May 25, 2018

STRATEGIES FOR LOOP PREVENTION FORWARDING LOGIC IN HYBRID INFORMATION CENTRIC NETWORKING

Scott Wainner

Luca Muscariello

Giovanna Carofiglio

STRATEGIES FOR LOOP PREVENTION FORWARDING LOGIC IN HYBRID INFORMATION CENTRIC NETWORKING

AUTHORS:
Scott Wainner
Luca Muscariello
Giovanna Carofiglio

ABSTRACT

Techniques are described herein that provide a forwarding strategy applicable to Information Centric Networking (ICN) networks and hybrid ICN (hICN) networks. In such networks, Interest requests are likely to be hair-pinned back to a Forwarder that is serving a dynamic cache system. When the dynamic cache system has no representation of the named data object, the Forwarder must recycle the Interest request back into the network. However, the Forwarding logic must avoid hair-pinning the Interest back to the same unpopulated cache.

DETAILED DESCRIPTION

Information Centric Networking (ICN) networks use named object identifiers to identify, locate, and retrieve data objects (e.g., code, data, process state, etc.) in accordance with a packet forwarding paradigm. The ICN paradigm uses an Interest packet sent by Consumers to express a request to obtain a named data object from a Producer via the network. The named data object is represented as a hierarchical path in much the same way that a Uniform Resource Locater (URI) is represented. A key difference between ICN data transfer and classic Internet Protocol (IP) data transfer is the absence of an end-to-end session state for ICN. The objective of the ICN network Forwarders is to direct the Interest request towards the nearest representation of the data by performing a longest path match. In many cases, there is a single source or Producer of the named data object; therefore, the network Forwarders have an obvious strategy to direct the Interest requests to the Producer.

The transition to ICN is complicated by the fact that all adjacent nodes must be capable of executing the ICN forwarding logic. This is often not the case in a real network. For this reason, a hybrid ICN (hICN) model was developed where named data objects could be represented as IP version 6 (IPv6) addresses. The same principles apply with regard to the Forwarding logic in ICN and hICN. An hICN network will forward an Interest in an

IPv6 named data object (represented by a /128 address) predicated on IPv6 routing. If there is a single Producer of the named data object, the network Forwarders may simply follow the most specific prefix match that leads to the object represented by the /128 address. In the case of a single Producer, there is a loop-free best path from any Consumer to the Producer.

The use of caching changes this paradigm by making multiple representations of the named data object. Effectively, there are multiple paths to the same object. Each representation of the named data object is advertised to the network such that the Forwarding nodes can choose the best path. If the named data object is present at every cache, the Forwarding nodes can successfully route the Interest request to a nearby named data object. This method is used by Domain Name Systems (DNS) that advertise an IP Anycast address that is flooded across the network. The routing infrastructure routes DNS requests to the nearest DNS server using the longest prefix match until the DNS request reaches a DNS server. The DNS server answers the client, redirects the client, or proxies the DNS request on behalf of the client to a different DNS server (i.e., IP address). A similar flooding approach is used for caching in hICN networks. However, the number of named data objects could easily consume all the memory of the routers if each named data object is represented by a /128 address. For this reason, a class of named data objects may be represented by a prefix (e.g., /64) that is flooded across the network. For example, a video asset or video channel might be represented by a /96 prefix. The hICN cache/forwarders would all present a /96 path to the routed network for a given set of named data objects. The routed network will still route /128 Interest requests on the longest prefix match to the nearest cache where a response is likely to occur.

One of the value propositions of dynamic caching is the ability to cache a subset of the entire library (most frequently accessed content) while enabling dynamic cache fill for requested content that has not already been cached. This poses a problem for hICN because an Interest request arrives at a cache where the named data object presumably exists. If the cache contains the named data object, a Response is provided and the Interest request is cleared from the Forwarders during the return of the Response. However, if the cache has not already dynamically cached the named data object, the cache must create an Interest request for the same object. If the adjacent router is also a Forwarder, the router/forwarder tables can be updated noting that the cache/forwarder does not have the object and the next best path to the prefix should be used on an alternate interface. In this example, the routing plane has not changed. If the adjacent router is not a Forwarder, that router's Forwarding Information Base (FIB) still presumes the cache should have the named data object because the router's longest prefix match for the set of named data objects still points to the cache/forwarder. An Interest request for the same named data object by the cache/forwarder will loop back to itself. This is effectively a micro-loop between the cache/forwarder and an adjacent non-Forwarder router.

The cache/forwarder cannot dynamically cache fill the named data object even though the objects of that set may exist somewhere else in the network (e.g., the Origin/Producer). One purpose of hICN is to facilitate a routed path between Forwarders where the Forwarders do not have to be adjacent, as the Forwarders may be reached via routers that simply perform the longest prefix match of IPv6 prefixes.

Accordingly, the techniques presented herein locate the nearest cache that may contain the named data object. A cache/forwarder that does not have the named data object may populate that cache by forwarding the Interest towards the next nearest representation of the named data object. Since the routed network's FIB has not changed, the Interest will likely return to the very same cache because it is advertising a representation of the named data object under the premise that it could have been cached at that location. This same problem exists as both micro-loops and macro-loops. Thus, the techniques presented herein address the problem of hICN Interest requests in named data objects where hair-pinning Interest requests is avoided by creating a set of conditional forwarding strategies.

