

HHS Public Access

Author manuscript *IEEE Netw.* Author manuscript; available in PMC 2021 November 01.

Published in final edited form as:

IEEE Netw. 2020; 34(6): 235-241. doi:10.1109/MNET.011.2000169.

NDNTP: A Named Data Networking Time Protocol

Abderrahmen Mtibaa,

University of Missouri-St. Louis, US.

Spyridon Mastorakis University of Nebraska, Omaha, US.

Abstract

Named Data Networking (NDN) architectural features, including multicast data delivery, stateful forwarding, and in-network data caching, have shown promise for applications such as video streaming and file sharing. However, collaborative applications, requiring a multi-producer participation introduce new NDN design challenges. In this paper, we highlight these challenges in the context of the Network Time Protocol (NTP) and one of its most widely-used deployments for NTP server discovery, the NTP pool project. We discuss the design requirements for the support of NTP and NTP pool and present general directions for the design of a time synchronization protocol over NDN, coined Named Data Networking Time Protocol (*NDNTP*).

I. INTRODUCTION

Information-Centric Networking (ICN) [1] and a prominent realization of its vision, the Named Data Networking (NDN) architecture [2], use application-defined names to retrieve secured data through a stateful forwarding plane. This waives the need for costly mappings between the application data and its physical location(s). The NDN design and its builtin features, such as multicast data delivery and in-network caching, have shown promise for applications and services such as video streaming, file sharing, and others [3], [4].

The ICN/NDN community has worked on adapting IP-based applications and services to run on top of NDN [3], including general guidelines on adapting non-NDN applications to function over NDN [5]. However, the NDN adoption for certain categories of applications remains largely unexplored. For instance, crowd-sensing and crowd-sourcing applications and services that enable multiple producers to return different content for the same Interest may introduce new challenges to the existing NDN design.

In this paper, we consider the Network Time Protocol (NTP) [6] as a use-case given the universal need of applications and devices for a time synchronization service. Through this use-case, we highlight new NDN design requirements and design an NDN-based service, called Named Data Networking Time Protocol (*NDNTP*), to provide time synchronization to applications, systems, and devices within islands of NDN connectivity. We present a preliminary investigation of how functionality equivalent to the IP-based NTP can be

Author Manuscript

provided over NDN. We further consider NTP pool [7] as one of the most widely-used NTP server discovery deployments. Motivated by the requirements of NTP and NTP pool, we discuss the *NDNTP* design as well as enhancements of the native NDN architectural features to support the NTP functionality.

As part of our work, we propose design directions for: (i) forwarding time synchronization requests towards different sets of servers and retrieving multiple time samples from each of these servers; (ii) fine-grained control over how far in the network time synchronization requests can travel; and (iii) controlling the NDN in-network caching and request aggregation mechanisms to enable the retrieval of up-to-date time samples. The design directions discussed throughout the paper can be leveraged by various NDN applications that share the same requirements as *NDNTP*.

The rest of this paper is organized as follows: Section II presents a brief background on NDN, NTP, and NTP pool, Section III discusses the design requirements of *NDNTP*, while Sections IV, V, and VI present alternative directions to satisfy each of the design requirement. Section VII briefly discusses *NDNTP* extensions, while Section VIII concludes the paper.

II. BACKGROUND

In this section, we first present an NDN primer. Then we present some background on NTP and NTP pool to prepare readers for the discussion in the rest of the paper.

A. NDN Background

NDN [2] features a data-centric, receiver-driven communication model, where each piece of data has an application-defined name. *Consumer applications* send requests for named data, called Interests, towards *data producer applications*. Each Interest carries the name of the requested data. Once a data producer receives an Interest, it sends back a Data packet, which is cryptographically signed by the data producer and contains the requested data.

