Uncertainty-Aware Reliability Analysis and Optimization

Faramarz Khosravi, Malte Müller, Michael Glaß, and Jürgen Teich
Hardware/Software Co-Design
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany
{faramarz.khosravi, glass, teich}@cs.fau.de, malte.mueller@fau.de

***Abstract*—Due to manufacturing tolerances and aging effects, future embedded systems have to cope with unreliable components. The intensity of such effects depends on *uncertain aspects* like environmental or usage conditions such that highly safety-critical systems are pessimistically designed for worst-case mission profiles. In this work, we propose to explicitly model the *uncertain characteristics* of system components, i. e. we model components using reliability functions with parameters distributed between a best and worst case. Since destructive effects like temperature may affect several components simultaneously (e. g. those in the same package), a *correlation* between uncertainties of components exists. The proposed uncertainty-aware method combines a formal analysis approach and a Monte Carlo simulation to consider uncertain characteristics and their different correlations. It delivers a holistic view on the system's reliability with best/worst/average-case behavior and also insights on variance and quantiles. But, existing optimization approaches typically assume design objectives to be single values or to follow a predefined distribution. As a remedy, we propose a dominance criterion for meta-heuristic optimization approaches like evolutionary algorithms that enables the comparison of system implementations with arbitrarily distributed characteristics. Our presented experimental results show that (a) the proposed analysis comes at low overhead while capturing existing uncertainties with sufficient accuracy, and (b) the optimization process is significantly enhanced when guiding the search process by additional aspects like variance and the 95 % quantile, delivering better system implementations as found by an uncertainty-oblivious optimization approach.**

I. Introduction

The continuous technology scaling allows the production of compact, high-performance components at low cost. But, their small device structures are increasingly susceptible to destructive effects like radiation or temperature—resulting in inherently unreliable components. This renders reliability one of today's main challenges in the design of embedded systems which requires both analysis and optimization techniques to compose reliable systems from such unreliable components.

Existing reliability analysis and optimization approaches typically assume that the parameters for the components' reliability are accurate and provide the optimization with a single value for the system's reliability. However, manufacturing tolerances and/or environmental and usage conditions have a growing and significant impact on the components' reliability, but may not be known a-priori—they are *uncertain*. In the design of highly safety-critical systems, worst-case mission profiles describe these conditions pessimistically such that the systems are pessimistically designed. To exemplify, an implementation of an H.264 encoder/decoder where hardware components (i. e. processing cores and accelerators) are exposed to varying conditions, modeled as Gaussian variations of their failure rates, is investigated. Figure 1 shows the resulting distribution of the system's Mean-Time-To-Failure (MTTF) including best/worst cases as well as the 95 % quantiles derived with our proposed uncertainty-aware analysis technique. For the dashed curve, there exists a significant difference between the bounds and the *relevant scenarios*—here given as 95 % quantiles, i. e. the interval around the mean where 95 % of all samples are located—shows a huge potential to design embedded systems less pessimistically with yet a high confidence. But, effects like temperature that cause aging in the form of, e. g. Negative-Bias Temperature Instability (NBTI) may affect several components at a time (e. g. those in the same package)—uncertainties are *correlated*. Consider again Fig. 1 and the solid curve: There, the components are exposed the same ambient temperature and the resulting distribution and quantiles are significantly different from those where no correlation in the uncertainty of the components' reliability was assumed (dashed).

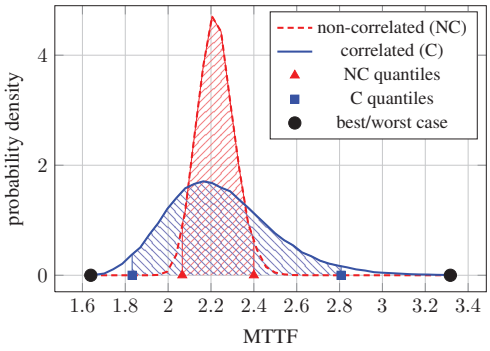

Fig. 1. The distribution of the MTTF of an H.264 implementation: The solid curve shows the reliability distribution for components with a correlated uncertainty vs. independent components given by the dashed curve. To highlight the significance of uncertainty and correlation consideration, theoretical bounds as well as 95 % quantiles are depicted.

In this work, we explicitly model the *uncertain characteristics* of system components. In particular, we describe the reliability of a component using arbitrary reliability functions (exponential, Weibull, custom) with parameters that are arbitrarily distributed between best- and worst-case bounds. These may be approximated by e. g. Gaussian variations around a mean, but also derived from measurements or field data. Moreover, a correlation between components can be specified to reflect that their uncertainty is induced by the same source like temperature, usage profile, or changing environments. We propose an uncertainty-aware reliability analysis approach that combines a formal analysis core with a Monte Carlo simulation, capable of considering uncertain characteristics and their different corre-

Supported in part by the German Research Foundation (DFG) as associated project *CRAU* (GL 819/1-2 & TE 163/16-2) of the priority program *Dependable Embedded Systems* (SPP 1500).

