



MAENAD



Grant Agreement 260057

Model-based Analysis & Engineering of Novel Architectures for Dependable Electric Vehicles

Report type	Deliverable D3.2.1
Report name	Analysis and Synthesis Concepts Supporting Engineering Scenarios
Dissemination level	PU
Status	Intermediate (2nd iteration)
Version number	1.1.0
Date of preparation	2012-03-30

Authors**Editor**

Martin Walker

E-mail

Martin.Walker@hull.ac.uk

Authors

Martin Walker

E-mail

Martin.Walker@hull.ac.uk

Sandra Torchiaro

Sandra.Torchiaro@crf.it

DeJiu Chen

chen@md.kth.se

Tahir Naseer

tnqu@md.kth.se

Mark-Oliver Reiser

Mark-Oliver.Reiser@tu-berlin.de

David Parker

D.J.Parker@hull.ac.uk

Henrik Lönn

Henrik.Lonn@volvo.com

David Servat

David.Servat@cea.fr

Reviewer

Andreas Abele

E-mail

Andreas.Abele@continental-corporation.com

The Consortium

Volvo Technology Corporation (S)

Centro Ricerche Fiat (I)

Continental Automotive (D)

Delphi/Mecel (S)

4S Group (I)

MetaCase (Fi)

Pulse-AR (Fr)

Systemite (SE)

CEA LIST (F)

Kungliga Tekniska Högskolan (S)

Technische Universität Berlin (D)

University of Hull (GB)

Revision chart and history log

Version	Date	Reason
0.1	2010-12-01	Initial outline.
0.9	2011-05-09	Final draft
1.0	2011-05-27	Final review
1.0.1	2011-07-1	Optimisation discussion
1.0.2	2011-07-6	Timing analysis update
1.0.3	2011-07-12	Added details on the optimisation architecture elements.
1.0.4	2011-08-31	Period 1 refinements
1.1.0	2012-03-30	Second iteration

List of Abbreviations

AADL	Architecture and Analysis Design Language
ARS	Assigned Requirement for Safety
ASIL	Automotive Safety Integrity Level
ASR	Assigned Safety Requirement
CCF	Common cause failure
COF	Constrained Output Failure
DRS	Driver Reaction System
DSR	Derived Safety Requirement
ETA	Event Tree Analysis
FEV	Fully Electric Vehicle(s)
FMEA	Failure Modes and Effects Analysis
FTA	Fault Tree Analysis
HAZOP	Hazard & Operability Analysis
HiP-HOPS	Hierarchically-Performed Hazard Origin & Propagation Studies
HRC	Heterogeneous Rich Components
MoC	Model of Computation
M(A)	Model (Airfix)
PROMELA	Process Meta-Language
RBD	Reliability Block Diagram
RAS	Replenishment-at-Sea
RSA	Required Safety Assignment
SAR	Search-and-Rescue
SPEEDS	Speculative and Exploratory Design in Systems Engineering
SPIN	Simple PROMELA Interpreter
SRA	Safety Requirement Allocation
SRD	Safety Requirement Derivation
UML	Unified Modeling Language

Table of contents

Authors.....	2
Revision chart and history log.....	3
List of Abbreviations.....	4
Table of contents	5
1 Introduction	7
1.1 Goals.....	7
1.2 Structure	8
2 Analysis and synthesis concepts to support ISO 26262.....	10
2.1 Summary of the ISO 26262 safety design workflow	10
2.2 Hazard analysis in ISO 26262.....	12
2.3 Risk assessment in ISO 26262.....	13
2.3.1 Evaluation Criteria.....	14
2.3.2 ASIL (Automotive Safety Integrity Levels)	15
2.4 ASIL decomposition and allocation in ISO 26262.....	17
2.4.1 ASILs in ISO 26262	17
2.4.2 ASIL Decomposition & Allocation Algorithm	22
2.4.3 HiP-HOPS and Tool Support for ASILs	34
2.5 Discussion & Further Work.....	36
2.5.1 Support for ASILs.....	36
2.6 Optimisation algorithm for ASIL Decomposition	38
2.6.1 The ASIL Optimisation Process.....	38
2.6.2 Encoding.....	38
2.6.3 Evaluation Heuristics	39
2.6.4 Optimisation algorithms	39
2.6.5 Implementation & Performance	40
3 Model-based system analysis techniques	41
3.1 Potential analysis support by EAST-ADL.....	41
3.2 Modelling support for behaviour analysis.....	42
3.2.1 A taxonomy of behaviours and behaviour models in system development.....	42
3.2.2 EAST-ADL Behaviour Annex.....	44
3.2.3 Related Approaches	46
3.3 Modelling support for dependability analysis	50
3.3.1 Dependability and error modelling in EAST-ADL.....	50
3.3.2 Safety Analysis with HiP-HOPS Technology	51
3.3.3 Multi-perspective analysis of EAST-ADL models.....	57
3.3.4 Tool specification and exchange formats.....	60

3.3.5	<i>Tool Integration</i>	61
3.3.6	<i>Semantic Mapping Transformation</i>	63
3.3.7	<i>Representation Transformation</i>	63
3.4	Timing Analysis	65
3.4.1	<i>EAST-ADL/TADL</i>	65
3.4.2	<i>MARTE</i>	67
3.4.3	<i>Summary</i>	69
3.4.4	<i>Timing analysis tooling</i>	69
3.5	Analysis and Synthesis concepts to support Electrical Vehicles	71
3.5.1	<i>Identified Needs</i>	71
3.5.2	<i>Component count and cost analysis</i>	71
3.5.3	<i>Cable Length</i>	75
3.5.4	<i>Power Consumption</i>	75
3.5.5	<i>Current analysis</i>	75
3.5.6	<i>Power Distribution Analysis</i>	76
4	Automatic multi-objective optimisation of system models	77
4.1	Brief description of optimisation definitions and concepts	77
4.2	Comparison of optimisation algorithms	78
4.2.1	<i>Genetic Algorithms (GAs)</i>	78
4.2.2	<i>Tabu Search</i>	80
4.2.3	<i>Ant Colony</i>	81
4.2.4	<i>Simulated Annealing</i>	81
4.3	Tool support for automatic optimisation	82
4.3.1	<i>Variability in HiP-HOPS</i>	82
4.3.2	<i>HiP-HOPS Optimisation Algorithm</i>	83
4.3.3	<i>Achieving full multi-objective optimisation in MAENAD</i>	85
4.3.4	<i>Optimisation Architecture Elements</i>	87
4.3.5	<i>Product Line Optimisation</i>	92
4.3.6	<i>The Plan</i>	93
5	Implementation support for variability management	94
6	Summary and conclusions	96
7	References	98

1 Introduction

1.1 Goals

The primary challenge of the MAENAD project is to develop methods and techniques to support the engineering and design of FEV (Fully Electric Vehicles). To achieve this, it is necessary to define language support in EAST-ADL for the required modelling, analysis, and optimisation concepts, and furthermore to develop algorithms and specify tools that can be implemented and used to analyse and optimise those models. The extensions to EAST-ADL to support the additional modelling concepts are covered in D3.1.1, while the specification of analysis and optimisation algorithms is described in this document, D3.2.1.

In particular, the D.3.x.1 deliverables are intended to support the following project objectives:

- O1-1: Support in EAST-ADL for the safety process of ISO 26262 and representation of safety requirements.

The ISO 26262 standard contains a number of different procedures and design techniques to ensure the safety of an automotive system. To achieve the project goals, EAST-ADL must be extended to support the different modelling approaches present in ISO 26262 and algorithms or tools should be developed to support the analyses required by ISO 26262. The additional analysis concepts and tool support required for ISO 26262 are described in this document.

- O1-2: Automatic allocation of safety requirements (ASILs).

One of the features of ISO 26262 is the concept of ASILs (Automotive Safety Integrity Levels) and the way they can be allocated to different system elements to meet safety requirements. The automatic allocation of safety requirements also has potential applications beyond simply supporting ISO 26262 and could be applied to similar types of qualitative requirements in other domains. While the language concepts necessary to support ASILs and safety requirements are described by D3.1.1, the algorithms for the analysis and allocation of ASILs are specified in this document.

- O2-1: Dependability analysis of EAST-ADL models (with capabilities for multi-mode and temporal analysis of failures, together with integrated assessment of HW-SW design perspectives).

To fully support ISO 26262, it must also be possible to perform various dependability analyses at different stages of the design process, ranging from more simple conceptual analysis to more detailed architectural analysis. Such analyses include hazard analyses, risk analyses, criticality analyses, Failure Modes and Effects Analysis (FMEA), and Fault Tree Analysis (FTA).

- O2-2: Behavioural Simulation of EAST-ADL models.

As well as dependability analysis, another goal of MAENAD is to introduce execution and compositional behaviour semantics into EAST-ADL to allow for behavioural simulation and analysis of EAST-ADL models. The modelling concepts necessary for behavioural simulation are described in D3.1.1 while the analysis and tool support necessary is specified in this document.

- O2-3: Timing Analysis of EAST-ADL models.

The third type of model-based analysis included MAENAD is timing and performance analysis. As with the other forms of analysis, the necessary language concepts are described in D3.1.1 while the algorithm and tool support is specified here in D3.2.1. In the case of timing analysis, this will mean developing ways of linking EAST-ADL models to external timing analysis tools like MAST or RT-DRUID.

- O3-1: Extension of EAST-ADL with semantics to support multi-objective optimisation for product lines.

In addition to analysis of EAST-ADL models of FEVs, another goal of MAENAD is to support model-based optimisation. This involves both extending existing or developing new concepts in EAST-ADL to allow for substitution of one design element for another to create a design space that can be explored by optimisation algorithms such as genetic algorithms. The notion of substitutability and the definition of alternative model elements that can be explored by the optimisation algorithm is achieved via extension of the variability capabilities of EAST-ADL.

The modelling concepts necessary for optimisation, including variability, are described in D3.1.1, while the algorithmic support is specified here in D3.2.1.

- O3-2: Definition of a library of standard architectural patterns that can be automatically applied on an un-optimised EAST-ADL model in order to improve dependability and performance.

Finally, one technique for supporting model-based optimisation is to allow different architectural patterns to be applied to a design model. Such patterns exhibit different dependability and performance characteristics, and the automatic selection and substitution of such patterns can allow different design candidates to be rapidly conceived and evaluated. The types of patterns and the methods with which they may be included in the optimisation process will be included in later versions of this document.

1.2 Structure

The above objectives are developed in MAENAD by several different working groups that each focus on different aspects of the project goals: ISO 26262 support (XG-I), system analyses (XG-A), optimisation (XG-O), and variability (XG-V). Both D3.x.1 deliverables are therefore organised along these lines, and therefore the overall structure of D3.2.1 is as follows:

Section 2 describes the requirement for additional analysis techniques and tools necessary to fully support the ISO 26262 design process, including a brief summary of the ISO 26262 workflow itself. ASIL decomposition is covered in this section.

Section 3 describes the concepts and algorithms behind the various system analyses and tools developed in MAENAD, including subsections for both dependability analysis to support ISO 26262 and each of the other types of analysis of EAST-ADL models, namely behavioural simulation and timing analysis. There is also a section on support for FEV analysis concepts. The specifications of the analysis concepts and algorithms in D3.2.1 is a valuable input to the tool implementation activities of WP5, as well as feeding back into development of EAST-ADL itself.

Section 4 covers the development of concepts and algorithms for the automatic optimisation of EAST-ADL models, particularly with respect to dependability and timing/performance objectives. Later in the project, it will also describe the development of different substitutable architectural patterns that can be used in the optimisation process, although these will be fully specified in an appendix instead.

Section 5 briefly covers the analysis and tool support developed in MAENAD for the modelling of variability. Some of the variability concepts are also inevitably discussed as part of the optimisation section.

Section 6 summarises the different analysis and synthesis concepts described in D3.2.1 and relates them to the objectives set out in this introduction, describing how each is fulfilled (or not) by the algorithms and tools developed in MAENAD.

2 Analysis and synthesis concepts to support ISO 26262

This section describes the requirements of the ISO 26262 standard and how we aim to support those requirements in MAENAD.

2.1 Summary of the ISO 26262 safety design workflow

The ISO/FDIS 26262 requires the application of a “functional safety approach”, starting from the preliminary vehicle development phases and continuing throughout the whole product life-cycle. This approach allows for the design of a safe automotive system. Furthermore, it provides an automotive specific risk-based approach for determining risk classes named ASILs (Automotive Safety Integrity Levels). The new standard uses ASILs for specifying the item's necessary safety requirements for achieving an acceptable residual risk, and provides requirements for validation and confirmation measures to ensure a sufficient and acceptable level of safety being achieved [1].

The main phases of ISO 26262 are:

- Concept phase (ISO 26262 - Part 3):
 - Item definition: during this phase the Item has to be identified and described. To have a satisfactory understanding of the item, it is necessary to know about its functionality, interfaces, and any relevant environmental conditions. Moreover the boundary of the item has to be defined.
 - Initiation of safety lifecycle: the safety lifecycle is initiated based on Item definition and by determining the development category (e.g. current production, new, modified, etc). In the case of modification, an impact analysis shall be performed.
 - Hazard analysis and risk assessment: to evaluate the risk associated with the item under safety analysis, a risk assessment is carried out. A risk assessment considers the functionality of the item and a relevant set of scenarios (operating conditions & environmental conditions). For each identified hazardous event the ASIL level has to be determined. Moreover, after the assessment of the risk, the safety goal (top level safety requirements) shall be defined.
 - Functional Safety Concept: based on the safety goals, the functional safety requirements have to be specified. The functional safety concept shall include:
 - Functional safety requirements → to define in a complete way the functional safety requirements, the following attributes have to be defined for each hazardous event: safety goals, operating modes, fault tolerant time interval, possible safe state, transitions to and from the safe state, emergency operation interval, functional redundancies, driver warning, degraded operation and driver's actions;
 - Item functional description and requirement allocation;
 - Interaction description with vehicle systems;
 - Functional specifications to achieve the safety goals;
 - Description of external measures to avoid or mitigate the effects of hazards.
- System level development – specification, (ISO 26262 - Part 4)
 - Initiation of product development at system level: the functional safety activities during the individual sub-phases of system development have to be determined and

planned. The safety plan shall include the Item integration and testing plan, the validation plan and the functional safety assessment plan.

- Specification of technical safety requirements: the technical safety requirements shall be specified by refining the functional safety concept, considering both the functional concept and the preliminary architectural assumptions. Moreover the consistency and traceability between functional safety requirements and technical safety requirements shall be verified.
- System design: the system design and the technical safety concept, that comply with the functional requirements and the technical safety requirements specification of the item, shall be defined. Methods and measures to avoid systematic failures have to be applied during system design. Furthermore, methods to avoid random hardware failures during system operation shall be applied in this phase.
- Comply with the technical safety requirements specification: it is necessary to verify that the system design and the technical safety concept comply with the technical safety requirements specification.
- Hardware and Software level development, (ISO 26262 - Part 5 and Part 6)
 - Hardware and Software development - safety technical design: based on the system design specification, the item is developed from the hardware and software level perspective.
- System level development – integration and validation (ISO 26262 - Part 4)
 - Item integration: the compliance with each safety requirement in accordance with its specification and ASIL classification shall be tested. Moreover it is necessary, during this phase, to verify that the "System design" covering the safety requirements are correctly implemented by the entire item. The item integration starts with the integration of the hardware and software of each element that the item comprises. After the integration of the elements into a full system, the final step is the integration of the item with other systems within a vehicle and with the vehicle itself.
 - Safety validation: the evidence of compliance with the safety goals and evidence that the functional safety concepts are appropriate for the functional safety of the item shall be provided. Moreover the evidence that the safety goals are correct, complete and fully achieved at the vehicle level shall be provided.
 - Functional safety assessment: it is necessary to assess that the functional safety is achieved by the item.
 - Release for production: the release for production criteria at the completion of the item development shall be specified. The release for production confirms that the item complies with the requirements for functional safety at the vehicle level.
- Production and Operation Level (ISO 26262 - Part 7)
 - Production: a production process for safety-related elements or items that are intended to be installed in road vehicles has to be developed and maintained. Functional safety during the production process by the relevant manufacturer or the person or organisation responsible for the process shall be achieved.
 - Operation, service and decommissioning: in order to maintain the functional safety after the item's release for production, the customer information, maintenance and repair instructions, and disassembly instructions regarding the item, system or element shall be specified.

- Management of functional safety (ISO 26262 - Part 2): in this phase the requirements for management of functional safety are specified. These requirements cover all the phases of the safety lifecycle, and are specified for:
 - Overall safety lifecycle management;
 - Safety management during the concept phase and the product development;
 - Safety management after the item’s release for production.

