The Freiburg Privacy Diamond: An Attacker Model for a Mobile Computing Environment

Alf Zugenmaier, Michael Kreutzer, Günter Müller
Institute for Computer Science and Social Studies
University of Freiburg
Friedrichstr. 50
D-79098 Freiburg
{zugenmaier|kreutzer|mueller}@iig.uni-freiburg.de

Abstract

An attacker model is a pre-requisite to make statements on the properties of security mechanisms. However, in order to be adequate, an attacker model should fit to the class of the considered security mechanisms and to the (computing) environment. This paper introduces a new attacker model, called the Freiburg Privacy Diamond (FPD), to evaluate anonymity mechanisms with regard to mobility, especially device mobility. This model takes into account the conditions of mobile computing. The FPD enables the analyser to adjust the strength of the attacker in a fine grained way and dependent on the (computing) environment.

1 Motivation and Overview

In a mobile computing environment information is processed on devices in constantly changing networks, where new devices join, others leave, thus changing the environment and the available services. A device may join a network automatically if the network is reachable through its radio interface. The operators and users of the other devices in this network may not be interested in protecting the security, especially protecting the privacy, of the user of this additional device.

Many mechanisms to achieve anonymity have been proposed, most of them requiring access to a sophisticated anonymizing infrastructure. In the case of mobile computing implementation of these mechanisms may lead to high computional overhead and network load. Evaluation of these mechanisms is mostly concerned with showing that "perfect" anonymity can be achieved, and analyzing resistance against attacks.

Using a analysis method that includes mobility, it is possible to derive anonymizing techniques that are better adapted to the needs of a mobile environment. In this paper such an analysis method is proposed.

This paper is structured as follows: The next section defines anonymity and refers how

Part of this work was supported by the Kolleg "Living in a Smart Environment" of the Gottlieb Daimler- and Carl Benz-Stiftung

other authors measure anonymity. Section three describes two approaches to model an attacker.The privacy diamond is introduced and formalized in section four, as well as applied to some examples. The paper concludes with remarks on the adequacy of the privacy diamond as a model.

2 Anonymity

The Oxford English Dictionary defines „anonymous“ as nameless, having no name. „Anonymity“ is the state of being anonymous, used of an author or his writings. A technical definiton is given in [PfKo01]: „Anonymity is the state of not being identifiable within a set of objects, the anonymity set.“ This definiton differs from the colloquial use of the word in that an environment, the anonymity set, is necessary for this type of anonymity. In [PfKo01], also a vague notion of the strength of the anonymity has been given, i.e. anonymity becomes stronger as the anonymity sets increases in size and the actions performed by the objects in the anonymity set are distributed more evenly.
In [ReRu98] Reiter and Rubin use probabilities to evaluate the anonymizing system „crowds“. They define degrees of anonymity by looking at the probability that a certain user requested a web page. The user is beyond suspicion if the probability that this user requested that web page is identical to the probability that any other user performed the action. Similarly probable innocence is defined as the probability of the user requestinng the web page is less than 1/2. For possible innocence this probability is greater than 1/2, but still less than one by a „non-trivial amount“.
These approaches for measuring the degree of anonymity lack the possibility of taking previous knowledge into account. E.g. if only one person of the anonymity set is known to be interested in a new job, the fact that one person of the anonymity set did look for a new job strongly supports that this person looked for the new job.

3 Modeling an Attacker

The attacker is usually characterized by pieces of information, such as protocol data units, that are known to him. The attacker’s capability to retrieve this information depends on the position of the attacker within the network, as well as the layer within the network protocol stack at which the attacker parses the protocol data units.
The best known attacker model is the omnipresent attacker model [DoYa83]. The omnipresent attacker is able to intercept all communications, be the legitimate recipient of a communication and initiate communications with any other participant. But the attacker is not able to break any of the cryptography used. Often this model is extended, the attacker may collaborate with some or all of the parties involved in the communication: the network provider, the anonymizing system provider and/or the service provider.
Another model is the partially present attacker who can intercept communications only at a limited number of points in the network. He may operate a part of the anonymizing system. Like the omnipresent attacker, the partially present attacker is not able break cryptography.
Example: Position of attackers inside a mobile computing setting (cf. Figure 1)
The danger of eavesdroppers on the link between a mobile device and the access point

