FIT: Future Internet Toolbox*

Thorsten Biermann, Christian Dannewitz, and Holger Karl

University of Paderborn, Research Group Computer Networks,
Pohlweg 47-49, 33098 Paderborn, Germany
{thorsten.biermann,christian.dannewitz,holger.karl}@upb.de

Abstract. Prototyping future Internet technologies is an important but complicated task, mainly caused by incompatibilities to existing systems and high implementation complexity. To reduce these problems, we have developed the *Future Internet Toolbox (FIT)*, consisting of four frameworks that cover data transport, information-centric networking, naming, and name resolution. These frameworks can be used separately or can be combined to a large testbed, covering many aspects at once. Experience has shown that FIT not only simplifies our own prototyping activities but is also useful for other projects due to its generality.

Keywords: Future Internet; network; testbed; framework; prototype; information-centric; data-centric; transport; naming; name resolution.

1 Introduction

Currently, a lot of research is done in the area of future Internet technologies. Hot topics include information-centric networking, data transport techniques and protocols, routing schemes, resource allocation, mobility, and cooperation/ coding techniques. After developing new concepts and protocols in these areas, they have to be evaluated. Often, it is desirable to do this via prototyping and testing under real-world assumptions in real networks. This procedure, however, is very difficult today and is often avoided because (1) the new concepts are incompatible with existing systems or (2) it is too costly to implement a whole prototype from scratch to evaluate a small component of an overall architecture.

To overcome these difficulties, we developed *FIT* – the *Future Internet Toolbox*. FIT simplifies developing future Internet prototypes and testbeds by providing a set of frameworks, covering the following aspects of networking:

- Data transport, including related aspects like routing, resource management, mobility, and cooperation techniques
- Naming and name resolution
- Information-centric networking, including search, publish/subscribe, caching/ storage, and information modeling

* The research leading to these results has been conducted in the 4WARD project and received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 216041.

T. Magedanz et al. (Eds.): TridentCom 2010, LNICST 46, pp. 442–453, 2011.
© Institute for Computer Sciences, Social Informatics and Telecommunications Engineering 2011

These frameworks solve both problems mentioned above as they (1) provide generic testbeds that integrate into today's network/system architectures and (2) provide a lot of ready-to-use building blocks that support and ease development. The downside of such high reusability is that frameworks always introduce a trade-off between reusability and flexibility. Having this in mind, the main goal was to mitigate this trade-off by designing the frameworks as generic as possible to allow implementing as many different concepts and protocols as possible.

The following sections introduce the frameworks that constitute FIT. Sec. 2 covers information-centric networking, Sec. 3 our naming framework. Data transport is addressed in Sec. 4 and Sec. 5 discusses generic name resolution.

An extended version of this paper, containing additional use cases and examples, is available in [1].

2 Information-Centric Networking Framework

This section describes the Information-Centric Network (ICN) framework. We will describe the framework itself and will illustrate how the framework can be used in specific information-centric architectures using the example of NetInf (and Content-Centric Networking (CCN) [2], see technical report [1]).

2.1 Overview

Several information/content/data-centric network architectures have been proposed lately [3,2,4]. To support the prototyping of ICN concepts, a framework providing the following aspects is desirable:

- A generic, adaptable *node structure* implementation
- A flexible, reusable implementation of various *architecture components*
- Support for defining *new services and protocols*, incl. especially communication

We are not aware of a framework that focuses on prototyping ICN architectures (see our technical report [1] for a discussion of related work). Therefore, we developed the ICN framework for rapidly and easily implementing and testing ICN designs as well as specific services and protocols.

The ICN framework solves the apparent conflict between flexibility and reusability by recursively addressing this conflict on four different levels: the *network architecture level* including communication between network nodes, the architecture of each *network node* that is composed of components, the architecture of each *component* containing multiple component services that implement architecture-specific services and protocols, and the design of those *component services*. On each level, we provide reusable structure and building blocks to accelerate the prototyping process while at the same time providing flexibility and extensibility based on a consistent plugin concept.

