A Survey of Mobility in Information-Centric Networks: Challenges and Research Directions

Gareth Tyson
Department of Informatics
King's College London,
London, UK
gareth.tyson@kcl.ac.uk

Nishanth Sastry
Department of Informatics
King's College London,
London, UK
nishanth.sastry@kcl.ac.uk

Ivica Rimac
Bell Labs
Alcatel-Lucent
Stuttgart, Germany
ivica.rimac@alcatel-lucent.com

Ruben Cuevas
Department of Telematic Engineering
Universidad Carlos III de Madrid
Madrid, Spain
cuevas@inv.it.uc3m.es

Andreas Mauthe
School of Computing and Communications
Lancaster University
Lancaster, UK
a.mauthe@lancs.ac.uk

ABSTRACT

In essence, an information-centric network (ICN) is one which supports a content request/reply model. One proposed benefit of this is improved *mobility*. This can refer to provider, consumer or content mobility. Despite this, little specific research has looked into the effectiveness of ICN in this regard. This paper presents a survey of some of the key ICN technologies, alongside their individual approaches to mobility. Through this, we highlight some of the promising benefits of ICN, before discussing important future research questions that must be answered.

Categories and Subject Descriptors

A.1 [**General Literature**]: Introductory and Survey; C.2.1 [**Computer Systems Organization**]: Computer Communications Networks

General Terms

Design, Documentation

Keywords

Information-Centric Networking, Mobility

1. INTRODUCTION

It has recently been observed that the Internet has evolved to be an infrastructure primarily used for content distribution [1][2]. In response to this, researchers have proposed re-architecting the Internet to be *information-centric*[1]. This involves replacing or augmenting the existing host-to-host routing infrastructure with a content-based one. Thus, applications generate content requests, which are routed using unique content identifiers to optimal sources. Therefore, in contrast to existing 'bolt-on' approaches such as content distribution networks (CDNs), an information-centric network (ICN) makes content an explicit first-class entity. One proposed benefit of this is superior *mobility* support; in this context, mobility involves the physical and/or topological re-location of a device in regards to its access network.

Of particular interest is mobility during active communications (e.g. audio or video streaming), since this can seriously affect a user's quality of experience. For the sake of clarity, we present a simple example: imagine a train containing a number of passengers. Some of them are watching video streams, hence acting as content consumers, whilst others are audio conferencing, thus acting as both consumers and providers. As the train moves, the access point (e.g. base station) of each device will change, requiring network re-configuration to reflect the device's new location. Clearly, to ensure a suitable quality of service, this re-configuration must take place in such a way that ensures (near) seamless communication. Ideally, mobility should therefore not result in lost data or extended periods of disconnection. A number of mobile architectures exist, however, in this paper, we focus on two general types of mobile networks:

- *Cellular Networks*: Mobile nodes connect to (multiple) Access Points (APs), changing their AP as their position changes. Note that an AP can be represented by a Base Station, a WiFi AP etc.
- *Mobile Ad-Hoc Networks (MANETs)*: Mobile nodes collaborate to build shared routing information, allowing message exchange between each other, as well as potentially external parties (via an AP).

[1]Synonymous terms include content-centric and data-oriented.

NoM'12, June 11, 2012, Hilton Head Island, SC, USA.

Handling the above mobility has been a significant challenge that has plagued location-oriented networks (e.g. IP) due to the complexity of managing changes in a node's geographic and topological location — a situation that was never anticipated when the principles of packet networking were first defined. Various 'bolt-on' protocols have been defined to enable mobility, including Mobile IP and the Host Identification Protocol; however, often these complicate existing protocols and avoid handling the core underlying issues. In contrast, from the outset, ICN designs have tried to move away from such location-oriented dependencies. However, in many ways, the core problems relating to this are not resolved; instead, the goal posts are simply shifted. As such, there are still a number of important research challenges that remain in this domain.

This paper aims to present a short survey of representative ICN designs' approaches to mobility. Through this, we extract some of the core benefits of deploying ICNs in mobile environments before exploring some of the remaining challenges that need to be address when moving forward in the field.

