Low Probability, High Stakes: A look at PKI

Alexander Lechner
Betreuer: Dipl.-Inform. Ralph Holz
Seminar Future Internet WS 2011/12
Lehrstuhl Netzarchitekturen und Netzdienste
Fakultät für Informatik, Technische Universität München
Email: lechneal@in.tum.de

ABSTRACT

This paper examines the risk assessment of low probability but extremely high stake events. When evaluating such risks, as for megaprojects or natural disasters, it is most often assumed that the underlying arguments are correct. Especially for events with very low probabilities this naïve approach is not sufficient as the theory and the applied model may be unsound and the calculations may be flawed. New and more complex approaches try to model the uncertainty of unsound arguments in order to provide a more accurate risk estimation of low probability events. One such approach, recently published by Ord *et al.*, is discussed in detail and is further applied to the security of a Public Key Infrastructure (PKI) for which an abstract and generic risk estimation model is developed. The analysis of the model shows that the impact of possible flaws in the argument of a PKI can be vast and must therefore not be neglected. Using several examples of disasters, the underestimation of risks connected to trivial error sources and the necessity of an advanced risk assessment methodology is emphasized.

Keywords

probability theory, risk assessment, security, high stake risks

1. INTRODUCTION

It is often required to assess the risk of highly unlikely events as they may hold extremely high stakes. Examples for such events could be cosmic incidents such as a meteorite colliding with earth, terrorist attacks causing thousands of deaths or success and failure of megaprojects where billions of dollars have been invested.

Based on the low probability of such events it is insufficient to use simple heuristics or the estimated likelihood of the occurrence in order to calculate the risk. It may be more probable that the argument is flawed than that the event actually occurs, which may lead to an underestimation of risk and consequently grave consequences. It is therefore necessary to apply a more sophisticated approach to such risk assessments.

This paper focuses on an advanced risk assessment of low probability high stake events, with an emphasis on the security of PKIs. In section 2 the approach proposed by Ord *et al.* [20] is explained and it is shown how the calculation of the probability can be improved by including the possibility of a flawed argument and by subdividing the argument into the three parts: theory, model and calculation. Further, several considerations for risk assessment are stated and an example of a well-done risk report is given.

Section 3 shows the application of the advanced risk assessment method to the security of Public Key Infrastructures (PKI). A PKI is a well established technique to ensure privacy and identification on the Internet. As such PKIs are widely used by governments, companies and individuals, a failure of such a system bears massive problems [1].

Some examples of well-known disasters in the history of mankind are shown in section 4. The failures are explained and examined with respect to the underlying theory, model and calculation. As many of those failures could have been avoided the necessity of advanced risk assessments becomes obvious.

Finally, section 5 offers a comparison with other approaches and related work, including human risk perception and decision making behaviour.

2. ADVANCED RISK ASSESSMENT

Let us assume the following: A risk report has to be written, assessing the risk of a devastating earthquake taking place in a certain area which has been earthquake-free for centuries. To calculate the probability of an earthquake occurring, a model has to be developed and assessed. Such a model is based on an underlying assumption which may include geodetic and seismologic information as well as geological theories on how earthquakes arise. As earthquakes are extremely rare in the given area, it is highly probable that the result of the model will state a low probability.

Let us use the following notation: "X" the earthquake occurs, "A" the underlying argument is sound and "L" the expected loss in case of the accident. The mathematical formula for calculating the risk is:

$$Risk = P(X) * L \tag{1}$$

A naïve approach would erroneously consider the outcome of the model to be $P(X)$ but in fact the result is $P(X|A)$, as the model can only calculate the probability of X given the underlying argument is valid. By using the naïve approach, the report completely ignores the possibility of the assumption not being sound and may therefore provide a false result. The probability of the argument being flawed may be actually higher than the probability of the earthquake arising.

doi: 10.2313/NET-2012-04-1_2

2.1 Ord et al.'s Approach

The assumption that an argument is entirely correct should never be made, as design and calculation errors are widely spread: section 4 provides several examples of disastrous failures. To address the problem of flawed arguments in risk assessment, Ord *et al.* [20] proposed a more sophisticated approach. Let "X" be the event happening and "A" the assumption being sound, then the probability of X is calculated as follows:

$$P(X) = P(X|A)P(A) + P(X|\neg A)P(\neg A) \tag{2}$$

The naïve approach only considers the first term: $P(X) = P(X|A)$, whereas Ord's approach also considers the case of the argument being unsound. If a naïve risk report states that X is impossible, it only means that $P(X|A) = 0$, yet the error term $P(X|\neg A)P(\neg A)$ may be bigger than zero, leading to a $P(X) > 0$ which means that X is possible after all. It may be impossible to acquire accurate values for $P(\neg A)$ and $P(X|\neg A)$, but coarse estimations can already lead to a change in our risk perception.
Ord's formula can be applied to all risk estimations, but it is of larger importance for low probability events. Having a large $P(X|A)$, the additional error term $P(X|\neg A)P(\neg A)$ is only of little significance in the overall probability estimation. However, for an extremely small $P(X|A)$, the error term may change the resulting $P(X)$ in several orders of magnitude.

