



City Research Online

City, University of London Institutional Repository

Citation: Gadala, M., Strigini, L. & Fujdiak, R. (2022). Authentication for Operators of Critical Medical Devices: A Contribution to Analysis of Design Trade-offs. In: Proceedings of the 17th International Conference on Availability, Reliability and Security. . ACM. ISBN 9781450396707 doi: 10.1145/3538969.3544474

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/28633/>

Link to published version: <https://doi.org/10.1145/3538969.3544474>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Authentication for Operators of Critical Medical Devices: A Contribution to Analysis of Design Trade-offs

Marwa Gadala
m.gadala@aston.ac.uk
Aston University
Birmingham, United Kingdom

Lorenzo Strigini
l.strigini@city.ac.uk
City, University of London
London, United Kingdom

Radek Fujdiak
fujdiak@vutbr.cz
Brno University of Technology
Brno, Czech Republic

ABSTRACT

Increasingly evident safety risks due to attacks on safety-critical devices are causing new requirements for authentication of these devices' human operators. These requirements have now extended to medical devices. However, authentication may also introduce new safety risks, reduce usability, cause delays, and/or encourage user behaviors that compromise the very security it should protect. Thus, design of authentication mechanisms needs to take on a holistic approach that considers such interrelationships, and the effects not just of the general method chosen (say, passwords vs. fingerprints), but also of its implementation details. We illustrate this problem on a medical case study. We report early steps in a trade-off analysis that captures interactions between safety, security, usability and performance issues, to assist designers in choosing and tuning viable solutions. A qualitative analysis to narrow down the field of possible solutions is followed by a probabilistic analysis. The analyses highlight non-obvious links between system attributes, especially links due to the complex way humans interact with, and adapt to, such devices. The probabilistic analysis systematically describes risk as a function of the authentication method and its design parameters. We show example results quantifying how some key design parameters produce opposite effects on risk due to accidental and malicious causes, requiring a trade-off: the quantitative model allows the designer to manage this trade-off to achieve an acceptable level of overall risk, taking into account environmental factors like the expected prevalence of certain attack types. Both the qualitative and quantitative approaches aim to help device designers make rational decisions about authentication options and the tuning of their design parameters.

CCS CONCEPTS

• **Security and privacy** → **Logic and verification**; Social aspects of security and privacy; *Vulnerability management*.

KEYWORDS

authentication, medical security, trade-offs, access control, medical devices, safety risk, usability, qualitative analysis, probabilistic model, critical use by multiple operators, medical environments present especially difficult challenges in the design of authentication, and in this work we focus on a medical case study.

ACM Reference Format:

Marwa Gadala, Lorenzo Strigini, and Radek Fujdiak. 2022. Authentication for Operators of Critical Medical Devices: A Contribution to Analysis of Design Trade-offs. In *The 17th International Conference on Availability, Reliability and Security (ARES 2022)*, August 23–26, 2022, Vienna, Austria. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3538969.3544474>

1 INTRODUCTION

The spread of computer control in all kinds of equipment, from factory machines to medical devices, increases interest in protecting against malicious use. True, computer control usually increases efficiency, convenience, and accuracy in the tasks performed, but often it also introduces security risks [Mehrfeld 2019]. Software control may allow subtler and more complex subversion than previously possible, due to more flexible operating modes, the interconnection of various devices, and, frequently, the presence of multiple operator roles with different privileges and modes of use. In particular, the issues posed by such security risks in medical devices have come to the forefront in recent years. Furthermore, while security vulnerabilities have been detected in medical devices, but with little evidence of exploitation yet reported, increasing security breaches in industry demonstrate the potential to cause physical harm or even death [Abraham et al. 2019].

To prevent malicious use of these potentially dangerous devices, it is natural - apart from limiting access to them by physical security measures - to use authentication, the automatic process of verifying the identity of a user. Authentication schemes already operate in many domains ranging from banking to border control [Palmer 2010]. Furthermore, a need for authentication functions is already acknowledged by some regulatory and standardisation bodies [Altinkemer and Wang 2011; Krol et al. 2015]. In the medical domain, FDA guidance for medical devices suggests that manufacturers consider, among security functions, "limiting access to devices through the authentication of users (e.g. user ID and password, smartcard, biometric)" [Center for Devices and Radiological Health 2014]. With many devices (a report indicates an average of more than 15 devices per hospital bed [Horblyuk et al. 2012]) in use by multiple operators, medical environments present especially difficult challenges in the design of authentication, and in this work we focus on a medical case study.

2 CURRENT CHALLENGES

Authentication addresses several types of risks:

- **Malicious Use.** The most obvious need for authentication is to prevent malicious use of a device. For example, preventing someone from changing settings on a life-supporting device in order to harm a patient.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ARES 2022, August 23–26, 2022, Vienna, Austria
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9670-7/22/08.
<https://doi.org/10.1145/3538969.3544474>

- **Well-Intentioned Abuse.** In most cases, use of critical functions must be reserved to staff with the required knowledge and authority, and authentication can be used to enforce these constraints. For instance, authentication on an infusion pump can enforce rules on which staff have authority to administer or change a dose, and perhaps which staff only have access to retrieve, but not change that information.
- **Accidental Misuse.** Authentication can also protect against unintentional command entry: for example, preventing a visitor who accidentally trips in a hospital room from changing parameters on a medical device.
- **Use as Means to Attack an Organization.** It is not just devices that perform critical functions that need to be secured; the trend towards integrated devices and interconnected infrastructure means that any device may be used to penetrate an organization's network. For instance, a TrapX report [Francis 2017] described medical devices being targeted by cybercriminals: blood gas analyzers and radiology equipment contained backdoors into hospital networks allowing attackers to send patient records to unknown locations in Europe and China. The report also highlighted an increasing trend: reported attacks with over 500 breached patient records rose more than 50% from 2015 to 2016. And although TrapX reported attacks to steal patient records for economic gain, not to physically harm patients, this is no justification for ignoring the latter until possibly catastrophic attacks.
- **Violations of Accountability and Corruption of Logs.** An auxiliary role of many devices is to keep logs of operations and commands received, and authentication can ensure that these logs link actions to the specific operator who performed them. Such logs can assist user training, handovers, and, in case of mishaps, incident investigations and liability tracking. But users may fear being accused of mistakes (or of violating rules, albeit for good reasons) and thus try to alter logs, prevent their identity from being recorded, or share their authentication credentials (intentionally or unintentionally).