Interests are forwarded based on their names towards producers by NDN forwarders. To achieve that, NDN forwarders make use of a Forwarding Information Base (FIB), which contains name prefixes along with a set of outgoing interfaces for each prefix. Forwarders also maintain a Pending Interest Table (PIT), where they store network state for each Interest that has been forwarded, but the corresponding data has not been received yet. A Data packet follows the reverse path of the corresponding Interest based on the network state in PIT and can be forwarded all the way back to a consumer only if there is a valid PIT entry for the corresponding Interest at each hop forwarder. Once a Data packet is retrieved in response to an Interest, the corresponding PIT entry at each hop forwarder will be consumed. This happens as part of the "flow balance" principle of NDN, which mandates that each Interest packet can bring back only one Data packets to be returned to a consumer, the "flow balance" principle is violated. Finally, each forwarder is equipped with a Contest Store (CS), where recently received Data packets are cached to satisfy Interests for the same data in the future.

B. NTP and NTP Pool Background

NTP [6] has been one of the longest running protocols on the Internet. It was created due to the need of applications and services for time synchronization over packet-switched networks. Discovering NTP servers and automating this process for clients has been essential for the deployment of NTP. Among the different service discovery schemes proposed (*e.g.*, broadcast/multicast/manycast servers), NTP pool has been one of the most widely-used ones [7]. NTP pool is a cluster of volunteer NTP servers used by large numbers of clients around the world (*e.g.*, it has been the default time synchronization option for most Linux distributions). NTP pool groups NTP servers based on their IP address geolocations into continental and country zones (Fig. 1).

Each NTP server in the pool is assigned a score, which reflects the accuracy of the provided time samples are This score is determined by a monitoring station. Once a new server joins the pool, it is assigned a low score by default. The monitoring station probes this newly joined (and every other) server over time, verifying the accuracy of its clock. As servers respond with accurate clock readings to the requests of the monitoring station, their scores improve. Once the score of a server reaches a certain threshold, the server will start receiving requests by NTP clients. NTP clients discover servers in the pool by querying DNS for the "pool.ntp.org" domain name. The DNS resolution of "pool.ntp.org" would usually return servers within or close to the client's country.

III. DESIGN REQUIREMENTS OF NDNTP

In this section, we first give an overview of the *NDNTP* network model. Then, we present *NDNTP* s design goals and requirements, focusing on the NDN features that need to be enhanced in the context of this design.

A. Network Model and Assumptions

We consider NDN islands, where consumers (*NDNTP* clients) are scattered geographically into multiple zones such as Europe, America, and Asia. In each zone, there is a number of *NDNTP* servers acting as NDN data producers. *NDNTP* clients are inter-connected with *NDNTP* servers through one or more network hops (NDN forwarders).

The *NDNTP* servers receive time requests (Interest packets) under the "/NDNTP/time" namespace sent by *NDNTP* clients, who initiate the time synchronization process. Servers respond to requests by generating and signing *NDNTP* responses (Data packets), which contain the current timestamp and other necessary fields as determined by the NTP specification [8]. To receive requests from clients, servers announce the "/NDNTP/time" namespace to the network. A name-based routing protocol, such as NLSR [9]), propagates the announcement and establishes routes towards servers (*i.e.,* install the proper forwarding information on the FIB of NDN forwarders).

B. NDNTP's Design Requirements

Deploying *NDNTP* in an NDN island can be challenging, since functionality requirements motivated by NTP and NTP pool require the augmentation of certain NDN architectural

Multi-server time sample fetching: The best current practices of NTP deployment [10] indicate that a client should contact more than one NTP servers in order to select accurate time sources and disregard unreliable ones. NDN inherently facilitates reaching multiple servers through its multipath/multicast nature. In this context, clients send requests for "/ NDNTP/time", which will be satisfied by any server running the time synchronization service. In *NDNTP*, we would like multicast requests in order to bring a response from each server they reached back to a client. However, NDN allows only a single Data packet to use the reverse path to the client–often the closest server will use the network state stored in PIT, thus responses sent by other servers will not find PIT state to be returned to the client. Fig. 2 illustrates such a scenario, where a client reaching three different servers would receive a single Data packet from the closest server (*i.e.*, *S*₁). This packet will consume the PIT entry at forwarder *F*₁, leaving data from other servers without a path back to the client.

