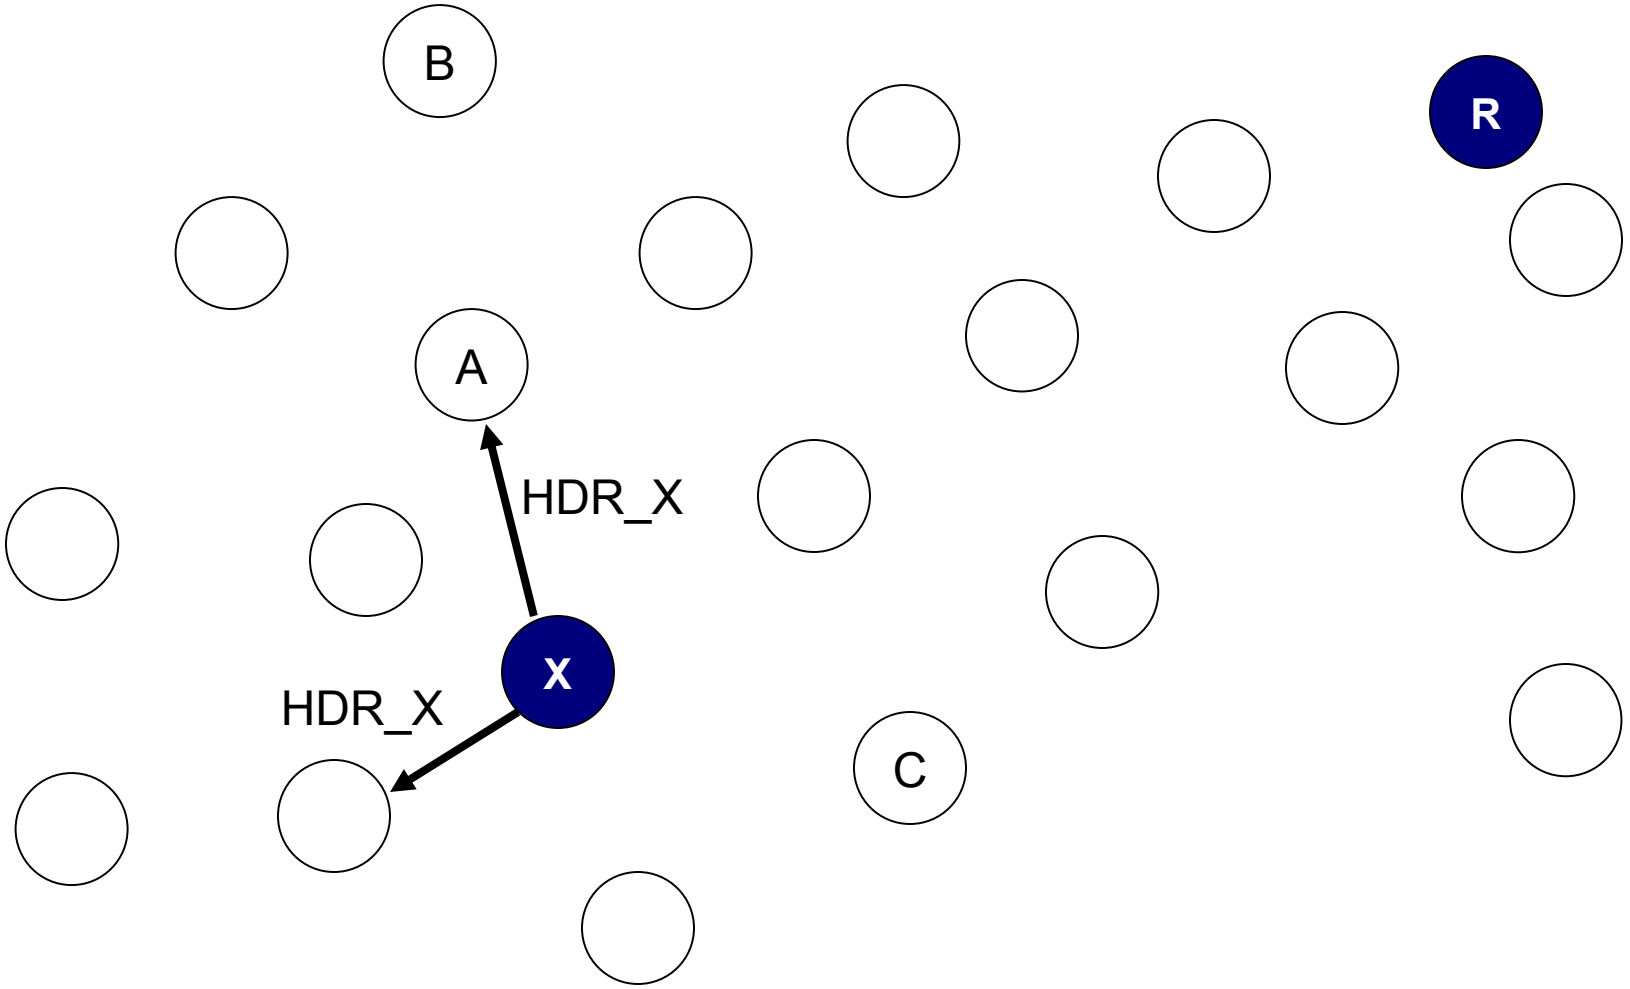

Anonymous Dissemination

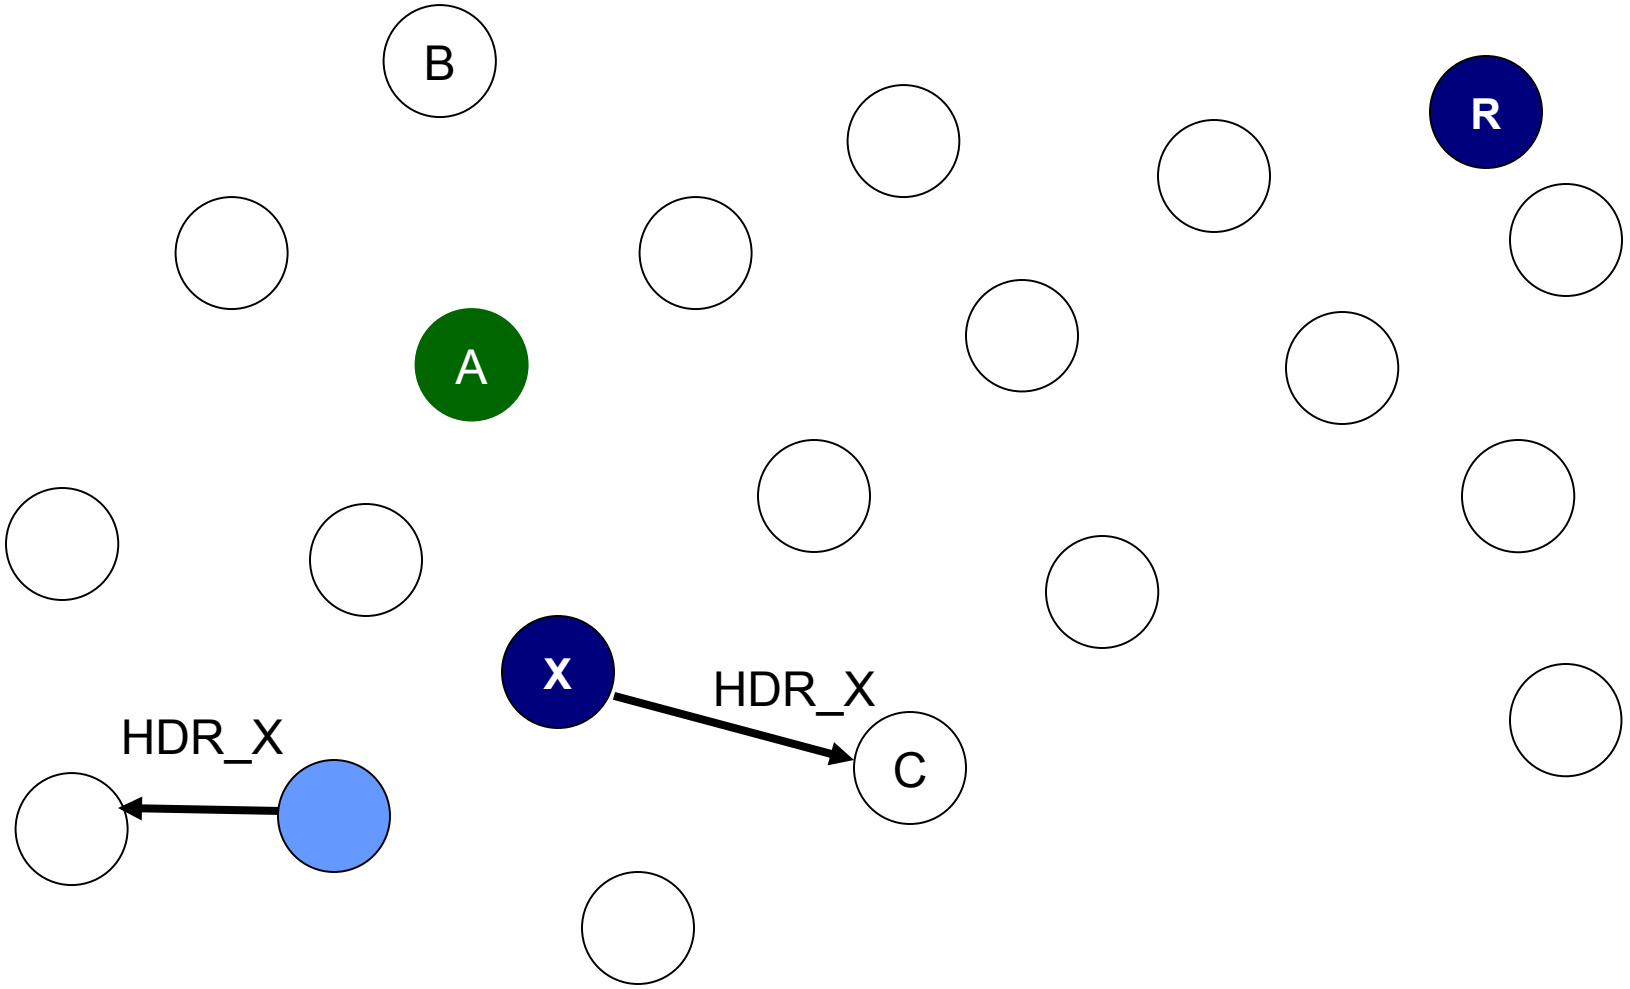

Anonymous Dissemination

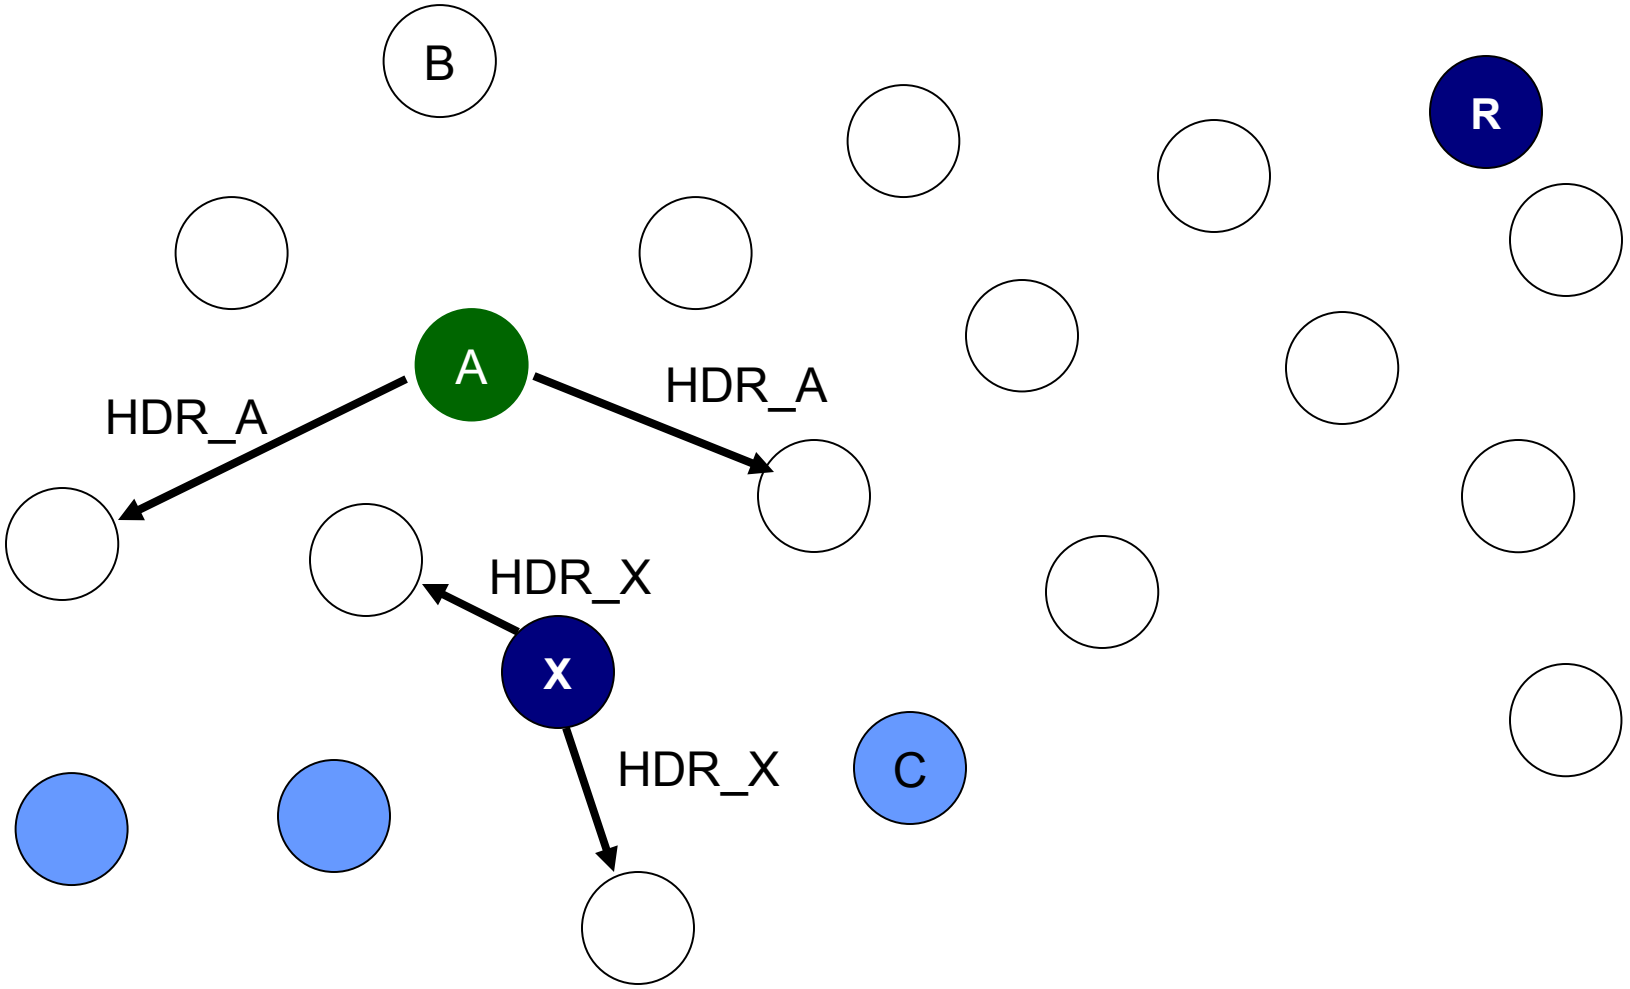

Anonymous Dissemination

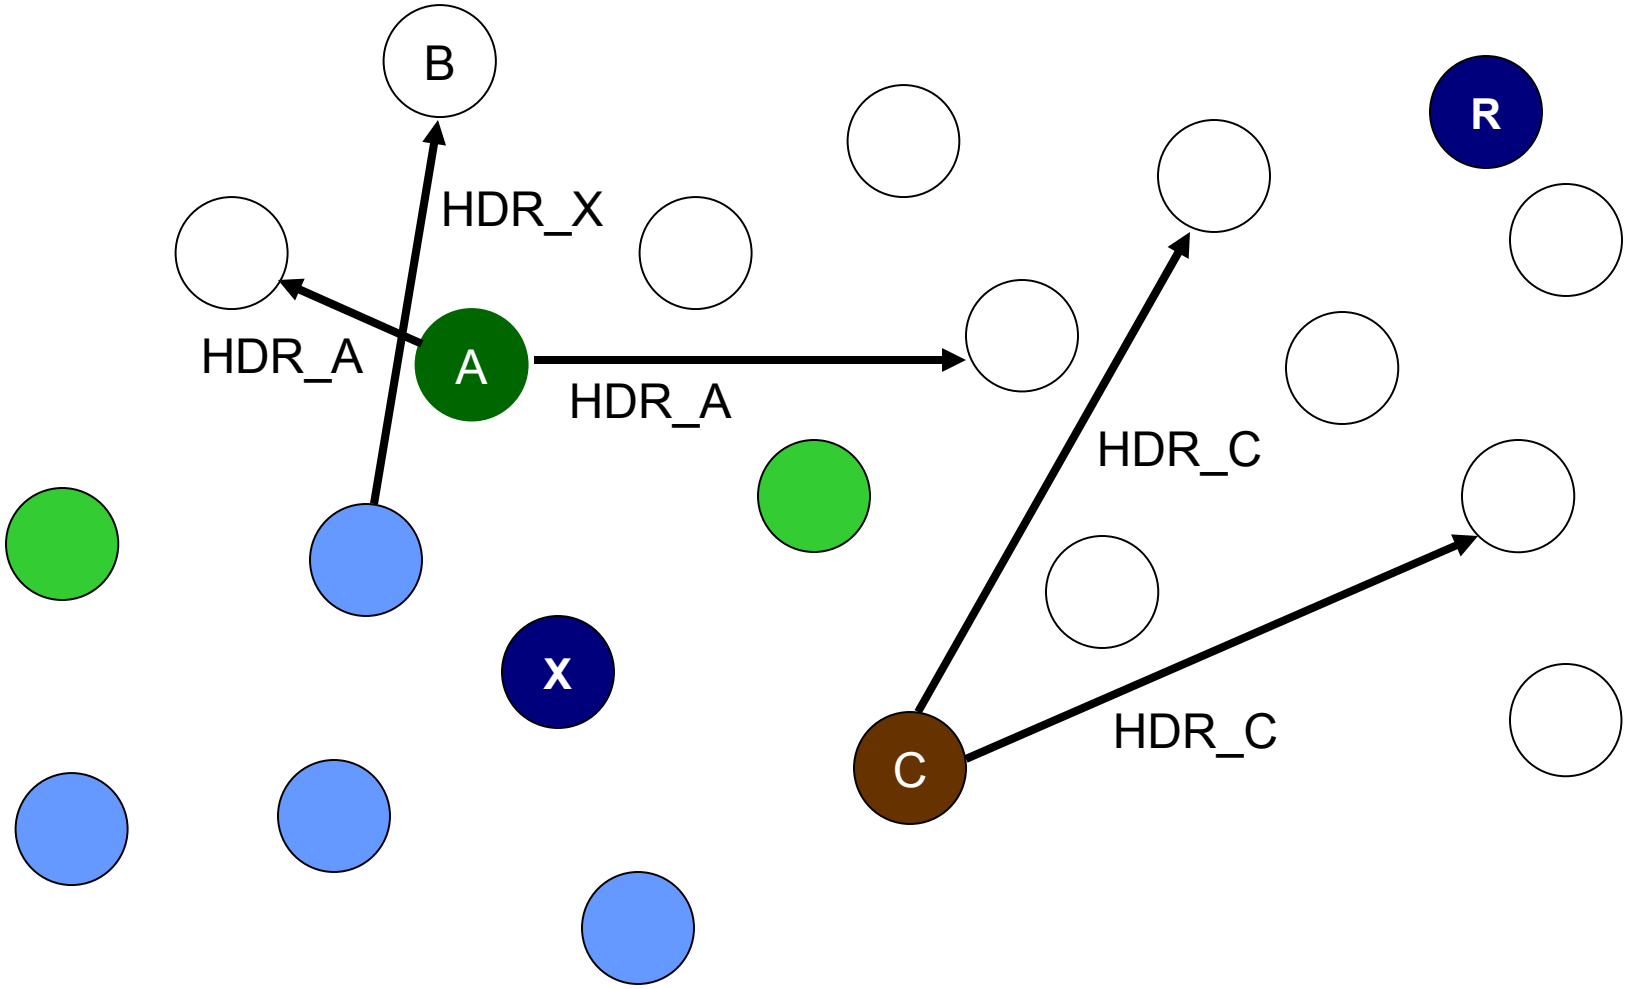

