

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. - (2013), pp. 1863-1871. ((Intervento presentato al convegno IEEE Infocom 2013 tenutosi a Torino nel April 2013 [10.1109/INFCOM.2013.6566980]).

Availability:

This version is available at: 11583/2505558 since: 2016-01-22T12:15:07Z

Publisher:

IEEE / Institute of Electrical and Electronics Engineers Incorporated:445 Hoes Lane:Piscataway, NJ 08854:

Published

DOI:10.1109/INFCOM.2013.6566980

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

©2013 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, Politecnico di Torino, Torino, Italy

[†] Dipartimento di Informatica, Università di Torino, 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 is varied 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.

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.

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. Another interesting 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].

II. SYSTEM ASSUMPTIONS

A. Network and mobility model

We consider a dense network comprising N nodes moving over a square region \mathcal{O} of area 1 with wrap-around conditions

G. Alfano was supported by MASP Project. M. Garetto was supported by AMALFI Project (Università di Torino/Compagnia di San Paolo). E. Leonardi was supported by EU through the FP7 FET Project STAMINA (GA 2645496).

(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 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 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 [17], [18]. This law implies that, having sorted the contents in decreasing order or 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

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 $g(n) = O(f(n))$.

²Our results extend easily to the case of variable-size contents, provided that the ratio between the largest and the smallest content size is bounded by a constant.

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.

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 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. We further assume that nodes can employ a power control strategy to compensate for the signal attenuation due to the distance. In particular, they can adapt the transmission power to the content being transmitted, assuming that a content m can be transmitted at most to a node located at distance R_m (the maximum transmission range for content m). To compensate for the signal attenuation, a node transmitting content m employs power $P_m = P \cdot R_m^\gamma$, where $\gamma > 2$ is the power loss exponent. By so doing, the useful signal arrives at the receiver with power at least equal to P , where P is a given constant.

It has been largely recognized that the throughput-delay performance achievable under the physical model is in order sense equivalent to that achievable by a simpler geometrical model (called protocol model) according to which the transmission of a content m from node i to node j is feasible if and only if the following conditions hold:

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

In our analysis we will consider the above protocol model, which is simpler to understand.

When a successful transmission occurs, we assume that the total amount of 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's file, as long as each content can be split into a bounded number of segments.

⁴The scaling order of our results do 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.

TABLE I

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

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 analysing the network performance achievable under the reshuffling mobility model. One important point to understand is that, in this case, the network performance cannot be improved by making nodes relay contents for other nodes, i.e., by delivering contents over multi-hop routes. This is a consequence of the fact that we jointly assume that: i) a message transmission occupy a finite 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.

In particular, assumption iii) above implies that we cannot perform a multi-hop route over multiple slots to progressively get closer and closer to the destination, since after each slot the destination moves to a totally different, arbitrary location. 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. At last, assumption ii) implies that we cannot even exploit the two-hop scheme proposed by Grossglauser-Tse to increase the transmission opportunities among the nodes, since 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. From the above discussion, 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 statically 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 $\rho(S)$ scales 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 $R = \sqrt{2S}$. 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 given by 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 interested to content m . Since both m -tx and m -rx scale linearly with the square area 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 $p_{\text{contact}}(m) = \Theta(\min(1, X_m R^2))$, where X_m is the number of replicas of content m .

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

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

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

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 assert follows applying Lemma 2. ■

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

Lemma 3: The aggregate transmission rate Λ (also called network capacity) 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 [19]). Essentially, the network area is divided into squarelets of area $S = R^2/2$. 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 subset. 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 A.

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 $\lambda = \frac{1}{N\hat{R}^2}$. 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 or equal to 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:

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

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 $\hat{R} = \sqrt{1/(N\lambda)}$) 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 the replication 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:

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

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 $\min(1, X_m \hat{R}^2)$ in the objective function tells us that it is useless to replicate any content more than $X^* = \lceil 1/\hat{R}^2 \rceil$ 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 $\Theta(N)$. 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 $m > m^*$:

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

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 $R = \Theta(1/\sqrt{N})$.
- For $1 < \alpha < 2$, the optimal throughput-delay tradeoff is $\bar{D} = \Theta(\lambda M^{2-\alpha})$. The minimum delay $\bar{D} = \Theta(1)$ can be achieved with $R = \Theta(M^{1-\alpha/2}/\sqrt{N})$, and the associated throughput is $\lambda = \Theta(M^{\alpha-2})$. The maximum throughput $\lambda = \Theta(M^{\alpha/2-1})$ can be achieved with $R = \Theta(M^{1/2-\alpha/4}/\sqrt{N})$, and the associated delay is $\bar{D} = O(M^{1-\alpha/2})$.
- For $\alpha < 1$, the optimal throughput-delay tradeoff is $\bar{D} = \Theta(\lambda M)$. The minimum delay $\bar{D} = \Theta(1)$ can be achieved with $R = \Theta(\sqrt{M/N})$, and the associated throughput is $\lambda = \Theta(1/M)$. The maximum throughput $\lambda = \Theta(1/\sqrt{M})$ can be achieved with $R = \Theta(M^{1/4}/\sqrt{N})$ and the associated delay is $\bar{D} = O(\sqrt{M})$.

