

Ubiquitous Computing and Its Influence on MSE

Max Mühlhäuser

Darmstadt University of Technology, Telecooperation Group
max@informatik.tu-darmstadt.de

Abstract

More than a decade after its invention, Mark Weiser's vision of ubiquitous computing finally seems to spark many research activities world wide. This article reviews the ideas behind ubiquitous computing and puts them into the context of the technical possibilities of today and tomorrow. The lack of appropriate software engineering approaches is identified as a major obstacle on the route. Following the discussion of open software engineering issues in general, multimedia related issues are emphasized. The five-step UbiMedia chain is proposed as a reference model for the reconciliation of ubiquitous computing and multimedia.

1. Introduction

The term *Ubiquitous Computing (UC)* was coined at Xerox PARC around 1988 by Mark Weiser [14] and his colleagues. Many believe that the advancement of technology has meanwhile prepared the ground for the UC vision to come true. On one hand, however, this UC vision is so comprehensive and far-reaching that it must be considered a collective effort and determination to move forward rather than a precise goal to attain (maybe comparable to the late fifth-generation program in Japan). On the other hand, the comprehensive, holistic nature of this vision makes it an issue of large-scale cooperative distributed software (and thus, of software engineering) in addition to an issue of base technology. This issue (or rather, large number of related issues) of software engineering will most likely take a long time to be even partly resolved. The present paper looks at multimedia-related issues of software engineering for UC in particular.

In the remainder, we will first try to make the fuzzy notion of UC as clear as possible, relating it to synonymous and dependent terms. As a contribution to this clarification of terms and concepts, we will introduce a three-layer model of UC. The fourth chapter will distill multimedia related concerns and point out open issues in multimedia software engineering as well as some early approaches to these. This chapter will be centered around a 'media chain' of UC multimedia.

2. The vision of ubiquitous computing (UC)

In order to get an idea about what different people mean by the term UC, we will list a number of characteristics of the envisioned UC world, and thereby hit synonymous or accompanying terms used in the literature. Fig. 1 arranges these terms in three layers to be explained at the end of this chapter.

| | |
|------------------|---|
| L.3: UC World | pervasive computing things that think post-PC era calm technology the invisible computer embodied virtuality anytime-anywhere /nomadic computing |
| L.2: Integration | UC engineering (service discovery, awareness management, AAAS, ...) scalable computing (MOM, Web caching, disconnected-op., components, AOP, ...) large-scale computing (organic computing cellular computing, active networks distributed AI, agents, ...) |
| L.1: Gadgets | smart devices (labels / badges /...) smart goods & spaces smart cars, web-on-wheels, ... smart paper / cloths / dust /... location / time / temp./... awareness ambient walls / offices / homes /... augmented reality wearable computers pico / ad-hoc networks Internet appliances |

Fig. 1: Three layers of UC-related terms

- * The mainframe era of computing made many people share one computer (N:1 relation), the PC era brought the 'personal' computer for each user (1:1). UC is the *post-PC era* [12] where billions of people will use trillions of devices (1:N).
- * These trillions of devices will mostly be dedicated and embedded in every-day objects (things that think [3]).

