Beyond Content Delivery: Can ICNs Help Emergency Scenarios?

Gareth Tyson, Eliane Bodanese, John Bigham, and Andreas Mauthe

Abstract

Natural disasters are becoming more and more prominent in our world today, increasing by over 400 percent in the last 20 years alone. To overcome the challenges brought about by such situations, it has been proven vital to ensure continued and effective communication. Unfortunately, however, this has often been difficult with network failures commonplace, particularly during large-scale emergencies, e.g. earthquakes. In this article we explore the potential of information-centric networks (ICNs) to provision highly resilient communications during disaster scenarios. We first provide background to the area, before highlighting the key characteristics of ICNs that are well suited to augmenting resilience in the domain. Following this, we explore remaining challenges of note, concluding that, despite its potential, several key obstacles must be overcome before a truly resilient ICN can be realized.

Natural disasters are becoming more and more prominent in our world today, increasing by over 400 percent in the last 20 years [1]. In 2009 alone, 335 natural disasters were recorded, causing $41.3 billion in economic damages and affecting approximately 120 million people. This trend is constantly expanding, with over 375 million people predicted to be affected by climate-related disasters per year by 2015 [2]. Beyond this, major social issues such as terrorist attacks (e.g. 9/11) continue to impact thousands. In light of this, it seems likely that nations will have to make increasing efforts to manage the challenges brought about by such devastating and unpredictable events.

Many technologies have been recently proposed to help in these circumstances. However, the majority rely on reliable underlying communications infrastructure. Take, for example, the Integrated Public Alert & Warning System, which was developed to integrate various existing warning systems in the U.S. It uses a variety of communications technologies (e.g. radio, Internet) to offer early warnings during disasters. Although sophisticated, it depends on communications infrastructure that is susceptible to damage. For example, natural disasters can result in days of network disruption due to physical damage (e.g. broken links) and power failures [3]. These problems are, of course, inescapable and therefore many working in the area have turned their attention to ensuring such infrastructure can always continue operation [4]. However, as of yet no solutions have been presented that have stopped the repeated infrastructural failings witnessed in the midst of disaster situations.

Gareth Tyson, Eliane Bodanese, and John Bigham are with Queen Mary University of London.

Andreas Mauthe is with Lancaster University.

Interestingly, when inspecting the types of communication that take place during disasters, we find that many are of an information-centric nature, e.g. warning dissemination, social messaging to loved ones, and retrieval of disaster information. However, due to the need for end-to-end paths in host-centric networks (HCN), these services frequently fail [3, 5], even when information sources are nearby with perfect reachability. We argue that Information-Centric Networks (ICNs) have the potential to redress this balance by enabling nodes to exploit local information sources. Further, its receiver-driven nature with self-securing information objects (IOs) offers the perfect mechanism to empower individuals in unpredictable scenarios.

In this article we explore the potential of ICNs to underpin future resilient information dissemination in emergency situations. Whereas recent ICN work (e.g. [6–8]) has focused on building systems to optimize content delivery, here we focus on sustaining a much lower level of service. We argue that, through the decentralized storage and management of information, ICNs are far more capable of serving the requirements of disaster scenarios. Specifically, we believe a shift should be made from traditional path resilience to information resilience, ensuring that all who need information can gain access to it.

Background

Information-Centric Networks

In essence, an ICN is a network with the sole purpose of delivering information (e.g. text, videos). As such, an ICN exposes a publish/subscribe style abstraction unlike the existing Socket API. Through this, information objects (IOs) are published to the network, allowing consumers to subsequently request them. The purpose of the network therefore becomes to bind consumers and publishers together, rather than to simply route packets between source and destination addresses. Typically, this is done by either offering information-based

