



Università degli Studi Mediterranea di Reggio Calabria
Archivio Istituzionale dei prodotti della ricerca

Forwarding in opportunistic information-centric networks: An optimal stopping approach

This is the peer reviewed version of the following article:

Original

Forwarding in opportunistic information-centric networks: An optimal stopping approach / Borrego, C., Amadeo, M., Molinaro, A., Mendes, P., Sofia, R.C., Magaia, N., Borrell, J.. - In: IEEE COMMUNICATIONS MAGAZINE. - ISSN 0163-6804. - 58:5(2020), pp. 56-61. [10.1109/MCOM.001.1900774]

Availability:

This version is available at: <https://hdl.handle.net/20.500.12318/76253> since: 2021-01-30T23:28:17Z

Published

DOI: <http://doi.org/10.1109/MCOM.001.1900774>

The final published version is available online at: <https://ieeexplore.ieee.org/document/9112743>

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website

Publisher copyright

This item was downloaded from IRIS Università Mediterranea di Reggio Calabria (<https://iris.unirc.it/>) When citing, please refer to the published version.

(Article begins on next page)

This is the post-print of the following article:

C. Borrego et al., "Forwarding in Opportunistic Information-Centric Networks: An Optimal Stopping Approach," in IEEE Communications Magazine, vol. 58, no. 5, pp. 56-61, May 2020

Article has been published in final form at:

<https://ieeexplore.ieee.org/document/9112743>

DOI: 10.1109/MCOM.001.1900774

© 2020 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 collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Forwarding in Opportunistic Information-Centric Networks: an Optimal Stopping Approach

Carlos Borrego, Marica Amadeo, Antonella Molinaro, Paulo Mendes, Rute C. Sofia, Naercio Magaia, Joan Borrell

Abstract—By natively supporting in-network caching and name-based forwarding, Information-Centric Networking (ICN) brings in features that are relevant to better support data transmission in opportunistic wireless networks. Such environments are highly challenged because of the node mobility and intermittent contacts. Therefore, forwarding packets to the right node at the right moment is critical to data retrieval performances. This paper proposes a novel ICN Forwarding Strategy that leverages notable metrics, such as centrality and reliability, and the optimal stopping theory to statistically select the best next-hop forwarders and the time to perform the forwarding decision. Simulations with realistic mobility traces show that the proposed forwarding strategy outperforms other state-of-the-art solutions by guaranteeing shorter retrieval time and less overhead in terms of packet replicas.

Index Terms—Information-Centric Networking, Optimal Stopping Theory, Opportunistic Networks

I. INTRODUCTION

The interest in opportunistic networks (OppNets) has revamped in recent years thanks to the rise of new research, such as Internet of Things (IoT) applications and swarms of Unmanned Aerial Vehicles, which can be characterized by intermittent connectivity, mobility and lack of a-priori knowledge of the network topology. In OppNets, nodes can communicate with each other even if an end-to-end path connecting them is hard to achieve: routes are built dynamically during the packet forwarding process aiming to move data closer to the destination; if no forwarding opportunity exists, then the data is stored in intermediate nodes waiting for future contacts [1].

A communication architecture that can effectively suit OppNets, by inherently implementing store-carry-and-forward delivery and name-based forwarding, is Information-Centric Networking (ICN) and, in particular, its Named-Data Networking (NDN) implementation [2]. NDN is centered around hierarchical content names, which are directly used at the network layer for data discovery and retrieval

and, unlike IP addresses, are not tied to a specific network location. A client, in the following denoted as *consumer*, sends the so-called Interest packets to ask for a given content by name, and the network nodes collaborate to forward the requests towards potential holder(s) of the content. The original content source, or any node maintaining a cached content copy, in the following denoted as a *provider*, can answer to the requests with the required Data packets.