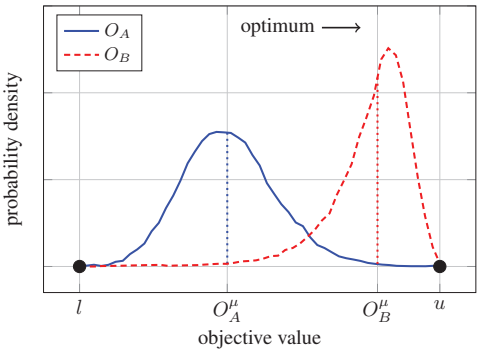

Fig. 2. Distribution of an uncertain objective for two implementations A and B. Note that both distributions have the same bounds, but different probability densities, and thus, different expected values and quantiles.

lations. While the analysis core can be realized by almost any existing technique (BDDs, fault trees, Markov chains, etc.), the simulation enables an extensible and affordable consideration of uncertainties with controllable execution time overhead.

The additional insight on the system's reliability delivered by our technique enables novel opportunities for Design Space Exploration (DSE) by providing notions on best/worst/average-case behavior and also variance and quantiles. Given reliability is not a single-valued design objective anymore, the proposed modeling requires an uncertainty-aware optimization technique in the DSE. Existing work on meta-heuristic optimization in the presence of uncertainty [1–3] typically assumes design objectives to be given as bounds and/or known distributions (e. g. uniform). Consider the two distributions given in Fig. 2: An optimization technique that only considers the boundary cases cannot differentiate between these distributions derived for two different implementations. In fact, investigating the distributions reveals one implementation (B) is significantly better in mean as well as quantiles such that an uncertainty-aware optimization technique should be able to differentiate the two. To realize this, the work at hand introduces a dominance criterion with a novel compare operator, applicable to various state-of-the-art meta-heuristic optimization approaches, that enables to consider arbitrarily distributed design objectives. The proposed approach can be applied adaptively: To keep the optimization fast while achieving the required accuracy, statistical properties are only considered when the comparison cannot find a preference based on the bounds. Moreover, overflows in the archive due to too many indifferent (*non-dominated*) implementations can be avoided by a proposed dynamic adaptation of the compare operator's sensitivity.

Our experiments will show that the proposed uncertainty-aware analysis can be implemented at reasonable overhead to be applied within a DSE even for complex embedded systems. The proposed uncertainty-aware DSE, considering mean values, quantiles, and/or variances as objectives, is superior to an existing DSE, delivering improved system implementations with comparable or better expected values and reduced uncertainty.

The rest of this paper is organized as follows: Section II reviews state-of-the-art. Section III explains the uncertainty model as well as the proposed uncertainty-aware reliability analysis while Sec. IV introduces the proposed uncertainty-aware DSE, in particular, the dominance criterion. Section V presents the experimental setup and evaluation results, and finally, Sec. VI concludes this work.

II. Related Work

Traditionally, uncertainty-aware analysis is performed mathematically to predict the distribution of uncertain objectives from variations in input parameters. Yin et al. [4] analyze parameter uncertainty in system reliability via Markov chains, and derive statistical properties like mean and quantiles using second-order and Gaussian approximations. But, their work does not consider arbitrariness and correlation among uncertainty distributions.

To deal with uncertain objectives in a multi-objective optimization, the work in [1] proposes a probabilistic dominance criterion to analytically compare implementation candidates whose objectives vary uniformly within given bounds. The work in [2] models uncertainty in reliability and cost objective functions as the lower and upper bounds of an interval with an unknown distribution, and use this boundary cases within a multi-objective optimization. The approach in [3] represents uncertainty in design objectives by expected value and bounds, and provides an uncertainty-aware optimization similar to the work in [1]. As discussed for Fig. 2, the consideration of bounds only may come at the drawback of pessimistic design and/or missing differentiation of implementation candidates with comparable bounds but different distributions.

