

Connecting the Edges: A Universal, Mobile-centric and Opportunistic Communications Architecture

Journal:	<i>IEEE Communications Magazine</i>
Manuscript ID	COMMAG-17-00325.R2
Topic or Series:	November 2017/Human-Driven Edge Computing and Communication
Date Submitted by the Author:	18-Sep-2017
Complete List of Authors:	<p>Sarros, Christos; Athena Research and Innovation Center in Information Communication and Knowledge Technologies, ENTA unit; Democritus University of Thrace School of Engineering, Electrical and Computer Engineering</p> <p>Diamantopoulos, Sotiris; Athena Research and Innovation Center in Information Communication and Knowledge Technologies, ENTA unit; Democritus University of Thrace, Electrical and Computer Engineering</p> <p>Reñé , Sergi; University College London, Electronic & Electrical Engineering</p> <p>Psaras, Ioannis; University College London, E&EE</p> <p>Lertsinsrubtavee, Adisorn; University of Cambridge, Computer Laboratory</p> <p>Molina-Jimenez, Carlos; University of Cambridge, Computer Laboratory</p> <p>Mendes, Paulo; COPELABS, University Lusofona, SITILabs</p> <p>Sofia, R; Senception, Lda</p> <p>Sathiaseelan, Arjuna; University of Cambridge, Computer Laboratory</p> <p>Pavlou, George; University College London, Electronic & Electrical Engineering</p> <p>Crowcroft, Jon; University of Cambridge, Computer Lab</p> <p>Tsaoussidis, Vassilis; Athena Research and Innovation Center in Information Communication and Knowledge Technologies, ENTA unit; Democritus University of Thrace, Dept. of Electrical & Computer Engineering</p>
Key Words:	Information-Centric Networking, Named Data Networking, Opportunistic networking, Delay-Tolerant Networking, Future Internet, Internet architecture

Connecting the Edges: A Universal, Mobile-centric and Opportunistic Communications Architecture

Christos-Alexandros Sarros^{1,2}, Sotiris Diamantopoulos^{1,2}, Sergi Rene³, Ioannis Psaras³, Adisorn Lertsinsrubtavee⁴, Carlos Molina-Jimenez⁴, Paulo Mendes⁵, Rute Sofia^{5,6}, Arjuna Sathiaseelan⁴, George Pavlou³, Jon Crowcroft⁴, Vassilis Tsaoussidis^{1,2}

Abstract

The Internet has crossed new frontiers with access to it getting faster and cheaper. Considering that the architectural foundations of today's Internet were laid more than three decades ago, the Internet has done remarkably well until today to cope with the growing demand. However, the future Internet architecture is not only expected to support the ever-growing number of users and devices but also a diverse set of new applications and services. Departing from the traditional host-centric access paradigm where access to a desired content is mapped to its location, an information-centric model enables the association of access to a desired content with the content itself, irrespectively of the location where it is being held. UMOBILE tailors the information-centric communication model to meet the requirements of opportunistic communications, integrating those connectivity approaches into a single architecture. By pushing services near the edge of the network, such an architecture can pervasively operate in any networking environment and allows for the development of innovative applications, providing access to data independently of the level of end-to-end connectivity availability.

INTRODUCTION

As technology moves forward, new devices and applications are increasingly emerging as diverse components of a giant puzzle, all to be supported by the same network architecture known as the Internet. The latter, until now, has coped more than well with the challenges. Indeed, an architecture designed for communication between fixed hosts and remote access of resources, has been accommodating a growing number of mobile devices, as well as applications that are increasingly directed on content dissemination.

This is becoming increasingly inefficient though, as these communication needs become the norm rather than the exception. Mobile users and traffic are now the dominant part of the network, while emerging paradigms such as the Internet of Things are bringing billions of distributed sensors into the picture, questioning the scalability of current solutions. While intermittent connectivity is naturally calling for support in such environments, other visions such as the Tactile Internet require extremely low latency, high availability, reliability – all of which are challenging the current Internet infrastructure.

In response to these ongoing developments, a lot of research has been invested into future Internet technologies. A first class of solutions are compatible with the current Internet architecture and devise suitable mechanisms to place the focus on the network edge, e.g. device-to-device opportunistic communications and edge caching in 5G mobile networks, realizing the fog and edge computing paradigm. The second class advocates alternative Internet architectures that solve the problems in a native way, such as the Information-Centric Networking (ICN) and Delay-Tolerant Networking (DTN) paradigms.

¹ Athena Research and Innovation Center

² Democritus University of Thrace

³ University College London

⁴ University of Cambridge

⁵ COPELABS, University Lusofona

⁶ Senception Lda

1 We build on all the latest advancements and develop an architecture that shifts the current focus of the architectural
2 design. Network periphery is our center of the attention, as the vast majority of traffic comes from the network
3 edges. We place special focus on localized communications and services at the edge, which will improve user QoE
4 and alleviate congestion in the core of the network.