Fig. 1a gives an example of a node structure for a generic ICN node, consisting of components. Each component can contain a single (e.g., *Naming*, *Information Model* component) or multiple different implementations of its component

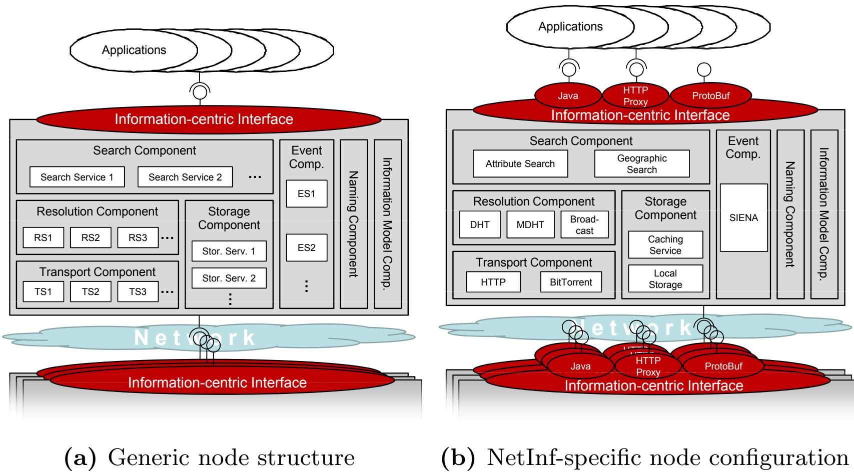

(a) Generic node structure **(b)** NetInf-specific node configuration

Fig. 1. Example testbed node configurations, consisting of adaptable components that each contain several service incarnations like *Resolution Services* (*RS*), *Transport Services* (*TS*), and *Event Services* (*ES*)

functionality (e.g., *Resolution* component). The API of the information-centric node as well as inter-node communication is combined in the same adaptable interface. To provide a broad and flexible mechanism to access the interface, we have implemented a layer of indirection on top of the interface (Fig. 1b) that provides access in different ways (e.g., via an HTTP proxy interface to connect legacy applications) and can easily be extended with new mechanisms.

Based on our NetInf prototype [5] and evaluation of related concepts like CCN [2], Content-Based Networking [6], and DONA [4], we identified the following main components of an ICN architecture: *Search*, *Naming*, *Name Resolution*, *Data Transport*, *Storage*, *Event Service*, and *Information Model*. The ICN framework provides ready-to-use implementations for those components.

Each component can be adapted with architecture-specific services and protocols based on a flexible plugin concept. Services and protocols are encapsulated in *component services*. To enhance flexibility, a component can contain multiple component services while each service fulfills the same interface but may implement and fulfill this service in a different way. A *component controller* is responsible for choosing between component services, and managing the order of execution. For example, the *Resolution Controller* manages several *Resolution Services* that implement specific name resolution mechanisms, e.g., via DHT or via DNS, that can be chosen depending on the type of namespace (flat/ hierarchical) to resolve. Adding a new service to a component is done via the simple plugin concept and reconfiguring the controller to include the new service.

Besides the architecture of a network node, the communication between nodes is the second main aspect of an ICN architecture. The ICN framework offers two distinct components for implementing and testing various communication protocols. First, the *Transport* component is used to implement and test communication

protocols in general. Implementations for legacy protocols are already provided by the ICN framework and new protocols can be implemented using the data transport framework (Sec. 4). Second, as the publish/subscribe paradigm often plays an important role in ICN [6,3], we provide a dedicated *Event* component to simplify the integration of different publish/subscribe mechanisms.

In summary, the ICN framework provides a testbed for building custom nodes and for interconnecting those nodes to implement and test different ICN architectures. It has a flexible and adaptable structure, and offers ready-to-use implementations for the main components of many ICN architectures as well as component services to accelerate the prototyping process.

2.2 Implementation

The ICN framework is consistently based on the interface design pattern for flexibility. It is implemented in Java for platform independence and a flexible choice of devices. The code runs on FreeBSD, Linux, Windows, and Android.

2.3 NetInf Use Case

Fig. 1b shows the configuration of a NetInf node created with the ICN framework. Each component contains several specific service implementations that represent the main services and protocols of the NetInf architecture. For example, the NetInf architecture contains multiple name resolution services in parallel that are each represented as a *Resolution Service* (*RS*) in the *Resolution* component.

