National and Kapodistrian
UNIVERSITY OF ATHENS

Advanced Networking Research (ANR) group

Outline of two research activities

Dr. Merkouris Karaliopoulos

Visiting Researcher (Marie Curie Fellow)

ANR in brief

- Set up in 1999 by Prof. Stavrakakis
- Comes under the division of Telecommunications and Signal Processing of the Department
- Currently 12 people
- http://anr.di.uoa.gr

ANR people

- Head
  - Professor Ioannis Stavrakakis
- Faculty Members
  - Ass. Prof. K. Oikonomou
  - Ass. Prof. C. Xenakis
  - Dr. G. Karagiorgos
- Senior Researchers/Post-Docs
  - Dr. Merkourios Karaliopoulos (Marie Curie Fellow)
  - Dr. Christoforos Ntantogian
- PhD students/Research Assistants
  - Ioannis Manolopoulos
  - Christoforos Panos
  - Panagiotis Pantazopoulos
  - Evangelia Kokolaki
  - Lesia Kozlenko
- Administrative Staff
  - Efi Tsoukali

Topics we carry out research on...

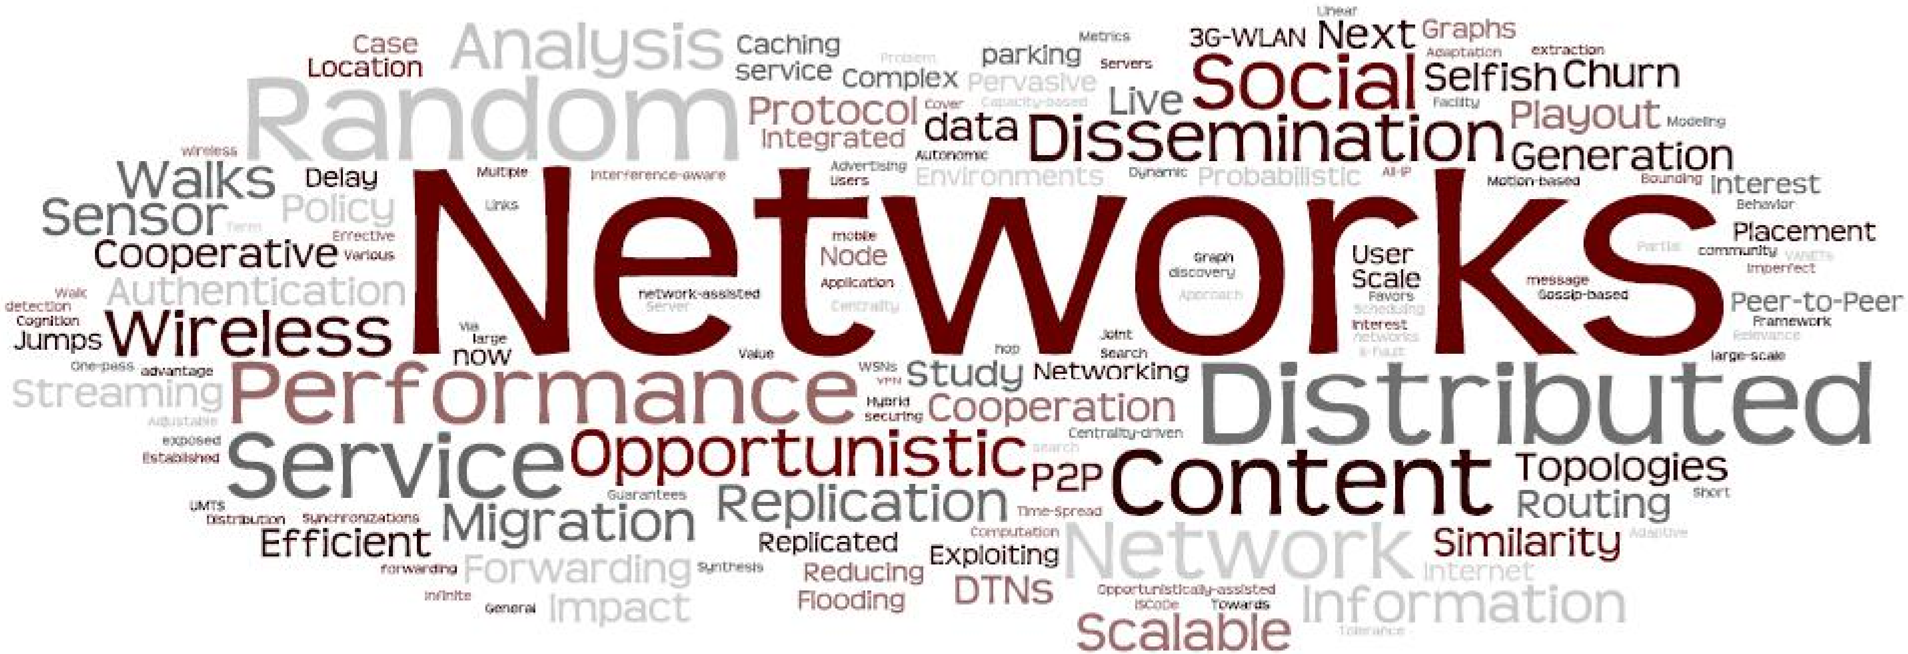

- Wordle on the papers of the last five years
  - http://anr.di.uoa.gr/index.php/publications