A common way of assessing the soundness of an argument is to distinguish between model and parameter uncertainty. Considering the possibility of failures in the background theory and in the calculation, Ord *et al.* [20] proposes a distinction between the argument's theory, model and calculation. Using this distinction, $P(A)$ can be estimated in a more precise way. For calculations their correctness is considered, for theories and models their adequacy. A theory is adequate if it is able to predict the required qualities with a certain precision. For instance Newton's mechanics is an adequate theory for aerodynamic or static problems but is surely inadequate for computations on a quantum mechanical level.
The distinction between theory, model and calculation results in an extended formula. Let "T" be the background theories are adequate, "M" the derived model is adequate and "C" the calculations are correct. As the model is dependent on the theory and the calculations and theories are independent, we get:

$$P(A) = P(T)P(M|T)P(C) \tag{3}$$

The estimation of the individual terms can facilitate the calculation of $P(A)$, but is still of significant difficulty. Below the three parts are separately described and common error sources are depicted.

2.1.1 Theory

The term "theory" used by Ord *et al.* refers to the argument's theoretical background knowledge. This does not only include established and mathematically elaborated principles, but also specific paradigms or best practices. Theories are defined on a high abstraction level, so they are associated with a higher confidence than models.
As many theories are well-established and underlie various models, a flaw in such theories may have an immense impact on the corresponding scientific area and render numerous models and results, at least partially, invalid. Especially in science fields where proofs and evidences are difficult to provide, such as sociology and psychology, or where empirical analysis is unethical, for instance the lethal dose of ionizing radiation, the theories are based on assumptions, extrapolations or even speculations. History shows that many, even well-established theories were flawed. Examples include the phlogiston theory which was used to explain the oxidation processes or Newton's corpuscular theory of light, in which he stated that light consists of small discrete particles. Another example of a flawed theory is the geocentrism: A theory which was assumed to be correct for nearly 2000 years and was the cause for Ptolemy inferring the Epicyclic model and Brahe the Tychonian system.

2.1.2 Model

A model is derived from the background theory and it models the examined phenomena. Various aspects influence the adequacy and the accuracy of such models. As a model is always an abstraction of reality, certain aspects are neglected whereas others are simplified. Too broad or too narrow models may subsequently lead to incorrect predictions. Further problems arise from *parameter uncertainty*, as the model has to be fed with several parameters in order to make a prediction. Those parameters can be imprecise or estimations by themselves. Furthermore, parameters are not always measurable, for instance if human values like trust or sincerity are part of the model. Finally, many phenomena are not well understood and the corresponding models are based on assumptions and approximations. A way to improve risk assessment for such ill-defined phenomena is to combine several models and theories.
The design and development of the model is a highly delicate procedure. Even if the theory and the formulas are correct, an imprecise parameter or an aspect neglected by the model may produce a significant deviation from the correct result. A small model modification, the literal *flap of a butterfly's wings*, can have a large impact on the outcome.

2.1.3 Calculation

The calculations are independent of the argument's theory and model. Still, calculation errors are more common than expected and can lead to the complete failure of the argument. Many different errors are made: accidental errors like forgetting a certain term, collaboration and communication errors when several teams work on the same problem, numerical errors as discretization and floating point errors.
A flawed calculation may have different consequences: There is the possibility that the error has only little effect on the result e.g. a slightly higher inaccuracy. An example of such a case is given in section 4: A bug in a Pentium processor series caused imprecise results for certain input values. Although this may not affect most customers, it may have consequences for scientific simulations and high performance computing. On the other hand it is also possible that a flawed calculation has a large impact on the whole model. For instance, a flaw in the implementation of a cryptographic algorithm may render a whole system insecure. In history, various accidents showed the impact of calculation errors. One prominent example is the Mars Climate Orbiter, which was lost due to a navigation error caused by a trivial miscalculation. The NASA worked with the metric system whereas the navigation software used the imperial system [6].

doi: 10.2313/NET-2012-04-1_2

Despite the fact that calculations are a frequent source of error, it is hard to provide reliable statistics of error rates. Reasons include that calculation errors can be caused by various factors such as hardware failure or human negligence.