5 Our goal is a tighter integration of opportunistic and delay-tolerant communications with the core of the network,
6 extending its reach towards universal coverage and enhancing operational reliability when the network
7 infrastructure is unavailable or impaired. We aim to further integrate social-awareness into the network by
8 leveraging human users' contextual information and create a truly pervasive network.
9

10 Towards this end, we are developing UMOBILE: a Universal, mobile-centric and opportunistic communications
11 architecture, which we present in the following sections. We first share our motivation and describe our vision.
12 Then, we provide some basic information about the key technologies enabling this vision before we finally unveil the
13 architecture by focusing on its main components and novel contributions.
14
15
16
17

18 MOTIVATION AND VISION

19
20 The miniaturization of computing and communication devices has culminated in smartphones becoming
21 ubiquitous and sensors being deployed seemingly everywhere. In this context, there are areas where technical
22 difficulties in communication are unavoidable; remote areas such as forests and mountains, or even places in the
23 heart of the cities such as the subway. Similar problems exist in crowded areas (e.g. stadiums) and multiply in
24 disaster cases, where the network might suffer from both fragmentation and extreme congestion, rendering the
25 backbone inoperable.
26
27

28 Notably, connectivity problems also arise due to economic constraints. Mobile users may find it costly to use
29 their provider's data services when roaming in a city, while sparse populations living in physically remote
30 locations are plagued by the fact that it is not cost-effective for Internet Service Providers (ISPs) to install the
31 required infrastructure.
32
33

34 One potential solution to these pertinent problems is to change how users communicate and access information,
35 by moving from the traditional host-centric access paradigm (where access to a desired content is mapped to its
36 location) to an information-centric model, where access to a desired content is mapped to the content itself
37 irrespectively of the location where it is being held [1]. This new model can also exploit the benefits of
38 opportunistic and delay-tolerant networking, cancelling the need of end-to-end connectivity and natively
39 supporting content caching to facilitate localized communications (e.g., data synchronization, group messaging,
40 content retrieval) even in disrupted environments. We extend the content-centric, opportunistic and delay-
41 tolerant networking concepts to also include computing, as opposed to static content transfer/delivery, and
42 exploit the huge amounts of computation available in handheld and edge-network devices. This is in stark
43 contrast to past approaches in the area.
44
45

46 Our main objective is to develop a mobile-centric, service-oriented architecture that efficiently delivers content
47 and services to the end-users, especially in challenging areas. UMOBILE decouples services from their origin
48 locations, shifting to a paradigm that incorporates aspects from both information-centric and opportunistic
49 networking, with the ultimate purpose of pushing network services (e.g., intermittent connectivity support) and
50 user services (e.g., pervasive content management) as close as possible to the edge. By pushing such services
51 close to the users, aspects such as bandwidth utilization and resource management can be optimized and service
52 availability in challenged network environments can be improved.
53
54
55
56
57

58 KEY ENABLERS

59 INFORMATION-CENTRIC NETWORKING

60

Information-centric networking is a novel approach to networking, which aims at solving inherent problems of the IP-based infrastructure by redesigning the network layer to focus on “what” is being transferred instead of “where”, effectively making the Internet a native content distribution network. The UMOBILE architecture is based on one of the most promising ICN implementations, Named Data Networking (NDN) [2], taking advantage of its inherent addressing and security mechanisms.

The NDN architecture provides features such as in-network content caching, facilitating content retrieval from the closest available source and providing load balancing without depending on add-ons such as DNS round robin. Hierarchical naming and routing is used, based on content identifiers instead of location identifiers. Consumer mobility, multicast and packet-level data security are integrated into the architecture by design.

In NDN, data consumers initiate communication, through the exchange of two types of packets: Interest and Data. A content request is issued by sending an Interest packet, which is routed through the network to the closest content instance. Routing is performed utilizing longest prefix matching with hierarchical aggregation to ensure scalability, using a Forwarding Information Base (FIB). The content is then sent back on Data packets, which follow the reverse path to the requester using “breadcrumbs” left in a Pending Interest Table (PIT) on each router. Along this path, copies are stored locally in each router’s Content Store to support subsequent requests.

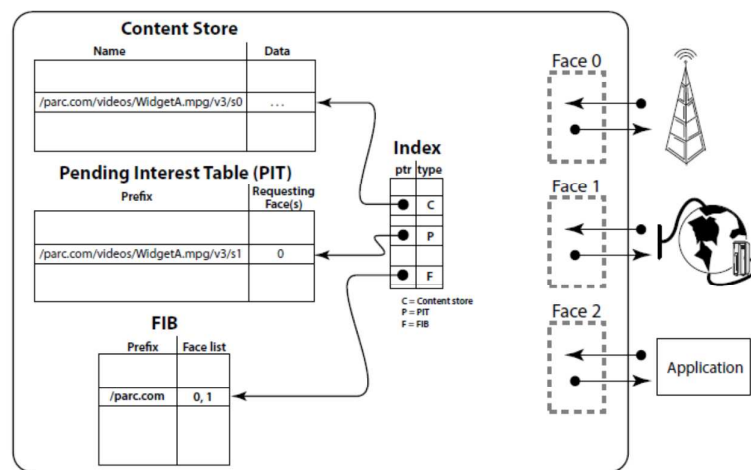


Figure 1: NDN Architecture

NDN provides a set of features that are especially beneficial for mobile and wireless networks, including:

- Host multihoming: NDN provides higher flexibility to use all interfaces available to retrieve data requested at the same time.
- Detachment of applications from location-oriented addresses: NDN allows applications to abstractly publish or consume content, without the need to store (or even know) their own network-layer addresses.
- No connection-oriented sessions: NDN allows the relocation of hosts without the necessity of re-establishing a connection.
- Resilience through replication: NDN allows local copies of the content to be retrieved, mitigating the effects of network failure in mobile nodes, and, thus, improving network performance.

Most of the NDN research has focused so far on core network issues, with some initial studies concerning MANETs and wireless sensor networks. In UMOBILE, we go one step beyond and extend the architecture to effectively exploit any communication opportunity to forward data towards receivers.

OPPORTUNISTIC NETWORKING

1 As UMOBILE primarily targets a wireless environment, opportunistic networking is considered an essential part of
2 our architecture. Recent standards in the context of wireless networking reinforce the relevancy of Device-to-Device
3 (D2D) communication, as complementary to existing access technologies. Although not fully mature yet, Wi-Fi Direct
4 is currently included in the Android OS and supported by most recent Android smartphones to support opportunistic
5 communications.

