Management and Service-aware Networking Architectures (MANA) for Future Internet

Position Paper: System Functions, Capabilities and Requirements

Alex Galis - UCL <a.galis@ee.ucl.ac.uk>, Henrik Abramowicz -ERICSSON <henrik.abramowicz@ericsson.com>, Marcus Brunner - NEC <Brunner@nw.neclab.eu>, Danny Raz -Technion <danny@cs.technion.ac.il>, Prosper Chemouil - Orange France Telecom <prosper.chemouil@orange-ftgroup.com>, Joe Butler – Intel <joe.m.butler@intel.com>, Costas Polychronopoulos – UoA <cpoly@di.uoa.gr>, Stuart Clayman - UCL <s.clayman@ee.ucl.ac.uk>, Hermann de Meer - University of Passau <hermann.demeer@uni-passau.de>, Thierry Coupaye - Orange France Telcom <thierry.coupaye@orange-ftgroup.com>, Aiko Pras - University of Twente <a.pras@utwente.nl>, Krishan Sabnani - Bell Labs Alcatel-Lucent <kks@alcatel-lucent.com>, Philippe Massonet – CETIC <philippe.massonet@cetic.be>, Syed Naqvi – CETIC <syed.naqvi@cetic.be>

Abstract— **Future Internet (FI) research and development threads have recently been gaining momentum all over the world and as such the international race to create a new generation Internet is in full swing: GENI [16], Asia Future Internet [19], Future Internet Forum Korea [18], European Union Future Internet Assembly (FIA) [8]. This is a position paper identifying the research orientation with a time horizon of 10 years, together with the key challenges for the capabilities in the Management and Service-aware Networking Architectures (MANA) part of the Future Internet (FI) allowing for parallel and federated Internet(s).**

Keywords- Position paper, Future Internet, Service-aware Networking, Management

I. INTRODUCTION

The paper is structured as follows: Section 1.A presents the FI overview and context, Section 1.B presents the scope of the MANA, Section II presents the MANA architectural model, Section III presents the research challenges for FI, Section IV presents the conclusions and integration paths towards FI, Section V presents the acknowledgement and contributors to the MANA paper[1].

A. Future Internet – Overview and Context

The current Internet has been founded on a basic architectural premise, that is: a simple network service can be used as a universal means to interconnect both dumb and intelligent end systems. The current Internet is centred on the network layer being capable of dynamically selecting a path from the originating source of a packet to its ultimate destination, with no guarantees of packet delivery or traffic characteristics. The often mentioned end-to-end argument has served to continue the desire for the simplicity in the network. The maintenance of this simplicity has pushed complexity into the endpoints, and has allowed the Internet to reach an impressive scale in terms of inter-connected devices. However, while the scale has not yet reached its limits, the growth of functionality and the growth of network size have both slowed down. It is now a common belief that the current Internet will soon reach both its architectural capability limits and its capacity limits (i.e. in addressing, in reachability, for new demands on QoS, the variation in Service and Application provisioning, etc).

The current Internet capability limit will be stressed further by the expected growth, in the next years, in order of magnitude of more services, the likely increase in the interconnection of smart objects and items (Internet of Things) and its integration with enterprise applications.

Although the current Internet, as a ubiquitous and universal means for communication and computation, has been extraordinarily successful, there are still many unsolved problems and challenges some of which have basic aspects. Many of these aspects could not have been foreseen when the first parts of the Internet were built, but these do need to be addressed now. The very success of the Internet is now creating obstacles to the future innovation of both the networking technology that lies at the Internet's core and the services that use it. In addition, the ossification of the Internet makes the introduction and deployment of new network technologies and services very difficult and very costly.

We are faced with an Internet that is good at delivering packets, but shows a level of inflexibility at the network layer and a lack of built-in facilities to support any non-basic functionality.

The aspects, which we consider to be fundamentally missing, are:

- Mobility of networks, services, and devices.
- Guaranteeing availability of service according to Service Level Agreements (SLAs) and high-level objectives.

[1] Disclaimer: The contents of this document include ideas of many individuals and may not be taken as the ultimate opinions of any those people exclusively, their employers, or the European Commission.

- Facilities to support Quality of Service (QoS) and Service Level Agreements (SLAs).

- Trust Management and Security; Privacy and data-protection mechanisms of distributed data.

- An adequate addressing scheme, where identity and location are not embedded in the same address.

- Inherent network management functionality, specifically self-management functionality.

- Cost considerations – the overhead of management should be kept under control since this is a critical part of life-cycle costs.

- Facilities for the large scale provisioning and deployment of both services and management; support for higher integration between services and networks.

- Facilities for the addition of new functionality, including capability for activating a new service on-demand, network functionality, or protocol (i.e. addressing the ossification bottleneck).

- Support of security, reliability, robustness, mobility, context, service support, orchestration and management for both the communication resources and the services' resources.

- Support of socio-economic aspects including the need for appropriate incentives, diverse business models, legal, regulative and governance issues.

- Energy awareness.

The current trend for networks is that they are becoming service-aware. Service awareness itself has many aspects, including:

- Delivery of content and service logic with consumers' involvement and control - a paradigm shift towards content-and-human centric networking as social, content and service networks.

- Fulfilment of business and other service characteristics such as Quality of Service (QoS) and Service Level Agreements (SLA) - a paradigm shift towards more intelligence within the network.

- Optimisation of the network resources during the service delivery - a paradigm shift towards communication resources as managed shared commodities and utilities.

- Composition and decomposition on demand of control and network domains – a paradigm shift towards cooperative managed networks with increase level of self-manageability.

- Interrelation and unification of the communication, storage, content and computation substrata - a paradigm shift from capacity concerns towards increased and flexible capability with operation control.

Conversely, services themselves are becoming network-aware. Networking-awareness means that the consumer-facing and the resource-facing services are aware of the properties, the requirements, and the state of the network environment, which enable services to self-adapt according the changes in the network context and environment. It also means that services are both executed and managed within network execution environments and that both the services and the network resources can be managed uniformly in an integrated way. Uniform management allows services and networks to harmonize their decisions and actions. The design of both networks and services is moving forward to include higher levels of automation, and autonomicity, which includes self-management.

B. *Scope*

This position paper identifies the research orientation, together with the key challenges for the capabilities and the systems in the Management and Service-aware Networking Architectures (MANA), as a stepping-stone towards the Future Internet.

In order to achieve the objective of having service-aware networks and network-aware services (that is, service and network resources must be aware of the relevant environmental conditions, as well as their own state: which is self-awareness), and to overcome the ossification of the current Internet, this position paper envisages various novel solutions for the FI. We begin with the conjecture that parallel Future Internets would co-exist with the current Internet. The FI(s) must be built as service-aware and as collaborative self-aware federated networks, which provide built-in and orchestrated operation aspects such as: context-awareness, reliability, robustness, mobility, security, efficient service support, and self-management of the communication, storage, content and computation resources and services. Such aspects suggest a transition from a service-agnostic Internet to a new service-aware and self-aware Internet, in which self-awareness is the knowledge of the network environment and operation without external intervention. This knowledge is supporting the communication and computation by means of enhanced in-network and in-service decisions, optimised for common goals.

