

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Content-centric wireless networks with limited buffers: when mobility hurts

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1506744> since 2016-03-14T13:53:02Z

Published version:

DOI:10.1109/TNET.2014.2361935

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Content-centric wireless networks with limited buffers: when mobility hurts

Original

Content-centric wireless networks with limited buffers: when mobility hurts / Alfano, Giuseppa; Garetto, M; Leonardi, Emilio. - In: IEEE-ACM TRANSACTIONS ON NETWORKING. - ISSN 1063-6692. - STAMPA. - 24:1(2016), pp. 299-311. [10.1109/TNET.2014.2361935]

Availability:

This version is available at: 11583/2570951 since: 2016-02-15T11:13:25Z

Publisher:

IEEE

Published

DOI:10.1109/TNET.2014.2361935

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Content-centric wireless networks with limited buffers: when mobility hurts

Giusi Alfano ^{*} Michele Garetto [†] Emilio Leonardi ^{*}

^{*} Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Italy

[†] Dipartimento di Informatica, Università di Torino, Italy

Abstract—We analyze throughput-delay scaling laws of mobile ad-hoc networks under a content-centric traffic scenario, where users are mainly interested in retrieving contents cached by other nodes. We assume limited buffer size available at each node and Zipf-like content popularity. We consider nodes uniformly visiting the network area according to a random-walk mobility model, whose flight size varies from the typical distance among the nodes (quasi-static case) up to the edge length of the network area (reshuffling mobility model). Our main findings are i) the best throughput-delay trade-offs are achieved in the quasi-static case: increasing the mobility degree of nodes leads to worse and worse performance; ii) the best throughput-delay trade-offs can be recovered by power control (i.e., by adapting the transmission range to the content) even in the complete reshuffling case.

Index Terms—Ad-hoc networks, content-centric networking, scaling laws, delay-throughput trade-off.

I. INTRODUCTION AND RELATED WORK

During the past several years, we have witnessed a gradual shift in the way users search and retrieve data from the Internet: the traditional host-to-host communication paradigm has evolved towards a new host-to-content kind of interaction, in which the main networking functionalities are directly driven by object identifiers, rather than host addresses. This change has been promoted by the great success obtained by Content Delivery Networks (CDNs), which represent nowadays the standard solution adopted by content providers to serve large populations of geographically spread users. The extreme of this new way of thinking about the Internet has been perhaps reached by recent Content-Centric-Networking proposals (CCNs), which aim at redesigning the entire Internet architecture, including core routers, with named data as the central element of the communication [1]. A key component of both CDNs and CCNs is the content replication strategy, i.e., how many copies of the available contents to put in the network, and where. High-performing, distributed and self-adapting caching solutions still represent one of the main challenges in this area.

It is inevitable that content-based networking will also affect the wireless domain, and this has already started in academic research. As observed in [2], the most celebrated results about the scalability of wireless networks (such as Gupta-Kumar [3], Grossglauser-Tse [4]) have pushed researchers to mainly consider the scenario in which n end-to-end flows are randomly established among the nodes. However, this (unicast) traffic

pattern is not suitable to describe content-centric networks, where users are primarily interested in retrieving objects: as contents may be cached in multiple nodes in the network, requests can be served from multiple locations (anycast), and they are typically directed to the closest node to save network resources and improve the user-perceived performance.

On the other hand, existing works departing from the assumption of unicast communications (i.e., those considering either multicast or anycast traffic) have mainly focused on the case of static networks [5], [6], [7]. We believe that a significant gap still exists in the asymptotic analysis of wireless networks, when we jointly consider anycast (content-centric) communications and node mobility. In this paper, we seek to partially fill this gap by considering a content-centric wireless network in which nodes are mobile. Given the tremendous number of different rules of the game that one could choose to study this problem, we decided to maintain the same assumptions adopted in the recent paper [2]. The two most important ones are: limited buffer size available at each node, and Zipf-like content popularity. Instead of the static grid topology considered in [2], we let the nodes independently and uniformly move over the network area. By varying the flight size of the random walk mobility model, we obtain a family of throughput-delay trade-offs, ranging from a quasi-static case to a fully mobile scenario similar to the reshuffling mobility model.

Interestingly, we discover that the best throughput-delay trade-offs are obtained in the quasi-static case: increasing the mobility degree of nodes leads to worse and worse performance. A rather surprising result is that the best throughput-delay trade-offs (i.e., those achievable under static or quasi-static conditions) can be recovered by power control (i.e., by adapting the transmission range to the content) even in the extreme case of the reshuffling mobility model.

Throughput-delay trade-offs in mobile networks under unicast traffic have been investigated in [8], [9], [10], [11], [12], [13] for various mobility models. Especially relevant to our work is [14], where authors show that if buffer sizes are not scaled appropriately, the scaling law for the throughput capacity of mobile networks is not significantly better than that for static networks.

A remarkable application of our theoretical analysis is the recent idea of exploiting device-to-device, opportunistic communications among mobile users to reduce the downlink traffic in cellular networks [15], [16], [17].

To the best of our knowledge, we are the first to analyse delay-throughput trade-offs in mobile wireless networks under content-centric (anycast) traffic. Our main contribution is to show that an apparently hard joint optimization problem, which combines interference, mobility, scheduling, routing, content replication, zipf-like content popularity, can nonetheless be reduced to simple constrained optimization problems that are solvable by standard Lagrangian relaxation. Although we rely on assumptions and techniques previously adopted in the literature, our results provide novel and interesting insights about the impact of mobility in content-centric scenarios.

II. SYSTEM ASSUMPTIONS

A. Network and mobility model

We consider an extended network comprising N nodes moving over a square region \mathcal{O} of area N with wrap-around conditions (i.e., a torus), to avoid border effects. Time is divided into slots of equal duration, which is normalized to 1. For what concerns nodes' mobility, we will first consider the simple case in which the position of every node is updated at the beginning of each slot by choosing a new location uniformly at random in the network area, independently of other nodes. Such a model has been called differently in the literature, as reshuffling model, or bi-dimensional i.i.d. mobility model [8], [9], [10]. In this work we will refer to it as the *reshuffling* model. This mobility pattern turns out to be very simple to analyze, although it is clearly unrealistic, as nodes are allowed to instantaneously jump to arbitrarily far positions in the network area. For this reason, we will later generalize our analysis to the case in which nodes move according to independent random walks with average flight size F .

B. Traffic model

We assume there are M contents available in the system, and we let M grow to infinite as the number of nodes increases. In particular, we will focus on the case¹ $M = \Theta(N^\beta)$, with $0 \leq \beta \leq 1$. We consider that all contents have the same (unit) size².

We assume that nodes have limited storage capacity. This turns out to be a crucial (but realistic) assumption, as explained later in Section III. In particular, let K be the storage capacity of each node, measured in number of (equal-size) contents. Similarly to [2], we assume that the set of contents stored by each node is a-priori, statically determined by the system, that can choose the number of replicas for each content (on the basis of its popularity) and pre-populate the caches of all

¹Given two functions $f(n) \geq 0$ and $g(n) \geq 0$: $f(n) = o(g(n))$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$; $f(n) = O(g(n))$ means $\limsup_{n \rightarrow \infty} f(n)/g(n) = c < \infty$; $f(n) = \omega(g(n))$ is equivalent to $g(n) = o(f(n))$; $f(n) = \Omega(g(n))$ is equivalent to $g(n) = O(f(n))$; $f(n) = \Theta(g(n))$ means $f(n) = O(g(n))$ and $g(n) = O(f(n))$; at last $f(n) \sim g(n)$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$.

²Our results extend easily to the general case of variable-size contents, provided that: i) the content size does not scale with N ii) the ratio between the largest and the smallest content size is bounded by a constant. In Sect V we address this point.

nodes. Notice that this assumption implies that we have a static set of contents with known popularity³.

We consider a Zipf's law for the content popularity distribution, which is frequently observed in traffic measurements and widely adopted in performance evaluation studies [18], [19]. This law implies that, having sorted the contents in decreasing order of popularity, a request is directed to content i with probability

$$p_i = \frac{H}{i^\alpha}, \quad 1 \leq i \leq M \quad (1)$$

where α is the Zipf's law exponent, and $H = (\sum_{i=1}^M i^{-\alpha})^{-1}$ is a normalization constant. We have:

$$H = \begin{cases} \Theta(1) & , \alpha > 1 \\ \Theta(1/\log M) & , \alpha = 1 \\ \Theta(M^{\alpha-1}) & , \alpha < 1 \end{cases} \quad (2)$$

We assume that users request contents according to the following sequential process: each node i) generates a content request according to the probability law (1); ii) it waits until it retrieves the requested content; iii) it further waits for a random idle time I with average \bar{I} ; iv) it generates another request; and so on. Idle times, which are assumed to form an i.i.d. sequence for each node, are introduced in the model to trade-off throughput and delay. Indeed, according to the above request generation process, node throughput λ (expressed in contents/slot) and average content transfer delay \bar{D} (expressed in slots) are tightly related by the following equation:

$$\lambda = \frac{1}{\bar{D} + \bar{I}} \quad (3)$$

as consequence of elementary renewal theory arguments. Note that each node has at most one pending content request at any given time⁴.

C. Communication Model

To account for interference among simultaneous transmissions, we adopt a generalized version of the classical protocol model, in which different transmission ranges are allowed (see [20]). More specifically, we assume that nodes can adapt their transmission range to the content being transmitted, i.e., content i is transmitted employing transmission range R_i , with $R_i = \Omega(1)$. Note that this model includes, as a special case, also the standard protocol model in which there is a unique transmission range R , by letting $R_i = R, \forall i$.

According to our generalized protocol model, the transmission of a content i from node s to node d is feasible if and only if the following conditions hold:

- 1) the distance between s and d is smaller than or equal to R_i , i.e., $d_{sd}(t) \leq R_i$.
- 2) for every other node k simultaneously transmitting, $d_{kd}(t) \geq (1 + \Delta) R_i$, being Δ a guard factor.

³We leave to future work the analysis of the case in which the set of available contents (and their popularity) can change over time, as well as other forms of run-time optimization.

⁴The scaling order of our results does not change when nodes are allowed to send out multiple contents requests in parallel, provided that the number of pending requests at each node is bounded.

We emphasize that, for a large class of power attenuation functions, the above generalized version of the protocol model can be easily shown to be in order sense equivalent to the physical model introduced in [3] for point-to-point communications over the Gaussian channel, according to which a transmission between two nodes is successful if and only if the signal-to-interference-and-noise ratio (SINR) at the receiver is larger than a given threshold.

