

# A Systematic Methodology to Migrate Complex Real-Time Software Systems to Multi-Core Platforms

Shaik Mohammed Salman<sup>a</sup>, Alessandro V. Papadopoulos<sup>b</sup>, Saad Mubeen<sup>b</sup>, Thomas Nolte<sup>b</sup>

<sup>a</sup>ABB AB, Västerås, Sweden

<sup>b</sup>Mälardalen University, Västerås, Sweden

---

## Abstract

This paper proposes a systematic three-stage methodology for migrating complex real-time industrial software systems from single-core to multi-core computing platforms. Single-core platforms have limited computational capabilities that prevent integration of computationally demanding applications such as image processing within the existing system. Modern multi-core processors offer a promising solution to address these limitations by providing increased computational power and allowing parallel execution of different applications within the system. However, the transition from traditional single-core to contemporary multi-core computing platforms is non-trivial and requires a systematic and well-defined migration process. This paper reviews some of the existing migration methods and provides a systematic multi-phase migration process with emphasis on software architecture recovery and transformation to explicitly address the timing and dependability attributes expected of industrial software systems. The methodology was evaluated using a survey-based approach and the results indicate that the presented methodology is feasible, usable and useful for real-time industrial software systems.

*Keywords:* Real-time systems, multi-core, software architecture, software migration, robotics.

---

## 1. Introduction

Software evolution has been a continuous process in industrial real-time embedded software systems with new functionality, performance improvements and bug fixes introduced with each new version, revision or release [1, 2]. Many of the industrial systems have been developed over the decades [3], undergoing major revisions due to technology shifts, changing customer requirements, improved development processes, among others. One constant factor associated with the evolution of such systems is that the software architectures and the implementations have focused on single-core computing platforms. Integrating new data-intensive and computationally demanding applications within the system, however, requires additional computational capacity. Moreover, with the decreasing availability of the single-core processors, migrating the existing software to multi-core computing platforms is becoming a

necessity. By *migration*, we refer to the modification of the existing software to execute on the multi-core platforms, while ensuring that the performance and quality attributes, such as dependability [4, 5], match the current system quality and more optimistically, improved much further. Such migration is essential since the long life-cycle of existing software systems has resulted in the creation of assets that have become critical for a business [6] and that a complete redevelopment may not be feasible.

Migrating existing real-time software systems towards multi-core systems requires (i) Identifying the timing requirements of the existing software systems and (ii) Identifying the technical solutions that can improve the performance, resource usage and the timing predictability of the software systems [7, 8, 9]. Invariably, any migration approach should also address the extra-functional attributes such as scalability, maintainability and portability of the software. Furthermore, the migration should consider maximum reuse of the existing software while minimizing the re-engineering efforts.

To address these aspects for the migration of a complex real-time software system with strict timing and

---

*Email addresses:* shaik.salman@se.abb.com (Shaik Mohammed Salman), alessandro.papadopoulos@mdh.se (Alessandro V. Papadopoulos), saad.mubeen@mdh.se (Saad Mubeen), thomas.nolte@mdh.se (Thomas Nolte)

43 dependability requirements, we used a focus group dis- 86  
44 cussion to formulate an open-ended Research Question 87  
45 (RQ), 88

46 **RQ: How to migrate a complex real-time software** 89  
47 **from a single-core to a multi-core architecture** 90  
48 **with maximum software reuse and minimal re-** 91  
49 **engineering effort?** 92

50 We further refined this question into the following sub- 93  
51 questions: 94

52 RQ1: Which migration methodology addresses the con- 95  
53 cerns of software reuse, dependability and timing 96  
54 requirements? 97

55 RQ2: How to evaluate and analyse the applicability of 98  
56 different multi-core solutions for embedded con- 99  
57 trol software? 100

58 RQ3: What are the tools that facilitate the migration pro- 101  
59 cess? 102

60 These questions were motivated by the need for mi- 103  
61 grating a configurable robot controller software [4] de- 104  
62 veloped at ABB Robotics<sup>1</sup>, with functionality ranging 105  
63 from motion control to cloud connectivity. The con- 106  
64 troller software has close to 140 tasks and 71,128 meth- 107  
65 ods, integrating real-time and non real-time functionali- 108  
66 ties with varying Quality of Service (QoS) requirements 109  
67 on a single-core platform. 110

68 To address the discussed questions, we used a mixed 111  
69 research methodology utilising discussions within a fo- 112  
70 cus group and subject experts, complemented with a 113  
71 review of the state-of-the-art literature, to identify key 114  
72 concerns and provide a systematic methodology to mi- 115  
73 grate industrial software with real-time requirements 116  
74 from single-core to multi-core platforms. Concretely, 117  
75 the paper provides the following contributions: 118

- 76 • A systematic methodology for migrating complex 119  
77 embedded software from single-core to multi-core 120  
78 platforms; 121
- 79 • A review of tools that facilitate the migration process; 122  
80 and 123
- 81 • A survey-based evaluation of the proposed methodol- 124  
82 ogy. 125

83 This paper reinforces the validity of the methodology 126  
84 presented in our previous work [10] by including a 127  
85 survey-based evaluation of the methodology. 128

<sup>1</sup>[https://new.abb.com/products/robotics/](https://new.abb.com/products/robotics/controllers)  
controllers 129

86 The rest of the paper is organised as follows. Sec- 87  
88 tion 2 provides an overview of a robotic system and 89  
89 its controller software. Section 3 reviews the exist- 90  
90 ing software migration methods. Section 4 provides an 91  
91 overview of the overall methodology. Section 5 includes 92  
92 a systematic approach focusing on architecture migra- 93  
93 tion, followed by implementation and verification of the 94  
94 migration in Section 6 and Section 7 respectively. A 95  
95 review of the tools facilitating the migration process is 96  
96 discussed in Section 8. Section 9 presents the evalua- 97  
97 tion of the proposed methodology. Finally, Section 10 98  
98 concludes the paper. 99

## 98 2. System Overview

99 The system corresponds to a typical robotic system 100  
100 consisting of a manipulator arm, a controller, and a 101  
101 graphical controller interface. The paper focuses on 102  
102 the software functionality of the controller, which can 103  
103 be divided into functions concerning (i) configuration, 104  
104 (ii) communication, and (iii) control. The configura- 105  
105 tion functions provide the robot programming interface that 106  
106 allows a user to configure and specify the runtime be- 107  
107 haviour of the manipulator. The user is also able to de- 108  
108 fine the robot environment such as additional sensors 109  
109 and actuators. The real-time communication functions 110  
110 allow the controller to interactThe communication func- 111  
111 tions provide a real-time networking capability to en- 112  
112 able the controller to interact with devices such as Pro- 113  
113 grammable Logic Controllers (PLCs). It also includes a 114  
114 non-real-time communication capability that allows the 115  
115 controller to interact with enterprise network including 116  
116 PCs and the cloud. The control functions generate the 117  
117 path the manipulator has to follow based on the user- 118  
118 defined configuration. The output of the control func- 119  
119 tions is used to drive controllers that manage the low- 120  
120 level motor actuation. 121

121 The controller software has different runtime modes 122  
122 and the available functions vary between the modes. 123  
123 The main modes include the “Initialisation mode”, 124  
124 “Safe-init mode”, “System update and configuration 125  
125 mode”, “Normal operation mode”, and “Fail-safe 126  
126 mode” [11]. The different modes and the transition be- 127  
127 tween the modes is shown in Fig. 1. At startup, the con- 128  
128 troller transitions into the initialisation mode. Here all 129  
129 the tasks are initialised with values based on the pre- 130  
130 viously saved configuration settings. The controller 131  
131 software is in the initialisation mode during startup. It 132  
132 enters the safe-init mode if there are errors during the 133  
133 startup. The behaviour of the controller software can 134  
134 be configured in the system update and configuration 135  
135 mode. ItOnce the required configuration has been set,

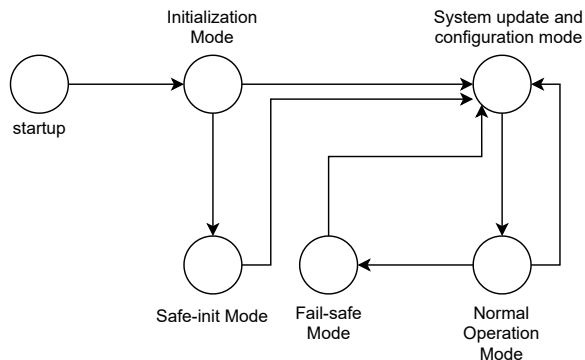


Figure 1: Main Modes in the System

the controller enters the normal operation mode. This is the operational mode of the controller, where the physical movement of the robot arm is enabled. It is in this mode that the controller executes the motion planning algorithms with real-time communication enabled during the normal operation mode for data exchange with external sensors and actuators. It transitions into a fail-safe mode from the normal operation mode if an unexpected error such as an unresponsive sensor, or detection of possible collision with unexpected objects occurs. During normal operation, the user-defined instructions from the robot programming interface provide input to the motion generation components of the software, which in turn generate the path to be followed by the manipulator. Simultaneously, the sensor information and actuator commands are read and written by the communication components based on the user configuration, as well as system configuration.

Timing related properties of a subset of the tasks that make up the robot controller is provided in the Table. 1. The RT communication component is responsible for ensuring real-time communication between the controller and the sensors and actuators. It consists of a network driver task along with a runtime middleware task that provides the necessary interface for data exchange with other tasks. There are two tasks, namely TS\_Ethercat and TS\_RT, that are responsible for real-time communication between the controller and the sensors and actuators. The TS\_Ethercat task comprises the network driver, whereas the TS\_RT task encapsulates the runtime middleware that provides the necessary interface for data exchange with other components. The two tasks are activated by periodic timers of 10 ms period each and their worst-case execution times (WCETs) are 120  $\mu$ s and 80  $\mu$ s respectively. The priorities of TS\_Ethercat and TS\_RT are 12 (highest) and 11 (second highest) respectively. Furthermore, the *utiliza-*

tion of these two tasks are 0.012 and 0.008 respectively. The utilization of a task represents the portion of CPU time required by the task and is calculated by dividing the WCET of the task by its period. The Non RT communication component provides web-based connectivity for communication with enterprise network and for uploading robot programs and managing and updating the controller configurations. It consists of a network driver task, a non real-time middleware task and the web server task, with the web-server task providing the interface for data exchange between the controller and external devices. The TS\_Ethernet, TS\_NRT and TS\_Web tasks are responsible for non real-time communication such as web-based connectivity for communication with enterprise network and for uploading robot programs and managing and updating the controller configurations. These tasks encapsulate the network drivers, non real-time middleware and web server providing an interface for data exchange between the controller and external devices respectively. The robot program interpreter component is responsible for converting the robot program into controller data structures that act as inputs for the trajectory generation component of the controller. It consists of two tasks, the TS\_RPI and the TS\_RPI\_Transform. The robot program interpretation is performed by the TS\_RPI and TS\_RPI\_Transform tasks. These tasks are responsible for converting the robot program into controller data structures that act as inputs for the trajectory generation functionality of the controller. The TS\_RPI task parses the robot program and validates its syntactical correctness. The TS\_RPI\_Transform task then converts the robot program into a data structure that can be used as input for the trajectory generation functionality, which allows planning of the robot motion and generating the required set-points for the controller task (TS\_Control). The trajectory generation functionality is realised with the tasks TS\_IPL\_Path and TS\_IPL\_JointPath. Further, the controller software includes the system state manager tasks, namely TS\_Sys\_Events and TS\_Sys\_Backup, that are responsible for managing different system level signals and generating events that define the behaviour of other tasks. For example, the system state manager task can observe a change in the state of the safety switch signal and generate an event that will trigger a mode change from normal operation mode to a fail-safe mode.