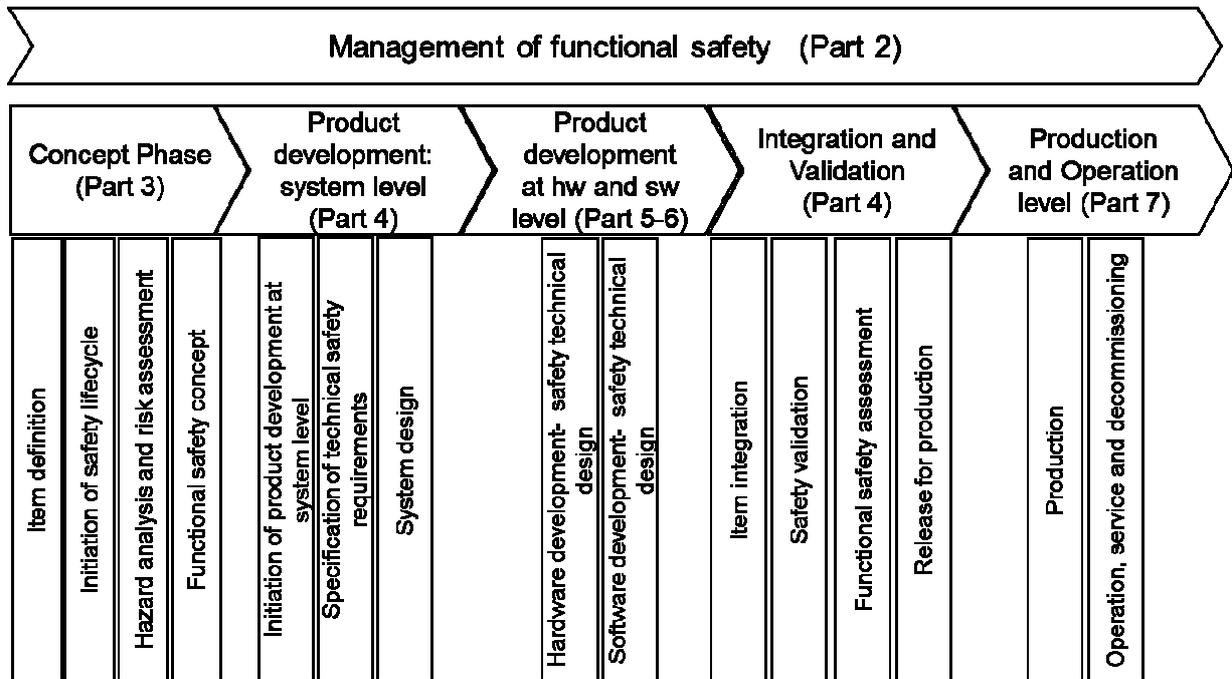


Figure 1 - ISO 26262 Safety lifecycle

2.2 Hazard analysis in ISO 26262

To identify the hazards, it is essential to define the malfunctions and eventual maintenance condition(s) related to the item.

All possible functional anomalies derivable from each foreseeable source shall be defined.

In general, for each *target function* (the function description in terms of output(s) behaviour) it is always possible to define, in an exhaustive manner, the malfunctions in terms of anomalies of function activation:

- Unwanted activation
- Missed activation
- Degraded activation (vs. value or timing)
- Incoherent activation (in presence of more than one function’s output)
- Unwanted deactivation
- Missed deactivation
- Degraded deactivation (vs. value or timing)
- Incoherent deactivation (in presence of more than one function’s output)

2.3 Risk assessment in ISO 26262

In order to perform a correct and complete risk assessment, a proper set of realistic scenarios (all variables and/or states that characterise the functions or affect them) must be defined.

A complete description of a scenario includes:

Operative conditions → selection of the factors related to vehicle usage scenarios to be considered for situation analysis and hazard identification. Examples in the vehicle dynamic control field are:

- *Dynamic driving state*
 - ✓ Speed
 - ✓ Longitudinal acceleration
 - ✓ Lateral acceleration
 - ✓ ...
 - *Manoeuvres*
 - ✓ Curve
 - ✓ Overtaking
 - ✓ Braking
 - ✓ Lane change
 - ✓ ...
 - *Driver condition*
 - ✓ Attention level
 - ✓ Hands on steering
 - ✓ ...
- Environmental conditions → selection of the factors related to environmental variables to be considered for situation analysis and hazard identification. Typical environmental conditions include:
- *Location*
 - ✓ Urban road
 - ✓ Suburban road
 - ✓ Motorway or highway (with lateral barriers)
 - ✓ ...
 - *Street condition*
 - ✓ Adherence
 - ✓ Road surface type (smooth, rough, etc)
 - ✓ Slope
 - ✓ ...
 - *Traffic situation*
 - ✓ Presence of other vehicle(s) (lateral, oncoming, following, etc)
 - ✓ Presence of cycles, pedestrians, etc.

The hazard under safety analysis, when applied to the operational situations (operative & environmental conditions), triggers the so called "hazardous event".

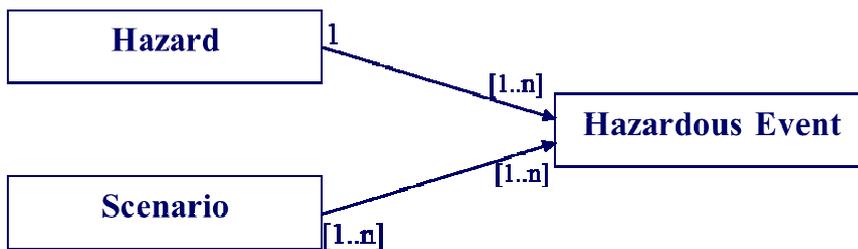


Figure 2 - Hazardous Events

For each identified hazardous event, the severity, controllability and exposure values should be ranked to determine the associated Automotive Safety Integrity Level (ASIL), representing the level of risk.

Parameters for risk determination according to ISO 26262:

- *Controllability [C]:*
Avoidance of a specified harm or damage through timely reactions of the persons involved.
- *Severity [S]:*
Measure of the expected degree of harm to an endangered individual in a specific situation.
- *Exposure time [E]:*
State of being in an operational situation that may be hazardous if coincident with the malfunction under consideration

2.3.1 Evaluation Criteria

Controllability

The controllability shall be assigned to one of the controllability classes C0, C1, C2 and C3 in accordance with the table below:

Class	Description
C0	Controllable in general
C1	Simply controllable
C2	Normally controllable
C3	Difficult to control or uncontrollable

Table 1 - Controllability

It is important to note that the controllability levels assigned to the various situations should be assessed through specific testing on the road, fault injection, etc.

Severity

The severity shall be assigned to one of the severity classes S0, S1, S2 or S3 in accordance with the table below.

Class	Description
S0	No injuries
S1	Light and moderate injuries
S2	Severe and life-threatening injuries (survival probable)
S3	Life-threatening injuries (survival uncertain), fatal injuries

Table 2 - Severity

Exposure time

The probability of exposure shall be assigned to one of the probability classes E1, E2, E3 and E4 in accordance with the table below.

Class	Description
E1	Very low probability
E2	Low probability
E3	Medium probability
E4	High probability

Table 3 - Exposure

2.3.2 ASIL (Automotive Safety Integrity Levels)

Safety Integrity Levels (SILs) are abstract classification levels that can be used to indicate the level of safety required of safety-critical systems (or elements thereof). SILs have been adopted as part of safety standards such as IEC 61508 and - in the automotive domain - ISO/FDIS 26262. In the context of the upcoming ISO 26262, SILs are known as ASILs - Automotive Safety Integrity Levels - and form a major part of the standard: ASILs are used to specify the necessary safety requirements for achieving an acceptable residual risk, as well as providing requirements for validation and confirmation to ensure the required levels of safety are being achieved.

Safety requirements in these standards are intended to ensure the system being designed is free from unacceptable risk (assuming the requirements are met) and are derived through a process of analysis and risk assessment. The aim of the process is to determine the critical system functions - those which have the potential to be hazardous in the instance of failure - and the requirements necessary to mitigate the effects or reduce the likelihood of those hazards. These safety requirements are often associated with integrity requirements that apply to those critical functions to indicate, in essence, what level of contribution they have towards the overall system safety and thus what level of safety they should implement to avoid system failures. A low ASIL therefore indicates that the element is not a major contributor to severe system failures, while a high ASIL indicates that it is potentially a major contributor, and this allows a means of verifying that system safety requirements are being achieved by ensuring that the ASILs allocated to system elements are also being met.

Therefore ASILs play a dual role in the development of safety-critical systems: they allow for top-down allocation of safety requirements to different elements of the system according to their

contribution to risk, and they allow for bottom-up verification to show that the safety requirements are being met by the developed system.

ASILs are divided into one of four classes classes to specify the item's necessary safety requirement for achieving an acceptable residual risk, with D representing the highest and A the lowest class. The ASIL-Level shall be determined for each hazardous event using the estimation parameters severity (S), probability of exposure (E) and controllability (C) in accordance with the table below. QM (Quality Management) can be applied to non-safety critical elements to indicate that there are no specific safety requirements in place.

		C1	C2	C3
S1	E1	QM	QM	QM
	E2	QM	QM	QM
	E3	QM	QM	A
	E4	QM	A	B
S2	E1	QM	QM	QM
	E2	QM	QM	A
	E3	QM	A	B
	E4	A	B	C
S3	E1	QM	QM	A
	E2	QM	A	B
	E3	A	B	C
	E4	B	C	D

Table 4 - Determining ASILs

ASIL = QM (Quality Management) → the function has no impact on safety - it is not necessary to define any safety requirement

ASIL = A → the function has a minimal impact on safety

ASIL = B → the function has an impact on safety – considerable damage

ASIL = C → the function has impact on safety – relatively high damage associated with medium-high probability of being in a situation of risk

ASIL = D → the function is critical - very significant damage associated with a high probability of being in a situation of risk

Top-level safety requirements

During the concept phase a *safety goal* shall be defined for each hazardous event. This is a fundamental task, since the safety goal is the top level safety requirement, and it will be the base on which the functional and technical safety requirements are defined. The safety goal leads to item characteristics needed to avert the hazard or to reduce risk associated with the hazard to an acceptable level. Each safety goal is assigned an ASIL value to indicate the required integrity level according to which the goal shall be fulfilled. For every safety goal a *Safe state*, if applicable, shall be identified in order to declare a system state to be maintained or to be reached when the failure is detected, so to allow a failure mitigation action without any violation of the associated safety goal. For each safety goal and safe state (if applicable) that are the results of the risk assessment, at least one safety requirement shall be specified.

2.4 ASIL decomposition and allocation in ISO 26262

One of the benefits of ASILs is that they can be decomposed across different elements and subsystems of the system architecture so that the highest constituent ASIL need not apply equally to the entire system; instead, those parts of the system that are most critical and that contribute most to potential failures are allocated higher ASILs and are thus subject to the strictest safety requirements, while less important parts that have little or no bearing on potential failures are allocated lower ASILs and are subject to lower safety requirements. This is important as high levels of safety are much more expensive to implement than lower levels of safety, and if the required level of safety can be achieved for a lower cost, then it is normally desirable to do so.

However, although ISO 26262 defines guidelines and procedures to be followed during ASIL decomposition and allocation, the process is generally a manual and unaided one. This can be problematic in situations where multiple safety-critical functions are delivered by complex, networked architectures. As part of the MAENAD project, we aim to provide both language and tool support for a more systematic approach to ASILs, including the potential for automatic decomposition of ASILs. Specifically, objective O1-1 states that EAST-ADL should be extended to support the safety process of ISO 26262 and provide representation of safety requirements, while objective O1-2 explicitly states a goal of developing support for the automatic allocation of ASILs. The intention is to ensure that the use of ASILs as described in ISO 26262 can be carried out more easily and efficiently even in larger and more complex systems.

2.4.1 ASILs in ISO 26262

ISO 26262 sets out a comprehensive methodological framework for the design and development of electronic engineering (EE) systems in the automotive domain. ASILs - and the associated acts of determining and upholding safety requirements - form a major part of this process as they are the mechanism by which the required levels of safety are represented and maintained throughout the design of the system. As such, ISO 26262 contains a considerable amount of information on ASILs and how they should be used. For detailed information, the reader should refer to ISO 26262 itself, but a general overview of the more relevant parts will be given here.

2.4.1.1 Hazard Analysis and definition of ASILs

ASILs are initially derived during the process of defining the safety requirements, i.e., what must be done to ensure a reasonable freedom from unacceptable risk. They are obtained through a combination of functional hazard analysis and risk assessment techniques, starting at an abstract, high-level concept stage (corresponding to EAST-ADL's Vehicle level). These are then developed further and verified at subsequent stages of the design, as represented in increasingly lower, more concrete design views (corresponding to EAST-ADL's Analysis and Design levels).



Figure 3 - Overview of the ISO 26262 safety workflow

The first step is the identification of hazards and the definition of hazardous events. This initially consists of assessing the 'target feature' of the item under safety analysis and determining its possible malfunctions, i.e., possible anomalies in its output (also known as *feature flaws*). All possible functional anomalies - whether internal, external (e.g. due to misuse), or as a result of maintenance (or the lack thereof) - are considered. These malfunctions form the basis of the hazards - the potential sources of harm that the item can cause. In addition, the various possible operating and environmental scenarios are also considered in which these hazards may occur; taken together, they form the definition of the *hazardous events*, i.e., the occurrence of a given hazard in a particular scenario.

Hazardous Events are then classified in terms of its risk by assigning it an ASIL. This is arrived at by considering the hazardous event from three perspectives: controllability, severity, and exposure, each of which is graded into four categories. The level of risk is determined as a function of the frequency and severity of the hazardous event, and the frequency in turn is determined by the controllability and exposure, so the ASIL can be derived from these factors, as shown in Table 4.

The ASIL itself is captured as part of the *safety goal* corresponding to the hazardous event. For all elements not classified as QM (and it is important to note that QM can only be assigned to elements that are not considered to be safety relevant), a safety goal should be defined. These can result in top-level (i.e., system-level) safety requirements that describe the characteristics the item needs to avert or otherwise mitigate the effects of the hazard in order to achieve an acceptable level of safety. In addition, for every safety goal, a safe state should be defined,

indicating what state the system should maintain or transition to when the hazardous event is detected in order to mitigate that event and achieve the safety goal.

It should also be noted that for a Safety Element out of Context (SEooC), much of this hazard analysis is not possible (due to the lack of contextual information), and so an assumed ASIL is typically assigned instead.

2.4.1.2 ASIL Decomposition

Once defined, the ASILs of the top-level safety goals are propagated throughout the development process as the safety requirements are refined. Each derived safety requirement inherits the ASIL from the safety goal or requirement it was derived from, and so the functional and technical safety requirements that are ultimately allocated to architectural elements all possess a corresponding ASIL.

It is also possible for ASILs to be defined independently of the safety goals when looking at Safety Elements out of Context (SEooC). A SEooC can be a system, subsystem, software or hardware component, or any other part of a system that is safety-critical and designed to be generic and used within more than one system. In this case it is possible to analyse the SEooC with regard to safety, without knowing the context in which the element may operate, by assuming ASIL values and using those values in later decomposition.

These ASILs can be tailored to fit the architecture of the design through the process of *ASIL decomposition*, in which different ASILs are allocated to different elements while ensuring that the overall ASILs as defined by the system-level safety requirements are still being achieved. This process takes advantage of architectural decisions - such as the inclusion of independent or redundant architectural elements - to implement the safety requirements across these elements and thus (potentially) assigning a lower ASIL to the decomposed safety requirements that apply to each individual element without affecting the ASIL that applies to those elements taken together. It is important to note that this can only be achieved where there is a sufficient degree of independence between elements and no critical dependencies (and thus dependent failures) to influence them; this can exclude instances of homogenous redundancy (situations where the same element is duplicated to achieve some measure of redundancy) where elements may suffer from common flaws even if no explicit dependencies exist between elements. ISO 26262 suggests that when deciding the matter of independence, the following should be considered:

- Similar or dissimilar redundant elements;
- Different functions implemented with identical software or hardware elements;
- Functions and their respective safety mechanisms;
- Partitions of functions or software elements;
- Physical distance between hardware elements, with or without a barrier;
- Common external resources.

ISO 26262 also distinguishes between independence (which requires freedom from both cascading/propagating failures as well as common cause failures, and which affects the viability of decomposition) and interference (which requires only freedom from cascading failures and which affects the allocation of QM to sub-elements).

Figure 4 shows how the ASILs can be decomposed such that two (or more) lower ASILs can be equivalent to a higher ASIL. Note that the original ASIL is retained in brackets afterwards. Decomposition can also be applied to elements that are not safety-critical, assuming there is no link between those elements and safety-relevant elements such that failures can cascade from one to the other; in these cases, the non-critical elements can be allocated QM while the safety-related element gains a true ASIL. Decomposition can be applied during multiple stages of the design process, including the concept phase (corresponding to EAST-ADL analysis level and

FAA), system phase (corresponding to the EAST-ADL design level, with FDA for functions and HDA for platform), and software and hardware phases (corresponding to Implementation level and AUTOSAR).

