

Introducing a Trust Metric Foundation and Deriving Trust-for-Buck

Scott Harper, Jonathan Graf,
Whitney Batchelor, Timothy Dunham
Graf Research Corporation
Blacksburg, Virginia 24060
Email: trust@grafresearch.com

Peter Athanas
Dept. of Electrical & Computer Engineering
Virginia Polytechnic Institute and State University
Blacksburg, Virginia 24060
Email: athanas@vt.edu

Abstract— This study defines a flexible quantitative metric for measuring trust-related aspects across a broad range of domains and a means of using that foundation to derive domain-specific measurements. A Trust Basis Metric is described here along with examples that build on its foundation to measure assurances and identify cost-effective trust-enhancing investments. Our primary motivation in performing this study was to quantitatively determine the best increase in trust per dollar (Trust-for-Buck) when investing in current device manufacture and distribution flows for microelectronic components.

Keywords— *Trust; Assurance; Risk; Microelectronics; Metrics*

I. INTRODUCTION

The driving force behind the study that developed the metrics presented below was a need to inform an approach ensuring the availability of advanced “Trusted” microelectronic components. Our goal was to *quantitatively* determine how to achieve the best increase in trust per dollar (*Trust-for-Buck*) when investing in current ASIC vendor device manufacture and distribution flows. In measuring this aspect, we considered *Trust* as a calculation with respect to an expectation. The approach that we describe below achieves this calculation by adding and subtracting from a zero-value baseline score using vendor responses to a survey covering device manufacture and distribution aspects. We call this approach a *Trust Basis Metric*. It is easily interpreted in its raw form, as a positive score indicates a vendor generally has done more to address trust issues than anticipated and a negative score means there is room for improvement. The unitless basis metric also provides a solid foundation for deriving measurements against some unit such as monetary cost. For example, Trust-for-Buck scores are Basis Metrics scaled by the cost of making a change to existing processes to produce a quantitative metric that allows for rapid identification of design and manufacturing aspects providing the best opportunities for increasing confidence in produced ASIC devices. This flexible quantitative metric approach allows for measurement of a broad range of trust-related concerns, and is generally applicable outside of the ASIC device manufacturing domain.

In developing this study, we interacted with a set of industry-leading commercial commodity ASIC device vendors to survey their current practices and assemble a data set for application of the various metrics. This data clearly has application across a range of topics. As documented below, the data formats and collection process used have both been designed to be extensible and maintainable. The approach to data collection and storage employed here were developed with the goal of transitioning into a living data set that can be used to support future work. As such, both aspects can be used to assess additional vendors, products, and product types as desired. By adopting a data foundation that is both human readable and easily machine parsable, the data set is provided in a way that can also support investigation, development, and automated calculation of alternate metrics for a variety of security concerns.

II. BACKGROUND

The term *trust* is often used in different ways by different entities. In the most fundamental sense, it defines a feeling of confidence in meeting expectations [1]. In other cases, it describes a level of guarantee or assurance that certain actions were taken [2]. Yet other usages find it encompassing a sense of the danger or risk in performing some action [3]. Here, we separate each of these usages for examination and define our terms for each as follows.

Trust :: a measure with respect to an expectation. In this case, the expectation is that a purchased ASIC device will operate as intended. This is our primary focus and is addressed by comparing manufacturing and distribution practices against a baseline expectation. The result is both a sense of how overall manufacturing practices compare to expectations and a means of identifying key candidate areas for potential improvement.

Assurance :: a measure of confidence that some specific actions were taken or concerns were addressed. This is often quantified as a binary checklist of actions, where checklist items are used as hurdles to define various levels of confidence. This is more specific than the Trust defined above. As such, it is less useful in assessing big picture manufacturing aspects, but quite useful for ensuring certain actions are taken (e.g. defining specific end user program operations). We consider this as a secondary focus of this study and describe an assurance metric derived from the Trust Basis Metric below.

This material is based upon work supported by the United States Air Force under Contract No. FA8650-17-C-5522. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited. (88ABW-2018-4470).