0890-8044/14/$25.00 © 2014 IEEE

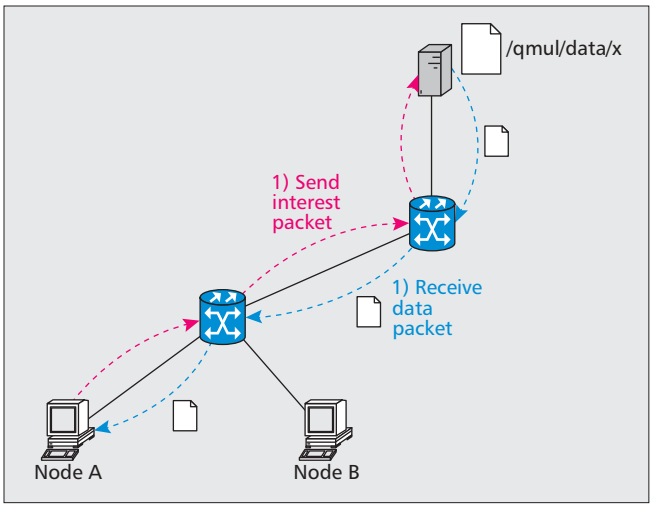

Figure 1. NDN routes Interest packets via (layer-3) routers to sources that reply with Data packets. Each router maintains a cache that can respond to Interest packets along the intermediate path.

routing of requests to sources (e.g. Named Data Networking [7]) or via a resolution service that maps requests to optimal sources (e.g. PURSUIT [6], NetInf [9]). Generally, all of these systems offer:

- Information addressing rather than host addressing.
- In-network information retrieval, allowing information to be accessed from the network without needing predetermined locators.
- Information security, allowing information to be independently verified, without needing to verify its source.

A particularly related system is MobilityFirst [10], which offers information-centric delivery specifically for mobile environments. It uses a global resolution service, similar to NetInf, to perform late binding of mobile hosts to content and services. To improve resilience and mobility support, it also integrates several functions, including push-based delivery, delay-tolerance, and hop-by-hop transport control. Examples of how all these systems work are presented in Figs. 1–4. In the rest of the article, our discussion centers on the use of the technologies, protocols, and concepts proposed by all of these key systems.

Little work has been performed specifically on ICN resilience as of yet, although recently a joint FP-7 NICT project has started, GreenICN, looking at using ICNs to assist in emergency situations. Antikainen [11] investigated various approaches to reliability in ICNs, showing how a cache-and-forward network could improve reliability by offering temporary storage of information on a hop-by-hop basis. In fact, a key goal of the caching facilities in NDN is packet recovery following loss [7]. NREP [12] is a mobile ICN, targeted at disaster scenarios, that allows message replication and routing based on priorities and geographical location. Oh *et al.* [13] also investigated the use of information-centric ad hoc networking in emergencies. These two systems highlight a particular strength of ICN, whereby hastily deployed ad hoc networks can be created without the need to secure host communications (just the information). As of yet, however, these concepts have not been tested in the extremely challenging situations that arise from disasters.

Emergency Networking

A variety of research efforts have attempted to understand and investigate techniques to maintain communications during emergencies. This has included studies into the behaviour of networks during disasters [3] as well as operator responses [5]. Worryingly, these have highlighted that even localized disasters can have a huge global impact (e.g. 9/11, the Taiwan Earthquake).

This work has shown that resilience typically falls to routing protocols (e.g. BGP) to handle. Unfortunately, it has also shown that these can often fail to provide sufficient uncongested links to handle the demand [5]. In fact, devastating physical damage often makes route re-convergence impossible [14]. Even explicitly designed resilience architectures (e.g. D^2R^2+DR [4]) are restricted by such physical network limitations. To handle these situations, alternative research efforts have focused on emergency ad hoc networking that can be rapidly deployed without requiring pre-existing infrastructure [15]. Particularly important are the various flavors of mobile delay tolerant networks (DTNs). In such architectures, nodes pass messages between each other, which are stored (temporarily) until a suitable next hop is found. This allows messages to be retained during disruptions — clearly, an ideal technology for disaster situations. Previous work has already identified the synergies and benefits of ad hoc DTNs and ICNs [15], offering strong motivation for its further exploration.

Emergency Applications

In this article we conjecture that a shared characteristic among the majority of emergency scenarios is that many of the fundamental communications required are *information-centric*. Users do not wish to communicate with individual hosts but, rather, they wish to exchange units of information. Common examples include:

- **Emergency warnings:** Early notification of a disaster can heavily mitigate human causalities by instructing people to avoid certain areas.
- **Evacuation Maps:** Civilians in distress, particularly those in unfamiliar areas, might wish to gain access to maps, guiding them to safety.
- **Beacons:** Civilians in distress would often wish to disseminate beacons requesting assistance from those nearby.
- **Messaging:** Civilians would often wish to discover the status of loved ones. Messaging services (e.g. Twitter) would allow civilians to share their status.

