Some Methods for Privacy in RFID Communication

K. Fishin, S. Roy, B. Jiang

IRS-TR-04-010

June 2004

Intel Research

Intel Research Seattle Tech Memo IRS-TR-04-010. ESAS 2004 – to appear

Some Methods for Privacy in RFID Communication

Kenneth P. Fishkin[1], Sumit Roy[2], and Bing Jiang[1, 2]

[1]Intel Research Seattle, 1100 NE 45th St,
Seattle, Washington, 98105 USA
Kenneth.p.fishkin@intel.com
[2]Department of Electrical Engineering, University of Washington,
Seattle, Washington, 98195 USA
{roy, bjiang}@ee.washington.edu

Abstract. For RFID tags to gain general acceptance, they will have to offer powerful and flexible privacy mechanisms. After reviewing existing and upcoming privacy mechanisms for RFID privacy, we propose that a key aspect of RFID communication with passive tags, namely its required energy transference from an external antenna, may offer promise when developing privacy mechanisms. We present two proposals for such mechanisms. In the first mechanism, analysis of the received signal by the tags can be used to estimate reader distance (and hence trust). We show that a simple metric analogous to signal to noise ratio correlates well with rough distance. In the second, antenna energy is used to power a tiered authentication scheme, in which tags reveal more information about themselves to more trusted and/or "energetic" readers.

1 Introduction

Radio Frequency Identification (RFID) technology is a technology which allows very cheap wireless tags to communicate identification to an interrogating reader located some distance. Conceptually, this is similar to a bar-code, but the wireless nature of the communication allows for significant qualitative and quantitative advances: the reader need not have line-of-sight to the tag, the tag can store and communicate many more bits of information, multiple tags can be interrogated by the same reader, the reader can be located to read passing tags automatically without explicit user action, and so forth. While the basics of this technology have existed for decades (with their traditional application domain being that of tagging and tracking livestock), recent advances in the capabilities of the tags (operating range, amount of memory, etc.,) coupled with drastic decreases in the cost of both readers and tags (at this writing, tags are available in the €0.30 range (and dropping rapidly), compared to €3.00 a few years ago) has changed the landscape of deployment. The ability to do much more at much less cost has combined to transform RFID from a "fringe" technology into something which is poised to be deployed in the billions per year [1]. However, with this explosion of deployment comes an explosion of responsibility – the responsibility to ensure that the data on the tag is only read by desired readers in desired ways. When RFID was a fringe technology, this was not a major issue. However, this issue

must now be addressed, or the RFID explosion may not occur, or may occur in a much more limited fashion.

The basic RFID scenario is as follows. An RFID reader emits a waveform of some energy at some frequency. That energy is then caught by an antenna on a distant, passive tag. The energy is used to both energize the computation on the tag and energize the returning waveform reply from the tag containing e.g. the tags globally unique ID. Slightly more sophisticated operations are possible, e.g. requests to read/write flash memory on the tag. Note that the reader signal has no disambiguation or identification associated with it.

The historic focus on livestock tracking as the main RFID app has had some positive influences, such as the existence of very physically robust tags (there are tags which can survive autoclaving, which can be pounded into tree trunks, etc.) which can be read without any special user action, it has had some unfortunate influences as well. In particular, there has historically been no emphasis on, or consideration of, privacy or security in the RFID communication: neither livestock nor its owners were concerned with this issue. Accordingly, in most existing RFID protocols, a tag will provide all information unquestioningly to any reader. The second main application domain which has emerged, that of supply chain management (tracking goods from the factory through the distribution center and on to the destination store), has also had only minimal privacy focus: in this realm, RFID tags are viewed as better barcodes, and require no more privacy or security than those do. The original RFID protocols for supply chain management, accordingly, had the same open access protocol as had been used by livestock, i.e. none. However, in introduction to protests from privacy advocates (cf [2]), the proposed protocol has been modified:

2 The Kill Switch

While the standard is still evolving, the predominant proposed privacy mechanism is the so-called "Kill Switch". While the details, again, are still evolving, the basic concept has remained the same (we base our discussion on the most recently published publicly available specification [3]). Each tag has a password (how many bits, and whether the password can be modified, is in flux). When a tag receives a "KILL" command from a reader, accompanied by the appropriate password, the tag essentially "Kills itself" – it sets an internal bit permanently, and so long as that bit has been set, the tag no longer responds to any interrogations from any readers. Conceptually, the tag has de-activated itself. The typical envisioned scenario is that a tag responds unquestioningly to all queries until the good it is attached to is purchased, at which point it shuts down. The idea is that the needs are met of both industry (which is largely focused on tracking the good up to the point of purchase) and privacy advocates (which is largely focused on post-purchase privacy), while requiring only very minimal changes to tag hardware and communication protocols.