2. MOBILITY IN ICN

We define a mobile ICN as an ICN that supports, at some element in the network path, consumer and/or provider[2] mobility. *Consumer mobility* allows consumers to re-configure their network location without disrupting connectivity, whilst *provider mobility* allows sources to re-locate without disrupting content availability. This section discusses some representative ICN designs, alongside their mobility management.

2.1 DONA

DONA [3] proposes introducing ICN in the form of a replacement (or supplement) to DNS. Content names are of the form *P:L*, where P is the cryptographic hash of the publisher's public key and L is a label that identifies the content. DONA requires each domain to deploy servers called Resolution Handlers (RH) that index content stored by authorised storage points. RHs are then structured into a tree topology that represents the BGP topology of the network. Lookups are performed by querying a consumer's local RH; if no reference is found, the query is forwarded up the tree until a source is discovered; an out-of-band delivery is then established by the source (over IP). DONA therefore follows an *early-binding* approach in which providers register $identifier \rightarrow locator$ mappings that must be resolved before deliveries can be performed.

DONA handles consumer mobility by simply changing a host's RH to that of the new network (this could be done through DHCP). If necessary, any existing content requests can then simply be re-issued to the new RH to locate the new optimal source. Unlike some other designs, however, DONA relies on out-of-band deliveries of content; although not stipulated, this would likely take place over TCP. Consequently, this would requires session re-establishment after consumer re-location (either to re-establish the same connection or one with a newly selected source), thereby complicating the process. Provider mobility would not create too great a challenge either because hosts could simply re-register their content with the new network's RH. It is important to note, however, that this would only maintain content availability for *new* requests — *existing* requests would either need to be re-sent or continued though a mechanism similar to Mobile IP. This is a notable limitation of approaches that use early-binding. An open question is therefore exactly what the delay and overhead involved in this would be. However, any delays associated with this process could perhaps be mitigated by cached copies in other networks.

[2]It is worth noting that this can also include intermediate routers.

2.2 CCNx

CCNx [2] is another prominent design. Content naming is based on hierarchical names encoded with the publisher, content identifier, data digest, version number and segment information (segments are equivalent to packets). Routing is performed using similar algorithms to current IP infrastructure, utilising longest prefix matching with hierarchical aggregation to ensure scalability. In CCNx, a content request is issued by sending an Interest packet, which is routed through the network to the closest instance of the content. Subsequently, if available, the source responds with a Data packet, which follows the reverse path back to the requester using 'breadcrumbs' left in a Pending Interest Table on each router. In contrast to DONA, CCNx therefore follows a *late-binding* approach in which each content request is only resolved to a specific location during the routing process (i.e. at the last hop). This means there is no direct $identifier \rightarrow locator$ mapping that must be handled before the request can take place.

Consumer mobility in CCNx is intrinsic due to its consumer-driven nature. When a consumer re-locates, it can simply re-issue any previously sent Interest packets that have not been satisfied yet. This can occur seamlessly because there is no need to perform any new registrations etc. Through this, is has been shown that CCNx can still handle up to 97% of requests during high mobility [4]. Provider mobility, however, is far harder as naming in CCNx is based on organisation names, which often map to relatively static locations. Consequently, to ensure routing aggregation, it is better for organisations to move as one, rather than having individual items of content move independently. Either way, however, it is necessary for routing updates to be distributed, leading to convergence delays. However, CCNx inherently supports multi-sourcing for consumers (e.g. from multiple caches), thereby mitigating these effects of provider re-location.

2.3 NetInf

NetInf proposes the use of a Name Resolution (NR) service rather than the routing-based approach taken by CCNx. As such, providers publish Information Objects (IOs) alongside their locator(s) to the NR service for later discovery by consumers; this is clearly another example of *early-binding*. Although IOs could be of any size, it is likely that they will be on a per-content granularity (rather than in segments). The NR service is underpinned by a Multi-level DHT (MDHT), allowing global content lookups on flat identifiers, whilst also supporting local resolution. Unlike many other designs, NetInf also supports content searching based on meta-data. Consumers therefore perform lookups, which are responded to with either a list of potential sources or a selected optimal source. Content can then be accessed using any supported transport protocol, allowing in-path caching (e.g. [5]).