All of the above applications involve the generation of self contained IOs. However, in our current host-centric network (HCN), all of these *information-centric* applications rely on establishing end-to-end paths with pre-determined hosts (that may fail during the disaster). This ignores the fact that, fundamentally, none of them actually need to do this; rather, they

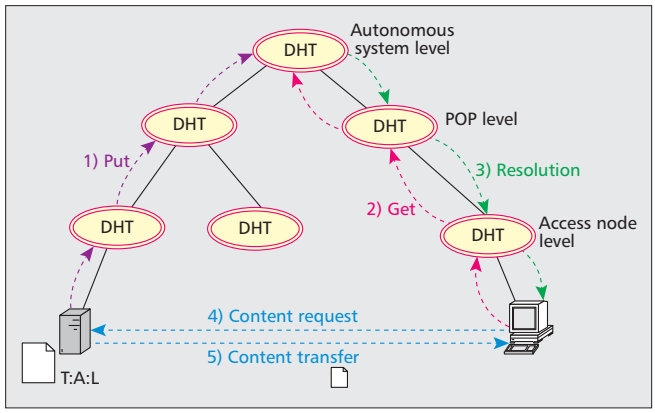

Figure 2. NetInf's global deployment uses hierarchical distributed hash tables to index information sources. Consumers use the index to lookup sources of information and then sent requests directly to the source.

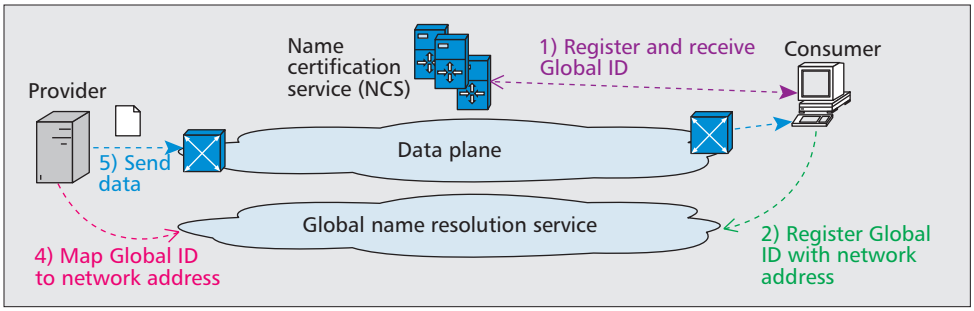

Figure 3. MobilityFirst uses a global name resolution service to map objects (services, content) to network addresses. Packets (e.g., content requests) can then be routed to these network addresses. MobilityFirst also supports in-network services such as caching.

only need to match sources of information with consumers of it. A source can be the original producer, an intermediate router or, perhaps, a nearby phone connected via ad hoc WiFi. In such cases, it seems that an ICN would therefore be far more suited to handling such applications. The rest of this article explores the benefits that ICN could bring to these circumstances, as well as the key challenges that must be surmounted before a real deployment could become possible.

ICN for Disaster Situations

The above section has highlighted a number of key characteristics of ICNs and disaster scenarios. This section explores the possible benefits that an ICN could bring to disaster situations.

Information Resilience

The concept of resilience in an ICN is revolutionized when compared to host-centric network (HCN) equivalents, because maintaining resilience in an ICN does not necessarily involve maintaining connectivity between devices (e.g. end hosts, routers). Instead, it involves maintaining connectivity between consumers and their desired information (possibly available in many locations). This is particularly powerful due to the easy migration and replication of information (a virtual entity) when compared to cables and bare metal (physical entities). These principles are also supported by findings from disasters such as the Great East Japan Earthquake and Tsunami, where local equipment was left unharmed but backbone links were severed [3]. In the simplest case, disconnected parts of a network may still be able to reach cached replicas of data, even though the data origin is no longer accessible. Solutions such as MobilityFirst, NetInf, and PURSUIT can exploit this by offering local discovery of content sources, as well as hash-based information verification that does not rely on a public key infrastructure. Clearly, sophisticated strategies could also be devised whereby data is explicitly maintained in locations to allow maximum resilience.

Another interesting point is that ICNs make no distinction between network and storage resources. This offers the new ability for anybody to simply "plug-in" new information using portable storage devices. For example, one could imagine emergency responders carrying backpacks with such equipment, allowing disconnected islands to more easily reach information from the wider world. Such facilities would be nearly impossible in HCNs without sophisticated application intelligence, yet would be seamless in an ICN.

