B-ICP: Backpressure Interest Control Protocol for Multipath Communication in NDN

Yuhang Ye†, Brian Lee‡, Ronan Flynn‡, Niall Murray†, Yuansong Qiao†
Software Research Institute, Athlone Institute of Technology, Athlone, Co. Westmeath, Ireland
†{yye, nmurray, ysqiao}@research.ait.ie, ‡{blee, rflynn}@ait.ie

***Abstract*—Named Data Networking (NDN) is a promising communication paradigm to support content distribution for the Future Internet. The objective of this paper is to maximize the consumer downloading rate by retrieving content via multiple paths concurrently. This is supported by adaptive forwarding in NDN. The majority of solutions for selecting the forwarding interfaces do so based on latency. However, this can overload the low-latency paths quickly. This can occur as users reduce requesting rate according to congestion signals, from the low-latency paths. Hence, the high-latency paths are not fully utilized. This paper solves this problem by introducing Backpressure Interest Control Protocol (B-ICP). In B-ICP, routers estimate the forwarding capabilities of interfaces based on congestion signals and limit the forwarding rates to interfaces accordingly. Thus, B-ICP avoids congestion in certain paths prematurely, with the aim being evenly distributed paths utilization. Simulation-based evaluations show that B-ICP improves throughputs, converges to equilibriums quicly and supports the producers that dynamically join the network in comparison with existing solutions.**

Keywords – Named Data Networking; Adaptive Forwarding Control; Concurrent Multipath Communication; Load balancing

I. Introduction

NDN [1] is a novel paradigm connecting consumers to location-agnostic content. In contrast to TCP/IP, NDN follows an information-centric and pull-based principle. In this system, consumers (i.e. users) send addressless requests (i.e. *interest*) for content (i.e. *data*) from producers (i.e. servers). In NDN, host information is removed from interests. An interest contains a *name* in the form of a Universal Resources Identifier (URI) that specifies the content required.

According to a URI, routers can forward the interest to any authorized producer to retrieve content. As in-network caches are deployed in NDN, interests can be satisfied by intermediate nodes. A recent study [2] suggests deploying a large amount of memory (cache) at the edge nodes enhances utilizations and reduces management overheads. Without intra-network caching, the bandwidth of intra-network routers can be insufficiently utilized, which is same as TCP/IP. Indeed, employing the Peer-to-Peer (P2P) communication concept, which enables consumers to share content with each other, can further improve the resource utilizations and enhance the users' downloading rates. By removing the host information, consumers are not able to set up visible connections to producers. Instead, routers deliver interest packets to different producers for load balancing. Thus, the peer selection in traditional P2P is mapped to the adaptive forwarding in NDN.

In literature, different strategies have been studied to balance a load of downstream paths via equalizing local metrics (e.g. round-trip time [3] and pending interests [4]) and perform the congestion control [4]–[7] at consumers solely to enhance the resilience [8]. However, they do not meet the requirement of maximizing resource utilizations. For instance, the strategy in [3] equalizes the RTT metric prefers to select the low-latency paths for interest forwarding. However, the RTT of a congested low-latency path can still be smaller than the RTT of a congestion-free high-latency path. Thus, a low-latency path can be congested earlier than a high-latency path. According to the congestion signals (e.g. packet loss or ECNs), consumers have to reduce the requesting rate, even though the high-latency paths are underutilized. Similarly, the solution based on the pending interests [4] is affected by the router's link-layer queue configuration and the path latency. If a router is equipped with a large FIFO Tail-Drop queue or is configured with a high Random Early Detection (RED) threshold, it is likely to have many packets delayed in the queue as well as the number of pending interests record in the routers. According to the same principle as RTT, the throughput is degenerated [7].

In the respect of maximizing network utilizations, we propose Backpressure Interest Control Protocol (B-ICP). B-ICP assumes that the optimal forwarding rate to an interface should fill the downstream bottlenecks critically. Thus, the router is responsible for controlling the forwarding rate to an interface according to the estimated capability. Three key concepts are at the core of the B-ICP design. Firstly, the congestion level of any bottleneck is encoded into a *2-bit ECN* [9] that is returned to the upstream nodes. Secondly, the routers and the consumers execute Multiplicative Increase (MI) to probe the capability if the downstream utilization is low and Additive Increase Multiplicative Decrease (AIMD) to balance the fairness if the utilization is high, via ECN. Thirdly, routers quantize the overall congestion level of the downstream paths and feedback new ECNs to control the forwarding rate from upstream nodes. Additionally, B-ICP employs an efficient flow-aware resource allocation scheme to share the bandwidth among different flows.