In Appendix A we prove that the above equivalence between protocol and physical models indeed holds also in our system, where we adopt more restrictive assumptions with respect to those considered in [20], where a similar proof was originally proposed.

First, to avoid well-known divergence problems at short distances, we assume the power attenuation at distance d to be $l(d) = \min(1, d^{-\gamma})$, with $\gamma > 2$. Note that previous expression, which has been largely used in previous work, guarantees that $l(d)$ is indeed an attenuation, for any value of d . Then, to compensate for the signal attenuation, we assume that a node transmitting content i employs power $P_i = P/l(R_i)$. By so doing, the useful signal arrives at the receiver with power at least equal to P , where P is a given constant. Note that the above power-control strategy does not perfectly compensate the power attenuation as a function of the specific distance between transmitter and receiver. This, together with our choice of power-attenuation function, are essentially the reasons why we need a different proof of the equivalence between protocol and physical model (see Appendix A). Having established this equivalence, in the following we will consider only the generalized protocol model.

When a successful transmission occurs, we assume that the total amount data transferred during the slot is large enough to permit the transfer of one content from the sender to the receiver. Although this assumption may appear to be simplistic, it is not a critical one: the same asymptotic results for throughput and delay are obtained for the case in which one successful transmission allows to transfer only one segment of the content file, as long as each content can be split into a bounded number of segments. The impact of variable-size contents will be discussed in more details in Section V.

On the other hand, previous work [11] has shown that by arbitrarily reducing the size of data segments exchanged between two nodes (up to the limit case in which the file can be considered as a fluid), one can achieve improved performance in order sense, since multi-hop communications become feasible during each slot. In our work, we do not consider this possibility, restricting our attention to the case in which data can be transmitted over a bounded number of hops during a slot. Table I summarizes the adopted notation.

III. RESHUFFLING MOBILITY MODEL

We start analyzing the network performance achievable under the reshuffling mobility model. One fundamental point to understand is that, in our considered system:

Proposition 1: Under the reshuffling mobility model, the network performance cannot be improved by making nodes

TABLE I
NOTATION

Symbol	Definition
N	number of nodes
M	number of contents
β	growth exponent of M : $M = \Theta(N^\beta)$, $0 \leq \beta \leq 1$
K	number of contents stored by each node, $K = \Theta(1)$
p_i	probability to request content i
α	Zipf's law exponent of content popularity
R_i	transmission range employed to transmit content i
λ	asymptotic per-node throughput (contents/slot)
\bar{D}	asymptotic average content transfer delay (slots)
F	average flight length

relay contents for other nodes, i.e., by delivering contents over multi-hop routes.

Proof: The assertion is a consequence of the fact that we jointly assume that: i) a message transmission occupies a non-vanishing fraction of each time slot (i.e., we cannot transfer arbitrarily small content pieces, like in the fluid limit); ii) nodes have a finite storage capacity; iii) the network topology is completely reshuffled at each step.

As shown in [14], assumption ii) implies that store-carry-and-forward schemes, such as the celebrated two-hops scheme proposed by Grossglauser-Tse [4], can not be exploited in our case to increase the transmission opportunities among the nodes. Indeed, as consequence of i) and ii) a node can only store packets destined to a finite number of destinations. Hence its transmission opportunities (which determine throughput and delay) scale in the same way as if it were responsible for transmitting only its own contents (we refer the reader to [14]–Section V, for a complete proof of this assertion).

Furthermore we cannot employ a multi-hop route over multiple slots to progressively move data closer and closer to the destination, as consequence of the fact that, after each slot, nodes move to totally different, arbitrary locations (assumption iii). This fact, which looks rather intuitive, can be formally derived from Lemma 6 in Section IV-A (setting $F = \Theta(\sqrt{N})$), whose proof is given in [11].

At last observe that we could, in principle, perform a multi-hop route within a single slot, but assumption i) implies that we can only make a finite number of hops, which does not improve the network performance in order sense. ■

From the above proposition, it follows that we can restrict ourselves to the case in which communications occur over just a single hop, i.e., when a node requesting a given content falls within the communication range of a node storing a copy of it. We will first consider in Section III-A the case in which the transmission range is the same for all contents. Later on, in Section III-B we will analyze the gains achievable by adapting the transmission range to the content.

A. Fixed transmission range

Let R be the common transmission range employed by all transmissions. We first introduce some definitions and existing results:

Definition 1: feasible tx-rx pair. A pair of nodes $\{i, j\}$ is defined to be a feasible transmitter-receiver pair (tx-rx pair) in a given time slot, if and only if the following conditions hold:

i) one node, say node j , has a pending request for a certain content m ; ii) node i stores content m in its cache. iii) the distance between i and j is smaller than or equal to R . Notice that a feasible tx-rx pair is not necessarily enabled to transmit by the scheduling scheme, i.e., it represents only a transmission opportunity. A feasible tx-rx pair $\{i, j\}$ is said m -feasible tx-rx pair if node i stores content m and node j is interested to m .

Definition 2: Active square. A square region of the network is defined to be active if it contains at least one feasible tx-rx pair, i.e., a feasible pair such that both transmitter and receiver lie in the considered square.

Notice that, if we want a square to be active with non vanishing probability (as we increase the number of nodes), a necessary condition is that the mean number of tx-rx pairs falling in it does not vanish. It would not be difficult to compute exactly the mean number $\rho(S)$ of tx-rx pairs falling in an arbitrary square of size S . Such mean depends on several factors⁵, including obviously the square size, the replication strategy, the transmission range R , and the probability that in an arbitrary slot a node has an active pending request for a given content. We do not present the exact expression of $\rho(S)$ as function of the above variables, because we do not really need it. For our purposes, we only need to establish an important property of $\rho(S)$, stated in the following lemma.

Lemma 1: Consider the case in which $S = \Theta(R^2)$. Then $\rho(S)$ increases quadratically (in order sense) with the square area S i.e., $\rho(S^2) = \Theta(S^2/S_0^2)$, for some S_0 such that $\rho(S_0) = \Theta(1)$.

Proof: Since by hypothesis $S = \Theta(R^2)$, without lack of generality we fix $R = \sqrt{S}$. Recall that every node has at most one pending request, hence: $\rho(S) = \sum_m \rho_m(S)$, where $\rho_m(S)$ is the average number of m -feasible tx-rx pairs falling in the considered square of area S . Indeed, by construction, exactly one content (the content the receiver is interested to) can be exchanged between any tx-rx pair contributing to $\rho(S)$.

Observe that $\rho_m(S)$ is proportional to the product of the number m -tx of nodes in the square storing a copy of content m and the number m -rx of nodes in the square interested to content m . Since in the considered square both m -tx and m -rx scale linearly with the area size S , $\rho_m(S)$ scales quadratically with S . The assert descends immediately, observing that by construction an $S_0 \leq 1$ can always be found such that $\rho(S_0) = \Theta(1)$. ■

Definition 3: contact probability. The contact probability $p_{\text{contact}}(m)$ associated to a given content m is defined as the probability that a node having a pending request for content m falls, in a given slot, within the transmission range of a node holding a copy of content m .

Lemma 2: The contact probability for content m satisfies $p_{\text{contact}}(m) = \Theta(\min(1, \frac{X_m R^2}{N}))$, where X_m is the number of replicas of content m .

Proof: Each of the X_m replicas of m falls in a disc of radius R around the requesting node with probability $\frac{\pi R^2}{N}$. Hence $p_{\text{contact}}(m) = 1 - (1 - \frac{\pi R^2}{N})^{X_m}$ which is in order

sense equivalent to $\min(1, \frac{X_m R^2}{N})$. ■

Corollary 1: Given the number of replicas X_m of content m , the average transfer delay \bar{D}_m associated to content m satisfies $\bar{D}_m = \Omega\left(\frac{1}{\min(1, \frac{X_m R^2}{N})}\right)$.

Proof: The delay associated to content m is lower bounded by the times it takes to a node requesting content m to come in contact with a node storing a copy of m , which is geometrically distributed with mean $1/p_{\text{contact}}(m)$. The results follows applying Lemma 2. ■

We now recall a basic result well known in the literature:

Lemma 3: The aggregate transmission rate (also called network capacity) Λ of a network of area A employing a protocol model with transmission range R , satisfies $\Lambda = O(A/R^2)$. Network capacity $\Lambda = \Theta(A/R^2)$ can be attained when the average number of tx-rx pairs in an arbitrary square of area $S = \Theta(R^2)$ is not vanishing.

We do not repeat the details of the scheduling scheme that allows to achieve (in order sense) the maximum network capacity under the protocol model, since such a scheme is well known in the literature (see for example [21]). Essentially, the network area is divided into squarelets of area $S = R^2$. The entire set of squarelets is then partitioned into a finite number of subsets, such that squarelets belonging to the same subset are sufficiently spaced apart (depending on the guard factor Δ) to permit scheduling an active transmission in each squarelet of the subset⁶. In any given slot, one subset is uniformly selected, and at most one tx-rx pair is enabled to transmit in each squarelet belonging to the selected subset. A network capacity in order sense equal to the number of squarelets can be achieved, provided that the average number of tx-rx pairs that can be enabled in an arbitrary squarelet is non vanishing (which implies that the probability that a squarelet is active is non vanishing).

Remark. Observe that, for the scheme we are designing, the network capacity equals the aggregate network throughput $\Lambda = N\lambda$, since contents are transferred over a single hop.

Previous lemmas allow us to establish our first fundamental result.

Theorem 1: Consider nodes generating requests according to the sequential model described in Section II-B, with specified average idle time \bar{I} . Given a replication strategy (i.e., given X_m for any m), the optimal network performance in terms of throughput and delay is achieved by selecting a transmission range \hat{R} such that the average number $\rho(\hat{R}^2)$ of tx-rx pairs in a square of area \hat{R}^2 satisfies $0 < c_1 < \rho(\hat{R}^2) < c_2$, where c_1 and c_2 are constants.

The proof is reported in Appendix B.

Remark. The optimal value \hat{R} for the transmission range characterized by previous Theorem depends on the chosen average idle time \bar{I} at nodes. Different trade-offs can be achieved by controlling \bar{I} , i.e. the interval between the reception of a content and the next content request. Indeed, observe that on the one hand from Theorem 1 and Lemma 3 the per-node throughput is tightly coupled to \hat{R} , by the relationship $\lambda = \frac{1}{\bar{R}^2}$.