Figures 1-4 below illustrate a first scenario involving the micro-loop across a direct interface between a router/forwarder and a cache/forwarder.

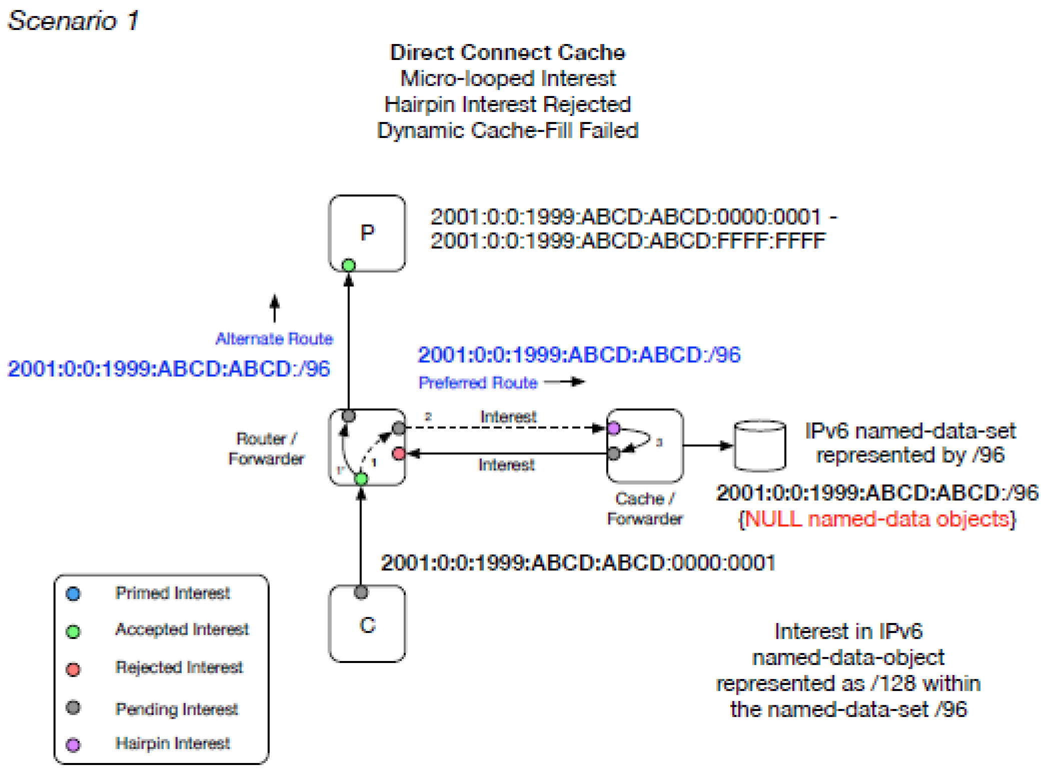

Scenario 1
Direct Connect Cache
Micro-looped Interest
Hairpin Interest Rejected
Dynamic Cache-Fill Failed
P
2001:0:0:1999:ABCD:ABCD:0000:0001 -
2001:0:0:1999:ABCD:ABCD:FFFF:FFFF
Alternate Route
2001:0:0:1999:ABCD:ABCD:/96
2001:0:0:1999:ABCD:ABCD:/96
Preferred Route
Router / Forwarder
Interest
Interest
Cache / Forwarder
IPv6 named-data-set represented by /96
2001:0:0:1999:ABCD:ABCD:/96
{NULL named-data objects}
2001:0:0:1999:ABCD:ABCD:0000:0001
C
Primed Interest
Accepted Interest
Rejected Interest
Pending Interest
Hairpin Interest
Interest in IPv6 named-data-object represented as /128 within the named-data-set /96