Moreover, clients typically request multiple time samples from the same set of servers to increase the accuracy of the time synchronization process. In NDN, clients rely on the network to guide requests towards the "best" (*i.e.*, often the closest) server, without being able to control whether subsequent time synchronization requests will reach the same or different servers. To this end, additional mechanisms are needed to let more than one requests be forwarded towards the same set of servers. All these challenges are discussed and different design directions are proposed in Section IV.

Distance-based server selection: Through the NTP pool, clients often select a set of NTP servers that are in the same country or continental zone as themselves. Clients discover different sets of servers over time in order to: (i) avoid the impact of malicious groups of neighboring servers colluding to return bogus timestamps to clients; and (ii) balance the load among servers in a zone.

NDN makes use of application-defined names for communication purposes, enabling access to services typically from the closest source/server that can offer them. However, to achieve distance-based server selection, *NDNTP* clients may need mechanisms for fine-grained control over how far their requests should travel and how these requests can be satisfied by servers that are potentially not the closest ones to clients. We further discuss this challenge and propose directions to address it by enhancing the NDN design in Section V.

Time synchronization freshness and accuracy: NTP clients aim to retrieve fresh (up-to-date) time samples, increasing the accuracy of the time synchronization process. However, due to NDN in-network caching, *NDNTP* client requests for time synchronization can be satisfied with outdated responses that have been cached in the network. Moreover, *NDNTP* requests (from the same or different clients) reaching a forwarder that has another *NDNTP* request pending in its PIT could result in request aggregation. That is, subsequent *NDNTP* requests may not be forwarded to a server, but will be satisfied when the response to the first pending request is received by the forwarder. This could skew the client round-

trip delay measurements, impacting the accuracy of the time synchronization process. We further discuss these challenges and propose directions to tackle them in Section VI.

IV. ENABLING MULTI-SERVER TIME SAMPLE FETCHING

In Section III, we present different design directions for the following *NDNTP* design goals/ objectives: (i) ensure that clients can reach multiple servers and collect multiple time samples for each of these servers; and (ii) enable a multicast request to bring a response from each server that it reached back to a client.

A. Gathering Multiple Samples From Multiple Servers

It is desirable for NTP clients to: (i) contact more than one NTP servers in order to identify accurate time sources and disregard inaccurate ones; and (ii) collect multiple time samples from each NTP server to enhance the accuracy of the synchronization process due to the dynamic nature of network conditions. The NDN communication model is inherently multipath/multicast supporting requirement (i), since an Interest can be forwarded towards multiple producers (*NDNTP* servers in our case). However, deterministically contacting the same server multiple times in the context of requirement (ii) is counter-intuitive to the purpose of NDN, where data can be retrieved from any party that can provide it. As a result, the NDN architecture itself does not provide explicit mechanisms to ensure that Interests will be forwarded along a specific path towards a certain server. Below, we present a wide spectrum of design directions to fulfill these requirements ranging from solutions of unicast nature to solutions of multicast nature.

The first design direction is inspired by source routing. Clients instruct the network about which paths their Interests should take through techniques to create path labels for Interest forwarding [11]. Each Interest carries such a label, which determines the next-hop that the Interest should be forwarded to. At each hop, NDN forwarders use this label for Interest forwarding bypassing the FIB lookup process. At the beginning of their operation, *NDNTP* clients will perform a path discovery process, where they will acquire labels for paths to different servers. This process will end when a client acquires labels for the desired number of servers. The client uses the same path label to collect different time samples from the same server, since Interests carrying the same label will be forwarded along the same path towards the same server.

