NDN's Stateful Forwarding Plane in the Presence of Ground-Satellite Handovers

Sirapop Theeranantachai
Computer Science Department
University of California, Los Angeles
Los Angeles, United States
stheera@g.ucla.edu

Beichuan Zhang
Computer Science Department
University of Arizona
Tucson, United States
bzhang@cs.arizona.edu

Lixia Zhang
Computer Science Department
University of California, Los Angeles
Los Angeles, United States
lixia@cs.ucla.edu

***Abstract*—Low-Earth-Orbit (LEO) satellite constellations provide Internet connectivity across the globe, but their network design poses significant challenges due to the large number of satellites and their fast movement. Named Data Networking (NDN) can bring many benefits to LEO satellite networks such as data-centric security, scalable content distribution, and intelligent data plane. A key enabler to these benefits is NDN's stateful forwarding, which unfortunately can be disrupted by the frequent handovers between the satellites and the ground terminals. In this paper, we investigate the impacts of these handovers on NDN's packet delivery and propose effective mitigation mechanisms. Using a newly developed packet-level simulator and in-depth analysis of packet forwarding behavior during handovers, we show that consumer handovers and producer handovers both lead to temporary packet losses but for different causes. To mitigate these problems, we design a new Interest retransmission strategy to handle consumer handovers, and a new forwarding strategy to handle producer handovers. Our evaluation shows that these solutions are effective in reducing packet losses and delivery time. This work sheds insights on how NDN can work in the presence of frequent satellite-ground handovers to enable data-centric communication in this new network environment.**

***Index Terms*—Low-Earth Orbit, Satellite Network, Named Data Networking**

I. Introduction

Recent years have witnessed rapid deployments in large Low-Earth-Orbit (LEO) satellite constellations [1]. LEO satellites orbit the Earth at an altitude below 2,000 km, which drastically reduces the communication delay compared to traditional Geosynchronous-Equatorial-Orbit (GEO) satellites at 35,786 km. Since each LEO satellite covers a relatively small area on Earth, a large number of them are needed to ensure global coverage. Technological advances in rocket reuse, satellite miniaturization, and phased array antenna have made it possible to deploy thousands of LEO satellites economically. For example, SpaceX's Starlink constellation is comprised of more than 6,000 satellites at the time of this writing [1].

Large LEO satellite constellations have the potential of bringing Internet to the three billion unconnected people in the world [2], offering low-latency communication to delay-sensitive applications [3], and providing redundant paths for increased throughput and resiliency. However, they also pose significant challenges to network protocol designs due to the sheer size of the network and fast movement of the satellites. Interconnecting this vast number of satellites and fully exploiting the LEO constellation's potential remains an open research question.

In this paper we explore the direction of applying Named Data Networking (NDN) [4]–[6] to LEO satellite networking. Unlike the TCP/IP network architecture, whose primitive is to deliver packets to their destinations identified by IP addresses, NDN secures data by cryptographic signatures and retrieves secured data by their names. Applications running over NDN networks enjoy many benefits, including end-to-end data security, built-in multicast delivery and scalable content distribution, and mobility friendliness.

A key enabler of NDN's architectural benefits is its stateful forwarding plane, which maintains in-network states to facilitate efficient data retrieval. When a ground terminal switches its connection from one satellite to another, the forwarding states established along the previous path may be outdated, leading to packet delivery disruptions. Since LEO satellites' fast movement leads to frequent handovers, understanding their impacts on NDN is a critical step before designing an NDN-based LEO satellite network.

In this paper, we analyze the impacts of LEO satellite handovers on NDN's stateful forwarding, and propose new mechanisms to mitigate the problems. We categorize the ground-satellite handovers into two types: consumer handover and producer handover. For consumer handovers, we propose an Interest retransmission strategy that can quickly rebuild the forwarding state along new paths to minimize packet losses and delays. For producer handovers, we design a new NACK-based forwarding strategy to rebuild the forwarding state along new paths after routing changes. Our analysis and evaluation are backed by packet-level simulations using a newly developed simulator, ndnSIM-leo, that integrates the Starlink constellation setup with the NDN simulator, ndnSIM [7].

The rest of the paper is organized as follows. §II offers an overview of LEO satellite constellations, §III provides a brief introduction to NDN and its stateful forwarding, §IV describes ndnSIM-leo and the simulation setup, §V and §VI presents the main results and analysis in consumer handovers and producer handovers, respectively, §VII summarizes related work, and §VIII concludes the paper.

II. LEO Satellite Constellations

Modern LEO satellite constellations can have thousands of satellites, grouped into *shells* according to their altitude and inclination [8]. Each shell is typically configured as a Walker constellation [1], with evenly spaced satellites on multiple orbits. For instance, Starlink's first shell consists of 72 orbits, each with 22 satellites at an altitude of 550km and inclination of 53°, totalling 1,580 satellites. [9], [10].

Satellites communicate via RF-based ground-satellite links (GSL) and laser-based inter-satellite links (ISL).

Utilizing ISLs is essential to maximizing the potential of LEO constellations. They are typically formed between satellites moving in the same direction. While the use of ISLs is still experimental, the latest Starlink satellites are believed to have four ISL interfaces [9], [10], capable of connecting neighboring satellites in both the same and adjacent orbits to form a grid topology (Fig. 1).

As a satellite moves along its orbit, it will travel northeast-bound on one side of the earth, then southeast-bound on the other side of the earth, and repeat the pattern. With many orbits of the same inclination surrounding the earth, when viewed from any particular point, half of the satellites in the sky are moving in one direction, and the other half moving in the other direction (Fig. 2). ISLs in the grid topology only exist between satellites moving on the same direction. This leads to an interesting phenomenon: As illustrated in Fig. 3, satellites A, B, and C are close to each other in space, but the network hops between them can be vastly different. A and B are one-hop away due to a direct ISL, but A and C can be tens of hops apart because they are moving in different directions. This has an important implication to the end-to-end delay.

GSLs are subject to frequent handovers due to the fast movement of LEO satellites. For example, at the altitude of 550 km, a satellite travels at 27,320 km/h, and can remain in the view of a ground terminal for about 5 minutes. GSLs are also prone to packet loss and bandwidth fluctuation due to factors such as weather, obstruction, interference, and inaccurate tracking of moving satellites [11]. Given the large number of satellites, depending on the location of a ground terminal, it may have multiple satellites in its view at a given moment. Existing work assumes that a ground terminal may connect to the satellite that leads to the shortest end-to-end path towards the destination. In reality, however, user traffic often goes to many different destinations and it is infeasible to switch satellite for each different destination. A more practical strategy is to have a ground terminal connect to the nearest satellite and use it for traffic going to all destinations until another satellite becomes closer. A measurement study [12] shows evidence that Starlink's ground terminals may switch between satellites once every 15 seconds.