⁵To simplify the notation we explicit only the dependency on S , characterized in Lemma 1.

⁶A similar scheme can be applied under the physical model, obtaining that squarelets belonging to the same subset can only weakly interfere with each other, guaranteeing one successful transmission in each squarelet.

On the other hand per-node throughput, average delay and average idle time are related by: $\lambda = \frac{1}{\bar{D} + \bar{I}}$. Thus, as long as the target throughput λ is feasible, i.e., it is smaller than the inverse of the target average delay \bar{D} , we can properly set \bar{I} so as to achieve the desired trade-off.

For what concerns the average delay \bar{D} , the most important consequence of Theorem 1 is stated in the following:

Corollary 2: Given a replication strategy (i.e., given X_m for any m), the average content transfer delay behaves asymptotically as:

$$\bar{D} = \Theta \left(\sum_{m=1}^M \frac{p_m}{\min(1, \frac{X_m \hat{R}^2}{N})} \right) \quad (4)$$

Proof: Theorem 1 guarantees that, by adopting the optimal transmission range \hat{R} , the delay experienced by any content transfer attains its lower bound in Corollary 1. Averaging over all contents, we obtain the provided expression for \bar{D} . ■

Let us now assume that a feasible per-node throughput λ (and the corresponding transmission range $\hat{R} = \sqrt{1/\lambda}$) has been chosen. Among all the possible replication strategies $\{X_m\}_m$, the optimal will be the one that minimizes the associated average delay \bar{D} in (4). Indeed, by selecting such replication strategy we achieve the best possible delay performance among all strategies guaranteeing the target throughput.

The optimal scheme can thus be found in two steps, by first identifying the minimum possible delay and the associated optimal replication strategy:

$$\begin{cases} \min_{\{X_m\}, m=1 \dots M} \sum_{m=1}^M \frac{p_m}{\min(1, \frac{X_m \hat{R}^2}{N})} \\ \text{s.t. } \sum_{m=1}^M X_m \leq KN \\ 1 \leq X_m \leq N \quad m = 1 \dots M \end{cases} \quad (5)$$

and then deriving the value of \bar{I} , so as to meet condition (3).

Focusing on the optimization problem (5), we observe that it is clearly better to allocate more replicas to the most popular contents, i.e., those having smaller index m . Hence the sequence $\{X_m\}$ should be non-increasing. However, the term $\min(1, \frac{X_m \hat{R}^2}{N})$ in the objective function tells us that it is useless to replicate any content more than $X^* = \lceil N/\hat{R}^2 \rceil$ times. Therefore, let $m^* \geq 0$ be the index such that all contents with index $m \leq m^*$ are replicated X^* times (if such contents do not exist, $m^* = 0$). These m^* most popular contents will consume $m^* X^*$ aggregate buffer space. We can assume that the remaining buffer space left for the least popular contents having index $m > m^*$ is still of order N . This assumption can be checked a-posteriori, but can be easily believed to be true by considering that the optimal delay in order sense should not be sensitive to the specific constant K . Hence we can always devote $K^* N$ aggregate buffer space, with $K^* \leq K$ independent of N , to the least popular contents without affecting the asymptotic results. The above considerations allow us to analyze the reduced optimization problem, valid for contents

of index $m > m^*$:

$$\begin{cases} \min_{\{X_m\}, m > m^*} \sum_{m > m^*} \frac{N p_m}{X_m \hat{R}^2} \\ \text{s.t. } \sum_{m > m^*} X_m \leq K^* N \\ 1 \leq X_m \leq N \quad m > m^* \end{cases} \quad (6)$$

We now have all ingredients to prove our main result for the considered scenario:

Theorem 2: The throughput-delay performance achievable under the reshuffling model with uniform transmission range depends on the Zipf's law exponent α :

- For $\alpha > 2$, it is possible to achieve the best possible throughput $\lambda = \Theta(1)$ and the best possible delay $\bar{D} = \Theta(1)$, using transmission range $R = \Theta(1)$.
- For $1 < \alpha < 2$, the optimal throughput-delay trade-off is $\bar{D} = \Theta(\lambda M^{2-\alpha})$. The minimum delay $\bar{D} = \Theta(1)$ can be achieved with $R = \Theta(M^{1-\alpha/2})$, and the associated throughput is $\lambda = \Theta(M^{\alpha-2})$. The maximum throughput $\lambda = \Theta(M^{\alpha/2-1})$ can be achieved with $R = \Theta(M^{1/2-\alpha/4})$, and the associated delay is $\bar{D} = \Theta(M^{1-\alpha/2})$.
- For $\alpha < 1$, the optimal throughput-delay trade-off is $\bar{D} = \Theta(\lambda M)$. The minimum delay $\bar{D} = \Theta(1)$ can be achieved with $R = \Theta(\sqrt{M})$, and the associated throughput is $\lambda = \Theta(1/M)$. The maximum throughput $\lambda = \Theta(1/\sqrt{M})$ can be achieved with $R = \Theta(M^{1/4})$ and the associated delay is $\bar{D} = \Theta(\sqrt{M})$.

The proof is reported in Appendix C. Note that Appendix C also includes an algorithm to distribute content replicas among the nodes, which guarantees that the buffer size needed at each node can be upper bounded by a constant.

Results obtained so far for the case of fixed transmission range are summarized in Table II.

B. Different transmission ranges

We now consider the case in which the transmission range can be adapted to the content being transmitted. The analysis goes along the same lines followed in Section III-A. We will consider only the case $\alpha < 2$, since for $\alpha > 2$ we already achieve the best possible performance in terms of both throughput and delay by employing a fixed transmission range for all contents (Theorem 2). We start with the following lemma:

Lemma 4: The contact probability for content m satisfies $p_{\text{contact}}(m) = \Theta(\min(1, \frac{X_m R_m^2}{N}))$, whose proof is analogous to that of Lemma 2. Similarly to before, it immediately follows that,

Corollary 3: The average delay \bar{D}_m associated to content m satisfies $\bar{D}_m = \Omega(\frac{1}{\min(1, X_m R_m^2/N)})$.

In the case of different transmission-ranges, the selection of the optimal set of feasible tx-rx pairs to be enabled in the network at a given time slot is not a trivial task. First, we characterize the maximum network capacity achievable by employing a given set of transmission ranges $\{R_m\}_m$, by the following result analogous to Lemma 3:

Lemma 5: The aggregate transmission rate Λ of a network of area A , such that contents of type m , transmitted

TABLE II
RESHUFFLING MOBILITY WITH FIXED RADIUS

Zipf's exponent	trade-off	$[\lambda_{\min}, \lambda_{\max}]$	$[\bar{D}_{\min}, \bar{D}_{\max}]$
$\alpha > 2$		$\lambda = \Theta(1)$	$\bar{D} = \Theta(1)$
$1 < \alpha < 2$	$\bar{D} = \Theta(\lambda M^{2-\alpha})$	$[\Theta(M^{\alpha-2}), \Theta(M^{\frac{\alpha}{2}-1})]$	$[\Theta(1), \Theta(M^{1-\frac{\alpha}{2}})]$
$\alpha < 1$	$\bar{D} = \Theta(\lambda M)$	$[\Theta(\frac{1}{M}), \Theta(\frac{1}{\sqrt{M}})]$	$[\Theta(1), \Theta(\sqrt{M})]$

TABLE III
RESHUFFLING MOBILITY WITH RADIUS ADAPTATION

Zipf's exponent	trade-off	$[\lambda_{\min}, \lambda_{\max}]$	$[\bar{D}_{\min}, \bar{D}_{\max}]$
$\alpha > \frac{3}{2}$		$\lambda = \Theta(1)$	$\bar{D} = \Theta(1)$
$1 < \alpha < \frac{3}{2}$	$\bar{D} = \Theta(\lambda M^{3-2\alpha})$	$[\Theta(M^{2\alpha-3}), \Theta(M^{\alpha-\frac{3}{2}})]$	$[\Theta(1), \Theta(M^{\frac{3}{2}-\alpha})]$
$\alpha < 1$	$\bar{D} = \Theta(\lambda M)$	$[\Theta(\frac{1}{M}), \Theta(\frac{1}{\sqrt{M}})]$	$[\Theta(1), \Theta(\sqrt{M})]$




Fig. 1. Illustration of the scheduling scheme in the case of variable transmission ranges, for $0 < \Delta < 1$ ($q = 3$).

with probability p_m , employ transmission range R_m , satisfies $\Lambda = O(A/\bar{S})$, where $\bar{S} = \sum_{i=1}^M p_i R_i^2$. Network capacity $\Lambda = \Theta(A/\bar{S})$ can be attained if the average number of m -feasible tx-rx pairs in a square of area R_m^2 is not vanishing.

Proof: Suppose that the network sustains a given network capacity Λ . Then, the average number of contents of type m that are sent in each slot is $\lambda_m = \Lambda p_m$. The transmission of a content of type m ‘consumes’ an area of size $Z_m = \pi(1 + \Delta)^2 R_m^2$, since we cannot put any other transmitter within the disc of area Z_m centered at the receiver. Therefore, considering the ideal (infeasible) case in which we can exploit the whole network area A to allocate transmissions (ideal packing), we obtain that Λ must satisfy the inequality $\Lambda \sum_{m=1}^M p_m Z_m \leq A$, from which we derive the upper bound $\Lambda = O(A/\bar{S})$.

A constructive scheduling scheme to achieve (in order sense) $\Lambda = \Theta(A/\bar{S})$ is the following. We partition the set of transmission ranges $\{R_m\}_m$ into a sequence of classes $i = 1, 2, \dots$ such that class i contains all transmission ranges R_m such that $R_{\max}/q^i < R_m \leq R_{\max}/q^{i-1}$, where R_{\max} is the maximum transmission range employed in the network,

and $q = \lceil 2 + \Delta \rceil$. Let \mathcal{M}_i denote the subset of indexes m such that R_m belongs to class i . Notice that transmission ranges falling in a given class have comparable sizes, meaning that there can be at most a factor q between the largest and the smallest transmission range in the class. The idea is to first allocate tx-rx pairs whose transmission range belongs to class 1. In the remaining network area, we proceed to allocate tx-rx pairs belonging to class 2, and so on. By so doing, we obtain a scheduling scheme with optimal (in order sense) spatial reuse, while achieving a feasible packing of transmissions employing different ranges⁷. More in detail, we first consider class 1, and partition the network area into squarelets of area R_{\max}^2 . Similarly to the traditional scheme for fixed transmission range, we can partition the squarelets into a finite number (q^2) of subsets, such that squarelets in each subset can be concurrently active. Fig. 1 illustrates the proposed scheme in the case of $0 < \Delta < 1$ ($q = 3$). In each slot, we first activate a number of squarelets of class 1 at most equal to $\Lambda_1 = \psi \sum_{m \in \mathcal{M}_1} \frac{A p_m}{\bar{S}}$, where ψ is an arbitrary constant smaller than 1. Constant ψ is introduced to guarantee that the priority assigned to tx-rx pairs employing larger ranges does not penalize tx-rx pairs employing smaller ranges. Notice that, if, in accordance with the Lemma assumptions, the average number of m -feasible tx-rx pairs in a square of area R_m^2 is not vanishing, we can surely activate $\Theta(\Lambda_1)$ squarelets, since the size of class-1 squarelets is bigger than or equal to R_m^2 , for all $m \in \mathcal{M}_1$. Notice that the precise positions of the class-1 squarelets to be activated are not important. After enabling class-1 squarelets, we remove them from the network, together with their guard zones, and divide the remaining network area into squarelets of edge R_{\max}/q , moving on to class 2, and so on for all classes. In the end, we obtain an aggregate rate $\Lambda' = \Theta(\frac{A}{\bar{S}})$, while fairly assigning transmission opportunities to contents in such a way that each content m achieves throughput $\Theta(p_m \frac{A}{\bar{S}})$. ■

⁷A selection of tx-rx pairs in arbitrary order, ignoring transmission ranges, does not in principle guarantee that enough tx-rx pairs employing large transmission ranges can be activated in the ‘holes’ not occupied by transmissions employing small transmission ranges.

The above Lemma allows us to establish a result similar to Theorem 1:

Theorem 3: Consider nodes generating requests according to the sequential model described in Section II-B, with specified average idle time \bar{I} . Given a replication strategy (i.e., given X_m for any m), the optimal network performance in terms of throughput and delay is achieved by selecting transmission ranges R_m such that the average number ρ_m of m -feasible tx-rx pairs in a square of area R_m^2 can be lower- and upper-bounded by two constants, for each m .

Proof:

First we observe that, if a set $\{\hat{R}_m\}_m$ of transmission ranges satisfying the above condition indeed exists (this will be proven later), we could achieve the network capacity $\Lambda = \Theta(N/\bar{S})$, with $\bar{S} = \sum_{i=1}^M p_i R_i^2$, applying the scheme presented in the proof of Lemma 5. The chosen set of transmission ranges, and the associated scheme, turn out to be optimal both in terms of throughput and in terms of delay. In terms of throughput, it is easy to see that we cannot achieve any higher throughput in order sense, by either increasing or decreasing any transmission range in the set. If we increase any transmission range, the throughput would decrease according to Lemma 5; if we decrease any transmission range (actually, all transmission ranges should be jointly reduced), the gain deriving from the potential higher spatial reuse cannot be exploited since the mean number of m -feasible tx-rx pairs decreases quadratically with $S_m = R_m^2$, for any m . In terms of delay, we observe that set $\{\hat{R}_m\}$ guarantees that the delay associated to any content achieves its lower bound as in Corollary 3. At last, the existence of a set $\{\hat{R}_m\}$ satisfying the requested constraint descends from the fact that $\rho_m(R_m^2)$ is an increasing function of R_m , for any m . ■

The most important consequence of Theorem 3 is stated in the following

Corollary 4: Given a replication strategy (i.e., given X_m for any m), the average network delay behaves asymptotically as:

$$\bar{D} = \Theta \left(\sum_{m=1}^M \frac{p_m}{\min(1, X_m \hat{R}_m^2 / N)} \right) \quad (7)$$

whose proof is identical to that of Corollary 2.

Let us now assume that a feasible per-node throughput λ (and the corresponding average square size $\bar{S} = 1/\lambda$) has been chosen. The associated average delay \bar{D} in (7) can be optimized in terms of both the number of replicas X_m and the transmission ranges R_m . We will actually optimize the performance in terms of square sizes $S_m = R_m^2$:

$$\left\{ \begin{array}{l} \min_{\{X_m, S_m\}} \sum_{m=1}^M \frac{p_m}{\min(1, \frac{X_m S_m}{N})} \\ \text{s.t.} \quad \sum_{m=1}^M X_m \leq K N \\ \quad \quad 1 \leq X_m \leq N \quad m = 1 \dots M \\ \quad \quad \sum_{m=1}^M p_m S_m = \bar{S} \end{array} \right. \quad (8)$$

Considerations analogous to those reported in Section III-A allow us to analyze the following reduced optimization problem, valid for contents of index $m > m^*$, where m^* is the

maximum index for which the delay D_m attains its minimum value of 1.

$$\left\{ \begin{array}{l} \min_{\{X_m, S_m\}, m > m^*} \sum_{m > m^*} \frac{N p_m}{X_m S_m} \\ \text{s.t.} \quad \sum_{m > m^*} X_m \leq K^* N \\ \quad \quad 1 \leq X_m \leq N \quad m > m^* \\ \quad \quad \sum_{m > m^*} p_m S_m = \bar{S}^* \end{array} \right. \quad (9)$$

where $\bar{S}^* = \bar{S} - \sum_{m \leq m^*} p_m S_m$.

We now have all ingredients to prove our main result for the considered scenario:

Theorem 4: By adapting the transmission range to the content, it is possible to improve the throughput-delay performance achievable under the reshuffling model:

- For $\alpha > 3/2$, it is possible to achieve the best possible throughput $\lambda = \Theta(1)$ and the best possible delay $\bar{D} = \Theta(1)$.
- For $1 < \alpha < 3/2$, the optimal throughput-delay trade-off is $\bar{D} = \Theta(\lambda M^{3-2\alpha})$, with $\bar{D} = \Omega(1)$ and $\bar{D} = O(M^{3/2-\alpha})$.
- For $\alpha < 1$, the optimal throughput-delay trade-off is $\bar{D} = \Theta(\lambda M)$, with $\bar{D} = \Omega(1)$ and $\bar{D} = O(\sqrt{M})$.

The proof is reported in Appendix D. Results for this case are summarized in Table III. We remark that the solution to (9) leads to an optimal replication strategy in which X_m is proportional to $p_m^{2/3}$, which is similar to the optimal replication strategy found in [2] in a totally different (static) scenario.

IV. RANDOM WALK MOBILITY MODEL

We now consider the case in which nodes move (independently of each other) according to a random walk mobility model. In particular, we consider a general class of random walks, in which the position $X(t)$ of a node at time slot t is updated according to the law, $X(t) = X(t-1) + Y_t$, where Y_t is a sequence of i.i.d., rotationally invariant random vectors describing the individual movements (referred to as flights) accomplished by the node. We denote by $f = \|Y_t\|$ the random variable describing the flight length, and by $F = E[f]$ its mean. In our analysis, we will consider for simplicity the case of bounded flight lengths $f \leq f_{\max}$, where $f_{\max} = \Theta(F)$. By letting F vary between the minimum value 1 (the typical distance between neighboring nodes) and \sqrt{N} (the edge of the network area) we vary the mobility degree of the nodes, obtaining a wide class of mobility patterns ranging from the quasi-static case ($F = 1$) to a fully mobile scenario ($F = \sqrt{N}$). For simplicity, we will assume that nodes are not enabled to communicate while moving. In other words, at any slot they can transmit or receive only from the position reached at the end of the flight. However, we emphasize that our analysis could be easily extended to the case in which nodes can communicate while moving, and that this possibility actually improves the network performance, in agreement with recent results [13].

A. Preliminary results

The following lemmas, taken from [11], provide the keys to analyze the case in which nodes move according to a random walk:

Lemma 6: Two nodes can effectively communicate over multi-hop routes if and only if $R = \Omega(F)$.

This is essentially due to the fact that, to effectively advance a message toward a far-away destination, nodes belonging to a multi-hop route should be considered as quasi-static at spatial scale R .

The following result derives from rather sophisticated properties of random walks reported in [11], [22].

Lemma 7: Consider two nodes a and b that independently move in a torus region of area A according to a random walk with flight size F . Assume that the two nodes are uniformly distributed over the region at time $t = 0$. The average first d -hitting time $T_{a,b}(d)$, defined as the infimum of $t > 0$ at which the distance between a and b is less or equal to d , is given by:

$$T_{a,b}(d) = \begin{cases} O\left(\frac{A}{d^2} \log\left(\frac{A}{F^2}\right)\right) \text{ and } \Omega\left(\frac{A}{d^2}\right) & \text{if } d = O(F) \\ O\left(\frac{A}{F^2 \log\left(\frac{d}{F}\right)}\right) & \text{if } d = \omega(F) \end{cases}$$

The above lemma can be exploited in our context to compute the average time $T(X_m, R)$ taken by a node requesting content m to fall within the transmission range R of a node holding a copy of it:

$$T(X_m, R) = \begin{cases} O\left(\frac{N \log N}{X_m R^2}\right) \text{ and } \Omega\left(\frac{N}{X_m R^2}\right) & \text{if } R = O(F) \\ O\left(\frac{N}{X_m F^2 \log\left(\frac{R}{F}\right)}\right) & \text{if } R = \omega(F) \end{cases} \quad (10)$$

Note that in the above equation we have assumed that $R = o(\sqrt{N/X_m})$. Otherwise the node is, with non vanishing probability, already within distance R from a node holding content m , hence for $R = \Omega(\sqrt{N/X_m})$ we have $T(X_m, R) = \Theta(1)$.

Before going on, it is useful to separately examine the case of a “quasi-static” network ($F = \Theta(1)$). The results for this preliminary case will shed light on the impact that mobility has in our system, and will come in handy later on.

B. Analysis of the quasi-static case $F = \Theta(1)$

In the case $F = \Theta(1)$, nodes can communicate with far away destinations using multi-hop routes as long as they employ any transmission range $R = \Omega(1)$, as immediate consequence of Lemma 6. In this section we analyze the performance of a scheme employing multi-hop routes to deliver contents to the nodes (from the closest source), instead of a single-hop communication scheme in which nodes wait until they come in contact with a node caching a copy of the requested content.

Considerations analogous to those in section III suggest that the optimal operating point for the network is when the average number of tx-rx pairs in a square of area R^2 is constant. It follows that $\lambda = \frac{1}{D+I} = \frac{1}{R^2 D}$ where D is the average number of hops.

