



Contents lists available at ScienceDirect

The Journal of Systems & Software

journal homepage: www.elsevier.com/locate/jss

Efficient synthesis of robust models for stochastic systems

Radu Calinescu^a, Milan Češka^b, Simos Gerasimou^{*,a}, Marta Kwiatkowska^c, Nicola Paoletti^d^a Department of Computer Science, University of York, UK^b Faculty of Information Technology, Brno University of Technology, Czechia^c Department of Computer Science, University of Oxford, UK^d Department of Computer Science, Stony Brook University, USA

ARTICLE INFO

Keywords:

Software performance and reliability engineering
 Probabilistic model synthesis
 Multi-objective optimisation
 Robust design

ABSTRACT

We describe a tool-supported method for the efficient synthesis of parametric continuous-time Markov chains (pCTMC) that correspond to *robust designs* of a system under development. The pCTMCs generated by our ROBust DEsign Synthesis (RODES) method are resilient to changes in the system's operational profile, satisfy strict reliability, performance and other quality constraints, and are Pareto-optimal or nearly Pareto-optimal with respect to a set of quality optimisation criteria. By integrating sensitivity analysis at designer-specified tolerance levels and Pareto optimality, RODES produces designs that are potentially slightly suboptimal in return for less sensitivity—an acceptable trade-off in engineering practice. We demonstrate the effectiveness of our method and the efficiency of its GPU-accelerated tool support across multiple application domains by using RODES to design a producer-consumer system, a replicated file system and a workstation cluster system.

1. Introduction

Robustness is a key characteristic of both natural (Kitano, 2004) and human-made (Phadke, 1995) systems. Systems that cannot tolerate change are prone to frequent failures and require regular maintenance. As such, engineering disciplines like mechanical and electrical engineering treat robustness as a first-class citizen by designing their systems based on established tolerance standards (e.g. International Organization for Standardization, 2010; International Organization for Standardization, 2013). By comparison, software engineering is lagging far behind. Despite significant advances in software performance and reliability engineering (Balsamo et al., 2004; Bondy, 2014; Becker et al., 2009; Fiondella and Puliafito, 2016; Stewart, 2009; Woodside et al., 2014), the quality attributes of software systems are typically analysed for point estimates of stochastic system parameters such as component service rates or failure probabilities. Even the techniques that assess the sensitivity of quality attributes to parameter changes (e.g. Gokhale and Trivedi, 2002; Lo et al., 2005; Huang and Lyu, 2005; Kamavaram and Goseva-Popstojanova, 2003; Filieri et al., 2016) focus on the analysis of a given design at a time instead of systematically designing robustness into the system under development (SUD).

To address these limitations, we propose a tool-supported method for the efficient synthesis of parametric continuous-time Markov chains (pCTMCs) that correspond to robust SUD designs. Our ROBust DEsign Synthesis (RODES) method generates sets of pCTMCs that:

- (i) are resilient to pre-specified *tolerances* in the SUD parameters, i.e., to changes in the SUD's operational profile;
- (ii) satisfy strict performance, reliability and other quality constraints;
- (iii) are Pareto-optimal or nearly Pareto optimal with respect to a set of quality optimisation criteria.

RODES comprises two steps. In the first step, the SUD design space is modelled as a pCTMC with discrete and continuous parameters corresponding to alternative system architectures and to ranges of possible values for the SUD parameters, respectively. In the second step, a multi-objective optimisation technique is used to obtain a set of low-sensitivity, Pareto-optimal or nearly Pareto-optimal SUD designs by fixing the discrete parameters (thus selecting specific architectures) and restricting the continuous parameters to bounded intervals that reflect the pre-specified tolerances. The designs that are slightly suboptimal have the advantage of a lower sensitivity than the optimal designs with similar quality attributes, achieving a beneficial compromise between optimality and sensitivity. A *sensitivity-aware Pareto dominance relation* is introduced in the paper to formally capture this trade-off.

Fig. 1 shows the differences between a traditional Pareto front, which corresponds to a fixed SUD operational profile, and a sensitivity-aware Pareto front generated by RODES, which corresponds to a SUD operational profile that can change within pre-specified bounds. Accordingly, the designs from the RODES sensitivity-aware Pareto front are bounded regions of quality-attribute values for the system. The size

* Corresponding author.

E-mail address: simos.gerasimou@york.ac.uk (S. Gerasimou).

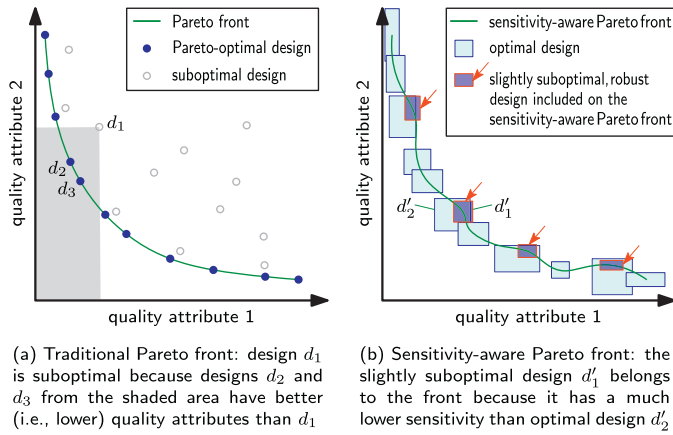


Fig. 1. Traditional Pareto front (a) versus sensitivity-aware Pareto front (b) for two quality attributes that require minimisation (e.g., response time and probability of failure).

and shape of these regions convey the sensitivity of the synthesised designs to parameter changes within the pre-specified tolerances. Small quality-attribute regions correspond to particularly robust designs that cope with variations in the system parameters without exposing users to significant changes in quality attributes. These designs require reduced maintenance, and can be implemented using high-variability components that are cheaper to develop or obtain off-the-shelf than low-variability components. Large quality-attribute regions from a RODES Pareto front—while still the most robust for the quality attribute trade-offs they correspond to—are associated with designs that are sensitive to SUD parameters variations. These designs may involve high maintenance and/or development costs, so they should only be used if justified by their other characteristics (e.g. desirable quality attribute trade-offs).

To the best of our knowledge, RODES is the first solution that integrates multi-objective stochastic model synthesis and sensitivity analysis into an end-to-end, tool-supported design method. As we show in detail in Section 7, the existing research addresses the challenges associated with design synthesis (e.g. Gerasimou et al., 2015; Martens et al., 2010) and sensitivity analysis (e.g. Gokhale and Trivedi, 2002; Lo et al., 2005; Huang and Lyu, 2005; Kamavaram and Goseva-Popstojanova, 2003; Filieri et al., 2016) separately. The main contributions of our paper are:

1. The extension of the notion of *parameter tolerance* from other engineering disciplines for application to software architecture.
2. The definitions of the parametric Markov chain synthesis problem and of the sensitivity-aware Pareto dominance relation for the synthesis of robust models for stochastic systems.
3. The RODES method for the generation of sensitivity-aware Pareto fronts by integrating multi-objective probabilistic model synthesis and precise *pCTMC* parameter synthesis.
4. A GPU-accelerated tool that implements the RODES method and is available preinstalled on an easy-to-use VirtualBox instance from our project website <https://www.github.com/gerasimou/RODES/wiki>.
5. A repository of case studies demonstrating the successful application of RODES to a replicated file system used by Google’s search engine, a cluster availability management system, and a producer-consumer system.

These contributions significantly extend our conference paper on robust model synthesis (Calinescu et al., 2017a) and the prototype probabilistic model synthesis tool (Calinescu et al., 2017b) in several ways. First, we provide a more detailed description of our solution, including a running example and new experimental results. Second, we

greatly improve the scalability of RODES by integrating the GPU-accelerated analysis of candidate designs into our prototype tool (Calinescu et al., 2017b). Third, we extend the experimental evaluation to demonstrate the impact of the GPU acceleration. Finally, we present an additional case study in which we apply RODES to a producer-consumer system, and we use the systems and models from our experiments to assemble a repository of case studies available on our project website.

The remainder of the paper is organised as follows. Section 2 introduces the RODES design-space modelling language and the formalism to specify quality constraints and optimisation criteria. Section 3 defines the sensitivity-aware dominance relation and introduces the parametric Markov chain synthesis problem. We then present our method for synthesising robust designs in the form of a sensitivity-aware Pareto set, and the GPU-accelerated tool RODES implementing the method in Sections 4 and 5, respectively. Finally, we evaluate our method within three case studies in Section 6, discuss related work in Section 7, and conclude the paper with a summary and future work in Section 8.

2. Modelling and specification language for probabilistic systems

This section formalises three key elements underpinning the formulation of the robust design problem: 1) the modelling of the design space of a SUD, 2) the specification of quality attributes and requirements, and 3) the sensitivity of a design.

2.1. Design space modelling

We use a *parametric continuous-time Markov chain* (*pCTMC*) to define the design space of a SUD. To this end, we extend the original *pCTMC* definition (Han et al., 2008), where only real-valued parameters determining the transition rates of the Markov chain are considered, and assume that a *pCTMC* also includes discrete parameters affecting its state space. Our definition captures the need for both discrete parameters encoding architectural structural information (e.g. by selecting between alternative implementations of a software component) and continuous parameters encoding configurable aspects of the system (e.g. network latency or throughput). As such, a candidate system design corresponds to a fixed discrete parameter valuation and to continuous parameter values from a (small) region.

Definition 1 (*pCTMC*). Let K be a finite set of real-valued parameters such that the domain of each parameter $k \in K$ is a closed interval $[k^{\perp}, k^{\top}] \subset \mathbb{R}$, and D a finite set of discrete parameters such that the domain of each parameter $d \in D$ is a set $T^d \subset \mathbb{Z}$. Let also $\mathcal{P} = \times_{k \in K} [k^{\perp}, k^{\top}]$ and $\mathcal{D} = \times_{d \in D} T^d$ be the continuous and the discrete *parameter spaces* induced by K and D , respectively. A *pCTMC* over K and D is a tuple

$$\mathcal{C}(\mathcal{P}, \mathcal{D}) = (\mathcal{S}, \mathcal{S}_{init}, \mathcal{R}, L), \quad (1)$$

where, for any discrete parameter valuation $q \in \mathcal{D}$:

- $\mathcal{S}(q) = S$ is a finite set of *states*, and $\mathcal{S}_{init}(q) \in S$ is the initial state;
- $\mathcal{R}(q): S \times S \rightarrow \mathbb{R}[K]$ is a parametric rate matrix, where $\mathbb{R}[K]$ denotes the set of polynomials over the reals with variables in K ;
- $L(q): S \rightarrow 2^{AP}$ is a labelling function mapping each state $s \in S$ to the set $L(q)(s) \subseteq AP$ of atomic propositions that hold true in s .

A *pCTMC* $\mathcal{C}(\mathcal{P}, \mathcal{D})$ describes the uncountable set of continuous-time Markov chains (CTMCs) $\{\mathcal{C}(p, q) \mid p \in \mathcal{P} \wedge q \in \mathcal{D}\}$, where each $\mathcal{C}(p, q) = (\mathcal{S}(q), \mathcal{S}_{init}(q), \mathbf{R}(p, q), L(q))$ is the instantiated CTMC with transition matrix $\mathbf{R}(p, q)$ obtained by replacing the real-valued parameters in $\mathcal{R}(q)$ with their valuation in p .

In our approach we operate with *pCTMCs* expressed in a high-level modelling language extending the PRISM language (Kwiatkowska et al., 2011) which models a system as the parallel composition of a set of

modules. The state of a module is encoded by a set of finite-range local variables, and its state transitions are defined by probabilistic guarded commands that change these variables, and have the general form:

$$[\text{action}]\text{guard} \rightarrow e_1: \text{update}_{e_1} + \dots + e_n: \text{update}_{e_n} \quad (2)$$

In this command, *guard* is a Boolean expression over all model variables. If the guard evaluates to true, the arithmetic expression e_i , $1 \leq i \leq n$, gives the rate with which the update_{e_i} change of the module variables occurs. When *action* is present, all modules comprising commands with this action have to synchronise (i.e., to carry out one of these commands simultaneously) and the resulting rate of such synchronised commands is equal to the multiplication of the individual command rates. Atomic propositions are encoded with label expressions of the form:

$$\text{label } \text{"id"} = b \quad (3)$$

where *id* is a string that identifies the atomic proposition and *b* is a Boolean expression over the state variables.

We extend the PRISM language with the following constructs (adopted from Gerasimou et al., 2015) for specifying the parameters $k \in K$ and $d \in D$ from Definition 1:

$$\begin{aligned} &\text{evolvedouble } k \text{ [min. max]} \\ &\text{evolveint } d \text{ [min. max]} \\ &\text{evolvemodule } \textit{ComponentName} \end{aligned} \quad (4)$$

where $N > 1$ instances of the last construct (with the same component name) define N alternative architectures for a component, introducing the index (between 1 and N) of the selected architecture as an implicit discrete parameter.

As per Definition 1, continuous parameters can only appear in the transition rates (expressions e_1, \dots, e_n above).

Explicit discrete variables (declared using *evolve int*) can instead appear in any type-consistent expression.

The translation of models expressed in the extended PRISM language into the corresponding pCTMC is fully automatic and follows the probabilistic guarded command semantics described above. The discrete state space \mathcal{L} results from all possible valuations of explicit discrete variables and implicit discrete variables (different implementations of a module). For a fixed valuation $q \in \mathcal{L}$, the parametric PRISM model describes a fixed set of modules with a fixed set of finite-range variables, and thus the state space $\mathcal{S}_S(q)$ is given by the Cartesian product of the value ranges for these variables. In contrast, q determines also the parametric rate matrix $\mathcal{S}_R(q)$ and atomic propositions $L(q)$, as q can affect guards and updates of PRISM commands, as well as label expressions.

Example 1 (Producer-consumer model). As a running example, we consider a simple producer-consumer system with a two-way buffering, illustrated in Fig. 2. The pCTMC PRISM model, extended with the evolvable constructs from Definition 4 is shown in Fig. 3. The system comprises a producer generating requests with rate p_rate . Each request is being transferred to a consumer either via a slow buffer or via a fast buffer with probabilities 0.6 and 0.4, respectively (lines 14 and 15




Fig. 2. Two-way producer-consumer system.

```

1  ctmc
   // buffer capacities
2  const int slow_max = 31;
3  const int fast_max = 21;
4  cons double p_rate = 40; //request production rate
5  cons double s_rate = 30; //request transmission rate
   // packet transmission rates
   // trans. rate for the fast buffer is r_slow_rate+delta_rate
6  evolve double r_slow_rate [5..30]; //slow buffer
7  evolve double delta_rate [0..30]; //fast buffer
8  const double c_rate = 40; //packet consumption rate
   // no redirection
9  evolve module Buffer
   // is request sent to fast/slow buffer?
10 fast : [0..1] init 0;
11 slow : [0..1] init 0;
   // buffers
12 buffer_s : [0..slow_max] init 0;
13 buffer_f : [0..fast_max] init 0;
   // has consumer received the packet?
13 consumer : [0..1] init 0;
   //produce
14 [] (fast=0) -> p_rate*0.4 : (fast'=1);
15 [] (slow=0) -> p_rate*0.6 : (slow'=1);
   //send
16 [] (slow=1) & (buffer_s < slow_max) -> s_rate :
   (slow'=0) & (buffer_s' = buffer_s + 1);
17 [] (fast=1) & (buffer_f < fast_max) -> s_rate :
   (fast'=0) & (buffer_f' = buffer_f + 1);
   //receive
18 [] (consumer=0) & (buffer_s > 0) -> r_slow_rate :
   (consumer'=1) & (buffer_s' = buffer_s - 1);
19 [] (consumer=0) & (buffer_f > 0) -> (r_slow_rate+delta_rate)*0.95 :
   (consumer'=1) & (buffer_f' = buffer_f - 1);
   // fast buffer loses the packet
20 [lost] (consumer=0) & (buffer_f > 0) -> (r_slow_rate+delta_rate)*0.05 :
   (buffer_f' = buffer_f - 1);
   //consume
21 [consume] (consumer=1) -> c_rate : (consumer'=0);
22 endmodule
   // redirection
23 evolve module Buffer
   :
   //send
24 [] (slow=1) & (buffer_s < slow_max) -> s_rate *(1-buffer_s/(10.0*slow_max)) :
   (slow'=0) & (buffer_s' = buffer_s + 1);
25 [] (slow=1) & (buffer_s > 0) & (buffer_f < fast_max) ->
   s_rate *buffer_s/(10.0*slow_max) :
   (slow'=0) & (buffer_f' = buffer_f + 1);
26 [] (fast=1) & (buffer_f < fast_max) ->
   s_rate : (fast'=0) & (buffer_f' = buffer_f + 1);
   :
27 endmodule

```

Fig. 3. PRISM-RODES encoding of pCTMC model of a producer-consumer system with two-way buffering and redirection. In the second module only the commands that differ from the first module are reported.

in Fig. 3). The fast buffer transmits requests to the consumer faster than the slow buffer, but it has smaller capacity and is less reliable, as it loses packets with a 5% probability (line 20). We consider two alternative designs of the producer-consumer model that differ in the way that the two buffers manage the pending requests. More specifically we consider

1. a no-redirection design in which once a request is sent to either buffer, the packet is transmitted by that buffer to the consumer (lines 9–22);
2. a redirection design that enables the slow buffer to transmit requests to the fast buffer with a probability proportional to its occupancy (lines 23–27). In particular, redirection is disabled when the slow buffer is empty and has maximum rate when it is full and is equal to $s_rate/10$, where s_rate is the request transmission rate without redirection. In addition to these two alternative designs, the model has two continuous parameters, the packet transmission rate for the

slow buffer, $r_{\text{slow_rate}}$, and delta_rate , i.e. the transmission rate difference between fast and slow buffers. Notably, the rate of packet loss by the fast buffer is proportional to its transmission rate, meaning that the buffer becomes less reliable as its rate increases. We formally capture the above system model with its continuous parameters and alternative designs by a pCTMC $\mathcal{C}_{\text{PC}}(\mathcal{P}, \mathcal{D})$, where $\mathcal{P} = [5, 30] \times [0, 30]$ defines the domains for the continuous parameters $r_{\text{slow_rate}}$, and delta_rate , respectively, and $\mathcal{D} = \{1, 2\}$ defines the domain for the discrete parameter corresponding to the two alternative designs (i.e. modules).