2.2 Considerations

Ord *et al.* provide an example of a plausible risk analysis which was performed by Giddings and Mangano in 2008 [8] and uses multiple sub-arguments: will the LHC cause black holes, which subsequently will absorb earth? This is a perfect example of a low probability but extremely high stake risk. To strengthen the argument Giddings and Mangano used 3 different sub-arguments:

- A_1: As a consequence of different physical theories, black hole decay is highly probable.
- A_2: If a non-decaying black hole will be created, it most probably will not be able to discharge itself, which is necessary in order to be harmful.
- A_3: If discharging black holes can be created, the time that is required to consume Earth would still be shorter than Earth's lifespan. This is valid under several assumptions on how black holes interact with matter.

One special problem of this example is that the underlying theories are highly theoretical and difficult to verify. Still, the combined argument consisting of A_1, A_2 and A_3 is significantly stronger than the sub-arguments themselves, as one argument comes into play if the previous one fails. Consequently, the combined error term $P(\neg A_1, \neg A_2, \neg A_3)$ is smaller than the error terms themselves: $P(\neg A_1)$, $P(\neg A_2)$ and $P(\neg A_3)$. Giddings and Mangano state that our current understanding of physics suggests black hole decay A_1 and the inability to discharge themselves A_2 as highly probable. Regarding the uncertainty of the parameters and the assumptions made for the theories, they did not provide a certain probability value for X, instead they concluded that there is no significant risk of black holes.

3. PKI SURVEY

Public Key Infrastructures (PKI) were developed as a consequence of public key cryptography. With the usage of public key cryptography, two parties can confidentially exchange messages and provide authentication by using signatures. For such tasks, both parties have to possess a public encryption key and a private decryption key. Sending a secure message requires the user to encrypt the message with the recipient's public key. The message can then only be decrypted by using the recipient's private key. One of the main issues is to ensure that the other party is the one it claims to be.
A PKI provides a framework for authentication and for the secure exchange of information over an insecure network. Therefore PKI uses certificates issued by Certificate Authorities (CA), which bind public keys to the verified user identities. Such certificates contain at least the identity information, an expiration date and the user's public key. In a hierarchical PKI the certificate of a small CA (e.g. a company CA) can be recursively signed by a larger CA, finally signed by a root CA. So called Registration Authorities (RA) are

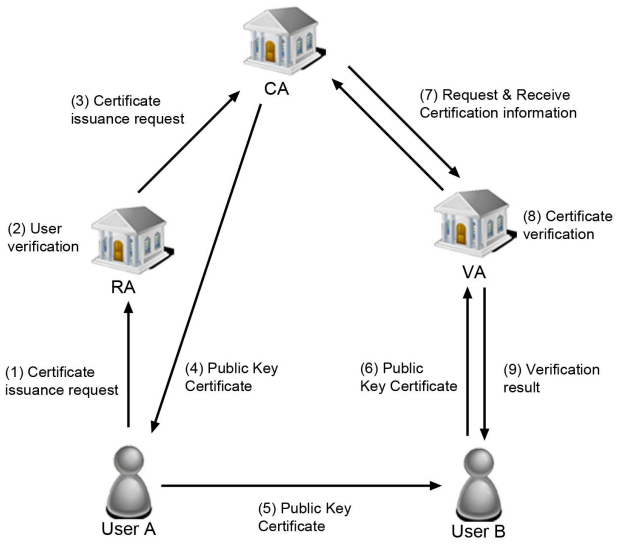

Figure 1: PKI: Certificate issuance and verification.

responsible for the verification of the user identities. The validation of the certificates, namely if a given public key corresponds to the intended user, is performed by Validation Authorities (VA). The issuance and validation procedure can be seen in figure 1. First, user A applies for a new certificate. The RA verifies the user's identity and forwards the request to a CA, which issues the certificate. User A can now proof his identity using the received certificate. Let us now assume that user B wants to check if user A is the person he claims to be. Therefore, user B forwards user A's certificate to a VA, which most often corresponds to the used software e.g. a web browser. The VA uses additional information from the CA, for instance if the certificate is still valid. If the validation process is successful, user B can be sure that user A is the person declared in the certificate. In the following analysis we focus on the X.509 standard, one of the most popular standards used for PKIs. This ITU-T standard specifies the standard formats of certificates, revocation lists and a certification path validation algorithm [4]. There are other approaches which are not covered by this survey. One example is the Web of Trust (WoT) where users express their trust by signing each others certificates and which has the popular implementation PGP.
As a high level of trust is required in many application areas, PKIs must provide an extremely strong security. In the following an advanced probability estimating model for the security of PKI is presented.