The average distance between a node requesting content m and the closest node holding a copy of it is $\sqrt{N/X_m}$. It

follows that the replication strategy that minimizes the delay is the solution to the following optimization problem:

$$\begin{cases} \min_{\{X_m\}, m=1 \dots M} \sum_{m=1}^M p_m \max\left(1, \frac{\sqrt{N}}{\sqrt{X_m} R}\right) \\ \text{s.t. } \sum_{m=1}^M X_m \leq KN \\ 1 \leq X_m \leq N \quad m = 1 \dots M \end{cases} \quad (11)$$

Similarly to the optimizations problems considered before, we can restrict ourselves to solving the following reduced optimization problem, for contents of index $m > m^*$, where m^* is the maximum content index for which the delay D_m attains its minimum value of 1.

$$\begin{cases} \min_{\{X_m\}, m > m^*} \sum_{m > m^*} \frac{p_m \sqrt{N}}{\sqrt{X_m} R} \\ \text{s.t. } \sum_{m > m^*} X_m \leq K^* N \\ 1 \leq X_m \leq N \quad m > m^* \end{cases} \quad (12)$$

After relaxing the condition $1 \leq X_m \leq N$ (which is verified by the solution), and applying the standard method of Lagrange multiplier, we obtain that X_m must be proportional to $p_m^{2/3}$ though a constant $C(N, M)$ possibly dependent on N and M . We obtain the following results:

Case $\alpha > 3/2$: In this case $C(N, M) = \Theta(N)$. The resulting delay is $\bar{D} = \Theta(1 + \frac{1}{R})$. By setting $R = 1$, we get the best possible performance $\bar{D} = \Theta(1)$, $\lambda = \Theta(1)$.

Case $1 < \alpha < 3/2$: Now $C(N, M) = \Theta(NM^{2\alpha/3})$. The resulting delay is $\bar{D} = \Theta(M^{3/2-\alpha}/R)$ and the general trade-off is $\bar{D} = \Theta(\lambda M^{3-2\alpha})$. The smallest possible delay $\bar{D} = \Theta(1)$ requires to reduce the throughput to $\lambda = \Theta(M^{2\alpha-3})$ by selecting $R = M^{3/2-\alpha}$, $\bar{I} = M^{3-2\alpha}$. The largest throughput $\lambda = \Theta(M^{\alpha-3/2})$ can be achieved with $R = 1$ (and $\bar{I} = 0$), and incurs a delay $\bar{D} = \Theta(M^{3/2-\alpha})$.

Case $\alpha < 1$. Here $C(N, M) = \Theta(NM^{1/3})$. The resulting delay is $\bar{D} = \Theta(\frac{\sqrt{M}}{R})$, and the trade-off is $\bar{D} = \Theta(\lambda M)$. The smallest possible delay $\bar{D} = 1$ requires to reduce the throughput to $\lambda = \Theta(1/M)$ by selecting $R = \Theta(\sqrt{M})$, $\bar{I} = \Theta(M)$. The largest possible throughput $\lambda = \Theta(1/\sqrt{M})$ can be achieved with $R = 1$ (and $\bar{I} = 0$), and incurs a delay $\bar{D} = \Theta(\sqrt{M})$.

Remarks. We can make the following fundamental observations. First, the above trade-offs for the quasi-static case (with common transmission range) are better than those achievable under the reshuffling mobility model with common transmission range (see Theorem 2). As expected, they are the same as those derived in [2] for the case in which nodes are statically placed on a regular grid. Second, they are, incidentally, exactly the same as those achievable under the reshuffling mobility model by adapting the transmission range to the content (see Theorem 4). In other words, by applying power control under the reshuffling mobility model, we achieve the same performance as that of a quasi-static network in which multi-hop transmissions are exclusively employed.

The above results already suggest the main findings of our work: the best performance is achieved under static (or quasi-static) conditions. Mobility negatively affects the achievable throughput-delay trade-offs, and the worst case is actually the reshuffling mobility model. However, even in this worst case

we can recover the optimal results of a quasi-static network by power control. In the next sections we will confirm that this intuition is correct.

C. Random walk mobility with fixed transmission range

We now consider the general case in which nodes move according to a random walk mobility model with mean flight size $F = \Omega(1)$, and employ a fixed transmission range R . Notice that F should be intended as an exogenous parameter, while R can be chosen to achieve a desired throughput-delay trade-off.

Theorem 5: The throughput-delay performance achievable under a random walk mobility model with average flight size F depends on the Zipf's exponent α :

- For $\alpha > 2$, it is possible to achieve optimal performance $\bar{D} = \Theta(1)$ and $\lambda = \Theta(1)$.
- For $3/2 < \alpha < 2$, we can achieve $\bar{D} = \Theta(1)$ jointly with throughput $\lambda = \Theta(1/F^2)$. If $F = \omega(M^{1/2-\alpha/4})$, we can achieve higher throughputs according to the trade-off $\bar{D} = \Theta(\lambda M^{2-\alpha} \log N)$, up to a maximum throughput of order $M^{\alpha/2-1}/\log N$.
- For $1 < \alpha < 3/2$, we can achieve trade-offs $\lambda = \Theta(M^{2\alpha-3}\bar{D})$, for $\lambda = \Omega(M^{2\alpha-3})$ and $\lambda = O(M^{\alpha-3/2}/F)$. If $F = \omega(M^{\alpha/2-1/2})$, we can achieve higher throughputs according to the trade-off $\bar{D} = \Theta(\lambda M^{2-\alpha} \log N)$, up to a maximum throughput of order $M^{\alpha/2-1}/\log N$.
- For $\alpha < 1$, we can achieve trade-offs $\lambda = \Theta(\bar{D}/M)$, with $\lambda = \Omega(1/M)$ and $\lambda = O(\sqrt{N}/(F\sqrt{M}))$. Higher throughputs can be achieved according to the law $\lambda = \Theta(\bar{D}/(M \log N))$, up to a maximum throughput of order $1/(\sqrt{M} \log N)$.

Proof: We provide only the main ideas behind the proof: details are based on the same steps adopted in the proof of Theorem 2 and in the derivation reported in Section IV-B. The trade-offs achievable for a given value of flight size F are essentially a combinations of the trade-offs achievable in a static network with those achievable under the reshuffling model. The main observations that allows us to identify the optimal communication strategy to be adopted in the network are the following:

- if we end up using a transmission range $R = \Omega(F)$, it is always more convenient to directly transfer the contents employing multi-hop communications, instead of waiting until nodes come in contact with the sources. Indeed, from the analysis of the static case we have found that multi-hopping provides largely better throughput-delay trade-offs even when the average hitting time for content m is $\Theta(N/(X_m R^2))$ (notice that when $R = \omega(F)$ the actual hitting time is larger than $N/(X_m R^2)$, see (10), reinforcing our claim);
- the only reason to use $R = o(F)$ would be to obtain a higher throughput than the maximum one achievable by multi-hopping with $R = \Omega(F)$. This, actually, is not always possible, but depends on F : only when F is larger than a given value (which is a function of M and N), it is possible to achieve a higher throughput by adopting a

single-hop scheme according to which nodes wait until they come in contact with a node holding a copy of the requested content. In this case the delay would be $D_m = T(X_m, R) = O(N \log N / (X_m R^2))$ (see (10)), which is essentially the same expression encountered under the reshuffling model increased by a factor $\log N$. Hence the optimization for $R = o(F)$ leads to the same trade-offs reported in Theorem 2, with the only difference that delays are increased by a factor $\log N$.

Notice that when the single-hop scheme indeed allows to achieve higher throughput than that achievable by multi-hopping (i.e., for F large enough), we get a discontinuity in the delay, as consequence of the fact that we switch from a multi-hop to a single-hop scheme. ■

D. Random walk mobility with different transmission ranges

At last, we consider the case in which the transmission range can be adapted to the content being transmitted, under a random walk mobility model with mean flight size F . This case turns out to be simple to analyze. Indeed, we have already seen that, by adapting the transmission range, one can essentially recover the trade-offs achievable in static or quasi-static conditions even under the extreme case of the reshuffling mobility model (Section III-B). Hence we can expect that the same is possible for intermediate degrees of mobility. This is actually the case, as stated in the following

Theorem 6: By adapting the transmission range to the content, it is possible to obtain the throughput-delay performance achievable under quasi-static conditions (i.e., $F = 1$).

Proof: One simple way to prove this result is the following. Given a desired (feasible under static conditions) throughput-delay trade-off, we compute the optimal transmission ranges R_m that should be adopted under the reshuffling mobility model to achieve the desired trade-offs, and the (fixed) transmission range \hat{R} that should be adopted under static conditions to achieve the same trade-off. Then we partition the contents in two subsets: the first subset includes all contents whose adapted transmission range $R_m = o(F)$, while the second subset includes contents for which $R_m = \Omega(F)$. If we schedule transmission belonging to the first subset in odd slots (applying for them the single-hop scheme), and transmission belonging to the second subset in even slots (applying for them the multi-hop scheme), we get at least half of the target throughput and at most twice the target delay, which is enough to establish our result in order sense. ■

V. EXTENSION TO VARIABLE-SIZE CONTENTS

As already anticipated in Sec. II-C, our results easily extend to the case of heterogeneous content sizes, as long as i) content sizes do not scale with N ; ii) the ratio between the largest and the smallest content size is bounded by a constant.

The natural way to tackle this case is to segment the contents into constant-size chunks, which are then independently transferred from sources to destinations using the previous schemes developed for constant content size. We can actually further generalize our analysis, accounting also for the probability that contacts among the nodes might not always provide a

reliable communication channel to successfully transfer even one chunk of a content.

The only assumption that we need is that, when two nodes fall within transmission range of each other, during the contact duration at least one chunk can be successfully transferred between them with non-vanishing probability $p > 0$. A sufficiently large target probability $p > 0$ can be achieved by properly setting the chunk size and/or the modulation/coding scheme.

Given assumptions i) and ii) above, any content is split into a number of chunks that is upper-bounded by a constant. Chunks are treated as independent objects that are replicated in the network and transferred between nodes according to the same schemes introduced before: when a contact occurs between, say, node a , which is requesting content m , and node b , which is storing one chunk of m not already collected by a (usual bitmap techniques can be employed for this), the chunk is transmitted from b to a . Chunks which fail to be transmitted correctly are simply discarded and retransmitted again in the future⁸.

