Navigation in Decentralized Online Social Networks

Sarunas Girdzijauskas, KTH

07 November 2016

Shortly About me

- **PhD from EPFL, Switzerland 2009**
  - Research Intern at IBM Haifa 2008
- **Post-doc at SICS 2009**
- **Senior researcher at SICS 2011**
- **Assistant Professor at LCN/EES 2012**
- **Associate Professor at SCS/ICT 2016**

Research Topics

- **Networks and Graphs**
  - *Graph algorithms for Cloud Computing, Distributed and Decentralized Systems, Social Networks, Privacy preservation.*
- **Research Areas**
  - *Partitioning and community detection (Fatemeh)*
  - *Storing Social Network (linked) data in Cloud Environments (Anis)*
  - *Streaming graph partitioning (Anis)*
  - *Anomaly detection in Social graphs (Amira)*
  - *Gossip learning for decentralized Online Social Networks (Amira)*
  - *NLP using Graph Partitioning (Kambiz)*
  - *VM consolidation in clouds using Gossip (w. University of Umeå)*

- **This talk: on Navigation in DOSNs**
  - *Decentralized Online Social Networks*
  - *Community Cloud*

DOSNs: Motivation Slide

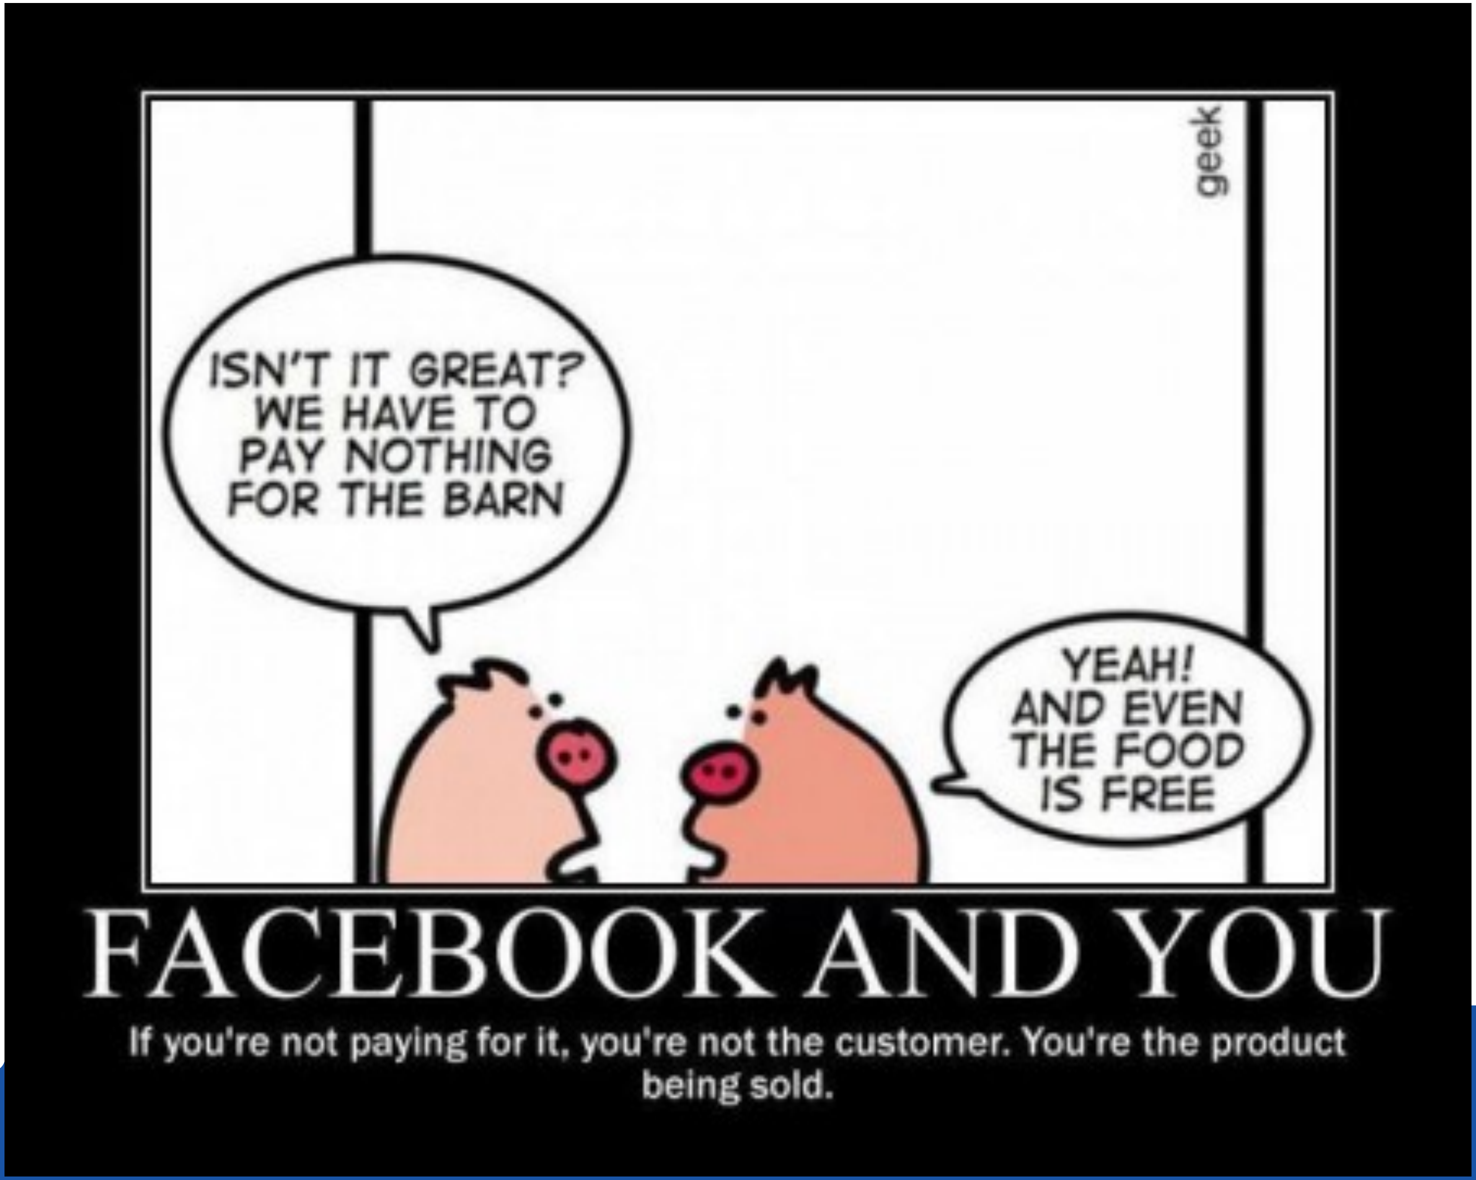

Why DOSNs?

Current Centralized Online Social Networks:

Give back the control to users:
Decentralized Online Social Networks

- **Challenging environments: Decentralized, highly heterogeneous (both resources and demand)**
- **Promises for: Scalability, availability, robustness, efficiency…**
  - *E.g,. DIASPORA, PeerSoN, Safebook, Vis-à-Vis, NetTube, DECENT, PrisM, GoDisco, SocialTube*

DOSN basic services

- **Data Dissemination/Search & Storage**
    - More advanced services: global aggregation/learning, analytics, recommendation etc.

- How do we arrange decentralized nodes together (i.e., design the **topology** of such decentralized system) where the above **services perform the best**.
    - I.e., Minimize traffic load/relaying and latency.

    - ***How can decentralized search even work?***
    - ***Milgram's Experiment***

Milgram’s Experiment

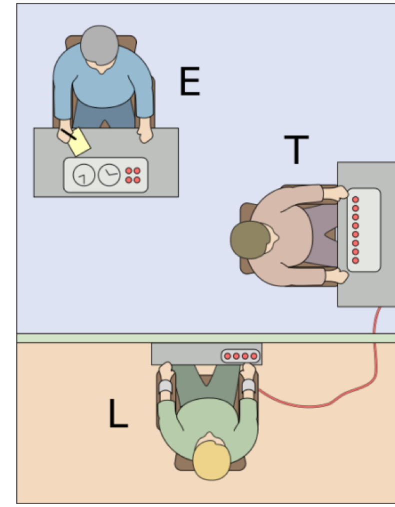

- **Not the one on obedience and authority!**
- **Milgram’s Experiment on the “topology of our Social Networks”**
  - *No Online Social Networks in 1960s*
  - *296 random people in USA forwarding a letter to a “target” person in Boston.*
  - *Personal info on the “target” (including address and occupation)*
  - *Forward only to a person known by first-name basis.*
  - *Result?*
  - *64 chains succeeded.*

Avg. number of steps?

