Sample work in Networking at UoA

NETWORK RESOURCE ALLOCATION

From bandwidth and storage allocation (80's, 90's)
to
networked content management (90's, 00's)
(placement/replication/caching/forwarding)
to
exploiting social layer (late 00's - present)

Ioannis Stavrakakis
Deputy Dept Chair
Director of Graduate Studies
http://www.di.uoa.gr/~ioannis
National & Kapodistrian University of Athens

Recent work supported by EU FET projects SOCIALNETS (2008-11) and RECOGNITION (2010-13)

Sample of recent Research Topics

- ❑ **Distributed Selfish Replication**
- ❑ **Content placement**
- ❑ **Opportunistic Content forwarding**
- ❑ **Resource Selection**

Social layer comes in through:

- ➢ self-awareness and selfishness – cooperation / competition
- ➢ interest- and location- based social structures / grouping
- ➢ centrality and other social metrics

Distributed Selfish Replication

Distributed Selfish Replication (DSR)

access cost: $t_l < t_r < t_s$

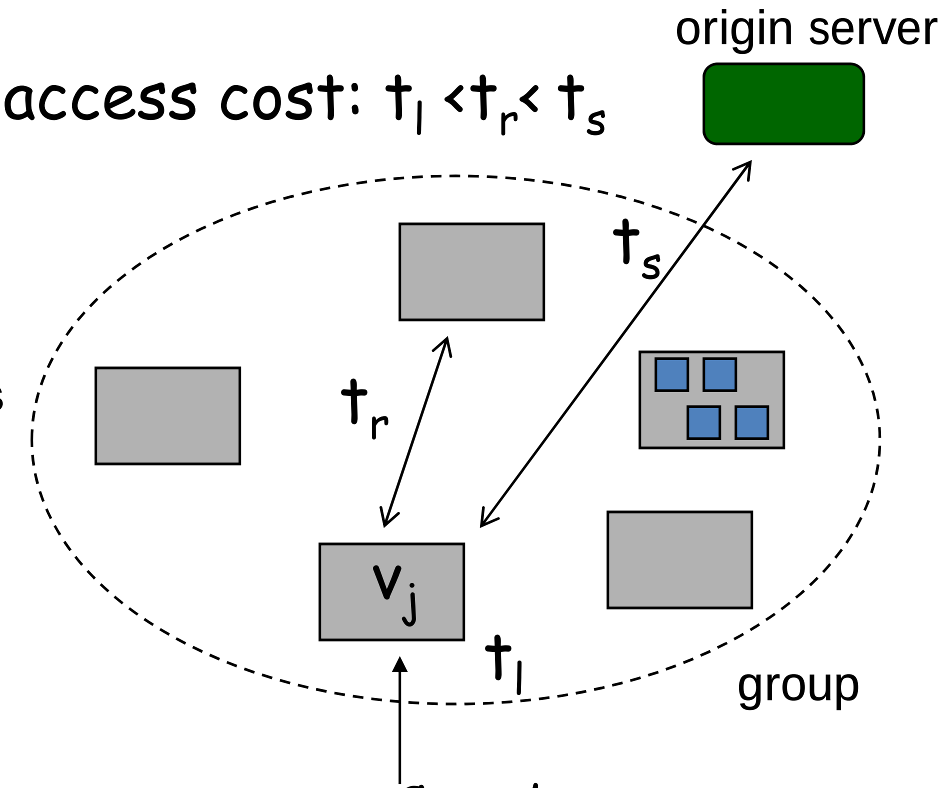

Applications
Content distribution
Shared memory
Network file systems

- n nodes
- N objects

C_j: v_j's storage capacity
r_{ij}: v_j's request rate for obj. o_i

DSR - contribution

- Formulated the *Distributed Selfish Replication* (DSR) game (*n*-**player, non-cooperative, non-zerosum game)**

- Obtained pure Nash equilibrium (EQ) replication strategies
  - a node is **guaranteed to be better off** than by following Greedy Local (isolated) - has no reason not to participate in EQ as it can only benefit from it
  - Call it: **Self-Aware Cooperative (SAC)** replication strategy

- Develop simple protocol for implementing our algorithm, which
  - is **distributed** and requires **only IDs of some of the stored objects**
  - Optimal solution would require transmission of request vectors (→ major OH plus unrealistic to know and hard to estimate)

N. Laoutaris, O. Telelis, V. Zissimopoulos, I. Stavrakakis, **"Distributed Selfish Replication,"** IEEE Transactions on Parallel and Distributed Systems, Dec., 2006

DSR - Other relevant social attributes

Besides Cooperation / Selfishness-Self-awareness (no mistreatment), consider the impact of other factors (behaviours)

A. Node uncertainties – churn

B. Node interest similarities

Effective content replication strategies should factor-in such attributes

(A) DSR - Impact of Node misbehaviour/churn

- Mistreatment-free property of SAC strategy lost under churn
- Design a churn-aware (CA-SAC) strategy:

  (nodes change their initial placements to minimize their imminent access cost, considering the availability probabilities (πi) of other nodes)
- Under an LRF (Less Reliable First) order, the churn-aware strategy **significantly improves performance** – Nash EQ not shown
- for N=2 nodes, CA-SAC is **mistreatment-free**

  (the 2nd node only evicts objects belonging to the 1st node, so the access cost of node 1 is not decreased) - for N≥3, mistreatment may occur)
- In homogeneous case (Cj = C, rij ~ ri for all i, j), under LRF
  - churn-aware strategy is **mistreatment-free under certain conditions.**

E. Jaho, I. Koukoutsidis, I. Stavrakakis, I. Jaho, "Cooperative Replication in Content Networks with Nodes under Churn", IFIP Networking 2008.

E. Jaho, I. Koukoutsidis, I. Stavrakakis, I. Jaho "Cooperative Content Replication in Networks with Autonomous Nodes", Computer Communications Journal, to appear

(B) DSR - Impact of group similarity

<u>**Q.**</u> How is Interest Similarity shaping benefits/costs of Selfish, Cooperative or Altruistic (social benefit max) behaviour?

Factor in implementation cost of each strategy:

- **Implementation cost** (altruistic → SAC → selfish)
- **Mistreatment (**optimally altruistic / SAC and selfish)
- **Overall Performance**

<u>**A.**</u>

For very tight (high degree of similarity) groups, no need for self-concerned. The <u>altruistic</u> behavior (yields the best performance for both the entire group and the individual nodes (no mistreatment))

When the group is made up of <u>**foreigners**</u> <u>(</u>almost no similarity)<u>,</u> no need to be socially-concerned. Altruism or cooperation cannot bring significant benefits to either the group or the individuals, and thus a <u>selfish</u> behavior would make sense due to its simplicity

(B) DSR - Impact of group similarity

Defining & Exploiting Interest Similarity / ISCoDE framework[1]