Risk :: a measure of potential for loss (financial, reputation, or other measure) given a set level of Trust or Assurance (as defined above) measured against an adversary scenario. Risk is calculated by considering the contributors to a Trust or Assurance measure and considering how those aspects interact with a given adversarial scenario. That is, it is the solution to a game played between two or more entities. A Risk metric is useful in planning and decision making for specific instances. It also requires more calculation overhead than Trust or Assurance. We discuss how a Risk metric might be derived from the Trust Basis Metric below, but did not use that metric to derive conclusions in this study.

A. Perspectives

Three perspectives govern trust views when considering commercial parts – 1) device, 2) system, and 3) the customer goals guiding system development. Looking at trust from each of these differing viewpoints, it is likely that rankings of issues will vary and top choices for funding trust related efforts will change. When generally considering the device irrespective of any system or specific program view, the key question of “*Can I trust this device?*” is answered by determining if the device itself contains malice or is exposed to malicious actions. An understanding of its ancestry is an important part of that answer. At the system level, the question shifts to “*Can this device help me build a trusted system?*”. Here, device features that do not necessarily impact trust in the device itself but can be leveraged to improve system trust become important. Things like Differential Power Analysis (DPA) resistant encryption cores are important aspects in systems that use encryption and expect power analysis attacks. At the system level, concerns regarding reduction of system exposure to malice can overtake general trust concerns with the device itself. As focus shifts to the program and the question “*Should I use this device to build my trusted system?*”, customer goals and deployment scenarios can once again shift the importance of various trust aspects.

The questions of increasing end user trust in purchased commodity chips or building more trust support into devices are also two different perspectives when considering the value of investment and may very well lead to different conclusions. The data collection approach and metrics system developed below support all of these viewpoints. Primary focus, however, is placed on the first item – how general trust in a device as produced by a commercial manufacturer can be improved.

B. Related Metrics

Several metrics approaches exist (or have been proposed) that have features of interest to this study. The key contenders in this arena are summarized here.

1) TRL/MRL

Technology Readiness Level (TRL) [4] and Manufacturing Readiness Level (MRL) [5] are scales used to measure aspects of technology. They are simple quantitative values that are easily written into program plans and conveyed to technology consumers. These metrics are easily written into system design

requirements and goals and are currently in wide use. The approach for setting a value, however, is qualitative in nature.

2) CWRAP – CWE/CVE

The Common Weakness Enumeration (CWE) is actively used to categorize software weaknesses [6]. Currently maintained by MITRE, it provides a foundation for discussing, finding, and dealing with causes of software security weaknesses. It is related to the Common Vulnerabilities and Exposures (CVE) list [7] and the Common Weakness Risk Analysis Framework [8] developed for software. These systems are actively in use in measuring and analyzing software issues. They are focused on weaknesses rather than strengths, but can be extended to support both, as their schemas provides standard representation of data in a way that allows for extension [9][10]. There are current efforts underway seeking to extend CWE/CVE-like concepts to hardware [11][12].

Two key aspects of the CVE metrics approach have informed our study: 1) the data storage framework is flexible and maintainable, as evidenced by the fact that it has been actively in use since 1999, and 2) it has been adopted as the foundation for a wide variety of applications related to software vulnerability, including mappings to Security Content Automation Protocol (SCAP) standards that are used for automated vulnerability management, measurement, and policy compliance [13]. This later aspect has led to an industrial support base in the form of tools that make use of the data and perform system assessments [14].

3) DMEA Trust Metric

In 2015, the Defense Microelectronics Activity (DMEA) presented a “comprehensive quantitative model for trust at the system level” at the GOMAC conference [15]. Their model focuses on understanding the relative risks involved in supply chains and the subsequent effectiveness of risk mitigation techniques. The result is a comprehensive model for system level trust (as defined by the paper) and a follow-on model for the integrity of an integrated circuit. A demonstration of how these approximations can improve hardware assurance during the systems engineering process is proposed in the publication.

