Multipath Transmission in Content Centric Networking Using a Probabilistic Ant-Routing Mechanism

Jonas Eymann and Andreas Timm-Giel

Hamburg University of Technology
Institute of Communication Networks
{jonas.eymann,timm-giel}@tuhh.de

Abstract. Content Centric Networking (CCN) is a new networking paradigm that names pieces of content rather than network nodes. It promises more efficient transmissions due to in-network caching and easier realization of mobile and multihomed devices. However, in order to leverage multipath transmission for multihomed devices, routing and forwarding mechanisms are needed that support this functionality. In this paper, we present a probabilistic ant-routing mechanism that enables multipath transmissions for CCN nodes. Using an OMNeT++ based simulation model, we show that our routing mechanism can support transmissions of data streams over multiple links to achieve higher throughput than any single link could provide.

Keywords: Content Centric Networking, CCN, Named Data Networking, NDN, Information Centric Networking, ICN, Routing, Forwarding, Multipath, Ant Colony Optimization.

1 Introduction

The architecture and core protocols of today's Internet, TCP/IP, were developed more than 30 years ago and have proven to be remarkably versatile. The Internet supports applications that were never foreseen and the number of users has exceeded all expectations, for which the exhaustion of IPv4 addresses is a very apparent sign.

However, besides the scarcity of IPv4 addresses, several other fundamental design decisions are no longer in line with today's requirements and prevalent usage of the Internet. Nowadays, users are mostly interested in content or information regardless of which specific server (i.e., which IP address) it is hosted on, yet the network still requires IP addresses to operate. The host-centric architecture with "fixed" IP addresses for hosts is becoming increasingly incongruous for mobile devices that feature more than one network interface (e.g., WiFi and UMTS) and hence have changing network addresses. Seamless mobility, while theoretically solved [1], has never reached a widespread adoption. For instance, switching from WiFi to the mobile phone network on a smartphone still breaks

D. Pesch et al. (Eds.): MONAMI 2013, LNICST 125, pp. 45–56, 2013.

existing TCP connections of running applications. Security was also not part of the original design, therefore numerous additional protocols were developed as an afterthought. Likewise, many services have to make a significant effort to overcome deficiencies in the original TCP/IP architecture. For example, the problems induced by Network Address Translation (NAT) require Skype and other applications to implement NAT traversal mechanisms like STUN [3,4] and due to the lack of in-network caching, services such as YouTube have to rely on Content Distribution Networks [2] to efficiently deliver its services.

Content Centric Networking, or CCN [5], is a new approach that sets out to close this widening gap between architecture and usage of the Internet by naming information chunks rather than addressing nodes.[1] Instead of sending packets to a destination specified by an IP address, CCN nodes send out requests (called Interests) for pieces of content (called Data) specified by a unique, hierarchically structured name. Any node that has matching Data to a traversing Interest can answer immediately because the name unambiguously identifies the Data. If an Interest cannot be answered locally at a node, it will be forwarded to one or more neighboring nodes. The unique name of Data packets enables in-network caching as the content can be stored by any node or router in the network. To prevent unauthorized alteration of content, each Data packet contains a digital signature. Optionally, Data packets can also be encrypted, hence security is a central part of the architecture though this is beyond the scope of this paper (see [6] for further information).

The absence of a fixed destination node address in CCN also facilitates multipath transmissions for multihomed devices with more than one network interface. Unlike IP packets, CCN Interests can be forwarded via several interfaces simultaneously to increase resilience to packet loss. For increased throughput, link bandwidths can be aggregated by forwarding Interests in an interleaving order. However, little research has been done so far on how such multipath forwarding mechanisms for CCN should be designed.

In this work we describe a probabilistic routing and forwarding mechanism for CCN to find multiple paths to sources of content. The mechanism is based on the idea how ant colonies in nature find paths to food sources. Using an OMNeT++ based simulation model of CCN, we demonstrate how such a probabilistic ant-routing mechanism makes it possible to use the paths to transmit constant bit rate data streams over multiple links to increase throughput.

The remainder of this paper is structured as follows: Section 2 discusses related work on probabilistic ant-routing. The basic forwarding mechanism of CCN and existing routing options are explained in Section 3. In Section 4 we present our probabilistic routing and forwarding mechanism for CCN. To test multipath transmission using probabilistic routing and forwarding, two different simulation scenarios are evaluated in Section 5. Finally, Section 6 gives some hints on future work and concludes the paper.

[1] Content Centric Networking is also known as Named Data Networking or NDN, see `http://www.named-data.net`

2 Related Work

The idea of applying the optimizing behavior of ant colonies (Ant Colony Optimization) to telecommunication networks has been studied in several previous works [9,11,12]. It has also been shown that due to their ability to adapt to changing environments, ant-like mechanisms are also suitable for dynamic networks such as MANETs (mobile ad-hoc networks) [13].

A common property in the publications above is the differentiation between "ants" or agent packets that explore the network and find paths, and data packets that transfer user data. As CCN operates by coupled pairs of Interest and Data packets, the existing mechanisms are not directly applicable. However, we can interpret Interests similar to "ants" which simplifies the design, as discussed later.

Ant-based routing and forwarding has also been proposed for CCN. In Services over Content Centric Routing (SoCCeR) the concepts of content centric networking and service centric networking are combined [14]. SoCCeR uses an ant-like approach as an additional routing mechanism for the communication of services, but not as a general routing mechanism for Interests to sources of static content. As all services are assumed to be known, special Interest ants are sent out at regular intervals towards randomly selected services in order to find and update paths. In our approach, we assume unknown destinations, thus use a broadcasting mechanism instead.

A second proposal for ant-routing in CCN is Greedy Ant Colony Forwarding (GACF [15]). GACF uses "Hello ants" created by routers to measure paths to all sources, thus, like SoCCeR, assumes known destinations. Client nodes (consumers) send out "Normal ants" for requesting user data packets. In contrast, our scheme only uses normal CCN Interest and Data packets, which in our view is sufficient and also more closely complies with the architecture of CCN.

Both SoCCeR and GACF showed that Ant Colony Optimization mechanisms are suitable for CCN and can achieve low delays. However, none of the related work on ant-routing investigates the aggregation of links. In this paper, we analyze the multipath behavior and aggregation of multiple links for higher throughput, which to our knowledge was not addressed before.

3 Forwarding and Routing in CCN

Similar to IP packet forwarding, the CCN forwarding process of Interests relies on a structure called Forwarding Information Base (FIB). The FIB contains name prefixes of content collections and one or more associated interfaces through which the content can be reached (Fig. 1). An interface can be either a physical link such as a network interface or a local application, hence the term "faces" to distinguish this generalized concept.

The names of incoming Interests that cannot be answered from the cache[2] are matched against the prefixes in the FIB using longest prefix matching and

[2] In Content Centric Networking, the cache is called Content Store.

then forwarded to potential sources of content. Additionally, the node also keeps the information from which interface the Interest arrived in a data structure called Pending Interest Table (PIT). If another Interest arrives for the same Data (i.e., same name), the node adds the arrival interface to the existing PIT entry but does not forward the Interest again. When the Data packet comes back in response to the Interest, the node caches the Data and sends it to all faces for the corresponding PIT entry, which is subsequently deleted. The aggregation of Interests thus saves Data packets from being transmitted more than once over the same link.

Prefix entries in the FIB can be configured manually, which is the method currently used for the prototype implementation of CCN that is developed in the project CCNx[3]. While manual configuration is sufficient for small testbeds with only a few nodes, larger networks will have to use other mechanisms that are also capable of dynamically reacting to changes in the network. It is possible to use existing routing protocols such as Open Shortest Path First (OSPF) and adapt them for CCN, which is the goal of OSPFN [7]. However, in order to

Prefix	Faces
/de/tuhh	2,3
/de/tuhh/comnets	3
/de/uni-bremen	4

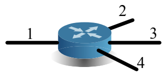

Fig. 1. Example of the Forwarding Information Base (FIB)

leverage cached data that is nearby a node but not on the direct path to the original server, routing protocols such as OSPFN cannot adapt quickly enough to capture the locations of highly dynamic cached data which changes at line speed. Previous research suggested that probabilistic routing and forwarding without algorithmically calculated shortest paths might therefore be beneficial for CCN [8].

4 A Probabilistic Ant-Routing Mechanism for CCN

Our probabilistic routing and forwarding scheme is based on the idea of Ant Colony Optimization [9] which is inspired by the way biological ant colonies find paths to sources of food. In nature, each single ant leaves a pheromone trail on the ground. Other ants usually follow trails with strong pheromone traces, thus further increasing the trail. With a some small probability, ants diverge from the trail looking for other paths to the known source or to discover new sources of food. As short paths to sources will generally be used more often (as an ant can cover a short path more often than a long path), there is an implicit optimization towards shorter paths [10]. Over time, pheromones evaporate and thus unused paths slowly fade.

[3] `http://www.ccnx.org`

Applying the idea of Ant Colony Optimization to CCN, we can interpret Interests as ants looking for food, and the corresponding Data packets as the ants carrying food back. The pheromone traces that real ants leave on the ground can be compared to the PIT entries and changes to the FIB that influence how Interests are forwarded. For this reason, we extended the FIB with additional information indicating the quality of a face for each prefix (i.e., the path to a source of content), as shown in Fig. 2. The algorithm of our probabilistic routing

Prefix	Face	q
/de/tuhh	2	0.7
	3	0.3
/de/tuhh/comnets	3	1
/de/uni-bremen	4	1

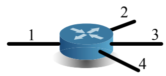

Fig. 2. Example of an extended Forwarding Information Base (FIB) augmented with information on the quality q of each face for a prefix

and forwarding mechanism consisting of three phases can then be described as follows:

1. In the beginning FIBs are empty, i.e., do not contain any entries. When an Interest arrives and there is not entry for the prefix, the node broadcasts the Interest to all faces except the one that it arrived on.
2. When a Data packet arrives and there is no entry in the FIB for the Data's prefix, a new entry is created and the quality indicator is set to 1 since there is only one face for the prefix. If there is already an entry for the prefix, the quality indicator q_i for the arrival face i is updated to

$$q_i \leftarrow q_i + r(1 - q_i) \quad (1)$$

where the factor r, $0 < r \leq 1$, influences the reinforcement of the quality value depending on the observed RTT (see below). For all other entries of faces $j \neq i$, the quality indicator is updated by

$$q_j \leftarrow q_j - rq_j \quad (2)$$

The quality indicator of the arrival face i is therefore increased while the quality indicators of all other faces are decreased. The factor r is calculated based on the round trip time d_{RTT} of the Data that has arrived according to

$$r = r(d_{\mathrm{RTT}}, \alpha) = e^{-\alpha d_{\mathrm{RTT}}} \quad (3)$$

For very small RTTs r approaches 1, indicating a very good connection and thus resulting in a larger reinforcement in Equation (1) and decrease in Equation (2). The additional parameter α in Equation (3) is used to adjust the influence of an arrived Data packet, i.e., how strongly new RTT values are taken into account to change the quality indicators of the links. A small value of α results in a faster reinforcement of good paths and evaporation of unused paths.

3. When Interests arrive at the node with a matching prefix in the FIB, the quality indicator determines the probability with which each face is used for forwarding. For example, in Fig. 2, for the first entry, face 2 is used with a probability of 0.7 while face 3 is used with a probability of 0.3. However, with a small probability p_b, the Interest is broadcast to all faces as if there was no entry. This way the node can discover new sources of Data and adapt to changes in the network. The normal probabilistic forwarding described above is therefore only applied with the probability $1 - p_b$.

The goal of the developed probabilistic ant-routing mechanism is to achieve high throughput using multipath transmission and at the same time minimize the RTTs of Interest/Data pairs. To test the mechanism, we simulated two different scenarios for different values of α, described in the following section.

5 Simulation Scenarios and Results

We implemented a model of CCN and our probabilistic ant-routing algorithm using the event-based OMNeT++ simulation framework[4]. In order to analyze the mechanism in changing network conditions, we simulated a simple dual path scenario where one path is deactivated during the simulation run and reactivated later. The behavior in a more complex network is simulated using a realistic backbone (NSFNET scenario) under static conditions.

5.1 Dual Path Scenario

The first scenarios is a network with one client that can receive a constant bit rate (CBR) stream from two different servers as shown in Fig. 3 (dual path scenario). To receive the streams, the client sends out an Interest every 20 ms.

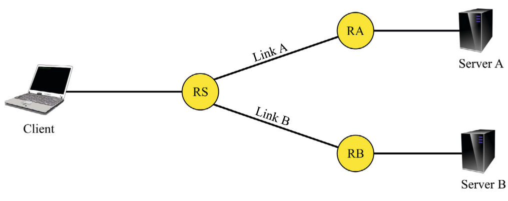

Fig. 3. Dual path scenario with two servers A and B offering the same streams. Link A and B are the bottleneck links with a transmission rate of 250 kbit/s and a delay of 100 ms.

The splitting router RS forwards the Interests either via link A, link B or both links at the same time in case of a broadcast (the broadcast probability is set to $p_b = 0.1$). Link A and B are configured as bottleneck links with 250 kbit/s.

[4] http://www.omnetpp.org

Both servers contain the same content and can answer to the Interests with Data packets. In one simulation configuration, the size of the Data packets is 500 B, resulting in a CBR stream of 200 kbit/s, which is less than the bottleneck transmission rate of one link. In the second configuration, Data packets of 1000 B result in a CBR stream of 400 kbit/s. This stream can only be successfully transmitted when both bottleneck links are aggregated and used simultaneously. In order to examine how the ant-routing algorithm adapts to changes in the network, link A is deactivated at $t = 50$ s, i.e., drops all packets. At $t = 150$ s, the link comes up again.

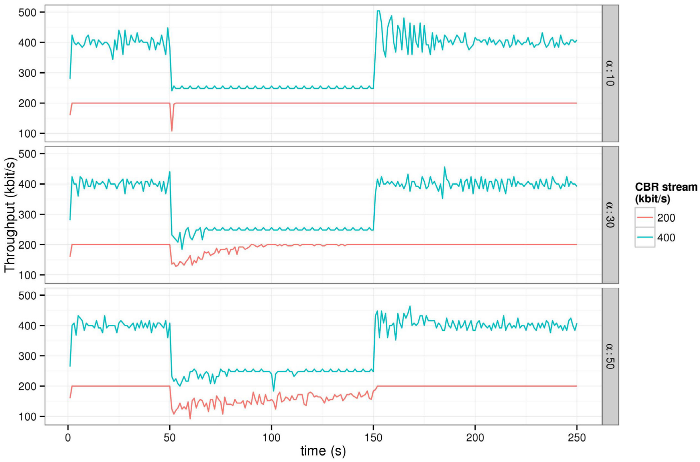

Fig. 4. Received throughput of CBR streams at the client in dual path scenario for α=10, 30 and 50 (single runs). Link A is deactivated between $t = 50$ s and $t = 150$ s.

Fig. 4 shows the achievable throughput for the dual path scenario for three different values of α. As expected, the lower CBR stream with 200 kbit/s can be transmitted, since the Data rate is lower than the bottleneck link speed. The higher CBR stream also achieves its nominal transmission rate, though the throughput fluctuates due to the queueing that occurs at the bottleneck links. The influence of α becomes visible at $t = 50$ s. For larger values of α, the low CBR stream needs much longer to shift to the one remaining link while for smaller α it adapts faster to the change. However, for small values of α, the routing does not stabilize well, as can be seen from the larger fluctuation for the higher CBR stream after the link comes up again at $t = 150$ s.

During the period that link A is deactivated, the high CBR stream can only transmit with the transmission rate of the bottleneck link of 250 kbit/s. Many Data packets are therefore queued and eventually dropped at router RB as the stream of Interests to the server still remains unchanged. After link A is reactivated, the 400 kbit/s stream can again achieve its full transmission rate as the ant-routing algorithm adapts to again using both links simultaneously.

5.2 NSFNET Scenario

The second scenario we simulated is based on the backbone of the National Science Foundation Network (NSFNET) as shown in Fig. 5. The NSFNET backbone consists of 14 routing nodes to which we connected a client and a server. All routing nodes are connected by links with a maximum transmission rate of 1.5 Mbit/s. The propagation delays are indicated by the values next to each link in Fig. 5. Since we are interested in the multipath transmission, the client and server links are faster (10 Gbit/s and a delay of 1 µs) to prevent that they are the bottlenecks. All simulations were repeated 10 times and the broadcast probability p_b was set to 0.1.

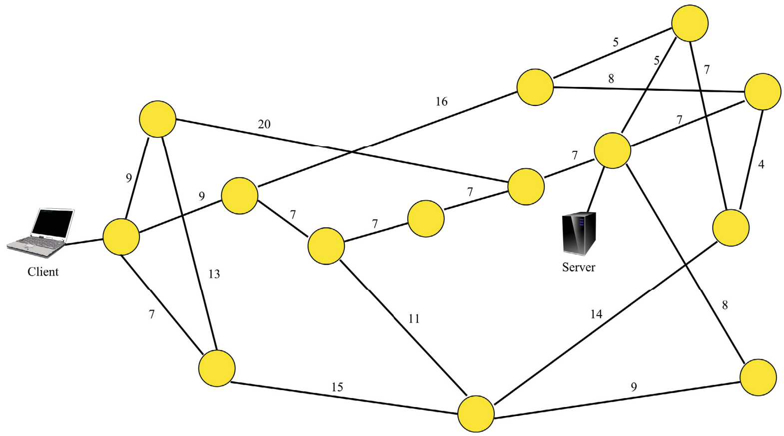

Fig. 5. NSFNET scenario with transmission rate of 1.5 Mbit/s for each link. The numbers indicate propagation delays in milliseconds [9].

Similar to the first scenario, the client sends out a constant stream of Interests. The inter-arrival times of the Interests are 12, 6, 4 or 3 ms. The server has matching Data to the Interests and answers with Data packets of 1.5 kB, resulting in CBR streams of 1, 2, 3 or 4 Mbit/s, respectively. Since any single link only has a data rate of 1.5 Mbit/s, the client can only fully receive the higher CBR streams when more than one link is used. The maximum aggregated link speed of the router to which the client is connected is 4.5 Mbit/s, thus all configured data streams can be transmitted in the ideal case.

Fig. 6 shows the measured throughput at the client for the four different CBR streams and values for α between 0.2 and 100. While the CBR stream with 4 Mbit/s never reaches its full throughput, the three other streams can be transmitted for values of α between 10 and 70.

The surprisingly low throughput of the 1 Mbit/s stream for small values of α can be explained by Fig. 7, which shows the loss rate (ratio of sent Interests to received Data packets) for the different streams. As can be seen in Fig. 7, the loss rate reaches over 10 % for $\alpha \leq 1$, thus resulting in the low throughput.

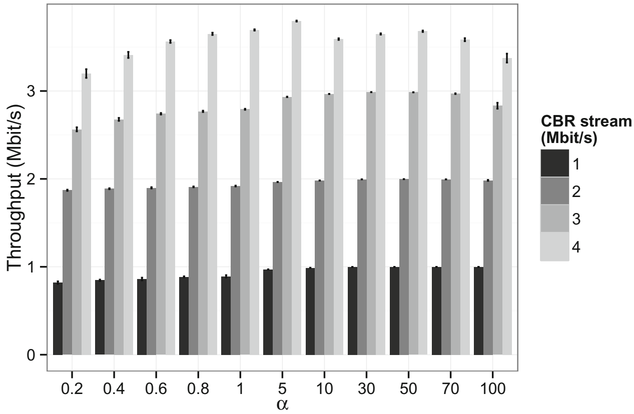

Fig. 6. Measured throughput of CBR streams at the client in NSFNET scenario for different values of α. Error bars indicate 95 % confidence intervals.

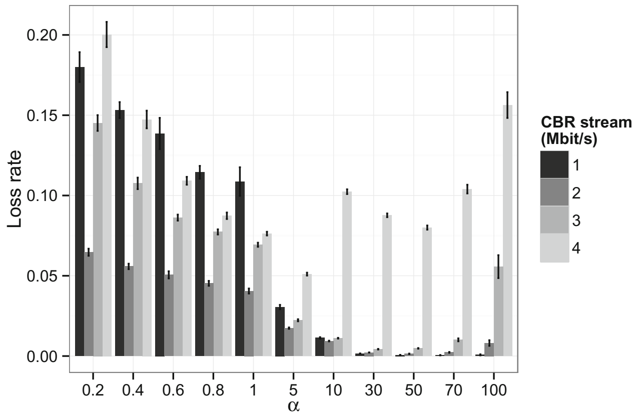

Fig. 7. Average loss rates at client in NSFNET scenario for different values of α. Error bars indicate 95 % confidence intervals.

For $30 \leq \alpha \leq 70$, the loss rates of the three lower CBR streams are very small. The highest CBR stream always experiences loss rates of more than 5 %, explaining the low throughput that is achieved.

Similar to the loss rates, the received delays also decrease for higher values of α (Fig. 8). As mentioned before, the value of α influences the impact the update of the quality indicator of an arrived Data packet. Small values of α therefore result in a highly fluctuating behavior. For $\alpha \approx 50$, we can observe the lowest delays and loss rates.

Fig. 9 shows the Cumulative Distribution Functions of the delay for $\alpha = 50$. As we can see, many packets of the 4 Mbit/s stream suffer from high delays. For the stream of 3 Mbit/s, over 75 % of all packets arrive within 100 ms

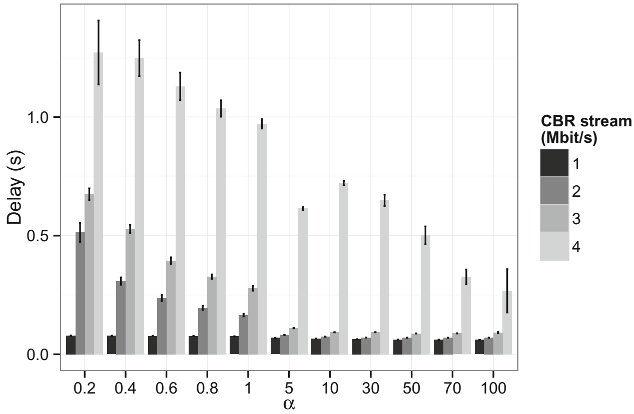

Fig. 8. Average delay of received data packets from server to client in NSFNET scenario for different values of α. Error bars indicate 95 % confidence intervals.

(the minimum delay for the shortest path is slightly above 60 ms). For the 2 Mbit/s stream, which still requires more than one link to achieve its full data rate, nearly all Data packets have a delay of less than 100 ms. The probabilistic routing is therefore able to maintain a low delay for most Data packets for the three lower CBR streams.

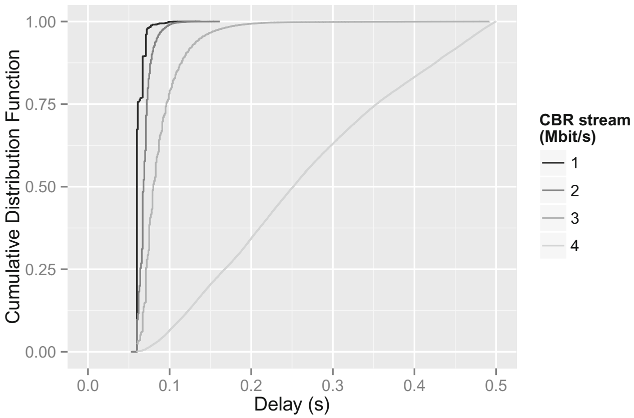

Fig. 9. Cumulative Distribution Functions of delay in NSFNET scenario for $\alpha = 50$ (single run)

6 Conclusion and Future Work

In this paper we presented a probabilistic routing and forwarding mechanism based on the idea of Ant Colony Optimization. Using an OMNeT++ based simulator, we demonstrated that such a mechanism can enable multipath transmission for CCN by using several paths simultaneously to aggregate link bandwidths. This way we can achieve higher transmission rates than any single link could support. Furthermore, the simulations showed that our ant-routing mechanism is capable to adapt to link failures. The time needed for adaptation depends on the parameter α.

Currently, the value for α has to be identified and adjusted manually to achieve good performance for the data transmission. In our future work we plan to investigate how the optimal value for α depends on the network topology and load. Subsequently, we will extend our ant-routing algorithm to dynamically adjust the value of α in order to minimize packet loss and delays and at the same time aim for the highest possible throughput.

References

1. Perkins, C.: IP Mobility Support for IPv4, Revised. RFC 5944 (Proposed Standard) (November 2010)
2. Rafetseder, A., Metzger, F., Stezenbach, D., Tutschku, K.: Exploring youtube's content distribution network through distributed application-layer measurements: a first view. In: Proceedings of the 2011 International Workshop on Modeling, Analysis, and Control of Complex Networks, Cnet 2011, pp. 31–36. ITCP (2011)
3. Rosenberg, J., Weinberger, J., Huitema, C., Mahy, R.: STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs). RFC 3489 (Proposed Standard) (March 2003), Obsoleted by RFC 5389
4. Baset, S.A., Schulzrinne, H.G.: An analysis of the skype peer-to-peer internet telephony protocol. In: Proceedings of 25th IEEE International Conference on Computer Communications, INFOCOM 2006, pp. 1–11 (2006)
5. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking Named Content. In: Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies, pp. 1–12 (2009)
6. Smetters, D.K., Jacobson, V.: Securing Network Content. Tech report, PARC (October 2009)
7. Wang, L., Hoque, M., Yi, C., Alyyan, A., Zhang, B.: OSPFN: an OSPF based routing protocol for named data networking. NDN Technical Report NDN-0003 (July 2012)
8. Rossini, G., Rossi, D.: Evaluating ccn multi-path interest forwarding strategies. Computer Communications 36(7), 771–778 (2013)
9. Di Caro, G.A.: Ant Colony Optimization and its application to adaptive routing in telecommunication networks. PhD thesis, Faculté des Sciences Appliquées, Université Libre de Bruxelles, Brussels, Belgium (November 2004)
10. Dorigo, M., Birattari, M., Stützle, T.: Ant colony optimization. IEEE Computational Intelligence Magazine 1(4), 28–39 (2006)
11. Di Caro, G.A., Dorigo, M.: AntNet: Distributed stigmergetic control for communications networks. Vivek, a Quarterly in Artificial Intelligence 12(3 & 4), 2–37 (1999), Reprinted from JAIR
12. Costa, A.: Analytic modelling of agent-based network routing algorithms. PhD thesis, University of Adelaide (2002)

13. Di Caro, G.A., Ducatelle, F., Gambardella, L.M.: AntHocNet: an adaptive nature-inspired algorithm for routing in mobile ad hoc networks. European Transactions on Telecommunications 16(5), 443–455 (2005)
14. Shanbhag, S., Schwan, N., Rimac, I., Varvello, M.: SoCCeR: services over content-centric routing. In: Proceedings of the ACM SIGCOMM Workshop on Information-centric Networking, ICN 2011, pp. 62–67. ACM, New York (2011)
15. Li, C., Liu, W., Okamura, J.: A greedy ant colony forwarding algorithm for named data networking. In: Proceedings of the APAN – Network Research Workshop (2012)