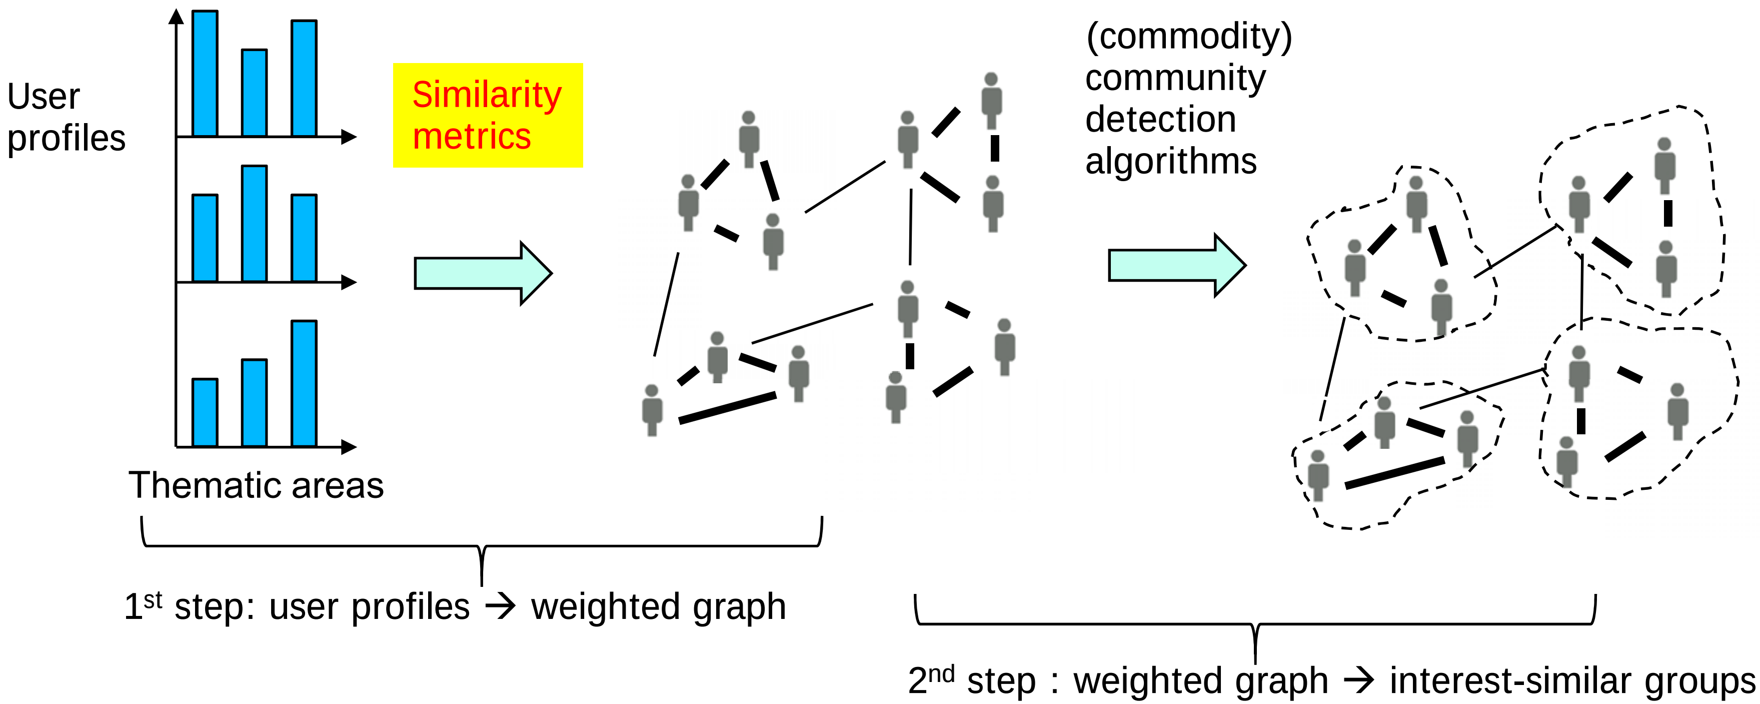

[1]E. Jaho, M. Karaliopoulos, I. Stavrakakis. ISCoDe: a framework for interest similarity-based community detection in social networks. Third International Workshop on Network Science for Communication Networks (INFOCOM-NetSciCom'11), Apr. 10-15, 2011, Shanghai.

Similarity metrics: PS and InvKL (Kullback Leibler)

- Proportional Similarity (PS)
  - PS : {F^i , F^j} → [0,1]

$$PS(F^i, F^j) = 1 - \frac{1}{2}\sum_{m=1}^{M}\left|F^i{}_m - F^j{}_m\right|$$

- Inverse symmetrized KL divergence
  - InvKL : {F^i, F^j} → (0,∞)

$$InvKL(F^i, F^j) = \frac{1}{\sum_{m=1}^{M} F^i{}_m \log\frac{F^i{}_m}{F^j{}_m} + \sum_{m=1}^{M} F^j{}_m \log\frac{F^j{}_m}{F^i{}_m}}$$

$F^n{}_m$, 1≤n≤N, 1≤m≤M : distribution of node n over interest class m

Example with M=2 interest classes and N=2 nodes

Low resolution for similar distributions

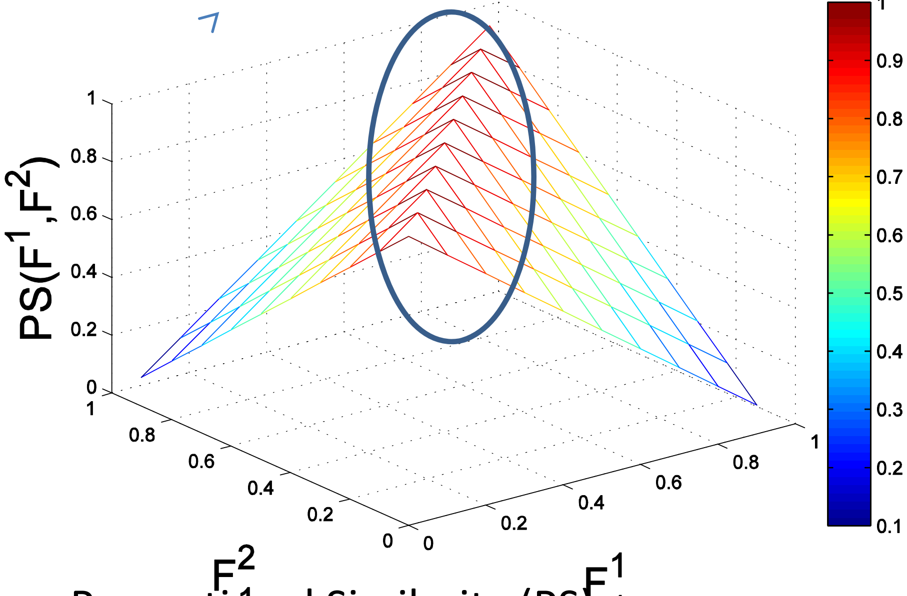

High resolution for similar distributions

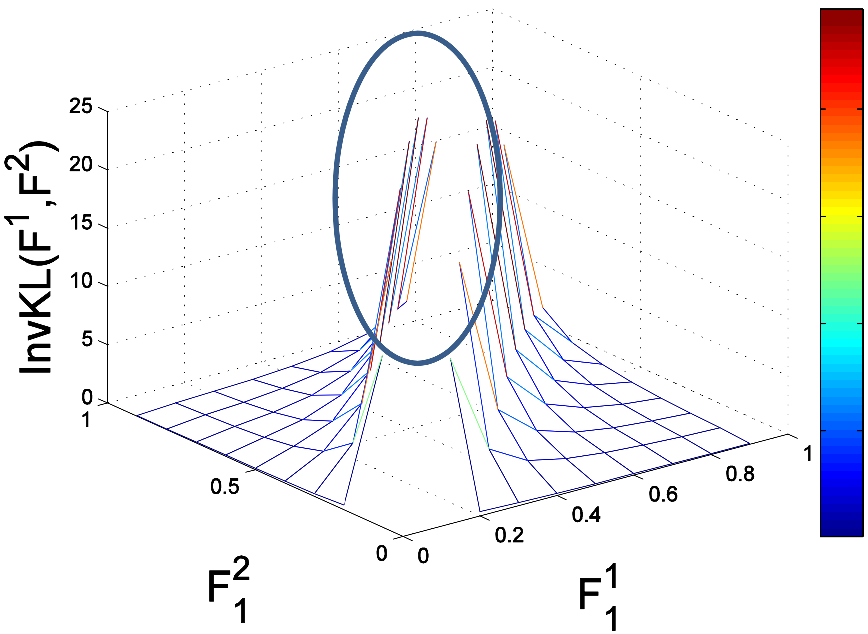

- Proportional Similarity (PS)
- Inverse symmetrized KL divergence (InvKL)

Resolution performance

(a) Similar nodes

PS			InvKL		
Q	C	partition	Q	C	partition
0.0215	2	{1..38} {39..80}	0.6740	5	{1..14} {15..28} {29..44} {45..61} {62..80}