MANA covers the management, the service-aware networking, plus the service platform technologies and systems, which form the critical infrastructure part of the FI(s). In this paper we also envisage capabilities spanning a range of technologies, including:

- Scalable and robust service-aware networking architectures, including:
  - ➢ Connectivity-to-network, network-to-network services, network service-to-service computing clouds, and other service-oriented infrastructures.
  - ➢ Cross-domain interoperability and deployment.
  - ➢ Optimal orchestration of available resources and systems; Interrelation and unification of the communication, storage, content and computation substrata.
  - ➢ Management systems covering FCAPS functionality, including increased levels of self-awareness and self-management (i.e. all self-* functions).
- Mobile, wireless and high function network core, edges and service nodes.

II. MANA Architectural Model

The current and well-established Internet architecture, is commonly presented as an hourglass shape. This hourglass is depicted on the left in Figure 1, and it shows the data plane functionality of the Internet, but omits the capabilities and the mechanisms needed for the control or management. In the last 40 years, while new network protocols and new service technologies were added straightforwardly into the data plane, it has become continually more problematic to include new control or management capabilities to the architecture. Such changes have created a control plane, which loses the simplicity of the data plane, leading to the "hefty waist" shape shown on the right in Figure 1.

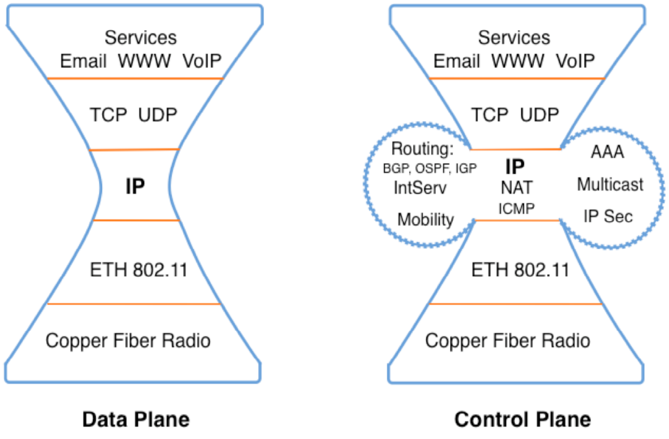

Figure 1 – Current Internet Architecture

The original motivation behind MANA for Future Internets came from the desire to overcome the existing problems and the observation that the monolithic and complex control architectures for the existing Internet could be restructured as a minimal set of components, allowing the services residing in each strata to be accessible through open programmable interfaces — providing the basis for easy prosumer-facing (producers and consumers) and resource-facing service creation, deployability, and manageability.

FI is a system of interoperable and interconnected systems. The following presents an architectural model applicable to FI, which aims to provide a general framework for mapping programming interfaces, operations, and interoperability of services and networks over any given resource technology. It represents a non-layered approach to provide a new control infrastructure, incorporating both management and service enablement functionality while keeping the current IP mainly for communications. The control functions of services and networks are harmonized towards common goals. It covers the inter-related and unified communication, storage, content, and computation sub-strata of FI. The development (refinement and validation) of such FI architectural model is one of the research challenges identified. The MANA architectural model, depicted in Figure 2, identifies the following four types of interfaces:

- α-interfaces: These provide a rich set of APIs to enable highly customized applications and software as a service entities.

- β-interfaces: These provide APIs to orchestrate and govern virtual systems and virtual resources that meet stated business goals having specific service requirements. They are responsible for orchestrating groups of virtual resources in response to changing user needs, business requirements, and environmental conditions.

- γ-interfaces: These mainly provide APIs that deal with virtual system setup and management issues. The APIs consist of methods for manipulating local network/service/storage resources abstracted as objects (i.e. as virtualualized resources) or directly into the real resources (i.e. with no virtualisation). The abstraction isolates upper layers from hardware dependencies or other proprietary interfaces. The γ-interfaces isolate the diversity of setup and management requests from the actual control loop that executes them. They are responsible for determining what portion of a component (i.e. a set of virtual resources) is allocated to a given task. This means that all or part of a virtual resource can be used for each task, providing an optimised partitioning of physical resources according to business needs, priority, and other requirements. Composite virtual services can thus be constructed using all or part of the virtual resources provided by each physical resource.

- δ-interfaces: These APIs provide access to lower level resources. It is a collection of protocols that enable the exchange of state and control information at a very low level between different types of resources and the external agents of the resources. These can aggregate resources into assurable pools of virtual resources. The resource types considered are: transport resources, forwarding resources, computation resources, storage resources, and content resources.

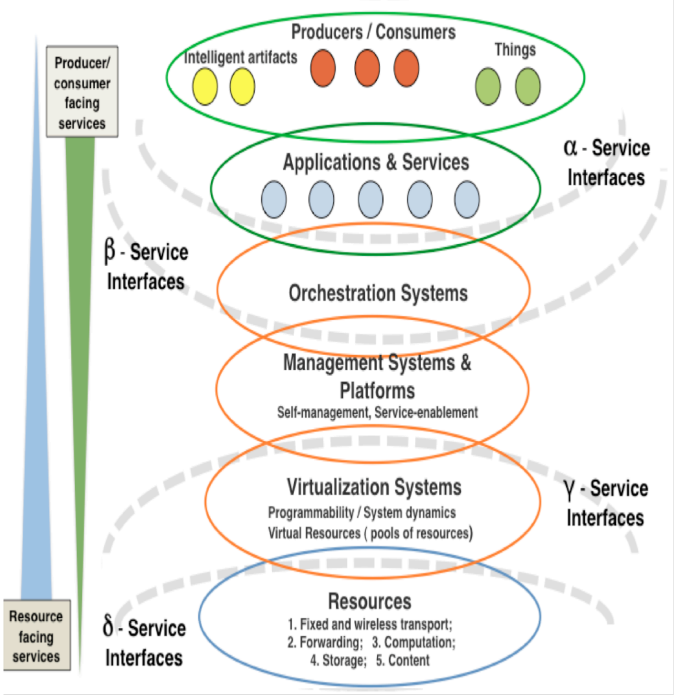

Figure 2 –architectural model

The functionality and capabilities relevant to the levels depicted in the MANA architectural model are described in the following Section III.