### 3. Related Work

Software migration is usually carried out when adopting a different architectural paradigm than the existing one, such as changing the programming language [12]

System Functions	Task functionality	Task	Task Trigger Type	Task Priority	Task Period (ms)	WCET (us)	Utilization
RT Comm.	Network driver	TS_Ethercat	timer	12	10	120	0.012
RT Comm.	Network middleware	TS_RT	timer	11	10	80	0.008
Non RT Comm.	Network Driver	TS_Ethernet	timer	5	10	75	0.0075
Non RT Comm.	Network Middleware	TS_NRT	timer	4	50	800	0.016
Non RT Comm.	Application	TS_Web	timer	2	100	200	0.002
Robot Program Interpreter	Parse robot program	TS_RPI	event from TS_NRT	3	50	4000	0.08
Robot Program Interpreter	Format data for trajectory generation	TS_RPI.Transform	event from TS_Sys_Events	6	20	200	0.01
System State Manager	Monitor and handle system state events	TS_Sys_Events	periodic	10	10	60	0.006
System State Manager	Create system backup	TS_Sys_Backup	event from TS_WEB	1	100	200	0.002
Trajectory Generation	Interpolate Cartesian Path	TS_IPL_Path	timer	7	20	2000	0.1
Trajectory Generation	Interpolate Joint Space Path	TS_IPL_JointPath	timer	8	20	200	0.02
Controller	Create setpoints and receive feedback for motor drivers	TS_Control	timer	9	2	100	0.05

Table 1: Subset of the tasks in the Robot Controller.

223 or when moving from native server deployments to 251  
224 cloud-based deployments [13, 14]. Sneed [15] proposed 252  
225 a five-step re-engineering planning process for legacy 253  
226 systems, covering *Project Justification*, *Portfolio Analy-* 254  
227 *sis*, *Cost estimation*, *Cost-benefit analysis* and *Contract-* 255  
228 *ing*. The author highlights the need for creating measur- 256  
229 able metrics to justify the effort and the improvements 257  
230 achievable with the migration. Erraguntla et al. [16] 258  
231 discussed a three phase migration method consisting of 259  
232 analysis, synthesis and transformation phases to migrate 260  
233 single-core to multi-core parallel environments. During 261  
234 the analysis and synthesis phase, the design of the ex- 262  
235 isting software is recovered while recommendations for 263  
236 the multi-core environment are made during the trans- 264  
237 formation phase of the migration method. They also 265  
238 provided a reverse engineering toolkit called *RETK* for 266  
239 the analysis and synthesis phases. Battaglia [17] pre- 267  
240 sented the *RENAISSANCE* method for re-engineering a 268  
241 legacy system. The method focuses on planning and 269  
242 management of the evolution process.

243 Menyctas et al. [18] presented a framework called 271  
244 *ARTIST*, a three-phase approach for software mod- 272  
245 ernization focusing on migration towards the cloud. 273  
246 They categorised the migration into three main phases, 274  
247 *Pre-migration*, *Migration and Modernisation* and *Post-* 275  
248 *migration*. During the pre-migration phase, they pro- 276  
249 posed a feasibility study to address the technical and 277  
250 economic points of view. During the migration and 278

modernisation phase, the actual migration is carried out 251  
and finally during the Post-migration phase, the system 252  
is deployed and validated. Forite et al. [19] proposed 253  
the *FASMM* approach to better manage the migration 254  
and to record and reuse the knowledge gained during 255  
the migration in other projects. More recently, Reuss- 256  
ner et al. [2] and Wagner [20] proposed model-driven 257  
approaches to software migration. The focus in these 258  
approaches is to reverse engineer the system using au- 259  
tomated tools and capture the information in modelling 260  
languages and then use the model-driven approach for 261  
further maintenance of the system. 262

Most of the works discussed so far focused on reverse 263  
engineering the existing system to get an understand- 264  
ing of the system, and then to use this information to 265  
model and transform the system based on the technical 266  
requirements. However, an important aspect we found 267  
lacking was emphasis on verification and validation of 268  
the reverse engineering processes. Additionally, while 269  
many of these works focused on architecture transfor- 270  
mation and implementation changes, emphasis on mi- 271  
gration of the testing methods was negligible. During 272  
our discussions in the focus group, testing was identified 273  
as an important domain which required investigation as 274  
multi-core architectures are more prone to concurrency 275  
issues, e.g., livelock, deadlock, race-conditions and data 276  
corruption along with the interference due to the con- 277  
tention for shared resources such as the caches affecting 278




Figure 2: Proposed migration workflow.

279 the timing predictability of the overall software system.

#### 280 4. Migration Methodology

281 Based on the reviewed methods and the extra-  
 282 functional requirements, we create a migration work-  
 283 flow as depicted in Fig. 2 and apply the *Analyze, Verify,*  
 284 *Transform* and *Validate* approach to this workflow. Es-  
 285 sentially, during analysis, the requirements for the mi-  
 286 gration process are established and the existing system  
 287 behaviour is recovered. Then the results of the analy-  
 288 sis are verified by the subject experts. New solutions  
 289 are identified and evaluated during the transformation  
 290 phase. Finally, the applicability of these solutions, along  
 291 with the migration process, is validated during the val-  
 292 idation phase. Additionally, we consider the migration  
 293 process to be iterative in the sense that each stage can  
 294 be revisited and decisions can be roll-backed or modi-  
 295 fied to address issues that may have been missed or if  
 296 they do not meet the objective of the migration. A brief  
 297 overview of the different stages of the proposed work-  
 298 flow is as follows:

- 299 1. During the first stage, we focus on the migration  
 300 of software architecture. In this stage, the goal is  
 301 to synthesize an abstract system model, validate its  
 302 accuracy and transform the model for the multi-  
 303 core environment.
- 304 2. In the second stage, the implementation and veri-  
 305 fication migration, the goal is to analyse the sys-  
 306 tem source code to identify potential concurrency  
 307 issues within the code and transform the code ac-  
 308 cording to the new multi-core architecture model.  
 309 Additionally, the existing verification techniques  
 310 are augmented with methods relevant for a multi-  
 311 core architecture.
- 312 3. In the third stage, we validate the migration process  
 313 by identifying the validation parameters and mea-  
 314 suring these parameters and then comparing them  
 315 with the values obtained before migration.




Figure 3: Various phases in the software architecture migration.

#### 316 5. Software Architecture Migration

317 Many of the real-time systems including the robot  
 318 controller software have a strong focus on timing, safety  
 319 and dependability requirements. Therefore, we need a  
 320 well-defined software architecture to support such re-  
 321 quirements. As there are significant differences in the  
 322 single-core and multi-core platforms, the existing soft-  
 323 ware architecture should be modified to address the con-  
 324 straints of multi-core platforms and make the best use  
 325 of the available resources. To approach this modifica-  
 326 tion systematically, the software architecture migration  
 327 stage is divided into five well-defined phases as shown  
 328 in the Fig. 3. The five phases are :

- 329 1. Architecture requirements specification;
- 330 2. Architecture abstraction and representation;
- 331 3. Architecture recovery;
- 332 4. Architecture transformation; and
- 333 5. Architecture verification.

##### 334 5.1. Architecture Requirements Specification

335 The architecture requirements specification is the  
 336 first phase of the architecture migration process. The  
 337 requirements are essentially high-level and the extra-  
 338 functional requirements of scalability, performance and  
 339 timing guarantees are the guiding principles for the  
 340 complete migration process. The more concrete re-  
 341 quirements are defined during the architecture recov-  
 342 ery phase of the migration process. We also include the  
 343 identification of a requirements specification and man-  
 344 agement process in this phase to better manage the re-  
 345 quirements for the rest the migration process.

## 5.2. Architecture Abstraction and Representation

In this phase, we seek to identify an abstraction level that can accurately represent the system behaviour. An abstraction level close to the implementation may be too detailed, while a higher abstraction level can miss critical information that may be necessary for assuring correct system behaviour. Therefore, to identify the right abstraction, we need to identify the system properties that can be affected when moving to the multi-core architectures. Further, a representation model that can sufficiently capture the system properties should be identified. The representation model should be easy to comprehend, and should act as a communication tool between different stakeholders such as the system architects and developers. To address these issues, we rely on expert interviews and the review of state-of-the-art literature related to multi-core in the real-time systems domain and the model-driven engineering domain to guide the selection of the abstraction level and for the identification of the representation tools.

*Software Architecture, Real-time Task Models and Representation Tools.* The system we considered provides multiple functionalities ranging from embedded control to cloud connectivity. Therefore, we relied on informal and open-ended interviews with the system software architects and domain architects to identify possible abstraction levels. From these discussions, we were able to identify that the task-level abstraction provides the necessary semantics to capture the system properties and therefore, can be used during the later stages of the migration process. Moreover, most of the literature in real-time systems uses the task-level abstraction for the system representation [21, 8].

Many modelling languages support the task level abstraction to represent the architecture of real-time systems. There are several modelling languages that allow modelling of software architectures and task-level abstraction models of real-time systems. The UML MARTE<sup>2</sup> profile [22], Rubus Component Model [23, 24], UPPAAL [25], MechatronicUML<sup>3</sup> [26], AUTOSAR [27], ART-ML Framework [28], are some of the possible modelling languages and frameworks that can be used to represent the system under discussion.