Anonymous Dissemination

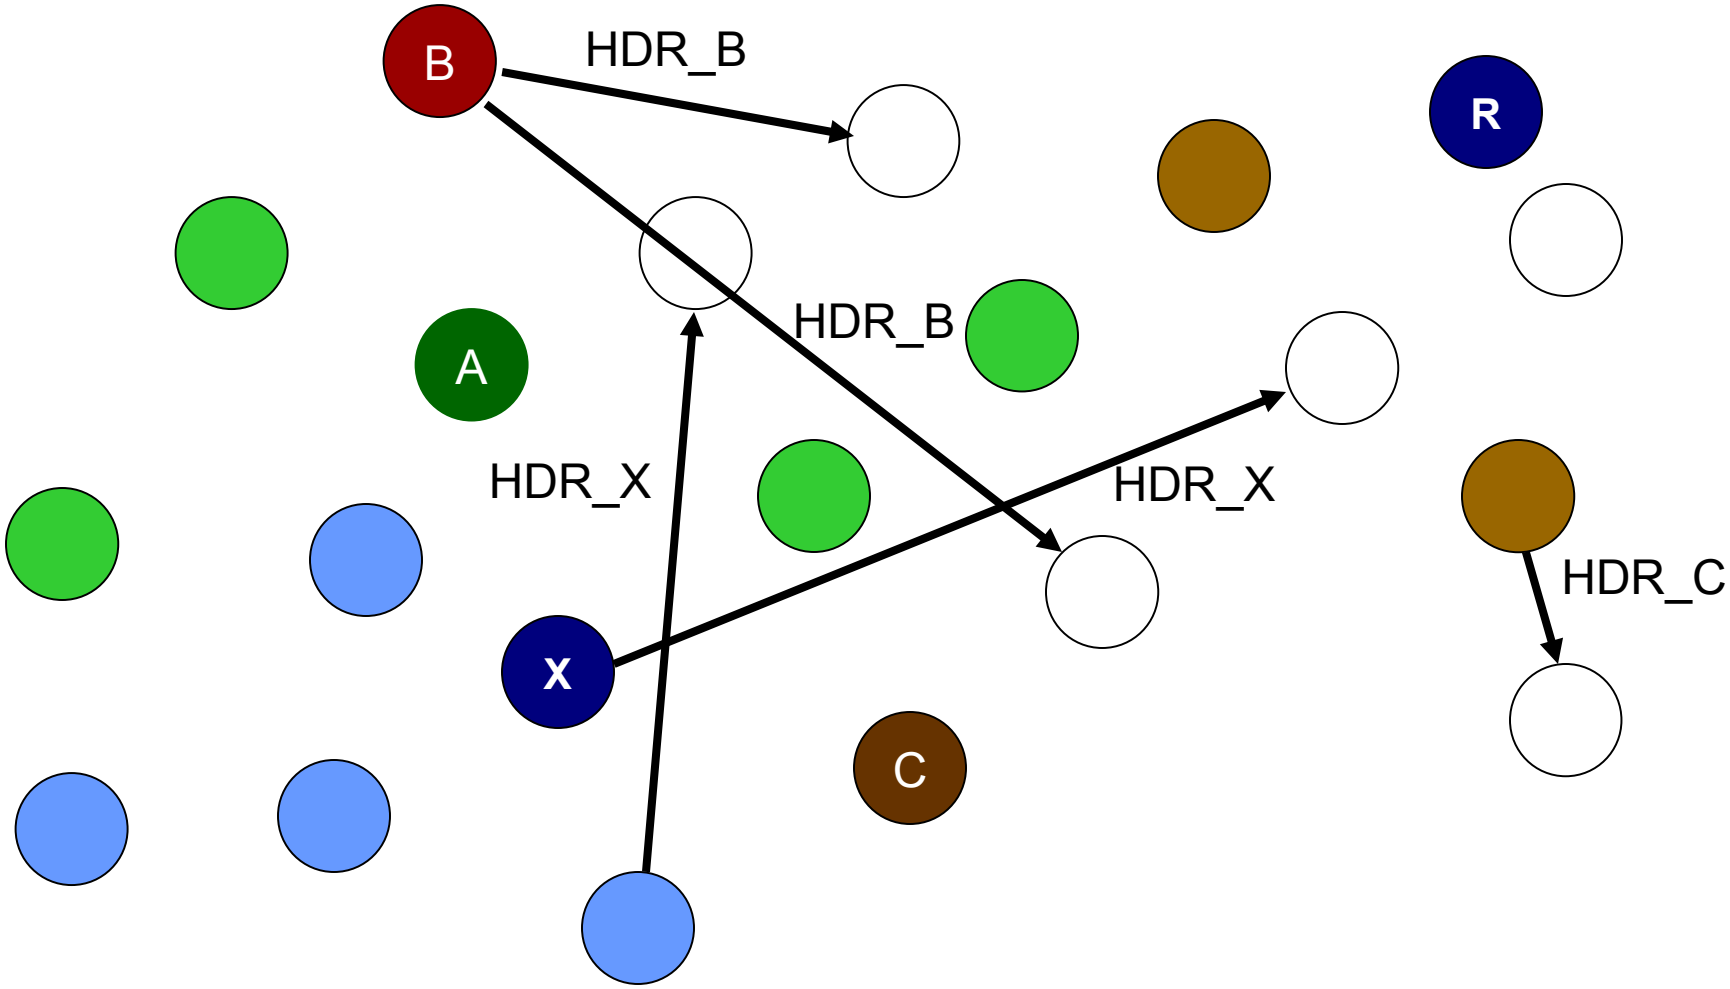

Anonymous Dissemination

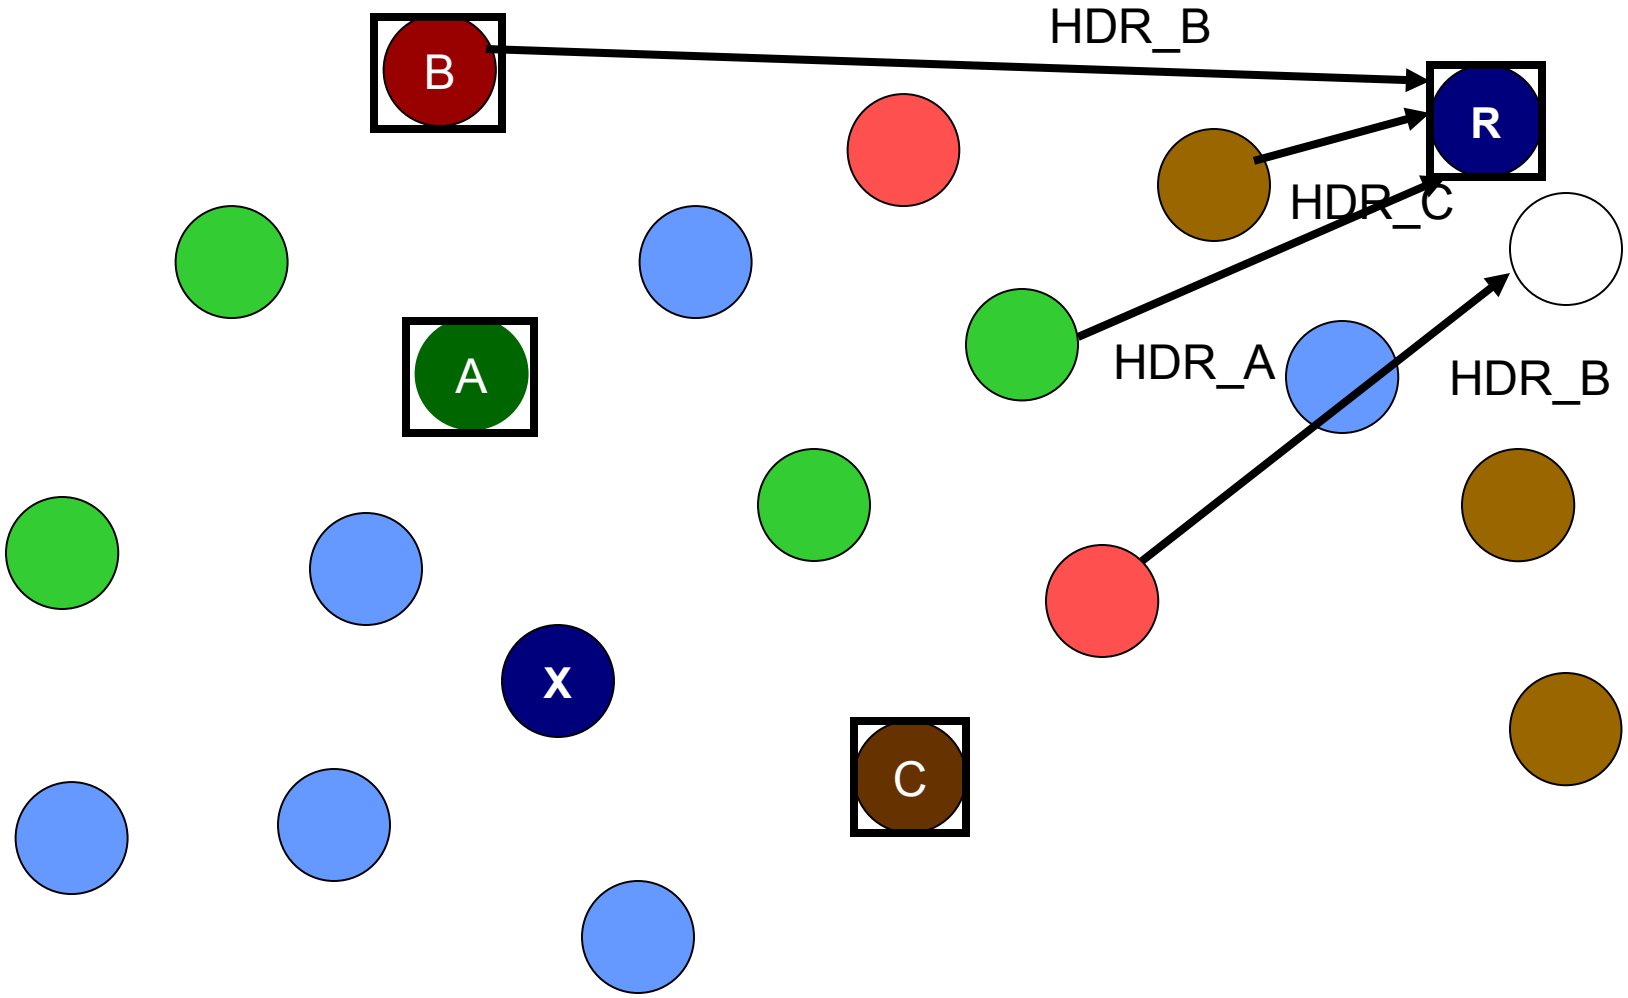

Performance Evaluation

- Study the protocol for large networks
  - Performance metrics should scale well