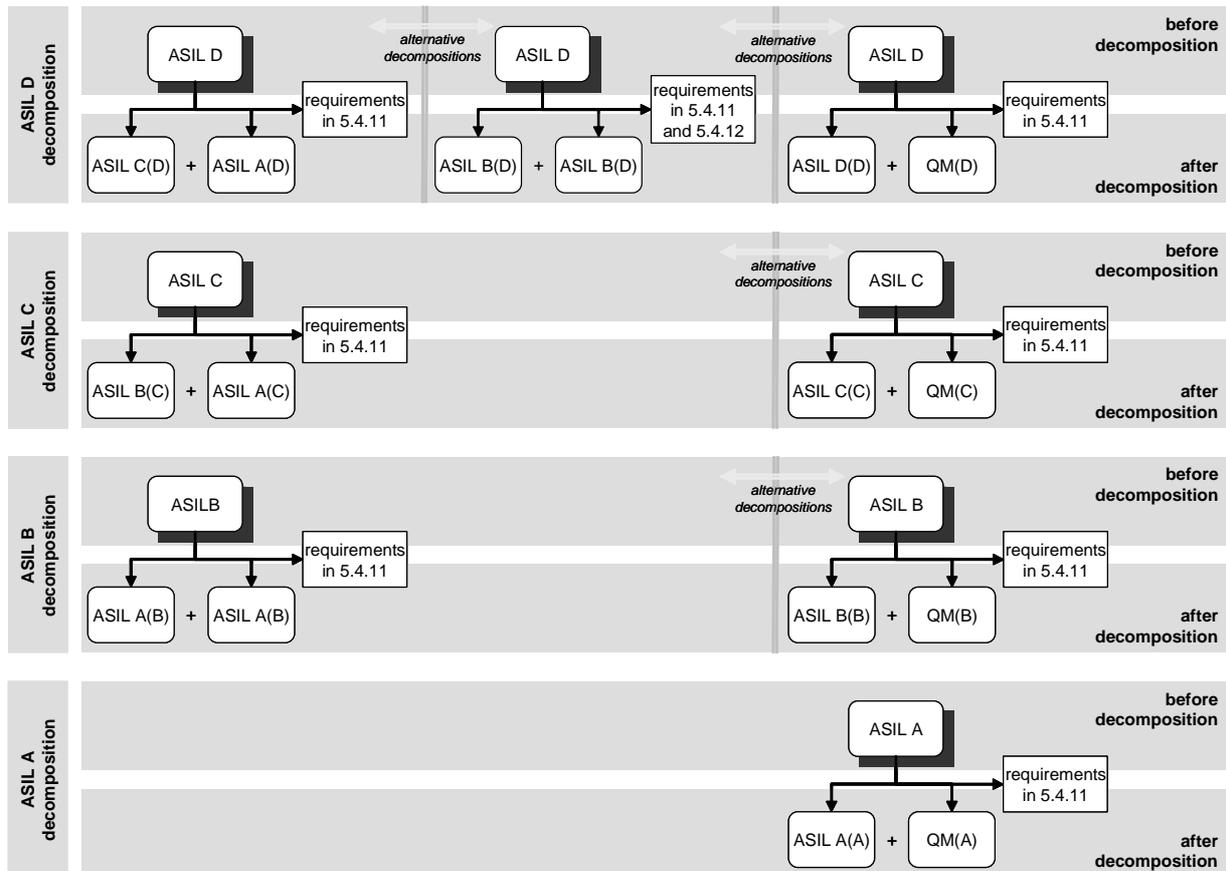


Figure 4 - Decomposition Algebra

Figure 5 shows how ASILs can be decomposed across independent channels in an architecture. The modelled system as a whole has ASIL C, indicating there is a relatively severe hazard that must be avoided. Parts of the system that contribute directly to the system failure (such as the Dedicated CAN and the outputs) receive the ASIL directly too. When there is a disjunction, e.g. two or more components can all cause the system failure individually, they all receive the full ASIL. Only when components must fail together in conjunction to cause the system failure is the ASIL diluted. This can be seen in the simple case for Macro-block3 in the diagram, which collectively receives an ASIL C; the subcomponents within (ultimately SWS1-4) may not have the full ASIL C, however, and there are a number of different possible combinations of ASILs that would fulfil the ASIL C for the subsystem as a whole (e.g. A + A + A + A).

A more complex case arises with Macro-block1 and Macro-block2. Both must fail in conjunction to cause system failure, thus both together receive a share of the ASIL C. Because of the AND failure logic (caused by multiplexers) in the "VCS" component, the received ASIL C can be decomposed over Macro-block 1 and Macro-block 2 such that Macro-block 1 receives ASIL A and Macro-block 2 receives ASIL B (which together meet ASIL C), since both are required to cause the VCS to fail. By contrast, Macro-block 3 inherits ASIL C directly since it is single-handedly capable of causing the VCS to fail. It is also worth noting that the ASILs of each of these three Macro-blocks could be decomposed further within those blocks too. For example, in Macro-block1, a failure of either output is all that is necessary, whereas in Macro-block2, a failure of both outputs is needed, so in Macro-block1, all components would require ASIL A (i.e., the ASIL of Macro-block1) since any individual failure is sufficient, but in Macro-block2, both originating contributors (NO and

NC in the EPB Button subsystem) could receive ASIL A (rather than ASIL B, the ASIL of Macro-block2) as both need to fail for Macro-block2 to fail.

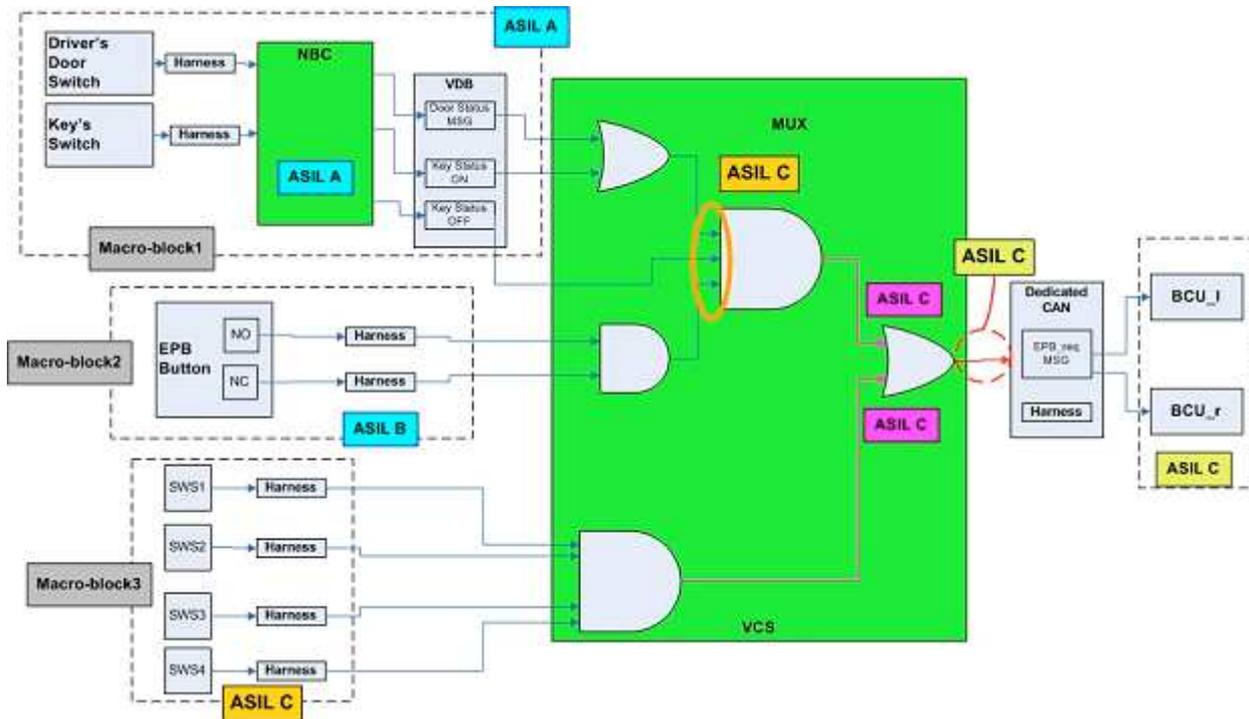


Figure 5 - ASIL Decomposition example

2.4.1.3 Safety Analysis in support of ASILs

An assessment of the failure logic of the system and thus a verification of whether elements of the system are indeed independent (without any dependent or common cause failures) can be carried out by a variety of safety analysis techniques. ISO 26262 suggests that qualitative methods should be applied during the abstract system phase and software phase, while quantitative methods can be applied during the hardware phase to allow for the quantification of random hardware failures (although it also suggests common cause failures should still be estimated on a qualitative basis due to the difficulty in accurately quantifying them). Any analysis should take into account:

- Random hardware failures (which may be quantified)
- Development/design faults (e.g. due to errors in design or requirements)
- Installation faults
- Repair/maintenance faults
- Environmental factors (e.g. temperature, pressure, corrosion, electromagnetism)
- Failures of common external resources (e.g. power, communication buses)
- Stress due to specific situations (e.g. wear & tear)

The objective of the safety analyses in these cases is to examine the consequences of the failures on the elements of the system, providing information on whether the safety goals & requirements are being met and also determining whether or not ASIL decomposition will be possible. Analysis may also discover new hazards. ISO 26262 suggests the following analysis techniques:

- Qualitative & quantitative FMEA
- Qualitative & quantitative FTA
- Qualitative & quantitative ETA
- HAZOP (qualitative)
- Markov models (quantitative)
- RBDs (quantitative)

FTA in particular is well suited to performing analysis for the purposes of ASIL decomposition since it can take into account the effects of multiple failures, which e.g. FMEA cannot. Thus it can determine whether or not two or more input failures must occur to cause an output failure (which may suggest the possibility for ASIL decomposition in that case).

In the case of random hardware failures, it is possible to apply quantitative analysis if the appropriate failure data is available. In these cases, quantitative target values for the maximum probability of the violation of each safety goal can be used; ISO 26262 provides a table with one such possible set of values:

ASIL	PROBABILITY
D	$< 10^{-8} \text{ h}^{-1}$
C	$< 10^{-7} \text{ h}^{-1}$
B	$< 10^{-7} \text{ h}^{-1}$

Table 5 - Possible target probabilities for different ASILs

However, it should be noted that these are only suggestions and can be tailored to fit the circumstances, particularly if e.g. trusted field data is available.

2.4.2 ASIL Decomposition & Allocation Algorithm

2.4.2.1 ASIL Decomposition with multiple hazards

As stated in the previous section, one of the primary objectives of MAENAD is not merely to provide support for the ISO 26262 safety design process but also to support the automatic decomposition of ASILs. Because of the large role ASILs play in ISO 26262, having tool support - and particularly a degree of automation - may be very important. While simple examples that deal with only a single hazard - like that in Figure 5 - can be decomposed manually as they do not have too many elements or too complex logic, when multiple hazards are in play, this becomes significantly more difficult, even for small, seemingly simple systems like that in Figure 6. In this example, there is a system (S) that provides two separate outputs, F1 and F2. The six sub-elements of the system (e.g. functions, components etc), E1 to E6, are the components necessary to provide these two outputs. Each output can fail in one of two ways - omission (lack of provision of output) and commission (unexpected or unwanted provision of output), and thus there are four hazards (omission of F1, commission of F1, omission of F2, and commission of F2) which interact at multiple points.

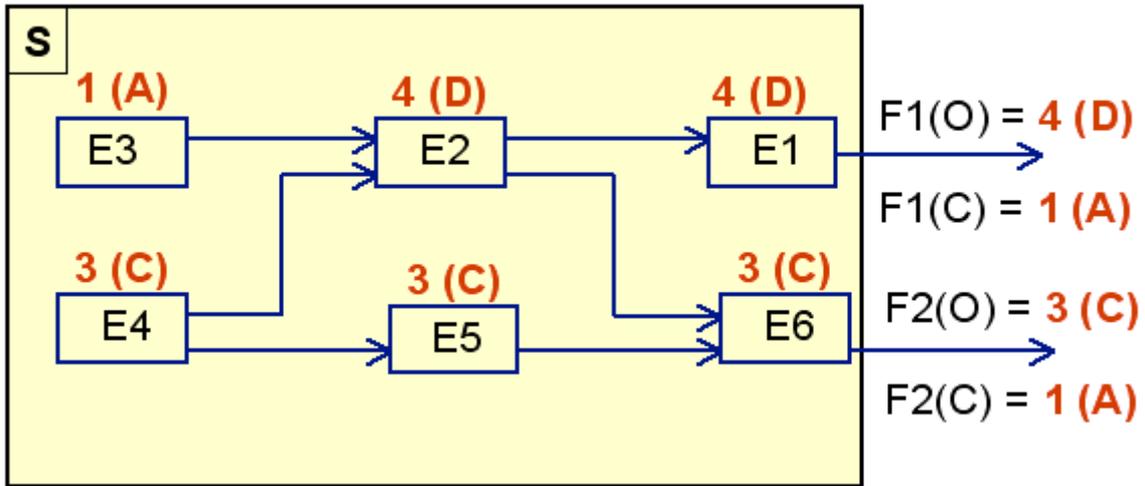


Figure 6 - Dealing with multiple hazard constraints

In most examples, there are many alternative ASIL assignment strategies that could be used (as in the case of Figure 5, for example). However, when there is more than one hazard and thus more than one system safety requirement, the competing demands of these multiple requirements can be used to reduce the number of possible assignments. In the more abstract example above, there is in fact only one possible assignment (there are other valid ones, but they are redundant), though this is definitely not apparent at first sight.

Firstly, each of the four hazards has its own assigned ASIL:

- F1(O) - Omission of output from F1, with ASIL D
- F1(C) - Commission of output from F1, with ASIL A
- F2(O) - Omission of output from F2, with ASIL C
- F2(C) - Commission of output from F2, with ASIL A

As stated above, due to the connecting failure logic amongst the six sub-elements in the system, it transpires that there is only one real possible ASIL assignment. To see why, it is necessary to go into more detail. First, the causes of each hazard must be explored. The failure logic of the system is better shown below in Figure 7. In this figure, the propagation of failure causes for the omission failures and propagation of failure causes for the commission failures are shown separately. The black dots represent internal failures of the components they are in and are numbered according to the element they are in, thus failure event FE3 is a failure of element E3. Furthermore, for the sake of simplicity, we assume that omission failures only cause other omission failures and commission failures only cause other commission failures, though in the general case this is not always true, and failures of one type can cause failures of another type.

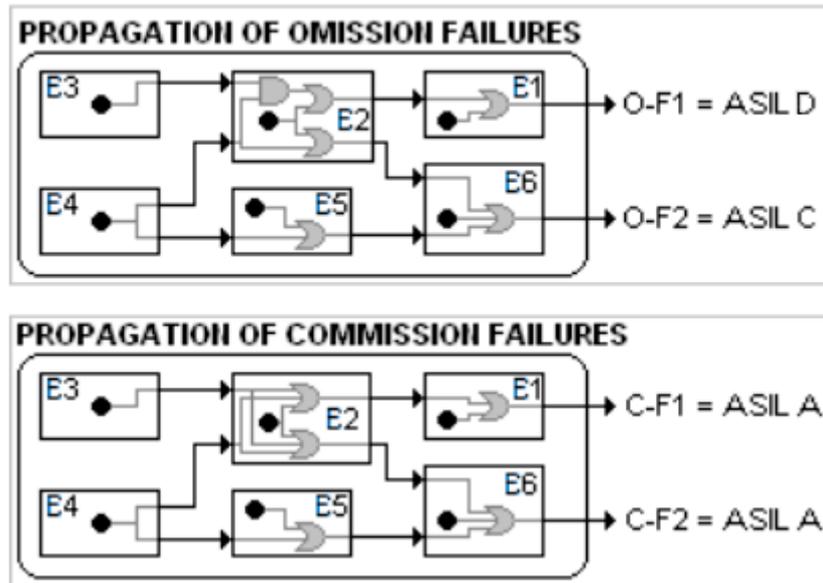


Figure 7 - Failure logic in the example system

As an example, take F1(O). Its initial cause is an omission from component E1, which uses OR logic - it can be caused either by an internal failure mode FE1 (the black circle) or by an omission of input. Since FE1 is directly capable of causing the output failure F1(O), it receives ASIL D. Omission of input to E1 is caused solely by an omission from E2, which in turn is caused either by an internal failure mode (FE2) or by an omission of input at *both* inputs. FE2 therefore receives ASIL D as well. However, omission of input to E2 is caused by a conjunction of omissions from both E3 and E4, so they each contribute some share of the ASIL D. There are five possible assignments for these two components:

- FE3 = QM (0), FE4 = D (4) Total = 4 (D)
- FE3 = A (1), FE4 = C (3) Total = 4 (D)
- FE3 = B (2), FE4 = B (2) Total = 4 (D)
- FE3 = C (3), FE4 = A (1) Total = 4 (D)
- FE3 = D (4), FE4 = QM (0) Total = 4 (D)

Assignments where ASILs for both FE3 and FE4 are *more* than sufficient to meet the safety requirement, e.g. FE3 = ASIL C and FE4 = ASIL C, are also possible, but not considered unless necessary (i.e. unless both FE3 and FE4 directly contribute to an ASIL C hazard somewhere). Furthermore, since both FE3 and FE4 can contribute to the system failure, neither should really be allocated QM according to ISO 26262. However, by assigning each ASIL an appropriate value (e.g. D = 4, A = 1), it can be seen that all of these alternatives are equivalent to ASIL D when summed and so all are at least theoretically valid.