NDN nodes are characterized by a stateful and adaptive forwarding plane [3] keeping state of pending Interests to guide the Data packets back to the consumer, recording the delivery performances, e.g., round-trip-time, and exploring multiple alternative paths in case of route failures. The so-called *Forwarding Strategy* module in an NDN node is in charge of deciding *if, where and when* Interest packets should be forwarded. Routing protocols are used for managing the configuration of forwarding tables in NDN nodes and for disseminating long-term changes in the network topology [4].

The NDN Forwarding Strategy has a key role in the case of highly dynamic networks like OppNets. There, a mobile node may have multiple candidate Interest forwarders in proximity, and it has to select the best one(s). Other times, it may not have suitable routing configuration in the forwarding table or it may experience route failures and new paths should be explored. So far, the dominant forwarding strategies in the literature have used controlled flooding or epidemic forwarding in the absence of routing information [5] [6]. Other strategies have promoted the use of forwarding metrics like social centrality [7] or the node Global Positioning System (GPS) position [8]. However, such solutions do not address the challenging task of identifying the *best moment* for transmitting an Interest packet to the next hop forwarder(s). Nodes take decisions according to defined forwarding metrics and do not consider if these

values are higher or lower in comparison to what they could find in future contacts. If the forwarding decision is taken too early, ignoring that better forwarders could be found in the future, then the delivery performance is not as good as it could have been. Similarly, if the forwarding decision is taken too late, then good past forwarding opportunities could have been missed and the data retrieval is uselessly delayed.

In this paper, we fill this gap by defining a novel Forwarding Strategy for NDN OppNets that considers two crucial steps. First, it identifies the suitable Interest forwarders in terms of a composite metric based on three attributes: centrality – how well connected a node is in the network; reliability – the likeliness of a node not dropping Interests; and reachability – how well connected a node is to the content providers. Second, it leverages the Optimal Stopping Theory (OST) [9] to choose the moment for applying the forwarding decision.

OST is a statistical solution for the problem of choosing the best moment to make a particular decision aiming to maximize a certain reward or to minimize a certain cost. In our context, the composite forwarding metric constitutes the reward of an *optimal stopping problem*, which aims at reducing the data retrieval latency and limiting the packet overhead by statistically finding the best forwarders, i.e., the ones with the higher forwarding metric.

II. NDN FORWARDING FABRIC

The forwarding plane of a NDN node includes three tables: the Content Store (CS), which caches incoming Data; the Pending Interest Table (PIT), which traces the received Interests that are not satisfied by the Data yet; and the Forwarding Information Base (FIB), which includes the outgoing interfaces for Interest forwarding. FIB entries are filled by the Routing Information Base (RIB), a data structure at the NDN control plane that is populated by routing protocols and applications [4].

After receiving an Interest packet, a node n first looks in the CS for a name matching. If it is found, then the Data is replied back through the same interface the Interest arrived from. Otherwise, node n looks in the PIT and, if a matching is found, it updates the existing entry with the new incoming interface of the Interest and then discards the packet. If no match is found in the PIT, n looks in the

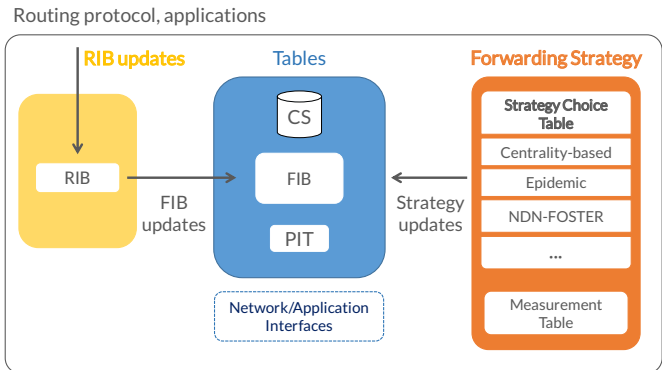


Fig. 1. NDN node architecture.