III. MANA RESEARCH ORIENTATION: GRAND CHALLENGES, CAPABILITIES AND REQUIREMENTS

This section summarises the important research challenges, requirements, and opportunities for the definition and design of Management and Service-aware Networking Architectures (MANA) for FI(s). The relationships and tradeoffs of such requirements and challenges are also part of the research area. We present the main areas where we feel research into the definitions and the design are required as a starting point. These are presented as 9 sub-sections, namely:

(I) General Capabilities

(II) Infrastructure Capabilities

(III) Control and Elasticity Capabilities

(IV) Accountability Capabilities

(V) Virtualization Capabilities

(VI) Self-management Capabilities

(VII) Service Enablement Capabilities

(VIII) Orchestration Capabilities

(IX) Overall Capabilities

and act as a starting point for the research work to be undertaken.

(I) General Capabilities:

Future Internet is seen as a common infrastructure for interconnecting networks, for interworking services, for interoperating computing machines, and for interflow of information objects would require improvements in its general capabilities and its core system components. It should inherently support a framework of general connectivity, mobility, security, and Quality of Service. As the network evolves to support a multitude of new devices, services, and applications, the general capabilities need to provide robustness and resilience, supporting evolvability, but also provide inherent management to simplify the handling of networks and for users.

The general capabilities encompass those aspects, which are fundamental for a FI to have, and also provide a basis for other capabilities that can be built on top. Regarding the Architectural Model as illustrated in Figure 1, the general capabilities are applicable to all layers. They include:

- Availability of services – anywhere anytime seamless migration all according to SLA and high- level objectives.
- Connectivity anywhere and anytime, meaning the possibility to connect everywhere.
- Manageability anywhere and anytime with an increase level of self-management - as the networked systems become more and more complex this is a necessity as well as an enabler for evolution.
- Mobility anywhere and anytime.
- Adaptability everywhere to changes in context and environment.
- Dependability, resiliency, and survivability to withstand threats and (D)DOS.
- Robustness and stability, including support for mission critical applications.
- Accountability anywhere and anytime - to ensure the possibility of tracking actions performed by a user or a management agent that might impact the networked systems and their performance.
- Evolvability as an inherent feature to ensure the possibility to evolve the networking systems in a smooth way without major disruptions.
- Scalability with respect to features and functions, as well as complexity.
- Trust and security ensuring that users make use of the networks and services in a secure environment.
- Multi-domains to allow for different administrations, technologies, and parallel or federated Internets.
- Support of heterogeneity for possible technology optimisation.
- Openness towards application and services enabling the Internet Openness.
- Energy efficiency of the systems architectures, protocols, and radio spectrum; the use of the networked systems for control of energy consumption.

(II) Infrastructures Capabilities:

Computing, networking, and storage elements represent the components of the MANA infrastructure. The MANA infrastructure consists of evolving and expandable clusters of computing, networking, and storage elements (e.g. deployed both on network systems and nodes, and on users' devices), which support the configuration and deployment of any real or virtual resource of both networks and services.

Proliferating facilities create new opportunities for designing and operating flexible networks, which are able to support new services and applications in a secure and efficient way. The diverse components offer a pool of resources, which could be seamlessly reached and used for various objectives such as resilience and differentiated QoS in a cost-effective way.

However the design and operations of these large-scale infrastructure call for the implementation of new capabilities, which allow them to efficiently share the infrastructure in order to accommodate the various demand of users and customers, while, at the same time, meeting the business expectations of network and service providers. As such, new concepts such as resource virtualization bring new opportunities to resource sharing and isolation for service deployment. In parallel to this, new approaches are also needed regarding network and service management, requiring the introduction of capabilities aimed at simplifying both configuration and control within the infrastructure, with particular focus on self-* capabilities. Regarding the Architectural Model as illustrated in Figure 2, the infrastructure capabilities are mainly concerned with the two lower layers.

Infrastructure includes:

Infrastructure Components:

- Core Nodes, for the provisioning of high-speed, high volume traffic flows for data processing functions (i.e. in the core we are moving from gigabit networks to terabit/s), including flexible control and management capabilities.

- Edge Nodes and Service Nodes, for the programmatic provisioning of the transport, computational, storage, and content resources needed to deploy wide-area services, plus new network functionality, including programmability of the network forwarding functions and flexible control and management capabilities.

- Mobile Nodes and Wireless Nodes, for the programmatic provisioning of the communication, forwarding, and computational resources needed to deploy wide-area services and new network functionality within a wireless or mobile network, including programmability of the network functions and flexible control and management capabilities. Access to wireless infrastructure will also require new, higher capacity radio technologies.

Infrastructure Virtualisation Components:

- Virtual Nodes, as packages of virtual resources, involved in the creation and management of a virtual slice of wired and wireless network, computing, and storage resources in support of a service.

- Programmable networks for the provision and control of networked resources for network clouds.

- Programmable data and service centres for the provisioning of networking computational resources for service clouds.

- Soft nodes with programmability of the control, management, and service logic.

In addition other crucial capabilities that are emerging are:

- Ubiquitous Connectivity, Computation, Storage, and Content infrastructures, together with the architectures, resources, self-management, and controls of such resources, including the assessment of infrastructure adaptations based on context-awareness.

- New globally accessible Infrastructure Services, including Information-centric, General Events, Object Directory and Context-centric networks.

- Ubiquitous connectivity and support infrastructure to autonomic objects (i.e. "Internet of Things"), which are context-aware and capable to generate code, services and human-controlled behaviours, using peer-to-peer communication models.

- Parallel Internet(s) enabling disruptive approaches to be deployed in parallel to the current and future legacy systems.

Apart from the above-mentioned technological issues, there exist some primary business and regulatory issues that are raised by the advent of a FI. From a more business-oriented viewpoint, stakeholders of FI(s)s can be roughly classified as follows: (i) Equipment and IT vendors; (ii) Network operators and service providers; (iii) Web players; (iv) Content Producers (movie companies, broadcasting companies,…); (v) Regulators and research agencies (e.g. IST, NSF, DARPA, Celtic, National agencies,…); and (vi) Academics (Public Research Labs and Universities).

The various stakeholders' interests can be observed through the so-called "net neutrality debate". Beyond the theoretical debate (e.g. in which many academics are involved), this discussion, which varies from country to country, stands as a major regulatory issue and hence may hinder future network deployment because of business uncertainty. Decisions that might come from regulators in coming years may thus have significant impact on the design and management of FI(s). This is why some activity related to business models and socio-economics issues are needed in addition to the deployment of new technical solutions. There is a need to develop a new business paradigm to accommodate the new interfaces and interactions between network/service providers, between service providers and Web players, and finally between providers and the final customers.

(III) Control and Elasticity Capabilities:

The current Internet is an evolution of the basic hour-glass model, in which the core of the network deploys an oblivious forwarding mechanism and most of the intelligence is in the routing protocols and in the end-to-end flow control mechanism. This has been proven to be a winning approach, which provides global connectivity based on a seamless combination of the distributed resources owned by different players with various commercial interests.

However, the FI is expected to move one step forward, from global connectivity towards efficiency. That is, availability and connectivity are now commodities, and the goal now is to provide it in a cost-effective and optimized way. For that, the old hour-glass model and its current derivatives are not sufficient any more and an in-network control mechanism is needed. Such a mechanism should be able to control configurations and manage resources in a way that could allow elastic capabilities, thus, the same resources could be consume by different services at different times (see Figure 2). The involved communication trade-offs need to be identified and managed.

Some of the capabilities needed are described below.

Cognitive Control:

- Uniform, open control frameworks for the FI. These have to be scalable and dynamic, yet be able to serve diverse operational and business requirements. Federation and composition of control frameworks for resources and systems are required.

- Explicit decoupling of the control (i.e. basic routing, content-based routing, source-influenced routing, and value added functions) and transport (i.e. forwarding) planes.

- Mechanisms for flexible data transport, including many relevant transport sub-layers between UDP and TCP; decoupling congestion control from the data transmission. The transport protocol functionality self-adaptation to the service requirements (e.g., level of reliability, QoS etc.).

- Mechanisms for a congestion control sub-layers with generalised fairness based on socio-economic models.

- Mechanisms for publish/subscribe - based inter-networking, aiming for a balance of network incentives and roles between the sender and the receiver. Information based publish / subscribe routing protocols are required.

- Uniform and self-configurable mobility frameworks for FI.

- New naming frameworks, including both channel identity and location, endpoints (source & destination points)-to-location resolution, identity/location splits, and support for addressing and observability of information, context objects and services at all relevant levels and layers as depicted in the MANA architectural model.

Control Operations:

- Systems and mechanisms for orchestration of all distributed control systems (i.e. an orchestration plane).

- An in-network control plane, where the distribution level can be tuned from a fully distributed scheme to a centralized scheme, with an option for intermediate ad-hoc control overlay.

- New tuneable protocols for different layers of the protocol stack in support of cleaner cross-layer interaction and dynamic service composition and collaboration.

- Flexible and cost effective operations of service platforms over core and edge transport networks.

- Mechanisms and interfaces to accommodate the conflicting interests of stakeholders in the FI architecture.

- Multiple and parallel paradigms: Anytime-Anywhere, Anytime-Somewhere, Sometime-Somewhere-When it is optimal (e.g. cheap, etc.), Sometime-Somewhere-As with required qualities (e.g. QoS, security, etc.).

- Interworking with the existing Internet.

(IV) Accountability Capabilities:

Whilst the need for accountability was known in the every early days of the Internet, it was safely omitted from the initial deployment stages. Each player knew the others, and all understood the limitation of the technical platform they were creating, so the impact of this decision was insignificant. Today the network is built from thousands of smaller networks, and they are supporting a much wider range of uses. This has led to tension and a tussle between all the different players. We aim at an “Accountable” Internet, where users are held accountable for any misbehaviour or congestion they cause - hence they are accountable for their impact on others. As such, we need an open delivery infrastructure that can accommodate innovation both at the network and service layer, including the aim to integrate both the technical and socio-economic aspects into a single solution. We need to address:

- Cross layer optimization, resources, network, transport and service layers - to enhance session-less application driven QoS approaches.

- Resource Pooling, for a cost effective way for the Internet to achieve high network utilization and secure future innovation where separate network resources behave like a single large pooled resource.

- Multi Transport Congestion Protocol, this combines multipath routing with congestion control and allows traffic to move away from congested links.

- Enhanced Service Control, enables increased control to the application when applications are best placed to choose the best path for transmission (e.g. low cost path) and manage mobility and multi-homing.

- Enhance Information exposure, where traffic carries info about its resource usage in such a way that the network can monitor the cost (e.g. impact on congestion) of a specific data flow but also the application can select one of the suggested paths from the network protocol to send specific traffic. The monitoring overhead may be traded to monitoring accuracy in case of limited resource availability.

- Lightweight Control Architecture, to avoid locating any mechanisms at network resources themselves for resolving usage conflicts with most of policing and management located at the ‘enforcement point’ – network ingress where customer attaches.

- Separate policy and mechanisms, which need common mechanisms across the infrastructure to control resource usage while the policy can be left under the control of the various stakeholders.

- Development of credible accountability mechanisms for various actors of the FI.

- Mechanisms for handling non-technical aspects of accountability such as legal, governance and ethical issues.

(V) Virtualisation of Resources, Virtual Infrastructures, Specific Network Clouds and Service Clouds Capabilities:

Due to the rise in hardware capabilities, virtualization has been rediscovered as a valuable tool for introducing an abstraction layer between software and the underlying hardware, as it is illustrated in Figure 2. The virtualization layer uses the δ-interfaces to control the physical resources and provides the γ-interfaces to the upper layers to allocate virtual resources to tasks. Virtualized resources are easier to manage and are not bound to specific physical hardware (servers, desktops, storage, or network). It becomes possible to use the same physical device for several virtual resources, to aggregate different physical resources, and to move virtual resources from one physical device to another one. In data centres especially, system virtualization is used popularly today to provide multiple services, which are in parallel and independent from each other on the same hardware, mainly to increase the utilization of resources. However, this concept is also useful in the context of networks. The main elements of a virtualized network infrastructure that is based on system virtualization are multiple virtual networks running in parallel, each consisting of virtual routers and virtual links. Such virtual networks form overlay structures that are not directly related to the underlying physical network. A virtual network has most of the ordinary properties of a physical network, but it also gains additional features inherited from system virtualization. The additional management functions provided by system virtualization allow autonomic network-level schedulers to set up and deploy different virtual networks in order to achieve goals like optimal resource usage, ensuring QoS or security levels, minimizing downtime arising from external influences, or energy-efficient operation of the network. Embracing this new kind of network model also opens up new business perspectives. The role of an ISP can be split up into two new (possibly independent) roles: hardware operators, who provide the physical devices, and service providers, who rent access to physical devices and deploy services on them in order to fulfil customer requirements. This allows two service providers to make use of the same physical network devices, with the virtualization layer providing proper encapsulation of each service, ensuring non-disruptive interoperability between services. Virtualisation capabilities include the Virtualization of Resources, the inclusion of Virtual Infrastructures, Specific Network Clouds, and Service Computing Clouds Capabilities. These capabilities encompass the following:

Virtual Resources:

- Ubiquitous Virtual Resources with integrated self-management of those resources. This allows for the integrated and flexible usage of heterogeneous and assumable virtual resources for wired and wireless networking, for computation, for storage, for content, and for mobility.
- Virtual assurable groups of resources, which do not necessarily correspond to administrative, topological, or geographical domains. This would take into account concerns such as confidentiality, availability, integrity, and safety; they can be used to enable collaborative groups of consumers to exchange information in pursuit of shared interests, services, or business processes.
- Resource allocation to virtual infrastructures or slices of virtual infrastructure.
- Auditability of virtual resource consumption. Virtual /real resource contracts, RLA – resource level agreements, will be constructed.
- Security concerns related to the use of virtual resource and their management.