Consumer mobility in NetInf is easily achieved through its indirection between identifiers and locators. The exact details of this vary based on the chosen locator selector mode [6]. In the requester-controlled mode, a consumer is provided with a list of potential sources, thereby allowing a node to select a new optimal source following re-location. In contrast, the MDHT-controlled mode results in a consumer only receiving a single source on each request, mandating a re-located node to contact the NR service again. Regardless of this, both modes should enable mobility, assuming fast lookups. Like DONA, however, it is likely that deliveries will take place over connection-oriented protocols, requiring session re-establishment. Provider mobility is more difficult as it requires the NR service to be updated; however, it is claimed that updates can be scalably handled.

2.4 PSIRP

PSIRP[7] revolves around the concept of publish/ subscribe (a closely linked paradigm to ICN). Providers publish content to the network, which can then be subscribed to by consumers using a human-readable Application Identifier (AI). These AIs can then be mapped to a Rendezvous Identifier (RI), which is used by a distributed Rendezvous System to generate a Forwarding Identifier (FI). Importantly, these are complemented by Scope Identifiers (SI) that restricts access to the content. Both RIs and SIs use similar cryptographic naming to DONA for securing content. After this, using bloom filter source routing, the FI can be used to route the content packets through the network (c.f. LIPSIN [8]).

Consumer mobility in PSIRP is relatively straight-forward to achieve. When a consumer re-locates, it simply re-subscribes to the content being accessed. This results in a new FI being computed for the host's new location. Clearly, the efficiency of consumer mobility is therefore dependent on the speed at which new FIs can be generated, however, it is claimed that PSIRP can lead to 50% less packet loss during mobility compared with Mobile IPv6 [9]. Provider mobility would likely be more complex as this would require updating routing information; unfortunately, it is unclear exactly how costly this process would be. It is likely that it could not be achieved at line speeds, however, replicated sources of the content could mitigate the effects of re-location.

2.5 Juno

Juno [10] proposes the placement of information-centric functionality in the middleware layer. Content is based on flat self-certifying identifiers that are indexed on a DHT called the Juno Content Discovery Service (JCDS). Like NetInf and DONA, content locators are resolved rather than routed to. Unlike other designs, however, Juno focusses on achieving backwards-compatibility by performing software re-configuration to interoperate with any sources that might offer the content, regardless of their delivery protocols. To achieve this, Juno attempts to discover as many content sources as possible by also probing third party indexing services such as Gnutella. By utilising dynamically attachable protocol plug-ins, each of these sources can then be interacted with to retrieve the content (e.g. if a HTTP source is located, a HTTP plug-in is dynamically attached to retrieve the content). Beyond this, it also proposes extending the ICN interface to include more sophisticated preferences regarding such things as performance, reliability, security, monetary cost etc. This interface is termed *delivery-centric*, which allows these preferences to be stipulated (and updated) on a per-request basis. These preferences are then used to shape the selection made by Juno regarding the source(s) used to access the content through.

Consumer mobility is easily achieved in Juno by simply re-selecting sources after host re-location. This can be done locally as Juno keeps a full list of sources from the resolution process. As this takes place in the middleware layer, however, it is likely that hand-off delays will be larger than in designs such as NetInf or PSIRP. Provider mobility is similarly possible by simply updating the JCDS; like NetInf, the performance of this depends on the DHT. This delay, however, is mitigated by Juno's inherent support for multi-sourcing.

3. THE POTENTIAL OF MOBILE ICN

The previous section has highlighted some prominent ICN designs, alongside their approaches to supporting mobility. This section now presents some key benefits of deploying ICN in mobile networks. Importantly, we aim to provide a conceptual overview of these benefits, rather than focussing on aspects supported by current designs.[3]

3.1 Host Multihoming

Within IP, many applications and protocols make the assumption that a host is uni-homed. Over the years, this assumption has been manifested in various infrastructural and software design decisions. For instance, by definition, a HTTP GET request is received over a single TCP connection, from a single source address. As such, it is difficult to exploit multiple potential network interfaces that might be available when using HTTP. Multi-path TCP [11] is one approach to addressing this, however, this is yet to be widely deployed.