- * The dream of intelligent computers shrinks to the modest notion of smart devices: they 'know' about part of their physical context (e.g., location, temperature); this awareness enables adaptation.
- * Of course, embedded computers have been around for a long time. Communication-enabled, they become Internet appliances, boosting the exponential growth of the Internet towards those trillions of nodes. In fact, the simplest Internet node has nothing but an identity and communication capabilities (smart label for goods, simplest form of smart badge for persons); memory and processing capability are optional.
- * Extrapolating the present 'non ease of use' of PCs to trillions of devices would yield a hostile world. Therefore, the UC vision emphasizes advancements in interaction technology. Dedicated smart devices are envisioned to be interwoven with the environment, so easy to use that they are not recognized as computers (thus the notion of the invisible computer [10]). The term pervasive computing also alludes to penetration (of everyday domains) plus interweaving. At the same time, the 'awareness' of smart devices should be exploited to make the environment as such adapt to the user, even without his or her explicit order (calm technology), yielding ambient environments such as ambient walls, houses, etc. [2].
- * Mark Weiser tried to position UC against the then-prominent virtual reality (VR) efforts. VR applications put the world into a computer (in fact, a 3D model – and more – of an excerpt of the world), they put the computer in the center and the human at the periphery (the human even becomes a computer peripheral, attached to it via data gloves, helmet etc.). Weiser made the quest for embodied virtuality as the contrary: he proposed to put the computer into the world (embedded), the human in the center (adapt technology to human needs, not vice versa) and the computer at the periphery (making it almost invisible, see above). While Weiser deliberately stirred up a controversy, it became soon clear that augmented reality is a kind of reconciliation between VR and UC.
- * The UC vision is by no means that of the ultimate wired Internet, but that of the ultimate wireless and mobile Internet [6]. Before Internet appliances will boost us to trillions of nodes, it is the Internet-enabled cell phone (WAP, DoCoMo i-node, etc.) that is about to lead to billions of nodes. The next step is ready to come, although future will have to tell in which form: it may make us carry, e.g., smart paper with digital ink, augmented-reality eye glasses, cloths with smart labels, smart shoes, or a form of wearable computing [13] for which either humans adapt to technology (cyborg look considered cool?) or technology becomes humane. One generation further, our cloths may be smart and communication-enabled devices

may be implanted in our bodies. As mentioned earlier: whatever the advancement of base technology, it is an issue of large-scale software integration whether nomadic computing comes true i.e. whether we can access and use UC technology on the move, making ourselves at home (as to how we use computing-related technology) anytime-anywhere. While mobile Internet protocols still depend a lot on the notion of a 'home base' for mobile devices and users, true nomadic computing is based on ad-hoc networks instead of mobile peripheral nodes attached to a fixed core. This is also relevant as devices get more decoupled from humans in the context of logistics (smart goods, smart spaces) and traffic (smart cars, web-on-wheels etc.). Briefly spoken, UC is intertwined with mobility, UC is the ultimate mobile computing.

- * The intertwining with mobile computing also calls for advancements in mobile communications, some of which are advancements of current wireless technologies (UMTS, new networks in the sky, WLAN advancements). Since most of the '1 person : N computers' communication will happen on a very short range, pico-networks have gained increased interest with and Irda-successors, Bluetooth™, etc. Note that ease-of-use in this context means auto-configuration i.e. the use of ad-hoc networking approaches.

3. UC and software engineering

Up to now, we have met some 20 terms synonymous or related to UC, all lively discussed in press and public media. According to fig. 1, all these terms have either described (characteristics of) the new gadgets which UC promises to proliferate – cf. layer 1 –, or (characteristics of) the new UC 'world' which is envisioned – cf. layer 3.

It is not by chance that most terms in the second layer of fig. 1 (the 'integration' layer which is basically a software engineering layer, dominated by the issue of scalability, see below) are much less popular. This layer advances slowly, leaves many question open yet, and (in cases of success) yields results which are difficult to convey to the public. Nevertheless, UC will not come true before considerable advancements are made in this layer. Since this layer is meant to make the envisioned trillions of 'gadgets' of layer 1 cooperate and work towards common goals, yet separate concerns (mainly, of different users and their organizations), it is a layer concerned with scalability. Thereby, we understand scalability as a broad term which includes the need to have objects (soft- or hardware) work in different, dynamically changing, and heterogeneous environments, as part of a system which, at the highest level, has the scope of (at least) the globe. In the remainder, we will reflect this emphasis on scalability by structuring the chapter into three parts:

- * Software engineering approaches to UC which attack selected problems arising from the UC vision described and which are not primarily scalability-related. Here, we will concentrate on three important examples
- * SE approaches to scalable computing which are rather in line with traditional SE.
- * More radical new approaches which try to overcome the inherent scalability constraints of classical SE; we will use the term 'large-scale computing' here.

