

Article

A Case for Security-aware Design-Space Exploration of Embedded Systems

Andy D. Pimentel

Parallel Computing Systems group
University of Amsterdam
The Netherlands
a.d.pimentel@uva.nl

Version July 10, 2020 submitted to *J. Low Power Electron. Appl.*

1 **Abstract:** As modern embedded systems are becoming more and more ubiquitous and interconnected,
2 they attract a world-wide attention of attackers and the security aspect is more important than
3 ever during the design of those systems. Moreover, given the ever-increasing complexity of the
4 applications that run on these systems, it becomes increasingly difficult to meet all security criteria.
5 While extra-functional design objectives such as performance and power/energy consumption are
6 typically taken into account already during the very early stages of embedded systems design, system
7 security is still mostly considered as an afterthought. That is, security is usually not regarded in the
8 process of (early) design-space exploration of embedded systems, which is the critical process of
9 multi-objective optimization that aims at optimizing the extra-functional behavior of a design. This
10 position paper argues for the development of techniques for quantifying the 'degree of secureness'
11 of embedded system design instances such that these can be incorporated in a multi-objective
12 optimization process. Such technology would allow for the optimization of security aspects of
13 embedded systems during the earliest design phases as well as for studying the trade-offs between
14 security and the other design objectives such as performance, power consumption and cost.

15 **Keywords:** Embedded computer systems; cyber security; system-level design and design-space
16 exploration; multi-objective optimization; system trade-offs

17 1. Introduction

18 Embedded computer systems are ubiquitous and have a major impact on our society. Examples
19 of such systems are close at hand: modern TVs contain one or multiple computer systems to
20 handle functionality such as decoding the input signal, performing various image enhancement
21 techniques as well as displaying and updating live information (e.g., program guide or weather
22 forecast). Smart-phones rely on embedded computer systems to allow users to make phone calls,
23 shoot photos and videos, perform GPS navigation, browse the Internet, execute apps, and so on. The
24 use of embedded computer systems is, however, by no means restricted to consumer electronics: in
25 industrial, medical, automotive, avionic, or defense applications they are equally pervasive.

26 The complexity of the underlying system architectures of modern embedded systems forces
27 designers to start with modeling and simulating (possible) system components and their interactions
28 in the very early design stages. This is often referred to as system-level design [1]. The system-level
29 models typically represent application workload behavior, characteristics of the underlying computing
30 platform architecture, and the relation (e.g., mapping, hardware-software partitioning) between
31 application workload(s) and platform architecture. These models are applied at a high level of
32 abstraction, thereby minimizing the modeling effort and optimizing the simulation speed. This is
33 especially needed for targeting the early design stages since many design decisions are still open

34 and, therefore, many design alternatives still need to be studied. High-level system modeling allows
35 for the early verification of a design and can provide estimates on the extra-functional properties
36 of a design such as system performance and energy/power consumption. The system-level models
37 are typically accompanied by a methodology for efficient design-space exploration (DSE) [2], which
38 is the process of assessing alternative design instances with respect to i) the platform architecture
39 that will be deployed (e.g., the number and type of processing elements in the platform, the type of
40 network to interconnect the processors, etc.) and ii) the mapping of application tasks to the underlying
41 platform components [3]. It is a multi-objective optimization problem that searches through the space
42 of different implementation alternatives to find optimal design instances. Exploration of different
43 design choices, especially during the early design stages where the design space is still at its largest,
44 is of eminent importance. Wrong decisions early in the design can be extremely costly in terms of
45 re-design effort, or even deadly to the product's success. Consequently, considerable research effort in
46 the embedded systems domain has been spent in the last two decades on developing frameworks for
47 system-level modeling and simulation that aim for early design-space exploration.

48 As embedded systems are becoming more and more ubiquitous and interconnected (illustrated
49 by, e.g., the strong trend towards the Internet of Things), they also attract a world-wide attention of
50 attackers. This makes the security aspect more important than ever during the design of these systems
51 [4]. Moreover, given the ever-increasing complexity of the applications that run on modern embedded
52 systems, it becomes increasingly difficult to meet all security criteria. While design objectives such as
53 performance and power/energy consumption are usually taken into account during the early stages
54 of design (as explained above), system security is still mostly considered as an afterthought. That
55 is, security is typically not regarded in the process of (early) design-space exploration of embedded
56 systems. However, any security measures that may eventually be taken much later in the design
57 process do affect the already established trade-offs with respect to the other extra-functional properties
58 of the system like performance, power/energy consumption, cost, etc. [4]. Thus, covering the security
59 aspect in the earliest phases of design is necessary to design systems that are, in the end, optimal
60 with regard to all extra-functional objectives. However, this poses great difficulties because unlike the
61 earlier mentioned conventional system objectives, like performance and power consumption, security
62 is hard to quantify: there exists no single metric with which one can measure the degree of secureness
63 of a design.

64 This position paper argues for the need for security-aware, system-level design-space exploration
65 methods and techniques for embedded systems. To this end, we will discuss a multifaceted,
66 scoring-based methodology for quantifying the degree of secureness of embedded system design
67 instances. This methodology allows for incorporating the secureness quantifications in a multi-objective
68 optimization process and would thus enable optimization of the security aspect during the earliest
69 phases of design. However, we want to emphasize the fact that this is a position paper and therefore
70 does not present an actual implementation of the proposed solution nor any experimental results.

71 The remainder of this paper is organized as follows. In the next section, we will provide a brief
72 introduction to the concept of design-space exploration. In Section 3, we will describe our proposal for
73 a security-aware DSE approach, focusing on a method to quantify the secureness of embedded system
74 design instances. Section 4 discusses related work, after which Section 5 concludes the paper.

75 2. Design-Space Exploration

76 During the design-space exploration (DSE) of embedded systems, multiple optimization *objectives*
77 – such as performance, power/energy consumption, and cost – should be considered simultaneously.
78 This is called multi-objective DSE [2]. Since the objectives are often in conflict, there cannot be a single
79 optimal solution that simultaneously optimizes all objectives. Therefore, optimal decisions need to be
80 taken in the presence of trade-offs between design criteria.

2.1. Multi-objective Optimization

Given a set of m decision variables, which are the degrees of freedom (e.g., parameters like the number and type of processors in the system, application mapping, etc.) that are explored during DSE, a so-called *fitness function* must optimize the n objective values [2]. The fitness function is defined as:

$$f_i : R^m \rightarrow R^1 \quad (1)$$

A potential solution $x \in R^m$ is an assignment of the m decision variables. The fitness function f_i translates a point in the solution space X into the i -th objective value (where $1 \leq i \leq n$). For example, a particular fitness function f_i could assess the performance or energy efficiency of a certain solution x (representing a specific design instance). The combined fitness function $f(x)$ subsequently translates a point in the solution space into the objective space Y . Formally, a multi-objective optimization problem (MOP) that tries to identify a solution x for the m decision variables that minimizes the n objective values using objective functions f_i with $1 \leq i \leq n$:

$$\begin{aligned} \text{Minimize } y &= f(x) = (f_1(x), f_2(x), \dots, f_n(x)) \\ \text{Where } x &= (x_1, x_2, \dots, x_m) \in X \\ y &= (y_1, y_2, \dots, y_n) \in Y \end{aligned}$$

Here, the decision variables x_i (with $1 \leq i \leq m$) usually are constrained. These constraints make sure that the decision variables refer to valid system configurations (e.g., using not more than the available number of processors, using a valid mapping of application tasks to processing resources, etc.), i.e., x_i are part of the so-called feasible set. In the remainder of this section, we assume a minimization procedure, but without loss of generality, this minimization procedure can be converted into a maximization problem by multiplying the fitness values y_i with -1 .

With an optimization of a single objective, the comparison of solutions is trivial. A better fitness (i.e., objective value) means a better solution. With multiple objectives, however, the comparison becomes non-trivial. Take, for example, two different embedded system architecture designs: a high-performance system and a slower but much cheaper system. In case there is no preference defined with respect to the objectives and there are also no restrictions for the objectives, one cannot say if the high-performance system is better or the low-cost system. A typical MOP in the context of embedded systems design can have a variety of different objectives, like performance, energy consumption, cost and reliability. To compare different solutions in the case of multiple objectives, the Pareto dominance relation is generally used. Here, a solution $x_a \in X$ is said to dominate solution $x_b \in X$ if and only if $x_a < x_b$:

$$\begin{aligned} x_a < x_b &\iff \forall i \in \{1, 2, \dots, n\} : f_i(x_a) \leq f_i(x_b) \wedge \\ &\quad \exists i \in \{1, 2, \dots, n\} : f_i(x_a) < f_i(x_b) \end{aligned}$$

Hence, a solution x_a dominates x_b if its objective values are superior to the objective values of x_b . For all of the objectives, x_a must not have a worse objective value than solution x_b . Additionally, there must be at least one objective in which solution x_a is better (otherwise they are equal).

An example of the dominance relation is given in Figure 1, which illustrates a two dimensional MOP. For solution H the dominance relations are shown. Solution H is dominated by solutions B , C and D as all of them have a lower value for both f_1 and f_2 . On the other hand, solution H is superior to solutions M , N and O . Finally, some of the solutions are not comparable to H . These solutions are better for one objective but worse for another.




Figure 1. A Pareto front and an example of the dominance relation (taken from [2]).

116 The Pareto dominance relation only provides a partial ordering. For example, the solutions A to F
 117 of the example in Figure 1 cannot be ordered using the ordering relation. Since not all solutions $x \in X$
 118 can be ordered, the result of a MOP is not a single solution, but a front of non-dominated solutions,
 119 called the *Pareto front*. A set X' is defined to be a Pareto front of the set of solutions X as follows:

$$\{x \in X' \mid \nexists x_a \in X: x_a < x\}$$

120 The Pareto front of Figure 1 contains six solutions: $A - F$. Each of these solutions does not
 121 dominate the other. An improvement on objective f_1 is matched by a worse value for f_2 . Generally, it
 122 is up to the designer to decide which of the solutions provides the best trade-off.

123 2.2. Search for Pareto optimal solutions

124 The search for Pareto optimal design points with respect to multiple design criteria entails two
 125 distinct elements [5]:

- 126 1. The evaluation of a single design point using the fitness function(s) $f(x)$ regarding all the
 127 objectives in question like system performance, power/energy consumption and so on. These
 128 evaluations are usually based on measurements using real systems or predictions from either
 129 analytical models or simulation models [2].
- 130 2. The search strategy for navigating through and covering the design space during the DSE process.
 131 Such search strategies can be based on exact, but typically unscalable, methods that guarantee
 132 finding the optimal solution(s). These exact methods can, for example, be implemented using
 133 integer linear programming (ILP) solutions (e.g., [6,7]) or branch & bound algorithms (e.g., [8]).
 134 Alternatively, so-called meta-heuristics, such as genetic algorithms (GA) or simulated annealing,
 135 can be used to search the design space for optimal solutions. They only perform a finite number
 136 of design point evaluations, and can thus handle larger design spaces. However, there is no
 137 guarantee that the global optimum will be found using meta-heuristics, and therefore the result
 138 can be a local optimum within the design space. GA-based DSE has been widely studied in
 139 the domain of system-level embedded design (e.g., [9–12]) and has demonstrated to yield good
 140 results.

141 In this paper, we focus on the fitness evaluation aspect of DSE. More specifically, we argue that
 142 while there are well-established techniques and metrics for the fitness evaluation of traditional design
 143 objectives such as performance, power / energy consumption, cost, and reliability, this is not the case
 144 for evaluating the fitness of design instances in terms of how secure they are. This lack of security
 145 fitness evaluation methods and metrics inhibits the use of system security as a first-class citizen in the
 146 process of early design-space exploration of embedded systems. As was indicated before, such design
 147 practice leads to suboptimal products because any security measures that may be taken later in the

148 design process do affect the already established trade-offs with respect to the other extra-functional
149 properties of the system like performance, power/energy consumption, cost, etc.

150 In the next section, we will therefore argue for the development of a security-aware DSE approach,
151 based on a multifaceted, scoring-based security quantification methodology. This methodology allows
152 for quantifying the degree of secureness of design instances such that these can be incorporated in the
153 DSE's multi-objective optimization process. Eventually, once such a security-aware DSE would have
154 been implemented, it would allow for optimization of security aspects of embedded systems in their
155 earliest design phases as well as for studying the trade-offs between security and the other design
156 objectives like performance, power consumption and cost. Evidently, such technology would provide
157 a substantial competitive advantage in the embedded systems industry.

158 3. Towards security-aware, system-level DSE

159 The envisioned approach for security-aware system-level design-space exploration, adopting
160 a multifaceted, scoring-based security quantification methodology, is illustrated in Figure 2. Below,
161 we will explain the different components of this approach. The blue parts of Figure 2 refer to the
162 methodology components that only need to be specified or performed once, whereas the red parts
163 refer to components that are dependent on the design-space exploration process and thus must be
164 revisited every time a new design instance is evaluated in terms of extra-functional properties such as
165 performance, power consumption, and of course, in the scope of this paper, also secureness. Before
166 describing our approach in detail, however, we will first discuss several assumptions that delimit our
167 proposed approach.

168 3.1. Assumptions

169 We focus on security threats in which the underlying embedded system architecture plays a
170 central role, and do not consider any security flaws that can be exploited purely at the application
171 level. This implies that we restrict ourselves to the following set of attack types:

- 172 1. *Side-channel attacks* like power analysis attacks, timing attacks such as the recent Spectre and
173 Meltdown attacks, scan attacks, differential fault analysis attacks and electromagnetic analysis
174 attacks (see [13,14] for an overview of these side-channel attacks);
- 175 2. *Denial of service attacks* [15,16];
- 176 3. *Software-based attacks* such as buffer overflows for which protection mechanisms may be available
177 at the system (architecture) level (e.g. [13]);
- 178 4. *Attacks directed towards breaking encryption algorithms* [17].

179 For each of the above attacks, we subsequently consider a range of protection mechanisms – derived
180 from literature – that can be applied to protect specific system components or the entire system against
181 these attacks.

182 Moreover, we consider system-level DSE in which both the platform architecture (e.g., selection
183 of platform components such processing elements, memories, and networking components) as well
184 as the mapping of application tasks and communications to the selected platform components are
185 optimized for traditional objectives such as performance, power consumption, and cost, but now
186 also for secureness. Such system-level design-space exploration is depicted in the top-middle part of
187 Figure 2 and could, for example, be performed with system-level DSE frameworks such as Sesame
188 [18,19] or a similar environment (e.g. [20–22]). Important to note here is that the performance, power
189 and cost models used in the DSE also need to account for the effects of deploying specifically selected
190 security protection mechanisms (as discussed above) inside a platform architecture.

191 3.2. A multifaceted, scoring-based methodology for secureness quantification

192 As a first step in our methodology, shown in the top left of Figure 2, the applications that need
193 to execute on the target embedded system together with their extra-functional requirements are

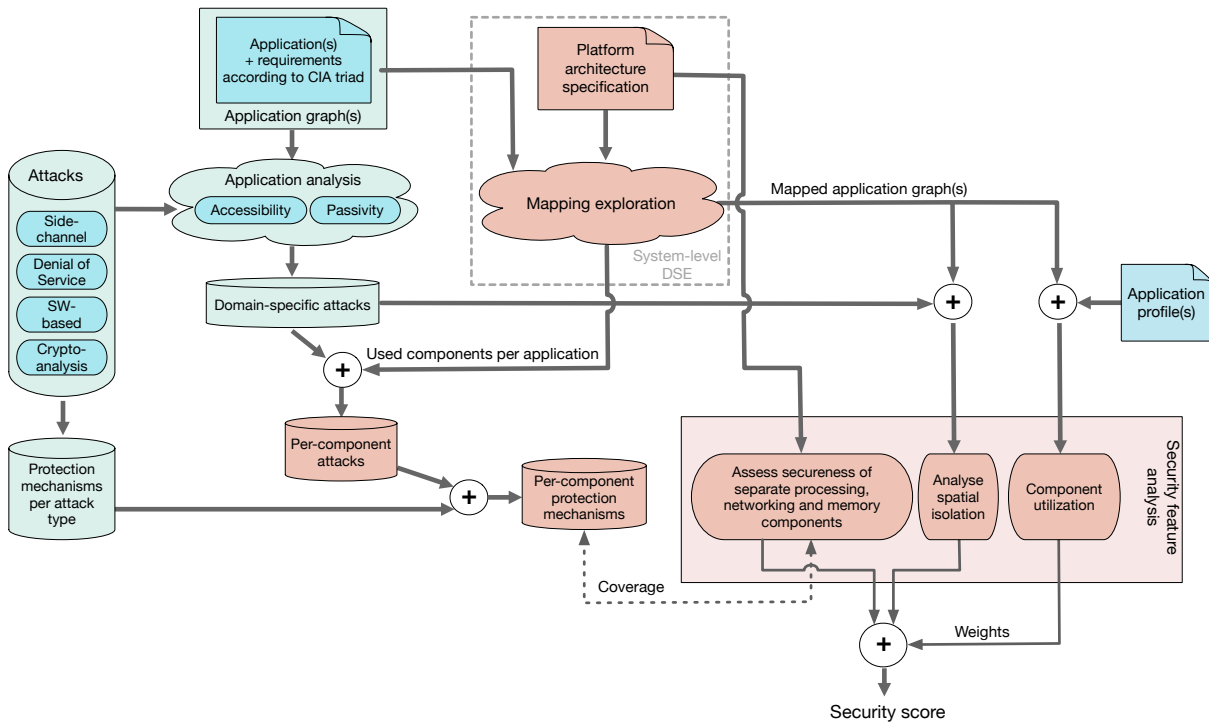


Figure 2. Proposed approach for security-aware system-level DSE using a multifaceted, scoring-based security quantification methodology.

194 identified and specified. The specified extra-functional requirements include the traditional ones
 195 such as performance and real-time requirements, power/energy consumption budgets, etc., but also
 196 requirements in terms of secureness. Regarding the latter, one can use the well-known CIA triad to
 197 indicate the needs with respect to the security aspects Confidentiality (preventing sensitive information
 198 from reaching the wrong people), Integrity (maintaining consistency, accuracy, and trustworthiness of
 199 data) and Availability (ensuring timely and reliable access to, and use of, information) [23]. Here, a
 200 domain-specific language (DSL) could be developed to specify these extra-functional requirements. The
 201 application workloads themselves can be specified using task or process graphs, explicitly describing
 202 application tasks and their interactions (communications).

203 Given the attack types we consider in our proposed approach, as discussed in Section 3.1 and
 204 shown at the left in Figure 2, only those attacks that are relevant for the embedded system under
 205 design need to be identified, which we refer to as the so-called domain-specific attacks. To this end,
 206 we need to consider the security requirements of the target embedded system as specified using the
 207 CIA triad as well as the characteristics of the specific attack types in terms of, e.g., passivity and
 208 accessibility. Here, passivity refers to what extent an attack manipulates the target system, either as a
 209 means or a goal of the attack. For example, a denial of service attack clearly is an active attack as its sole
 210 aim is to manipulate the system, whereas a side-channel attack based on power analysis is a passive
 211 attack. Accessibility refers to the access level that is required for an attack to be performed. Revisiting
 212 the example of a side-channel attack via power analysis, such an attack obviously requires physical
 213 access to the embedded system, whereas e.g. a software-based attack does not require this. If we now
 214 consider, for example, an anti-lock braking system, then confidentiality is not a major concern as such
 215 a system does not process sensitive information. This makes passive attacks such as side-channel
 216 attacks not relevant and can therefore be excluded from the set of domain-specific attacks. However,
 217 the braking system may not be disrupted or manipulated (i.e., integrity and availability are crucial
 218 CIA elements) thereby making active attacks highly relevant. Table 1 provides an overview of the
 219 required access level and passivity of the attack types we consider in this paper. Here, virtual access

Attack	Sub-type	Access level	Passivity
Side-channel	Power analysis	Physical	Passive
	Timing attack	Virtual	Passive
	Scan attack	Physical	Active
	Fault Analysis	Physical & Virtual	Passive & Active
	Electromagnetic Analysis	Physical	Passive
Denial of Service		Virtual	Active
Software	Buffer overflow	Virtual	Active
Cryptanalysis		None	Passive

Table 1. Required access level and passivity of various types of attacks on embedded systems.

level attacks require access to one process that runs within an application on the embedded system in question. Cryptanalysis attacks often do not require access to the system. For example, in the case of public key cryptography, the public key is distributed to other systems and therefore freely available.

Once the set of domain-specific attacks has been determined, those attacks that are relevant to the different components in the underlying platform architecture can be determined: for example, a networking component is not susceptible to a software-based buffer overflow attack, whereas a microprocessor component is. To do so, we also need the mapping information (i.e., which application tasks and communications are mapped onto what platform components) of the design instance(s) that are currently being explored by the system-level DSE process. Subsequently, for each component in the platform architecture, we can now determine the set of possible security protection mechanisms that can be deployed to effectively increase its secureness ('Per-component protection mechanisms' in Figure 2). These sets of possible per-component protection mechanisms are an important ingredient of our envisioned scoring-based security quantification methodology: they allow for determining the coverage with respect to the protection mechanisms that are actually deployed in the design instances being explored by the system-level DSE. To achieve this, a scoring technique would be needed that can capture binary coverage relationships (i.e., a certain protection mechanism is available or not) as well as numerical coverage relationships. The latter applies in cases where, for example, a certain amount of random noise is added to a system component to disguise real power behavior in order to complicate or even prevent side-channel attacks based on power analysis [24] (here, the amount of noise forms a power / security trade-off) or when the strength of a cryptographic processor is identified as a function of the key-size it uses.

Besides the coverage of protection mechanisms deployed in the platform architecture components, one can take two other facets into account to determine the security score of a particular design instance. First, the spatial isolation realized in design instances can be considered. That is, reducing the amount of resource sharing between applications or even between tasks from a single application will increase the secureness of the system, as this will complicate certain types of attacks such as side-channel attacks. Therefore, a proper technique for quantifying the spatial isolation (using the mapping information from the system-level DSE) would be required such that it can be used for security scoring purposes. As a final ingredient of our anticipated security score, the platform component utilization can be used. The rationale behind this is that higher utilized components typically are more prone to certain attacks. Moreover, higher utilized components possibly also play a more important role in achieving the CIA-triad system requirements. To include the platform component utilization, we need to profile the application(s) to measure the activity of application tasks and communications, i.e., the degree to which they utilize the underlying resources. Hereafter, this information is related to the mapping of these application tasks and communications onto the platform architecture. The resulting component utilization can then be used to weight the protection mechanism coverage and spatial isolation of design instances in the final security scoring (as shown at the bottom of Figure 2).

257 3.3. Scoring the security of design instances

258 Above, we described the ingredients of our envisioned security scoring methodology. We
 259 do realize that we have not provided any details on how such security scoring could actually be
 260 implemented. Actually, this remains a topic for future research, which will hopefully also be picked up
 261 by the community. Nevertheless, in this section, we would like to provide a rough sketch of a fairly
 262 simple approach to do such scoring.

263 Given a mapping of a (set of) application(s) to the underlying resources of a possible platform
 264 architecture, which includes the mapping of application tasks to computational resources as well as
 265 the mapping of inter-task communications to network and memory resources. Then, for each utilized
 266 component in the platform, we could calculate a security score along the following lines. Let AT_x be
 267 the set of Attack Types (see e.g. the second column of Table 1) to which component x is susceptible:

$$AT_x = \left\{ \bigcup_{\forall t_i|c_i \text{ mapped on } x} \text{Attacks}_{t_i|c_i} \right\}$$

268 For AT_x , we only consider the application tasks t_i or inter-task communications c_i (dependent
 269 on whether x is a processing or communication component) that are mapped to component x .
 270 $\text{Attacks}_{t_i|c_i}$ refers to the possible attacks for task t_i or communication c_i , taking into account its security
 271 requirements according to the CIA triad as well as taking into account the access level and passivity of
 272 the various attacks (see Table 1). To determine a security score S_x for a component x , one could then
 273 perform the following calculation:

$$S_x = \frac{\sum_{p \in P_{AT_x}} \text{Protection-level}(p, x)}{|AT_x|} \cdot \frac{1}{\text{Utilization}_x}$$

274 Here, p is a particular protection mechanism and is part of the set P_{AT_x} that consists of the
 275 possible protection mechanisms for the attacks listed in set AT_x for component x . The function
 276 $\text{Protection-level}(p, x)$ returns a value that indicates the extent to which component x implements
 277 protection mechanism p . This function could, for example, return a value between 0 and 1: The
 278 value 0 would mean that the component does not implement the protection mechanism, implying
 279 that component x would be fully susceptible to the associated attack type. The value 1, on the other
 280 hand, would refer to an available implementation of protection mechanism p such that component x
 281 is fully protected against attacks of the associated type. Evidently, the returned value may also be
 282 between 0 and 1, indicating partial protection. For example, in the case protection mechanism p adds a
 283 certain amount of random noise to a component to disguise real power behavior in order to prevent
 284 side-channel attacks based on power analysis [24], the level of added noise (which is a trade-off
 285 between power consumption and security) would determine the returned value of the function
 286 $\text{Protection-level}(p, x)$. To calculate S_x , we also consider the reciprocal of the utilization of component
 287 x . Here, the utilization refers to the fraction the computing/communication component x contributes
 288 to the overall computing or communication time of an application. The rationale behind this is that
 289 one could argue that those components (both processing and communication components) that are
 290 less active over time will be less susceptible to certain types of attacks, like side-channel attacks. To
 291 determine the overall security score of a specific design instance (i.e., a particular application mapping
 292 to a selected platform architecture), one could simply accumulate the S_x scores for all components x
 293 used in the design instance.

294 The above scoring example is by no means meant to be a complete and fully-fledged solution to
 295 the security scoring problem. It is merely meant to act as an illustration for the direction of thought
 296 as presented in this paper. Moreover, in the above scoring example, we also do not consider spatial

297 isolation as part of the security score. One direction to accomplish this, would e.g. to penalize the
298 mapping of multiple application tasks or communications to a single platform component.

299 The multifaceted, scoring-based security quantification methodology as outlined in this section
300 could provide a real innovation to system-level embedded system design as it would facilitate designers
301 to study the trade-offs between the performance, power consumption, cost, and secureness of design
302 instances during the early stages of design.

303 4. Related work

304 The need for recognizing security as a first-class citizen, next to traditional design objectives such
305 as performance, cost and power consumption, in the design of embedded systems is not new. For
306 example, quoting from [4]: "However, security is often misconstrued by embedded system designers as
307 the addition of features, such as specific cryptographic algorithms and security protocols, to the system.
308 In reality, it is a new dimension that designers should consider throughout the design process, along
309 with other metrics such as cost, performance, and power.". Nevertheless, the integration of security
310 aspects in the process of system-level design-space exploration of embedded systems has never really
311 got off the ground and is still a largely uncharted research ground. Only a few efforts exist that address
312 this problem but, at best, most of them provide partial solutions or solutions to very specific security
313 problems. For example, in [25], the evaluation of security protocols is integrated in the design process.
314 For instance, it rates the security of a system based on the probability of a hash collision. However, it
315 does not cover other types of attacks, such as timing attacks and power analysis. The authors of [26]
316 try to neutralize several types of side-channel attacks by means of spatial isolation in a DSE setting
317 but, again, they do not consider any other types of attacks / protection mechanisms. In [27], a small
318 number of attacker capabilities and corresponding requirements that refer to the CIA triad are defined
319 in the context of DSE. The problem of this approach is that it is not trivial to relate types of attacks to
320 those capabilities and requirements. The work of [28] incorporates security in system-level DSE by
321 first generating potential architecture configurations, after which an automated security analysis is
322 performed to check the generated configurations against designer-specified security constraints.

323 In all the above works of [26–28] security is modelled as a requirement in the DSE process, which
324 does not allow for studying actual trade-offs between performance, power consumption or cost in
325 relationship to secureness of a design.

326 An alternative approach for quantifying security is by means of a security risk assessment using
327 a specific attack model [29]. For example, [30], proposes an attack tree model to evaluate the user's
328 privacy risks associated with an Internet-of-Things eco system. They evaluate the potential risks
329 based on varying attack attributes, the probable considerations or preferences of an adversary, and
330 the varying computational resources available on a device. Research efforts like this are, however,
331 typically not focused on the process of (early) DSE.

332 To the best of our knowledge, only the works of [31,32] and [33] are similar to what we propose
333 in terms of aiming at incorporating security as an objective that can be traded off with other objectives
334 in the process of early DSE. In [31,32], the authors introduce an UML-based approach in which
335 application security requirements can be described together with security 'capabilities' – in addition to
336 other extra-functional aspects such as performance and power consumption – of system components
337 stored in a library. This then allows for a DSE process during which the application requirements
338 are matched with the component capabilities. The very recent work of [33] introduces a novel DSE
339 framework that allows for considering security constraints, in the form of attack scenarios, and attack
340 mitigations, in the form of security tasks. Based on the descriptions of the system's functionality
341 and architecture, possible attacks, and known mitigation techniques, the framework tries to find the
342 optimal design for an embedded system.

5. Conclusions

As embedded systems are becoming increasingly ubiquitous and interconnected, they attract a world-wide attention of attackers. This makes the security aspect during the design of these systems more important than ever. However, state-of-the-art design tools and methodologies for embedded systems do not consider system security as a primary design objective. This is especially true for the early design phases in which the process of design-space exploration is of eminent importance for performing trade-off analysis. Any security measures that may eventually be taken much later in the design process will then affect the already established design trade-offs with respect to the other, and more traditional, design objectives like system performance, power consumption and cost. It goes without saying that such a design practice leads to suboptimal products.

In this position paper, we therefore argued for security-aware design methods for embedded systems that will allow for the optimization of security aspects of embedded systems in their earliest design phases as well as for studying the trade-offs between security and the other design objectives such as performance, power consumption and cost. To this end, we proposed a multifaceted, scoring-based methodology for quantifying the degree of secureness of embedded system design instances, which would allow for incorporating these secureness quantifications in early design-space exploration of embedded systems. The proposed methodology has not yet been implemented, and would require further research to do so. We do hope, however, that this position paper will be a trigger for more wide-spread research on techniques that allow for incorporating security as a first-class citizen in the process of early design-space exploration of embedded systems.

References

1. Keutzer, K.; Newton, A.; Rabaey, J.; Sangiovanni-Vincentelli, A. System-level design: orthogonalization of concerns and platform-based design. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems* **2000**, *19*, 1523–1543.
2. Pimentel, A.D. Exploring Exploration: A Tutorial Introduction to Embedded Systems Design Space Exploration. *IEEE Design & Test* **2017**, *34*.
3. Singh, A.K.; Shafique, M.; Kumar, A.; Henkel, J. Mapping on multi-/many-core systems: survey of current and emerging trends. Proc. of the Design Automation Conference (DAC), 2013, pp. 1–10.
4. Ravi, S.; Raghunathan, A.; Kocher, P.; Hattangady, S. Security in embedded systems: Design challenges. *ACM Transactions on Embedded Computing Systems (TECS)* **2004**, *4*.
5. Gries, M. Methods for evaluating and covering the design space during early design development. *Integration, the VLSI Journal* **2004**, *38*, 131–183.
6. Niemann, R.; Marwedel, P. An Algorithm for Hardware/Software Partitioning Using Mixed Integer Linear Programming. *Design Automation for Embedded Systems* **1997**, *2*, 165–193.
7. Lukasiewicz, M.; Glass, M.; Haubelt, C.; Teich, J. Efficient symbolic multi-objective design space exploration. Proc. of the Asia and South Pacific Design Automation Conference (ASP-DAC), 2008, pp. 691–696.
8. Padmanabhan, S.; Chen, Y.; Chamberlain, R.D. Optimal design-space exploration of streaming applications. Proc. of the IEEE Int. Conference on Application-specific Systems, Architectures and Processors (ASAP), 2011, pp. 227–230.
9. Palesi, M.; Givargis, T. Multi-objective Design Space Exploration Using Genetic Algorithms. Proc. of the Int. Symposium on Hardware/Software Codesign (CODES), 2002, pp. 67–72.
10. Erbas, C.; Cerav-Erbas, S.; Pimentel, A.D. Multiobjective Optimization and Evolutionary Algorithms for the Application Mapping Problem in Multiprocessor System-on-Chip Design. *IEEE Trans. on Evolutionary Computation* **2006**, *10*, 358–374.
11. Thompson, M.; Pimentel, A.D. Exploiting Domain Knowledge in System-level MPSoC Design Space Exploration. *Journal of Systems Architecture* **2013**, *59*, 351–360.
12. Quan, W.; Pimentel, A.D. Towards Exploring Vast MPSoC Mapping Design Spaces using a Bias-Elitist Evolutionary Approach. Proc. of the Euromicro Digital System Design Conference (DSD), 2014, pp. 655–658.