It is worthwhile to mention that although many of these languages, frameworks and supporting tools offer detailed semantics for capturing multiple viewpoints which are essential for managing real-time systems,

the learning curve for many of these tools is however, rather steep, especially when being used for representing task-level abstraction of existing systems. To demonstrate the software architecture abstraction in the proposed methodology, we model the software architecture of the robot controller using the Rubus Component Model as shown in Fig. 4. Note that the Rubus Component Model and its runtime environment consider a one-to-one mapping between a *software component* and a task. A software component is the lowest-level hierarchical element in a component model that is used to model the software architecture of a system. The software component is a design-time entity that may correspond to one or more tasks at runtime. For example, the model of a software component that conforms to the Rubus Component Model (RCM) [23, 24] is shown in Fig. 5. A software component communicates with other components by means of input and output data and trigger ports. The trigger ports indicate when the task (corresponding to the software component) is activated for execution. A software component can be triggered by an independent source (e.g., a periodic clock) or by another software component. The properties of the software component such as their execution times, activation periods and priorities are specified using the values from Table 1. Note that there are two timing constraints, namely Age (50ms) and Reaction (50 ms), that are specified on a chain of software components within the software architecture in Fig. 4. These timing constraints conform to the AUTOSAR standard and are supported by several other modelling languages and methodologies for real-time systems [29].

## 5.3. Architecture Recovery

We need to have a better understanding of existing architecture to be able to modify and adapt it to new platforms. However, in many cases, the documented architecture or the intended architecture does not represent the actual implementation. Such deviations can be attributed to multiple reasons. For example, many of the software systems are developed using a top-down development approach. As a result, implementation level changes are not propagated back to the architectural documents resulting in inconsistencies. Recovering the architecture, therefore, is an essential step for the migration. While many useful architecture visualisation tools such as CodeSonar<sup>4</sup> and Imagix<sup>5</sup> analyse the

<sup>2</sup><https://www.omg.org/omgmarte/>

<sup>3</sup><http://www.mechatronicuml.org/en/index.html>

<sup>4</sup><https://www.grammatech.com/products/code-visualization>

<sup>5</sup><https://www.imagix.com/index.html>

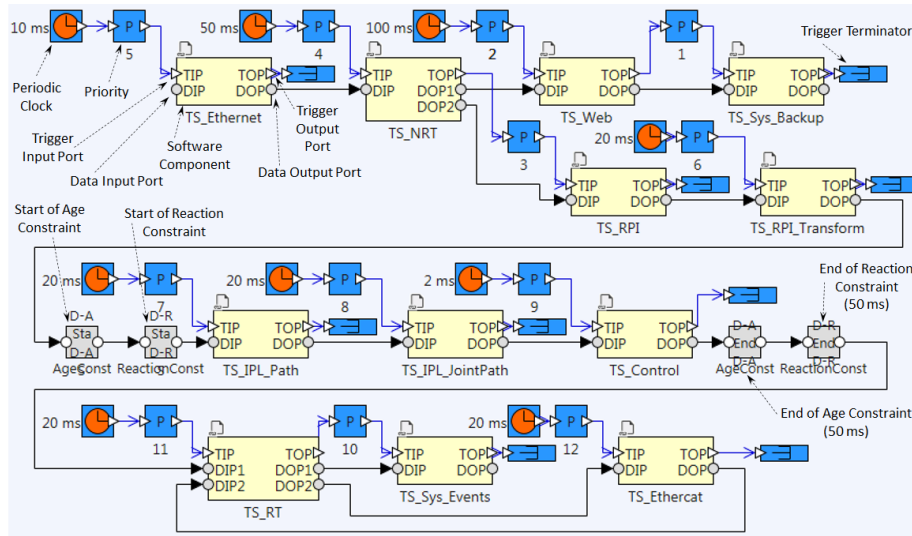


Figure 4: Software Architecture Representation.

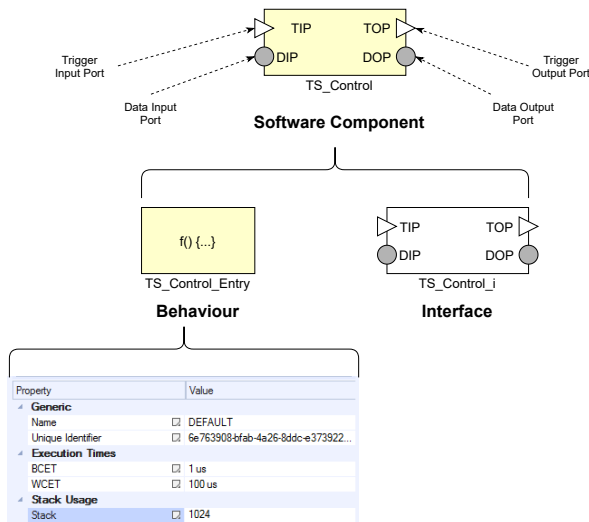


Figure 5: Properties of a Software Component.

439 source code to provide architecture visualisation, they  
 440 only provide information on the logical structure of the  
 441 software and additionally, they may not be able to detect  
 442 faulty architectural patterns within the recovered archi-  
 443 tecture.

444 In this phase, Since the transition to multi-core plat-  
 445 forms in general affects the timing behaviour of the  
 446 system, we focus primarily on extracting the temporal  
 447 properties of the system. , which can manifest them-  
 448 selves in different forms such as deadlines or message  
 449 buffer sizes. For example, a timing requirement can be  
 450 derived based on the communication between TS\_IPL

451 Path and the TS\_IPL\_JointPath. Here, one job of the  
 452 TS\_IPL\_Path generates data for n jobs of the TS\_IPL  
 453 JointPath. The next instance of the TS\_IPL\_Path task  
 454 should complete its execution before the nth job of  
 455 the TS\_IPL\_JointPath is executed. Further, we con-  
 456 sider the system to be modelled with cause-effect task  
 457 chains [30, 31], which implicitly consider maintaining  
 458 the causality in the underlying communication. These  
 459 chains are constrained by the timing constraints similar  
 460 to that of the AUTOSAR standard.

461 At the task-level abstraction, each task can be repre-  
 462 sented in terms of its period, worst-case execution time  
 463 and various types of timing requirements such as dead-  
 464 line, data age, and data reaction constraints [32]. Note  
 465 that the tasks and their corresponding software compo-  
 466 nents at the software architecture abstraction have the  
 467 read-execute-write semantics, which allow them to be  
 468 adapted to comply with the Logical Execution Time  
 469 (LET) model [33]. In addition to these, there can be  
 470 indirect temporal requirements such as the number of  
 471 messages in a message queue should not be less than a  
 472 specific value during a certain operating mode, which  
 473 then requires that the task producing the messages for  
 474 the queue can be blocked only for a duration that does  
 475 not violate this requirement. Therefore, we need a com-  
 476 prehensive multi-dimensional software comprehension  
 477 and reverse engineering approach to extract such infor-  
 478 mation from the existing software architecture, specifi-  
 479 cally, the timing properties and constraints, which are  
 480 crucial in verifying timing predictability of the sys-  
 481 tem [32].

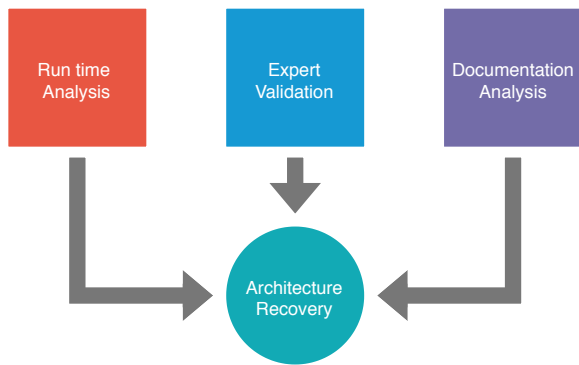


Figure 6: Architecture analysis.

482 To extract the necessary timing requirements, such as  
 483 the periodicity, execution times and deadlines, we re-  
 484 quire analysis of multiple data sources. We identified  
 485 that the architecture documentation, the run-time ex-  
 486 ecution logs and expert validation of the analysis are es-  
 487 sential resources for the architecture recovery phase of  
 488 the migration process, also shown in Fig. 6.

489 **Documentation Analysis.** The architecture of large  
 490 software intensive systems is normally documented ac-  
 491 cording to the “4+1” architectural view model [34] or  
 492 an enhanced variant. The format for architecture doc-  
 493 umentation can vary depending on the internal pro-  
 494 cess and industry-relevant certification requirements.  
 495 SysML [35] and UML models are some of the formal  
 496 description formats for documentation used in the in-  
 497 dustry. Complementing such formal description formats  
 498 are the textual documents explaining the architecture in  
 499 natural language as a part of the documentation. These  
 500 high-level architectural models and documents identify  
 501 the different components of the system and the inter-  
 502 action between components, summarise the design pat-  
 503 terns and technologies employed in the implementation  
 504 and provide a concise overview of the functions of these  
 505 components. By analysing the documentation, it should  
 506 be possible to identify chains of dependent components,  
 507 the tasks associated with these components and the ex-  
 508 pected timing behaviours. The system we considered  
 509 was documented both in UML models, as well as tex-  
 510 tual documents. However, during our analysis, we found  
 511 that existing documentation did not contain any infor-  
 512 mation mapping different tasks to their respective com-  
 513 ponents and there was limited information on expected  
 514 timing behaviours either unavailable or was incomplete  
 515 available in the architecture documents, necessitating  
 516 other analysis approaches such as run-time analysis and  
 517 expert validation.

518 **Run-time Analysis.** While the high-level documents  
 519 are good sources of information, the information pro-  
 520 vided by such documentation may either be incomplete  
 521 or may not reflect the actual implementation. One rea-  
 522 son for such an inconsistency is due to the structure of  
 523 the development process, where the information flow  
 524 is usually top-down, and the changes made at the im-  
 525 plementation level are not propagated back to the ar-  
 526 chitecture documents [36]. Additionally, these indus-  
 527 trial software systems have been incrementally devel-  
 528 oped over many years with the addition of new func-  
 529 tionality, bug fixing, and other optimisations in each  
 530 increment. Therefore, due to the accumulation of un-  
 531 documented changes made during implementation over  
 532 the years, relying solely on high-level documentation as  
 533 the only source of information for modelling the sys-  
 534 tem can result in an inaccurate representation of the ex-  
 535 pected system behaviour. This makes it necessary to  
 536 consider the run-time logs as complementary sources  
 537 of the system information. One approach to under-  
 538 standing the run-time behaviour of the system is the  
 539 tracing and measurement-based approach [37]. Using  
 540 this approach, information such as number of context  
 541 switches, response times, execution times, number of  
 542 task instances, periodicity of the tasks, among others  
 543 can be collected. By using dynamic analysis and visual-  
 544 isation tools such as Tracealyzer [37], additional infor-  
 545 mation such as the communication flow between differ-  
 546 ent tasks, identification of shared resources, task chains  
 547 and precedence constraints between the tasks can be ob-  
 548 tained. The information gained from the run-time anal-  
 549 ysis can be used to refine and enhance the model.