3.1 Selected UC-related problems

Service Discovery: Ad hoc networking requires nodes to get acquainted and to cooperate without a-priori knowledge about one another. As to mutual 'introduction', service discovery (the only well-established notion from layer 2 of fig. 1) is the term denoting the present state of the art [11]: it describes approaches like Jini, uPnP, Salutation, or Bluetooth-SDP, all destined to allow devices to announce their presence in an ad-hoc network and to learn about other devices. Usually, a kind of directory is maintained by (at least) one of the devices present (called, e.g., lookup service), and mobile devices register as either clients or servers. A critical analysis shows that service discovery approaches are just a natural augmentation of 'traders' known from conventional middleware (which are, in turn, augmentations of name servers). In fact, mutual introduction without any a-priori knowledge is utopia, even more so than among humans: at least a common language, better a common ontology must be predefined; to this end, wide-spread ontologies of smart devices may help in the future; in the worst case, competing 'standards' (such as the above cited ones) will lead to the lack of a common language and will render truly nomadic computing impossible.

Awareness Management: as discussed, smart devices tend to be fairly dedicated and will thus usually realize awareness for one or few out of many possible contexts. Since many applications and devices may want to access the same context (i.e. exhibit the same awareness), since several smart device may be needed to generate a more complex kind of awareness (such as group awareness), and since many kinds of context exist, context or awareness management in a large open system is a software engineering issue to be addressed. To this end, little is found in the literature about generic approaches – there is still a lot of research dedicated to understanding, making available, and exploiting specific kinds of context. In the future, software engineering may contribute substantially to making awareness a commodity.

AAAS: Privacy and security issues are recognized as primary concerns as the Internet goes business. The UC vision, however, leads to a drastically aggravated degree of concern here: i) For nomadic users, it is the standard

case to be working in a foreign, non-trusted network. Thereby, both the hosting environment and the nomadic guest have similar concerns: both want to be protected against one another without restricting the degree of cooperation, both want dependable accounting of the resource usage (note that in a UC world, it can no longer be a 'sign of friendship' to have a foreigner attach to one's network). Concerns differ in many details, of course (e.g., guests will want to make sure to be accounted for nothing more than the true 'successful' use of the host's resources, and to do so at possibly minimal cost; hosts will want to assure comprehensive accounting; etc.). ii) The UC vision aggravates this problem as it drastically raises some privacy concerns such as blending the access and mobility patterns (and thus, information which is dreadful in need, e.g., for location awareness in applications).

The term AAAS as used in this context stands for accounting, authorization, authentication, and secrecy, and shall denote, beyond these terms, the comprehensive accounting and privacy/security concerns in a UC world. The rapid evolution, cost (in the broadest sense), and diverging application needs in this respect call for a software engineering solution to application level AAAS.

As we move on to scalability as the most important software engineering issue, one should keep in mind that the three areas discussed above are exemplary rather than comprehensive.

3.2 Scalable computing

Of course, there have always been researchers in the past who put a focus on scalability and remained within the boundaries of conventional software engineering. For instance, the successful growth of the Internet from 4 to some 400 million nodes is based on the rigorous strive of protocol designers for scalability. Nevertheless, the IP routing protocols had to be improved a lot since the early times – mostly because of lacks of scalability. Today, the (again, not well established) term scalable computing – with its focus on scalable middleware - may be used to describe work destined to augment traditional approaches towards scalability. For instance,

- * the exponential growth of the Web made researchers emphasize on Web caching;
- * Scalability issues were the main reason for which middleware has evolved from RPC (or remote method invocation) and client-server computing to 'scalable middleware', described with terms such as MOM (message oriented middleware), event (or 'push') based mechanisms, and the publish/subscribe notion.
- * Scalability as an issue of heterogeneity is currently most visible as the Internet goes wireless: disconnected operation denotes approaches to support devices with unreliable and changing connectivity; transcoding [4,7] is a common term used for an (in-

sufficient) attempt to automatically reformat Web pages for different terminal devices (mapping HTML onto markup languages such as WAP for 'cell phone size' displays, for voice-based I/O, etc.).