Connectivity Resilience

Resilience in HCNs centers on the maintenance of N^2 connectivity between all device pairs. Thus, traditionally, measuring resilience often involves the removal of links and verifying the existence of this N^2 property. As discussed above, this is no longer a valid measure in an ICN. Despite this, there is obviously always the need to maintain some level of physical connectivity, even in ICNs. Interestingly, however, the novel characteristics of ICNs also have the potential to improve connectivity resilience. Examples of how this is possible include the following.

Multihoming: Multihoming has the potential to improve connectivity resilience by allowing multiple interfaces to be exploited during failures. Although traditionally a challenge, ICNs allow multihoming to easily occur by multiplexing requests (e.g. for small information chunks) over multiple interfaces without needing to bind communications to the interface for longer than the individual request. The load on a failed interface could therefore instantly be shifted to another interface. This is particularly well suited to routing ICN architectures (e.g. NDN), which do not require a resolution procedure.

Connectionless interactions: ICNs propose a receiver-driven request/response interaction paradigm, where there is no need to establish long-term connections (unlike TCP). This means that common TCP issues (e.g. time outs, discarding out-of-order packets) are removed. It also avoids the need to waste needless round trips setting up connections, time that could be better spent utilizing fleeting connectivity. Although UDP has a similar connectionless nature, most application-layer protocols running over UDP will still require connection-like behavior, whereby they interact with a specific host for an extended period of time. Further, many UDP-based protocols (e.g. RTP) require supporting connection-oriented protocols (e.g. RTSP) to operate correctly. An ICN has no such constraints, creating an entirely connectionless environment.

Hastily deployed communications: Often emergency situations result in the use of back-up communications, e.g. ad hoc WiFi. This is difficult to manage in HCNs as it becomes necessary for the back-up communications to offer global reachability (to allow them to reach the hosts they need). In contrast, ICNs can easily create ad hoc islands of connectivity, offering reachability to information of interest without any complex changes to requesting hosts [13].

In essence, these properties of ICN remove various points of weakness that arise in HCNs. They allow more resilient interactions to occur by changing the fundamental manner in which nodes interact with the network.

Resilient Collaboration

A key element of HCNs is their ignorance of the higher-layer activities going on in their infrastructure (e.g. overlay networks). This, on the one hand, is a strength that improves scalability and extensibility. However, on the other hand, it also limits the cross-layer cooperation that can occur between the different stakeholders operating in the ecosystem (e.g. publishers, consumers, network operators).

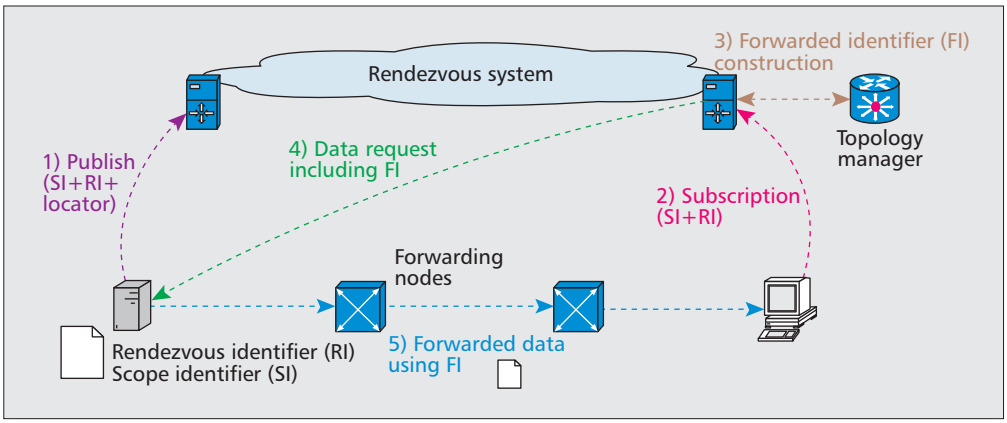

Figure 4. PURSUIT's global deployment uses a Rendezvous Service to forward information subscriptions to appropriate providers, who then publish information to the consumers when available.