The proposal is simple and effective in this domain, but has some weaknesses:

1. It is an "all or nothing" privacy mechanism – the tag responds to everyone until the kill switch is set, and then responds to no-one. There is no way to have finer-grained disclosure, e.g. to disclose the expiration date on a pill bottle to a reader in a person's medicine cabinet, but not to anyone else.
2. The user has no way to know whether the tag has actually received the KILL command, let alone that the command was interpreted successfully.
3. It appears that the tag will reveal its password to anyone who asks. Therefore it is very easy for malicious readers to KILL tags prematurely.

A very recent, as of yet unpublished, proposed extension adds a "CONCEAL" command. As long as the CONCEAL bit is set (and unlike KILL, this is a reversible state), the tag will still respond to all queries, but with a random ID. The idea is that this allows a count metric (how many tagged items are in the truck?) without revealing detailed information as to the nature of the items. While this is a positive step away from the "all or nothing" paradigm, it still has the same underlying weaknesses of the KILL command.

3 The Blocker Tag

To address some of the weaknesses and simplicities of the "Kill Switch", a second recent proposal that has received much attention is the "Blocker Tag" proposal of Juels et al. [4].

The Blocker tag exploits a characteristic of RFID communication: the tree-walking protocol a reader uses to determine which tags it sees. When a reader sends a signal out into space, looking for tags, it asks the question: "are there any tags out there whose ID starts with <B>"? (<B> is some bit string). There are three possible answers the reader can receive:

1. No answer.
2. Exactly one answer: the full ID of a sensed tag.
3. More than one answer. In this case, the reader gets a "jumble" of potentially overlapping signal responses – it cannot parse all the responses at one time. Instead, therefore, it recurses, asking for tags whose ID begins with <B>0, and then for tags whose ID begins with <B>1. (Technically, for speed purposes, this can be optimized, for example some current readers immediately issue 8 queries, recursing to depth 3 at one step, but the principle is the same).

The blocker tag responds by *always responding* to the reader query – essentially, it ignores <B>. In this way, it can serve to passively jam the reader, forcing the reader to fruitlessly chase down very long trees.

The blocker tag can act either reflexively or transitively. By a "reflexive" tag, we mean a blocker tag which prevents itself from being read. In this case, the blocker tag is conceptually equal to the "Kill switch", although it is implemented in a very different manner. By a "transitive" tag, we mean that the main purpose of the blocker tag can be to prevent the reader from reading *other* tags nearby. For example, a privacy-conscious consumer might scatter blocker tags about their house or person, to ensure that any other RFID tags in the neighborhood are jammed via their proximity to the blocker tag.

A blocker tag can choose when to act in this hostile manner. For example, there might be "danger zones", so that a reader will only be jammed if it enters a certain "forbidden" area of the tree. This is normally conceived of in a breadth-first way : you can read everything from the tree rooted at <X>, but nothing from the tree rooted at <Y>. For example, the envisioned check-out scenario is that the tag receives a command asking it to switch its in from <X> to <Y> - this functionally acheieves the similar effect of the KILL switch. Note that this could also be implemented in a depth-first manner. A blocker tag might for example allow its first 14 bits to be queried, yielding information roughly equivalent to a bar code, but then begin jamming if a reader pries beyond that point.