Figure 1

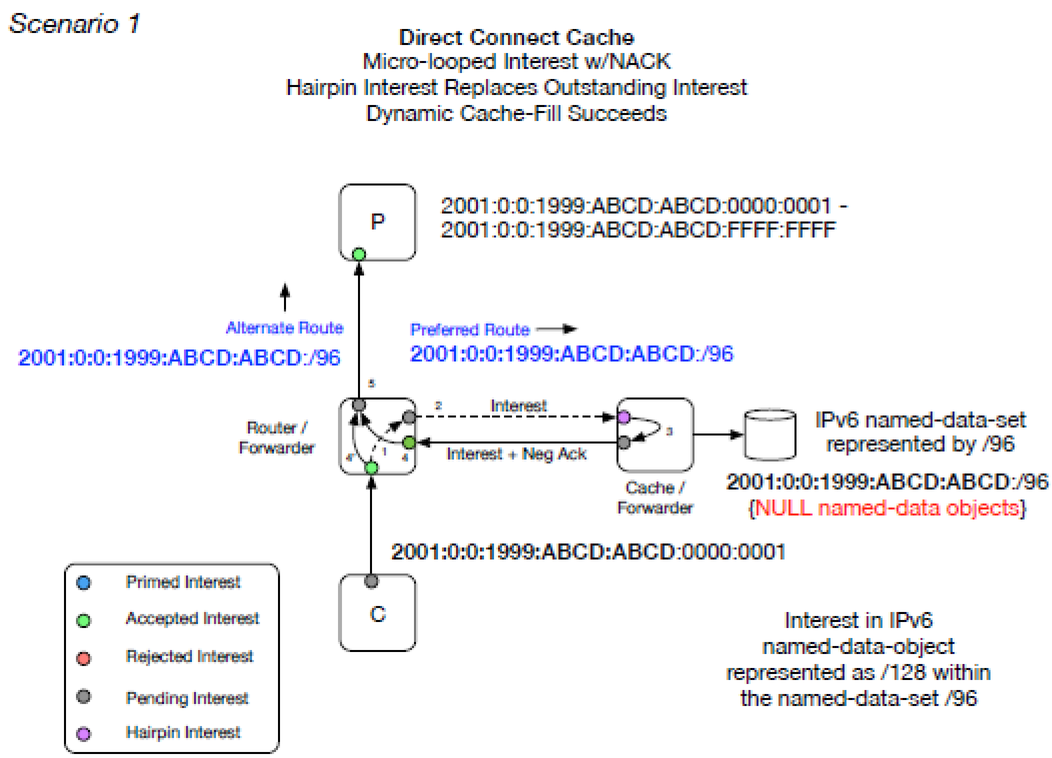
Scenario 1
Direct Connect Cache
Micro-looped Interest w/NACK
Hairpin Interest Replaces Outstanding Interest
Dynamic Cache-Fill Succeeds
P
2001:0:0:1999:ABCD:ABCD:0000:0001 -
2001:0:0:1999:ABCD:ABCD:FFFF:FFFF
Alternate Route
2001:0:0:1999:ABCD:ABCD:/96
Preferred Route
2001:0:0:1999:ABCD:ABCD:/96
Router / Forwarder
Interest
Interest + Neg Ack
Cache / Forwarder
IPv6 named-data-set represented by /96
2001:0:0:1999:ABCD:ABCD:/96
{NULL named-data objects}
2001:0:0:1999:ABCD:ABCD:0000:0001
C
Primed Interest
Accepted Interest
Rejected Interest
Pending Interest
Hairpin Interest
Interest in IPv6 named-data-object represented as /128 within the named-data-set /96

Figure 2

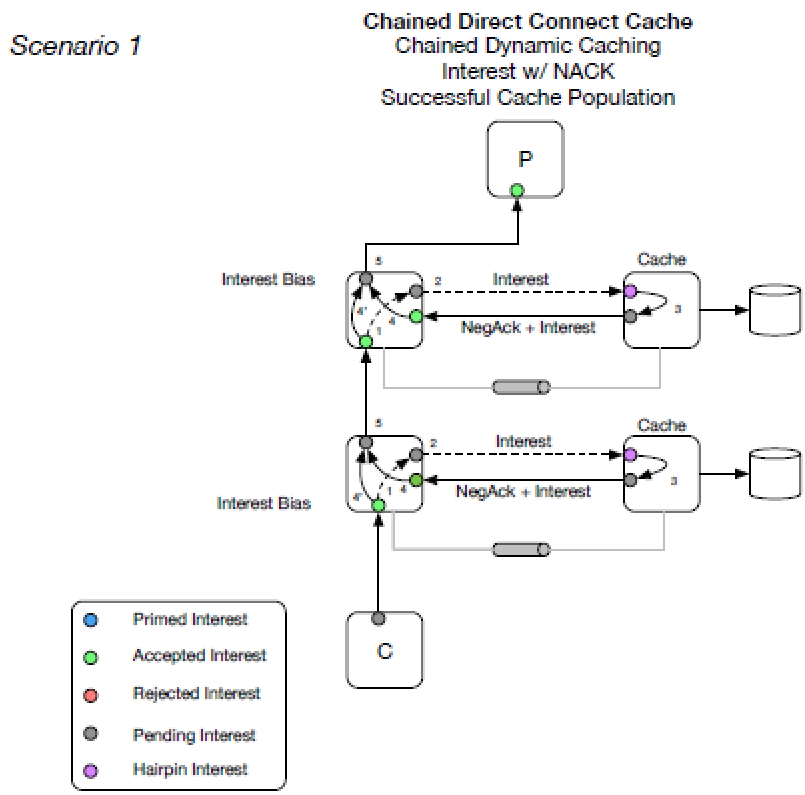
Scenario 1
Chained Direct Connect Cache
Chained Dynamic Caching
Interest w/ NACK
Successful Cache Population
P
Interest Bias
Interest
Cache
NegAck + Interest
Interest Bias
Interest
Cache
NegAck + Interest
C
Primed Interest
Accepted Interest
Rejected Interest
Pending Interest
Hairpin Interest