6 This ever-growing interest in the opportunistic networking paradigm is backed up by different opportunistic routing
7 solutions, ranging from simple epidemic approaches to more elaborate ones relying on social awareness. As far as
8 different opportunistic routing approaches are concerned, social-aware solutions have shown great potential given
9 their more stable approach to data transmission dynamics and lower impact in terms of data replication. Contact
10 duration and frequency can be employed to identify communities [3], as the time that nodes spend together has
11 been used as a good indication of social interaction [4]. Based on that observation, some approaches propose the
12 adaptation of routing calculations to a node's behavior [5]. Furthermore, it has been identified that nodes with
13 similar interests make good next-hops for message delivery [6].

14 A significant weakness of current opportunistic solutions is the fact that their usefulness in realistic environments is
15 not clear, as they have been extensively evaluated only through simulators and emulators. On the other hand, most
16 opportunistic networking solutions rely on a host-centric approach, transmitting packets between nodes regardless
17 of the data being exchanged. Information-centric networks offer a huge potential in disrupted environments that
18 hamper normal communication between hosts.

22 DELAY-TOLERANT NETWORKING

23 To support communications at the edges of the network, where frequent disruptions in communication may take
24 place, we also enhance our platform with Delay/Disruption Tolerant Networking (DTN) functionality. Being
25 originally conceived as an architecture for the Interplanetary Internet, the DTN architecture was soon generalized
26 to include challenged terrestrial networks [7]. It relies on the store-and-forward networking paradigm and is
27 compatible with diverse underlying technologies, with the support of multiple convergence layer protocols and
28 provide interoperability among different platforms even in disruptive environments.

29 By utilizing these features, DTN can bind different internetworks and incorporate devices and applications with
30 limited and/or local functionality, which require a form of internetworking capability. We leverage these
31 advantages and utilize DTN to achieve network operation in challenging environments where connectivity is
32 intermittent, explicitly incorporating concepts such as data muling into our architecture. DTN features such as
33 custody transfer, also offer an increased reliability fabric to our architecture.

38 THE UMOBILE ARCHITECTURE

41 ARCHITECTURAL ADVANCEMENTS

42 In order to cater for the deficiencies of the NDN architecture, we modify and enhance it to enable support for the
43 edge of the network, especially for mobile opportunistic wireless environments. We have identified the relevant
44 requirements [8] and are developing several mechanisms that range from new forwarding strategies tailored to
45 these needs, to an opportunistic routing engine for NDN. To this end, we exploit all available communication
46 opportunities by utilizing an Ethernet interface in the fixed part of the network, and Wi-Fi/Wi-Fi Direct and LTE
47 interfaces in the mobile part. Furthermore, we introduce context awareness into our architecture, allowing the
48 network to adapt to different user behaviors by leveraging contextual information available through the Bluetooth
49 and Wi-Fi/Wi-Fi Direct interfaces. We, finally, expand application capabilities by providing new APIs and design
50 new QoS mechanisms to enable service deployment even on remote, challenging regions with intermittent
51 connectivity.

52 The architectural diagram of the UMOBILE platform is provided in Figure 2. Red modules depict the new
53 components that have been developed in the context of UMOBILE, whereas blue ones depict the original NDN
54

framework. UMOBILE is currently deployed over both fixed Linux-based and mobile Android-based devices. Our main technical advancements are presented below, while the full implementation details can be found in [9].

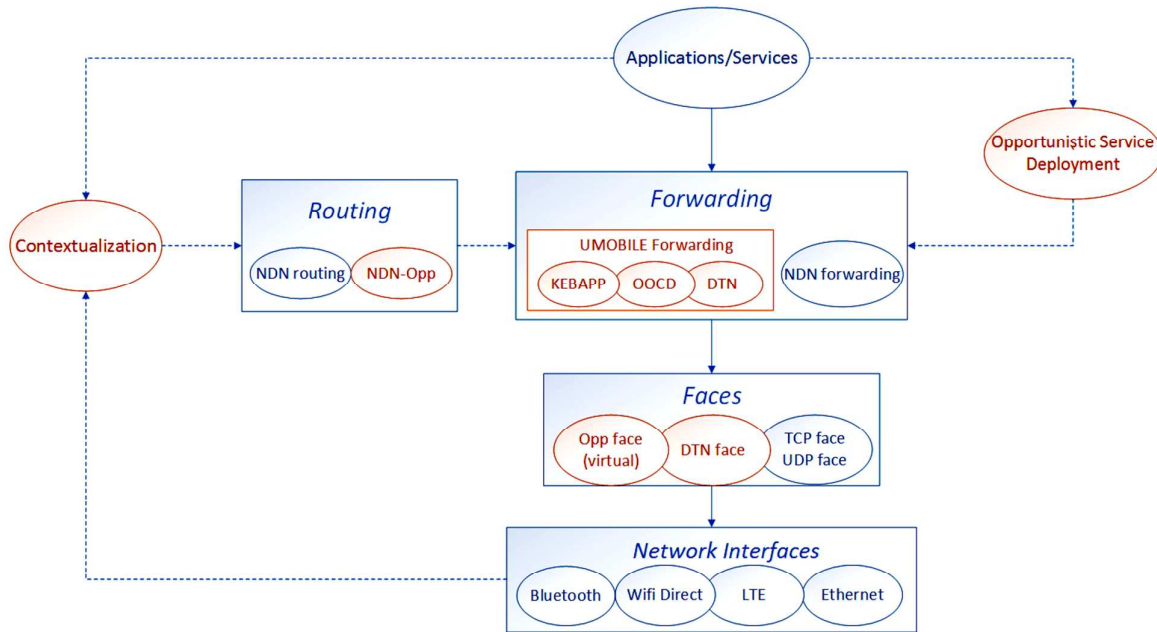


Figure 2. The UMOBILE architecture

FORWARDING

The core of the proposed information-centric communication model lies in the usage of Interests issued by clients (pull-based model) or content providers (push-based model). To be successfully delivered, a decision must be made on whether, when and to whom an Interest should be forwarded. Given the diverse application requirements, along with networking conditions that may span from continuous high-speed connectivity to intermittent disruptive communications, a collection of different forwarding strategies needs to be supported by the UMOBILE architecture.