FIB and, if no matching entry is found, the Interest can be dropped or forwarded based on a specific strategy, such as a (controlled) flooding technique. Forwarding strategies are included in the so-called Strategy Choice Table (SCT), while performance metrics are included in the so-called Measurement Table (MT), see Fig. 1.

Data packets in NDN are not routed, they follow the chain of PIT entries back to the consumer(s) and can be cached by any node located in the reverse path. As a result, content retrieval in NDN takes place into two steps: *content discovery*, which includes the forwarding of the Interest towards a potential provider, and *content delivery*, which includes the Data transmission over the reverse path.

III. FORWARDING IN NDN OPPNETS

Implementing NDN in OppNets requires some changes in the forwarding fabric, in order to manage transmissions over the wireless channel in the presence of mobile nodes and intermittent connectivity. In the following, we present the most representative changes and review the related literature.

Long Lived Interests (LLIs). Interests are maintained in the PIT for a certain lifetime (e.g., about 4s in the vanilla NDN implementation) and discarded if they are not consumed by the Data. However, OppNets are characterized by unpredictable delays in the data retrieval that depend on the contact opportunities. Potentially, Interests could be satisfied minutes, hours or even days after their generation. LLIs have been introduced in NDN in order to deal with situations where requests need a longer lifetime to be satisfied, which is the case of OppNets [4].

Mobility and forwarding decision. In wireless ad hoc and opportunistic networks, consumers, forwarders and producers are mobile and, generally,

there are no contacts with infrastructure nodes. Specific forwarding strategies have been proposed in the NDN-related literature [10] that cope with mobility and intermittent connectivity, but none of them leverages OST. Most approaches can work even without a routing protocol, sometimes considered infeasible in highly dynamic scenarios [8], and they leverage Interest broadcasting to maximize the benefits of in-network caching and limit single-link failures. They range from simple epidemic flooding [6] to more advanced schemes that steer the Interests towards the locations where the data are hosted [5] [8]. Conversely, other approaches leverage unicast Interest transmissions towards next-hop forwarders that are selected according to specific metrics. For instance, in [11] Interests are forwarded to nodes with the higher delivery predictability, which is based on the history of encounters. In [7], instead, nodes maintain a social-tie table that measures the social relationships and is utilized to compute a centrality-based hierarchy of nodes. Then, Interests are forwarded towards the most popular node in the centrality hierarchy. The rationale behind is that the higher the centrality of a node, the higher the chance that this node will find the content.

Transmission scheduling. In ad hoc environments, the NDN Forwarding Strategy may defer the Interest/Data transmissions in order to limit the collision probability over the channel or to prioritize certain packets (e.g., in emergency scenarios). Packet deferral techniques are usually coupled with overhearing and counter-based suppression in order to limit the number of replicas [5] [8]: if the node overhears the same packet transmitted by other nodes for a certain number of times, then it cancels its own transmission.

IV. OPTIMAL STOPPING THEORY

It is very prevalent in computer communications that a node (we call it *decider*) needs to select one or more nodes among different *candidates*. This decision is typically made in terms of a certain metric (e.g., centrality) in the context of a particular network action (e.g., forwarding decision). For these difficult choices OST [9] has proven to be an efficient statistical solution. OST works in two phases. During the first one, the *explore phase*, the decider keeps track of the maximum value of the studied metric for the different observed nodes, but

still, no decision is made. Because of the dynamicity of the considered metric, re-encountered nodes can be considered new nodes after a certain period of time (that is tuned depending on the scenario), otherwise they are excluded from the analysis because already observed. After a fixed number (traditionally denominated *the stopping value*) of candidates has been examined, the *decision phase* starts, where the decider has acquired enough knowledge about the observed nodes in terms of the forwarding metric, and it is ready to make a decision. It will select the first candidate, among the new encountered nodes, whose metric outperforms all the previously observed values. If this node is not encountered, the last suitable candidate will be selected. Optimal stopping values, which influence the final decision, depend on the *maximum number of candidates a decider could eventually observe* that we call *node horizon (NH)*. A notable example where OST is applied is the classical secretary problem (CSP) [12]. There, an administrator must interview a group of NH candidates, which can be ranked from the best to the worse, with the aim of selecting the best one. If the decision could be deferred to the end of all the interviews, the problem would be easily and optimally solved with a simple maximum selection algorithm. The difficulty here is that the decision must be made immediately: once a candidate is not selected, he/she cannot be called again. Therefore, the target is to maximize the probability of selecting the best candidate. In this context, it has been proved that the optimal stopping value is equal to the number NH over the constant e , the Euler's number. The CSP has been generalized in [13] to find a set k of candidates to maximize a certain reward/metric; in this case, k stopping values are to be identified.

In this paper, we introduce OST in the NDN Forwarding Strategy to select the best moment to choose the forwarders that maximize a reward metric. The decision is modelled at each decider with a generalized secretary problem, where the node horizon is computed according to the estimated contact opportunities, as it will be clarified in the following. To the best of our knowledge, this is the first work that introduces OST in NDN OppNets.

V. NDN FORWARDING BASED ON OPTIMAL STOPPING THEORY (NDN-FOSTER)

The proposed NDN-FOSTER is implemented in mobile nodes equipped with an opportunistic wire-

less access interface, e.g., IEEE 802.11 in ad hoc mode. It is based on the following main design principles:

1. *Compliance with the NDN primitives.* Content retrieval is based on the exchange of NDN Interest/Data packets and consists of content discovery and delivery phases. To accommodate potentially long retrieval times, consumer applications start the discovery with an LLI.

2. *To limit the number of Interest forwarders during the content discovery phase.* Similarly to [5] [8], NDN-FOSTER leverages one-hop broadcast Interest transmissions during the discovery phase, to maximize the probability of finding the content in the neighbourhood. When receiving the Interest, a neighbour node can reply with a Data packet, if available in the CS. If a provider is not available in proximity, the request can be further forwarded but, to limit the packet redundancy, each sender has to explicitly choose the next hop forwarder(s).

3. *To statistically select the best Interest forwarders with OST.* NDN-FOSTER leverages a composite metric that allows each Interest sender (i.e., the decider node) to identify the suitability of a neighbour to act as a forwarder. The suitability metric of a node is defined as a linear combination of the following three attributes (equal weights of 1/3 are considered after a tuning simulation campaign).

- *Centrality.* This attribute measures how well connected an NDN node is. The more connected the node is, the more likely it will discover the content [7]. It is computed as the normalized betweenness egocentric centrality, a centrality metric with a small computational cost.

- *Reliability.* This attribute captures how reliable a node is when it acts as a forwarder; it is computed as a function of non-dropped Interests in an observation interval, i.e., Interests that are not consumed by the Data packets within their lifetime are finally dropped from the PIT. Therefore, the higher the non-dropped Interest rate, the more reliable a node is.

- *Reachability.* This attribute, expressed in terms of number of active FIB entries normalized with respect to the number of encountered nodes, catches how likely connected a node is with content producers. The higher the reachability factor of a node, the higher the chance that this node will find a content provider.

Each decider records the recently encountered nodes and the suitability metric in the MT and

applies the OST to select the best forwarder(s).

A. NDN-FOSTER behaviour

A consumer application wishing to retrieve a content issues an LLI. In addition to the *Name* field, it carries a *Lifetime*, indicating the time remaining before the request expires (most probably higher than the default value of 4s), and a newly defined field called *MaxF*, indicating the maximum number of nodes that can forward it. Parameters *Lifetime* and *MaxF* are application specific and prevent the unrestrained propagation of Interests.

When receiving the LLI from the application interface without a matching route towards a provider in the FIB, the NDN Forwarding Strategy at the consumer side starts the discovery phase. In particular, while the LLI stays in the PIT, the strategy autonomously (re)-broadcasts standard Interests that we call DISC-INT (Discovery Interest), carrying the content name and the field *HopLimit* set to 1 to discover a provider in the neighbourhood or, alternatively, the best next-hop forwarders. Candidate forwarders are called *deterministic forwarders* (D-FORW), when they maintain a route (i.e., a FIB entry) for the requested content. Conversely, they are called *probabilistic forwarders* (P-FORW).

DISC-INT broadcasting is performed periodically to exploit all the contact opportunities. If the Data is not retrieved within the LLI lifetime, the LLI is removed from the PIT and, depending on the application, the strategy can start a new discovery phase or finally give up.

When receiving the DISC-INT, a neighbour node n looks in the CS. If a matching is found, n can broadcast the Data. A random deferral transmission technique with overhearing, similar to the one in [5], is implemented to avoid collisions between potential multiple senders. At the Data reception, the consumer creates a new FIB entry that binds the content name with the MAC address of the discovered provider. The cost of the route is set in terms of the measured Round-Trip-Time (RTT). Further Data packets composing the content can be retrieved with standard Interests sent in unicast towards the discovered node.

If the CS matching fails, then n looks in the FIB. If a match is found, it means that a provider has been already discovered. Therefore, node n replies with a DATA-ACK, a newly defined packet that does not carry the content payload, but reports the node

forwarding status, D-FORW, and the route cost, as available in the FIB entry. If no matching is found in the FIB, n replies with a DATA-ACK reporting its P-FORW status, and the value of its forwarding suitability metric.

If DATA-ACK packets are received from deterministic forwarders, the Forwarding Strategy updates the FIB accordingly, it selects as next hop the node with the lowest forwarding cost and issues the LLI transmission in unicast towards it. Otherwise, the strategy can re-broadcast the DISC-INT to collect further suitability metrics from encountered nodes until the optimal stopping value is reached, and it applies the OST-based decision to select the probabilistic forwarders to which the LLI is transmitted.

B. OST-based forwarding

When modelling the forwarding decision with OST, each LLI sender acts as a decider that autonomously applies a secretary problem to select a set of k forwarders, given a certain node horizon NH , with stopping values defined according to [13]. In the following, we explain how the forwarder selection works in practice and analyse it from an optimal stopping perspective.

LLI dissemination. In our scenario, parameter k depends on the value $MaxF$ carried in the LLI, which is decremented every time a new forwarder is selected. More specifically, when $MaxF$ is higher than 1, a decider node n_d can re-broadcast the DISC-INT packets and select new LLI forwarder(s). It halves the $MaxF$ value and transmits the LLI to the first discovered forwarder. Then, n_d repeats the procedure by considering the new value of $MaxF$. When $MaxF$ is lower or equal to 1, the decider can further re-broadcast the DISC-INT but it cannot select a new forwarder. By following this algorithm, the total number of selected forwarders is equal to the logarithm (base 2) of $MaxF$. As a result, the decider can establish multi-path routes towards possible content holders. Each forwarder can apply the same routine to further discover other nodes.

Node horizon estimation. In our scenario, the node horizon is the number of nodes NH a sender can meet before the LLI lifetime expires. This value is generally not known in advance in OppNets, due to the unpredictability of the network topology. In order to estimate it, we consider the *inter-contact time*, that is, the time between two successive node

contacts. We leverage the study in [14], where the authors show that past values of the inter-contact time are correlated with future observed values and compute the inter-contact time in terms of the exponential moving averages of the historical values. Once the value of the inter-contact time is estimated, the node horizon is calculated by dividing the LLI lifetime by the inter-contact time: the longer the lifetime the larger the node horizon, and vice versa.

OST Problem formulation. For a given LLI at a decider n_d , we define as NH the list of n_d 's future contacts before the LLI lifetime expiration. In an unlikely situation where these nodes are met all together, they can be ordered from the best to the worst in terms of their suitability metric. In a more realistic scenario, n_d will successively meet different nodes until it forwards the LLI to a subset k of such nodes that is equal to the logarithm (base 2) of $MaxF$. The objective is to forward the LLI to those k nodes that have the higher suitability metric. This corresponds to an optimal stopping problem where the decision is whether a set of nodes should or should not be chosen as new forwarders. If the first k encountered nodes are chosen, it is quite likely that these nodes will not be the k best nodes in terms of the suitability metric. In the same way, by waiting for the last k contacted nodes, it will be equally likely that the best suitable nodes will be discarded. This situation can be modelled with a generalized secretary problem, and we follow the resolution methodology in [13], where for different values of k , a set of optimal stopping values are provided. These values are such that the p^{th} forwarder (with p ranging from 1 to k) of a certain LLI is the first node which is the best of all the previously seen nodes, if the decider has encountered the first stopping value. If this node is not found, then the second best node of all the previous encountered nodes will be selected, if the decider has met the second stopping value, and so on. Otherwise, if these nodes are not found, then the p^{th} forwarder will be the last p^{th} node before encountering the last node.

A toy example. Fig. 2 shows how the proposed OST decision strategy works with 4 forwarders. The consumer application at node n_1 generates the LLI with parameter $MaxF$ equal to 4. Let us assume that the neighbours of n_1 do not hold a content copy and do not have a FIB entry, hence OST must be applied. The Forwarding Strategy at n_1 realizes that a number k of two forwarders must be selected.




Fig. 2. OST-based forwarding: example with $MaxF$ equal to 4.

When k is equal to two, [13] demonstrated that the two stopping values are equal to NH multiplied by a factor 0.22 and 0.6, respectively. For every new encountered node time, n_1 updates a moving average of inter-contact times (t_i in the figure) that allows to infer the maximum number of nodes n_1 could eventually observe before the LLI lifetime expiration (n_1 's horizon, NH), as in [14]. n_1 collects the DATA-ACKS from encountered nodes and, once the optimal stopping values have been reached (i.e., the 3rd and the 7th encountered nodes), it selects as first and second forwarders, respectively, the first nodes (n_2 and n_3) whose metric outperforms all the previously observed values. For example, at time instant t_6 , n_1 has fed its inter-contact time moving average with values t_1 , the difference between inter-contact times t_2 and t_1 and so on, until the difference between inter-contact times t_6 and t_5 , and it has realised that (i) it has already reached the first stopping value and, (ii) node n_2 outperforms all the previously contacted nodes. In the very same way, at time t_9 , n_1 has already reached the second stopping value and node n_3 outperforms all the previously contacted nodes. Therefore, n_2 receives the LLI with $MaxF$ equal to 2, while n_3 receives the LLI with $MaxF$ equal to 1. Finally, n_2 applies the same procedure and selects n_4 , which will get the LLI with $MaxF$ equal to 1. The rest of the nodes receiving the LLI can only perform the lookup in their CS.

VI. EVALUATION AND DISCUSSION

In this section, we present an experimentation using NDN-ONE-Sim, a customized simulator that integrates the NDN forwarding with the Opportunistic Network Environment (ONE) [15]. We have

enhanced the ONE simulator to include the proposed NDN-FOSTER strategy and, for comparison purposes, four ICN forwarding proposals, i.e., the epidemic forwarding [6], the probabilistic name-based forwarding [11], a copy-limited forwarding similar to the classical Spray and Wait routing protocol [1], and the centrality-based forwarding [7]. These approaches are referred in the following as “*epid*”, “*prob*”, “*sw*” and “*central*”, respectively.