- * Component-based programming has gained a lot of attention with technologies like JavaBeans or variants and descendents of OLE/COM, trying to improve reusability and high-level composability, and to push the stage of dynamic composition as late as possible in the life cycle. Terms like aspect-oriented programming (AOP), on the other hand, denote efforts to separate concerns which are usually interwoven in traditional programs, such as core functionality and operational issues (security, reliability, etc.).

3.3. Large scale computing

Looking back at the scope of service discovery, mutual introduction of ad-hoc networking devices is one thing, cooperation is a different story. Years of research about cooperating mobile agents and distributed AI [15] have shown that exchange formats like KIF are only a small step towards making agents pursue a common goal.

Smart devices concentrate on a dedicated task and are thus rather dumb compared to the 'intelligence' expected to live in the overall network. Yet, where could this 'intelligence' reside, actually? Protagonists of distributed AI may argue that it may (more or less) evolve from the mere coupling of smart devices, just like cells self-organize to form a complex intelligent human – counter arguments being that i) (at least) the DNA contains a lot of information about the entire 'system' in organic life, ii) evolution plays a major role in organic life and has characteristics considered inappropriate for computing (time factor, co-existence of extremely inferior and superior species, etc.).

In search of new approaches to large-scale computing, the ever increasing importance of decentralized approaches is not seriously questioned. The above counter-arguments show, however, that necessary research must concentrate on questions of aggregation, separation, and hierarchies, on how local and global 'knowledge' (such as local / global goals) should move bottom-up and top-down in a hierarchy, on what or who is 'the top', and on the extent to which human-set goals (see counter-argument ii) should dominate the evolution of the distributed system.

In this context, authors like Kevin Kelly [5] support the view that control structures and software engineering approaches known from distributed (also: concurrent, sequential) computing do not scale to the UC world described. They believe in the need to 'learn' from organizational, structural, and 'coopetitive' (cooperation blended with competition) approaches as studied in biology, economics, and sociology. If we interpret the (not well established) term *organic computing* in this way, than this is a field of primary importance for UC – but a field that will

revolutionize software engineering! Organic computing may be described as an attempt to reach from 'the humanities' towards computing (just like genetic and neural computing do) in search of large-scale software engineering approaches. There are also attempts to break the limits of the fields which dominated large-scale software in the past, such as concurrent programming (which is put in perspective with cellular programming) and the 'network world' of communication protocols (put in perspective with active networks).

4. Relation to MSE

Up until now, we have not discussed anything about multimedia. This has two reasons: i) this article is destined for readers with a strong MSE background and (at best) marginal UC background; ii) it was necessary to go step by step from explaining the UC vision to deriving general software engineering concerns to deriving MSE concerns. Again, we will have to concentrate on few essential points for brevity.

4.1. The ubiquitous media (UbiMedia) vision

This section is centered around the five-step approach to Ubiquitous Media as depicted in fig. 2.

With the predicted ubiquitous availability of computing resources, either interwoven with our environment or carried by mobile goods and humans, it is natural to expect ubiquity of media, too, for reasons listed below.

Five trends towards ubiquitous media:

1. Ubiquitous computers must be much more easy and natural to deal with: mouse and keyboard must more and more be replaced by voice, gesture, and other modalities.
2. Ubiquitous computers must generally require much less interaction; based on awareness and cooperation, they must act much more autonomously. Thereby, awareness must be based on 'classical sensory' (temperature, acceleration, ...), on 'UC sensory' (cooperative location computation, biometrics, ...), and on media based 'computer perception' (vision, handwriting recognition etc.).
3. The Internet is about to become the fourth mass medium (in addition to press, radio, and TV) [8], making personalized multimedia delivery a ubiquitous application.
4. Consumer electronics for digital video and digital audio advance rapidly, bringing ubiquitous 'commodity media' into reach: the average UC user will use computer-based personal multiple media just as naturally as s/he uses email and personal notes today.