Figure 3

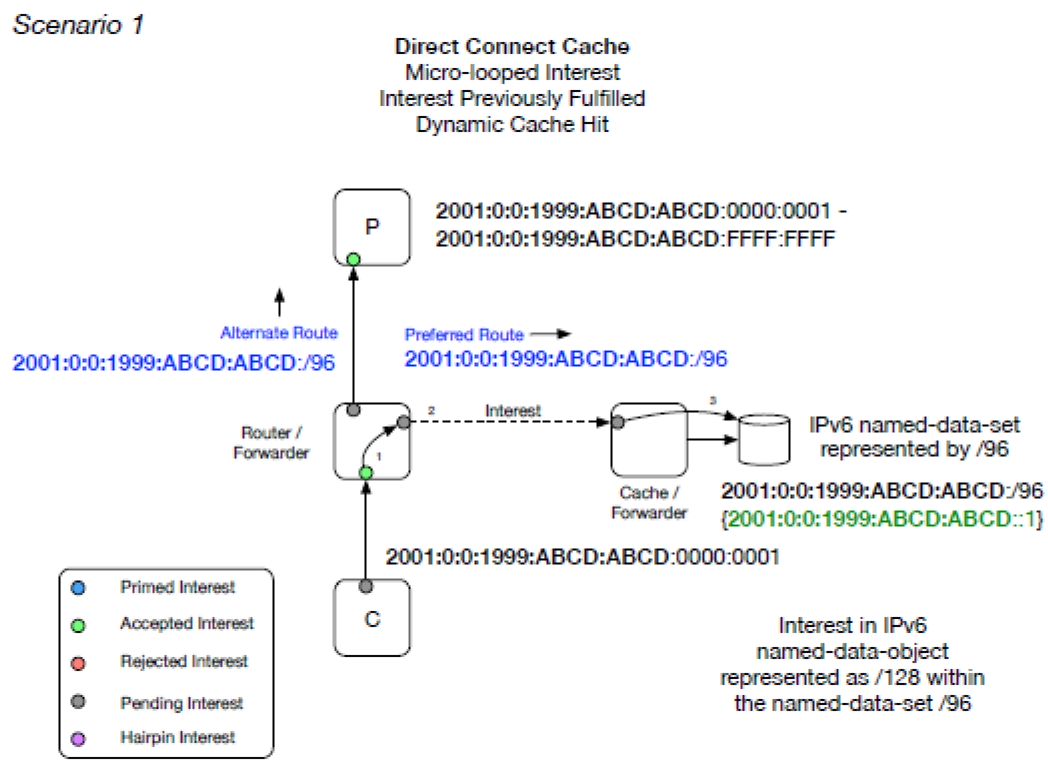

Figure 4

In this case, the router/forwarder may direct the Interest request of a /128 named data object across the direct interface because it has the longest prefix match. The longest prefix match can be set statically in the router/forwarder where an assumption is made that the set of named data objects for the prefix exists on the cache. Alternatively, the cache/forwarder may advertise the IPv6 prefix to the router/forwarder because the cache/forwarder is authorized to cache that set of named data objects. A cache that has been pre-populated with the set of named data objections will simply respond to the Interest request. However, a dynamic cache/forwarder that has yet to cache fill the named data object will fail to answer the Interest request. In fact, the dynamic cache/forwarder should become a Consumer and create an Interest request for the same named data object. The adjacent router/forwarder may recognize the Interest request from the cache/forwarder and discard it because it has a pending Interest targeting that same cache/forwarder. The first mechanism invoked is an Interest with a Negative Acknowledgement returned by the cache/forwarder. Such a packet may add a pending Interest on the router/forwarder's direct interface toward the cache/forwarder and negate the creation of an Interest back to the cache/forwarder as the preferred path. Now the router/forwarder may direct the new Interest from the cache/forwarder via an alternate interface along the next best path from the router's FIB. Presumably, the Interest now proceeds to an alternate cache or the Origin serving as a Producer. The response back to the router/forwarder may be copied to the cache/forwarder such that the named data object can be cached and the router/forwarder responds back to the original Consumer Interest request. Any subsequent requests for the same named data object may be answered by the cache/forwarder because it has now populated its cache. The Interest with the Negative Acknowledgement replaces the outgoing pending Interest toward the cache/forwarder with an incoming pending Interest. The router/forwarder may record the cache/forwarder interface as a Consumer with the pending Interest of the named data object and avoid the micro-loop.

Figures 5 and 6 below illustrate a second scenario involving two adjacent router/forwarders are also adjacent with the same cache/forwarder.

