Information Exposure Control through Data Manipulation for Ubiquitous Computing

Boris Dragovic
Computer Laboratory
University of Cambridge
Boris.Dragovic@cl.cam.ac.uk

Jon Crowcroft
Computer Laboratory
University of Cambridge
Jon.Crowcroft@cl.cam.ac.uk

ABSTRACT

The vision of Ubiquitous Computing [22] creates the world in which information is omnipresent, migrating seamlessly through the environment to be accessible whenever and wherever needed. Such a vision poses substantial challenges to information security and privacy protection.

Unlike in traditional, static, execution environments, information in the Ubiquitous world is exposed, throughout its lifetime, to constantly varying security and privacy threats caused by the inherent dynamicity and unpredictability of the new computing environment and its mobility. Existing data protection mechanisms, built for non- or predictably slowly-changing environments, are unable to strike the balance in the information availability vs. security and privacy threat trade-off in the Ubiquitous world thus hindering the feasibility of the overall vision.

In this paper, we present our initial work on a novel paradigm for information security and privacy protection in the ubiquitous world. We model security and privacy threats through sets of contextual attributes and mitigate the projected risks through proactive and reactive data format transformations, subsetting and forced migrations while trying to maximize information availability. We also try to make the approach flexible, scalable and infrastructure independent, as required by the very vision of the Ubiquitous Computing.

1. INTRODUCTION

Traditional computer applications expect a static execution environment. Such environments imply non- or slowly-evolving information security and privacy threat models. Existing security models and mechanisms have been built on the assumptions of such environments. Ubiquitous computing is based on a fundamentally different vision [22], aiming for computation being unobtrusively and indistinguishably embedded in the environment around us, providing us with information whenever and wherever we need it. The inherent dynamicity and unpredictability of such an environment poses fundamental challenges to information security and privacy protection [19].

NSPW 2004 Nova Scotia Canada

Information, represented by data objects, is, throughout its lifetime in a Ubiquitous system, whether contained on a device or within a communications channel, exposed to constantly changing set of security and privacy threats. This is due to the data objects migration, containing devices' migration or other environmental, context, changes. To facilitate the vision of information omnipresence we need novel security paradigms which will ensure maximum information availability while limiting its exposure to the threats. Considering the mere complexities involved in reasoning about information security risks and the fact that one of the aims of the Ubiquitous computing is for the computation to be transparent and disappear into the periphery of our mental activity it is unfeasible to expect humans to be able to reason and act effectively to protect the information themselves.

Past research in the field of Ubiquitous systems security has focused mostly on adapting the existing security models and mechanisms to the new environment. One of the focal points has been the recognition of the importance of the context information as the means of adapting authentication and authorization mechanisms to suit the Ubiquitous Computing requirements (e.g. [5, 18, 21, 4, 17] etc.). Adapting existing security mechanisms for application in Ubiquitous systems certainly provides a sound foundation, however, it does not address the specific issues of the Ubiquitous computing.

In the presented work, still in its very infancy, we propose a novel, Ubiquitous computing specific, information security and privacy protection paradigm. In [12] Myers and Liskov note that security models have two goals: to prevent accidental or malicious destruction of information and to control the release and propagation of information. The paradigm we propose falls into the latter category. We address the problem of controlling information exposure to the surrounding while it is being legitimately accessed by a authenticated and authorized user.

For every data object existing in the Ubiquitous world, we assess security and privacy risks in its environment, depending on its sensitivity, type, access method etc. and try to mitigate the risks by: manipulating its format and changing the environment to a less risky one. The former relies on the fact that inherent in the data format is the level of quality of represented information. The latter tries to directly avoid the risk itself.

The model assumes a *cooperating user* scenario i.e. it is built as an aid for users to effectively protect and reason about security and privacy of the information in their possession.

2. MOTIVATION

The advent of mobile computing, in the form of laptops, Personal Digital Assistants (PDAs), mobile phones etc. together with the advances in communications technologies has enabled us to access information, and thus work, on the move and has increased our efficiency and quality of life. As a consequence, we break well established security perimeters by taking wealth of sensitive information with us, on our mobile devices.

Consider how many times in the past several years has media reported about confidential information (trade secrets, intelligence information, personal information etc.) leaked by laptops, PDAs, mobile phones etc. being stolen from pockets, handbags, cars etc. in public places. Imagine how many times such information has leaked by conversations being overheard, displays being peeked at over shoulders, data transmissions over publicly open communications links, residual data left on public output devices such as displays or printers, etc. In all of the situations the information was legitimately stored on the devices and accessed by authorized users. The existing security mechanisms were unable to prevent the exposure, mitigate the risks in the environment or simply warn the users.