550 The run-time analysis comes with its own set of con-  
 551 undrums. As the system under consideration is config-  
 552 urable, i.e., the user can configure and specify the run-  
 553 time behaviour, it is difficult to identify a configura-  
 554 tion that can be a single representative of possible con-  
 555 figurations for run-time analysis. One possible approach  
 556 to address this issue is to use the “maximum load” ap-  
 557 proach. We consider the system to be in “maximum  
 558 load” state, if under normal operation mode, all sys-  
 559 tem tasks are active and that each task is executing its  
 560 most computationally heavy or memory intensive jobs.  
 561 Relying on a single configuration, however, is not suffi-  
 562 cient to make any statistically reliable conclusions about  
 563 the measurements. Therefore, another argument would  
 564 be to gather run-time behaviour from as many possi-  
 565 ble configurations as feasible. Again, identifying this  
 566 “feasible” number is not straight forward. This is made  
 567 even more complicated by the continuous development  
 568 process, where code is modified and new builds gener-  
 569 ated daily. Identifying a fixed version of the software



570 for analysis becomes non-trivial for such cases. Further, since the controller software operates under different modes, the “maximum load” approach could be pessimistic. Depending on the system under migration, we will need to identify an appropriate configuration and analyse the run time behaviour of each mode independently. For the controller software considered, the “normal operation mode” had the highest resource demand and since all the other modes run only a subset of the “normal operation mode” tasks, we use the maximum load configuration of the “normal operation mode” and ensure that all the required system software components are active during the trace period. Note that we rely on the latest released version of the software.

584 During the run-time analysis of our system, we found that there were inconsistencies between the expected and observed behaviours. A few of the inconsistencies were a result of incorrect configuration of the instrumented code, while others were actual deviations from the expected behaviour. For example, the incorrect configuration resulted in the trace logs showing multiple instances of the jobs of a task as a single job of the same task. This observation highlights the fact that relying on a single source for information is not only ineffective but also error-prone. This necessitates the need for expert validation of the collected information to create a sufficiently accurate system model.

597 **Expert Validation.** Architectural design decisions are made by analysing multiple factors such as domain requirements, dependencies on services provided by the operating systems and the underlying hardware platform, among others. However, the high-level architectural models and documents do not describe the rationale behind the design decisions and even if they do, such information is limited. Moreover, in legacy systems, such documents do not completely reflect the implementation [36]. Furthermore, as the information from the run-time analysis is quantitative and statistical in nature, it is possible to misinterpret any deviation from a commonly occurring pattern as an inconsistency whereas this could have been a design decision. To avoid such misinterpretations and improve system model accuracy, discussions with domain experts are mandatory during the architecture analysis. These discussions will be used to understand the rationale behind the design decisions, and to validate the observations of the documentation and the run-time analysis phases. In our work, we were able to validate the inconsistencies such as the deviation from a commonly occurring pattern as a design decision and also mark some of the observed results as an outcome of incorrect code instru-

621 mentation configuration. For example, due to incorrect configuration of the code instrumentation library, the periodicity of the TS\_RPI observed during run-time analysis phase did not match the values expected by the experts. The functional behaviour however, was accurate, prompting a separate analysis. This analysis identified incorrect configuration of the code instrumentation as the root cause for observed deviation in the periodicity.

#### 5.4. Architecture Transformation

631 As discussed earlier, the architecture transformation phase focuses primarily on evaluating potential solutions and identifying the most appropriate ones for the final implementation. Before we evaluate any solution, we need to identify the system requirements that need to be considered to identify, evaluate and qualitatively rank possible solutions. Since in our case, the migration to multi-core will primarily affect the runtime behaviour, we focus on the explicit temporal requirements, implicit requirements such as the number of messages in a queue and assigned QoS levels to different functional domains. An important requirement here is to ensure that this transformation results in improved system predictability, performance and that the architecture is scalable in terms of the number of cores and new functionality that needs to be integrated into future versions of the software. Since the terms predictability, performance, and scalability are generic in nature, we need to ensure that we have measurable definitions for these terms. For example, we use scalability to refer to the capability of the controller software to control more than one manipulator on the same hardware platform. Once we define the evaluation criteria, we then move towards the evaluation process itself. The evaluation can be carried out in various ways depending on the evaluation metric and the solution being considered, such as simulation, model-checking and analytical calculations. Once the evaluation of possible solutions is complete, we rank these solutions based on an agreed evaluation metric and based on these rankings, we select the solutions for the final implementation phase. To ensure that this transformation is systematic, we divide the transformation phase into the following steps:

- 664 1. identification of potential solutions;
- 665 2. evaluation of the solutions;
- 666 3. ranking of the solutions;
- 667 4. selection of the solutions.

668 **Identification of potential solutions.** Identification of  
669 potential solutions can be done in many different ways.  
670 Although we don't make any specific recommendations,  
671 we would like to point out that the number of po-  
672 tential solutions could be infinitely many and we hy-  
673 pothesize that evaluating each solution will be impossi-  
674 ble. Especially in the case of real-time systems, where  
675 the search space in terms of near-optimal solutions is  
676 large [8, 9, 38, 39]. Therefore, a good starting point in  
677 this stage are the domain experts. Also, the information  
678 from the architecture abstraction and recovery phases  
679 can be a useful guide in reducing the search space. In  
680 our case, we use expert interviews and review the state-  
681 of-art in the real-time systems domain to identify poten-  
682 tial solutions. Another important consideration is that  
683 since application developers are focused primarily on  
684 the application functionality, they rely on the operat-  
685 ing systems to provide support for real-time properties.  
686 This implies that in many cases, only those mechanisms  
687 supported by an operating system can be considered as  
688 part of the potential solution set.

689 As highlighted earlier, the purpose of an abstract sys-  
690 tem model is to capture all the relevant properties of the  
691 system but without the functional complexity. This en-  
692 ables creation of synthetic tasks for simulation and veri-  
693 fication of new design solutions. These abstract task sets  
694 can be modified and verified in short time spans when  
695 compared to modification of the actual implementation  
696 of the system. Many of the real-time workload models  
697 such as those reviewed in [21] have been successfully  
698 used to represent practical systems such as in the avionics  
699 domain as well as in the automotive domain. While  
700 many of these workload models consider the tasks to be  
701 independent, we found that the system under study vi-  
702 olates this assumption and that new jobs of tasks are  
703 triggered by jobs of other tasks. Also, the presence  
704 of event triggered components within the system along  
705 with multi-rate task chains implementing a single func-  
706 tionality, requires that the precedence constraints as well  
707 as task chains be considered when considering potential  
708 solutions [30].

709 Some of the relevant issues that should be addressed  
710 by the potential solutions for transitioning from single  
711 core to multi-core platforms were highlighted by  
712 Macher et al. [40], and Nemati et al. [41]. For exam-  
713 ple, use of single-core hardware implies that the system  
714 tasks execute in sequential manner. If run on multi-core,  
715 the task precedence constraints may not be maintained  
716 affecting system dependability. Additionally, systems  
717 designed for single-core do not require any mapping  
718 of software and multiple compute resources. However,  
719 predictable execution on multi-core is provided by parti-

tioned scheduling approaches [39]. Ad hoc partitioning  
can affect system performance and scalability. Multi-  
level caching can cause data inconsistencies when tasks  
sharing a variable are executing on different cores [42].  
In the case of fixed-priority scheduling, priority assign-  
ment can impact response times [38].

**Evaluation of the solutions.** Once the potential solu-  
tions have been identified, the next step is to evaluate  
these solutions. By evaluation, we refer to the applica-  
tion of the potential solutions from the previous step to  
the abstract model from the architecture recovery stage  
and measurement of the identified metrics. The eval-  
uation can be done in different ways as already high-  
lighted earlier such as simulation in the case of ART-  
ML framework [28] or the Cheddar tool [43], analytical  
calculations if using techniques such as those identified  
in [39], or model-checking if using the timed automata  
approach specified in [44]. For the system described in  
Section 2, one strategy could be to allocate the parts of  
the system that are constrained by the timing constraints  
to one core and rest of the software components to  
other cores (e.g., TS\_IPL\_Path, TS\_IPL\_JointPath, and  
TS\_Control to one core and the rest of the components  
to the other core(s)). Another strategy could be to allo-  
cate the software components to the cores such that the  
specified age and reaction delays are minimized. An-  
other strategy could be based on precedence constraints  
between the software components, which should be on  
the same core (e.g., TS\_Web and TS\_Sys\_Backup have  
an implicit precedence constraint as the latter is trig-  
gered by the former, hence both should be on the same  
core). Similarly, another allocation strategy could be  
based on the criticality levels associated to the software  
components so that non safety-critical software cannot  
interfere with the safety-critical software as proposed  
in [45]. We would like to point out that given the safety-  
critical nature and complexity of the system, we hypoth-  
esise that the potential solution identification and eval-  
uation steps are rather time consuming and are critical  
in the migration process. The time spent during these  
phases can potentially result in practical solutions that  
ensure that the migration process is successful in meet-  
ing the extra-functional requirements.

Moving forward, we return to the question of iden-  
tifying the best solution among the many evaluated so-  
lutions. To guide in this direction, we use the ranking  
approach as follows.