- *We live in a* ***Small-World:*** *average length of the chains that were completed lied between* ***5 and 6 steps;***
- *Coined as "****Six degrees of separation****" principle.*
- *Similar results have been found in many other social networks*

Topology of Social Graph

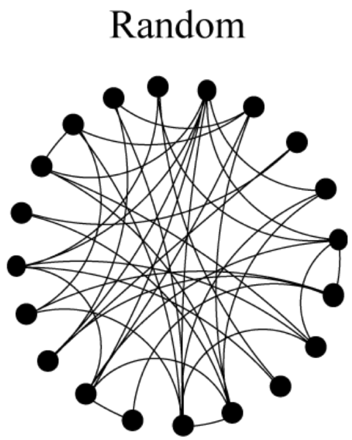

- **How to interpret such network?**
- **Maybe it is a Random Graph?**
  - *low diameter!*
  - *But very little clusteristaion i.e., very few triangles (common friendships)*
- **Watts-Stogatz Small World Model**

Weak links

Strong Links

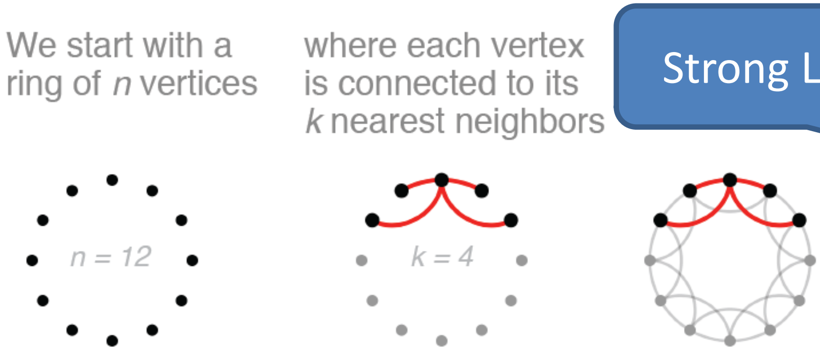

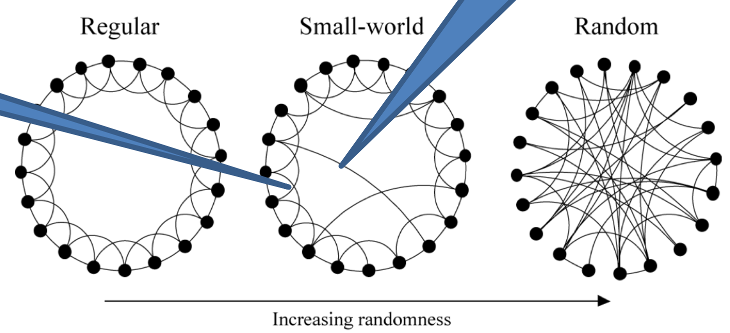

High clusteristaion but large diameter!

High clusteristaion and Short diameter!

Milgram's Small-World Experiment (cont.)

- **In 2011 Facebook shrunk 'Six Degrees Of Separation' to just 4.74 (721m users, 69b friendships)**
  - *Twitter's 5,91 of 12,8M friendships*
- ***Does the low diameter of the SN graph answer our question?***
  - *Surprise is that we can find these paths in a decentralized manner, i.e.,* ***navigate successfully*** *with no "map", no central authority, no "Big Brother".*
- ***Why it Works?***

So why Milgram’s experiment worked?

- **Social network is not a bare graph of vertices and edges**
  - *Nodes come with certain implicit “labels” representing various dimensions of our life*
  - *Hobbies, work, geographical distribution etc.*
- **There is (are multiple) “concept” space(s)**
  - *E.g., Geographical, Occupational, Hobby etc.*
  - *Each with a explicit or implicit* ***distance metric****!!!*
  - *We can greedily minimize the distance!!*
  - ***Decentralized search****: a greedy distance minimizing routing algorithm*

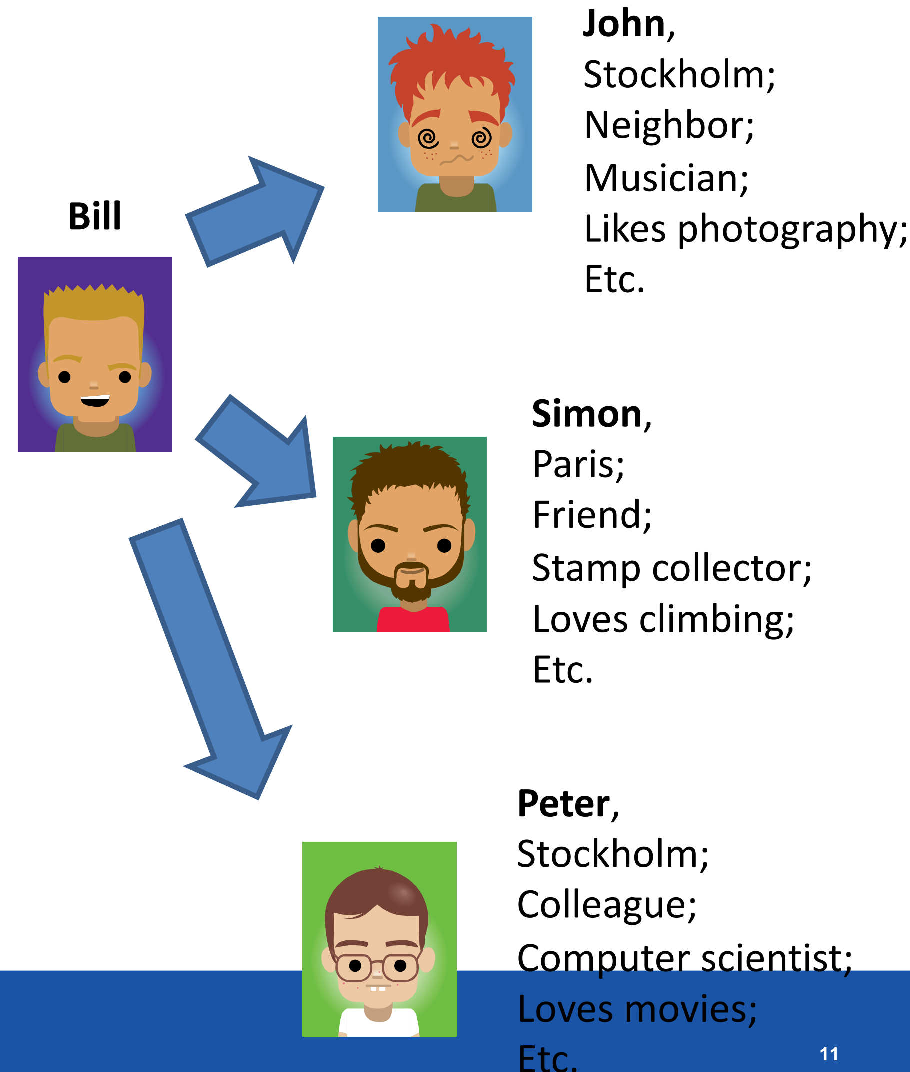

Fundamental Navigability Rules

- **Kleinberg’s Navigable Small-World model**
  - *Very rough insight (”ru :*
    - Connect to nodes tha ***ely proportional*** to the distance from you.