Forwarding in UMOBILE follows typical NDN forwarding (Best route, Broadcast, Client control, Ncc), enhanced with new forwarding strategies and mechanisms that allow for delay/disruption tolerant and opportunistic communications. A per-namespace selection of the appropriate strategy is available, to fulfill the needs for different applications that utilize different naming schemes. On the following subsections, we present the new forwarding options developed in the context of the project.

DTN TUNNELING

An important forwarding option developed within the framework of the project is DTN tunneling. To achieve this, we utilize the IBR-DTN implementation of the Bundle protocol and enhance NDN by creating a novel DTN face. By tunneling NDN packets through the newly-created DTN face, we extend NDN operation and provide reachability to remote areas where no typical Internet connectivity exists, as well as reliability to services in challenged environments.

Tunneling via DTN-enabled devices, also solves the issue of NDN breadcrumb-based Data forwarding in opportunistic mobile networks. The latter would require the same node that carries an Interest packet, to return the Data packet as well – a requirement that significantly inhibits performance in such networks where nodes are highly volatile. In our approach, forwarding is handled to the DTN layer inside highly-opportunistic “DTN islands” that may exist between NDN nodes. An edge UMOBILE node (e.g. an access point) maintains a single FIB/PIT entry towards the next edge node, while Interest/Data packets are encapsulated and forwarded transparently by the intermediate DTN mobile nodes. This enables data forwarding between two remote, intermittently-connected

UMOBILE nodes by using different upstream and downstream routes/nodes, without modifications to NDN semantics.

Furthermore, DTN can also be used to support various QoS schemes. Along these lines, we are using DTN to enable service deployment in challenging regions. We build on the data muling concept to transfer not only single requests, but entire services (i.e. execution code that can run on hotspots) to remote regions. This way, service availability and latency are improved as the requests do not have to be periodically transferred back to the core of the network and basic computation is offloaded to the hotspots. While the DTN face can be readily used by the existing NDN forwarding strategies, it is in our future plans to design strategies that fully leverage its potential, as well as investigate its use as a supportive mechanism to congestion control and avoidance.




Figure 3. Main DTN functionality in the context of UMOBILE.
Left: DTN tunneling. Right: Data muling.

OPPORTUNISTIC OFF-PATH CONTENT DISCOVERY

Another mechanism that we use to facilitate packet forwarding at the edge, is Opportunistic Off-path Content Discovery (O OCD). The Information- or Content-Centric Networking paradigm was originally proposed as a paradigm shift to transform the Internet from a communications system to a *native content distribution network*. However, due to the scalability and other challenges of routing and forwarding found in ICN architectural proposals, most architectures have focused on scaling content resolution towards the main origin (or CDN surrogate server). That said, all current architectural proposals, including the most prevalent one (i.e., NDN) focus on the optimization of “how to route requests towards the core of the network”. We argue that this goes against the original vision of a “native content distribution network” and prevents requests from discovering nearby content, unless the content is on the shortest path to the core of the network.

With a view to improving the content discovery capabilities of the NDN architecture, we have proposed an enhancement to the routing fabric, which keeps track of successful (i.e., served) content requests in a separate routing table, called “Downstream FIB” (D-FIB). D-FIB effectively acts as a FIB table that points to downstream nodes that have recently received (and therefore, cached) requested content. D-FIB uses trails left behind by data packets from the content origin to the sources in order to discover off-path cached content. Our results [10] indicate that the opportunistic discovery mechanism can significantly increase cache hit rate compared to NDN’s default forwarding strategy, while limiting the overhead at acceptable levels. This enables a native content discovery and distribution system, enables a native P2P network, but also serves well the purposes of communications in case of network fragmentation. That is, when the path to the origin server is broken (due to link or node failures), the D-FIB approach discovers content off-path, possibly stored/cached in end-user devices.

ROUTING

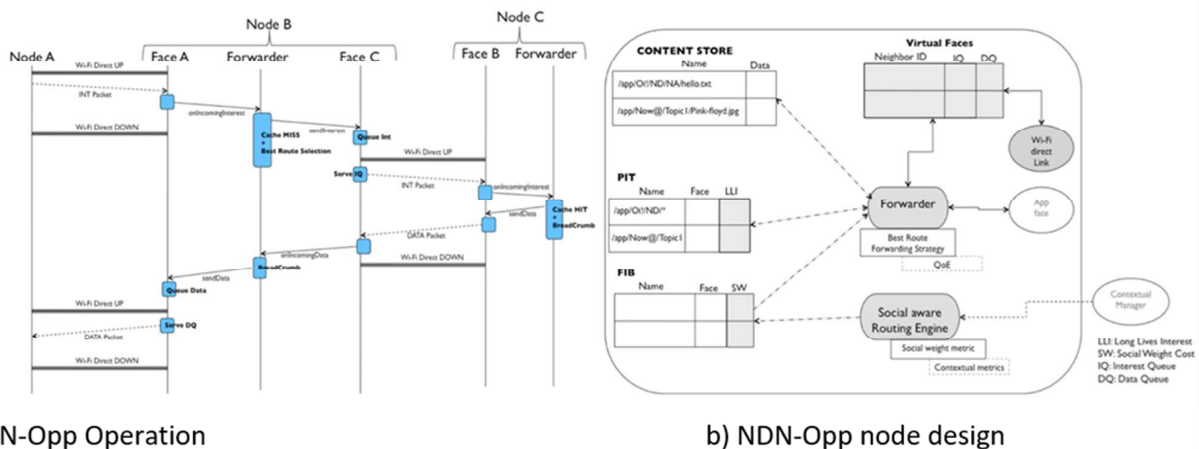
Routing in UMOBILE primarily targets opportunistic wireless environments. Thus, in the context of our architecture we have developed a relevant routing module, *NDN framework for Opportunistic Networks* (NDN-Opp), that is able

to exploit wireless communication opportunities. NDN-Opp provides an Android implementation of NDN for opportunistic networks.

NDN-Opp handles the dynamics of opportunistic networks by forwarding interest packets towards neighbors with high probability to meet nodes carrying the interested data. To be compatible with NDN, NDN-Opp uses the NDN best route forwarding strategy to deliver interest packets, and the NDN “breadcrumb” approach to deliver data packets based on the information stored on the PIT.

NDN-Opp implements the NDN pull communication model, supporting data sharing applications, such as *Now@*, and a push communication model used by interactive applications, such as the short message application, *Oi!*. To handle intermittent connectivity, NDN-Opp includes the concept of Opportunistic faces (*Opp Face*), relying on Wi-Fi Direct for all device-to-device communications. An *Opp Face* can be in two states (ON and OFF), and is named after an identified neighbor. Since forwarded packets may not be immediately dispatched, *Opp Faces* implement two queues: i) Interest Queue (IQ), storing Interest packets to be sent; ii) Data Queue (DQ) storing pointers to the Data block (in the Content Store) to be sent.

As shown in Figure 4a, the operation of an NDN-Opp node is similar to the NDN operation, as the node state is changed by the arrival of Interest or Data packets. Figure 4a also shows that contrary to the faces used by NDN, *Opp Faces* create an indirection: NDN operation finishes with a packet being sent; NDN-Opp operation ends with a packet being stored in an *Opp Face*. Packets are transmitted by serving interest and data queues when *Opp Faces* get up.



a) NDN-Opp Operation

b) NDN-Opp node design

Figure 4. NDN-Opp illustration

Figure 4b shows the NDN-Opp design, based on a novel forwarder and routing engine. Upon reception of an Interest packet, the NDN state machine is used, but the PIT also stores information related to the Interest duration (Long Lived Interests - LLI). Interest packets are forwarded based on the NDN best route forwarding strategy, as the cost of name prefixes (stored in the FIB) is related to social weights computed by a social-aware routing module. Data packets are forwarded following the normal NDN operation.

The social-aware routing engine computes the costs of using *Opp Faces* to reach data related to specific name prefixes. The proposed routing engine can run any type of social-aware opportunistic routing algorithm; NDN-Opp currently makes use of the Time-Evolving Contact Duration algorithm (TECD) [11] used by routing protocols such as dLife [4] and SCORP [6]: computing routing information based on social interactions has great potential as less volatile proximity graphs are created.

CONTEXTUALIZATION

UMOBILE also addresses social-awareness via its context plane, with the purpose of improving data dissemination. Social awareness has been on the rise, particularly when considering the capability of exploiting personal device

mobility to reduce the need for data muling, as well as leveraging traffic locality to improve service/content delivery.

The UMOBILE contextual plane is defined by the Contextual Manager (CM, Figure 5), a service that runs in background on an end-user device or on an Access Point. The CM captures information concerning the device affinity network (roaming patterns and peers over time and space) as well as usage habits and interests (internal device information). Costs derived from such contextualization are then passed, upon demand or periodically, to other UMOBILE modules to assist in different network operational aspects with the ultimate goal of achieving more efficient data dissemination.

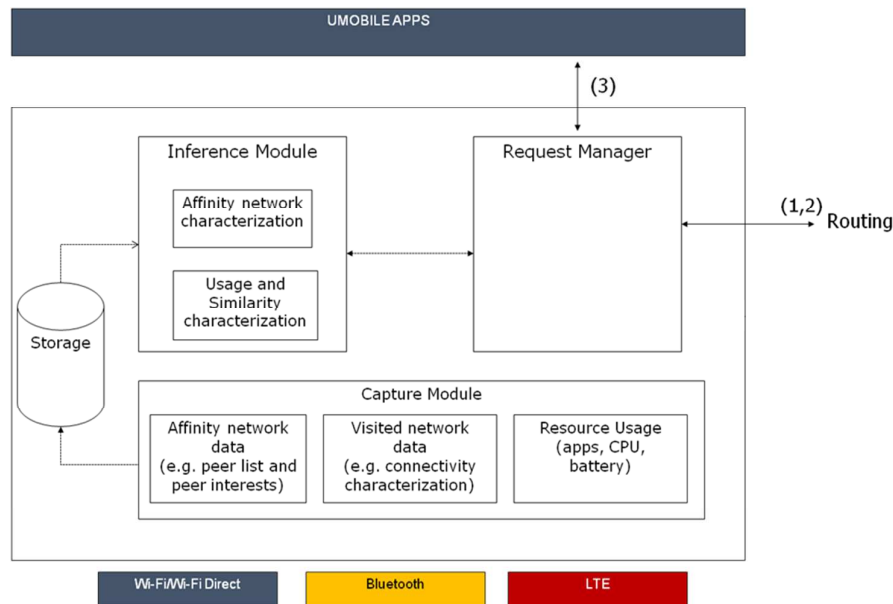


Figure 5. Contextual Manager, high-level scheme.

The CM holds three different interfaces towards other modules (two interfaces for the routing module, (1), (2); one interface for native applications (3)) allowing other modules or applications to query or to obtain information from the CM. The type of information that any CM interface provides can be categorized into two main sets: i) **affinity network characterization data**; ii) **usage and similarity characterization data**. Affinity network information concerns, e.g., peer status over time and space as well as affinities (matches) between source nodes and peers. Examples of indicators that are considered in the characterization of affinity networks and are provided via the different CM interfaces to other UMOBILE modules are peer lists, interests associated with each peer and indicators of peer status (e.g., battery) and connectivity.

Indicators that can be provided and that concern usage and similarity characterization are built upon data collected internally (in the device). Examples of indicators that fit this category are, for instance, geo-location, as well as categories of preferred applications.

NORTHBOUND APIs

Applications interact with the network in various ways. On UMOBILE, we provide a framework to cover different application needs, as well as provide APIs that are general enough to open new design spaces for application designers. In the following section, we describe the relevant enhancements to ICN:

PUSH SERVICES

NDN natively supports the pull-based communication model where the transfer of named content is initiated by the consumer. This model is suitable for applications that involve the transfer of large and persistent content, but not for ones involving dissemination of frequently produced ephemeral contents [12]. For this class of

1 applications, the push-based communication is more convenient. In this model, the producer takes the initiative
2 of pushing the content to consumers. To account for generality, we have enhanced the UMOBILE architecture with
3 a push-based communication model by implementing three different variations [9].

4 **Interest polling:** Suitable for applications in which the consumer does not know when the content is made
5 available by the producer. It relies on a polling mechanism that allows the consumer to issue periodic request
6 against the producer. The producer places the content in a Data packet and pushes it to the consumer. The
7 drawbacks of polling are that it induces network overhead and that the freshness of the received data depends on
8 the polling frequency.

9 **Interest notification:** Useful when the content is a piece of text that is small enough to be appended in an Interest
10 name by the producer. Upon the reception of the Interest message, the receiver can optionally send an ack Data
11 to acknowledge reception.

12 **Publish data dissemination:** Convenient for transferring content of any type (including text) and size. When the
13 producer makes the content available, it sends a notification Interest to the consumer who replies with an Interest
14 message requesting the content. The first chunk of the content is sent together with information about the
15 remaining chunks which are transferred in subsequent exchanges.

21 KEBAPP APPLICATION SHARING PLATFORM

22 Tens of applications are being introduced every day in the application markets. Given the technology progress and
23 the fact that mobile devices are becoming strong computing devices, mobile *applications* are expected to follow suit
24 and become computation-heavy, bandwidth-hungry and latency-sensitive. UMOBILE introduces a new mobile
25 computing paradigm to alleviate some of the network stress that mobile applications are already putting into the
26 network, *e.g.*, in case of crowded areas and events, where the mobile network effectively collapses.

27 We take a step further from content sharing, host-centric communications and fog computing and focus on the
28 prevailing application-centric computation and communication model. The proposed framework [13] explicitly
29 enables access to the desired *processed and non-personalized* information through the concept of *application*
30 *sharing*, effectively leveraging on a *pool of application resources*. Namely, we leverage application-centrism to
31 facilitate information discovery through application-driven and application-defined, hierarchical namespaces. Given
32 the ad-hoc nature of the proposed computation framework, our approach further extends these namespaces by
33 introducing the concept of *keywords*, *i.e.*, free text or application-driven (GUI) parameters used to enable the
34 invocation of applications at co-located mobile devices.

35 This enables the description, discovery and retrieval of *processed* information, further supporting variable accuracy
36 results, instead of only exact matches, *e.g.*, a search result that does not contain all search terms. Note that the
37 invocation of remote processing (in co-located smartphone or Wi-Fi AP devices) is central to our framework, as
38 opposed to previous work on retrieving static content from nearby devices. Our *keyword-based mobile application*
39 *sharing framework (KEBAPP)*, manages connectivity in an application-centric way, *i.e.*, coupling connectivity options
40 and opportunities to applications and their namespaces. By leveraging KEBAPP, UMOBILE extends existing ICN
41 primitives, namely CCN/NDN, thus resulting in a generic solution across different applications and overcoming the
42 pitfalls of IP. Initial results, showing that the average response time for a request can be minimized to less than a
43 couple of seconds with only a small percentage of users in a subway station sharing their application data, look
44 promising [13].

QUALITY OF SERVICE

In UMOBILE, we place particular focus on QoS at the edges. To achieve this goal, we have developed an opportunistic service deployment platform that is used to transfer services to the edge, reducing traffic at the core and enabling service deployment at remote regions where they would otherwise be unavailable.

OPPORTUNISTIC SERVICE DEPLOYMENT

We use opportunistic service deployment as an application-level mechanism to overcome latency and availability constraints. We assume a business model where the network provider is in possession of network, computational and storage resources that it is willing to virtualize to host services [14]. In this model, the content producer delegates the responsibility of deploying QoS-constrained services to the network provider. A conceptual view of the architecture is shown in Figure 6.