Another direction would be to utilize the stateful NDN forwarding plane. Forwarding modules (strategies) can be created and deployed on NDN forwarders, so that clients can discover multiple *NDNTP* servers and collect a number of time samples from each of these servers. Specifically, *NDNTP* clients create a "session-like" paradigm with a server, so that they request multiple time samples from it. An example namespace to achieve that is illustrated in Fig. 3. The name prefix includes the *NDNTP* name prefix ("/NDNTP/time"), followed by: (i) *a random hash* to by-pass the aggregation of requests from different clients as further discussed in Section VI (each different hash identifies a set of sessions initiated by a client–each session is with a different server); (ii) *a session number* that identifies a specific time sample requested from a server.

Page 6

When an NDN forwarder receives an *NDNTP* request, it dispatches it to the corresponding forwarding module, which identifies requests from the same client but from different sessions. To this end, time sample requests for the same session number are forwarded along the same path, so that they reach the same *NDNTP* server. At the same time, the stateful network forwards requests for different session numbers along different paths towards different servers.

Multicast Support: We note that both solutions mentioned above can be extended to support multicast Interest forwarding. For instance, during the discovery process of paths performed by clients, the path labels can include multiple next hops at each forwarder. As a result, client Interests will be forwarded along different paths, potentially reaching different servers. In a similar manner, in the second approach, client Interests can include a list of session numbers instead of a single session number in order to enable multipath stateful forwarding. Clients will re-use the same list of sessions to ensure that subsequent requests for time samples reach the same set of servers. However, directions of multicast nature require multiple responses (one from each server that received a request) to be forwarded back to clients. We discuss solutions to that below.

B. Receiving Multiple Responses Upon Sending Multicast Requests

To address this challenge, one direction would be to enhance the NDN communication model, so that forwarders can accept multiple Data packets (one per interface that the corresponding Interest was forwarded through). These Data packets will be aggregated into a single packet, which will be forwarded back to the forwarder's previous hop [12]. In this way, forwarders can multicast *NDNTP* requests towards multiple servers, receive a response from each of these servers, and aggregate the responses into a single response. This "aggregated" response will eventually be received by the requesting *NDNTP* client. This enhancement does not invalidate NDN's principle of "flow balance, since a single response will be returned to clients. However, it can take arbitrarily long for a forwarder to collect multiple responses through different paths, since the length of each path and the network conditions might differ. This could impede the time synchronization accuracy by skewing the round-trip delay measurements of clients.

An alternative direction consists of consuming a PIT entry only after all the expected responses have been received. Specifically, a PIT entry will stay alive and bring back one response for each outgoing interface that the corresponding *NDNTP* request was forwarded through. This would help clients to collect accurate round-trip delay measurements and improve the time synchronization accuracy. However, this direction would invalidate the NDN "flow balance" principle, since corresponding PIT entries may stay open and accept several Interests.

V. DISTANCE-BASED SERVER SELECTION

While NDN inherently supports reaching the closest server offering the time synchronization service, there might be cases where reaching distant servers within a zone may be desirable in order to: (i) reduce targeted attacks, where entities controlling a group of neighboring servers collude to return bogus timestamps; and (ii) enable load balancing

among nearby and distant servers. Reaching servers within similar distances from a client may also be desirable in order to reduce jitter in round-trip delays.

Preliminary Experiments: To investigate whether NTP pool shares the above objectives and better understand how it operates, we performed a set of experiments. Specifically, we deployed NTP clients–two in North America (in Missouri, US and in New Mexico, US) and two in Europe (in Zurich, Switzerland and in Athens, Greece)–and we configured them to use NTP pool for time synchronization. Each client queries four NTP pool servers and collects four time samples from each of these servers (*i.e.*, the default configuration options for NTP) every minute over a 24-hour period¹.