To ensure data availability, information omnipresence, to legitimate users on the move, and still balance it with the security and privacy risks present in the environment we need security paradigms designed to balance the information content provided to a legitimate user while maintaining the risk of the information exposure to the surrounding at an acceptable level.

Part of the motivation for proposing such an approach stems from the observation of human everyday behavior. How many times have we lowered the volume of our voice, switched from a speakerphone to telephone headset or changed topics when we realized our conversation could be overheard? This is nothing else but matching the form and characteristics of information to the perceived security and privacy risks in the environment. In this project, we try to mimic that behavior in the Ubiquitous computing arena and thus maximize the information availability versus security and privacy risk exposure trade-off.

The most obvious application is within the area of user interfaces. Projects like the Personal Servers [20], Steerable User Interfaces [8], Virtual Network Computing (VNC) [14] etc. emphasizes the use of environmental output devices (displays, audio interfaces etc.) for accessing information stored on Ubiquitous devices. However, the approach is applicable to all data objects, throughout their lifetime in the Ubiquitous world, as, depending on their sensitivity, characteristics of the devices they are stored on and environmental attributes, they are constantly exposed to certain security risks. For example, a mobile phone containing data is under a higher risk of being stolen in a public place than in an individual's office, let alone if the owner is not in its immediate proximity.

We also draw from the access control paradigm. However, the proposed approach differs in three fundamental ways: it controls adverse information exposure level to the surrounding while it is being accessed by legitimate, access control authorized, user - thus operating at a different level; we drift away from binary access decisions and create an continuum between granting and refusing access by performing fine-grained information content exposure control through proactive and reactive protective actions; data is protected throughout its lifetime, by being constantly tracked and security risks reevaluated triggered by context changes. The proposed approach can be seen as complementary to traditional notion of access control and is not intended to replace it as it addresses a different issue.

Traditional security mechanisms were designed to operate in environments with well established data security and privacy threats and within secure perimeters. Their main task was granting or refusing access to information. Ubiquitous computing vision breaks this model by making the notion of a secure perimeter deprecated and requiring information to be available where and when the users need it. To ensure the information availability we need mechanisms to protect it in the environments in which it exists.

3. CONTAINMENT - DATA OBJECT CENTRIC MODEL OF THE WORLD

Our work focuses on assessing information security and privacy risks in the Ubiquitous Computing environments and providing adequate protection. The term used for a particular information representation throughout this text is *data object*. As data objects represent a central focal point in our research we model the world adequately.

3.1 The Data Model

When describing a particular information representation we use the term *data object* as suggested by Policroniades [13]. A data object is not an equivalent to the traditional notion of a *file*, although it can be regarded as such. Data objects bind data of certain common attribute, e.g. within a HTML file, a picture can be an data object, each paragraph can be a separate data object etc. One of the advantages of such data model is a high degree of flexibility in data manipulation it provides. In our research, data objects represent collections of data of the same security sensitivity, as determined by a security policy. For example, traditionally, a classification level of a document containing information of heterogeneous individual sensitivities is dictated by the most restrictive constituent label. By regarding the document as a collection of data objects, we can provide higher degree of information availability and finergrained access controls by matching individual data objects' classifications, within a document, to the threat level they are potentially exposed to e.g. in some situations, we would be able to grant access to certain paragraphs of text, classified at secret, omitting satellite images classified at top secret. The data model resembles research efforts into multilevel database management systems [1].

3.2 Containers and Containment

We define a *container* to be physical or virtual enclosure in which an data object or a *lower level* container exists, either fully or partially. Examples of containers are storage devices (e.g. a device's memory, hard drive, etc.), displays, audio devices, communications links, virtual circuits etc. Further, containers are PDAs, mobile phones etc. just as well as physical spaces in which these reside.

Naturally, containers can be nested in a container hierarchy. However, unlike in the location hierarchies, elements, nodes, in container hierarchies are not necessarily unique. Containers are identified by their classification based on a set of characteristics, depending on which multiple instances of a particular container may exist concurrently. For example, we may define a container to be determined by a set of cryptographic protocols available over a communications channel. Thus, any communications link providing the specified services represents an instance of the container. In

this work, we define containers as with respect to perceived threat model related attributes. Container nesting is equivalent to physical nesting, e.g. a storage device within a mobile node etc.

Every container has an internal state, for example the traditional notion of context applicable to physical spaces. Container state is inherited down the container hierarchy as constrained by *container transparency* rules. In other words, every container creates an environmental state for its contents. A state is expressed in therms of a set of state attributes and their respective values.