Definition 2 (Candidate design). A candidate design of the pCTMC $\mathcal{C}(\mathcal{P}, \mathcal{D})$ from (1) is a pCTMC

$$\mathcal{C}(\mathcal{P}', \{q\}) = (\mathcal{D}'_S, \mathcal{D}'_{\text{init}}, \mathcal{D}'_R, L') \quad (5)$$

where $\mathcal{P}' = \times_{k \in K} [k^{\perp}, k^{\top}] \subseteq \mathcal{P}$, $q \in \mathcal{D}$, $\mathcal{D}'_S(q) = \mathcal{D}_S(q)$, $\mathcal{D}'_R(q) = \mathcal{D}_R(q)$, $\mathcal{D}'_{\text{init}}(q) = \mathcal{D}_{\text{init}}(q)$ and $L'(q) = L(q)$. The *tolerance* of the candidate design with respect to the real-valued parameter $k \in K$ is defined as

$$\gamma_k = \frac{k^{\top} - k^{\perp}}{2(k^{\top} - k^{\perp})}, \quad (6)$$

in line with the fact that the design restricts the value domain of k to the interval $[\bar{k} - \gamma_k(k^{\top} - k^{\perp}), \bar{k} + \gamma_k(k^{\top} - k^{\perp})]$, $\bar{k} = \frac{k^{\perp} + k^{\top}}{2}$.¹ For convenience, we will use the shorthand notation $\mathcal{C}(\mathcal{P}', q) \equiv \mathcal{C}(\mathcal{P}', \{q\})$ in the rest of the paper.

Example 2 (Candidate design). Consider the pCTMC $\mathcal{C}_{\text{PC}}(\mathcal{P}, \mathcal{D})$ from Example 1 and a single tolerance value $\gamma = 0.005$ for both continuous parameters $r_{\text{slow_rate}}$ and delta_rate . By (6), candidate designs have continuous parameter ranges of size $2\gamma(k^{\top} - k^{\perp}) = 0.25$ for $r_{\text{slow_rate}}$ and of size 0.3 for delta_rate . Two examples of valid candidate designs for the second module (redirection), obtained using our RODES synthesis method (see also results in Fig. 7), are pCTMCs $d_1 = \mathcal{C}_{\text{PC}}(\mathcal{P}', 2)$ and $d_2 = \mathcal{C}_{\text{PC}}(\mathcal{P}'', 2)$ where $\mathcal{P}' = [15.02, 15.27] \times [1.93, 2.23]$, $\mathcal{P}'' = [13.2, 13.45] \times [3.51, 3.81]$. The pCTMCs $d_3 = \mathcal{C}_{\text{PC}}(\mathcal{P}', 1)$ with is instead a valid candidate design for the first module (no redirection).

2.2. Quality attribute specification and requirements

We specify quality attributes over pCTMCs-defined design spaces using *continuous stochastic logic* (CSL) extended with reward operators (Kwiatkowska et al., 2007). Our focus is on timed properties of pCTMCs expressed by the time-bounded fragment of CSL with rewards comprising state formulae (Φ) and path formulae (ϕ) with the syntax:

$$\begin{aligned} \Phi &::= \text{true} | a | \neg \Phi | \Phi \wedge \Phi | P_{\sim r}[\phi] | R_{\sim r}[C^{\leq t}] \\ \phi &::= X \Phi | \Phi U^I \Phi \end{aligned} \quad (7)$$

where a is an atomic proposition evaluated over states, $\sim \in \{<, \leq, \geq, >\}$ is a relational operator, r is a probability ($r \in [0, 1]$) or reward ($r \in \mathbb{R}_{\geq 0}$) threshold², $t \in \mathbb{R}_{\geq 0}$ is a time bound, and $I \subseteq \mathbb{R}_{\geq 0}$ is a bounded time interval. The ‘future’ operator, F , and ‘globally’ operator, G , are derived from U in the standard way³. As briefly discussed in Section 4.2, our approach can be extended to unbounded CSL.

Traditionally, the CSL semantics is defined for CTMCs using a satisfaction relation \models . Intuitively, a state $s \models P_{\sim r}[\phi]$ iff the probability of the set of paths starting in s and satisfying ϕ meets $\sim r$. A path $\omega = s_0 t_0 s_1 t_1 \dots$ satisfies the formula $\Phi U^I \Psi$ iff there exists a time $t \in I$ such that $(\omega @ t \models \Psi \wedge \forall t' \in [0, t]. \omega @ t' \models \Phi)$, where $\omega @ t$ denotes the state in ω at time t . A state $s \models R_{\sim r}[C^{\leq t}]$ iff the expected rewards over the path starting in s and cumulated within t time units satisfies $\sim r$, where the

¹ In other words, the tolerance of parameter k , γ_k , measures the extent to which k can be perturbed from its reference (midpoint) value.

² For simplicity, we use $\sim r$ to denote the threshold for both probability and reward quality attributes.

³ $P_{\sim r}[F^I \Phi] = P_{\sim r}[\text{true} U^I \Phi]$ and $P_{\sim r}[G^I \Phi] = P_{\sim r}[\text{true} \wedge F^I \neg \Phi]$.

rates with which reward is acquired in each state and the reward acquired at each transition are defined by a *reward structure*.

In line with our previous work (Češka et al., 2017), we introduce a *satisfaction function* $\Lambda_{\phi}: \mathcal{P} \times \mathcal{D} \rightarrow [0, 1]$ that quantifies how the satisfaction probability associated with a path CSL formula ϕ relates to the parameters of a pCTMC $\mathcal{C}(\mathcal{P}, \mathcal{D})$, where, for any $(p, q) \in \mathcal{P} \times \mathcal{D}$, $\Lambda_{\phi}(p, q)$ is the probability that ϕ is satisfied by the set of paths from the initial state $\mathcal{D}_{\text{init}}(q)$ of the instantiated CTMC $\mathcal{C}(p, q)$. The satisfaction function for reward CSL formulae is defined analogously.

Quality requirements. We assume that the quality requirements of a SUD with design space given by a pCTMC $\mathcal{C}(\mathcal{P}, \mathcal{D})$ are defined in terms of:

- 1) A finite set of objective functions $\{f_i\}_{i \in I}$ corresponding to quality attributes of the system and defined in terms of a set of CSL path formulas $\{\phi_i\}_{i \in I}$, such that for any $i \in I$ and $(p, q) \in \mathcal{P} \times \mathcal{D}$,

$$f_i(\mathcal{C}(p, q)) = \Lambda_{\phi_i}(p, q); \quad (8)$$

- 2) A finite set of Boolean constraints $\{c_j\}_{j \in J}$ corresponding to the set of CSL path formulas $\{\psi_j\}_{j \in J}$ and thresholds $\{\sim r_j\}_{j \in J}$, such that for any $j \in J$ and $(p, q) \in \mathcal{P} \times \mathcal{D}$,

$$c_j(\mathcal{C}(p, q)) \Leftrightarrow \Lambda_{\psi_j}(p, q) \sim r_j. \quad (9)$$

Note that quality requirements (8) and (9) are defined over (non-parametric) CTMCs, but, in order to compare candidate designs with respect to some objective function, we need to interpret quality requirements over pCTMCs. Indeed, due to the continuous parameter space, a single candidate design induces an infinite number of objective function values, from which the designer must choose a representative value. For a candidate design $\mathcal{C}(\mathcal{P}', q)$ and objective f_i , this is typically identified as one of the minimum, maximum and mid-range value of $f_i(\mathcal{C}(p, q))$ over all $p \in \mathcal{P}'$, as illustrated in Table 1.

On the other hand, constraints have a unique interpretation because they must be met for any parameter value of a candidate design. Formally, for candidate design $\mathcal{C}(\mathcal{P}', q)$ and constraint c_j , we define

$$c_j(\mathcal{C}(\mathcal{P}', q)) \Leftrightarrow \forall p \in \mathcal{P}'. c_j(\mathcal{C}(p, q)).$$

Without loss of generality, we will assume that all objective functions $\{f_i\}_{i \in I}$ in Sections 3 and 4 should be minimised and that all thresholds $\{\sim r_j\}_{j \in J}$ are upper bounds of the form of $\leq r_j$.

Example 3 (Quality requirements). Below we define quality requirements for the producer-consumer model of Example 1. We consider two maximisation objectives and one constraint:

f_1 : $R\{\text{“consume”}\}_{=?} [C^{\leq 25}]$, a cumulative transition reward describing the number of requests transferred to the consumer within 25 time units (line 21 in Fig. 3);

f_2 : $P_{\geq 0.5} [G[20, 25]((\text{buffers} \geq \text{slow.max}/2) \& (\text{bufferf} \geq \text{fast.max}/2))]$, which calculates the probability that the utilisation of both buffers is at least 50% of their respective capacities;

c_1 : $R\{\text{“lost”}\}_{\leq 10} [C^{\leq 25}]$, a cumulative transition reward that limits the number of packets lost within 25 time units (line 20 in Fig. 3). With these quality requirements, we seek to maximise the system throughput (objective f_1), expressed as the number of requests transferred to the consumer, and also to maximize the probability that both buffers are sufficiently utilised after an initial period (objective f_2).

Finally, constraint c_1 imposes a reliability requirement by restricting the number of packets lost to be less than 10 within 25 time units of operation.

Table 1
Alternative definitions for objective functions $\{f_i\}_{i \in I}$ over candidate designs.

Type	Notation	Definition
Lower bound	$f_i^{\perp}(\mathcal{C}(\mathcal{P}, q))$	$\inf_{p \in \mathcal{P}} \Lambda_{\phi_i}(p, q)$
Upper bound	$f_i^{\top}(\mathcal{C}(\mathcal{P}, q))$	$\sup_{p \in \mathcal{P}} \Lambda_{\phi_i}(p, q)$
Mid-range	$f_i^*(\mathcal{C}(\mathcal{P}, q))$	$(f_i^{\perp}(\mathcal{C}(\mathcal{P}, q)) + f_i^{\top}(\mathcal{C}(\mathcal{P}, q)))/2$

2.3. Sensitivity of candidate designs

Quantifying the sensitivity of candidate designs is a crucial step in our robust synthesis method. Intuitively, the sensitivity of a design $\mathcal{C}(\mathcal{P}, q)$ captures how the objective functions $\{f_i\}_{i \in I}$ change in response to variations in the continuous parameters $k \in K$. The variation of each objective f_i is measured by the length of the interval $[f_i^{\perp}(\mathcal{C}(\mathcal{P}, q)), f_i^{\top}(\mathcal{C}(\mathcal{P}, q))]$, describing the range of admissible values for f_i and $\mathcal{C}(\mathcal{P}, q)$ (cf. Table 1). The degree of variation for multiple objectives is given by the product of interval lengths, i.e., the volume of the corresponding quality-attribute region. The sensitivity takes also into account the size of the underlying parameter region, in order to account for designs with different tolerance values. For instance, a design with a large quality-attribute volume and high tolerance (large parameter region volume) must be considered *more robust* (less sensitive) than another design with comparable quality-attribute volume but lower tolerance.

Definition 3 (Sensitivity). For a set of objective functions $\{f_i\}_{i \in I}$ and tolerances $\{\gamma_k\}_{k \in K}$, the *sensitivity* of a feasible design $\mathcal{C}(\mathcal{P}, q)$ is defined as the volume of its quality-attribute region over the volume of \mathcal{P} :

$$\text{sens}(\mathcal{C}(\mathcal{P}, q)) = \frac{\prod_{i \in I} (f_i^{\top}(\mathcal{C}(\mathcal{P}, q)) - f_i^{\perp}(\mathcal{C}(\mathcal{P}, q)))}{\prod_{k \in K} 2\gamma_k(k^{\top} - k^{\perp})}. \quad (10)$$

Example 4 (Sensitivity). Consider the candidate designs d_1, d_2, d_3 with tolerance $\gamma = 0.005$ from Example 2, and the objective functions f_1 (number of “consumed” packets) and f_2 (probability of buffers being sufficiently used) introduced in Example 3. Assume the following ranges for f_1 and f_2 :

$$\begin{aligned} [f_1^{\perp}(d_1), f_1^{\top}(d_1)] &= [416.94, 439.65] \\ [f_2^{\perp}(d_1), f_2^{\top}(d_1)] &= [0.8977, 0.9809] \\ [f_1^{\perp}(d_2), f_1^{\top}(d_2)] &= [407.11, 423.10] \\ [f_2^{\perp}(d_2), f_2^{\top}(d_2)] &= [0.891, 0.9621] \\ [f_1^{\perp}(d_3), f_1^{\top}(d_3)] &= [384.81, 413.09] \\ [f_2^{\perp}(d_3), f_2^{\top}(d_3)] &= [0.7501, 0.8225]. \end{aligned}$$

Recall that the three designs have the same tolerance, thus yielding the same parameter region volume

$$\prod_{k \in K} 2\gamma_k(k^{\top} - k^{\perp}) = 0.25 \cdot 0.3 = 0.075$$

The resulting sensitivities are:

$$\begin{aligned} \text{sens}(d_1) &= (439.65 - 416.94)(0.9809 - 0.8977)/0.075 = 25.19 \\ \text{sens}(d_2) &= (423.10 - 407.11)(0.9621 - 0.891)/0.075 = 15.16 \\ \text{sens}(d_3) &= (413.09 - 384.81)(0.8225 - 0.7501)/0.075 = 27.3 \end{aligned}$$

indicating that d_2 is the most robust design (with the smallest sensitivity value). The three designs can be visualised in the *quality-attribute space* (i.e. the objective space), as shown in Fig. 4, providing a direct and intuitive way to assess robustness.

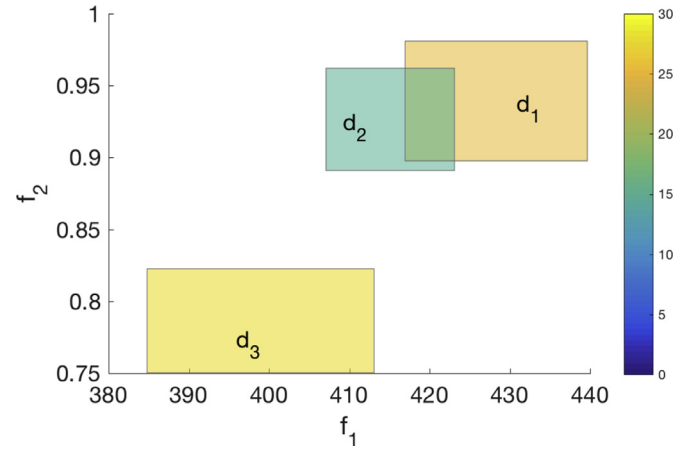


Fig. 4. Candidate designs of Example 4 represented in the quality-attribute space and coloured by sensitivity. Designs d_1 and d_2 were synthesised using RODES (full results are reported in Fig. 7 on a different scale).

3. Sensitivity-aware Pareto dominance relation

In this section, we introduce a novel dominance relation that adequately captures tradeoffs between the sensitivity and optimality of candidate designs with respect to given quality requirements, and that enables to formulate the robust design problem as an optimisation problem.

Consider a system with design space $\mathcal{C}(\mathcal{P}, \mathcal{Q})$, quality requirements given by objective functions $\{f_i\}_{i \in I}$ and constraints $\{c_j\}_{j \in J}$, and designer-specified tolerances $\{\gamma_k\}_{k \in K}$ for the continuous parameters of the system. Also, let \mathcal{F} be the set of feasible designs for the system (i.e., of candidate designs that meet the tolerances $\{\gamma_k\}_{k \in K}$ and satisfy the constraints $\{c_j\}_{j \in J}$):

$$\begin{aligned} \mathcal{F} &= \{\mathcal{C}(\mathcal{P}, q) \mid \mathcal{P} = \mathbf{X}_{k \in K} [k^{\perp}, k^{\top}] \subset \mathcal{C} \wedge q \in \mathcal{Q} \wedge \\ &\forall k \in K. k^{\top} - k^{\perp} = 2\gamma_k(k^{\top} - k^{\perp}) \wedge \forall j \in J. c_j(\mathcal{C}(\mathcal{P}, q))\}. \end{aligned} \quad (11)$$

Definition 4. A sensitivity-aware Pareto dominance relation over a feasible design set \mathcal{F} and a set of minimisation objective functions $\{f_i\}_{i \in I}$ is a relation $< \subset \mathcal{F} \times \mathcal{F}$ such that for any feasible designs $d, d' \in \mathcal{F}$

$$\begin{aligned} d < d' &\Leftrightarrow \\ &(\forall i \in I. f_i(d) \leq f_i(d') \wedge \exists i \in I. (1 + \epsilon_i)f_i(d) < f_i(d')) \vee \\ &(\forall i \in I. f_i(d) \leq f_i(d') \wedge \exists i \in I. f_i(d) < f_i(d') \wedge \\ &\text{sens}(d) \leq \text{sens}(d')). \end{aligned} \quad (12)$$

where the objective functions $\{f_i\}_{i \in I}$ are calculated using one of the alternative definitions from Table 1 and $\epsilon_i \geq 0$ are sensitivity-awareness parameters.