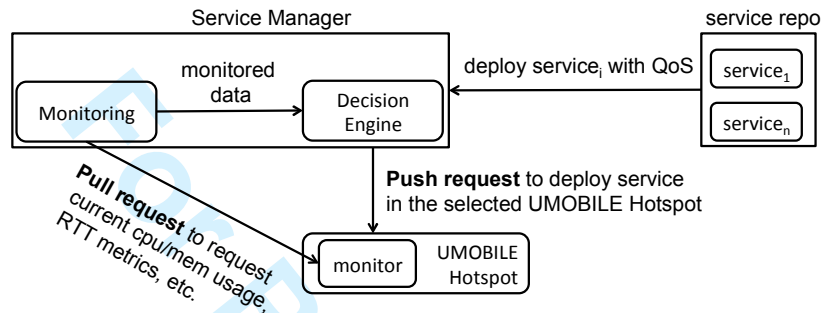


Figure 6. Opportunistic service deployment architecture.

In our implementation, we use the abstractions offered by NDN and take advantage of monitoring and light-weight virtualization (e.g. dockerizing) to deploy services at the edge of the network, on demand. We exploit UMOBILE APIs, such as the push-based communication model (publish data dissemination).

The *Service Manager* implements all the functionalities that a service provider needs to deploy his services. It includes a service repository (*service repo*) where dockerized compressed images ($service_1, \dots, service_n$) of the services are stored. It also includes a *Decision Engine* that is responsible for deciding what services to deploy and in which UMOBILE Hotspot. The DE makes its decisions based on monitored information collected by the *Monitoring* component, which is responsible for placing *Pull requests* against the *monitor* deployed in each *UMOBILE Hotspot* to collect metrics about their resource consumption status.

The *Service Manager* uses *Push requests* to transfer the dockerized service image to the UMOBILE Hotspot selected by the DE and instantiate it. Our implementation benefits from the facilities of the UMOBILE architecture in terms of efficient data dissemination. For instance, if a service is already deployed in neighbor UMOBILE Hotspots, the UMOBILE Hotspot selected by the DE can opportunistically retrieve the service from a neighbor.

Results [15] indicate that service deployment in Raspberry-Pi hotspots can be a viable approach to instantiate services at the edge of the network when compared against metrics such as CPU utilization and load, or server response time.

CONCLUSION

In this article, we have introduced the UMOBILE architecture. Drawing from diverse developments in the ICN, DTN and opportunistic networking community we have envisioned a new, pervasive architecture that is designed for the edge of the network and leverages users and their context to improve its efficiency. We have described our advancements towards realizing this vision and provided an overview of the mechanisms we utilize to achieve it.

1 Research is still being carried on many levels. As D2D communications and information-centric primitives are
2 drawing increasing attention in next-generation networks, the integration of all relevant aspects in the UMOBILE
3 architecture is crucial to support emerging services and applications towards extending the reach of future
4 Internet.
5
6

7 REFERENCES

- 10 1 G. Xylomenos, et al, "A survey of information-centric networking research." IEEE Communications Surveys &
11 Tutorials 16.2, 2014, pp. 1024-1049.
- 12 2 L. Zhang, et al, "Named data networking." ACM SIGCOMM Computer Communication Review 44.3, 2014, pp. 66-
13 73.
- 14 3 P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: social-based forwarding in delay tolerant networks." Mobile
15 Computing, IEEE Transactions on, vol. 10, November, 2011, pp. 1576–1589.
- 16 4 W. Moreira, P. Mendes, and S. Sargento, "Opportunistic routing based on daily routines." in Proceedings of the
17 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012, pp.
18 1–6.
- 19 5 R.I. Ciobanu, C. Dobre, D.G. Reina, and S. L. Toral, "A dynamic data routing solution for opportunistic networks."
20 In Telecommunications (ConTEL), 2017 14th International Conference on, IEEE, 2017, pp. 83-90.
- 21 6 W. Moreira, P. Mendes, and S. Sargento, "Social-aware opportunistic routing protocol based on user's
22 interactions and interests." In International Conference on Ad Hoc Networks, Springer, Cham, 2013, pp. 100-115.
- 23 7 K. Fall, "A delay-tolerant network architecture for challenged internets." In Proceedings of the 2003 conference
24 on Applications, technologies, architectures, and protocols for computer communications, ACM, 2003, pp. 27-34.
- 25 8 UMOBILE project "D2.3 System and network requirements specification" Available: [http://umobile-](http://umobile-project.eu/deliverables)
26 [project.eu/deliverables](http://umobile-project.eu/deliverables) [Accessed: 18- Sep- 2017]
- 27 9 UMOBILE project "D3.2 UMOBILE architecture report (2)" Available: <http://umobile-project.eu/deliverables>
28 [Accessed: 18- Sep- 2017]
- 29 10 O. Ascigil, V. Sourlas, I. Psaras and G. Pavlou, "Opportunistic off-path content discovery in information-centric
30 networks." 2016 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN), Rome,
31 2016, pp. 1-7.
- 32 11 W. Moreira, P. Mendes, "Impact of human behavior on social opportunistic forwarding." Elsevier Ad Hoc
33 Networks, Volume 25, February 2015, Part B, pp 293-302.
- 34 12 A. Carzaniga, M. Papalini, and A.L. Wolf, "Content-based publish/subscribe networking and information-centric
35 networking." in Proceedings of the ACM SIGCOMM workshop on Information-centric networking. ACM, 2011, pp.
36 56-61.
- 37 13 I. Psaras, et al, "Keyword-based mobile application sharing." in Proceedings of the Workshop on Mobility in the
38 Evolving Internet Architecture. ACM, 2016, pp. 1-6.
- 39 14 B. Schwarz, "Content delivery networks 3.0" Available: [bpl-business.com/pdf/exhibitors/broadpeak/Live-webTV-](http://bpl-business.com/pdf/exhibitors/broadpeak/Live-webTV-Whitepaper.pdf)
40 [Whitepaper.pdf](http://bpl-business.com/pdf/exhibitors/broadpeak/Live-webTV-Whitepaper.pdf) [Accessed: 18- Sep- 2017]
- 41 15 UMOBILE project "D4.4 Set of QoS interfaces and algorithms" Available: <http://umobile-project.eu/deliverables>