**Ranking of the solutions.** The ranking step of the  
transformation phase orders the evaluated solutions in  
terms of certain criteria. For example, the evaluated so-

770 lution may be required to adhere to safety and security 815  
771 requirements of the domain. Further it may be possi- 816  
772 ble that the extra-functional properties such as portabil- 817  
773 ity between different hardware platforms may be priori- 818  
774 tised over performance improvement on a single hard- 819  
775 ware device. To address such requirements in a system- 820  
776 atic manner, we propose to use the following multi-step 821  
777 approach:

- 778 • identify parameters to rank potential solutions; 822
- 779 • provide measurable definitions to the identified pa- 823  
780 rameters; 824
- 781 • arrive at a consensus on measurement methods for the 825  
782 parameters; 826
- 783 • prioritize or assign weights to the parameters for 827  
784 trade-off analysis; 828
- 785 • rank the evaluated solutions. 829

786 We believe that this approach provides a systematic 830  
787 way to measure effectiveness of the evaluated solutions 831  
788 and guide in selection of the final solution. By identi- 832  
789 fying measurable parameters, the methods to measure 833  
790 them, and prioritize them if a trade-off is necessary, we 834  
791 can remove any ambiguity associated with the perceived 835  
792 effectiveness. To identify these parameters, we propose 836  
793 focus group discussions involving the different domain 837  
794 experts. 838

795 **Selection of the solutions.** Once the potential solutions 843  
796 have been evaluated and ranked, the selection of final 844  
797 solutions should be rather straight forward. However 845  
798 we would like to point out the fact that there could be 846  
799 solutions that may optimize one requirement while neg- 847  
800 atively affecting another requiring a trade-off analysis to 848  
801 select a final solution. 849

#### 802 5.4.1. Architecture Verification

803 The last step in the architecture transformation phase 852  
804 is the verification of the transformed architecture. Here 853  
805 we essentially verify if the transformed architecture 854  
806 complies with requirements from the architecture re- 855  
807 quirements specification phase and the recovery phase. 856  
808 The verification stage is rather simple and straight for- 857  
809 ward since the different steps in the transformation 858  
810 phase involve verification in the evaluation stage with 859  
811 the systematic ranking and selection approach. 860

## 812 6. Implementation Migration

813 So far, we discussed the transformation at the archi- 864  
814 tecture level of the system in our migration process. We 865

now discuss the processes necessary to implement the 815  
transformed architecture at the source code level. Al- 816  
though not directly related to the migration process it- 817  
self, we consider that some form of refactoring at the 818  
source-code level may be necessary prior to the mi- 819  
gration process. Depending on the existing logical ar- 820  
chitecture and the quality of the software, the refactor- 821  
ing may address different concerns. For example, re- 822  
moval of duplicate and dead code, creating components 823  
based on functionality, adoption of a layered architec- 824  
ture among others. For further discussion, we assume 825  
that the system has a layered architecture with well- 826  
defined components, that the logical architecture is ca- 827  
pable of handling new components and modifications in 828  
the abstraction layers, and that the source code is sepa- 829  
rated according to the components. 830

Further, we classify the architecture solutions as ab- 831  
stract component level or functional component level 832  
solutions. For example, if the solution is a new priority 833  
order for the tasks, then it is functional component level 834  
solution if the tasks are associated with the component 835  
and that the priorities can only be changed in the compo- 836  
nent files. If it is a new synchronisation protocol, then it 837  
is an abstract level solution, which is used by all compo- 838  
nents and may need a new implementation. Therefore, 839  
before we make the changes, we identify components 840  
that need to be modified, map solutions that need new 841  
components and then implement the changes. 842

### 843 6.1. Component Identification and Creation

844 The solutions selected during the transformation 845  
846 phase may require that changes be made to the exist- 847  
848 ing components in the system. For example, if the com- 849  
850 ponents use nested semaphores and if the identified so- 851  
852 lution does not support nested semaphores, then such 853  
854 nested semaphores need to be removed. To do this in a 855  
856 systematic manner, we index and categorise the trans- 857  
858 formed solutions, review the solutions with the domain 859  
860 experts and component owners and associate each com- 861  
862 ponent with the solution that requires that component 863  
864 to be modified. For example, the trajectory generation 865  
component may require that its source code be modi-  
fied to accommodate the changes necessary to migrate  
to multi-core platform. We then review the solution with  
the owners of the trajectory generation component. Fur-  
ther, if there are solutions that are classified as abstract-  
level solutions or which could not be mapped to exist-  
ing components, we create new components for such  
changes. For example, if a new real-time middleware,  
that will provide a common inter-task communication  
mechanism is to be implemented, then a new compo-  
nent will be created.

## 6.2. Implementation

Once all components have been identified for modification and new components created, the necessary changes are implemented in the source code. Although the concurrency related issues are addressed during the architecture transformation phase, it is possible that they could manifest during the implementation stage. Therefore, coding guidelines that address these issues are provided to the developers to minimise the manifestation of these issues during the implementation.

## 7. Verification Migration

The system verification and validation stage is the final stage of the migration process. Typically, for the system such as the one being considered, a reliable verification process is already in place. This includes the usual verification approaches such as unit testing, functional testing, and system integration tests. Since the architectural transformation is primarily related to the runtime behaviour and performance, we expect that most, if not all existing tests related to functional behaviour to be valid. Therefore, we hypothesise that any failures here could be related to the concurrent execution of the system tasks. To maintain the quality of the system software, we focus on augmenting the existing tests with concurrency related testing approaches along with performance verification. Again, to approach this enhancement in a systematic way, we divide the verification migration process into concurrency testing and the migration validation phase.

### 7.1. Concurrency Testing

The goal during this phase is to augment the existing verification process to identify concurrency related issues. These include race conditions, atomicity violations and deadlocks. A comprehensive review can be found in the work by Bianchi et al. [46]. We propose the analysis of solutions during the architecture transformation phase to identify scenarios that could lead to potential concurrency issues. This way, it will be possible to create tests for those specific scenarios. Additionally, static code analysis that identifies concurrency bugs is added to enhance the verification process.

### 7.2. Migration Validation

During this phase, we focus on validation of the migration process itself. We begin by identifying the parameters to qualitatively validate the outcome of the process. We use two metrics for this purpose: (i) results of the functional and system integration tests, and (ii)

performance related parameters such as response times. In the first case, no new failures should be introduced after the migration. In the second, the values of the performance parameters should not be less than those measured with the pre-migration version. We point out here that although the validation is the last step, depending on the development process, this validation can be applied to each build prior to release. By using the results of the validation with each build, the pace of the migration process can be measured.

## 8. Tools for Migration

Software migration from single-core to multi-core architectures is a complex process and requires the use of different tools at different stages of the migration process. Here, we review some of the tools that can be used during the different phases of the migration process.

### 8.1. Architecture Representation

Software requirements and the architecture can be described in natural language and as models using different modelling languages such as the UML. For embedded systems with timing requirements, there exist many tools that allow modelling and specification of different views of the system. The APP4MC tool<sup>6</sup>, allows modelling and specification of the hardware as well as software components and provides support for scheduling algorithms. Another tool is the MARTE [47] profile for UML. The MARTE profile extends the UML models to include description of timing requirements. The MAST tool-suite<sup>7</sup> allows for modelling as well as performing automatic schedulability analysis and supports many of the common scheduling algorithms for single-core as well as multi-core architectures. UPPAAL [25] is another tool for modelling the software as timed-automata and it supports model checking for formal analysis and verification. A few concerns with many of these tools are that some have steep learning curves, while others such as UPPAAL are not scalable to large systems and almost all lack support for automatic conversion of existing source code to abstract models.

### 8.2. Architecture Recovery

For architecture recovery, static code visualization tools such as CodeSonar and Imagix could be used. For dynamic analysis, tools which provide visualization of

<sup>6</sup><https://www.eclipse.org/app4mc/>

<sup>7</sup><https://mast.unican.es/>

956 the run-time behaviour along with statistical informa- 1003  
957 tion on timing properties can be effective. For exam- 1004  
958 ple, Tracelyzer allows visualization of the run-time be- 1005  
959 haviour and provides different views to analyse this in- 1006  
960 formation. 1007

## 961 **9. Evaluation** 1008

962 We chose a survey-based approach to evaluate the 1010  
963 proposed methodology. We followed the guidelines 1011  
964 provided by Kitchenham et al. [48] for survey-based re- 1012  
965 search and the discussion of the results. We begin by 1013  
966 describing the design of the survey and then discuss the 1014  
967 results of the survey. 1015

### 968 *9.1. Survey Design* 1017

969 As a first step in the survey-based evaluation, we 1019  
970 identified (i) feasibility, (ii) usability and, (iii) use- 1020  
971 fulness as the evaluation objectives for the migration 1021  
972 methodology. Next, we identified the target population 1022  
973 for the evaluation to be those organisations that develop 1023  
974 complex real-time software systems such as industrial 1024  
975 automation systems and construction vehicles. We iden- 1025  
976 tified a sample from the target population in a non- 1026  
977 probabilistic manner through convenience and judge- 1027  
978 ment based sampling. We created the survey instrument 1028  
979 in the form of online questionnaire that included both 1029  
980 close and open ended questions. The close ended ques- 1030  
981 tions were designed to verify the generalisation of the 1031  
982 observations and the applicability of the different steps 1032  
983 in the methodology. The open ended questions required 1033  
984 the respondents to provide their opinion in a textual for- 1034  
985 mat on feasibility and usefulness of the methodology. 1035  
986 The complete questionnaire was piloted by requesting 1036  
987 colleagues not involved in the study to ensure clarity of 1037  
988 language before it was shared with the respondents. The 1038  
989 questionnaire was made available digitally and included 1039  
990 a brief overview of the purpose of the questionnaire. 1040  
991 The respondents were requested to read about the pre- 1041  
992 sented methodology before they answered the survey. 1042  
993 The received responses were then analysed to evaluate 1043  
994 the methodology. 1044

#### 995 *9.1.1. Evaluation Objectives* 1045

996 As previously mentioned, we identified three key ob- 1045  
997 jectives for the evaluation, namely feasibility, usability 1046  
998 and usefulness of the methodology. For each of these 1047  
999 objectives, we adopt the definitions used by Adesola 1048  
1000 et al. [49] to evaluate their business improvement pro- 1049  
1001 cess methodology. Briefly, we use *feasibility* to imply 1050  
1002 that all the steps in the methodology can be followed in 1051

practice. We use the term *usability* to refer to the ease 1003  
of applicability of the methodology steps and the tools 1004  
mentioned therein. We use *usefulness* to refer to the 1005  
outcome of applying the methodology to relevant sys- 1006  
tems by an organisation. Furthermore, we also included 1007  
the objective of validating the possibility of generalising 1008  
key observations in the methodology. 1009

#### 1010 *9.1.2. Target Population and Sampling Strategy*

To address the evaluation objectives, the target popu- 1011  
lation was identified as organisations developing com- 1012  
plex real-time systems. As for the sample, we iden- 1013  
tified 2 different departments within the same organi- 1014  
sation working on independent and unrelated products 1015  
and also two other organisations. We then identified 9 1016  
expert practitioners from the sample group as the most 1017  
relevant for the evaluation. The participants were cho- 1018  
sen based on their experience in managing and develop- 1019  
ing software(10+ years) for industrial systems and for 1020  
background in multi-core technologies and their knowl- 1021  
edge of the application domains. 1022

#### 1023 *9.1.3. Instrument Design*

The survey was designed in the form of a question- 1024  
naire, combining nominal, close-ended questions, and 1025  
the open-ended questions requiring textual input from 1026  
the respondents. The questionnaire was designed to ad- 1027  
dress two different aspects, (i) problem relevance and 1028  
(ii) methodology evaluation. For the problem relevance, 1029  
we developed six questions to verify if the respondents 1030  
were considering multi-core platforms for their prod- 1031  
ucts. The rest of the questionnaire was focused on 1032  
methodology evaluation. We classified the evaluation 1033  
related questions as either implicit or explicit. The im- 1034  
plicit questions required the respondents to reflect on 1035  
the overall feasibility, usability and usefulness of the 1036  
methodology. The explicit questions were designed to 1037  
validate the generalisation of some of the observations 1038  
made in the methodology. Table 2 shows the mapping 1039  
among the different steps of the methodology, the eval- 1040  
uation type for each of the step and the associated ques- 1041  
tion IDs. Appendix A.3 shows the questionnaire. 1042

### 1043 *9.2. Survey Results and Discussion*

As mentioned previously, the questionnaire was 1044  
shared with nine carefully identified participants from 1045  
the sample population. Of the nine participants invited, 1046  
five respondents participated in the survey. We use the 1047  
labels A,B,C,D and E to refer to each of the respondent 1048  
individually. We discuss the results for the objectives 1049  
of problem relevance, generalisation, overall feasibility, 1050  
overall usability and the overall usefulness. 1051

Table 2: Mapping among the different steps of the methodology, the evaluation type for each of the step and the associated question IDs.