III. NDN and Its Stateful Forwarding

Every piece of data in an NDN network is identified by a unique name, and data delivery is carried out by two types of packets: *Interest* and *Data*. *Producers* publish data by signing the Data packets and make them available. *Consumers* retrieve data by sending Interest packets that carry the names of desired data. Network nodes forward Interests according to their names, and return the matching Data when found. NDN networks offer several key features typically absent in today's IP networks:

- Data-centric Security: NDN's data signature protects data integrity and provenance. Applications can make use of semantically meaningful data names and certificate names to define fine-grained security policies [13].
- Scalable Content Distribution: NDN's built-in multicast delivery and in-network caching make it ideal for large-scale content distribution without additional infrastructure. Multiple consumers can quickly fetch the data when the cache is stored nearby. It also reduces bandwidth consumption as data do not have to be fetched from the original producer.
- Intelligent Data Plane: NDN's *stateful* forwarding empowers its data plane with new capabilities, including loop detection, multi-path forwarding, built-in feedback loop to detect network anomalies, and various forwarding strategies to optimize network performance [14].
- Mobility Friendliness: Because data names are decoupled from topological locations and do not change when network connectivity changes, NDN makes it easy to support mobile and intermittent connectivity.
- Fast Loss Recovery: As returning data are cached along the path, if the data is lost before reaching the consumer, retransmitted interests can quickly retrieve the data from cache.

For the purpose of this paper, we will focus on explaining NDN's stateful forwarding and how it may be affected by LEO satellite handovers. Fig. 4 shows the overview of how packets are processed within an NDN node. When an Interest packet arrives, its name will be used to look up three tables in the order of Content Store (CS), Pending Interest Table (PIT), and Forwarding Information Base (FIB). The CS is essentially a cache of previously received data, and if a match is found, the Data will be returned. The PIT records the Interests that have been forwarded to the nexthop but haven't received the returning data yet. If a match is found, the incoming face of the Interest will be added to the PIT, so that when the data returns, it will be sent to that face, and the Interest does not need to be forwarded. The FIB is a name-based routing table, and the Interest will be forwarded to the nexthop supplied by the FIB and create a new entry in the PIT. If the node does not know where to forward the Interest or detects errors in the forwarding process, it can drop the Interest and optionally send a *NACK* back to notify the downstream node about the problem.

When a Data packet is received from the upstream, its name will be used to look up the PIT. A matching PIT record will have one or more faces associated, and the Data will be forwarded to all of them. The data will also be added to the CS.

To summarize, Interests are routed based on FIB and

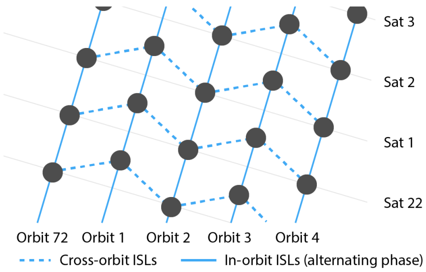

Fig. 1. The Grid Topology

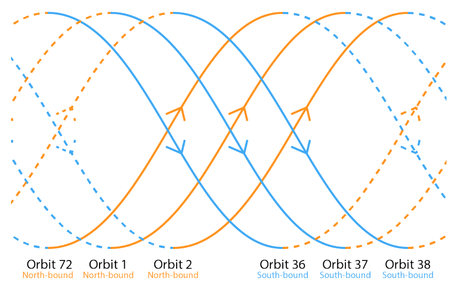

Fig. 2. Different orbit orientations

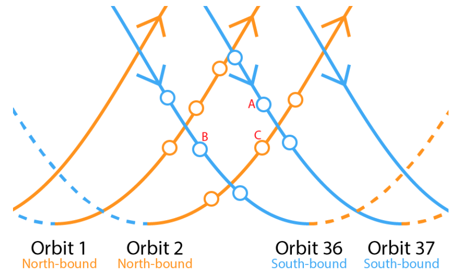

Fig. 3. Euclidean Distance vs. Network Hops

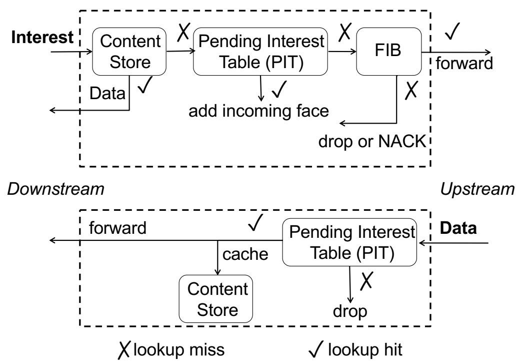

Fig. 4. Packet Processing in an NDN node

Interest names. As they traverse the network, Interests leave a "breadcrumb" in the PIT of each router they pass. Data packets are not *routed* but *forwarded* based on the PIT, taking the exact same path as the Interests but in the opposite direction. Though Interests usually are forwarded *towards* the producer, they do not have to reach the producer; they can bring back data from any node in the network as long as the Data name matches the Interest's name.

It is the PIT that makes NDN's forwarding stateful and also powerful. Traditional network's data plane relies solely on the control plane for correct forwarding decisions. In NDN, the data plane can detect problems and find solutions on its own without waiting for the control plane to converge. After an NDN node forwards an Interest, it will know the forwarding decision is correct when the Data returns, or there is a problem when NACK is received or a timeout occurs. In the latter case the node can retransmit the Interest to other faces attempting to get the Data. It does not have to worry about loops because the PIT will detect them when the same Interest comes back to the same node. Therefore NDN can monitor its forwarding decisions and quickly explore multiple paths if there is any problem. The specific mechanisms of exploring different paths and optimizing forwarding decisions are called *forwarding strategy*, which can be customized to specific network environments.

The PIT state is susceptible to link failures. As a link fails, in-flight data packets will reach the upstream node of the failed link but cannot get to the downstream node. The common solution is to have the consumer or the downstream nodes retransmit the Interest after timeouts, which will establish a new path to retrieve the data. NDN's in-network caching, even with a small cache size, can help with *fast loss recovery* as a retransmitted Interest can retrieve the lost data packet from some router cache in the network, as opposed to TCP/IP's end-to-end retransmission. However, it is unclear how well NDN's stateful forwarding will perform in a highly dynamic environment like LEO satellite networks, and given the frequent GSL handovers, some mitigation mechanisms are likely needed to minimize packet losses and additional delay, which are the focus of this work.

IV. The ndnSIM-leo Simulator

To examine NDN's packet forwarding behavior in LEO satellite networks, we develop a new simulator, ndnSIM-leo, by integrating two existing simulators: ndnSIM [7] for NDN functionality, and Hypatia [10] for LEO satellite constellations, and both are based on ns-3 [15]. Compared with a previous work [16], our simulator offers better integration with ndnSIM and better satellite support from Hypatia. Specifically, the differences are as follows:

1) The previous work assumes fixed, constant propagation delays for all ISL links throughout the simulation, with network snapshots taken before and after handovers. Our simulator calculates propagation delays in real time and updates routes every 100ms, similar to Hypatia [10].
2) Multiple consumer capabilities: The previous work models the GSL connectivity as permutations of point-to-point links, making it difficult to utilize NDN's data multicast and suffering a significant performance degradation in large networks. The new simulator implements a custom GSL link layer that adds flexibility for more realistic GSLs and improves overall performance.
3) Prolonged Scenarios: The new simulator can simulate a dynamic satellite network for longer period of time.

In our simulation model, each ground terminal has a single GSL face, which connects to the *nearest* satellite among all the satellites in the view of the terminal at that moment. As satellites move over time, a GSL handover occurs when a new satellite becomes the nearest. Each satellite has one GSL face and four ISL faces. The GSL is used to communicate with one or more ground terminals, and the ISLs connect to other satellites to form a grid topology (Fig. 1). We use Starlink's first shell of satellites in all our simulations, which has 72 orbits at 550 km altitude and 53 degree inclination, and 22 satellites per orbit. In this setting, there is no ISL handover because the neighboring satellites are within range all the time. The distance between the satellites, however, changes as they move. The simulator keeps track of these changes and updates the link delay accordingly for every packet transmission. As in Hypatia and other existing work, routing is done by computing shortest paths offline and installing them to satellites every 100 millisecond. Link delay is used as link cost in shortest-path computation. To focus on the forwarding behavior, we assume routing changes are done instantly.

We used several different ground locations for consumers and producers in the simulations. Without loss of generality, we will present a few representative cases and offer in-depth analysis. In each simulation, the consumer sends Interests at the rate of one packet per millisecond, and the producer replies with corresponding data. To precisely identify the behavior of packet transmission during the handover, each link is configured with enough bandwidth and no errors, so that the only reason for packets to be dropped or delayed is the satellite movement.

V. Consumer Handover

From NDN's point of view, a GSL handover can happen between the consumer and its satellites, or between the producer and its satellites. It turns out that these two cases have different impacts and require different solutions. We will examine consumer handover first, then the producer handover in the next section.

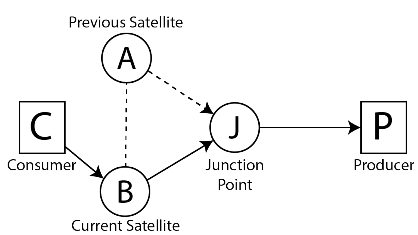

Fig. 5. Routing Paths after Consumer Handover

A. Impacts from Consumer Handover

Fig. 5 illustrates the case of consumer handover, where the consumer's ground terminal switches its GSL connection from the *previous* satellite A to the *current* satellite B. The problem caused by handovers is that in-flight data packets from the producer to the consumer will still go to A because of existing PIT states in the network. Moreover, in-flight interest from the consumer to the producer will also retrieve data back to the previous satellite. As a result, one round-trip time (RTT) of data is lost during a consumer handover.

Note that since Data packets are not *routed*, a consumer handover does not require any routing changes. Once the handover is complete, Interests sent from the current satellite will be able to retrieve data back to the consumer without any updates in the control plane, which is a nice mobility-friendly feature of NDN. Traditional networks would need complicated signaling to update the control plane in order to get packets to the consumer's new location.

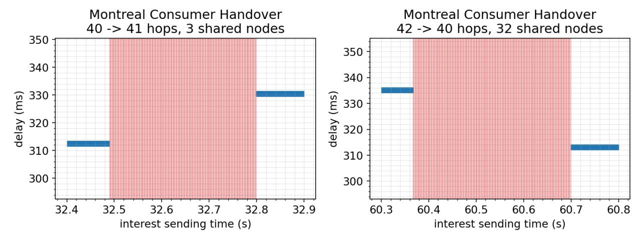

Fig. 6. Losses when the consumer do not re-transmit after a handover

Fig. 6 shows the packet loss in two different consumer handover scenarios. X-axis represents the time when the an interest was transmitted, the Y-axis represents the time that it takes to retrieve the data, and the red vertical lines represents data losses. The two blue flat lines reflect the RTTs before and after the handover, while the red vertical lines in between are the one-RTT worth of data loss caused by the handover. It is clear that a mitigation mechanism is needed to recover the lost data.

B. Data Recovery Solutions

In order to retrieve the data dropped during a consumer handover, the general solution is to have the consumer, after detecting the handover, to retransmit all *pending* Interests through the current satellite, which will establish new PIT states in the network and get data to the consumer. This retransmission is triggered by the handover event detected at the link layer. Without loss of generality, our analysis assumes this handover at the link layer is instantaneous and does not include any additional delay it might cause.

As proposed by previous work [16], the consumer has two options: (1) simply retransmit the pending interests, which will be forwarded towards the producer, or (2) retransmit the pending interests towards the previous satellite using a forwarding hint. While [16] proposed these solutions and simulated the performance improvement, we will provide in-depth analysis of exactly how they work, and under which conditions these solutions will be effective or not. Furthermore, we will propose a new solution that combines the two approaches to achieve minimize data retrieval delay and network bandwidth consumption.

C. Analysis of Interest Retransmission

One of NDN's architectural benefits is in-network caching. Not only does it make large-scale content distribution more efficient, it also speeds up packet recovery, e.g., after packet loss or link handover. The former usually requires lots of memory, but the latter only needs to cache data of the most recent RTT, for which most today's network devices have the memory capacity.