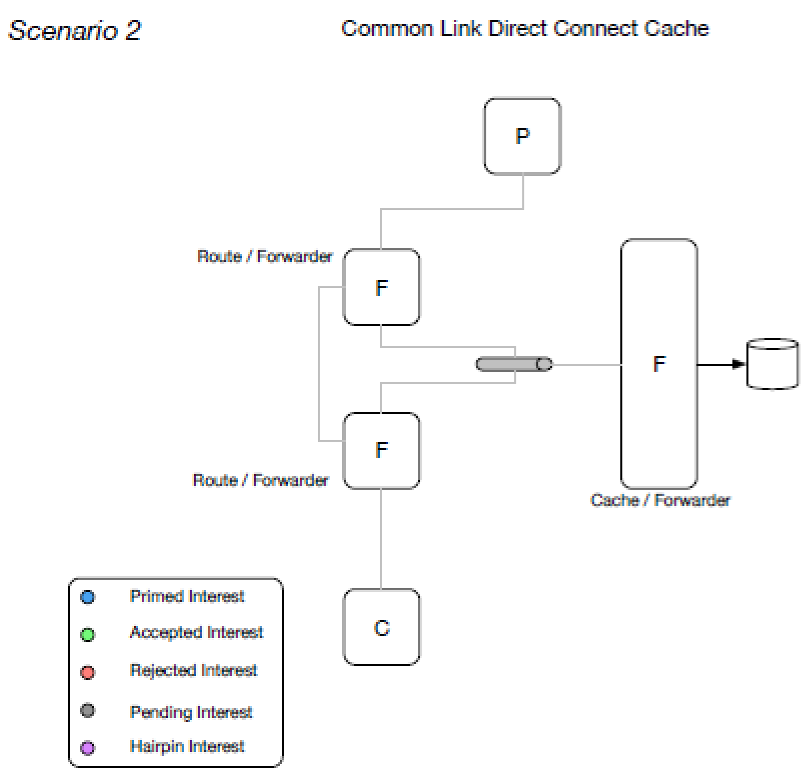

Figure 5

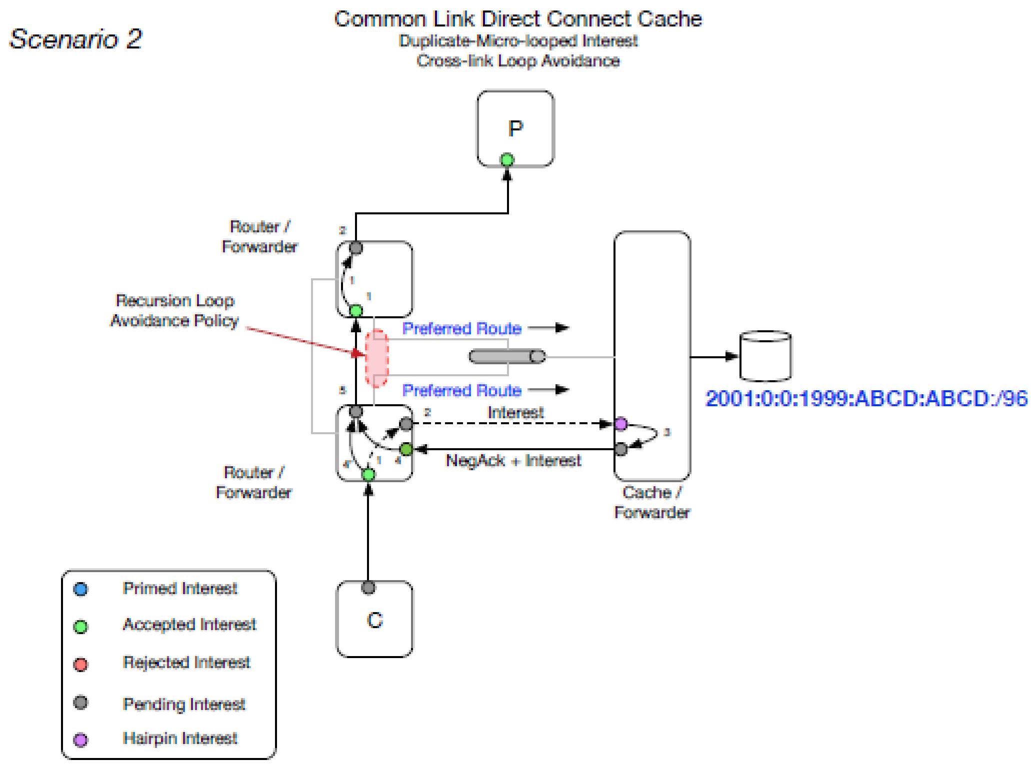

Figure 6

This architecture assumes that a router/forwarder-1 and router/forwarder-2 share a common link to the cache/forwarder and they have a cross-link between them. This is a common architecture where the cache/forwarder is provided two egress paths to the network at large. In this example, an Interest request is first directed by the router/forwarder-1 to the cache/forwarder as described in Figures 1-4 above. The cache/forwarder that is executing the dynamic caching algorithm may respond to the router/forwarder-1 as before with an Interest with a Negative Acknowledgment. Most likely, the router/forwarder-1 will find that the next best path to the same named object is via router/forwarder-2 (because the cache/forwarder advertised the prefix to both router/forwarder-1 and router/forwarder-2). Router/forwarder-1 may not send an Interest via the same common shared link to router/forwarder-2; however, router/forwarder-2 identifies an alternate path in the FIB for the cross-link. Thus, router/forwarder-1 directs an Interest request to router/forwarder-2 using the longest prefix match via the cross-link. Router/forwarder-2 shows the longest prefix match is toward the same cache/forwarder on the shared link, which is directly connected. The router/forwarder-2 may send the Interest request on the common link to the cache/forwarder because that is the best path to the named data object. The dynamic cache fill Interest request of the cache/forwarder now arrives back at itself, creating another micro-loop (the cache/forwarder has not yet received a response to its Interest via router/forwarder-1). The cache/forwarder should not create a second Interest with a Negative Acknowledgement as the network (both router/forwarder-1 and router/forwarder-2) returns duplicate responses for the same object. The cross-link between router/forwarder-1 and router/forwarder-2 may be marked such that the Interest with the Negative Acknowledgement requests arriving from the cache/forwarder on a common shared network can only be directed via the cross-link if the router/forwarder-1 has no alternative path. This is a conditional micro-loop avoidance strategy that is invoked on the router/forwarders that supersedes the FIB's default routing path (shortest path to named object via cross-link). An Interest request arriving from a downstream Consumer for the named data prefix may transit the cross-link as a viable path to the cache/forwarder. This is important because the common link between one of the router/forwarders to the cache/forwarder may have failed. The loop-avoidance strategy in this second scenario is a conditional forwarding of Interest requests via the cross-link towards the cache/forwarder where the Interest does not originate from a common shared link. An anti-affinity property of the named data object set is used to override the default FIB entries. A micro-loop is still possible in this second scenario.