Currently, when two hosts interact to exchange data, the network sees this as a packet flow. The only facility offered by the network is ensuring that packets reach their destination (even Quality of Service differentiation is poorly supported). By integrating information-centric knowledge into the network, the network operators can begin to play a far more proactive role in the delivery process. For example, they can re-configure cache placement, as well as pro-actively retain important information within their domains to improve resilience (in principle, all ICNs could support this). These decisions could be informed by collaborative information flowing from other stakeholders, informing each other of the respective importance of different IOs and how they should be treated (e.g. live content should be treated differently than web pages). Importantly, by making these explicit network operations, this could be performed automatically without human negotiation. The importance of this is perhaps best highlighted by the Taiwan Earthquake, which occurred over the Christmas holidays while most staff were on holiday [5].

Superior Disruption Tolerance

The integration of information-centric principles into networks also offers new opportunities in terms of the features provided. Most notably, resilience features such as disruption tolerance could be integrated [15]. Disruption tolerance is usually achieved using a cache-and-forward mechanism by which routers temporarily retain packets until they can be forwarded. This allows periods of disruption (e.g. link failure) to be handled without discarding data. Whereas this is difficult in an HCN due to the network layer's ignorance of packet payloads, this is far easier in an ICN, where all packets are uniquely identified IOs. This is particularly the case for static, immutable content that users could wait to receive (e.g. a webpage). This could be introduced in specific sub-domains or, alternatively, globally deployed by default [11]. MobilityFirst, for example, explicitly supports a cache-and-forward architecture, which can adapt to link failures. Importantly, ICN routers are already equipped with the necessary caches to provide delay-tolerant storage. Whereas caching is typically used for performance reasons (storing the most popular objects), emergency scenarios could trigger alternative strategies that retain the most critical information.

Future Research Directions

The previous section explored some of the key benefits that an ICN could offer resilient communications during emergency situations. Here we explore some of the most exciting future research directions and challenges that remain.

Discovering Information and Sources

A key problem with host-centric networks (HCNs) is the need to discover the locations (e.g. via URLs) of information, typically achieved using a search engine. Before any network interaction can take place, a node must know where it wishes to send a packet to (i.e. the IP address). This complicates network interactions and delays retrievals. In ICNs, the act of discovering these hosts is handled by the network; however, the need to discover the appropriate information identifier still exists. For example, a node desiring a warning announcement needs to know the information identifier of it before a request can be issued. What if the mapping engine (human-readable term → information identifier) became unavailable? In such a case, only users with locally known information identifiers would be able to use the network. This is roughly equivalent to the failure of today's DNS infrastructure: even if underlying connectivity were available, most services would cease to operate if DNS were not operational. Clearly, some mechanism would need to be devised to disseminate up-to-date information identifiers during emergency situations. This could be particularly challenging when needing to disseminate self-certifying identifiers that require trusted sources (a malicious node could provide an identifier for the wrong information). We believe that more decentralized solutions would be appropriate here, whereby search-based functionality could be located locally on clients (using cached search data). For example, it would be easily feasible to locally store a gigabyte search index. Search data could then be automatically updated in local caches to ensure high availability. This would, of course, be accompanied by search infrastructure located in each autonomous system. Critically, this would need to move away from the relatively centralized approaches offered by systems like Google and Bing.

Settling for "Second Best"

All mainstream ICNs currently offer exact matching on information identifiers. Requests for information not referenced in the resolution or routing infrastructure are discarded. However, in many cases, particularly emergencies, users would be perfectly satisfied by alternative IOs. For example, a user in need of an evacuation map might be largely satisfied by a textual evacuation description. While this compromise is usually not necessary, it is far more important in emergency situations, when the "ideal" information might not be available. This places far greater complexity on the routing/resolution infrastructure, requiring the graph-based modelling of information relationships. Despite its complexity, the benefits for an ICN in the setting could be significant. It is therefore

important to develop far more sophisticated resolution schemes capable of handling these mappings (it is unlikely such complexity could be placed in routers). We posit that functionality must therefore be embedded within a supplementary resolution infrastructure that is only invoked when absolutely necessary. Once again, this would need to operate in a relatively decentralized way (unlike cloud-based graph processing systems, e.g. Apache Giraph). Gossip protocols seem particularly well suited to this challenge, where people could share object descriptions and similarities via ad hoc links.

