

UC Davis

UC Davis Previously Published Works

Title

Trust of Medical Devices, Applications, and Users in Pervasive Healthcare

Permalink

<https://escholarship.org/uc/item/2bh935hx>

Authors

Clifford, Michael

Bishop, Matt

Publication Date

2011-05-01

Peer reviewed

Trust of Medical Devices, Applications, and Users in Pervasive Healthcare

Michael Clifford
Computer Security Laboratory
University of California, Davis
mclifford@ucdavis.edu

Matt Bishop
Computer Security Laboratory
University of California, Davis
bishop@cs.ucdavis.edu

ABSTRACT

In the future, patients may be able to receive health care through the use of pervasive medical devices, sensors, and applications, even outside of hospitals. These data sources monitor and assist patients, aid in treatment, and notify doctors of problems as they develop, allowing them to send help and prepare for emergency treatment faster than otherwise possible. But such data sources are subject to attack or failure. Current trust models guarantee control over access to patient data, but not to determine the trustworthiness of the sources of that data, or of the data itself. This paper shows how the Solar Trust Model can be used to evaluate the trustworthiness of data and data sources in networks of pervasive healthcare devices, sensors, and applications.

Categories and Subject Descriptors

J.3 [Life and Medical Sciences]: Health; C.2.0 [Computer-Communication Networks]: General – *Security and protection*.

General Terms

Security

Keywords

Trust, security, pervasive healthcare, pervasive computing

1. Introduction

Providing health care in the future may rely upon data from a wide variety of devices, such as sensors, monitoring devices, and smartphone applications. These will provide more timely and accurate data about an outpatient's health than is currently feasible. Some devices may behave autonomously, responding immediately to a patient's changing health by using data from sensors and the patient, and following policies set by the patient's doctors. Because many medical devices are highly specialized, sensors and devices with different specializations may monitor different aspects of a patient's health, and provide that data not only to doctors but also to the other devices in the network.

Such a network is not without risk. Sensors and devices may be lost, stolen, attacked, function unreliably, or fail. Applications may be compromised. Untrusted people may access the medical devices and view, alter, or block access to the data. Some sensors

may provide data more reliably and accurately than others. The value and trustworthiness of the data from each source is context dependent, and will be different from the perspective of each entity, such as a person or device, that consumes it.

The Solar Trust Model [1,2,3,4] represents trust relationships between different entities as a dynamic social network. We show how the model can be used to determine which data is sufficiently trustworthy for use by each healthcare device or person, in a particular context, in a pervasive healthcare network.

2. Related Work

There is a large body of work on trust models [10-22]. However, few of these models have been applied to healthcare related problems. A number of security models have been developed for, or applied to, health care. These models focus on who can access or alter data, and under what conditions they can do so. Role Based Access Control (RBAC) [5,6] grants or denies permission for users to access objects based on the role that they are acting in, such as "doctor" or "patient". RBAC has been extended [7] to add context sensitivity to access control in health care. Another approach [8] determines whether a user is sufficiently trusted to access specific data in a distributed healthcare system. Other work [9] uses a set of principles designed to ensure privacy and integrity of patient data by limiting who is allowed to access the data, preventing deletion of existing data, requiring patient notification upon certain kinds of data accesses, and in some cases obtaining patient consent for those accesses. These approaches all determine whether or not someone is trusted to access medical data. They do not examine whether the *source* of the data, or the data itself, should be trusted. Our work examines whether or not sources of medical data are themselves trustworthy.

3. A Pervasive Healthcare Scenario

Consider a patient who, with the aid of remote monitoring, can receive some of their care outside the hospital. This patient could be monitored by remote sensors, which would send data back to the hospital. One nurse might monitor many remote patients simultaneously, proactively alerting both doctors and patients to any problems and required actions. Automated systems could even continuously analyze the data, identifying problems using signatures or heuristics before they would be evident to a human.

We examine trust in pervasive devices using the following example. Alice is a doctor at a hospital. Her patient, Bob, is recovering at home from a major operation. Alice orders that he

This material is based upon work supported by the National Science Foundation under Award CCF-0905503. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation.

be monitored with a suite of sensors, medical devices, and interactive smartphone applications, in order to detect and respond to any complications. If a complication occurs, an ambulance will automatically be called to take him back to the hospital while Alice prepares emergency treatment, thus cutting precious minutes off the time for him to receive care. The sensors used to monitor Bob are pervasive. Some may attach to his body, while others are distributed around his car and office. Bob may even use sensors that are set up to authenticate and monitor any patient in the area, not just him. Bob is also able to take advantage of a suite of pervasive devices that take data from the sensors, and use it to make automated decisions, such as how much medication to dispense, or whether to apply a life-saving therapy until a physician can respond. The devices may share information with each other to coordinate their responses to provide better overall care. Like the sensors, these devices may also be located anywhere, and thus may serve many potential patients, not just Bob.

Applications running on Bob's phone behave like both sensors and devices. They collect data about Bob, authenticate him, and giving him important treatment information.

Data from the sensors, devices, and applications may be routed to Bob's doctors, to other clinicians, or to other devices, through any available network, such as WiFi or cellular networks. This must be done in a way that protects Bob's security and privacy. Further, the network involved is not static. New doctors, devices, sensors, and applications may join or leave at any time, as Bob's needs change, and as he moves from one location to another.

Treatment may require devices to act on data in some sequence in order to determine the correct course of action. For example, data from temperature and heart rate sensors may be sent to a device that computes and dispenses the correct dosage of medication. Dosage data may be sent to another device that alters its behavior to allow the medications to work more effectively. Data from that device may be routed to Alice. Each of these relationships between producers and consumers of data forms a path, and the union of all such paths forms a dynamic medical data network.

In a pervasive medical network, data may come from many sources, and may follow many paths. Different entities may control different sources. For example, Bob, Alice, other doctors, and third parties including attackers may each control one or more devices in the network.¹ Data routed along one path may pass only through devices trusted by the consumer of that data, while data routed along another path may pass through one or more untrusted devices. Data may also be received from more than one path simultaneously. Thus, each consumer of data, whether a device, application, or person, needs to be able to determine the extent to which it can trust the data that it is receiving, and which sets of data it can consider most trustworthy in a given context.

4. Trust

Trust is a subjective, relative measure of the degree to which some entity "believes" that another entity will exhibit a certain set of properties in a certain context. For example, you might have great confidence in a particular doctor to diagnose and treat most illnesses, but little confidence in that doctor's ability to perform brain surgery because they were not trained as a brain surgeon.

Context is the set of information that a specific entity uses in making a trust judgment, and a set of constraints on the applicability of the scope of that trust judgment. For example, if Bob is young and healthy, his benefits from taking a certain medication may be outweighed by the potential harm from its side effects. Conversely, if Bob is likely to die without taking the medication, then the risk posed by taking the drug may be small compared to the consequences of not taking it. Likewise, the answer to "Who is the best doctor?" depends on the context of the problems that are being treated, as different doctors specialize in treating different kinds of problems. Also, as different people will apply different preferences, knowledge, and experiences to determining the correct answer, that answer depends on who is asking the question.

Trust is subjective because what makes something trusted is specific to the individual making the trust judgment, and is based on that person's knowledge and experience. For example, when Bob chooses a doctor, he may have many criteria in mind, such as the doctor's qualifications, experience, success rate, cost, and ease of access. Bob weights different criteria differently, based on their importance to him. If another patient, Rachel, looks for a doctor, she may rely on a different set of criteria and weights.

Because trust is subjective, it doesn't make sense to measure it using a fixed scale. For example, it does not make sense to say Alice can be trusted to a degree of .6 units, because each person views trust differently, and that interpretation of trust varies based on context. So we use a notion of relative trust. Relative trust describes the trust that some entity places in some object in relation to the trust that they place in another object, viewed within the constraints of a specific context. For example, Alice is a brain surgeon and Charlie is a heart surgeon. In the context of performing brain surgery, Bob will probably trust Alice more than Charlie, because Alice is trained in brain surgery and Charlie is not. On the other hand, Bob will probably trust Charlie more than Alice to save his life after a heart attack. Because relative trust does not require users to agree on a common scale for rating trust, they can communicate about the degree to which they trust specific objects without agreeing upon criteria for evaluating that trust. Relative trust also makes possible defining arbitrary limits on whether some object is trusted, such as "Bob trusts Charlie to perform brain surgery only if Bob trusts Charlie at least as much as he trusts Alice to perform brain surgery."

5. The Solar Trust Model

The Solar Trust Model provides a way to model trust relationships between a set of nodes as a dynamic social network of trust relationships. The low-level design of the model has been extensively discussed in prior work. Here we simply describe the model at a high level, and then show its applicability for use in networks of pervasive healthcare devices.

To see how the model works, imagine that Bob has arrived in a new city and needs a doctor. Bob could look in the phone book and pick a name at random, but this gives Bob no information about how good the doctor might be at addressing his problem. Bob might ask his friend Danica whom to see, and Danica might recommend Alice. If Danica is a doctor, then Bob might trust her recommendation highly. If Bob knows that Danica knows little about doctors, he might not trust her recommendation, and might ask Charlie, whose recommendations he trusts more, instead.

Note that Bob has a different kind of relationship with Danica than he has with Alice. Bob knows Danica directly, so we say

¹ We assume that any device or sensor could fall victim to a successful attack.

that Bob has a direct trust relationship with Danica. This means that he can use his experiences with her, and any knowledge about her, to aid in deciding how much to trust her recommendation. If Danica makes a bad or inaccurate recommendation to Bob, he may trust her future recommendations less. If his experience matches her recommendation, he may trust her more. If Bob has never interacted with Alice, he cannot judge how much he should trust her based on direct experience. He can make a judgment about Alice based on his indirect relationship with her through Danica. This is called an indirect trust relationship.

The extent to which Bob trusts information about Alice depends on the path that that information takes from Alice to Bob (potentially traversing other nodes in the process). Bob trusts some paths a lot, and others very little. If a path is insufficiently trusted, Bob will not trust information sent along it. If there are multiple, sufficiently trusted paths, Bob will trust most highly the information that followed the most trusted path. Note that Bob can order all of his paths to Alice based on their degree of trust.

To learn of paths to Alice (and other entities), Bob keeps track of his direct trust relationships using a “solar system”. A solar system consists of a star (the user, or a server acting as a proxy), and a set of orbits. Each orbit contains entities that are trusted to the same extent in the same context. Orbits are ordered from most trusted to least trusted. Users can use any policy to map entities into orbits, so long as they have a direct trust relationship with those entities. If a user changes its relationship with an entity, the entity moves to an orbit reflecting the new relationship.

Given the model’s representation of a user’s direct relationships, one can find all the sufficiently trusted indirect relationships with other entities using the Solar Trust Model’s path finding algorithm [2,3,4]. To do this, each Solar Trust server sends out a path query to each entity in a sufficiently trusted orbit. A copy of the query is propagated along each sufficiently trusted direct relationship in the solar system of each subsequent entity until either the query cannot be forwarded to a sufficiently trusted entity, or until adding the next entity would result in the maximum path length for at least one solar system along the path being traversed. The complete path followed by the query (including the last entity) is digitally signed and sent back to each entity along the path for which the path has just terminated.

The trust network may not remain static. When a new person or device joins the network, it runs the path finding algorithm, finding all new paths of trust, and updating any existing paths that may have changed as a result of the new entity joining. When a person or device leaves the network, they do so gracefully (by informing the affected upstream nodes of the change, thus causing the network to automatically update), or through path aging (paths slowly become less trustworthy over time if not refreshed). If an existing node changes one of its relationships, it runs the path update algorithm, which alerts all affected nodes to the effects of the changes. In this way, each participating entity can maintain an up-to-date set of direct and indirect trust relationships with every other entity that it trusts sufficiently.

The result of this process is a dynamic network of paths of trust from each entity to the other entities it trusts sufficiently. Each trust relationship reflects the preferences and experiences of the entity at which the relationship originates, and the knowledge and experiences of any intermediate nodes between that entity and the object it is evaluating.

Once the network has been established, messages can be sent from a sending entity to a receiving entity along a path that the

receiving entity trusts sufficiently. Messages are digitally signed to ensure that they can be trusted to the same extent as the path to the private key used to sign them.

6. Applying the Solar Trust Model to Pervasive Healthcare

We now apply the Solar Trust Model to our pervasive healthcare example. In our example, Alice, other doctors, Bob, and each medical device are all users, with their own solar systems. If they receive data directly from a source (another user or sensor), that source is placed into the appropriate orbit in their solar system. For example, Alice would place data she receives from Bob, along with data received directly from any of the medical devices monitoring him, into her solar system.

To determine how to assign different entities, such as Bob, sensors, devices, and other doctors, to different orbits, Alice needs to define a context and a policy associated with that context. In the context of information about patient symptoms, such as “My chest hurts!” Alice would likely place Bob in an orbit with a high degree of trust, because she knows that Bob is likely to report his symptoms correctly.² Alice might also assign information from doctors who have recently examined Bob, and from devices and sensors that can monitor Bob’s health directly, into high orbits, and map information from doctors who have not examined Bob, but who have experience with similar patients, into a slightly less trusted orbit. On the other hand, because she doesn’t find them credible, Alice might place articles from health blogs into untrusted orbits. Conversely, in the context of treatment options, Alice may assign information from doctors with expertise in Bob’s symptoms to the highest orbit in her solar system, while information from Bob himself might be mapped in a low orbit.

Bob may set up his own solar system using a different set of policies, representing that his experiences and needs are different than Alice’s. In the context of which treatment advice is the best, he might map information from Alice into a high orbit because he trusts Alice’s medical expertise. He may also place information from health blogs in a high orbit because he believes that blogs may provide him with useful information. Conversely, he may assign information from devices and sensors to low orbits because he does not understand the information that they provide, and therefore cannot judge if that information is likely to be correct.

To determine how to assign orbits to information from other devices and sensors, devices may follow some set of heuristics, or consult a database. They may also make different decisions based on context. Suppose multiple sensors monitor the level of a certain hormone in Bob’s body. Sensor A may provide the most accurate measurement, but only provide a reading once an hour. Sensor B may provide less accurate readings, but do so every second. If a device must decide how to behave based on the best available sensor data, it may place sensor A in a more trusted orbit than sensor B if high accuracy is more important than frequently updated results, and sensor B in a more trusted orbit than sensor A if data freshness is more important than accuracy. This approach allows the device to fall back on data from the less trusted sensor if the sensor that is more trusted in that context fails.

Considering several cases will show how this affects the pervasive medical network. Imagine that a sufficiently trusted path exists from Alice to Device A to Device B to Sensor C. Device A

² If Bob has a history of inventing symptoms, Alice might place information from Bob into a lower orbit.

determines that the data from Device B is no longer being delivered reliably. Device A will switch to the next most trusted path to Sensor C (say, through Device D), and will inform Alice's server (and any other servers that have paths passing through B) that paths with the subpath A→B are no longer valid. Because she no longer has access to data from B, Alice may attempt to find another path that provides equivalent data, or may be forced to make decisions without that data being available.

Now assume that Emily is a doctor, and Alice decides to consult with her about Bob's condition. Because of her expertise, Alice places Emily in a highly trusted orbit, and finds sufficiently trusted paths through Emily. Emily may not be able to access data about Bob directly, but can access them indirectly if she establishes a sufficiently trusted path through Alice. So Emily trusts the data about Bob to the extent that she trusts Alice to provide the most relevant data about Bob, combined with Alice's own evaluation of how much she trusts the data she is receiving about Bob. Alice trusts Emily's recommendations about Bob as far as she trusts Emily to make recommendations, combined with the extent to which Emily trusts the data about Bob.

Finally, suppose an entity in the network fails to respond in a timely manner. Device B monitors Bob's blood pressure, but generates new data only once an hour. Sensor C monitors Bob's heartbeat continuously. Bob's heart stops. Data from the heartbeat monitor will detect this quickly, and therefore be highly trusted in this context. But the data from the blood pressure monitor will become untrusted because it will not reflect the current state of Bob's health; to do so would require an immediate response.

7. Conclusion

Pervasive healthcare relies on the availability of data, and on devices that can produce, analyze, and use that data in an unconstrained environment such as a home. Different devices and sensors provide different specialized functions. Their combination may provide better health care information than each individual, isolated device can. These devices and sensors may join or leave a network at random, may be subject to attack, and may have variable reliability or accuracy. We showed how the Solar Trust Model can determine the relative trustworthiness of data from many potential sources, thus giving entities the ability to select the source best suited for that particular context, thereby enabling the improvement of data gathered for pervasive healthcare.