- Study protocol for dynamic networks
  - Performance metrics should be sustained for varying degrees of network disturbance

Reliability

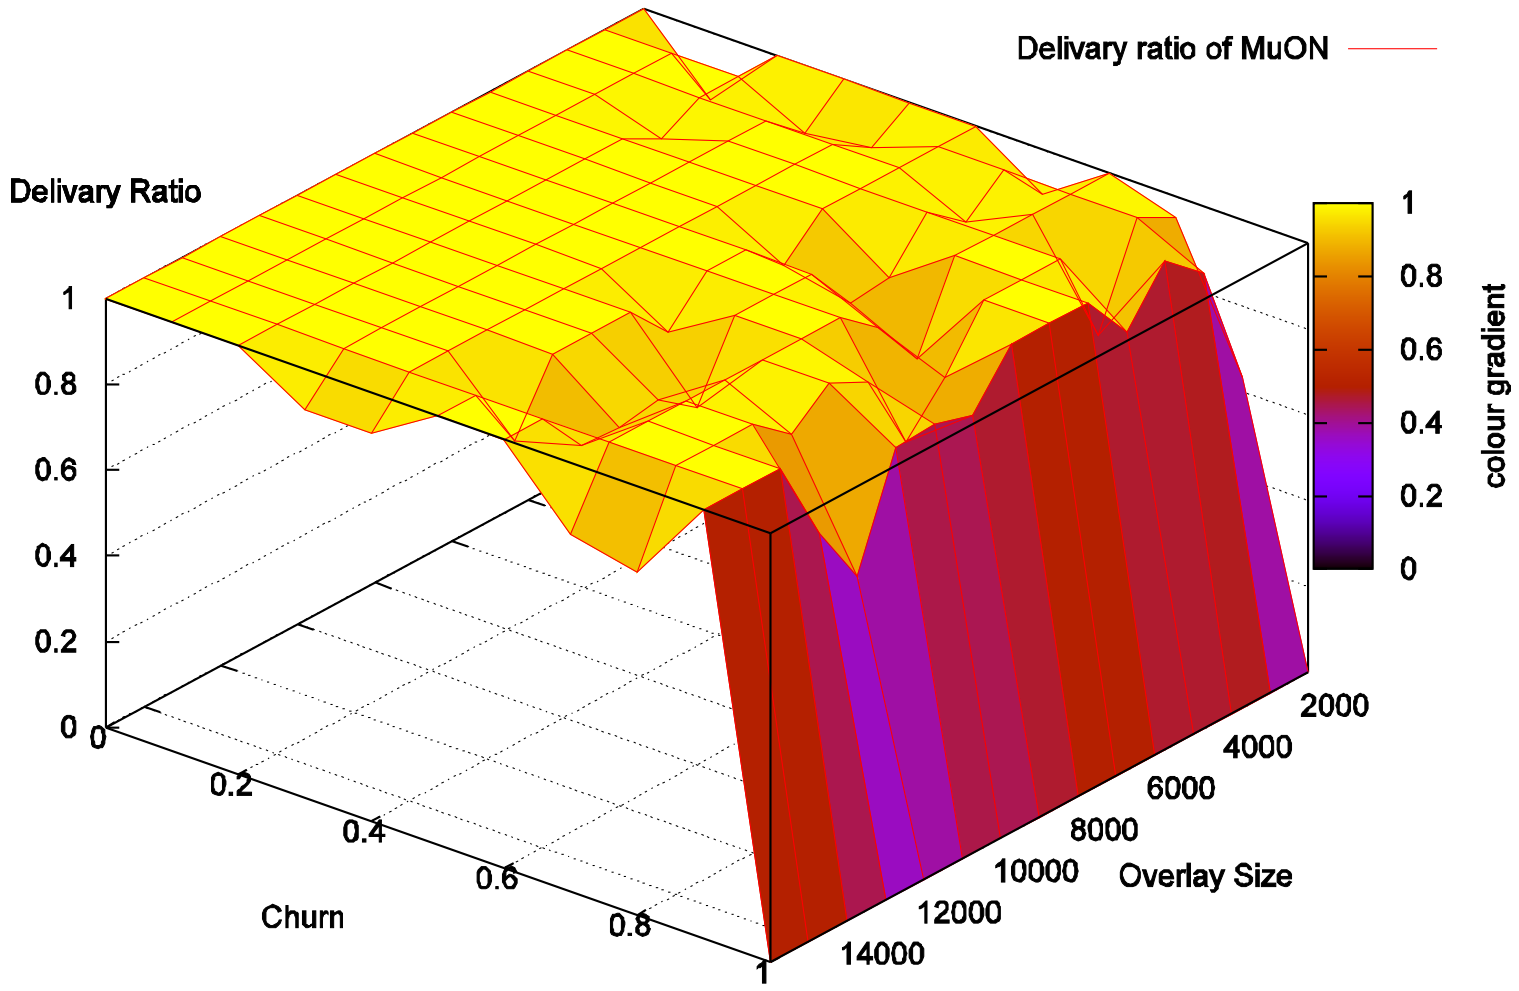

Latency

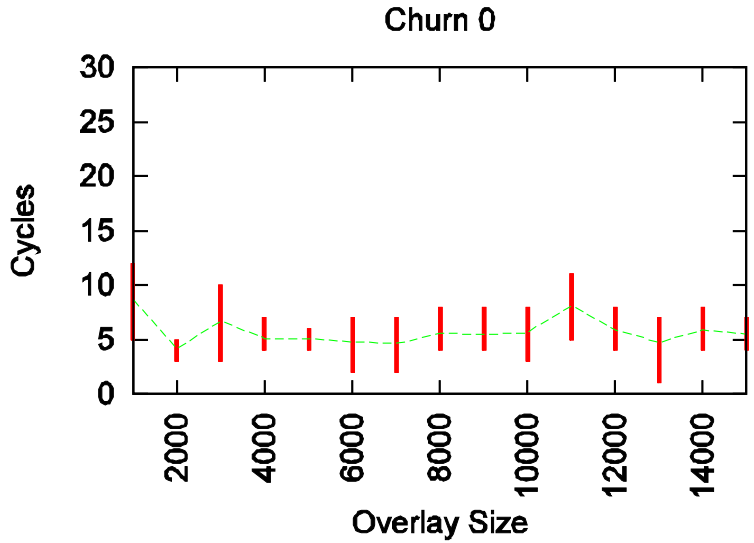