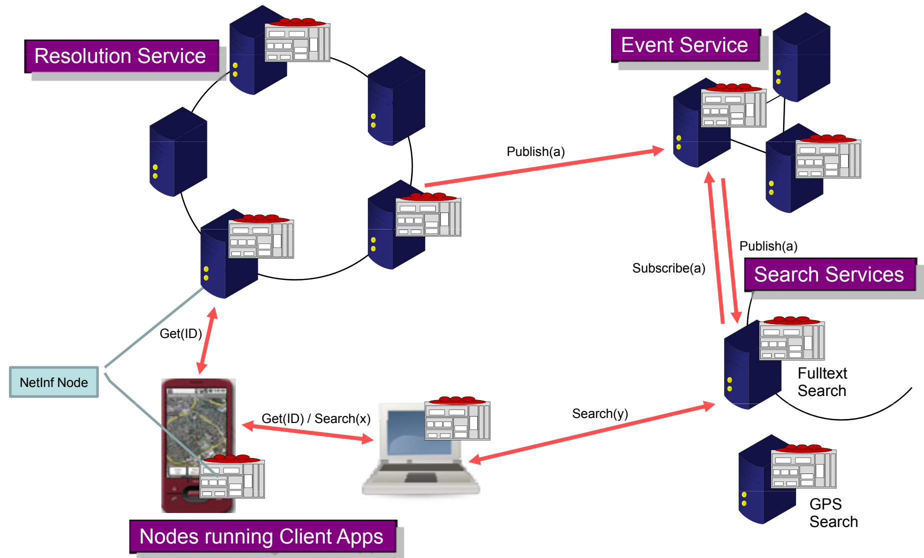

Fig. 2. Network of information-centric nodes

The NetInf prototype also illustrates the generality of the ICN framework by integrating the NetInf architecture with another content-based network architecture, SIENA [6]. SIENA is used to realize the publish/subscribe communication, integrated with the NetInf information model.

Our NetInf prototype extends the information-centric node interface with three different mechanisms to access the node interface (Fig. 1b): Java applications can simply use the provided Java interface. For applications that talk HTTP (e.g., a Web browser plugin), we provide an HTTP proxy interface. For inter-node communication, we provide access to the node interface via Google Protocol Buffers (Protobuf) [7]. Google Protobuf enables the simple and fast definition of custom protocol messages and provides efficient data transfer.

Fig. 2 shows an example of several NetInf nodes that form an ICN. Some are running client applications while others provide services like global name resolution and search services. Communication between those nodes is based on the information-centric node interface (implemented with NetInf-specific primitives like *Get(ID)*, *Publish(a)*) and Google Protobuf.

3 Naming Framework

A naming framework should provide mechanisms that support a wide variety of different naming schemes while at the same time minimizing the implementation overhead. As a result of those considerations, our naming framework [8] is based on a simple, yet flexible and powerful mechanism: names are composed of a sequence of labels, each of which is a 'labelName=labelValue' pair. There are two ways to handle the ordering of labels: *labelNames* are either themselves part of the name, e.g., 'label1=value1 & label2=value2 & label3=value3', thereby explicitly assigning *labelValues* to *labelNames*. Alternatively, *labelNames* are not part of the name, e.g., 'value1 & value2 & value3', which then requires to define the ordering of labels in advance.

Including the *labelNames* in each name allows for flexibility. This makes it possible to define naming schemes with a flexible set of labels as well as a flexible ordering of labels. On the other hand, namespaces with compact names can be achieved by excluding *labelNames* from the names and predefining an explicit label ordering. Via those two mechanisms, our naming framework supports common naming schemes like IP addresses and URIs as well as complex naming schemes like the NetInf naming scheme [3].

Some more complex naming schemes, especially in the area of information-centric networking [3,2], involve features like information-centric security that go far beyond common naming schemes. To support such features, we optionally combine names with a flexible set of metadata that can, e.g., contain security-related data like a hash value of the content. In addition, our framework integrates implementations for security-related algorithms like symmetric and asymmetric encryption, (self-)certification, and public/private-key-based authentication to simplify the implementation of complex naming schemes which include such security features.

4 Data Transport Framework

4.1 Overview

The main observation that triggered our work to develop a data transport framework was the difficulty to introduce new functionality/protocols into today's network stacks. One reason for this is that there is no coherent way to identify entities and to manipulate them as today's network architecture is based on a mainly statically layered stack and functionality is located in end systems.