Whatever the purpose(s), deciding an appropriate authentication solution for a given device and use environment presents a complex challenge. Importantly, this challenge is not limited to technological aspects, and this section aims to highlight various aspects of the problem.

2.1 Effects of Authentication on Safety, Usability and Performance

One of the main difficulties in choosing an appropriate authentication method stems from the multiple, and sometimes negative, effects that can result from introducing authentication. If not designed carefully, authentication measures may slow down user performance, be a nuisance to some users, or be undermined by user workarounds. In extreme cases, authentication might even increase the overall risk in the operation of a device. For example, a patient may suffer harm if authentication delays, or wrongly prevents, a clinician's order to an infusion pump to administer a drug in an emergency. Increasing security via authentication has

effects on other inter-related attributes of the socio-technical system, such as safety and performance: the problem requires a careful and holistic trade-off analysis, in order to avoid unintended and counter-productive side effects of selecting an unsuitable authentication scheme [Palmer 2010]. Below, we consider a few examples of such design trade-offs:

1. In one-time authentication methods, users often remain authenticated until a timeout occurs, the user intentionally disconnects or the session is otherwise broken. But users often do not log out, which gives rise to security risks. The most common solution, an inactivity timeout (sometimes referred to as authentication expiry time) is inevitably a trade-off: shorter timeouts reduce security risks at the cost of user effort and possibly safety, while longer timeouts improve usability and performance, by reducing the frequency of authentication requests. Other solutions include implementation of continuous authentication methods, such as those based on user proximity to the device. However, these may trigger privacy concerns of authentication without consent [Krol et al. 2015; Schwartze et al. 2014], issues of reliability and accuracy [Bonneau et al. 2012; Halunen et al. 2017; Krol et al. 2015], and also only confirm whether the user is nearby but not whether the user is actually using the device.

2. Although token- or biometric-based methods may have speed or convenience advantages over knowledge-based authentication methods, their reduced reliability in some situations may introduce new safety hazards. For example, a card reader can fail to read and some biometric methods still have high false reject rates, complicated by day-to-day user and environmental variations. In a safety-critical environment, these drawbacks can pose serious safety hazards by causing delays in emergency situations.

3. The "redundancy" of multi-factor authentication augments security, but at the same time often "diminishes the user experience" [Braz and Robert 2006]. From the user's perspective, multi-factor authentication imposes additional physical and cognitive burdens (e.g., requiring them to carry an additional token) [Krol et al. 2015], increases the time needed to authenticate [Altinkemer and Wang 2011; Krol et al. 2015], increases the number of errors and lockouts [Krol et al. 2015] and raises issues regarding the use of additional user information collected [Altinkemer and Wang 2011; Braz and Robert 2006]. Furthermore, security benefits of two-factor authentication may be less than desired. It is to be noted that the presence of a second factor can cause users to choose weaker passwords than if passwords alone are used to protect an account. In other words, especially where user choice is involved, there can be erosion of the efficacy of one protection method when users know that there is a second one as well.

Users are normally aware of some of the above-mentioned harms to usability, efficiency and safety that can be caused by authentication schemes, and often seek to mitigate them. Their "workarounds" can then contribute to new, "second-order" security risks. A heavily studied and well-known example of this concerns the use of passwords. It is hard to remember many long, complex passwords. Thus, users sometimes seek workarounds that can undermine security, e.g.: (i) having one password for many devices, (ii) choosing a password that is easy to remember thus also easy to crack, (iii) sharing passwords, (iv) keeping passwords on paper or plaintext

files, or (v) reusing or recycling old passwords [Zhang-Kennedy et al. 2016].

This is not to say that users are to blame for such workarounds. A “user’s capacity for effort (basically a combination of time and energy) is one of the most valuable and scarce resources available in the information security field. If a user is expected to spend his or her resources inefficiently or ineffectively on security it can only be expected in return that security instructions will be ignored or circumvented” [Kiljan et al. 2018]. However, by deviating from mandated practices, users can make even the best protected devices vulnerable.

Currently, an all-too-common understanding is that the more burdensome security measures are on authorized users, the more secure a device is against unauthorized users; in other words, that usability must be sacrificed in order to achieve significant gains in security [Cranor and Buchler 2014]. However, studies are finding that this is not the case, and that this is often counterproductive [Kroeger et al. 2013]. Stringent security policies that make people’s work difficult can become self-defeating. The National Academy of Engineering (NAE) has noted that “if security systems are burdensome, people may avoid using them, preferring convenience and functionality to security” [Cranor and Buchler 2014]. Ultimately, the security of a scheme very much depends on the end users, as they choose whether to adhere to it [Cranor and Buchler 2014; Zimmermann¹ et al. 2018]. Thus, it is not enough to add security measures unless designers can also ensure that users will embrace them. This is not a straightforward task. For example, using passwords that are more difficult to crack may make them more difficult to remember. Due to such considerations, the U.S. National Institute of Standards and Technology and the U.K.’s National Cyber Security Centre recently reversed their long-standing advice on password policies, acknowledging that policies previously considered “most secure” (complex passwords, changed frequently) caused users to invent workarounds that undermined authentication [Grassi et al. 2020]. Thus, more consideration of the effect of security requirements, especially in embedded systems, on other system attributes such as safety or usability is needed during the whole development process, including in relevant standards. Careful trade-off analysis is essential as the consequences of selecting an unsuitable scheme can result in users choosing alternative ways to achieve their aims or circumventing the scheme [Palmer 2010].