In contrast, an ICN does not bind a particular flow with a specific interface (i.e. individual network address). This is because an interface is very much a location-oriented concept, which can be abstracted away from in a content request/reply model. Specifically, an application only expects the network stack to reply with a unique data item, without needing information regarding which interface was used. Consequently, in an ICN, a request can easily be multiplexed over a number of different interfaces without the need for application-level awareness. This means that applications (both providers and consumers) running on a multihomed ICN node can seamlessly exploit these different interfaces without needing to understand which interface has actually been used.

3.2 Abstraction from Network Addresses in Applications

Currently, many mobility mechanisms also attempt to maintain consistency in the node's network address. This is vital for many applications that may utilise a node's IP address for long-term usage. A typical example is BitTorrent, which will see a node's IP address being registered with a tracker for future discovery. Mobile IP, for instance, introduces the concept of a Home Agent to allow hosts to change their physical address, whilst still maintaining a constant public address. This, however, is frequently criticised

[3]Not all designs presented in Section 2, for instance, support all these features.

due to the costs of tunnelling data. Unfortunately, the alternative involves placing greater intelligence in applications to make them aware of mobility, thereby allowing them to update their location information. This could occur at the consumer end but also, more worryingly, at the provider end. Whilst this approach can increase performance, it results in a greater load on the developer.

In contrast, an ICN does not force applications to take on location-oriented information. Instead, it detaches the application from such concerns. This allows the application to abstractly publish or consume content, without the need to store (or even know) its own network-layer address. In essence, it promotes content, which is already an explicit application-layer element, to an explicit network entity as well, thereby requiring the application to maintain only knowledge that does not deviate from its own traditional knowledge base.

3.3 Removal of Connection-Oriented Sessions

A key problem with mobility in location-oriented networks is their frequent dependency on connection-oriented protocols. In fact, ≈90% of all Internet traffic is TCP-based [12]. This prevalence has arisen because often sessions must be maintained to ensure that interactions take place in a proper fashion. This is largely due to the inability of both interacting parties to explicitly understand what *should* be sent and received at any given time. Without this, reliability, for example, cannot be achieved. Consequently, supplementary connection-oriented protocols are required (e.g. TCP). This means that mobility can often require the re-establishment of the session so that both parties are aware of the up-to-date network addresses, as well as any pertinent parameters. Often this is not a significant challenge when dealing with non real-time applications, however, this can cause severe disruption for live video or audio streams.

In contrast, an ICN does not necessarily suffer from such constraints. Instead, communications are made explicit at the network level — when a consumer generates a request for a given item of content, it knows exactly what it should receive in return. For instance, when a host generates a request for Content C, the kernel knows precisely what must occur to validate a successfully executed network interaction, i.e. it must receive C. Importantly, this ability does not require cooperation from the provider; as such, the use of persistent sessions established between both parties become less necessary.[4] This removal of connection-oriented protocols therefore means that the re-location of a host does not necessitate the re-establishment of a connection. A cost that can be significant (e.g. in long fat networks).

3.4 Scoping of Content and Location

Currently, consumers are generally identified by their location (IP address). Often, however, this is incorrectly used for scoping purposes, i.e. information is interpretted from the address that should not be. For instance, the BCC iPlayer service can only be accessed from UK IP addresses; consequently, this makes mobility difficult for legitimate UK residents who may temporarily utilise connectivity abroad. A similar problem emerges when attempting to utilise IP addresses for selecting optimal content replicas (e.g. from a CDN). This is because (at request time) the CDN will utilise a node's location to resolve an optimal source, even though the node may later change its location.

In contrast, an ICN makes an explicit split between location and identifier, thereby offering an elegant mechanism by which nodes can re-locate without having to artificially manipulate their persistent address. Thus, a node's location can seamlessly change whilst still maintaining a consistent identifier. Consequently, the above situations can be avoided; the scoping of iPlayer's access control would not require the consumer's physical location, whilst the scoping of a CDN's name resolution could seamlessly change as the node re-locates.

3.5 Resilience through Replication

Information exchange in a location-oriented network is usually based on some concept of location (e.g. a URL). As such, generally, access to an item of content is dependent on achieving a successful interaction with that stipulated location. This means that any failure in the path between the consumer and provider will result in serious problems. MANETs, for example, are particularly vulnerable to this because paths can frequently fail due to network partitions [13]. Unfortunately, there is little that can be done to address this, unless complex (layer-7) redirections can somehow be performed.