In the context of the issues highlighted above, the contributions of this paper are: 1) a distributed forwarding controller to maximize network utilizations for multipath communications; 2) An MIAIMD-based control scheme to accelerate the convergence of network utilization; 3) A flow-aware interest shaping scheme to support the early congestion detection and notification [10] for fast convergence of inter-flow fairness.

The rest of the paper is structured as follows. Section II discusses the high-level notion of B-ICP. Section III shows the deployment and architecture of B-ICP. Section IV illustrates the detail design of the B-ICP components. Section V verifies the effectiveness of B-ICP by comparing it with the conventional protocols. Section VI concludes the paper.

II. Assumption and Principle

A. Assumption

Theoretically, the optimal traffic allocation for each path is a multi-commodity flow problem (MCF) [4]. However, as the demands of consumers keep changing, the cost of solving the MCF problems becomes expensive, either in a distributed or a centralized manner. In practice, the directed acyclic graph generated by routing protocols satisfies most applications and is convenient for B-ICP. The first assumption considers the directed acyclic routing paths are given by routing protocols, which enables the B-ICP to estimate the available bandwidth of downstream paths in a recursive manner. The second B-ICP assumption is *interest congestion*. The one-interest-one-data nature in NDN enables routers to predict data congestion based on the interest forwarding rate [10]. In order to avoid data congestion, Hop-by-hop Interest Shaping (HIS) [10] was proposed to limit the interest rate in advance, such that the bi-direction throughputs of a link are maximized. The interests that exceed the limitation are delayed in the interest shaping queue, where the queue length indicates the congestion of interest packets. In contrast to data-based congestion detection, the interest-based shaping scheme allows a flexible and fast interest re-distribution. For the sake of convenience, the "congestion" in the following content is equivalent to "interest congestion" unless with a special instruction.

B. Principle: Recursive Bandwidth Estimation

Apart from consumers and producers, the intermediate network nodes are classified into two types, namely *unipath routers* and *multipath routers*. A unipath router only maintains a single forwarding interface to access the content of each flow. A multipath router has multiple interfaces to access the content of a flow. Based on the multipath routers, the directed acyclic graph can be split into smaller link-structure sub-graphs, so called sub-path. For each sub-path, the head node denotes a consumer or a multipath router that forwards interests to the sub-path. Each head node is responsible for estimating the available bandwidth of its sub-path to avoid congestion. This bandwidth is defined as the smallest value of the bottleneck bandwidth of the sub-path and the *digesting ability* of the tail node. The tail node can be a producer or a multipath router. Its digesting ability is determined as follows: 1) if the tail node is also a head node (i.e. multipath router), the digesting ability is the sum of the available bandwidth of all its downstream sub-paths, 2) if the tail node is a producer, the digesting ability is assumed to be infinite as the producing ability of applications is typically as fast as required.

 # Example: recursive bandwidth estimation

For the topology in Figure 1, Node *Z* is the consumer, nodes *F*, *G*, *H* are producers providing the same content and nodes *A* and *D* are multipath routers. The graph is first split into the sub-paths *ZA*, *ABD*, *DEF*, *DG*, and *ACH*. For each sub-path, congestion can be avoided if the forwarding rate of the head node never exceeds the forwarding capability of unipath routers or the digesting ability of tail nodes. For instance, the sub-path *ABD* is cut via the head node *A* and the tail node *D*. Firstly, the available bandwidth of *A* to *ABD* must

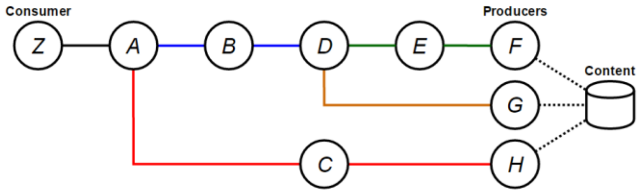

Figure 1 Topology segmentation via multipath routers

not exceed the forwarding capability of *AB* and *BD*. Secondly, the rate must be smaller than the digesting ability of *D* which is equal to the sum of the available bandwidth of *DEF* and *DG*. Note that the recursive bandwidth estimation design has a limitation on the topologies (e.g. diamond topology) that contains aggregating sub-paths [11], which affects the estimations and is solved by the following error-correction approach.