- Ubiquitous availability of media capturing devices will be used for background (ambient) recording of the 'lifestream': intuitively speaking, such recording memorizes both what the UC user sees/hears on the move and what s/he could have seen/heard if s/he had been at other places (the recording could of course be augmented beyond human sense, cf. infrared / 3D vision etc.).

A quest for substantial changes: The fifth item cited makes most obvious what the entire list above suggests: substantially new approaches in most major components of networking and computing are required in order to cope with the enormous amount and diversity of information. A sketch of where these changes could lead is given in the remainder, called *UbiMedia* vision (cf. fig. 2). This vision is based on three assumptions:

- The above list of five trends will at least gradually and partially come true: *we have to face UbiMedia*
- Any claim that there will be enough storage and communication capacity to permit coping with UbiMedia the way we cope with digital media today is a myth: *the UbiMedia famine will outpace the advancements in storage and communication capacity.*
- Not only do we have to change the way in which we represent and handle the digital media themselves; we also have to change the rules and organization of their storage and communication throughout the system, and of their organization and representation on the user side: *we need substantial changes both with respect to the media themselves and with respect to computer and network organization.*

4.2. Zero-Bandwidth representations (ZBR)

At the dawning of multimedia in computer science, there was nothing but the digitized (sampled, quantized, and coded) representation of analog signals. Meanwhile, there have been enormous efforts towards making digital media 'understood' by the computers. Image recognition, video scene analysis, and other automatic indexing techniques represent one of the most active and wide-spread research fields in computer science. Ontologies and standards for the description of media contents, in particular the MPEG-7 standardization work [9], and work on dynamic geometric models for humans (cf. model based videoconferencing) represent important steps towards semantic description of media.

In this context, we can understand the origin of the term 'zero-bandwidth': once the basic model for a scene is known to the receiving end, only operations on this model have to be transmitted in order to provide enough information for the replay of time-based media on the receiving side. This is considered the ultimate degree of compression, hence the joking term 'zero-bandwidth' (which is joking since the necessary bandwidth for

joking since the necessary bandwidth for transmission is of course still larger than zero). Note that the term 'model for a scene' above is to be understood in the widest sense i.e. not restricted to video; 3D geometric model data (e.g., of a videophone participant) are just one example, audio-related and alternative video-related models may be imagined.

In the remainder, we will use the term 'zero-bandwidth representation' (ZBR) in the most general sense as a purely semantic description of (in particular, time-continuous) media. We will discuss a *UbiMedia* world in which all media captured are associated with a (possibly multi-resolution) ZBR.

Note that the ZBR approach has some negative and some positive connotations:

- + ZBR have by definition the smallest size of all representations, minimizing the required storage and transmission capacities.
- + ZBR are fully machine readable and thus offer the maximum support for software-based manipulation (browsing, query, analysis, etc.).
- + ZBR are the ideal basis for translation into alternative modalities (voice2text, video2animation, etc.)
- ZBR are difficult to achieve (cf. the myriad of research programs for all aspects of media recognition). If it were easy, it would be wide spread already.
- ZBR are abstractions i.e. considerably reduce fidelity; for many purposes, they are not (and will probably never be) sufficient (cf. cinema).
- ZBR impose very high Quality-of-Service requirements, especially for real-time transmission (no error tolerance, very low jitter despite considerably varying data rate); this is the 'price' for the low overall data rate, due to the elimination of redundancy.

4.3. UbiMedia as a 5-step media chain

Based on the core concept of ZBR, the *UbiMedia* vision can be explained by looking at the five steps from media creation to media use as depicted in fig. 2.