[Accessed: 18- Sept- 2017]

AUTHOR BIOGRAPHIES

Christos-Alexandros Sarros obtained his diploma in Electrical and Computer Engineering from Democritus University of Thrace, Greece in 2016. Currently, he is pursuing a PhD on the same university under the advisory of Prof. Vassilis Tsaoussidis and is working as a research assistant in Athena Research and Innovation Center. His research interests lie in the area of Information-Centric Networking and Delay/Disruption Tolerant Networking.

Sotiris Diamantopoulos received a Diploma in Electrical and Computer Engineering from Democritus University of Thrace, Greece in 2008, and a M.Sc. in Computer Networks from the same University. He is currently a PhD candidate under the advisory of Prof. Vassilis Tsaoussidis. His research interests lie in the fields of novel networking technologies and, more specifically, in transport and routing mechanisms.

Sergi Reñé (s.rene@ucl.ac.uk) received his B.S. in electrical engineering and M.S. in network engineering from the Technical University of Catalonia (UPC), Barcelona, Spain, in 2006 and 2010, respectively. In 2015, he also received his Ph.D. from the Network Engineering Department at UPC. In 2015, he joined UCL, carrying out research on information-centric networks, congestion control, and mobility.

Ioannis Psaras (i.pсарas@ucl.ac.uk) is an EPSRC Fellow at the Electrical and Electronic Engineering Department of UCL. He is interested in resource management techniques for current and future networking architectures with particular focus on routing, replication, caching, and congestion control. Before joining UCL in 2010, he held positions at the University of Surrey, United Kingdom, and Democritus University of Thrace, Greece, where he also obtained his Ph.D. in 2008.

Adisorn Lertsinsruttavee received the M.E. degree in ICT from Asian Institute of Technology, Thailand and the M.S. degree in Communication Networks and Services from Telecom Management SudParis, France in 2009. He received the Ph.D. degree in computer science from UPMC Sorbonne Universities, France in 2013. He is currently a research associate at the Computer Laboratory, University of Cambridge, UK. His research interests include cognitive radio networks, future internet architectures, and information centric networks.

Carlos Molina-Jimenez received the MS degree in computer science from the Universidad Nacional Autonoma de Mexico (UNAM) and the Ph.D. degree in computing science from Newcastle University, United Kingdom. He is currently a senior research associate at Computer Laboratory, University of Cambridge. His research interests lie in distributed systems and are currently focused on middleware for supporting contract-regulated business applications, policy enforcement, and formal validation of business processes.

Paulo Mendes is associated professor at University Lusofona, and director of the Ph.D program on Informatics. He is co-founder of the COPELABS research center, where he leads the SITI research group and co-founder and CTO Senception Lda. In 2004, Paulo Mendes got his Ph.D. degree in Informatics Engineering from the University of Coimbra, while being a visiting scholar at Columbia University. He participates regularly in IETF working groups, and is an ACM and IEEE member.

Rute Sofia (PhD, 2004) is the CEO of Senception Lda, startup focused on network mining, and an Associate professor of University Lusofona as well as Senior Researcher of the R&D unit COPELABS, Lisbon, Portugal. Her research interests concern networking architectures; mobility modelling and management; network mining. She holds several patents and publications in her field of expertise. She participates actively in several IETF working groups, and is an ACM and IEEE member.

Arjuna Sathiseelan is a Senior Research Associate at the Computer Laboratory, University of Cambridge. He leads the Networking for Development (N4D Lab). He is the Chair of IRTF Global Access to the Internet for All (GAIA) research group and a member of the Internet Research Steering Group (IRSG). He has a Ph.D. in Networking and M.Sc. in Computing and Internet Systems from Kings College London in 2001 and 2005.

1 **George Pavlou** is Professor of Communication Networks at the Dept. of Electronic Engineering, University College
2 London, UK, where he coordinates networks and services research activities. His research interests focus on
3 networking and network management, including traffic engineering, autonomic networking, information-centric
4 networking and software-defined networks. He has been instrumental in a number of research projects that
5 produced significant results with real-world uptake and has contributed to standardization activities in ISO, ITU-T
6 and the IETF.

7
8 **Jon Crowcroft** received the B.S. degree in physics from the University of Cambridge in 1979, and the M.Sc. degree in
9 computing and Ph.D. degree from University College London in 1981 and 1993, respectively. He has been a Professor
10 with the University of Cambridge since 2001, where has been involved in Internet support for multimedia
11 communications. Among others, he is a Fellow of the Royal Society, the Association for Computing Machinery, and
12 the IET.

13
14 **Vassilis Tsaoussidis** is Professor at DUTH and adjunct Professor at Athena RC, holding degrees in Applied
15 Mathematics (Aristotle University) and Computer Science (Ph.D. in Computer Science, Humboldt University, Berlin).
16 Vassilis joined the research community of Rutgers in 1996 and later the faculty communities of Stony Brook and
17 Northeastern. He also joined MIT as visiting Professor in 2009. Vassilis was scientific leader for ESA, FP-7 and
18 HORIZON 2020 projects, including the UMOBILE project.
19

20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Review Only