Fig. 4 presents the CDF of the unique NTP pool servers sorted by popularity (*i.e.*, number of times a server is selected) in the four traces (results collected from each client). We notice that both North American traces have very similar trends characterized by almost identical CDF distributions, while the European traces have different properties compared to the traces from North America. Based on the European traces, the unique servers discovered are 15 to $20\times$ fewer than those discovered in North America (*e.g.*, 14 unique servers in the trace from Athens compared to 300 in the Missouri trace). Our understanding is that this is due to the difference in the zone sizes for the US and the European countries (Greece and Switzerland). For example, at the time of the experiments, there were 774 NTP servers in the zone of the US, while only 15 NTP servers in the zone of Greece.

Moreover, we found that the NTP pool servers in Europe are closer to the clients than the servers in the US. For example, more than 80% of the servers discovered from Athens are within 50 kilometers from the client, while less than 10% of the servers discovered from New Mexico are within 1000 kilometers from the client (*i.e.*, $250 \times$ further than the 50-kilometer neighboring circle in Athens). This is due to the different physical sizes of the countries–in terms of landmass, the US is 9,833,000 *km*², while Greece is only 131,957 *km*².

In Fig. 5, we illustrate the average distances (in kilometers) and the standard deviations per run of the NTP protocol (four servers are selected and four times samples are collected from each server per run) for the (a) Athens and the (b) New Mexico traces respectively². First, we notice that the four servers selected in each run tend to have similar distances from the client as shown by the lower values in standard deviation distributions of Fig. 5. Furthermore, we observe that in 40% of the runs for the Athens trace, requests reached four servers within the same distance of 45 to 50 kilometers. On the other hand, the four servers selected in each run of the New Mexico trace have different distances from the client, resulting in larger standard deviation values. Our understanding is that this is due to the larger physical surface of the US compared to Greece.

¹The datasets consisting of the raw timestamps and servers probed are available at https://github.com/amtibaa-cmu/NTP-Data-Traces. ²For better readability, we selected one trace from Europe and another from the US. We have verified the validity of our conclusions for all collected traces.

A. Lessons Learned

The lessons we learned from these experiments fall into three main categories: (i) *zone probing*: NTP pool employs a zone probing scheme that consists of reaching (nearby and occasionally distant) NTP servers from within the client's (continental and/or country) zone; (ii) *similar distance probing*: most servers selected in each run of the NTP protocol are within similar distances from the client, thus responses return to the client with similar delays; and (iii) *reaching a different set of NTP servers in each run*: NTP clients that make use of NTP pool reach, in general, a different set of servers for every run. Our data showed that the same US servers are rarely selected more than a few times in the 24-hour period. However, in zones, such as in Greece and Switzerland, where the number of NTP servers are selected, although infrequently (as shown in Fig. 5), to reduce the reliance on a small set of servers that may form a collaborative malicious group (collusion attacks).

B. Distance-Based Server Selection in NDNTP

According to the lessons we learned from our experiments, *NDNTP* needs to be able to reach servers within a given "*zone*" and within controlled distances from clients. While a zone in an IP-based NTP pool design refers to a country or a continent, which may have largely variable sizes, we define a zone as the region within a certain number of hops from a client. For instance, a client in New Mexico will be able to reach servers within *h* hops making the case of reaching servers in a different country likely in addition to reaching servers in other US states. This can be achieved through the deployment of stateful forwarding modules that allow clients to reach not only the closest but also distant *NDNTP* servers. We propose two client-based directions that enable fine-grained control over how far a request can travel.

The first design direction consists of taking advantage of the *hop limit* field in Interest packets. Clients that would like their requests to reach servers at least *h* hops away will set the hop limit value of their requests to *h*. Each forwarder decrements the hop limit value of *NDNTP* Interests and forwards them through the outgoing interface with the highest cost–in NDN, outgoing interfaces for a given name prefix are associated with a cost, thus the lower the interface cost is, the closer a server that can satisfy Interests for this name prefix is. When the hop limit value of an Interest reaches a given threshold, the forwarding plane switches to using the outgoing interface with the lowest cost for this Interest. As a result, after this point, the Interest will be forwarded with the goal of reaching the closest server(s). Through this approach, all the selected servers are at least *h* hops away and within similar distances from clients, so that round-trip delays for requests are similar.

The second direction would be to enable clients to include a probability to the name of their *NDNTP* requests. The lower this probability is, the higher will be the cost of the outgoing interface chosen by forwarders. For instance, a request with a name "/NDNTP/time/P=0.3" would be handled by forwarders so that they choose the outgoing interface with the highest cost with a probability of 0.7 and the lowest cost interface towards a server with probability of 0.3. This method allows clients to have fine-grained control over when requests should

reach closer versus more distant servers, however, it does not ensure reaching servers at certain distances (*e.g.*, at least *h* hops away) from clients.

VI. IN-NETWORK CACHING AND REQUEST AGGREGATION IN NDNTP

In this section, we discuss how to achieve time synchronization in the face of in-network caching and request aggregation.

In-network caching:

To avoid fetching outdated cached responses, the requests of *NDNTP* clients need to be satisfied directly by *NDNTP* servers rather than in-network caches. To this end, servers can set the value of the *Freshness Period* field of the responses they generate to a reasonably low value (or even 0), so that their responses become (almost) instantly stale when they are cached in the network [13]. Subsequently, clients send *NDNTP* Interests that contain the *MustBeFresh* field to avoid retrieving stale cached *NDNTP* responses.

An additional challenge arises when servers misbehave by ignoring this guideline and assign a large Freshness Period value to the generated responses. To this end, appropriate cache management policies can be deployed in the network to prevent forwarders from caching *NDNTP* responses for a long period of time or entirely avoid caching *NDNTP* responses. Note that this solution assumes trusted forwarders–we discuss security considerations and potential solutions in cases of malicious *NDNTP* servers and forwarders in Section VII-C.

NDNTP request aggregation in PIT:

NDNTP requests may be aggregated in PIT when a forwarder has a PIT entry for a request with the same name. As a result, subsequent *NDNTP* requests with the same name may not be forwarded to a server, but will be satisfied when the response to the first pending request is received by a forwarder. This can skew the client round-trip delay measurements posing another challenge for clients to accurately infer the current time. To address this challenge, *NDNTP* clients can randomize the name of their Interests, so that they avoid PIT aggregation. This can be achieved by attaching a random hash in the Interest name. Given that the *NDNTP* request name prefix is used for forwarding purposes, the randomization should happen in the request name suffix. We presented a mechanism to achieve name randomization in Section IV-A.

VII. DISCUSSION

In this section, based on the fundamental *NDNTP* design requirements and goals mentioned in Section III, we highlight and discuss potential extensions of the *NDNTP* design as well as the different research directions that may be pursued as a continuation of this work.

A. In-network Time Synchronization

Given that NDN forwarders are aware of the communication context, they can identify whether a specific Interest is an *NDNTP* request and whether a specific Data packet is an *NDNTP* response. This can be particularly useful in the following cases: (i) forwarders that have an up-to-date estimation of the current time (*e.g.*, they recently received a timestamp

from an accurate *NDNTP* server) can directly respond to *NDNTP* client requests with their own time. These responses are signed by the forwarder itself. This paradigm may constitute a distinct mode of operation that forwarders can enable/disable based on their load, available resources, and management policies; and (ii) forwarders can utilize ongoing exchanges between *NDNTP* clients and servers to satisfy their own time synchronization requirements. Specifically, forwarders can set their own clock based on the content of *NDNTP* responses forwarded by them back to clients.

B. Strata Organization and Synchronization

NTP servers are organized in strata that determine their distance from a reference clock. The larger the stratum number, the further away a server is from the reference clock. Servers that belong to stratum N+ 1 synchronize their clocks with servers that belong to stratum N, servers that belong to stratum N synchronize their clocks with servers of stratum N-1, etc. NTP servers of the same stratum can also peer with each other to address clock inconsistencies and achieve reliability. To achieve this synchronization/peering process, NDNTP servers that belong to a specific stratum can announce a name prefix "/NDNTP/ time/stratum=<% for peering purposes with other servers of the same stratum. To synchronize their clocks with servers of stratum 1, they use the namespace "/ NDNTP/time/stratum=1".