C. Approach: Error-Correction

In practice, an *error-correction* approach is used to accomplish bandwidth estimation. Here, the *error* is defined as the congestion level (e.g. number of delayed interests) detected by routers, and *correction* denotes that the router adjusts the forwarding rate to downstream paths to reduce the *error*, which is continuously executed unless the bottleneck is critically loaded. In order to feedback errors without adding expensive overheads [9]. B-ICP employs a 2-bit ECN to encode the *error* and trigger the MIAIMD *correction*.

 # Example: error-correction bandwidth estimation

Bandwidth estimation using the error-correction approach is also explained using the topology in Figure 1. For each sub-path, the ECNs carried by data packets are returned from downstream (e.g. the tail node) to upstream (e.g. the head node). In order to avoid congestion anywhere, the ECN received by the head node must be the heaviest congestion level that reflects the bottleneck status of the sub-path. Using the sub-path *ABD* as an example, the tail node *D* estimates the congestion level by comparing its digesting ability (i.e. the sum of the available bandwidths of *DEF* and *DG*) and the interest arriving rate and then attaches proper ECNs to the data packets. Thereafter, the ECN is updated along the sub-path, only if the congestion level of *AB* and *BD* is heavier that the indicated in the ECN generated by *D*. After ECNs are received by the head node *A*, *A* adjusts the available bandwidth to the sub-path *ABD*. Note that as the available bandwidth of *ABD* has been changed, and because node *A* is also the tail node of sub-path *ZA*, it will need to update its digesting ability (i.e. the sum of the available bandwidths of *ABD* and *ACH*).

III. Deployment and Architecture Design

A. Deployment

B-ICP is a cross-layer design which deploys the functional modules to the face layer, the strategy layer and application layer in the NDN stack. The face layer is responsible for detecting congestion level and feedback notifications. The strategy adjusts the forwarding rate to the interface according to notifications, aggregates congestion levels of downstream paths and returns notifications. The application layer adjusts the forwarding rate to the downstream paths as same as the strategy layer. However, it does not aggregate congestion levels and returns notifications, as the consumer is the head node of the graph.

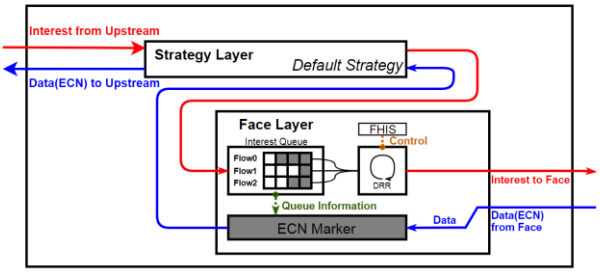

Figure 2 Unipath Router Design

B. *Architecture: Unipath Router*

As shown in Figure 2, B-ICP introduces three components to the face layer of the unipath routers, namely, a Flow-aware Hop-by-hop Interest Shaper (FHIS), an Interest Queue and an ECN Marker. FHIS estimates the optimal interest forwarding rate according to HIS [10] and schedules interests for each flow according to max-min fairness. As flows are managed independently, each flow maintains a separate queue. The interests that exceed the forwarding limits are buffered in the Interest Queue. The length of the Interest Queue indicates the congestion level of the corresponding flow. The ECN Marker encodes the congestion level and attaches it to the returning data packets of the same flow.

C. *Architecture: Multipath Router*

For the multipath routers, B-ICP is deployed to both the strategy layer and the face layer, as shown in Figure 3. The components on the face layer are identical to that of unipath routers. Moreover, the strategy layer introduces three new components, namely a Sub-path Interest Controller (SIC), an Interest Queue and an ECN Marker. Different from unipath routers, the SIC adjusts the interest forwarding rate to an interface according to the received ECNs. The interests that *exceed the total forwarding ability to all eligible interfaces* are delayed and buffered in the Interest Queues. Similarly, the ECN Marker utilizes queue lengths as congestion levels to generate ECNs and attach them on the returning data packets. Note that the dual-layer (i.e. face and strategy layers) queue deployment at multipath routers is essential. The length of the Interest Queue at the face layer denotes congestion level of an interface, while the queue length at the strategy layer indicates the overall congestion level of all downstream sub-paths.