...and where we publish our results

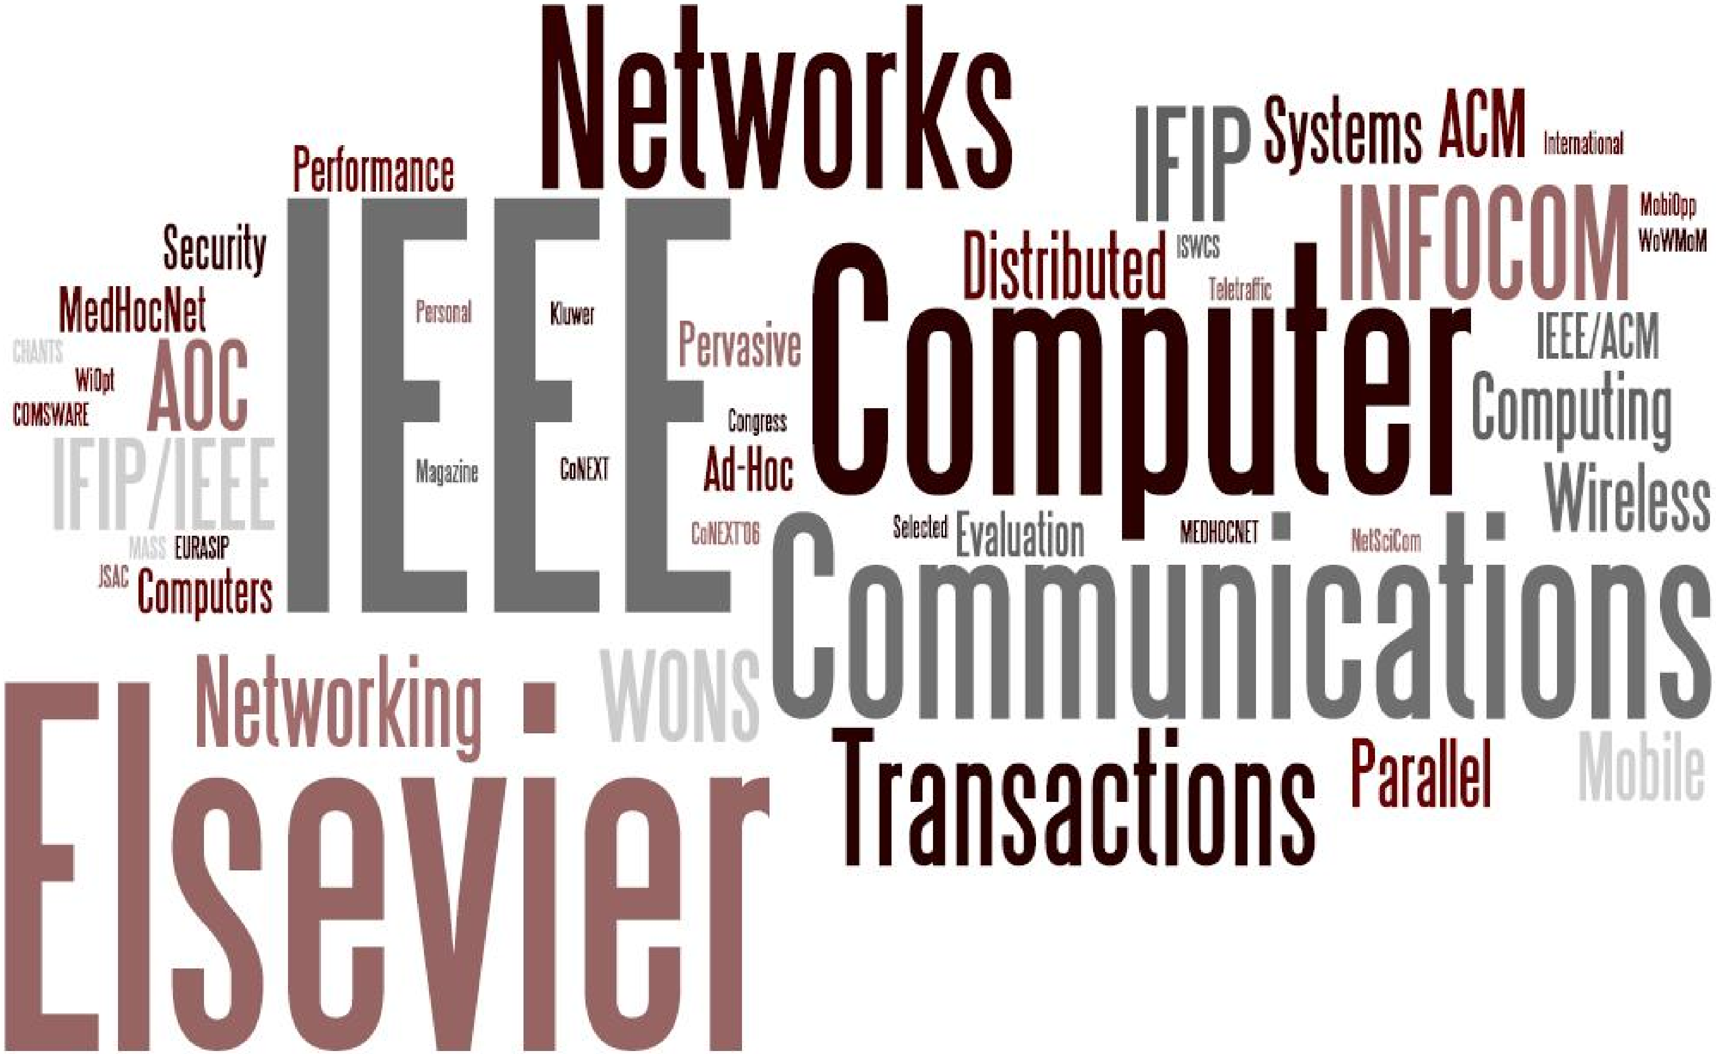

- Wordle on the venues of our papers over the last five years
    - http://anr.di.uoa.gr/index.php/publications

Two research work items in 25'

- Service placement
  - group research thread[1,2,3,4]
  - distributed socio-aware heuristic for the facility location problem

- Perfomance analysis of opportunistic forwarding over mobility traces (a.k.a dynamic networks)
  - Combination of graph expansion techniques with simple path-finding algorithms

[1]K. Oikonomou, I. Stavrakakis, "**Scalable Service Migration in Autonomic Network Environments**", special issue of IEEE Journal on Selected Areas in Communications (JSAC) on "Recent Advances in Autonomic Communications", vol. 28, no. 1, pp. 84-94, Jan. 2010

[2]G. Smaragdakis, N. Laoutaris, K. Oikonomou, I. Stavrakakis, and A. Bestavros, "**Distributed Server Migration for Scalable Internet Service Deployment**", IEEE/ACM Transactions on Networking. (To Appear)

[3]P. Pantazopoulos, M. Karaliopoulos, I. Stavrakakis, "**Centrality-driven scalable service migration**", in Proc. 23rd International Teletraffic Congress, Sept. 6-8, 2011, San Francisco, USA

[4]P. Pantazopoulos, M. Karaliopoulos, I. Stavrakakis, "**Distributed placement of Internet services**", submitted to IEEE Transactions on Parallel and Distributed Systems

Service placement : description/ motivation

- User generated content (UGC) → User generated services (UGS) trend
  - service facilities generated almost anywhere across the network
    - many in number, often of local (small-scale) demand (replication: not preferred option)
    - lack of central coordination

- in-network storage
  - content- and information-centric networking

- storage task distribution towards many lighter storage platforms
  - nanodatacenters aim to offload power-hungry data centers

Objective

- scalable distributed mechanisms for "optimally" placing service components
- "optimally" ≡ minimize aggregate service access costs