In contrast, an ICN does not bind content to a specific location; it, in fact, liberates content from such a burden and makes it an interpretable network level entity. A key advantage of this is the ability to perform seamless replication and in-router caching, potentially allowing local copies to be retrieved. In essence, this is achieved through content multihoming (as opposed to host multihoming). On the one hand, this can significantly improve performance [14]. However, beyond this, the effects of failures in a mobile network can also be mitigated [15]. This is because ICN caching can increase the number of potential end points for each request, thereby adding redundancy in the face of failures.

4. RESEARCH CHALLENGES

This section now explores some of the core research challenges that remain in the area of mobile ICN. Many challenges extend to ICN as a general technology (a good overview can be found in [16]), however, many interesting mobile specific aspects emerge when looking at this area.

4.1 Provider Mobility

Broadly speaking, consumer mobility is a well handled phenomenon due to the consumer-driven nature of most ICN designs. However, a larger challenge is maintaining routing consistency during provider mobility. This is because whenever a provider re-locates, it is clearly necessary to update (global) locator information. This is heavily exacerbated by the obvious increase in the number of content objects when compared to hosts (an ICN must be able to deal with at least 10^{12} objects). The effects of this perhaps can be mitigated via caching and replication but unpopular content is still likely to suffer greatly if high speed provider hand-offs cannot be achieved.

The precise focus of this challenge varies with the different naming and discovery techniques employed. CCNx, for instance, uses hierarchical naming and route aggregation to improve scalability. However, because naming is based

[4]Although, benefits such as lower overheads (e.g. by performing per-session authentication) and improved Quality of Service could perhaps be better achieved through sessions.

on organisations (which are often statically positioned), this creates significant challenges when dynamically re-locating nodes to other organisations. This is because it clearly undermines the hierarchy of the address space; in fact, using any content that is cached off-path introduces a similar challenge. Unfortunately, line-speed switching relies heavily on this aggregation, meaning that mobility will introduce significant scalability challenges when using packet-sized objects. In this routing paradigm, the precise mechanisms for mobility and content multihoming are yet unclear, however, it is clear that long-term re-locations could significantly undermine routing scalability. A good starting point looking at this area can be found in [17].

In contrast to CCNx's approach, many challenges also arise in early-binding systems such as NetInf and Juno. This is because any provider mobility must be reported to the resolution service; clearly, high levels of mobility could result in phenomenal loads. As such, hand-off delays could be significant. In [6], the authors discussed the handling of 1% of churn in registrations, however, mobility could increase this greatly. An interesting solution proposed is to chain locators together in the resolution service to point to the latest locations. This, however, creates undesirable overheads in a similar way to Home and Foreign Agents in Mobile IP.

From the above, it is clear that achieving provider mobility is an important next step. Particularly, it is important that mechanisms are put in place to ensure that real-time content is not compromised during hand-offs. Interestingly, this could result in a variety of routing configurations for each item of content, based on a provider's content and mobility characteristics.

4.2 Pairwise Path Routing

Due to the (attempted) removal of location from the concept of ICN, many approaches utilise pairwise hop-by-hop knowledge to ensure that data can find its way back to consumers. Unfortunately, however, this approach can result in problems in the face of route dynamics, as it requires knowledge either about or in the core of the network. This is particularly the case in late-binding approaches that require request and response packets to follow the same path. Consequently, this means making a truly location-agnostic mobile ICN (e.g. sending a request from one location but receiving the reply at another) is yet to be achieved.

PSIRP, for instance, encodes per-hop source routing in its Forwarding Identifier; however, if the computed route changes due to mobility (e.g. in a MANET), then this per-hop knowledge will become out-dated. Attempts such as encoding redundant virtual links [8] cannot even address this due to the potential mobility of these redundant alternatives. A similar example is that of CCNx, which introduces the concept of using breadcrumbs along the paths of Interest packets to allow reverse routing of Data packets. Once again, however, if any pairwise links between routers change, the Data packet will be lost. Clearly, if ICNs are to be deployed in dynamic mobile environments, this is an extremely important research issue to address.