Churn 0
Cycles
30
25
20
15
10
5
0
2000
4000
6000
8000
10000
12000
14000
Overlay Size

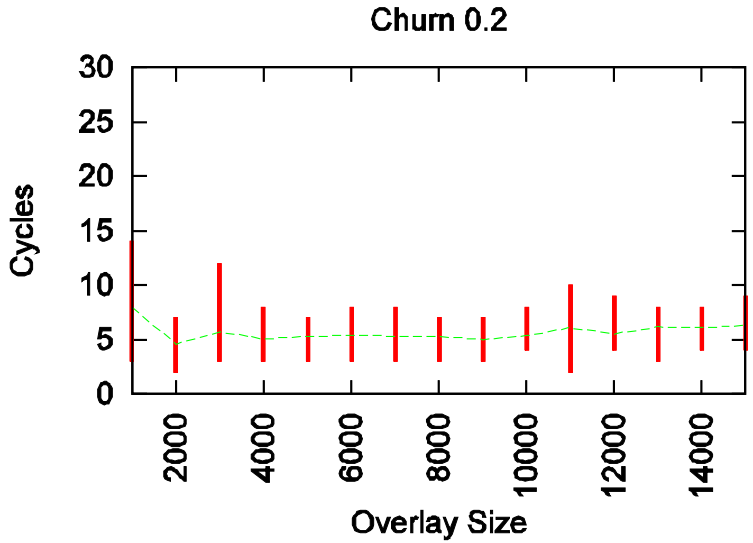
Churn 0.2
Cycles
30
25
20
15
10
5
0
2000
4000
6000
8000
10000
12000
14000
Overlay Size