2.2 The Articulation and Agreement of Requirements from Different Stakeholders

Even during the early stages of requirement elicitation and early design development, there is often difficulty in “articulating and agreeing requirements from different stakeholders” regarding an appropriate authentication scheme [Palmer 2010]. Stakeholders are often driven by different, and possibly conflicting interests and motivations (e.g. mitigating risks versus improving ease of use). For instance, [Krol et al. 2015] describe a strong correlation between users’ preferred authentication scheme and perceived convenience, “but only a weak correlation between perceived security and perceived convenience/usability”. Even within the user community, “feelings, competencies and preferences may vary considerably” [Palmer 2010]. For example, a relevant distinction is between the

needs and requirements of users preferring convenience and those prioritizing privacy.

These differences are further complicated by the users’ possibly complex mental models that affect how they perceive a given authentication scheme. For example, in the usability study presented in [Krol et al. 2015], researchers noted that some users felt their bank account was already secure before the need for a hardware token, because they had not experienced any fraud and/or because they believed their password was already difficult to crack. Some users believed that authentication was faster in the morning than in the evening. These observations reveal that users, based on their personal experiences and understanding, develop mental models about how an authentication scheme works as well as its perceived reliability. This may create even more diversity in the user population’s needs and requirements.

“An evaluation of suitability needs to encompass the diverse perspectives and values of the stakeholder groups involved, whether directly in the use of the [scheme] or indirectly as a consequence of the [scheme’s] failure” [Palmer 2010]. To achieve this, it is essential to fully engage with the stakeholders in order to obtain their commitment for the scheme. For example, running usability tests not only to validate task efficiency, but also, to assess user confidence in the ability of the authentication scheme to protect their interests [Palmer 2010]. Also important is clarity in identifying stakeholder requirements and preferences, from the early stages of development as well as participatory involvement throughout the development life cycle in order to iteratively refine user needs and preferences [Palmer 2010]. Simply imposing an authentication solution can have serious consequences for the stakeholders involved and may result in some users, when possible, choosing to take their business elsewhere, as found in [Krol et al. 2015].

2.3 Fine-Grained Design Decisions

Another difficulty in choosing an appropriate authentication method is that authentication choices are not just about general approaches to authentication, but also about detailed design of how a protected device will use authentication. Examples of key factors that a designer must consider include:

- **System Boundaries.** It is important to decide where authentication is required: e.g., at the entrance of a control room versus on a device itself. For example, medical devices in an operating room may rely on physical security of the room to avoid the need for authentication on each device. Another type of boundary concerns which operator commands and use modes require authentication; these could be limited to ones that could lead to high severity consequences (e.g., monitoring patient parameters need not require the same strength of authentication as infusing a patient with a drug).
- **User Roles.** Devices are often used by a range of users with different privileges and “authorization is the concept of specifying what a user or entity is allowed to do once they have authentication” [Guel 2002]. For example, a nurse may be allowed to silence an alarm, but not to set certain parameters affecting a patient’s treatment. The scalability and manageability of such authorization systems can be a big concern,

especially for large organizations [Guel 2002]. Enforcing these rules through authentication mechanisms makes them harder to bend or break for special cases and emergencies; this may require a reassessment of the potential role of intentional violations of rules in ensuring patient safety in such special cases.

- **Workflow and Environment.** It is important to understand the setting in which authentication will happen. For example, fingerprint authentication may not be practical in a setting where gloves are required.
- **Authorization Expiry Time or Grace Period.** Choosing an appropriate "authentication expiry time" is a decision that requires device designers to take into account many factors, including the expected frequency of attacks. As discussed in the previous section, allowing extra time before initiating authentication requests can reduce the burden of authentication and improve user compliance, but also increase security risks.
- **Break-Glass Considerations.** To mitigate safety risks due to authentication, e.g. delayed operator intervention in emergencies, one may consider "break-glass" options, e.g. multiple modality options so that users can access a device using various methods (card and password) in case one (e.g., card) fails. Alternatively, high-priority alarms could override the need for authentication: e.g., if a patient's vital signs drop below a critical level, a clinician may be allowed to infuse a drug without authentication. This design precaution could alleviate some safety concerns, but also raise other concerns about patient safety from attacks during an emergency.

2.4 Related Work Comparing Authentication Methods

Thus, for a chosen authentication approach, numerous, interleaved design factors and policy variations need to be considered, to optimize its effects, as far as feasible, and mitigate new risks. In the literature, several articles focus on comparison and selection criteria or decision frameworks for authentication schemes. According to our own search and a recent 2018 review [Velásquez et al. 2018], not many articles related to this were found. In these papers, each authentication scheme (sometimes referred to as method or proposal) is evaluated against a set of criteria (sometimes referred to as properties or benefits or metrics). These studies all highlight that no scheme examined is perfect. The studies are defined by (1) the criteria/viewpoints they consider, (2) the context of use they study, (3) the schemes they evaluate, and (4) their comparison method.

In our study of these works, we identify the following key gaps that motivate our work and set it apart from previous studies: (1) The criteria often focus only on usability, deployability/cost and security/privacy [Velásquez et al. 2018]. There is a need, especially in certain contexts, to consider other attributes such as safety. (2) The contexts considered in comparison articles we found include: banking, wireless, mobile, cloud, gadget-free technology [Halunen et al. 2017], multimedia communications [Eliasson et al. 2009], and clinical documentation workflow [Schwartz et al. 2014]. The context of use is as an important element, as the articles either consider this as one of the decision criteria or the article's proposal itself is

directed to a specific context. Different applications/domains (e.g., banking vs. gaming), have different requirements and priorities. Thus, there is likely not one generic solution/optimal scheme or framework that fit all environments, and we note a lack of studies in healthcare contexts, possibly because the idea of authentication is not as mature as in other domains. (3) Certain schemes have not been thoroughly considered; for example, continuous authentication. "Continuous authentication is a critical component to any resilient solution. Such approaches move beyond traditional passwords, cryptocards, and smart badges, which only provide a simple instant of trust in the current context" [Kroeger et al. 2013]. (4) There is often a focus either on qualitative or quantitative methods, rather than a combined approach.

3 A QUALITATIVE ANALYSIS OF THE AUTHENTICATION ISSUE

We have so far referred to problems of authentication for medical devices in general. We now introduce a concrete example. One of the "use cases" in project AQUAS (Aggregated Quality Assurance for Systems)¹ concerned extensions of an existing device for monitoring blood pressure and neuromuscular transmission, enabling it to control an infusion pump and perform closed-loop control of these physiological parameters². Clinicians provide inputs such as an initial infusion dose and target values for the physiological parameters; the device then calculates and infuses appropriate doses of drugs to maintain the parameters within the given targets. Our risk analysis (including a partial Hazard and Operability analysis [IEC 2016] with the designers and other AQUAS researchers [Gadala et al. 2019]) identified, among others, some risks associated with: malicious use, unintentional misuse, and unauthorized access to the hospital system, suggesting a need for user authentication.

3.1 Aim

Both the qualitative analysis presented in this section and the probabilistic analysis presented in the next section aim, using the specific case study described above as an example, to:

- (1) Help device designers inform their decisions on whether or not to implement authentication for a certain medical device, and if so, to identify an appropriate authentication scheme for their specific application scenario, by considering the multiple, direct and indirect effects of such decisions. Authentication scheme providers and proponents, perhaps subconsciously, may sometimes have an optimistic or incomplete view of potential risks.
 - (2) Inform the decision of device designers on the detailed implementation (especially the protocol of use) and the specific parameter values to use in their chosen authentication scheme.
 - (3) Capture the multi-faceted, rather than one-dimensional, nature of usability, safety, security and performance, as the situation may be much more complex than simply a linear trade-off between the system attributes.
 - (4) Recognize high-level patterns that might otherwise be missed.
- We aim to achieve these goals by (1) systematically describing the problem of authentication of the operator(s) of the medical device, (2) qualitatively analysing the problem from the combined

¹<https://cordis.europa.eu/project/id/737475>, <https://aquas-project.eu/>

²<https://aquas-project.eu/use-cases/#medical-devices>

viewpoints of safety, security, performance and usability, and finally (3) describing potential design alternatives and the key design parameters that designers can tune, so that risk levels resulting from each potential design can be estimated and/or compared.

The long-term scientific value of our contribution will lie not as much in the raw data we present, but more in the methodology which we propose. In essence, bringing a team of experts to a shared understanding of the different angles of the problem is more valuable than any specific ranking of authentication schemes or a specific choice of scheme.

3.2 Method

To decide the most appropriate authentication method and the best design options, we need to start by clarifying design trade-offs, and describing the risk associated with each design solution, so that designers can choose on a rational basis. Analyses can be divided into two main stages: a qualitative analysis (described here in Section 3) followed by a probabilistic analysis (described in Section 4). The descriptive analysis helps to decide which attributes of the socio-technical system matter, and clearly define them. For example, which specific aspects of authentication performance need to be considered (average time to authenticate, probability of failure on first attempt, etc.)? Equally important is to identify minimum requirements for each of these aspects, which allows a pre-selection/elimination of methods. For example, password-based methods were excluded from further analyses as they were considered too slow and disruptive. These steps help the designer to compare basic designs and exclude those that are inferior from all the important viewpoints. This leaves a shortlist of solutions of which none dominates the others: in comparing any pair, each has advantages and disadvantages.

At this point, the probabilistic analysis comes into play: it supports comparing the most viable solutions in the shortlist from the viewpoint of the overall risk they present (in various operation and threat scenarios), or performance penalty they impose; and calculating the effects of tunable design parameters of each solution, so that the designer can select appropriate or optimal values.

We begin the descriptive analysis by qualitatively describing various authentication methods (knowledge, token, and biometric-based) from the different viewpoints (security, performance, usability, cost, and safety). One of the solutions compared is "no authentication" as there is no a priori certainty that authentication will reduce overall risk. In particular, apart from the unintended effects mentioned in Section 2, an authentication mechanism is an attractive target for denial of service (DoS) attacks, meant to make the critical device unavailable to its operator. Each authentication method was described against each of the system attributes based on discussions between security, usability and stakeholder viewpoints. These descriptive analyses were then used to reduce the list of potential schemes to be further refined by the quantitative analysis.

3.3 Preliminary Results

The descriptive analysis revealed interesting trade-offs between cost, security, safety, performance, and usability. The analysis, represented in a table, helped capture the complexity of the issue. For

instance, it highlighted advantages of biometric methods compared to alternatives, in terms of: (i) convenience (no need for users to remember or carry anything special), and (ii) reduction of certain security risks, especially malicious use of a stolen token. However, this is not without disadvantages: the high reject rates in some biometric methods bring serious safety concerns related to timely patient care, and may cause user frustration and loss of focus. The descriptive analysis revealed an intricate web of factors affecting the authentication decision. To clarify their relationships and describe the risk associated with each authentication solution, we created a dependency diagram (Fig. 1) that represents an incomplete yet complex account of how patient harm can be caused.

Of utmost interest in the trade-off analysis is to describe the risk associated with each alternative solution. In the dependency diagram, this is mapped as all paths leading to the "patient harm" node, including, for instance, scenarios where: a pump administers an incorrect dosage, or a clinician fails to respond to an alarm from the device. We note in the diagram four input arrows to "patient harm". These represent four main conjectures about how authentication issues may trigger patient harm in this use case:

- **Malicious Use.** An attacker uses access to the device to (a) directly alter the target parameters or dosage, or (b) tamper with the alarm functions and the clinician's adaptation to them (e.g. by producing excessive false alarms for patient conditions, or omitting alarms for exhaustion of drug supply in the pump), or (c) deny service by refusing access to legitimate users.
- **Accidental Use.** Consider, for example, a higher-privilege user close enough to the device to trigger automatic RFID authentication, so that a slip by a lower-privilege user in entering a potentially harmful command is accepted by the device.
- **False Reject.** The authentication method falsely rejects a legitimate user, thus preventing them from assisting the patient when needed. For example, this could be due to a card reader failure.
- **User Concentration.** The authentication method harms users' concentration so as to distract them or hamper their response to emergencies. This highlights the important area of second-order effects of adding a feature, which risk being neglected when focusing only on the intended effect of the feature. A similar example is alarms, meant to improve response to danger, but sometimes causing distraction, or introducing, "automation-induced errors" (or "automation bias") that can even result in worse overall performance.

Besides describing different ways that patient harm may occur, the dependency diagram can reveal unexpected links. For instance, an organization may enforce stricter protocols to reduce the risk of malicious use due to theft of the key (i.e., "Stricter protocols" mitigation node); but various arrows leaving this mitigation strategy lead to increased "P(Forgetting the key)", increased "P(Sharing the key)", decreased "ease of use", etc. – all of which can lead to patient harm through a different chain.

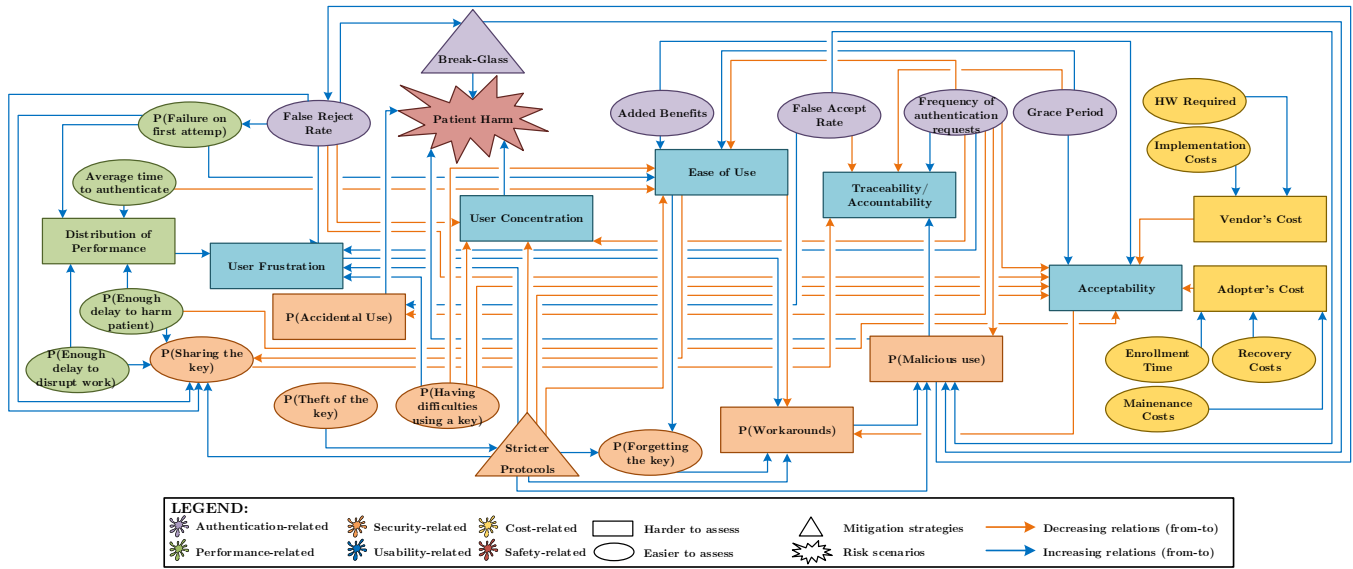


Figure 1: Relationship between Various Factors in the Decision to Implement Authentication.

4 A PROBABILISTIC ANALYSIS OF THE AUTHENTICATION ISSUE

The qualitative analysis described in Section 3 allowed an initial comparison between authentication solutions, but it highlights the complex dependencies that make necessary a risk-based, quantitative comparison. In theory, a perfect quantitative risk analysis would turn the design problem into a numerical constrained optimization problem, minimizing a risk variable by tuning certain design parameters. In practice, it will at least allow the designer to identify large differences in overall risk levels, check the sensitivity of risk levels to assumptions made and to specific parameters of the design and environment, and identify acceptable solutions or ranges of solutions.

Furthermore, once an appropriate authentication method and protocol is chosen, a quantitative analysis gives a reliable method for choosing appropriate values or ranges for key design parameters such as (our example presented later in this section) authentication expiry time: how long a user is allowed to operate the device after successfully authenticating him/herself.

To achieve these goals, we build a probabilistic model, and the dependency diagram in Fig. 1 suggests important factors to model and output variables to measure.

4.1 Scope

The qualitative analysis helped to narrow the scope of the problem and limit choices to the most viable solutions: (1) No authentication, (2) Authentication by a smart card that needs to be tapped (RFID device) on the authentication device, (3) Continuous authentication using a RFID card or voice that authenticates users within the range of a few meters.

Also important to define is the type of attacks considered. Many potential attacks can be imagined for this specific case study, including “man in the middle” attacks against the registration process,

server attacks, denial of service attacks, remote sniffing, jamming attacks, identity fraud, among others. We start by studying a specific type of attack (we intend later to extend the model with other attack types): an attacker waits for a moment when the clinician is not engaged using the device, approaches the device, and proceeds to perform some task on the device that harms the patient. This specific attack type encompasses attackers who have illegitimately obtained an authentication ‘token’ (e.g. smartcard), or attackers without a token who wait for a chance when the clinician is not engaged with the device, but authorization has not expired, in order to give their malicious command.

The model concerns use of the device in an Intensive Care Unit (ICU) and not in an Operating Room, because discussions with device designers determined that (1) the expected level of physical security in an Operating Room reduces the probability of external visitors and thus most attacks considered, and (2) users of the device during an operation are not likely to change, as is the case instead in the ICU, where a patient may be present through multiple staff shifts.

4.2 Method

Our probabilistic model uses the formalism of “stochastic activity networks” (SANs) using the software tool Mobius, developed and maintained by the University of Illinois³.

- **The Model.** Two of the main sections of the model concern modelling of the tasks and the authentication device.
- **The Task.** In the model, tasks arise that a clinician must perform. Tasks represent any actions using the device, or changes to the device, that are deemed critical and thus require authentication. Only when a user becomes authorized can they start work on a task. The user is allowed to continue performing the task so long as they stay authorized.

³<https://www.perform.illinois.edu/>

Depending on the duration of the task and the authentication expiry time, the user may become unauthorized while performing the task. Also, a new task cannot start until the clinician becomes available to initiate it (i.e., the clinician may be busy attending to other tasks or be in a different room when the need for a new task arrives).

- **The Authentication Device.** This part of the model captures the authentication process. It includes detailed design choices like two different setups for the authentication expiry time: either as soon as a task is completed the device returns to a non-authorized state, or the device only returns to a non-authorized state when the authentication expiry time, a design parameter, ends – independent of the completion of a task. In the intended environment the latter arrangement was deemed more realistic/desirable by designers, as it allows clinicians to perform more than one task without the need to re-authenticate, and as it avoids the need to implement a way to detect the completion of a task in order to trigger the non-authorized state.
- **Inputs.** In the model, we can manipulate several variables representing: (i) key parameters in the control of the designer, such as authentication expiry time; (ii) characteristics of the authentication method; for example, the false rejection rate (probability that the authentication method or hardware will falsely reject a legitimate user), or the time needed for a user to authenticate using a specific authentication method; (iii) variables on which the designer has little or no control, as they relate to the environment in which the device is deployed; for example, duration of a user task, time elapsing from when a task becomes necessary to when a clinician becomes available to start work on it, or time limit for the clinician to complete the task before the delay may cause harm to the patient; and finally, (iv) a complex set of variables that attempts to model certain user behaviors, such as how many times a user will retry to authenticate before giving up.
- **Outputs.** Some of the interesting outputs computed using the model include: the ratio of tasks completed by staff to needed tasks, the number of authorizations needed per task (ideally 1, but re-authentication difficulties can make the average greater than 1), the fraction of tasks that result in patient harm, and the number of malicious tasks successfully completed during a specific duration.

4.3 Example Results

The model outputs are valuable information about the influence of different protocols of use on safety, security, and usability. Designers can compare different authentication methods/vendors, and identify desirable values for key design parameters. For example, in Fig. 2 we see the effect of the false rejection rates of different authentication methods on patient harm, a key measurable outcome of the reliability of the authentication solution. The graph shows a direct effect of the reliability of the security method on patient safety, which is not surprising, but it also quantifies it.

A more complex result is one where to choose the value of a design parameter it is necessary first to analyse a trade-off between




Figure 2: Illustrative results of the modelling: effect of varying the False Rejection Rate on frequencies (as a fraction of the total number of tasks started) of failed tasks and tasks that cause patient harm.

its effects. For example, in Fig. 3 we can see that by increasing the value of a key design parameter, authentication expiry time, we can decrease the safety risk caused by tasks that fail to complete due to failed authentications; but at the same time we increase the security risk caused by attackers taking advantage of the window between when a clinician moves away from the device and when the authentication actually expires. In the graph, we thus see a range of acceptable values of the authentication expiry time, such that the overall patient harm caused by both the safety and security risks is close to a minimum. These curves are calculated, for the sake of illustration, using plausible values of the various model parameters. For concrete decisions about these settings, one would replace these parameter values with values measured in the specific type of hospital environment of interest. For some parameters, which may have even significant effects on results, such as the expected rate of attacks in this specific environment, it is quite possible there would only be rough conjectures. The model will show how this uncertainty affects the risk associated with a choice of authentication expiry time value.

5 DISCUSSION

The analyses described are steps towards helping designers to make well-informed decisions about user authentication. They highlight an important, generalizable point: by introducing security measures, designers may end up introducing new safety risks (e.g. unreasonable delays that pose harm to end users) and security risks (e.g. denial of service), sometimes increasing the very risks they were designed to mitigate. Hence, analysts and designers must:

- **Take a holistic approach.** The analyses presented exemplify how (i) security controls meant to preserve safety of operation conflict with safety and operation performance




Figure 3: Illustrative results of the modelling: effects of varying Authentication Expiry Time on potential safety and security risks: patient harm occurrences due to attacks or due to the authentication mechanism intended to protect against them.

goals and (ii) users' attempts to preserve safety and performance in spite of security controls may impair security. Some risks may even go unnoticed in the grey areas between specialisms: e.g., a safety analyst may conclude that a device has sufficient availability, and risk of misuse is eliminated by authentication; a security analyst, while aware of the risk of a device being disabled by a DoS attack on authentication, may only be asked to assess how well authentication prevents misuses, its original motivation. Standards and guidelines also need to consider the need for a holistic approach to address trade-offs, especially since different attributes (security, usability, etc.) tend to be covered in separate documents. The holistic approach needs to consider indirect negative effects which are often neglected, since a solution considered least risky on the basis of its direct effects may actually be riskier than some alternative. For example, in the dependency diagram, we note how patient harm may be caused not just by direct, expected causes such as malicious use, but also through reduced user concentration, triggered by factors such as frequency of authentication requests, and by well-meaning mitigation attempts such as stricter authentication policies.

- **Beware of introducing mitigation strategies "in isolation" without considering their effects from all viewpoints.** These effects may introduce new hazards. For instance, requesting authentication more frequently may seem logical, to reduce the probability of malicious use. However, its direct connection to user frustration can activate links leading to patient harm. Likewise, multi-factor authentication may reduce direct security risks, but the inconvenience of an extra authentication level may lead users to invent

workarounds (a classical example: hanging an authentication badge on the device), thus jeopardizing safety and security.

- **Clearly identify the attributes to be considered in the decision, at an early stage.** For example, further dividing "usability" into: ease of use, user concentration, user frustration, traceability/accountability, etc. This helps capture the multi-dimensionality of the attributes, without which the analysis will be simplistic. For example, a longer authentication expiry time does not affect usability in a single direction: it improves acceptability and ease of use but decreases accountability/traceability.
- **Use probabilistic analysis of the authentication issue.** Designers may be loath to use probabilistic assessment of the authentication problem, through lack of familiarity and because there are so many uncertainties about parameters and even the model itself. Despite such uncertainties, the modelling approach promises to help device designers quantify risk in a way that informs rational decisions. Even in the simple examples shown here, the trade-off we presented between risk due to accidental causes (unreliability of authentication) and malicious causes (intruders operating the device) cannot be decided rationally unless one quantifies the overall effect on risk. This quantification of course depends on assumptions (for instance, how often intruders will enter the ICU) but a rational decision must acknowledge these assumptions, selecting plausible ones or verifying that the results are robust to assumption errors. When we move to more complete descriptions of the possible behavior of users and attackers, the overall effects on risk can no longer be estimated intuitively, and computer-supported modelling as we have outlined here becomes essential for insight on the direction and magnitude of the effects of variations in the design and environment.

6 CONCLUSIONS AND FUTURE WORK

This work is one step towards helping medical device designers make rational decisions about authentication of operators of critical devices by: representing the design problem systematically, capturing its complexity, and analyzing it from different viewpoints. This preliminary work – albeit limited to a single device, rather than sets of devices – offers a few useful insights on: the complexity of the issues that may be hidden by the simple requirement for authentication that is both secure and usable; the usefulness of structured analysis methods; the idea that some of the trade-offs can be resolved with reasonably simple quantitative calculations, allowing for ranges of uncertainty about the parameters.

Further work on the quantitative analysis has to ensure that the model captures, as closely as feasible, the main real-life phenomena in the socio-technical system. Apart from modelling various types of attacks, the main next step we see now is to add descriptions of the unintended effects on user behavior, and especially user workarounds.

We have modelled a single attack type; a complete study needs to take into account other, main realistic attack types. The overall risk of each option will change based on prevalent attack modes, whose probabilities are difficult to estimate, vary depending on

users and use environments, and are dynamic - changing over time even within the same organization. Scenarios with different dominant attacks, e.g., attempts to input harmful commands to specific patients versus attempts to harm patients at random via DoS, will result in different optimal solutions. So, designers need to consider whether solutions are robust over the range of such assumed threat environments; or, consider tunable authentication options, which bring their additional security and safety concerns.

Finally, it would be desirable to package the model in a user-friendly and generalizable way that could be used by designers not just of medical devices, but a wide category of similar critical devices to conveniently inform decisions on the issue of authentication. For these useful analysis methods to be successful in industry, practitioners should not need to build models from scratch in a specialized mathematical language but to configure a general-purpose model described in familiar terms.

In this work we found that to assist practitioners in the choice of authentication methods it is necessary to have a good practical checklist of aspects to check, so that authentication products can be easily characterized according to the various important criteria. We are thus working on an analysis that extends the comparative framework approach of authors cited above (section 2.4) to medical or, more generally, safety-critical, real-time devices.

ACKNOWLEDGMENTS

This work was supported in part by the ECSEL Joint Undertaking under grant agreement No 737475, project AQUAS (Aggregated Quality Assurance for Systems), and by the Technology Agency of the Czech Republic, project No. FW01010200. The authors wish to thank the AQUAS partners who contributed to the medical “use case” in the project, from which this study originated, and especially RGB Medical Devices and Prof. Ricardo Ruiz who led that work.