The parametric Markov chain synthesis problem consists of finding the Pareto-optimal set PS of candidate designs (5) (i.e. p CTMCs) with tolerances $\{\gamma_k\}_{k \in K}$ that satisfy the constraints $\{c_j\}_{j \in J}$ and are non-dominated with respect to the objective functions $\{f_i\}_{i \in I}$ and the sensitivity-aware dominance relation ' $<$ ':

$$PS = \{\mathcal{C}(\mathcal{P}, q) \in \mathcal{F} \mid \nexists \mathcal{C}(\mathcal{P}', q') \in \mathcal{F}. \mathcal{C}(\mathcal{P}', q') < \mathcal{C}(\mathcal{P}, q)\}, \quad (13)$$

Before discussing the rationale for this definition, we show that the sensitivity-aware Pareto dominance relation is a strict order like the classical Pareto dominance.

Theorem 1. *The sensitivity-aware Pareto dominance relation is a strict order.*

Proof. See Appendix A \square

The classical Pareto dominance definition can be obtained by setting

$\epsilon_i = 0$ for all $i \in I$ in (12). When $\epsilon_i > 0$ for some $i \in I$, dominance with respect to quality attribute i holds in our generalised definition in two scenarios:

- 1) when the quality attribute has a much lower value for the dominating design, i.e. $(1 + \epsilon_i)f_i(d) < f_i(d')$;
- 2) when in addition to a (slightly) lower quality attribute value, i.e. $f_i(d) < f_i(d')$, the sensitivity of the dominating design is no worse than that of the dominated design, i.e. $sens(d) \leq sens(d')$.

These scenarios are better aligned with the needs of designers than those obtained by using sensitivity as an additional optimisation criterion, which induces Pareto fronts comprising many designs with low sensitivity but unsuitably poor quality attributes. Similarly, each objective function definition from Table 1 captures specific needs of real-world systems. Thus, using the “upper bound” definition (f_i^\top) in (12) supports the synthesis of *conservative designs* by comparing competing designs based on the worst-case values of their quality attributes. This is suitable when the worst-case performance, reliability, etc. must be specified for a system, e.g. in its service-level agreement. In contrast, the “lower bound” definition from Table 1 (f_i^\perp) can be used when design selection must be based on the best expected quality values of a system. Finally, the “mid-range” definition (f_i^*) may be useful—in conjunction with the actual sensitivity (10)—to compare and select designs based on their reference midpoint quality values.

Importantly, for $\epsilon_i > 0$ our generalised definition induces Pareto fronts comprising designs with non-optimal (in the classical sense) objective function values, but with low sensitivity. We call such designs *sub-optimal robust*. Thus, ϵ_i can be finely tuned to sacrifice objective function optimality (slightly) for better robustness. Below we formally characterize the set of robust sub-optimal designs and provide an example of the sensitivity-aware dominance relation.

Definition 5 (Sub-optimal robust design). Let PS be a Pareto-optimal set defined as per (13). A design $d' \in PS$ is called robust sub-optimal if $\exists d \in PS$ s.t.:

$$(\forall i \in I. f_i(d) \leq f_i(d') \wedge \exists i \in I. f_i(d) < f_i(d'))$$

Example 5 (Sensitivity-aware Pareto dominance relation). Consider the quality-attribute regions of Fig. 4 induced by designs d_1, d_2, d_3 of the producer-consumer model introduced in Examples 1–4, and the objective functions defined as $f_i = f_i^\perp$ for $i \in 1, 2$. Visually, f_i^\perp corresponds to the lower-left corners of the regions in Fig. 4. Since we maximize both objectives, for clarity, we report below the dominance relation for maximisation:

$$\begin{aligned} d > d' \Leftrightarrow & (\forall i \in I. f_i(d) \geq f_i(d') \wedge \exists i \in I. f_i(d) > (1 + \epsilon_i)f_i(d')) \vee \\ & (\forall i \in I. f_i(d) \geq f_i(d') \wedge \exists i \in I. f_i(d) > f_i(d') \wedge \\ & \quad sens(d) \leq sens(d')). \end{aligned}$$

The designs d_1, d_2, d_3 have identical parameter tolerances and thus, same parameter space volume V . We have that $d_1 > d_2 > d_3$ when $\epsilon_1 = \epsilon_2 = 0$ (classical dominance) because for $i = 1, 2$, $f_i^\perp(d_1) > f_i^\perp(d_2), f_i^\perp(d_2) > f_i^\perp(d_3)$. Further, we have that $d_1 \succ d_2$ when $\epsilon_1 = \epsilon_2 = 0.05$, implying that d_2 is robust sub-optimal, i.e., is retained in the sensitivity-aware Pareto-optimal set, because $f_1^\perp(d_1) \succ > 1.05 \cdot f_1^\perp(d_2)$, $f_2^\perp(d_1) \succ > 1.05 \cdot f_2^\perp(d_2)$, and $sens(d_1) \succ \leq sens(d_2)$. Design d_3 is not included in the front ($d_1, d_2 > d_3$) because $f_1^\perp(d_1), f_1^\perp(d_2) > 1.05 \cdot f_1^\perp(d_3)$ for $i = 1, 2$.

4. Synthesis of sensitivity-aware Pareto sets

In this section, we describe our method for computing sensitivity-aware Pareto sets. The method employs genetic multi-objective optimisation algorithms for generating candidate designs and a precise parameter analysis of p CTMCs for evaluating the candidate designs. We

```

1: function SYNTHESIS( $C(\mathcal{P}, \mathcal{Q}), \{f_i\}_{i \in I}, \{c_j\}_{j \in J}, \{v_k\}_{k \in K}$ )
2:    $\overline{PS} \leftarrow \emptyset$ 
3:   while  $\neg$ TERMINATE( $C(\mathcal{P}, \mathcal{Q}), \overline{PS}$ ) do
4:      $CD \leftarrow$  CANDIDATEDESIGNS( $C(\mathcal{P}, \mathcal{Q}), \{v_k\}_{k \in K}, \overline{PS}$ )
5:     for all  $d \in CD$  do
6:        $(\{f_{i,d}^\top\}_{i \in I}, \{f_{i,d}^\perp\}_{i \in I}, \{c_{j,d}\}_{j \in J}) \leftarrow$ 
         ANALYSEDESIGN( $d, \{f_i\}_{i \in I}, \{c_j\}_{j \in J}$ )
7:       if  $\bigwedge_{j \in J} c_{j,d}$  then
8:          $dominated = \text{false}$ 
9:         for all  $d' \in \overline{PS}$  do
10:          if  $d' < d$  then  $dominated = \text{true}$ ; break
11:          if  $d < d'$  then  $\overline{PS} = \overline{PS} \setminus \{d'\}$ 
12:          end for
13:          if  $\neg$ dominated then  $\overline{PS} = \overline{PS} \cup \{d\}$ 
14:          end if
15:        end for
16:      end while
17:    return  $\overline{PS}$ 
18: end function

```

Algorithm 1. Parametric Markov chain synthesis.

start with a method overview, then we describe the two components the method builds on.

4.1. Method overview

Computing the Pareto-optimal design set (13) using exhaustive analysis is very expensive and requires a significant amount of computational resources as the design space $\mathcal{C}(\mathcal{P}, \mathcal{Q})$ is extremely large due to its real-valued parameters. Also, every candidate design $\mathcal{C}(\mathcal{P}, \mathcal{Q})$ consists of an infinite set of CTMCs that cannot all be analysed to establish its quality and sensitivity. To address these challenges, our p CTMC synthesis method combines search-based software engineering (SBSE) techniques (Harman et al., 2012a) with techniques for effective p CTMCs analysis (Češka et al., 2017; 2016), producing a close approximation of the Pareto-optimal design set.

Algorithm 1 presents the high-level steps of our p CTMC synthesis method. The approximate Pareto-optimal design set \overline{PS} returned by this algorithm starts empty (line 2) and is assembled iteratively by the while loop in lines 5–15 until a termination criterion $\text{TERMINATE}(\mathcal{C}(\mathcal{P}, \mathcal{Q}), \overline{PS})$ is satisfied. Each iteration of this while loop uses an SBSE metaheuristic to get a new set of candidate designs (line 4) and then updates the approximate Pareto-optimal design set \overline{PS} in the for loop from lines 5–15. This update involves analysing each candidate design $d = \mathcal{C}(\mathcal{P}, \mathcal{Q})$ to establish its associated objective function and constraint values in line 6, where we use the shorthand notation $f_{i,d}^\top \equiv f_i^\top(\mathcal{C}(\mathcal{P}, \mathcal{Q}))$, $f_{i,d}^\perp \equiv f_i^\perp(\mathcal{C}(\mathcal{P}, \mathcal{Q}))$ and $c_{j,d} \equiv v_j \in \mathcal{P}$. $c_j(\mathcal{C}(p, q))$ for all $i \in I, j \in J$. If the design satisfies all constraints (line 7), the for loop in lines 9–12 finds out if the new design d is dominated by, or dominates, any designs already in \overline{PS} . Existing designs dominated by d are removed from \overline{PS} (line 11), and d is added to the Pareto-optimal design set if it is not dominated by any existing designs (line 13).

The elements below must be concretised in the synthesis algorithm, and are described in the next two sections:

- 1) The ANALYSEDESIGN function for establishing the quality attributes and constraint compliance of a candidate design;
- 2) The CANDIDATEDESIGNS SBSE metaheuristic and the associated TERMINATE criterion.

The time complexity of Algorithm 1 is *linear* with respect to the overall number of optimisation objectives and constraints and the time required to analyse one quality attribute of a candidate design. The

complexity is further affected by the SBSE metaheuristic setting, namely by the number of generations k (i.e. the number of iterations of the while loop) and the size of the candidate design population $N = |CD|$. Increasing the total number of design evaluations (i.e. $k \cdot N$) typically improves the Pareto optimality of the generated design set, but also slows down the synthesis process. We provide a detailed complexity analysis of the synthesis process in [Appendix B](#).

4.2. Computing safe property bounds for pCTMCs

To establish the quality attributes and sensitivity of candidate designs, ANALYSEDESIGN uses precise parameter synthesis techniques ([Češka et al., 2017](#)) to compute safe enclosures of the satisfaction probability of CSL formulae over pCTMCs. Given a pCTMC $\mathcal{C}(\mathcal{P}, q)$ and a CSL path formula ϕ , these techniques provide a safe under-approximation Λ_{\min}^q and a safe over-approximation Λ_{\max}^q of the minimal and maximal probability that $\mathcal{C}(\mathcal{P}, q)$ satisfies ϕ :

$$\Lambda_{\min}^q \leq \inf_{p \in \mathcal{P}} \Lambda_{\phi}(p, q) \quad \text{and} \quad \Lambda_{\max}^q \geq \sup_{p \in \mathcal{P}} \Lambda_{\phi}(p, q).$$

This supports the safe approximation of the bounds $\{f_i^{\perp}, f_i^{\top}\}_{i \in I}$ of the objective functions and of the constraints $\{c_j\}_{j \in J}$. As shown in [Češka et al. \(2017\)](#), the over-approximation quality improves as the size of \mathcal{P} decreases. Therefore, the precision of the approximation can be effectively controlled via parameter space decomposition, where \mathcal{P} is decomposed into subspaces $\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_n$ and Λ_{\min}^q (Λ_{\max}^q) is taken as the minimum (maximum) of the bounds computed for these n subspaces. Although this refinement step improves the precision of bounds, it also increases the complexity of ANALYSEDESIGN n -fold ([Češka et al., 2017](#)).

The satisfaction function Λ_{ϕ} is typically non-monotonic (and, for nested properties, non-continuous), so safe bounds cannot be obtained by simply evaluating Λ_{ϕ} at the extrema of parameter region \mathcal{P} . Accordingly, our technique builds on a parametric backward transient analysis that computes safe bounds for the parametric transient probabilities in the discrete-time process derived from the pCTMC. This discretisation is obtained through standard uniformisation, and through using the Fox and Glynn algorithm ([Kwiatkowska et al., 2007](#)) to derive the required number of discrete steps for a given time bound. Once the parametric discrete-time process is obtained, the computation of the bounds reduces to a local and stepwise minimisation/maximisation of state probabilities in a time non-homogenous Markov process. Presenting the technique in detail as well as the analysis of the approximation error is outside the scope of our paper, but the interested reader can find a complete description in [Češka et al. \(2017\)](#).

Our approach can be easily extended to also support time-unbounded properties by using the method of [Quatmann et al. \(2016\)](#) for parameter synthesis of discrete-time Markov models and properties expressed by time-unbounded formulae of probabilistic computation tree logic.

4.3. Metaheuristic for parametric CTMC synthesis

To ensure that CANDIDATEDESIGNS selects suitable candidate designs, [Algorithm 1](#) is implemented as a *multiobjective optimisation genetic algorithm* (MOGA) such as NSGA-II ([Deb et al., 2002](#)) or MOCeL ([Nebro et al., 2009](#)). MOGAs are genetic algorithms specifically tailored for the synthesis of close Pareto-optimal set approximations that are spread uniformly across the search space. As with any genetic algorithm ([Koza, 1992](#)), possible solutions—candidate designs in our case—are encoded as tuples of *genes*, i.e. values for the problem variables. In particular, any candidate design $\mathcal{C}(\mathcal{P}, q)$ that satisfies a fixed set of tolerances $\{\gamma_k\}_{k \in K}$ is uniquely encoded by the gene tuple (p, q) , where $p \in \mathcal{P}$ is the centre point of the continuous parameter region \mathcal{P} . The structure of the gene tuple (p, q) for any pCTMC $\mathcal{C}(\mathcal{P}, \mathcal{D})$ is automatically extracted through parsing the evolvable constructs (4). This feature enables to conveniently encode the pCTMC parameters into a

representation suitable for the MOGAs.

Example 6 (Candidate design encoding). Consider the candidate designs d_1, d_2, d_3 with tolerance value $\gamma = 0.005$ from [Example 2](#). The gene tuple (p, q) of a candidate design $\mathcal{C}(\mathcal{P}, q)$ has the structure (rslorate, deltarate, moduleidx), where moduleidx $\in \{1, 2\}$ is the index of the Buffer module used by the candidate design. Thus, the designs d_1, d_2, d_3 have gene tuples given by (15.145, 2.08, 2), (13.325, 3.66, 2) and (17.365, 2.93, 1), respectively.

The first execution of CANDIDATEDESIGNS from [Algorithm 1](#) returns a randomly generated *population* (i.e. set) of feasible designs (11). This population is then iteratively evolved by subsequent CANDIDATEDESIGNS executions into populations of “fitter” designs through MOGA *selection*, *crossover* and *mutation*. Selection chooses the population for the next iteration and a *mating pool* of designs for the current iteration by using the objective functions $\{f_i\}_{i \in I}$, the sensitivity-aware dominance relation (12) and the distance in the parameter space \mathcal{P} between designs to evaluate each design. Crossover randomly selects two designs from the mating pool, and generates a new design by combining their genes, and mutation yields a new design by randomly modifying some of the genes of a design from the pool.

The evolution of the design population terminates (i.e. the predicate $Terminate(\mathcal{C}(\mathcal{P}, \mathcal{D}), \overline{PS})$ returns true) after a fixed number of design evaluations or when a predetermined number of successive iterations generate populations with no significantly fitter designs.

The implementation of the selection, crossover and mutation operations is specific to each MOGA. For instance, [Deb et al. \(2002\)](#) presents these features for the NSGA-II MOGA used in our experimental evaluation from [Section 6](#).

5. RODES: a robust-design synthesis tool

Our GPU-accelerated RODES tool synthesises sensitivity-aware Pareto sets by implementing the process described in [Algorithm 1](#). In this section, we first present the architecture of RODES, and then describe how we achieved significant performance and scalability improvements through the use of a two-level parallelisation for the synthesis process.

5.1. RODES architecture

As shown in [Fig. 5](#), the operation of RODES is managed by a *Robust-design synthesis engine*. First, a *Model parser* (built using the Antlr parser generator, www.antlr.org) preprocesses the design-space pCTMC model. Next, a *Sensitivity-aware synthesiser* uses the jMetal Java framework for multi-objective optimisation with metaheuristics (jmetal.github.io/jMetal) to evolve an initially random population of *candidate designs*, generating a close approximation of the sensitivity-aware Pareto front. This involves using a *Candidate design analyser*, which invokes the probabilistic model checker PRISM-PSY ([Češka et al., 2016](#)) to obtain the ranges of values for the relevant quality attributes of candidate designs through precise parameter synthesis. The Pareto front and corresponding Pareto-optimal set of designs are then plotted using MATLAB/Octave scripts, as shown in [Fig. 7](#).

A key feature of RODES is its modular architecture. The Sensitivity-aware synthesiser supports several metaheuristics algorithms, including variants of genetic algorithms and swarm optimisers. Furthermore, the sensitivity-aware Pareto dominance relation can be adapted to match better the needs of the system under development (e.g., by comparing designs based on the worst, best or average quality attribute values). Finally, different solvers could be used for the probabilistic model checker component, including the parameter synthesis solvers for discrete-time Markov chains and time unbounded properties ([Quatmann et al., 2016](#)) implemented in the tools PROPhESY ([Dehnert et al., 2015](#)) and STORM ([Dehnert et al., 2017](#)).

The open-source code of RODES is available on our project website

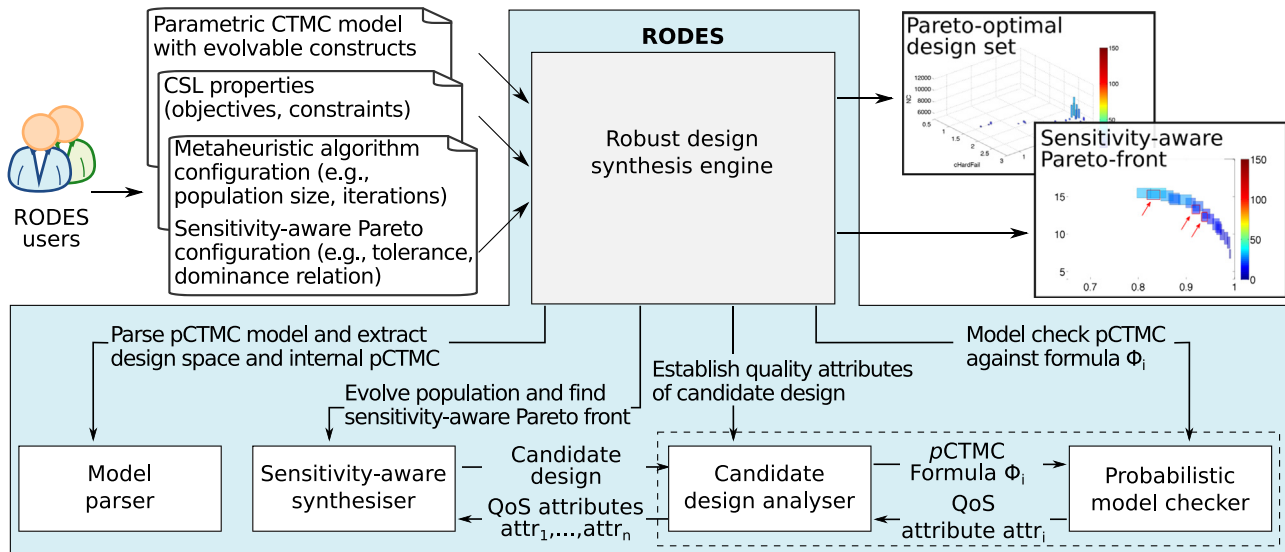


Fig. 5. High-level RODES architecture.

<https://www.github.com/gerasimou/RODES>.

5.2. Two-level parallelisation

Synthesising sensitivity-aware Pareto sets is a computationally expensive process. To mitigate the performance issues that could arise due to the increased total number of evaluations ($k \cdot N$) or the complexity of evaluating candidate designs (t), we employ a two-level parallelisation.

At the first level, we exploit the fact that the evaluations of particular candidates within a single population are independent and thus they can run in parallel (line 6 in Algorithm 1). A synchronisation is required only after all candidates are evaluated to update the approximate Pareto-optimal set \overline{PS} and to generate new candidates. This granularity of parallelism allows us to efficiently utilise both multi-core and multi-processor architectures. In particular, we can span in parallel a number of tasks that is equal to the population size N and thus significantly alleviate the complexity corresponding to the total number of design evaluations per MOGA generation. We can further increase the parallelisation at this level given that the evaluation of quality attributes for each design is independent. Thus, we can span up to $N \cdot (|I| + |J|)$ tasks to evaluate these attributes in parallel and reduce the computation time. The current RODES implementation supports parallelisation at the population level but not at the level of quality attributes, which we plan to add in future tool releases.

The second level of parallelisation aims at accelerating the evaluation of a single candidate over a single quality attribute. The key factor affecting the time t required to analyse a quality attribute of a candidate design is the size of the candidate, namely, the number of non-zero elements M in the rate matrix of the underlying pCTMC. This number is proportional to the number of states in the pCTMC, and reflects the complexity of the candidate designs. To ensure that RODES supports robust design synthesis for complex systems comprising up to tens thousands of states, our second-level parallelisation improves scalability with respect to the number of states. In particular, we build on our previous work (Češka et al., 2016) to integrate a GPU acceleration of the pCTMC analysis into RODES.

This parallelisation is much more involved, since the computation for individual states is not independent. As such, the pCTMC analysis is formulated in terms of matrix-vector operations, making it suitable for effective data-parallel processing. Accordingly, RODES implements a state space parallelisation, where a single row of the parametric rate matrix (corresponding to the processing of a single state) is mapped to a single computational element. As the underlying pCTMCs typically have

a balanced distribution of the state successors, this mapping yields a balanced distribution of non-zero elements in the rows of the matrix. The outcome is a good load balancing within the computation elements, leading to significant acceleration. In contrast to the parallelisation proposed in Češka et al. (2016), RODES is designed to leverage the computational power of modern GPUs, which provide hundreds of computational elements and can schedule thousands of active threads in a different way. In particular, RODES can evaluate on a single GPU several candidate designs (that can differ both in their discrete and in their continuous parameters) in parallel, provided that the underlying pCTMCs can fit in the GPU memory. This enables an efficient and flexible utilisation of the available computation power for complex robust design synthesis problems (see performance evaluation results in Section 6.4).

6. Evaluation

We evaluate the effectiveness of RODES using three systems from different application domains. Also, we assess the performance and scalability of RODES including the impact of the two-level parallelisation. We conclude our evaluation with a discussion of threats to validity.

6.1. Research questions

The aim of our experimental evaluation was to answer the following research questions.

RQ1 (Decision support): Can RODES support decision making by identifying effective tradeoffs between the QoS optimality and the sensitivity of alternative designs? To support decision making, RODES must provide useful insights into the robustness of alternative system designs. Therefore, we assessed the optimality-sensitivity tradeoffs suggested by RODES for the software systems used in our evaluation.

RQ2 (Performance): Does the two-level parallelisation improve the efficiency of RODES? Since the synthesis of robust models is a computationally expensive process, we examined the change in performance thanks to the two-level parallelisation architecture described in Section 5.2.

RQ3 (Metaheuristic effectiveness): How does our RODES approach perform compared to random search? Following the standard practice in search-based software engineering Harman et al. (2012b), we assessed if the stochastic models synthesised

by RODES “comfortably outperform” those synthesised by a random search approach.

6.2. Analysed software systems

We performed a wide range of experiments to evaluate our RODES approach and tool using three software systems from different application domains:

- a producer-consumer (PC) software system described in Examples 1–5;
- a replicated file system used by Google’s search engine Baier et al. (2013);
- a cluster availability management system Haverkort et al. (2000).

We have already presented the PC system in Examples 1–5. In the following paragraphs, we introduce the other systems, provide a description of their stochastic models and present the objectives and constraints used to synthesise robust Pareto optimal designs. Further information about these systems are available on our project website at <https://www.github.com/gerasimou/RODES/wiki>.

Google file system (GFS). GFS partitions files into chunks of equal size, and stores copies of each chunk on multiple *chunk servers*. A master server monitors the locations of these copies and the chunk servers, replicating the chunks as needed. During normal operation, GFS stores CMAX copies of each chunk. However, as servers fail and are repaired, the number c of copies for a chunk may vary from 0 to CMAX.

Previous work modelled GFS as a CTMC with fixed parameters and focused on the analysis of its ability to recover from disturbances (e.g. $c < \text{CMAX}$) or disasters (e.g. master server down) (Baier et al., 2013). In our work, we adapt the CTMC of the lifecycle of a GFS chunk from Baier et al. (2013) by considering several continuous and discrete parameters that a designer of the system has to decide. Fig. 6 shows the resulting model, encoded in the PRISM modelling language extended with the evolve constructs from (4). As in Baier et al. (2013), we model separately the software and hardware failures and repairs, for both the master server (lines 22–25) and the chunk servers (lines 26–31), and assume that loss of chunk copies due to chunk server failures leads to further chunk replications, which is an order of magnitude slower if $c = 0$ and a backup of the chunk must be used (line 32).

To evaluate RODES, we assume that GFS designers must select the hardware failure and repair rates $c\text{HardFail}$ and $c\text{HardRepair}$ of the chunk servers, and the maximum number of chunks NC stored on a chunk server within the ranges indicated in Fig. 6. These parameters reflect the fact that designers can choose from a range of physical servers, can select different levels of service offered by a hardware repair workshop, and can decide a maximum workload for chunk servers. We consider an initial system state modelling a severe hardware disaster with all servers down due to hardware failures and all chunk copies lost, and we formulate a $p\text{CTMC}$ synthesis problem for quality requirements given by two maximising objective functions and one constraint:

- f_1 : $P_{\geq ?} [\neg \text{SL1 } U^{[10,60]} \text{SL1}]$, where $\text{SL1} = \text{Mup} \wedge c > 0$ holds in states where *service level 1* (master up and at least one chunk copy available) is provided;
- f_2 : $R\{\text{“active”}\}_{=?} [C \leq 60]$, where a reward of 1 is assigned to the states with a number of running chunk servers of at least 0.5M (i.e., half of the total number of chunk servers);
- c_1 : $R\{\text{“replicates”}\}_{\leq ?} [C \leq 60]$, where a transition reward of 1 is assigned to each chunk replication transition.

Objective f_1 maximises the probability that the system recovers service level 1 in the time interval [10,60] hours. Objective f_2

```

1  ctmc
   // Failure rates
2  const double mSoftFail = 0.000475; // master software
3  const double mHardFail = 0.000025; // master hardware
6  const double cSoftFail = 0.475; // chunk server software
7  evolve double cHardFail [0.25..4.0]; // chunk server hardware

   // Repair rates
4  const double mSoftRepair = 12; // master software
5  const double mHardRepair = 6; // master hardware
8  const double cSoftRepair = 12; // chunk server software
9  evolve double cHardRepair [0.5..4.0]; // chunk server hardware

10 const int N=100000; // total number of GFS chunks
11 const int M=20; // number of chunk servers
12 evolve int NC [5000..20000]; // max chunks per chunk server
13 const int CMAX=3; // optimal number of chunk copies

14 module GFS_Chunk
15   M_up : bool init false; // master is up
16   M_sdown : bool init false; // master is down with SW problem
17   M_hdown : bool init true; // master is down with HW problem
18   Cup : [0..M] init 0; // number of chunk servers up
19   Csdown : [0..M] init 0; // number of chunk servers down (SW problem)
20   Chdown : [0..M] init 20; // number of chunk servers down (HW problem)
21   c : [0..CMAX] init 0; // number of chunk copies available

   // Master server failure and repair
22 [] M_up → mSoftFail : (M_up' = false) & (M_sdown' = true);
23 [] M_up → mHardFail : (M_up' = false) & (M_hdown' = true);
24 [] M_sdown → mSoftRepair : (M_up' = true) & (M_sdown' = false);
25 [] M_hdown → mHardRepair : (M_up' = true) & (M_hdown' = false);

   // Chunk servers failure and repair
26 [] Cup > 0 & c > 0 & Csdown < M → (c/Cup)*cSoftFail :
      (Cup' = Cup-1) & (Csdown' = Csdown+1) & (c' = c-1);
27 [] Cup > 0 & Cup > c & Csdown < M → (1-(c/Cup))*cSoftFail :
      (Cup' = Cup-1) & (Csdown' = Csdown+1);
28 [] Cup > 0 & c > 0 & Chdown < M → (c/Cup)*cHardFail :
      (Cup' = Cup-1) & (Chdown' = Chdown+1) & (c' = c-1);
29 [] Cup > 0 & Cup > c & Chdown < M → (1-(c/Cup))*cHardFail :
      (Cup' = Cup-1) & (Chdown' = Chdown+1);
30 [] Cup < M & Csdown > 0 → Csdown*cSoftRepair :
      (Csdown' = Csdown-1) & (Cup' = Cup+1);
31 [] Cup < M & Chdown > 0 → cHardRepair :
      (Chdown' = Chdown-1) & (Cup' = Cup+1);
32 [] M_up & c < CMAX & Cup > c & Cup*NC >=(c+1)*N →
      ((c>0)?20:2):(c'=c+1);

33 endmodule

```

Fig. 6. $p\text{CTMC}$ model of the Google file system.

maximises the expected time the system stays in (optimal) states with at least 0.5M chunk servers up in the first 60 hours of operation. Finally, constraint c_1 restricts the number of expected chunk replications over 60 h of operations.

Workstation cluster (WC). We extend the CTMC of a cluster availability management system from Haverkort et al. (2000). This CTMC models a system comprising two sub-clusters, each with N workstations and a switch that connects the workstations to a central backbone. For each component, we consider failure, inspection and repair rates (where repairs are initiated only after an inspection detects failures), and we assume that designers must decide these rates for workstations—i.e., the real-valued parameters wsFail , wsCheck and wsRepair for our $p\text{CTMC}$, respectively. Additionally, we assume that designers must select the sub-cluster size N , and must choose between an expensive repair implementation (i.e., $p\text{CTMC}$ module) with a 100% success probability and a cheaper repair module with 50% success probability—i.e., two discrete parameters for the $p\text{CTMC}$. We made this model available on our repository of case studies.

For an initial system state with 5 workstations active in each sub-cluster and switches and backbone working, we formulate a $p\text{CTMC}$ synthesis problem for quality requirements given by two maximising objective functions and one constraint:

- f_1 : $P_{\geq ?} [\neg \text{premium } U [20, 100] \text{premium}]$, where premium denotes a system service where at least 1.25N workstations are connected and operating;
- f_2 : $R\{\text{“operational”}\}_{=?} [C \leq 100]$, where a reward of 1 is assigned to

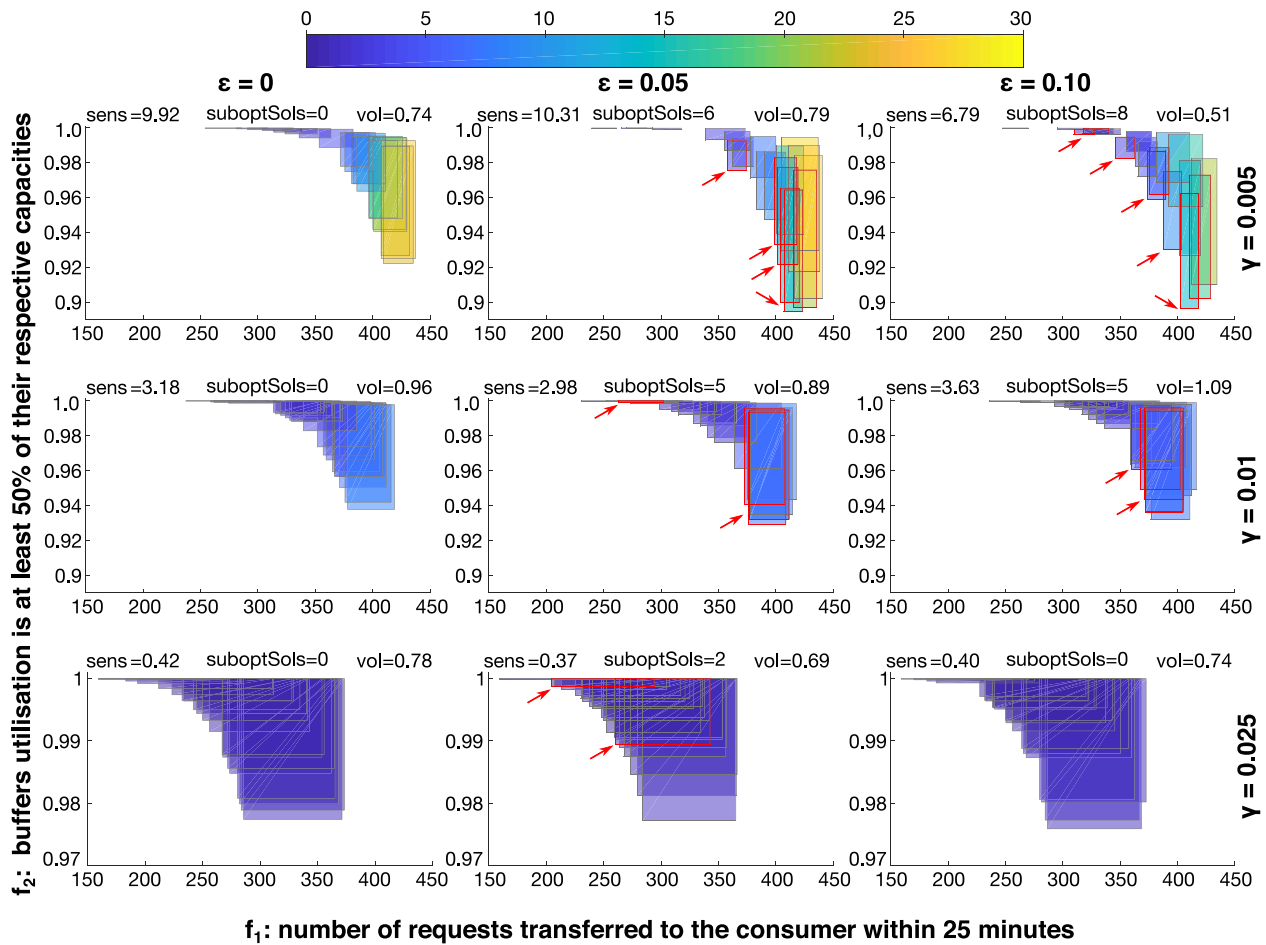


Fig. 7. Sensitivity-aware Pareto fronts for the producer-consumer model. Boxes represent quality-attribute regions, coloured by sensitivity (yellow: sensitive, blue: robust). Red-bordered boxes indicate sub-optimal robust designs. Designs are compared based on the worst-case quality attribute value (i.e. lower-left corner of each box). Statistics are: sens, average sensitivity of the front; suboptSols, number of suboptimal solutions; vol, average volume of the front. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

states with a number of operating clusters between $1.2N$ and $1.6N$; $c_1: R\{\text{“repair”}\} \leq 80 [C \leq 100]$, where transition rewards are associated with repair actions of the workstations, backbone and switches.

Objective f_1 maximises the probability that the system recovers the premium service in the time interval $[20,100]$ hours. Objective f_2 maximises the expected time the system spends in cost-optimal states during the first 100 hours of operation. Constraint c_1 restricts the cost of repair actions during this time (the definition of the cost is provided on our project website).