Previous scaling order analysis can be directly applied also to this more general case, thanks to the fact that contents are divided into a bounded number of independent chunks. While the average content throughput can be easily expressed in terms of the average chunk throughput and the successful chunk transmission probability p , some care is needed when we consider the delay. Indeed, we must consider the fact that a content is fully transferred in the network when all chunks of it are successfully received by the requesting node. The delay analysis of a specific chunk can be carried out repeating similar arguments as before. Note that the fact that a chunk is successfully transmitted only with probability $p > 0$ does not affect results in order sense, since the considered chunk is successfully retrieved after a geometrically distributed number of contacts. The total content transfer delay is equal, by construction, to the maximum delay incurred by the constituting chunks. As long as the number of chunks is bounded, we can easily conclude that the average content transfer delay scales in the same way as the average chunk transfer delay, in light of the following general property: given $\{Z_i\}_{i \in I}$ i.i.d. non-negative random variables (belonging to a set I of cardinality $|I|$), and given $Y = \max_i(Z_i)$, we have $\mathbb{E}[Y] \leq \mathbb{E}[\sum_i Z_i] = |I| \cdot \mathbb{E}[Z_1]$.

VI. CONCLUSIONS

We have established, for the first time to the best of our knowledge, asymptotic delay-throughput trade offs for a mobile ad-hoc network operating in a content-centric scenario under the same assumptions adopted in previous work in the case of a static grid topology. Our results show that mobility tends to worsen the system performance, as the best throughput-delay trade-offs are achieved in a quasi static case. The adoption of smart power control techniques permits to fully recover the optimal performance also in scenarios characterized by a high degree of mobility. In all considered

cases, the size of the content catalog, and the content popularity profile, both have a dramatic impact on the system performance.

REFERENCES

- [1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, R. L. Braynard, "Networking named content", *CoNEXT '09*.
- [2] S. Gitis, G. S. Paschos, L. Tassiulas, "Asymptotic Laws for Joint Content Replication and Delivery in Wireless Networks," *IEEE Trans. on Inf. Theory*, 59(5), pp. 2760–2776, 2013.
- [3] P. Gupta, P.R. Kumar, "The capacity of wireless networks", *IEEE Trans. on Information Theory*, 46(2), pp. 388–404, 2000.
- [4] M. Grossglauser, D.N.C. Tse, "Mobility increases the capacity of ad hoc wireless networks", *IEEE/ACM Trans. on Netw.*, 10(2), August 2002.
- [5] X.-Y. Li, "Multicast capacity of wireless ad hoc networks", *IEEE/ACM Trans. Netw.*, 17(3), pp. 950–961, June 2009.
- [6] U. Niesen, P. Gupta, D. Shah, "The balanced unicast and multicast capacity regions of large wireless networks", *IEEE Trans. on Inf. Theory*, 56(5), pp. 2249–2271, 2010
- [7] U. Niesen, D. Shah, and G. Wornell, "Caching in wireless networks", *ISIT 2009*
- [8] M. J. Neely, E. Modiano "Capacity and delay trade offs for ad-hoc mobile networks" *IEEE Trans. on Inf. Theory*, 51(6), pp. 1917–1937, 2005.
- [9] S. Toumpis, A. J. Goldsmith, "Large Wireless Networks under Fading, Mobility, and Delay Constraints," *IEEE INFOCOM '04*.
- [10] X. Lin, N.B. Shroff, "The Fundamental Capacity-Delay Tradeoff in Large Mobile Ad Hoc Networks," *MedHoc '04*.
- [11] A. El Gamal, J. Mammen, B. Prabhakar and D. Shah, "Throughput and Delay in wireless networks – Part I: Fluid case," Special joint issue of *IEEE Transaction on Information Theory* and *IEEE/ACM Transaction on Networking*, 52(6), pp. 2568–2592, 2006.
- [12] X. Lin, G. Sharma, R. Mazumdar, N. B. Shroff, "Degenerate delay-capacity tradeoffs in ad-hoc networks with Brownian mobility", *IEEE/ACM Trans. Netw.*, vol. 14, Issue SI, pp. 2777–2784, June 2006.
- [13] Y. Kim, K. Lee, N. B. Shroff, and I. Rhee, "Revisiting delay- capacity tradeoffs for mobile networks: The delay is overestimated," *IEEE INFOCOM 2012 mini-conference*, Orlando, FL.
- [14] J.D. Herdtnier, E.K.P. Chong, "Throughput-storage tradeoff in ad hoc networks", *INFOCOM '05*.
- [15] K. Doppler, M. Rinne, C. Wijting, C. Ribeiro, K. Hugl, "Device-to-device communication as an underlay to LTE-advanced networks," *IEEE Communications Magazine*, 47(12), pp. 42–49, Dec. 2009.
- [16] Y. Li, G. Su, D. Jin, L. Su, L. Zeng, P. Hui, "Multiple Mobile Data Offloading Through Delay Tolerant Networks," *ACM Mobicom, CHANTS Workshop*, 2011.
- [17] X. Bao, Y. Lin, U. Lee, I. Rimac, R.-R. Choudhury, "DataSpotting: Exploiting Naturally Clustered Mobile Devices to Offload Cellular Traffic," *IEEE INFOCOM '13*.
- [18] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: evidence and implications," *IEEE INFOCOM 99*.
- [19] C. Fricker, P. Robert, J. Roberts, and N. Sbihi, "Impact of traffic mix on caching performance in a content-centric network," *IEEE NOMEN Workshop 2012*.
- [20] F. Xue and P. Kumar, "Scaling Laws for Ad Hoc Wireless Networks: An Information Theoretic Approach," *Foundations and Trends in Networking*, Now Publishers, 2006.
- [21] D. Ciullo, V. Martina, M. Garetto, E. Leonardi, "Impact of Correlated Mobility on Delay-Throughput Performance in Mobile Ad Hoc Networks," in *IEEE/ACM Trans. Netw.* 19(6), pp. 1745–1758, 2011.
- [22] H. Cai, D.Y. Eun, "Toward stochastic anatomy of inter-meeting time distribution under general mobility models", *MobiHoc '08*.

APPENDIX A

EQUIVALENCE BETWEEN PHYSICAL AND PROTOCOL MODEL

It is rather immediate to see that, if a given set of tx-rx pairs satisfy the physical model (i.e., they can be activated simultaneously producing a SINR at each receiver greater than a desired threshold σ), the same set is feasible also under

⁸Of course, more than one chunk transmission can be attempted during the same contact, if the contact duration is large enough.

a protocol model with a suitable choice of guard factor Δ . Note, indeed, that physical model constraints are in general more stringent than those imposed by the protocol model, as direct consequence of the fact that while the physical model explicitly accounts for the total aggregate interference at the receiver (sum of all interfering contributions), the protocol model essentially approximates the cumulative interference with the strongest interfering signal alone. The challenging part of the order-sense equivalence proof between protocol and physical model is to show the opposite, i.e., that a configuration of tx-rx pairs satisfying the protocol model, also meets the SINR constraint imposed by the physical model, for a sufficiently large guard factor Δ .

We start establishing the following minimum distance property between simultaneous transmitters under the protocol model:

Property 1: According to the considered protocol model, given two nodes s_1 and s_2 that are simultaneously transmitting to node d_1 and d_2 , respectively, employing transmission range R_1 and R_2 , respectively, then the distance between the two transmitters satisfies:

$$d_{s_1, s_2} \geq \Delta \max(R_1, R_2) \quad (13)$$

In addition, whenever $\Delta R_2 > R_1$:

$$d_{s_2, d_1} \geq \Delta R_2 - R_1 \quad (14)$$

Proof: The proof descends immediately from the triangular inequality, i.e.;

$$d_{s_1, s_2} \geq \begin{cases} d_{s_2, d_1} - d_{s_1, d_1} & \geq (1 + \Delta)R_1 - R_1 \\ d_{s_1, d_2} - d_{s_2, d_2} & \geq (1 + \Delta)R_2 - R_2 \end{cases}$$

Similarly:

$$d_{s_2, d_1} \geq d_{s_1, d_2} - d_{s_1, d_1} - d_{s_2, d_2} \geq (1 + \Delta)R_2 - R_2 - R_1$$

From previous property, it immediately descends that:

Property 2: If we draw a ball of radius $\frac{\Delta}{2}R_i$ around every transmitter i (employing transmission range R_i), the balls do not intersect.

Proof: To have an intersection between any two balls of radii R_1 and R_2 , the distance between their centers (i.e., the distance between s_1 and s_2) must satisfy: $d_{s_1, s_2} < \frac{\Delta}{2}R_1 + \frac{\Delta}{2}R_2 < \Delta \max(R_1, R_2)$, but this contradicts (13). ■

Let us now consider a tagged receiver, whose transmitter is employing transmission range R_0 . We have the following property:

Property 3: The number $N(R_i, r)$ of interfering transmitters employing transmission range R_i and lying within distance r from the tagged receiver satisfies:

$$N(R_i, r) \leq \begin{cases} 0 & r \leq \max[(1 + \Delta)R_0, \Delta R_i - R_0] \\ 4 \frac{\beta_i \pi r^2}{\pi \frac{\Delta^2}{4} R_i^2} & r > \max[(1 + \Delta)R_0, \Delta R_i - R_0] \end{cases} \quad (15)$$

for some set of constants $\beta_i > 0$ such that $\sum_i \beta_i = 1$.

Proof: The proof descends from the fact that i) no transmitter employing transmission range R_i can be found within distance $\max[(1 + \Delta)R_0, \Delta R_i - R_0]$ from the tagged

receiver, as consequence of the definition of protocol model and of (14); ii) all potential transmitters residing in the disc of radius r centered at the tagged receiver satisfy property 2. Thus, if we neglect for the moment border effects, we can easily bound $N(R_i, r)$ by computing an upper bound to the number of balls of radius $\frac{\Delta}{2}R_i$ that can be packed into a disc of radius r . If we denote by β_i the fraction of the disc of radius r covered by balls of radius R_i we can write:

$$N(R_i, r) \leq \frac{\beta_i \pi r^2}{\pi \frac{(\frac{\Delta}{2} R_i)^2}{4}}$$

In the above expression, factor 4 takes into account possible border effects, i.e., the fact that not necessarily an entire ball of radius $\frac{\Delta}{2}R_i$ must lie within the considered disc of radius r , given that the ball center falls in it. Elementary geometrical considerations, however, allow us to say that at least 1/4 of any considered ball must lie within the disc of radius r , since by construction i) the ball center is located inside the disc; ii) radius R_i of the ball satisfies $\Delta R_i < r + R_0$, as immediate consequence of property 1 and i).⁹ ■

The total interference I_0 at the tagged receiver can be computed as:

$$\begin{aligned} I_0 &= \sum_i \left[\int_0^1 P_i dN(R_i, r) + \int_1^\infty P_i r^{-\gamma} dN(R_i, r) \right] = \\ &= \sum_i P_i N(R_i, 1) + \sum_i \gamma \int_1^\infty P_i r^{-(\gamma+1)} N(R_i, r) dr \end{aligned} \quad (16)$$

where the second sum that appears in the last expression is obtained integrating by parts. Now, recall that, without loss of generality, all transmission ranges can be assumed to be larger than 1. As consequence, using property 3, the first sum that appears in the last term of (16) is identically equal to zero.