REFERENCES

- Chon Abraham, Dave Chatterjee, and Ronald R Sims. 2019. Muddling through cybersecurity: Insights from the US healthcare industry. *Business horizons* 62, 4 (2019), 539–548.
- Kemal Altinkemer and Tawei Wang. 2011. Cost and benefit analysis of authentication systems. *Decision Support Systems* 51, 3 (2011), 394–404.
- Joseph Bonneau, Cormac Herley, Paul C Van Oorschot, and Frank Stajano. 2012. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *2012 IEEE Symposium on Security and Privacy*. IEEE, IEEE, San Francisco, CA, USA, 553–567.
- Christina Braz and Jean-Marc Robert. 2006. Security and usability: the case of the user authentication methods. In *Proceedings of the 18th Conference on l'Interaction Homme-Machine*. ACM, Montreal, Quebec, Canada, 199–203. <https://doi.org/10.1145/1132736.1132768>
- Center for Devices and Radiological Health. 2014. Content of Premarket Submissions for Management of Cybersecurity in Medical Devices, Guidance for Industry and FDA Staff. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/content-premarket-submissions-management-cybersecurity-medical-devices>
- Lorrie Faith Cranor and Norbou Buchler. 2014. Better together: Usability and security go hand in hand. *IEEE Security & Privacy* 12, 6 (2014), 89–93.
- Charlott Eliasson, Markus Fiedler, and Ivar Jørstad. 2009. A criteria-based evaluation framework for authentication schemes in IMS. In *2009 International Conference on Availability, Reliability and Security*. IEEE, IEEE, Fukuoka, Japan, 865–869.
- R Francis. 2017. Hospital devices left vulnerable, leave patients at risk. *Retrieved on June 30 (2017)*, 2019.
- Marwa Gadala, Lorenzo Strigini, et al. 2019. *Combined Safety, Security and Performance Analysis and Assessment Techniques - Preliminary version*. AQUAS Project Deliverable No. D3.2. http://aquas-project.eu/wp-content/uploads/2019/11/D3.2_CombinedAnalysisPreliminary_v2_obis.pdf
- Paul Grassi, James Fenton, Elaine Newton, Ray Perlner, Andrew Regenscheid, William Burr, Justin Richer, Naomi Lefkowitz, Jamie Danker, Yee-Yin Choong, Kristen Greene, and Mary Theofanos. 2020. Digital Identity Guidelines: Authentication and Lifecycle Management [includes updates as of 03-02- 2020]. <https://doi.org/10.6028/NIST.SP.800-63b>
- Michele D Guel. 2002. A framework for choosing your next generation authentication/authorization system. *Information Security Technical Report* 1, 7 (2002), 63–78.
- Kimmo Halunen, Juha Häikiö, and Visa Vallivaara. 2017. Evaluation of user authentication methods in the gadget-free world. *Pervasive and Mobile Computing* 40 (2017), 220–241.
- Ruslan Horblyuk, Kristopher Kaneta, Gary L McMillen, Christopher Mullins, Thomas M O'Brien, and Ankita Roy. 2012. Out of Control: How clinical asset proliferation and low utilization are draining healthcare budgets. *GE Healthcare* 19, 6 (2012), 64–68.
- IEC. 2016. *IEC 61882, Hazard and operability studies (HAZOP Studies)-Application guide*. <https://webstore.iec.ch/publication/24314>
- Sven Kiljan, Harald Vranken, and Marko van Eekelen. 2018. Evaluation of transaction authentication methods for online banking. *Future Generation Computer Systems* 80 (2018), 430–447.
- Thomas M Kroeger, Sean Peisert, and Edward B Talbot. 2013. *Laws of Authentication*. Technical Report. Sandia National Lab.(SNL-CA), Livermore, CA (United States).
- Kat Krol, Eleni Philippou, Emiliano De Cristofaro, and M Angela Sasse. 2015. "They brought in the horrible key ring thing!" Analysing the Usability of Two-Factor Authentication in UK Online Banking. In *Proceedings of the 2015 Network and Distributed System Security (NDSS) Symposium*. Internet Society, San Diego, California, 1–10. <https://doi.org/10.48550/arXiv.1501.04434>
- Jens Mehrfeld. 2019. Cyber security and Industry 4.0. *at-Automatisierungstechnik* 67, 5 (2019), 361–363.
- Anthony J Palmer. 2010. Approach for selecting the most suitable Automated Personal Identification Mechanism (ASMSA). *computers & security* 29, 7 (2010), 785–806.
- J Schwartz, B Haarbrandt, D Fortmeier, R Haux, and C Seidel. 2014. Authentication systems for securing clinical documentation workflows. *Methods of information in medicine* 53, 01 (2014), 3–13.
- Ignacio Velásquez, Angélica Caro, and Alfonso Rodríguez. 2018. Authentication schemes and methods: A systematic literature review. *Information and Software Technology* 94 (2018), 30–37.
- Leah Zhang-Kennedy, Sonia Chiasson, and Paul van Oorschot. 2016. Revisiting password rules: facilitating human management of passwords. In *2016 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Toronto, ON, Canada, 1–10. <https://doi.org/10.1109/ECRIME.2016.7487945>
- Verena Zimmermann¹, Nina Gerber, Marius Kleboth¹, Alexandra von Preuschen¹, Konstantin Schmidt¹, and Peter Mayer. 2018. The quest to replace passwords revisited—rating authentication schemes. In *Proceedings of the Twelfth International Symposium on Human Aspects of Information Security & Assurance (HAISA 2018)*. IFIP, IFIP, Dundee, Scotland, 38.