To design a more flexible, powerful, and reusable data transport architecture, we need an approach to model, design, identify, and use data flows. This approach, however, needs to be generic enough to stretch over a wide range of technology levels and should encompass a wide range of data processing and forwarding functions in end systems as well as in intermediate nodes.

With such a set of requirements, it is impossible to come up with *the, single* solution for a *one-size-fits-all* flow type. Therefore, we decided to choose a design process that combines (1) a uniform appearance and interface for all different flow types and (2) flexibility in supporting a wide range of flow types, in as many different environments as possible. A network architecture that fulfills these requirements is the *Generic Path (GP) architecture* [9]. We use its concepts as basic building blocks for our data transport framework.

We chose an object-oriented approach to design network components while keeping them coherent in their interfaces and basic structures. This allows to incorporate new networking techniques more flexibly than in today's network architectures as networks can be arbitrarily composed of the components. Furthermore, networks can easily be adapted according to any cross-layer information during runtime, thanks to the unified interfaces. Examples for data transport aspects that have been modified/integrated into our framework are routing, mobility [10], cooperation and coding techniques [11], and resource allocation.

A discussion on related work in this area can be found in [1, Sec. 4.5].

4.2 Framework Components

This section introduces the components that constitute the data transport framework: Generic Paths (GPs), Entities, Compartments (CTs), Endpoints (EPs), and Hooks. An example of their interactions is given in Fig. 3. The scenario consists of six Entities (E1-E4, two Cores), four CTs (C1, C2, N1, N2), four EPs (EP1-EP4) forming two GPs, and two Hooks. The GP between EP1 and EP2 in C1 (e.g., IP) is realized by the GP between EP3 and EP4 in C2 (e.g., Ethernet).

Entity. An Entity is the generalization of an application that takes part in any kind of communication. Depending on the implementation, this can be a process, a set of processes, a thread. Communication between Entities is realized via GPs.

Furthermore, an Entity keeps state information that is shared among multiple GPs and runs processes or threads that manage this state. Examples for such state information are routing tables, resolution tables, and access control tables.

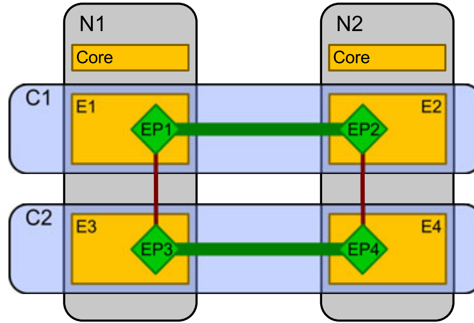

Fig. 3. Overview and interaction of GP architecture components. Entities are drawn as rectangles, CTs as rectangles with rounded corners, EPs as squares, GPs as horizontal lines, and Hooks as vertical lines.

Generic Path. A Generic Path (GP) is an abstraction of data transfer between communicating Entities located in the same or in remote nodes. The actual data transfer, including forwarding and manipulation of data, is executed by EPs.

Endpoint. An Endpoint (EP) keeps the local state information of a specific GP instance, i.e, it is a thread or process executing a data transfer protocol machine and doing any kind of traffic transformation. EPs are created by an Entity and may access shared information of that Entity.

Usually, GPs require other GPs to provide their service. E.g., a TCP/IP GP requires another GP that provides *unreliable* unicast, like an Ethernet GP, to provide a *reliable* unicast service at the end. Therefore, EPs are bound via Hooks to other EPs within the same node. Besides exchanging data, Hooks also hide names from other CTs to permit changing a GP's realization later on.

Compartment. A Compartment (CT) is a set of Entities that fulfill the following requirements:

– Each entity carries a name from a CT-specific name space (e.g., MAC addresses in the Ethernet CT). These names can be "empty" and do not need to be unique. Rules how names are assigned to entities are specific to each CT.
– All entities in a CT *can* communicate, i.e., they support a minimum set of communication primitives/protocols for information exchange. These protocols are implemented as different GP types. Hence, for joining a CT, an Entity must be able to instantiate the EP types required by the CT.
– All entities in a CT *may* communicate, i.e., there are no physical boundaries or control rules that prohibit their communication.

A special CT is the Node CT (N1 and N2 in Fig. 3). It corresponds to a processing system, i.e., typically an operating system that permits communication

between different processes (e.g., by using Unix domain sockets). By means of virtualization, multiple Node CTs can be created on one physical node.

An Entity is typically member of at least two CTs, the “vertical” Node CT and a “horizontal” CT. Furthermore, the Entity has a (possibly empty) set of names from each of the respective CT name spaces.

Note that GPs cannot cross CT boundaries due to the possibly different name spaces, protocols, etc. GPs always reside within a single CT.

Core. The Core is a special Entity. It exists once per Node CT, is only member of the Node CT, and is responsible for node-wide management. The Core’s mainly supports other Entities in cross-CT (i.e., cross-layer) issues, like, name resolution, mobility, managing/controlling Hooks, and service discovery. E.g., in the example in Fig. 3, when setting up the GP between EP1 and EP2, it is the Core that tells E1 that E3 is able to provide the service required to realize its GP.

Note that the Core is also able to reconfigure the realization of existing, i.e., already established, GPs during runtime. This is done transparently to the involved Entities by moving a Hook from one EP to another, possibly in another CT. This feature is required, e.g., to realize mobility or to switch between different transmission modes, like cooperative and direct transmission.

4.3 Testbed Implementation

To be able to use Entity, EP, and Core implementations in various environments, like Linux, Windows, and embedded systems, we used C++ for efficiency and strictly separated the implementation of the logic parts from the environment-specific parts. In detail, this separation means that our testbed abstracts execution environment functions, e.g., by mapping Hooks to available IPC mechanisms. Entities, EPs, and Cores just use these abstractions, which enables to use their content (transport protocols, routing strategies, mobility schemes, data encoding, etc.) in different environments without changing code.

We realized this separation by inheritance, provided by the object-oriented programming paradigm. For this, we implemented all abstractions, like Hook, timeout, and callback handling, in a root class, called `AbstractEntity` (and analogously for EPs and Cores). From this root class, wrapper classes like `PosixEntity` and `OmnetEntity` inherit to map the abstractions to the appropriate execution environment APIs. Thereafter, an `Entity` class is derived from *one* of these wrappers (chosen during compile time). Own, environment-independent Entity classes inherit from this class. Environment-specific entities directly inherit from a wrapper class. Fig. 4 illustrates this.

Currently, we have ready-to-use wrappers for POSIX-compliant systems, like Linux, BSD, and Darwin (Mac OS X) and for OMNeT++ [12], an open-source discrete event simulator. The wrapper for OMNeT++ is especially useful during the development phase and for demonstrating scalability while at the same time using the implementation in real-world scenarios via the POSIX wrapper. Additionally, we are currently working on wrappers for Windows and OpenFlow [13].

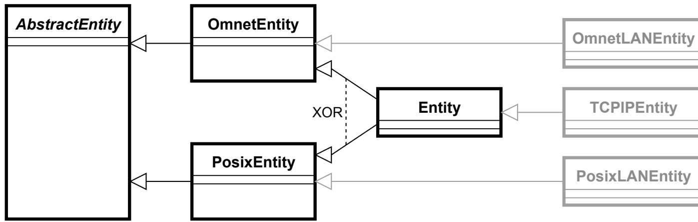

Fig. 4. Inheritance graph for Entity implementations. `Entity` is usually the base class from which user-specific Entities (gray) are derived.

5 Name Resolution Framework

Resolving "names" into "addresses" is used in everyday networking at various layers and abstraction levels. It is realized by a patchwork of individual techniques and concepts. In current networking practice, there is no clear consensus on what a "name" or an "address" really is and how they should be mapped to each other between different layers of a system. Closely linked to this confusion is a lack of a clear concept how the entities inside the individual layers in a system refer to each other and what is necessary to identify the mapping between such two entities; only patchwork solutions for typical combinations of layers (e.g., ARP, DNS) exist. These issues make it difficult to introduce a new protocol or a new layer, as this likely to break existing name resolution schemes.

We propose a flexible yet unified name resolution framework that has two advantages: (1) With a very small set of primitives, a vast range of resolution mechanisms, like ARP or DNS, can be captured. (2) Introducing new layers is much easier. We will discuss information-centric networking as an example.