3.1 Need for Risk Assessment

Many companies, Web services and individuals rely on PKIs nowadays. Several PKIs exist and application areas include online banking, secured communication and access management. But as PKIs and public key cryptography became widely used, various issues emerged. Such issues range from protocol attacks to the question of confidence in the involved authorities. The end of PKI was often predicted, as it has been criticized by several quarters. However, PKI is not

doi: 10.2313/NET-2012-04-1_2

dead and the criticism helped to evolve the technology [1].
Assuming an ideal world, a highly simplified model and ignoring its fallibility, PKI can be considered highly secure. In fact, in our real world those assumptions do not apply. Perfect cryptographic algorithms and protocols do not exist, governments and politics may influence the trustfulness of CAs and infrastructure components may malfunction.
Ellison and Schneier [7] state 10 major risks of PKIs including "How secure is the verifying computer?", "Is the CA an authority?" and "Is the user part of the security design?". Other issues are mentioned by Adams and Just [1], who describe the evolution and history of PKIs and Boldyreva *et al.*[3], who analyze the security and efficiency of PKIs and provide considerations for building encryption and signature schemes. Besides the numerous cryptographic and computational security issues, the trust in, and the reliability of authorities plays a major role. Authorities do not always deploy PKI correctly. A recent report at the Defcon 18 [21] showed several problems concerning CAs: Certain CAs issued certificates with weak keys, they did not revoke invalid certificates and they signed unqualified names like "localhost" multiple times. Further problems arise from the usage of certificate revocation lists (CRL). If a certificate is no longer valid as it was lost or stolen, the CA has to revoke it. Best practices require to check such CRLs when validating a certificate. However, this practice is often ignored, as the recent DigiNotar hack showed [18]. A poor deployment of PKI results in a higher number of security leaks, facilitating possible attacks by adversaries and substantially undermining the security of the entire infrastructure.
There are obviously many issues concerning PKI but what are the stakes? The stakes highly depend on the usage of PKI. Users may simply use a PKI to communicate with friends, not wanting anyone to be able to eavesdrop. For such a scenario the stakes would be relatively small. More sophisticated usages of a PKI could include classified government information or military communication. The stakes for such scenarios are obviously high. Attackers can use flaws in a PKI for espionage, online fraud or even identity theft. A recent example of a PKI related issue is the Dutch certificate authority DigiNotar, which detected and revoked multiple fraudulent certificates in 2011 [18]. Subsequently the main browsers removed DigiNotar from the list of trusted authorities. As a result the chain of trust for the certificates of the Dutch government's PKI "PKIoverheid", which used DigiNotar certificates for different government operations and agencies, was broken.
As the stakes, depending on the individual usage of PKI, can be extremely high a way of assessing the risk is required. Instead of assuming infallible theories and models and in order to handle the complexity of PKI, the approach introduced by Ord *et al.* [20] was used.

3.2 Probability Model

Let us use the following notation: "A", the argument is sound and "X", a PKI related operation is successfully attacked. To better illustrate the model, we focus on the issuance of a certificate as "X" in this section. The argument is then subdivided in theory, model and calculation.

Theory: The main theory used in a PKI and in cryptographic systems overall is that certain mathematical problems are intractable. Deduced from this theory it is assumed, that the algorithms used for the public key cryptography and protocols in a PKI are secure. However, the security of the public key cryptosystems cannot be guaranteed, as feasible solutions to one of those problems may simply not have been found yet. Furthermore, faster computers and new technologies may influence the infeasibility assumption: Quantum algorithms as for instance Shor's algorithm are able to efficiently solve the discrete logarithm problem.
As the mathematical problems depend on the used public key system, there can be large differences between their complexity and solvability. For instance the discrete logarithm problem on elliptic curves is much more difficult to solve than the factorization of large integers. Some of the most famous existing public key cryptosystems exploit such problems, as for instance RSA, ElGamal and Elliptic curve cryptography (ECC). RSA relies on the infeasibility of factoring large composite integers, ElGamal and ECC rely on the difficulty of computing the discrete logarithm for certain groups.
If the cryptosystem theory is flawed, the impact on $P(X)$ and consequently on PKI will be huge as the PKI can no longer provide the security of its services.