When pending interests are retransmitted after a consumer handover, they will retrieve data from somewhere in the middle of the path because of NDN's caching and data multicast. This is faster compared to retrieving from the producer as conventional network would do. Exactly how much the latency is depends on where the data is cached. Considering the paths before and after the handover, the junction point (J in Fig. 5) of the two paths is the most important to our analysis. Intuitively, under the nearest-satellite strategy, the previous satellite and the current satellite are at most two hops away, thus their paths to the producer should have a significant overlap. However, as Fig.3 shows, if two satellites are moving on different directions, even if they are close to each other, the number of network hops between them can be large, depending on the latitude of the consumer. Therefore, the benefit of in-network caching is highly dependable on satellite movement. Here we will analyze two cases to explain the detailed behavior of data retrieval.

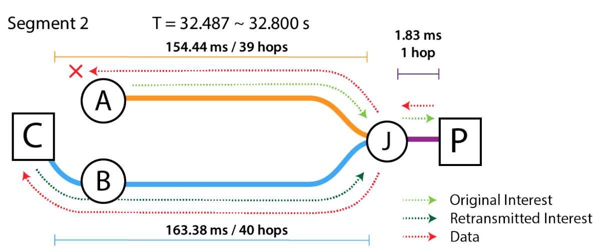

Fig. 7. Segment 2 of Fig. 8. The re-transmitted interest is fetching from the cache stored at the junction

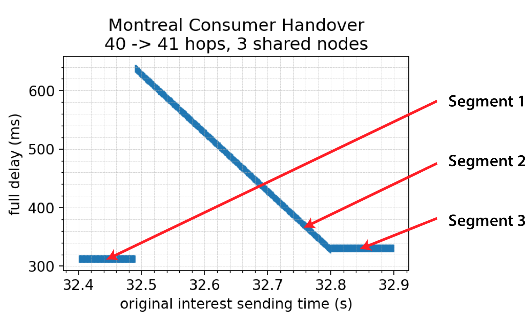

Fig. 8. Data retrieval time of the *minimal path overlap* case study

1) Minimal path overlap: Fig. 7 illustrates a consumer handover with minimal path overlap. The handover happens at time = 32.8s, when the consumer switches from a satellite (on orbit-52) moving north-bound to a satellite (on orbit-66) moving south-bound. The junction point of the old and new paths is at the last hop satellite near the producer.

Figure 8 plots the data retrieval time for all the Interests. Here, data retrieval time is defined as the time between the first transmission of an Interest and the receipt of the corresponding data. The result curve has three distinct segments. Segments 1 and 3 are straightforward: they are the path RTTs before and after the handover, respectively. Segment 2 are the packets affected by the handover, and it lasts one RTT of the old path (in orange).

When a handover happens, data retrieved by existing Interests will continue to follow the obsolete PIT entries to go to the previous satellite. They will have to be retrieved by the retransmitted Interests later, thus higher delay. The very first data packet impacted by handover was dropped by the previous satellite at the time of handover, 32.8s, therefore we can calculate its corresponding Interest was sent at $32.8 - 0.312 = 32.49$s, where 0.312s is the RTT of the orange path. This matches the result in the figure. This Interest will be retransmitted right after the handover (t = 32.8s), and it takes $2 * 163.38$ms to retrieve the data from the junction point, which has already cached the data. Here 163.38ms is the one-way delay to the junction point satellite. Thus the total data retrieval time for this data packet is $312\text{ms} + 163.38\text{ms} * 2$. That is the y-axis value of the starting point (and the highest point) of segment 2.

This analysis applies to later packets impacted by the handover as well. Their data retrieval times decrease because the original interests were paced at one per millisecond, but the retransmitted interests were sent out all at once. Therefore all data arrive at the consumer at the same moment ($t = 32.8\text{s} + 163.38\text{ms} * 2$), but their original Interest transmissions were spread out evenly at the rate of one packet per millisecond. This explains the straight decreasing line of segment 2. In practice, if the consumer decides to pace the Interest retransmission, the data retrieval time will be affected by when an Interest is retransmitted; the later the retransmission, the longer the data retrieval time will be.

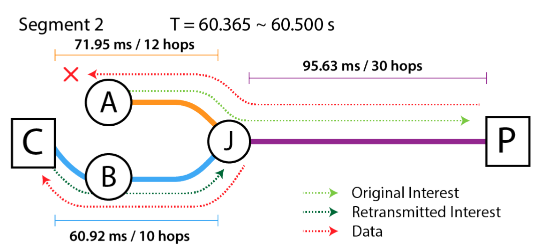

Fig. 9. Segment 2 of figure 10. The re-transmitted interest is fetching from the cache stored at the junction

2) Significant path overlap: Fig. 9 is a case with significant path overlap. The handover happens at time = 60.7s, when the consumer switches from a satellite (on orbit-66) moving north-bound to a satellite (on orbit-67) moving north-bound as well. The two paths merge at a junction point about $1/4$ into the path, and overlap for the remaining $3/4$ of the path.

Fig. 10 shows the data retrieval time around this handover. It can be explained by the same analysis based on NDN's forwarding mechanism. It has 5 segments, more complicated than the previous example, because the significant path overlap amplifies some details of in-network caching.

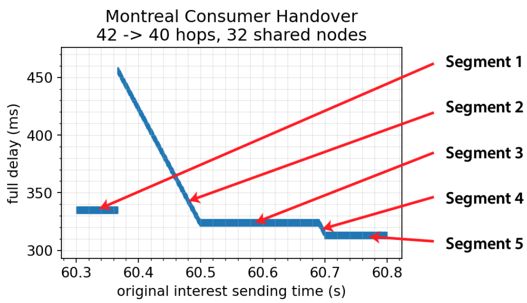

Fig. 10. Data retrieval time of the *significant path overlap* case study

Segments 1 and 5 are straightforward: they are packets not affected by the handover, thus taking the RTT of the path before and after the handover, respectively.

Segment 2 can be explained the same way as we explained segment 2 of the first example. These are the data packets that, at the moment of handover, have already passed the junction point on their way to the previous satellite along the orange path. Their retransmitted Interests will recover the data from the junction point. The very first data packet dropped will experience the highest data retrieval time, which is the orange RTT (to the producer) + blue RTT (to the junction point). Here caching at the junction point is able to reduce the data retrieval time significantly due to the path overlap.