Service placement : problem formulation

- Facility-location problem
- INPUT

  V : set of nodes

  w_n : demand generated by node n

  $d(x_j,n)$: distance between nodes x_j and n

  OUTPUT

  F : placement

- 1-median: minimize the access cost of a service located at node k

$$Cost(k) = \sum_{n \in \mathcal{V}} w(n) \cdot d(k,n).$$

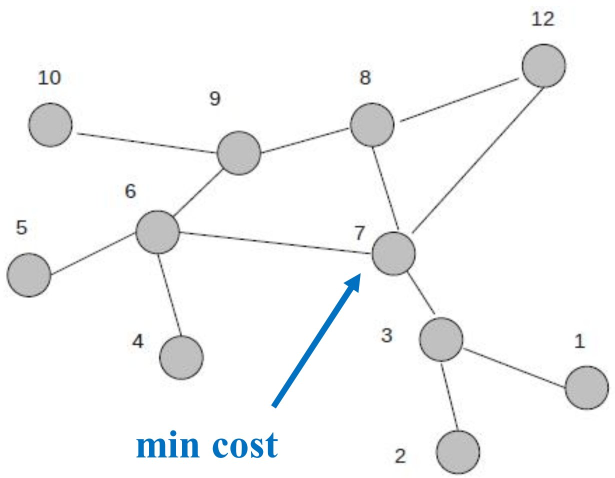

Service placement : our solution

- service migrates towards the optimum host (opt) in a finite number of steps

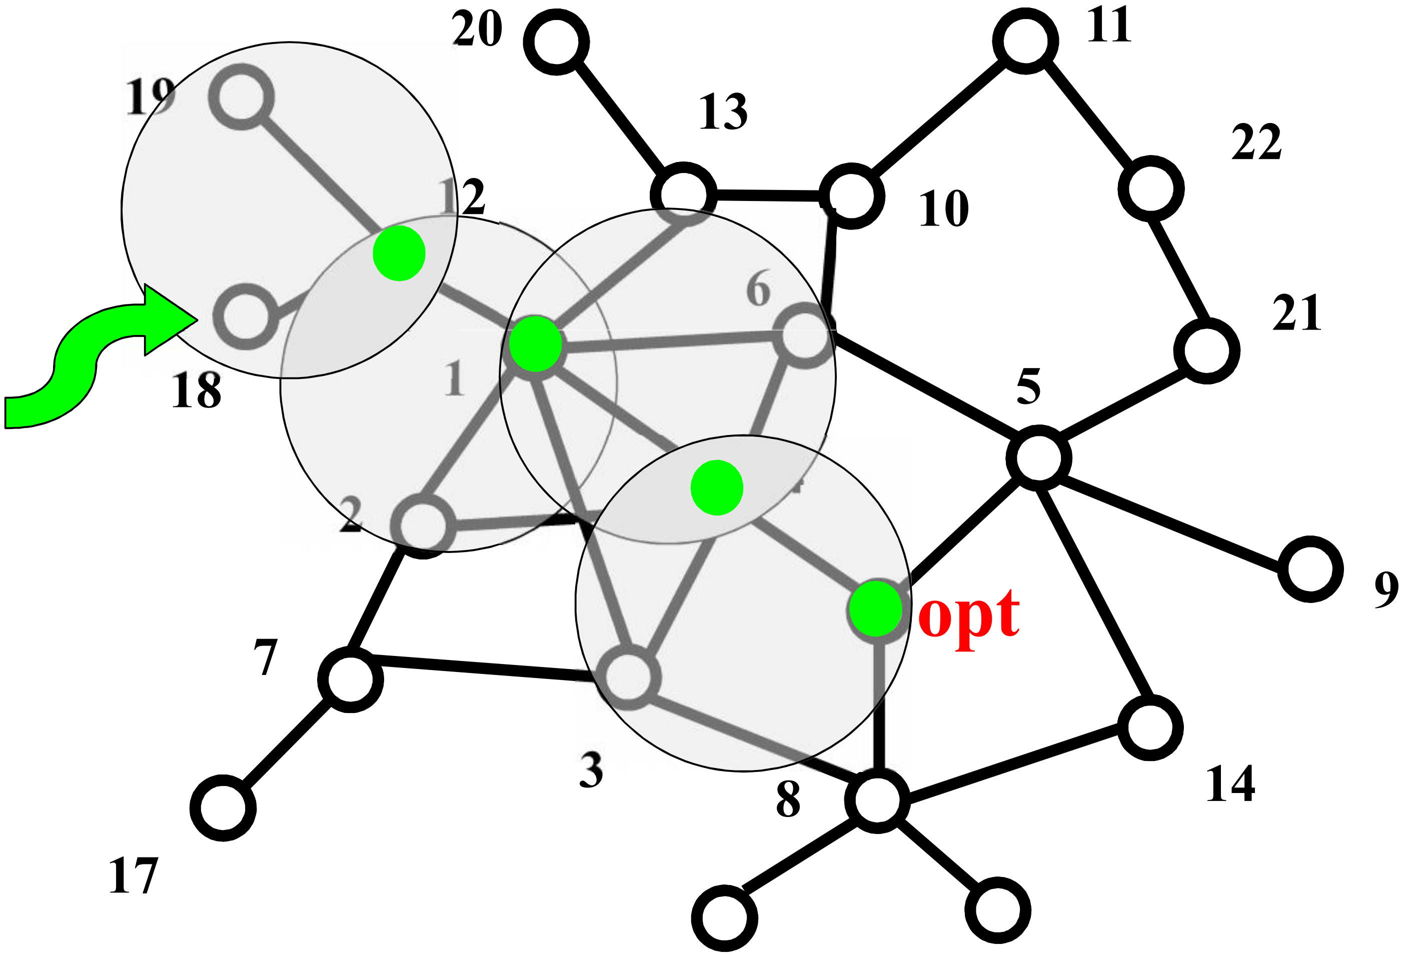

The R-balls heuristic*

- Reduce the original k-median to multiple smaller 1-median problems
  - solved within a limited neighborhood of R-hops around current facility

- Demand generated by outer nodes is mapped to the nodes at the outer shell of the R-hop neighborhood

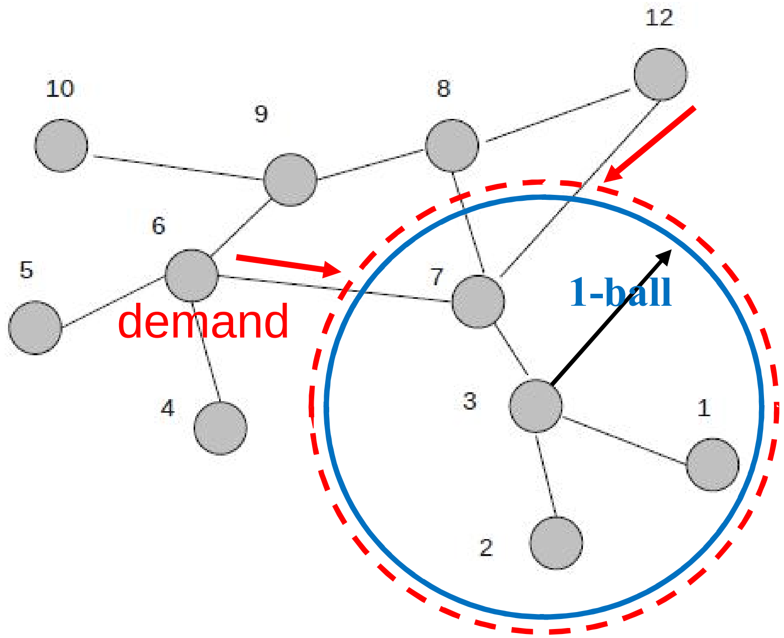

$$Cost(\mathcal{F}) = \sum_{n \in \mathcal{V}} w_n \cdot min_{x_j \in \mathcal{F}} \{d(x_j, n)\}$$

* G. Smaragdakis, N. Laoutaris, K. Oikonomou, I. Stavrakakis, A. Bestavros, "Distributed Server Migration for Scalable Internet Service Deployment," to appear in IEEE/ACM ToN

Service placement : wCBC / demand mapping

- a high number of shortest paths through the node u (e.g. node 8) does not necessarily mean that equally high demand stems from their sources!

- weighted conditional BC

$$wCBC(u;t) = \sum_{s \in V, u \neq t} w(s) \cdot \frac{\sigma_{st}(u)}{\sigma_{st}}.$$

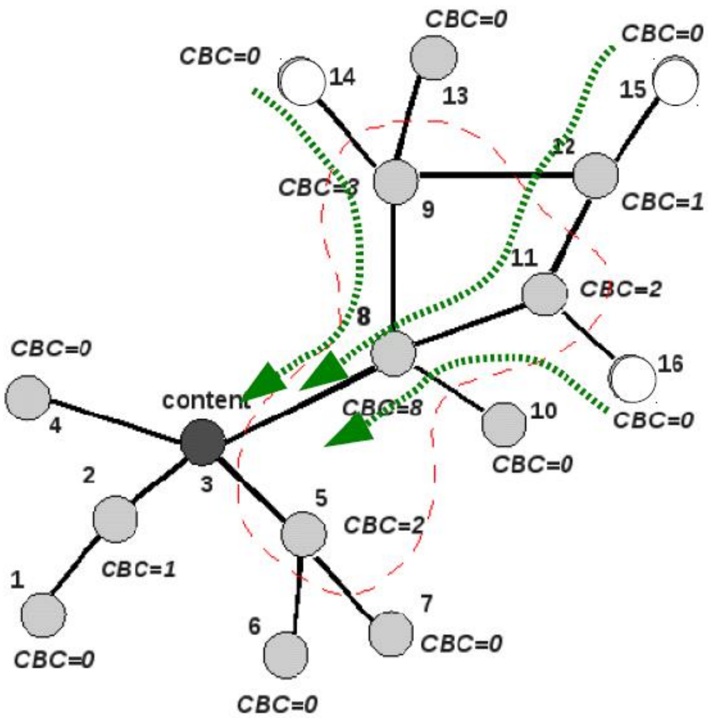

- *wCBC* assesses to what extent a node can serve as demand concentrator towards a given service location
  - The **top a% wCBC-valued** nodes are included in the 1-median subgraph