C. Security

NDNTP utilizes NDN's network-layer security principles for response authentication purposes. Specifically, clients based on the signature of responses can: (i) verify that the responses have not been spoofed (*e.g.*, due to man-in-the-middle attacks); and (ii) decide whether they trust the servers that generated responses based on a pre-established set of trust anchors [14].

NDN also provides a solid foundation for the mitigation of DDoS attacks directly at the network layer. The NDN stateful forwarding plane can limit/throttle DDoS traffic close to its source(s) on a per name prefix basis. In packet delay attacks, adversaries delay time synchronization requests and responses between clients and servers in an asymmetric manner in order to skew the round-trip delays measured by clients [15]. Such attacks can be mitigated through mechanisms for forwarding *NDNTP* requests over different network paths towards time servers. Furthermore, a threshold value for acceptable round-trip delays between time synchronization requests and the corresponding responses can be introduced– clients discard responses with round-trip delays larger than the threshold.

None of the mechanisms above protects against servers that turn malicious over time (*e.g.*, due to a security breach) or rogue forwarders that cache server time samples in order to satisfy client requests with outdated time samples. To address such cases, we can use a Distributed Ledger (DL) to log transactions (*i.e.*, timestamps received from servers) and verify if specific servers have sent inaccurate/bogus timestamps to clients. This DL can be implemented through an NDN distributed dataset synchronization protocol. Clients and a number of Verifier Nodes (VNs)–for example, a group of clients or trusted entities–form a

synchronization group, so that transactions added to the DL are received by every party in the group. The verification process can be: (i) proactive, where VNs mine the DL to verify the transaction correctness; or (ii) reactive, where clients report inconsistent/inaccurate timestamps triggering the verification process (assuming that the majority of selected servers is legitimate).

VIII. CONCLUSION AND FUTURE WORK

In this paper, we presented the challenges of designing *NDNTP*, a time synchronization protocol that provides functionality analogous to NTP over NDN, and we discussed general directions for its design. Our concluding remarks indicate that the legacy NDN architectural design needs to be augmented to support the requirements of an NTP-like protocol, such as reaching multiple time servers at the same time and returning multiple responses (one from each time server) back to clients, requesting multiple time samples from a set of selected servers, and reaching time servers within certain distances from clients. Future directions include designing in detail, implementing, and evaluating an *NDNTP* prototype that takes advantage of the NDN feature enhancements proposed in this paper for accurate time synchronization. We also plan to extend this work by performing an analytical study to improve the server selection process using a multi-objective function.

ACKNOWLEDGEMENTS

This work is partially supported by a pilot award from the Center for Research in Human Movement Variability and the NIH (P20GM109090), a planning award from the Collaboration Initiative of the University of Nebraska system, and the Nebraska Tobacco Settlement Biomedical Research Development Funds.

BIOGRAPHIES

Abderrahmen Mtibaa (amtibaa@umsl.edu) is currently an Assistant Professor at Department of Computer Science in the University of Missouri–St. Louis. Prior to that he has occupied several research positions including a visiting assistant professor at the Computer Science department in New Mexico State University; a Research Scientist at Texas A&M University; and a Postdoc in the School of Computer Science at Carnegie Mellon University. His current research interests include Information-Centric Networking, Networked Systems, Social Computing, Personal Data, Privacy, IoT, mobile computing, pervasive systems, mobile security, and mobile opportunistic networks/DTN.

Spyridon Mastorakis (smastorakis@unomaha.edu) is an Assistant Professor in Computer Science at the University of Nebraska Omaha. He received his Ph.D. in Computer Science from the University of California, Los Angeles (UCLA) in 2019. He also received an MS in Computer Science from UCLA in 2017 and a 5-year diploma (equivalent to M.Eng.) in Electrical and Computer Engineering from the National Technical University of Athens (NTUA) in 2014. His research interests include network systems and protocols, Internet architectures (such as Information-Centric Networking and Named-Data Networking), and edge computing.

REFERENCES

- [1]. Xylomenos G et al., "A survey of information-centric networking research," IEEE Communications Surveys & Tutorials, vol. 16, no. 2, pp. 1024–1049, 2013.
- [2]. Zhang L et al., "Named data networking," ACM SIGCOMM Computer Communication Review, vol. 44, no. 3, pp. 66–73, 2014.
- [3]. Mastorakis S, Afanasyev A, Yu Y, and Zhang L, "nTorrent: Peer-to-peer file sharing in named data networking," in 2017 26th International Conference on Computer Communication and Networks (ICCCN). IEEE, 2017, pp. 1–10.
- [4]. Gusev P and Burke J, "NDN-RTC: Real-time videoconferencing over named data networking," in Proceedings of the 2nd ACM Conference on Information-Centric Networking, 2015, pp. 117– 126.
- [5]. Liang T, Pan J, and Zhang B, "Ndnizing existing applications: research issues and experiences." in ICN, 2018, pp. 172–183.
- [6]. Mills DL, "Internet time synchronization: the network time protocol," IEEE Transactions on communications, vol. 39, no. 10, pp. 1482–1493, 1991.
- [7]. "NTP Pool Project," https://www.ntppool.org, 2019, Accessed on May 21, 2020.
- [8]. Mills D, Martin J, Burbank J, and Kasch W, "Network time protocol version 4: Protocol and algorithms specification," RFC, 2010.
- [9]. Hoque A, Amin SO, Alyyan A, Zhang B, Zhang L, and Wang L, "Nlsr: named-data link state routing protocol," in Proceedings of the 3rd ACM SIGCOMM workshop on Information-centric networking. ACM, 2013, pp. 15–20.
- [10]. Reilly D, Stenn H, and Sibold D, "Network time protocol best current practices," Work in Progress, draft-ietf-ntp-bcp-00, 2017.
- [11]. Moiseenko I and Oran D, "Path switching in content centric and named data networks," in Proceedings of the 4th ACM Conference on Information-Centric Networking. ACM, 2017, pp. 66–76.
- [12]. Mastorakis S, Mtibaa A, Lee J, and Misra S, "ICedge: When Edge Computing Meets Information-Centric Networking," IEEE Internet of Things Journal, 2020.
- [13]. Mastorakis S, Gusev P, Afanasyev A, and Zhang L, "Real-Time Data Retrieval in Named Data Networking," in 2018 1st IEEE International Conference on Hot Information-Centric Networking (HotICN). IEEE, 2018, pp. 61–66.
- [14]. Zhang Z et al., "An overview of security support in named data networking," IEEE Communications Magazine, vol. 56, no. 11, pp. 62–68, 2018.
- [15]. Sibold D, Roettger S, and Teichel K, "Network time security," Internet Engineering Task Force, Internet Draft version 15, 2016.





Zones of NTP servers based on their geolocations in NTP pool; clients reach servers in the same or a nearby zone.



Fig. 2:

Challenge to collect multiple Data packets for a single multicast Interest. The client sends an Interest that reaches all servers S_1 , S_2 , S_3 . Link delays indicate that S_1 will receive and send its data back earlier than the other servers.



Fig. 3:

Namespace design enabling clients to: (i) reach multiple *NDNTP* servers; and (ii) collect multiple time samples from a specific *NDNTP* server



Fig. 4:

Distribution of server occurrences; servers are sorted from the most to the least selected in the New Mexico (NM), Missouri (MO), Zurich (Zur), and Athens (Ath) traces. The inside plot is in log scale zooming into the distribution of the most popular servers.





CDF of the average distances (AVG) and standard deviations (STD) per run for the (a) Athens and the (b) New Mexico traces