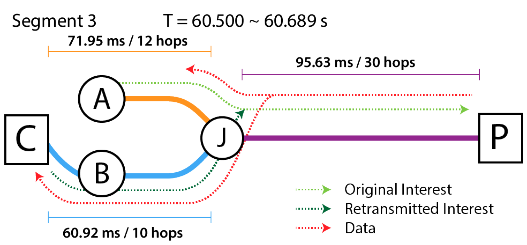

Fig. 11. Segment 3 of figure 10. The re-transmitted interest is now aggregated at the junction, waiting for the data

Segment 3 are the packets that by the time their retransmitted Interests arrive at the junction point, the corresponding data (which were retrieved by the original Interest) are still on their way to the junction point. In this case, the retransmitted Interests will be aggregated at the junction point and not be forwarded. They simply sit in the PIT and wait for the data. When the data arrive, they will be forwarded along both the orange and blue paths, which is exactly how NDN's multicast works. Therefore the data retrieval time of these packets is the one-way delay of the old path to the producer plus the one-way delay of the new path to the consumer. This behavior happens in the first example too, but because the junction point in the first example was so close to the producer, the effect was not discernible in the figure. Segments 2 results from NDN's in-network caching, and segment 3 results from NDN's multicast capability.

Segment 4 are the packets whose retransmitted Interests

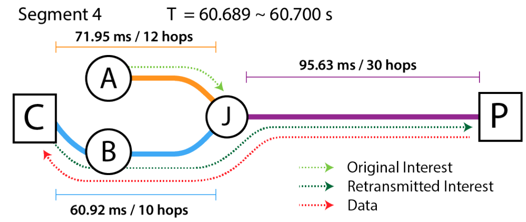

Fig. 12. Segment 4 of figure 10. The re-transmitted interest arrives at the junction before the original interest does.

arrive at the junction point before the original interests get there. This is possible when the current satellite is closer to the junction point than the previous satellite. In this particular example, the difference is about 10ms. All data packets will be retrieved by the retransmitted Interests directly from the producer, which takes the RTT of the new path. Since the original interests were sent out 10ms to 0ms before the handover, it will be added to the total data retrieval time. Therefore we see a small decrease in the data retrieval time lasting for about 10ms for segment 4.

In summary, interest retransmission can recovery data that would be dropped after a consumer handover. NDN's in-network caching and native multicast speed up such data recovery, and the extent of such speedup depends on the junction point of the old and new paths.

D. Analysis of Retransmission with Forwarding Hint

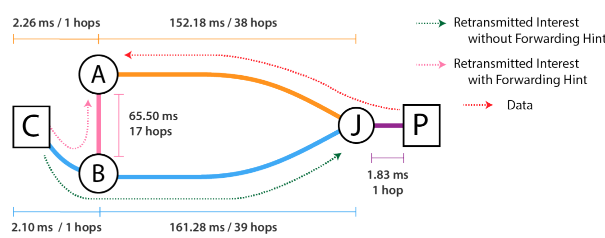

Fig. 13. Consumer handover with Forwarding Hint: small path overlap

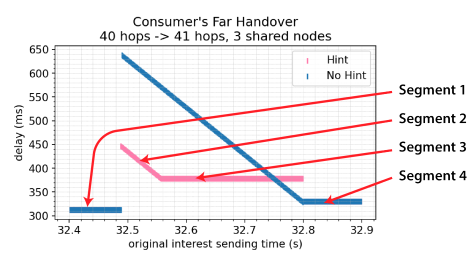

Fig. 14. Data retrieval time of small overlapping path, with and without using forwarding hint

Knowing that after the consumer handover, data will continue to follow the existing PIT entries and arrive at the previous satellite anyway, why not send Interests to the previous satellite to retrieve these data? Teng et. al. [16] proposed to use forwarding hint to do exactly the same. Forwarding hint is an

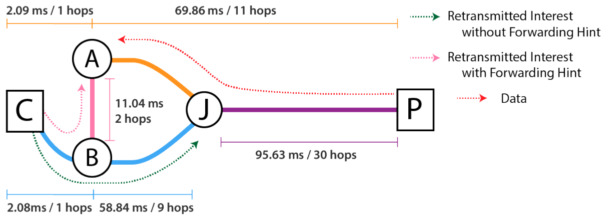

Fig. 15. Consumer handover with Forwarding Hint: large path overlap

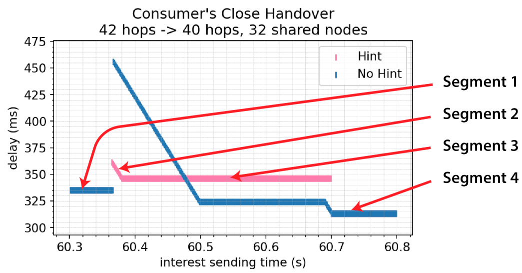

Fig. 16. Data retrieval time of large overlapping path, with and without using forwarding hint

additional name (i.e., delegate name) attached to the Interest packet. NDN nodes will forward these Interests towards the delegate name first. If the node that owns the delegate name does not have the data, the Interest will then be forwarded towards the producer. Note that in each step of the forwarding, NDN node will always want to match the data and interest based on the interest/data name, and the delegate name is only used in looking up the routing table. Forwarding hint gives the consumer limited control in directing the interest in forwarding. In this case of consumer handover, the intention is to have the retransmitted interests go to the previous satellite first, and if data is not there yet, continue towards the producer.

Following our analysis of the previous examples, we can see that the most delay overhead in data retrieval occurs in segment 2 (Figures 8 and 10), where data is retrieved from the junction point without forwarding hint. Using forwarding hint may be able to reduce this delay if the previous satellite is closer than the junction point to the current satellite. This is indeed the case. In the first example, the previous satellite is 17 hops away from the current satellite 13, while the junction point is 39 hops away. In the second example, the previous satellite is only 2 hops away from the current satellite, while the junction point is 9 hops away 15.

Figures 14 and 16 compare the data retrieval time with and without forwarding hint for our two handover examples. The forwarding hint result curves are similar. They both have 4 segments, where segment 1 and segment 4 reflects the RTT before and after the handover. Segment 2 is the data that are already cached at the previous satellite. The duration of this segment is the latency of the path between previous and current satellites. Segment 3 are the data packets that are still in flight when retransmitted interests arrive at the previous satellite. Similar to our previous analysis, these interests will be aggregated at the previous satellite and wait for the data to arrive. These data's retrieval time is the RTT of the path from the current satellite to the previous satellite and go on to the producer.