(b) Dissimilar nodes

PS			InvKL		
Q	C	partition	Q	C	partition
0.7860	10	{1..8}..{73..80}	0	1	{1..80}

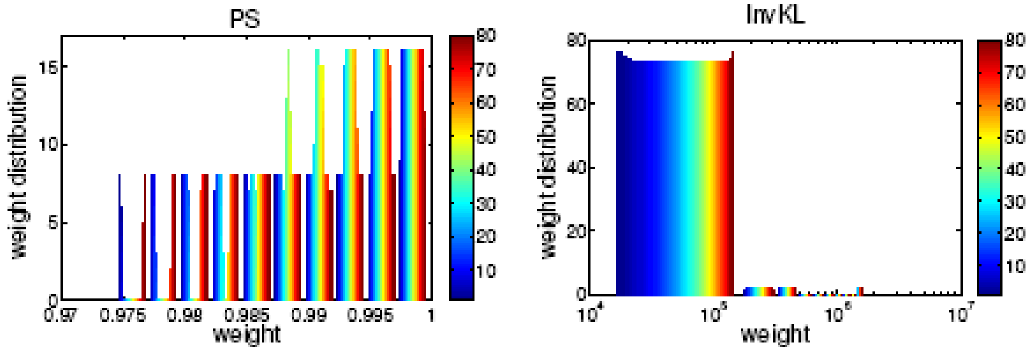

(a) Similar nodes

InvKL can identify smaller and more similar communities than PS, in a highly similar network

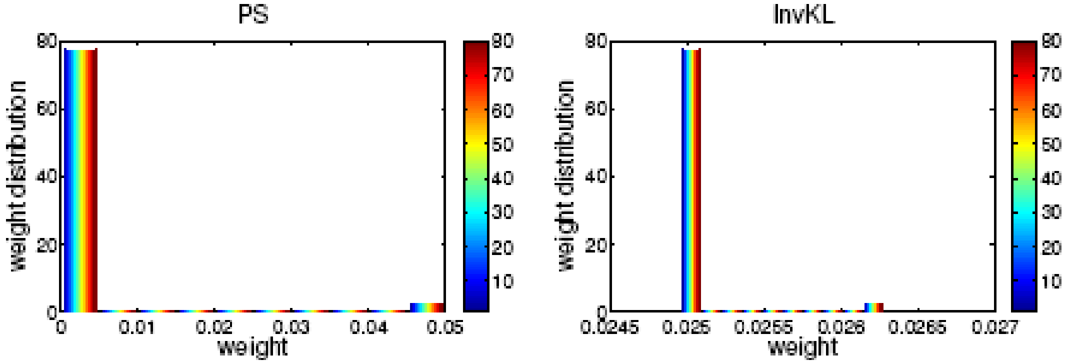

(b) Dissimilar nodes

PS can identify smaller communities than InvKL, in a highly dissimilar network

(could argue that this is not very useful)

DSR - Impact of group similarity

the benefits of cooperation increase with the group tightness, therefore tightness should be used as a decision criterion:

- when choosing content placement strategies under given group membership; or,
- for carrying out performance-driven group management operations such as group formation/merging/splitting.

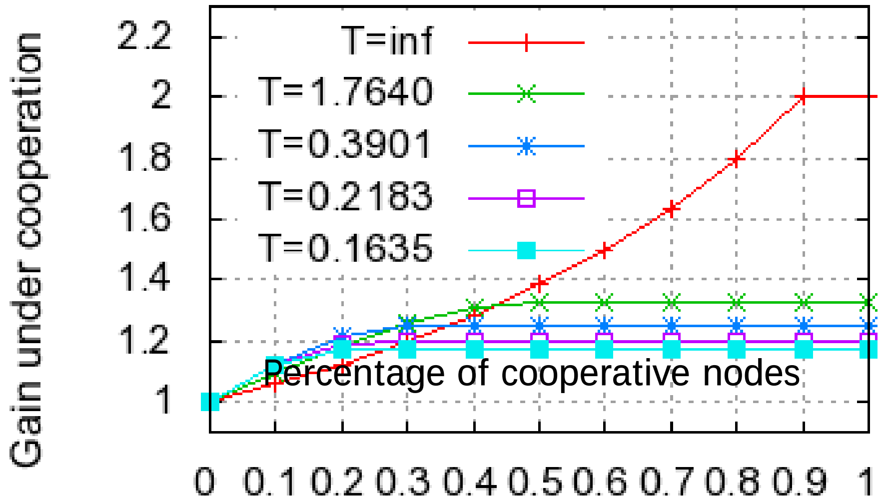

- T : tightness metric (=mean invKL), measuring interest similarity across group members

[1] E. Jaho, M. Karaliopoulos, I. Stavrakakis, “Social similarity as a driver for selfish, cooperative and altruistic behavior”, in Proc. AOC 2010 (extended version submitted to IEEE TPDS)

Distributed Content Placement

The classic facility location problem*

K-median problem formulation:

Given a fixed number of facilities, minimize the total service cost

$$Cost(\mathcal{F}) = \sum_{n \in \mathcal{V}} w_n \cdot min_{x_j \in \mathcal{F}} \{d(x_j, n)\}$$

Topology represented by ***G(V,E)***

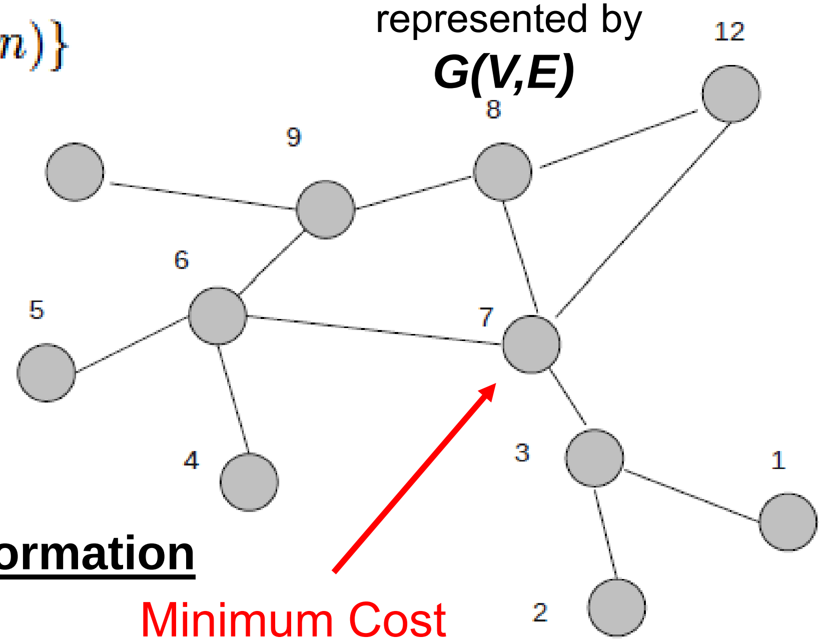

d **(*k, n*) : cost path between nodes *k* , *n***

w_n : demand generated by node n

<u>Requires full topology and demand information</u>

---

* **P. Mirchandani and R. Francis. *Discrete Location Theory*. John Wiley and Sons, New York, NY, 1990**

Approach 1: neighbor-hopping service migration

- <u>**Exploits strictly local information**</u>
- Bases service facility movement on the relative difference of aggregate service demands

Monotonic cost reduction for any (connected) topology
- Response to dynamic changes

Optimum is guaranteed
- for one service facility and in unique shortest path tree topologies (e.g., trees)