Containers can be classified orthogonally by their *type* and their *class*. The latter denotes the container's primary functionality: a room, a display, a storage device, a communications channel etc. We say that two containers are of a type if they are transparent to the same set of state attributes and their values when exposed to equivalent set of state attributes and their values from a parent container.

Containment denotes the state of a container or a data object being within a container together with any relevant state attributes, their values, and any applicable rules, such as the transparency rules.

By definition, the leaf nodes of the container hierarchy represent data objects. A data object may migrate among the containers but must remain at the bottom of the hierarchy. Containers that are the direct parents of data objects are called *first-level* containers. For example, storage devices, displays, communications channels can be classified as firs-level containers. In simple terms, a device or medium on which a data object directly resides and can be extracted from.

4. THREAT MODELING

4.1 Threats - an Informal Definition

In Section 3 we specified the role of state attributes and their values in our model of the world. The attributes that we use are chosen to describe potential threats as perceived at the data object level in the environment. At any point in time, security and privacy threat a data object is exposed to is represented by a set of set state attributes and their values the respective first-level container is transparent to. This may be illustrated by an example where contents of a display residing within a glass-walled room are visible outside the room as opposed to a display within a room with solid concrete walls.

Now, we can provide a higher-level definition of a first-level *container type*. Two first-level containers are of the same type if a same data object is exposed to the equivalent security and privacy threat in both containers under the same environmental states (security relevant state attributes and their values the first level-container is affected by).

4.2 Levels of Exposure (LoEs)

Security and privacy savvy users often find asking themselves: is the information displayed on my screen visible to anyone apart from me? Can the audio I am listening to, or the conversation I am conducting, be overheard by someone? What is the risk of adverse information exposure, or simply exposure in this context in the rest of the paper, if a particular communications link is used? etc.

More generally, *Levels of Exposure* (LoEs) quantify and qualify the extent to which a piece of information is accessible to its surrounding at any particular moment i.e. the degree of possible information

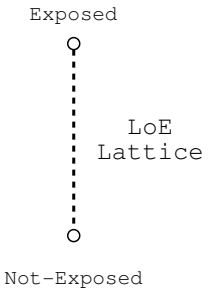

Figure 1: Two-level LoE model.

leakage. LoEs do not account for type of degree of data access as exercised by a legitimate, as determined by an complementary authorization mechanism, data user.

LoEs model is defined at an organizational level, along a wider security policy. Its semantics is uniform across the ubiquitous system. LoEs apply to all the data objects existing within the system and depend directly on each individual data object sensitivity level, as determined by a security model employed. For every instance of a data object, throughout its lifetime in a ubiquitous system, there may be only one LoE *active* at any particular moment. Each of the LoEs is associated with a set of proactive actions to mitigate the implied security and privacy risks.

The simplest, and possibly sufficient for majority of applications, Levels of Exposure model is a "not-exposed", "exposed" two-level model. The levels denote cases in which there is none or a credible possibility of information leakage to any third party respectively. The two levels can be regarded as the two extremes of a LoE lattice representing a finer-grained model (as depicted in Figure 4.2). This simplistic model shall be used in all further discussion about LoEs models.

4.3 LoE Modeling

Consider a piece of classified data being displayed on an overhead screen. Depending on the contained information sensitivity, if the screen is within a possible visibility field of a third party, the data would be labeled at the "exposed" LoE. An instance of the same information, contained on a storage device would be unaffected under the same circumstances. However, should the proximity of the device to the owner decrease, the LoE of the stored data object would be changed to "exposed". Unclassified data would not be affected in either case and would permanently remain at "not-exposed" LoE.

The example illustrates that environmental state triggering any particular LoE for a data object depends directly on: the data object's security sensitivity, as determined by a wider security policy; and on the container type the data object resides on. Therefore, for each sensitivity level and for each container type, we define one or more sets of, possibly overlapping, state attributes and their values that trigger every applicable LoE. Depending on the LoE model, not all exposure levels need to be defined for all first-level container types or data-object sensitivity levels.

Figure 4.3 depicts a mapping of the two-level LoE model to a lattice-based security model and the influence of container types on state attributes and their values triggering each of the exposure levels. Representation of the trigger attributes and their values in

the figure is rather simplistic whereas practice it would consist of first-order logic expressions.

5. CONTEXT RELATED ISSUES

The previous Sections have shown high dependence of the proposed model on the context-related information for establishing container attributes for building containment hierarchies and evaluating security and privacy threats. To provide continuous model operation and its graceful degradation we have to insure the continuity of context-information availability with guaranteed minimums.

