The Virtue of Patience: Offloading Topical Cellular Content through Opportunistic Links

IEEE MASS 2013

Wei Peng[1] Feng Li[1] Xukai Zou[1] Jie Wu[2]

[1]Indiana University-Purdue University Indianapolis

[2]Temple University

15 October 2013

growing mobile traffic

smartphones drove a **200-fold** wireless traffic increase
for AT&T between 2007 and 2011

mobile data offloading

goal

- alleviate pressure of growing mobile traffic
- an alternative to mobile infrastructure channel

technical readiness

- increasing infrastructure-less proximity-channel bandwidth at little cost
  - NFC, Wi-Fi Direct, Bluetooth 3
- more intuitive interface
  - contact-less transfer

idea

- offload **cellular** traffic through the **proximity** channel

problem formulation

high-level overview

problem

- a piece of content
- **some** users are interested in it. . .
- . . . within some finite time
- delivery alternatives
  - cellular channel
    - instant. . .
    - but costly
  - proximity channel (NFC, Wi-Fi Direct, Bluetooth 3)
    - cheap/free. . .
    - but with uncertain delay

goal

- balance cost and delay
- without central coordination

model scope

factors included

- users' interest in content
  - in a large network, nobody desires (or is able) to consume all generated content
  - this lies behind the quest for better search engines...
  - ...and the rise of social taxonomy, or folksonomy, in tagging content
- bounded delivery-delay tolerance
  - i.e., soft real-time constraint on content delivery
  - allows *some* delay in delivering content (so users can carry the content around)...
  - ...but *not too much*, lest it becomes stale

factors not included and left for future work

- incentive: why users should participate
- privacy: minimize identifying information sharing
- enforcement: why users abide by protocol, detect black hole
- packetization, buffer, churning: all the networking details

users' interests complicate offloading strategy

shaded nodes: interested users

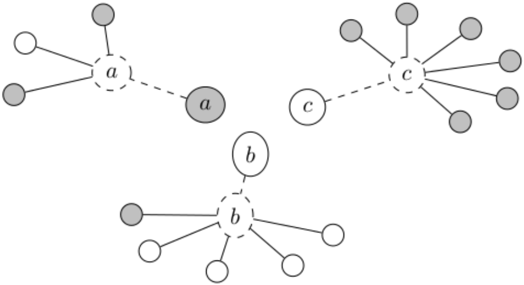

if a, b, and c meet

who shall cellular-download and **who** shall proximity-download-and-carry?

who. . . and **when**

- “**who**” was formulated in previous works as a target-set problem
  - solutions require central knowledge of users’ opportunistic topology
- . . . “**when**” is equally important

compare these offloading strategies:

- diligent: everyone cellular-downloads ASAP
  - essentially no offloading
  - no delay, but large costs
- lazy: no one cellular-downloads until someone does near deadline
  - perhaps smaller costs, but with a large delay
- interest-and-time aware: socially interested and/or little-time-left ones cellular-download
  - balance between costs and delay

the goal, the means, and the result

the goal:

interest-and-time aware + no central coordination

the means:

- users estimate their **relative social importance** with **weighted ego-centric betweenness centrality**
- users estimate their (and their acquaintances') **aggregated interests** based on **their likelihood of meeting each other**
- users consolidate *relative social importance* and *aggregated interests* in **patience**
- patience determines **cellular-download probability** over **time**

the result:

- social, content/interest, and situation awareness
- involving topologically important, but otherwise disinterested, users helps reduce cellular traffic...
- ...while satisfying users' content demand

model

elements

- content tagged by multiple tags (topics)
- I_u: tags interested by smartphone user u
- f_g: content g's freshness/expiration date
- after content is centrally released, users choose from either:
  - cellular download (instant but costly)
  - waiting for proximity-download (free but with an uncertain delay)

assumptions

- proximity links are free
  - epidemic propagation of content on proximity links
  - ignore packetization and buffer management
- users follow the protocol
  - honestly share their interests with neighbors
  - cellular download even it is only for the greater good
  - about privacy, incentive, and enforcement