Making the search “more informed” : cDSMA

R-ball heuristic

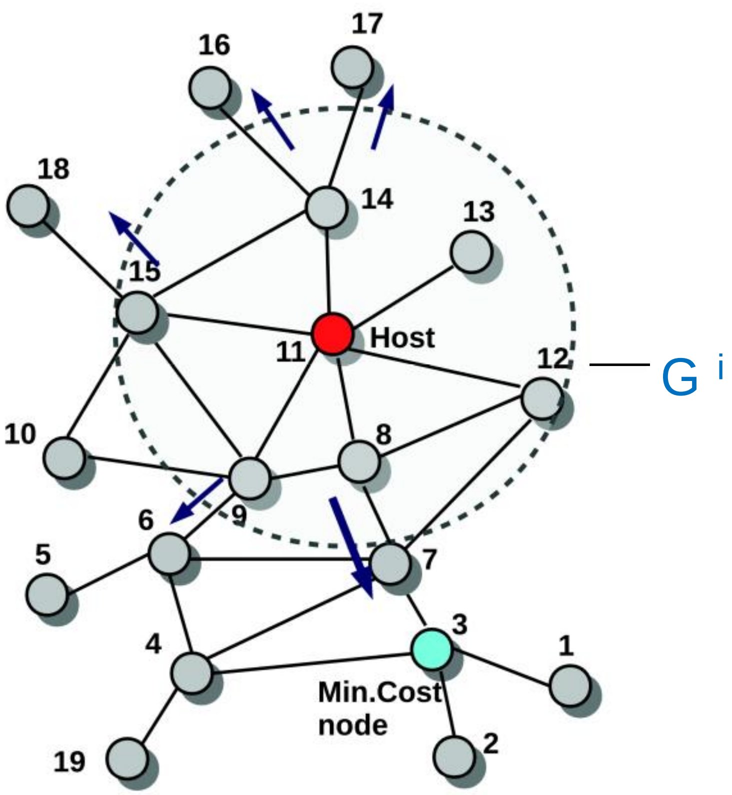

cDSMA

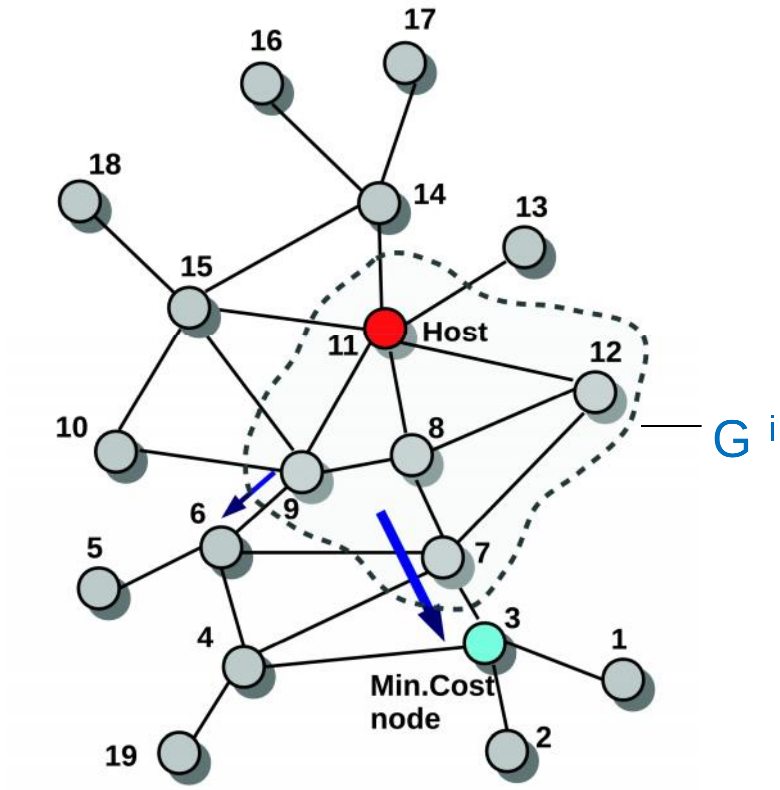

$$Cost(k) = \sum_{n \in \mathcal{V}} w(n) \cdot d(k, n).$$

Approximation ratio : a negative result…

- N nodes, service generated at node B
- A heavy hitter node A with demand W, all other nodes pose equal unit demand

$$\frac{C_{cDSMA}(B)}{C_{OPT}} = 1 + 2\frac{W-1}{N}$$

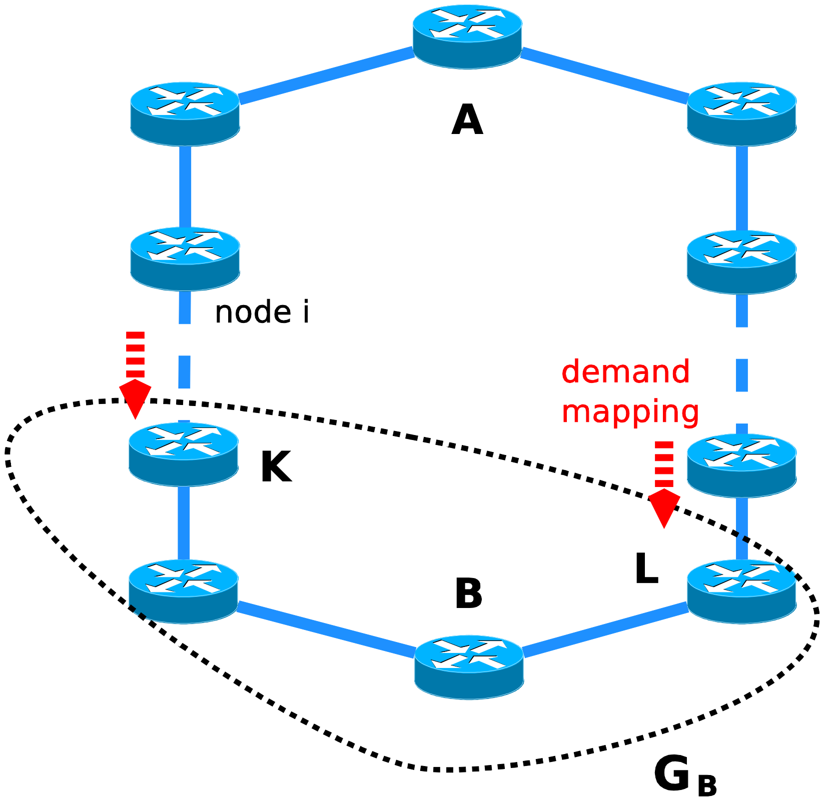

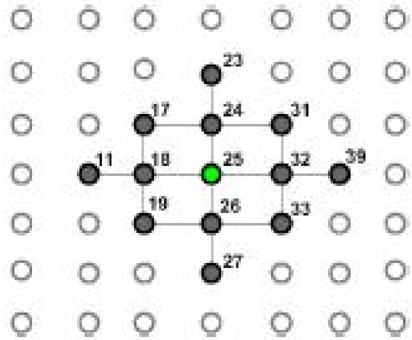

a. $|G_{Host}| = 5$

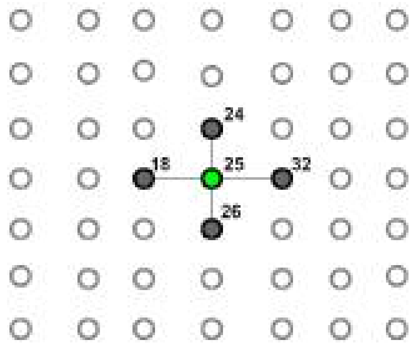

c. $|G_{Host}| = 12$

Yet…

Performance is excellent in realistic network topologies

- Datasets correspond to different snapshots of 7 ISP topologies

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

					s=0		s=1		s=2	
ISP	Dataset id/AS#	mCC nodes	Diameter	<Degree>	$\alpha_{0.025}$	$\lceil \lvert G^s \rvert \rceil$	$\alpha_{0.025}$	$\lceil \lvert G^s \rvert \rceil$	$\alpha_{0.025}$	$\lceil \lvert G^s \rvert \rceil$
type: **Tier-1**										
Global Crossing	36/3549	76	10	3.71	0.047±0.001	4	0.047±0.002	4	0.046±0.001	4
-//-	35/3549	100	9	3.78	0.045±0.002	5	0.045±0.001	5	0.043±0.001	5
NTTC-Gin	33/2914	180	11	3.53	0.024±0.002	5	0.022±0.002	4	0.019±0.002	4
Sprint	23/1239	184	13	3.06	0.019±0.002	4	0.018±0.002	4	0.017±0.002	4
-//-	21/1239	216	12	3.07	0.016±0.002	4	0.016±0.002	4	0.014±0.003	4
Level-3	27/3356	339	24	3.98	0.018±0.002	7	0.017±0.002	6	0.014±0.003	5
-//-	13/3356	378	25	4.49	0.012±0.002	5	0.012±0.002	5	0.011±0.002	5
type: **Transit**										
TDC	46/3292	71	9	3.30	0.033±0.003	3	0.027±0.004	2	0.026±0.003	2
DFN-IPX-Win	41/680	253	14	2.62	0.019±0.003	5	0.015±0.003	4	0.015±0.003	4
JanetUK	40/786	336	14	2.69	0.012±0.003	5	0.012±0.002	5	0.013±0.002	5

Less than half a dozen nodes suffice in almost all cases, even under uniform demand

P. Pantazopoulos, M. Karaliopoulos, I. Stavrakakis, "**Centrality-driven scalable service migration**", in Proc. 23rd International Teletraffic Congress, Sept. 6-8, 2011, San Francisco, USA