In the following, we use the framework component definitions introduced in Sec. 4.2 to describe our name resolution framework.

5.1 Resolution Process

When resolving a name, an Entity knows the name of its desired peer Entity and the CT to which itself and this name belongs. The objective of name resolution is to find the following additional information for such a name:

- The name of a CT via which the peer Entity can be reached (e.g., "WLAN").
- An Entity inside this "lower CT", which can handle the communication on behalf of the originator Entity (typically, by means of sharing a Node CT).
- The name of a remote Entity in the lower CT that can pass on data to the actual peer Entity (typically, by means of sharing a Node CT).

The core point in designing a unified name resolution system is to avoid spreading knowledge of how to interpret a name outside of its CT. Neither does the upper CT understand names of the lower CT nor vice versa. Hence, the only thing an Entity can do to resolve a name (in absence of further knowledge) is to

contact all other Entities in its own CT and ask which Entity *has* this name (for optimizations, see [1]) – a WhoHas message is sent *inside* its own CT.

Horizontal CTs usually do not have direct communication means, i.e., spreading a WhoHas message requires assistance of suitable lower CTs (which CTs are suitable depends on the required communication service). A lower CT Entity, however, needs to be told where to send this message; it needs a remote address, which has to be provided by the higher CT Entity that initiates the resolution. Note that this lower CT remote address is, from the perspective of the higher CT Entity, a configuration parameter (opaque string) which it needs to provide but not to understand. Hence, the resolving Entity sends its WhoHas message to all local Entities in all suitable CTs, with the remote address (in the *lower* CT) being looked up, e.g., from a configuration file. This address could be a unicast, broadcast, or even anycast address inside the lower CT.

In the lower CT, the WhoHas message is distributed to its remote address, possibly to many Entities in this CT. The receiving lower-level Entity will receive the WhoHas message and will distribute this message to *all* entities in the original CT. It does *not* have to process names of the original CT.

Inside the original CT, Entities understand names contained in the WhoHas message. Each Entity checks whether it matches the desired name (it does not have to *be* the named Entity, cp. ARP). If no, it can silently discard the message. If yes, it answers with an IHave message, containing (1) the original CT name, (2) the name to be resolved in the original CT, (3) the lower CT name, and (4) the lower CT address over which the name in the original CT can be reached.

5.2 Testbed Implementation

We implemented the name resolution framework directly in conjunction with the data transport framework (Sec. 4.3). In consequence, the transport testbed is able to deal with any naming scheme and any name resolution implementation.

Our name resolution framework implementation spans the Core and Entities. When an Entity wants to resolve a name, it creates a WhoHas message, containing the following information: (1) its own name, (2) the name to be resolved, and (3) the CT name to which these two names belong. Thereafter, this message is passed to the Core, along with a descriptor that specifies which service will be

Table 1. Example `ResolvConf`. SrcCT is the CT from which a name is resolved, DstCT the CT to which the resulting address belongs, ResolverCT the CT in which resolution takes place, and Resolver the name of the resolver Entity.

SrcCT	DstCT	ResolverCT	Resolver
`TCP_IP`	`LAN_A`	`LAN_A`	`00:00:00:00:00:00`
`TCP_IP`	`LAN_B`	`LAN_B`	`00:00:00:00:00:00`
`NetInf`	`NetInfTransport`	`NetInfRes`	`12345`

Table 2. Example `ResolutionTable`. SrcEntity is the Entity that starts the resolution, DstEntity the name to be resolved (both within CT). LowerSrcEntity/LowerDstEntity are the Entity names in LowerCT that were found during resolution.

CT	SrcEntity	DstEntity	LowerCT	LowerSrcEntity	LowerDstEntity
TCP_IP	1.2.3.4	5.6.7.8	LAN_B	45:67:89	89:AB:CD
TCP_IP	1.2.3.4	2.3.4.5	LAN_A	45:67:AB	23:45:67

required for data transfer later on (after a successful name resolution). The Core uses this descriptor to determine for which of the available (horizontal) CTs a resolution is performed. For each of these suitable CTs, the Core checks if an entry exists in the `ResolvConf`. An example is illustrated in Tab. 1; it contains two entries for configuring name resolution within the CTs `TCP_IP` and `NetInf`. Resolving a `TCP_IP` name in the `LAN_B` CT also takes place in the `LAN_B` CT; resolving from `NetInf` to `NetInfTransport` is done in the `NetInfRes` CT.