D. *Architecture: Consumer*

The ECNs will finally arrive at the consumers. Following the same operating logic as SIC, consumers should adjust requesting rate according to ECNs. If some consumers are not cooperating, the traffic of other flows will not be affected as routers reserve resources for each flow using FHIS.

IV. Design in Detail

A. *Flow-aware Hop-by-hop Interest Shaper (FHIS)*

At the face layer, the Flow-aware Hop-by-hop Interest Shaper is an integration of Hop-by-hop Interest Shaping (HIS) [10] and flow-aware resources allocation [12]. HIS prevents data congestion of the opposite direction by limiting the interest outgoing rate in advance. However, HIS does not allocate resources between flows explicitly. It has been suggested [6], [12] that a light-weight flow-aware resource allocation can bring significant advantages. Thus, FHIS

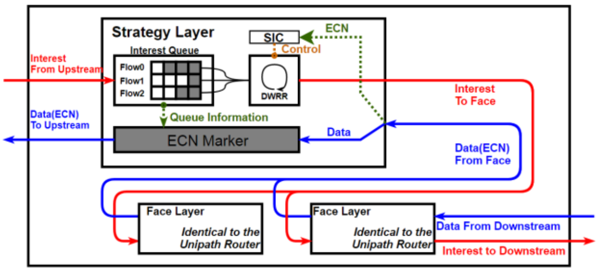

Figure 3 Multipath Router Design

employs a low-complexity scheduler, Deficit Round Robin [13] (DRR) to serve each flow with a fair (i.e. max-min) amount of resources. As each flow is served separately, independent Interest Queues are deployed to store the delayed interests for each flow.

B. *ECN Marker*

At the face layer and the strategy layer, ECN Marker detects the congestion level via the length of the Interest Queue and attaches ECNs to data packets. After receiving a data packet, the ECN Marker will first locate the Interest Queue of the corresponding flow and then sample the length of the queue. Due to the bursty nature of Internet traffic, the router smoothes the samples using an exponentially weighted moving average (EWMA) model in equation (1).

$$\overline{q_f}(t) = (1-\eta)\overline{q_f}(t-\Delta t) + \eta q_f(t) \tag{1}$$

Where $\overline{q_f}(t)$ is the smoothed sample of the queue length $q_f(t)$, η is a constant that defines the weight assigned to current sample values and previous estimations. As shown in Figure 4, three thresholds are defined (q_0, q_1, q_2), which split the congestion level into four stages, defined as follows,

Idle: $\overline{q_f}$ is smaller than q_0, the router encourages the upstream routers to radically increase the forwarding rate by a multiplicative increase (ECN: 0b01).

Lax: $\overline{q_f}$ is within the range between q_0 and q_1, the router requires the upstream routers to gradually increase the forwarding rate by an additive increase (ECN: 0b10).

Active: $\overline{q_f}$ is in the range between q_1 and q_2, the router returns ECNs of additive increase (ECN: 0b10) with probability $p = (q_2 - \overline{q_f})/(q_2 - q_1)$ and multiplicative decrease (ECN: 0b11) with probability $p = (\overline{q_f} - q_1)/(q_2 - q_1)$ to converge the queue length.

Overloaded: $\overline{q_f}$ is larger than q_2, the router inhibits the forwarding rate of upstream routers by a multiplicative decrease (ECN: 0b11).

In the face layer, the ECN is updated only if the detected congestion level is greater than the one carried by the data. This ensures the forwarding rate to be no larger than the minimum bandwidth of a sub-path. For multipath routers, the ECN received from an interface will trigger SIC, and it is deleted before the ECN marking in strategy layer. This is because that the congestion of a single sub-path does not represent the congestion of others. The ECN generated by strategy layer is a global view of the congestion level of downstream paths. The functionality of SIC will be discussed in the next section.