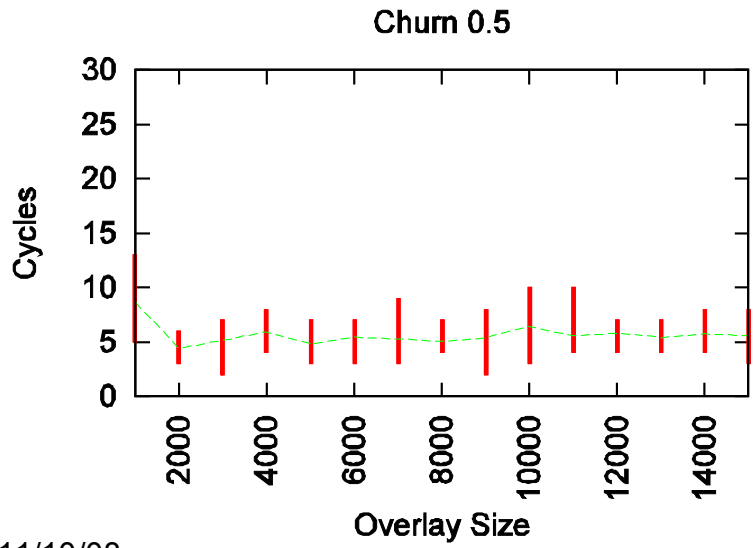
Churn 0.5
Cycles
30
25
20
15
10
5
0
2000
4000
6000
8000
10000
12000
14000
Overlay Size

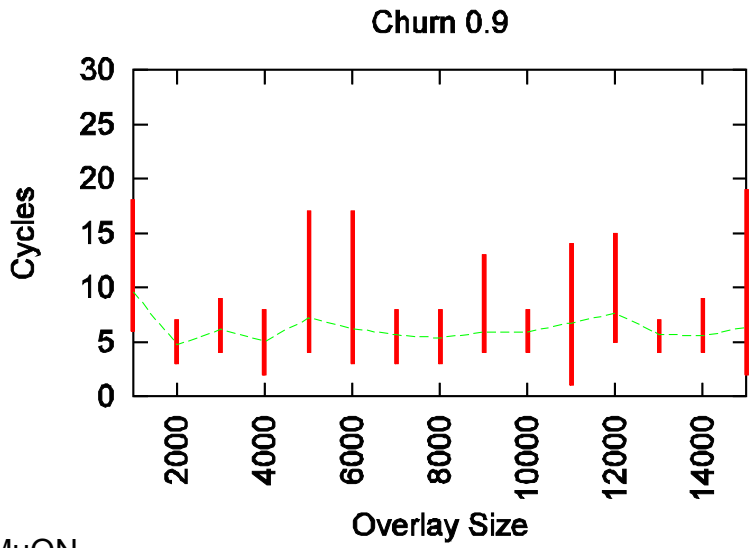
Churn 0.9
Cycles
30
25
20
15
10
5
0
2000
4000
6000
8000
10000
12000
14000
Overlay Size

Bandwidth Consumption

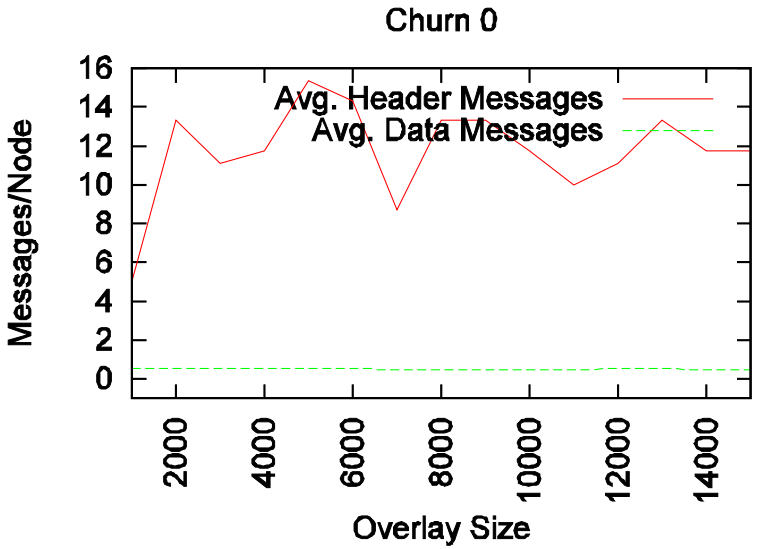

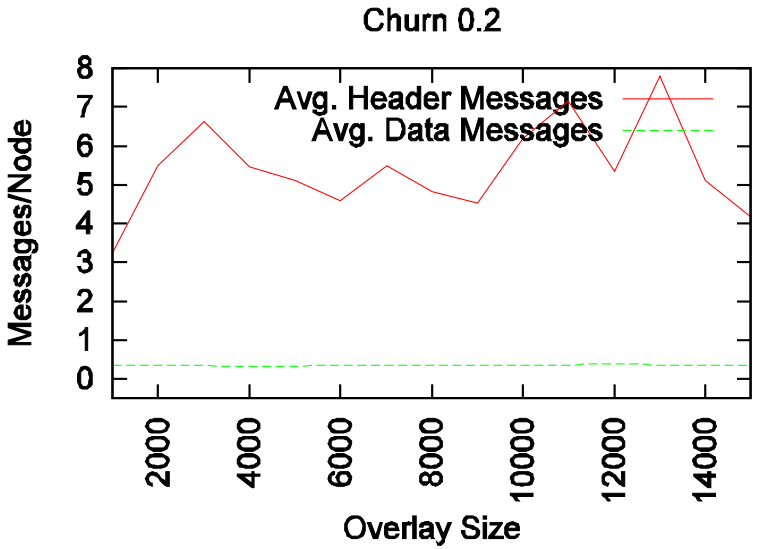

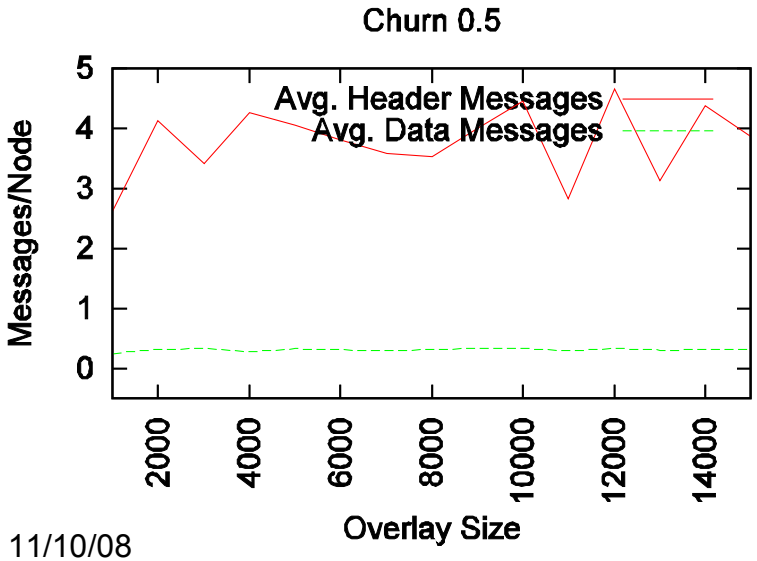

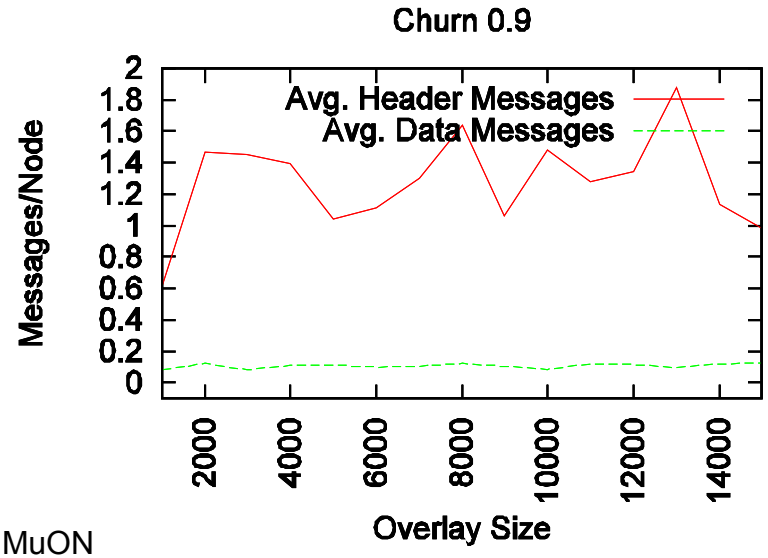

MuON: Message format

- Sending message from I to S
  - MSG=$\{r_1, id, data\}k_{session}$
    - Data encrypted with $k_{session}$ for confidentiality
    - Contains nonce r_1 and identifier for integrity
  - Profile (hdr)=$\{r_1, k_{session}, k_I^+, \{H(D)\} k_I^-\} k_S^+$
    - D=$\{r_1, k_{session}, k_I^+, MSG\}$
  - Message identifier is H(hdr)