STEP 1: media creation. This end is the origin of *UbiMedia*. We have to distinguish computer-generated media (probably subject to comparatively easy ZBR creation), legacy media (cf. the gigantic video archives of broadcast companies and news agencies), and live media. The latter can be associated with the five trends towards *UbiMedia* cited earlier, i.e. they range from **commodity media** (background) recordings to professional mass media contents.

STEP 2: semantics generation: the second step yields, in essence, the ZBR. Advancements in all areas of media indexing can be expected to improve the feasibility of the middle arrow depicted in fig. 2, step 2: the extraction of semantic information from signal-based representations of

the media. Another processing path becomes possible along with the proliferation of UC devices. We call this principle **tagging**. Thereby, media capturing devices cooperate with smart devices. The cooperative goal is to associate the media captured with as much semantic information as available 'outside' the physical signals captured. Such tagging may range, e.g., from timestamps (coming from a time server) to location information (coming from a location-aware smart device in proximity) to related events (coming from the calendar software of the user(s) associated with the captured media) to the identities of the humans captured on audio / video (coming from their active badges), etc. Obviously, a substantial part of the ZBR may be generated by tagging. Given sophisticated tagging, one can imagine queries like "give me all videos where Jack and Tom met in the lobby to hand over documents" to become feasible without applying any conventional video indexing techniques. As to the third box in fig. 2, tagging for computer-generated media is a matter of proper programming in the generating software and should be rather straightforward.

STEP 3: broadspectrum representation. The ZBR is of course at the heart of this step. As mentioned, however, the ZBR will not be sufficient for many purposes. On the other hand, ZBR provides an excellent basis for automated translation into other modalities and for the generation of representations with variable fidelity (based on, e.g., the degree of DCT quantization, the sampling rate, etc.). In other words, ZBR is only one out of infinitely many possible representations. (In fact, ZBR may itself provide for a variable degree of detail and fidelity). At least the ZBR plus the full rate capturing in the original modality must be retained in order to preserve the full spectrum of possible representations. Note, however, that this approach is not feasible for UbiMedia since it does not help to master the problems on the storage side (the highest-fidelity representation will always consume the largest amount of storage). We will resume this issue below. The multi-fidelity multi-modality representations depicted in fig. 2 are closely related to the ideas brought

forth in IBM's transcoding projects and products; in fact, this approach of variable representations is based to IBM's info pyramid concept. The reader may consult [4,7] for more insight into how different degrees of fidelity and different modalities can be mapped onto terminal devices with varying capabilities (e.g., WAP-enabled cell phones vs. multimedia PCs).

STEP 4: propagation in time and space. As mentioned several times above, novel approaches to the handling of media must be developed in order to cope with the massive amount of multimedia data to be expected from the UbiMedia vision. We use the term **UbiMedia-Net** here to denote novel approaches to storage and dissemination. We propose to base these approaches on the principles of information handling inherent in human behavior: humans make very heavy use of principles like those of locality and aging. This means that, *for instance*,

- * humans memorize information as they perceive it, distinguishing 'common' from 'unusual' events, events in proximity or in 'regions of interest' from 'remote events', etc.
- * memory fades gradually i.e. becomes less and less accurate, in particular if not refreshed
- * as humans interests focus or dissolve, so does the density distribution of information gathering
- * fading is not a simple time-linear process but follows more sophisticated rules (short vs. long term memory, much improved retention after just one reminder, etc.)
- * humans disseminate information in proximity with higher fidelity than remotely (cf. "family talk" vs. "smalltalk with acquaintances").
- * In addition to locality and aging, some characteristics of human information handling exhibit similarity to other natural laws such as the laws of gravity ('gravitation forces' attach information to both the human and the location of origin; these forces 'compete' as the human changes location).

Obviously, such principles can be translated into the UbiMedia-Net approach as rules for the capturing and erasure, dissemination and migration of media.