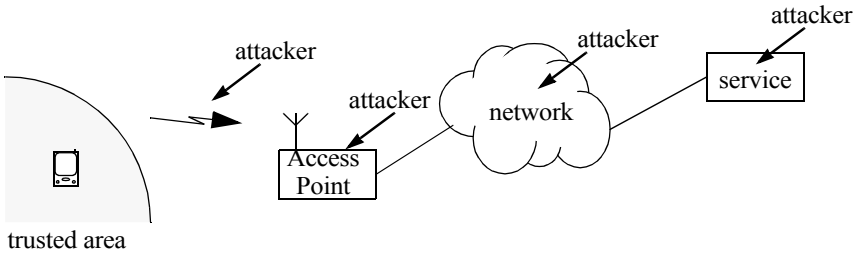

FIGURE 1. Possible locations for the eavesdropper(s) (arrows) of the attacker model inside a mobile computing setting.

or base station is far greater in a mobile wireless environment than in a fixed network environment, as access to the transmission medium cannot be controlled. The access point itself may already be considered an adversary, as the number of access providers increases and become more complex. The attacker could also be listening to the traffic in the network. Finally the service itself may not be trusted to handle personally identifying information.

Current attacker models only describe possible sources of information for the attacker. They do not describe the inference rules the attacker may use to deduce the identity of the user from the information he gathers. That is done in the privacy diamond, described below.

4 The Freiburg Privacy Diamond (FPD)

4.1 The Deduction Principle

The privacy diamond, which is introduced in this paper, is a model making explicit assumptions about deduction rules of an attacker. How and where the attacker gathers this information was detailed in the previous section and is not within the scope of this part of the attacker model.

The privacy diamond in Figure 2 represents relations between the action, the user, the location of the user and the operating device of the user. With this completely connected directed graph it is possible to determine information that can be deduced from other information. The use of the privacy diamond is illustrated in a very simplified fashion by the following examples: The information where a user is located can be deduced, if the location of the device of the user is known. If the device used for the transaction is concealed, e.g. using a mix network, this deduction is not possible. But it may also occur that the device and the location of the device are known, if the user goes to an Internet-Cafe. However, there is no a priori knowledge of the user of the device. This knowledge can only be gained, if the user reveals her or his identity directly.

4.2 Model Assumptions

The privacy diamond is used to represent the knowledge of the attacker in the following situation: A user operates a device at a certain location to initiate an action. Four entities are necessary to model the situation: the user, the action, the location, and the device. These will be addressed below. Time will only be considered as an implicit parameter.

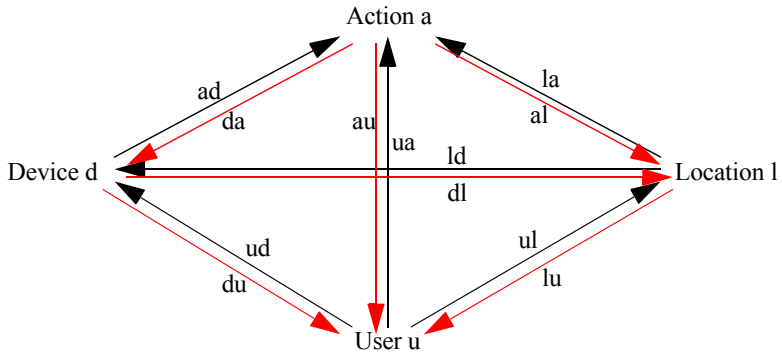

FIGURE 2. The privacy diamond showing relations and possible inferences between the user, the action, the device and the location of the user and the device.

Mobile users use a device to perform actions. These actions are considered to be atomic; during an action neither the user, the device of the user nor the location of the user changes. The action is also instantaneous, it is carried out while the user uses the device.
To model location information, the world is divided into cells. The size of these cells determines the maximum resolution to which a device or user can be located. This simplification helps to avoid mathematical difficulties with continuous variables.
Users perform actions using a single device from a set of devices. The device is located at the same position of the user. This assumption is realistic as the user has to be in the proximity of the device to operate it.
The scope of the definition of the device, abbreviated d, includes all software on this device. If the software is able to migrate from device to device, like mobile agents, this node of the graph would have to be replaced by several nodes. This situation is not considered in this paper.