In [5], the authors investigate the effects of uncertainty during reliability analysis on a redundancy-hardening allocation problem in series-parallel systems. They linearize the system's reliability function and use integer linear programming to maximize the expected value and minimize the variance of the system reliability. But, this work ignores the arbitrariness and correlation in the uncertainty distributions, and is only applicable to series-parallel systems. The authors of [6] deal with uncertainty as the lack of knowledge about the exact effects of architectural decisions on objective values, and represent it using a triangular fuzzy representation including its anticipated, optimistic, and pessimistic values to compare different implementations. Their approach, however, may consider two non-overlapping implementations being indifferent, resulting in archive overflow. Meedeniya et al. propose a reliability evaluation method to deal with non-correlated uncertainty distributions in [7], and integrate it within a multi-objective optimization in [8]. Their uncertainty analysis is based on sampling from uncertain parameters and evaluating the system with sampled parameters to obtain statistical information, e. g. percentiles, as design objectives. During the optimization, each objective function, based on its criticality, is statically represented by a single statistical property, e. g. the 5th percentile. Compared to their approach, the work at hand considers uncertainty correlation during the analysis and supports an adaptive dominance criterion to avoid archive overflow and performance overhead during the optimization.

III. Uncertainty-Aware Reliability Analysis

In this section, we introduce the proposed uncertainty-aware reliability analysis approach. After giving an outline of the overall flow, we present the modeling of components' uncertainties and their correlations before the section is concluded by an investigation of the confidence level of the employed Monte Carlo simulation.

A. Overall Analysis Flow

The general idea of the proposed analysis is given in Fig. 3. At the core, the approach may employ various existing relia-

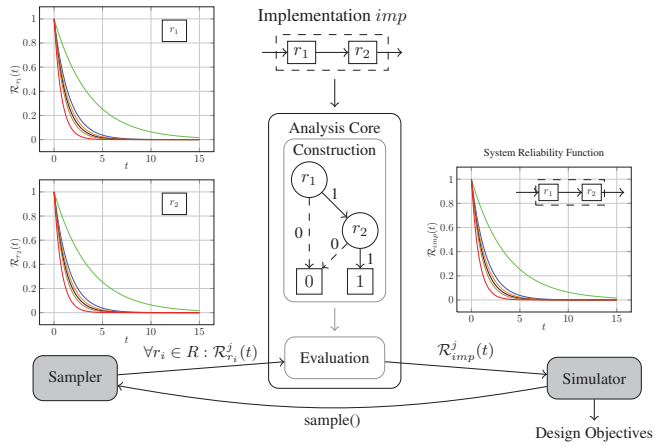

Fig. 3. The proposed uncertainty-aware reliability analysis: The *sampler* applies samples of reliability functions $\mathcal{R}(t)$ for each component r_i in the implementation imp to the existing reliability analysis core, here, a BDD. The core constructs a structure function representing imp, and re-evaluates it for the sampled reliability functions of the components, i. e. $\mathcal{R}^j_{r_i}(t)$s, in order to provide samples for the reliability functions of imp, i. e. $\mathcal{R}^j_{imp}(t)$. The desired statistical properties are then obtained for the design objectives and sent to the optimization.

bility analysis techniques that typically require the reliability function $\mathcal{R}_r(t)$ of each component $r \in R$ at time t, and deliver the system implementation's reliability function $\mathcal{R}_{imp}(t)$.[1] In the concrete case, a BDD-based approach is shown, modeling an implementation consisting of two components in series. This formal analysis core is employed by a Monte Carlo-based *sampler* which samples reliability functions from each component and applies them to the analysis core. The gathered sample reliability function $\mathcal{R}^j_{imp}(t)$ of sample j is collected by a statistical *simulator* that constructs the reliability distribution of the system implementation as a set of sampled reliability functions, i. e. $\Phi_{imp} = \bigcup_{j=1}^{n} \{\mathcal{R}^j_{imp}(t)\}$. The statistical simulator determines the number of required samples n to later obtain desired statistical properties like mean and quantiles from Φ_{imp} with a guaranteed confidence level which we will employ during the DSE presented in Sec. IV. To exemplify, for each sample $\mathcal{R}^j_{imp}(t)$ in Φ_{imp}, the MTTF can be calculated via the following integration:

$$\text{MTTF} = \int_0^\infty t \cdot f^j(t)\, dt \text{ with } f^j(t) = \frac{d(1 - \mathcal{R}^j_{imp}(t))}{dt} \quad (1)$$

Based on each sample's MTTF, design objectives like best/worst-case MTTF, expected MTTF, etc. can be derived.

At this point, we need to introduce a model for the uncertain reliability functions of components that enables a sampling. Afterwards, we introduce the proposed modeling of correlation.