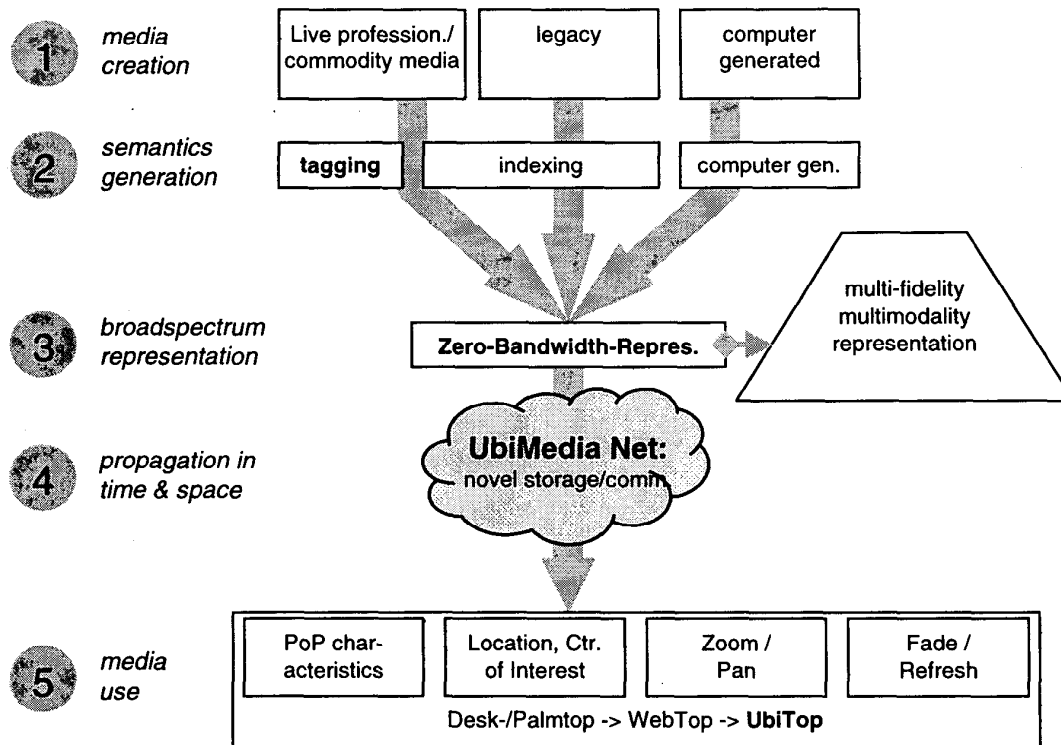


Fig. 2: The five-step media chain

In our envisioned UbiMedia-Net network, as time goes by, 'memory fades' i.e. hi-fidelity and storage-intensive non-original modalities are gradually erased (if not sufficiently refreshed i.e. used by humans or applications). In the end, only the ZBR is retained (possibly even reduced in degree of details). As users move and as remote applications connect to (real-time i.e. conferencing, or stored) media, the ZBR is replicated, accessed, or transmitted first. Only when users or applications focus on (parts of) media, more detailed and higher fidelity representations are accessed.

STEP 5: Media Use. The final step in the media chain is where media are actually demanded by humans and / or applications. We use the term **UbiTop** here to refer to the UC compliant successor to DeskTops and WebTops (a term used to denote desktops reaching out to the Internet). UbiTops are envisioned model both the nomadic user and the footsteps which s/he leaves at places visited (remember that true UC is independent of a physical location serving as 'home base'). UbiTops also cope for varying and cooperating terminal devices used. Finally, the UbiTop model is envisioned to model organization of and access to a user's (multimedia) information in compliance with the above-mentioned UbiMedia-Net vision. As fig. 2

depicts, this means, in particular, that the modality and fidelity of media as transported to the (distributed) UbiTop depends on and changes with

- * the characteristics of the PoP (point-of-presence), i.e. the current cooperating terminal devices (headset? silent room with co-presence? video capabilities (which resolution)? bandwidth cost/availability? etc.)
- * the current physical location and virtual locations (points of interest) of the user
- * the highly dynamic degree of focus on subject matters (zoom/pan)
- * the rules for erasure (of different modalities and different fidelity levels) of stored media (fade/refresh).