Maintaining Management Functionality

As with any HCN, future ICNs will need a sophisticated management infrastructure to handle such things as authentication, charging, accounting, and identifier control. The key benefit of an ICN is the ability to continue operation in disconnected islands through the replication of IOs. However, this assumes that these islands can continue operation as independent units: the presence of IOs is only one part of the recipe. For example, typical home broadband users will always require access to authentication servers before being allowed access. This problem is likely to be minor when network failures happen remote to the consumers (i.e. at the provider side) but could become significant during failures at the consumer side. If a user cannot perform the necessary management tasks, they will be blocked from the network. Similarly, approaches that rely on public key infrastructures for information verification will suffer from similar issues. Of course, an easy solution would be to disable all such requisites; however, this would likely generate many problems of its own (e.g. security issues). Consequently, we argue that future ICN designs should support varying modes of operation, whereby disaster situations could initiate new access rights and modes of behavior (bypassing the usual management tasks). ICNs offer a particularly powerful abstraction for this, as people could be limited to accessing information deemed highly critical. Further, requisites like digital certificates should always be locally stored, to ensure that hosts can verify important emergency information.

Resilient Routing/Resolution Convergence

One of the biggest challenges facing resilient networks is the reaction of the various embedded routing (BGP, OSPF) and resolution (DNS) algorithms to rapid and significant changes, as caused by many types of disasters, e.g. earthquakes. For example, the failure of a link must be detected and appropriate routing information propagated so that future traffic is routed around the area. Unfortunately, this problem remains in ICNs; in fact, it could even be exacerbated due to the far greater number of IOs when compared to hosts. Re-converging routes (using BGP) often takes in the orders of minutes; as of yet, ICN route convergence times are unknown on a large-scale. A particularly significant problem may arise if it is impossible to aggregate information identifiers (e.g. when using a flat identifier space). In this circumstance, it becomes necessary to perform routing/resolution updates for all objects, rather than the aggregated updates performed using BGP. Such costly operations could be dangerous when dealing with ephemeral connectivity. We therefore argue that in future ICNs, routing protocols should be developed that can prioritize the dissemination of particular information routing/resolution updates. Clearly, this would therefore require classification of such information, allowing only the most important objects to remain indexed in the system during disaster situations. Perhaps more importantly, an accompanying management infrastructure must ensure that this is done correctly due to the security challenges that may arise.

Push-Based Information Dissemination

Primarily, the communications model proposed in ICNs is pull-based. This is attractive for resilient networks because it allows receivers to take control of communications without needing complex control protocols. Instead, a receiver can simply re-request information that has not been received yet. Whereas this is appropriate in many circumstances, it is also inappropriate in others. Many communications during disaster situations are actually push-based (e.g. pushing information to civilians). To address this, polling mechanisms could be used, where receivers periodically request potential publications (using pre-defined or algorithmic identifiers). However, this costly process can quickly become undesirable. Rather, it is more suitable to also offer a push-based mechanism, where publishers can select who receives information. Although some solutions, such as PURSUIT and MobilityFirst, offer push-based dissemination, this style of interaction is yet to receive widespread use in all ICNs. Further, there is not yet any standardized way to describe users/hosts to better facilitate push-based interactions. Instead, users must subscribe to particular streams of objects. This might be difficult in disaster situations, where individuals are unlikely to know what exactly to subscribe to. It would therefore be better if such subscriptions could automatically be managed at the publisher's side, by describing the characteristics of individuals who should receive information (e.g. all civilians over 18 in a given area). We therefore argue that ICNs should start developing more sophisticated descriptions of users, as well as content. These descriptions could then be used to dynamically match providers and consumers in a richer way than simple identifier matching. Although a potentially high-overhead process, this would be extremely beneficial in dynamic emergency scenarios.

Conclusion

This article has explored the possible resilience benefits that ICNs could bring during emergency situations. We have argued that ICNs could offer significant advancement over host-centric networks (HCNs) by detaching the prerequisite of global reachability from the successful operation of the network. As such, these considerations make a strong motivating case for the deployment of ICNs, particularly in countries stricken by frequent disasters. We have already begun work in this area by integrating delay-tolerant support for ICNs, supporting resilient object placement, and performing disaster application development using ICN principles (e.g. early warnings). However, many avenues of future work remain. Of particular interest is the design of more decentralized ICN architectures that can dynamically adapt their routing and resolution infrastructure to reflect connectivity problems. By pushing the communications abstraction up to the information layer, it seems that ICN systems are well placed to introduce these new types of network adaptation without the need to maintain traditional end-to-end connectivity.