Multi-facility extension
- Monotonic cost reduction
- Optimality not guaranteed

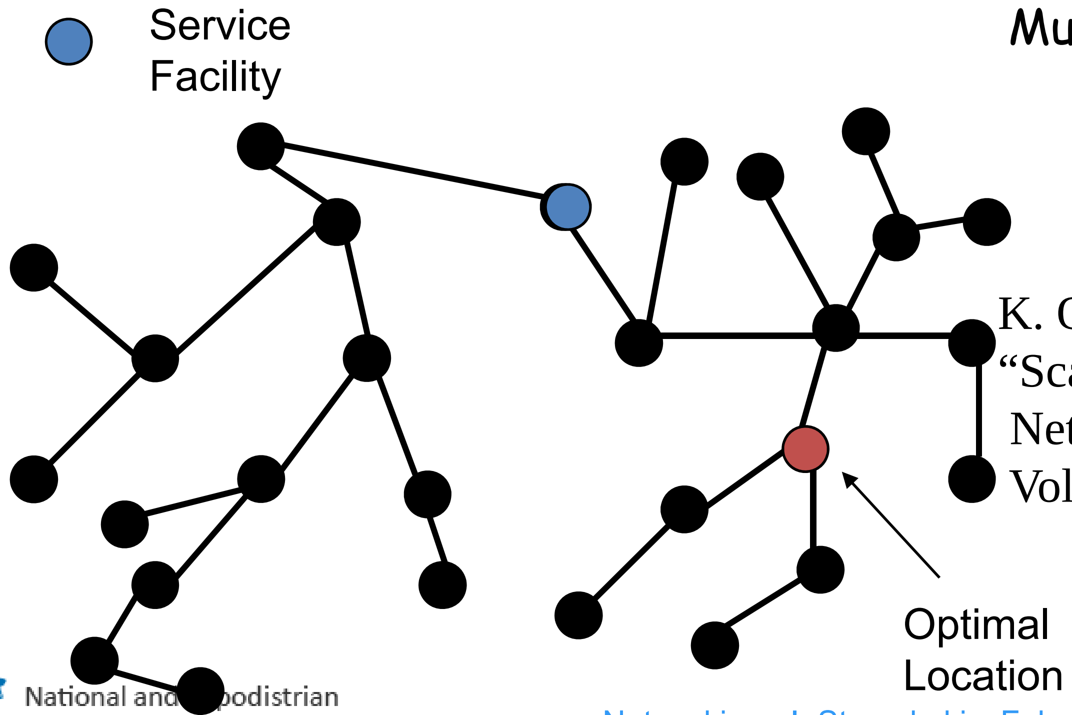

K. Oikonomou, I. Stavrakakis, "Scalable Service Migration in Autonomic Network Environments," IEEE JSAC, Vol. 28, No. 1, January 2010, pp. 84-94

Approach 2: R-neighborhood service migration

- Solve iteratively small-scale optimization problems **<u>within a limited neighborhood</u>** of the current facility location

- **Map the outside demand on the surface of the R-ball.**

- Real ISP traces used for validation

- Good approximation and fast convergence *

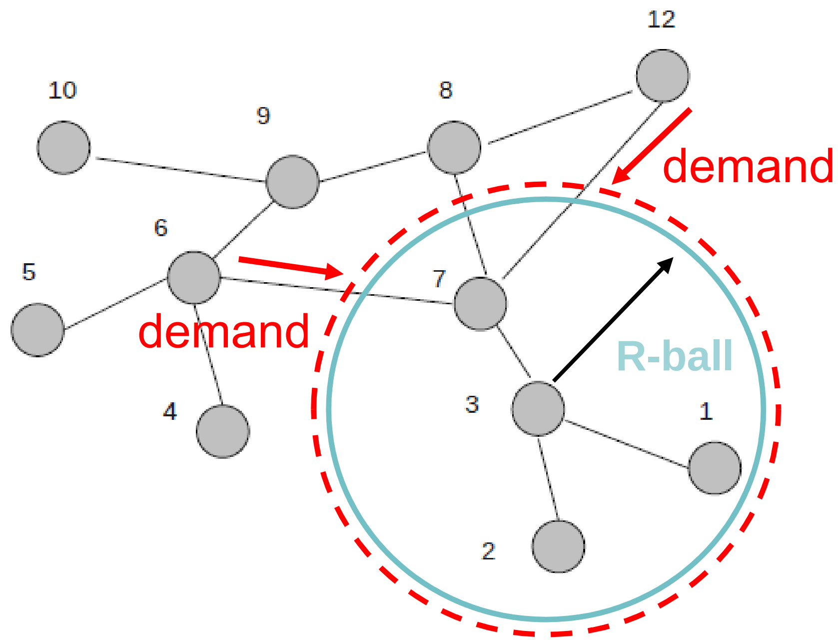

* G. Smaragdakis, N. Laoutaris, K. Oikonomou, I.Stavrakakis, A. Bestavros, “Distributed Server Migration for Scalable Internet Service Deployment”, to appear in IEEE/ACM T-Net. (to apear) , also in INFOCOM2007

Approach 3: Generalized neighborhood service migration

Since physical and topological localities become less relevant, consider Generalized neighborhood: set of nodes with highest values <u>wrt specific social metric</u>

- Solve *iteratively* **small-scale** ones, on subgraphs G^i ∈ G, around the current facility location of $host^i$ containing the **top nodes based on a social-inspired metric related to $host^i$**

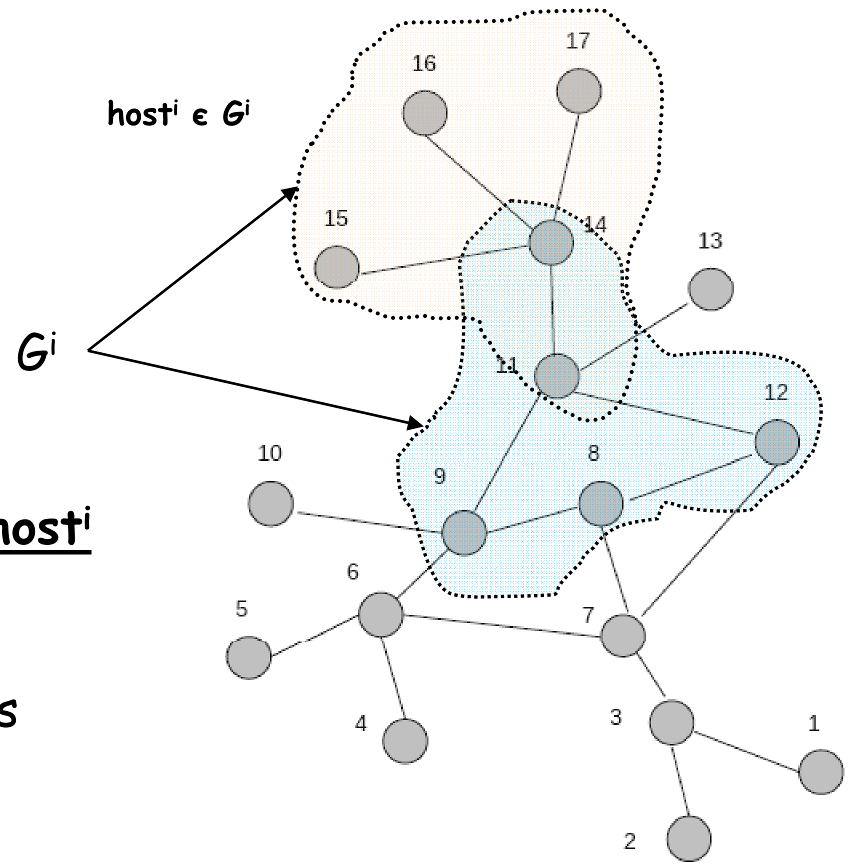

- Map the outside demand properly on nodes in subgraphs G^i I

The social metric

a measure of the importance of node's *u* social position : lies on paths linking others

Betweenness Centrality (*u*): portion of all pairs shortest paths of G that pass through node *u*

$$BC(u) = \sum_{s=1}^{|V|} \sum_{t=1}^{s-1} \frac{\sigma_{st}(u)}{\sigma_{st}}$$

Conditional **Betweenness Centrality (*u, t*)** : portion of all shortest paths of G from node *u* to *target t,* that pass through node *u*

$$CBC(u;t) = \sum_{s \in V, u \neq t} \frac{\sigma_{st}(u)}{\sigma_{st}}$$

a measure of the importance of node's *u* social position : ability to control information flow towards *target* node

Projecting the "world outside" on the selected nodes

- *wCBC* metric facilitates the demand mapping of the (world outside) $G\backslash G^i$ nodes, on the selected G^i ones

- Nodes in G^i exhibit an *effective demand :*

$$w_{eff}(n;Host) = w(n) + w_{map}(n;Host)$$

- *Where*

$$w_{map}(n;t) = \sum_{s \in \{G \backslash G^i_{Host}\}} w(s) \frac{\sigma'_{st}(n)}{\sigma_{st}}$$

- The influence of node z є $G\backslash G^i$ is "credited" *only to the first* G^i node encountered on each shortest path from z towards the host.

- the demand of G^i nodes is subtracted when computing the effective demand, with which each node participates in the of the *1-median* solution.