With O(log(N)) links we can nav gate in O(log(N)) hops

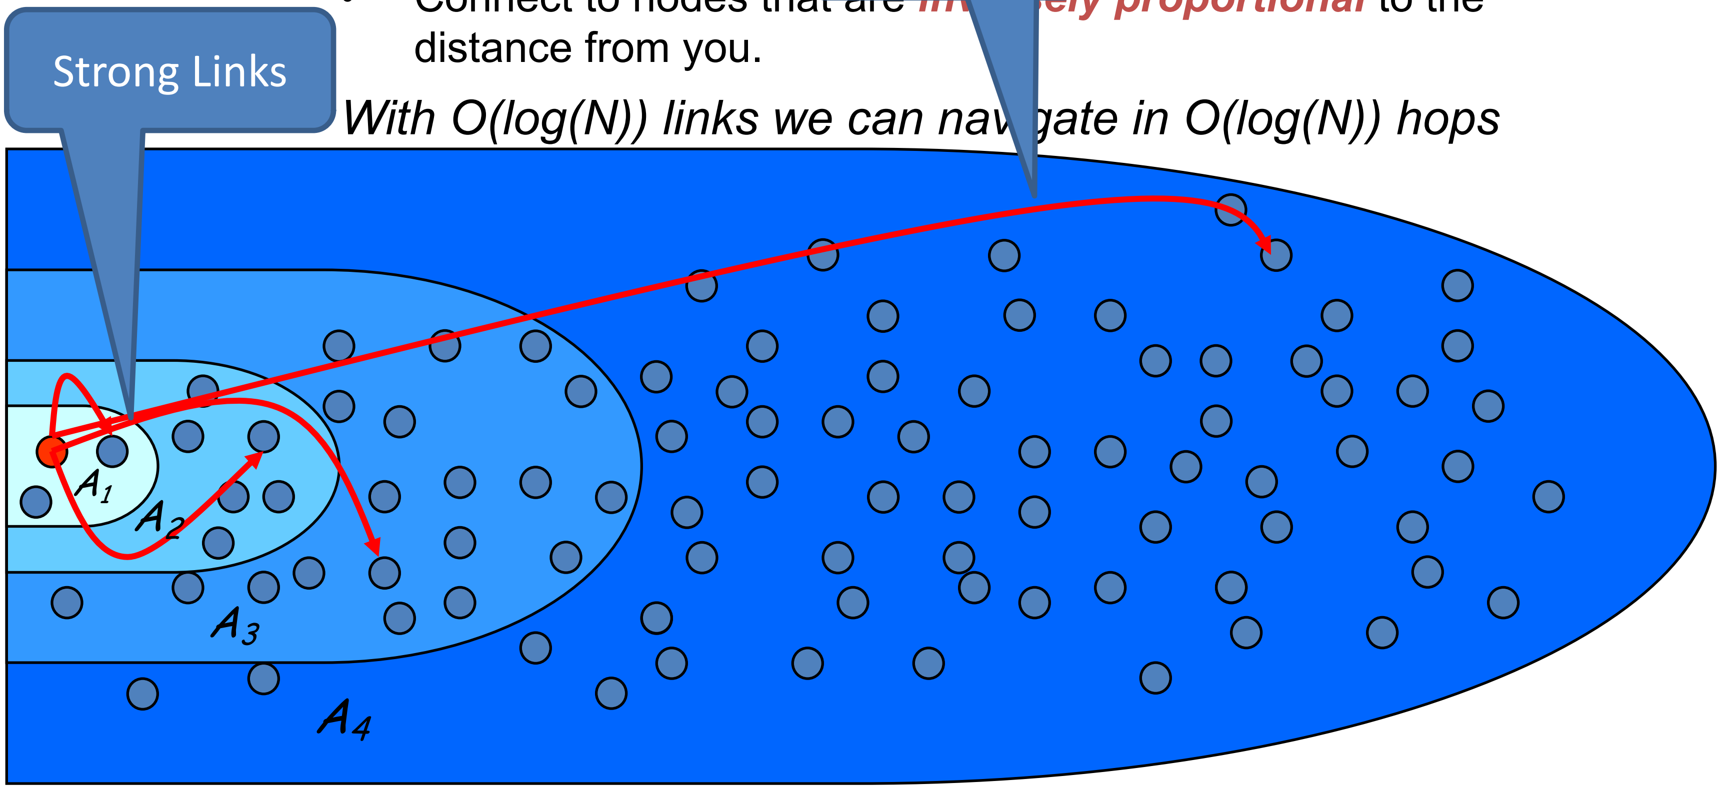

Traditional DHTs and Kleinberg model

- P2P networks, DHTs
- Kleinberg's model

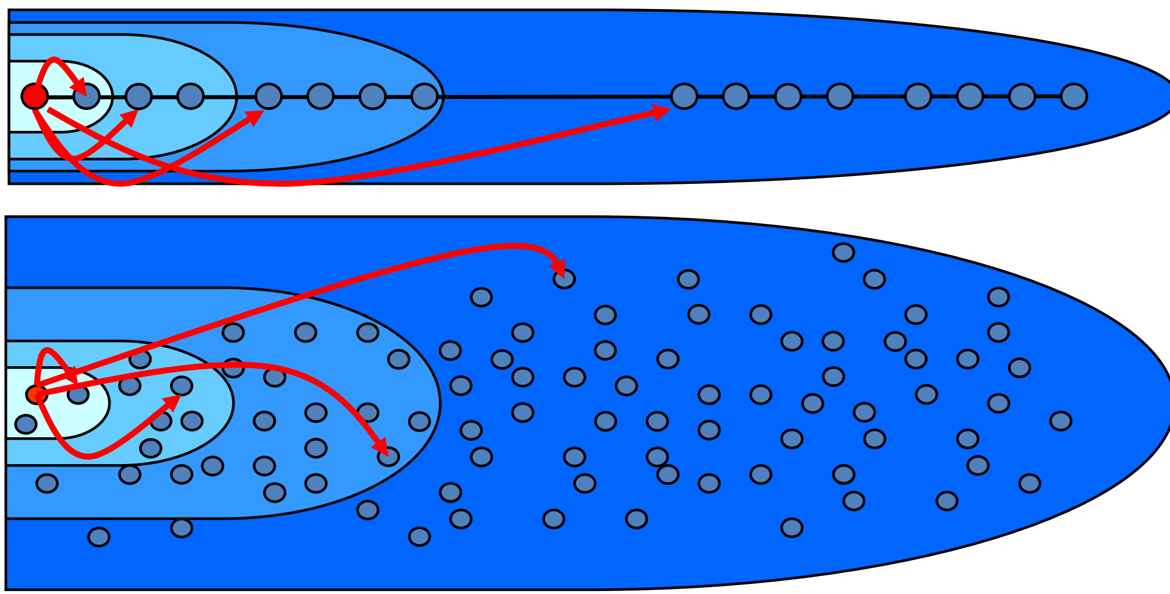

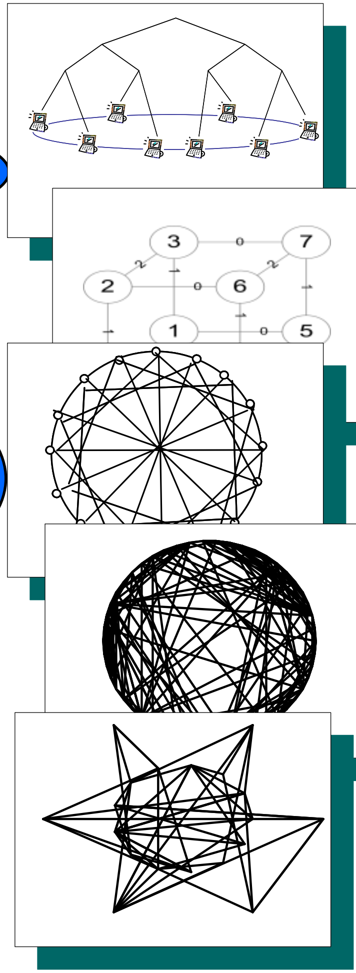

How do we build Structured P2P (DHT)? (recap)

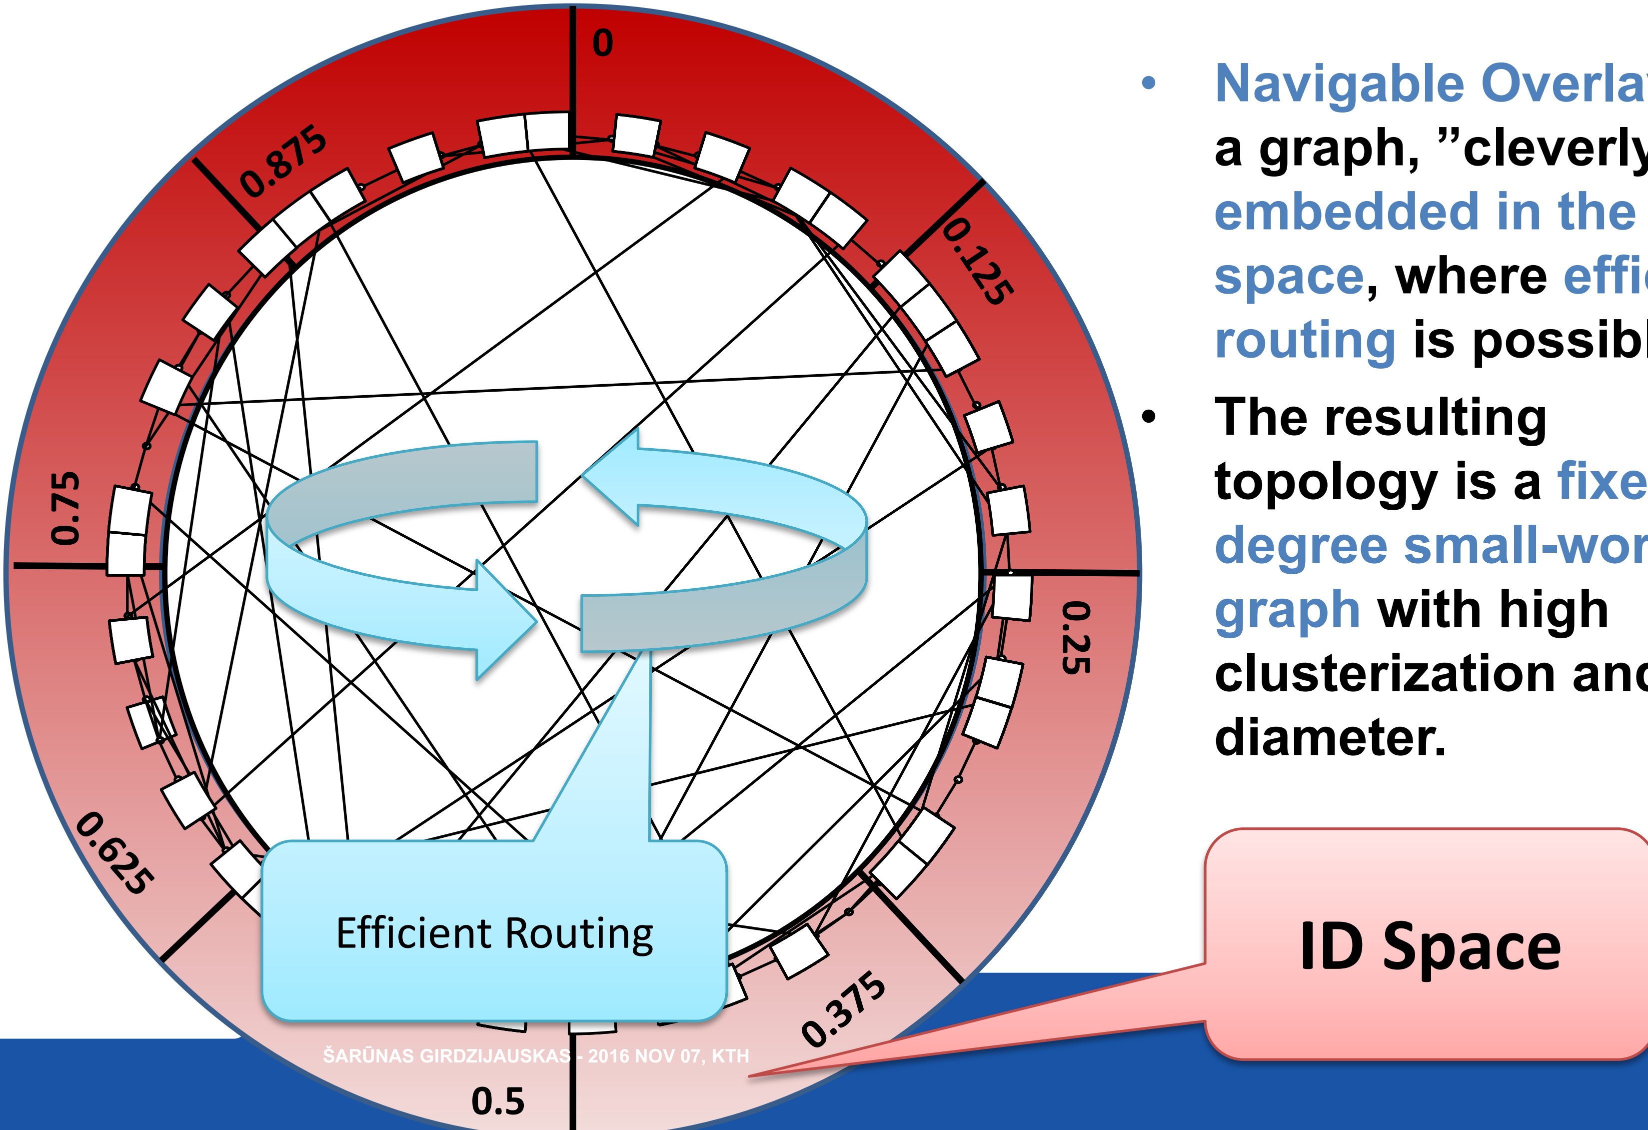

- **Navigable Overlay is a graph, "cleverly" embedded in the ID space, where efficient routing is possible.**
- **The resulting topology is a fixed degree small-world graph with high clusterization and low diameter.**

Choice of Topology for DOSNs?