Virtual Infrastructure, Operation and Systems:

- Dynamic creation and management of virtual infrastructures/slices of virtual infrastructure across diverse resources.
- Dynamic mapping and deployment of a service on a virtual infrastructure/slices of virtual infrastructure.
- Inter-working, inter-operability, and federation of virtualised infrastructures.
- Inter-cloud trading and brokering of virtual resources.
- Self-Management and manageability of Virtual Clouds (Network Clouds, Service Clouds, Virtual Infrastructures).
- Composition / decomposition of Virtual Clouds (Network Clouds, Service Clouds, Virtual Infrastructures).
- Programmability and cross-layers programmability of Virtual Clouds (Network Clouds, Service Clouds, Virtual Infrastructures).
- Secure and on-demand virtual infrastructure provisioning (programmatic access, sustainable federation, automated system management).
- Mechanisms for managing trust between the virtualised infrastructure and the users.
- Virtual resource-facing services enabling flexible usage of the physical resources.
- Increased level of service-aware virtual/real resource control.
- Agility in virtual/real resources; including dynamic re-negotiation of service configuration.
- Real-time service computing clouds and virtual-private service clouds, integrating the necessary storage, networking, and service resources.
- Ubiquitous light-weight virtual channels for integrating an Internet of Things into a service-aware network infrastructure.
- Service Clouds viewing the virtual and real network as a service.
- Service Clouds: application as service in a Cloud, platforms in the Cloud, Infrastructure Clouds, network infrastructure as a service in the Cloud; Federated Clouds with Networks for business applications.
- Increased level of automation and autonomicity in the Service Clouds.
- Overlays for enabling decentralized component interactions and for the provisioning of virtualisation of the infrastructure resources; overlays for creating a topology of nodes for the interactions of different components.

(VI) Self-management Capabilities:

The area of FI is considered as a representative example of a complex adaptive organization, where the involved partners have conflicting goals and tension to maximize their gains.

This evolution renders imperative the need for adaptable, stable, and scalable systems that operate in unpredictable environments, having self-management features and the ability to handle complexity. FI designers are required to conceive new network architectures that are flexible, ubiquitous, and self-manageable. In FI environments, mobility becomes a critical part of the technological landscape, while protocols should operate efficiently both in the wired and the mobile wireless world. Quality of Service and security mechanisms should also be integrated. Furthermore, networks are required to be service-aware through continuous flow observation or application signalling while featuring inherently functional componentisation principles and reconfiguration capabilities. These features will support dynamic optimisation techniques, while services will continue evolving, being more adaptable and aware of user context and preferences. Network elements should support autonomous decision-making mechanisms like, for instance, having the ability to decide in an intelligent way the path followed by the traffic, taking into account the capabilities of the underlying technologies, the type of the information being transferred, as well as user's preferences (e.g., presence awareness). Therefore, a key challenge of the FI is to provide means that will enable cognitive network management through dynamic, ad hoc, and optimized resource allocation and control, fault tolerance and robustness associated with real-time trouble-shooting capabilities.

The ability to have self-management is another important aspect of a FI. The self-management (or autonomic) capabilities are applicable to all levels with the hub at the γ-level of the architectural model (Figure 2): These are the capabilities:

Self-functionality Mechanisms:

- Cross-domain self-management functions, for networks, services, content, together with the design of cooperative systems providing integrated management functionality of system lifecycle, autonomicity, SLA, and QoS.
- Embedded and inherent management functionality in most systems in the FI, such as in-infrastructure management, including in-network management and in-service management.
- Mechanisms for dynamic deployment on-the-fly of new management functionality without running interruption of any systems. The operations required are: Plug-and-Play, Unplug-and-Play, and (re)programmability of the forwarding and control planes.
- Mechanisms for dynamic deployment of measuring and monitoring probes for service and network behaviours, including traffic. Mechanisms for monitoring algorithms and frameworks. SLA-aware sensing and continuous monitoring of systems' adaptations. Adaptive SLA-aware infrastructure. Use of monitoring services in support of the self-management functionality.
- Mechanisms for high performance distributed triggering frameworks and event management (transport, correlation/composition).
- Mechanisms for distribution and use of monitoring probes information; configurable and programmable distributed real-time monitoring of all subsystems.
- Mechanisms for conflict and integrity-issues detection and resolution across multiple self-management functions and policies.
- Mechanisms for optimising tradeoffs between the requirements of multiple systems.
- Mechanisms for intelligent and efficient decision-making where there are multiple participating entities.
- Mechanisms, tools, and methodology for the verification and assurance of different self-capabilities that are guiding systems and their adaptations correctly.
- Mechanisms for allocation and negotiation of different resources. High flexibility in resource control.
- Mechanisms for unified information modelling and storage as a support to context building.
- Mechanisms for support of new/enchanted information modelling of MANA nodes or elements
- Mechanisms for fault diagnosis and possibly self-repair able to cope with incomplete or erroneous management information.
- Mechanisms for self-adaptation of management functions.
- Mechanisms for context-awareness of cross-stratum (communication, storage, content, and computation sub-strata) interaction.
- Mechanisms for socio-economic model based management, which enable control and optimisation of systems life costs.
- Mechanisms for use and development of appropriate ontologies for self-management and orchestration systems.
- Mechanisms for controlling and stabilizing the behaviour of nodes and systems in the context of continuous triggers and changes made autonomously, or in response to inputs (events or programming). Detection and management of normal /abnormal behaviour (i.e. security, intrusion, resources failure and/or malfunction). Explicit relationship between behaviour management, socio-economics and uncertainty.

Self-functionality Infrastructure and Systems:

- Increased level of self-awareness, self-stability, self-configuration, self-organisation, self-optimisation, self-healing, self-protection, self-adaptation, self-contextualisation, self-assessment and self-management capabilities for all FI systems, services, and resources.
- Increased level of self-adaptation and self-composition of resources to achieve effective, autonomic and controllable behaviour.
- Increased level of self-contextualisation and context-awareness for network and service systems and resources.

- Efficient resource management frameworks, including discovery, configuration, deployment, utilization, control and maintenance.

- Automated auditing and traceability of the decisions and changes triggered by the management systems.

- Increased level of cost effectiveness of resources' usage, of system operations and of management operations (monitoring, computations, control, change) and of self-awareness.

- Self-awareness capabilities to support system-level objectives of minimizing system life-cycle costs and energy footprints.

