FIA Principal Investigators' Meeting

November 14-15, 2013
UCSD, San Diego, Ca

Meeting summary

This FIA meeting addressed four topics: fault management, routing, Internet of things, and performance. The format of the meeting dedicated a half day to a discussion of how the four FIA projects related to each of these schemes, and what lessons could be learned about the overall problem from looking at the four cases.

Fault management:

This topic addressed management of faults, which includes detection, localization, isolation/avoidance/working around them, and repair/recovery. This is an aspect of security and also management, and was selected because it seems focused enough for discussion. One of the problems with the current Internet is that while certain obvious faults, such as a failed router, can be detected by its neighbors, and potentially corrected by the dynamic routing algorithms if there is adequate redundancy, there is no general way to detect or localize more complex or malicious failures. Certain failures, such as a router that advertises one route to a destination but actually uses another, cannot be detected in the current Internet. Many of the FIA proposals have something to say about fault management, but (as with routing) the points of view may be very different. Schemes that include some sort of explicit user choice over routes must be able to feed back failure information (e.g. location of the fault) to the element making the choices.

The Nebula project, which has a rich and rather complex set of mechanisms to prove that a packet has taken a certain route, can detect these sorts of problems. If an AS is mis-directing a packet, the neighbor ASes can detect and localize this behavior. So the mechanisms of Nebula may provide a much richer foundation for fault isolation in a multi-AS network. However, the signed routing controls of Nebula do not detect other sorts of failures, such as malicious corruption of the packet contents.

The XIA design has several mechanisms that can contribute to better availability and resilience in the face of faults. The multiple *principal types* (e.g., destination, service, content) can exploit different routing schemes, so a failure in one may not disable others. The communicants can switch among them in an attempt to mitigate a fault. The Content principal type (the CID) offers an intrinsic “anycast” delivery model, which may allow data to be retrieved even if one part of the net is misbehaving. The SCION scheme for control of forwarding prevents certain classes of bad behavior, such as announcement of invalid paths, and supports end-point controlled multipath routing. At the time of this meeting, the team had a fault localization protocol under development.

MobilityFirst has the capability of re-routing a packet in transit by looking up the Global UID (a "flat" end-point identifier) in the Global Name Resolution Service to try to find a new network location (NA) for the GUID. The ability for either the sender or an element in the network to associate the GUID with alternate/multiple NAs allows for a natural form of multi-path routing. At the same time, the GNRS is a new service component, which itself must be designed to deal with faults, malicious components and the like. To deal with faults and impairments, especially at the edge of the network (potentially a wireless network supporting mobility) the design allows end-noted to query management data from any network, to support multi-homing and routing path preferences.

The very different base architecture of Named Data Network reshapes the landscape of faults and their detection. Since a router forwarding an interest packet makes a record of that event, any router can detect whether an outgoing path yields a data packet in response to an interest. Paths that do not yield any responses can be flagged as potentially flawed. This local knowledge can drive the local routing/forwarding decisions. Routing protocols thus need not perform highly dynamic failure detection, but can run with a longer time-constant.

These various schemes thus display a range of approaches. Nebula and Scion rely on a very expressive packet header to provide detailed route control. In contrast, NDN relies on a very constrained packet header but per-packet state in the routers. Several of the schemes exploit one or another form of multi-path routing, which will sometimes bypass a fault along one of the paths. However, few of these schemes provide specific feedback to the sender (or to intermediate nodes) that is designed to lead to fault isolation and remediation.

Routing:

The question of interest here is not the complete details of how the schemes are currently doing routing, but what the different designs provide as foundations for routing, and what they require from a routing protocol. Issues of mobility, ICNs, storage in the net, etc. have important implications for what routing must do. The contrast between proposals with a "source route" flavor and those with a more traditional IGP/EGP approach may be revealing, in terms of the required architecture support.

MobilityFirst provides both a "flat" GUID and a network identifier (an NA, similar to an AS number in the current Internet). This richer packet header decouples routing in the core (based on the mesh of NAs) and at the edge (routing on the GUIDs). The design of MF does not define a specific core routing protocol, but the current proposal is some sort of path vector scheme. Packets can include more than one NA associated with the same GUID, which allows for a form of multicast or multi-path routing. However, these more expressive forwarding features do not require major innovation in how NA-level routing is done. Since the GUID do not capture any form of location or clustering, routing within an NA is "flat", which will limit the size of an NA to a size where routing on flat identifiers is practical.

Nebula, as discussed above with respect to fault isolation, uses a very expressive header that provides something similar to an AS-level source route across the Nebula network. The routes are computed in advance, using a control plane that obtains consent from all the forwarding regions—these permissions are encoded in the packet. Within a region, some sort of local routing (perhaps similar to that of MF) is used.