To reduce the number of combinations, it is necessary to look at the other hazards and their ASILs. It is determined through decomposition that commission of F1 (i.e., F1 (C)), with ASIL A, is caused by any internal failure mode of E1, E2, E3, and E4; thus each of those failure modes (FE1-FE4) must be at least ASIL A. That firmly eliminates two possibilities for FE3 and FE4 (namely, where one is ASIL D and the other is QM). Next, decomposition of F2(O) – with ASIL C – shows that there are four possible causes: FE2, FE4, FE5, or FE6. Each therefore requires a minimum of ASIL C. This removes two more possibilities from the list, leaving only one possible allocation: FE3 = ASIL A, FE4 = ASIL C.

In this way, ASIL decomposition using multiple hazards allows us to determine which failure modes contribute to which hazards, and by process of elimination, determine what ASIL each

failure mode should have (or, more frequently, determine a set of many possible ASIL allocations for those failure modes). The new algorithm developed in MAENAD (and originated in ATESS2) supports this kind of multi-hazard ASIL decomposition process.

2.4.2.2 Using FTA as the basis of ASIL decomposition

Before ASILs can be decomposed, it is necessary to understand which parts of the system are independent and whether or not failures can cascade from one element to another. FTA - particularly component-based FTA like that performed by the HiP-HOPS tool (see Section 3.3.2) - is ideal for this, as a fault tree naturally analyses the propagation and combinations of failures throughout a system. The failure logic in the previous example, as shown in Figure 7, is effectively a set of interconnected fault trees: there are four top events (one for each hazard), but only six basic events (the internal failures of the six components). Each failure can therefore contribute to more than one top event, and it is this fact that is the basis of the multiple-hazard ASIL decomposition. The actual fault trees used are presented (in a rather unconventional interconnected form) in Figure 8.

Looking at the figure, it is possible to see how FE2 for example contributes to all four top events; in this case it would have to receive the highest of the top event ASILs (in this case, ASIL D). FE4 contributes to all four events, but one of them (O-F1) is via an AND gate, meaning that it is now subject to multiple constraints:

- It must have at least the same ASIL as C-F1 (A)
- It must have at least the same ASIL as C-F2 (A)
- It must have at least the same ASIL as O-F2 (C)
- Together with FE3, it must add up to the same ASIL as O-F1 (D)

Taken together, the most likely assignment here is to give FE4 ASIL C and then FE3 can have ASIL A, assuming it is not subject to any higher ASIL from different branches.

This system of overlapping ASIL constraints both adds complexity to the process and also adds value to it, since it allows us to reduce the total number of possible ASIL assignments (and in general, there are a lot). Furthermore, fault trees are one of the best ways of being able to model these overlapping constraints. Other techniques that deal only with single failures, like FMEAs, would be unsuitable for this as they cannot show the effects of conjunctions of failures, which is what makes decomposition possible.

In particular, compositional FTA (the HiP-HOPS approach) - in which the system fault trees (e.g. O-F1, C-F2 etc) are synthesised by connecting the individual fault trees of the separate components (E1, E2 etc) - is ideally suited for this and can readily create such fault trees.

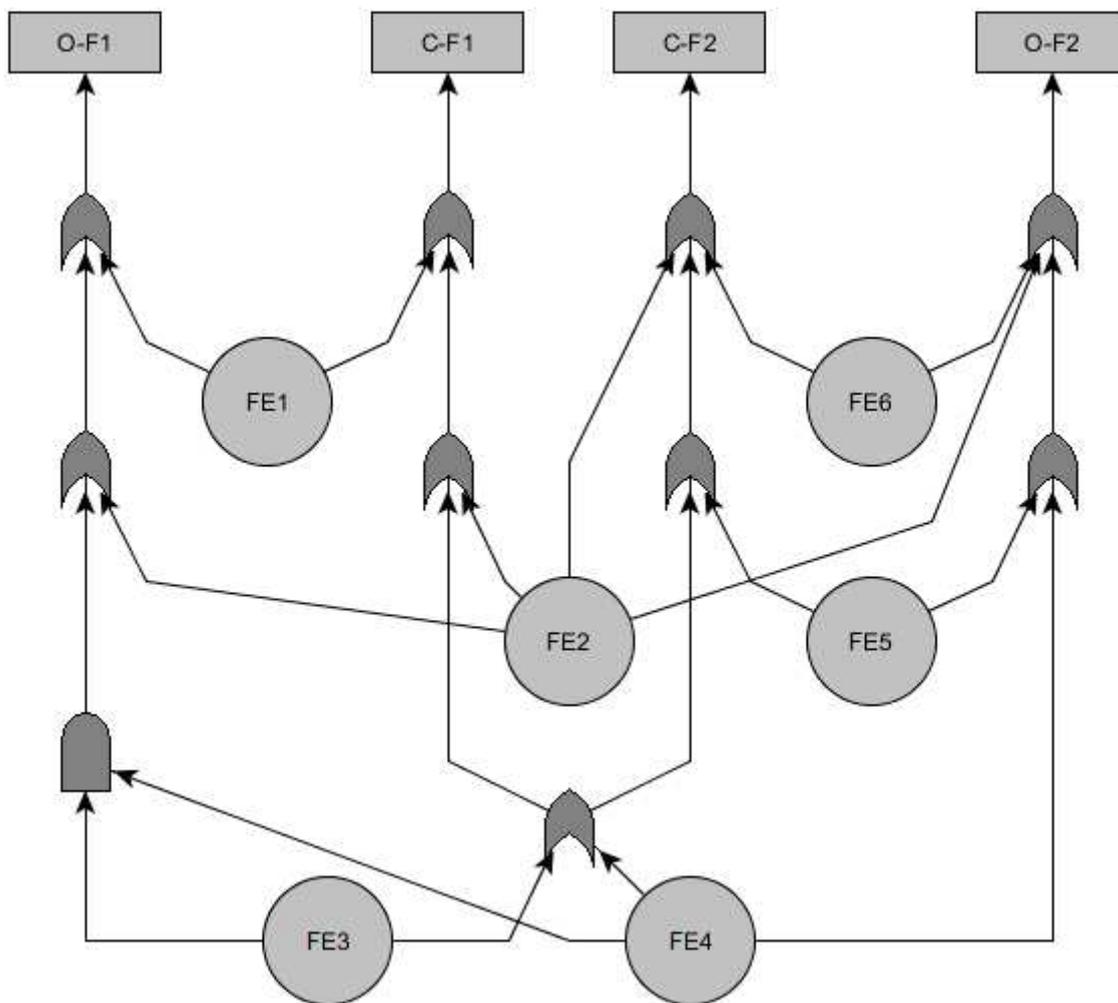


Figure 8 - Interconnected FTs

In fact, FTA has an even bigger advantage in that what really matters are the minimal cut sets of the fault trees, i.e., the set of smallest possible combinations of basic events capable of causing the system failures (or hazards in this case). Since these are calculated directly by FTA, it is much easier to see the relationships between the basic events and the hazards they cause. For example (where the black dot represents a conjunction):

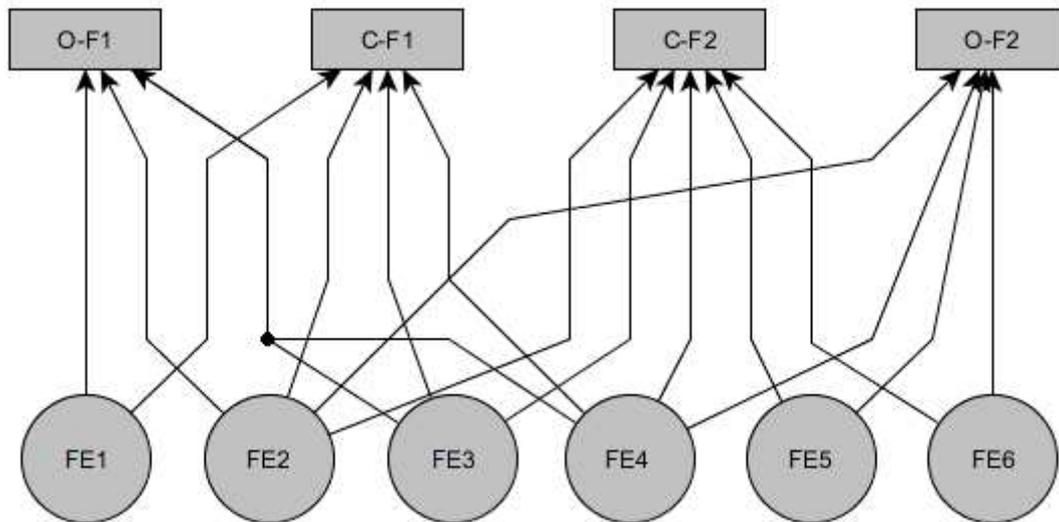


Figure 9 - Direct connections from basic events to hazards

On the basis of these minimal cut sets, we can allocate the hazards' ASILs directly to the cut sets, and then (if applicable) decompose them across the basic events that comprise those minimal cut sets.

2.4.2.3 Decomposition Algorithm

The first step in the algorithm is to ensure that a set of *safety constraints* (i.e., ASILs applying to one or more output failures in the system) have been defined. These safety constraints, and their assigned ASILs, may have been produced as a result of a prior hazard analysis of the system (in which case the ASILs are derived from the hazardous events and safety goals) or may simply be assumed in the case of a SEooC. These safety constraints (and their ASILs) constitute one input to the algorithm; the other is a logical representation of the failure behaviour of the system, e.g. as described by the results of a fault tree analysis. Typically, one fault tree must be generated for each safety constraint, modelling the causes of the *constrained output failure* (COF).

Thus:

Inputs:

- Set of safety constraints, each of which should have an ASIL and a fault tree generated for its constrained output failure.
- Set of minimal cut sets (from the fault trees modelling the logical failure behaviour of the system) that cause the constrained output failures.

Outputs:

- Set of possible ASIL assignments for all basic events in the form of *Safety Requirement Assignments*; there will typically be many of these.

The important entities in the algorithm are as follows:

- **Safety Constraint (SC)**

Represents an ASIL applied to an output failure (malfunction) of an element of the system (or of the system itself). The output failure becomes the top event of a fault tree and its assigned ASIL serves as the basis of the decomposition.

- **Cut Set (CS)**

Cut sets are the results of the FTA. They contain sets of basic events (BEs) and are responsible for causing the COFs.

- **Basic Event (BE)**

A basic event (or BE) is an individual component or function failure in the system. Individually it may or may not have any effect on the system, but if it is a member of a cut set, it will contribute to at least one COF. Basic events are the primary targets for assignment of ASILs from the safety constraints they affect (whether directly or indirectly). When allocated an ASIL from a Derived Safety Requirement (or DSR - see next), the DSR also 'locks' them at that ASIL so that only DSRs with higher ASILs can change them until the original DSR releases the lock. This prevents BEs being incorrectly allocated lower ASILs, so a BE can always be allocated a higher ASIL, but never a lower one.

- **Derived Safety Requirement (DSR)**

In the context of the algorithm, a Derived Safety Requirement describes an ASIL value that a set of basic events (i.e, a minimal cut set) must implement to achieve required system risk levels. DSRs are produced by the ASIL decomposition and allocation process. More than one DSR may apply to the same basic event, in which case the highest ASIL from all the DSRs that apply to that BE is used.

- **Safety Requirement Allocation (SRA)**

A Safety Requirement Allocation is the output of the algorithm and represents a set of possible ASIL assignments to the basic events of the system that will meet the initial safety requirements as defined by the Safety Constraints. Typically many SRAs will be generated.

The algorithm itself is described in pseudocode below.

ASIL Decomposition Algorithm

```
for each safety constraint (SC):
{
  for each cut set (CS) in SC:
  {
    Generate new 'Derived Safety Requirement' or DSR (contains events and ASIL from SC)
  }
  Sort the DSRs - least basic events first, then highest ASIL first.
}
```

Recurse for every derived safety requirement (DSR):

```
{
  while there remain unallocated DSRs:
  {
    if this DSR applies to only one basic event:
    {
      Assign DSR's ASIL to the only BE
      If the assignment succeeded and BE's ASIL is >= DSR's ASIL:
      {
        BE is locked at that ASIL by this DSR {*1*}
        Recurse for next DSR
      }
    }
    else
    {
      Recurse for next DSR
    }
    Unlock BE
  }
  else (DSR applies to more than one BE)
  {
    Iterate through all possible assignments of DSR's ASIL for these BEs {*2*}
    {
      If this assignment meets the DSR's ASIL:
      {
        Attempt to assign ASILs to BEs {*1*}
        If it worked and the DSR is achieved, recurse for next DSR
      }
    }
    Unlock all BEs
  }
}
```

```

else (there are no remaining DSRs)
{
    Store the current configuration of BEs & ASILs as one possible assignment (SRA) {*3*}
    Backtrack up recursion stack to try new combinations
}
}

```

The above algorithm describes the overall process, but it is worth describing certain aspects (numbered) in more detail:

1) Allocating ASILs to Basic Events

The key to the algorithm is the priority locking mechanism used by the Basic Events. When a Derived Safety Requirement attempts to allocate an ASIL to a Basic Event, this operation only succeeds if one of the following is true:

- The Basic Event is not locked
- The BE is locked, but the DSR is the one who holds the most recent lock
- The DSR's ASIL is higher than the ASIL of the current lock holder

When the assignment succeeds, the DSR that allocated it locks the Basic Event and registers itself with the Basic Event as the current lock holder (the BE retains a stack of lock holders). This ensures that a Basic Event can only be given an equal or higher ASIL than the one it currently holds, never a lower one (for this purpose, an ASIL is equated to a numeric value, e.g. A = 1, B = 2, C = 3 etc). If the DSR's ASIL is lower, then the new assignment fails and the algorithm will potentially backtrack to try a new combination of assignments.

2) Iterating through possible combinations of assignments for a set of BEs

There can be many possible combinations of assignments for any given set of BEs. The theoretical maximum number of permutations is given by the formula:

$$p = (m + 1)^n$$

where p is the number of permutations, m is the maximum ASIL value (4 in this case, but then we add 1 because we include 0/QM), and n is the number of basic events constrained by the SR. So for a 2 event cut set, there is a maximum of 25 possible assignments; for 3 events, there are 125.

Each assignment is determined mechanically, so it will start with e.g. 0/0/0, then 0/0/1, then 0/0/2 etc, and assign each of these values to the targets. However, not all of these permutations are used. Firstly, those that do not meet the current DSR are immediately discarded; for example, if the required ASIL value is 3 and the current permutation is 0/0/1 (sum = 1, which is less than 3), this would not even be tested. Secondly, it is possible for one or more of the BEs to already have been locked by earlier DSRs; in these cases, permutations in which those BEs have lower ASIL values are not even enumerated and are thus skipped. This helps to improve the efficiency of the process. Note however that assignments that *more* than meet the DSR are kept at this stage, so 1/2/3 (total = 6) would be kept even if the required total was only 3.

3) Storing possible assignments

Once all safety requirements have been examined and all of their ASILs assigned to basic events in such a way that all of the DSRs are still being achieved (i.e., no combination of basic event ASILs add up to less than the ASILs of any of the hazards they cause), that assignment of ASILs

and corresponding basic events is stored as a Safety Requirement Allocation (SRA). Safety Requirement Allocations are nothing more than dictionaries using basic events as keys and allocated ASILs as their corresponding values.

However, there can be a great many Safety Requirement Allocations and it is important to cut down the number wherever possible. Therefore, we can do some redundancy checking at this stage. The main form of checking is to see whether one allocation renders another redundant. For example, if basic event BE1 is assigned ASIL C and BE2 is assigned ASIL B, and that is a valid allocation (i.e., all DSRs are met), then an allocation in which BE1 is assigned ASIL D and BE2 is assigned ASIL C would be redundant. To borrow an optimisation term, it is *dominated* by the first allocation. It is important to note that this is only the case if *all* of the ASILs are higher than (or equal to) another allocation; if BE1 was ASIL B and BE2 was C, then it would not be redundant.

To extend the example, assuming two BEs in total, BE1 and BE2, and four allocations added in the following order:

	BE1	BE2	TOTAL	
1)	C (3)	B (2)	D (5)	
2)	C (3)	A (1)	D (4)	- Makes #1 redundant, so #1 is removed
3)	D (4)	A (1)	D (5)	- Is made redundant by #2, so not added
4)	B (2)	B (2)	D (4)	- Is <i>not</i> redundant

In this case there are two non-redundant allocations (#2 and #4) and two redundant ones (#1 and #3). By keeping the number of allocations to a minimum throughout the process, we reduce the amount of checking we need to do when we add later SRAs. Even so, this can still be an expensive process, as every new SRA needs checking against every existing SRA. There are potential optimisations possible here, however.

2.4.2.4 Example

To see how the algorithm works when applied to the previous 6-component example in Figure 6, we can work through the algorithm manually (albeit skipping the sorting for illustrative purposes).

First, we need a list of minimal cut sets, together with the hazards they cause and associated ASILs:

O-F1 - ASIL D (4)

{FE1} {FE2} {FE3.FE4}

C-F1 - ASIL A (1)

{FE1} {FE2} {FE3} {FE4}

O-F2 - ASIL C (3)

{FE2} {FE4} {FE5} {FE6}

C-F2 - ASIL A (1)

{FE2} {FE3} {FE4} {FE5} {FE6}

We can now generate a set of Derived Safety Requirements, one for each cut set. This tells us what different ASILs apply to each cut set. We can represent these in the form of equations, for clarity:

{FE1=4}	{FE2=4}	{FE3 + E4=4}		
{FE1=1}	{FE2=1}	{FE3=1}	{FE4=1}	
{FE2=3}	{FE4=3}	{FE5=3}	{FE6=3}	
{FE2=1}	{FE3=1}	{FE4=1}	{FE5=1}	{FE6=1}

Next we have to recurse through all of these DSRs, attempting to lock the Basic Events at each point. To make it clear what's happening, the status of the BEs (and their locks) is shown at each point. Bold event values indicate a successful lock; italic values indicate failure. We shall only examine the first full recursion.

	FE1	FE2	FE3	FE4	FE5	FE6
1. FE1=4	4 (#1)	-	-	-	-	-
2. FE2=4	4 (#1)	4 (#2)	-	-	-	-
3. For 3, we have two events, so we need to iterate through all permutations that add up to at least 4. We shall consider these separately.						
3a: FE3=0, FE4=4:						
	4 (#1)	4 (#2)	0 (#3a)	4 (#3a)	-	-
4. FE1=1	4 (#1)	4 (#2)	0 (#3a)	4 (#3a)	-	-
5. FE2=1	4 (#1)	4 (#2)	0 (#3a)	4 (#3a)	-	-
6. FE3=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	-	-
7. FE4=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	-	-
8. FE2=3	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	-	-
9. FE4=3	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	-	-
10. FE5=3	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	-
11. FE6=3	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	3 (#11)
12. FE2=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	3 (#11)
13. FE3=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	3 (#11)
14. FE4=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	3 (#11)
15. FE5=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	3 (#11)
16. FE6=1	4 (#1)	4 (#2)	1 (#6)	4 (#3a)	3 (#10)	3 (#11)

We have now run out of DSRs and so can store this as one possible SRA for the system:

$$FE1 = 4, FE2 = 4, FE3 = 1, FE4 = 4, FE5 = 3, FE6 = 3$$

As it turns out, the next recursion gives us a better solution. When we backtrack, we go back to the last time we have more than one basic event in a SR and were iterating through possible permutations. In this case, that SR was #3, so we can skip back to there and continue with #3b instead:

3b: FE3=1, FE4=3:

	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
4. FE1=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
5. FE2=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
6. FE3=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
7. FE4=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
8. FE2=3	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
9. FE4=3	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	-	-
10. FE5=3	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	-
11. FE6=3	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	3 (#11)
12. FE2=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	3 (#11)
13. FE3=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	3 (#11)
14. FE4=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	3 (#11)
15. FE5=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	3 (#11)
16. FE6=1	4 (#1)	4 (#2)	1 (#3b)	3 (#3b)	3 (#10)	3 (#11)

Notice how this time, no overriding took place. Our allocation this time is:

$$FE1 = 4, FE2 = 4, FE3 = 1, FE4 = 3, FE5 = 3, FE6 = 3$$

and when we compare that to the previous one, we find the only difference is that this one has a lower ASIL for FE4. Therefore, this new allocation renders the previous one redundant, so the old one is removed and the new one is kept. As it turns out, in this example, this is the only undominated allocation, and it would render all other allocations (like FE1=4, FE2=4, FE3=2, FE4=3, FE5=3, FE6=3) redundant.

In cases where there are more results, these will typically be sorted (e.g. by total ASIL value, or by weighted total where A=1, B=10, C=100, D=1000 etc) to make it easier for the analyst to absorb the results by presenting the most attractive possibilities first. However, only the analyst can decide for sure what is a good allocation and what is not.

2.4.2.5 Scope of the decomposition

Although the algorithm decomposes and allocates ASILs directly to internal faults of system functions/components, it is also possible - once this has been done - to assign ASILs to other things, including propagated failures (or deviations), the ports that define the interface of a function/component, or even the function/component as a whole (e.g. to represent a process fault). This is done simply by taking the highest ASIL that propagates through a given port or the highest ASIL that passes through a component. The raw information provided by ASILs allocated to basic events is very flexible and this information can be used to derive allocations to virtually any other attribute or part of the system model.

2.4.3 HiP-HOPS and Tool Support for ASILs

2.4.3.1 Current HiP-HOPS Support

The algorithm described in Section 2.4.2 is already implemented as part of HiP-HOPS, though only in prototype form and likely not bug-free. It can be activated by including a new command line option: "decomposeSILs=true".

The main input to the process is the Hazard, which is included as part of the new XML interface to HiP-HOPS. It is intended for Hazards to replace the older System Output style of system failure definition, as they are more flexible and more useful (e.g. to support ASIL decomposition).

The XML *Model* element contains an optional *Hazards* element, which in turn contains any number of *Hazard* elements defining the system-level hazard failures. Each Hazard contains:

- Name (string)
- Description (string, optional)
- Safety Requirement (integer, optional) - represents the ASIL. Uses integers, so ASIL D = 4, ASIL C = 3, ASIL B = 2, ASIL A = 1, QM = 0.
- Severity (double, optional)
- A Hazard Expression (string) - a logical expression consisting of one or more output deviations in any perspective of the model connected by AND or OR operators. Advanced operators (NOT, temporal) are not applicable here (yet).

This is all the information that HiP-HOPS needs to commence ASIL decomposition, as it obtains the fault trees and cut sets itself by analysing the model it is given. The algorithm has not been tested on large models and it is anticipated that the process will be relatively slow for larger numbers of basic events and safety requirements (due to the combinatorial explosion nature of enumerating all possible valid allocations of ASILs). More control over the process (e.g. timeouts, logging information) could be provided by adding further command line parameters.

At the other end, the output can be included (assuming ASIL decomposition takes place) as part of the HTML-style output currently produced by HiP-HOPS:

Safety Allocations			
A:failed	B1:failed	B2:failed	C:failed
4	0	4	4
4	1	3	4
4	2	2	4
4	3	1	4
4	4	0	4

Figure 10 - Example ASIL decomposition output

The output shows each basic event in a different column and then each row contains a different ASIL allocation that should meet the original ASILs defined in the Hazard. However, it may be that

a more formalised output needs to be decided on, particularly if the results are to be imported back into the model (see below).

2.4.3.2 Plugin Extension

To provide rudimentary support for ASIL decomposition, the EAST-ADL/HiP-HOPS plugin should be extended to support the output of hazards in the newer HiP-HOPS XML input format. This information is derived from the SafetyConstraints in the EAST-ADL model. It also needs to be possible to activate the process by sending the new command line parameter (when appropriate).

2.4.3.3 Expanding the scope of decomposition

As currently implemented, HiP-HOPS theoretically allows ASILs to be assigned to any cut set-capable event, thus including things like input and output deviations. However, due to the way HiP-HOPS handles these extended node types (they do not ordinarily appear in the results and in fact usually get stripped out early in the analysis process), actually performing decomposition and allocation with them is not necessarily a straightforward proposition.

In practice, the simplest solution may in fact be to perform decomposition as it stands and then go back to the original trees and propagate the ASILs upwards from the basic events. This would ensure that all input/output deviations receive the correct ASILs. In case of conflicts, the highest ASIL received would be retained and, in theory, the ASILs of the system output deviations should match the ASILs of the hazards they cause. Doing this would also allow us to assign ASILs to ports relatively easily by taking the highest ASIL of any deviation at that port.

Allocation of ASILs to components and subsystems is substantially easier since we already know which basic events belong to which components. By simply taking the highest ASIL of all the basic events and deviations of a component, we can 'allocate' that ASIL to the component. Or, in the case of subsystems, do the same thing but take the highest ASIL of its subcomponents instead.

2.4.3.4 Storing the results

During the York meeting in May 2011, it was decided that the goal for the first year of MAENAD would be to present the ASIL decomposition results in HiP-HOPS form only: the aim is to present a list of possible Safety Requirement Allocations (as in Figure 10) and allow the user to select one; this would then present a new display showing a tree-based view of the system and how the ASILs are assigned to its constituent elements (specifically, the basic events, the input/output deviations, and the component/process fault ASIL).

More substantial work would be needed to be able to import the results back into the EAST-ADL model. To begin with, the decision would have to be made whether *all* possible SRAs would be stored, or whether the user would select the most appropriate allocation and only that one would be stored. In the latter case, the existing EAST-ADL language elements, e.g. the SafetyConstraint, would likely be sufficient to store the result (unless we also wished to be able to allocate ASILs to other parts of the model too), but some kind of selection dialogue would be necessary. However, in the former case, we would need to create some kind of storage mechanism to be able to map the different possible ASIL assignments to their targets without interfering with each other.

On a technical level, we would also need to import the results from the HiP-HOPS output files back into whichever modelling tool (e.g. Papyrus) was being used. This has not yet been done even for ordinary fault tree results, however. It would likely involve a consolidation of HiP-HOPS output into a single XML file to make it easier for the file to be imported back into the EAST-ADL model. The plugin may also need to perform some further model transformation on these results to know where they should be stored.

2.5 Discussion & Further Work

2.5.1 Support for ASILs

ASIL decomposition and allocation is an important objective for MAENAD and a major requirement in order to be able to fully support ISO 26262-compatible safety-driven design. This section details a newly developed algorithm that enables the automatic decomposition and allocation of ASILs across independent elements of the system by building upon earlier work on FTA; ASILs assigned to hazards can then be decomposed to the minimal cut sets that cause those hazards. By enumerating the different permutations of those ASILs assigned to multi-event cut sets in a recursive process, it is possible to determine all possible valid ASIL allocations for the basic events of a system while ensuring that the resulting allocations are still capable of meeting the original safety requirements. EAST-ADL language support is relatively mature and language elements for both hazard analysis and ASILs are present. However, it may be that these need tweaking or extending to streamline the process in response to practical experience.

2.5.1.1 Existing Support for ASIL decomposition

At present, much of the infrastructure required to support ASIL decomposition and analysis is already present in both EAST-ADL and HiP-HOPS. In particular, EAST-ADL supports:

- Hazard analysis and definition of Hazards, HazardousEvents, and SafetyGoals. HazardousEvents and SafetyGoals can both store ASIL values.
- SafetyConstraints can be used to assign ASILs to elements of the error model.
- SafetyConstraints may also provide a mechanism for linking the resulting ASIL allocations back to the faults & failures of the error model after decomposition.

The automatic ASIL decomposition and allocation algorithm has been implemented in HiP-HOPS, along with preliminary support for input and output. This may need extending in the future but is sufficient for early testing and experimentation. The EAST-ADL/HiP-HOPS plugin requires extending to support this capability but the impact on the input interface to HiP-HOPS is hopefully relatively minor. HiP-HOPS can also produce output in its existing HTML-style format.

2.5.1.2 Further work and potential changes

Although an algorithm has been developed to enable automatic ASIL decomposition, and although a lot of the foundation for supporting that algorithm now exists both in the EAST-ADL language and in the HiP-HOPS tool, there is still plenty of scope for further work in all areas. Much of that work revolves around the following issues.

2.5.1.3 Linking the language support to the tools

EAST-ADL has sufficient language support to enable the decomposition and the algorithm has been implemented in the tool. The main obstacle at present is the link between the two: namely, the Papyrus plugin, which needs extending to enable the output of hazards and ASIL information to HiP-HOPS and allow the starting of the decomposition process. This will enable us to begin testing the decomposition of actual EAST-ADL models with the algorithm, which will highlight bugs to be fixed and other areas for further work (e.g. any streamlining or clarification of the existing language elements, any additions necessary to the language, what sort of bugs exist in the tools, and how efficient and scalable the algorithm is).

2.5.1.4 Presentation and storing of the results

Currently the results are presented as part of the standard HTML-style HiP-HOPS results. It is also sorted in a rudimentary fashion. However, some additional forms of filtering and/or sorting could be applied to the results as well, e.g. using a point-based heuristic that favours more lower ASILs over fewer higher ASILs. Also, the results currently only display what is essentially the raw assignment information: the direct assignment of ASILs to basic events. More development is necessary to be able to display the results in a more refined form, or display assignments of ASILs to other elements (e.g. output/input failures).

A secondary issue is that of storing the results back in the model. This is a general unsolved issue not specific to ASILs (e.g. the fault tree results are currently completely external as well) but is something to be investigated further.

2.5.1.5 Scope of decomposition

At present, the algorithm focuses solely on allocating the ASILs to the basic events (component/function failures) of the system. This is done partly because it links in most closely with the way FTA and its analysis of system failure propagation works, and partly because it also offers a very sound foundation for further expansion into assignment to other model or system elements (e.g. ports or entire components).

The goal for MAENAD is to expand this to also allow assignment to input/output failures and entire components (to represent process faults). This will require an extension of the prototype implementation in HiP-HOPS as well as a new method of displaying the results. However, at the same time, it should not impede preliminary testing with existing support for ASILs.

2.5.1.6 Scalability of the algorithm and optimisation

Due to the nature of the problem - the number of possible permutations increases very rapidly with the number of AND gates and basic events, leading to combinatorial explosion - it is unlikely that the algorithm is very scalable. Quite how scalable it is has yet to be fully tested. In any case, it can almost certainly be made more efficient in certain areas and there may even be other approaches that help overcome its difficulties. This does warrant some further investigation.

Furthermore, a more radical alternative may be to re-frame the process of ASIL allocation as an optimisation problem and employ genetic algorithms to solve it. This would enable us to employ existing HiP-HOPS technology and may result in a more efficient process. However, it would obviously also mean investigating how best to go about doing this and would entail both theoretical and practical work on HiP-HOPS and its optimisation capabilities.

At the York meeting in May 2011 it was decided to attempt to develop prototype ASIL optimisation support by September 2011 to determine the feasibility of the process. This algorithm is described below in section 2.6.

2.5.1.7 Examples and Case Studies

As always, it would be highly beneficial to have some larger scale examples to work with, both for bug hunting purposes and to ensure that we are properly supporting an ISO 26262 compatible ASIL decomposition. Furthermore, small but realistic examples would be very valuable as case studies to use in papers etc for dissemination purposes.

At the York meeting in May 2011, VTEC agreed to develop a brake-by-wire example for use with ASIL decomposition & allocation.

2.6 Optimisation algorithm for ASIL Decomposition

The ASIL allocation and decomposition problem can also be solved as an optimisation process. This differs to architectural optimisation (discussed in section 4) in that the domain is the set of possible ASIL allocations to the basic events (or other elements) of the system, rather than the set of possible architectural configurations in the model. Furthermore, the scope of the evaluation is much more limited since the other factors that count are whether or not an ASIL allocation is valid, i.e., whether or not it meets its requirements, and how 'expensive' the allocation is.

2.6.1 The ASIL Optimisation Process

The inputs to the optimisation are the same as to the standard ASIL allocation/decomposition algorithm: namely, the set of constrained output failures (COFs) and the safety requirements derived (DSRs) from them that apply to the causes of those output failures:

Inputs:

- Set of safety constraints, each of which should have an ASIL and a fault tree generated for its constrained output failure.
- Set of minimal cut sets (from the fault trees modelling the logical failure behaviour of the system) that cause the constrained output failures.

Similarly, the output is also the same – a set of safety requirement assignments (SRAs):

Outputs:

- Set of possible ASIL assignments for all basic events in the form of *Safety Requirement Assignments*.