Note that UbiMedia-Net and UbiTop mutually influence each other, e.g., in that information access from the UbiTop triggers a refresh operation which limits fading.

5. Summary and Conclusion

Following an introduction to the terms and visions of ubiquitous computing (UC), we have identified three major challenges for software engineering, coined as UC engineering (service discovery, awareness management, and AAAS), scalable computing, and more advanced ap-

proaches denoted as 'large-scale computing'. In this broad context of software engineering challenges, we took a closer look at 'UbiMedia', our term for the envisioned multimedia challenges in the UC era. We identified the need for not only for 'marginal' adaptations of software engineering methods or models, but for substantial changes with respect to the overall understanding of how computers and networks should work in a UbiMedia based world. The major suggestions towards such a change were explained based on a five-step 'media chain'. Major ideas introduced as part of this chain were explained, called commodity media, tagging, ZBR, UbiMedia-Net, and UbiTop. Since the UbiMedia challenges and the general software engineering challenges for Ubiquitous Computing are fairly substantial, the article is to be considered a long-term call for ubiquitous multimedia research.

6. References

- [1] E. Freeman, D. Gelernter. "Lifestreams: A Storage Model for Personal Data", *ACM SIGMOD Bulletin*, March, 1996
- [2] H.W. Gellersen, M. Beigl: "Ambient Telepresence: Colleague Awareness in Smart Environments", *P. Nixon et al. (Eds.) Managing Interactions in Smart Environments*, Springer Verlag Heidelberg etc., 1999, pp 80-88
- [3] N. A. Gershenfeld, *When Things Start to Think*, Henry Holt & Co., Inc.; New York, NY, USA, 1999
- [4] "IBM Transcoding Solution and Services", *IBM Whitepaper*, www.research.ibm.com/networked_data_systems/transcoding/transcoddef.pdf
- [5] K. Kelly, *Out of Control*, Perseus Books, Reading, MA; USA, 1994.
- [6] T. Lewis, "UbiNet: The Ubiquitous Internet Will Be Wireless", *Computer*, Vol.32 No. 10 Oct. 1999, pp. 128ff
- [7] R. Mohan, J. Smith, C.-S. Li, "Adapting Multimedia Internet Content For Universal Access" *IEEE Transactions on Multimedia*, March 1999, pp. 104-114.
- [8] M. Mühlhäuser, "Content Development for the Internet as a Mass Medium", *Proc. MSE'98 Multimedia Software Engineering*, Kyoto, Japan, April 20-21 '98. IEEE CS Press, Los Alamitos, 1998 pp. 2-9
- [9] F. Nack, A. Lindsay, "Everything you wanted to know about MPEG-7", *IEEE Multimedia*, Vol. 6, No. 3, July-Sep. 1999, pp. 65-77 (Part II in Vol. 6, No. 4).
- [10] D. A. Norman: *The Invisible Computer*. MIT Press, Cambridge, MA, USA, 1998
- [11] R. Pascoe, "Service Discovery: The Missing Link for Networked Systems", *Embedded Systems Development*, Vol. 3, No. 4, Apr. 2000, pp. 31-39.
- [12] L. Press, "The Post-PC Era", *CACM* Vol. 42, No. 10, Oct. 1999, pp. 21-24.
- [13] B. Rhodes, N. Minar, J. Weaver, "Wearable Computing Meets Ubiquitous Computing", *Proc. 3rd International Symposium on Wearable Computers (ISWC '99)*, San Francisco, CA, October 18-19 1999, pp. 141-149.
- [14] M. Weiser, "The Computer for the 21st Century". *Scientific American*, September 1991, pp. 94-104.
- [15] G. Weiss, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, Cambridge, MA, USA, 1999