According to the nature of the contextual attributes we can roughly divide containers into two categories: first-level containers (Section 3); and higher-level containers. First-level containment is determined by tracking data objects and its attributes are expected to be pre-set e.g. through a certification process. Higher-level containment, on the other hand, along with the security attributes determining threat models need to be "sensed".

We envisage three ways of establishing context-related information:

We envisage three ways for establishing context:

- Ubiquitous Unit's individual capabilities.
- Via trust-based collaboration groups.
- Use of dedicated infrastructure.

5.1 Ubiquitous Units' Context Provision

A *ubiquitous unit* is defined as any computationally capable individual entity in a ubiquitous system e.g. a PDA, a mobile phone etc. A unit may comprise several containers e.g. a storage device, a display etc. Although very few devices today have built-in dedicated context sensing capabilities they can provide a guaranteed minimum of context-information necessary for the model operation.

Firstly, a minimum of the first-level in the containment hierarchy is determined form a pre-set specifications, in terms of attribute sets, transparency rules etc. of all containers within the device. To track data objects' migrations among the first-level containers system-level mechanisms are utilized. This provides the most coarse grained model operation.

Secondly, much of the today's ubiquitous devices' built-in functionality can be used for limited context inference. For example, a reachability of local wired LAN may mean physical presence within a secure perimeter, similar applies to the visibility of *landmarks* via some communications technology; Bluetooth connectivity to a tag user wears, e.g. a personal mobile phone most of us constantly carry around, may be used to determine owner's proximity; audio analysis combined with high level diary information or a local mobility model may yield the fact of an activity taking place, etc.

5.2 Collaboration Groups

We expect users to carry multiple ubiquitous units at any one time each of which will possibly have different capabilities commensurate with their primary functionality. Collaboration groups, based on trust [18], can be formed among such sets of units for aided context awareness or simply increased confidence. The above example where device proximity, through short-range Bluetooth visibility, is used to determine owner's presence illustrates this point.

In infrastructure-rich environments, dedicated, high-confidence, context awareness services can be used, by joining a unit's collaboration group. For example, while in an Active Bat [9] enabled environment, a unit may form a collaboration group with the location service to obtain accurate and high-confidence location information.

5.3 LoE Establishment Confidence

As specified in Section 4 every LoE is defined by a one or more sets of attributes and their values that trigger it. With every attribute value a capturing confidence is associated. Only when the confidence of all the values in a set are above a threshold is the particular LoE triggered.

Furthermore, every ubiquitous unit has to be certified for each LoE it can establish for any information sensitivity level at any available container type within the unit. This mechanism may be used to determine the maximum sensitivity of the data a unit may accommodate - should it not be certified to establish LoEs above a certain sensitivity level.

5.4 Context Abstraction and Modularity

Abstracting away low-level sensory information from multiple sources possibly spanning multiple devices, capturing and reasoning about errors and confidence is outside the scope of this project. For this purpose we refer to the work done on the Context Toolkit [7]. Context Toolkit provides a framework for building a modular and flexible context mapping from local and remote low-level sensory information to a higher level context descriptions through a set of: *widgets*, abstracting away the notion of sensors; *interpreters*, raising the level of abstraction of a piece of context information; and *aggregators*, collecting multiple pieces of logically related context information. The framework was designed with to aid rapid prototyping of context aware applications. Context Toolkit does not support the notion of trust, confidence or dynamically, run-time, resizing context-establishment collaboration groups in the form we require them. We intend to adapt the Context Toolkit approach to our needs. Figure 3 depicts schematically the process of abstracting low-level contextual information both locally and form a remote sensing unit.

6. PROTECTIVE ACTIONS

Once a credible risk of a threat for a data object within a container is established, resulting in a LoE activation, protective actions need to be taken to mitigate the implied security and privacy risks. The result of performing the actions is lowering the LoE of the data object. Possible protective actions are classified into three categories:

- Data object's format transformations.
- Subsetting.
- Container *hopping*.

By manipulating a data object's format we exploit the fact that different forms in which data exists in different contexts provides varying levels of information content to its surrounding. We can divide the transformations into two orthogonal sets of categories

