The Dynamic Behavior of a Data Dissemination Protocol for Network Programming at Scale

Jonathan W. Hui, David Culler

Presented by:

Thomas Repantis

`trep@cs.ucr.edu`

Overview

Deluge, a reliable data dissemination protocol for propagating large data objects from one or more source nodes to many other nodes over a multihop, wireless sensor network.

1. Introduction
2. Protocol Description
3. Experimental Evaluation
4. Conclusions
5. Comparison with P2P

Network Programming

- Deluge's goal is the propagation of complete binary images to the nodes of the network.
- Network programming (the programming of nodes by disseminating code over the network) is useful:
  - Application requirements can change.
  - Thousands of nodes embedded.
  - Also useful for debugging and testing.

Requirements

Issues in reliable dissemination of large data objects from few to many nodes of a multi-hop WSN:

- Large data objects vs constrained storage hierarchy: Packet (32 bytes) « RAM (4K) « program (128K) < external flash (512K)
- High node density (hinders transmission).
- Complete reliability required, even with highly lossy links.
- All nodes need to receive the newest code version.
- Fast propagation is desired.

Deluge Protocol Overview

Nodes follow strictly local rules:

- Advertise
- Request data
- Broadcast

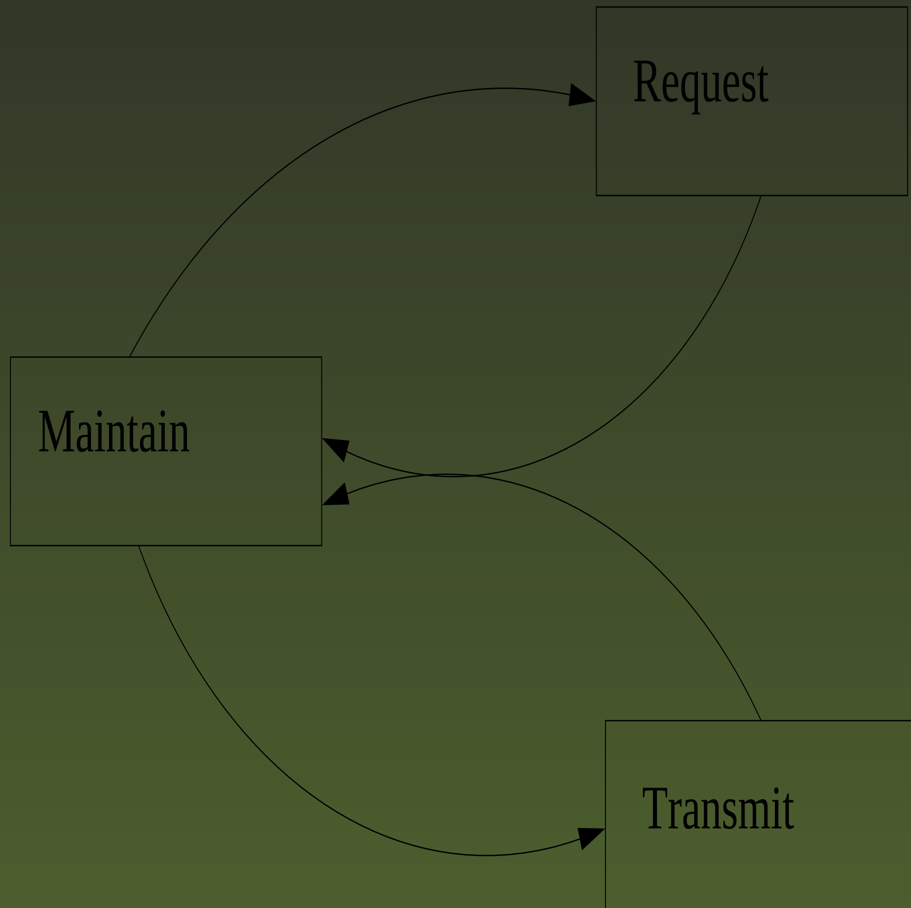

Epidemic Protocol

- Each node occasionally advertises (by broadcasting) the most recent version of the data object it has available.
- If a node receives an advertisement from an older node, it responds with its object profile.
- The older node determines from the object profile which portions of its data need updating and requests them from any neighbor that advertises the availability of the needed data.
- Nodes receiving the request then broadcast any requested data.
- Nodes then advertise newly received data in order to propagate it further.

Protocol Properties

- Density-aware (redundant advertisements and requests are suppressed).
- Robust to asymmetric links (three-phase handshaking).
- Switches between quick propagation and quite phases (dynamically adjustable advertisement rate).
- Minimizes the set of nodes concurrently broadcasting data within a cell.
- Allows parallel transfers of data (spatial multiplexing).

Data Representation 1/2

CRCs at page and packet level ensure correct transmission.

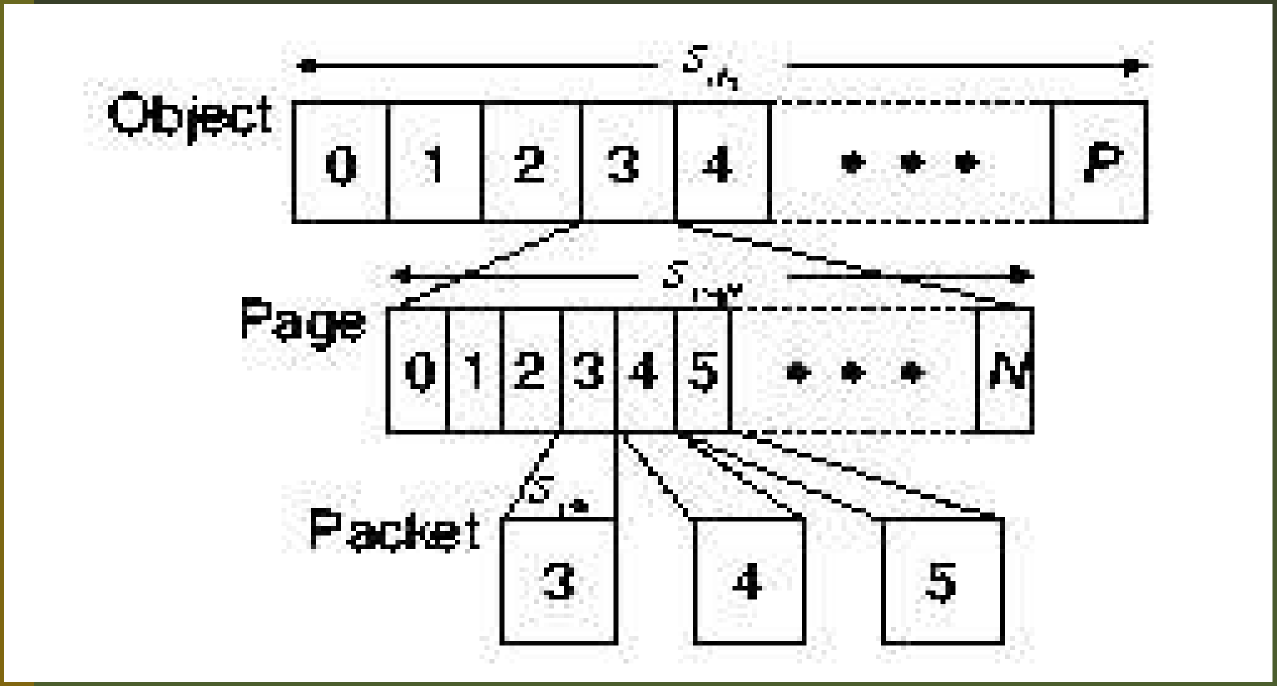

Data Representation 2/2

Advantages of fixed-size pages:

- Reduced memory requirements for maintaining state about which packets are needed.
- Enables efficient incremental updates using:
  - A monotonically increasing version number v for every update.
  - An age vector a describes how old a page is.
  - An object profile is defined by the tuple (v,a).
- Allows for spatial multiplexing.

Spatial Multiplexing

- Propagation in “waves” (pipelining the page transfers).
- Exploits the limited range of radio to allow for concurrent broadcasts.
- Reduces the completion time from o(d*Sobj) to o(d+Sobj).

