# NSFA: Nested Scale-Free Architecture for Scalable Publish/Subscribe over P2P Networks

Huanyang Zheng and Jie Wu
Dept. of Computer and Info. Sciences
Temple University

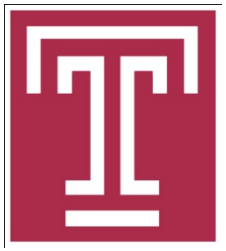

# Road Map

- Introduction
- Nested Scale-Free Architecture
- Publish/Subscribe System
- Experiments
- Conclusion

# 1. Introduction

## Publish/Subscribe System

Design paradigm for Content Delivery Network (CDN)

Three roles: subscribers, publishers, and brokers

Subscribers express their interests via system subscriptions, in order to receive interested events

Events are issued by publishers and are delivered to subscribers via brokers

Examples: Yahoo Message Broker, Global Data Synchronization Network, and SuperMontage

# 1. Introduction

Publish/Subscribe System

- Implemented as overlay networks on the Internet
- Usually expensive infrastructure costs
- Akamai's pub/sub: 56,000 dedicated servers

P2P-based Publish/Subscribe System

- Inexpensive and highly scalable
- A peer can be a subscriber, a publisher, or a broker

# 1. Introduction

## P2P network classification

Peer churn refers to peer arrivals and departures

| | Unstructured | Structured |
|---|---|---|
| Example | Gnutella, Kazaa, and Bitcoin | Chord, P-Grid, and Pastry |
| Peer connections | "Random" | Given topology |
| Advantage | Low construction overheads and peer churn robust | Efficient routing |
| Disadvantage | Inefficient routing | High construction overheads and not churn robust |

# 1. Introduction

Key observation

Unstructured P2P networks are not truly random

| | Unstructured | Structured |
|---|---|---|
| Example | Gnutella, Kazaa, and Bitcoin | Chord, P-Grid, and Pastry |
| Peer connections | "Random" | Given topology |

They are scale-free (existing literature)

They are nested scale-free (our result)

New system is designed based on nested scale-free

# 2. Nested Scale-Free Architecture

## Scale-Free Architecture (SFA)

Definition: A network satisfies SFA, if its node degree distribution follows power-law.

Let $P_d$ denote the fraction of peers with a degree of d, the power-law means that

$$P_d = \frac{\alpha - 1}{d_{\min}} \cdot \left(\frac{d}{d_{\min}}\right)^{-\alpha}$$

α is a constant power-law exponent ranging from 2 to 3. $d_{min}$ is also a constant parameter from which the power-law degree distribution holds.

# 2. Nested Scale-Free Architecture

## Nested Scale-Free Architecture (NSFA)

Let $G_s$ denote the set of subgraphs generated by iteratively removing the lowest-degree node and its connections in a network, $G$.

Definition: $G$ satisfies NSFA if

- (i) $G$ and all the subgraphs in $G_s$ satisfy SFA
- (ii) the standard deviation of their power-law exponents (i.e., $\alpha$), is $o(1)$.

# 2. Nested Scale-Free Architecture

Nested Scale-Free Architecture (NSFA)

NSFA refers to onion architecture

Nodes with small degrees:

Outer network/onion layer

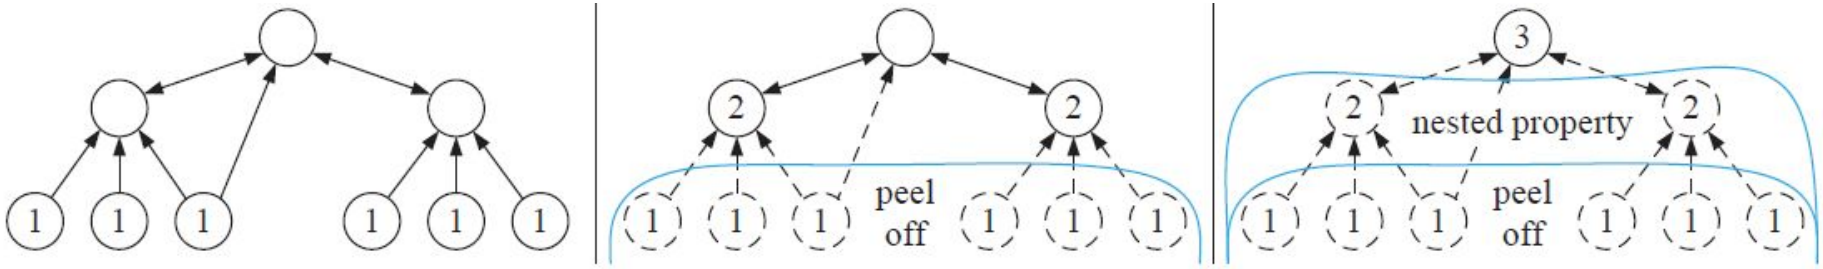

# 2. Nested Scale-Free Architecture

NSFA Verification Experiments

Gnutella dataset:10,876 peers with 39,994 connections

Bitcoin dataset: 4,579 peers with 18,667 connections

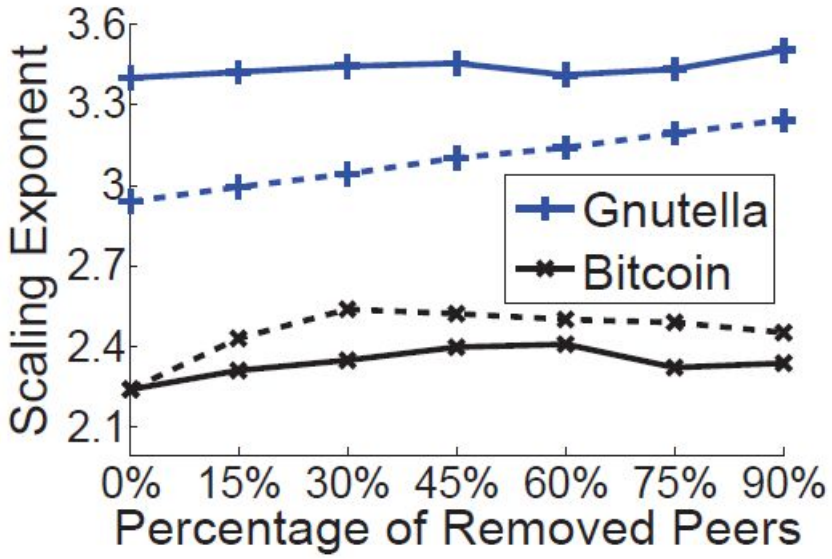


(a) Power-law scaling exponent.

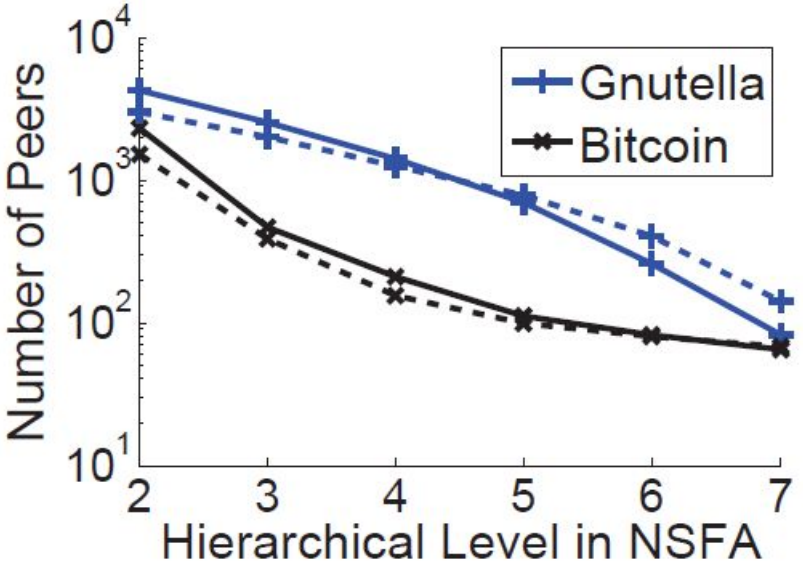


(b) Peer distribution in NSFA.

# 2. Nested Scale-Free Architecture

NSFA Verification Experiments

Gnutella dataset:10,876 peers with 39,994 connections

Bitcoin dataset: 4,579 peers with 18,667 connections

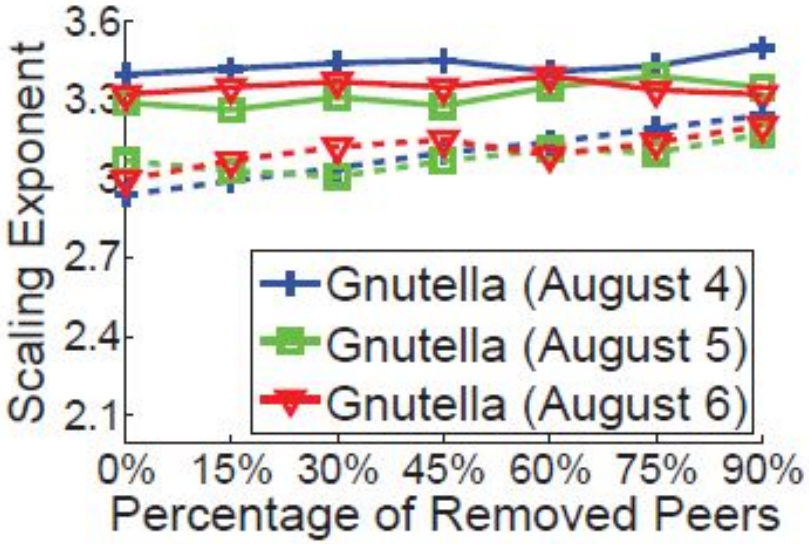


(c) Peer arrival and departure.

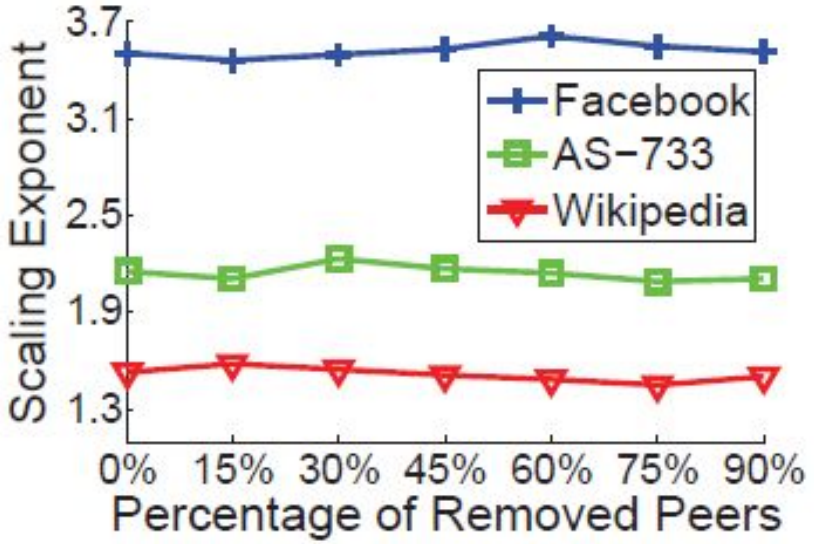


(d) Other NSFA networks.

# 2. Nested Scale-Free Architecture

Distributed Algorithm to reveal NSFA

Effective degree of a peer is defined as its number of connections to its unlabeled neighbors

1: Initialize peer v as unlabeled.
2: **repeat** in a round-by-round manner **do**
3: **if** v has the smallest effective degree among all its unlabeled neighbors (include tie) **then**
4: Set v's label to be the largest label among its labeled neighbors plus one (if v does not have a labeled neighbor, it is one).
5: **return** v's label as its hierarchical level.

# 2. Nested Scale-Free Architecture

Key Result

> Theorem: Suppose that an unstructured P2P network has V peers and satisfies NSFA. The distributed labeling algorithm is expected to terminate within $\Theta(\ln V)$ rounds of synchronous peer iterations. The maximal peer label is also $\Theta(\ln V)$.