- Structured P2P networks (DHTs)

Event Dissemination in DOSNs

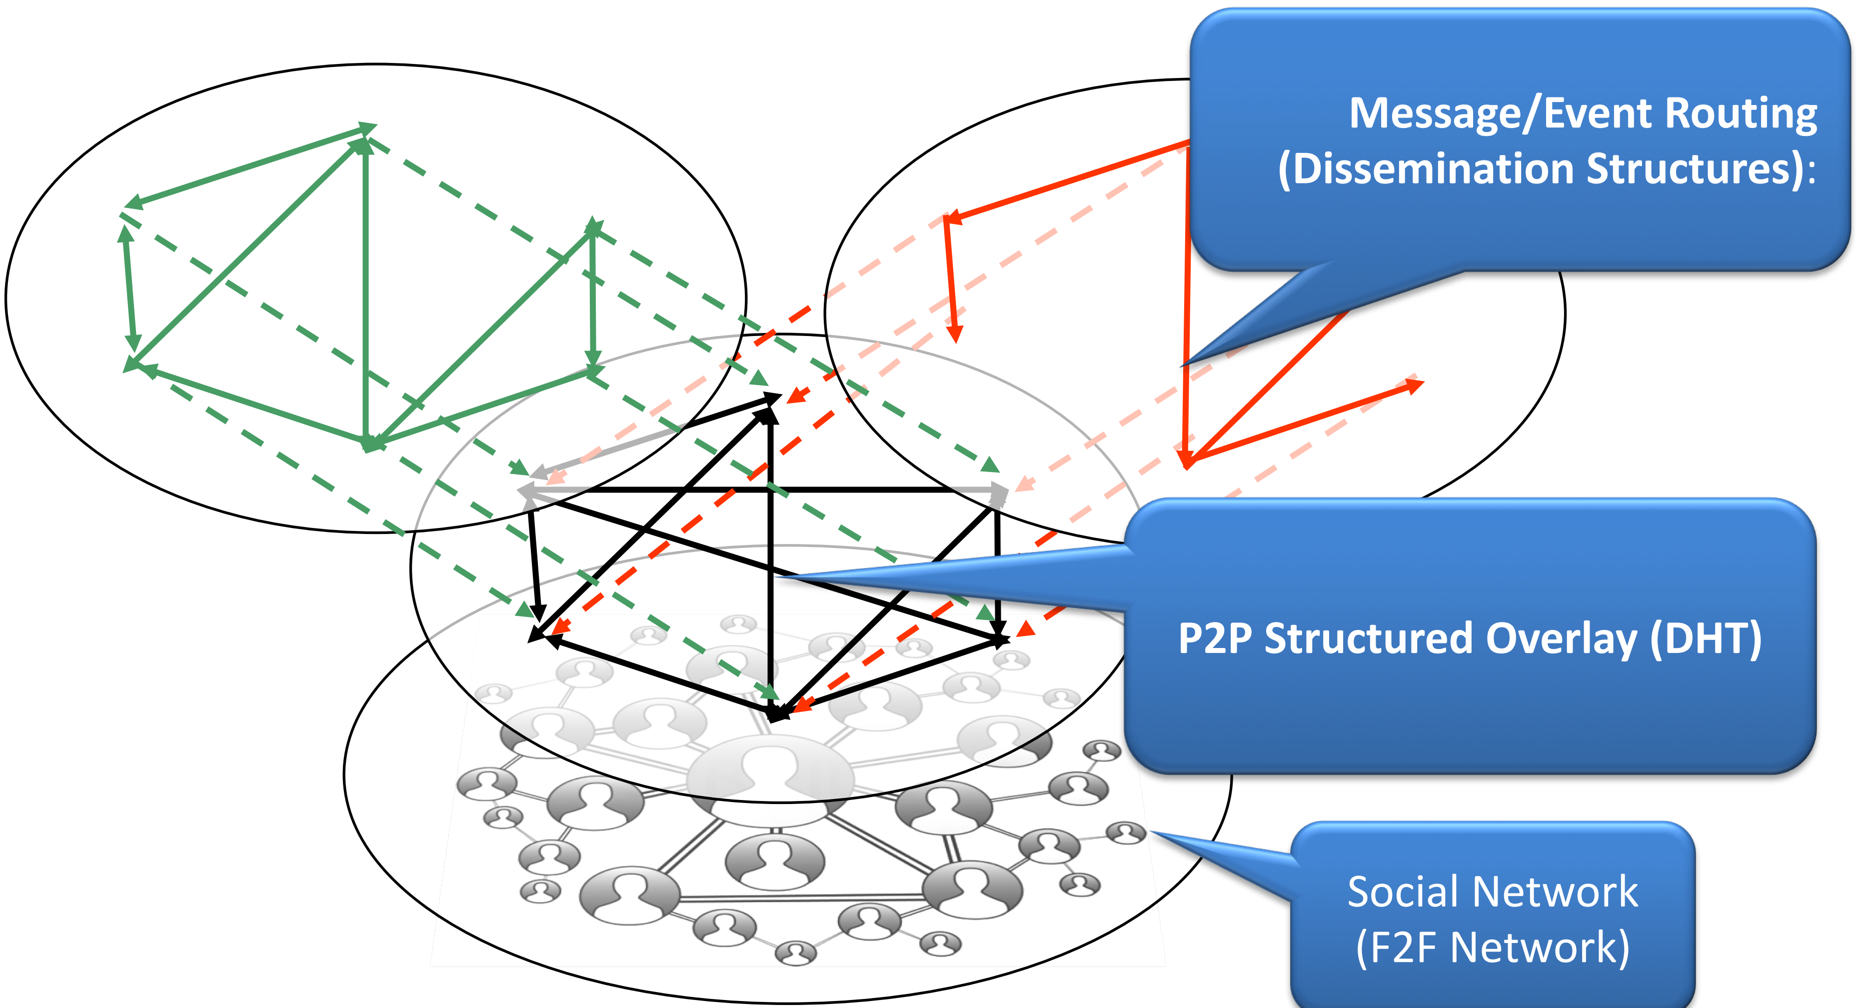

Navigable Overlays (DHTs) as the backbone for DOSNs

- **Navigable** (e.g., . Search in O(logN) number of hops) **and has very low degree** (O(logN)) graph
- Each node gets **uniform random IDs** (e.g., on a unit ring)
- Connect by some **predefined rules to *k* other** nodes based on their IDs
  - *As discussed in previous slides*
  - *E.g, Chord, Pastry, Symphony etc.*

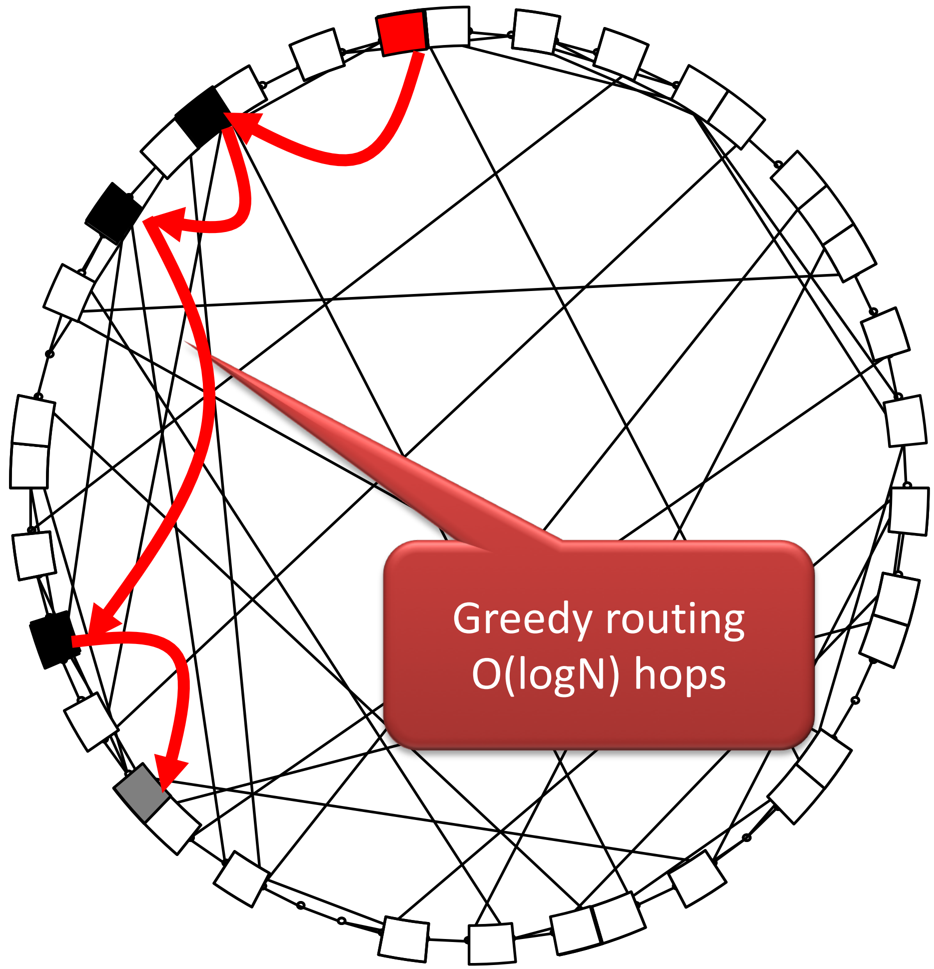

Event Publishing on Navigable Overlays

- **Building** **efficient** **data dissemination structures**
  - Creating a dissemination tree with a fanout max of DHT degree O(logN), and a depth max of expected search cost O(logN)
- **Scalable**
- **Robust**
- **Bounded delay**
- **Message overhead and relays!**
- **Privacy!**
- **Services on top! (e.g,. Storage)**