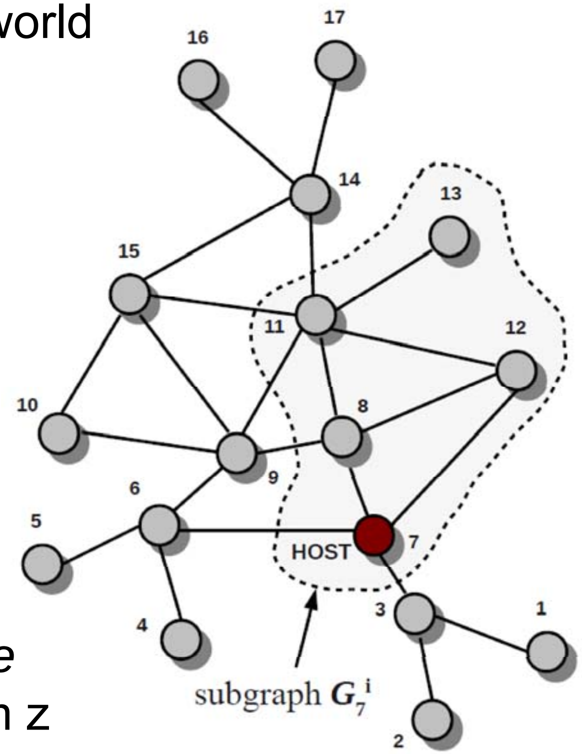

The R-neighborhood VS the generalized neighborhood

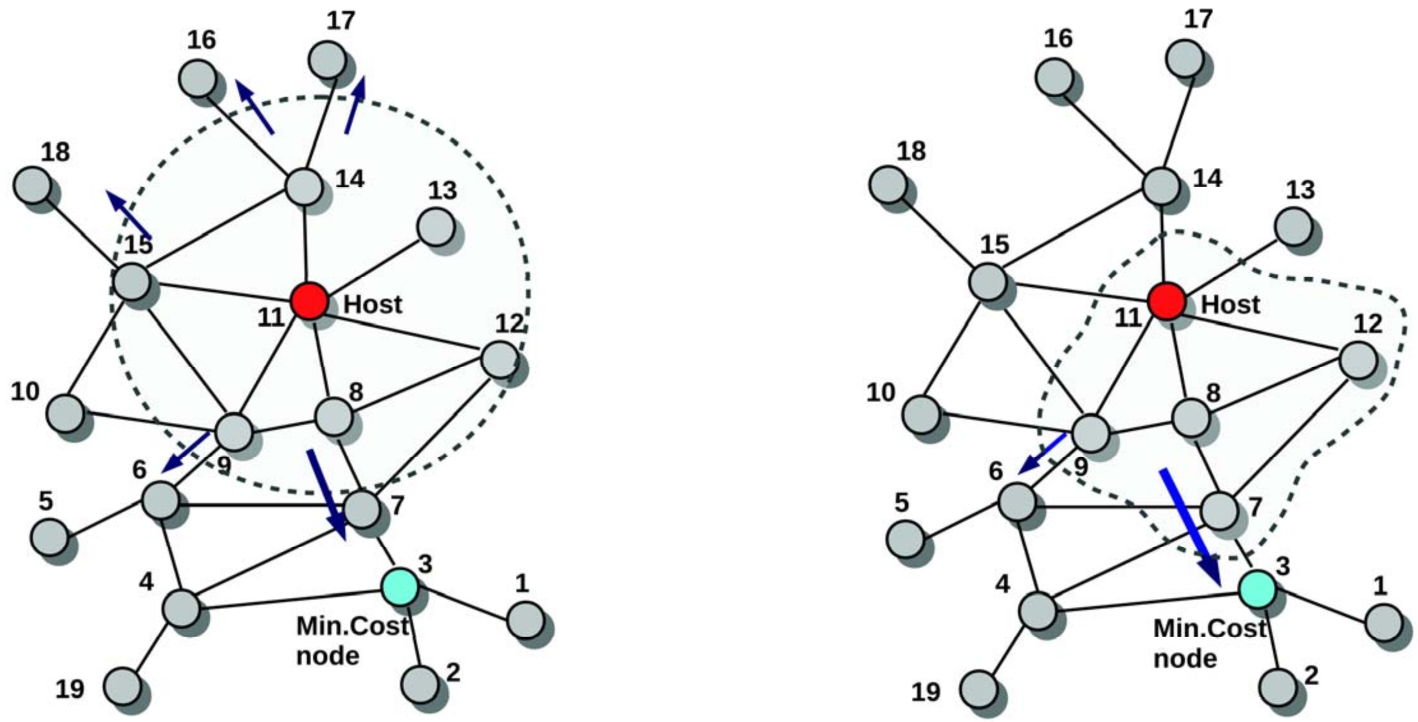

“Detect” the pulling forces while looking for next-best location

simulation results: ISP topologies / non-uniform load

Less than a dozen of nodes is enough!

Demand load : Zipf distribution (with skewness *s*)

$$\beta_{alg}(\alpha; G, \overline{w}) = E[\ \frac{C_{alg}(\alpha; G, \overline{w})}{C_{opt}(\alpha; G, \overline{w})}\ ]$$

Datasets correspond to different snapshots of 7 ISPs collected by `mrinfo` multicast tool *

$$\alpha_\epsilon = argmin\ \{\alpha | \beta_{alg}(\alpha) \leq\ (1+\epsilon)\}$$

Table 1.2 Results derived by the Generalized Neighborhood Service Migration Strategy

Size of physical topology	Min. Subgraph size for solutions within 2.5% of the optimal $s=0$	$s=1$
76	4	4
100	5	5
180	5	4
184	4	4
216	4	4
339	7	6
378	5	5

* J.-J. Pansiot, P. Mérindol, B. Donnet, and O. Bonaventure, “Extracting intra-domain topology from mrinfo probing,” in Proc. Passive and Active Measurement Conference (PAM), April 2010.

Opportunistic Content forwarding

Opportunistic, encounter-based Content forwarding

Problem: efficient routing of UGC/message in unstructured, time-varying nets (via mobility-enabled node encounters)

Epidemic routing: (unlimited copies in net)

- A copy is forwarded without erasure upon encounter
- Optimal wrt delay at high (BW/memory/…) overhead

Two-hop routing/relaying: (limited copies in net)

- A limited (fixed, N) number of copies generated by source
- One copy is forwarded with erasure upon node encounter
  - only the source relays (is responsible for distributing) more than one copy

Binary spray-and-wait routing (BSW): (limited copies in net)

- A limited (fixed, N) number of copies generated by source
- Half of the copies are forwarded with erasure upon node encounter
  - intermediate nodes and source can relay (are responsible for distributing) more than one copy

Opportunistic Content forwarding

A. Panagakis, A. Vaios, I. Stavrakakis, “On the Performance of Two-Hop Message Spreading in DTNs”, Ad-Hoc Networks Journal, Vol. 7, Iss. 6, pp. 1082-1096, Aug. 2009.

- ✓ Analytic, closed-form for delay distribution (via bounding performance auxiliary systems)
- ✓ Analysis of spreading-stopping reverse process

Opportunistic Content forwarding

A. Panagakis, A. Vaios, I. Stavrakakis, “On the Effects of Cooperation in DTNs”, IEEE COMSWARE, Jan. 8-12, 2007, Bangalore, India