4.3 Formalizing the FPD

To formalize the privacy diamond, the sets involved have to be defined, as well as possible relations, and semantics for these relations.
The following disjunct finite sets define the basic sets of the privacy diamond:

$$a \in A \subseteq \text{Set of all possible actions}, \quad d \in D \subseteq \text{Set of all devices},$$
$$l \in L \subseteq \text{Set of all spatial cells}, \quad u \in U \subseteq \text{Set of all users}$$

The assumption that set A is finite is no restriction as the privacy diamond cannot produce knowledge about unknown actions. Generally, it can be assumed that the number of known actions is finite.

Knowledge is then expressed in the form of relations:

da which device d or set of devices could have been used
ad which action could have been initiated from the device
la from where (location l) could the action a have been initiated
al which action could have been initiated from location l
ud which user could have operated the device
du which device could have been operated by the user

ua which user performed action a
au which action could have been performed by user u
ul which user could have been at the location
lu where could the user have been
ld where could the device have been
dl which device could have been at the location

The arrows (c.f. Figure 3) are interpreted with a calculus based on relations, which can also be weighted by probabilities.

4.4 Interpretation using Relations

Relations can be used if the information gathered by the attacker can not be weighted. The confidence in all pieces of information is the same.

The privacy diamond is formalized by 12 relations,

$R_{DA} \subset D \times A, R_{AD} \subset A \times D, ..., R_{LD} \subset L \times D, R_{DL} \subset D \times L$.

The statement for the relation R_{DA} on arrow 1 means: Which device d initiated action a? If (d,a) is contained in the relation R_{DA}, then the interpretation is that device d could have initiated action a. This relation does not exclude the possibility that another device d' could have initiated this action. If d_1 or d_2 could have initiated action a, the relation R_{DA} would be $R_{DA} = \{(d_1, a), (d_2, a)\}$.

Making a distinction between R_{XY} and R_{YX}, where $X, Y \in \{D, U, L, A\}$ and $X \neq Y$, may seem artificial at first. An example helps to clarify this distinction: Considering $R_{UA} = \{(u, a_1), (u, a_2)\}$ and $R_{AU} = \{(a_1, u), (a_2, u)\}$, R_{UA} is interpreted that actions a_1 and a_2 occurred, and user u is the sole suspect for having done these. Conversely, R_{AU} leads to the interpretation, that user u is known, and he could have done either one or both of the actions a_1 and a_2, or none of them.

Combining the knowledge represented by two relations, $R_{XY} \subset X \times Y$ and $R_{YZ} \subset Y \times Z$, where $X, Y, Z \in \{D, U, L, A\}$ and $X \neq Z$, new pairs of relation $R_{XZ} \subset X \times Z$ can be deduced following the „transitivity" rule:

$$R_{XZ,\,new} = R_{XZ} \cup \{(x, z) | \exists y \in Y : (x, y) \in R_{XY} \wedge (y, z) \in R_{YZ}\} \qquad \text{(Eq 1.1)}$$

For example, if we know that device d initiated action a and that user u used device d at the same time, i.e. $R_{DA} = \{(d, a)\}$ and $R_{UD} = \{(u, d)\}$, we can deduce $R_{UA} = \{(u, a)\}$, that is, user u initiated action a.

$View_O = R_{DA} \cup R_{AD} \cup \ldots \cup R_{DL} \cup R_{LD}$ is the initial set of relations that are known to an observer O. The closure of $View_O$ under the transitivity rule will be denoted by $\overline{View_O}$.

The following two examples illustrate the attacker model and show how anonymizing systems can be analyzed.

Example 1: This example shows how the privacy diamond can be applied to an anonymizing system like "crowds" [ReRu98], with n computers participating in the crowd. Because crowds treats all computers as stationary, crowds hides the location of these computers as well.

To model the attacker with the FPD, we first define the four sets $A = \{a_1, a_2, \ldots, a_N\}$, $U = \{u_1, u_2, \ldots, u_M\}$, $L = \{l_1, l_2, \ldots, l_L\}$, and $D = \{d_1, d_2, \ldots, d_K\}$, with $K, L, M, N \in \mathbb{N}$, representing the actions, the users,

the locations and the devices the attacker knows of. The computers $d_1, \ldots, d_n$, $n < K$ participate in the crowds system, therefore a specific request can not be related to a single computer, instead it must be related to all participating computers. We assume the observer knows which computers participate in the crowds system. The knowledge of the observer O ($View_O$) logging the requests of the web server for page 1(a_1), coming from the crowds system, and page 2 (a_2), coming from a computer (d_m, $m < N$) without anonymizing techniques, is given by:

$$R_{DA} = \{(d_1, a_1), (d_2, a_1), \ldots, (d_n, a_1), (d_m, a_2)\}$$

Additionally, the observer may know which user uses a certain computer, giving rise to:

$$R_{UD} = \{(u_1, d_1), (u_2, d_2), \ldots, (u_n, d_n), (u_m, d_m)\}$$

As all devices are located at well known fixed locations, the relations involving the

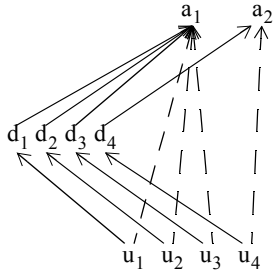

FIGURE 3. Relations of the initial knowledge of the observer in Example 1 with n=3 and m=4: $View_O$ **(solid arrows) and the inferred relations (dashed arrows)**

location do not offer additional insight and are omitted in this example.
These two can be combined to give knowledge about R_{UA} by applying Equation (Eq 1.1):

$$R_{UA} = \{(u_1, a_1), (u_2, a_1), \ldots, (u_n, a_1), (u_m, a_m)\}$$

The observer still does not know which of the users $u_1, \ldots, u_n$ requested page 1 (a_1).
Example 2: The Internet Cafe is a simple method to achieve anonymity on the Internet. A user can go to the Internet Cafe (location l) and use the freely accessible computer (device d) to access the Internet and request a web page (action a). $View_O$ consists of the relations $R_{DL} = \{(d, l)\}$, $R_{LD} = \{(l, d)\}$ and $R_{DA} = \{(d, a)\}$. Unless users provide their names, there is no knowledge about the identity of possible users. Therefore $R_{XU} = R_{UX} = \varnothing$ for X = A,D or L. Thus $\overline{View_O} = View_O \cup \{(l, a)\}$. No deductions about the user who requested the web page is possible for the observer evaluating the server's log files.

4.5 The Intersection Attack

The intersection attack [BeFK00] can be modeled with the notation of the privacy diamond. If an attacker knows that a certain relation in a privacy diamond corresponds to a relation in another privacy diamond, the attacker is able to link the two privacy diamonds. It could be that the attacker knows that two actions a_1 and a_2 were carried out using the same device. If the attacker now assumes that all possible device to action re-

lations are contained in both privacy diamond's relations R_{DA} and R'_{DA}, these relations can be intersected to give rise to the relation $R''_{DA} = \{(d, a_{12}) | (d, a_1) \in R_{DA} \wedge (d, a_2) \in R'_{DA}\}$ with a_{12} the action combining a_1 and a_2. The intersection attack is carried out with the relations in $\overline{View_O}$ where R_{DA} is joined with R''_{DA}. Then the transitive closure is applied again. Likewise the relation R'_{DA} in $\overline{View'_O}$ is joined with R''_{DA} and the transitive closure calculated. The intersection attack adds information to both privacy diamonds.

Example: Similar to example 3 of the previous section, several computers use a crowds network. At one time computers d_1, d_2, and d_3 participate in the crowds network, and page 1 is requested (a_1). Later computers d_3, d_4, and d_5 are part of the crowds network and page 2 is requested (a_2). Each computer is used by one user only, d_1 by u_1, d_2 by u_2, and so on. The relations the attacker knows $View_O$ are: $R_{DA} = \{(d_1, a_1), (d_2, a_1), (d_3, a_1)\}$, $R_{UD} = \{(u_1, d_1), (u_2, d_2), (u_3, d_3)\}$ leading to
$\overline{View_O} = R_{DA} \cup R_{UD} \cup \{(u_1, a_1), (u_2, a_1), (u_3, a_1)\}$.
Similarly,
$\overline{View'_O} = \{(d_3, a_2), (d_4, a_2), (d_5, a_2), (u_1, d_1), (u_2, d_2), (u_3, d_3), (u_4, a_2), (u_5, a_2)\}$
Because a cookie is set by page 1 and retrieved by page 2, it is possible to link a_1 and a_2.
Applying the rules for the intersection attack we have the privacy diamonds $\overline{View_O} \cup \{(d_3, a_{12}), (u_3, a_{12})\}$ and $\overline{View_O'} \cup \{(d_3, a_{12}), (u_3, a_{12})\}$, both stating that the sole suspect for having done both actions (a_{12}) is u_3.

4.6 A Probabilistic Attacker Model

The interpretation of the privacy diamond given in the previous section has one major drawback: It is not possible to grade information, i.e., how certain is a specific information. It could be that a device is used by several users, but is most likely used by one primary user. A measure of confidence in a piece of information is desirable, i.e. in an element of a relation. Within the context of trust relationships this has been done in [Maur96]. In this section the same measure of confidence is applied to the privacy diamond.
The use of the "random-worlds" method of [Nils86] is best clarified by an example. Suppose we have the statement: "User u initiated action a." We can imagine two sets of possible worlds, one containing the worlds where the statement holds, and another one, containing the worlds where the statement is false. To model the uncertainty about the reality, we ascribe a probability p_1 to the reality being in the first set and a probability $p_2 = 1 - p_1$ to the reality belonging to the second set. We can say that the probability of the statement, i.e. the probability that the statement is true, is p_1.
If more than one statement is considered, more sets of possible worlds are required. In general, one set of possible worlds is necessary for every subset of statements that are true in this world.
To model the privacy diamond with this approach, the deterministic view of the observer is replaced by possible worlds, each with a view, and a probability distribu-

tion over these worlds. In every world the "transitivity" rule is applicable in the same way as in the deterministic model.
The initial view of the observer has to be replaced by a probabilistic initial view. If R_O denotes the set of all elements of the relations considered by the observer, the probabilistic initial view is a probability function $P : \wp(R_O) \to [0, 1]$ over sets of possible relations of the FPD. This probability function implies a probability measure for the set of possible relation sets. $View_O : \Omega \to \wp(R_O)$ now is a random variable, as is $\overline{View_O} : \Omega \to \wp(\overline{R_O})$. Following the notation defined in [Maurer1996], for all $\nu \in \wp(R_O)$, $P(\nu) \equiv P(View_O = \nu)$ gives the probability of the elementary event ν. The abbreviation $\nu \subseteq View_O$ is used to denote $\{\mu \subseteq R_O : \nu \subseteq \mu\}$, the set of subsets of R_O that contain the elementary event ν.
The confidence value for a pair of relations, $r \in \overline{R_O}$, is the probability that it is contained in the closure of R_O:

$$conf(r) = P(r \in \overline{View_O}) = \sum_{\nu \subseteq R_O \,:\, r \in \bar{\nu}} P(\nu) \qquad \text{(Def 1.2)}$$

The probability is the sum of the probabilities of all elementary events ν of the initial view of the observer, from which the pair r can be deduced using the transitivity rule.
Using definition (Def 1.2) all confidence values can be calculated. It is only necessary to specify the initial view, that is the probability function, for all ν. However, this means specifying $P(\nu)$ at $2^{|R_O|}$ points. One way would be to assign values for a small set of ν where $P(\nu)$ can be determined and calculate the limits of possible confidence values for the remaining ν [Nils86]. Instead of giving the maximum and minimum values for the set of ν where $P(\nu)$ is not known in advance, those values of $P(\nu)$ can be provided that add the least information. This maximum entropy method relies on the fact, that the greatest number of possible worlds is placed near the state of maximum entropy. In our approach "independent initial confidence parameters" $p(r)$ are introduced, similar to the way described in [Maur96] for authentication networks. The initial confidence parameters assigned to the pairs of the relations in R_O are assumed to be independent. If no $r \in R_O$ is contained in $\overline{R_O - \{r\}}$, R_O is called minimal and $conf(r) = p(r)$. For given initial confidence parameters $p : R_O \to [0, 1]$, the probability measure can be calculated:

$$P(\nu \subseteq View_O) = \prod_{r \in \nu} p(r) \qquad \text{(Eq 1.3)}$$

The confidence value for an elementary event is:

$$P(View_O = \nu) = \prod_{r \in \nu} p(r) \prod_{r \notin \nu} (1 - p(r)) \qquad \text{(Eq 1.4)}$$