Anonymity in MuON

P_{inter}= 0.1

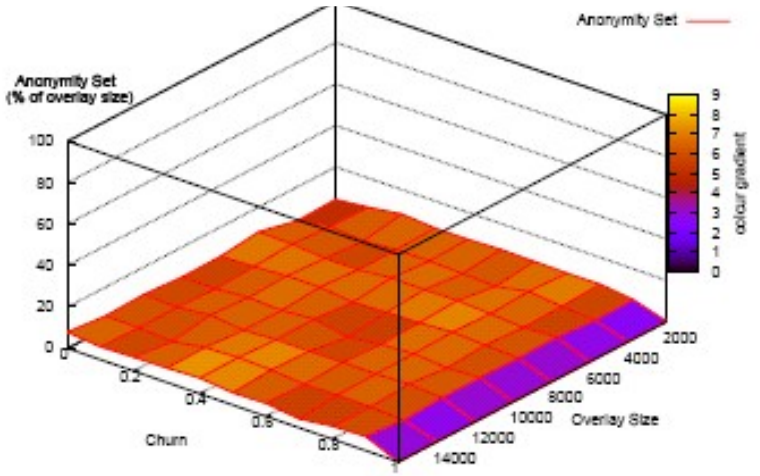

P_{inter}= 0.3

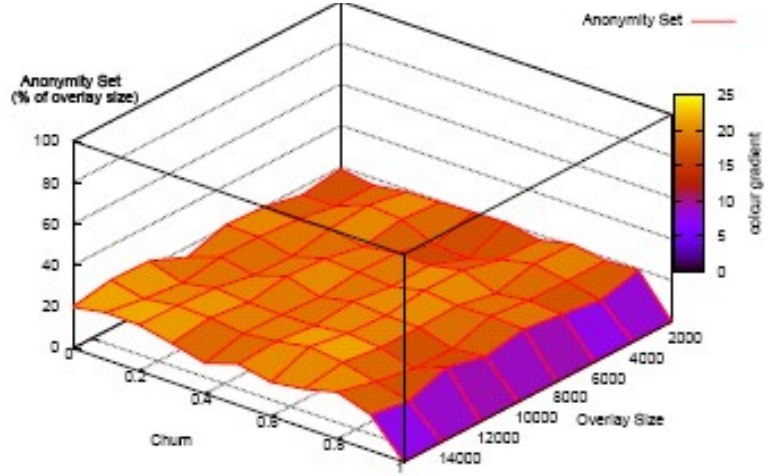

P_{inter}= 0.5

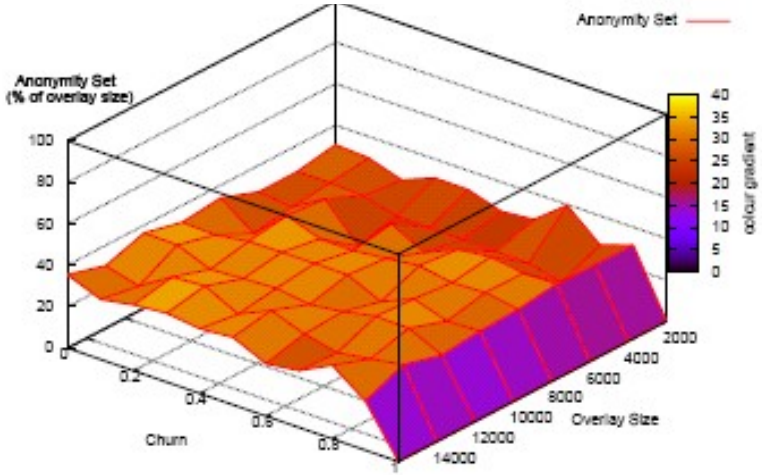

P_{inter}= 0.8

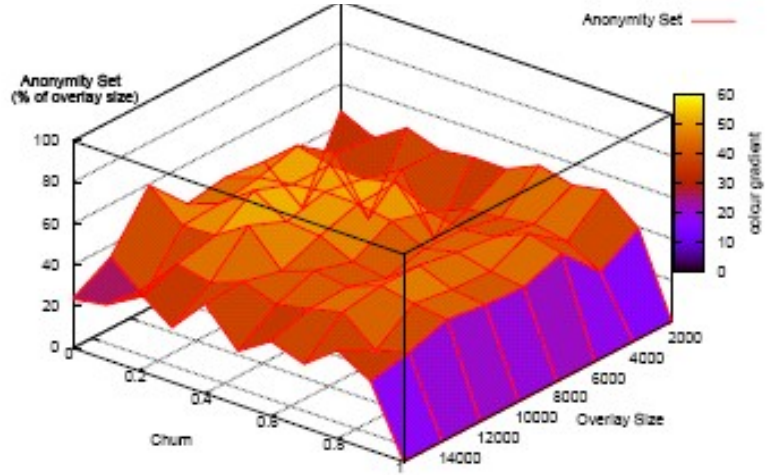

$FanOut = lg_2(overlay\_size)$, $GC = lg_2(overlay\_size)$, UDP losses = 10%

Anonymity in MuON

Churn 0

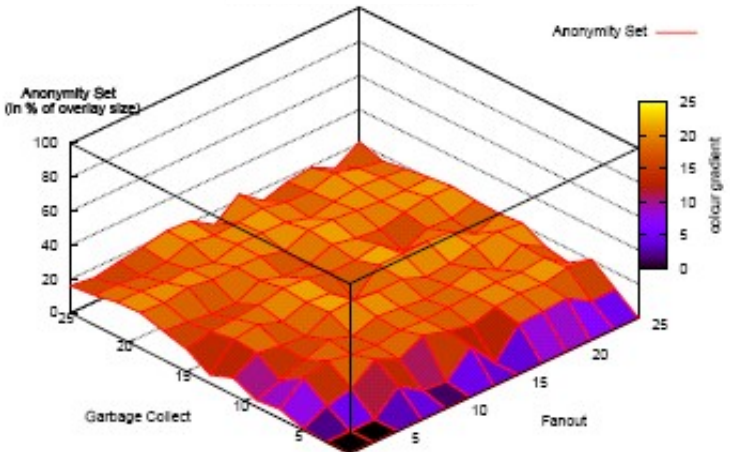

Churn 0.3

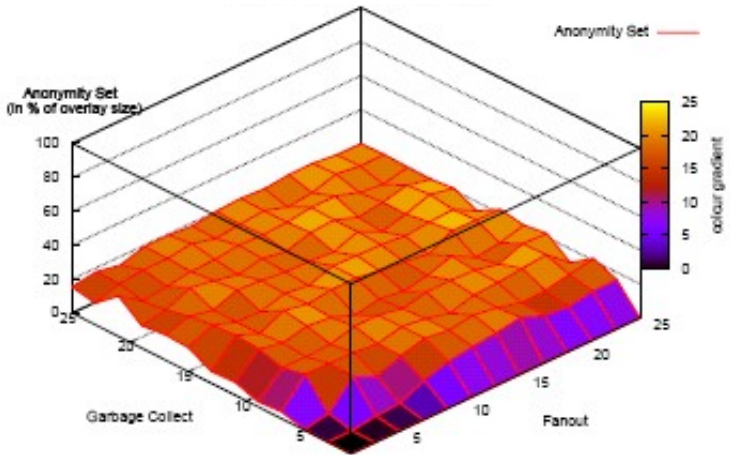

Churn 0.5

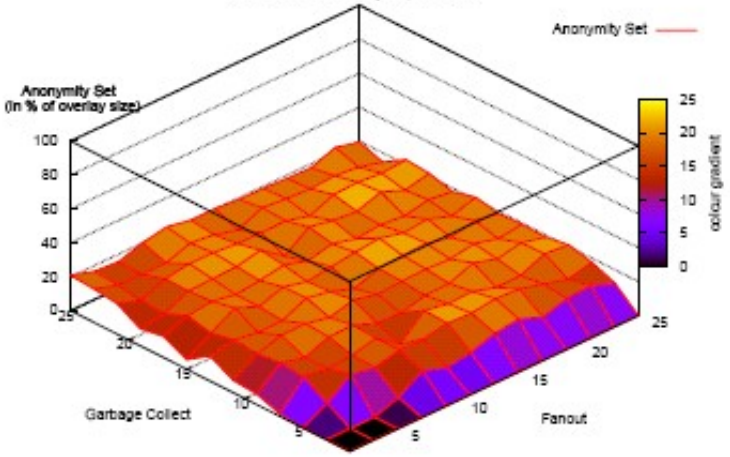

Churn 0.9

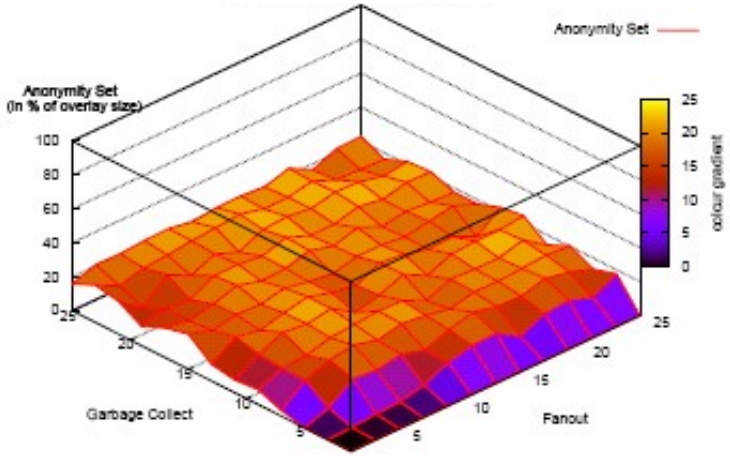

$overlay\_size = 10000$, UDP losses $= 10\%$, $p_{inter} = 0.3$,

Anonymity analysis

Local eavesdropper	Withstands attack
Collusion attack	Withstands unless all nodes are malicious
Timing attack	Withstands unless global adversary
Traceback attack	Withstands unless global adversary
Predecessor attack	Withstands attack
Intersection attack	Withstands unless global adversary
Message volume attack	Withstands unless global adversary

Conclusion and future work

- Contributions of our protocol (MuON)
  - Searchable mutual anonymity
  - Low overheads and latency
  - Performance scales well with large dynamic overlay (P2P)

- Future work
  - Two assumptions of epidemic protocols
    - Random selection of log(N), value of “N”
  - Multi-group anonymity system
    - Intra-group communication protocol, efficient and unobservable
    - Inter-group routing and forwarding protocols, efficient and anonymous
  - Anonymity vs. anonymity-breaking technologies
  - Implementation and testbed
    - Planetlab
    - MTU and any voluntary participants
  - Social networks?