Methodology Stage (Step)	Evaluation Type/ No. Of Questions	Question ID
Architecture Abstraction and Representation (General)	Explicit : 1 Implicit : 6	11 27-32
Architecture Abstraction and Representation (Expert Interviews)	Implicit : 6	27-32
Architecture Abstraction and Representation (State-of-art in Real-time Systems)	Implicit : 6	27-32
Architecture Abstraction and Representation (State-of-art in Model-Driven Engineering)	Implicit : 6	27-32
Architecture recovery (Documentation Analysis)	Explicit : 3 Implicit : 6	13-15 27-32
Architecture recovery (Runtime Analysis)	Explicit : 7 Implicit : 6	12, 16- 21 27-32
Architecture Recovery (Expert Validation)	Implicit : 6	27-32
Architecture Transformation (Identification of Potential Solutions)	Implicit : 6	27-32
Architecture Transformation (Evaluation of the Solutions)	Implicit : 6	27-32
Architecture Transformation (Ranking of the Solutions)	Explicit : 3 Implicit : 6	22- 24 27-32
Architecture Transformation (Selection of the solutions)	Implicit : 6	27-32
Architecture Verification	Implicit : 6	27-32
Implementation Migration (Component Identification and Creation)	Implicit : 6	27-32
Implementation Migration (Implementation)	Implicit : 6	27-32
Verification Migration (Concurrency Testing)	Implicit : 6	27-32
Verification Migration (Migration Validation)	Explicit : 2 Implicit : 6	25-26 27-32
Tools for Migration (Architecture Representation)	Implicit : 6	27-32
Tools for Migration (Architecture Recovery)	Implicit : 6	27-32

1052 **Problem Relevance.** From the problem relevance per- 1076  
1053 spective, 4 of the 5 the respondents, (A,B,C and E) said 1077  
1054 that their applications were not designed for multi-core. 1078  
1055 Respondent *D* said that their applications were designed 1079  
1056 for multi-core but they have been developed from the 1080  
1057 scratch with only limited reuse of existing code. Re- 1081  
1058 spondents *C* and *E* confirmed that they are planning 1082  
1059 to migrate to a multi-core platform while the rest of 1083  
1060 the respondents did not provide any information. Ad- 1084  
1061 ditionally, the same four respondents chose the option 1085  
1062 of redesigning the application while reusing the exist- 1086  
1063 ing code over developing the application from scratch. 1087  
1064 The responses indicate that migration to multi-core plat- 1088  
1065 forms is being considered in the industry and at the same 1089  
1066 time, the respondents prefer reusing the existing code 1090  
1067 over the development of the applications from scratch. 1091

1068 **Generalisation and Feasibility.** Since the methodol- 1093  
1069 ogy was developed based on observations of one sys- 1094  
1070 tem, we created the questionnaire to verify if the ob- 1095  
1071 servations made in different steps can be generalised 1096  
1072 for other complex real-time software systems as well. 1097  
1073 This was done by asking directed nominal questions fo- 1098  
1074 cused on architecture representation, architecture recov- 1099  
1075 ery (runtime analysis and documentation), architecture 1100

transformation (ranking of solutions), and verification migration. For the architecture representation, the results indicate that only parts of the application can be described by timing properties such as worst-case execution times, periods and deadlines.

Similar to the observations about lack of information in the documentation, 4 of the 5 the respondents, (A,B,C and E) said that the application design was not fully documented. Further, only one respondent said that the timing properties were discussed in the design documentation while the rest of the respondents said that the timing properties of only a few critical parts of the application were discussed in the documentation.

The methodology relies on the presence of diagnostic information such as execution times and periodicity for architecture recovery. All the respondents said that their systems provide such diagnostic information. Furthermore, all the respondents mentioned that their applications had multiple configurations and that the runtime behaviour depended on the configuration. None of the respondents said that they tested all possible configurations but only a few. Four out of five respondents (A, B,C and D) said they tested average-case configurations. Furthermore, respondents A and E said that they test the worst-case configurations while respondent D

1101 said that they test the best-case, average-case as well as 1151  
1102 the worst-case configurations. This indicates that iden- 1152  
1103 tifying a representative configuration for architecture is 1153  
1104 not straight forward and can depend on individual ap- 1154  
1105 plication requirements. 1155

1106 Evaluation and ranking of solutions is an important 1156  
1107 step in the methodology. Here we assumed that it will be 1157  
1108 possible to identify and provide measurable metrics for 1158  
1109 ranking possible multi-core solutions. To verify if the 1159  
1110 assumptions are valid, the respondents were explicitly 1160  
1111 asked if they can provide measurable parameters and 1161  
1112 also prioritise them. Four out of five respondents (A, B 1162  
1113 D and E) agreed that they can define as well as prioritise, 1163  
1114 while respondent C answered negatively. 1164

1115 For the verification migration stage of the method- 1165  
1116 ology, a key assumption is that the complex real-time 1166  
1117 systems such as the one discussed in this paper have a 1167  
1118 robust testing mechanism in place for verifying func- 1168  
1119 tional correctness. All the respondents agreed that they 1169  
1120 do have such a mechanism in place. Further, all respon- 1170  
1121 dents agreed that they will reuse the existing tests to ver- 1171  
1122 ify the behaviour of the systems after migration, which 1172  
1123 is consistent with the assumptions made in the proposed 1173  
1124 methodology. 1174

1125 The results of the questionnaire so far indicate that 1174  
1126 much of the observations can be generalised to other 1175  
1127 complex real-time systems. One key observation how- 1176  
1128 ever, is that describing all of the application components 1177  
1129 with timing properties may not be possible. For the 1178  
1130 steps not discussed in generalisation, we address them 1179  
1131 from the overall feasibility perspective discussed next. 1180

1132 **Overall Feasibility.** In order to validate the feasibility 1182  
1133 of the methodology, i.e., to verify if all the steps of 1183  
1134 the methodology can be followed, the respondents were 1184  
1135 asked to answer if they found the methodology feasible 1185  
1136 and to describe the rationale behind their choice. Four 1186  
1137 out of five respondents (A B D and E) considered the 1187  
1138 methodology to be feasible while respondent C consid- 1188  
1139 ered otherwise. When describing the rationale, respon- 1189  
1140 dent C said that they needed more information and the 1190  
1141 correct answer would actually be that they are not sure. 1191  
1142 Respondents B and E did not explain the rationale. 1192  
1143 Respondent A and D agreed that it is possible to repre- 1193  
1144 sent the architecture at a feasible abstraction level and 1194  
1145 that the methodology covered all the critical steps. One 1195  
1146 concern however was that the industrial applications are 1196  
1147 rather big, and therefore we need to address the migra- 1197  
1148 tion in parts and avoid a “big bang” approach. 1198

1149 **Overall Usability.** The survey also included questions 1199  
1150 to evaluate the overall usability of the methodology, i.e., 1200

to verify if the steps in the methodology are workable  
and are easy to apply in practice. Similar to the question  
of feasibility, four out of five respondents (A B D and E)  
answered positively while respondent C said no. When  
describing the rationale, respondent C said that their  
correct answer would actually be that they are not sure.  
Respondent A and B said that the transformation phase  
was uncertain but the steps are general enough to be fol-  
lowed and that the difficulty in following the steps may  
depend on the “architecture, requirements and availabil-  
ity of tools”. Similar response was provided by respon-  
dent D who said that the level of modelling may vary  
depending on the company. Based on the responses it  
can be observed that the steps in the proposed methodol-  
ogy can be followed in general but the overall usability  
is dependent on individual applications.

**Overall Usefulness.** Another objective of the evalua-  
tion is to assess overall usefulness of the methodol-  
ogy for the target population. To address this, the re-  
spondents were asked to evaluate “Usefulness: if the  
methodology can produce results that the organisation  
will find useful?”. Two out of five respondents (A and  
B) consider the methodology to be useful for the indus-  
try, whereas the remaining three respondents consider  
the methodology to be “partially” useful. Respondent B  
justified their choice by highlighting the general appli-  
cability of the steps and respondent A said that having  
such a methodology will create a “common understand-  
ing” between the different stakeholders and the devel-  
opers, thus *increasing the possibility of success and de-  
creasing risks*. Respondent C said the it may not be  
possible to follow the steps completely, but the ideas  
can be “useful”. A similar observation was made by re-  
spondent D who said it will be necessary to consider the  
product to see if the methodology fits the product being  
considered for migration. Although it is not possible  
to draw a straight forward conclusion about the useful-  
ness of the methodology, we can observe from the re-  
sponses that having a methodology can reduce the risks  
of migration projects but the methodology will have to  
be adapted to suit individual application needs in the in-  
dustry.