Using definition (Def 1.2) and Equation (Eq 1.4) the confidence value $conf(r)$ for a pair $r \in \overline{R_O}$ is available. [Maur96] also describes a more straightforward method of calculating $conf(r)$. Determination of the minimal sets ν_i where $r \in \overline{\nu_i}$,

$i = 1, \ldots, k$, and calculation of the probability of the union of the events $v_i \subseteq View_O$ for $i = 1, \ldots, k$:

$$conf(r) = P\left(\bigcup_{i=1}^{k} (v_i \subseteq View_O)\right) \quad \text{(Eq 1.5)}$$

The probability of a union of k events can be calculated by summing up their probabilities, then subtracting the pair wise intersections, adding the intersections per 3 events, etc.

$$conf(r) = \sum_{i=1}^{k} P(v_i \subseteq View_O) - \sum_{1 \le i_1 < i_2 \le k} P((v_{i_1} \cup v_{i_2}) \subseteq View_O) + \ldots \quad \text{(Eq 1.6)}$$

The numbers for the initial confidence values may be obtained using statistical information or estimated values by guessing. In the case that several pairs are contained in one relation without additional information, one assumes identical initial confidence values adding up to one.

The following two examples illustrate the application of the probabilistic model.
Example 1: A web server's log shows that a certain page has been requested (action a) from a certain IP address. This IP address is usually assigned to a certain computer (device d). Thus we have a confidence parameter $p((d, a) \in R_{DA})$ for the relation (d,a). The computer is mostly used by user u, implying another confidence parameter $p((u, d) \in R_{UD})$ for the relation (u,d). The confidence parameter for the resulting relation (u,a) is the product of the initial confidence parameters $conf((u, a) \in R_{UA}) = p((u, d) \in R_{UD})p((d, a) \in R_{DA})$. The initial confidence parameters are annotated to the edges of the privacy diamond in Figure 4. The numbers

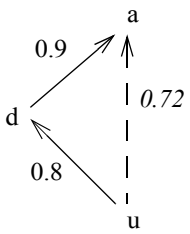

FIGURE 4. A privacy diamond with initial confidence parameters and the resulting confidence in the relation (u,a) (set in italics)

that are set in normal type face are the initial confidence parameters, the number in italics is the calculated confidence value.
Example 2: The crowds system that already served as an example in section 4.4 as interpreted using probabilities. The probability that a certain request for a page (action a_1) actually originates from one of the computers (devices $d_1, \ldots, d_n$) is $1/n$. For this example $n = 4$ is assumed. The computer d_1 is used only by user u_1, while the three other computers are only used with a probability of 80% by users u_2, u_3, and u_4 respectively. The confidence values are calculated as in the previous example by multiplying the appropriate initial confidence parameters:

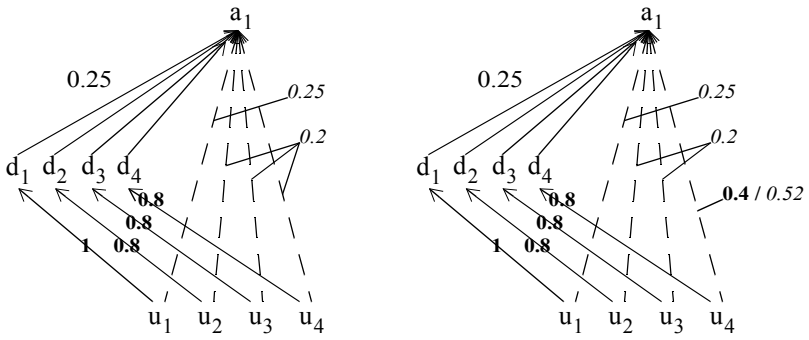

FIGURE 5. Relations of the initial knowledge of the observer in Example 2: View_O (solid arrows) with initial confidence parameters in bold and the inferred relations (dashed arrows) with calculated confidence values in italics. In the right graph the attackers context information regarding the surfing behaviour is included.