- Self-awareness capabilities for managing operations in time of crisis

- Orchestration as a system of management systems (i.e. bootstrapping, workflow of control, interactions and update of the management systems). Service driven dynamic orchestration. Programmability of the orchestration plane.

- Capabilities for the control relationships between Self-Management and Self-Governance of the FI.

- (Re)establish fundaments of the management of FI by revisiting the science and the mathematics.

- Several degrees of freedom to the design of management functionality for FI (degrees of embedding, degrees of autonomicity, degrees of abstractions, degrees of costs, degrees of manageability; allow clean slate and migration paths). Allow only degrees of freedom that are associated with guaranteed stability.

- Trust in self-management systems.

- Assessment/proof methodologies, mechanisms and technologies of individual self-* capabilities – aposteriori (i.e. benchmarking) and possibly a priori (e.g. by means of simulation or emulation).

(VII) Service Enablement Capabilities:

A service-aware network is an abstract landscape of network services, which can be discovered, negotiated, and contracted with by higher level consuming services at the application level. Offerings, which are exposed as services, are network configuration options, which also map to the requirements of the external services. They need to be discoverable and be able to describe attributes such as capacity, throughput, QoS, latency, protocol support, availability, security, etc., in a consistent format. They need to express cost and availability, scalability, and potentially elasticity and support for usage variations. They need to be supported by a negotiation service, which can implement contracts with consumers. In order to support the SLAs implemented with consuming services, they need to support logging and exception handling. Additionally, autonomic capability within the network needs to be wired to the contracts and policies associated with SLA negotiation such that SLAs in place are enforceable.

For internal efficiency, resource management within the network needs to be aware of SLAs currently in place, both for the negotiation of incoming SLAs, and to ensure that existing contracts are supported.

The level of service awareness will vary, depending on the level. Low level utility resources such as transport and storage in Figure 2, will present minimal descriptive interfaces to services above, will have simplistic or no negotiation, and will be unaware of service concerns at the level above. Moving up the stack however, δ-interfaces will present a somewhat richer level of functional and non-functional property descriptions and programmability. The introduction of management capability at the γ-interface, enables basic service awareness but with very limited or no SLA assurance and enforceability. Only at the α-interface where orchestration-based horizontal scaling, failover, and migration is enabled, can full service awareness be implemented including SLA based service negotiation with enforceability and violation penalties. This approach gives flexibility to the networked environment without risking the stability of the system. Key capabilities include:

- Network services exposed for consumption are virtual, enabling them to be:

- instantiated at run-time over physical resources based on negotiated features (or requirements) such as bandwidth/throughput, security, spatialness, etc.

- managed at runtime with SLA compliance as an objective.

- torn down upon termination of SLA, freeing up physical resources for new use.

- Network service interfaces discoverable by consuming services using standard languages and protocols.

- All relevant service parameters detailed in the service interface.

- A negotiation service, which supports SLA contracting with consuming services.

- Transparent monitoring, logging, and exception handling to track potential SLA violations.

- Network accounting tracks service violation penalties.

- Details of service contracts available to network autonomics so that SLAs can be enforced hierarchically at runtime.

- Run-time network management comprehending details of SLAs in place, when making decisions on infrastructure allocation, as well as negotiating incoming SLAs.

(VIII) Orchestration Capabilities:

The purpose of the Orchestration capabilities is to govern the integrated behaviour and operations of FI system-of-systems and to dynamically adapt and optimize network and service resources in response to changing context and in accordance with applicable business goals and governance policies. It supervises and it integrates all other system behaviour insuring integrity of the FI operations.

These capabilities, which are shown at the β-level or the architectural model (Figure 2), can be thought of as a control framework into which any number of components can be plugged into in order to achieve the required functionality. These components could orchestrate the control algorithms, situated in the control plane of the Internet (i.e., to govern the real-time reaction of the control algorithms), and interwork with other management and service functions (i.e., to provide itself near real-time reaction). Together these distributed systems form a software-driven control infrastructure that will run on top of all current network and service resources.

Some of the capabilities needed are described below:

- Mechanisms for controlling workflow for all systems of all FI system-of-systems, ensuring bootstrapping, initialisation, dynamic reconfiguration, federation, adaptation and contextualisation, optimisation, organisation, and closing down of service components.
- Mechanisms to control co-existence of multiple and parallel FI(s) based on multiple socio-economies matrices and measures.
- Mechanisms for distributed governance.
- Mechanisms to control the sequence and conditions in which one service component invokes other service components in order to realize some useful function.
- Mechanisms for negotiation in order to solve conflicts among FI systems. Negotiation can also occur between different domain systems.
- Mechanisms for allowing conflicting interests (the so called "tussle networking" introduced by D. Clark) such as conflicting policies, traffic patterns, different compensation approaches and different operations.
- Mechanisms for the dissemination of knowledge regarding the Orchestration Plane.
- Mechanisms for FI federation: these control the union/separation of network and service resources having different autonomic management domains. They identify the steps necessary to compose/decompose different federated domains, triggering actions to change the networks and services.
- Mechanisms for controlling the information flow. They define the "What, When and Where" of the information: What information to collect, when to collect, and from whom (where). They supervise the storage of information.
- Mechanisms for cognitive control. They define system data collection, management and decision making, which enable the Internet infrastructure to learn about its own behaviour, to tune its operation, and to enforce its decisions on data manageability.
- Mechanisms for bootstrapping and initialisation systems under supervision.
- Mechanisms for dynamically reconfiguring and adapting of other systems under supervision.
- Mechanisms for dynamically optimising and organising other systems under supervision.
- Mechanisms for dynamically closing down of other systems under supervision.
- Mechanisms for supervision of QoS controllers, triggering an instantaneous modification of the configuration. For example, when following a failure, an instantaneous reconfiguration of the virtual systems is necessary.
- Mechanisms for supervision of resource allocation in several virtual systems. For example, this capability would trigger a change in resource allocations following changes in the context.
- Mechanisms and ontologies that describe the functionalities and enable dynamic discovery, understanding and interaction with the respective offered capabilities.
- Mechanisms to create holistic network view from separate views of the elements in all network level and in all virtualization levels.
- Mechanisms to allow nesting of different control loops with respects to the same objective or the same set of resources.

(IX) Overall Capabilities:

Research in the area of FI should address all the capabilities, their associations and interactions at all the relevant levels and layers as depicted in the MANA architectural model (Figure 2). The eight groups of challenges, capabilities, and requirements of MANA address different parts of the proposed architectures and are encapsulated at a high level in the Figure 3. This figure depicts how each of the eight capabilities relates to each other, with each capability, I to VIII, being labeled.