- u estimates frequency of meeting its neighbors U_u based on historic encounters
- $\hat{s}_u(v)$: average consecutive-encounter delay between nodes u and v
  details
- **temporal tie strength** (tie) $s_u(v) \in [0, 1]$:

$$s_u(v) = \begin{cases} \exp(-\alpha_s \hat{s}_u(v)) & s_u(v) \in [0, +\infty), \\ 0 & s_u(v) = +\infty, \end{cases} \quad (1)$$

  1 $\Rightarrow$ strong tie; 0 $\Rightarrow$ weak tie
- $\alpha_s > 0$: a scaling parameter to prevent $s_u(v)$ from dropping too fast from increasing $\hat{s}_u(v)$

- u measures its own social importance among its neighbors U_u
- G_u: u's neighborhood weighted by $\hat{s}_u(v)$
- **weighted ego-centric betweenness centrality** $\beta_u \in [0, 1]$—the portion of shortest path passing through u:

$$\beta_u = \begin{cases} \frac{\sum_{v,w \in U_u, v \neq w} [p(v,w)]}{2\binom{|U_u|}{2}} & |U_u| \geq 2, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

- $p(v, w)$: "(v, u, w) is a shortest path between v and w"

$$[p] = \begin{cases} 1 & p \text{ is true,} \\ 0 & p \text{ is false.} \end{cases}$$

design elements
interest aggregation

- u aggregate its and its neighbors' interests on content with tag g
- I_v: v's interested tags (reported to u upon their encounters)
- u's **aggregated interest** $i_u(g) \geq 0$ on tag g:

$$i_u(g) = [g \in I_u] + \sum_{v \in U_u} s_u(v)[g \in I_v]. \quad (3)$$

- $i_u(g) < 1$ only if $g \notin I_u$.

design elements

patience and probabilistic cellular downloading strategy

- u's **patience** $p_{u,g} : [0,1] \rightarrow [0,1]$ for tag g:

$$p_{u,g}(x) = \begin{cases} \left(1 - e^{-\alpha_i i_u(g)}\right) x^{\alpha_\beta^{(1-2\beta_u)}} & g \in I_u, \\ \left(1 - e^{-\alpha_i i_u(g)}\right) (1-x)^{\alpha_\beta^{(1-2\beta_u)}} & g \notin I_u. \end{cases} \quad (4)$$

- $\alpha_i > 0$ and $\alpha_\beta > 1$: scaling parameters for $i_u(g)$ and β_u

At the moment $t + x \cdot f_g$ $(x \in [0,1])$ between:

- the time t that u first learns about a piece of content with tag g and
- the time $t + f_g$ that the content becomes stale for u

u cellular-downloads the content with a probability of:

$$p_{u,g}(x).$$

Property

If u has a higher chance of serving users (possibly including itself) before content expiration, the maximal probability that u will download the content in one round is higher.

Property

Other things being equal, more socially important users have higher cellular downloading probabilities.

Property

If u is not interested in a tag g, u's downloading probability will decrease over time; otherwise, u's downloading probability will increase over time.

analysis
patience is flexible

$$p_{u,g}(x) = \begin{cases} \left(1 - e^{-\alpha_i i_u(g)}\right) x^{\alpha_\beta^{(1-2\beta_u)}} & g \in I_u, \\ \left(1 - e^{-\alpha_i i_u(g)}\right) (1-x)^{\alpha_\beta^{(1-2\beta_u)}} & g \notin I_u. \end{cases}$$

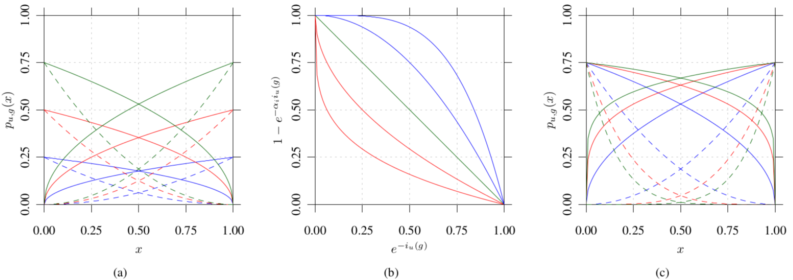

Fig. 2: The patience function $p_{u,g}$ and the scaling parameters α_i (interest $i_u(g)$) and α_β (centrality β_u). (a) Given the scaling parameters α_i and α_β, the patience $p_{u,g}$ function is jointly determined by the aggregated interest $i_u(g)$ and the centrality β_u. For $\alpha_i = 1$ and $\alpha_\beta = 2$, the patience functions corresponding to the 12 $(3 \times 2 \times 2)$ combinations $i_u(g) = 0.29, 0.69, 1.39$ (corresponding to $1 - e^{-\alpha_i i_u(g)} = 0.25, 0.50, 0.75$; blue, red, green), $\beta_u = 0, 1$ (dashed, solid), and the cases $g \in I_u, g \notin I_u$ (increasing, decreasing) are plotted for comparison. (b) The effect of the (interest) scaling parameter α_i. The maximum of the patience function $(1 - e^{-\alpha_i i_u(g)})$, which corresponds to the maximal probability that u will download the content through the cellular link in one decision, are plotted against the inverse exponential of the aggregated interest $(e^{-i_u(g)})$ with $\alpha_i = 0.25, 0.5, 1, 2, 4$ (greater than 1: blue; less than 1: red; equal to 1: green) for comparison. (c) The effect of the (centrality) scaling parameter α_β. For (interest) scaling parameter $\alpha_i = 1$ and aggregated interest $i_u(g) = 1.39$ (corresponding to the maximal cellular downloading probability $1 - e^{-\alpha_i i_u(g)} = 0.75$), the patience functions $p_{u,g}$ corresponding to the 12 $(2 \times 2 \times 3)$ combinations $\beta_u = 0, 1$ (dashed, solid), $\alpha_\beta = 2, 4, 6$ (blue, red, green), and $g \in I_u, g \notin I_u$ (increasing, decreasing) are plotted for comparison.

evaluation

dataset

- Haggle INFOCOM 2006
  - 78 attendees and 20 stationary nodes
  - conference venue in 3 days
  - time resolution: 1 second
- NUS contact
  - synthesized from the class schedules and rosters
  - students attending same session are considered to have contacts with each other
  - 1,000 students who share at least one class with others
  - time resolution: 1 hour

- 3 variants of the patience-based strategy

		eager	moderate	lazy
Haggle	α_i	0.5	0.1	0.05
	α_β	2		
	α_s	0.01		
NUS	α_i	0.05	0.03	0.01
	α_β	2		
	α_s	0.01		

  - localized collection and adaptive decision
- a previous target-set strategy (Han et al. [2012])
  - central collection of and training over user encounter traces