$\mathrm{conf}((u_i, a_1) \in R_{UA}) = p((u_i, d_i) \in R_{UD})p((d_i, a_1) \in R_{DA})$, $i = 1 \ldots 4$.
Therefore the highest confidence is in u_1 initiating the request for the web page (c.f. Figure 5).

If additional information about the surfing behavior of the individual users is known, this can be specified by giving additional confidence parameters. Assume that a confidence parameter of 0.4 is already assigned to user u_4 initiating the action. $\mathrm{conf}(((u_4, a_1)) \in R_{UA})$ can be determined with Equation (Eq 1.6) to be:
$\mathrm{conf}((u_4, a_1) \in R_{UA}) = p((u_4, d_4) \in R_{UD})p((d_4, a_1) \in R_{DA})$
$+ p((u_4, a_1) \in R_{UA}) - p((u_4, d_4) \in R_{UD})p((d_4, a_1) \in R_{DA})p((u_4, a_1) \in R_{UA})$

The fact that someone using the crowds system initiated action a_1 and the fact that u_4 could have used the crowds system increases the initial confidence in the assumption that u_4 requested page a_1. The initial confidence parameter, which was still in the range of "probable innocence" as defined by [ReRu98], combined with the knowledge that someone using the crowds system did in fact perform action a_1 edges the confidence that u_4 did it into the range of "possible innocence", which may also be seen as possible guilt (c.f. Figure 5).
This example shows that $\mathrm{conf}(r)$ is not normalized, i.e. $\sum_r \mathrm{conf}(r) \neq 1$ in the general case. This fact is consistent with the interpretation of section 4.3. Only the confidence in statements whether it is possible that the user initiated the action is expressed. The model does not require that all users are specified in advance, it also does not exclude collaborative actions, that is more than one user initiated an action.

4.7 Introducing Time

The most straightforward way of introducing time into the privacy diamond is to have time dependent initial confidence parameters. This allows inclusion of knowledge about dependencies in the statistical distribution of events upon time. It may be known that one user is more likely to be active than the other at one time, while it can be different at a later time.

5 Conclusions

The privacy diamond models the information processing of the attacker. It can be used to analyze anonymity mechanisms and leads to a measure of the degree of anonymity, to the confidence that a specific user performed a certain action. The privacy diamond can include additional knowledge about the distribution of actions, and can thus be used to model an imperfect anonymity mechanism, i.e. a non-uniform distribution of confidence values for all possible users performing that action. Determination of initial confidence parameters still poses a problem: Unless a uniform distribution is assumed, only relative statements are possible, i.e. the confidence in this statement is twice as high as the confidence in another one.

For a real attacker, using this model would have a major drawback: because no exclusion of possibilities is possible in the model, the information sought after can get lost in the mass of inferences.

As future work, this attacker model will be evaluated whether it can be used to construct new anonymity mechanisms.

References

[BeFK01] Oliver Berthold, Hannes Federrath, and Stefan Köpsell: Web MIXes: A System for Anonymous and Unobservable Internet Access. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies*, LNCS 2009, pp 115-129, 2001

[DoYa83] Danny Dolev, Andrew Yao: On the Security of Public Key Protocols; in: *IEEE Transactions on Information Theory*, Vol. IT-29, 2, pp 198-208, 1983

[Maur96] Ueli Maurer: Modelling a Public-Key Infrastructure, in *Proceedings of 1996 European Symposium on Research in Computer Security (ESORICS '96)*, E. Bertino (Ed.), Lecture Notes in Computer Science, Berlin: Springer-Verlag, 1996

[Nils86] Nils Nilsson: Probabilistic Logic, in *Artificial Intelligence*, 28, pp 71-87, 1986

[PfKo01] Andreas Pfitzmann, Marit Köhntopp: Anonymity, Unobservability, and Pseudonymity - A Proposal for Terminology; in: Hannes Federrath (Ed.): *Designing Privacy Enhancing Technologies; Proc. Workshop on Design Issues in Anonymity and Unobservability*; LNCS 2009; 2001; pp. 1-9

[ReRu98] Michael Reiter, Aviel Rubin: Crowds: Anonymity for Web Transactions; *ACM Transactions on Information and Systems Security* Vol. 1, No. 1, November 1998, pp. 66-92