Service placement: distributed protocol

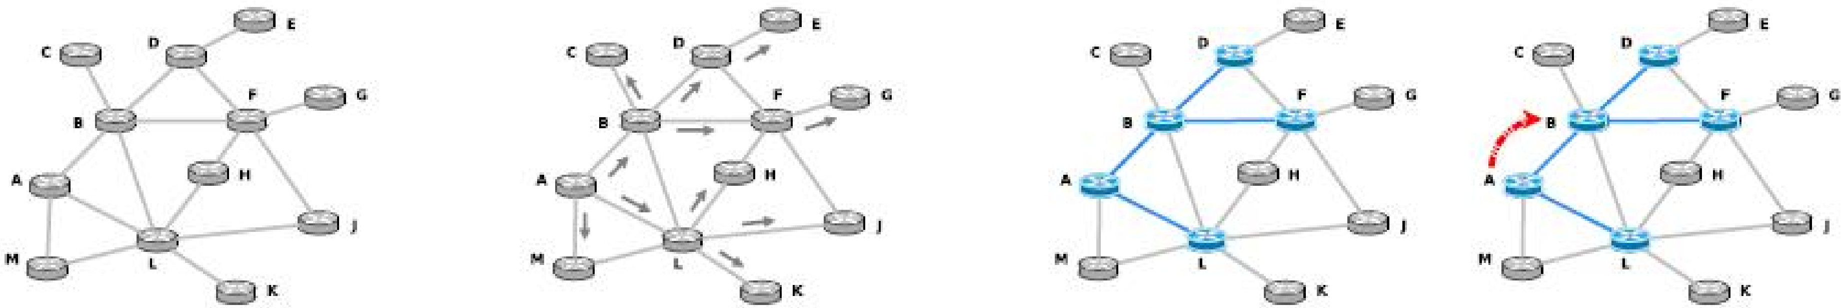

a. The graph $G(E, V)$ b. Advertisement phase (service at A) c. Host induces G_A d. Host solves 1-median

Header of

msg_D :	D	B	A
msg_F :	F	B	A
msg_B :		B	A
msg_L :		L	A

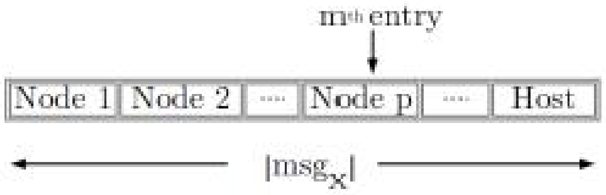

Algorithm 2 Message header parsing and demand mapping under SP (DeMaSP)

---

input: *set of selected nodes in* G_{Host},
  $\{msg_u\}$ $\forall u \in G_{Host}$
output: *vector* $w_{eff}(u)$ $\forall u \in G_{Host}$
Initialization
for all $x \in G_{Host}$ **do** $w_{eff}(x) = w\hat{CBC}(x)$
vector $B \leftarrow$ *sort all* msg_x *in decreasing order of* $|msg_x|$

for $i = 1$ **up to** $Len(B)$ **do**
  parse $B(i) = msg_x$
  for $m = 1$ **up to** $|msg_x| - 1$ **do**
    if $msg_x(m)$ *is marked* **then**
      if $m > 1$ **then**
        $k = msg_x(m-1)$, $l = msg_x(m)$
        $w_{eff}(l) = w_{eff}(l) - w\hat{CBC}(k)$
      end
      drop msg_x
    else
      if $m > 1$
        $k = msg_x(m-1)$, $l = msg_x(m)$
        $w_{eff}(l) = w\hat{CBC}(l) - w\hat{CBC}(k)$
      end
      mark msg_x *as read*
    end if
  end for
end for
$w_{eff}(Host) = w(Host) + (w_{trans}(Host; Host) - \sum_{z \in Ch(Host;T)} w\hat{CBC}(z))$

P. Pantazopoulos, M. Karaliopoulos, I. Stavrakakis, “Distributed placement of Internet services”, submitted to IEEE Transactions on Parallel and Distributed Systems

Two research work items in 25'

- Service placement
  - Distributed heuristic for the facility location problem

- Perfomance analysis of opportunistic forwarding over mobility traces (a.k.a dynamic networks)
  - Combination of graph expansion techniques with simple path-finding algorithms[5,6]

[5] M. Karaliopoulos, C. Rohner, "**Trace-based Performance Analysis of Opportunistic Forwarding under Imperfect Cooperation**", IEEE INFOCOM 2012 Mini-Conference program, Orlando, FL, USA, March 25-30, 2012

[6] M. Karaliopoulos, C. Rohner, "**Trace-based Performance Analysis of Opportunistic Forwarding**" under submission to IEEE Transactions on Mobile Computing

Opportunistic forwarding

- Leverage node mobility and pairwise encounters for disseminating data
  - Store-carry-and-forward principle