Model: In order to reduce complexity we further subdivided the model into the following parts: Cryptographic Security, Infrastructure Security and Authority Trust.
Cryptographic Security: The cryptographic security refers to the security of the used cryptosystems. As one theory underlies many different models respectively cryptographic algorithms, they differ in many characteristics. Depending on the used parameters, hash functions and (pseudo-) random number generators the security may vary strongly. Attacks exploiting weak conditions or certain circumstances are numerous. For instance plain RSA offers many security leaks which enabled several attacks, such as Coppersmith's attack and chosen-ciphertext attacks [5]. One example of a parameter related issue is the key length of private and public keys. As computing power constantly increases and brute force attacks get feasible, formerly secure keys become insecure. To address this problem, longer keys can be chosen, multiple encryption can be performed (Triple DES) or more secure methods than ECC can be used.
Infrastructure Security: Besides the security of the cryptographic algorithm, the infrastructure bears several additional risks and may render a given PKI insecure. With infrastructure security, the security of all computers, networks and protocols involved in PKI operations as well as their interaction is meant. Especially the interaction between the different components may cause several dangers: as Boldyreva *et al.* [3] state, the mixing of proven cryptographic methods with key registration protocols does not make the system automatically secure. Other security issues concern the involved protocols and servers. Popular attacks, exploiting an insufficient infrastructure security include side channel attacks such as timing and man-in-the-middle attacks.
Authority Trust: One main aspect concerning the security of a PKI is the trust in the involved authorities. Three types of authorities exist: Certificate Authorities (CA) issue digital certificates, Validation Authorities (VA) verify given certificates and Registration Authorities (RA) identify and register certificate holders. The important question is: can an authority be trusted and how can this trust be ensured?

doi: 10.2313/NET-2012-04-1_2

Multiple risks mentioned by Ellison and Schneider [7] address this problem. Further problems arise as a consequence of CAs issuing certificates for other CAs, forming a chain of trust. This procedure requires a global root CA, but in reality there is no such authority. One existing solution is to trust multiple top level CAs, which can be seen as a Bridge CA. This corresponds to the "root store" of web browsers, which consists of a list of trusted CAs. But as there is no hierarchical structure and therefore no root CA the question of trust remains: *Quis custodiet ipsos custodes?* Who can watch the watchmen? Additionally, authorities may be influenced by local laws and governments and be forced to permit them access to their data and to cooperate. Another point is that authorities are no social services but profit oriented companies and therefore interested in a high number of certificates and users. As a result, CAs may not always be an authority on what the certificate contains and the strictness of the user identification and the PKI deployment may be too low [9].

Calculation: The calculation refers to the implementation of every piece of software involved in a certain PKI. This includes the implementation of cryptographic algorithms, protocols and end-user programs such as browsers or mail clients. All of those implementations hold the possibility of errors which can be used to successfully attack a PKI operation. Especially the interaction of different applications, coming from different sources, bears a high risk potential. Due to the mere amount of involved software and the far reaching consequences, it is required to take calculation errors into account: a flaw in the signing implementation may render the whole PKI insecure. Moreover, calculation errors are much more common than one may think. For instance the X.509 standard for PKI had and still has to deal with several implementation issues which can be exploited: Dan Kaminsky demonstrated at 26C3 how to include an unknown attribute into a certificate signing request by using illegal padding or integer overflows [12].

Using the subdivision into theory T, model M and calculation C the probability of the argument A can then be calculated using equation 3. In order to calculate $P(M|T)$, the model has to be adapted to the assessed PKI operation. Depending on the operation, different cryptographic algorithms, protocols and authorities have to be considered.

3.3 Probability Calculation

Having developed a PKI model, the probability of a successful attack $P(X|A)$, given that argument A is sound, can be calculated. But for the calculation of $P(X)$, the possibility of the argument being unsound $\neg A$ has to be taken into account as well. As explained in section 2, the probability of X can then be calculated by using equation: $P(X) = P(X|A)P(A) + P(X|\neg A)P(\neg A)$.
To gain accurate values for $P(A)$ and likewise $P(\neg A)$ is extremely difficult, regarding the argument's complexity and model's the level of detail. Instead, coarse estimations are already sufficient to improve the result significantly, as in the case of PKI there is always a fair chance that the argument is unsound. Possible sources of error in the argument include miscalculated parameters, under-/overestimations of the security and neglected but important factors. As the possibility of a flawed argument is obvious, a reasonable value for $P(X|\neg A)$ is required. This term is even more difficult to calculate, so a rough approximation has to be sufficient [20]. In order to propose a generic approach which is independent of the particular user performing a PKI action, human factors were excluded. If it is desired to calculate the probability of a successful attack depending on the individual, such factors must be taken into account. Examples of human factors influencing the individual security using a PKI are numerous: users may ignore invalid certificates for SSL connection while browsing the Internet, choose easy-to-remember but insecure passwords and, intentionally or unintentionally, not keep their passwords secret. The exploitation of psychology instead of technology can be a much easier strategy for a potential attacker.

3.4 Impact on PKI Security

For PKI and many other cryptographic systems it can be said that the security is only as strong as the weakest link. Despite of having correct cryptographic algorithms, an error in the registration process may still render the whole infrastructure insecure. Already for one single flaw in a PKI there is a relatively high probability that the security of the whole infrastructure can not be longer guaranteed.
In the case of PKI the probability $P(X|\neg A)$ may be much higher than $P(X|A)$, which can lead to far-reaching consequences. Therefore, it is not only recommended but moreover mandatory to consider the fallibility of the argument. As a result, Ord *et al.*'s method provides a much more reliable result compared to the naïve approach, which only considers $P(X|A)$.
In summary it can be said that the security and the risk assessment of a PKI is a complex task. The building of a perfect model can be considered impossible and incorrect models can lead to an extreme underestimation of the risk. As PKI is only a framework, the security should be improved by combining secure and well-proven components. Special attention must be paid to the interoperability of those components, which is a frequent source of error. Although the disproof of the PKI's theory would have the strongest impact on PKI security, a flaw in the model and in the calculations, a poor PKI deployment and the misbehavior of the involved users is by far more probable.

4. FAILURE EXAMPLES

In this section several examples of famous accidents are presented. It can be seen that flaws in design and calculation are wider spread than one may think and the value of the advanced approach by Ord *et al.* becomes clear. The disasters are often accompanied by poor and insufficient risk reports, together with problems in human risk estimation and probability evaluation.

4.1 Therac-25

From 1985 to 1987 the radiation therapy machine *Theriac-25* caused the death of 3 persons and severe injuries of a further 3. Several software errors caused a malfunction of the medical device, giving the patients massive overdoses of radiation. An investigative commission [14] stated, that the prime reason was not a certain software error, but a combination of bad software design and development practices. Those flaws included incorrect synchronization, missing quality reviews and unrealistic risk assessments. The

doi: 10.2313/NET-2012-04-1_2

analysis contained several independent assumptions, such as specific hardware errors with quite low probabilities, but completely ignored software and coding errors as an error source. Even after the first accidents the investigators failed to eliminate the root causes. Instead they continuously fixed individual software flaws which did not solve the main problem.
Splitting this failure up into theory, model and calculation errors it can be said that both model and calculations were flawed. On the calculation side several crucial coding errors were made. Even after the first accidents had happened, the coding errors were not found. This was the result of an overconfidence in software, the lack of independent software code review and no testing of the final software and hardware. On the model side, multiple design errors occurred during the development process. Self-checks, error-detection and error-handling were entirely missing. There was no feature for detecting massive overdoses, the patients reaction was the only indicator on the seriousness of the problem.
Besides the actual errors in design and software, human risk perception and handling was a major issue. Furthermore inadequate software engineering practices were used: software quality assurance practices were violated and a highly complicated design was used.

4.2 Mars Climate Orbiter

The Mars Climate Orbiter (MCO) and the Mars Polar Lander (MPL) were both part of the Mars Surveyor Program established by the NASA in 1994. The aim of the MCO was the study of the Martian climate and atmosphere processes and to act as a communication relay for the MPL. In 1999, after reaching a loose elliptic orbit, the MCO was planned to approach a final 400 km circular orbit. During this orbital insertion maneuver the communication was lost and could not be reestablished. The Mars Polar Lander was destroyed several months later in the course of the landing procedure. The disaster caused a loss of $327.6 million in total [11].
The Phase 1 report released by the Mars Climate Orbiter Mishap Investigation Board [22] concluded that the primary cause of the failure was human error. The navigation software of the MCO, developed by Lockheed Martin, used imperial units (pound-seconds) for the calculation of the momentum whereas the NASA used metric units (newton-seconds). Consequently, the effect on the spacecraft trajectory was underestimated by a factor of 4.45.
The error causing the destruction of the MPL could not be determined. The most likely causes include incorrect sensor feedback and a failure of the heat shield.
The destruction of the MCO can be seen as both a calculation and a model error. The main error was the usage of imperial instead of metric units in a software file. Still, Euler *et al.*[6] state that the error was caused by mistakes in judgment and poor application of standard practices. As an example he names the "Faster, Better, Cheaper" philosophy of the NASA at that time, when all non-value adding activities were eliminated and development steps were reduced.