**Discussion.** The proposed methodology was evaluated  
for feasibility, usability and usefulness by expert prac-  
titioners via a questionnaire. From the feasibility per-  
spective, the analysis of the questionnaire responses in-  
dicate that the methodology covers the critical steps  
necessary for a software migration. From the usability  
perspective, the analysis of the responses shows that the  
different steps can be applied in practice but depending

1201 on the application, the abstraction level and the mod- 1250  
1202 elling requirements will depend on individual applica- 1251  
1203 tions. From the usefulness perspective, the responses 1252  
1204 show that following the methodology steps can decrease 1253  
1205 the risks associated with the migration. From the Gen- 1254  
1206 eralisation perspective, the response show that the ob- 1255  
1207 servations made in the methodology can be extended to 1256  
1208 systems other than the robotic system considered, while 1257  
1209 highlighting the fact that it may not always be possible 1258  
1210 to describe the timing properties for all of the applica- 1259  
1211 tion components. 1260

1212 **Threats to Validity.** Since the evaluation of the method- 1261  
1213 ology has been carried out using a survey, we include a 1262  
1214 discussion on the validity of the results. Kitchenham et 1263  
1215 al. [48] advocates that a survey is reliable if it has been 1264  
1216 administered multiple times and if we get similar results 1265  
1217 each time. In our case, the survey was administered 1266  
1218 only once. This implies that the results may vary if the 1267  
1219 respondents were to answer questionnaire at different 1268  
1220 times. However, much of the questionnaire had nomi- 1269  
1221 nal questions and the number of options provided were 1270  
1222 binary but with an additional option to provide textual 1271  
1223 information thereby limiting the possibility of variabil- 1272  
1224 ity in the responses. Furthermore, although the sample 1273  
1225 group was carefully chosen in a non-probabilistic man- 1274  
1226 ner, it is possible that a different sample of respondents 1275  
1227 may have provided different responses, affecting the va- 1276  
1228 lidity of the conclusions drawn from the survey results. 1277  
1229 While the survey included questions relating to general- 1278  
1230 isation of the observations, not all of the methodology 1279  
1231 steps were explicitly considered but were included un- 1280  
1232 der the general questions of overall feasibility, usability 1281  
1233 and usefulness. Explicit questions may have lead to a 1282  
1234 different conclusion from the one discussed in the pa- 1283  
1235 per. 1284

## 1236 10. Conclusion 1285

1237 Migration of complex embedded software from 1286  
1238 single-core to multi-core computing platforms is non- 1287  
1239 trivial. To ensure a successful migration of these soft- 1288  
1240 ware systems, a systematic approach is needed that 1289  
1241 takes multiple software engineering perspectives into 1290  
1242 account such as software processes, software architec- 1291  
1243 tures, requirements engineering, reverse engineering, 1292  
1244 model-based development, real-time scheduling and 1293  
1245 schedulability analysis. In this paper, we presented a 1294  
1246 systematic multi-stage methodology for migrating real- 1295  
1247 time industrial software systems from single-core to 1296  
1248 multi-core computing platforms. In this regard, we stud- 1297  
1249 ied a complex real-time software system from the au-

1250 tomation industrial domain that requires such a migra-  
1251 tion. We used focus group discussions, expert inter-  
1252 views and reviewed the literature to guide the develop-  
1253 ment of the migration strategy. We identified the soft-  
1254 ware architecture transformation as the main phase in  
1255 the migration process and presented a systematic ap-  
1256 proach to perform the transformation with emphasis on  
1257 the architecture recovery and an evaluation mechanism  
1258 for possible multi-core solutions. We used task-level ab-  
1259 straction of the system to drive the transformation and  
1260 associated timing properties to task-level models and  
1261 proposed their use as input for the evaluation of multi-  
1262 core solutions. To select suitable solutions from the set  
1263 of evaluated approaches we proposed ranking of these  
1264 solutions based on measurable parameters for the final  
1265 implementation and we reviewed some of the tools that  
1266 can be used during the migration process. We evalu-  
1267 ated feasibility, usability and usefulness of the method-  
1268 ology using a survey-based approach. Majority of the  
1269 respondents agreed that the methodology is feasible, us-  
1270 able and useful in general for the industrial applications.  
1271 The evaluation also revealed that the methodology will  
1272 have to be individually adapted to each system under  
1273 migration.

## 1274 Acknowledgements 1285

1275 The research leading to these results has received  
1276 funding from the European Union’s Horizon 2020 re-  
1277 search and innovation programme under the Marie  
1278 Skłodowska-Curie grant agreement No. 764785,  
1279 FORA—Fog Computing for Robotics and Industrial  
1280 Automation, the Swedish Governmental Agency for  
1281 Innovation Systems (VINNOVA) through the projects  
1282 DESTINE and PROVIDENT, and the Swedish Knowl-  
1283 edge Foundation (KKS) through the projects HERO, FI-  
1284 ESTA and DPAC. We would also like to thank our in-  
1285 dustrial partners for providing valuable feedback via the  
1286 survey.

## 1287 References 1288

- 1289 [1] N. Chapin, J. E. Hale, K. M. Kham, J. F. Ramil, W.-G. Tan,  
1290 Types of software evolution and software maintenance, *Journal*  
1291 *of Software Maintenance* 13 (1) (2001) 3–30.
- 1292 [2] Ralf Reussner, Michael Goedicke Wilhelm Hasselbring, Birgit  
1293 Vogel-Heuser, Jan Keim, Lukas Martin (Ed.), *Managed Soft-*  
1294 *ware Evolution*, Springer Nature Switzerland AG, 2019.
- 1295 [3] J. Kraft, Y. Lu, C. Norström, A. Wall, A Metaheuristic Approach  
1296 for Best Effort Timing Analysis Targeting Complex Legacy  
1297 Real-Time Systems, in: 2008 IEEE Real-Time and Embedded  
Technology and Applications Symposium, pp. 258–269.



- [4] G. Mustapić, J. Andersson, C. Norström, A. Wall, A Dependable Open Platform for Industrial Robotics – A Case Study, in: R. de Lemos, C. Gacek, A. Romanovsky (Eds.), *Architecting Dependable Systems II*, Springer Berlin Heidelberg, 2004, pp. 307–329.
- [5] A. Avizienis, J. Laprie, B. Randell, C. Landwehr, Basic concepts and taxonomy of dependable and secure computing, *IEEE Transactions on Dependable and Secure Computing* 1 (1) (Jan 2004).
- [6] G. Mustapic, A. Wall, C. Norstrom, I. Crnkovic, K. Sandstrom, J. Froberg, J. Andersson, Real world influences on software architecture - interviews with industrial system experts, in: *Proceedings. Fourth Working IEEE/IFIP Conference on Software Architecture*, 2004.
- [7] S. Mubeen, E. Lisova, A. V. Feljan, Timing predictability and security in safety-critical industrial cyber-physical systems: A position paper, *Applied Sciences* 10 (9) (2020) 3125.
- [8] C. Maiza, H. Rihani, J. M. Rivas, J. Goossens, S. Altmeyer, R. I. Davis, A survey of timing verification techniques for multi-core real-time systems, *ACM Comput. Surv.* 52 (3) (2019) 56:1–56:38.
- [9] R. I. Davis, L. Cucu-Grosjean, A survey of probabilistic schedulability analysis techniques for real-time systems, *LITES* 6 (1) (2019) 04:1–04:53.
- [10] S. M. Salman, A. V. Papadopoulos, S. Mubeen, T. Nolte, A systematic migration methodology for complex real-time software systems, in: *2020 IEEE 23rd IEEE International Symposium on Real-Time Distributed Computing*, pp. 192–200.
- [11] G. Mustapić, J. Andersson, C. Norström, A. Wall, A dependable open platform for industrial robotics – a case study, in: R. de Lemos, C. Gacek, A. Romanovsky (Eds.), *Architecting Dependable Systems II*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 307–329.
- [12] Leszek Wlodarski, Boris Pereira, Ivan Povazan, Johan Fabry, Vadim Zaytsev, Qualify First! A Large Scale Modernisation Report, in: *SANER, IEEE, 2019*, pp. 569–573.
- [13] P. Church, H. Mueller, C. Ryan, S. V. Gogouvitis, A. Goscinski, Z. Tari, Migration of a SCADA system to IaaS clouds – a case study, *Journal of Cloud Computing* 6 (1) (2017) 256.
- [14] K. Plakidas, D. Schall, U. Zdun, Software Migration and Architecture Evolution with Industrial Platforms: A Multi-case Study, in: C. E. Cuesta, D. Garlan, J. Pérez (Eds.), *Software Architecture, Vol. 11048 of Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2018, pp. 336–343.
- [15] H. M. Sneed, Planning the reengineering of legacy systems, *IEEE Software* 12 (1) (1995) 24–34.
- [16] Ravi Erraguntla, Doris L. Carver, Migration of sequential systems to parallel environments by reverse engineering, *Information & Software Technology* 40 (7) (1998) 369–380.
- [17] M. Battaglia, G. Savoia, J. Favaro, Renaissance: a method to migrate from legacy to immortal software systems, in: *Proceedings of the Second Euromicro Conference on Software Maintenance and Reengineering*, 1998, pp. 197–200.
- [18] A. Menychtas, K. Konstanteli, J. Alonso, L. Orue-Echevarria, J. Gorronogoitia, G. Kousiouris, C. Santzaridou, H. Bruneliere, B. Pellens, P. Stuer, O. Strauss, T. Senkova, T. Varvarigou, Software modernization and cloudification using the ARTIST migration methodology and framework, *Scalable Computing: Practice and Experience* 15 (2) (2014).
- [19] L. Forite, C. Hug, FASMM: Fast and Accessible Software Migration Method, in: *2014 IEEE Eighth International Conference on Research Challenges in Information Science*, IEEE, pp. 1–12.
- [20] C. Wagner, *Model-Driven Software Migration: A Methodology*, Springer Fachmedien Wiesbaden, Wiesbaden, 2014.
- [21] M. Stigge, W. Yi, Graph-based models for real-time workload: a survey, *Real-Time Systems* 51 (5) (2015) 602–636.
- [22] F.Herrera, H. Posadas, P.Peñil, E.Villar, F.Ferrero, R.Valencia, G.Palermo, The COMPLEX methodology for UML/MARTE Modeling and design space exploration of embedded systems, *Journal of Systems Architecture* 60 (1) (2014) 55–78.
- [23] K. Hänninen, J. Mäki-Turja, M. Sjödin, M. Lindberg, J. Lundbäck, K.-L. Lundbäck, The Rubus Component Model for Resource Constrained Real-Time Systems, in: *3rd IEEE International Symposium on Industrial Embedded Systems*, 2011.
- [24] S. Mubeen, H. Lawson, J. Lundbäck, M. Gälnder, K. L. Lundbäck, Provisioning of predictable embedded software in the vehicle industry: The rubus approach, in: *IEEE/ACM 4th International Workshop on Software Engineering Research and Industrial Practice (SER&IP)*, 2017.
- [25] K. G. Larsen, P. Pettersson, W. Yi, Uppaal in a nutshell, *Int. J. Softw. Tools Technol. Transf.* 1 (1-2) (1997) 134–152.
- [26] W. Schäfer, H. Wehrheim, *Model-Driven Development with Mechatronic UML*, Springer Berlin Heidelberg, 2010, pp. 533–554.
- [27] The AUTOSAR Consortium, Autosar technical overview, in: Version 4.3., 2016, <http://autosar.org>.
- [28] A. Wall, *Architectural Modeling and Analysis of Complex RealTime Systems*, PhD thesis, Mälardalen University, Västerås Sweden (2003).
- [29] L. Lo Bello, R. Mariani, S. Mubeen, S. Saponara, Recent advances and trends in on-board embedded and networked automotive systems, *IEEE Transactions on Industrial Informatics* (2019).
- [30] M. Becker, D. Dasari, S. Mubeen, M. Behnam, T. Nolte, End-to-end timing analysis of cause-effect chains in automotive embedded systems, *Journal of Systems Architecture* 80 (2017) 104 – 113.
- [31] M. Becker, S. Mubeen, D. Dasari, M. Behnam, T. Nolte, A generic framework facilitating early analysis of data propagation delays in multi-rate systems (invited paper), in: *2017 IEEE 23rd International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, 2017, pp. 1–11.
- [32] S. Mubeen, T. Nolte, M. Sjödin, J. Lundbäck, K.-L. Lundbäck, Supporting timing analysis of vehicular embedded systems through the refinement of timing constraints, *Software & Systems Modeling* 18 (1) (2019) 39–69.
- [33] T. A. Henzinger, B. Horowitz, C. M. Kirsch, Giotto: a time-triggered language for embedded programming, *Proceedings of the IEEE* 91 (1) (2003) 84–99.
- [34] P. B. Kruchten, The 4+1 View Model of architecture, *IEEE Software* 12 (6) (1995) 42–50.
- [35] E. Andrianarison, J. Piques, SysML for embedded automotive systems: a practical approach., in: *Conference on Embedded Real Time Software and Systems.*, IEEE, 2010.
- [36] U. Eliasson, R. Heldal, P. Pelliccione, J. Lantz, *Architecting in the Automotive Domain: Descriptive vs Prescriptive Architecture*, in: *12th Working IEEE/IFIP Conference on Software Architecture*, 2015.
- [37] J. Kraft, A. Wall, H. Kienle, Trace recording for embedded systems: Lessons learned from five industrial projects, in: *Proceedings of the First International Conference on Runtime Verification (RV 2010)*, Springer-Verlag (Lecture Notes in Computer Science), 2010.
- [38] R. I. Davis, L. Cucu-Grosjean, M. Bertogna, A. Burns, A review of priority assignment in real-time systems, *Journal of Systems Architecture* 65 (2016) 64–82.
- [39] B. B. Brandenburg, M. Gul, Global Scheduling Not Required: Simple, Near-Optimal Multiprocessor Real-Time Scheduling with Semi-Partitioned Reservations, in: *IEEE Real-Time Sys-*