- 393 13. Ravi, S.; Raghunathan, A.; Chakradhar, S. Tamper resistance mechanisms for secure embedded
394 systems. 17th International Conference on VLSI Design. Proceedings., 2004, pp. 605–611.
395 doi:10.1109/ICVD.2004.1260985.
- 396 14. B. Yang.; K. Wu.; R. Karri. Scan based side channel attack on dedicated hardware implementations
397 of Data Encryption Standard. 2004 International Conference on Test, 2004, pp. 339–344.
398 doi:10.1109/TEST.2004.1386969.
- 399 15. Papp, D.; Ma, Z.; Buttyan, L. Embedded systems security: Threats, vulnerabilities, and attack
400 taxonomy. 2015 13th Annual Conference on Privacy, Security and Trust (PST), 2015, pp. 145–152.
401 doi:10.1109/PST.2015.7232966.
- 402 16. Koopman, P. Embedded systems security. *IEEE Computer* **2004**, *37*.
- 403 17. Swenson, C. *Modern Cryptanalysis: Techniques for Advanced Code Breaking*; Wiley Publishing, 2008.
- 404 18. Pimentel, A.D.; Erbas, C.; Polstra, S. A Systematic Approach to Exploring Embedded System
405 Architectures at Multiple Abstraction Levels. *IEEE Trans. on Computers* **2006**, *55*, 99–112.
406 doi:http://doi.ieeecomputersociety.org/10.1109/TC.2006.16.
- 407 19. Erbas, C.; Pimentel, A.D.; Thompson, M.; Polstra, S. A Framework for System-level Modeling and
408 Simulation of Embedded Systems Architectures. *EURASIP Journal on Embedded Systems* **2007**.
- 409 20. Mohanty, S.; Prasanna, V.K.; Neema, S.; Davis, J. Rapid Design Space Exploration of Heterogeneous
410 Embedded Systems Using Symbolic Search and Multi-granular Simulation. Proc. of LCTES+SCOPES,
411 2002, pp. 18–27.
- 412 21. Mariani, G.; Brankovic, A.; Palermo, G.; Jovic, J.; Zaccaria, V.; Silvano, C. A correlation-based design
413 space exploration methodology for multi-processor systems-on-chip. Proc. of the Design Automation
414 Conference (DAC), 2010, pp. 120–125.
- 415 22. Jia, Z.J.; Bautista, T.; Núñez, A.; Thompson, M.; Pimentel, A.D. A System-level Infrastructure for
416 Multi-dimensional MP-SoC Design Space Co-exploration. *ACM Trans. on Embedded Computing Systems*
417 **2013**, *13*.
- 418 23. Whitman, M.E.; Mattord, H.J. *Principles of Information Security*, 4th ed.; Course Technology Press: Boston,
419 MA, USA, 2011.
- 420 24. Benini, L.; Macii, A.; Macii, E.; Omerbegovic, E.; Pro, F.; Poncino, M. Energy-Aware Design Techniques for
421 Differential Power Analysis Protection. Proceedings of the 40th Annual Design Automation Conference
422 (DAC), 2003, pp. 36–41.
- 423 25. Lin, C.W.; Zheng, B.; Zhu, Q.; Sangiovanni-Vincentelli, A. Security-Aware Design Methodology and
424 Optimization for Automotive Systems. *ACM Transactions on Design Automation of Electronic Systems* **2015**,
425 *21*.
- 426 26. Weichslgartner, A.; Wildermann, S.; Götzfried, J.; Freiling, F.; Glaundefined, M.; Teich, J.
427 Design-Time/Run-Time Mapping of Security-Critical Applications in Heterogeneous MPSoCs.
428 Proceedings of the 19th International Workshop on Software and Compilers for Embedded Systems
429 (SCOPES), 2016, pp. 153–162.
- 430 27. Stierand, I.; Malipatlolla, S.; Fröschle, S.; Stühling, A.; Henkler, S. Integrating the Security Aspect into
431 Design Space Exploration of Embedded Systems. Proceedings of the IEEE International Symposium on
432 Software Reliability Engineering Workshops, 2014, pp. 371–376.
- 433 28. Tan, B.; Biglari-Abhari, M.; Salcic, Z. An Automated Security-Aware Approach for Design of Embedded
434 Systems on MPSoC. *ACM Transactions on Embedded Computing Systems* **2017**, *16*.
- 435 29. Rocchetto, M.; Ferrari, A.; Senni, V. Challenges and Opportunities for Model-Based Security Risk
436 Assessment of Cyber-Physical Systems. In *Resilience of Cyber-Physical Systems. Advanced Sciences and*
437 *Technologies for Security Applications*; Flammini, F., Ed.; Springer, 2019.
- 438 30. Asif, W.; Ray, I.G.; Rajarajan, M. An Attack Tree Based Risk Evaluation Approach for the Internet of Things.
439 Proceedings of the 8th International Conference on the Internet of Things, 2018.
- 440 31. Ferrante, A.; Milosevic, J.; Janjšević, M. A security-enhanced design methodology for embedded systems.
441 Proc. of the International Conference on Security and Cryptography (SECRYPT), 2013, pp. 1–12.
- 442 32. Ferrante, A.; Kaitovic, I.; Milosevic, J. Modelling Requirements for Security-Enhanced Design of Embedded
443 Systems. Proc. of the Int. Conference on Security and Cryptography (SECRYPT), 2014, pp. 315–320.
- 444 33. Gressl, L.; Steger, C.; Neffe, U. Security Driven Design Space Exploration for Embedded Systems. Forum
445 for Specification and Design Languages (FDL), 2019, pp. 1–8.

446 © 2020 by the author. Submitted to *J. Low Power Electron. Appl.* for possible open access
447 publication under the terms and conditions of the Creative Commons Attribution (CC BY) license
448 (<http://creativecommons.org/licenses/by/4.0/>).