evaluation

Haggle results

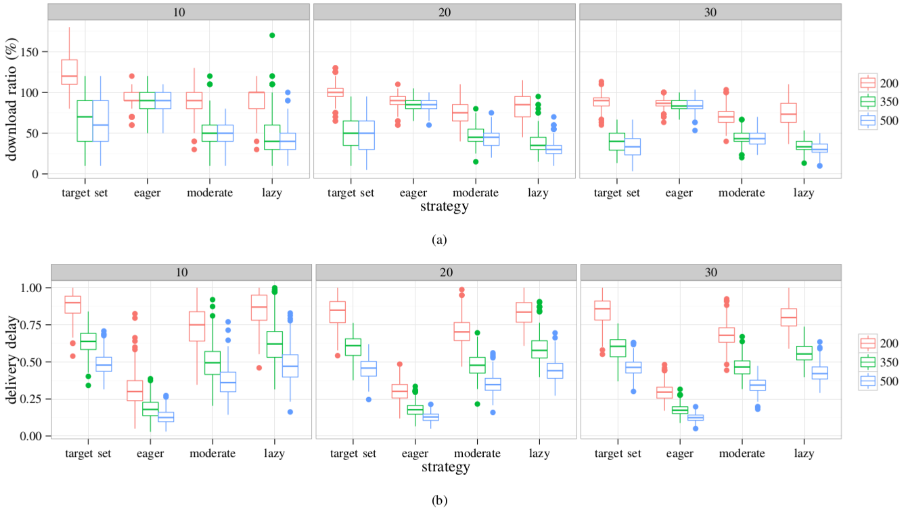

Fig. 3: Haggle: download ratio and (normalized) delivery delay. Results with different numbers of interested users (10, 20, and 30 interested users out of the 98 nodes) and content delivery deadline (200, 350, and 500 seconds) are compared. For the patience strategies, a downloading decision is made every 50 seconds.

evaluation

NUS results

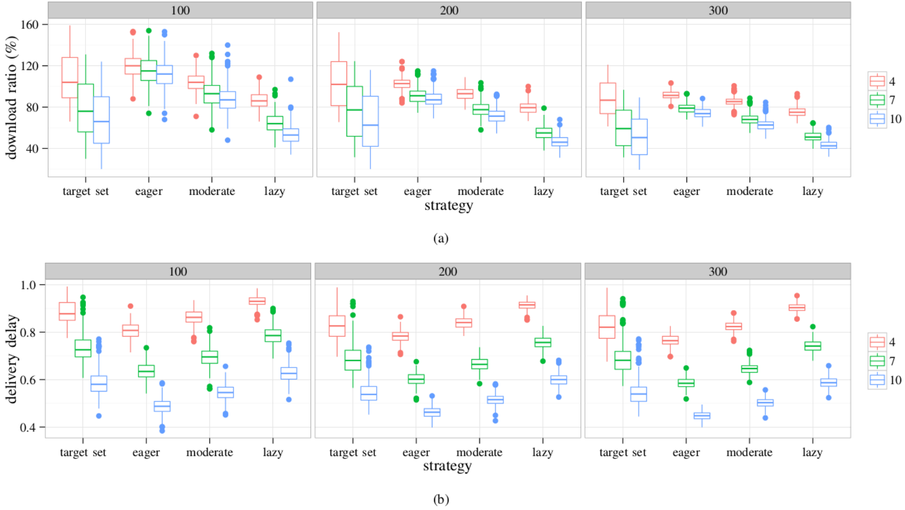

Fig. 4: NUS: download ratio and (normalized) delivery delay. Results with different numbers of interested users (100, 200, and 300 interested users out of the 1,000 nodes) and content delivery deadline (4, 7, and 10 hours) are compared. For the patience strategies, a downloading decision is made every 1 hour.

take-aways

- in offloading topical cellular content, the **virtue of patience** is to allow **the more capable** to have **better chances** of **serving the common good**
- patience function shows one approach to **locally** synthesizing **topological importance** and **content demand** for better offloading efficiency
- properly involving **topologically important**, but **otherwise disinterested**, users in downloading and forwarding content helps in reducing cellular traffic

thank you

average consecutive-encounter delay

U_u: node u's acquaintances
chronologically ordered encounters between u and $v \in U_u$:
$[a_1, b_1], \ldots, [a_k, b_k]$
current time: t
the average interval between consecutive encounters $\hat{s}_u(v)$ is defined as:

$$\hat{s}_u(v) = \begin{cases} \frac{(t - b_k) + \sum_{i=1}^{k-1}(a_{i+1} - b_i)}{k} & u \text{ and } v \text{ have met.} \\ +\infty & \text{otherwise.} \end{cases}$$

back to "temporal tie strength"