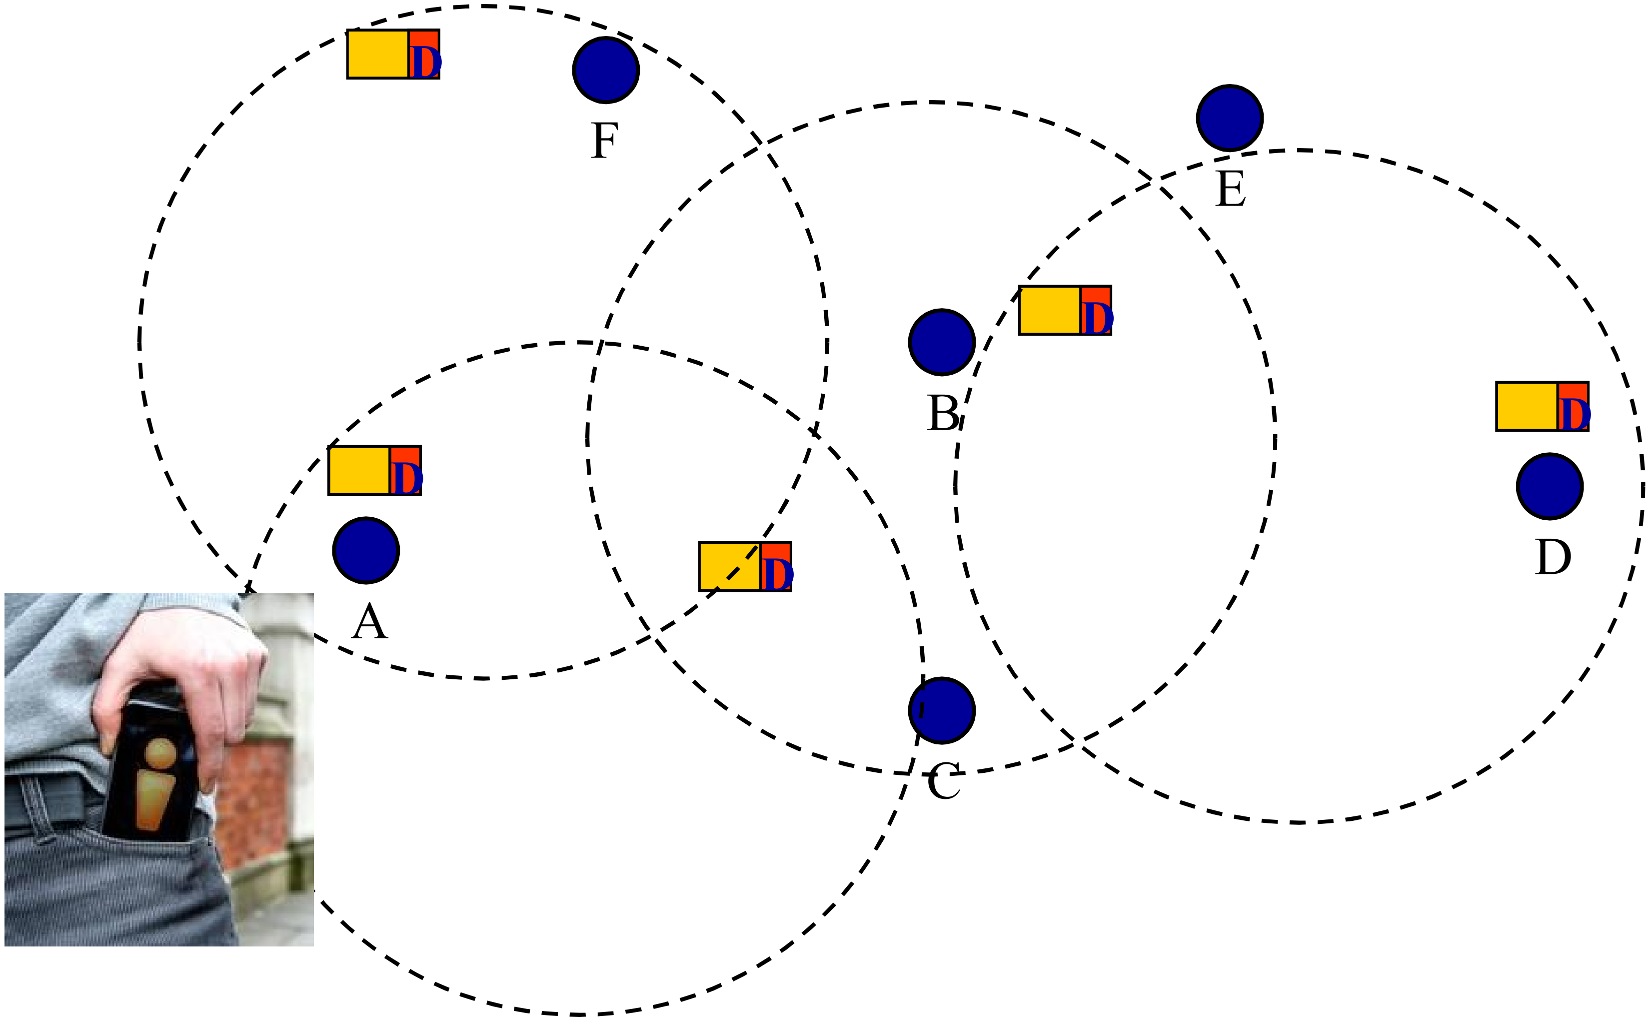

The problem

The input

contact id	involved nodes		contact start time	contact end time
…	…	…	…	…
C_0	1	3	589	940
C_1	2	3	589	619
C_2	6	7	639	699
C_3	1	5	700	816
C_4	5	6	702	818
C_5	2	7	816	1185
C_6	4	8	819	909
C_7	8	6	938	1182
…	…	…	…	…

The output

- Protocol performance metrics
  - Message delivery probability/delivery delay/replicas, (space-time) path hopcounts
- ...for various forwarding protocols
  - randomized and utility-based
- ...for different communication modes
  - point-to-point/multipoint
- ...under full/partial node cooperation

How we go about it

- Work on per-message basis
  - message m = $\{s,d,t_g\}$

- Three main processing steps
  - Contact filtering : Filter the original trace for relevant forwarding entries
  - Graph construct derivation : Build the forwarding contact graph Gc=(Vc, Ec)
    - graph expansion technique resulting in a (sparse) directed acyclic graph (DAG)
  - Space-time path computation : Run standard path finding algorithms over the DAG
    - e.g., Dijkstra runs in O(|Vc|+|Ec|) time over the linearized Gc

Contact-filtering step : terminology

- Forwarding contacts : encounters that result in message copying/forwarding

- Forwarding list : list storing nodes that hold a message copy

- Depending on the current status of the forwarding list, parsed encounters are split into three categories
  - 0-entry contacts : neither node is listed in the forwarding list
  - 1-entry contacts : one of the two nodes is listed in the forwarding list
  - 2-entry contacts : both nodes are listed in the forwarding list

⟹ The treatment of the three types of contacts depends on the actual forwarding protocol rules and the performance metric

Contact-filtering step : epidemic spread

- msg = $(n1,n4,t_g)$, $t_g=t_{0,s}-\varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	…
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Original trace

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	…
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

filtered trace : epidemic spread

Contact-filtering : Two-hop forwarding

- msg = $(n1,n4,t_g)$, $t_g=t_{0,s}-\varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	…
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Original trace

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	…
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

filtered trace : two-hop forwarding

Graph construct extraction : epidemic

- msg = (n1,n4,t_g), $t_g=t_{0,s}-\varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	…
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Forwarding list = {n1,n2}

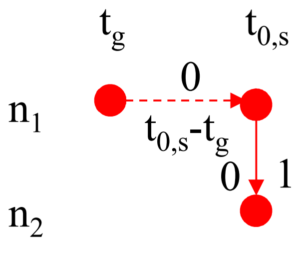

Graph construct extraction : epidemic

- msg = (n1,n4,t_g), $t_g = t_{0,s} - \varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
~~C_1~~	~~n3~~	~~n4~~	~~$t_{1,s}$~~	~~$t_{1,e}$~~	…
~~C_2~~	~~n4~~	~~n5~~	~~$t_{2.s}$~~	~~$t_{2,e}$~~	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Forwarding list = {n1,n2,n5}

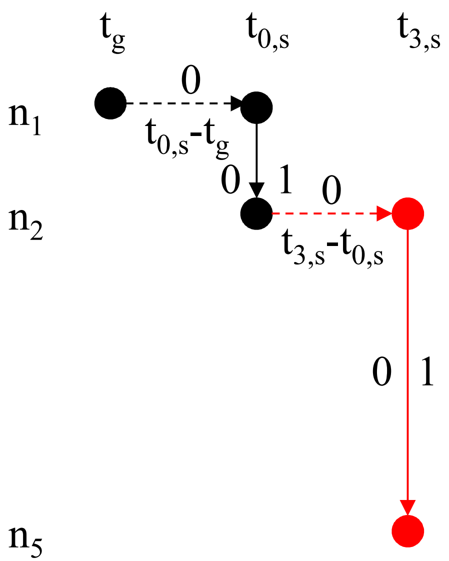

Graph construct extraction : epidemic

- msg = $(n1,n4,t_g)$, $t_g=t_{0,s}-\varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
~~C_1~~	~~n3~~	~~n4~~	~~$t_{1,s}$~~	~~$t_{1,e}$~~	~~…~~
~~C_2~~	~~n4~~	~~n5~~	~~$t_{2,s}$~~	~~$t_{2,e}$~~	~~…~~
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4,s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6,s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Forwarding list = {n1,n2,n5,n6}

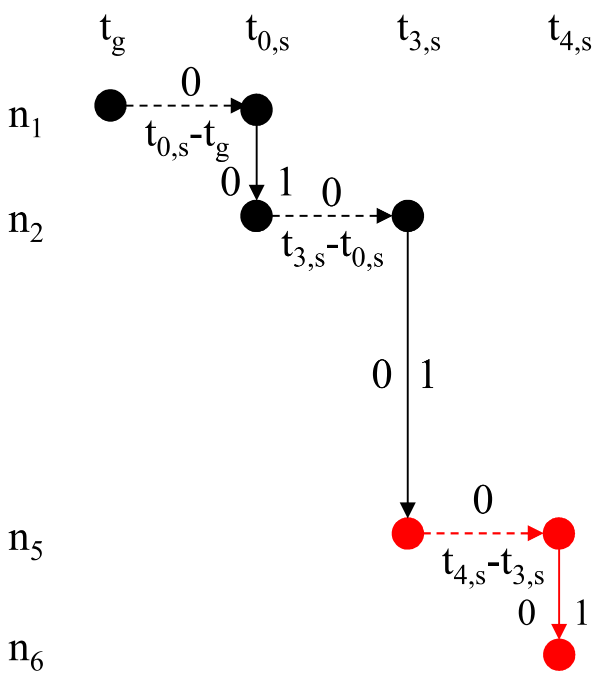

Graph construct extraction : epidemic

- msg = $(n1,n4,t_g)$, $t_g = t_{0,s} - \varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
~~C_1~~	~~n3~~	~~n4~~	~~$t_{1,s}$~~	~~$t_{1,e}$~~	…
~~C_2~~	~~n4~~	~~n5~~	~~$t_{2,s}$~~	~~$t_{2,e}$~~	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4,s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
C_6	n1	n6	$t_{6,s}$	$t_{6,e}$	…
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Forwarding list = {n1,n2,n5,n6,n3}

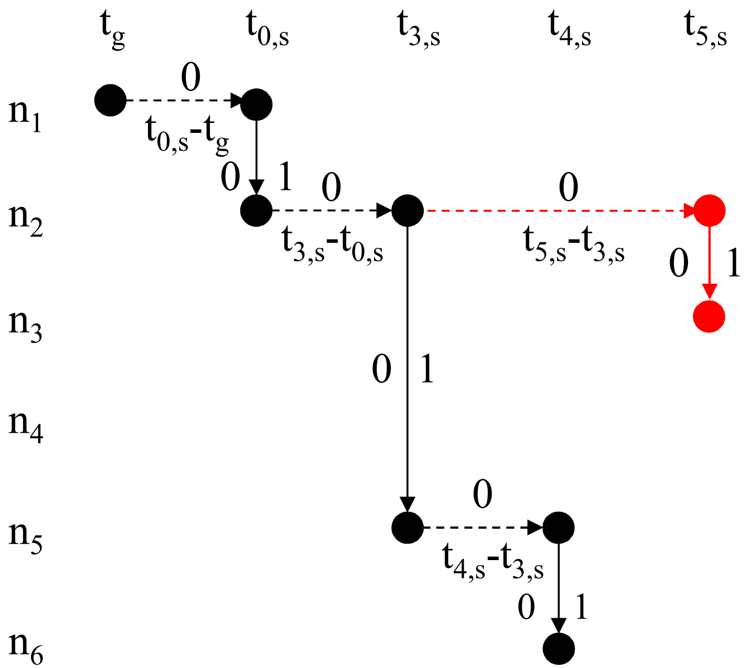

Graph construct extraction : epidemic

- msg = (n1,n4,t_g), $t_g = t_{0,s} - \varepsilon$, $\varepsilon \rightarrow 0$

contact id	involved nodes		contact start time	contact end time	additional fields
…	…	…	…	…	…
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	…
~~C_1~~	~~n3~~	~~n4~~	~~$t_{1,s}$~~	~~$t_{1,e}$~~	…
~~C_2~~	~~n4~~	~~n5~~	~~$t_{2,s}$~~	~~$t_{2,e}$~~	…
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	…
C_4	n5	n6	$t_{4,s}$	$t_{4,e}$	…
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	…
~~C_6~~	~~n1~~	~~n6~~	~~$t_{6,s}$~~	~~$t_{6,e}$~~	…
~~C_7~~	~~n2~~	~~n6~~	~~$t_{7,s}$~~	~~$t_{7,e}$~~	…
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	…
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	…
…	…	…	…	…	…

Forwarding list = {n1,n2,n5,n6,n3,n4}

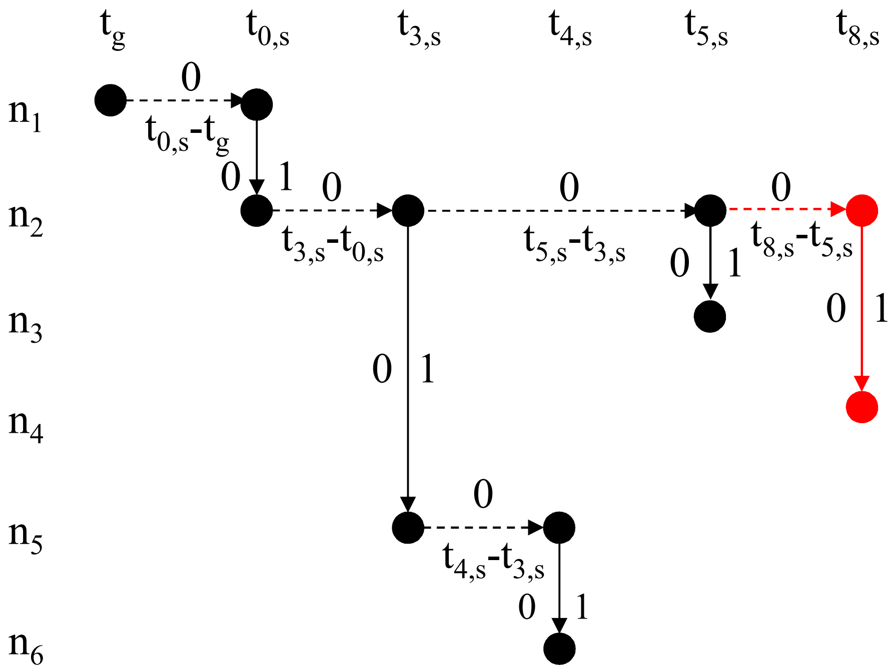

Stop parsing

Graph construct examples

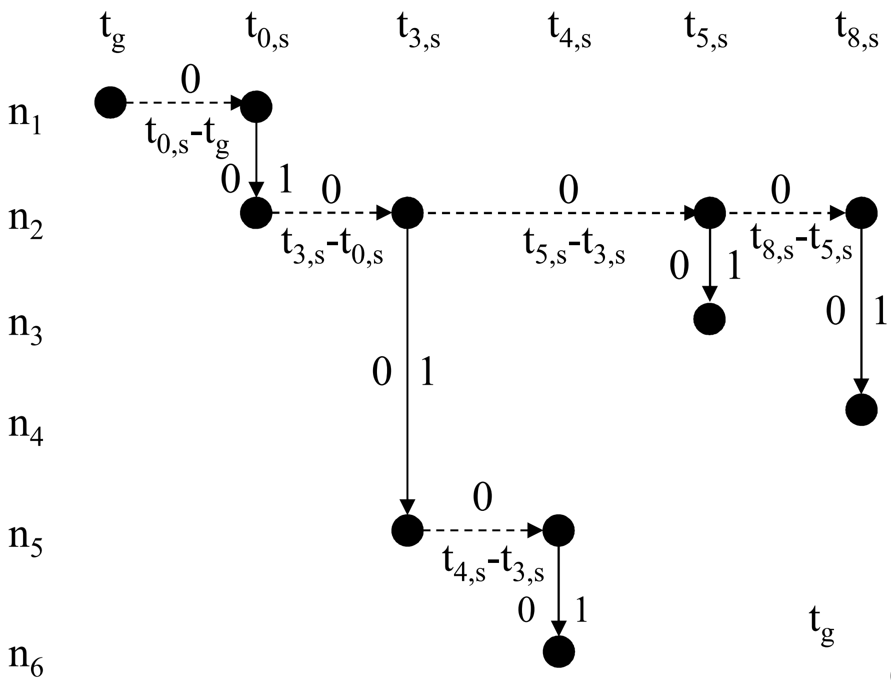

Epidemic forwarding

Two-hop forwarding

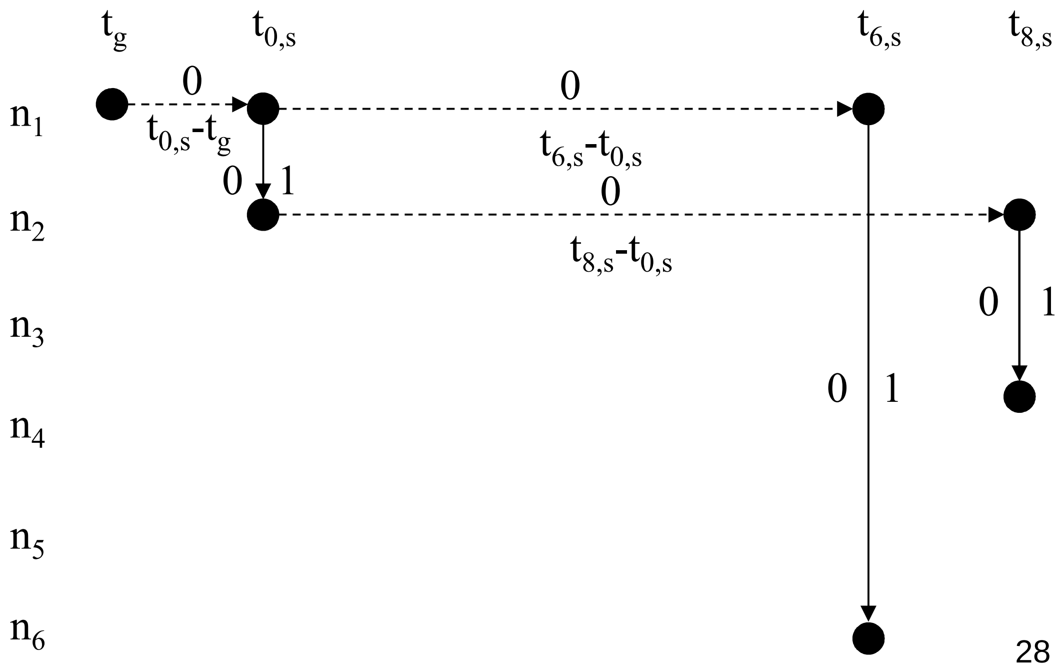

Shortest path computation

- Run DFS and linearize the DAG

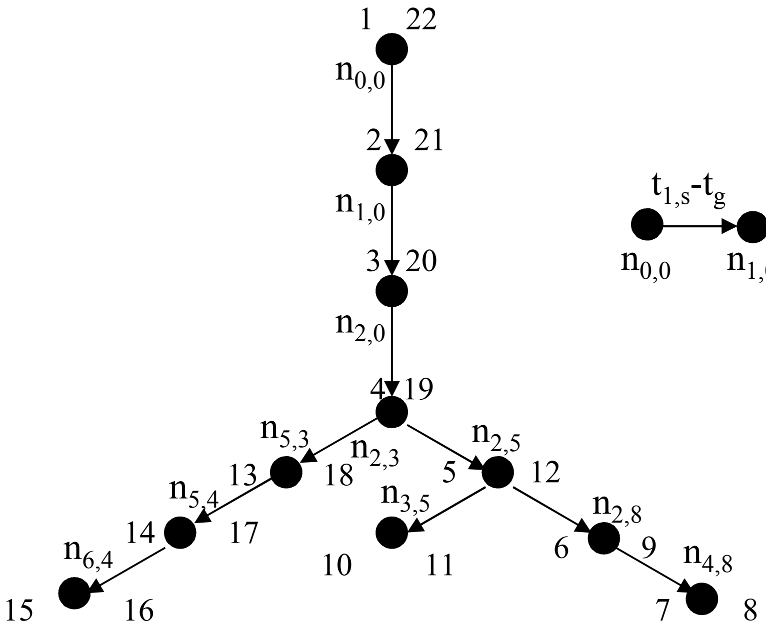

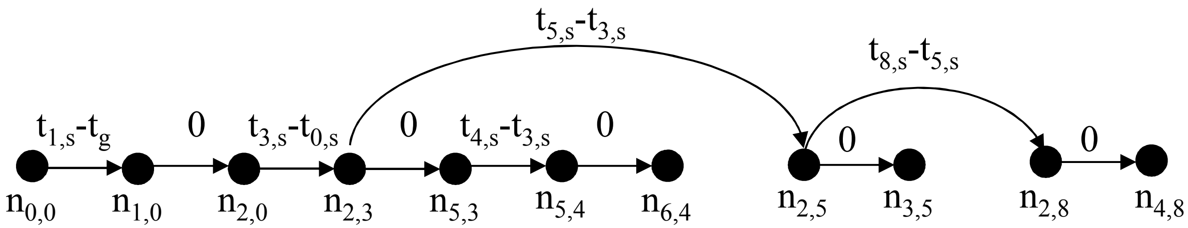

- Dijkstra runs in O(|Vc|+|Ec|) time
  - |Vc| = O(|V|), |Ec| = O(|V|) and Dijkstra complexity : O(|V|)

Extensions for addressing node misbehaviors

- Deferral from message copying/forwarding (encounter-agnostic)
  - if a node defers from both, all encounters involving it can be filtered out offline
  - if a node defers from forwarding only, its contacts are filtered online
    - additional state : list of misbehaving nodes

black list = {2}, msg = (n1,n6,t_g), $t_g = t_{0,s} - \varepsilon$

contact id	involved nodes		contact start time	contact end time	additional fields
...	...	...	...	...	...
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	...
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	...
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	...
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	...
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	...
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	...
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	...
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	...
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	...
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	...
...	...	...	...	...	...

filtered out due to 2's misbehavior

inflated message delivery time (instead of $t_{4,s}$)

Extensions for addressing node misbehaviors

- Encounter-dependent misbehavior (e.g. social selfishness)
  - per-node black-hole lists
  - may be aggregated offline into disjoint or overlapping communities

comms = {1,2,3,6} {3,5,7}   msg = (n1,n6,t_g), $t_g=t_{0,s}-\varepsilon$

contact id	involved nodes		contact start time	contact end time	additional fields
...	...	...	...	...	...
C_0	n1	n2	$t_{0,s}$	$t_{0,e}$	...
C_1	n3	n4	$t_{1,s}$	$t_{1,e}$	...
C_2	n4	n5	$t_{2.s}$	$t_{2,e}$	...
C_3	n2	n5	$t_{3,s}$	$t_{3,e}$	...
C_4	n5	n6	$t_{4.s}$	$t_{4,e}$	...
C_5	n3	n2	$t_{5,s}$	$t_{5,e}$	...
C_6	n1	n6	$t_{6.s}$	$t_{6,e}$	...
C_7	n2	n6	$t_{7,s}$	$t_{7,e}$	...
C_8	n2	n4	$t_{8,s}$	$t_{8,e}$	...
C_9	n3	n5	$t_{9,s}$	$t_{9,e}$	...
...	...	...	...	...	...

} filtered out due social community bias (C_3, C_4)

→ inflated message delivery time (instead of $t_{4,s}$) (C_6)

Comparison with simulations runs in ONE

Run times

		black hole selfishness						social selfishness					
		Intel		Cambridge		Infocom05		Intel		Cambridge		Infocom05	
	msgs	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4
epid	trace	14	216	43	585	1027	10853	21	212	100	969	883	8487
	ONE	3109		12021		53168		4848		9615		53190	
2hop	trace	8	104	30	360	545	5250	29	286	130	1253	1269	12436
	ONE	2699		5978		23615		2717		4435		18508	
ssaw	trace	8	105	30	368	585	5615	29	287	129	1262	1390	12879
	ONE	2366		4886		8625		1885		2652		2834	

- ONE run times 2-3 orders of time higher for 10^3 messages
  - Simulator run times prohibitive for 10^4 messages

Performance metrics

		black hole selfishness									social selfishness								
		Intel			Cambridge			Infocom05			Intel			Cambridge			Infocom05		
	metric	P_D	D_{avg}	h_{avg}	P_D	D_{avg}	h_{avg}	P_D	D_{avg}	h_{avg}	P_D	D_{avg}	h_{avg}	P_D	D_{avg}	h_{avg}	P_D	D_{avg}	h_{avg}
epid	trace	976	71474	1.66	619	44066	2.42	884	28101	3.82	750	82506	2.13	615	69802	2.64	888	30438	4.32
	ONE	975	62874	1.67	619	43373	2.43	887	27081	3.6	750	78270	2.15	613	68210	2.65	889	29073	4.35
2hop	trace	976	71579	1.6	604	49751	1.71	883	30900	1.85	626	82333	1.59	427	56154	1.61	848	34443	1.88
	ONE	975	60810	1.61	604	49013	1.7	885	29867	1.85	631	78772	1.59	425	54980	1.59	849	33703	1.87
ssaw	trace	976	71579	1.6	602	49444	1.71	876	32930	1.79	626	82558	1.59	423	55219	1.6	810	36253	1.83
	ONE	975	60741	1.62	602	48703	1.7	881	32006	1.80	626	82558	1.59	423	55219	1.6	810	36253	1.83

- Results' accuracy very good (within 5% in almost all cases)

Extensions

- utility-based forwarding
  - per-node additional state: time-varying node-dependent utilities for other nodes

- multicasting
  - stand-alone opportunistic networks
  - coexistence with terrestrial networks : cellular offloading scenarios

M. Karaliopoulos, C. Rohner, "**Trace-based Performance Analysis of Opportunistic Forwarding**" under submission to IEEE Transactions on Mobile Computing

EU-funded collaboration 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)
  - Autonomic Network Architecture
- (EU-IST-FET/FIRE) BIONETS (06-10) with INRIA (Eitan Altman)
  - Biologically inspired networks and services
- (EU-IST-FET/FIRE) CASCADAS (06-09)
  - Component-ware for Autonomic Situation-aware Communications, and Dynamically Adaptable Services
- (EU-IST) CONTENT (06-09)
  - Overlay networks and services for audio-visual services