The difference is in how these SRAs are obtained. Whereas the standard algorithm essentially performs a brute force search through the entire optimisation space to find all valid solutions, the optimisation will perform a more randomised search, guided by heuristics, in an attempt to obtain a reasonable set of solutions without having to search through every possible valid solution.

The same genetic algorithm technology described in Section 4 is also applied in this case. The principle is that the optimisation follows a process similar to evolution: candidates evaluated to have good fitness (in this case, they meet the DSRs and are not too expensive) will be retained and will pass on its 'genes' to the next generation, whereas candidates with poor fitness (e.g. they do not meet the safety requirements, or are very expensive) will be rejected and will die out.

There are two main elements to consider when devising a genetic optimisation algorithm in this way: the encoding format and the evaluation heuristics.

2.6.2 Encoding

The encoding used for the ASIL decomposition optimisation is much simpler than that used for the architectural optimisation. Rather than have a hierarchical tree structure to represent the makeup of the system, the ASIL optimisation simply uses the set of relevant basic events and assigns each of them an ASIL value between 0 (QM) and 4 (ASIL D). Thus each encoding can in fact be represented as a string of numbers between 0 and 4. This effectively serves as the 'genes' of each individual; successful individuals will have some or all of their encoding string preserved in the next

generation (meaning that some or all of their ASIL assignments remain the same), whereas unsuccessful allocations of ASILs will tend to die out instead.

The format of the encoding is very similar to that of the SRA, which means they can be output upon completion of the optimisation almost directly.

2.6.3 Evaluation Heuristics

The other aspect of the ASIL optimisation algorithm is the evaluation heuristics. These are used to evaluate each individual in a population to determine their “fitness” scores – high-scoring individuals are kept, poor-scoring ones are not. Typically the optimisation problem itself is defined in terms of *objectives* – e.g. minimise cost, maximise safety – and thus we need to be able to evaluate according to those objectives. In a single objective approach, multiple evaluation scores are combined into one value, which is then to be either maximised or minimised. In multiple objective approaches, the scores remain separate and can be maximised or minimised individually.

In the case of the ASIL optimisation, there are two objectives to consider. The first is cost: how expensive a particular ASIL allocation might be to actually implement. This need not be particularly accurate or realistic – it only matters that it helps the algorithm to differentiate between different solutions. In our case, we chose a simple heuristic that scores each ASIL in an allocation like so:

- QM = 0 points
- ASIL A = 1 point
- ASIL B = 10 points
- ASIL C = 100 points
- ASIL D = 1000 points

Other schemes are also possible, including linear scales (e.g. A = 1, B = 2, C = 3, D = 4) and powers of two (2, 4, 8, 16 etc).

The ASIL values in an SRA are then summed together to find a total cost for that particular SRA. Thus an SRA of AABD will get a cost of 1012, whereas AACC would have a cost of 202.

The other heuristic to consider is the number of invalid (failed) allocations. Each cut set in the FTA results is subject to one or more DSR – derived safety requirement – that indicates the minimum ASIL that has to apply to that cut set in order to meet the overall system safety requirements. Unlike the exhaustive ASIL decomposition algorithm, which guarantees that all SRAs will be valid, the optimisation makes no such guarantees (and in fact doing so could hamper the algorithm’s performance and success). Thus we award 1 penalty point for every cut set that does not meet its DSRs; the goal is obviously to obtain an SRA with 0 penalty points.

2.6.4 Optimisation algorithms

Two genetic optimisation algorithms have been tested thus far. One is a penalty-based single-objective algorithm, whereas the other is a multi-objective algorithm similar to that used for architectural optimisation in Section 4.

2.6.4.1 Penalty-based algorithm

The penalty-based algorithm tries to obtain a single fitness score for each individual by combining the cost heuristic with a penalty for each invalid cut set. The penalty is scaled by a user-defined factor (e.g. x2, x3 etc) and added to the overall cost; furthermore, the penalty is also dynamically scaled by the algorithm depending on the progress through the optimisation itself. Early on in the

process, the penalty is reduced, to encourage early experimentation, whereas later in the process, the penalty is increased to encourage the optimisation to converge on current solutions rather than trying new ones. The objective is to minimise the combined cost.

So far the penalty-based algorithm has worked reasonably well for small examples (such as the 6-component example given earlier). For larger examples with larger search spaces, however, it often fails to converge on a valid solution, suggesting that further experimentation and refinement is needed.

2.6.4.2 Multi-objective algorithm

The multi-objective algorithm is based on the architectural optimisation approach described later in Section 4. It attempts to balance the two different objectives: minimising cost while also minimising the number of invalid solutions. These can be thought of as conflicting, because reducing the cost also makes it harder to meet the derived safety requirements, thus the process is always a trade-off between the two objectives.

Although this algorithm works for small examples, the main problem with this approach is that there is insufficient incentive to have zero invalid cut sets. The optimisation attempts to minimise them, but to the algorithm, it is the same difference between 100 and 101 invalid cut sets as it is between 0 and 1 invalid cut sets. As a result, it struggles to find many valid solutions for larger examples. However, by changing the priority of the heuristic to focus more on having zero invalid solutions, and by tweaking the parameters a bit, the performance of the algorithm can hopefully be improved.

2.6.5 Implementation & Performance

The optimisation is implemented in the HiP-HOPS tool, building on its existing support for optimisation and the prototype ASIL decomposition algorithm described earlier. As a result, performance between the standard exhaustive algorithm and the optimisation approaches can be directly compared. The central advantage of the optimisation approach is that it does not require HiP-HOPS to consider every possible solution, meaning that it performs much, much quicker; the trade-off is that it will not discover every solution in each case. The aim is simply to obtain a reasonable number of good, valid solutions.

For the small six component example described earlier, both the exhaustive algorithm and the optimisation algorithms are very quick (<1 second), and in the case of the optimisation, it also finds the valid solution. For larger examples, the exhaustive algorithm performs much worse, but will usually find all valid solutions eventually (unless the algorithm runs out of memory or is cancelled early). The optimisation algorithms scale far better but are not guaranteed to find any valid solutions, although the penalty-based approach has thus far been more successful. In one medium-sized example, the penalty-based optimisation took 18 seconds versus the exhaustive approach's 18 *hours*, but the exhaustive algorithm did find all 50 valid solutions whereas the optimisation found only one.

Continued testing and development is required to try to improve the success rate of the optimisation algorithms. Their advantage in performance is clear as optimisation is far more scalable than standard exhaustive approaches; however, to be more practical, they need tweaking further to yield more valid solutions.

In addition, further developments of the standard exhaustive algorithm may also produce good results, although it is doubtful that it will ever be fast enough to be applied to large systems.

3 Model-based system analysis techniques

This section presents an overview of potential model-based analysis techniques in the context of the MAENAD project and related analysis techniques.

3.1 Potential analysis support by EAST-ADL

Figure 11 shows the overall landscape of potential architectural analysis support by EAST-ADL, structured according to IEEE1471/ISO42010. As shown in the figure, each analysis is motivated by certain engineering concerns of stakeholders. According to the particular process of system development, an analysis may focus on certain levels of abstraction or particular design decisions. To carry out an analysis, the system properties of concern must be described with EAST-ADL in such a way that the conformity with underlying analysis techniques and tools in regard to the syntax and semantics can be guaranteed. For example, the analysis of safety constraints (Analytical Viewpoint) is motivated by the concern of failure modes and their consequences (Concern), which in turn constitutes the basis for deriving component safety requirements. This particular analysis can be based on some formal or semiformal model (Description) of system operational situation, errors and error propagations, and carried out on a particular level of abstraction (Scope) from the view point of a safety engineer (Stakeholder).

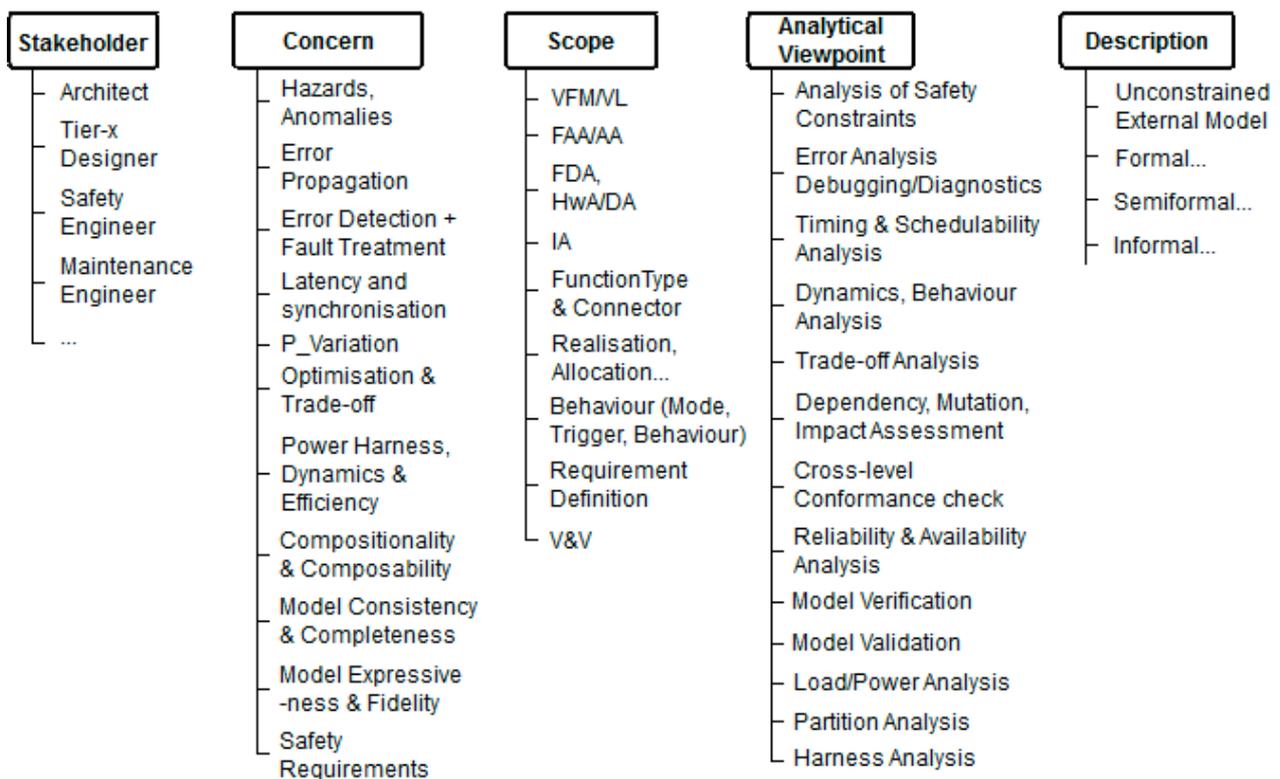


Figure 11 - Landscape of potential analysis support by EAST-ADL.

As a prerequisite to clarify the analysis support of EAST-ADL to be enhanced or provided through MAENAD, the engineering concerns that are of particular interest from an industrial point of view, as well as the prioritisation for FEV (Fully Electrical Vehicle), have to be defined. For example, the stakeholders requirements from WP2 of MAENAD need to be refined into a set of well structured language requirements, explaining the stakeholders, their concerns, expected viewpoints and related system scope. Such stakeholder viewpoints in turn will constitute the basis for specifying

the expected analytical viewpoints, expected EAST-ADL modelling support, and the specifications of related analysis engines.

Two categories of analysis can be benefited by EAST-ADL:

- **Class I – Logical analysis of topology-centric properties:** this type of analysis is directly based on the specifications of architectural artefacts declared with EAST-ADL. The assessments predict for certain quality attributes by composing related component attributes according to their topological relations in the architecture through some given Boolean/Arithmetic logics. Analyses belonging to this category include architectural change impact assessment, partition-isolation analysis, fault tree based error analysis, preliminary safety analysis (PHA), utilisation and cost estimation, etc.
- **Class II – Static and dynamic analysis of behaviour-centric properties:** this type of analysis is based on the specifications of behaviours of architectural artefacts declared with EAST-ADL. The assessments predict for the compositionality and composability of behaviours in requirements, design solutions and their nonfunctional constraint specifications, based on certain behaviour semantics such as the formalisms for discrete-event, continuous- and discrete-time models. Due to the complexity, the support of external analysis engines is normally preferred. Analyses belonging to this category include dynamic behaviour analysis, cross-level behaviour conformance check (e.g., behaviour refinement), model validation in regard to use cases.

3.2 Modelling support for behaviour analysis

3.2.1 A taxonomy of behaviours and behaviour models in system development

The nominal behaviours in an embedded system can be categorised into

- **Computational behaviour** – related to the transfer functions of functional blocks and their corresponding modes. It can be captured as constraints on the block I/O boundary or internal computation, such as in terms of parameter invariants, state-machines, control and data-flows, detailed logical and mathematical operations.
- **Communication behaviour** – related to the interactions between different functional components and the corresponding modes. It corresponds to the network protocol or other synchronisation mechanisms. It can be captured as constraints on the sending & receiving synchronisation as well as actions on the message buffers and critical sections, such as in terms of parameter invariants, state-machines, control and data-flows, detailed logical and mathematical operations.
- **Execution behaviour** – related to the execution of the computation and communication behaviours on a simulation or implementation platform. It corresponds to the models of computation and models of communication, specifying the simulation semantics, or alternatively the data precedence, the function triggering and scheduling schemes, specifying the implementation semantics. For example, the execution of transfer functions of functional blocks can be run-to-completion, implying that a new input to a function will be processed only when the earlier is processed.

For behaviour description and related analysis support of embedded systems, several frameworks are of interest, such as SPEEDS, Ptolemy, MILAN and Metropolis. In the following text, we summarise a few relevant behaviour formalisms and their analysis leverages.

Timed Automata:

A formalism used for modelling and verification of real-time systems. It is basically a finite automaton with real-valued variables [2]. The variables are used to model clocks. Clock constraints or guards are used to restrict the behaviour of the automata. Mathematically, a timed automaton A is a tuple $\langle L, l_o, C, A, E, I \rangle$ where,

- L is finite set of finite locations or nodes (e.g. states or a discrete system, modes etc.)
- $l_o \in L$ is the initial location
- C is the set of clocks
- A is the set of actions, co-actions and the internal τ -action
- $E \subseteq L \times A \times B(C) \times 2^C \times L$ is the set of edges (transitions) between locations with an action, a guard and a set of clocks to be reset.
- $I : L \rightarrow B(C)$ assigns invariants to locations

For most of the cases timed automata are often composed of network of timed automata which can communicate with each other with the help of synchronization channels. Timed automata can be used for analyzing safety, reachability, liveness and equivalence between different components. A timed automata can be used at all levels of abstraction especially Analysis and Design levels in terms of EAST-ADL.

Hybrid Automata:

A hybrid automaton is a system representation consisting of both transition systems with continuous dynamics. Mathematically, a hybrid automaton H is a tuple

$$H = (Q, V, f, Init, Inv, \Theta, G, R, \Sigma, \lambda)$$

where:

- $Q = \{q_1, \dots, q_k\}$ is a finite set of discrete states (modes, control locations etc.);
- $V = \{x_1, \dots, x_n\}$ is a finite set of continuous variables (relate to dynamics modelled by laws of physics);
- $f : Q \times \mathfrak{R}^n \rightarrow \mathfrak{R}^n$ is an activity function;
- $Init \subset Q \times \mathfrak{R}^n$ is the set of initial states;
- $Inv : Q \rightarrow 2^{\mathfrak{R}^n}$ describe the invariants of the locations;
- $\Theta \subseteq Q \times Q$ is the transition relation;
- $G : \Theta \rightarrow 2^{\mathfrak{R}^n}$ is the guard condition;
- $R : \Theta \rightarrow 2^{\mathfrak{R}^n} \times 2^{\mathfrak{R}^n}$ is the reset map;
- Σ is a finite set of synchronisation labels;
- $\lambda : \Theta \rightarrow \Sigma$ is the labelling function.

The automaton H describes a set of (hybrid) states $(q, x) \in H = Q \times \mathfrak{R}^n$.

A hybrid automaton is one of the most common formalisms for modelling hybrid systems. A few of the examples include aircraft control, control of vehicular caravan etc. For embedded systems this formalism can be used to specify the complete behaviour. In addition, tools such as Matlab, Modelica, and Ptolemy can be used to perform analysis of systems whose behaviour can be represented by a hybrid automaton.