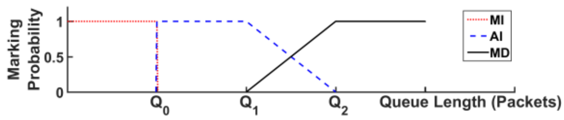

Figure 4 ECN Marking Functions

C. Sub-path Interest Controller

For the multipath routers, the forwarding rate to the forwarding interface is controlled by the received ECNs, which follows the MIAIMD algorithm [9] in equation (2).

$$\begin{aligned} MI \quad & L_{f,p} := L_{f,p} + \gamma \\ AI \quad & L_{f,p} := L_{f,p} + \frac{\alpha}{L_{f,p}} \\ MD \quad & L_{f,p} := L_{f,p} \times \beta \end{aligned} \tag{2}$$

$L_{f,p}$ denotes the interest forwarding rate of a certain flow f at interface p. α, β and γ are the parameters that affect the aggressiveness of the algorithm. Specifically, if the strategy layer of a multipath router receives an ECN from a sub-path, it adjusts the forwarding rate of the corresponding flow. If an ECN of 0b01 (MI) is received, the router increases the forwarding rate by γ to aggressively fill in the empty sub-path. If an ECN of 0b10 (AI) is received, the router increases the rate by α / $L_{f,p}$ to enhance utilization. In contrast, the router will reduce the forwarding rate by multiplying the forwarding rate by β, once an ECN with 0b11 (MD) is received.

As high bandwidth networks may cause the routers to generate a vast amount of congestion signals (either packet drops or congestion ECNs) simultaneously, the continuous deduction will drain the multipath routers' interest requesting rates instantaneously. In SIC, a single multiplicative decrease with an *ignoring timer* is employed to ensure stability. The timer makes the router ignore all the received MD ECNs for a certain period τ after one MD procedure is executed.

V. Evaluations

B-ICP is evaluated via three metrics, namely, steady-state downloading rate, utilization convergence time and fairness convergence time. As many of the existing traffic control solutions for NDN (i.e. CHoPCoP [6], ECP [14]) lacks the discussion of the multipath forwarding design, this makes the comparative study difficult. B-ICP is compared with the state-of-the-art approaches which have considered multipath communications: 1) Optimal Multipath Congestion Control and Request Forwarding [4] (OMCC-RF) and 2) Dynamic Interest Limiting [15] (DIL). Specifically, OMCC-RF formulates a global optimization problem with a two-fold objective that minimizes cost and maximizes utilizations. As OMCC-RF requires the consumers to know the content origin and forwarding path via route-labelling, it is open to security weaknesses [7]. DIL makes the router detect congestions via the congested data packets and notifies upstream nodes via NACK packets. DIL suffers from two issues: 1) the data-based congestion detection causes a long control loop which slows down the reaction of upstream nodes to congestion signals and exacerbates the congestion [10], 2) DIL does not support slow start, which takes a long time to converge. Additionally, for the last experiment, an unipath interest shaping approach - HIS [10] is compared with B-ICP for verifying the fairness convergence of bidirectional flows. Here, bidirectional flows denote that a link is bearing interests from both directions simultaneously. In this case, the data packets will compete for the bandwidth with interests in the same direction. Without a proper interest/data allocation, the flows of a certain direction may be starved by others. As HIS does not consider the multipath forwarding, it is only compared using the unipath scenario.

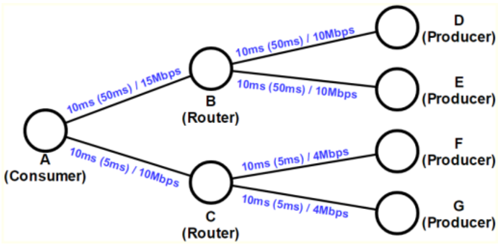

Figure 5 Topology for Multi-source Scenario

TABLE I LINK UTILIZATIONS FOR DIL, OMCC-RF AND B-ICP

CASE I: SMALL LATENCY VARIANCES / CASE II: LARGE LATENCY VARIANCES

	Link Utilizations (in Mbps)					
	Case I			Case II		
Link	**B-ICP**	**DIL**	**OMCC-RF**	**B-ICP**	**DIL**	**OMCC-RF**
(A, B)	14.91	14.96	14.32	14.63	13.43	11.06
(B, D)	7.49	7.44	7.22	7.31	6.71	5.52
(B, E)	7.42	7.52	7.10	7.32	6.72	5.54
(A, C)	7.89	7.94	7.22	7.88	7.94	7.41
(C, F)	3.94	3.96	3.66	3.91	3.97	3.71
(C, G)	3.95	3.98	3.56	3.97	3.97	3.70
Throughput (Mbps)	22.8	22.9	21.54	22.51	21.31	18.47
Utilization	99.1%	99.6%	93.7%	97.9%	92.7%	80.3%

A. Simulation and Parameter Setup

B-ICP is implemented through the NDN Forwarding Daemon and simulated via ndnSIM [16]. The MTU size is set to 1500bytes. If a node can cache content, its buffer aims to hold the data packets for at least 20s (i.e. the cache is never hit by new interests). For the EWMA model, we select $\eta = 0.875$. For ECN Marker, the parameters are set as follows: $q_0 = 0$, $q_1 = 20$ and $q_2 = 50$. The parameters for SIC are configured as, $\gamma = 1$, $\alpha = 2$, $\beta = 0.875$, $\tau = 1$. The re-transmission timer for B-ICP consumers is set as 1.25 times the upper bound of 100ms historical RTT samples to avoid re-transmitting too frequently.