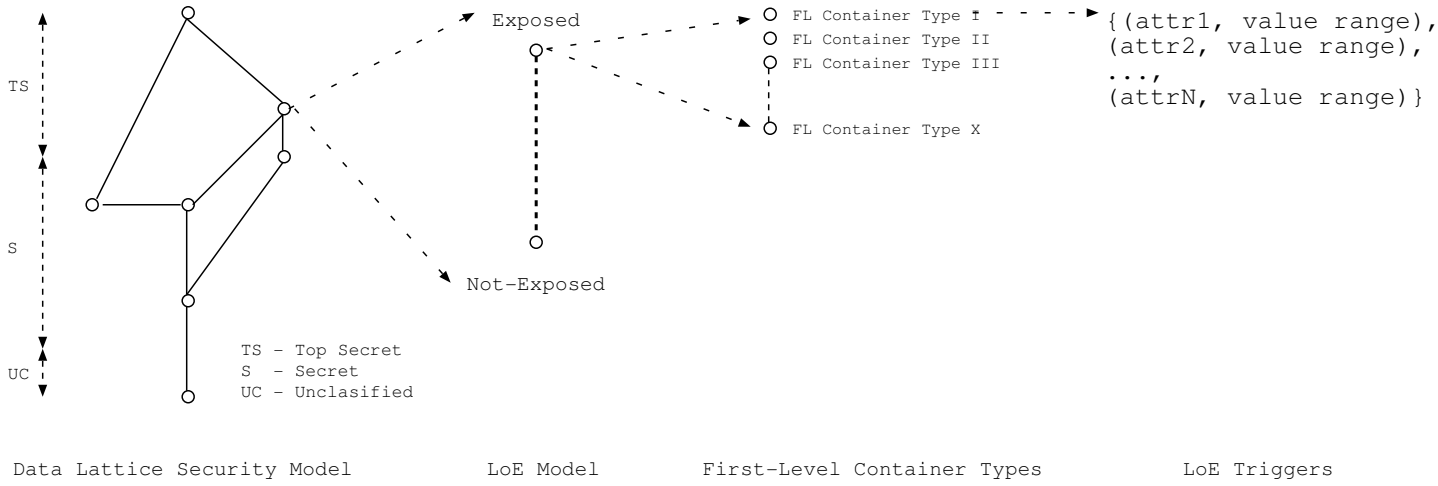

Figure 2: LoE mapping.

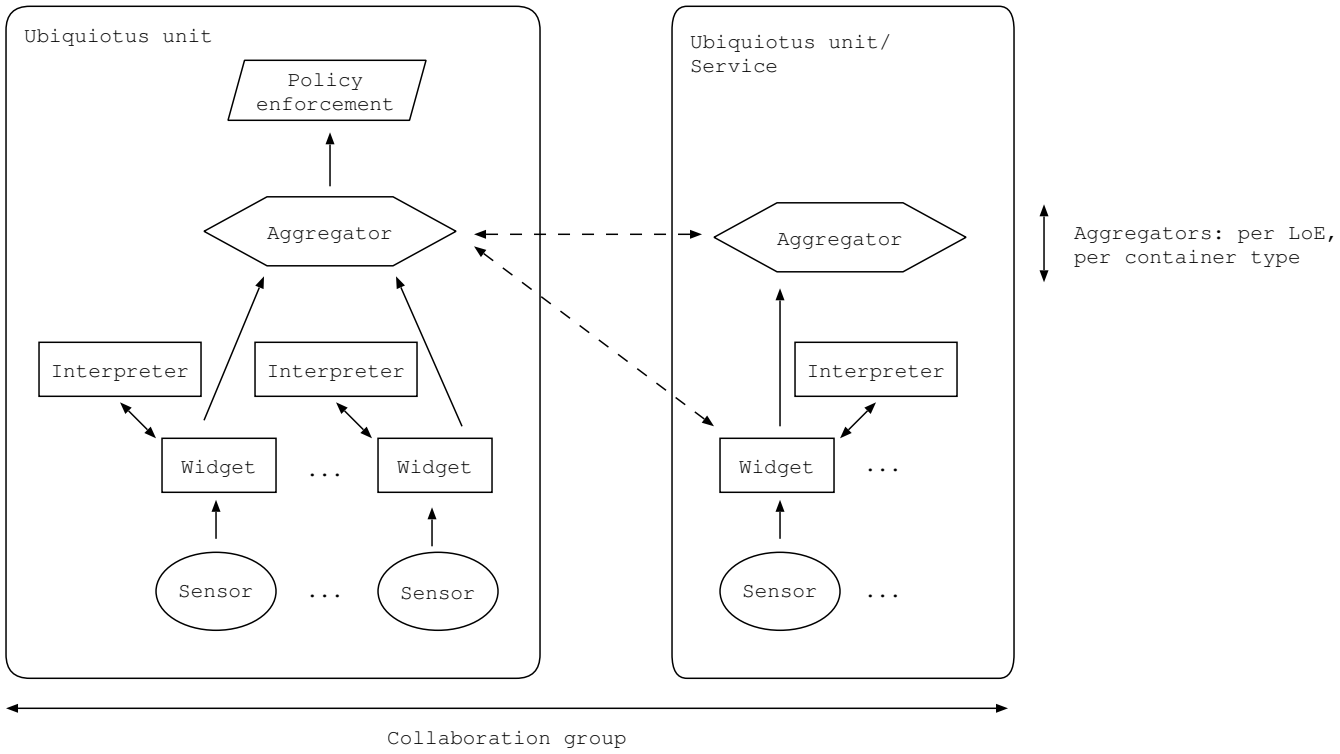

Figure 3: Context establishment w/o collaboration group.