The Blocker tag can therefore be viewed as an extension of the KILL switch, with much more flexibility as to whether it responds, and the ability to "kill by association" nearby tags. It does have some problematic aspects, however:

1. While the blocker tag can control whether it responds, again its only response is to respond with either everything or nothing.
2. The decision as to whether to respond is invariant with respect to the particular reader: all readers who request the same information will retrieve the same result.
3. With time, the jamming can be overcome. Suppose the tree-walking mechanism is in some area of the tree where only the blocker tag exists. Even if the response has been more circumspect such that it only says "I exist", and does not return its ID, the tree-walker can determine that it is in the presence of a blocker tag whenever only one tag responds to bit-string <B>, and yet it receives exactly one response to both <B>0 and <B>1. Therefore, when walking a deep tree, many parts of the tree will, in practice, be quickly pruned significantly.

Our first proposal is similar to the blocker tag, in that we again return to an examination of the idiosyncratic nature of the RFID communication protocol, to see if there are privacy-enhancing techniques available within the RFID environment that are not available in more standard security domains. We combine that concept with another existing concept, that of location-sensitive verification (see e.g. [5]), where the tag seeks to distinguish between different requests for the same information by using knowledge about the spatial location/orientation of the requestor.

We wish to maintain the desirable properties of the KILL switch and blocker tag: a solution which can work with little or no change to tag or reader hardware, little or no change in the communications protocol, and no battery or energy source in the tag. This drastically reduces the space of appropriate techniques. We propose, however, that at least one such technique may be of use which fits those constraints, one which ties the level of information revealed to the distance between the tag and the reader.

4 Distance Implies Distrust

Assume a scenario in which some hostile RFID reader wishes to interrogate (or even change) the information on an RFID tag. Note that often (though not always) such scenarios involve a reader which is physically distant from the receiving tag. This is because the closer the reader is, the more subject it is to scrutiny by the wearers, owners, and/or users of the tagged object. If the tag is located inside a house, for example, it is far more likely that an attack will come from outside the house, than from an intruder inside the house. If the tag is located on a person, it is much more likely that an unwanted, unwarranted request will come from far enough away that the recipient does not see the reader. We therefore propose that RFID privacy mechanisms can and should use the physical distance between the information requestor and the information owner as part of their algorithms.

However, how is a tag to infer the distance between itself and an RFID reader? Existing algorithms such as the ECHO protocol mentioned earlier [5], which use RF and ultrasound transmission to echo-locate, cannot be employed here, due to the extremely difficult RF-only environment. The cost and power requirements of ultrasound are prohibitive in this domain, and pure RF echo measurements (i.e. using the radio both ways) would incur several challenging problems; namely the presence of multipath that could be effectively exploited by attackers and thus necessitate significant advances over the current ECHO protocol. The question is, then, whether we can do rough distance inference given only the RF signal coming from the reader. Several approaches are possible:

4.1 Triangulation via time-of-arrival analysis

One common technique for distance inference is triangulation: we could imagine solutions in which a set of tags compare their received signals, perform cross-correlation, infer time-of-difference in time-of-arrival, and from that infer a fairly accurate distance measurement. However, this solution violates our requirement for minimal cost and infrastructure, and even then, the variations in the RF field might well make the calculation very unreliable.

4.2 Triangulation via signal strength analysis

Another common technique, which requires no inter-tag communication, is to look at the amount of energy received by the tag. By comparing the amount of energy received to that known to be originally sent, a distance estimate can be derived. While this solution requires no infrastructure, and would be very cheap to implement, it will not work in this RF environment, for two reasons. First, there is no way to know how much energy the reader sent out: different readers of different sizes and ranges differ in their energy output, and of course a hostile distant reader could simply increase its energy sufficiently to "mimic" the energy level of a weaker, closer reader. Second, the RFID energy field is notoriously irregular, full of local variations, "nulls", whorls, and so forth, and affected by many different environmental conditions: Fishkin et al. [6], for example, describe 10 different variables that can all affect the signal received by a tag from a reader, and that is not an inclusive list: one simply cannot reliably infer distance from signal strength.

4.3 Noise Analysis

We have found that a minor variation on the signal strength analysis *does* correlate tag distance, fairly well, to received energy. The variation is to look at the *noise* in the received signal, in addition to the strength of the signal. Current RF tags and readers do not allow easy access to their returned signals. We therefore use an approximation:

current RF readers do support a "POLL" command, whereby a reader polls for tags in range. If one does a number of poll operations (e.g. 20), we find that the number of responses serves as a reasonable proxy for the signal. The figure below shows a representative example, where the tag is responding to slightly less than 50% of the POLL commands. Note the noise in this signal, showing that even when the environment is fixed, the RF response still fluctuates significantly.