6.3. Evaluation methodology

We used the following configuration to evaluate RODES: NSGA-II MOGA, 10,000 evaluations, initial population of 20 individuals, and default values for single-point crossover probability $p_c = 0.9$ and single-point mutation probability $p_m = 1/(|K| + |D|)$, with $|K| + |D|$ the number of (continuous and discrete) design-space parameters. We examine the behaviour of the sensitivity-aware Pareto dominance relation using different combinations of tolerance values $\gamma \in \{0.005, 0.01, 0.025\}$ and sensitivity-awareness coefficients $\epsilon_i \in \{0.00, 0.05, 0.10\}$.

For each experiment, we report the sensitivity-aware Pareto fronts (Figs. 7, 9, 12 and 14). The Pareto-optimal designs are depicted as boxes in the quality-attribute space and coloured by sensitivity, using the same representation as in Figs. 1 and 4. We also show the synthesised designs in the design space, given by the continuous and discrete

parameters of the system. In this case, designs are represented as boxes in the continuous parameter space, representing the extent of the parameter variation under the given tolerance. The third dimension (vertical axis) in Figs. 10 and 13 gives the value of the discrete parameter.

6.4. Results and discussion

RQ1 (Decision support). We analysed the designs synthesised by RODES in order to identify actionable insights regarding the tradeoffs between the QoS attributes and sensitivity of alternative architecture designs. For each system, we present our findings independently.

Producer-consumer system (PC). First, we present the results for the producer consumer system introduced in Examples 1–5, obtained by running our RODES tool with tolerances $\gamma \in \{0.005, 0.01, 0.025\}$ for both continuous parameters ($r_{\text{slow_rate}}$ and δ_{rate}). The resulting Pareto fronts are shown in Fig. 7, for objectives f_1 (number of requests transferred to the consumer within 25 minutes) and f_2 (probability of adequate buffer utilization) and sensitivity-awareness parameters $\epsilon_1 = \epsilon_2 = \epsilon \in \{0, 0.05, 0.1\}$. The corresponding synthesised designs are presented in Fig. 8.

These Pareto fronts provide a wealth of information supporting the evaluation of the optimality and robustness of alternative designs. In particular, the Pareto front for $\epsilon = 0$ and $\gamma = 0.005$ contains several large (yellow) boxes that correspond to highly sensitive designs.

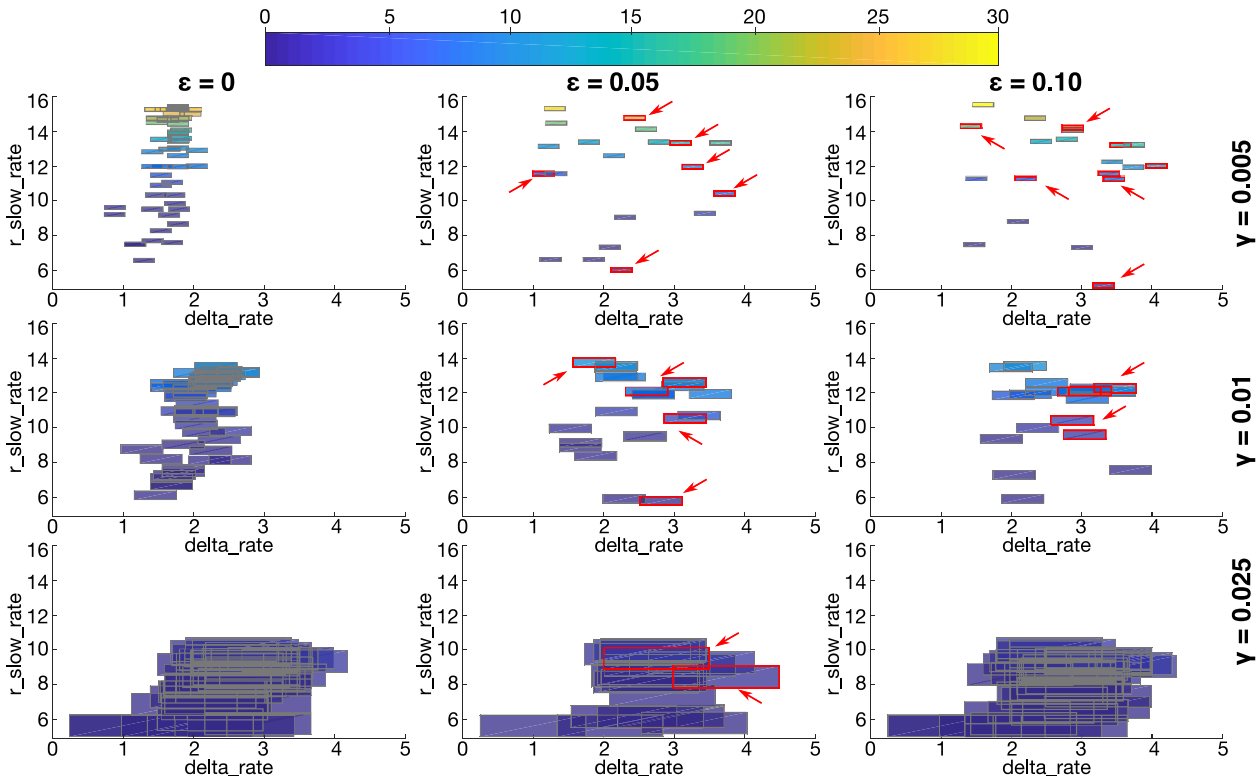


Fig. 8. Synthesised Pareto-optimal designs for the producer-consumer model and experiments from Fig. 7. Rectangles in x-y plane correspond to the continuous parameter regions. The discrete parameter (module - ‘mod’) is omitted since RODES synthesised solutions using only the redirection module (‘mod2’). Boxes are coloured by sensitivity.

Increasing ϵ produces a number of robust sub-optimal designs (red-bordered) with slightly sub-optimal quality attributes but improved robustness. Such designs represent valuable alternatives to the highly sensitive solutions obtained using the classical, sensitivity-agnostic, dominance relation. This ability to identify poor (i.e. highly sensitive) designs and then alternative robust designs with similar quality attributes is a key and unique benefit of our design synthesis method. Consider for instance the results for $\epsilon = 0.05$ and $\gamma = 0.005$. There are several sensitive designs at high f_1 values (see Fig. 7), which correspond to designs with $r_{slowrate}$ above 15 and low values δ_{rate} (below 2.5), see Fig. 8. Through our method, we found that there exist alternative sub-optimal designs with improved robustness (highlighted green boxes), corresponding to higher δ_{rate} and lower $r_{slowrate}$ values, i.e. to designs with a slower slow buffer and a faster fast buffer.

Furthermore, we observe that the overall sensitivity improves as the tolerance γ increases, meaning that the uncertainty (volume) of the quality attribute regions grows proportionally smaller than the uncertainty of the corresponding parameter regions, see (10). This explains why we observe fewer sub-optimal robust designs for higher tolerances ($\gamma = 0.01, 0.025$). Increasing parameter tolerances also affects the quality attribute profiles as it leads to larger ranges for objective f_1 (i.e., more sensitive) and to smaller ranges for f_2 (i.e., more robust). As a consequence, RODES tends to favour Pareto-optimal solutions with better f_2 and worse f_1 values as the tolerance increases. In particular, for $\gamma = 0.025$ all designs with $f_1^+ \geq 300$ are excluded (corresponding to the most sensitive designs for $\gamma = 0.005, 0.01$), which yields regions with average volume comparable to those for $\gamma = 0.025$.

The synthesised parameter regions (Fig. 8) indicate that redirection (second module - ‘mod2’) is always preferred to non-redirection. Also, the generated designs select values for the continuous parameters from the lower-end of their respective range, with $r_{slowrate} \in [5.00, 15.650]$ and $\delta_{rate} \in [0.242, 4.489]$. In other words, our algorithm found Pareto-optimal designs where both buffers have slow transmission rates

(with the fast buffer being slightly faster), while solutions where the fast buffer has a sensibly higher transmission rate, but a proportional packet loss rate, are excluded. In particular, configurations with slow transmission rates have associated good robustness, with very little ranges for objective f_2 .

We also observe an interesting relationship between the Pareto-optimal fronts and the Pareto-optimal designs for different values of the sensitivity-awareness parameter $\epsilon \in \{0, 0.05, 0.1\}$. The average values for both objectives f_1 and f_2 experience only little variation as ϵ increases for a fixed tolerance value. For instance, when $\gamma = 0.01$, on average $f_1 \in [369.37, 372.40]$ and $f_2 \in [0.974, 0.98]$, and when $\gamma = 0.05$, $f_1 \in [284.94, 285.48]$ and $f_2 \in [0.995, 0.996]$. Conversely, the average values for the continuous parameters $r_{slowrate}$ and δ_{rate} experience more significant variation and present an interesting negative relationship. More specifically, for any γ value and as the ϵ parameter becomes larger, $r_{slowrate}$ shows a decreasing trend while δ_{rate} shows an increasing trend. We used the Pearson correlation test to analyse this observation and received a strong negative correlation with the coefficient $R \in [-0.992, -0.988]^4$. This result indicates that as ϵ increases, the sensitivity-aware Pareto-optimal set includes designs in which the transmission rate difference between the slow and fast buffers grows. Although unexpected, this observation is very useful.

Producer-consumer variant. We further analyze a variant of the producer-consumer model, illustrated in Fig. 11. In this version, we assume a different redirection strategy (lines 10 and 11) that yields a 100% probability of redirection when the slow buffer is full, while in the original variant the maximum redirection probability is limited to 0.1. We also consider different continuous parameters: the request

⁴ This result should not be confused with the correlation between the continuous parameters $r_{slowrate}$ and δ_{rate} for fixed γ and ϵ values which ranges from zero to weak, i.e., $R \in [0, 0.3]$.

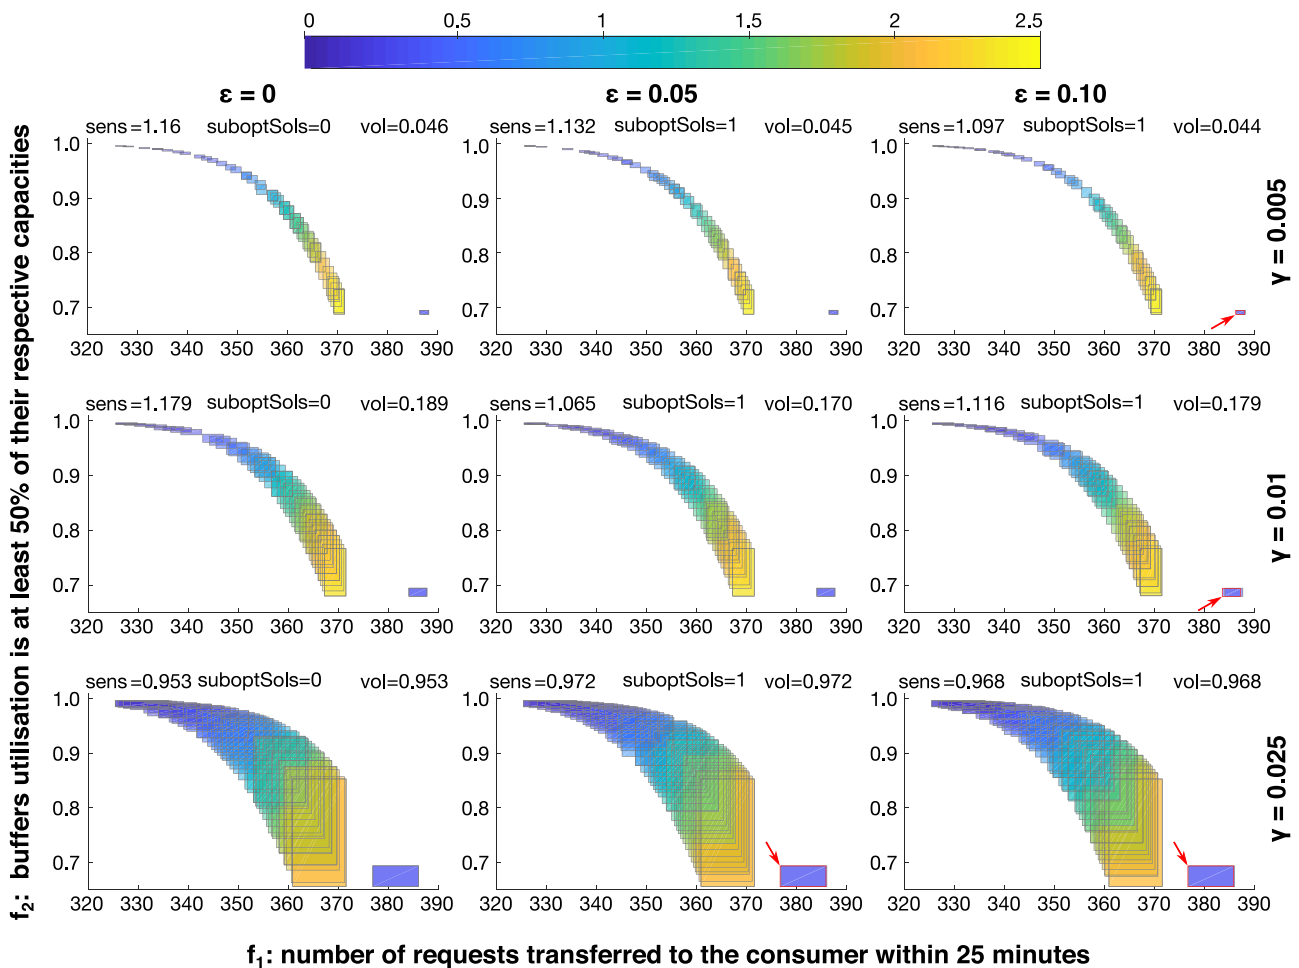


Fig. 9. Sensitivity-aware Pareto fronts for the second variant of the producer-consumer model. Legend and colour code are as in Fig. 7. Designs are compared based on the worst-case quality attribute value (i.e. lower-left corner of each box).

production rate (p_rate) and the packet transmission rate for the fast buffer (r_fast_rate). The synthesized Pareto fronts and designs are reported in Figs. 9 and 10, respectively.

We observe that the obtained Pareto-optimal set is substantially different from the one obtained in the first variant of the model (Fig. 7). Solutions in this variant are generally more robust, demonstrated by the fact that at most one suboptimal solution is synthesised for each configuration. A common trait is that favouring objective f_2 leads to robust designs, while robustness is penalized for high f_1 values. Comparing the two PC variants, whose $pCTMC$ models are shown Figs. 3 and 11, we observe that most of the solutions of the second variant are dominated by the Pareto front of the first variant for $\gamma \in \{0.005, 0.01\}$ and all ϵ values, which therefore provides the best performance.

The synthesized parameter regions (Fig. 10) confirm the results of the first variant: redirection is always preferred (for all but one design), and the fast buffer rate is not too far from that of the slow buffer ($rfastrate = 13.03$). Similarly, all synthesized values for parameter p_rate are very close to the fixed value (40) used for the same parameter in the first variant of the model. In the Pareto front, we can observe an outlier yielding the highest system throughput (f_1). This design is obtained when redirection is disabled (see Fig. 10). Notably, no other designs with no redirection are present in the Pareto front which provides evidence that redirection is essential to achieve a well-balanced utilisation of the buffers.

Google file system (GFS). Given the $pCTMC$ model, the two maximisation objectives and one constraint of the GFS system, we used RODES to generate Pareto-optimal design sets with tolerances

$\gamma \in \{0.005, 0.01, 0.025\}$ for both continuous parameters ($cHardFail$ and $cHardRepair$) of our $pCTMC$. Fig. 12 shows the Pareto fronts obtained using the “lower bound” definition from Table 1 for the objective functions f_1 and f_2 over candidate designs, and parameters $\epsilon_1 = \epsilon_2 = \epsilon \in \{0, 0.05, 0.1\}$ for the sensitivity-aware Pareto dominance relation (12). The design-space representation is given in Fig. 13. We observe that the Pareto front for $\epsilon = 0$ and $\gamma = 0.005$ contains several large (yellow) boxes that correspond to highly sensitive designs. For $\epsilon \in \{0.05, 0.1\}$ and $\gamma = 0.005$, these poor designs are “replaced” by robust designs – surrounded by red borders – with very similar quality attributes but slightly sub-optimal. The same pattern occurs for $\gamma = 0.01$ and (to a lesser extent because of the overall lower sensitivity) for $\gamma = 0.025$. For instance, consider the sensitive design obtained for $\epsilon = 0.1$ and $\gamma = 0.005$ characterized by low hardware fail and repair rates and high number of chunks (yellow bar on Fig. 13). Our method found that a more robust solution is possible (highlighted green region), with lower NC and higher $cHardFail$ and $cHardRepair$.

We also observe that favouring objective f_1 over f_2 generally yields more robust designs (i.e., smaller quality-attribute regions towards the right end of the Pareto fronts) for all combinations of ϵ and γ .

The design-space view of Fig. 13 evidences a trade-off between $cHardFail$ and $cHardRepair$, i.e., optimal designs tend to have either high failure rates and high repair rates, or low failure and repair rates. Results for $\gamma = 0.025$ reveal that there is actually an ideal ratio between the two parameters as the corresponding optimal design appear to keep a relatively constant proportion between $cHardFail$ and $cHardRepair$. This result was unexpected, yet very useful, since it indicates that designs not satisfying this trade-off yield excessively fast or slow recovery

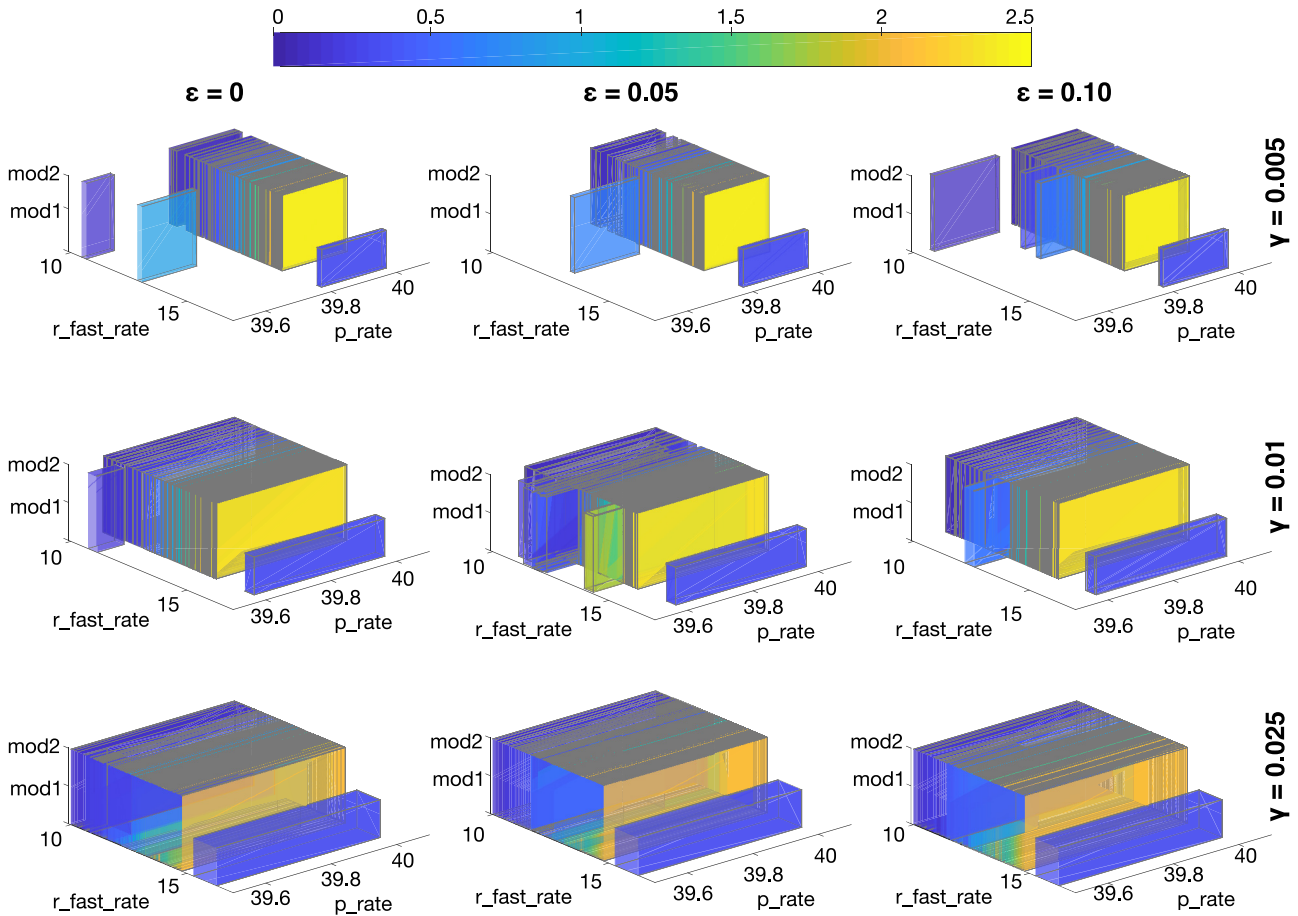


Fig. 10. Synthesised Pareto-optimal designs for the second variant of the producer-consumer model and experiments from Fig. 9. Rectangles in x-y plane correspond to the continuous parameter regions Boxes are coloured by sensitivity.

```

1  ctmc
   // buffer capacities
2  const int slow_max = 31;
3  const int fast_max = 21;
4  evolve double p_rate [20..40]; //request production rate
5  const double s_rate = 30; //request transmission rate
   // packet transmission rates
6  const double r_slow_rate = 10; //slow buffer
7  evolve double r_fast_rate [10..30]; //fast buffer
8  const double c_rate = 40; //packet consumption rate
   :
   // redirection
9  evolve module Buffer
   :
   //send
10 [] (slow=1) & (buffer_s<slow_max) -> s_rate *(1-buffer_s/slow_max) :
    (slow'=0) & (buffer_s' = buffer_s +1);
11 [] (slow=1) & (buffer_s>0) & (buffer_f<fast_max) -> s_rate *buffer_s/slow_max :
    (slow'=0) & (buffer_f' = buffer_f +1);
12 [] (fast=1) & (buffer_f<fast_max) -> s_rate :
    (fast'=0) & (buffer_f' = buffer_f +1);
   :
13 endmodule

```

Fig. 11. Variant of the producer-consumer model introduced in Section 2.

times, and thus are far from the optimal f_1 values.

Further, we observe that the maximum number of chunks per server, NC, has a major influence on the design robustness, with high NC values leading to highly sensitive designs. These designs should be avoided in favour of the alternative designs with low NC values depicted in Fig. 13 (for $\epsilon > 0$).

Workstation cluster (WC). Fig. 14 depicts the Pareto fronts obtained for all γ, ϵ combinations of the WC pCTMC model. These Pareto fronts show again how the large quality-attribute regions (corresponding to high-sensitivity designs) obtained for $\epsilon=0$ are “replaced” by much smaller quality-attribute regions on the Pareto fronts obtained for both $\epsilon > 0$ values. For instance, the fronts produced for $\gamma=0.005$ and $\epsilon \in \{0.05, 0.10\}$, include sub-optimal robust designs in the objective space $[0.6, 0.8] \times [40, 50]$ that do not exist for $\epsilon=0$. Further, the Pareto front for $\gamma=0.005, \epsilon=0.10$ includes a sub-optimal robust design in the objective space $[0.3, 0.5] \times [45, 70]$ to support the Pareto-optimal but volatile (i.e., highly sensitive) designs within that space. Similar observations can be made for other γ values.

With respect to the system dynamics, our sensitivity-aware synthesis method reveals that the most robust solutions correspond to the objective-function “extrema” from the Pareto front, i.e., to quality-attribute regions in which either f_1 is very high and f_2 is very low, or vice versa. In particular, solutions in the middle of quality-attribute regions are highly sensitive as indicated by the yellow-green boxes for $\gamma = 0.005$ and $\epsilon \in \{0.00, 0.05, 0.10\}$. The equivalent solutions are absent from the Pareto fronts for $\gamma = 0.01$ indicating that they are replaced by more robust solutions whose quality attributes are close to the low- and high-end of their respective ranges. Thus, if designers seek robust solutions they need to select designs that favour one of the quality attributes, since solutions with balanced trade-off between the quality attributes lead to either sensitive or sub-optimal robust designs.

We also identified an interesting property of the synthesized designs. Although they cover the entire design space for the real-valued parameters wsFail, wsCheck and wsRepair, the synthesized designs select very few values for the sub-cluster size N . In particular, in more

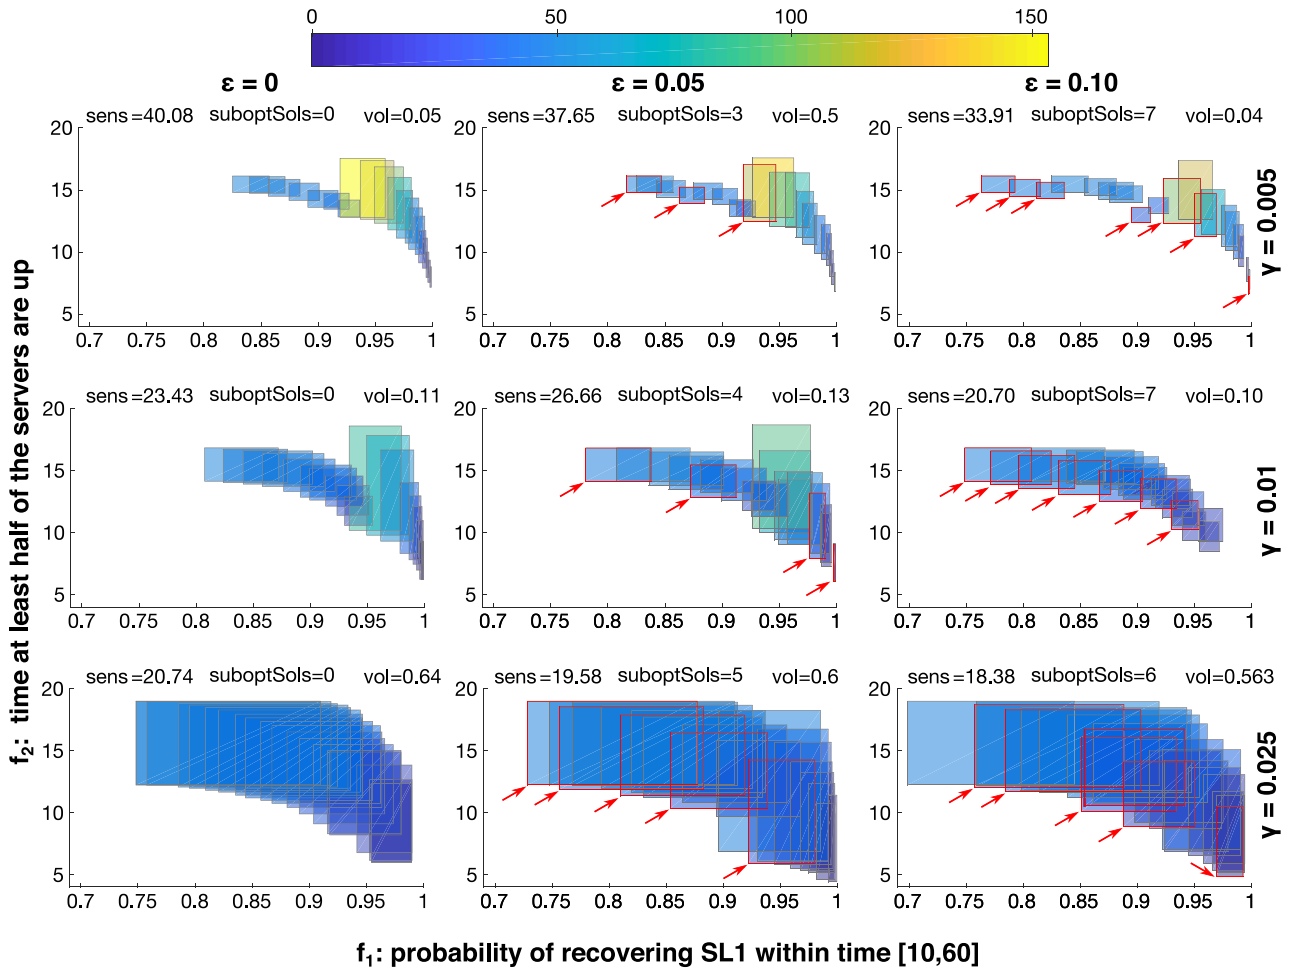


Fig. 12. Sensitivity-aware Pareto fronts for the GFS model. Legend and colour code are as in Fig. 7. Designs are compared based on the worst-case quality attribute value (i.e. lower-left corner of each box).

than 95% of the experiments $N \in \{10, 15\}$ and in the remaining $N \in \{9, 12\}$. We analysed further this observation and ran another experiment by setting the possible range for sub-cluster size $N \in \{11, \dots, 14\}$. Table 2 compares the average sensitivity between these two experiments for all γ, ϵ combinations. Our results validate that the ‘ideal’ values of the parameter N for the synthesised robust designs are 10 or 15. This finding demonstrates an unexpected and interesting relationship between the size of the cluster and robustness, impossible to derive through existing analysis methods.

RQ2 (Performance). Since the synthesis process is computationally demanding (see Appendix B), we evaluated the performance of RODES to analyse multiple candidate designs in parallel using the two-level parallelisation architecture described in Section 5.2. By employing the two-level parallelisation, we are able to partially alleviate the CPU overheads incurred not only due to the complexity of evaluating a candidate design but also due to the high number of evaluations. All experiments were run on a CentOS Linux 6.5 64bit server with two 2.6GHz Intel Xeon E5-2670 processors and 64GB memory. For the experiments involving GPU parallelisation, we used two nodes using either an nVidia K40 GPGPU card or an nVidia K80 GPGPU card.

The key results of our performance evaluation are described in Tables 3 and 4. The tables show the design synthesis run-times for $k = 500$ and $N = 20$ (i.e. for $kN = 10,000$ design evaluations), for our three case studies. Run-time statistics are computed over more than 30 independent runs, obtained using all combinations of $\epsilon \in \{0, 0.05, 0.1\}$ and $\gamma \in \{0.005, 0.01, 0.025\}$. Note that 10,000 evaluations, for which we obtained high quality sensitivity-aware Pareto fronts, are still

negligible with respect to the size of the design space that an exhaustive search would need to explore (theoretically the design space is uncountable). To demonstrate this difference, we list the number of candidate designs required to ‘cover’ the design space for a given tolerance value γ (this number is indeed much smaller than the total number of candidate designs). For PC model ($\gamma = 0.005$) it is around 20,000 designs, but for WC and GFS ($\gamma = 0.01$) it is more than 3 millions designs.

Results in Table 3 confirm that performance of the synthesis process is affected mainly by the size of the underlying pTCMC and by the average number of the discretisation steps required to evaluate particular quantitative attributes (around 4000 steps are required for WC and PC, 160,000 for GFS v1, and 46,000 for GFS v2). Note that this number depends on the highest time bound appearing in the properties and on the highest rate appearing in the transition matrix. This observation also explains the significant slowdown of the synthesis process when switching from v1 to v2 of GFS.

First, we evaluate the performance of CPU-only parallelisation at different numbers of cores. The results clearly confirm the scalability with respect to the number of cores. We can also observe that a better scalability is obtained for more complicated synthesis problems (i.e. 5.5-times speed for 10 cores on GFS v1 versus 7.9-times speed up for 10 core on GFS v2).

Second, we evaluate the performance of the two-level parallelisation. Table 4 compares the run-times for different number of CPU cores and GPU devices. In this configuration, we obtain a significant reduction of runtimes, e.g. for GFS v2 we obtain 8.4-times speedup with one GPU and one CPU core, and 7.6-times speedup with two GPUs and two

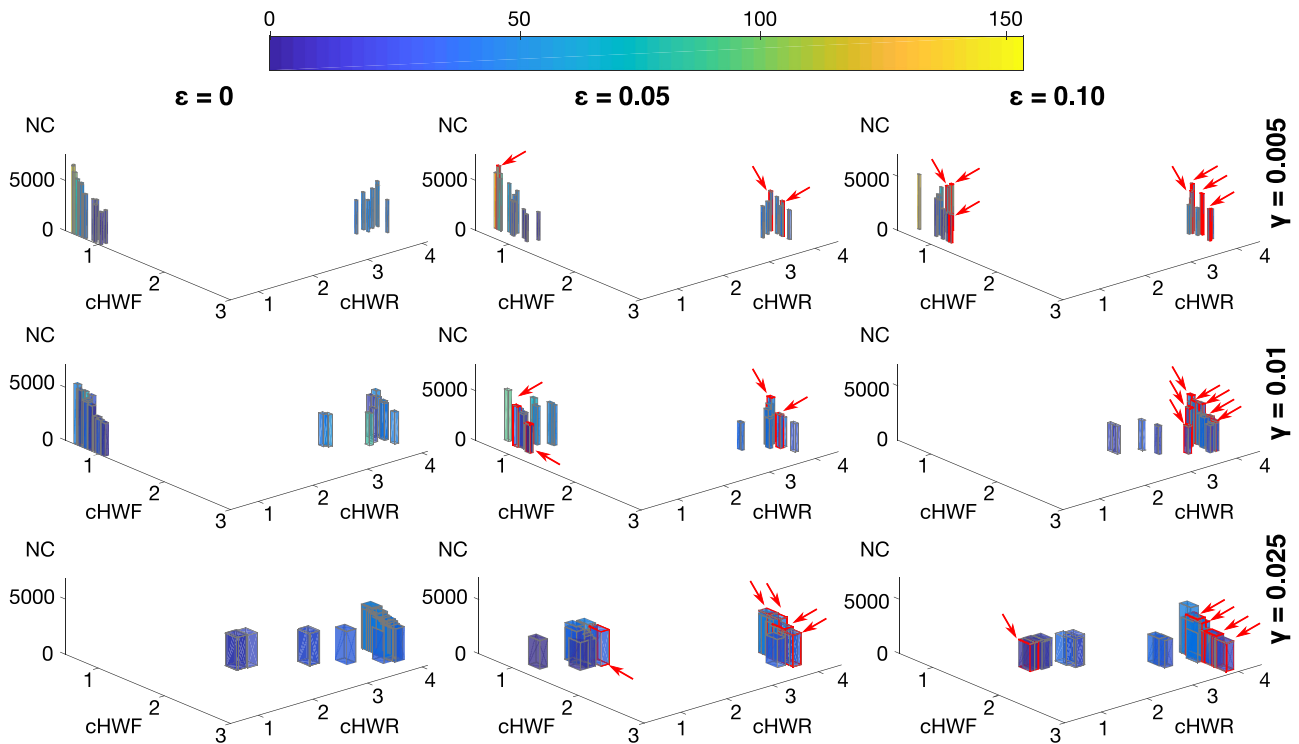


Fig. 13. Synthesised Pareto-optimal designs for the GFS model and experiments from Fig. 12. Rectangles in x-y plane correspond to the continuous parameter regions.

CPU cores. The slightly worse speedup observed in the latter case is due to the increased CPU-GPU communication overhead when more devices are employed.

Finally, we see that evaluating more than one candidate solutions (generated using several CPU cores) on a single GPU further improves the performance until the GPU is fully utilised (i.e. the maximal number of active threads that can be dispatched is reached and thus some parallel evaluations has to be serialised). The performance is also affected by the memory access pattern that depends on the concrete candidate solutions evaluated in parallel. In particular, the performance degrades when the memory access locality is decreased. Note that the maximal number of candidate solutions that can be evaluated in parallel on a single GPU is also limited by the GPU memory that has to accommodate the underlying pCTMC.

RQ3 (Metaheuristic effectiveness). To answer this research question, we analysed the goodness of the Pareto-optimal designs of the GFS model obtained with our NSGA-II-based RODES against a variant that uses random search (RS). For each variant and combination of $\epsilon \in \{0, 0.05, 0.10\}$ and $\gamma \in \{0.005, 0.01\}$ we carried out 30 independent runs, in line with standard SBSE practice (Harman et al., 2012b). As building the actual Pareto front for large design spaces is challenging and computationally expensive (GFS has $|\mathcal{P} \times \mathcal{A}| > 24E10$ assuming a three-decimal precision for continuous parameters), we again followed the standard practice and combined the sensitivity-aware Pareto fronts from all 60 RODES and RS runs for each ϵ, γ combination into a *reference Pareto front* (Zitzler et al., 2003). We then compared the Pareto fronts achieved by each variant against this reference front by using the metrics

$$M_1 = wI_{\epsilon_{norm}} + (1 - w)\overline{sen}\overline{s}_{norm}$$

and

$$M_2 = wI_{IGD_{norm}} + (1 - w)\overline{sen}\overline{s}_{norm}$$

which use a weight $w \in [0, 1]$ to combine normalised versions of the established (but sensitivity-agnostic) Pareto-front quality metrics I_{ϵ} and I_{IGD} (Zitzler et al., 2003) with the normalised design sensitivity. The

unary additive epsilon (I_{ϵ}) gives the minimum additive term by which the objectives of a particular design from a Pareto front must be altered to dominate the respective objectives from the reference front. The inverted generational distance (I_{IGD}) measures the shortest Euclidean distance from each design in the Pareto front to the closest design in the reference front. These indicators show convergence and diversity to the reference front (smaller is better).

Fig. 15 compares RODES and RS across our ϵ, γ combinations using metrics M_1 and M_2 with $w=0.5$. The RODES median is consistently lower than that of RS for all ϵ, γ combinations with the exception of $\epsilon=0, \gamma=0.01$ (which ignores design sensitivity) for M_2 . For a given γ , RODES results improve as ϵ increases, unlike the corresponding RS results. Thus, the difference between RODES and RS increases with larger ϵ for both metrics. This shows that RODES drives the search using sensitivity (10), and thus it can identify more robust designs. We confirmed these visual inspection findings using the non-parametric Mann-Whitney test with 95% confidence level ($\alpha=0.05$). We obtained statistical significance (p -value < 0.05) for all ϵ, γ combinations except for $\epsilon=0, \gamma=0.005$, with p -value in the range $[1.71E-06, 0.0026]$ and $[1.086E-10, 0.00061]$ for M_1 and M_2 , respectively.

Considering these results, we have sufficient empirical evidence that RODES synthesises significantly more robust designs than RS. These results are also in line with our previous work which demonstrated through extensive evaluation that probabilistic model synthesis using MOGAs achieves significantly better results than RS (Gerasimou et al., 2015). Hence, the problem of synthesising sensitivity-aware Pareto optimal sets (13) is challenging, as expected for any well-defined SBSE problem.

6.5. Threats to validity

Construct validity threats may arise due to assumptions made when modelling the three systems. To mitigate these threats, we used models and quality requirements based on established case studies from the literature (Ghemawat et al., 2003; Haverkort et al., 2000).

Internal validity threats may correspond to bias in establishing cause-

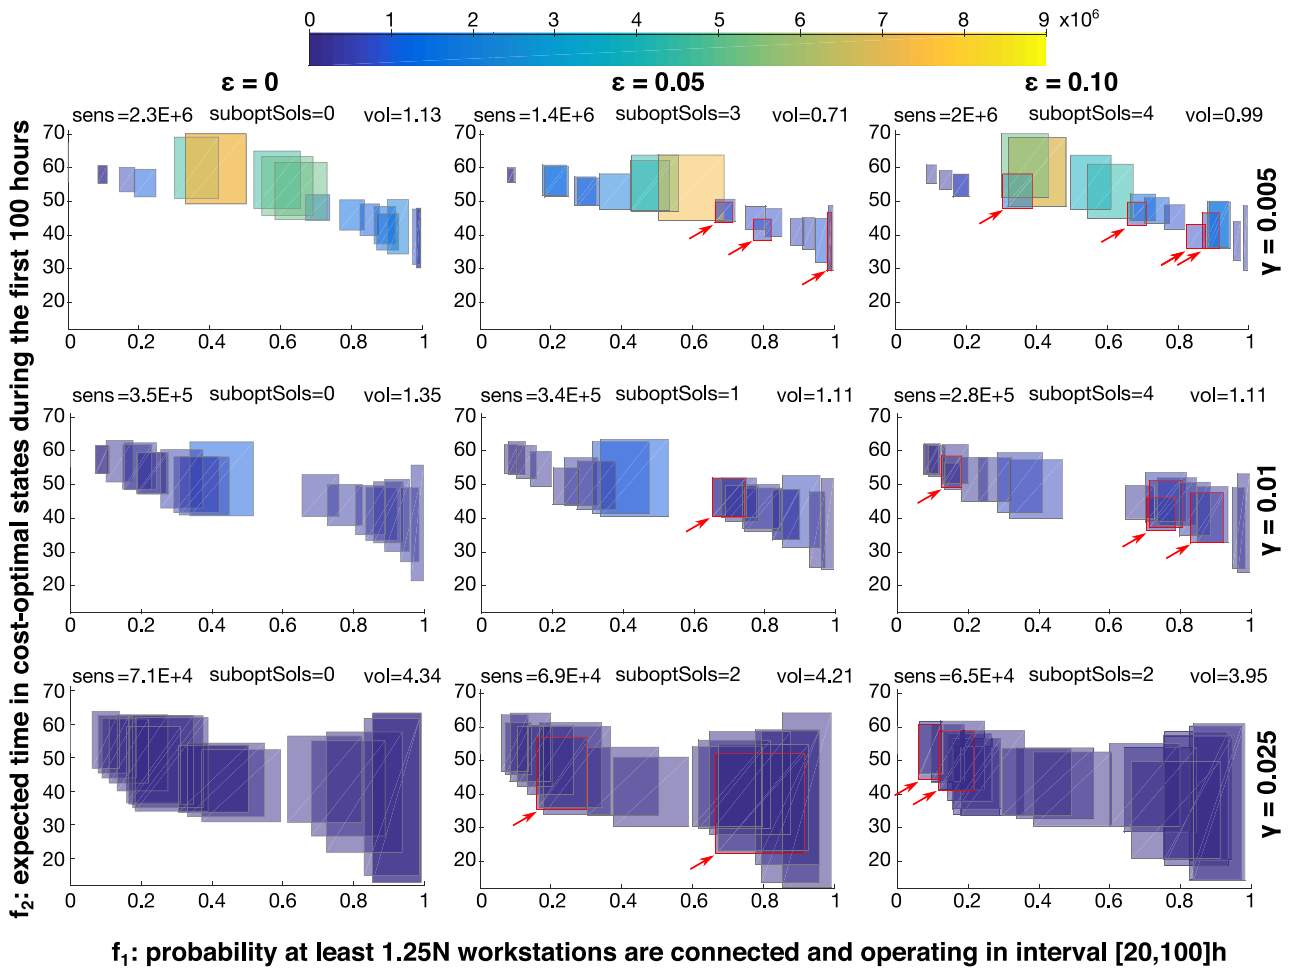


Fig. 14. Sensitivity-aware Pareto fronts for the workstation cluster model. Legend and colour code are as in Fig. 7. Designs are compared based on the worst-case quality attribute value (i.e. lower-left corner of each box).

Table 2

Average design sensitivity for two variants of the workstation cluster synthesis problem, given by different ranges for parameter N . Sensitivity-aware designs (i.e. where $\epsilon > 0$) for $N \in \{10..15\}$ have lower sensitivity than for $N \in \{11..14\}$.

N	Average sensitivity								
	$\gamma = 0.005, \epsilon = 0.00$	$\gamma = 0.005, \epsilon = 0.05$	$\gamma = 0.005, \epsilon = 0.10$	$\gamma = 0.01, \epsilon = 0.00$	$\gamma = 0.01, \epsilon = 0.05$	$\gamma = 0.01, \epsilon = 0.10$	$\gamma = 0.025, \epsilon = 0.00$	$\gamma = 0.025, \epsilon = 0.05$	$\gamma = 0.025, \epsilon = 0.10$
{10..15}	1.6E6	7.86E5	6.58E5	2.1E5	2.49E5	2.19E5	6.45E4	6.68E4	7.56E4
{11..14}	1.33E6	1.3E6	1.22E6	5.2E5	5.28E5	4.77E5	2E5	1.93E5	1.87E5

effect relationships in our experiments. We limit them by examining instantiations of the sensitivity-aware Pareto dominance relation (12) for multiple values of the sensitivity-awareness ϵ_i and tolerance level γ_k . To alleviate further the risk of biased results due to the MOGAs being stuck at local optimum and not synthesising a global optimum Pareto

front, we performed multiple independent runs. Although this scenario never occurred in our experiments, when detected, it can be solved by re-initialising the sub-population outside the Pareto front. Also, Algorithm 1 ensures that the Pareto front monotonically improves at each iteration. Finally, we enable replication by making all

Table 3

Time (mean \pm SD) in minutes for the synthesis using 10,000 evaluations for one-level CPU parallelisation. #states (#trans.): number of states (transitions) of the underlying pCTMC. |K|: number of continuous parameters.

Model	#states	#trans.	CPU (#cores)			
			1	2	5	10
WC (K = 3)	3440–8960	18,656–49424	394 \pm 25	217 \pm 29	118 \pm 14	68 \pm 8
PC (K = 2)	5632	21,968–24572	251 \pm 46	131 \pm 33	50 \pm 2	31 \pm 5
GFS v1 (K = 2)	1323–2406	7825–15545	390 \pm 27	267 \pm 49	125 \pm 19	71 \pm 10
GFS v2 (K = 2)	21,606	145,335–148245	19,011 \pm 400	8207 \pm 361	4562 \pm 36	2399 \pm 9

Table 4
Time (mean \pm SD) in minutes for the synthesis using 10,000 evaluations for two-level CPU+GPU parallelisation.

Model	CPU (#cores)			CPU (#cores)/GPU (#devices)				
	1	2	5	1/1	2/1	5/1	2/2	5/2
GFS v2	19,011 \pm 400	8207 \pm 361	4562 \pm 36	2264 \pm 33	1736 \pm 8	1625 \pm 16	1082 \pm 3	1043 \pm 22

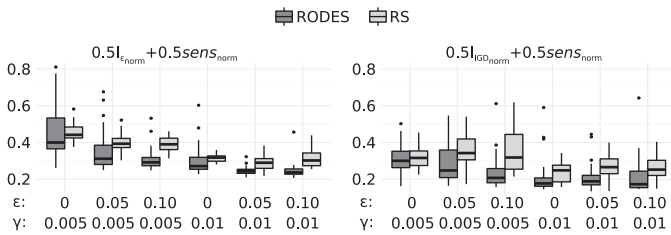


Fig. 15. RODES vs. random search (RS) comparison for combinations of $\gamma \in \{0.005, 0.01\}$ and $\epsilon \in \{0, 0.05, 0.10\}$, over 30 independent GFS runs. For both metrics – I_ϵ indicator and sensitivity (left) and I_{GD} indicator and sensitivity (right) – smaller is better.

experimental results publicly available on the project webpage.

External validity threats might exist if the search for robust designs for other systems cannot be expressed as a pCTMC synthesis problem using objective functions (8) and constraints (9). We limit these threats by specifying pCTMCs in an extended variant of the widely used modelling language of PRISM (Kwiatkowska et al., 2011), with objective functions and constraints specified in the established temporal logic CSL. PRISM parametric Markov models are increasingly used to model software architectures, e.g. in the emerging field of self-adaptive software (Calinescu and Kwiatkowska, 2009; Calinescu et al., 2015; Moreno et al., 2015; Gerasimou et al., 2014). Another threat might occur if our method generated a Pareto front that approached the actual Pareto front insufficiently, producing only low quality designs or designs that did not satisfy the required quality constraints. We mitigated this threat by using established Pareto-front performance indices to confirm the quality of the Pareto fronts from our case studies. Nevertheless, additional experiments are needed to establish the applicability and feasibility of the method in domains with characteristics different from those used in our evaluation.

7. Related work

RODES builds on the significant body of *software performance and reliability engineering* research that employs formal models to analyse the quality attributes of alternative software designs (e.g. Balsamo et al., 2004; Bondy, 2014; Becker et al., 2009; Fiondella and Puliafito, 2016; Stewart, 2009; Woodside et al., 2014). Approaches based on formal models such as queueing networks (Balsamo et al., 2003), Petri nets (Lindemann, 1998), stochastic models (Calinescu et al., 2016; Sharma and Trivedi, 2007) and timed automata (Hessel et al., 2008; Larsen, 2014), and tools for their simulation (e.g. Palladio (Becker et al., 2009)) and verification (e.g. PRISM (Kwiatkowska et al., 2011) and UPPAAL (Hessel et al., 2008)) have long been used for this analysis. However, unlike RODES, these approaches can only analyse alternative models through a tedious iterative process carried out manually by experts.

Performance antipatterns can be used to speed up this process by avoiding the analysis of poor designs (Arcelli et al., 2012; Smith and Williams, 2000; Cortellessa et al., 2010), but approaches that automate the search for correct or optimal designs have only been proposed recently. Three types of such approaches are related to RODES. Given a Markov model that violates a quality requirement, the first approach—called *probabilistic model repair* (Bartocci et al., 2011; Chen et al., 2013)—automatically adjusts its transition probabilities to

produce a “repaired” model that meets the requirement. The second approach is called *precise parameter synthesis* (Češka et al., 2017), and works by identifying transition rates that enable continuous-time Markov models to satisfy a quality requirement or to optimise a quality attribute of the system under development. Finally, our previous work on *probabilistic model synthesis* (Gerasimou et al., 2015) applies multi-objective optimisation and genetic algorithms to a design template that captures alternative system designs, and generates the Pareto-optimal set of Markov models associated with the quality optimisation criteria of the system. While these approaches represent a significant advance over the previously manual methods of alternative design analysis, they do not take into account the robustness of their repaired or synthesised models. Likewise, the approach from Martens et al. (2010) employs evolutionary algorithms to search the configuration space of Palladio Component Models, but the synthesis process does not reflect the sensitivity of the candidate models.

Syntax-guided synthesis has been used to find probabilistic programs that best match the available data (Nori et al., 2015), including synthesis from “sketches”, i.e. partial programs with incomplete details (Solar-Lezama et al., 2005). In Solar-Lezama et al. (2006), counter-example guided inductive synthesis (CEGIS) has been introduced as an SMT-based synthesiser for sketches and, due to the enormous improvement of SMT solvers in the last decade, CEGIS is currently able to find deterministic programs for a variety of challenging problems (Solar-Lezama et al., 2005; 2008). Very recently, the concept of meta-sketches introducing the “optimal synthesis problem” has been proposed (Bornholt et al., 2016) and adapted for synthesis of stochastic reaction networks (Cardelli et al., 2017). These solutions are complementary to RODES, as they explore other aspects of design alternatives, and do not take robustness into account.

Methods that rigorously evaluate how the transition probabilities affect the satisfiability of temporal properties (expressed as probabilistic temporal logic formulae) have been studied in the context of parameter synthesis. The methods either construct symbolic expressions describing the satisfaction probability as a function of the model parameters (Dehnert et al., 2015; Hahn et al., 2011), or compute—for given intervals of parameter values—safe bounds on the satisfaction probability (Quatmann et al., 2016). In contrast to this work, our robust design synthesis directly integrates sensitivity analysis into the automated design process.

Another research area related to RODES is *sensitivity analysis*, which analyses the impact of parameter changes on the performance, reliability, cost and other quality attributes of the system under development (e.g. Gokhale and Trivedi, 2002; Lo et al., 2005; Huang and Lyu, 2005). However, sensitivity analysis typically operates by sampling the parameter space and evaluating the system quality attributes for the sampled values. As such, the result is not guaranteed to reflect the whole range of quality-attribute values for the parameter region of interest. RODES does not have this drawback, as it operates with close over-approximations of the quality-attribute regions for the synthesised robust designs. The perturbation theory for Markov processes has been applied to analysing the sensitivity of software operational profiles (Kamavaram and Goseva-Popstojanova, 2003). However, this approach quantifies the effect of variations in model transition probabilities without synthesising the analysed solutions. Furthermore, RODES supports a wide range of continuous and discrete parameters that cannot be used with the approach from Kamavaram and Goseva-Popstojanova (2003). Stochastic analysis of architectural models was

used for early predictions of system component reliability and sensitivity with respect to different operational profiles (Cheung et al., 2008). Unlike RODES, the research from Cheung et al. (2008) focuses on exploiting different architectural models and associated analysis techniques, and is therefore complementary to the work presented in our paper.

The *smoothed model checking* technique from Bortolussi et al. (2016a) computes an analytical approximation of the satisfaction probability of a formula over a parametric CTMC. While not providing the same guarantees as the safe over-approximation method from RODES, the technique was experimentally shown to be highly accurate, so it can be used to estimate the sensitivity of a probabilistic temporal logic property to variations in the CTMC parameters.

Finally, the problem of parameter synthesis of stochastic reaction networks with respect to multi-objective specification has been recently considered in Bortolussi et al. (2016b). The authors employ statistical methods to estimate how kinetic parameters affect the satisfaction probability and average robustness of Signal Temporal Logic properties. In contrast to our approach, a candidate solution from Bortolussi et al. (2016b) has all parameters fixed and the robustness captures how far the candidate is from violating the particular properties.

8. Conclusion

Robustness is a key and yet insufficiently explored characteristic of

Appendix A. Proof of theorem 1

We show that the sensitivity-aware Pareto dominance relation defined in Definition 4 is a strict order.

Proof. We need to show that relation $<$ from (12) is irreflexive and transitive. For any $d \in \mathcal{F}$, $d < d$ would require that $f_i(d) < (1 + \epsilon_i)f_i(d)$ or $f_i(d) < f_i(d)$ for some $i \in I$, which is impossible. Thus, $<$ is irreflexive. To show that $<$ is transitive, consider three designs $d, d', d'' \in \mathcal{F}$ such that $d < d'$ and $d' < d''$. According to (12), we have $\forall i \in I. f_i(d) \leq f_i(d')$ and $\forall i \in I. f_i(d') \leq f_i(d'')$, so $\forall i \in I. f_i(d) \leq f_i(d'')$ due to the transitivity of \leq . Furthermore, at least one half of the disjunction from definition (12) must hold for each of $d' < d''$ and $d' < d''$. We have three cases. Assume first that the left half holds for $d < d'$, i.e. that $(1 + \epsilon_{i_1})f_{i_1}(d) < f_{i_1}(d')$ for some $i_1 \in I$; as $f_{i_1}(d') \leq f_{i_1}(d'')$, we also have $(1 + \epsilon_{i_1})f_{i_1}(d) < f_{i_1}(d'')$, so $d < d''$ in this case. Assume now that left half of disjunction (12) holds for $d' < d''$, i.e., that $(1 + \epsilon_{i_1})f_{i_1}(d') < f_{i_1}(d'')$ for some $i_1 \in I$; as $f_{i_1}(d) \leq f_{i_1}(d')$, we again have $(1 + \epsilon_{i_1})f_{i_1}(d) < f_{i_1}(d'')$ and $d < d''$. Finally, consider that only the right half of disjunction (12) holds for both $d < d'$ and $d' < d''$. In this last case, $sens(d) \leq sens(d') \leq sens(d'')$ and there is an $i_1 \in I$ such that $f_{i_1}(d) < f_{i_1}(d') \leq f_{i_1}(d'')$, so also $d < d''$, and therefore $<$ is transitive. \square

Appendix B. Complexity analysis

The time complexity of Algorithm 1 representing the synthesis process is

$$\mathcal{O}(k \cdot N \cdot (|I| + |J|) \cdot t + k \cdot |I| \cdot N^2),$$

where k is the number of iterations of the (MOGA) while loop (i.e. the number of generations); $N = |CD|$ is the size of the candidate design population; $|I| + |J|$ is the overall number of objective functions and constraints; and t is the time required to analyse a quality attribute of a candidate design. The term $k \cdot N \cdot (|I| + |J|) \cdot t$ quantifies the overall complexity of evaluating candidate designs, while $k \cdot |I| \cdot N^2$ corresponds to comparing designs and building the front in lines 7–14 of Algorithm 1.

The factor t depends on the size of the underlying state space and on the number of discrete-time steps required to evaluate the particular quality attributes. As shown in Češka et al. (2017), $t = \mathcal{O}(t_{CSL} \cdot t_{pCSL})$. The factor $t_{CSL} = |\phi| \cdot M \cdot q \cdot t_{\max}$ is the worst-case time complexity of time-bounded CSL model checking Kwiatkowska et al. (2007), where $|\phi|$ is the length of the input CSL formula ϕ , t_{\max} is the highest time bound occurring in it, M is the number of non-zero elements in the rate matrix and q is the highest rate in the matrix. The factor t_{pCSL} is due to the parametric analysis of the design and depends on the form of polynomials appearing in the parametric rate matrix $\mathcal{D}'_{\mathbf{R}}$. Models of software systems typically include only linear polynomials, for which $t_{pCSL} = \mathcal{O}(n)$, where n is the number of continuous parameters.

References

- Alur, R., Bodik, R., Juniwal, G., et al., 2013. Syntax-guided synthesis. FMCAD'13. pp. 1–8.
- Arcelli, D., Cortellessa, V., Trubiani, C., 2012. Antipattern-based model refactoring for software performance improvement. QoSA'12. pp. 33–42.
- Baier, C., Hahn, E.M., Haverkort, B., et al., 2013. Model checking for performability. Math. Struct. Comp. Sci. 23 (4), 751–795.
- Balsamo, S., Di Marco, A., Inverardi, P., Simeoni, M., 2004. Model-based performance prediction in software development: a survey. IEEE Trans. Softw. Eng. 30 (5), 295–310.
- Balsamo, S., Personè, V.D.N., Inverardi, P., 2003. A review on queueing network models with finite capacity queues for software architectures performance prediction. Perform. Eval. 51 (2), 269–288.
- Bartocci, E., Grosu, R., Katsaros, P., et al., 2011. Model repair for probabilistic systems. TACAS'11. pp. 326–340.
- Becker, S., Koziolok, H., Reussner, R., 2009. The palladio component model for model-driven performance prediction. J. Syst. Softw. 82 (1).
- Bondy, A.B., 2014. Foundations of Software and System Performance Engineering. Addison Wesley.
- Bornholt, J., Torlak, E., Grossman, D., Ceze, L., 2016. Optimizing synthesis with metasketches. Proc. of POPL'16. ACM, pp. 775–788.

- Bortolussi, L., Milios, D., Sanguinetti, G., 2016. Smoothed model checking for uncertain continuous-time markov chains. *Inf. Comput.* 247, 235–253.
- Bortolussi, L., Policriti, A., Silveti, S., 2016. Logic-based multi-objective design of chemical reaction networks. *HSB'16*. Springer, pp. 164–178.
- Calinescu, R., Gerasimou, S., Banks, A., 2015. Self-adaptive software with decentralised control loops. *FASE*. pp. 235–251.
- Calinescu, R., Ghezzi, C., Johnson, K., et al., 2016. Formal verification with confidence intervals to establish quality of service properties of software systems. *IEEE Trans. Rel.* 65 (1), 107–125.
- Calinescu, R., Kwiatkowska, M., 2009. CADs*: Computer-aided development of self-* systems. *FASE*. pp. 421–424.
- Calinescu, R., Češka, M., Gerasimou, S., Kwiatkowska, M., Paoletti, N., 2017. Designing robust software systems through parametric Markov chain synthesis. *ICSA. IEEE*, pp. 131–140.
- Calinescu, R., Češka, M., Gerasimou, S., Kwiatkowska, M., Paoletti, N., 2017. RODES: A robust-design synthesis tool for probabilistic systems. *QUEST*. pp. 304–308.
- Cardelli, L., Češka, M., Fränzle, M., Kwiatkowska, M., Laurenti, L., Paoletti, N., Whitby, M., 2017. Syntax-guided optimal synthesis for chemical reaction networks. *CAV'17*. Springer, pp. 375–395.
- Češka, M., Dannenberg, F., Paoletti, N., et al., 2017. Precise parameter synthesis for stochastic biochemical systems. *Acta Informatica* 54, 589–623.
- Češka, M., Pilař, P., Paoletti, N., Brim, L., Kwiatkowska, M., 2016. PRISM-PSY: Precise GPU-accelerated Parameter Synthesis for Stochastic Systems. *TACAS'16*. Springer, pp. 367–384.
- Chen, T., Hahn, E.M., Han, T., et al., 2013. Model repair for Markov decision processes. *TASE'13*. pp. 85–92.
- Cheung, L., Roshandel, R., Medvidovic, N., Golubchik, L., 2008. Early prediction of software component reliability. *ICSE. ACM*, pp. 111–120.
- Cortellessa, V., Martens, A., Reussner, R., Trubiani, C., 2010. A process to effectively identify 'guilty' performance antipatterns. *FASE'10*. pp. 368–382.
- Deb, K., Pratap, A., Agarwal, S., Meyarivan, T., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comp.* 6 (2), 182–197.
- Dehnert, C., Junges, S., Jansen, N., et al., 2015. PROPHECY: a probabilistic parameter synthesis tool. *CAV'15*. Springer, pp. 214–231.
- Dehnert, C., Junges, S., Katoen, J.-P., Volk, M., 2017. A STORM is coming: a modern probabilistic model checker. *CAV'17*. Springer, pp. 592–600.
- Filieri, A., Tamburrelli, G., Ghezzi, C., 2016. Supporting self-adaptation via quantitative verification and sensitivity analysis at run time. *IEEE Trans. Softw. Eng.* 42 (1), 75–99.
- Fiondella, L., Puliafito, A., 2016. Principles of Performance and Reliability Modeling and Evaluation. Springer Series in Reliability Engineering.
- Gerasimou, S., Calinescu, R., Banks, A., 2014. Efficient runtime quantitative verification using caching, lookahead, and nearly-optimal reconfiguration. *SEAMS*. pp. 115–124.
- Gerasimou, S., Tamburrelli, G., Calinescu, R., 2015. Search-based synthesis of probabilistic models for QoS software engineering. *ASE'15*. pp. 319–330.
- Ghemawat, S., Gobioff, H., Leung, S.-T., 2003. The google file system. *SOSP'03*. pp. 29–43.
- Gokhale, S.S., Trivedi, K.S., 2002. Reliability prediction and sensitivity analysis based on software architecture. *ISSRE'03*. pp. 64–75.
- Hahn, E.M., Hermanns, H., Zhang, L., 2011. Probabilistic reachability for parametric markov models. *STTT* 13 (1), 3–19.
- Han, T., Katoen, J., Mereacre, A., 2008. Approximate parameter synthesis for probabilistic time-bounded reachability. *RTSS*. pp. 173–182.
- Harman, M., Mansouri, S.A., Zhang, Y., 2012. Search-based software engineering: trends, techniques and applications. *ACM Comp. Surv.* 45 (1), 11:1–11:61.
- Harman, M., McMinn, P., de Souza, J., Yoo, S., 2012. Search based software engineering: techniques, taxonomy, tutorial. *Emp. Softw. Eng. and Verif.* pp. 1–59.
- Haverkort, B.R., Hermanns, H., Katoen, J.-P., 2000. On the use of model checking techniques for dependability evaluation. *SRDS'00. IEEE*, pp. 228–237.
- Hessel, A., Larsen, K.G., Mikucionis, M., et al., 2008. Testing real-time systems using UPPAAL. *Formal Methods and Testing*. Springer, pp. 77–117.
- Huang, C.-Y., Lyu, M.R., 2005. Optimal testing resource allocation, and sensitivity analysis in software development. *Trans. Reliab.* 54 (4), 592–603.
- International Organization for Standardization, 2010. ISO 286–1:2010, Geometrical product specifications (GPS) – ISO code system for tolerances on linear sizes.
- International Organization for Standardization, 2013. ISO general purpose metric screw threads – Tolerances.
- Kamavaram, S., Goseva-Popstojanova, K., 2003. Sensitivity of software usage to changes in the operational profile. *NASA Soft. Eng. Workshop*.
- Kitano, H., 2004. Biological robustness. *Nat. Rev. Genet.* 5, 826–837.
- Koza, J.R., 1992. Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press.
- Kwiatkowska, M., Norman, G., Parker, D., 2007. Stochastic model checking. *SFM'07*. pp. 220–270.
- Kwiatkowska, M., Norman, G., Parker, D., 2011. PRISM 4.0: verification of probabilistic real-time systems. *CAV'11*. pp. 585–591.
- Larsen, K.G., 2014. Verification and performance analysis of embedded and cyber-physical systems using uppaal. *MODELSWARD'14*.
- Lindemann, C., 1998. Performance modelling with deterministic and stochastic petri nets. *Perform. Eval. Rev.* 26 (2), 3.
- Lo, J.-H., Huang, C.-Y., Chen, I.-Y., et al., 2005. Reliability assessment and sensitivity analysis of software reliability growth modeling based on software module structure. *J. Syst. Softw.* 76 (1), 3–13.
- Martens, A., Koziolok, H., Becker, S., Reussner, R., 2010. Automatically improve software architecture models for performance, reliability, and cost using evolutionary algorithms. *WOSP/SIPEW*. pp. 105–116.
- Moreno, J.C., Lopes, A., Garlan, D., Schmerl, B., 2015. Impact models for architecture-based self-adaptive systems. *FACS*. pp. 89–107.
- Nebro, A.J., Durillo, J.J., Luna, F., et al., 2009. MOCeL: A cellular genetic algorithm for multiobjective optimization. *J. Intell. Syst.* 24 (7), 726–746.
- Nori, A.V., Ozair, S., Rajamani, S.K., Vijaykeerthy, D., 2015. Efficient synthesis of probabilistic programs. *PLDI'14*. pp. 208–217.
- Phadke, M.S., 1995. Quality Engineering using Robust Design. Prentice Hall PTR.
- Quatmann, T., Dehnert, C., Jansen, N., et al., 2016. Parameter synthesis for Markov models: faster than ever. *ATVA'16*. pp. 50–67.
- Reyes-Sierra, M., Coello, C.C., 2006. Multi-objective particle swarm optimizers: a survey of the state-of-the-art. *Int. J. Comput. Intell. Res.* 2 (3), 287–308.
- Sharma, V.S., Trivedi, K.S., 2007. Quantifying software performance, reliability and security: an architecture-based approach. *J. Syst. Softw.* 80 (4), 493–509.
- Smith, C.U., Williams, L.G., 2000. Software performance antipatterns. *Workshop on Software and Performance*. 17. pp. 127–136.
- Solar-Lezama, A., Jones, C.G., Bodik, R., 2008. Sketching concurrent data structures. *Proc. of PLDI'08. ACM*, pp. 136–148.
- Solar-Lezama, A., Rabbah, R., Bodik, R., Ebcioğlu, K., 2005. Programming by sketching for bit-streaming programs. *Proc. of PLDI'05. ACM*, pp. 281–294.
- Solar-Lezama, A., Tancu, L., Bodik, R., Seshia, S., Saraswat, V., 2006. Combinatorial sketching for finite programs. *Proc. of ASPLOS'06. ACM*, pp. 404–415.
- Stewart, W.J., 2009. Probability, Markov Chains, Queues, and Simulation: The Mathematical Basis of Performance Modeling. Princeton University Press.
- Woodside, M., Petriu, D., Merseguer, J., Petriu, D., Alhaj, M., 2014. Transformation challenges: from software models to performance models. *J. Softw. Syst. Model.* 13 (4), 1529–1552.
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C., da Fonseca, V., 2003. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Trans. Evol. Comp.* 7 (2), 117–132.

Radu Calinescu is a senior lecturer within the Department of Computer Science at the University of York, UK. His main research interests are in formal methods for self-adaptive, autonomous, secure and dependable software systems, and in performance and reliability software engineering. He is an active promoter of formal methods at runtime as a way to improve the integrity and predictability of selfadaptive and autonomous software systems and processes.

Milan Češka has completed his Ph.D. thesis in the area of data-parallel algorithms for model checking at the Faculty of Informatics, Masaryk University, Brno, Czech Republic in Jun 2012. Afterwards, he was a research assistant in the Systems Biology Laboratory at the same faculty under supervision of prof. Lubos Brim. From February 2014, he was a postdoctoral researcher at Department of Computer Science at Oxford University in the group led by prof. Marta Kwiatkowska. In May 2016, he returned to Czech Republic as an assistant professor at Faculty of Information Technology, Brno University of Technology. His current research interests include formal methods for automated design of probabilistic and approximate systems, and formal analysis of biochemical systems.

Simos Gerasimou is a research associate within the Department of Computer Science at the University of York. He received his BSc in Computer Science from the University of Cyprus, Cyprus, in 2010, and his MSc in Software Engineering and PhD in Computer Science from the University of York, UK, in 2011 and 2017, respectively. His research interests lie in the area of self-adaptive and autonomous systems with a focus on methods that enable dependable system adaptation. Other research interests include runtime quantitative verification, search-based software engineering, model-driven engineering, and artificial intelligence.

Marta Kwiatkowska is Professor of Computing Systems and Fellow of Trinity College, University of Oxford. Prior to this she was Professor in the School of Computer Science at the University of Birmingham, Lecturer at the University of Leicester and Assistant Professor at the Jagiellonian University in Cracow, Poland. Kwiatkowska has made fundamental contributions to the theory and practice of model checking for probabilistic systems, focusing on automated techniques for verification and synthesis from quantitative specifications. She led the development of the PRISM model checker (www.prismmodelchecker.org), the leading software tool in the area and winner of the HVC Award 2016. Probabilistic model checking has been adopted in many diverse fields, including distributed computing, wireless networks, security, robotics, game theory, systems biology, DNA computing and nanotechnology, with genuine flaws found and corrected in real-world protocols. Kwiatkowska awarded an honorary doctorate from KTH Royal Institute of technology in Stockholm in 2014 and the Royal Society Milner Medal in 2018. Her recent work was supported by the ERC Advanced Grant VERIWARE “From software verification to ‘everyware’ verification” and the EPSRC Programme Grant on Mobile Autonomy. She is a Fellow of ACM and Member of Academia Europea.

Nicola Paoletti is a postdoctoral associate at the Department of Computer Science at Stony Brook University. He obtained an MSc in Computer Science in 2010 and a Ph.D in Information Sciences and Complex Systems in 2014 at the University of Camerino, Italy. From 2014 to 2016, he was a post-doctoral researcher at the Department of Computer Science, University of Oxford, UK. His research interests are in the verification, control, and synthesis of stochastic and hybrid systems, with application to biological and biomedical systems.