The overall DMEA measure is the product of individual probabilities of each system component, including ICs, printed circuit boards, and embedded software components, conforming to a Trust expectation. Each component level score is broken down into three variables: confidentiality, integrity, and availability. In their formulation, integrity represents the confidence in the ability of the system to meet the specifications of a given application and confidentiality of a system is the degree to which the system is protected and secure from malice. Confidentiality depends upon the system/component integrity. Lastly, availability is defined in the formulation as the level of protection or how available the system component may be provided the component integrity and confidentiality. It is important to note that availability depends upon both of these factors.

While this probabilistic formulation of system trust is not well suited to the goal of revealing opportunities for ASIC vendor investment to improve trust in commercial devices, it does contain some ideas that can inspire the metrics for this

research. Ideally, the data storage mechanism developed for this study will be able to provide information needed by the DMEA metric as well as metric solutions developed here.

4) OpTrust

Game theory is an emerging research approach to production of trust (risk) metrics. An example of this type of solution is the OpTrust strategy optimization tools that prescribe optimal verification techniques for Trojan detection [16]. That approach takes into account hardware Trojan strategies, design criticality, and threat environment, then uses game theory to prescribe optimal hardware Trojan defense strategy.



Figure 1. Ideal functionality of the OpTrust approach

OpTrust guides developers seeking to select optimal mitigation strategies. It solves the games on behalf of the designer, can handle the entire ASIC device attack surface, and can be adapted for many adversary/defender interactions in microelectronics. While not directly applied, the components that comprise the utility functions used to set up the games solved by OpTrust are, in many cases, considered by the metrics formulation of this study.

III. TRUST BASIS METRIC

In deriving a quantitative foundation for measuring trust, we considered *trust* as defined above – a calculation with respect to an expectation. The approach to trust calculation defined in this study is rooted in quantitative assessment of gathered vendor data. In answering multiple choice questions, vendors describe their processes as falling within a defined spectrum of trust related options. The options for each question are designed to range from a selection that implies little or no trust effort to a selection that implies maximum effort expended toward trust. The same survey process can be used to establish a baseline expectation of trust, and vendor responses can be measured against this baseline to determine a quantitative trust score. We call the outcome of this approach a *Trust Basis Metric*.

Our example use of a Trust Basis score provides a measure of how well ASIC vendor device development, production, and distribution flows meet expectations for producing a trusted product. Scores are calculated by measuring vendor deviation from a baseline for a set of questions (Figure 2). If the product line being considered does more than required to address the concern expressed in the question, a positive score is attained for that concern. If it does less, the result is a negative score. Scores for all questions are accumulated into five domains of commodity device trust to produce a measure of trust for each

domain. Domain scores are accumulated to result in a single value Trust Basis Score for a device product line.




Figure 2. Scoring Multiple Choice Answers Against a Baseline Expectation

Note that the accumulated score is an assembly of indicators providing a quantified sense of trust. The final accumulated value should be taken as guidance rather than specific instruction, however, given the data aggregation. Negative overall scores imply room for improvement and positive scores indicate more is being done overall to address trust that expected. A positive score may still have negative contributors that could be improved.

For the purposes of this study, no scaling was used in score accumulation for the raw basis results. This is the recommended use of the basis metric, as it allows for subsequent derivation of metrics targeted at specific goals such as identifying investment opportunities or locating prime opportunities for reducing risk within a given program. The raw score assembly formula is performed as per (1) and (2).

$$\text{Question Score} = \text{Vendor Answer} - \text{Baseline Answer} \quad (1)$$

$$\text{Trust Basis Score} = \sum [\text{Question Scores}] \quad (2)$$

IV. CASTING THE BASIS

The scoring flow and the data formats described here support weighting of contributors and domains during score assembly. A variety of targeted trust metrics can therefore be derived from the trust basis by applying different scaling factors to the score assembly process. Financial costs required to change answers can be used to produce investment insights, program-specific risks introduced by selecting certain answers can be used to identify opportunities for improving a program stance, and other modifiers can be used individually or in combination to produce desired insights.

A. Trust-for-Buck

A *Cost Weighted Metric* was derived in this study to recognize investment opportunities. Our goal was to identify prime candidate areas for trust-enhancing investments with commercial vendors. As such, the vendor survey included a mechanism for returning the relative costs of changing operations across the five domains. This relative cost relationship was used to scale Trust Basis scores within domains and extract a Trust-for-Buck view of potential investment with respect to trust concerns in vendor operations.

When viewed as a set of individual scored responses or accumulated domain scores, this scaling allows for the identification of top investment candidates across and within domains, respectively. Accumulation of cost weighted scores into single values produces an overall Trust-for-Buck score that provides a measure of potential return on investment with a given vendor (for a select device) across all domains. This accumulation process is shown in (3).

$$\sum_{\text{all domains}} \left(\frac{\text{Question Scores}}{\text{Relative cost of changing operations in domain}} \right) \quad (3)$$

The accumulated Trust-for-Buck metric can provide a general comparison of overall trust investment payoff across various vendors and product lines. This summary value only provides rough insight for high-level product comparisons, however. A more insightful view of the Trust Basis and Cost Weighted metrics is seen at the question and domain level. Here, specific concerns can be compared and the relative payoff of investments within domains can be assessed.

Accumulating concerns across all vendors into more general domain scores for visualization is useful in understanding how well the industry as a whole is addressing trust issues in various phases of device production. This accumulation can provide visualizations similar to Figure 3. Here, one can see that the domain with the most room for improvement on average is Silicon Design. The range of the Silicon Fabrication domain responses, however includes both concerns that have the most room for improvement (the lowest bar) and concerns for which the most is being done beyond expectations (the highest bar). Finally, the Chip Supply domain is on average showing more than expected being done to address trust concerns.




Figure 3. Visualizing aggregate Trust-for-Buck scoring (screen captured from interface developed here, but not necessarily representative of actual survey results)

A more granular view of a dataset is depicted in Figure 4. This chart is an example graphical depiction of basis scores for individual concerns weighted by the cost of making changes within each domain. In this view, individual question response

scores are represented as bars rising or dropping from the baseline expectation. The thick bars at the top of the figure represent the relative costs of making changes within the domain, and those weights were used to scale the individual question variance from the baseline. This representation allows for rapid identification of individual outliers as well as a general sense of how concerns vary across domains. For example, the item in this graphic for which the least is being done to address trust (the longest negative bar) lies in the Silicon Fabrication domain, and it is clear that the bulk of the concerns within the Chip Supply domain are being addressed appropriately with respect to expectations.




Figure 4. Example Graphical View of Concern Scores (screen captured from interface developed here, but not necessarily representative of actual survey results)

B. Quantified Risk

The concept of weighting can be extended beyond the trust concern survey data to merge that information with other data sets. For example, a *Quantified Risk Metric* can be produced by scaling the Trust Basis with potential gains and losses for individual risk concerns or domain deviations from a baseline. Game theory can provide a good source for this type of scaling information, and programs like the Sandia National Labs *PRactical Evaluation and Synthesis of Trust In Government systEms* (PRESTIGE) [17] or Graf Research *OpTrust* [16] can be integrated with the Trust Basis approach to develop this type of risk-based metric. This idea of merging results can be extended further to combine both weighting based on the cost of changing operations with risk weighting, resulting in a measure of the cost of addressing risks within a program. Figure 5 illustrates this concept of deriving metrics for various purposes from the Trust Basis.




Figure 5. Merged Metric Derivation

Flexibility in the score accumulation and trust concern data set also allows for more complex derivation of metrics from the Trust Basis. A simple adjustment is to emphasize some aspects over others. For example, if *DPA resistant encryption* is more of a concern than having a *Hardware ID* on a device, the *DPA resistant core* score can be multiplied during accumulation. Similarly, if the *Chip Supply* domain is more of a concern than the *User Circuit Design* domain, all concerns within the domain could be multiplied by a scaling factor. Given the defined data formats, it is also possible to alter more complex aspects of the score accumulation. For example, the step size between answers for any individual question can be adjusted so that the variation between answers is counted as more or less than unity. Equation (4) describes the process of applying emphasis at a variety of points during score accumulation.

C. Pillars of Assurance

The Trust Basis score is well suited to providing a measure of trust against an expectation but does not convey assurance that select operations occurred in the course of device design, fabrication, or distribution. A measure of Assurance that certain questions were answered in select ways can, however, be derived from Trust Basis survey data by tagging questions of interest. Here, we define a quantitative metric using this approach called *Pillars of Assurance*.

The Pillars of Assurance metric defines Assurance Levels that are attained as certain actions are taken. For example, in the Silicon Design domain an end user may feel that general company staff are sufficient for some programs, but vetted individuals are required for others. A question might ask, “Characterize the vendor influenced people involved in each of the following processes:” with the options of 1) “Unknown”, 2) “All general staff or contractors”, 3) “All company vetted staff”, and 4) “All hold clearances”. The end customer can set a requirement that answer 3 or above be selected for Assurance Level 0-3 and answer 4 must be selected to reach Assurance Level 4. This expectation is set by tagging the question within the data set. Figure 6 provides a graphical sense of this type of evaluation.

While the number of pillars can be varied, this approach is likely most useful with a fairly limited number of levels (e.g. no more than 10). Similarly, the contributing factors can be varied. The set of contributions selected for this study was designed to be broadly applicable and somewhat intuitive in application. It divides assurances by domain and addresses consistent concerns across those domains. Within each, it establishes assurance that certain actions were taken in first-party IP development, third-party IP assessment, verification of results and materials, controls on data and material access, and material transit operations. In measuring these concerns against vendor operations, the assurance level of each concern is depicted to provide a means for identifying underperforming areas. Domain assurance is set as the minimum concern assurance and Overall level is defined as the minimum of the domain levels. Grey items in the visualization indicate no questions were tagged for assurance checks in that category.




Figure 6. Measuring Assurance Levels via Pillars of Assurance (screen captured from interface developed here, but not necessarily representative of actual survey results.)

V. SUMMARY AND FUTURE WORK

Building an understanding of ASIC device trust includes many aspects spanning production, supply, design, deployment, and sustainment. Similarly, the System-on-Chip (SoC) development approaches used in many modern chips include devices, user circuit designs, and user software designs. This study weights all these aspects against each other, letting

$$\sum_{\text{domains}} \text{DomainWeight} \times \frac{\sum_{\text{categories}} (\text{CategoryWeight} \times \sum_{\text{questions}} (\text{QuestionWeight} \times \text{Question Score}))}{\text{Relative cost of changing operations in domain}} \quad (4)$$

metrics emerge from the data. The goal is to avoid squeezing information into a one-dimensional framework by simply choosing a closest-fit existing metric. Several existing approaches were considered as providing potential starting points. In the end, however, it was determined that the most suitable metric approach is to create a numerical survey-based “basis metric” that can be used to form other metrics exploring different aspects of the trust space. The flexibility of this approach permits the derivation of cost-weighted metrics related to potential return on investment as well as use of the trust concern survey results to inform metrics for use in other trust-related spaces. In the course of this study, we defined a Trust Basis Metric and derived three targeted metrics from that basis as outlined in TABLE I.

TABLE I. MEASUREMENT PERSPECTIVES SUMMARY

Perspective	Description	Fundamental Measure	Quantitative Metric
Trust	Expectation with respect to concerns	Comparison of practices to baseline expectations	Trust Basis Metric <i>(How much is vendor effort addressing concerns?)</i> Cost Weighted Metric <i>(Which trust investments provide the most return? Trust-for-Buck)</i>
Assurance	Guarantee that specific actions were taken, or certain properties hold	Binary comparison of actions to guarantees	Pillars of Assurance (coarse-grained measure indicating desired actions have been taken)
Risk	Potential for loss given a level of Trust or Assurance measured against a scenario	Potential for loss given a level of Trust or Assurance measured against a scenario	Trust Basis Metric + Program-Specific Question Weighting

The driving force behind this study was a need to inform funding of new efforts to be led by commodity ASIC device manufacturers to help ensure the availability of advanced “Trusted” microelectronic components. In addition to the metrics discussed above, the study also developed a range of underlying capabilities that will be useful in related work.

The data storage mechanism employed for the survey is based on a machine-parsable JSON template for concern (question) data. The format is similar to that used in the Common Vulnerabilities and Exposures (CVE) database [18]. It includes unique IDs for survey questions, a plain language description of why the issue is a concern, and references regarding the issue, as well as categorization and tagging of entries. The format also supports aspects related to basis metric assembly including concern/category weighting, specification of contribution step size at a per-question level, and definition of both the multiple-choice answer sets used for concern scoring and relational answer sets used for generating relative costs (weights) across categories.

An interactive html survey interface allows for the survey data to be consumed and processed in either a server-based or local store environment. The interface takes advantage of tagging within the dataset to permit assembly of surveys based on a subset of documented concerns. This was used extensively in this study, as consumption of industry guidance documents led to some concerns being repeated and other concerns being entered that do not relate to commercial manufacturing. Rather than leave this information out of the data set, it was simply not tagged as relevant to this particular study. Similar interfaces supported both survey takers and survey consumers. Some of the consumer end is represented in the screen captures shown in Figure 3, Figure 4, and Figure 6 above.

The approach of using a quantified but unitless basis metric for deriving measures of interest has shown itself to be useful in addressing a range of trust, risk, and assurance concerns. We anticipate this approach, as well as the material developed to support its deployment, will continue to be useful in assembling other measures of trust-related concern.

REFERENCES

- [1] B. Barber, “The Logic and Limits of Trust,” Rutgers University Press, 1983
- [2] Northern Lincolnshire and Goole, “Trust Assurance Framework (TAF),” NHS Foundation Trust, Directorate of Performance Assurance, October 2015
- [3] J.R. Eiser, et al., “Risk interpretation and action: A conceptual framework for responses to natural hazards,” International Journal of Disaster Risk Reduction, vol 1, October 2012.
- [4] Defense Acquisition Guidebook, 2013.
- [5] Manufacturing Readiness Level (MRL) Deskbook, 2016.
- [6] MITRE Corp, “Common Weakness Enumeration (CVE),” <https://cwe.mitre.org>
- [7] MITRE Corp, “Common Vulnerabilities and Exposures (CVE),” <http://cve.mitre.org/>
- [8] MITRE Corp, “Common Weakness Risk Analysis Framework (CWRAF),” <https://cwe.mitre.org/cwraf/>
- [9] MITRE Corp, “Common Weakness Enumeration Schema,” <https://cwe.mitre.org/documents/schema/>
- [10] MITRE Corp, “Common Vulnerabilities and Exposures Schema,” <https://cve.mitre.org/data/downloads/#xml>
- [11] B. Cohen, “Harmonizing Software, Firmware, and Hardware Assurance,” Software and Supply Chain Assurance Winter Working Group, 2016
- [12] <http://www.counterfeit-ic.org/>
- [13] <https://cve.mitre.org/about/>
- [14] <https://nvd.nist.gov/scap/validated-tools>
- [15] D. Pentrack, N. Levine, J. Lloyd, A. Gahoonia, “Quantifying System Trust and Microelectronics Integrity,” Government Microcircuit Applications & Critical Technology Conference (GOMACTech), March 2015
- [16] J. Graf, “OpTrust: Software for Determining Optimal Test Coverage and Strategies for Trust,” Government Microcircuit Applications & Critical Technology Conference (GOMACTech), March 2017.
- [17] M. Galiardi, E. Vugrin, B. Eames, et al, “On Modeling Detection for Quantitative Trust Analysis,” Government Microcircuit Applications & Critical Technology Conference (GOMACTech), March 2018
- [18] MITRE Corp, Common Vulnerabilities and Exposures Information Format, https://cve.mitre.org/cve/cna/rules.html#Appendix_B_format