These examples show that using forwarding hint can improve the worst-case data retrieval time significantly. In general, as long as the previous satellite is closer than the junction point to the current satellite, forwarding hint will be able to improve the worst-case data retrieval time. The average case, however, is less clear. Consider segment 3 in Fig. 16, using forwarding hint causes longer data retrieval time for some packets due to the detour the Interests take via the previous satellite. The overall result depends on how the delays of the three path segments compare to each other: from previous satellite to junction point (A-J), from current satellite to junction point (B-J), and from current satellite to previous satellite (B-A).

E. Optimal Interest Retransmission

Based on the analysis, we propose a new retransmission strategy for optimal data recovery time. Since the satellite movement is predicable within a certain time frame, the consumer can predict the handover and compute the paths before and after the handover. Using this information, the consumer can retransmit some interests using forwarding hints, and the rest directly to the producer. Taking Figures 14 and 16 as examples, the consumer can retransmit the Interests using forwarding hints until the intersection point of the blue and pink curves, then the remaining interests directly to the producer. That turning point happens around 60.48s in Fig. 16 and around 32.75s in Fig. 14. This way, the consumer can achieve the best of both worlds, resulting in optimal data retrieval time.

In addition to delay, using forwarding hint may also reduce the total traffic in the network. In Fig. 5, the bandwidth along the path A-J has been committed with or without forwarding hint. Path A-B will be used with forwarding hint, and path B-J will be used without forwarding hint. If A-B has fewer hops than B-J, which is quite common in our simulations, using forwarding hint will reduce the total traffic in the network. Again, assuming the consumer can calculate the paths before and after the handover, it can take the total network resource consumption into consideration when deciding when to use forwarding hint.

In summary, the forwarding hint allows the consumer to influence where Interests are forwarded. Given the predictable nature of satellite movement, consumers may choose to calculate the breakpoint for when to use the forwarding hint – trading computational resource to achieve the optimal data retrieval time

VI. Producer Handover

Producer handover differs from consumer handover in NDN. On one hand, the handover only disrupts the PIT states at the producer, affecting a relatively small number of data packets. On the other hand, since the producer has changed its GSL, the system has to update routing so that future interests can be routed towards the producer. As we will show later in this section, while new interests sent out after the handover will be routed correctly, some of the interests sent out before the handover may run into a particular forwarding problem after the routing change. To focus on the forwarding plane behavior, we assume the routing updates are done instantaneously, i.e., after the handover, routing tables at all nodes are updated immediately with the new nexthops towards the producer.

A. Impacts from Producer Handover

1) Packet loss at the producer: At the time of producer handover, for data packets that have already reached the previous satellite, they will be forwarded according to existing PIT states in the network and not be affected by the handover or the routing change. For a small number of pending Interests, however, their corresponding data have not been able to cross the previous GSL and reach the satellite. These data will be lost because the producer cannot send them to the new satellite, which doesn't have any PIT entry for these data (Fig. 17).

Fig. 18 is one simulation result to illustrate this type of packet loss, represented by the red vertical lines. In this case, the previous satellite and the new satellite are on adjacent orbitals, making the data retrieval delays very similar before and after the handover.

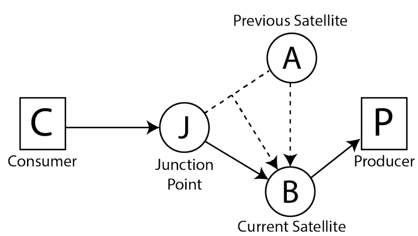

Fig. 17. Routing Paths after Producer handover

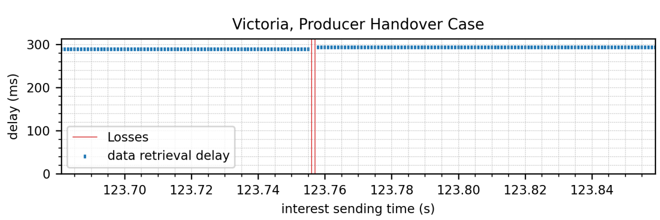

Fig. 18. Data retrieval time, without the new strategy, handover to a nearby satellite.

2) Packet loss in the network: At the time of the producer handover, many Interests are still making their way from the consumer to the producer. After the handover and the routing change, these Interests will be forwarded according to the new routing table. Depending on which satellite along the path an Interest is at, its new nexthop has three possibilities (Fig. 19):

- Unchanged: The new next-hop is the same as the previous next-hop.

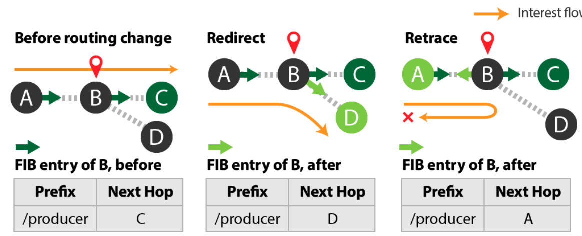

Fig. 19. Routing Changes

- Redirect: The new next-hop is not the same as the previous next-hop, but also not the incoming face of the Interest.
- Retrace: The new next-hop is the same as the incoming face of the Interest.

In the case of Unchanged, Interests will be forwarded exactly the same without any impact. In the case of Redirect, Interests will still be forwarded towards the producer, but they may take a detour and experience longer delay. For example, Fig. 20 shows the delay and loss when the producer switches to a new satellite that is 35 hops away from the previous satellite. Similar to the previous case, a small number of packet are lost at the producer. Furthermore, the data retrieval time jumps immediately after the handover. This is because those Interests were forwarded towards the previous satellite before the handover, but need to change direction to go to the new satellite after the handover, resulting in higher data retrieval time.

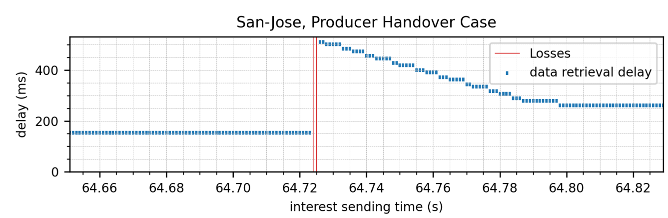

Fig. 20. Data retrieval time, without new strategy, handover to a satellite far away

Retrace, however, would not be able to retrieve data under NDN's default settings and forwarding strategy. When an Interest is sent back to where it came from, i.e., the incoming face, it will be detected by the receiver as a looping Interest and be dropped. This leads to data retrieval failure. Fig. 21 illustrates this problem with another producer handover case. In addition to the small number of loss at the producer (around 192.68ms), there are two large sequences of packet loss due to retrace. Between the two sequences, there are some packets that get delivered with increased data retrieval time. This can be explained using Fig. 22. For Interests at RD, the new next-hop points to Z, which is a Redirect. For Interest at RT, the new next-hop points to X, which is a Retrace. Thus, depending on the routing change at each node, Redirect and Retrace nodes may be mixed on the path. The former will lead to longer delay but no loss, but the latter will cause data loss.