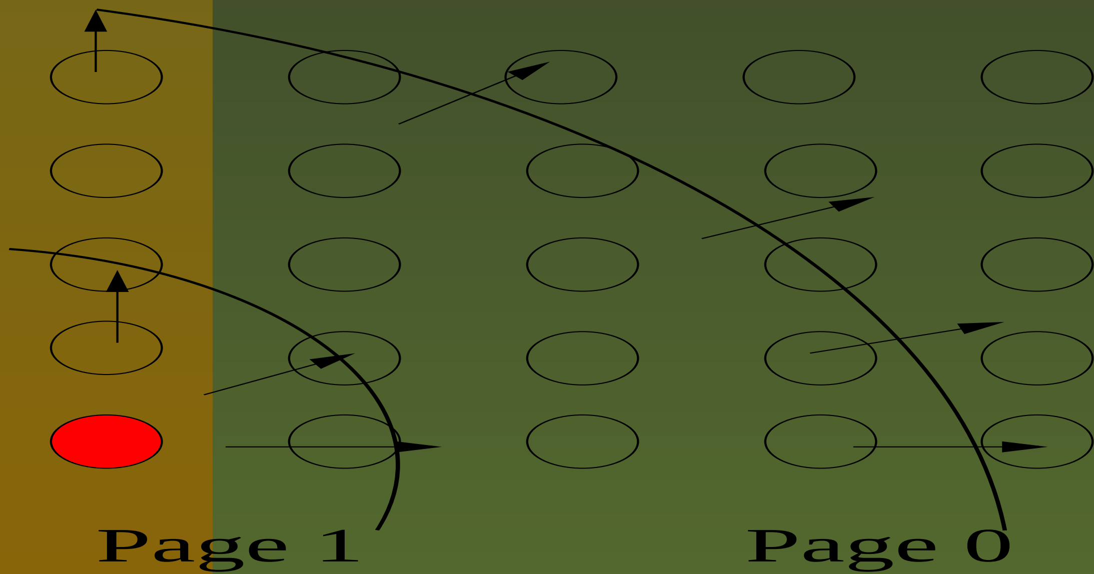

Maintain State

- Advertise version and fraction of image that is complete (version number and highest page available, since nodes request pages sequentially).
- Deluge uses Trickle to control the advertisement broadcasts.
- Uses suppression to decrease the advertisement rate as neighbors increase.
- Bounds advertisement rate independent of node density (by suppressing the transmission of redundant advertisements).
- Allows for quick propagation during an upgrade and low resource consumption in the steady state by decreasing and increasing the advertisement rate respectively.

Maintain State Transitions

- Transition to Transmit:
  When receiving a request.
- Transition to Request:
  When receiving an advertisement with newer data, unless a request or data packet was recently overheard.

Request State

- Transmit a request after random backoff time.
- Suppress if any similar requests of data packets are overheard during the backoff period.
- Minimize senders by unicasting requests to the node that advertised.
- Transition to maintain after receiving all packets of a page or after k requests (to protect against asymmetric lossy links).

Transmit State

- Transmit all requested packets.
- Packets are sent in round-robin order to provide fairness among requesters.
- Transition to maintain after all requested packets have been transmitted and no new packets have been requested.

Experimental Evaluation

- Real-world deployment (77 Mica2-dot nodes, in a building).
- Simulation (TOSSIM, up to 800 nodes, in a square grid).

Metrics:

- Complete Reliability
- Completion Time
- RAM Usage
- Energy Consumption

Real World Completion Time

Completion time is linear with the object size.

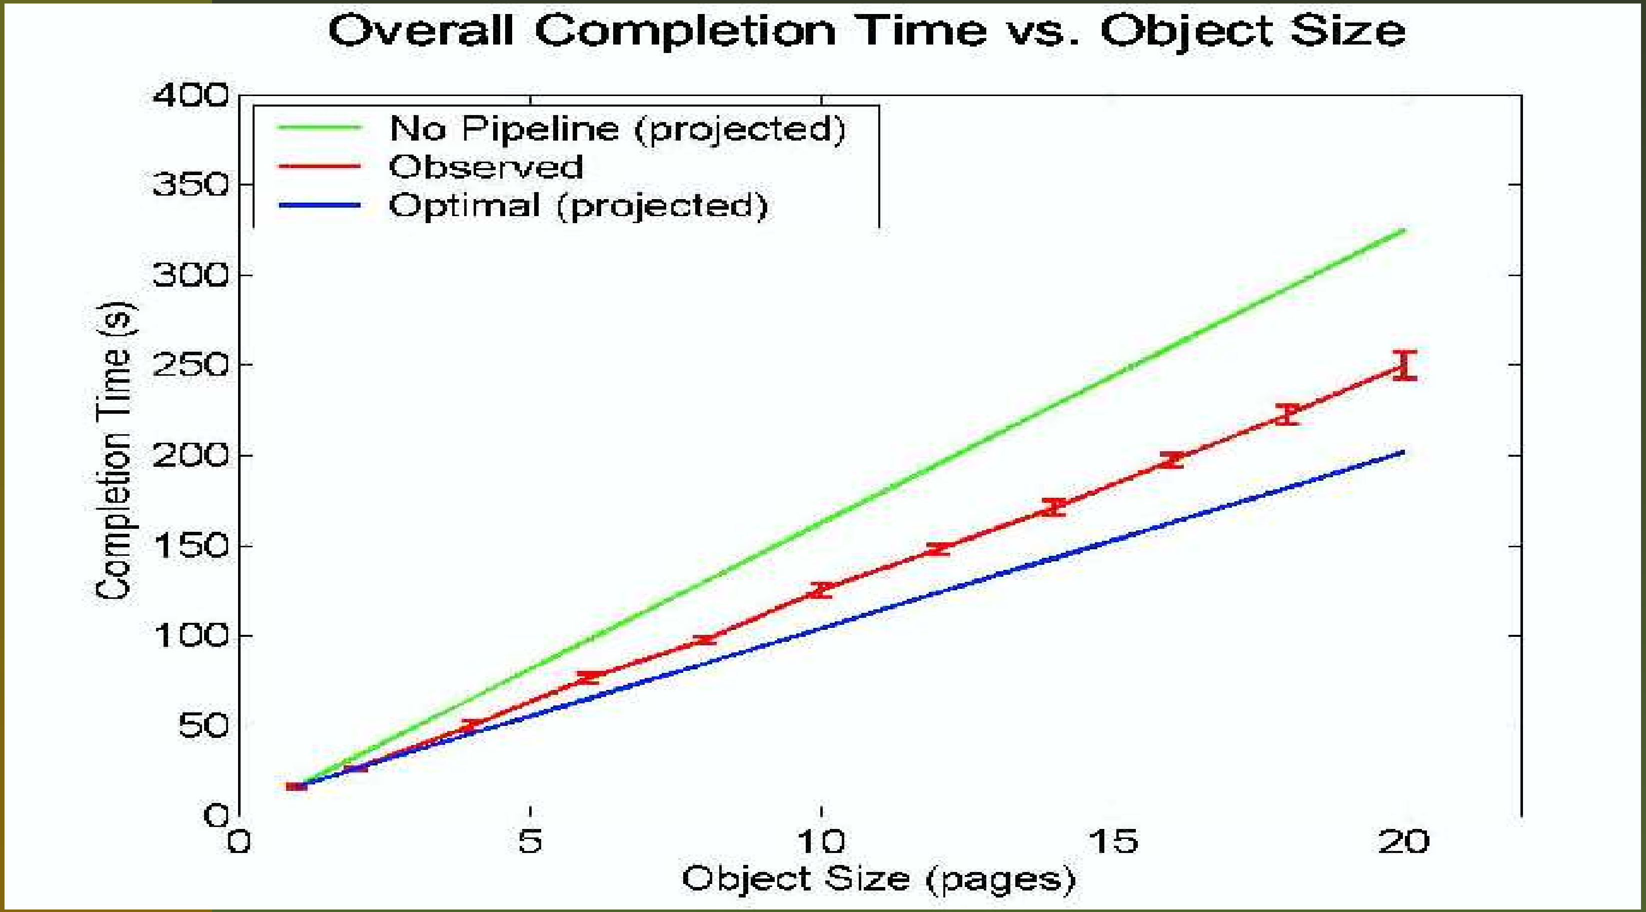

Real World Advertisement Rate

Without suppression it would have been ten times more.