Figures 7-9 below illustrate a third scenario.

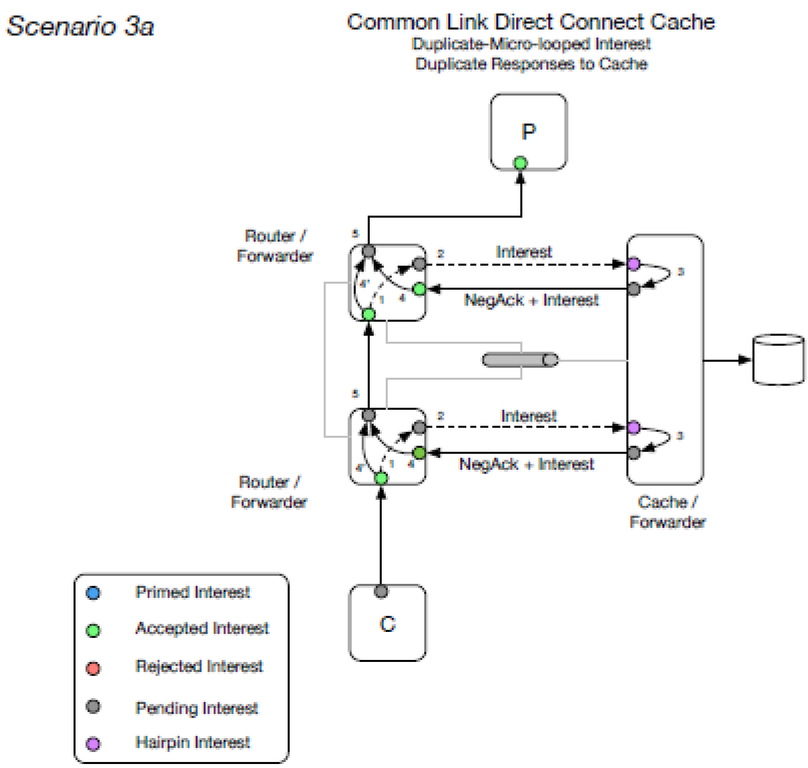

Figure 7

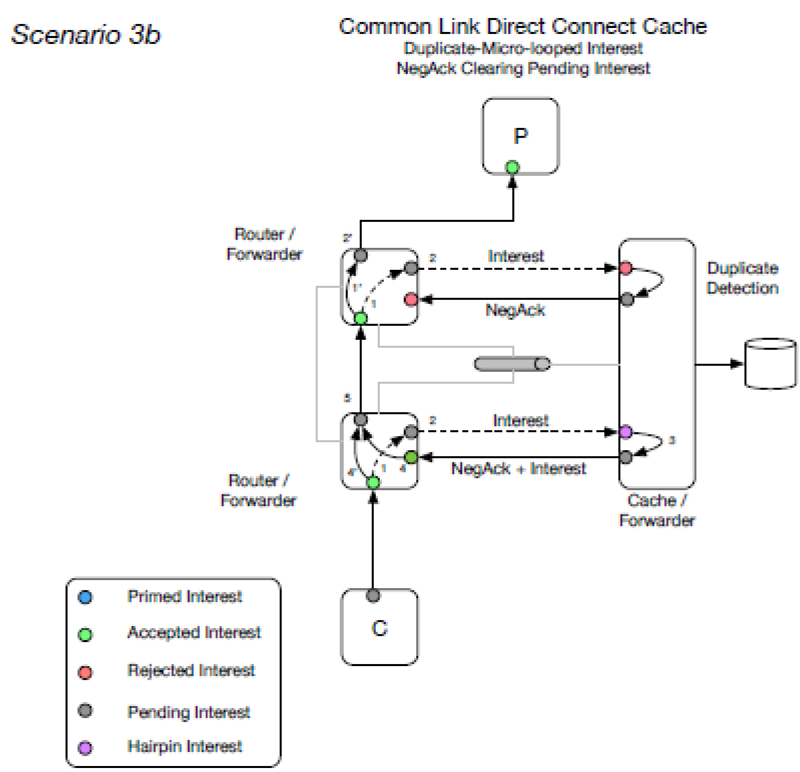

Figure 8

11 5607X

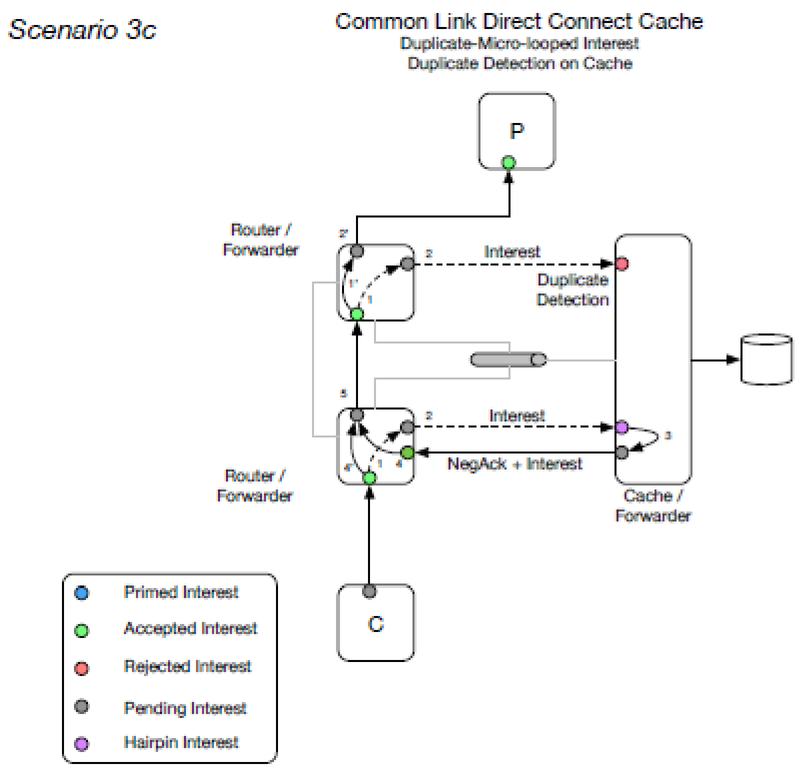

Figure 9

If a Consumer interest arrives on router/forwarder-1 and is directed to the cache/forwarder on the common shared network, the cache/forwarder may respond with an Interest with a Negative Acknowledgement for the newly requested object. The router/forwarder-1 with no alternative interfaces except the cross-link forwards the Interest request via the cross-link to router/forwarder-2 as a last resort. Router/forwarder-2 may forward the Interest requests on the common shared link because the FIB indicates that the prefix is directly connected and the cache/forwarder is the best path. In this case, the cache/forwarder will receive a duplicate Interest request for the same object. The cache/forwarder cannot answer its own request because it does not yet have the named object.

In one example, as illustrated in Figure 7 above, the cache/forwarder may create an Interest with a Negative Acknowledge as in the first scenario. However, the cache/forwarder may obtain two answers to the same Interest. Presumably, router/forwarder-2 forwards the Interest via a different Wireless Area Network (WAN) interface and receives a response that is copied to the interface of the cache/forwarder and the cross-link. Router/forwarder-1 copies the response for the Interest on the same shared common interface toward the cache/forwarder as well as the original Consumer. This method is therefore unsatisfactory.

In another example, as illustrated in Figure 8 above, the cache/forwarder may create a Negative Acknowledgement without an Interest, thus clearing the pending Interest request on router/forwarder-2 toward the cache/forwarder. This Negative Acknowledgement may only clear the pending Interest. The Negative Acknowledgement may not be cached because routing failure may require a subsequent Interest for the same object to be answered via the router/forwarder-2's interface towards the