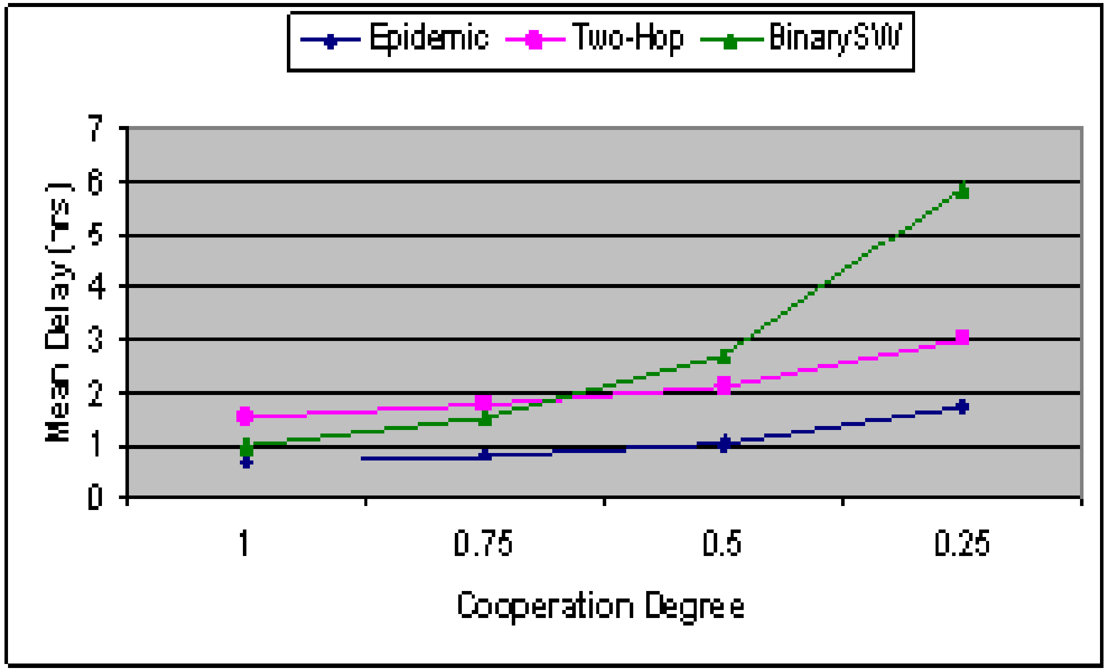

2-hop relaying outperforms BSW for low cooperation degree
(BSW’s copy distribution process involved heavier the intermediate, potentially non-cooperating nodes)

Assessing centrality-driven routing in opportunistic nets

(SimBetTS / BubbleRap use BC values of encounters for content forwarding)

How is performance of centrality-based routing affected by

- ❑ Adding or not, destination awareness to BC (BC vs CBC)
- ❑ Working with ego-centric vs socio-centric BC values
- ❑ Type of contact graph (unweighted vs. weighted) ?

P. Nikolopoulos, et.al. “How much off-center are centrality metrics for opportunistic routing?”, CHANTS 2011 Workshop (in MobiCom), Sept 23, 2011, Las Vegas

Datasets

5 well-known iMote-based real traces available from the Haggle Project at CRAWDAD.

CHARACTERISTICS OF EMPLOYED DATASETS

Configuration	Intel	Cambridge	Infocom05	Content	Infocom06
Device type	iMote	iMote	iMote	iMote	iMote
Network type	B/T	B/T	B/T	B/T	B/T
Duration (days)	6	6	4	24	4
Scan time (sec)	5-10	5-10	5-10	5-10	5-10
Granularity (sec)	120	120	120	120-600	120
Mobile Devices	8	12	41	36	78
Stationary Dev.	1	0	0	18	20
External Dev.	119	211	233	11368	4421
Average internal contacts/pair/day	9.09	12.09	8.60	0.66	9.03
# of Contacts	2766	6732	28216	41330	227657

BC vs CBC

- opt ➔ optimal routing through knowledge of contact sequences.
- BC/CBC ➔ up to 30% of messages never reach their destination
  ➔ about 5 times more hops and 1 day of additional delay

BC outperforms CBC in delay
(due to zero CBC values when destination in an unconnected cluster)

CBC outperforms BC in hops (up to 50% shorter paths, due to selecting more proper nodes to forward to)

<u>Since delay is not the concern, CBC better as it conserves resources (BW, memory, energy)</u>

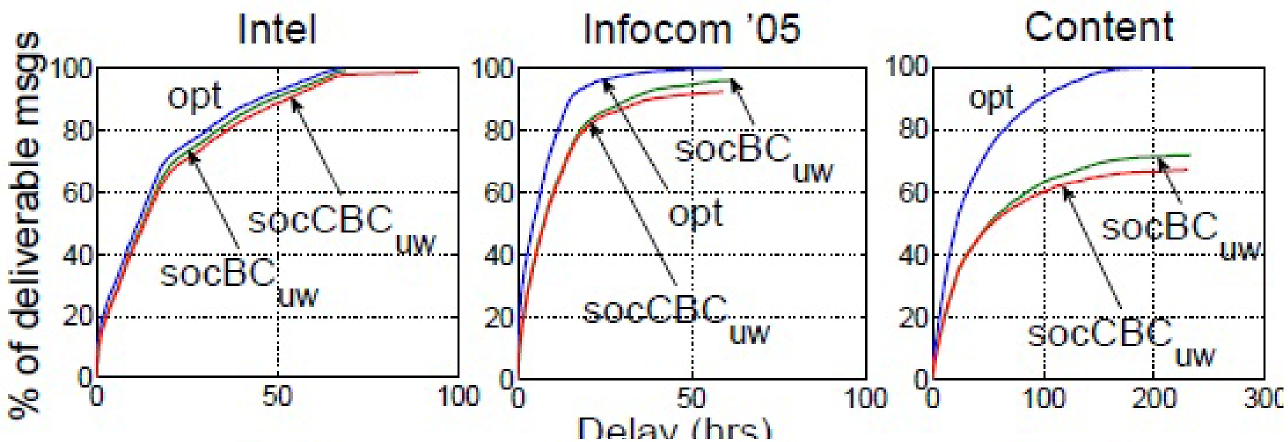

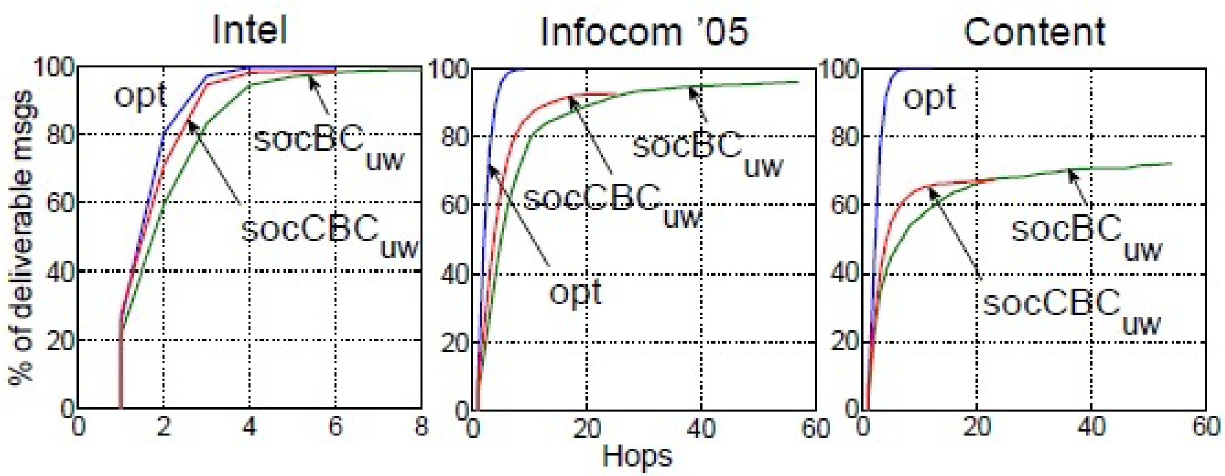

National and Kapodistrian UNIVERSITY OF ATHENS

Resource Selection Problem

The parking spot selection problem

- The environment:
  - Scarce cheap **on-street public** parking spots
  - (more) Expensive **private** parking facilities
  - Selfish cost-minimizers in search of parking space

⇩

Tragedy of commons effects

- **Game-theoretic investigation*** :
  - fully rational drivers with
    - -complete knowledge of the parking supply & the charging policy
    - -high, limited or no uncertainty about the overall parking demand

E.Kokolaki, M.Karaliopoulos , I.Stavrakakis, ”Game-theoretic approaches to the parking spot selection problem,” submitted.