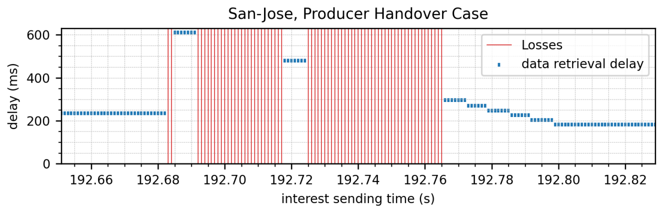

Fig. 21. Data retrieval delay, no new strategy, Retrace causes packet loss.

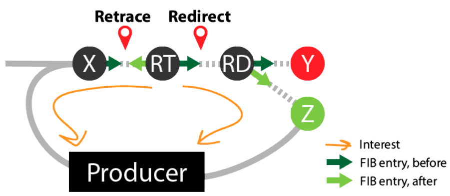

Fig. 22. Retrace situation: If interest is due at Node RD, it's faster to go through Z. If interest is due at Node RT, it's faster to go through X

B. A New Forwarding Strategy

The solution to producer handover has two pieces. First, to recover data loss caused by PIT state disruption at the producer, the previous satellite need to retransmit its pending Interests upon the handover. Second, to recover data loss caused by NDN's loop detection, the forwarding strategy needs to be revised to allow Interests to retrace after a routing change. Fortunately, NDN supports the customization of forwarding behavior via different forwarding strategies. To this end, we propose the following forwarding strategy for LEO satellite networks to support Interest retrace:

1) When an Interest is to be transmitted but the next-hop is the incoming face of the Interest (i.e., Retrace), the NDN node (including the previous satellite of the handover) should transmit a NACK on that face, drop the Interest, and remove it from the PIT. Otherwise (i.e., Unchanged or Redirect), the Interest is sent out to the next-hop.
2) Upon receiving a NACK, an NDN node should retransmit the corresponding Interest to its best next-hop following step 1.

In this strategy the upstream node detects the potential loop because the Interest would be sent to where it comes from. Instead of sending the Interest, the node will send a NACK so that the downstream node has a chance to send the Interest to a different face. Therefore Interests in the network will be able to retrace in the form of NACK until they arrive at a node that can redirect the Interests to the producer. This strategy works for both the retransmitted interests by the previous satellite and the in-flight interests by other satellites on the path. We plan to incorporate this strategy into NDN's default forwarding strategy since it does not conflict with the default strategy and will only enhance it.

Figures 23 and 24 illustrate the behavior of this strategy. In both cases the Interests will retrace to node A, which is the first node that can redirect the Interests to the producer. Returning data will follow the new PIT entries to reach A then the consumer. They will not to go B because the NACK has cleared the PIT states between A and B. The difference between the two figures is that the NACK is triggered by an arriving Interest at node B in Fig. 23, but a producer handover at the previous satellite in Fig. 24. Both cases are handled by the same strategy.

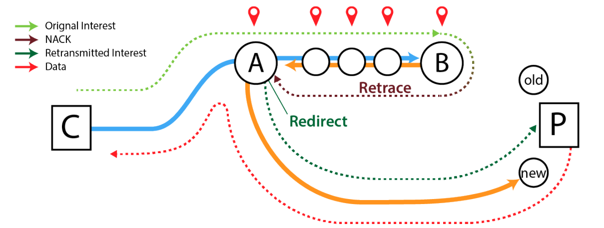

Fig. 23. A received interest triggers NACK

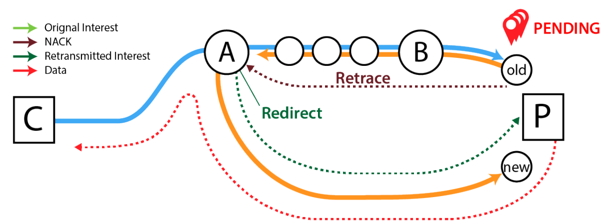

Fig. 24. Handover event triggers NACK from PIT records

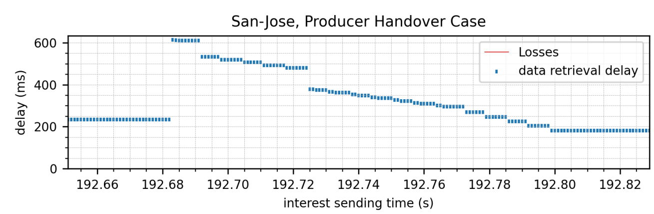

Fig. 25. Custom forwarding strategy with NACK-triggered retransmission.

We implemented this forwarding strategy and evaluated its effectiveness using the same case in Fig. 21. The result is presented in Fig. 25. With the interest retransmission by the previous satellite and the Interest retrace strategy in the network, all data have been successfully retrieved without a loss.

In summary, producer handovers are different from consumer handovers in that it will cause routing changes. It can lead to small packet loss at the producer GSL but more loss on other links due to the lack of support for interest retrace. The solution of Interest retransmission and a new forwarding strategy is able to fix these issues, demonstrating the flexibility and capability of NDN's forwarding plane.

VII. Related Work

The mobile handover problem has been studied extensively in conventional network settings, such as cellular networks and WiFi networks. When a mobile node hands off from an old base station to a new base station, the network must ensure ongoing communication can continue with minimal disruption. The main issue is how to update the *control plane* so that packets will be routed to the new base station. Although the design and implementation of various solution tend to be specific to the link layer and/or the media technology, their basic approach involves signaling between the two base stations and switching centers to set up a new path. Note that the old and new base stations tend to be physically close to each other,

On the other hand, a ground terminal's handover in a LEO network can have the old and new satellites far apart topologically even if they are close in physical space. This makes the existing handover solutions no longer applicable. Furthermore, the handovers in an NDN network make the most impact at the stateful *data plane*, and this work develops data-plane solutions to address data-plane issues.