Discrete Event System Specifications (DEVS):

This formalism was introduced by Zeigler [3] for specification, modelling and simulation of discrete event systems. The main features of DEVS are set-theory as the basis, system-theoretic representation and hierarchical and modular representation. A DEVS model can be atomic (non-decomposable) or coupled (a collection of atomic or coupled model). Mathematically, an atomic DEVS AM is defined as follows:

$$AM = \langle X, S, Y, \delta_{ext}, \delta_{int}, \lambda, ta \rangle$$

where:

- X is a set of input events;
- S is a set of sequential states;
- Y is a set of output events;

with the following constraints:

- $ta : S \rightarrow R$ (non-negative real number), time advance function.
- $\delta_{int} : Q \rightarrow Q$, an internal transition function;

where $Q = S \times R = \{(s, e) \mid s \in S \text{ and } 0 \leq e \leq ta(s)\}$ states set of M ;

- $\lambda : Q \rightarrow Y$, an output function
- $\delta_{ext} : Q \times X \rightarrow Q$, an external transition function.

With the help of DEVS, it is possible to verify a desired behaviour and evaluate the performance of a given system. The behavioural properties which can be verified by DEVS are liveness, safeness, desired states / event sequence etc. With the help of the concept of system morphisms, DEVS can also be used for representing continuous time dynamics.

3.2.2 EAST-ADL Behaviour Annex

An update proposal with the objective to allow explicit declarations of key behaviour properties with EAST-ADL and thereby to constrain behaviour definitions in external models and formalisms (e.g., Simulink, UML, etc) has been proposed in ATESS2. The provision of such behavioural declarations in EAST-ADL is also considered necessary to support the formalisation of textual requirements as well as for the reasoning of operational behaviours of vehicle features and environmental situations in early development stages. The behaviour modelling extension is given as a set of behaviour constraints. See Figure 12 for an overview of these constructs and their relations to other EAST-ADL definitions.

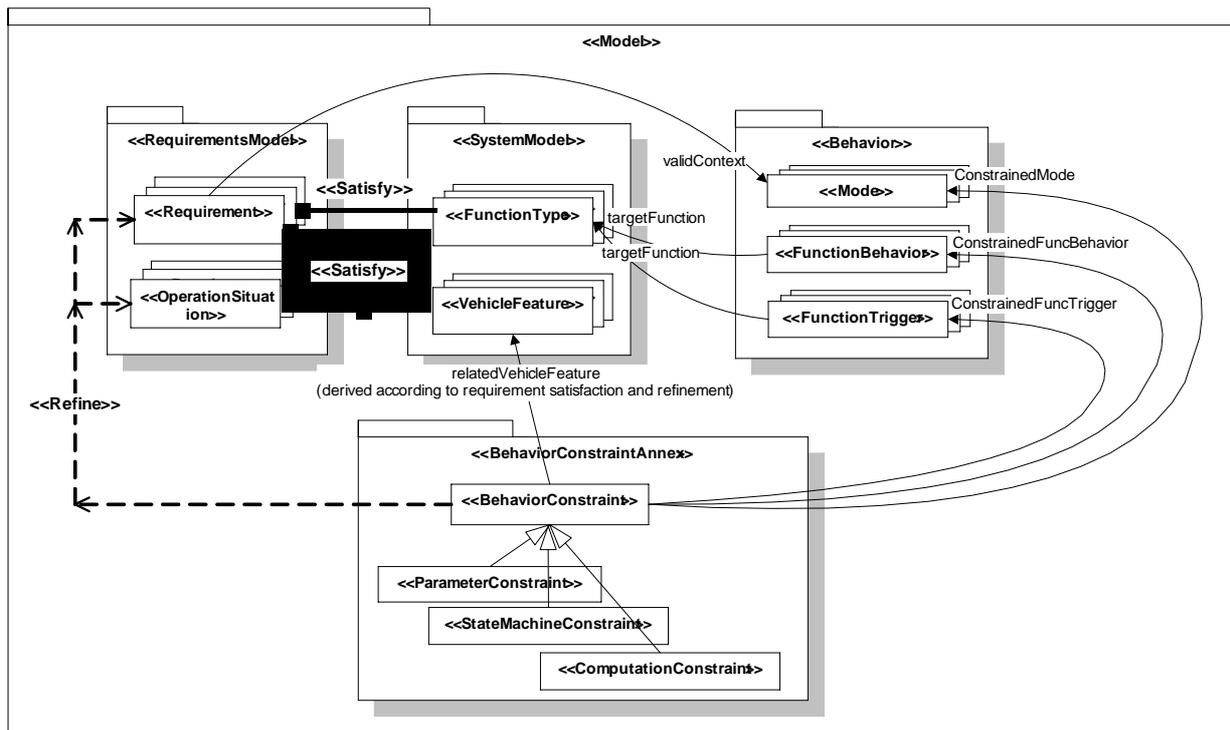


Figure 12 - An overview the domain model update proposal for extended behaviour modelling support.

The behaviour constraints are managed within the behaviour constraint annex module. They are integrated into the EAST-ADL metamodel in a similar way as timing, dependability, and other non-functional constraints. This means that the behaviour constraints target directly `Behavior::Mode`, `Behavior::FunctionTrigger`, `Behavior::FunctionBehavior`, and thereby any external behaviour models linked by these EAST-ADL definitions. Through the references from function behaviours and function triggers to their targets, the system artefacts under constraint can be derived. Depending on the system artefacts under constraint, a behaviour constraint can be applied either for *FAA* (Functional Analysis Architecture), or for *FDA* (Functional Design Architecture), or for environment specification.

For functional requirements and operation situations, the introduction of behaviour constraints allows the related textual descriptions in the requirement model to be refined through the existing requirement *refine* relationship in EAST-ADL as shown in Figure 12. The refinements in terms of behaviour constraints provide declarations of behaviour elements, such as the expected parameters, states and transitions, of one or multiple functional requirements. This facilitates the analysis and validation of requirements in regard to the semantics and consistency of textual descriptions, as well as the reuse and management of requirements and their implied behaviours.

In EAST-ADL, the definitions of system end-to-end functionalities are supported by the `VehicleFeature` constructs. Such features *satisfy* some functional requirements and are in turn *realized* by some system artifacts in the *FAA* (Functional Analysis Architecture). The definitions of such end-to-end functionalities are currently textual. There are no links from such functionality definitions to the definitions of modes and function behaviours. With the behaviour constraints, the behavioural restrictions assigned to such end-to-end functionalities are derived through the *Satisfy* and *Refine* links. In other words, a vehicle feature satisfying a particular functional requirement normally also possesses all behaviours refining that functional requirement.

Current EAST-ADL supports the definitions of error behaviours and anomalies through external logics. The behaviour constraints can be used to provide explicit language support for defining the

semantics of an anomaly (i.e., fault condition), and fine-grained integration error behaviours to nominal state machine and computation behaviours.

The readers are referred to ATESS2_Deliverable_D2.1_A3.4 for more details about the extension. Current work in MAENAD aims to develop a behaviour core with hybrid automata semantics that can be extended to EAST-ADL native representation of models based on external formalisms for

1. formalising requirement and constraint statements, attribute relations, behavioural constraints, and
2. defining mapping rules and algorithm to transform the behaviour model to various analysis tools.

This is illustrated in Figure 13. The related approaches are discussed in the following section.

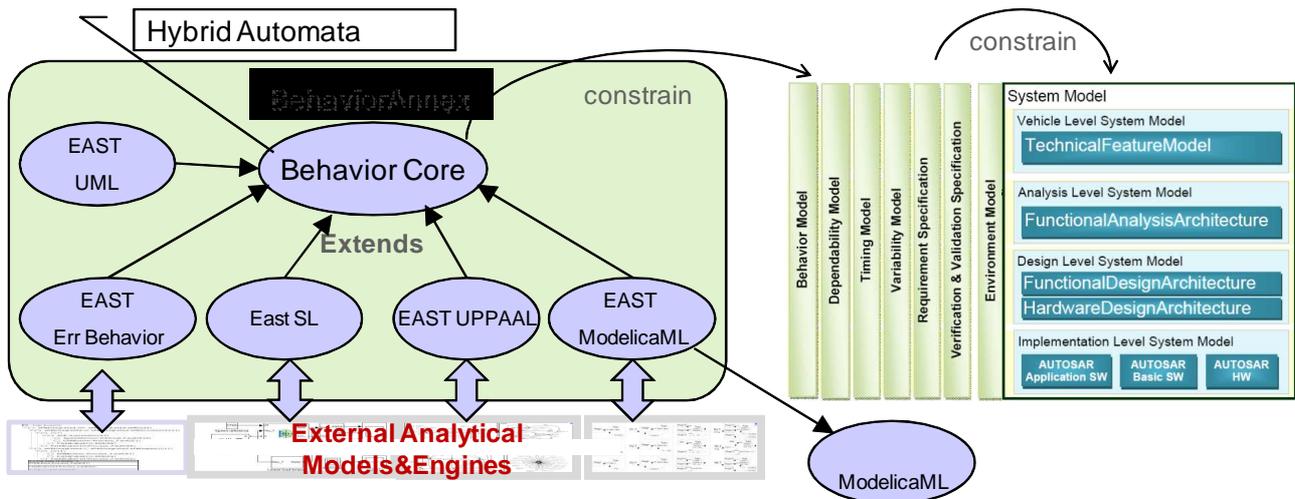


Figure 13 - Proposed approach in MEANAD to language enhancement for behaviour centric analysis.

3.2.3 Related Approaches

Unified Modelling Language (UML)

UML is a general purpose language created by the OMG to support model-driven architecture development. The language was originally developed as an abstraction language for representing software systems but was later adopted for other systems, such as enterprise systems and embedded systems. This was made possible due to the extension possibilities that stem from using profiling mechanisms.

UML provides a graphical representation of different aspects of a system. This representation is categorised into three types: structure, behaviour, and interactions. Different kinds of diagrams can be used at different development phases and for different analysis. The following briefly describes this usage:

- Use-case diagrams describe the features (end-to-end functionality) of a system. A feature may include actors, their goals (represented as use cases), and any dependencies among those use cases. These diagrams are useful at an early stage of a development process.
- State machine diagrams are used to represent different states, the transitions between different states, any guards to prevent an unspecified transition, as well as the actions performed during the transition or in a particular state. This diagram can be used to specify and analyse behaviour at all the higher levels of abstractions.

- Activity diagrams are used to represent the workflow of a system or its components. More specifically it is useful for specifying and analysing control flow at different levels of abstractions.
- Interaction diagrams are used to specify data flow between different architectural entities.
- Composite structure diagrams provide an overall structure view of a system.
- Class diagrams show different classes in a system, their attributes, and any relationship between the different classes. These kind of diagrams can either be used for information specification or (with some additional support) used to generate an executable implementation together with the behavioural properties from state machines of activity diagrams.

Matlab/Simulink

Matlab and Simulink are widely used tools for embedded systems development provided by MathWorks®. The two tools and their libraries (such as Stateflow, SimEvents, and Control System Toolbox) enable users to model, analyse, and simulate a large variety of systems ranging from discrete event, discrete time, continuous time, and hybrid systems. Due to its strong modelling and analysis support, Simulink can be used to model and simulate system behaviour under different conditions and hence is a strong candidate for FEV (Full Electric Vehicle) analysis.

A prototype plug-in to transfer information to and from EAST-ADL already exists. Just like Uppaal (see below), direct or indirect support (via some mapping or tool extension) for a timed automaton is required from the behaviour modelling part of EAST-ADL. With direct support it would be possible to specify Simulink functions in EAST-ADL and with indirect support, this is possible with an extension of the current plug-in.

Modelica

Modelica is a modelling language that allows specification of mathematical models for the purpose of computer simulation of dynamic systems. The Modelica language is a “free object-oriented modeling language with a textual definition to describe physical systems in a convenient way by differential, algebraic and discrete equations” [modelica.org]. One key feature of the Modelica language is that it enables physical modelling, i.e., modelling the topology of a system with physical “acausal” connectors between components, using the built-in connect statement. Using the connect statement, simulation components can be connected like the real components would be connected physically, which makes it intuitive to model these systems. Modelica’s object-oriented software component model, which includes e.g. inheritance, allows definition of reusable components. Hence, a large library of open-source components, the Modelica standard library, has been developed by the Modelica community.

Current ModelicaML research activities at the EADS IW and Linköping University aim at converging the UML/SysML and Modelica in order to take advantage of both: the descriptive power of the UML/SysML graphical notations for software and systems modelling and the power of the Modelica language of simulating time-continuous as well as time-discrete/event-based system behaviour.

Ptolemy

Ptolemy is a framework and simulation environment developed by the Ptolemy Project at Berkeley to support the modelling and simulating of hybrid embedded systems. The focus is on application level design with heterogeneous Models of Computation (MoCs) on the simulation.

Ptolemy employs a unified semantics, referred to as *TSM* (Tagged Signal Model), to interpret or encode other specific models of computations, such as synchronous languages, synchronous dataflow process networks¹, discrete-event systems, etc. *TSM* is an *abstract formalism* (also called a “meta model”) proposed by Lee and Sangiovanni Vincentelli. According to *TSM*, the behaviour of a system is determined by a set of *signals*, i.e., sequences of *events* exchanged among its processes. The timing and ordering of events from concurrent processes is captured through *tags*. Each event is a *value-tag* pair. Tags can form partially ordered or totally ordered sets, depending on the firing rules of processes for event producing and consuming and the communication protocols for event communication. For example, the tags of events from discrete-time models are totally ordered, while the tags of events from discrete-event models only share the same tag if the events are synchronised. Different MoCs are distinguished by the differences in their firing rules and communication protocols.

Besides this abstract semantics, Ptolemy also has an operational semantics underlying the Ptolemy II simulator. Each model is a hierarchical composition of components/processes that are connected with communication channels via component ports. For each MoC of components/processes, there is a *director* controlling its behaviour. The behaviour of a communication channel is controlled by its *receiver*, such as when components of two separate MoCs are connected.

SPIN

The SPIN (Simple PROMELA Interpreter) model checker is a popular and very efficient modelling and computation tool for logic model checking, which aims to find the logic bugs in a simplified process model through characterising all possible executions, rather than a small subset explored by simulation and testing. As the name suggests, logic model checking does not address real-time properties, but considers only the temporal sequence of actions. It is very helpful for designing concurrent software, because the number of interactions is often exponential in the number of processes and traditional means of debugging and testing cannot effectively explore the huge possibilities.

SPIN works on the process model written in PROMELA (Process Meta-Language) language. A SPIN model consists of asynchronous processes connected through global variables and message channels. A message channel is either a synchronisation point between two processes or a message queue. For efficient computation, SPIN converts the PROMELA model to Büchi automata, i.e., the automata with finite states but accepting infinite transition sequences. The correctness claims to be verified by SPIN can be invariant assertions, temporal logic claims, or Büchi automata. They can be either about safeness, i.e., something bad can never happen, or liveness, i.e., something good will eventually happen.

PROMELA defines precise execution semantics on process execution flow and communication. Of particular importance is the rich syntax and semantics for process communication. Compared to the EAST-ADL framework, PROMELA semantics belong to the Functional Design Architecture and encompass the limited behaviour semantics on component communication. Referring to the PROMELA semantics, this report will propose an extension to the behavioural modelling in EAST-ADL. Consequently, the EAST-ADL model will be applicable to a larger scope of software systems

¹ A synchronous dataflow process network is a special type of untimed discrete-event model for signal processing applications characterised by data insensitive computation and transformation. It extends the Kahn process network characterised by process communication via FIFO channels with firing rules to specify the ordering of communication activities by a component. The Kahn process network is used for example for modelling multi-media application based on streaming.

that require complicated communication protocols. Moreover, the alignment with SPIN also allows one to verify and improve the EAST-ADL model at an earlier design time.

SPEEDS

SPEEDS (**Speculative and Exploratory Design in Systems Engineering**) is a European FP6 project aiming at providing tool integration, support for modelling of complex embedded systems using heterogeneous subsystem models, usage of formal analysis tools in architecture modelling, and component-based development of embedded systems. SPEEDS is based around the assumption on usage of *model-based design*, in its broadest meaning. The vision includes seamless usage of models throughout the design process for several purposes, including specification, simulation, and formal analysis. Issues such as interchange of models, version and configuration management are also approached.

To tackle the issues implied by these concerns, SPEEDS supplies an appropriate *component model* – called HRC (heterogeneous rich components) – which is supposed to be expressive enough to cover the entire development cycle and also addresses both functional and extra-functional aspects of the development. HRC is based on a hybrid automata variant (the combination of a finite automaton, using discrete states, with additional time-continuous state variables that may be used as e.g. a basis for guard conditions for transitions). In addition to using this basis, a formally founded model for contracts based on the notion of assume/guarantee relations to the environment is deployed, making decoupling between different component instances possible.

Uppaal

Uppaal is a tool developed for modelling, simulation, verification and validation of real-time systems behaviour. Based on the concepts of timed automata, the tool is supported by its own description language, a simulator and a model checker. Some of the industrial use cases where the tool has been used include model-based testing of WAP gateways and formal modelling and analysis of audio/video protocols. The tool is well suited for testing and analysis of safety critical systems. It has already been used for systematic evaluation of fault trees and the analysis can also be used for design optimisation. This makes it a prospective candidate for MAENAD.