B. Scenario 1: Multi-source Communications

In practice, the "distances" between peers can be much different, where the factor of heterogeneous latency must be considered for evaluating the robustness. The first experiment validates the effectiveness of B-ICP in the multi-source transmission scenarios that the latencies of paths are different. The topology is shown in Figure 5. Two cases of link settings (with small and large latency differences) are considered. The latency settings of Case II are inside parentheses. The caching is disabled for the experiments. Table I reports the average link utilization of Case I and Case II. For Case I, DIL, OMCC-RF and B-ICP have near-perfect performances. As DIL utilizes the link-layer queue length to detect congestions, the performance is slightly better than the others [15]. For Case II, only B-ICP can achieve near-perfect use of the overall resources. The RTT-biased metric of OMCC-RF [7] makes nodes distribute unbalanced interests to paths, and incorrectly reduce the forwarding rate if any path is congested. DIL can equally distribute interests to paths with heterogeneous latencies. However, as DIL detects congestion based on data queues, the larger latency in Case II makes DIL over-oscillating, which reduces the downloading rate.

C. Scenario 2: Content Caching

Caching enables routers to store content temporally, which prevents duplicated requests to producers with the same content. The second experiment verifies the effectiveness of

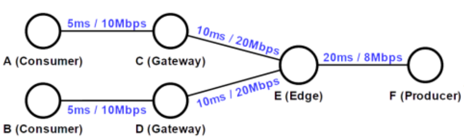

Figure 6 Topology for Content Caching (Scenario 2)

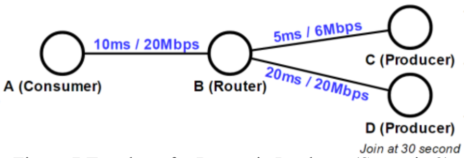

Figure 7 Topology for Dynamic Producer (Scenario 3)

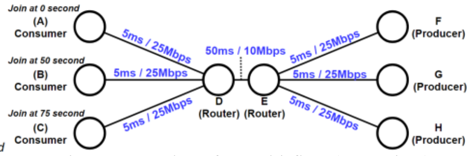

Figure 8 Topology for Multi-flow (Scenario 4)