4.3 Pentium FDIV Bug

The Pentium FDIV bug was a bug in the Intel P5 Pentium floating point unit which led to incorrect results of floating point divisions with certain input values. The error occurred extremely rarely and remained therefore undetected for 1.5 years. At that time, approximately 3-5 million copies of such flawed Pentium processors were in use. Later it came to light that Intel had already been aware of the flaw for several months. In the end, Intel was forced by public pressure to offer a replacement of all affected Pentium processors, resulting in an additional financial burden of $475 million [19].
The flaw was detected by Thomas Nicely [19] in 1994 and is a perfect example of a calculation error, namely a clerical mistake: 5 entries in the look-up tables used for the division algorithm contained corrupt values: a "0" instead of a "2". As a result the error only occurred when the corresponding table entries were used for the calculation. Intel states in its white paper that the worst possible inaccuracy occurs it the 4th significant decimal digit and the probability of a flawed divide instruction with two random input values is 1 in 9 billion [10]. One example of a flawed calculation is given by Nicely [19]:

$$\frac{4195835.0}{3145727.0} = 1.3338204491362410025 \text{ (Correct value)} \quad (4)$$

$$\frac{4195835.0}{3145727.0} = 1.3337390689020375894 \text{ (Flawed Pentium)} \quad (5)$$

Although the error occurs extremely rarely and has nearly no impact on an ordinary end-user, high-level applications which require a high precision may be severely influenced. Encountering the flaw in a scientific, financial or military application, where high precision is mandatory, can entail massive consequences. The stakes of a miscalculation in a financial simulation or a military device can be imagined by anyone.

5. RELATED WORK

In this section we present some of the related research literature which addresses the security of PKIs, risk assessment of low probability high stake events and human risk perception.
One model of assessing the trust in a PKI was developed by Marchesini and Smith [16]. Instead of assuming an ideal world, they adapted the model to the real, imperfect world. To do so, they had to handle many real-world PKI concepts, such as authorization, trust management and certificate revocation. By applying their calculus to several PKI systems they showed its ability to reason about real-world PKIs.
A document issued by the Department of Finance and Deregulation of the Australian Government provides a valuable overview of a PKI risk assessment methodology [2]. They arrange possible threats in 7 different categories, not only referring to infrastructure failures but also to the individual usage of certificates and to social engineering. They further provide a template for risk assessment: Each encountered risk is assigned a likelihood indicator, from "rare" to "almost certain" and a consequence indicator, ranging from "insignificant" to "catastrophic". By combining the likelihood with the impact, a risk analysis matrix is created, showing the significance of the risk. This is a highly simplified and subjective way of performing a risk assessment, but nonetheless valuable as many security leaks are considered and subsequently managed.
As risk analysis is vitally important for companies and scientific projects, it is crucial for risk assessments to consider the psychology of high stakes decision making and human risk perception. Individuals are influenced in many different

doi: 10.2313/NET-2012-04-1_2

ways by social and private factors. This is especially the case for low probability events where the necessary experience is missing. Kunreuther *et al.*[13] discusses low probability and high stake decision making. He states several problems in human risk assessment, like people who insufficiently use or even ignore available information. Further, the likelihood of an event happening is considered so low, that the event is neglected although the stakes are high. Another problem of human risk perception is the focus on short time horizons. Instead of taking the long term consequences into consideration, people tend to focus on a short subsequent time frame. Another important point is the strong influence of feelings and emotions. Having a certain personal experience or a certain aversion can severely change one's own risk estimation. As people miss the experience with low probability decisions, they also lack experience in handling them. As a consequence they may be strongly influenced by social norms or simply decide to take no decision at all. Finally Kunreuther makes several proposals on how to improve human risk perception. He gives two considerations: First, the usage and understanding of prescriptive heuristics has to be thought to humans so that extremely low probabilities can be better interpreted. Secondly, financial incentives should be developed in order to attract companies and governments to consider long term impacts of their actions.

A popular example of low probability, high consequence events are aviation accidents. Although they occur very infrequently, extensive effort is dedicated to reduce and eliminate the probability. In order to assess the risk of such accidents, Luxhoj and Coit [15] presented an Aviation System Risk Model (ASRM). Their model was designed to cover multiple causalities of known and unknown risk factors, as they are often neglected by similar models.

Finally, the book by Morgan and Henrion [17] gives a detailed overview on uncertainty analysis and presents several approaches on how to model and assess probability events. Apart from the philosophical background they also cover mathematical models of uncertainty calculation and propagation and provide information on human judgement about and with uncertainty.

6. CONCLUSION

As the number of megaprojects and the synergy between different areas increases, viable risk assessment is required. Therefore it is mandatory to cover the chance of an unsound assumption by the risk analysis. The argument used for the assessment can be further strengthened by using several independent arguments based on different models and theories.

The methodology, proposed by Ord *et al.*, was applied to PKI in order to provide an overview of its security and to depict possible error sources. It was shown that the advanced risk analysis considerably differs from the naïve approach. Although it is most difficult to gather reliable values for the different probabilities, the consideration of flawed arguments improves the reliability of the resulting probability evaluation significantly.