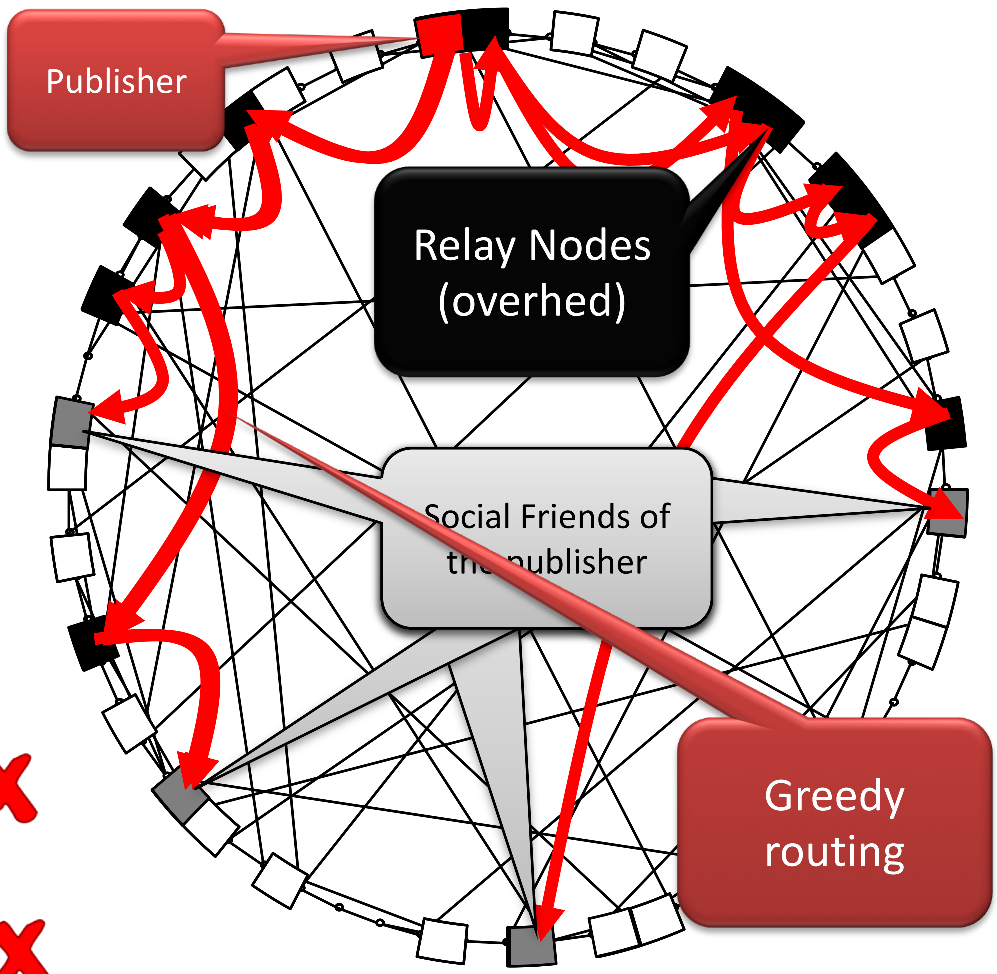

Why is DHT so inefficient in messages and dissemination cost?

- **We build a very efficient (navigable) overlay that is based ONLY on node IDs and is completely oblivious to Friend-to-Friend (F2F) network and node localities**
  - **Friends that are close in "social graph" are uniformly distributed in the DHT ring – scrambled forever**
  - **Any group communication induced by social activities on expectation will not be local and will induce O(log N) communication cost (likely non-interested relays).**
  - **In effect we end with:**
    - ***No locality,***
    - *(almost) **no direct friend-to-friend communication***
    - *Dissemination structures (trees) that we build will have on expectation around **O(logN) relay nodes** for every friendship*
      - i.e., vast majority of relays for each "newsfeed" or other action.

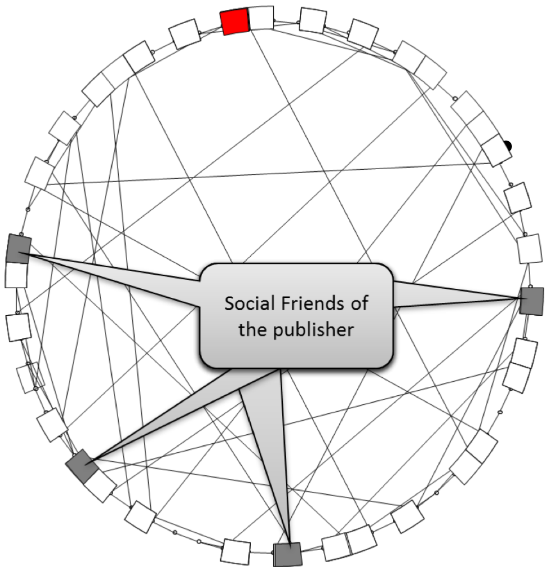

Choice of Topology for DOSNs?

- ~~Structured P2P networks (DHTs)~~
- Unstructured P2P
  - *Friend-to-Friend based networks*

Overlay Design for DOSNs (2nd attempt)

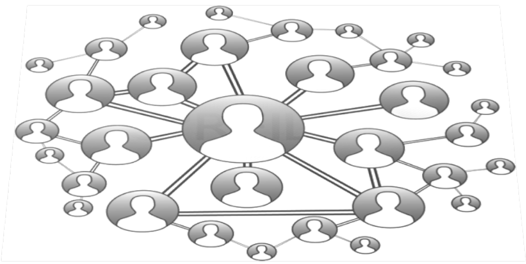

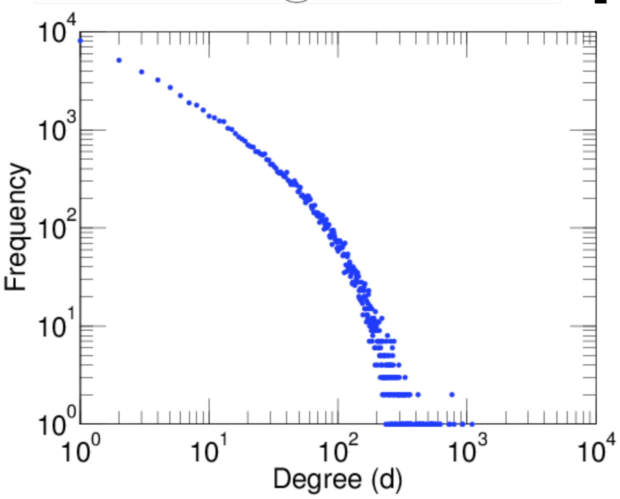

- **Let's connect them all: DOSN overlay that mirrors "social network", i.e., build Friend-to-Friend overlay.**
- **Low latency, Low communication cost**
- **Still global search might be problematic…**
- **Most of social networks have relatively high avg. degree and have power-law degree distributions.**
  - *Issue of handling most popular nodes!*
  - *E.g., tens of millions of links in Twitter and even 5000 friendship links in Facebook!*
  - *E.g, no go for WebRTC*

Challenge: Keep the node degree fixed!

Dilemma…

- Using only Unstructured Overlays (F2F based):
  - Too large node degrees
  - Search/routing is not effective
    - *hard to build other services on top: i.e., storage, recommendation, analytics*
- Using Structured Navigable overlays (DHT based):
  - locality is lost,
  - high relay traffic.
- *Is there a way out?*

Locality Aware Structured P2Ps?

- **Naïve solutions**
  - *Use **DHT** solutions that provide some ”**degrees of freedom**” while selecting neighbors (e.g., Pastry, Symphony etc)*
    - *Castro et al. ”**Topology-aware routing in structured peer-to-peer overlay networks**” 2003*
    - *Antaris et al. “**A Socio-Aware Decentralized Topology Construction Protocol**” HotWeb2015*
    - *Chen et al. “**Design of Routing Protocols and Overlay Topologies for Topic-based Publish/Subscribe on Small-World Networks**” Middleware 2015*
  - *Stick to the ”**rules**”, but choose only those **peers that are more ”useful”***

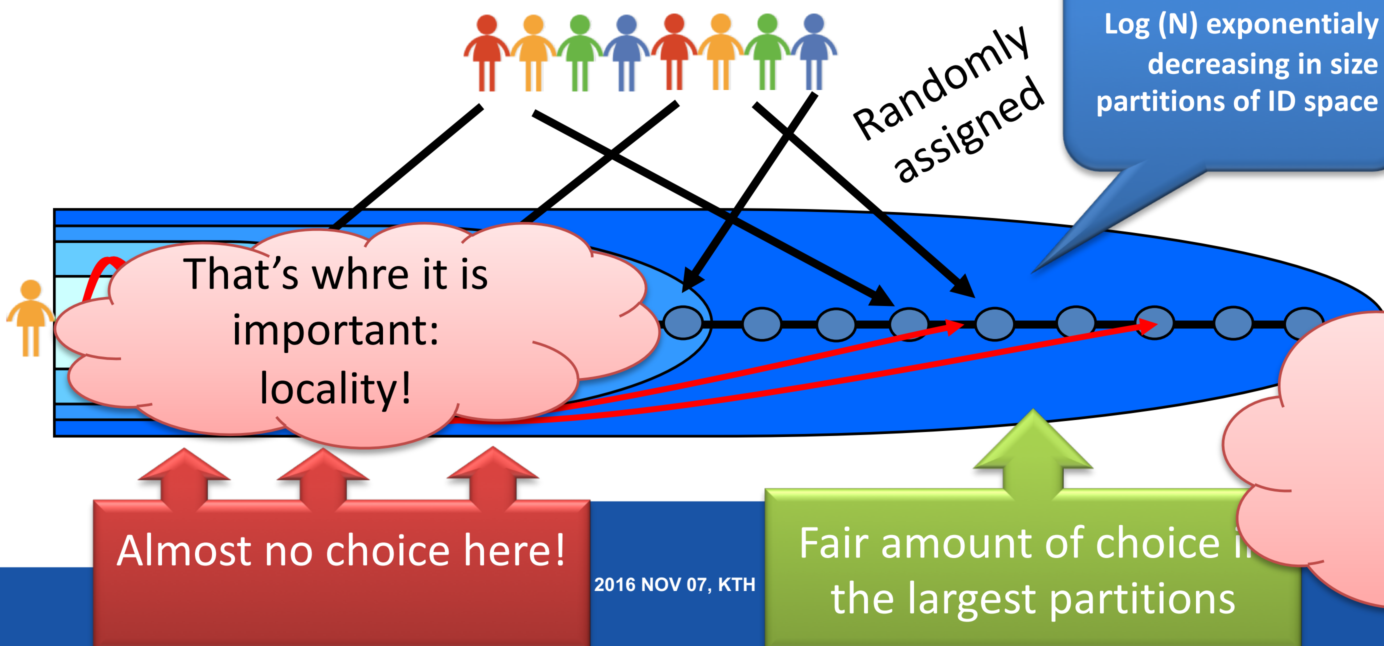

- **Take 1M network and 20 friends**
- **~20 partitions**
  - *1st partition: ~10 friends*
  - *2nd partition : ~5*
  - *~2.5*
  - *on*

F2F social graph vs. Navigable P2P Overlay

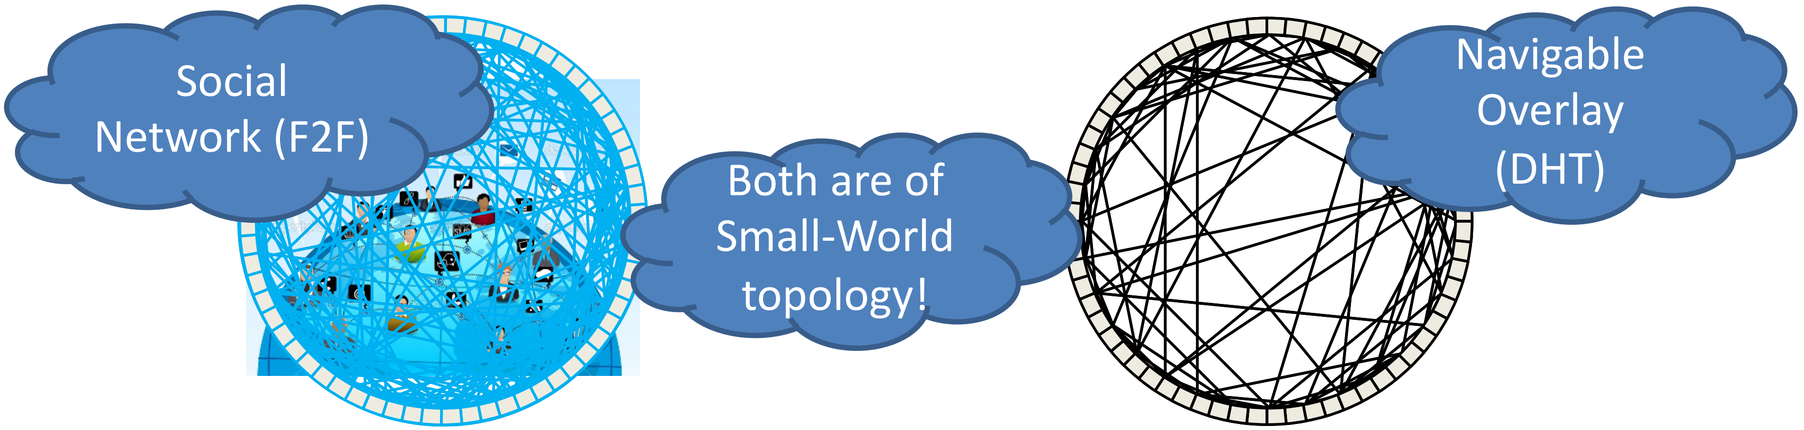

- **Both networks are non-random, but ”small-world” like.**
  - ***High*** **clusterisation,** ***low*** **diameter.**
  - **Differences:**
    - *F2F:* ***power-law degree*** *distribution*
    - *Navigable Overlay:* ***fixed degree*** *distribution*

Making F2F Network Navigable

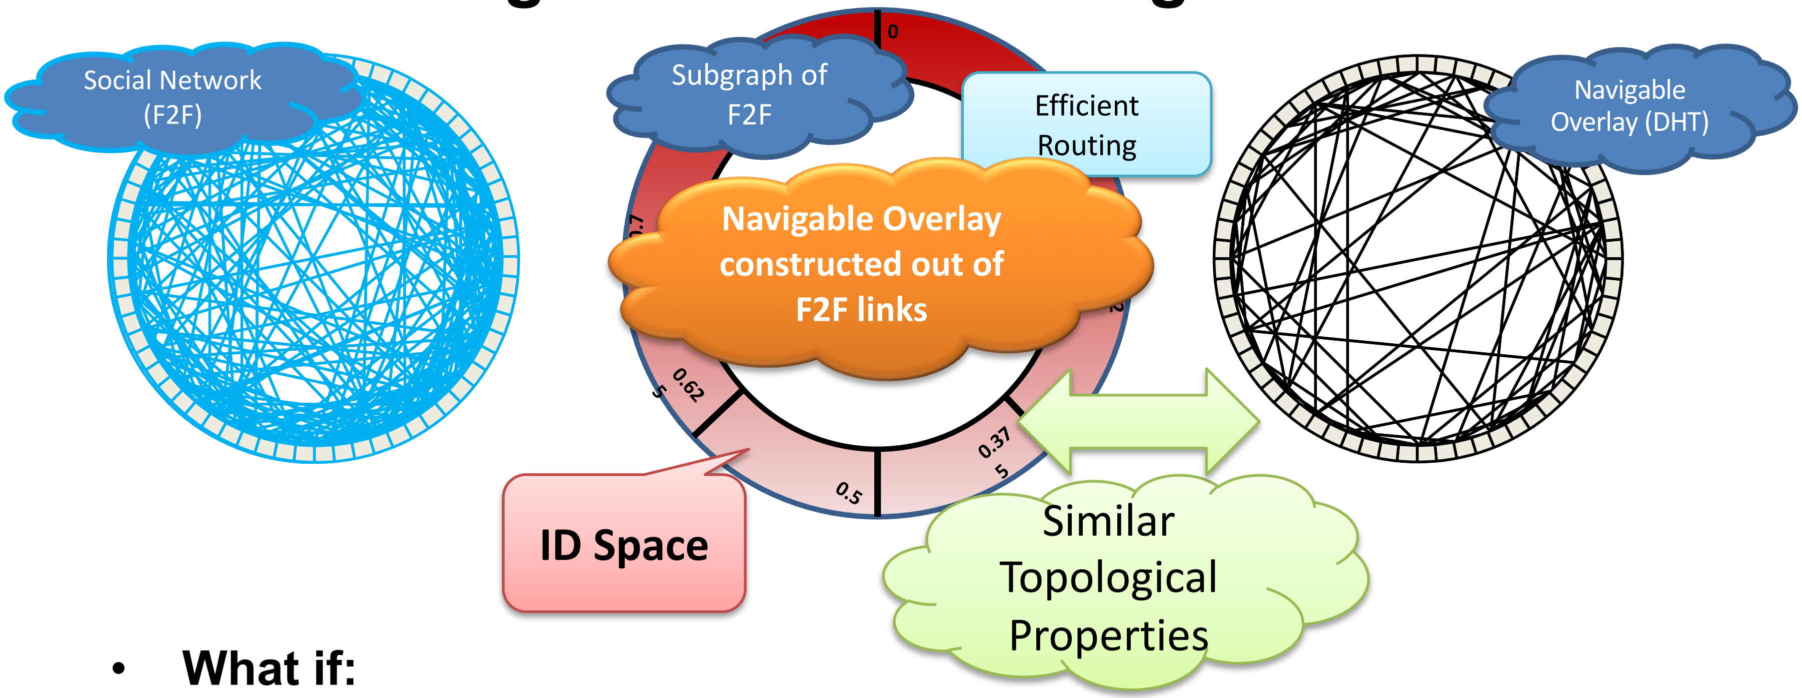

- **What if:**
  - **1) we take a subgraph of F2F with fixed degree that is topologically similar to the graph induced by the Navigable Overlay?**
  - **2) embed it into ID space so that the routing is efficient?**
    - *i.e., ”cleverly assign IDs for each node”.*

How to assign IDs to F2F Network ?