cache. Separate Consumer requests may have arrived at router/forwarder-2 where the Interest should be forwarded towards the cache/forwarder if the named data object exists on the cache/forwarder. The cache/forwarder may send a Negative acknowledgement for each of these Interest requests. This method is therefore unsatisfactory because it creates a flooding effect when the named data object is popular. The cache/forwarder must repeatedly answer with a Negative Acknowledgement.

In another example, as illustrated in Figure 9 above, the cache/forwarder may ignore the Interest request as a duplicate that arrived on the shared common interface even though the Interest had a different source. Typically, the cache/forwarder would answer all pending Interest with unique sources (router/forwarder-1 and router/forwarder-2). It is important to service requests from both router/forwarder-1 and router/forwarder-2 as Consumers may be attached to both and they used the shortest path of the prefixes to reach the cache. In this case, the router/forwarder-2 may have a pending Interest toward the cache/forwarder such that subsequent Interest requests for the same object are not repeated. When the router/forwarder-2 receives a response for the Interest via a WAN link, it may clear the pending Interest towards the cache/forwarder. The response to the Interest may traverse the cross-link to router/forwarder-1, where it is copied to both the original Consumer and the cache/forwarder. The router/forwarder-2 may have cleared its Interest request towards the cache/forwarder and router/forwarder-1 may have answered the Interest requests from the cache/forwarder. Now subsequent requests for the same object arriving on router/forwarder-1 or router/forwarder-2 may build Interest requests toward the cache/forwarder where both are answered without having to traverse the cross-link. This method provides the most optimal results for the third scenario.

Figures 10 and 11 below illustrate a fourth scenario in which adjacent routers to the cache/forwarder are not enabled with hICN forwarding logic.