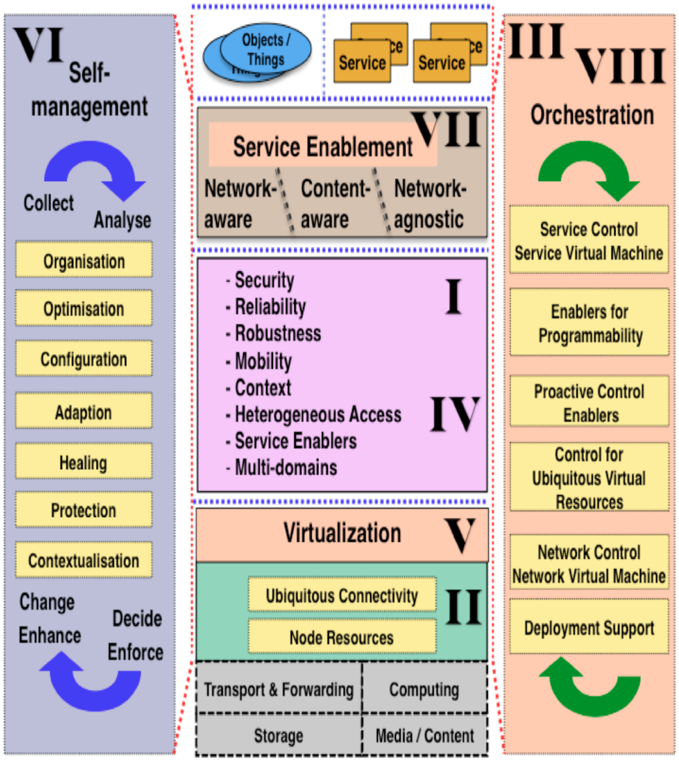

Figure 3 – High-level Future Internet Capabilities

A Future Internet will have to cope with a vast number of highly heterogeneous devices. Computing capabilities are expected to range from supercomputers over commodity workstations, down to highly mobile devices and finally smart dust. While some of these devices can still be considered to be general-purpose computers, many will be highly specialized. Moreover, devices will be restricted not only in their computing capacity, but also in storage space, connectivity, and power. FI architectures have to take this into account and offer communication services targeted to special functions in order to enable optimal operation of these devices. In Future Internet several protocols are expected to exist, especially for the integration of the low-end specialized devices, however these will tap to the global Future Internet via the usage of gateways or service mediators.

On the other hand, the current Internet is already experiencing a huge rise in services offered online. In a FI this is expected to extend even further. However, this multitude of services causes a multitude of different service requirements, like Quality of Service (QoS) or security requirements, which have to be taken into account. Having multiple services run over the same hardware leads to possible service conflicts, which have to be resolved. The current one-size-fits-all TCP/IP architecture is not able to solve these problems. A new Internet architecture has to be found, which is flexible enough to accommodate all of the services and provide interconnectivity without stringent restrictions on services.

Yet a third requirement for a FI architecture is to take into account global social and political factors. The slow and gradual adoption of new architectures also has to be taken into account. We are already experiencing very low adoption rates for the evolutionary approach of IPv6, even though it is many years old. Clean-slate architectures will have an even harder time to succeed and therefore concepts have to be developed, allowing for co-existence and, if possible, co-operation of different network architectures, including both clean-slate and evolutionary, possibly over several decades.

To reach such goals, several tools available today have to be investigated and evaluated. One technology is resource virtualization. By providing abstraction from the underlying physical devices, resource virtualization decouples services and hardware allowing on the one hand to multiplex different services onto one physical device and on the other hand to move a service from one device to another one, once service requirements demand different hardware. Another emerging technology is service-oriented architectures (SOA). Taking the SOA approach to the network level enables elimination of existing redundancies in the current network architecture and instead provides an easy way to combine networking primitives to new and innovative protocols. SOA can be also embedded in low-end devices such as sensors in order to fully integrate their capabilities in progressive Enterprise applications.

The FI needs to be based on an energy-efficient infrastructure and on energy-efficient protocols and services. Raising energy-costs, increasing energy-consumption, and the desire do reduce the world-wide CO2 emissions require ICT for energy-efficiency on the one hand, and energy-efficient ICT on the other hand. Energy awareness at first step and energy optimisation has to be considered from the beginning in FI infrastructures, it is not feasible to adopt it afterwards. The current key technologies of the FI (virtualization of resources and service oriented architectures) are also key-technologies to achieve an energy-efficient Future infrastructure.

Mobility will be a very important part of the Future Internet. Cellular networks with billions of end points will migrate to an Internet core. The main mobility protocol being used today, Mobile IP is a patch on the original IP with a flawed design, which needs to be fixed in order to make mobility an integral part of Future Internet. Consider that Mobile IP traffic sent to a mobile client is first sent to a home agent, which in turn tunnels it to the client in its current location. This design should be corrected to avoid triangulation. Currently, an IP address is used for identifying a mobile terminal as well locating it. The location and terminal identification must be separate. A mobile terminal should be able to get a local IP address at its present location; this information – terminal identical and its location - should be available to any other Internet host.

The facilities for orchestration of trust, security, and privacy for communication and service resources, and mechanisms to protect distributed data are dealt with in several of the FI capabilities, yet they interlock in a coherent overall security capacity. The use of virtualisation in services and networks aims at improving security levels. The self-management capabilities include self-protection mechanisms for virtualised services and networks, as well as intrusion detection mechanisms. Security is also one of the service enablement and orchestration capabilities and orchestration.

V. Conclusions

This paper aims at identifying the research orientation with a time horizon of 10 years, together with the key challenges for the capabilities in the Management and Service-aware Networking Architectures (MANA) part of the Future Internet (FI) allowing for parallel and federated Internet(s). It is also aimed at identifying the research priorities for the future European Union research projects [8].

Further work is envisaged including: a) Analysis of the problems and bottlenecks of the current Internet, leading to a basis for research papers; b) Proposals for the development of the MANA architectural model and systems for evolutionary and clean-slate approaches aligned with visions of other cross-domain topics; c) Proposals for engineering multiple MANA system of systems for parallel FIs, which include layered and non-layers approaches to provide the new control infrastructures; d) Proposals for mapping existing IP overlays, inlays, and underlays into the new control infrastructure; e) Proposals for Integration, Interoperability, Evaluation, Demonstrations, and Testbeds.

Acknowledgment

We acknowledge significant contributions from the following group: Henrik Abramowicz, *Ericsson Research, Sweden*; Nancy Alonistioti, *University of Athens, Greece*; Andreas Aurelius, *Acreo AB, Sweden*; Pablo Arozarena, *Telefonica, Spain*; Alessandro Bassi, *Hitachi Europe SAS, France*; Marcus Brunner, *NEC Research, Germany*; Albert Banchs, *Universidad Carlos III de Madrid, Spain*; Joe Butler, *Intel, Ireland*; Moiso Corrado , *Telecom Italia, Future Center, Italy*; Paolo Bellavista, *Università degli Studi di Bologna, Italy*; Stuart Clayman, *University College London, United Kingdom*; Ranganai Chaparadza, *Faunhofer FOKUS, Germany*; Jeff Chase, *Duke University, USA*; Georg Carle, *Technical University of Munich, Germany*; Piero Castoldi, *Scuola Superiore Sant'Anna, Italy*; Franco Callegati, *Università degli Studi di Bologna, Italy*; Prosper Chemouil, *Orange Labs, France*; Thierry Coupaye, *Orange Labs, France*; Alan Davy, *Waterford Institute of Technology, Ireland*; Steven Davy, *Waterford Institute of Technology, Ireland*; Simon Dobson, *University College Dublin, Ireland*; Henk Eertink, *Telematica Institut, The Netherland*; Olivier Frestor, *INRIA, France*; Alex Galis, *University College London, United Kingdom*; David Hutchison, *Lancaster University, United Kingdom*; Jose Jimenez, *Telefonica, Spain*; Stamatis Karnouskos, *SAP Research, Germany*; Michael Kleis, *Fraunhofer FOKUS, Germany*; Theo Kanter, *Mid-Sweden University, Sweden*; Raimo Kantola, *Helsinki University of Technology, Finland*; Apostolis Kousaridas, *University of Athens, Greece*; Miguel Ponce de Leon, *Waterford Institute of Technology, Ireland*; Antonio Manzalini, *Telecom Italia, Italy*; Lefteris Mamatas, *University College London, U.K.*; Philippe Massonet, *CETIC Belgium*; Panagis Magdalinos, *University of Athens, Greece*; Hermann de Meer, University *of Passau, Germany*; Syed Naqvi, *CETIC, Belgium*; Angel Ferreiro Olivo, *Telefonica, Spain*; Costas Polychronopoulos, *University of Athens, Greece*; Eleni Patouni, *University of Athens, Greece*; George Pavlou, *University College London, U.K.*; Bernhard Plattner, *ETH Zürich, Switzerland*; Aiko Pras, *University of Twente, Holland*; Danny Raz, *Technion University, Israel*; Krishan Sabnani, *Bell Labs Alcatel-Lucent, U.S.A*; Antonio F. Gómez Skarmeta, *Universidad de Murcia, Spain*; Mikhail Smirnov, *Faunhofer FOKUS, Germany*; Rao Sathya, *Telscom, Switzerland*; Patrick Strating, *Telematica Institut, The Netherland*; Andrea Soppera, *British Telecom, United Kingdom*; Makis Stamatelatos, *University of Athens, Greece*; James Sterbenz, *Lancaster University, United Kingdom*; Martin Vigoureux, *;* Mark Wusthoff, *SAP Research, U.K.*; Yang Yang, University College London, U.K.; Jens Zander, *KTH, Sweden*; Tanja Zseby, Faunhofer FOKUS, Germany;

Some of the results of the Dagstuhl seminar on ”Management of the Future Internet”; (27 -30 January 2009; http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=09052), as edited by Alex Galis, have being incorporated in the MANA position paper. In addition some of the results of the Dagstuhl seminar on ”Perspectives Workshop: Architecture and Design of the Future Internet”; as edited by Danny Raz and Alex Galis, (14-17 April 2009; http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=09162), have being also incorporated in the MANA position paper.

Part of the work has been carried out in the framework of the Autonomic Internet (AutoI) project [1], which is partially funded by the Commission of the European Union.

The views expressed in this paper are solely those of the contributors and do not necessarily represent the views of their employers, the EU projects, or the Commission of the European Union.

References

[1] Autonomic Internet (AutoI) FP7 project: http://ist-autoi.eu/autoi/

[2] 4WARD FP7 project: http://www.4ward-project.eu/

[3] TRILOGY FP7 project: http://trilogy-project.org/

[4] PSIRP FP7 project: http://www.psirp.org/

[5] RESERVOIR FP7 project: http://www.reservoir-fp7.eu/

[6] SLA@SOI FP7 project: http://sla-at-soi.eu/

[7] FP7 projects: http://cordis.europa.eu/fp7/home_en.html

[8] Future Internet Assembly (FIA) position papers on (i). Future Content Networks, (ii). Management and Service-aware Networking Architecture, (iii). Trust and Identity, (iv). Future Internet Service Offer, (v). Real World Internet, (vi). Socio-economics, (vii). Future Internet Research and Experimentation - December 2008;

http://www.future-internet.eu/home/future-internet-assembly/madrid-dec-2008.html

[9] MANA Scenarios for Future Internet - http://www.reservoirfp7.eu/twiki/bin/view/FIA_MANA_Members/WebHome

[10] NESSI Strategic Research Agenda – February 2008,

http://www.nessiEurope.com/Nessi/Publications/NESSIDocuments/tabid/590/Default.aspx.

[11] eMobility Strategic Research Agenda – December 2008;

http://www.emobility.eu.org/

[12] NEM Strategic Research Agenda– September 2008;

http://www.nem-initiative.org/

[13] Future Internet, The Cross-ETP Vision Document- January 2008;

http://www.futureinternet.eu/fileadmin/documents/reports/FI_Rep_final__281108_.pdf

[14] Fireworks - Future Internet Research and Experimentation – December 2008;

http://cordis.europa.eu/fp7/ict/fire/

[15] ISTAG Report on Revising Europe's ICT Strategy- February 2009;

ftp://ftp.cordis.europa.eu/pub/ist/docs/istag-revising-europes-ict-strategy-final-version_en.pdf

[16] Global Environment for Network Innovations (GENI) –

http://www.geni.net/

[17] Global Trends 2025: A transformed world, US National Intelligence Council, November 2008;

http://www.dni.gov/nic/PDF_2025/2025_Global_Trends_Final_Report.pdf

[18] Future Internet Forum Korea: http://fif.kr/

[19] Asia Future Internet: http://www.asiafi.net/

[20] EIFEL white paper, December 2006;

http://www.fp7-eiffel.eu/fileadmin/docs/EIFFEL-FINAL.pdf

[21] Dagstuhl seminar on ”Management of the Future Internet”; 27 -30 January 2009;

http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=09052

[22] Dagstuhl seminar on ”Perspectives Workshop: Architecture and Design of the Future Internet” 14-17 April 2009;

http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=09162

[23] “Key Research Challenges in Network management” – A. Pras, J. Schönwälder, M. Burgess, O. Festor, G. M. Pérez, R. Stadler, B. Stiller - IEEE Communications Magazine; October 2007;

http://www.comsoc.org/dl/commag/

[24] “Tussle in cyberspace: defining tomorrow's internet” – D. Clark, J. E. Wrocklawski, K. R. Sollins, R. Braden- ACM SIGCOMM Computer Communication Review, Volume 32, Issue 4 October 2002