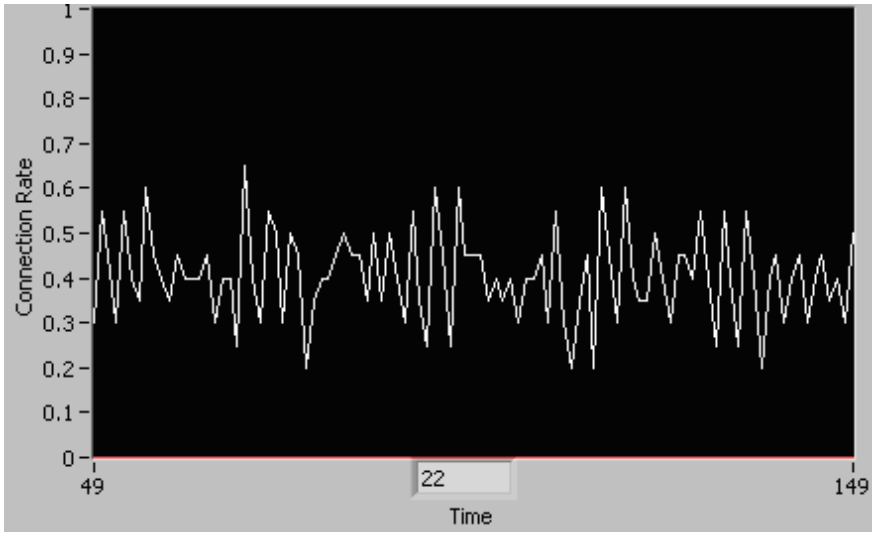

Figure 1: POLL responses over time

We then performed a series of measurements, using the most common commercially available (the Alien technology 915 MhZ reader), with a common commercially available tag, also from Alien technology. As above, the Y-axis represents the response rate (the percentage of responses to a POLL), over 20 polls. The X axis represents time, with 128 such response rate measurements being obtained. Each such X vs. Y plot was then analyzed for mean, standard deviation, and the ratio of standard deviation to mean *FF* (the Fano factor [7], used to approximate signal-noise):

Table 1. The mean (μ) and standard deviation (σ) of response rate, and the ratio *FF* of the two, as tag distance and orientation with respect to a reader varies

Distance (ft)	Angle (deg.)	μ	σ	*FF*
3	0	1.0	0.0	0
3	30	1.0	0.0	0
3	60	1.0	0.0	0.0
3.5	0	1	0	0

3.5	30	1.0	0.0	0.0
3.5	60	0.99867	0.021929	0.021958
4	0	9,8843	0.010754	0.111834
4	30	0.9957	0.0154	0.015467
4	60	0.448047	0.118289	0.26401
4.5	0	0.72226	0.07705	0.106679
4.5	30	0.48085	0.0748	0. 0.1555
4.5	60	0.05898	0.0534	0.905392
5	0	0.70117	0.078485	0.111934
5	30	0.88632	0.5602	0.632052
5	60	0.7957	0.9114	0.114541
5.5	0	0.5105	0.06826	0.133712
5.5	30	0.54296	0.09108	0.167747
5.5	60	0	----	----
6	0	0	----	----
6	30	0.9957	0.0154	0.015467
6	60	0.28906	0.08033	0.277901
6.5	0	0.99453	0.018	0.018099
6.5	30	0.99648	0.01283	0.012875
6.5	60	0.96446	0.0372	0.038571
7	0	0.4914	0.0905	0.184168
7	30	0.69726	0.05929	0.085033
7	60	0.4207	0.06714	0.159591
7.5	0	0	----	----
7.5	30	0.0371	0.0522	1.407008
7.5	60	0.04335	0.03457	0.797463
8	0	0.05	0.05058	1.0116
8	30	0	----	----
8	60	0	----	----
8.5	0	0.100039	0.060018	0.599946
8.5	30	0.08398	0.06241	0.743153
8.5	60	0	----	----
9	0	0	----	----
9	30	0.01757	0.03111	1.770632
9	60	0	----	----
9.5	0	0.64922	0.08716	0.134253
9.5	30	0.6875	0.081085	0.117942
9.5	60	0.63167	0.67495	1.099858

To make the data in this table clearer, the figure below graphs μ only: the average signal strength with distance. Note how it demonstrates the idiosyncracies of RFID communication: signal strength sometimes increases with distance (sometimes dramatically), there are often "nulls" (dead spots) between viable distances, and sometimes increaseng the angle of the tag to the antenna, which should serve to strictly

reduce the amount of energy received and re-transmitted, instead serves to increase it (probably due to reflections off chairs, metal wall beams, etc):

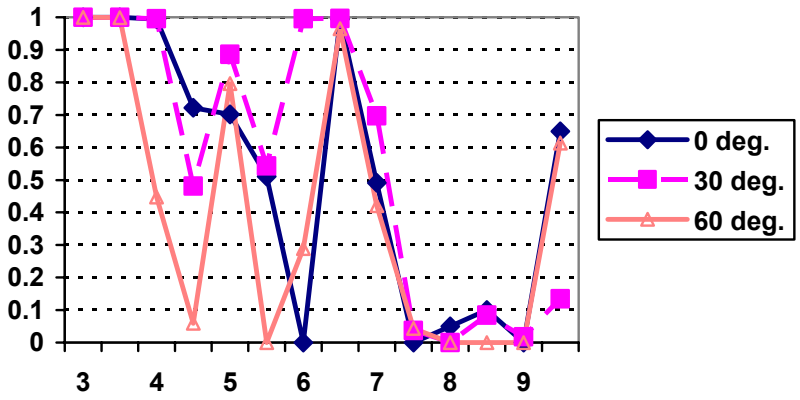

Figure 2: mean signal strength as a function of distance

As we can see, there is no reliable correlation between signal strength and distance in practice, especially as a tag cannot know its orientation relative to the reader.

However, if we instead graph *FF* as a function of distance, the graph changes significantly, for the better:

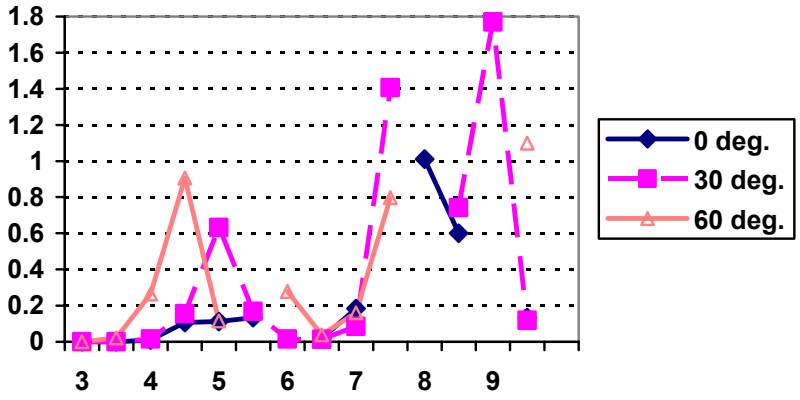

Figure 3: noise ratio as a function of distance

While the correlation is not perfect, particularly at long ranges, it is generally fairly reliable. An important special case here is that *FF* = 0 at all orientations up to around 3.5 feet. Interestingly, this is also when the Fraunhofer "far field" effect kicks in [8]. This effect begins roughly at a distance of 2 *L*L* / λ, where *L* is the diameter of the reader antenna, and λ is the wavelength of the signal, and is when the energy wave from the antenna can be considered in the far field (conceptually, when the curvature of the energy wave is nearly zero). In our case, *L*, the antenna aperture, is roughly 30 cm, and the wave frequency is 915 MhZ, giving us an approximate far-field effect at roughly 1.09m, or 3.5 feet. This means that a tag can not only get a rough estimate of difference (becoming less precise as distance increases) from *FF*, but that it can find a "cliff" between the near and not-so-near distances. A tag on a medicine bottle could, then, for example, return all its information to a reader in the near field (the reader on the shelf in the medicine cabinet should be very near), while returning no or limited information to any reader with a significant *FF*. It is also encouraging that *FF* holds roughly similar across tag orientations, making it a more generally robust and reliable indicator.