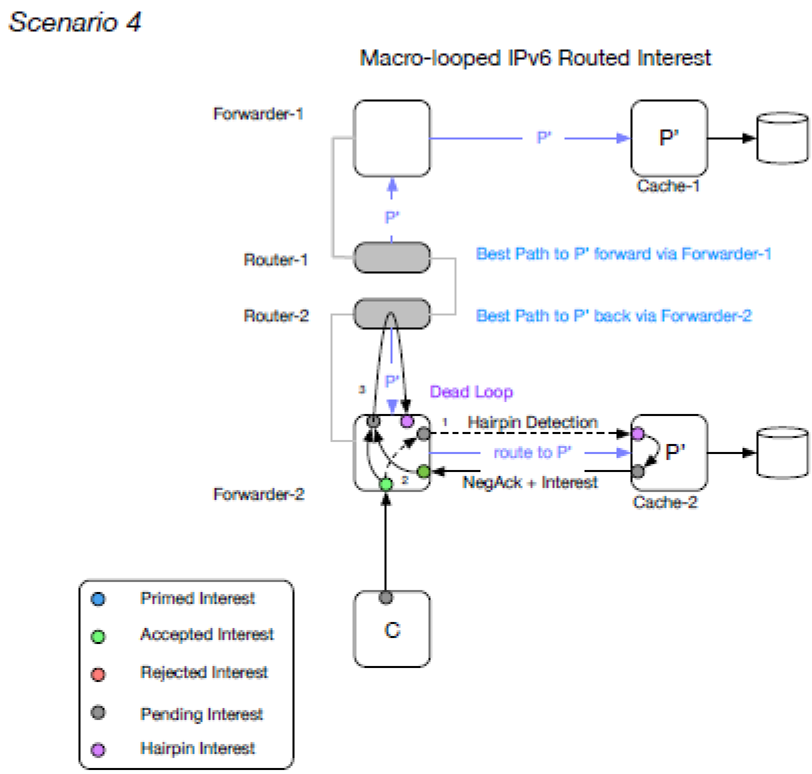

Figure 10

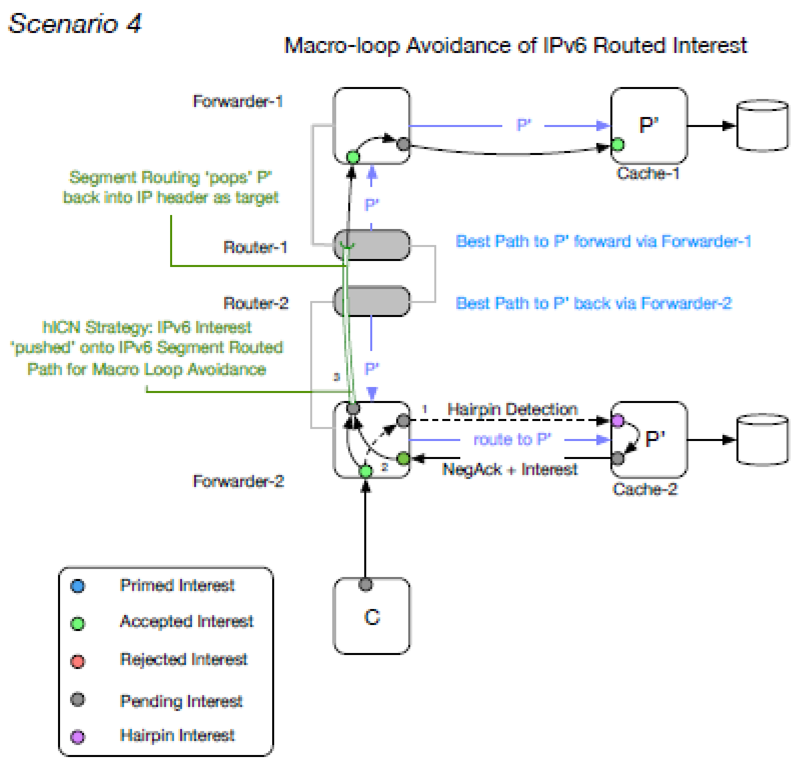

Figure 11

The previous scenarios assumed that hICN forwarding logic existed on router/forwarder-1 and router/forwarder-2, and thus the Pending Interest Table could be used to avoid micro-loops. In this forth scenario, the FIB in the adjacent router-1 and router-2 always point to the cache/forwarder. In this case, the Interest request arrives via router-1 at the cache/forwarder, where the object has not been dynamically cached. The cache/forwarder may create an Interest for the same object (same prefix) that is sent to either router-1 or router-2 according to the FIB. Router-1 and router-2 have no knowledge of the Interest requests and simply refer to their FIBs. They drop the Interest packet because its destination is on the same subnet from which the packet was received. The Interest request will never be answered. The cache/forwarder needs to know the best path to the logically adjacent hICN node such that intermediate routers (without hICN forwarding logic) will not loop the Interest back to the cache/forwarder. Therefore, a new strategy is defined for multi-hop Interest. The cache/forwarder may build an IPv6 Segment Routed path toward the logically adjacent hICN node for the target prefix. The Segment Routed path may cause the Interest request to traverse all the non-hICN routers without macrolooping repercussions. The logically adjacent hICN node may now build a pending Interest request from the cache/forwarder, thereby avoiding routing the Interest back on the same path from which the Interest arrived and using its FIB to forward the Interest on the next best prefix. Using the IPv6 Segment Routing strategy, the hICN nodes may build an optimal search path for content while following the FIB's shortest path to the IPv6 prefix representing the object. Other tunneling techniques may be used to achieve similar effects (e.g., Multiprotocol Label Switching (MPLS), Generic Routing Encapsulation (GRE), etc.).

The mechanisms described herein apply for an unspecified number of paths. The Negative Acknowledgement does not remove the pending Interest(s) but rather revokes the strategy decision. Thus, when the data packet returns, all cached interests must be satisfied. These mechanisms avoid looping back toward the same forwarding choice from which a new valid interest is received because the end cache wants to fill in with the data. This valid Interest is then cached as well such that if a valid data packet returns, all cached interests are satisfied and a data transmission is transmitted downstream.

In summary, techniques are described herein that provide a forwarding strategy applicable to ICN networks and hICN networks. In such networks, Interest requests are likely to be hair-pinned back to a Forwarder that is serving a dynamic cache system. When the dynamic cache system has no representation of the named data object, the Forwarder must recycle the Interest request back into the network. However, the Forwarding logic must avoid hair-pinning the Interest back to the same unpopulated cache.