depending on their granularity and reversibility. With respect to the granularity we differentiate among the bulk transformations such as encryption, image blurring, information hiding etc.; and fine-grained operations such as anonymization, feature selection etc. Choice of appropriate set of format transformations is influenced by the LoE they apply to, data object type and current format and first-level container class the data object is contained within. For example, data encryption makes little sense if the data is contained within a display. A more appropriate choice in this case might be simple screen blanking.

Subsetting, introduced in [2], can be regarded as a coarse-grained format manipulation where data objects can either exist in full or not at all within a container based on the (would-be) experienced LoE.

Container *hopping* represents switching data objects between available containers of the same class but different type (Section 3) and thus lower the experienced LoE. Unlike the format transformations this approach mitigates the perceived risks while maintaining the information availability level. For example, while at a "not-exposed" LoE a sensitive document may be displayed at an overhead display, once a "exposed" LoE becomes active for the document, at the container, the sensitive information, e.g. access codes, may be transfered to a mobile phone for displaying while they are obfuscated at the overhead display. Another compelling example would be in cases of data transmission across communications channels when there are multiple technologies available in the environment. Sensitivity of each of the data object's transmitted could be matched to security and trust characteristics of the available links taking into account issues such as performance, transmission urgency, resource consumption etc. Container hopping can be, as hinted in the example, performed among collaborating units in a similar way to the context establishment, as outlined in Section 5.

Every protective action, or a set of actions, comes with various resource costs. We envisage multiple alternative choices to be available in any single situation. The question arises on the choice of the "best" (set of) actions. There are multiple factors involved in making the decision ranging from device capabilities to the characteristics (temporal, spatial etc.) of the respective threat model. We intend for every unit to have the ability of making an adequate choice among the alternatives.

Protective actions are another aspect of the project where collaboration groups (Section 5) come into play. As outlined previously, container hopping may be performed among containers spanning a collaboration group. Furthermore, we expect more-capable group members to be able to perform certain required protective actions on behalf of weaker members.

7. RESEARCH DIRECTIONS AND ISSUES

At the presented, initial, state of the project there is a number of possibilities for the direction of the research. In this Section we give a hint at what we plan the main direction to be outlining some immediate issues.

7.1 The Model of the World

In Section 3 we presented our idea of the world model that satisfies the project requirements, namely modularity, flexibility, continuous operation in cases of context-information quality and quantity unpredictability and general heterogeneity. One of the primary steps is to provide a formal description of the model including a supporting logic and an inference algebra alike Egenhofer's [15] container surface algebra. The issues of dynamic model formation and reasoning within a unit and among members of a collaborating group, in a distributed fashion, need to be resolved.

7.2 Policy

Throughout the paper, at several points, we have assumed existence of several types of policies, most notably the LoE definition policies binding contextual attributes, containment hierarchy and data object sensitivity to perceived threat models and protective actions. Furthermore, we will need formal descriptions of the protective actions and policies for discriminating between the alternatives (Section 6); policies for formation of trust-based collaboration groups; etc. For binding LoEs to protective actions we are investigating policy languages able to express obligations, such as [6].

There is a forseeable cost associated with administering and maintaining policies. It has been shown in the past that practical acceptance of new security mechanisms has been significantly impacted by the ease of necessary policy related work. The more so will be the case in a Ubiquitous Computing system where the operational transparency is of a highest value.

7.3 System Architecture

The operation of the proposed model largely depends on the ability to continuously track data objects as they move among containers. To provide this ability in the traditional system architectures would require a complex mechanism spanning operating system privileged layer and application layer alike to be able to link any piece of data to a specific policy at any point in time. Performing protective actions complicates this further as it requires an additional level of application awareness or hugely increased complexity at the operating system level. It is our opinion that the effort required to implement the proposed approach fully within a traditional system architecture would be of limited value and result in an patchwork of application-specific solutions.

The advent of Ubiquitous computing requires a change in which we think about computing and thus the way in which we design and build systems and applications. Therefore, we are set to investigate system architecture features that would facilitate seamless integration of the proposed model.