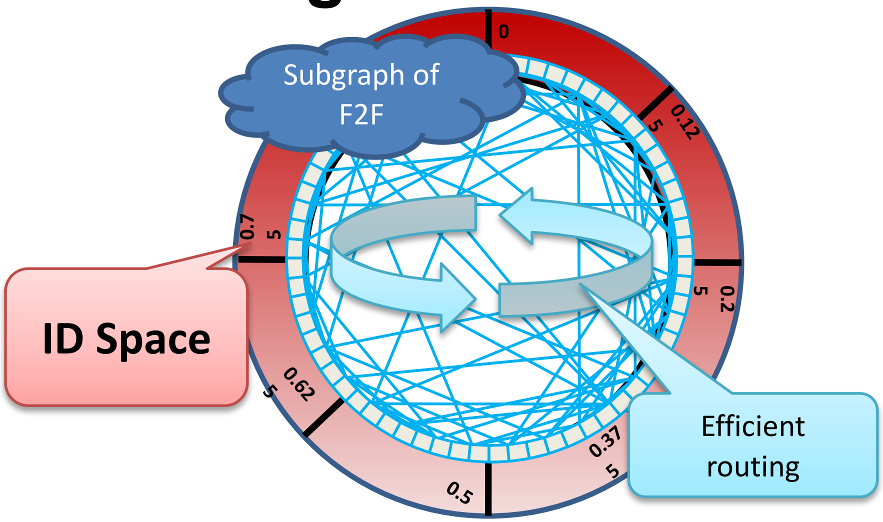

- **Option 1:**
  - **Start with random ID assignment**
    - *Expected poor initial routing performance*
  - **Keep on exchanging IDs between pairs of nodes to "improve" the routing.**
- **Option 2:**
  - **Try to "infer" what are the topological "clusters" in the graph and allocate similar IDs for the nodes in those clusters.**

Option 1: random ID assignment

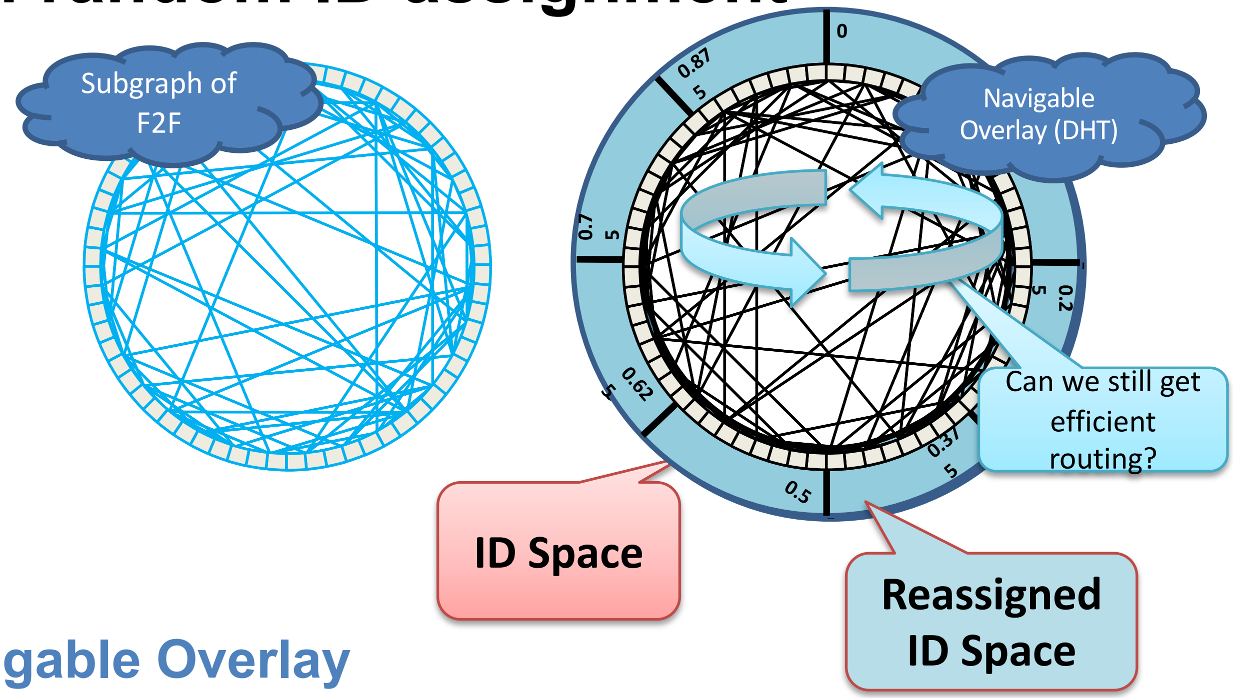

- **Let's start "easier":**
  - **Take existing Navigable Overlay**
  - **"forget" ID space (remove IDs)**
  - **Try to reassign IDs just by looking into the topology of the graph.**
- **NP-hard problem – at least as hard as community detection,**

Proposed Solutions

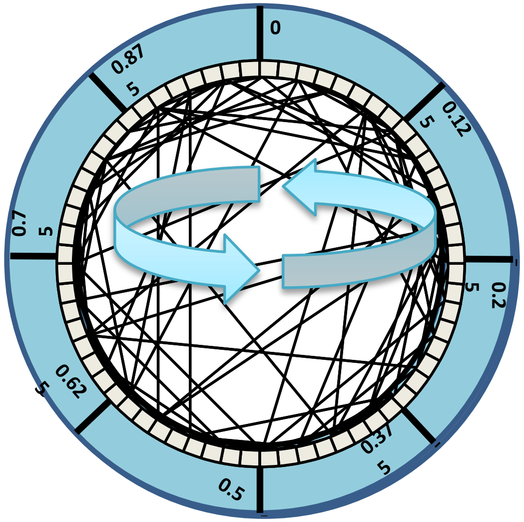

Sandberg et al. "Distributed Routing in Small-World Networks" ALENEX2006

- Each peer gets a random IDs
- Each peer periodically exchanges info of their IDs with a random peer and decides whether to swap the IDs.
- Better than random, but largely fails to "discover" right ID allocation.
- For larger graphs (100k nodes) up to *50% of queries failed* to reach the destination
- Reason: All links are of the same "importance"
  - *Too many "**degrees of freedom**", too **many dependencies**: position improvements in one pair "damages" many other positions.*
- **Our Solution**:
- **The links are not the same** (remember strong and weak links)**?! Treat each link differently!**

Weak and Strong Ties

- **Each node orders all the neighbors by the "strength of their ties"**
  - *Weak vs. Strong links*
  - *E.g., counting friendship triangles, gossip to detect local communities etc.*
- **Graph pruning**:
  - *Consider only **top k strongest** links (representing communities), **ignore weakest links**.*

  - *A graph with **large diameter** emerges,*
    - i.e, less "degrees of freedom" -> **easier to converge** to local optima

Socially-Aware Distributed Hash Tables (Nasir et al, P2P2015)

- **Each peer gets a random IDs**
- **Each peer periodically exchanges info of their IDs with a random peer and decides whether to swap the IDs.**
  - ***Decision***: *based on a cost function that (locally) improves the positions of the two nodes as compared to their neighbors*
  - ***Cost function***: *biases the **preference towards the strong neighbors** (prefer to have strong links with IDs that are as close as possible in the ID space, while disregarding the weak links)*
  - ***Gossip based,*** *Integrated with **Symphony Overlay.***
  - ***Reduces lookup latency by ~30%***
- ***Can we do even better?***
  - *Antaris et al. "SELECT: A Distributed Publish/Subscribe Notification System for Online Social Networks" (collaboration with University of Cyprus)*

Option 2: ID assignment on the fly

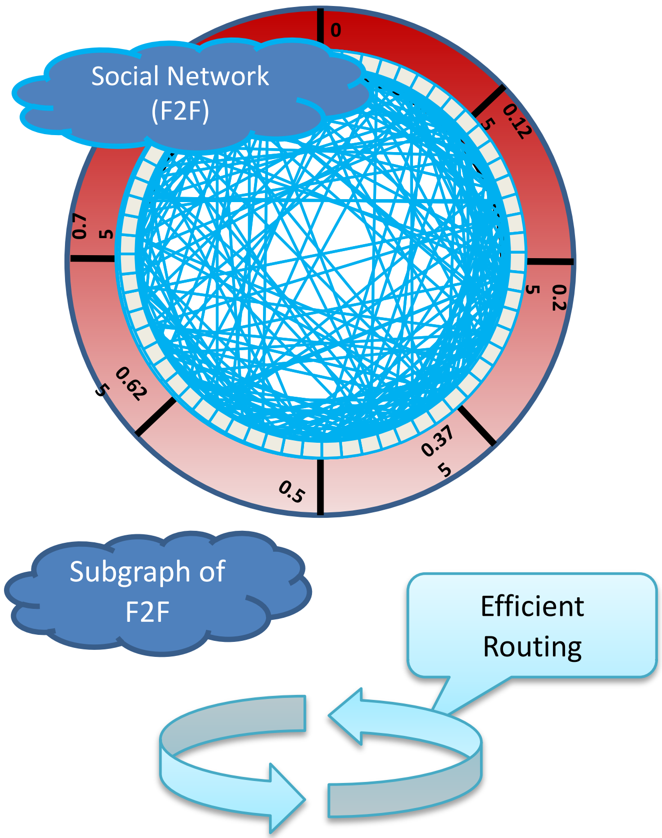

- **F2F graph: nodes arrive with associated edges.**
  - *Bootstrapping: Only first nodes get **random IDs**.*
- **Subsequent arriving nodes identify the “strongest” existing friendship communities and “join” them by selecting ID centered in these communities.**
  - *Arriving node **assumes ID** as a centroid between k strongest links (e.g., nodes that share most of the friends)*
  - *ID ranges identify communities*
- **We have to cap the degree (take a subgraph)**
  - Following Kleinbergian rule: identify most dissimilar communities/regions to point to
  - +Bias toward more reliable nodes
  - +Bias toward particular workloads in F2F graph.

Few extra links: to maintain the ring, and fill up the finger table for nodes with low social degree (people with few friends)

Option 2: ID assignment on the fly (cont.)