The proof is reported in Appendix B.

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 $p_{\text{contact}}(m) = \Theta(\min(1, X_m R_m^2))$, 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 $\bar{D}_m = \Omega(\frac{1}{\min(1, X_m R_m^2)})$.

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 with probability p_m , employ transmission range R_m , satisfies $\Lambda = O(A/\bar{S})$, where $\bar{S} = \sum_{m=1}^M p_m R_m^2$. Network capacity $\Lambda = \Theta(A/\bar{S})$ can be attained if and only if the average number of m -feasible tx-rx pairs in a square of area R_m^2 is $\Omega(p_m)$.

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 $R_{\max}/2^i < R_m \leq R_{\max}/2^{i-1}$, where R_{\max} is the maximum transmission range employed in the network. 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 of two between the largest and the smallest transmission range in the class. The idea is to first allocate enough 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. More in detail, we first consider class 1, and partition the network area into squarelets of area $R_{\max}^2/2$. Similarly to the traditional scheme for fixed transmission range, we can partition the squarelets into a finite number of subsets, such that squarelets in each subset can be concurrently active. In each slot, we activate an average number Λ_1 of squarelets of class 1 equal to $\Lambda_1 = \sum_{m \in \mathcal{M}_1} \frac{\Lambda p_m}{2\bar{S}}$. If the average number of m -feasible tx-rx pairs in a square of area R_m^2 is $\Omega(p_m)$, we can surely activate the requested number of Λ_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 we enable Λ_1 class-1 squarelets, we remove them from the network, and divide the remaining network area into squarelets of edge $R_{\max}/2$, moving on to class 2, and so on for all classes. In the end, we obtain an aggregate rate $\Lambda = \frac{A}{2\bar{S}}$. Factor 1/2, which does not affect the achievable throughput in order sense, guarantees that the priority assigned to contents with larger transmission range does not affect the transmission opportunities of contents with smaller transmission range (which is also important for the resulting delay). Indeed notice that any squarelet is covered by exactly one of all larger squarelets belonging to higher-priority classes, and the probability that none of them is activated is, by construction, larger than 1/2. Moreover, the transmission opportunities given to contents belonging to the same class are fairly assigned. ■

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 satisfies $\rho_m(R_m^2) = \Theta(p_m)$, for each m .

Proof: Due to lack of space the proof is reported in [21].

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:

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

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 $\bar{S} = 1/(N\lambda)$) 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 $S_m = R_m^2$:

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

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:

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

where \bar{S}^* is a constant independent of transmission ranges of index $m > m^*$.

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 tradeoff 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 tradeoff is $\bar{D} = \Theta(\lambda M)$, with $\bar{D} = \Omega(1)$ and $\bar{D} = O(\sqrt{M})$.

Proof: Due to lack of space the proof is reported in [21]. Here we just mention that the solution to (9), again performed by applying the standard method of Lagrange multipliers, 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 $F = \mathbb{E}[f]$ 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 $1/\sqrt{N}$ (the typical distance between neighboring nodes) and 1 (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 ($F = 1/\sqrt{N}$) to a fully mobile scenario ($F = 1$).

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].

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

A. Preliminary: the quasi-static case

In a quasi-static network, nodes can communicate with far away destinations using multi-hop routes. Moreover, nodes can employ a minimum transmission range⁶ $R = 1/\sqrt{N}$, which allows to obtain a network capacity $\Lambda = N$. More in general, let us assume that nodes employ a common transmission range $R = \Omega(1/\sqrt{N})$. 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 $\lambda = \frac{1}{\bar{D}+1} = \frac{1}{N\bar{R}^2\bar{D}}$ where \bar{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 $1/\sqrt{X_m}$. It follows that the replication strategy that minimizes the delay is the solution to the following optimization problem:

$$\left\{ \begin{array}{l} \min_{\{X_m\}, m=1 \dots M} \sum_{m=1}^M p_m \max \left(1, \frac{1}{\sqrt{X_m} R} \right) \\ \text{s.t.} \quad \sum_{m=1}^M X_m \leq KN \\ 1 \leq X_m \leq N \quad m = 1 \dots M \end{array} \right. \quad (10)$$

Similarly to the optimizations problems considered before, we can restrict ourselves to solving a 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. 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

⁶It can be shown that for any $F = \Omega(1/\sqrt{N})$ the network is mobile enough to smooth out the effects due to the randomness of the topology. In particular, we do not need the minimum transmission range $R = \sqrt{\log N/N}$ necessary to guarantee connectivity in a static random network.