Further, the analysis of the PKI security showed the value of Ord's *et al.* approach when dealing with low probability, high stake events. The possibility of a flawed argument must not be neglected by PKI risk assessments, as the security primarily relies on the correctness of the theory, the model and the calculations. PKIs are used for highly classified information and are involved in various application areas. The failure of such a PKI may have disastrous impacts. As many critical issues and possible error sources were shown, further research in this area is required to help the existing frameworks and implementations to evolve.

7. REFERENCES

[1] C. Adams and M. Just. PKI: Ten years later. In *3rd Annual PKI R&D Workshop*, pages 69–84, 2004.

[2] Australian Government: Department of Finance and Deregulation. Gatekeeper PKI Framework, Threat and Risk Assessment Template. http://www.finance.gov.au/e-government/security-and-authentication/gatekeeper/index.html, 2010. [Online; accessed 29-August-2011].

[3] R. Boldyreva, M. Fischlin, A. Palacio, and B. Warinschi. A closer look at PKI: Security and efficiency. In *Proc. 10th international conference on Practice and theory in public-key cryptography*, pages 458–475, 2007.

[4] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. RFC 5280 - Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. Technical report, May 2008.

[5] D. Coppersmith. Small Solutions to Polynomial Equations, and Low Exponent RSA Vulnerabilities. *Journal of Cryptology*, 10:233–260, 1997.

[6] S. D. J. Edward A. Euler and H. H. Curtis. The Failures of the Mars Climate Orbiter and Mars Polar Lander: A Perspective from the People Involved. In *Proc. of Guidance and Control 2001, American Astronautical Society (AAS)*, pages 01–074, 2001.

[7] C. Ellison and B. Schneier. Ten Risks of PKI : What You're not Being Told about Public Key Infrastructure. *Computer Security Journal*, XVI(1):1–8, 2000.

[8] S. B. Giddings and M. L. Mangano. Astrophysical implications of hypothetical stable TeV-scale black holes. *Physical Review D*, 78(3), Aug. 2008.

[9] R. Holz, L. Braun, N. Kammenhuber, and G. Carle. The SSL landscape - a thorough analysis of the X.509 PKI using active and passive measurements. In *Proc. 11th Annual Internet Measurement Conference (IMC '11)*, 2011.

[10] Intel Corporation. Statistical analysis of floating point flaw. http://www.intel.com/support/processors/pentium/fdiv/wp. [Online; accessed 29-August-2011].

[11] D. Isbell. 1998 Mars Missions. ftp://ftp.hq.nasa.gov/pub/pao/presskit/1998/mars98launch.pdf, 1998. [Online; accessed 29-August-2011].

[12] D. Kaminsky. Black Ops of PKI, Talk at 26th Chaos Communication Congress. http://events.ccc.de/congress/2009/Fahrplan/events/3658.en.html, 2009. [Online; accessed 29-August-2011].

[13] H. Kunreuther, R. Meyer, R. Zeckhauser, P. Slovic, B. Schwartz, C. Schade, M. F. Luce, S. Lippman, D. Krantz, B. Kahn, and R. Hogarth. High Stakes Decision Making: Normative, Descriptive and

Prescriptive Considerations. *Marketing Letters*, 13:259–268, 2002.

[14] N. G. Leveson. An Investigation of the Therac-25 Accidents. *IEEE Computer*, 26(7):18–41, 1993.

[15] J. T. Luxhoj and D. W. Coit. Modeling low probability/high consequence events: an aviation safety risk model. In *Proc. of the annual Reliability and Maintainability Symposium (RAMS)*, pages 215–221. IEEE Computer Society, 2006.

[16] J. Marchesini and S. W. Smith. Modeling Public Key Infrastructure in the Real World. In *2nd European PKI Workshop: Research and Applications*, pages 1–17, 2005.

[17] M. Morgan and M. Henrion. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, 1992.

[18] Mozilla Security Blog. DigiNotar removal follow up. http://blog.mozilla.com/security/2011/09/02/diginotar-removal-follow-up, 2011. [Online; accessed 25-October-2011].

[19] T. R. Nicely. Pentium FDIV flaw. http://www.trnicely.net/pentbug/pentbug.html, 2011. [Online; accessed 29-August-2011].

[20] T. Ord, R. Hillerbrand, and A. Sandberg. Probing the improbable: methodological challenges for risks with low probabilities and high stakes. *Journal of Risk Research*, 13(2):191–205, March 2010.

[21] J. B. Peter Eckersley. Is the SSLiverse a safe place? Talk at 27C3. http://www.eff.org/files/ccc2010.pdf, 2010. [Online; accessed 29-August-2011].

[22] A. Stephenson. Mars Climate Orbiter Mishap Investigation Board Phase 1 Report. *NASA*, 10, November 1999.

doi: 10.2313/NET-2012-04-1_2