Of course, like many Ubicomp sensors, we can see that this is an imprecise sensor: in practice one would have to balance confidence in the sensor value with expected consequences when implementing a range of policies.Our point is that with this mechanism, which requires no change to the RFID communications protocol, no change to the RFID reader, and only minimal change to the RFID tag, we can now enable a set of such policies. A tag might, for example, decide to turn itself into a blocker tag if it believes its reader to be uncomfortably distant, or only respond with a few bits of its ID, or even respond with false data: the CONCEAL command discussed earlier, for example, is a special case of such a policy: tags could seamlessly report only their existence to a reader located a distance from a truck, but then when the factory worker climbs into the truck with a handheld reader, the tags start responding with more detailed and accurate information. We repeat that we do not claim that this measurement, even if it was perfect, could address all (or even most) of the dystopic privacy scenarios, only that it could address a good chunk of them, and do so with virtually no changes in infrastructure, protocol, or hardware.

Finally, note that a distant antenna cannot imitate a closer antenna, as it could with a metric which relied on signal strength alone, as the *FF* metric would still reveal the difference. A close antenna could imitate a more distant antenna by deliberately introducing noise into its signal, but this serves only to signal greater distance (and hence less trust) than is actually the case: spoofing to pretend less distance (and hence more trust) is not possible; a desirable feature.

5 Tiered Revelation

As we mentioned earlier, whatever the policy to decide *whether* to reveal information, nearly all mechanisms make no control over the *level* of information so revealed: either everything is revealed, or nothing. The CONCEAL switch offers a 1-bit intermediary, but that is only a first step. In this section we sketch a design that incorporates recent proposals in Ubicomp sensor revelation (cf [9,10]) in which data provid-

ers *blur* their information to a variable degree depending on the bona fides of the requestor.

Our proposed revelation scheme accordingly moves from a "flat" data space (where either all bits are transmitted or none), to a tiered data space. Every level of disclosure is assigned to a certain tier. When an antenna requests information, it requests a certain tier level. All information at that level and below is transmitted in response. In this way, different readers can be provided with differing amounts of detail, the detail required for their functionality. The more detail the reader requests, the greater its authentication burden, as discussed in the next section.

For example, to show how this tiered revelation might work in practice, let's consider an RFID tagged object such as a shirt made by Benetton. Its tiered information might be structured as follows:

- Level 0 – reports that it is an object. This is useful for baseline testing of a functioning reader, and is equivalent to the CONCEAL level.
- Level 1 – additionally reports that it is a shirt, its fabric, and its color. This is useful for a reader integrated with a washer or dryer – with this information it can tailor its behavior depending on the set of clothes placed in it.
- Level 2 – additionally reports its purchase cost. Now we start to enter the realm of more skeptically granted information. This level would be useful to, for example, an insurance adjuster. By walking through a house with an RFID reader equipped with level 2 authority, it could quickly ascertain the proper amount of insurance needed to cover the objects in the home. By looking at the *FF* in the request earlier, the tag might be able to distinguish between the insurance adjuster and a distant burglar.
- Level 3 – additionally reports which factory it was made at, and at which date. This level is useful to determine if the shirt requires a recall.
- Level 4- additionally reports which store it was bought at, and at which date. This level is useful to determine if the shirt qualifies for a refund/return.

The details of course will vary in practice; this is simply intended as an illustrative example that different levels of information disclosure are needed in different scenarios for the same object.

5.1 Tiered Authentication

This tiered information structure is naturally married to a tiered authentication structure; a reader is provided the information at a given level if and only if it passes an authentication protocol for that level – the higher the level, the more rigorous the authentication protocol.

One method (though by no means the only method) is for the tag and reader to communicate using public-key cryptography, where the number of bits employed is variable, and a function of the desired level of information. In this way, higher and higher amounts of revelation are naturally associated with higher and higher amounts of computation and data transmission.