$p_m^{2/3}$ through 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 $\bar{D} = \Theta(1 + 1/(\sqrt{N}R))$. By setting $R = 1/\sqrt{N}$, 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 $\bar{D} = \Theta(M^{3/2-\alpha}/(R\sqrt{N}))$ 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 $R = M^{3/2-\alpha}/\sqrt{N}$, $\bar{I} = M^{3-2\alpha}$. The largest throughput $\lambda = \Theta(M^{\alpha-3/2})$ can be achieved with $R = 1/\sqrt{N}$ (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 $\bar{D} = \Theta(\sqrt{M}/(\sqrt{N}R))$, 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 $R = \Theta(\sqrt{M}/N)$, $\bar{I} = \Theta(M)$. The largest possible throughput $\lambda = \Theta(1/\sqrt{M})$ can be achieved with $R = 1/\sqrt{N}$ (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.

The above results already suggests one of 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.

B. Fixed transmission range

We now consider nodes moving according to a random walk with mean flight length F , and employing 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. The following lemmas, taken from [11], provide the keys to analyse this case:

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], [20].

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:

$$T(X_m, R) = \begin{cases} O\left(\frac{\log N}{X_m R^2}\right) \text{ and } \Omega\left(\frac{1}{X_m R^2}\right) & \text{if } R = O(F) \\ O\left(\frac{1}{X_m F^2 \log\left(\frac{R}{F}\right)}\right) & \text{if } R = \omega(F) \end{cases} \quad (11)$$

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

The above results allow us to prove the following

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 $\lambda = \Theta(1/NF^2)$. If $F = \omega(M^{1/2-\alpha/4}/\sqrt{N})$, 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 $\lambda = O(M^{\alpha-3/2}/(\sqrt{N}F))$. If $F = \omega(M^{\alpha/2-1/2}/\sqrt{N})$, 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 $\lambda = O(1/(F\sqrt{MN}))$. 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-A. The trade-offs achievable for a given value of flight size F are essentially a combinations of the trade-offs achievable in a quasi-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 quasi-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 $\Theta(1/(X_m R^2))$ (notice that when $R = \omega(F)$ the actual hitting time is larger than $1/(X_m R^2)$, see (11), 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 $D_m = T(X_m, R) = O(\log N/X_m R^2)$ (see (11)), 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. ■

C. 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., $F = 1/\sqrt{N}$).

Proof: One simple way to prove this result is the following. Given a desired (feasible under quasi-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 quasi-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. 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 fully recovering the optimal performance also in scenarios characterized by a high degree of mobility. In all considered cases, the size of the contents' catalogue, and the contents' 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 INFOCOM '12*.
- [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*, Vol. 52, No. 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. Herdtner, 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*, vol. 47, no. 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] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: evidence and implications," *IEEE INFOCOM 99*.
- [18] C. Fricker, P. Robert, J. Roberts, and N. Sbihi, "Impact of traffic mix on caching performance in a content-centric network," *IEEE NOMEN 2012*.
- [19] 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.
- [20] H. Cai, D. Y. Eun, "Toward stochastic anatomy of inter-meeting time distribution under general mobility models", *MobiHoc '08*.
- [21] Companion technical report, available at: <http://www.telematica.polito.it/~emilio/papers/CCNtechrep.pdf>

APPENDIX A
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 $\Lambda = \Theta(1/\hat{R}^2)$ applying the standard scheme recalled in Lemma 3, according to which the network is partitioned into squarelets of area $\hat{R}^2/2$, each guaranteed to be active with non-vanishing probability. The transmission range \hat{R} , 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 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 $\Lambda = 1/R^2$, 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 $p_{\text{contact}}(m) = \Theta(X_m R^2)$. 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 $R = \sqrt{2}$ 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 B
PROOF OF THEOREM 2

Case $\alpha > 2$. Consider the following replication strategy:⁷

⁷For simplicity we disregard the fact that contents should be replicated an integer number of times. Rounding effect do not affect asymptotic results.

$X_m = \max(1, \frac{N}{2m^{\alpha/2}})$, $\forall m$, combined with the choice of transmission range $R = \sqrt{1/N}$. 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, $X_m R^2 \leq 1$, $\forall m$. 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 $\min(1, X_m R^2)$ saturates to 1. By convention, $m^* = 0$ if $X_m R^2 \leq 1$, $\forall m$. The average delay of any content $m \leq m^*$ is $\Theta(1)$, hence the average overall delay is given in order sense by

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

We obtain a family of delay-throughput trade-offs by varying R . In particular, the minimum possible delay $\Theta(1)$ is attained by choosing $R = \Theta(\frac{M^{1-\alpha/2}}{\sqrt{N}})$. Notice that this quantity is $o(1)$, 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 $\frac{M^{2-\alpha}}{NR^2} = NR^2$, which provides $R = \Theta(\frac{M^{1/2-\alpha/4}}{\sqrt{N}})$. 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 $X^* = \lceil 1/R^2 \rceil$ 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. Details are reported in [21]. ■