8. References

- [1] Clifford, M., Lavine, C. and Bishop, M. 1998 "The Solar Trust Model: Authentication Without Limitation." In *Proceedings of the 14th Annual Computer Security Applications Conference*, pp. 300-307.
- [2] Clifford, M. 2002. "Networking in The Solar Trust Model: Determining Optimal Trust Paths in a Decentralized Trust Network." In *Proceedings of the 18th Annual Computer Security Applications Conference*, pp. 271-281.
- [3] Clifford, M. 2006. "An Implementation And Performance Analysis Of The Solar Trust Model, A Dynamic, Distributed Trust And Authentication Model." Masters' Thesis, The George Washington University.
- [4] Clifford, M. *Attacks and Defenses on the Solar Trust Model*, Ph.D. Thesis, University of California, Davis. In preparation.
- [5] Ferraiolo, D.F. and Kuhn, D.R. 1992. "Role Based Access Control." In *Proceedings of the 15th National Computer Security Conference*. pp. 554-563.
- [6] Sandhu, R.S. Coyne, E.J. Feinstein, H.L. and Youman, C.E. 1996. "Role-Based Access Control Models." In *Computer*, vol. 29, no. 2, pp. 38-47.
- [7] Hu, J. and Weaver, A.C. 2004. "A Dynamic, Context-Aware Security Infrastructure for Distributed Healthcare Applications". In *Proceedings of First Workshop on Pervasive Security, Privacy and Trust*.
- [8] Weaver, A.C. Dwyer, S.J. III. Snyder, A.M. Van Dyke, J. Hu, J. Chen, X., Mulholland, T. and Marshall, A. 2003. "Federated, Secure Trust Networks for Distributed Healthcare IT Services." In *Proceeding of the International Conference on Industrial Informatics*, pp. 162- 169.
- [9] Anderson, R. 1996. "A Security Policy Model for Clinical Information Systems," In *Proceedings of the 1996 IEEE Symposium on Security and Privacy* pp. 30-43.
- [10] Abdul-Rahman, A. 1997. "The PGP Trust Model." *EDI-Forum: the Journal of Electronic Commerce*.
- [11] Biba, K. 1977. *Integrity Considerations for Secure Computer Systems*. U.S. Air Force Electronic Systems Division Technical Report 760372.
- [12] Blaze, M. Feigenbaum, J. and Lacy, J. 1996. "Decentralized Trust Management." In *Proceedings of the 17th IEEE Symposium on Security and Privacy*, pp. 164-173.
- [13] Bell, D. and La Padula, L. 1973. *Secure Computer Systems: Mathematical Foundations*. MITRE Corporation.
- [14] DeFigueiredo, D., Barr, E. and Wu, S.F. 2009. "Trust Is in the Eye of the Beholder." In *Proceedings of the 2009 IEEE International Conference on Privacy, Security, Risk and Trust*.
- [15] International Telecommunication Union. 1995. *Information Technology—Open Systems Interconnection—The Directory: Authentication Framework*. ITU-T Recommendation X.509, pp. 7-17
- [16] Mayer, R., Davis, J., and Schoorman, F. *An Integrative Model of Organizational Trust*. The Academy of Management Review, Vol. 20, No. 3 pp. 709-734, July 1995.
- [17] Network Associates. 1999. *How PGP Works*. PGP 6.5.1 documentation. <http://www.pgpi.org/doc/pgpintro/>.
- [18] Linn, J. 1993. *Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures*. RFC 1421.
- [19] Abdul-Rahman, A. and Hailes, S. 1997. "A Distributed Trust Model. In *Proceedings of the 1997 Workshop on New Security Paradigms*, pp. 48-60.
- [20] Spear, M., Lu, X., Matloff, N. and Wu, S.F. 2009. "KarmaNET: Leveraging Trusted Social Paths to Create Judicious Forwarders." In *Proceedings of the First International Conference on Future Information Network*.
- [21] Wang Y. and Vassileva J. 2003. "Trust and Reputation Model in Peer-to-Peer Networks." In *Proceedings of the IEEE Conference on P2P Computing*.
- [22] Xiong, L. and Liu, L. 2003. "A Reputation-Based Trust Model for Peer-to-Peer E-Commerce Communities." In *Proceedings of the IEEE International Conference on E-Commerce* pp. 275-284.