Therefore, under this protocol, a reader changes its initial request from a blanket request, to a request which indicates:

1) The desired level of revelation
2) The reader public key
3) The number of bits of encryption desired, C
4) The amount of energy received by the tag.

If C (the number of bits of encryption desired) is less than the number of bits the tag requires for that level of revelation, the tag makes no response. Similarly, if the amount of energy received is insufficient, again the tag makes no response; as mentioned below, this provides the advantage of automatically increasing the burden on more physically distant readers. Assuming both of these tests have been passed, the tag responds with a cryptographic challenge-response protocol, keyed to a C-bit encryption.

5.2 Energy-Sensitive Authentication

Our proposed algorithm requires that the reader provide the tag with a minimum amount of energy, and that that amount of energy can be a function of the level of information disclosure required. This requirement can exist even for tags which are *not* passive, i.e. ones which don't need that extra energy. By still requiring a certain minimal amount of energy to reach them, we can again tie the notion of trust to the notion of physical proximity. Further readers, being less trusted by nature, will have to transmit a greater amount of energy than nearby readers, *even if* they pass cryptographic muster on the other criteria. Therefore, a distant hostile interrogator, even one which has cracked the authentication mechanism, may have to increase their energy level to unattainable levels, or at least levels which can be easily detected by mechanisms such as that discussed in the previous section. Again, the essential idea is that the tag reveals the desired information only if the energy signal passes cryptographic muster, *and* if the received energy is sufficient.

In this protocol, then, a reader conceptually offers up a 3-tuple: the desired level of revelation R, the encryption level C, and the energy level E. We point out two special-cases of this:

1. ***C=0***. In this case, the tag is not being asked to perform any encryption whatsoever. However, we are still achieving at least *some* security, by using the required energy level E. By requiring higher levels of energy, we may at least partially satisfy the goal of requiring greater proximity for greater revelation, without requiring any sort of encryption/decryption circuitry on the tag itself.
2. ***E=0***. In this case, the tag is not requiring any particular energy level. In this special case, tag authentication collapses to existing standard methods for tiered information interchange.

6. Conclusions

In this paper, a novel privacy-enhancing technique for RFID is put forward. Based on energy analysis, the mechanism can be divided into two steps. In the first step, an easy-to-compute metric similar to the signal-to-noise ratio is used to estimate the distance between an RFID reader and a tag. In the second step, this distance information is used as a variable in a tiered authentication scheme, where tags reveal more information about themselves to more trusted readers. Trust is a function of (a) perceived distance, (b) level of cryptographic assurance, and (c) level of information desired.

There is a great deal of future work in this area. Given the great requirements for consumer and individual privacy, and the limited and idiosyncratic nature of RFID communication, there are a host of technical, social, cultural, and political issues which will have to be considered to develop satisfactory privacy-preserving RFID deployments. We hope that this paper can help as a stepping-stone along that path.

7. Acknowledgements

We thank Alien Technology ™ Corporation for providing their antenna data.

8. References

1. http://www.rfidjournal.com/article/articleview/796/1/2/
2. http://www.spychips.org/press.html
3. auto-id center. “Draft protocol specification for a 900 MhZ class 0 Radio Frequency Identification Tag”, 23 Feb 2003.
4. Juels, A., Rivest, R., and Szydlo, M. The Blocker Tag: Selective Blocking of RFID Tags for Consumer Privacy. Proceedings of CCS 2003, Washington D.C., pp. 103-111.
5. Sastry, N., Shankar, U.l and Wagner, D.. Secure verification of Location Claims. ACM Workshop on Wireless Security (WiSe 2003). September 2003, pp. 1-10
6. Fishkin, K., Jiang, B., Philipose, M., and Roy, S. I Sense a Disturbance in the Force: Unobtrusive Detection of Interactions with RFID-tagged Objects. Ubicomp 2004 (to appear).
7. Fano, U. (1947). Ionization yield of rations. II. the fluctuations of the number of ions. Phys. Rev., 72:26-29
8. Kraus, J. Antennas. McGraw-Hill. 1988.
9. Gruteser, M., and Grunwald, D. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. MobiSys 2003
10. Hong, J.I., and Landay, J. Architecture for Privacy-Sensitive Ubiquitous Computing. Mobisys 2004. to appear.