Acknowledgments

This work is funded by the EPSRC IU-ATC project grant. We would also like to thank the referees for their helpful feedback.

References

[1] "Natural Disasters Up More Than 400 Percent in Two Decades," http://www.naturalnews.com/023362_natural_disasters_floods.html.

[2] S. Ganeshan and W. Diamond, "Forecasting the Numbers of People Affected Annually by Natural Disasters Up to 2015," Oxfam Report, 2009.

[3] K. Cho *et al.*, "The Japan Earthquake: the Impact on Traffic and Routing Observed by A Local ISP," *Proc. Special Wksp. Internet and Disasters*, 2011.
[4] P. Smith *et al.*, "Network Resilience: A Systematic Approach," *IEEE Commun. Mag.*, vol. 49, no. 72011, pp. 88–97.
[5] Y. Kitamura *et al.*, "Experience with Restoration of Asia Pacific Network Failures from Taiwan Earthquake," *IEICE Trans. Commun.*, vol. 90, no. 11, 2007, pp. 3095–103.
[6] D. Trossen and G. Parisis, "Designing and Realizing an Information-Centric Internet," *IEEE Commun. Mag.*, vol. 50, July 2012, pp. 60–67.
[7] V. Jacobson *et al.*, "Networking Named Content," *Proc. ACM CoNEXT*, 2009.
[8] G. Tyson *et al.*, "Juno: A Middleware Platform for Supporting Delivery-Centric Applications," *ACM Trans. Internet Technology*, vol. 12, no. 2, 2012, pp. 4:1–4:28.
[9] C. Dannewitz *et al.*, "Network of information (NetInf) an Information-Centric Networking Architecture," *Comp.Commun.*, vol. 36, no. 7, 2013, pp. 721–35.
[10] D. Raychaudhuri, K. Nagaraja, and A. Venkataramani, "MobilityFirst: A Robust and Trustworthy Mobility-Centric Architecture for the Future Internet," *ACM SIGMOBILE Mobile Computing and Commun. Rev.*, vol. 16, no. 3, 2012, pp. 2–13.
[11] M. Antikainen, "Reliable Data Delivery for Publish/Subscribe Networks," Technical Report, Aalto University, 2010.
[12] I. Psaras *et al.*, "Name-based Replication Priorities in Disaster Cases," *Proc. IEEE INFOCOM Wksp. Name Oriented Mobility (NOM)*, 2014.
[13] M. G. Soon-Young Oh and Davide Lau, "Content Centric Networking in Tactical and Emergency MANETs," *Proc. IFIP Wireless Days Conf.*, 2010.
[14] J. P. Sterbenz *et al.*, "Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines," *Comp. Networks*, vol. 54, no. 8, 2010, pp. 1245–65.
[15] G. Tyson, J. Bigham, and E. Bodanese, "Towards an Information-Centric Delay-Tolerant Network," *Proc. IEEE INFOCOM Wksp. Emerging Design Choices in Name-Oriented Networking (NOMEN)*, 2013.

Biographies

GARETH TYSON (gareth.tyson@qmul.ac.uk) is a post doctoral researcher at Queen Mary University of London. He has a Ph.D. from Lancaster University and has worked on several national and international projects. He has (co-)authored more than 30 papers and is involved with a number of prominent conferences and journals, as well as operating as a reviewer for both national and international funding bodies. His current interests relate to information-centric networks, multimedia content delivery and network measurements.

ELIANE BODANESE joined Queen Mary University of London in 2003. She is a lecturer at the School of Electronic Engineering and Computer Science. Her research interests include distributed systems middleware, QoS provisioning in cellular and ad hoc networks, intelligent mechanisms for integration of heterogeneous wireless environments, and context-aware sensor networks.

JOHN BIGHAM received his undergraduate degree in mathematics from the University of Edinburgh, his M.Sc. degree in cybernetics from Kings College University of London, and his Ph.D. in artificial intelligence from Queen Mary University of London. His research interests include resource management in wireless networks, the resilience and QoS of middleware systems, and security.

ANDREAS MAUTHE is a Reader in Networked Systems at the School of Computing and Communications, Lancaster University, UK. His research focus is in two areas within the networking and systems domain: Network Management, and Multimedia Systems. He has (co-)authored more than 80 papers published in international journals, conferences and workshops. He is also author of a textbook on Professional Content Management Systems. Andreas worked in industry for more than four years in various management and research positions.