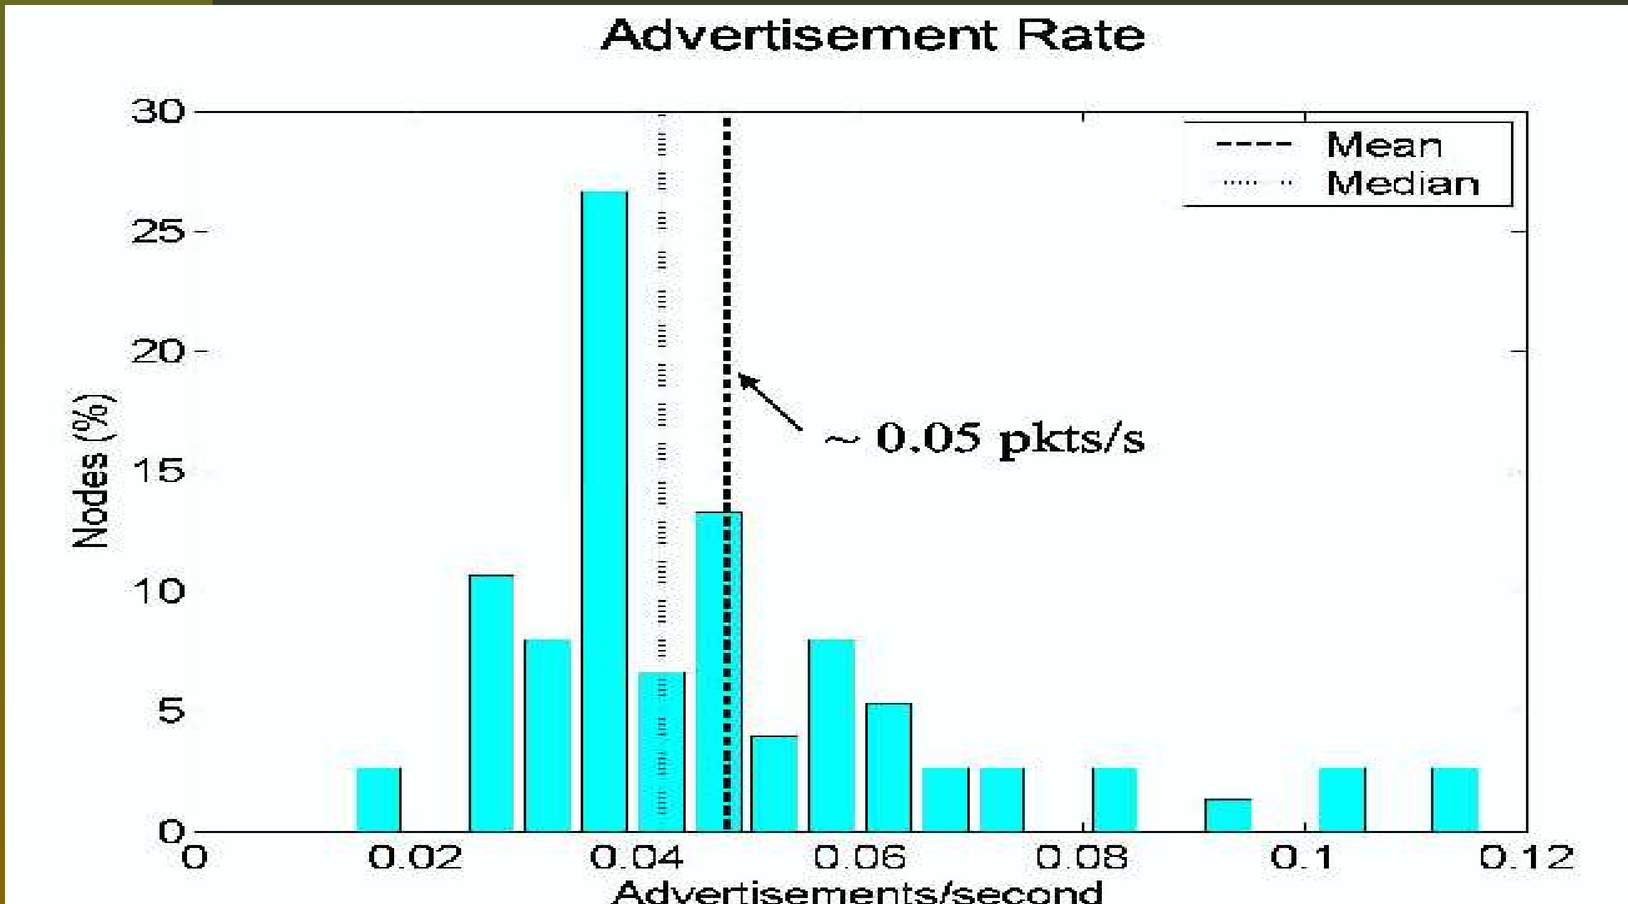

Real World Transmitted Requests

Many nodes took advantage of the broadcast channel and did not need to transmit requests to receive entire pages.

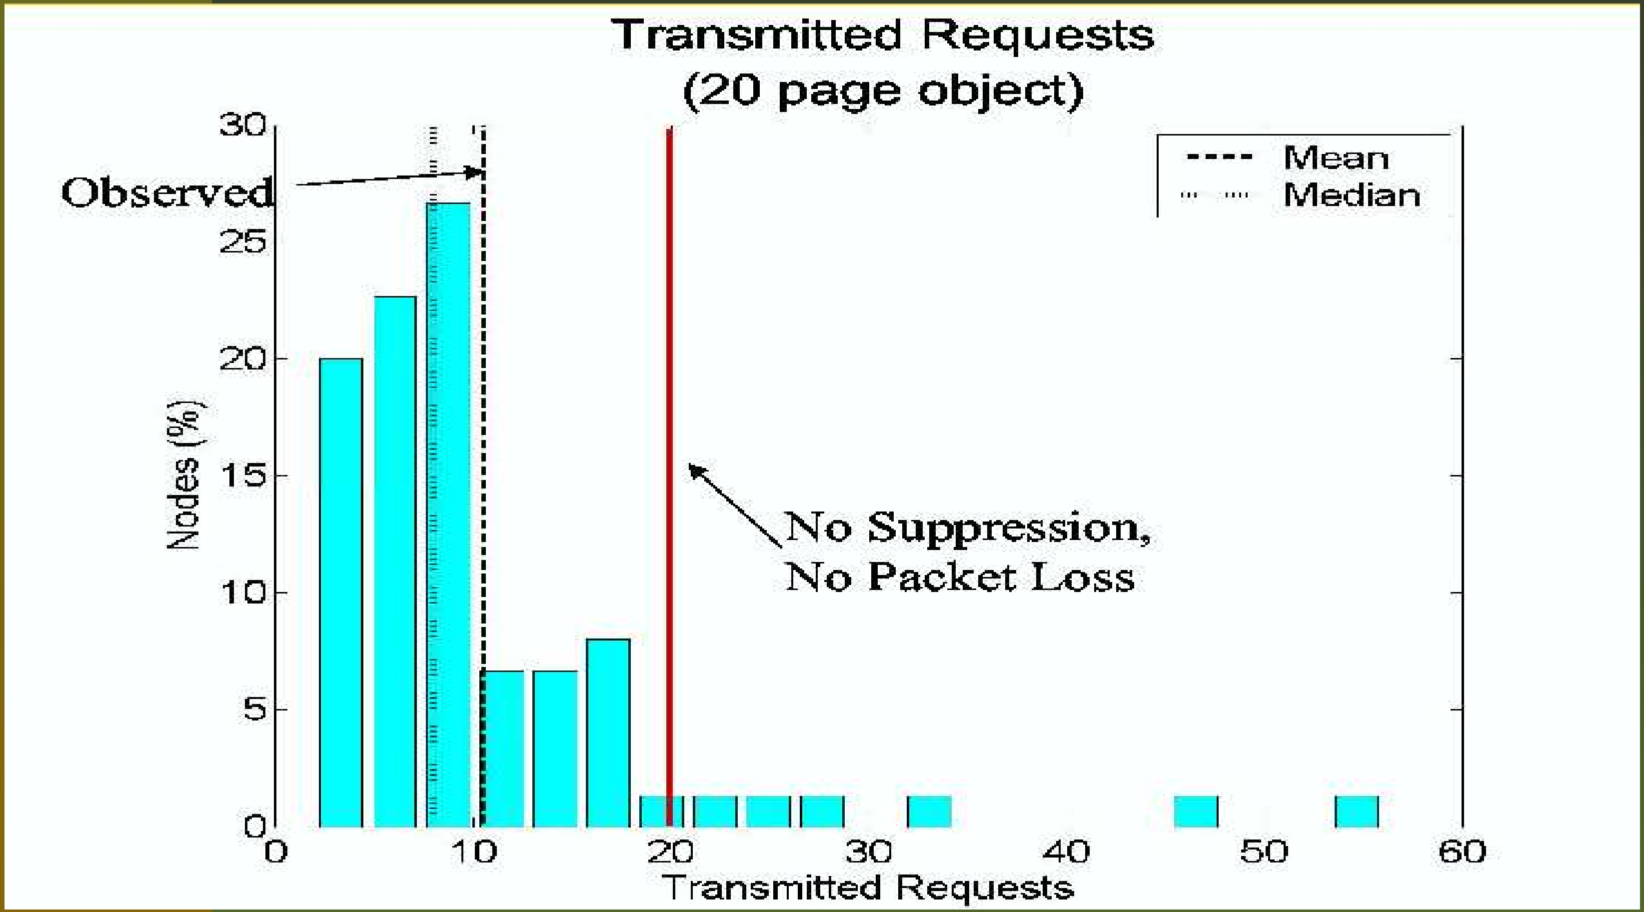

Real World Data Redundancy

On average a node receives about 3.35 times the minimum number of required data packets, due to the single-channel, broadcast network.

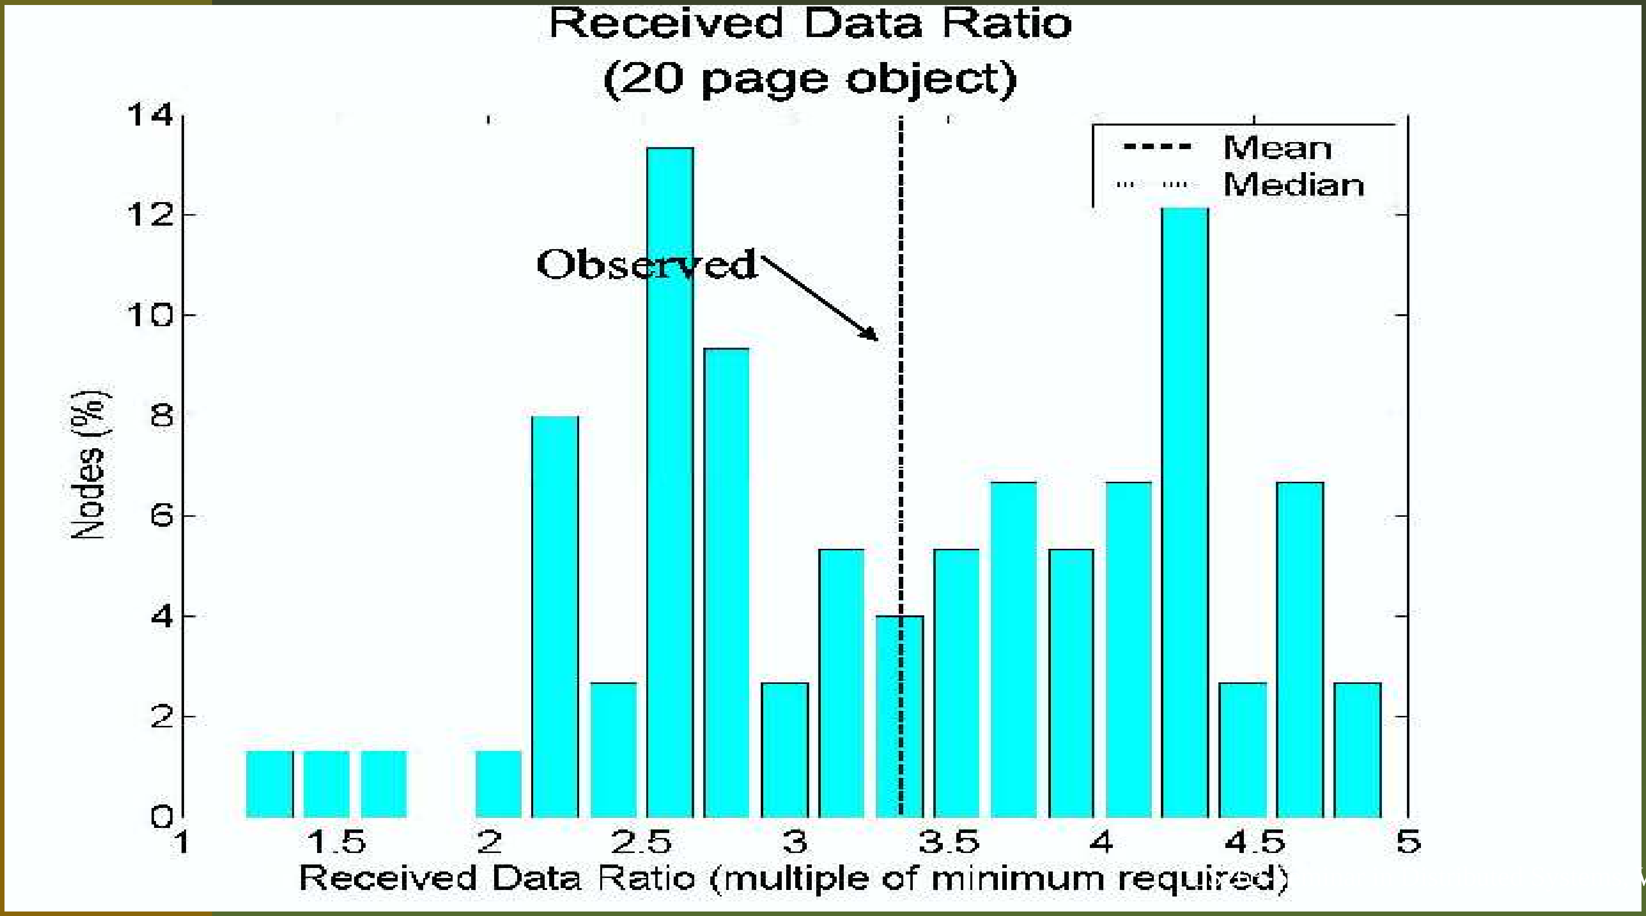

Simulated Completion Time for Varying Network Diameter

Spatial multiplexing is effective (propagation time is the sum of the network diameter and object size).

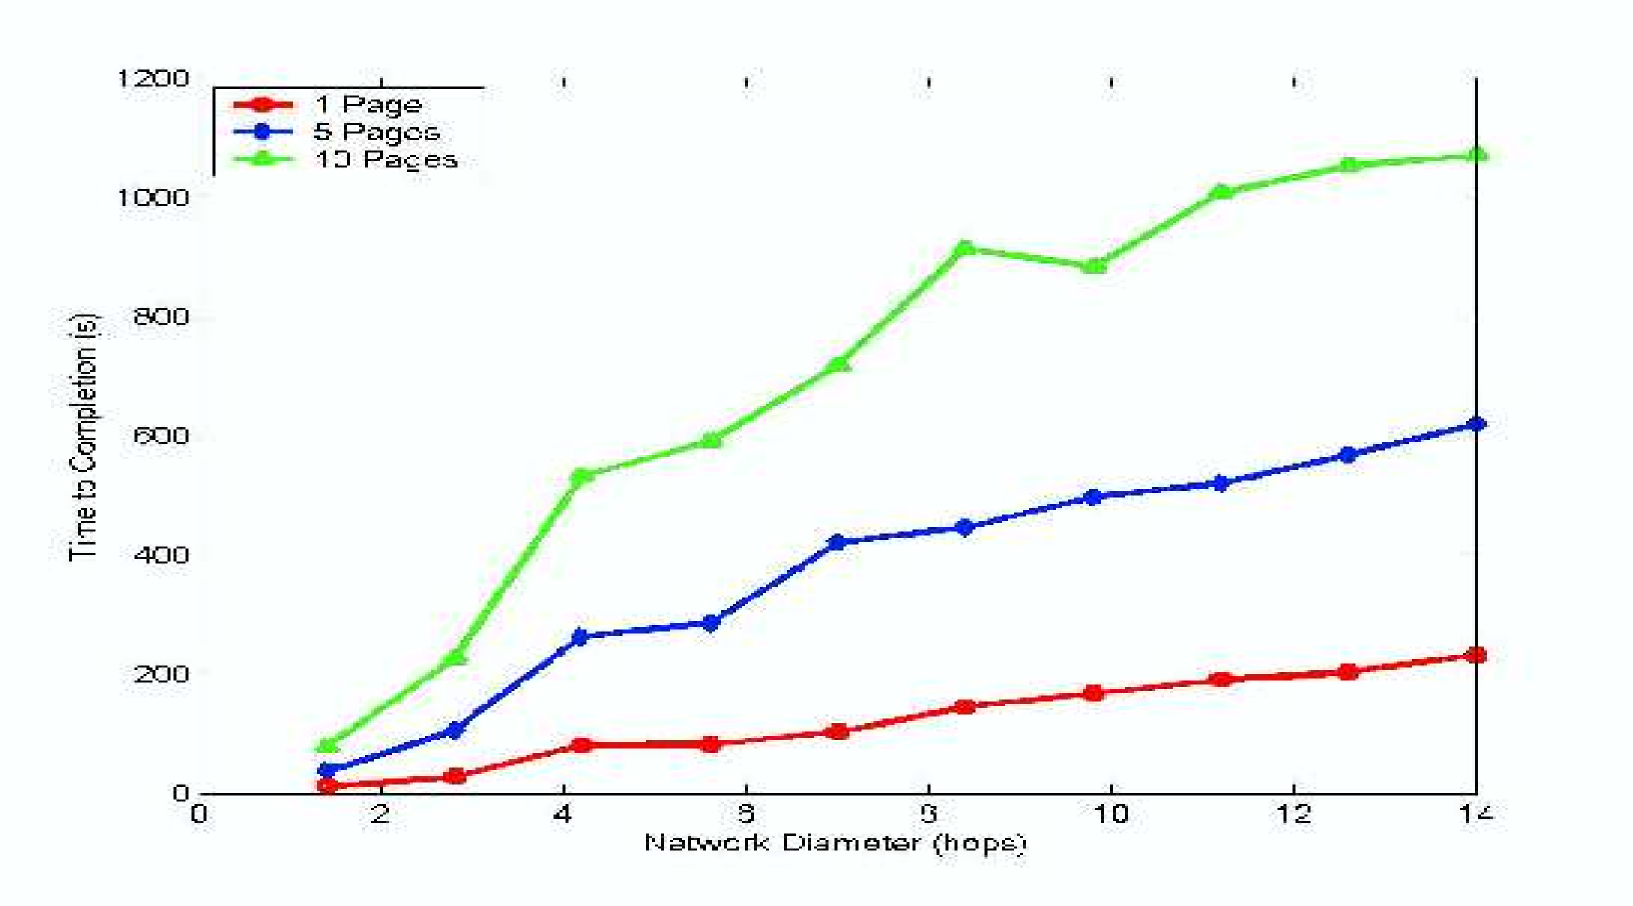

Simulated Completion Time for Varying Object Size

Completion time is linear with the object size.

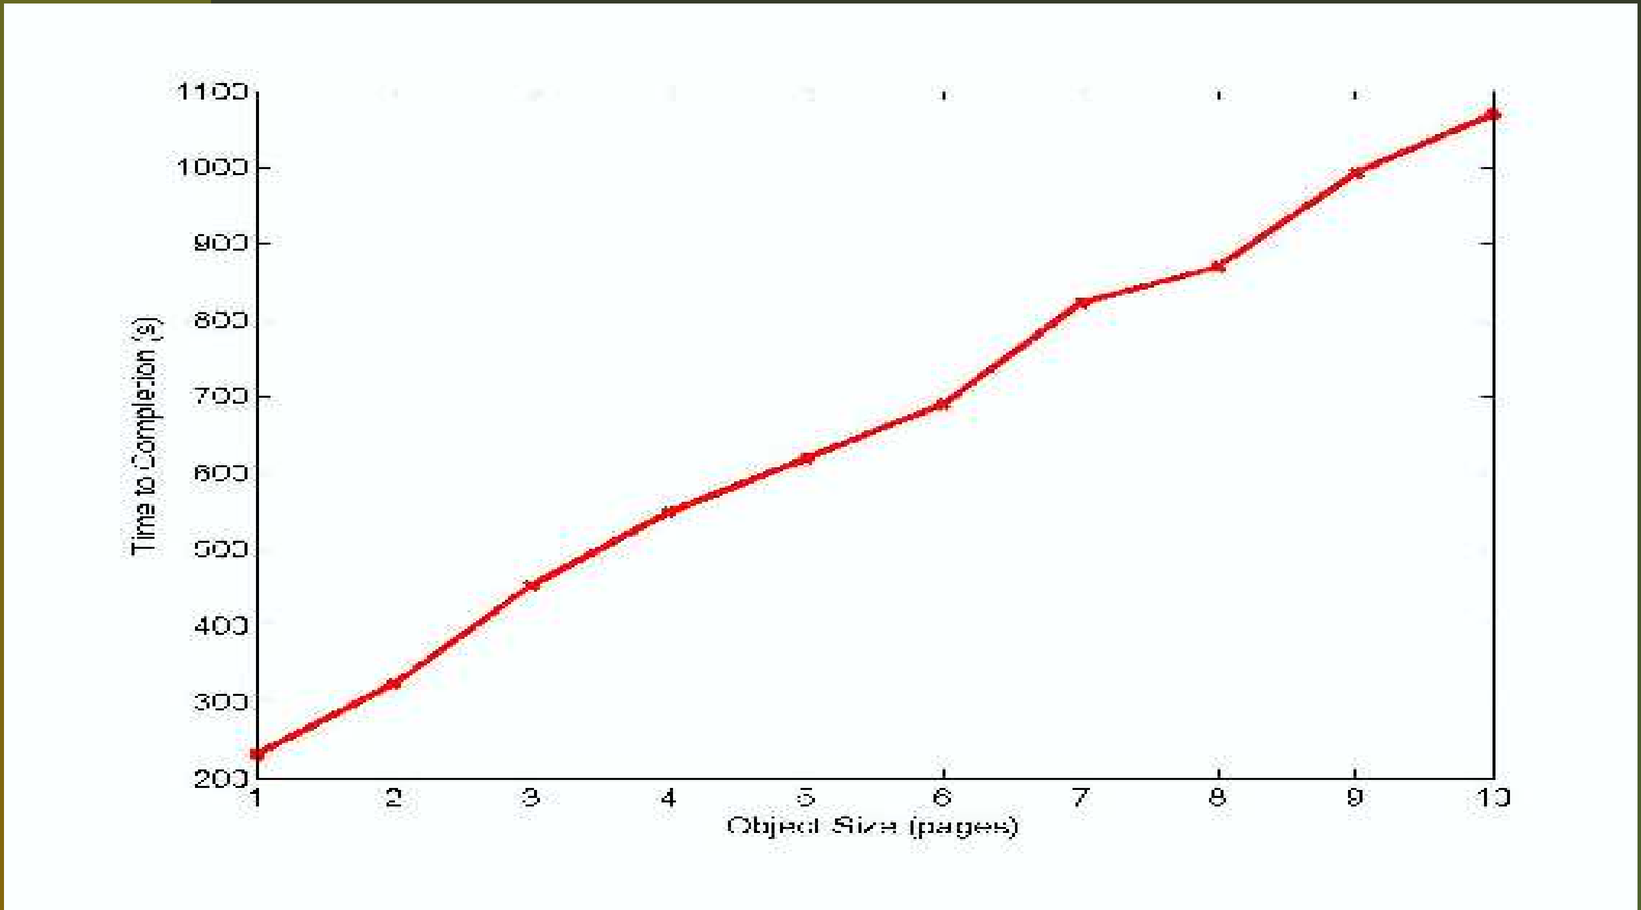

Simulated Completion Time for Varying Density

Completion time increases with density.

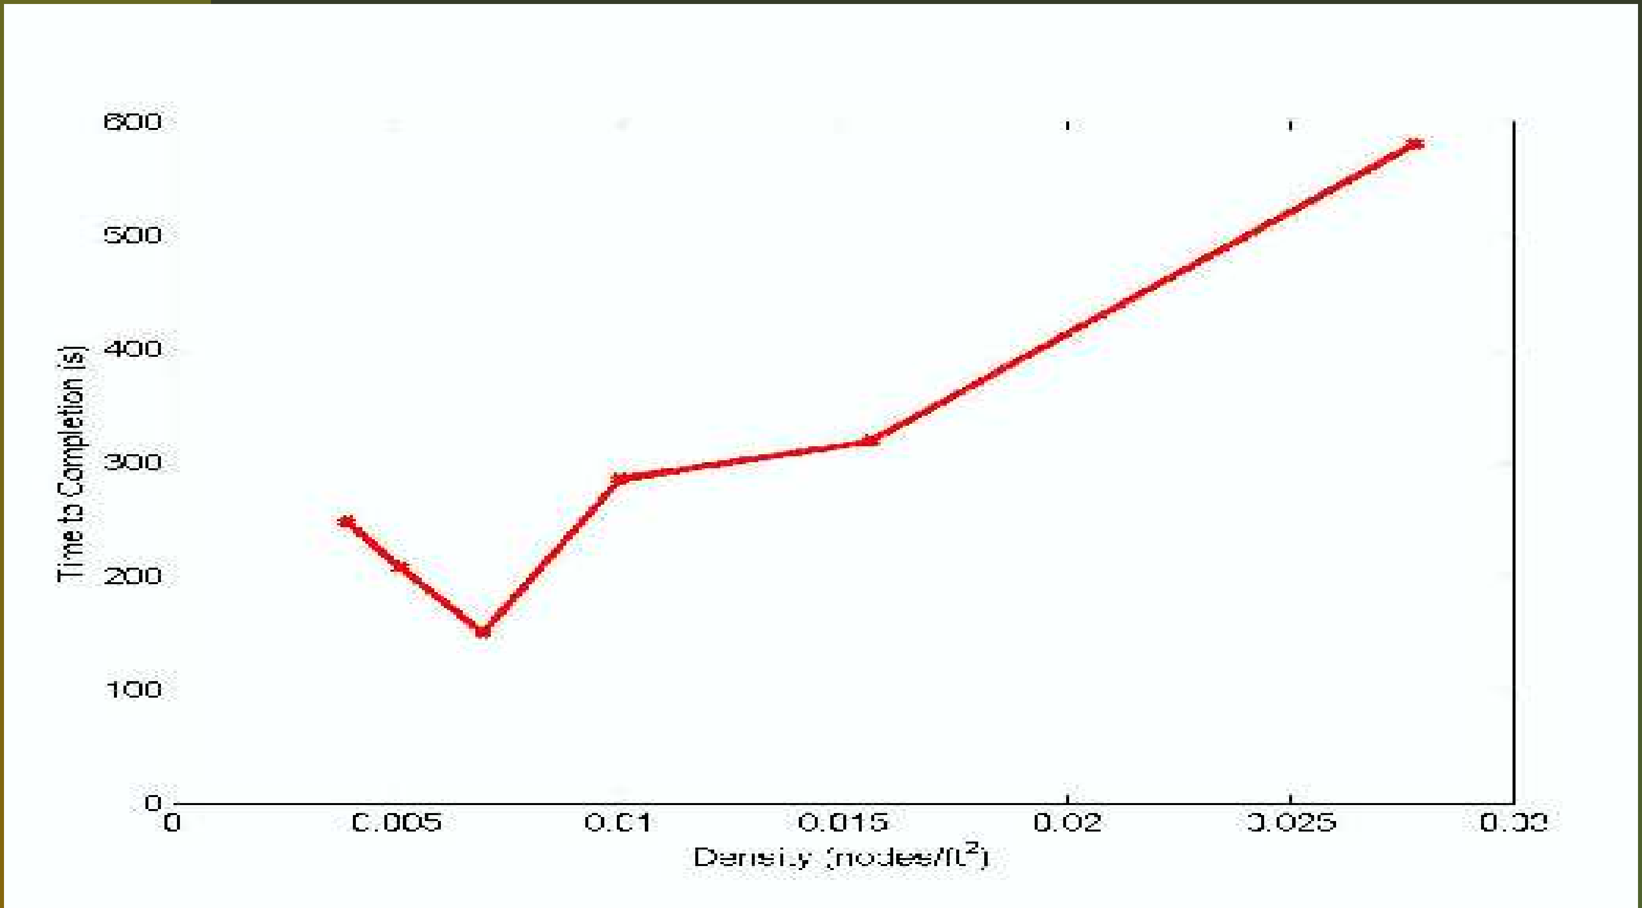

Simulated Propagation Time for a Sparse Network

Constant rate, wavefront pattern.

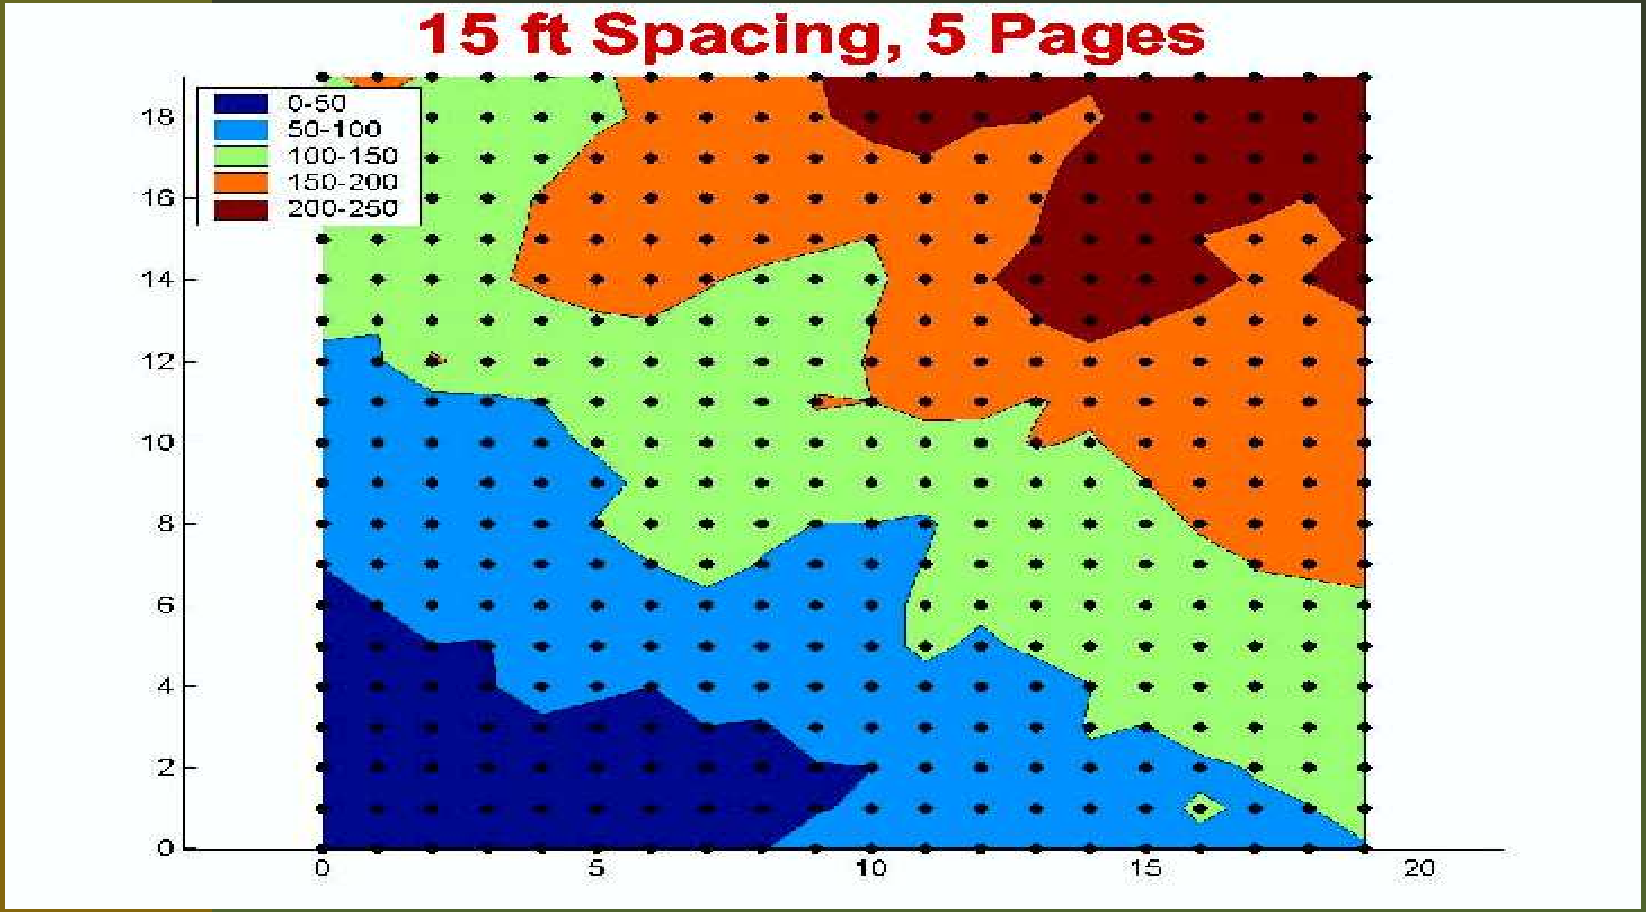

Simulated Propagation Time for a Dense Network

Propagation fast on the edges, slow in the middle.

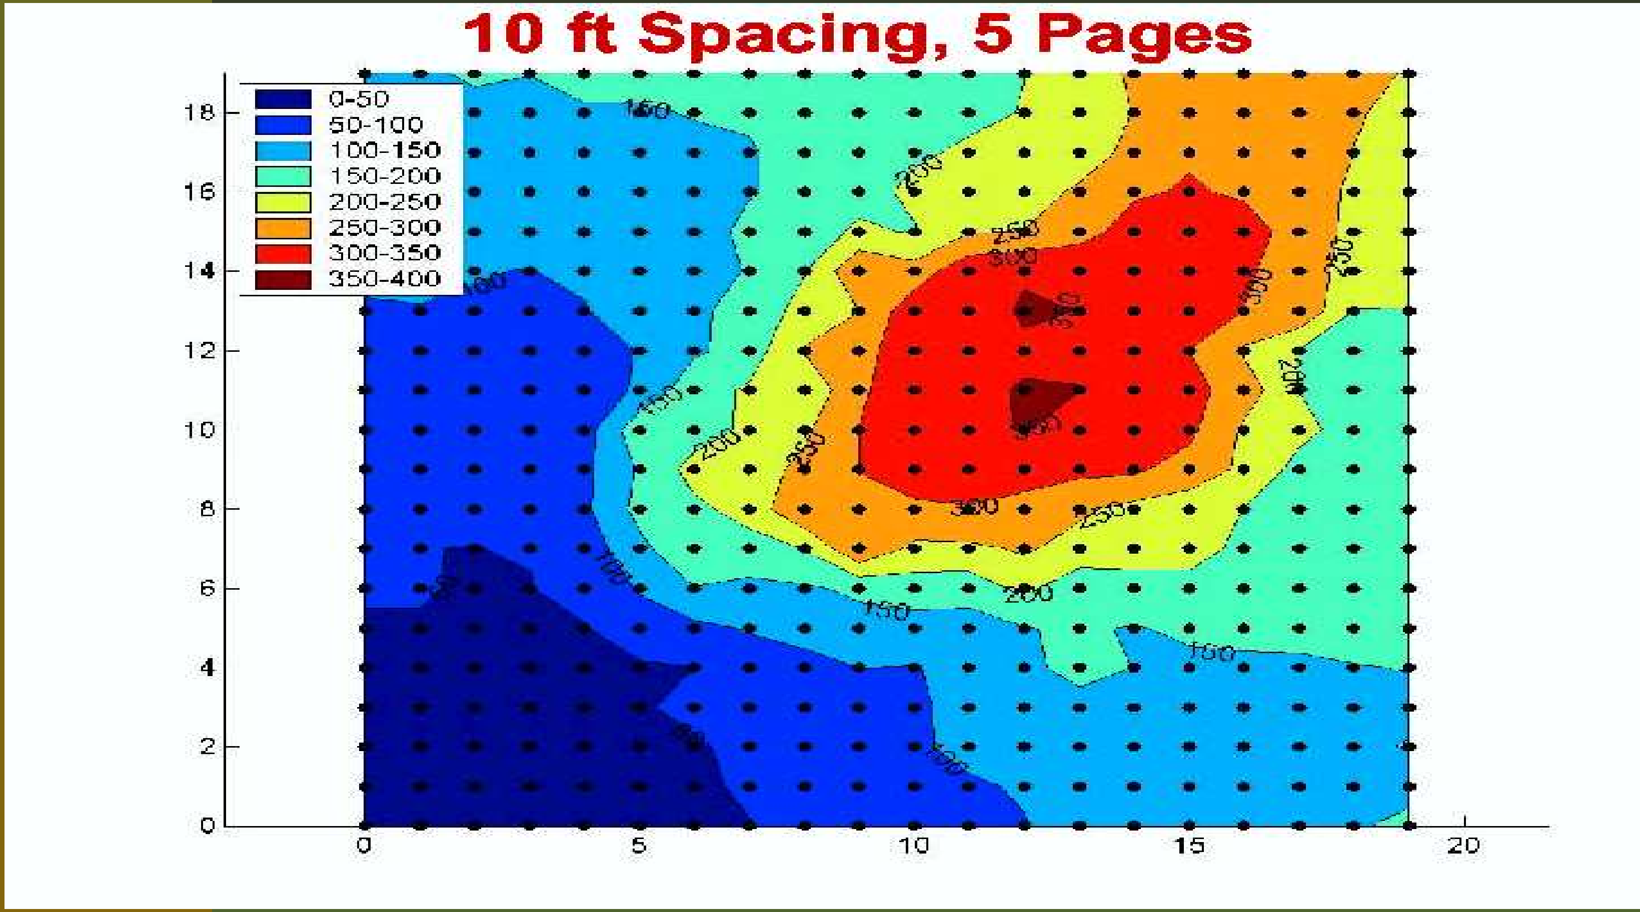

The Culprit is the Hidden Terminal Problem

Nodes in the center have more neighbors and are more prone to collisions.

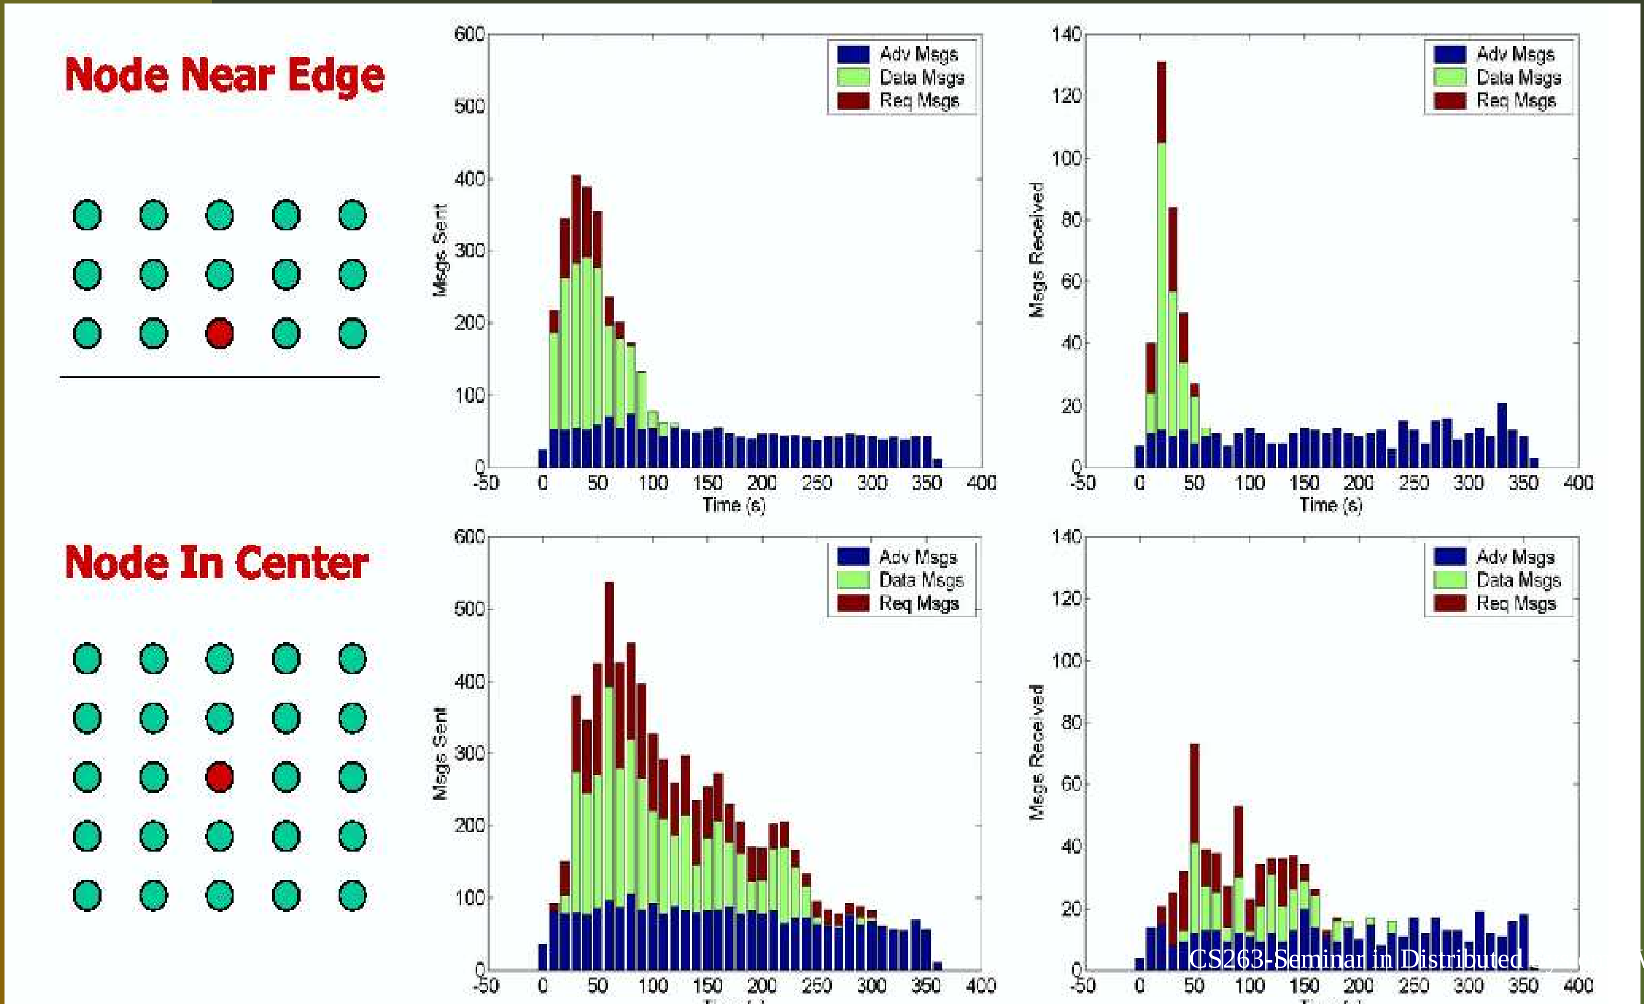

Lower Bound on Dissemination Rate

- Dissemination is slower than routing a single message across the network.
- Spatial multiplexing provides great improvement.
- Empirically data redundancy ratio is usually less than 5.

Conclusions

- Deluge, a reliable data dissemination protocol for propagating large data objects from a few nodes to many nodes of a multi-hop WSN.
- Epidemic protocol, achieves 90 bytes/second dissemination rate.
- Part of TinyOS 1.1.8.

Mobile P2P vs WSNs

Mobile peer-to-peer networks differ from WSNs in a number of ways:

- Built for different purposes.
- Nodes are more powerful, allowing for more complicated protocols.
- Direct connections between the nodes are possible.

Data Dissemination in Mobile P2P

- Nodes build content synopses of their data based on Bloom Filters.
- Adaptively disseminate them to the most appropriate nodes of the network.
- Based on the synopses they forward queries to the nodes with the highest probability of providing the desired results.

Dissemination Strategies

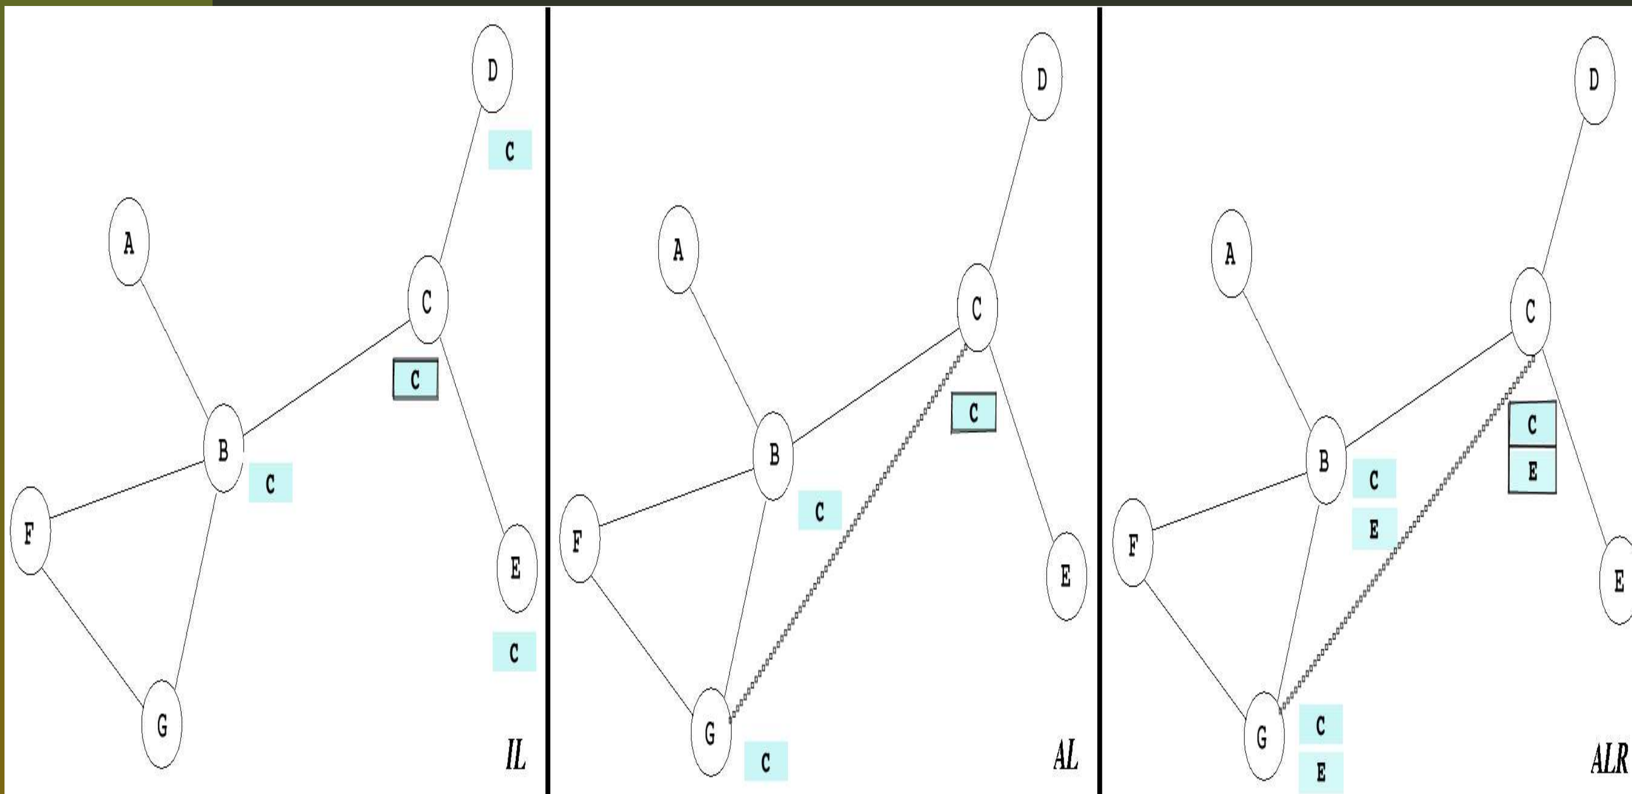

Thank you!

Questions/comments?

References:

- The Dynamic Behavior of a Data Dissemination Protocol for Network Programming at Scale, Sensys 2004
- Trickle: A Self-Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks, NSDI 2004
- Data Dissemination in Mobile Peer-to-Peer Networks, MDM 2005