4.3 Discovering Local Cached Content

As previously discussed, one particularly beneficial property of mobile ICNs is the increased resilience attained through replicating content throughout unreliable networks (e.g. via caching). Unfortunately, however, this creates significant content discovery challenges in MANETs due to the far higher level of route churn that will be observed when compared to performing host-to-host routing.

The specifics of this challenge will vary with the particular content binding approach employed. CCNx would likely suffer heavily due to the increased overhead of maintaining routing information; this would be further exacerbated by the complexity of routing on hierarchical identifiers in an unstructured (changing) topology. In contrast, approaches such as PSIRP or DONA would only be able to exploit this local caching if the local Rendezvous Point or Resolution Handler were accessible, thereby negating the greater potential resilience.

A prominent research challenge that therefore remains is to build naming, resolution and routing schemes that can handle this type of unpredictable re-location of content. Perhaps a more worrying problem arises, however, when looking at content that is unpopular, which doesn't benefit from caching (i.e. accessed only once). For instance, when dealing with smaller MANETs (e.g. <300), Varvello et al. [15] found that the performance benefits of using more sophisticated structured routing protocols (e.g. GHT) were dwarfed by their overheads due to the presence of unpopular content.

4.4 Request Staleness

Another interesting property of ICN mobility emerges from the connectionless approach taken in some implementations, notably CCNx. This is because it becomes difficult to revoke previously issued content requests. For example, during mobility, CCNx can leave a potentially large number of Interest packets in the network, each following breadcrumbs to out-of-date locations (via routers' Pending Interest Tables). To achieve high performance, the window size used could be very large; in TCP, for instance, window scaling is frequently used, potentially leading up to window sizes of approximately a gigabyte [12]. As such, to utilise this data, it would be necessary to use some form of forwarding information to redirect the content. This, however, would involve greater knowledge of locations and topology; something which ICNs try to avoid. It is less clear how such staleness will affect other systems such as DONA and NetInf because the content is not generally packet sized, however, any approaches using connectionless windowing will need to address this problem.

4.5 Security and Privacy

Security in open mobile systems has been a long-term challenge, with many possible attacks. Unsurprisingly, a key research challenge is therefore handling these types of concerns in a mobile ICN. In principle, ICNs primarily focus on securing the content itself, i.e. guaranteeing a content item is what it claims to be. This is generally achieved through naming; Juno, for instance, uses data hashes for names. This, however, introduced severe privacy issues, as it requires content names to be included in requests. The level of concern will vary based on the individual setup of the mobile host. Connectivity provided over 802.11 could, for instance, offer an acceptable level of privacy for users who trust their access points. However, those accessing content through MANETs would reveal a huge amount about themselves to their neighbouring peers. This is particularly challenging in densely connected networks that easily permit packet sniffing. Further, attempts to address privacy problems can even be undermined when operating in such environments. For example, Arianfar et. al. [18] propose segmenting content into multiple blocks, which can then be retrieved (and combined) by many cooperating nodes. If, for instance, malicious nodes could acquire a sufficiently large-view of a wireless network to observe this behaviour, then privacy could be weakened. To make ICNs a viable technology for deployment, it is evident that such issues will need to be addressed.

Beyond this, other security issues relating to such things as routing are yet to be adequately explored. In theory, any node can publish the ability to serve an item of content, thereby empowering malicious nodes to manipulate routing. This can be an issue in traditional fixed infrastructure; however, it is particularly prevalent in networks such as MANETs, which have extremely open routing policies. A good example of this is black hole routing in which malicious nodes advertise routes to all content, thereby allowing them to receive (and destroy) all requests. Such capabilities could cripple an open mobile network. Alongside this, a variety of alternate security problems exist; a good taxonomy of such attacks can be found in [13].

5. SUMMARY

This paper has presented a survey of ICN solutions, alongside their individual approaches to supporting mobility, looking at promising benefits as well as future challenges. As discussed in Section 3, it is evident that many benefits could be gained through the information-centric paradigm, offering strong motivation for future work on these remaining issues. It is important, however, to note that these are not necessarily weaknesses in ICN. Instead, they are exciting topics deserving future attention. Such promising research has already begun to develop (e.g. [19][20]), however, we feel it is an important topic that the ICN community should begin to focus on.

ACKNOWLEDGMENTS

We would like to acknowledge the EPSRC Project: A Framework for Innovation and Research in MediaCityUK (FIRM). We would also like to thank Yehia El-khatib, Simon Miles and Peter McBurney for their interesting discussions on the topic of this paper.

6. REFERENCES

[1] H. Schulze and K.Mochalski, "Ipoque internet study," tech. rep., ipoque GmbH, 2008/2009.

[2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *In Proc. 5th ACM CoNEXT*, 2009.

[3] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 181–192, 2007.

[4] J. Wang, R. Wakikawa, and L. Zhang, "DMND: Collecting data from mobiles using named data," in *Proc. IEEE Vehicular Networking Conference (VNC)*, 2010.

[5] J. Ott, K. Budigere, P. Sarolahti, and C. Perkins, "Poor man's content centric networking (with TCP)," Tech. Rep. Aalto-ST 5/2011, Aalto University, 2011.

[6] M. D'Ambrosio, C. Dannewitz, H. Karl, and V. Vercellone, "MDHT: A hierarchical name resolution service for information-centric networks," in *Proc. ACM SIGCOMM Workshop on ICN*, 2011.

[7] D. Lagutin, K. Visala, and S. Tarkoma, "Publish/subscribe for internet: PSIRP perspective," *Valencia FIA book*, vol. 4, 2010.

[8] P. Jokela, A. Zahemszky, C. Esteve Rothenberg, S. Arianfar, and P. Nikander, "LIPSIN: line speed publish/subscribe inter-networking," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 195–206, 2009.

[9] N. Fotiou, P. Nikander, D. Trossen, and G. C. Polyzos, "Developing Information Networking Further: From PSIRP to PURSUIT," in *Proc. Intl. Conference on Broadband Communications, Networks, and Systems (BROADNETS)*, 2010.

[10] G. Tyson, A. Mauthe, S. Kaune, P. Grace, and T. Plagemann, "Juno: An adptive delivery-centric middleware," in *Proc. 4th Intl. Workshop on Future Media Networking (FMN)*, 2012.

[11] H. Han, S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley, "Multi-path TCP: A joint congestion control and routing scheme to exploit path diversity in the internet," *IEEE/ACM Trans. Netw.*, vol. 14, pp. 1260–1271, Dec. 2006.

[12] J. Wolfgang and S. Tafvelin, "Analysis of internet backbone traffic and header anomalies observed," in *Proc. 7th ACM SIGCOMM Conference on Internet measurement (IMC)*, 2007.

[13] D. Djenouri, L. Khelladi, and A. Badache, "A survey of security issues in mobile ad hoc and sensor networks," *IEEE Communications Surveys & Tutorials*, vol. 7, no. 4, 2004.

[14] G. Tyson, S. Kaune, S. Miles, Y. El-Khatib, A. Mauthe, and A. Taweel, "A trace-driven analysis of caching in content-centric networks," in *Submitted to Intl. Conference on Computer Communication Networks (ICCCN)*, 2012.

[15] M. Varvello, I. Rimac, U. Lee, L. Greenwald, and V. Hilt, "On the design of content-centric MANETs," in *Proc. Intl. Conference on Wireless On-Demand Network Systems and Services (WONS)*, 2011.

[16] A. Ghodsi, T. Koponen, B. Raghavan, S. Shenker, A. Singla, and J. Wilcox, "Information-centric networking: Seeing the forest for the trees," in *Proc. ACM Workshop on Hot Topics in Networks*, 2011.

[17] A. Ghodsi, T. Koponen, J. Rajahalme, P. Sarolahti, and S. Shenker, "Naming in content-oriented architectures," in *Proc. ACM SIGCOMM Workshop on ICN*, 2011.

[18] S. Arianfar, T. Koponen, B. Raghavan, and S. Shenker, "On preserving privacy in content-oriented networks," in *Proc. ACM SIGCOMM Workshop on ICN*, 2011.

[19] M. Meisel, V. Pappas, and L. Zhang, "Ad hoc networking via named data," in *Proc. MobiArch*, 2010.

[20] M. Amadeo and A. Molinaro, "CHANET: A content-centric architecture for IEEE 802.11 MANETs," in *Intl. Conference on the Network of the Future (NOF)*, 2011.