Header overhead

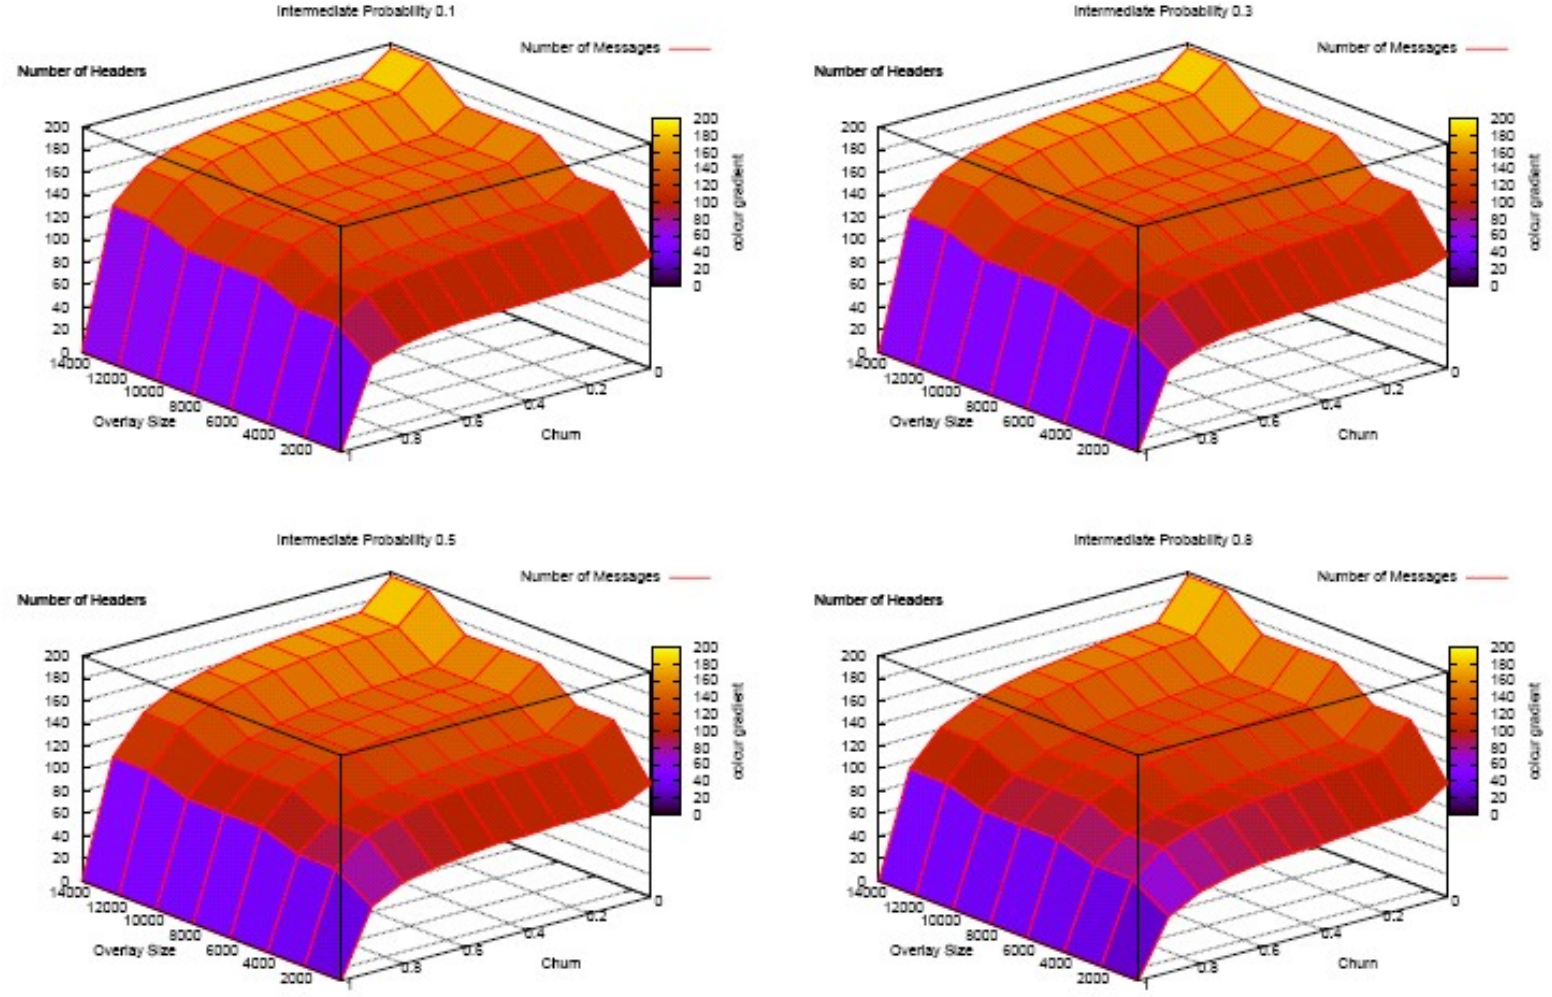

$FanOut = lg_2(overlay\_size)$, $GC = lg_2(overlay\_size)$, UDP losses $= 10\%$, $p_{inter} = 0.3$