Node contacts are modelled through real mobility traces from the Crawdad database, a community resource for collecting wireless data at Dartmouth College. Two scenarios are considered: the first one, named *Cisco*, is based on traces retrieved from 500 taxi cabs in San Francisco, USA, and includes 449226 contacts. The second scenario, named *Cambridge*, is based on 10641 contacts from 51 students from the System Research Group of the University of Cambridge carrying small devices for six days. For both scenarios, during 24 hours of time, every node in the network is randomly selected as consumer requesting contents in a catalog of 1000 items, with an average frequency of 1 Interest/second and according to a Zipf-like distribution with skewness parameter α equal to 0.8. At the beginning, each content is randomly associated to a node acting as the original provider. The content size varies from 1KB to 20KB. Re-encountered nodes are considered new nodes after 300s.

To assess the effectiveness of NDN-FOSTER, we compute two metrics: (i) the *retrieval latency*, which is the average time for retrieving the content; (ii) the *discovery ratio*, which is the ratio between the number of Interests (LLI or INT-DISC) reaching a provider and the number of all Interests (LLI and INT-DISC) generated in the network. In addition, to assess the efficiency of NDN-FOSTER, we consider (i) the *replication degree*, computed as the average number of total Interest replicas (i.e., LLI and INT-DIST) per each LLI created by the consumers, and (ii) the *dissemination efficiency*, computed as the discovery ratio divided by the replication degree.

In Fig. 3, we compare NDN-FOSTER against *epid*, *prob*, *sw* and *central* approaches in terms of the above mentioned metrics. It can be observed that NDN-FOSTER guarantees the lowest data retrieval latency and a discovery ratio almost equal to the *epid* approach. Moreover, NDN-FOSTER replicates less Interests than the other solutions and also shows higher efficiency. Such advantages are obtained

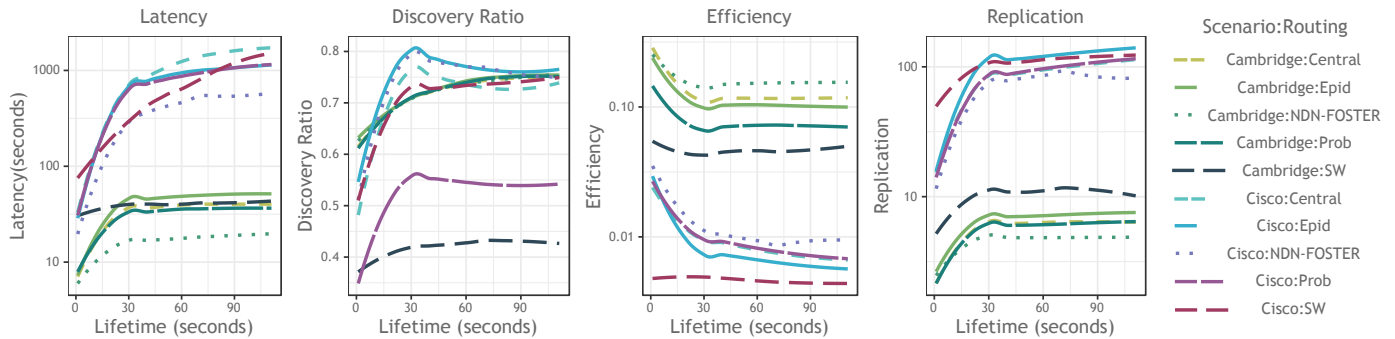


Fig. 3. Retrieval latency, discovery ratio, dissemination efficiency and replication degree as a function of the LLI lifetime.

thanks to the OST-based strategy that couples the selection of the best forwarders with the best time to perform the forwarding action, and therefore it is able to take wiser decisions.

VII. CONCLUSIONS

In this paper, we have presented an OST-based forwarding strategy for NDN opportunistic networks that outperforms state-of-the-art approaches in terms of latency and dissemination efficiency.

ACKNOWLEDGEMENTS

This work was supported by FCT through the LASIGE Research Unit, ref. UIDB/00408/2020 and ref. UIDP/00408/2020.