tems Symposium, 2016.

[40] G. F. H. Macher, Andrea Höller, Eric Armengaud, C. J. Kreiner, Automotive Embedded Software: Migration Challenges to Multi-Core Computing Platforms, in: Proceedings INDIN 2015, 2015, pp. 110–118.

[41] F. Nemati, M. Behnam, T. Nolte, Efficiently migrating real-time systems to multi-cores, in: 2009 IEEE Conference on Emerging Technologies & Factory Automation, pp. 1–8.

[42] S. A. Asadollah, H. Hansson, D. Sundmark, S. Eldh, Towards Classification of Concurrency Bugs Based on Observable Properties, in: 2015 IEEE/ACM 1st International Workshop on Complex Faults and Failures in Large Software Systems (COUF-LESS), pp. 41–47.

[43] F. Singhoff, J. Legrand, L. Nana, L. Marcé, Cheddar: a flexible real time scheduling framework, in: SIGAda, ACM, 2004, pp. 1–8.

[44] C. Norstrom, A. Wall, W. Yi, Timed automata as task models for event-driven systems, in: Proceedings Sixth International Conference on Real-Time Computing Systems and Applications, 1999, pp. 182–189.

[45] A. Bucaioni, S. Mubeen, F. Ciccozzi, A. Cicchetti, M. Sjödin, Modelling multi-criticality vehicular software systems: evolution of an industrial component model, International Journal on Software and Systems Modeling 19 (2020) 1283–1302. URL <http://www.es.mdh.se/publications/5787->

[46] F. A. Bianchi, A. Margara, M. Pezzè, A survey of recent trends in testing concurrent software systems, IEEE Transactions on Software Engineering 44 (8) (2018) 747–783.

[47] The UML profile for MARTE: Modeling and analysis of real-time and embedded systems., OMG Group, 2010.

[48] B. A. Kitchenham, S. L. Pfleeger, Personal opinion surveys, in: Guide to Advanced Empirical Software Engineering, Springer London, London, 2008, pp. 63–92.

[49] Adesola Sola, Baines Tim, Developing and evaluating a methodology for business process improvement, Business Process Management Journal 11 (1) (2005) 37–46.

## Appendix A.

Table A.3 summarises the survey questionnaire and shows the mapping between the questions and the different stages of the methodology.

. Alessandro Vittorio Papadopoulos is an Associate Professor of Computer Science at Mälardalen University, Västerås, Sweden. He is part of the SPEC Research Group on cloud computing, with a leading role in the definition of methodologies for benchmarking and performance evaluation in cloud deployments. He was elevated to Senior Member of IEEE in 2019. He received his B.Sc. and M.Sc. (summa cum laude) degrees in Computer Engineering from the Politecnico di Milano, Milan, Italy, and his Ph.D. (Hons.) degree in Information Technology from the Politecnico di Milano, in 2013. From 2014 to 2016, he was a Post-Doctoral Researcher with the Department of Automatic Control, Lund University, Lund, Sweden, and he was also a member of the Lund Center for Control of Complex Engineering Systems, Linnaeus Center, Lund University. He was a Postdoctoral Research Assistant at the Dipartimento di Elettronica, Informazione e Bioingegneria at the Politecnico di Milano (2016). His research interests include robotics, control theory, real-time and embedded systems, and autonomic computing.

. Thomas Nolte: Professor Nolte is leading the Complex Real-time Embedded Systems research group at Mälardalen University, Sweden. He is a senior member of IEEE. He was awarded a B.Eng., an M.Sc., a Licentiate, and a Ph.D. degree in Computer Engineering from the same University in 2001, 2002, 2003, and 2006, respectively. He has been a Visiting Researcher at University of California, Irvine (UCI), Los Angeles, USA, in 2002, and a Visiting Researcher at University of Catania, Italy, in 2005. He has been a Postdoctoral Researcher at University of Catania in 2006, and at Mälardalen University in 2006-2007. He has co-authored over 300 research publications in peer-reviewed conferences, workshops, books and journals.

. Saad Mubeen is an Associate Professor at Mälardalen University, Sweden. He has previously worked in the vehicle industry as a Senior Software Engineer at Arcticus Systems and as a Consultant for Volvo Construction Equipment, Sweden. He is a Senior Member of IEEE and a Co-chair of the Subcommittee on In-vehicle Embedded Systems within the IEEE IES Technical Committee on Factory Automation. His research focus is

1511 on model- and component-based development of pre-  
1512 dictable embedded software, modeling and timing anal-  
1513 ysis of in-vehicle communication, and end-to-end tim-  
1514 ing analysis of distributed embedded systems. Within  
1515 this context, he has co-authored over 135 publications  
1516 in peer-reviewed international journals, conferences and  
1517 workshops. He has received several awards, including  
1518 the IEEE Software Best Paper Award in 2017. He is  
1519 a PC member and referee for several international con-  
1520 ferences and journals respectively. He is a guest editor  
1521 of IEEE Transactions on Industrial Informatics (TII),  
1522 Elsevier's Journal of Systems Architecture and Mi-  
1523 croprocessors and Microsystems, ACM SIGBED Re-  
1524 view, and Springer's Computing journal. He has or-  
1525 ganized and chaired several special sessions and work-  
1526 shops at the international conferences such as IEEE's  
1527 IECON, ICIT and ETFA. For more information see  
1528 *http : //www.es.mdh.se/staff/280 – SaadMubeen.*

Table A.3

No.	Methodology Stage/ Purpose	Question
1	Participant information	Before you proceed, please take the time to read the paper describing the methodology.
2	Participant information	Name of the organization:
3	Participant Relevance	Does your application have real-time components?
4	Participant Relevance	Is your application designed to run on multi-core platforms?
5	Participant Relevance	Have you in the past, migrated your application to a multi-core platform?
6	Participant Relevance	Are you considering migrating the application to a multi-core platform?
7	Exploratory	Did you follow any specific methodology or guidelines to migrate the application to a multi-core platform?
8	Exploratory	Will you recommend the existing approach to others?
9	Exploratory	If you'd like to provide more information about the used methodology, please do so here.
10	Problem relevance	Select the preferred option : a) I prefer to redesign and redevelop the application from scratch for a multi-core platform. b) I prefer to redesign but also reuse the existing code for a multi-core platform.
11	Architecture Abstraction and Representation	Can your application be described with timing properties such as "worst case execution times", "period", "deadlines"?
12	Architecture Recovery (runtime analysis)	Is it possible to identify a particular build of the application that can be used to recover the timing requirements of the application and can such timing requirements be used to create a model of the application?
13	Architecture recovery (documentation)	Is the design of your application documented?
14	Architecture recovery (documentation)	Does the application design documentation contain timing properties?
15	Architecture recovery (documentation)	The behaviour of your application can be:( choose one) a) accurately inferred from the design documentation b) cannot be accurately inferred from the design documentation
16	Architecture Recovery (runtime analysis)	Does your application provide diagnostic logs of runtime behaviour?
17	Architecture Recovery (runtime analysis)	The code instrumentation : a) is fully reliable. b) may not be fully reliable.
18	Architecture Recovery (runtime analysis)	Does your application have multiple configurations?
19	Architecture Recovery (runtime analysis)	Does the runtime behaviour of the application depend on the configuration?
20	Architecture Recovery (runtime analysis)	Do you test all possible configurations of the applications?
21	Architecture Recovery (runtime analysis)	Which configuration do you test
22	Architecture Transformation (Ranking of solutions)	Do you have any existing process/guidelines in place to evaluate and choose between different solutions that may be specific to multi-core platforms?
23	Architecture Transformation (Ranking of solutions)	Is it possible to define measurable parameters that will suit your application's timing requirements to choose one solution over the other?
24	Architecture Transformation (Ranking of solutions)	Is it possible to prioritize the measurable parameters that will suit your application requirements to choose one solution over the other?
25	Verification Migration	Does your application have a verification and validation process in place for checking functional correctness?
26	Verification Migration	Will you reuse the existing tests to verify the behaviour on multi-core platforms?
27	Feasibility	Feasibility: Can the methodology described be followed?
28	Feasibility	Please briefly describe the reason behind your answer here:
29	Usability	Usability: Is the methodology workable? Are the steps and tools easy to use and apply?
30	Usability	Please briefly describe the reason behind your answer here:
31	Usefulness	Usefulness: Is the methodology worth following? Does the methodology produce results that the business will find helpful?
32	Usefulness	Please briefly describe the reason behind your answer here:
33	Overall comments	Which part of the methodology will you like to improve? (you can choose multiple options)
34	Overall comments	Please provide any suggestions and improvements you want to see in the methodology here: