

Future Internet Public Private Partnership

*2nd Future Internet PPP
Usage Area Workshop
21 – 22 June, Brussels*

Position Paper



National Technical University of Athens
Distributed, Knowledge and Media Systems Group

*contacts: Prof. Theodora Varvarigou, dora@telecom.ntua.gr
Dr. Angelos Yannopoulos, ang@telecom.ntua.gr*

June 2010

Summary Statement

The basis of this position paper is our belief (founded on our own recent research achievements) that the Future Internet vision should be enhanced with

- greatly increased end-user participation in the customisation of applications and services, and
- greatly expanded contribution of machine-processable content by end-users through a conscious, explicit process where the end-users can express their ideas in a format equally useful for communication to other users and to application software

The enabler for such a breakthrough is the creation of *natural-feeling, linguistic formal interaction mechanisms* for human-computer interaction, both for users to *create content* (and so communicate with other users) and for users to *interact with applications* (when they express ideas to which an application is intended to respond).

Background – State-of-the-Art Research & Development Context

The Future Internet vision encompasses hardware and software infrastructures of extreme power and flexibility. The resulting applications will achieve tremendous impact – but their interactions with end-users are not being presented in current research agendas as sufficiently allowing the users to control these applications.

The prevalent basic philosophy is for “very simple graphical user interfaces”, which essentially means one-size-fits-all interfaces: however powerful the delivered functionality may be, average users get little opportunity to customise it. While simplicity and ease of use are paramount requirements, letting end users lose control of application behaviour is a serious compromise.

In practice today, the danger of in-flexibility of application behaviour is counterbalanced by a variety of techniques that invariably involve automated decision-making of an application in order to assist a user. User profiles provided by users are augmented with histories, annotated according to social graphs, enriched by preference elicitation methods, and so on; the information about each user is exploited for explicitly programmed automated application customisation, for automatically managed content-based data-mining processes, for collaborative-filtering-based or other social recommender-systems, and so forth. As the Future Internet adds capabilities such as integrating sensor networks into services and applications, the capability for software to automatically adjust to user context will continue to increase at a considerable pace.

All these tools – current and future – result in systems that are very impressive at transparently adapting themselves to users’ needs and interests. However, they are not performing well at something much more basic and important: simply *asking* users what they actually want! In order for this to be achieved, the systems must be capable of handling a sophisticated and complicated response (because users’ needs and interests are very sophisticated and complicated), while all the time allowing users to easily express themselves to the systems (because users’ technical capability to manage complicated software systems is on average not at all sophisticated or inclined to handle complicated software systems).

Two possible solutions to this problem are not delivering the required successes:

✎ Natural Language Processing (NLP)

- the users simply say what they want in their own language
- if it worked, this would be the perfect system, since, if the user can conceptualise a desired application behaviour, the full conceptualisation can be exposed to the application to handle
- however, NLP is itself, as a technology, simply not capable of delivering the necessary human-computer communication quality in the general case
- in any case, each individual user’s command of a natural language would be a limiting factor in such a scenario

✎ Human Computer Interaction (HCI) research towards the development of *simple* graphical user interfaces (GUIs) for enabling *advanced* application customisation

- users are becoming more and more accustomed to manipulating computer applications, and progress in HCI is making GUIs simpler and simpler to understand
- average users (without technical knowledge) could one day control applications just as effectively as computing experts do today
- however, the prevalent trend both in popular applications and in research agendas indicates that the improvements being achieved are nowhere near the ideal state of (desired, currently not practically imaginable) NLP-driven capability for users to express to applications exactly what they want
- therefore, application customisation and custom content creation and description capabilities are becoming more and more hidden in practice, rather than more and more exposed

Proposed Innovation & Answers to the five “open questions to be addressed during the Workshop”

Future research

The enabler for the breakthrough that we envision is the creation of *natural-feeling, linguistic formal interaction mechanisms* for human-computer interaction, both for users to *create content* (and so communicate with other users) and for users to *interact with applications* (when they express ideas to which an application is intended to respond).

Software is not capable of processing natural human language, while humans are not comfortable with adjusting the technical parameters of software systems. Therefore, our position is that what is required is: a formal interaction mechanism that is perceived by humans as a linguistic system, while it is processed by software as a formal information representation system.

We call this a “linguistic system” to avoid confusion with “language”, as in “natural language”. A linguistic system is one that possesses structure similar to human languages – like our ongoing background research & development as described in the neighbouring section of this document “Current research of our group, supporting this vision”.

The system is designed and implemented through the cooperation of application domain experts, human computer interaction experts, and software engineers. The concurrent work of these different experts ensures that the system is both perfectly natural feeling to end-users, and fully formally modelled and capable of being processed for all required purposes by a software system.

(1) *What use case and scenario in your area would you consider the most appropriate and representative one for large-scale experimentation with the Future Internet platform to be built starting from 2013?*

Multimedia content creation by non-professional-multimedia-authors is a very important and broad

Current research of our group, supporting this vision

capturing professional content creators’ domain knowledge through formal user-centric symbolic notation languages – first application: film directing

Domain experts’ tacit knowledge of an application is notoriously difficult to elicit and model in a formal system, giving rise to the so-called “knowledge acquisition bottleneck”. In the project ANSWER (ICT-2007.4.2-216489), NTUA has tackled this problem in a completely original way. The chosen application domain is that of professional film directing. In the project, NTUA is developing a system called DirectorNotation, a formal user-centric symbolic notation language, which allows the director to express creative choices for a film much like notes and choreography (music notation and dance notation) have already allowed in their respective domains. Musicians and choreographers have long been able to express their intentions using logical symbolic structures, yet those working in the movie industry have to rely on cartoons and verbal description, and the only record of their artistry is the result itself. The notation system being developed is also being formally modelled, in an iterative process. The

application area. Creating multimedia is an ideal way to communicate, but currently supporting tools are mostly lacking. Users of blogs, social networking sites and even email and mobile messaging applications are still swapping pictures or mostly-poor-quality amateur videos, and writing text that is only loosely connected to the visual media, when they could be combining pictures, video segments, audio, and linguistic content into highly expressive messages. We stress that “atomic” content items, self-contained messages with a specific expressive purpose, can be built by combining different media “segments” that, held together by formal linguistic content in which the content-creator explicitly authors the structure of the message and states its abstract ideas, deliver the intended meaning. The applications that would support such content creation will of course benefit from the advanced customisation capabilities discussed previously in this document, so this use case combines our proposals for content creation and for interaction with software and services in the context of the Future Internet. Finally, the massive-scale creation of multimedia-content that is structured based on a formally modelled language, fully amenable to semantic analysis, would then offer tremendous potential for content-based search, data mining, collaborative authoring with re-purposing of existing media, and a myriad other applications.

(2) What innovative Internet functionality and technologies would you consider important for your suggested use case and scenario (e.g. context awareness, sensor networks, advanced real time processing capabilities handling huge volume of data, ad hoc service composition and mash-up, managed broadband connectivity and services, embedded media support for interfaces easing the interpretation of processed contextual data, etc.)?

The suggested research & development innovation focused on in this document does not directly depend on other innovative Future Internet functionality and technologies, rather it is a new tool in the Future Internet toolbox. However, the use case suggested shows how many synergies exist between the already considered functionality and technologies and our suggestion. Multimedia creation, communication,

chosen workflow ensures that the system is capable of serving the artist’s expressive needs, and also is completely modelled technically, so that input provided by a user in notation can be completely processed by software tools. Example tools are, currently, the automatic generation of animated pre-visualisations, capturing metadata for content-based access in post-production, and, in the future, automatic budget estimation of a projected film production, augmented-reality presentation of documented film directing choices at the set itself and in real-time, in preparation for or even during a shoot, and many more. As a methodological result, we have demonstrated that a deep involvement in the artistic creative process allows the development of an artistic means of expression, which is also controlled to be, technically, a knowledge representation with desirable computational characteristics.

Software tools for the film industry

We are developing specific tools for the film industry. The use of DirectorNotation at the artistic level assists the film director in documenting creative choices. Documentation is an important enabler for self-reflection, for the analysis of intended choices, and for iterative improvement of creative design. The notation

consumption, archiving and re-use is a lifecycle that naturally makes extensive use of all the functionality and technologies mentioned in the question. The capability of the end users to easily and precisely bend all such functionality and technologies to their will is our fundamental contribution.

(3) *Which of the identified functionalities would you expect the Future Internet core technology platform to deliver to support your and other usage area scenarios?*

In our view, it is very difficult to prioritise between the examples of core functionalities identified by typical documents on the Future Internet vision, such as the Final Report on the Usage Areas Workshop for the Future Internet - Public Private Partnership. (Such examples include: data management; service/content handling; rating and billing; identity and access management; rights management; and lifecycle management.) In our view, there will certainly be many important applications and services that will require all of these core functionalities.

Instead of choosing between the identified functionalities, we stress the need for tools that support the seamless integration between applications built atop the Core Platform and all of these core functionalities. Recent advances in Semantic technologies facilitate the provision of such integration capabilities, but do not radically solve the overall challenge of integrating heterogeneous tools and functionalities. The problem is that current semantic technologies are themselves technically challenging to understand and use (let alone to develop or modify). The challenge that we believe should be taken up is to create an easily understandable and usable system of semantic formalisation of software services – technically, this would be a meta-modelling system, built over technical layers using e.g. Description Logics, but exposing semantic models to software developers and even, possibly, to end users, in non-technical, easy to understand terms. The idea of *natural-feeling, linguistic formal interaction mechanisms* proposed in this document tackles more specific problems than this general challenge, but is an example of an innovative way to approach improving the usability of

that the user writes is also useful input for important software tools. The automatic generation of animated pre-visualisations is important because, today, sketched storyboards insufficiently visualise some film aspects such as time-related elements of rhythm and dynamics, while the manual creation of animated storyboards, or “animatics” is very expensive. The application of content-based access in post-production is particularly important, because the ability of the director to access footage produced so far during a shoot is of critical importance to decision-making about the continuation of production activities. We are also exploring the improved film directing workflow made possible with these tools, in the context of computer games production, which more and more employ cinematic presentation techniques.

Knowledge acquisition capabilities

The work of NTUA in this area results in a unique capability to capture expert knowledge, formalise it, create an actionable semantic model of the application domain, and create software tools to support the user. The work paradigm employed brings domain experts and software engineers into very close co-operation. The development

semantic technologies.

(4) What kind of experimentation environment would you consider necessary for broad large scale testing of the platform to be developed in your use area? What would be needed to experiment new services and applications cutting across use areas (services and application mash-up) and building a new services and application ecosystem around the prototype implementations of the platform?

The key to successful experimentation in this area is the provision of an advanced innovative Future Internet Core Platform that has been implemented so that it can be introduced incrementally into use by applications that also exploit massively popular real-world deployment environments, such as applications embedded in social networks. The new applications will be instantly capable of reaching a large audience through their use of the state-of-the-art infrastructure, but will also offer improved functionality over competing applications by making use of the Future Internet Core Platform.

(5) How do you see the potential role of your organisation in the FI-PPP, in the context of Usage areas taking a prominent role in the Initiative, to ensure an appropriate application driven approach?

The role of our group, as presented in this position paper, is focused on the Content usage area. Our vision contributes to a transformation of the way content is created and consumed by end-users, in a carefully planned technical approach that creates a massive technical synergy: content creation and consumption according to our suggested paradigm is achieved in a user-friendly way, but also follows strict semantic formalisation. Thus, software tools will possess complete technical information describing what and how content is produced and consumed. The technical improvements to application capabilities that follow from this innovation are innumerable, since today semantically modelling content is one of the most important challenges (and one that does not improve simply when available resources such as network capacity increase).

of a formal user-centric symbolic notation language progresses with regular meetings of domain experts and software engineers, incremental testing of the notation on real-world scenaria, and incremental formal modeling ensuring that notation usage follows an exact formal semantics. As a result, a system is developed that allows an end-user to express ideas, in a language whose semantics are captured by a complete formal specification. Usefully, this formal specification in fact insulates the remaining software engineering tasks from the application domain: The software tools must be able to process information delivered in the formally specified language. Since software engineers can work with the formal specification of the input they will receive, they do not need to be in contact with domain experts when developing the software tools. Thus, the knowledge acquisition process is made much more effective, with domain experts and software engineers collaborating closely, while the software development work beyond the knowledge acquisition phase is insulated from non-technical concerns.