REFERENCES

- [1] V. F. Mota, F. D. Cunha, D. F. Macedo, J. M. Nogueira, and A. A. Loureiro, "Protocols, mobility models and tools in opportunistic networks: A survey," *Computer Communications*, vol. 48, pp. 5–19, 2014.
- [2] G. Xylomenos *et al.*, "A survey of information-centric networking research," *IEEE Communications Surveys and Tutorials*, vol. 16, no. 2, pp. 1024–1049, 2014.
- [3] C. Yi *et al.*, "A case for stateful forwarding plane," *Elsevier Computer Communications*, vol. 36, no. 7, pp. 779–791, 2013.
- [4] P. Mendes, R. Sofia, V. Tsaoussidis, and C. Borrego, "Information-centric routing for opportunistic wireless networks," *IETF Internet Draft*, Sept 2019.
- [5] M. Amadeo, C. Campolo, and A. Molinaro, "Forwarding strategies in named data wireless ad hoc networks: Design and evaluation," *Elsevier Journal of Network and Computer Applications*, vol. 50, pp. 148–158, 2015.
- [6] G. Tyson, J. Bigham, and E. Bodanese, "Towards an information-centric delay-tolerant network," in *IEEE INFOCOM Workshops*, 2013, pp. 387–392.
- [7] Y. Lu, X. Li, Y.-T. Yu, and M. Gerla, "Information-centric delay-tolerant mobile ad-hoc networks," in *IEEE INFOCOM Workshops*, 2014, pp. 428–433.
- [8] G. Grassi, D. Pesavento, G. Pau, L. Zhang, and S. Fdida, "Navigo: Interest forwarding by geolocations in vehicular named data networking," in *IEEE WoWMoM*, 2015, pp. 1–10.

- [9] A. N. Shiryaev, *Optimal stopping rules*. Springer Science & Business Media, 2007, vol. 8.
- [10] X. Liu, Z. Li, P. Yang, and Y. Dong, "Information-centric mobile ad hoc networks and content routing: a survey," *Ad Hoc Networks*, vol. 58, pp. 255–268, 2017.
- [11] M. W. Kang, Y. Kim, and Y. W. Chung, "An opportunistic forwarding scheme for ICN in disaster situations," in *IEEE ICTC*, 2017, pp. 948–950.
- [12] P. Freeman, "The secretary problem and its extensions: A review," *International Statistical Review/Revue Internationale de Statistique*, pp. 189–206, 1983.
- [13] K. Ano, "Optimal selection problem with three stops," *J. Oper. Res. Soc. Japan*, vol. 32, pp. 491–504, 1989.
- [14] C. Borrego, E. Hernández-Orallo, and N. Magaia, "General and mixed linear regressions to estimate inter-contact times and contact duration in opportunistic networks," *Elsevier Ad Hoc Networks*, vol. 93, 2019.
- [15] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE simulator for DTN protocol evaluation," in *EAI SIMUtools*, 2009.

BIOGRAPHIES

CARLOS BORREGO is an assistant professor at the Department of Mathematics and Computer Science of the University of Barcelona. He teaches and researches on computer networks and cryptography.

MARICA AMADEO is an assistant professor at the University Mediterranea of Reggio Calabria. Her main research interests are in the field of Information-Centric Networking and Internet of Things.

ANTONELLA MOLINARO is an associate professor at the University Mediterranea of Reggio Calabria. Her current research focuses on 5G, vehicular networking and future Internet architectures.

NAERCIO MAGAIA is an invited assistant professor at the University of Lisbon. His current research interests cover the Internet of Vehicles, Network Security, and Artificial Intelligence.

JOAN BORRELL is an associate professor of the Department of Information and Communications Engineering of the Autonomous University of Barcelona.

RUTE C. SOFIA, IEEE Senior Member, is the IIoT Head at fortiss, and an Invited Associate Professor at University Lusófona. Her research interests are IoT, network architectures and protocols.

PAULO MENDES, IEEE Senior Member, is Senior Scientist at Airbus, and Invited Associate Professor at University Lusófona. His major research interest is on information-centric wireless networks.