B. Uncertainty Modeling

We propose to model the reliability characteristics in the presence of uncertainty $\mathcal{U}_r$ of a component r using arbitrary reliability functions $\mathcal{R}_r(t)$ that are distributed within a given lower $\mathcal{R}^l_r(t)$ and upper bound $\mathcal{R}^u_r(t)$ reliability function, i. e.

$$\mathcal{U}_r = \left[\mathcal{R}^l_r(t), \mathcal{R}^u_r(t)\right]. \quad (2)$$

[1]The standard definition of $\mathcal{R}(t)$ with $\mathcal{R} : \mathbb{R}^+_0 \to [0,1]$, $\mathcal{R}(0) = 1$, $\lim_{t\to\infty} \mathcal{R}(t) = 0$, and $\forall t, t' \in \mathbb{R}^+_0, t \le t' : \mathcal{R}(t) \ge \mathcal{R}(t')$ is assumed.

In particular, the outlined *sample* function takes $\mathcal{U}_r$ as input and delivers a sampled reliability function $\mathcal{R}^j_r(t)$ for which holds $\forall t \in \mathbb{R}^+_0 : \mathcal{R}^l_r(t) \le \mathcal{R}^j_r(t) \le \mathcal{R}^u_r(t)$. Note that the *sample* routine has to ensure that the sampled reliability functions follow the intended distribution within the given bounds. The sampler may consider arbitrary distributions, in particular the well-known, continuous distributions also used in [7] such as uniform, Gaussian, or Beta distributions to model the uncertainty induced by ambient temperature.

In practice, many component reliability functions are derived from measurements that are fitted to closed-form reliability functions with exponential ($\mathcal{R}_r(t) = e^{-\lambda_r \cdot t}$) and Weibull ($\mathcal{R}_r(t) = e^{-\lambda_r \cdot t^{\beta_r}}$) distributions being frequently used. As can be seen, these distributions are parametrized e. g. with the failure rate λ. Here, the proposed uncertainty model $\mathcal{U}_r$ can be refined to employ a set of uncertain parameters P_r, distributed within the bounds $[P^l_r, P^u_r]$. Given $[P^l_r, P^u_r]$, the sampler takes a sample from each parameter $p_r \in P_r$ and—together with the assumed closed-form reliability function—constructs a sample reliability function $\mathcal{R}^j_r(t)$. To exemplify, assuming an exponential distribution with given bounds $[\lambda^l_r, \lambda^u_r] = [0.01, 0.0075]$[2], a sample j of 0.0088 could be derived which delivers an $\mathcal{R}^j_r(t) = e^{-0.0088 \cdot t}$.

C. Consideration of Uncertainty Correlations

To model correlation among the uncertainty in components' reliability functions, we first determine if different functions are exposed to the same sources of uncertainty and are, thus, subject to correlative variations. For the example implementation in Fig. 3, the reliability functions $\mathcal{R}_{r_1}(t)$ and $\mathcal{R}_{r_2}(t)$ are considered correlated if they are exposed to the same ambient heat (e. g. they are in the same package); otherwise, they will be treated as independent distributions. Therefore, we define a *correlation group* which is a simple set of components being affected by the same source of uncertainty. Since we have full control over the sampling routine, we can even consider $\mathcal{U}_r$s being included in multiple correlation groups.[3] This is achieved by the following routine: Whenever a sample for Φ_{imp} is required, the samples from independent reliability functions are constructed using independently sampled parameters first. Then, the sample for a reliability function in a correlation group g is constructed using the q^{th}_g quantile of its uncertain parameters, where $q_g \in [0,1]$ is a uniformly distributed random number and applies to all reliability functions in g. Therefore, among different samples for Φ_{imp}, the uncertain parameters of the reliability functions in g vary together, and their variations are independent of those of other correlation groups.

D. Confidence Adaptation

The derived system's reliability distribution should guarantee a certain level of confidence since we have to compare different implementation candidates during a DSE. Given a required confidence $1 - \alpha$ and a minimum number of samples n, the maximum estimation error e can be calculated:

$$e = z_{(1-\frac{\alpha}{2});m} \times \frac{\sigma}{\sqrt{n}} \quad (3)$$

[2]Note that a larger value for the failure rate λ delivers a lower (worse) reliability function, thus the lower bound is given as the higher λ value.

[3]However, in this work we assume each $\mathcal{U}_r$ to be either independent or in only one correlation group.

where $z_{(1-\frac{\alpha}{2});m}$ is the $1-\frac{\alpha}{2}$ quantile of the Student's-T distribution function with $m = n-1$ degrees of freedom. Thus, a confidence interval $I = [\hat{\mu} - e, \hat{\mu} + e]$ for the precise mean value μ can be determined such that $P(\mu \in I) \geq 1 - \alpha$, where $\hat{\mu}$ is the estimated expected value. If the minimum number of samples $n_{\min}$ is unknown in advance, it can be obtained for a maximum acceptable error $e_{\max}$ and $\alpha_{\max}$ as follows:

$$n_{\min} = \frac{z^2_{(1-\frac{\alpha_{\max}}{2});m} \times \sigma^2}{e^2_{\max}} \tag{4}$$

We determine the number of samples based on the overheads of the Monte Carlo simulation on the time complexity of the analysis. Given the average execution times for the construction t_c and evaluation t_e of the analysis core, a single analysis takes $t_a = t_c + t_e$. For n samples, the analysis time equals $t_a(n) = t_c + n \times t_e$ since the analysis core typically employs structures (BDDs, fault trees, Markov chains) that can be re-evaluated by assigning new values to its variables. Using the inequation $t_a(n) \leq (1 + \delta_a) \times t_a$ and the upper bound δ_a for the analysis time overhead, the maximum n can be calculated as follows:

$$n_{\max} \leq 1 + \frac{\delta_a \times t_a}{t_e} \tag{5}$$

Our experiments in Sec. V show that we can achieve an acceptable confidence level for uncertain objectives at an analysis overhead suitable for DSE. Moreover, it might not be necessary to have precise evaluation and comparison of design objectives in early DSE steps while the accuracy of the archive may be crucial at the end of optimization, thus, the sample size can be dynamically adapted to further reduce the analysis overhead.

IV. Uncertainty-Aware Optimization

In this section, we propose an uncertainty-aware framework which enables existing optimization approaches to consider uncertainty in the design objectives of system implementations. We follow the state-of-the-art in embedded system DSE and employ multi-objective meta-heuristic optimization techniques as the underlying optimization core. These heuristics rely on so-called *dominance criteria* that compare different implementation candidates to (a) select which candidates to keep and vary for the next iteration and (b) store the best implementation candidates in an *archive*. To enable an uncertainty-aware optimization with design objectives (reliability, costs, etc.) being arbitrarily distributed within bounds, we propose a novel compare operator which compares uncertain objective values in stages, combining the consideration of bounds and statistical properties like expected values and quantiles. After introducing the proposed dominance criterion, we present a dynamic adaptation technique for the compare operator's sensitivity applied during the optimization which adapts itself to the current number of implementation candidates in the archive. The latter avoids that too many indifferent implementation candidates crowd the archive and, thus, deteriorate the optimization quality.

A. Uncertainty-Aware Dominance Criterion

For the case of each design objective being a single value included in the objective vector $\mathcal{O}(imp) = (O_1, \ldots, O_v)$ of an implementation imp with v objectives, the dominance of two implementation candidates imp and imp' is given as follows:

$$imp \succ imp' \text{ iff } \begin{array}{l} \forall i = 1, \ldots, v : O_i(imp) \geq O_i(imp') \\ \wedge\, \exists j, 1 \leq j \leq v : O_j(imp) > O_j(imp') \end{array} \tag{6}$$

Without loss of generality, we assume the objective to be maximized. Here, the compare operator $>$ in $O_i(imp) > O_i(imp')$ simply compares two exact values of an objective. The work in [1], extends the compare operator of a specific objective value to consider lower and upper bounds and basically judges the *overlap* of the resulting intervals. This, however, has the problem outlined in Fig. 2 where intervals with a large overlap cannot be differentiated although their distribution may significantly differ. Since the proposed techniques are applied to each objective, we will—for the sake of simplicity—refer to the objective value $O_i(A)$ of objective O_i of an implementation A simply as O_A and O, respectively.

We extend the idea in [1] towards a *three-stage compare operator* for each uncertain design objective: (1) If the intervals of an objective O (given by the lower bound O^l and upper bound O^u) of two implementations do not overlap, one is trivially better ($>$) than the other. (2) If the intervals overlap, an *average criterion* based on the estimated mean, mode, or median is used, judging whether one average O^{avg} is preferred with respect to a configurable sensitivity threshold ε. (3) If the average criterion does not differentiate, a *spread criterion* compares objectives with respect to their deviation, e. g. standard deviation, variance, or quantiles and judges whether one is better. This three-stage compare operator is detailed in Algorithm 1 for the case of comparing an uncertain objective O of two implementations A and B. In case none of the three stages determines that one objective is better, the objectives are considered equal ($=$).

For the average criterion, a configurable threshold value ε determines if the difference of the considered objective values is significant with respect to the given objective bounds. This threshold enables to control the sensitivity of the comparison. The value for ε may be adjusted to (a) the criticality of design objectives and, thus, required precision, (b) the optimization phase, e. g., in accepting smaller differences (smaller ε) in the beginning due to the *explorative* character while demanding more significant differences (greater ε) in the later *exploitation* phase where implementation candidates with comparable objectives should not be discarded recklessly, and (c) the number of non-dominated implementation candidates to avoid crowds in the archive. To establish a range-independent ε, normalized copies of objective values should be used. Thus, the value for objective O_A is considered better than O_B if the following holds:

$$\frac{O^{avg}_A - O^{avg}_B}{max(O^u_A, O^u_B) - min(O^l_A, O^l_B)} > \varepsilon \tag{7}$$

Here, ε is initialized to a $\varepsilon_{min} \geq 0$ and increased stepwise along with the optimization phase up to a $\varepsilon_{max} < 1$ to provide a more rigorous comparison.[4] In case of an overflow in the archive, ε is decreased and is used to update the archive (all its candidates are compared with each other) until the archive fits its capacity—some candidates are found dominated. The values for ε_{min} and ε_{max} is determined for each design objective individually based on its criticality.

The proposed spread criterion determines that the objective value with lower spread is better than the other:

$$dev(O_A) < dev(O_B) \Rightarrow O_A > O_B \tag{8}$$

[4]Note that $\varepsilon = 0$ always prefers the objective with better expected value, whereas $\varepsilon = 1$ would render the average criterion ineffective since the left-hand side of Eq. (7) is always less than one.

Algorithm 1 The Proposed Three-Stage Compare Operator

procedure COMPARE(O_A, O_B)
Require: Maximization objective functions O_A and O_B
Ensure: If O_A is better ($>$)/worse ($<$) than, or equal ($=$) to O_B.
  if $O_A^l > O_B^u$ **then return** $O_A > O_B$
  else if $O_A^u < O_B^l$ **then return** $O_A < O_B$
  else
    if $\frac{O_A^{avg} - O_B^{avg}}{max(O_A^u, O_B^u) - min(O_A^l, O_B^l)} > \varepsilon$ **then**
      return $O_A > O_B$
    else if $\frac{O_B^{avg} - O_A^{avg}}{max(O_A^u, O_B^u) - min(O_A^l, O_B^l)} > \varepsilon$ **then**
      return $O_A < O_B$
    else
      if $dev(O_A) < dev(O_B)$ **then return** $O_A > O_B$
      else if $dev(O_A) > dev(O_B)$ **then return** $O_A < O_B$
      end if
    end if
  end if
  return $O_A = O_B$
end procedure

Note that all these comparison criteria are calculated using the samples from each objective's distribution. Given a set of n samples for an uncertain design objective O, its variance can be calculated as follows:

$$O^{\sigma^2} = \frac{1}{n}\sum_{i=j}^{n}\left(O^j - O^\mu\right)^2 \text{ where } O^\mu = \frac{1}{n}\sum_{i=j}^{n} O^j \qquad (9)$$

With O^μ denoting the expected value of the distribution. Furthermore, to calculate the q^{th} quantile of the samples, we use the *inverse empirical distribution function* which traverses the samples in the ascending order and returns the very first sample after the $q\%$ smallest samples.

V. Experimental Setup and Evaluations

In this section, we investigate the overhead imposed by the uncertainty-aware reliability and the achievable enhancements of the optimization quality. Table I summarizes the characteristics (number of resources, tasks, mapping, and resource variants) of each investigated test problem, i. e., a real-world specification of an H.264 encoder/decoder and two synthetic specifications of different size and complexity. The variants of a resource model so-called *hardening levels* where at higher cost, reliability-enhancing techniques are already inbuilt in the resource. Moreover, the number of correlation groups, i. e. the number of groups wherein resources are subject to the same uncertainty correlation, is given in parentheses in the resources column. All experiments are carried out on a desktop computer with a 3.40 GHz CPU.

A. Analysis Overhead

Figure 4 shows the average construction time t_c and evaluation time of one sample t_e for our test cases on a logarithmic scale. As can be seen, t_c is significantly larger than t_e, particularly when the degree of redundancy increases. This results from the BDD requiring exponential space in the number of variables (worst-case) while a single evaluation (sample) can be done in linear time. While in existing approaches the BDD is typically constructed and evaluated once, we reuse the same BDD for multiple samples. For simple test cases, t_c is about one to two orders of magnitude larger than t_e such that e. g. $1,000$ samples increase the analysis time by a factor of 10 x—with an overall analysis time around 1 s, this is completely feasible for DSE. For more complex examples, t_c is more than three orders of magnitude larger than t_e such that e. g. $1,000$ samples do not even double analysis time. From these results, it can be concluded that due to the construction overhead of the formal analysis core, the proposed analysis is still suitable for DSE while for complex examples, the overhead is minor compared to the gained insight.

TABLE I. TEST PROBLEMS

problem	resources	tasks	mappings	variants
H.264	15 (3)	66	275	4–6
Synth-I	25 (5)	56	261	4–6
Synth-II	50 (5)	101	592	4–6

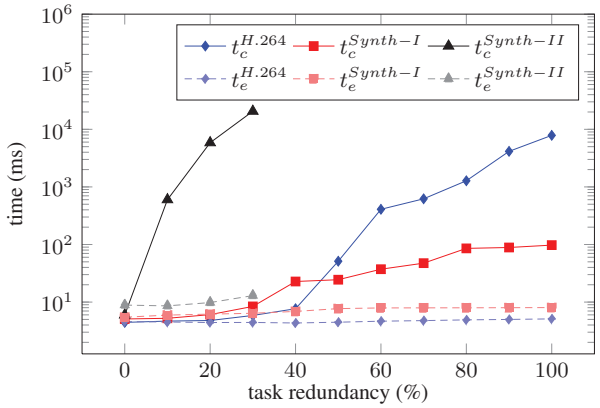

Fig. 4. Construction (t_c) and evaluation times (t_e) of the BDDs for our test problems and implementations with different redundancy levels. Note that the times are given in log-scale.

B. Optimization Quality

The proposed uncertainty-aware reliability analysis and optimization are integrated in a DSE which uses the OPT4J optimization framework [9]. We use the well-known *Nondominated Sorting Genetic Algorithm II* (NSGA-II) [10] with a population size of 25 and 100 generations and we take 1000 samples[5] during each uncertainty-aware analysis. The considered system characteristics are reliability (i. e. MTTF)—to showcase that our optimization approach is suitable for multiple objectives—monetary costs with component costs being subject to uncertainty. For comparison, we consider a common approach (*comm.*) that uses the mean value of the MTTF and costs with a standard dominance criterion. For the proposed uncertainty-aware optimization, we use the mean value as the average criterion and investigate (I) the variance (σ^2) and (II) the 95 % quantiles ($q_{0.95}$) as the spread criteria. Note that both approaches require our proposed analysis to determine the correct mean MTTF even for the single-value objective approach due to the effect of correlated uncertainties. Thus, the common approach comes at the same analysis overhead as our proposed approach. Table II summarizes the optimization results for the test cases using the well-known ϵ-dominance [11]—not to be confused with the proposed ε-threshold—as well as the average standard

[5]The resulting maximum estimation error is $e < 0.001$ with the confidence level of 95 % ($z_{(1-\frac{\alpha}{2});m} \approx 1.96$ in Eq. (3)).

TABLE II. THE COMPARISON OF THE PROPOSED DOMINANCE CRITERION WITH THE SECONDARY DECISIONS BASED ON 95 % QUANTILE ($q_{0.95}$) AND VARIANCE (σ^2), TO THE MEAN-BASED CRITERION (COMM.) BASED ON THE ϵ-DOMINANCE AND AVERAGE STANDARD DEVIATION FOR THE IMPLEMENTATION CANDIDATES FOUND BY THE OPTIMIZATION.

problem	criterion	ϵ_{dom}	σ_{avg}^{MTTF}	σ_{avg}^{cost}
H.264	comm.	0.246	0.141	0.041
	$q_{0.95}$	**0.020**	0.057	0.016
	σ^2	0.261	**0.049**	**0.012**
Synth-I	comm.	0.025	0.043	0.003
	$q_{0.95}$	**0.013**	**0.026**	**0.001**
	σ^2	0.023	0.027	0.002
Synth-II	comm.	0.128	0.016	**0.001**
	$q_{0.95}$	0.139	**0.015**	**0.001**
	σ^2	**0.020**	0.017	0.002

deviation of the objectives for the implementation candidates. It shows that the proposed uncertainty-aware dominance criterion outperforms the uncertainty-oblivious criterion with respect to the quality as well as the confidence, i. e., reduced uncertainty and enhanced worst cases, of the found implementations. To illustrate the effect of the proposed optimization with respect to the design objectives, Fig. 5 shows the final implementation candidates for H.264 (top) and Synth-I (bottom), comparing the comm.- and $q_{0.95}$-based optimizations. Here, the points and the boxes respectively indicate the mean value and the best/worse cases for each objective value of an implementation. As shown, the proposed approach delivers in all areas comparable or enhanced mean values or even finds implementation candidates of very high reliability or very low costs that the common approach does not find at all. In all test cases, the proposed approach significantly reduces the level of uncertainty for each implementation candidate: This not only enhances the worst cases—connected by dashed lines to highlight this effect in the figures—but may also enhance the confidence with which such system implementations may be integrated in larger systems.

VI. CONCLUSION

The intensity of destructive effects, e. g. aging, on device reliability cannot be precisely predicted since it depends, e. g., on environmental conditions and usage profiles, and thus, has to be considered an uncertain parameter in the design phase. Such uncertainties in different parts of the system can be correlated because destructive effects such as temperature may affect several system components simultaneously. This paper presents an approach for uncertainty-aware reliability analysis which explicitly models uncertainty in system components as reliability functions with parameters distributed within best/worst-case bounds. It combines a formal analysis core with a Monte Carlo simulation to consider correlated and arbitrarily distributed uncertain parameters and delivers the distribution of the system's reliability. To enable the comparison and optimization of arbitrarily distributed objective values, a novel dominance-criterion using a three-stage compare operator is presented that benefits from statistical properties such as quantiles. The experimental evaluation shows that while imposing acceptable computational overhead to the analysis, the proposed approach enhances the optimization regarding the expected values, variances and worst cases of the optimized objectives.

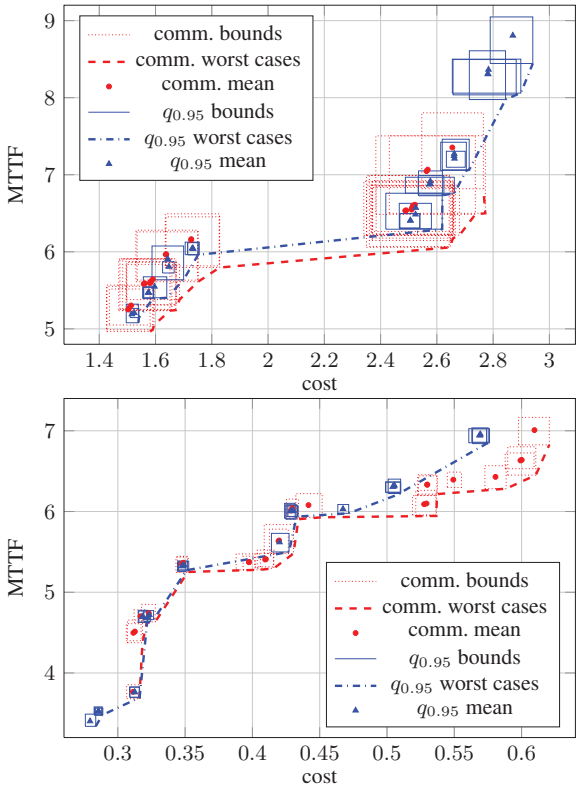

Fig. 5. Resulting Pareto fronts for optimizing MTTF and cost for H.264 (top) and Synth-I (bottom) using the proposed criterion based on mean and 95 % quantiles, and the simple mean criterion. Shown are mean values (dots), the best/worst cases as boxes, as well as lines outlining the worst-case fronts.

REFERENCES

[1] J. Teich, "Pareto-front exploration with uncertain objectives," in *Proc. of EMO*, 2001, pp. 314–328.

[2] P. Limbourg, "Multi-objective optimization of problems with epistemic uncertainty," in *Proc. of EMO*, 2005, pp. 413–427.

[3] N. Esfahani and S. Malek, "Uncertainty in self-adaptive software systems," in *Software Engineering for Self-Adaptive Systems II*, 2013, pp. 214–238.

[4] L. Yin, M. Smith, and K. Trivedi, "Uncertainty analysis in reliability modeling," in *Proc. of RAMS*, 2001, pp. 229–234.

[5] H. Tekiner-Mogulkoc and D. Coit, "System reliability optimization considering uncertainty: Minimization of the coefficient of variation for series-parallel systems," *IEEE Transactions on Reliability*, vol. 60, no. 3, pp. 667–674, 2011.

[6] N. Esfahani, K. Razavi, and S. Malek, "Dealing with uncertainty in early software architecture," in *Proc. of FSE*, 2012, pp. 1–4.

[7] I. Meedeniya, I. Moser, A. Aleti, and L. Grunske, "Architecture-based reliability evaluation under uncertainty," in *Proc. of QoSA*, 2011, pp. 85–94.

[8] I. Meedeniya, A. Aleti, and L. Grunske, "Architecture-driven reliability optimization with uncertain model parameters," *Journal of Systems and Software*, vol. 85, no. 10, pp. 2340–2355, 2012.

[9] M. Lukasiewycz, M. Glaß, F. Reimann, and J. Teich, "Opt4j - a modular framework for meta-heuristic optimization," in *Proc. of GECCO*, 2011, pp. 1723–1730.

[10] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II," in *Proc. of PPSN*, 2000, pp. 849–858.

[11] E. Zitzler, L. Thiele, M. Laumanns, C. Fonseca, and V. Da Fonseca, "Performance assessment of multiobjective optimizers: An analysis and review," *IEEE Transactions on Evolutionary Computation,*, vol. 7, no. 2, pp. 117–132, 2003.