For the other summation, using again property 3, we obtain,

$$\begin{aligned} I_0 &\leq \sum_i 4\gamma \int_{\max[(1+\Delta)R_0, \Delta R_i - R_0]}^\infty P_i r^{-(\gamma+1)} \frac{\beta_i \pi r^2}{\pi \frac{\Delta^2}{4} R_i^2} dr = \\ &= \frac{16\gamma}{(2-\gamma)\Delta^2} P \sum_i \beta_i R_i^{\gamma-2} (\max[(1+\Delta)R_0, \Delta R_i - R_0])^{2-\gamma} \end{aligned}$$

Now, by making Δ sufficiently large we can achieve a full control of the interference (indeed $I_0 \rightarrow 0$ when $\Delta \rightarrow \infty$), thus we can always set Δ in the protocol model so as to meet the desired constraint on the SINR imposed by the physical model.

APPENDIX B PROOF OF THEOREM 1

First we observe that, if a value \hat{R} satisfying the requested condition on $\rho(R^2)$ indeed exists (this will be proven later), we could achieve the network capacity $\Lambda = \Theta(N/\hat{R}^2)$ applying the standard scheme recalled in Lemma 3, according to which the network is partitioned into squarelets of area \hat{R}^2 , each guaranteed to be active with non-vanishing probability. The transmission range \hat{R} , and the associated scheme, turn out

⁹Without lack of generality we assume the arbitrary guard factor $\Delta \geq 1$.

to be optimal both in terms of throughput and in terms of delay. In terms of throughput, it is easy to see that we cannot achieve any higher throughput in order sense, by either increasing or decreasing the transmission range: if we increase R , the maximum network capacity Λ and the corresponding per-node throughput λ , would decrease according to Lemma 3; the network capacity could be in principle increased by augmenting the spatial reuse, i.e., by reducing R , according to the formula $\Lambda = N/R^2$, but values $R = o(\hat{R})$ would lead to a vanishing number of tx-rx pairs in a square of area R^2 i.e., to a vanishing probability that the square is active, which totally offsets the achievable gain. Indeed, according to Lemma 1, the mean number of tx-rx pairs decreases quadratically with $S = R^2$, hence the average number of simultaneously active squarelets (equal to Λ) decays as R^2 , for $R = O(\hat{R})$.

In terms of delay, the value \hat{R} guarantees that nodes having a pending request for an arbitrary content can obtain it after a delay that equals (in order sense) the time needed to come in contact with a node holding a copy of the requested content: indeed, when this condition occurs, the two nodes form a tx-rx pair which has a constant probability to be immediately enabled to transmit. This because the tx-rx pair falls in a squarelet with constant probability, and the average number of tx-rx pairs in the squarelet is bounded. Hence the average delay \bar{D}_m associated to content m achieves (in order sense) the lower bound $1/p_{\text{contact}}(m)$. We cannot achieve any better delay by either increasing or decreasing the transmission range: if we select $R = o(\hat{R})$, the contact probability can only decrease (and the corresponding delay increases accordingly). It would make sense to increase the transmission range only if $p_{\text{contact}}(m) = o(1)$, which occurs when $p_{\text{contact}}(m) = \Theta(X_m R^2/N)$. However, the gain achievable by increasing the contact probability would be totally offset by the contention arising by the fact that the number of tx-rx pairs in a squarelet increases quadratically with $S = R^2$, according to Lemma 1.

At last, the existence of a value \hat{R} satisfying the requested condition on the number of tx-rx pairs falling in it follows from the fact that $\rho(R^2)$ increases monotonically with R , and in the extreme case of $R = \sqrt{2N}$ coincides with the total number of nodes having a pending content request, which can be reasonably assumed to be larger than 1 (at least, larger than one with non vanishing probability).

APPENDIX C PROOF OF THEOREM 2

Case $\alpha > 2$. Consider the following replication strategy:¹⁰ $X_m = \max(1, \frac{N}{2m^{\alpha/2}})$, $\forall m$, combined with the choice of transmission range $R = 1$.

It can be verified that conditions $\sum_{m=1}^M X_m \leq KN$ (for $K > 2$) and $1 \leq X_m \leq N$ are both satisfied. Moreover, $X_m R^2 \leq N$, $\forall m$. It follows that the average delay is $\bar{D} = \sum_{m=1}^M \frac{H}{m^\alpha} m^{\alpha/2} = \Theta(1)$. In any square of area R^2 we have a bounded mean number of nodes (and thus a bounded

mean number of tx-rx pairs). Moreover, considering bounded idle time \bar{I} , in any square of area R^2 we find with non vanishing probability a node requesting content $m = 1$ jointly with another node holding a copy of content $m = 1$. Hence the network capacity is $\Lambda = \Theta(N)$, and the per-node throughput is $\lambda = \Theta(1)$. Since we cannot have any better performance (in order sense) than $\Theta(1)$ for either throughput or delay, the chosen scheme is enough to establish the results for this case.

Case $1 < \alpha < 2$. We first consider the reduced optimization problem (6), and show that the optimal solution to it, for $1 < \alpha < 2$, satisfies $X_m = \Theta(\frac{NM^{\alpha/2-1}}{m^{\alpha/2}})$, $m > m^*$, where m^* is (for now) an arbitrary index. We solve it by relaxing the condition $1 \leq X_m \leq N$ (which is verified by the found solution) and applying the standard method of Lagrange multiplier. We obtain that the ratio p_m/X_m^2 must be the same for all contents, i.e., the number of replicas X_m should be proportional to $\sqrt{p_m}$, through a constant $C(N, M)$ possibly dependent on N or M . Indeed, by imposing that $\sum_{m>m^*} \frac{C(N, M)}{m^{\alpha/2}}$ equals K^*N , we obtain the proportionality factor $C(N, M) = \Theta(NM^{\alpha/2-1})$. Let now m^* be the index (if any) such that for all $m \leq m^*$ quantity $\min(1, X_m R^2/N)$ saturates to 1. By convention, $m^* = 0$ if $\hat{X}_m R^2 \leq N$, $\forall m$. The average delay of any content $m \leq m^*$ is $\Theta(1)$, hence the average overall delay is given in order sense by

$$\bar{D} = \Theta \left(\sum_{m=1}^{m^*} \frac{H}{m^\alpha} + \sum_{m>m^*} \frac{H}{m^\alpha} \frac{m^{\alpha/2}}{M^{\alpha/2-1} R^2} \right) = \Theta \left(1 + \frac{M^{2-\alpha}}{R^2} \right) = \Theta(\lambda M^{2-\alpha})$$

We obtain a family of delay-throughput trade-offs by varying R . In particular, the minimum possible delay $\Theta(1)$ is attained by choosing $R = \Theta(M^{1-\alpha/2})$. Notice that this quantity is $o(\sqrt{N})$, i.e., we do not need to make the transmission range comparable to the network edge to obtain bounded delay. Indeed, we can obtain an associated throughput $\lambda = \Theta(M^{\alpha-2}) = \omega(1/N)$, choosing $\bar{I} = M^{2-\alpha}$.

On the other extreme, the maximum throughput $\lambda = 1/\bar{D}$, achievable with $\bar{I} = 0$, is obtained by solving for R the identity

$$\frac{M^{2-\alpha}}{R^2} = R^2$$

which provides $R = \Theta(M^{1/2-\alpha/4})$. With this choice, we obtain $\lambda = 1/\bar{D} = M^{\alpha/2-1}$. In addition, we observe that the choice of R determines the value m^* such that all contents with index $m \leq m^*$ can be replicated just $X^* = \lceil N/R^2 \rceil$ times. Indeed, it is sufficient to compute the minimum m' such that $X'_m = \Theta(\frac{NM^{\alpha/2-1}}{m'^{\alpha/2}}) > X^*$, and set $m^* = m' - 1$. At last, we can check a posteriori that the additional constraint $1 \leq X_m \leq N$ is satisfied by our solution.

Case $\alpha < 1$. The analysis of this case is similar to the previous one. The optimal solution to the reduced optimization problem is still $X_m = \Theta(\frac{NM^{\alpha/2-1}}{m^{\alpha/2}})$. Indeed, the number of replicas X_m should again be proportional to $\sqrt{p_m}$, through a constant $C(N, M)$ possibly dependent on N or M . This time we have $C(N, M) = \Theta(N/\sqrt{M})$. Notice that, for $\alpha < 1$,

¹⁰For the moment we disregard the fact that contents should be replicated in an integer number of times. As we will show rounding effects do not affect the asymptotic results.

we need to take into account also the fact that $H = M^{\alpha-1}$. Similarly to before, we can express the average delay as:

$$\bar{D} = \Theta \left(\sum_{m=1}^{m^*} \frac{M^{\alpha-1}}{m^\alpha} + \sum_{m>m^*} \frac{M^{\alpha-1}}{m^\alpha} \frac{m^{\alpha/2}}{M^{\alpha/2-1}R^2} \right) = \Theta \left((M/m^*)^{\alpha-1} + \frac{M}{R^2} \right) = \Theta(1 + \lambda M)$$

In the last passage, we have considered that the delay cannot be lower than $\Theta(1)$, when $\lambda M = o(1)$. We obtain a family of delay-throughput trade-offs by varying R . In particular, the minimum possible delay $\Theta(1)$ is attained by choosing $R = \Theta(\sqrt{M})$. Notice that this quantity is $o(\sqrt{N})$, i.e., we do not need to make the transmission range comparable to the network edge to obtain bounded delay. Indeed, we can still obtain an associated throughput $\lambda = \Theta(1/M) = \omega(1/N)$, with $\bar{I} = M$.

On the other extreme, the maximum throughput $\lambda = 1/\bar{D}$, achievable with $\bar{I} = 0$, is obtained by solving for R the identity

$$\frac{M}{R^2} = R^2$$

which provides $R = \Theta(M^{1/4})$. With this choice, we obtain $\lambda = 1/\bar{D} = \Theta(\frac{1}{\sqrt{M}})$.

To conclude the proof we show that contents' replicas can be distributed among the nodes in such a way that the buffer size needed at each node is bounded by a constant, while respecting the constraint that different replicas of the same content are stored by different nodes. In particular, we describe a simple algorithm that permits distributing the set of all replicas evenly among the nodes, guaranteeing that the number of contents to be stored by each node is upper bounded by $K' = 2K$, which is enough for our purposes.

To satisfy our content replication requirements, the proposed algorithm deterministically distributes $\lceil X_m \rceil$ replicas of each content m to different nodes, being X_m the (real) value resulting from the solution of the optimization problem (5).

First observe that, by construction, $\sum_m (\lceil X_m \rceil - X_m) \leq M$, and since $\sum_m X_m \geq M$ (being $X_m \geq 1$, $\forall m$) we have $\sum_m \lceil X_m \rceil \leq 2 \sum_m X_m \leq 2KN$. To distribute content replicas, our algorithm implements a simple water-filling-like strategy, in which contents are considered in sequence. In particular, the algorithm is decomposed into M steps, each in charge of distributing the replicas of a given content to different nodes. At the m -step, the m -th most popular content (content m) is considered¹¹. The basic idea is that, at the generic step m , a set \mathcal{N}_m of (distinct) nodes (of cardinality $\lceil X_m \rceil$) is selected and a replica of content m is assigned to every node in \mathcal{N}_m . Nodes in \mathcal{N}_m are preferentially chosen among those having the smaller number of replicas assigned to them after the execution of the first $m-1$ steps of the algorithm (possible ties are broken at random). Note that, by construction, our algorithm guarantees that replicas of the same content are distributed to different nodes, and that, after

the execution of step M , all replicas of all contents have been assigned to nodes, and the algorithm can terminate. The correctness of the algorithm is based on the following invariance property:

Lemma 8: Let $K_n(m)$ be the load of node n after step m (i.e., $K_n(m)$ is the number of contents assigned to node n after the execution of the first m steps). We have: $\max_n K_n(m) - \min_n K_n(m) \leq 1$, $\forall m$ (i.e., at any step the load is almost balanced among the nodes).

Proof: The proof is easily obtained by induction. In the first step ($n = 1$), necessarily $\max_n K_n(1) = 1$, whereas we either have $\min_n K_n(1) = 0$ (if $\lceil X_1 \rceil < M$) or $\min_n K_n(1) = 1$ (if $\lceil X_1 \rceil = M$). Thus, in both cases, $\max_n K_n(1) - \min_n K_n(1) \leq 1$.

Now assume that the property holds after executing the generic step m . This means that for some $h \geq 0$, $\min_n K_n(m) = h$, while $\max_n K_n(m) \leq h+1$ (i.e., $\max_n K_n(m) - \min_n K_n(m) \leq 1$). Let $N_{h+1}(m)$ and $N_h(m)$ denote the number of nodes with load $h+1$ and h , respectively, after the first m steps (by construction $N_{h+1}(m) + N_h(m) = N$). At step $m+1$, the algorithm has to assign a copy of content $m+1$ to a set of nodes \mathcal{N}_{m+1} of size $\lceil X_{m+1} \rceil$. It does so by first using the pool of nodes having load h , breaking ties at random (if any). There are three possible cases: i) if $\lceil X_{m+1} \rceil < N_h(m)$, only a subset of nodes with previous load h are assigned a copy of content $m+1$. Hence we still have $\min_n K_n(m+1) = h$, whereas $\max_n K_n(m+1) = h+1$. In particular, $N_{h+1}(m+1) = N_{h+1}(m) + \lceil X_{m+1} \rceil$ and $N_h(m+1) = N_h(m) - \lceil X_{m+1} \rceil$; ii) if $\lceil X_{m+1} \rceil = N_h(m)$, we just use all of the nodes having previous load h . Hence $N_h(m+1) = 0$, and consequently $N_{h+1}(m+1) = N$. This implies that $\max_n K_n(m+1) = \min_n K_n(m+1) = h+1$; iii) if $\lceil X_{m+1} \rceil > N_h(m)$, we have to use, in addition to all nodes having previous load h , also some nodes with previous load $h+1$. Hence, we will have $\max_n K_n(m+1) = h+2$, while $\min_n K_n(m+1) = h+1$. In particular, $N_{h+2}(m+1) = \lceil X_{m+1} \rceil - N_h(m)$ and $N_{h+1}(m+1) = N_h(m) + N - \lceil X_{m+1} \rceil$. We can observe that, in all three cases above, $\max_n K_n(m+1) - \min_n K_n(m+1) \leq 1$, hence the property is preserved at step $m+1$.

When the algorithm terminates, at most $2KN$ replicas (see previous upper bound on the total number of replicas to be distributed) are nearly equally distributed among the nodes, and we can conclude that the load of each node is bounded by $K' = 2K$. This concludes the proof of Theorem 2. ■

APPENDIX D PROOF OF THEOREM 4

Recall that we can restrict ourselves to $\alpha < 2$. We have three cases:

Case $3/2 < \alpha < 2$: By selecting $X_m = \Theta(\frac{N}{m^{2\alpha/3}})$, $S_m = m^{\alpha/3}$, for all $m \geq 1$, one we can achieve the best possible performance $\bar{D} = \Theta(1)$, $\lambda = \Theta(1)$.

Case $1 < \alpha < 3/2$: We relax the condition $1 \leq X_m \leq N$ (and later check that this condition is verified by the obtained solution), and apply the standard method of Lagrange multipliers to the other two constraints related to

¹¹Our algorithm does not require that contents are considered in decreasing order of their popularity (actually, any order is fine). This particular choice just simplifies the algorithm description.

X_m and R_m , respectively. We obtain that the number of replicas X_m should be proportional to $p_m^{2/3}$, through a constant $C(N, M)$ possibly dependent on N or M . By imposing that $\sum_{m > m^*} \frac{C(N, M)}{m^{2\alpha/3}}$ equals K^*N , we obtain $C(N, M) = \Theta(NM^{2\alpha/3-1})$. Moreover, squares sizes S_m should be proportional to $p_m^{-1/3}$, through a constant $B(N, M)$ possibly dependent on N or M . Since transmission ranges increase with m , the value $\bar{S}^* = \Theta(\bar{S}) = \Theta(N/\Lambda)$, and we obtain $B(N, M) = \Theta(NM^{2\alpha/3-1}/\Lambda) = \Theta(M^{2\alpha/3-1}/\lambda)$. The resulting throughput-delay trade-off is $\bar{D} = \Theta(\lambda M^{3-2\alpha})$. The smallest possible delay $\bar{D} = 1$ requires to reduce the throughput to $\lambda = \Theta(M^{2\alpha-3})$ by selecting $\bar{S} = \Theta(M^{3-2\alpha})$, $\bar{I} = \Theta(M^{3-2\alpha})$. The largest possible throughput $\lambda = \Theta(M^{\alpha-3/2})$ can be achieved with $\bar{S} = \Theta(M^{3/2-\alpha})$ (and $\bar{I} = 0$), and incurs a delay $\bar{D} = \Theta(M^{3/2-\alpha})$. At last, for indexes $m < m^*$ we can use transmission ranges $S_m = B(N, M)p_m^{-1/3}$, and set $X_m = N/S_m$.

Case $\alpha < 1$. In this case, the constraints on X_m and S_m lead to $C(N, M) = \Theta(NM^{-1/3})$ and $B(N, M) = \Theta(\bar{S}M^{-1/3})$. The resulting throughput-delay trade-off is $\bar{D} = \Theta(\lambda M)$. The smallest possible delay $\bar{D} = 1$ requires to reduce the throughput to $\lambda = \Theta(1/M)$ by selecting $\bar{S} = \Theta(M)$, $\bar{I} = \Theta(M)$. The largest possible throughput $\lambda = \Theta(1/\sqrt{M})$ can be achieved with $\bar{S} = \Theta(\sqrt{M})$ (and $\bar{I} = 0$), and incurs a delay $\bar{D} = \Theta(\sqrt{M})$. Indexes $m < m^*$ are treated as in case $1 < \alpha < 3/2$. ■

PLACE
PHOTO
HERE

Emilio Leonardi (M'99, SM'09) is an Associate Professor at the Dipartimento di Elettronica of Politecnico di Torino. He received a Dr.Ing degree in Electronics Engineering in 1991 and a Ph.D. in Telecommunications Engineering in 1995 both from Politecnico di Torino. In 1995, he visited the Computer Science Department of UCLA, Los Angeles; in 1999 he joined the High Speed Networks Research Group, at Bell Laboratories/Lucent Technologies, NJ; in 2001, the Electrical Engineering Department of the Stanford University, in 2003, the IP Group at Sprint, Advanced Technologies Laboratories, CA, and finally in 2012 the Network Research Division at NEC Labs Europe (Germany). His research interests are in the field of performance evaluation of wireless networks, P2P systems, packet switching.

PLACE
PHOTO
HERE

Giusi Alfano was born in Naples, Italy, on March 22, 1978. She received the Laurea degree in communication engineering from the University of Naples Federico II, Naples, Italy, in 2001. From 2002 to 2004, she was involved in radar and satellite signal processing studies at the National Research Council and University of Naples. From April 2004 to October 2007, she was a Ph.D. student in information engineering at the University of Sannio, Benevento, Italy. She currently holds a post-doctoral position with Politecnico di Torino, Torino, Italy. Her research work lies mainly in the field of random matrix theory applications to MIMO wireless communications and sensor networks and to the characterization of physical layers of random networks.

PLACE
PHOTO
HERE

Michele Garetto (M'04) received the Dr.Ing. degree in Telecommunication Engineering and the Ph.D. degree in Electronic and Telecommunication Engineering, both from Politecnico di Torino, Italy, in 2000 and 2004, respectively. In 2002, he was a visiting scholar with the Networks Group of the University of Massachusetts, Amherst, and in 2004 he held a postdoctoral position at the ECE department of Rice University, Houston. He is currently assistant professor at the University of Torino, Italy.