7.4 Socio-Technological and Acceptance Issues

As stated in the introduction, the proposed model aims at aiding the users in protecting the privacy and security of their data in face of increasing technological complexities involved in an Ubiquitous Computing environment. It has been proved in the past that humans see security mechanisms as obstacles and try to circumvent them. This is even more emphasized in the Ubiquitous world where the computation needs to disappear in the background and so do the security mechanisms. On the other hand, users need to be aware of the reasons causing certain security decisions [3] as otherwise they are likely to feel uncomfortable.

The cooperating user scenario puts us in the position to seek user feedback in situations of insufficient context information or ambiguity. We also expect to have multiple models operating simultaneously enforcing multiple policies e.g. a corporate policy and a user's private policy. Users should be able to specify their preferences with respect to protective actions in line with respective policy.

Another question that poses itself is of the general cost of practically rolling out the model and the "critical mass" required for the model to be useful. With the previously outlined system architecture issues arising and considering real commercial interests and market forces involved, the cost of a full, generalized, model deployment would be significant. However, we envisage application-specific fieldings of the model as well as its employment in closed-environments. The model itself can be sufficiently useful even on an individual user's basis for personal information protection, let alone at an organisational level.

7.5 Computational complexity and resource requirements

It is unclear, at this stage, how the complexity and resource requirements of operating the model will range based on its granularity and representation, differing policies, levels of context awareness, structure of the protective actions and other factors. This directly impacts the feasibility of the practical deployment of the model in the world characterized by a high degree of heterogeneity in all aspects. We realize that there will always be "smaller" and "weaker" devices than the minimum required and that is one of the reasons of introducing the notion of the trust-based collaboration groups. Building a prototype will aid us in assessing these factors which may impact on the overall model design.

7.6 Scalability

Possible scalability issues with the proposed approach lie within the policy related complexities and dependence on an organizational-wide security policy, container-type classifications and ubiquitous units certification etc. These factors may confine deployment of a single model within one administrative entity. There are, however, no limiting factors in infrastructural dependence or confinement within a geographical boundaries. Furthermore, we envisage the possibility of multiple model instances operating in parallel.

8. RELATED WORK

Generalized Role Based Access Control (GRBAC) [5] represents the most prominent effort to extend an access control mechanism with context awareness. In addition to traditional subject and object roles, GRBAC defines environmental roles which is used to capture security relevant aspects of the environment. Although GRBAC policies are useful in restricting access to data based on environmental attributes such as e.g. time of day, day of week, system load, session duration, etc. it is, in its essence, an access control mechanism and can thus be regarded as complementary to our paradigm (Section 2).

Scott et al. have, in their work on spatial policies for sentient mobile applications [16], addressed the issue of controlling execution behavior of mobile applications through location-aware policies. The work provides useful insights into formal development of context-aware policies, their expression and reasoning about them. It also gives an example of associating policies with mobile entities, in this case active, and their enforcement.

Work that comes closest to ours, in terms of the problem addressed, the target environment and the generality, is by X. Jiang et al. [10, 11]. They divide Ubiquitous world in a set of *information spaces*, delimited by physical, social and activity-based boundaries, with the aim of protecting privacy by controlling information flow across the spaces. Three main privacy properties of data are identified as: persistence, observational accuracy and observational confidence. By manipulating various aspects of the three, as data crosses information space boundaries, the proposed model controls *information asymmetry* between the communicating ends - which the authors identify as crucial aspect in protecting privacy. The notion of information spaces resembles the one of containers, and manipulating data properties is a superset of data format transformations. However, we do not try to model or affect information flow between entities, we try to mitigate security and privacy risks present for data objects in the environment in which they exist while providing the maximum information content availability to legitimate users.

A number of research efforts have been, and still are, aimed at restricting the availability of data to certain computing environments such as e.g. firewalls or proxy servers. More recently, we have seen a proliferation of Digital Rights Management solutions, through various trusted computing platforms. These are also, in a more general way, aimed at confining the availability of data to predefined, licensed, computing environments. Research in the area of information flow control, such as e.g. [12], has, again, has the same aim. The concept we propose in this work provides a more general, not application specific, solution in the abstraction of a container as a data object's physical enclosure, be it a communications link, a storage device, a display or any other. It then puts this into a big picture by considering issues in the physical surrounding. And, finally, by understanding the threat models and data sensitivity maximizes information availability while protecting its security, steering away from the traditional binary decision models.

9. CONCLUSION

Ubiquitous Computing vision promises disappearance of computing into the periphery of our mental activity and its full embedding into the environment around us. Information will be at our disposal wherever and whenever we are. This vision represents a substantial departure from the notion of computing as we now it. Traditional information security and privacy protection mechanisms, although portable, as demonstrated by numerous research efforts, merely extend the functionality into the new environment but are unable to address the Ubiquitous Computing specific challenges.