XIA provides a rich packet header that can include a number of identifiers—host, service, content and the like. Conceptually, the XID identifies the destination, the full header with its DAG of identifiers serves as the locator. These different sorts of IDs depend on suitable routing protocols to provide a matching forwarding table, but XIA, like the current Internet, does not dictate how these routing protocols are implemented.

The content names of NDN raise a challenge for the underlying routing protocols to scale to the number of name prefixes that could be expected, which might be orders of magnitude more than the regions of the network. However, a traditional link-state routing protocol can be used to compute routes to these name prefixes. NDN, like most of the FIA systems and the current Internet, does not require a new generation of routing protocol to work, but it is possible that a scheme called Hyperbolic Routing, where content (and each user) is positioned in a hyperbolic space, and interests are forwarded based on the distance to the coordinates of the content. Another mechanism that can reduce the scaling demands of NDN is an encapsulation scheme that allows a name to be prefixed with a name that maps efficiently to a location.

The Internet of Things:

The term Internet of Things (IoT) refers to a future where "everything" is network-enabled in some form. Small devices (sensors and actuators) raise issues of scale, power, security, management and the like. Several of the projects address one or another aspect of the IoT vision.

NDN allows an application to create a flexible name space of devices (as well as users). This flexibility is exploited to reduce the complexity of system configuration and management. Compared to IP, where devices must be issued IP addresses that reflect topology, the name structure of NDN eases system setup. The data plane of NDN also provides a very different way of conceptualizing communication with sensors and actuators. NDN has been augmented with an enhanced interest packet that can carry a small amount of data (IoT control messages), so NDN supports both the use of interest packets with their soft state in routers and direct control using an interest packet.

The Nebula project is focused on the core of the network, and is thus less concerned with how edge-devices are actually connected, and more concerned with the relationship between edge devices and cloud, where data from devices is aggregated and processed. The resilience and control provided by the Nebula architecture should be of direct benefit to the IoT ecosystem, once it is remembered that the IoT ecosystem includes the cloud.

The work in MobilityFirst stresses issues of power and battery life, which might lead to designs that are "transmit only", and issues of naming. In their conception, devices will have a low-level flat name (like a GUID), and this name can be translated into a structured name (a URI) for higher-level processing.

Overall, the landscape of IoT is not yet fully mapped, and there is probably no single architecture for IoT that will prevail. Issues of power will not apply to many powered industrial controllers, ease of configuration will matter to both commercial and consumer devices, but perhaps in different ways—the smart home may not benefit from the same approach as health instrumentation. These talks only scratched the surface of how basic network architecture will shape the ecosystem of IoT. However, there was a strong suggestion that the issues will not be centered on the data plane and packet forwarding, but on issues of configuration, security, power management and the like.

Performance:

Most of the FIA proposals do not distinguish themselves from the current Internet on the basis of simple performance improvements (e.g., efficient use of links). However, the different schemes raise different and interesting questions about aspects of performance.

A key performance issue for XIA is whether the complex header raises computational challenges for the router. The expressive power of the header can allow for potentially beneficial forwarding options (e.g., high performance paths with slower fallback alternatives), but packet processing overhead could be an issue. The XIA team reported an early result (NSDI 2012) in which a Click-based implementation paid a 23% penalty forwarding small packets with an XIA header with two fallback options compared to IP. With respect to overall network performance, an issue of current importance is control over source selection for video cached in a CDN. XIA can provide control over this case both by giving different clients different DAGs for retrieval, and by doing sophisticated route selection for content IDs (CIDs).

For MobilityFirst, a key issue for performance is the GNRS, which must provide translation of GUIDs to NAs rapidly enough to allow packets to be redirected while in flight. Of course, there could be multiple schemes designed with different performance implications, but the project intends to provide one or more proof of concept, showing that reasonable performance is obtainable. One approach is described in DMap: A Shared Hosting Scheme for Dynamic Identifier to Locator Mappings in the Global Internet, Vu, et al, ICDCS 2012. Another is described in Venkataramani, Sharma, Tie, Uppal, Westbrook, Kurose, Raychaudhuri, Design requirements of a global name service for a mobility-centric, trustworthy internetwork, IEEE COMSNETS 2013.

The ChoiceNet project represents a very different view of performance: they do not define a data plane with forwarding mechanisms, but rather a scheme by which performance impairments can be localized to a specific region of the network, so that the user has enough information to hold the provider accountable for the impairments (loss of business or payment penalties, for example).

NDN, because of its intrinsic caching capability, has the potential to improve the delivery of content. This improvement depends on cache replacement strategies, distribution of retrieval requests, and so on, and thus the architecture itself does not specify the potential improvement; this depends on the specifics of the implementation and deployment. But the capabilities of the design suggest new dimensions of performance optimization.

Both for ChoiceNet and for Nebula, performance has a strong component of economics: route selection (and thus performance) is a result of route setup negotiation, which will involve negotiation over price. These designs express most explicitly what is sometimes only implicit—performance is tied of necessity to cost and thus price.