Liang et al. [16] conducted the first investigation into mitigating the impact of *consumer handovers* on NDN's stateful forwarding plane. Their work utilized the Interest retransmission approach proposed in [17] to minimize the impact from consumer handovers. Liang et al. noticed that simply retransmitting Interests upon handover is not always effective due to potential mismatches in the Interest forwarding paths before and after the handover, resulting in under-utilization of NDN's in-network caching. To address this limitation, they proposed an alternative solution. Upon detecting a consumer handover, an NDN forwarder retransmits pending Interests to the previous satellite using forwarding hint. Their evaluation, however, is inconclusive in that using forwarding hint on average does not yield an improved data retrieval time. The performance reported in [16] are shown as bar graphs of average delay and throughput, and are not fully explained due to the limitation of the simulator.

Our work extended Liang's work in several major ways. First, we studied both consumer and producer handovers, the latter have not been studied before. Second, we provide an in-depth analysis of exactly what happens at the data plane after a handover at packet movement level, to fully understand the interaction of NDN's stateful forwarding and LEO satellite handovers. Third, we propose a new consumer handover solution that can optimize the data retrieval time and bandwidth usage. Finally, we devised a new forwarding strategy to minimize packet losses caused by producer handovers.

VIII. Conclusion and Future Work

NDN networking over LEO satellites is an interesting design space to explore for future space-based networks. However, the high dynamics in LEO connectivity raises a question on the viability of NDN's stateful forwarding plane. In this paper, we provide a systematic and detailed analysis on the impact of frequent satellite handovers on NDN's stateful forwarding, and develop effective solutions that can mitigate the negative impacts of consumer and producer handovers.

The solution development in this work makes some important assumptions, including instant routing changes at all satellite nodes and absence of ISL failures. In the future work, we plan to relax these assumptions and investigate how to enhance the current solution to make it work in more realistic environments with routing propagation delays and ISL failures. In particular, we will look into how link failures may interact with NDN's stateful forwarding, how the stateful forwarding may facilitate LEO network routing, and how NDN's multi-path forwarding capability may help support multi-path forwarding in LEO constellations. The earlier work by Yi et. al [18] showed that NDN's stateful forwarding could help reduce both routing overhead and convergence delay, it would be interesting to see whether the same holds true in LEO network setting.

Acknowledgment

This work is partially supported by Cisco Systems. The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the official views or position of the sponsors.

References

[1] C. Westphal, L. Han, and R. Li, "Leo satellite networking relaunched: Survey and current research challenges," 2023.

[2] D. York and G. Huston, "Leo satellites for internet access," *The Internet Protocol Journal*, no. 2, 2023. [Online]. Available: https://ipj.dreamhosters.com/wp-content/uploads/2023/10/262-ipj.pdf

[3] M. Handley, "Delay is not an option: Low latency routing in space," ser. HotNets '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 85–91. [Online]. Available: https://doi.org/10.1145/3286062.3286075

[4] V. Jacobson, D. K. Smetters, J. Thronton, M. F. Plass, N. H. Briggs, and R. Braynard, "Network Named Content," *CoNEXT*, 2009.

[5] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, P. Crowley, C. Papadopoulos, L. Wang, B. Zhang *et al.*, "Named data networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 66–73, 2014.

[6] A. Afanasyev, T. Refaei, L. Wang, and L. Zhang, "A Brief Introduction to Named Data Networking," in *Proc. of IEEE MILCOM*, Oct. 2018.

[7] S. Mastorakis, A. Afanasyev, and L. Zhang, "On the Evolution of NdnSIM: An Open-Source Simulator for NDN Experimentation," *Computer Communication Review*, vol. 47, no. 3, 2017. [Online]. Available: https://doi.org/10.1145/3138808.3138812

[8] R. Morales-Ferre, E. S. Lohan, G. Falco, and E. Falletti, "Gdop-based analysis of suitability of leo constellations for future satellite-based positioning," in *2020 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*, 2020, pp. 147–152.

[9] FCC, "Request for orbital deployment and operating authority for the spacex gen2 ngso satellite system," Federal Communications Commission, Tech. Rep. FCC 22-91, 2022. [Online]. Available: https://docs.fcc.gov/public/attachments/FCC-22-91A1.pdf

[10] S. Kassing, D. Bhattacherjee, A. B. Águas, J. E. Saethre, and A. Singla, "Exploring the "internet from space" with hypatia," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 214–229. [Online]. Available: https://doi.org/10.1145/3419394.3423635

[11] S. Ma, Y. C. Chou, H. Zhao, L. Chen, X. Ma, and J. Liu, "Network characteristics of leo satellite constellations: A starlink-based measurement from end users," 2022.

[12] G. Huston, "Using leos and geos," 2022. [Online]. Available: https://www.potaroo.net/ispcol/2022-04/leogeo.html

[13] Y. Yu, A. Afanasyev, D. Clark, k. claffy, V. Jacobson, and L. Zhang, "Schematizing trust in named data networking," in *Proceedings of the 2nd ACM Conference on Information-Centric Networking*, ser. ACM-ICN '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 177–186. [Online]. Available: https://doi.org/10.1145/2810156.2810170
[14] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang, "A case for stateful forwarding plane," *Computer Communications*, 2013.
[15] The University of Washington NS-3 Consortium, "ns-3 simulator." [Online]. Available: https://www.nsnam.org/
[16] T. Liang, Z. Xia, G. Tang, Y. Zhang, and B. Zhang, "Ndn in large leo satellite constellations: A case of consumer mobility support," in *Proceedings of the 8th ACM Conference on Information-Centric Networking*, ser. ICN '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 1–12. [Online]. Available: https://doi.org/10.1145/3460417.3482970
[17] G. Carofiglio, L. Muscariello, M. Papalini, N. Rozhnova, and X. Zeng, "Leveraging icn in-network control for loss detection and recovery in wireless mobile networks," in *Proceedings of the 3rd ACM Conference on Information-Centric Networking*, ser. ACM-ICN '16. Association for Computing Machinery, 2016. [Online]. Available: https://doi.org/10.1145/2984356.2984361
[18] C. Yi, J. Abraham, A. Afanasyev, L. Wang, B. Zhang, and L. Zhang, "On the role of routing in Named Data Networking," in *Proceedings of ACM SIGCOMM ICN Conference*, 2014.