The Core extends the WhoHas message with the information from the ResolvConf and sends it to all Entities in the own Node CT that are member of the CT defined by ResolverCT. These Entities send the WhoHas message to the resolver Entity (ResolverName) where it is processed and answered.

When the resulting IHave message arrives back at the originating Entity that sent the WhoHas in the CT defined by ResolverCT, it is passed to the Core. The Core adds the contained information to its `ResolutionTable`. An example is shown in Tab. 2. It holds two results; one for 5.6.7.8 in the `TCP_IP` CT that succeeded in the `LAN_B` CT and one resolution that was answered in `LAN_A`.

The Core now informs the Entity that requested the resolution (1.2.3.4 in the example) about the success and returns a reference to the `ResolutionTable` entry. The Entity cannot read the `ResolutionTable` content. It uses the pointer to reference the name resolution result when sending data to the resolved destination name. This strictly keeps names in their own CTs and makes common misuse impossible (e.g., to put "lower addresses" like IP into the payload).

To our knowledge, this is the first approach towards a unified name resolution framework that captures a wide range of existing and future resolution services and spans over all technological levels, i.e., network layers.

6 Conclusion

We presented four frameworks for prototyping future Internet techniques related to data transport, information-centric networking, naming, and name resolution. These frameworks can be used on their own to prototype individual concepts or can be combined to complex testbed implementations, like we did it for the integrated NetInf prototype.

Experience during our work confirmed that the frameworks simplify prototyping a lot. Due to their generality, they will also be useful for other projects

as they reduce redundant implementation of basic testbed functions and provide ready-to-use building blocks to rapidly complement new, small components with an overall network architecture. We will publish our code as open-source project.

Acknowledgments

We gratefully thank the project groups AugNet I and II for valuable discussions and their excellent implementation support. Furthermore, we would like to thank our colleagues in the 4WARD project for their input and fruitful discussions.

References

1. Biermann, T., Dannewitz, C., Karl, H.: FIT: Future Internet Toolbox — extended report. Technical Report TR-RI-10-311, University of Paderborn (February 2010)
2. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M., Briggs, N., Braynard, R.L.: Networking named content. In: Proc. ACM CoNEXT (December 2009)
3. Ahlgren, B., D'Ambrosio, M., Dannewitz, C., Marchisio, M., Marsh, I., Ohlman, B., Pentikousis, K., Rembarz, R., Strandberg, O., Vercellone, V.: Design considerations for a network of information. In: Proc. ReArch 2008 (December 2008)
4. Koponen, T., Chawla, M., Chun, B.G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I.: A data-oriented (and beyond) network architecture. In: Proc. ACM SIGCOMM 2007, pp. 181–192. ACM Press, New York (2007)
5. Dannewitz, C., Biermann, T.: Prototyping a network of information. Prototype demonstration at the IEEE LCN (October 2009)
6. Carzaniga, A., Wolf, A.L.: Forwarding in a content-based network. In: Proc. ACM SIGCOMM 2003, pp. 163–174. ACM Press, New York (2003)
7. Google: Google protocol buffers – Protobuf, Open source project (July 2008)
8. Dannewitz, C., Golic, J., Ohlman, B., Ahlgren, B.: Secure naming for a network of information. In: Proc. 13th IEEE Global Internet Symposium 2010 (March 2010)
9. Biermann, T., et al.: D-5.2.0: Description of Generic Path mechanism, Project deliverable (January 2009)
10. Bertin, P., Aguiar, R.L., Folke, M., Schefczik, P., Zhang, X.: Paths to mobility support in the future Internet. In: Proc. IST Mobile Comm. Summit. (June 2009)
11. Biermann, T., Polgar, Z.A., Karl, H.: Cooperation and coding framework. In: Proc. IEEE Future-Net (June 2009)
12. Varga, A., et al.: OMNeT++ discrete event simulation system Open source project
13. McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., Turner, J.: Openflow: Enabling innovation in campus networks. SIGCOMM Comput. Commun. Rev. 38(2), 69–74 (2008)