The parking spot selection game

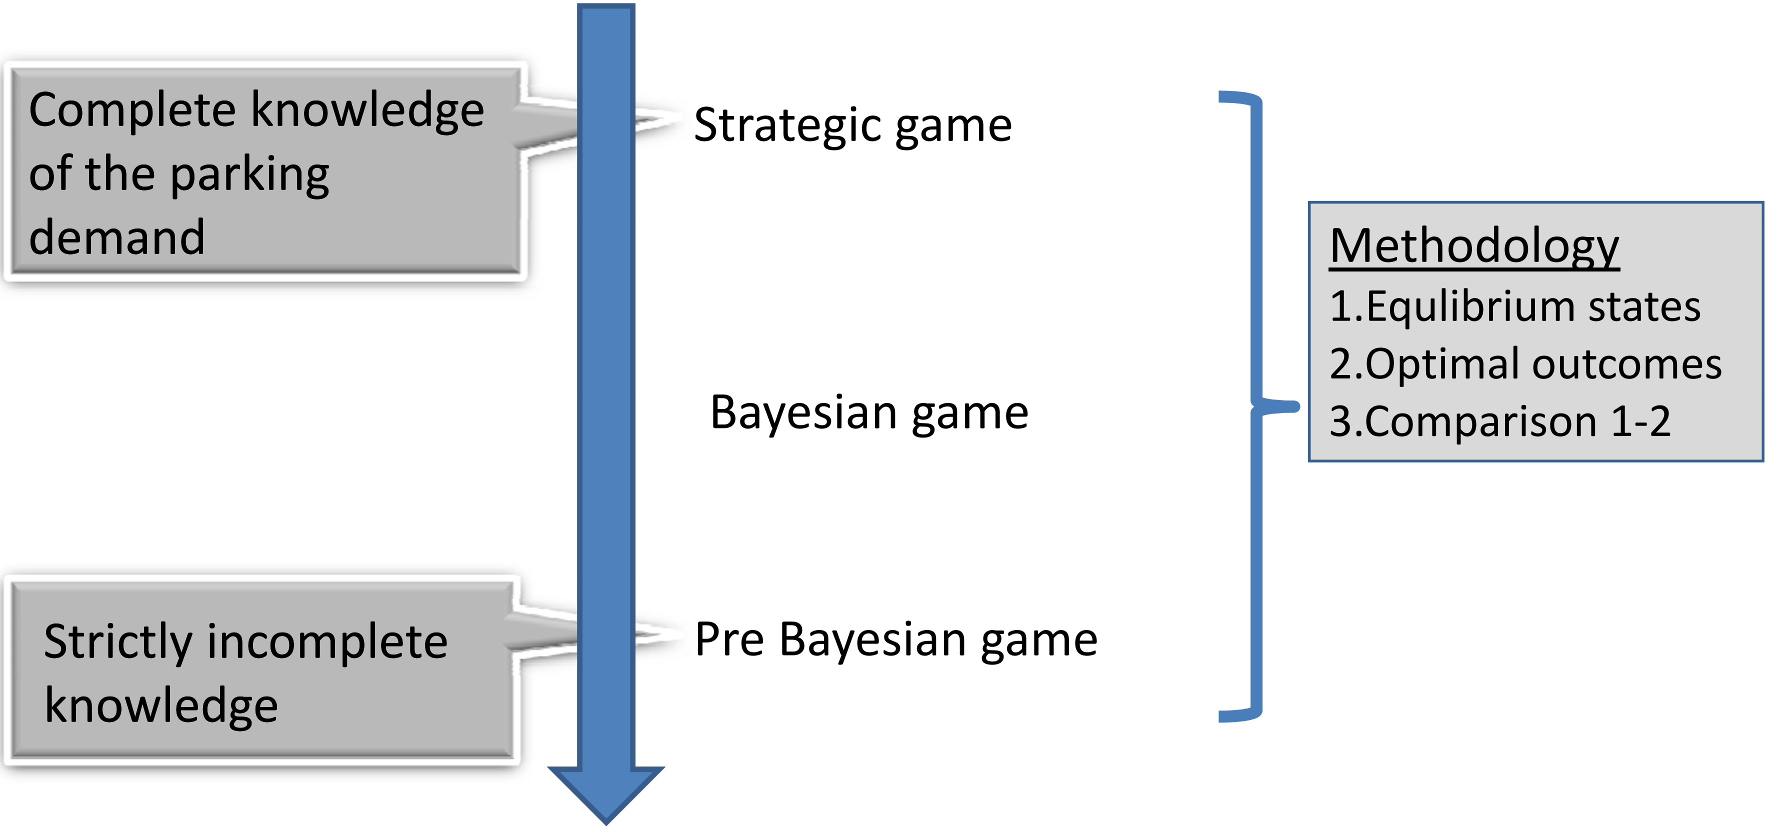

Complete knowledge of the parking demand

- The strategic parking spot selection game:
  - Set of *drivers*: N ={1,…,N}, N>1
  - Set of *parking spots*: $\mathcal{R}$=Rpub□Rpriv, R=|Rpub|□1 and |Rpriv|□N
  - Action set: $A = \{\text{public}, \text{private}\}$
  - Cost functions:

$$w_{pub}(k) = \min\left(1, \frac{R}{k}\right) c_{pub,s} + \left(1 - \min\left(1, \frac{R}{k}\right)\right) c_{pub,f}$$

$$w_{priv}(k) = c_{priv}$$

where

- □$c_{pub,s}$: cost of successfully competing for public parking
- □$c_{pub,f} = \gamma \cdot c_{pub,s}, \gamma > 1$: cost of competing, failing and paying for private parking
- □$c_{priv} = \beta \cdot c_{pub,s}, 1 < \beta < \gamma$: cost of private parking
- □$\delta \cdot c_{pub,s} = (\gamma - \beta) \cdot c_{pub,s}$: cruising cost

Numerical results

❑Worst-case equilibria:
If N>R, some drivers pay the price of the lack of coordination

❑Optimal case (centralized parking spot allocation):
Number of competing drivers=number of public parking resources

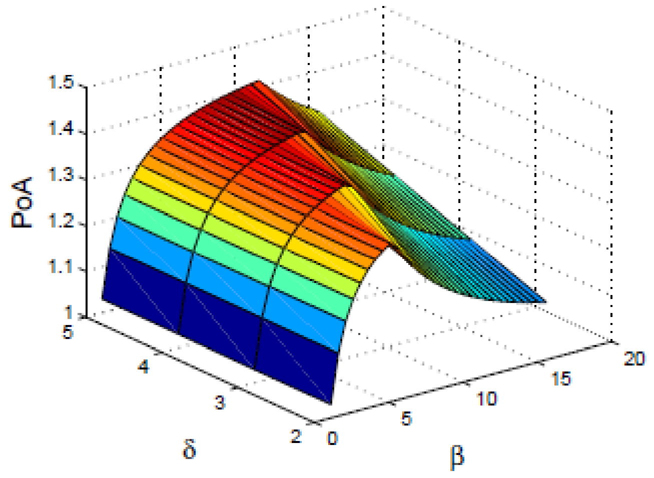

N=500,R=160,cpub,s=1€

❑Social cost=total cost paid

$$\text{Price of Anarchy} = \frac{\text{social cost in worst-case equilibrium}}{\text{optimal (minimum) social cost}}$$

$$When \;\; PoA = \frac{C_{eq}}{C_{opt}} \rightarrow 1?$$

Less-is-more phenomena under uncertainty

1. General trends for N>R:

↑ total number of drivers,N ⟹ ↓ Prob to compete

↓ Prob to compete ⟹ ↓ Social cost

2. Drivers decide their action, assuming the total number of drivers, N, in pre-bayesian games or the number of active players, K<N, under complete information.

By [1,2] the social cost under strictly incomplete information is less that the total cost paid under complete information for the parking demand.

Sample of recent EU-funded Projects (last 5 years)

- (EU-IST-FIRE) EINS (11-15)
  Network of Excellence in Internet Science
- (EU-IST-FET) RECOGNITION (11-14)
  Relevance and cognition for self-awareness in a content-centric Internet
- (EU-Marie_Curie) Retune (10-12)
  Resilience of Opportunistic Networks to Node Misbehaviors
- (EU-IST-FET) SOCIALNETS (08-11)
  Social Networking for Pervasive Adaptation

- (EU-IST-FET/FIRE) ANA (06-10)
- (EU-IST-FET/FIRE) BIONETS (06-10) with INRIA (Eitan Altman)
- (EU-IST-FET/FIRE) CASCADAS (06-09)
- (EU-IST) CONTENT (06-09)
- (EU-Marie_Curie) CoMig (05-06)
- (EU-IST_NoE) E-NEXT (04-06)

Sample of recent Networking Dissertations

- Eva Jaho, **"Cooperative mechanisms for information dissemination and retrieval in networks with autonomous nodes"** (2011)
- Leonidas Tzevelekas, **"Energy efficient algorithms for topology control and information dissemination/ retrieval in large scale Wireless Sensor Networks"** (2010)
- Dimitris Kogias, "**Study and Design of Algorithms for Information Dissemination in Unstructured Networking Environments**", (2010)
- Constantinos Vassilakis , "**Content distribution support in modern wireless and wired networks**", (2008, University of Athens). Employment: Greek NREN, GRNet.
- Athanasios Vaios, "**Short-range multi-hop communication extension to Wireless LANS**", (2008, University of Athens). Information Systems Division, Univ. Athens
- Pandelis Balaouras, "**Rate Adaptation Schemes for Continuous Media Steams**", (2005, University of Athens) – Video-conference and Multimedia Services, Univ. Athens.
- Nikolaos Laoutaris , "**Modeling and Optimization of Content Networks**", (2004, University of Athens) - currently with Telefonica Research, Barcelone (post-doc fellow at Harvard, Marie-Curie Post-Doc Fellow at Boston Univ. and Univ. Athens).
- Konstantinos Oikonomou , "**Topology Unaware MAC Protocols for Ad Hoc Networks**", (2004, University of Athens) – Lecturer, Ionian Univ.
- Antonis Panagakis , "**Study of the capacity of Static Weight Allocation Scheduling schemes and Wireless Multi-hop Networks**", (2004, University of Athens) – Hellenic Telecommunications and Post Commission.

Networking Faculty at DIT

Prof. Lazaros Merakos (1994)
CNL (Communications Networks Lab – wireless)

Prof. Ioannis Stavrakakis (1999)
ANR (Advanced Networking Research)

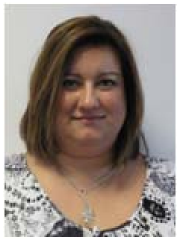

Lect. Nancy Alonistioti (2009)
SCAN (Self-evolving Cognitive and Autonomic Networking)

Networking at DIT

30 Ph.D. dissertations at UoA completed

(21 in Last 5 Years & 17 in progress)

Over 120 Undergraduate and M.Sc. Theses in Last 5 Years

Working for ….

Academia (U. Peloponnese, Ionian, AUBE, U. Piraeus, UoA, …)

R&D (GRNET, GUNET, UoA-DIT, UoA-NOC, BU, Harvard, Telefonica Research, Intracom, Thales, …)

Organizations (EETT, ….) / Higher Education & Gov

- Over 10 EU projects in the last 5 years
- Numerous publications

Building up Networking Science & Engineering
The funding – (1.1.2006 - today)

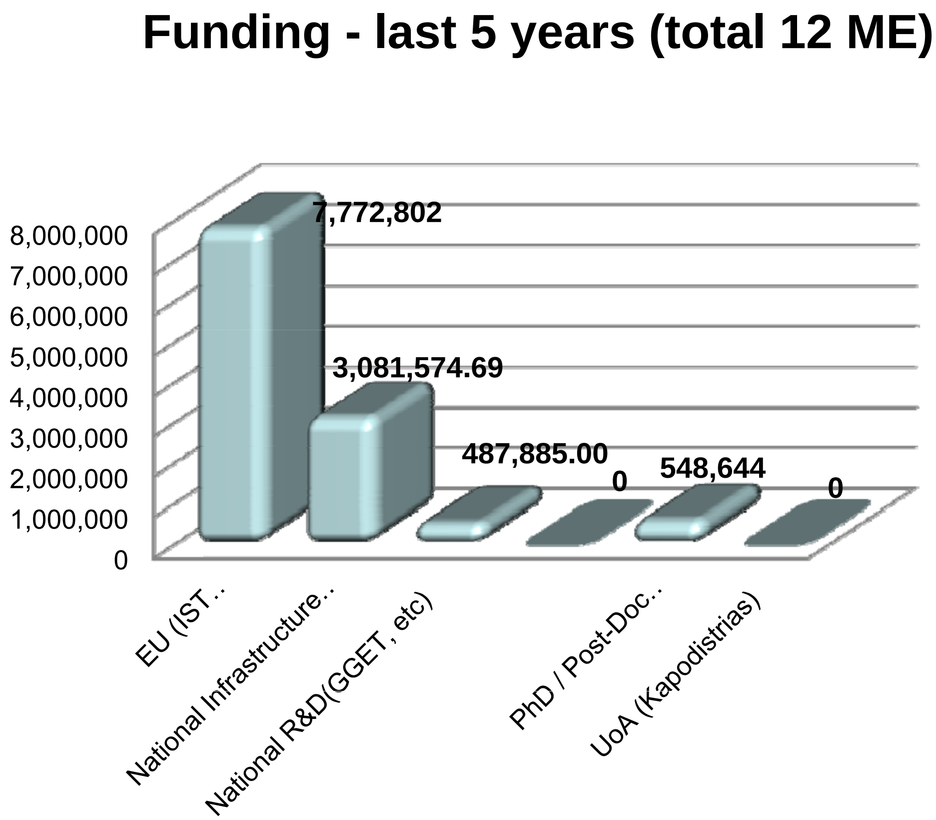

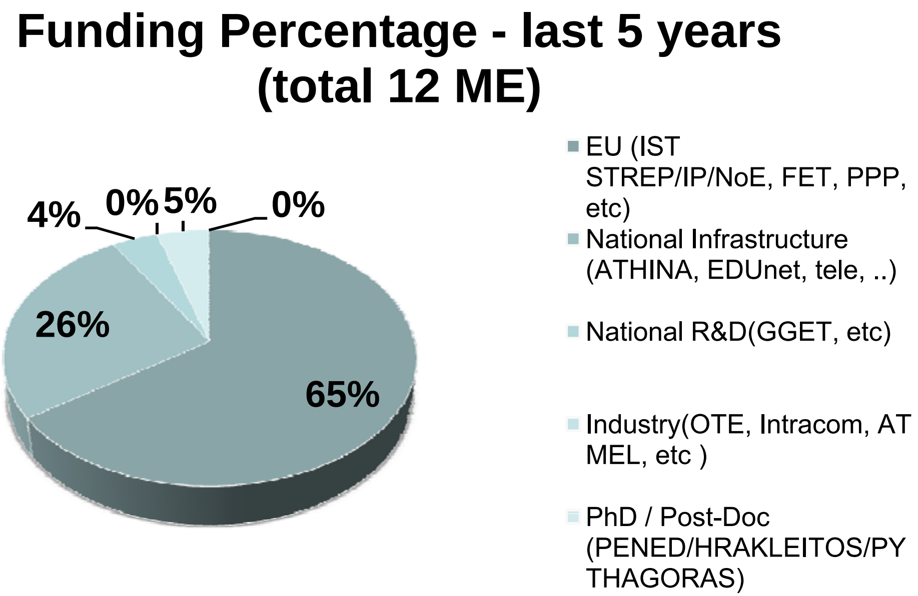

Building up Networking Science & Engineering
The International collaborations - L5Y (1.1.2006 - today)

- **Operators:** more than 10

  (e.g., Telecom Austria, Vodafone, OTE, FranceTelecom, Orange, Telefonica, TelecomItalia, TNO, DOCOMO)
- **Industry**: more than 15

  (e.g., ALCATEL LUCENT, Thales Communications SA, Toshiba, INTEL, HUAWEI, NEC NTUK, ATOS, IMEC, DICE, ATB, Nokia, SIEMENS, Unilever)
- ***Academics / Research Institutes:*** more than 15

  (e.g., UPC, UC3M, U. Coimbra, UPMC-Lip6, INRIA, Eurecom, U. Cambridge, U. Oxford, Cardiff U., U. Surrey,, King's College, UCL, U. Lancaster, U. Trento, CREATE-NET, CNR-Pisa, U. Napoli, U. Florence, TU Delft, IBBT, VUB, VTT, U. Limerick, U. Karlsruhe, Fraunhofer Institute, Darmstadt U., Boston U.)

“Networking” supporting disciplines in DIT

- Ass. Prof. Stathes Hadjiefthymiades
  - The Pervasive Computing Research Group
- Algorithms and optimization
- Distributed Systems
- Game Theory

Division of Telecommunications and Signal Processing

- 16 Faculty members / 18 positions

- Signal processing
- Telecommunications and Networking
- Hardware (electronics – embedded systems – photonics)