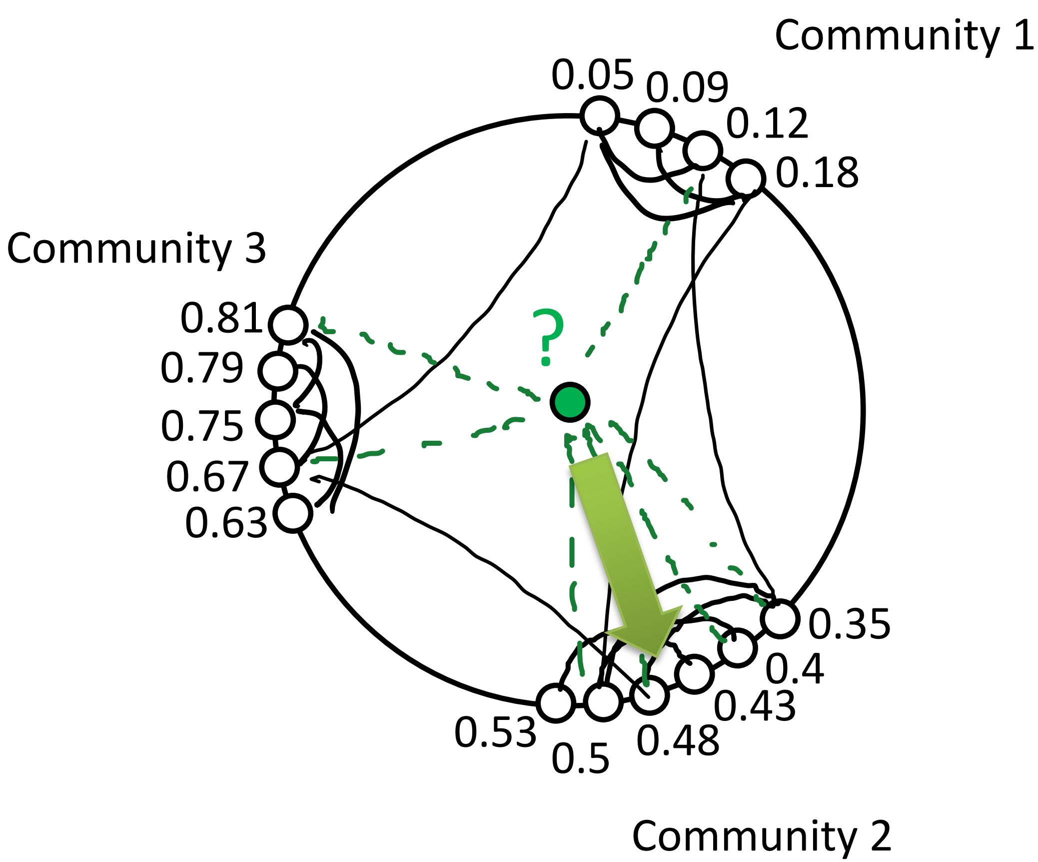

- **Which community should the newcomer join?**
  - *Which ID should it take?*
- **Join between "strongest" existing friendships in a particular community**
  - *e.g., id - a centroid of k-strongest links (nodes that shares most of the friends)*
  - *Communities and Strength discovered by gossiping*

Capping the degree

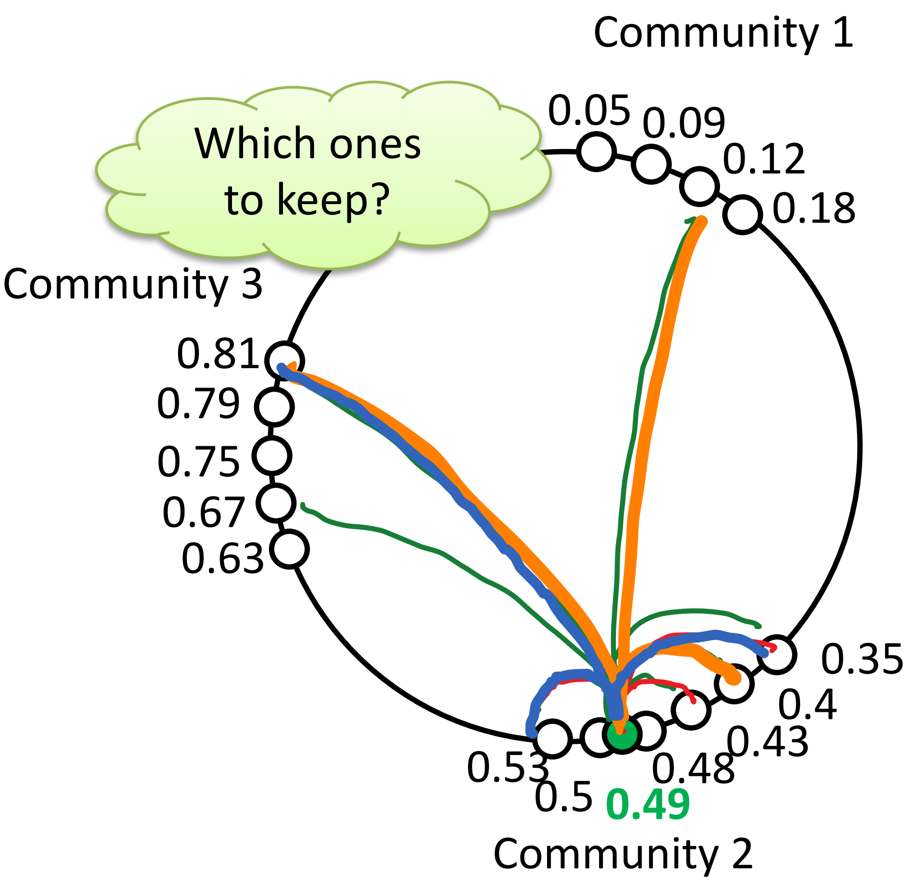

- **Usually a node has way more social friends than "connection quota".**
  - *E.g., max 20 out of 500 friends.*
- **Which ones to keep?**
  - *Only the closest friends?*
    - We lose long range links
  - ***Random friends?***
    - Similar to Watts&Strogats model:
    Small-World but not navigable...
  - ***Ideas from Kleinberg?***
    - Can not apply directly since we do **not want to create new links** and **ID space is not uniform**!
      - Serch performance should be biased toward Friends nodes
      - Detect k friends, representing k-most dissimilar regions (Kleinbergian partitions, e.g., using gossiping or LSH technique) and establish connections to them
- **Few extra links if necessary**

SELECT: A Distributed Publish/Subscribe Notification System for Online Social Networks

- Some Results:

More shallow dissemination trees (up to 3x reduction on Delay)

Almost no relay nodes (Improvement on message overhead)

	…le+ (110k nodes, 12M …es, avg degree 157)	T… …s,
	…ize of dissemination tree	
DHT (Symphony)	…468	869
SELECT	169 (↓64%)	…3 (↓84%)
	Average Relay nodes	
DHT (Symphony)	311	769
SELECT	12 (↓96%)	16 (↓98%)

Take Aways

- **The World is not only Small (6-degrees of separation), but also navigable in a completely decentralized fashion.**

- **Community aware assignment of IDs and selection of links enables efficient navigation in F2F networks**
  - *in turn improving **privacy/security**, getting rid of majority of **relay traffic** and allowing **locality aware** services*

Bibliography

- **References**

Overlay navigation, dissemination and pub/sub

1. Muhammad Anis Uddin Nasir, Nicolas Kourtellis, Sarunas Girdzijauskas: Highly Scalable, Heterogeneous and Reliable Overlays for Decentralized Online Social Networks. P2P 2015
2. Mansour Khelghatdoust, Sarunas Girdzijauskas: Short: Gossip-Based Sampling in Social Overlays. NETYS 2014: 335-340
3. Fatemeh Rahimian, Thinh Le Nguyen Huu, Sarunas Girdzijauskas: Locality-Awareness in a Peer-to-Peer Publish/Subscribe Network. IFIP International Conference on Distributed Applications and Interoperable Systems, DAIS 2012: 45-58
4. Fatemeh Rahimian, Sarunas Girdzijauskas, Amir H. Payberah, Seif Haridi: Vitis: A Gossip-based Hybrid Overlay for Internet-scale Publish/Subscribe Enabling Rendezvous Routing in Unstructured Overlay Networks. IPDPS 2011: 746-757
5. Fatemeh Rahimian, Sarunas Girdzijauskas, Amir H. Payberah, Seif Haridi: Subscription Awareness Meets Rendezvous Routing, The Fourth International Conference on Advances in P2P Systems, AP2PS 2012.
6. Sarunas Girdzijauskas, Gregory Chockler, Ymir Vigfusson, Yoav Tock, Roie Melamed: Magnet: practical subscription clustering for Internet-scale publish/subscribe. DEBS 2010: 172-183
7. Sarunas Girdzijauskas, Gregory Chockler, Roie Melamed, Yoav Tock: Gravity: An Interest-Aware Publish/Subscribe System Based on Structured Overlays, [fast abstract] DEBS 2008.
8. Fabius Klemm, Sarunas Girdzijauskas, Jean-Yves Le Boudec, Karl Aberer: On Routing in Distributed Hash Tables. Peer-to-Peer Computing 2007: 113-122
9. Anwitaman Datta, Sarunas Girdzijauskas, Karl Aberer: On de Bruijn Routing in Distributed Hash Tables: There and Back Again. Peer-to-Peer Computing 2004: 159-166
10. Sarunas Girdzijauskas, Wojciech Galuba, Vasilios Darlagiannis, Anwitaman Datta, Karl Aberer: Fuzzynet: Ringless routing in a ring-like structured overlay. Peer-to-Peer Networking and Applications 4(3): 259-273 (2011)

- **References**

Overlay Construction

1. Sarunas Girdzijauskas, Anwitaman Datta, Karl Aberer: Oscar: A Data-Oriented Overlay For Heterogeneous Environments. ICDE 2007: 1365-1367
2. Sarunas Girdzijauskas, Anwitaman Datta, Karl Aberer: Oscar: Small-World Overlay for Realistic Key Distributions. DBISP2P 2006: 247-258
3. Sarunas Girdzijauskas, Anwitaman Datta, Karl Aberer: On Small World Graphs in Non-uniformly Distributed Key Spaces. ICDE Workshops 2005: 1187
4. Karl Aberer, Luc Onana Alima, Ali Ghodsi, Sarunas Girdzijauskas, Seif Haridi, Manfred Hauswirth: The Essence of P2P: A Reference Architecture for Overlay Networks. Peer-to-Peer Computing 2005: 11-20
5. Sarunas Girdzijauskas, Anwitaman Datta, Karl Aberer: Structured overlay for heterogeneous environments: Design and evaluation of oscar. TAAS 5(1) (2010)

Thank you!