Header overhead

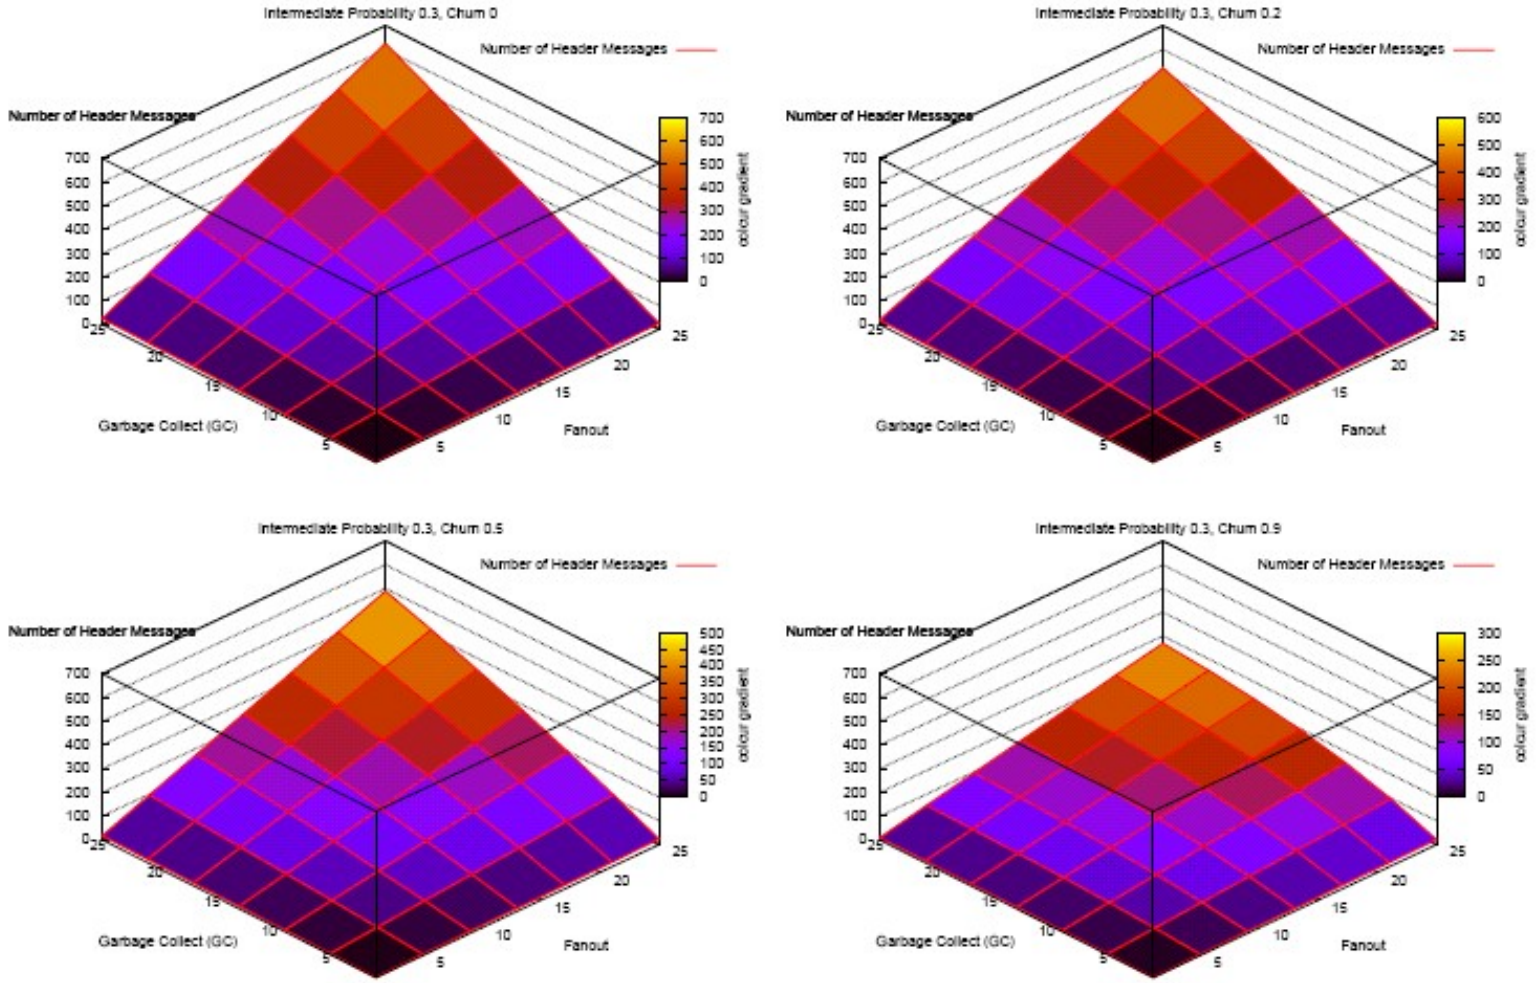

$overlay\_size = 10000$, UDP losses $= 10\%$, $p_{inter} = 0.3$

Data overhead

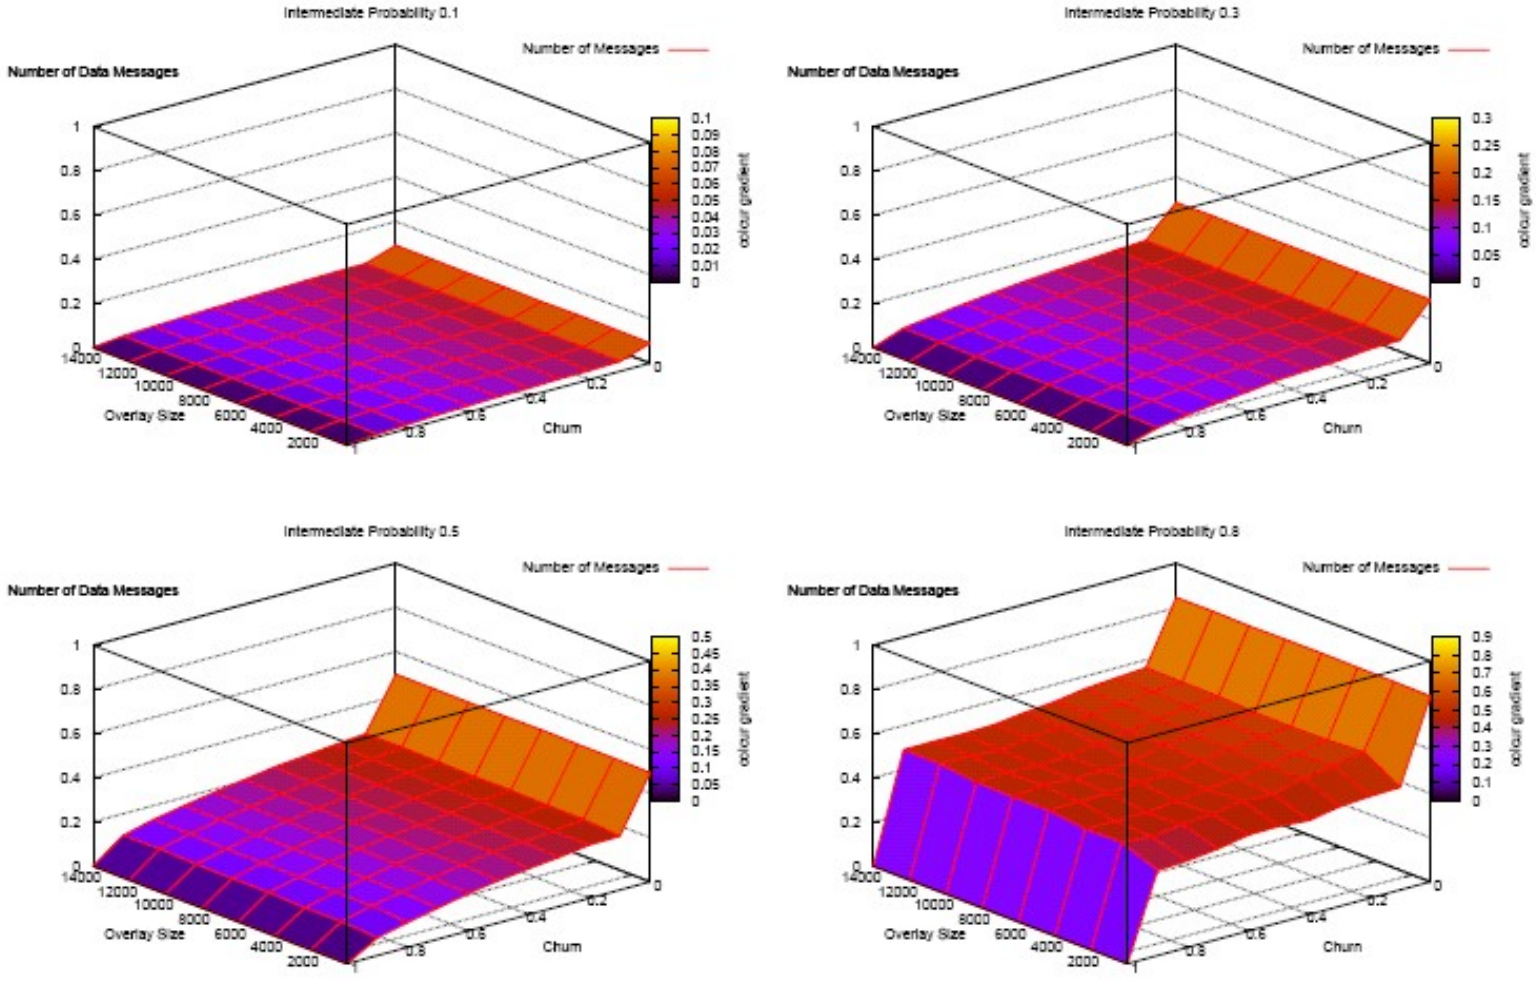

$FanOut = lg_2(overlay\_size)$, $GC = lg_2(overlay\_size)$, UDP losses $= 10\%$, $p_{inter} = 0.3$

Data overhead

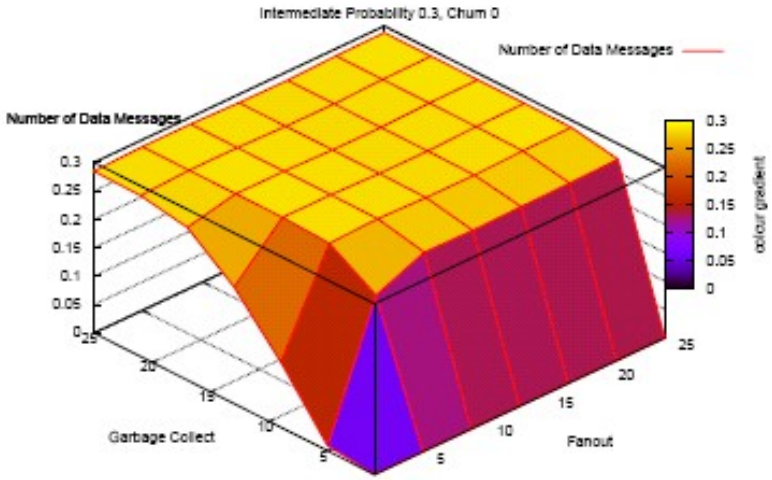

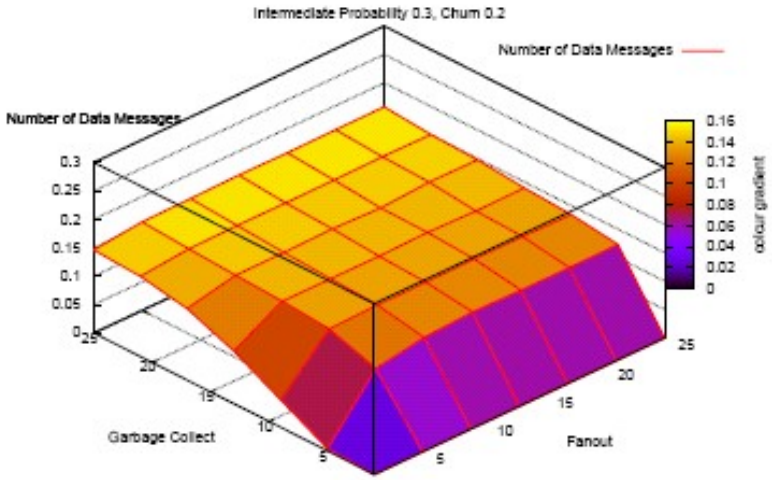

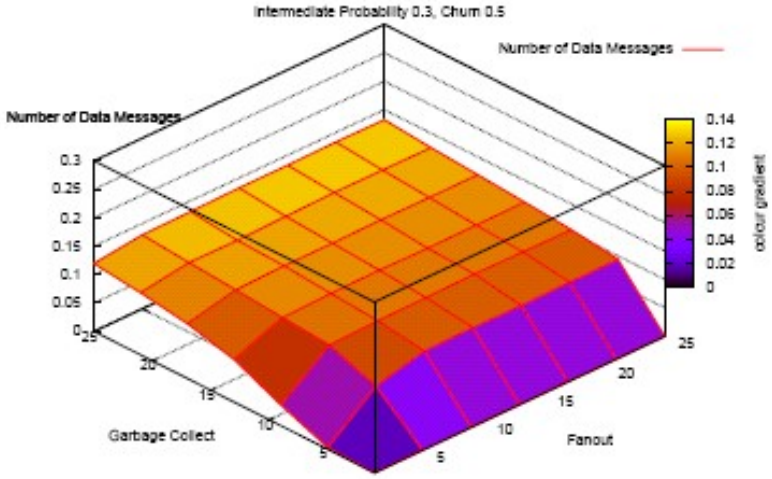

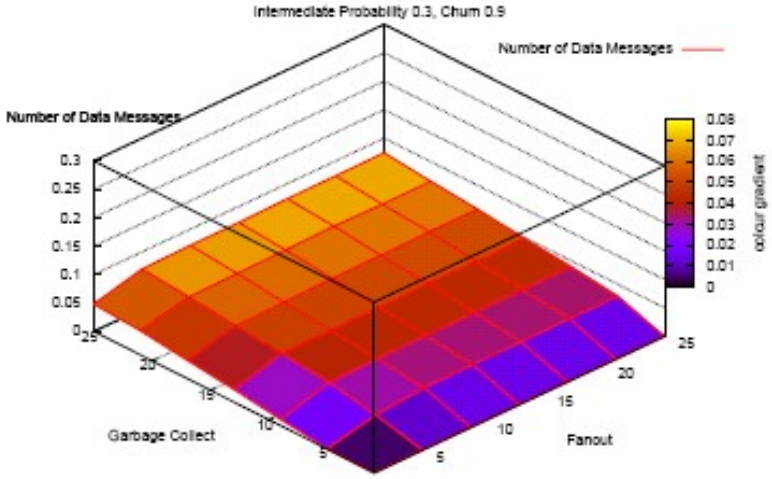

$overlay\_size = 10000$, UDP losses $= 10\%$, $p_{inter} = 0.3$