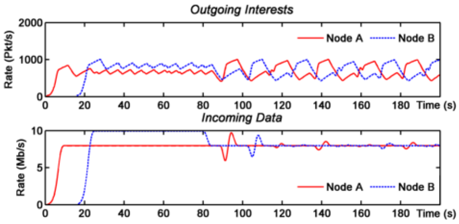

Figure 9 Interest Rate and Data Downloading Rate of Scenario 2

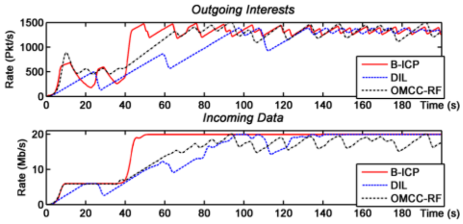

Figure 10 Interest Rate and Data Downloading Rate of Scenario 3

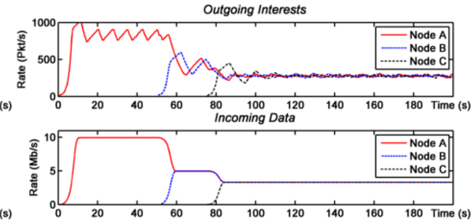

Figure 11 Interest Rate and Data Downloading Rate of Scenario 4

B-ICP when the cache is available. A simplified topology with the cache deployed at the edge nodes is shown in Figure 6. Two consumers (*A* and *B*) of the same flow download the content from a unique producer *E*. To prevent routers aggregating the duplicated interests [1], consumer *A* starts requesting content at time 0s, while *B* starts at time 15*s*. The interest sending rates and downloading rates of the two consumers are shown in Figure 9. Initially, the consumer *A* downloads the content from the producer F, and the content is cached at the edge *E*. The bottleneck link *FE* (8Mbps) is fully utilized. At 15s, consumer *B* begins to send interests. As the content has been stored at the edge *E*, the interests that hit the cache at *E* will fetch the data packets directly. It is easy to figure out that, the path of *FECA* is not shared with *EDB*, the downloading rate of *B* increases to 10Mbps (*DB* becomes the new bottleneck) immediately. After about 80 seconds, the interests from *B* begin to become synchronized with the interests from *A*. Thereafter, edge router *E* aggregates the interests that come from the both consumers and provides the data to them together. The result shows the effectiveness of B-ICP to support content caching.

D. *Scenario 3: Dynamic Producers*

The multipath nature of NDN allows consumers (e.g. peers in P2P) to publish local content to the network. Additionally, NDN also supports the CDN-like content caching at specific nodes [16]. The network resources will fluctuate while producers join or leave. In this scenario, the experiments evaluate B-ICP's convergence speed for tracing dynamic network resources. The detailed topology is shown in Figure 7 and the experimental results are shown in Figure 10. Initially, producer *C* is connected to router *B*. Meanwhile consumer *A* starts to download content. At time 30s, a dynamic producer – *D* joins the network, which is connected to the router *B*.

According to the graph, we can clearly figure out that, the proposed B-ICP (red line) spends the least time ($\approx$7.5s) to fully utilize the bandwidth provided by producer *D* without introducing instability. Although OMCC-RF (black line) enables the slow-start at the beginning, after congestions,

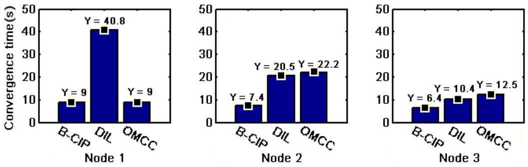

Figure 12 Convergence Time Comparison (smaller better)

OMCC-RF needs linear time ($\approx$40s) to probe the added resources. Moreover, as DIL (blue line) does not implement slow-start to probe the available bandwidth, it takes even longer time (>40s) to achieve a full utilization.

E. *Scenario 4: Inter-flow Fairness and Convergence Speed*

In this scenario, we investigate how B-ICP behaves when multiple flows co-exist. The experiments are conducted on the topology defined in Figure 8, which consists of three consumers (*A*, *B*, and *C*) and three producers (*F*, *G*, and *H*). Each consumer downloads different content from the corresponding producer (i.e. *A* to *F*, *B* to *G* and *C* to *H*), which results in three flows in the network. In order to demonstrate the fairness convergence among the competing flows clearly, *A* is started at time 0s, *B* at time 50s and *C* at 75s. Figure 11 presented the fairness convergence of B-ICP. Before *B* starts, the only flow from A tries to fill up the bottleneck link exclusively. This results in an overall downloading rate of 9.94Mbps to *A*. As soon as *B* starts, each flow converges to the optimal the downloading rate: 4.97Mbps for both *A* and *B*. As soon as *C* starts; the bandwidth is equally divided into 3.32Mbps for *A*, *B*, and *C*. Figure 12 presents the fairness convergence time for each flow. Obviously, B-ICP outperforms the other solutions, as it reserves the bandwidth for each flow proactively.

F. *Scenario 5: Bi-directional Inter-flow Fairness*

The traffic in NDN is inherently bidirectional. Specifically, a link undertaking interests from both directions inevitably cause the competition between data and interests. An optimal allocation is essential, such that the link's utilization is maximized. Although HIS has considered a reasonable interest/data allocation, it does not proactively allocate the resources for each flow in the same direction. This will lead to problems when consumers are non-cooperative [12], [17]. The experiment described here aims to verify the effectiveness of the B-ICP protocol if non-cooperating consumers exist. The topology is shown in Figure 13, where four flows exist. Node *A* downloads content from node *D* (starts at 80s); node *B* downloads content from node *C* (starts at 40s); node *C* downloads content from *A* (starts at 0s); node *D* downloads content from *B* (starts at 120s). The link *R1R2* is the bottleneck. Figure 14reports the interest rate and content downloading rate of nodes *A*, *B*, *C* and *D* for HIS and B-ICP respectively. Node *B* is a non-cooperative consumer that sends interests with a constant rate (500 packets per second). As