Proof idea: a constant percentage of unlabeled peers are expected to label themselves in each round of the distributed labeling algorithm

# 3. Publish/Subscribe System

Pub/Sub System based on NSFA

Label node hierarchy (i.e., a forest) based on NSFA

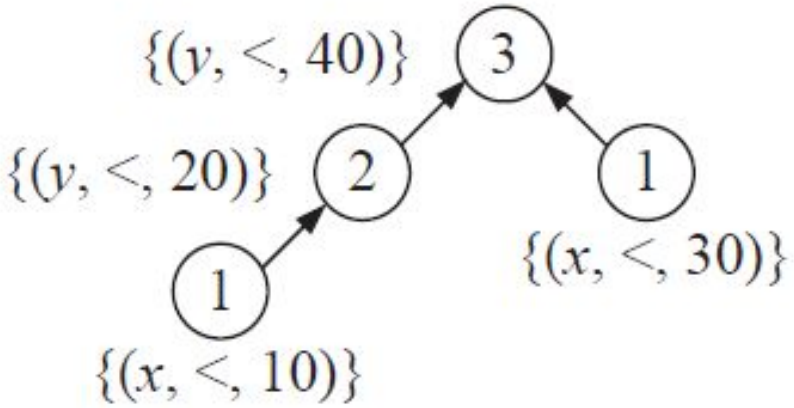


(a) Peer interests.

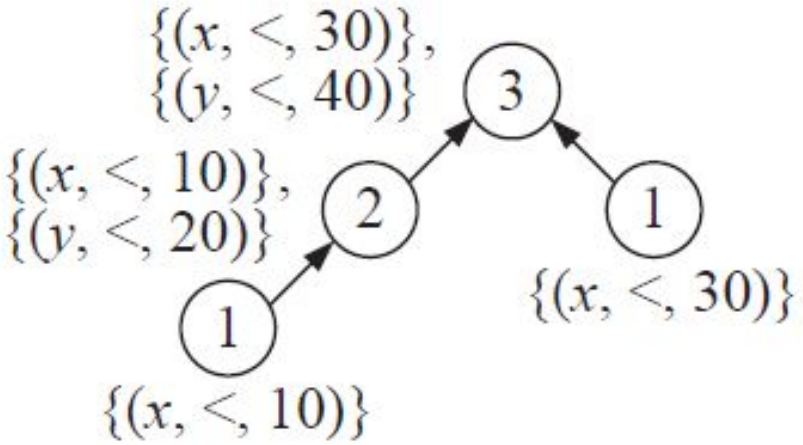


(b) Subscriptions.

Note that the tree depth is bounded by $\Theta(\ln V)$ according to the previous theorem

# 3. Publish/Subscribe System

Pub/Sub System based on NSFA

Up-and-Down routing for event delivery

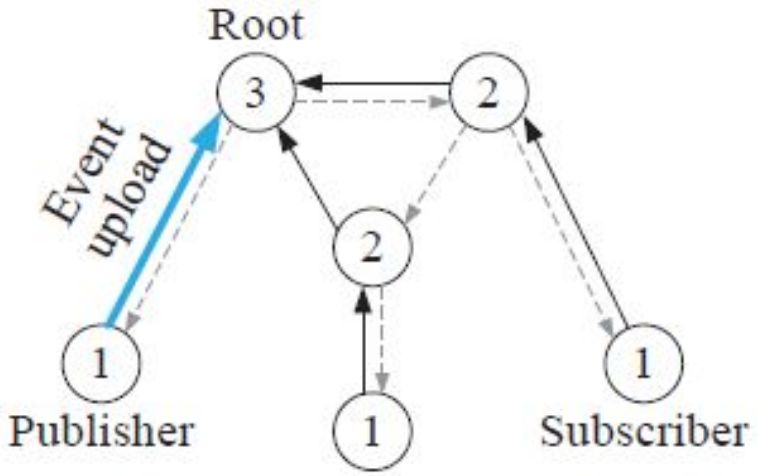


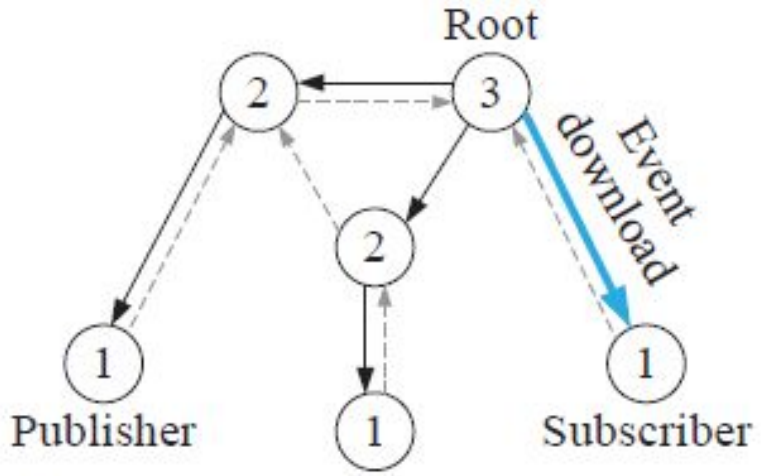


Use hierarchical routing for event delivery,
which includes two stages of upload and download.

# 3. Publish/Subscribe System

## Pub/Sub System based on NSFA

| Metrics | Structured P2P-based Pub/Sub | | |
|---|---|---|---|
| | Terpstra [25] | Meghdoot [9] | PastryStrings [1] |
| Event routing | $O(\ln \lvert V \rvert)$ | $O(\tau \lvert V \rvert^{\frac{1}{\tau}})$ | $O(\log_\mu \lvert V \rvert)$ |
| System robustness | $O(\ln \lvert V \rvert)$ | $O(\tau \lvert V \rvert^{\frac{1}{\tau}})$ | $O(\log_\mu \lvert V \rvert)$ |
| Overhead | $O(\ln \lvert V \rvert)$ | $O(\tau)$ | $O(\mu \log_\mu \lvert V \rvert)$ |

| Metrics | Unstructured P2P-based Pub/Sub | | | |
|---|---|---|---|---|
| | Sub-2-Sub [26] | Vitis [17] | Poldercast [20] | NSFA-based |
| Event routing | $O(\lvert V \rvert)$ | $O(\ln^2 \lvert V \rvert)$ | $O(\ln \lvert V \rvert)$ | $O(\ln \ln \lvert V \rvert)$ |
| System robustness | $O(1)$ | N/A | $O(\ln \lvert V \rvert)$ | $O(\ln \ln \lvert V \rvert)$ |
| Overhead | $O(1)$ | $O(1)$ | $O(1)$ | $O(1)$ |

# 4. Experiments

Event routing efficiency in Gnutella and Bitcoin

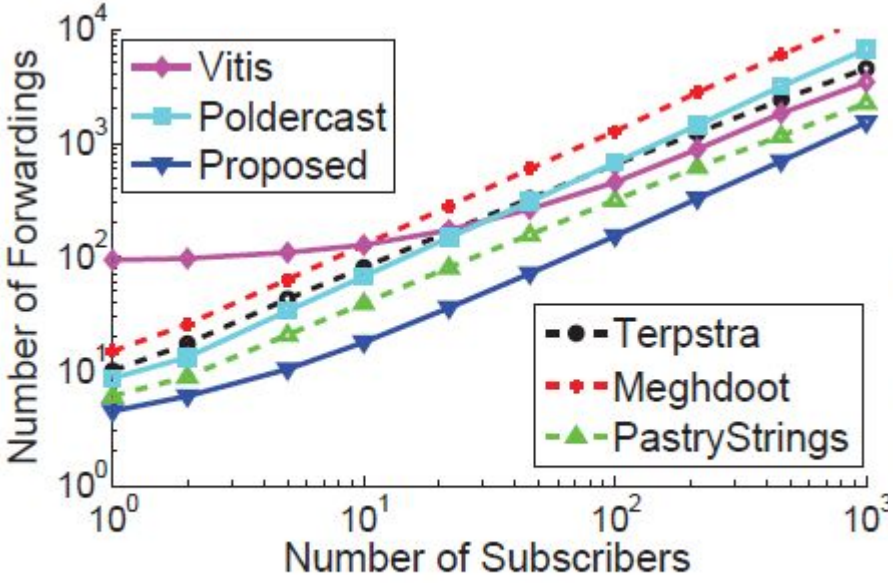


(a) Gnutella.

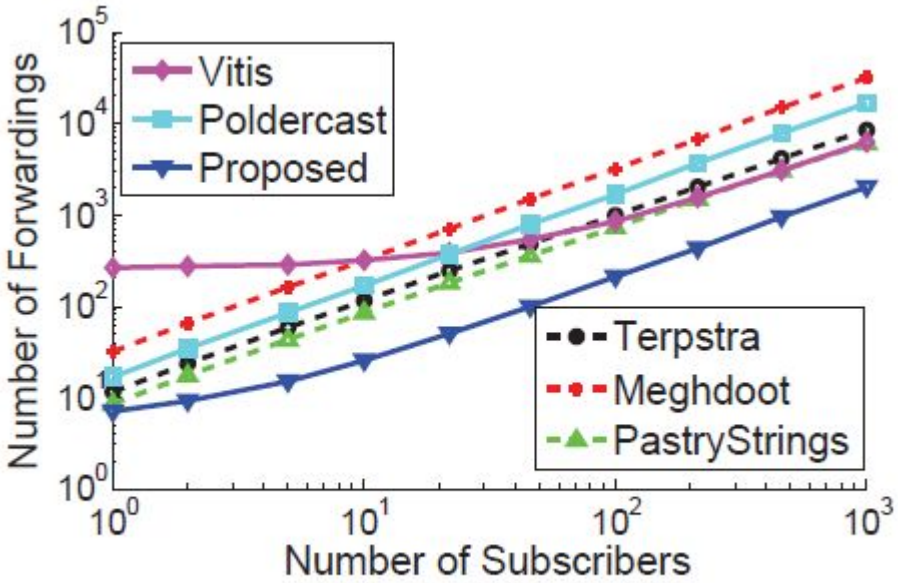


(b) Bitcoin.

NSFA-based pub/sub system outperforms the others since it has the best asymptotical bound of $O(\ln \ln |V|)$

# 4. Experiments

Peer arrival and departures in Gnutella

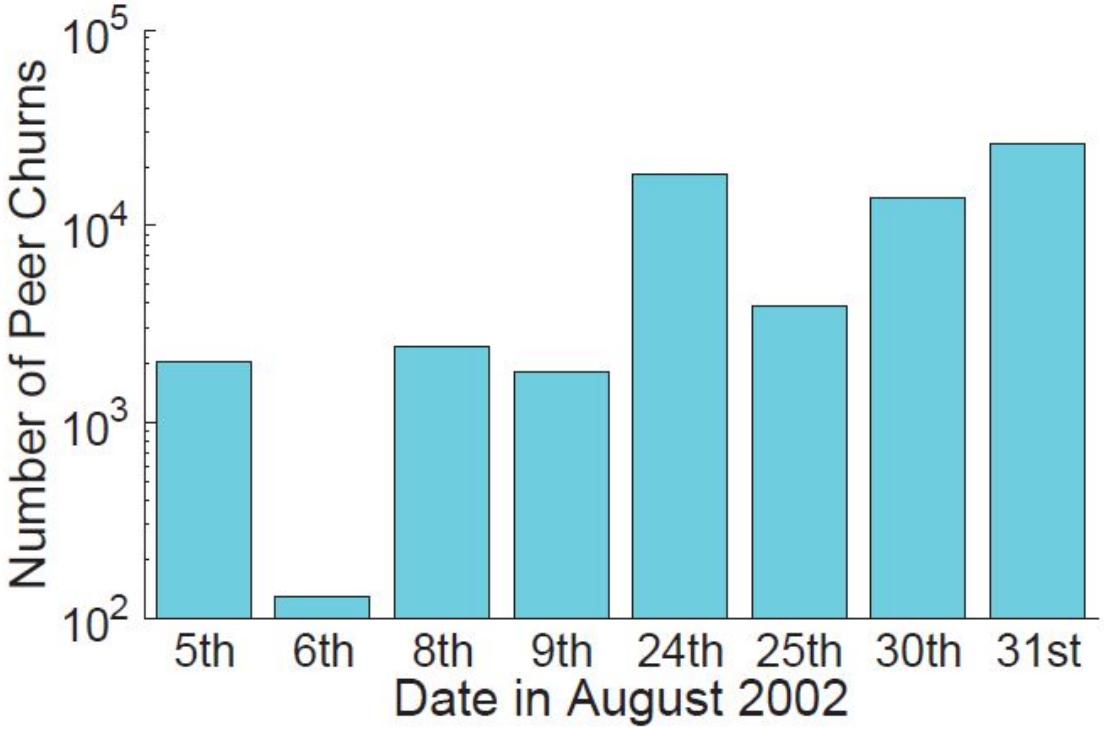


(a) Peer churn statistics.

# 4. Experiments

Peer arrival and departures in Gnutella

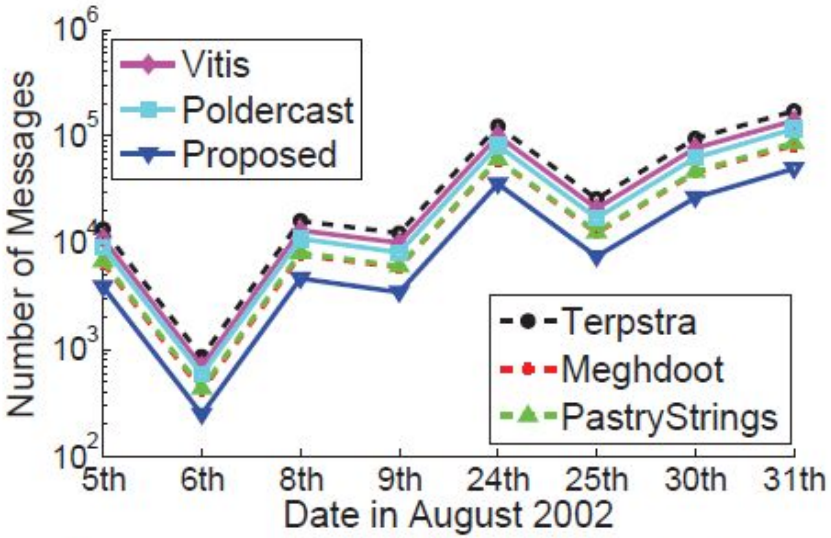


(b) Message complexity (no shared interest).

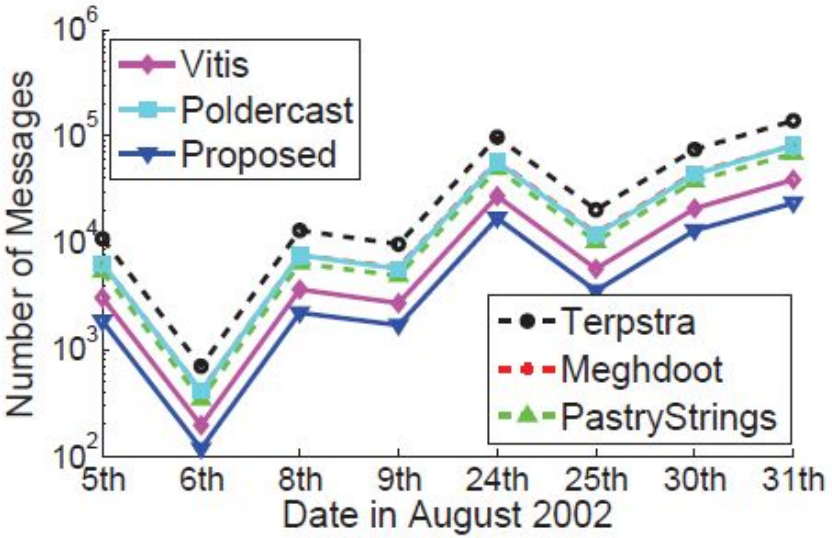


(c) Message complexity (shared interests).

Number of messages scales with numbers of peer churns

NSFA-based pub/sub system has asymptotical advantage

# 5. Conclusion

Nested Scale-Free Architecture (NSFA)

- Onion-like network hierarchy

NSFA-based Pub/Sub system

- Peers are not randomly connected
- NSFA leveraged to achieve better performance
- Asymptotical advantage over existing systems