In this work, we identified a need for a novel security paradigm to protect information security and privacy while maximizing its availability throughout its lifetime in the Ubiquitous system. The paradigm is aimed at controlling the level of information exposure to the surrounding while accessed and in possession of a legitimate user or simply migrating through the environment. In other words, continually throughout its lifetime. The core motivation lies in the realization that, in the Ubiquitous world, information is constantly exposed to variable levels of security and privacy risks, caused by the requirement of its omnipresence. Information protection is accomplished through: transformation of the format in which the information exists in the environment i.e. choosing appropriate information representation restrictiveness; and by modifying the threat model information is exposed to by switching the environment.

As the work is still in its very infancy, the proposed approach is presented on a high-level with the main aim being exposing the idea and motivation and identifying the issues present at this stage.

10. REFERENCES

[1] S. Castano and et al. *Database Security*. ACM Press, New York, NY, 1995.

[2] D. Chalmers, M. Sloman, and N. Dulay. Map adaptation for

users of mobile systems. In *Proceedings of the 10th International World Wide Web Conference (WWW-10)*, 2001.

[3] M. Chalmers. Seamful design and ubicomp infrastructure. In *Proc. Ubicomp 2003 Workshop 'At the Crossroads: The Interaction of HCI and Systems Issues in UbiComp'*, 2003.

[4] M. J. Covington, P. Fogla, Z. Zhan, and M. Ahamad. A context-aware security architecture for emerging applications. In *Proceedings of the ACSAC'02*, 2002.

[5] M. J. Covington, M. J. Moyer, and M. Ahamad. Generalized role-based access control for securing future applications. In *Proceedings of the National Information Systems Security Conference (NISSC)*, October 2000.

[6] N. Damianou, N. Dulay, E. Lupu, and M. Sloman. The ponder policy specification language. *Lecture Notes in Computer Science*, pages 18–39, Jan 1995.

[7] A. K. Dey, D. Salber, and G. D. Abowd. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction Journal*, 16(2-4):97–166, 2001.

[8] G. S. P. et al. Steerable interfaces for pervasive computing spaces. In *PerCom'03: Proceedings of the IEEE International Conference on Pervasive Computing and Communications*, 2003.

[9] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. In *Proceedings of the 5th annual ACM/IEEE Conf. on Mobile Computing and Networking*, pages 59–68, 1999.

[10] X. Jiang, jason I. Hong, and J. A. Landay. Approximate information flows: Socially-based modeling of privacy in ubiquitous world. In *Proceedings of The UbiComp 2002*, 2002.

[11] X. Jiang and J. A. Landay. Modeling privacy control in context-aware systems. *IEEE Pervasive Computing*, 1(3):59–63, July 2002.

[12] A. Myers and B.Liskov. Complete, safe information flow with decentralized labels. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*, 1998.

[13] C. Policroniades, R. Chakravorty, and P. Vidales. A data repository for fine-grained adaptation in heterogeneous environments. In *Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access*, pages 51–55. ACM Press, 2003.

[14] T. Richardson, Q. Stafford-Fraser, K. R. Wood, and A. Hopper. Virtual network computing. *IEEE Internet Computing*, 2(1), January 1998.

[15] A. Rodr?guez and M. Egenhofer. A comparison of inferences about containers and surfaces in small-scale and large-scale spaces. *Journal of Visual Languages and Computing*, 11(6):639–662, 2000.

[16] D. Scott, A. Beresford, and A. Mycroft. Spatial policies for sentient mobile applications. In *Proceedings of the FASE 2003 (ETAPS 2003)*, 2003.

[17] D. Scott, A. Beresford, and A. Mycroft. Spatial security policies for mobile agents in a sentient computing environment. In *Proceedings of The FASE 2003*, 2003.

[18] F. Stajano. *Security for Ubiquitous Computing*. Wiley, 2002.

[19] F. Stajano and J. Crowcroft. *The Butt of the Iceberg: Hidden Security Problem of Ubiquitous System*. Kluwer, 2003.

[20] R. Want, T. Pering, G. Danneels, M. Kummar, M. Sundar, and J. Light. The personal server - changing the way we think about ubiquitous computing. In *Proceedings of the UbiComp 2002*, LNCS, pages 194 – 209, October 2002.

[21] H. F. Wedde and M. Lischka. Role-based access control in ambient and remote space. In *Proceedings of the SACMAT 2004*, 2004.

[22] M. Weiser. The computer for the 21st century. *Human-computer interaction: toward the year 2000*, pages 933–940, 1995.