In order use the tool for analysis, direct or indirect support (via some mapping or tool extension) for a timed automaton is required from the behaviour modelling language.

AADL

AADL is an architecture description language for embedded and safety critical control systems. The approach has its focus on software and hardware integration. It provides modelling support for software, platform, and implementation design, but does not have inherent language support for other system aspects such as for environment, requirements, and functions. Instead, the AADL relies on external domain specific CASE tools such as Matlab/Simulink and ControlH to cover such aspects.

Behaviours in AADL are fundamentally imperative, i.e., based on control flow description and source code modules (e.g., in Ada). The concurrency and timing aspect is supported by multitasking (e.g., preemptive RM scheme) on a platform. AADL naturally specifies only the non-deterministic behaviour of a thread as a set of subprogram calls. For sophisticated state transition and communication, the behavioural annex is introduced to allow behaviour modelling using the non-deterministic hierarchical automaton.

For safety analysis and reliability assessment, the Error Model Annex provides standard annotation for possible errors of the components and their connections in the AADL model. The

error models of low-level components typically capture the results of failure modes and effects analysis (FMEA). The system error model is a composition of the error models of its components, derived from the system hierarchy, the interaction error models, and the shared computing platform resources. The system error model reflects error propagation between components based on the component dependencies in the AADL architecture model. It typically captures the results of system hazard analysis (HA). Moreover, the Error Model Annex can also model the risk mitigation methods that improve safety and reliability by specifying how components detect and mitigate errors.

The resulting annotated AADL model can then be used as an input to various dependability analysis techniques for fault forecasting during different phases of the development cycle. It also enables an integrated approach that ensures consistency and completeness between hazard analysis and FMEA and the safety and reliability analyses that associate them together.

For safety design, AADL/MetaH provides specific language constructs for specifying safety levels, criticalities, fault tolerance, and error handling. There is also a separate error model describing errors and their propagations.

3.3 Modelling support for dependability analysis

3.3.1 Dependability and error modelling in EAST-ADL

As part of its support for dependability engineering, EAST-ADL allows the modelling of system error behaviours for safety analysis (e.g. FTA) and allows safety constraints to reference the faults/failures under consideration. Figure 14 shows the related key language concepts, which are explained below.

An *ErrorModel* is a constrained model providing the analytical information about the error logic of a target component/system, i.e., which faults may occur and how they propagate. Associations to the nominal architecture identify the target components/systems and the modelling contexts. While the fault ports (*FaultInPort*) declare which incoming faults the target can receive from its environment, the failure ports (*FailureOutPort*) declare which local failures may occur and propagate to the environment. A *PropagationLink* connects an output port with an input port for error propagation across different target components/systems. Such analytical ports and links can have references to the corresponding nominal ports and connectors (e.g. the data flow ports and connectors) when a fine-grained traceability is needed.

An *InternalFault* represents an internal condition of the target component/system, which can cause errors when become activated. A *ProcessFault* refers to a systematic fault introduced due to design or implementation flaws (e.g. incorrect requirements or buffer size configuration). Because any of the output failures of an error model may be caused by process flaws, a process fault affects all failures that the component may exhibit.

By means of the *FaultFailure* construct, the actual value of an anomaly can be declared (i.e. one of the possible values of an incoming fault, an internal fault, a process fault, or an outgoing failure) based on an enumeration type, such as *{omission, commission}* or *{brake-loss, brake-fade, brake-too-much}*. Within an error model, the syntax and/or semantics of existing external formalisms can be adopted for a precise description of the *FailureLogic*, capturing which output failures are caused by which input and internal faults within the target component/system. This, together with the error propagation links, makes it possible to perform safety analyses and derive the related safety constraints.

A *safety constraint* defines the qualitative bounds on the fault and failure values in terms of the safety integrity level (ASIL) for avoiding its occurrence (as defined in ISO 26262). For an incoming fault or an internal fault, a safety constraint means that the error model assumes this integrity level on the fault. For process faults and output failures, it means that the error model propagates failures at this integrity level.

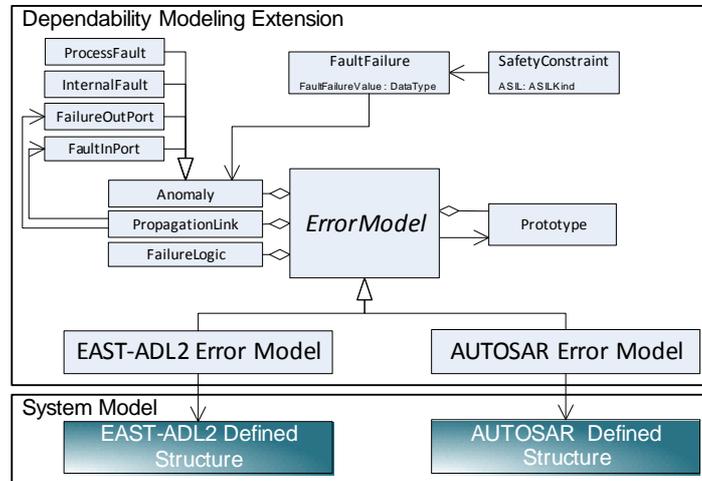


Figure 14 - Error modelling in EAST-ADL

3.3.2 Safety Analysis with HiP-HOPS Technology

Hierarchically Performed Hazard Origin & Propagation Studies, or "HiP-HOPS", is a comprehensive safety analysis methodology with a great degree of automation. In this technique, fault trees and FMEA are automatically constructed from topological models of the system that have been augmented with appropriate component failure data [4].

3.3.2.1 Overview of the methodology

A HiP-HOPS study of a system under design has four phases:

- System modelling and failure annotation
- Fault Tree synthesis
- Fault Tree analysis & FMEA synthesis
- Optimisation

The first phase is executed manually while the later three phases are fully automated.

The first phase consists of developing a model of the system (hydraulic, electrical or electronic, mechanical systems, and conceptual block and data flow diagrams) and then annotating components in that model with failure data. In this phase, modelling can be carried out in a modelling tool like Matlab Simulink as it would normally have been carried out for the purposes of simulation. Failure annotations are then added to components of the model.

The second phase is the fault tree synthesis process. In this phase, an automated algorithm is applied to the annotated system model to create a set of fault trees that define system failures and their causes in the architecture. The algorithm works by taking failures of system outputs and progressively working backwards through the model to determine which components caused those failures. System failures and component failures are then joined together using the appropriate logical operators to construct fault trees with the failures at the system outputs as the top events and the root causes as basic events. The concept is illustrated in Figure 15.

In the third phase, fault trees are analysed and an FMEA is constructed which combines all the information stored in individual fault trees. The FMEA is presented in the form of a table that lists, for each component, the effects of each component failure on the system. As part of this process, both qualitative (logical) and quantitative (numerical-probabilistic) analyses are carried out on the fault trees. These analyses provide both the minimal cut sets of each fault tree and the unavailability (i.e. failure probability) of top events.

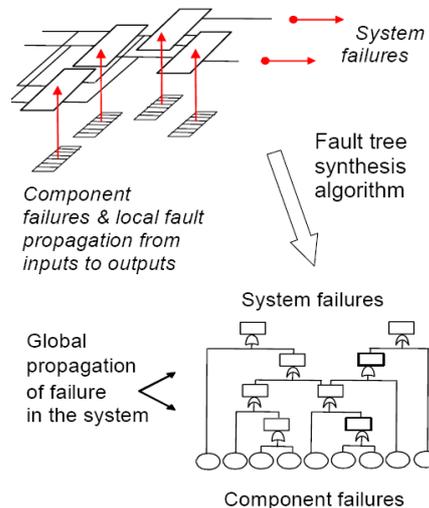


Figure 15 - The synthesis of fault trees from the system model

If the system is to be optimised, then the fourth phase is applied. This is an iterative optimisation of the system model, using genetic algorithms to continually evolve the model to find the best configuration of components on the basis of safety and cost.

3.3.2.2 Modelling Phase

HiP-HOP studies can be performed on any model of a system that identifies components and the material, energy or data transactions among components. Such models can be hierarchically arranged, to manage complexity, if necessary. The basic idea of HiP-HOP is that an output failure of a component can either be caused by an input failure, an internal failure, or some combination of both. The local component output deviations and topology information are used to determine the relation between local deviations and top events.

For the purpose of the analysis, each component in the model must have its own local failure data, which describes how the component itself fails and how it responds to failures propagated by other components in the vicinity. Essentially, this information specifies the local effects that internally generated or propagated failures have on the component's outputs. This is achieved by annotating the model with a set of failure expressions showing how deviations in the component outputs (output deviations) can be caused either by internal failures of that component or corresponding deviations in the component's inputs. Such deviations include unexpected omission of output or unintended commission of output, or more subtle failures such as incorrect output values or the output being too early or late. This logical information explains all possible deviations of all outputs of a component, and so provides a description of how that component fails and reacts to failures elsewhere. At the same time, numerical data can be entered for the component, detailing the probability of internal failures occurring and the severity of output deviations. This data will then be used during the analysis phase to arrive at a figure for the unavailability of each top event. Once done, the component can then be stored together with the failure data in a library, so that other

components of the same type can use the same failure data or this type of component can be re-used in other models with the same failure data. This avoids the designer having to enter the same information many times.

For the specification of the components' failure modes (which are the effects by which the component failures are observed), a generic and abstract language was developed. There are different ways of classifying failure modes, e.g. by relating them to the function of the component, or by classifying according to the degree of failure – complete, partial, intermittent etc. In general, however, the failure of a component will have adverse local effects on the outputs of the component, which, in turn, may cause further effects travelling through the system on material, energy or data exchanged with other components. Therefore in HiP-HOPS, effects are generally classified into one of three main classes of failure, all equally applicable to material, energy or data outputs. These are: the omission of an output, i.e. the failure to provide the output; a commission of an output, i.e. a condition in which the output is provided inadvertently and in the wrong context of operation; and an output malfunction, a general condition in which the output is provided but at a level which deviates from the design intention. Since this classification adopts a functional viewpoint which is independent of technology, it could provide a common basis for describing component failures and their local effects. However, HiP-HOPS can work with any classification of failure modes as long as it is consistent from one component to the next.

Components do not only cause failures; they also detect and respond to failures caused by other components or transfer failures of other components. In addition to generating, mitigating, or propagating failures, they may also transform input failures to different types of output failure. For instance, a controller may be designed to respond to detected sensor failures by omitting any further output to ensure that hazardous control action is avoided. In this case, malfunctions at the input are intentionally transformed into omission failures. To capture those general patterns of behaviour of a component in the failure domain, manual analysis is performed at component level and focuses on the output ports through which a component provides services to other components in the system. In the course of the analysis, each output port is systematically examined for potential deviations of parameters of the port from the intended normal behaviour, which generally fall into the following three classes of failure:

- (a) Omission: failure to provide the intended output at the given port
- (b) Commission: unintended provision of output
- (c) Malfunction: output provided, but not according to design intention.

Within the general class of malfunctions, analysts may decide to examine more specific deviations of the given output which, in most applications, will include conditions such as the output being delivered at a higher or lower level, or earlier or later than expected. As an example, Figure 16 shows the analysis of a two-way computer controlled valve. The figure shows the valve as it would typically be illustrated in a plant diagram and records the results of the local safety analysis of the component in two tables that define valve malfunctions and output deviations respectively.

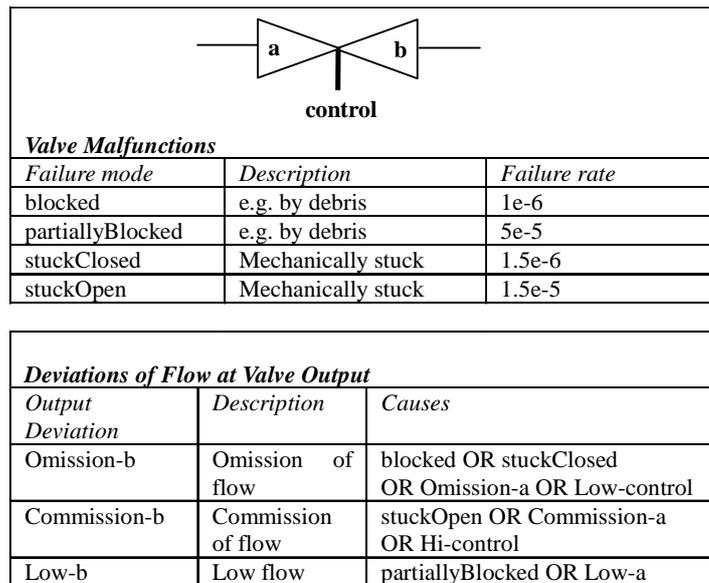


Figure 16 - Failure annotations of a computer-operated two-way valve

In normal operation, the valve is normally closed and opens only when the computer control signal has a continuously maintained value of a logical one. Valve malfunctions include mechanical failures such as the valve being stuckOpen or stuckClosed, and blockages caused by debris such as blocked and partiallyBlocked. For each malfunction, the analysis records an estimated failure rate while the effects of those malfunctions on the output of the valve can be seen in a second table that lists output deviations.

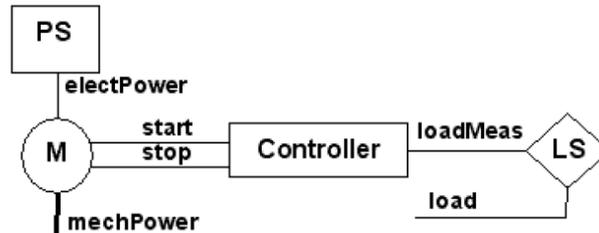
This specification of failure modes is generic in the sense that it does not contain references to the context within which the valve operates. Failure expressions make references only to component malfunctions and input/output ports of the component. The failure behaviour described in these expressions has been derived assuming a simple operation that the component is expected to perform in every application (valve is normally closed unless the value of control signal is 1). For these reasons, the specification of Figure 16 provides a template that could be re-used in different models and contexts of operation, perhaps with some modifications, e.g. on failure rates, to reflect a different environment.

3.3.2.3 Synthesis Phase

As seen in Figure 16, component failure data relate output deviations to logical expressions that describe the causes of those deviations as component malfunctions and deviations of the component inputs. Each such expression is effectively a mini fault tree that links a top event (the output deviation) to leaf nodes, some of which may represent input deviations. When a component is examined out of system context, input and output deviations represent only potential conditions of failure. However, when the component is placed in a model of a system, the input deviations specified in the analysis can actually be triggered by other components further upstream in the model and the specified output deviations can similarly cause more failures further downstream.

This mechanism by which output failures of a particular class at one end of a connection trigger input failures of the same class at the other end results in a global propagation of failures through the system which may ultimately cause significant hazardous failures at the outputs of the system. Given a model of the system and the local safety analyses of its components, it is possible to capture this global propagation of failure in a set of fault trees. These fault trees are mechanically constructed by traversing the model and by following the propagation of failure backwards from the final elements of the design (e.g. electromechanical actuators) towards the system inputs (e.g.

material/energy resources, operators and data sensors). The fault tree is generated incrementally, as the local safety analyses of the components are encountered during the traversal, by progressively substituting the input deviations for each component with the corresponding output failures propagated by other components. Figure 17 illustrates the principle that underpins this process of fault tree synthesis. The figure shows a hypothetical motor and its starter circuit as a unit (M) that transforms electrical power provided by a power supply (PS) to mechanical power on a rotating axis.



<i>Component</i>	<i>Output Deviation</i>	<i>Description</i>	<i>Causes</i>
M	OE-mechPower	Omission of mechanical power	motorFailed OR OE-electPower OR OS-start OR CS-stop
PS	OE-electPower	Omission of electrical power	powerSupplyFailed
Controller	OS-start	Omission of start signal	controllerFailed
	CS-stop	Commission of stop signal	elecMagnInterference OR HV-loadMeas
LS	HV-loadMeas	Hi value in load measurement	sensorBiased

Figure 17 - Example system and fragments of local safety analyses

The motor starter receives normal start and stop commands from a control computer (Controller). As a safety feature the controller is also connected to a sensor (LS) that monitors the load connected to the axis of the motor and when the measurement exceeds a specified load it issues a stop command to protect the motor. To illustrate the fault tree synthesis, the figure also provides a table that contains, for simplicity, only fragments from the local analyses of the Motor, the Power Supply, the Controller and the Load Sensor. For simplicity, deviations in this table refer to names of connections in the model rather than local names of component ports. Collectively, these deviations and their causes define the propagation of failures that result in an omission of mechanical power at the output of the motor.

The local analysis of the motor, for example, defines that this event can be caused by a failure of the motor, an omission of electrical power at the input, an omission of the start signal or, interestingly, a commission of the stop signal. The causes of some of those events can in turn be explored further in the local analyses of the components connected to the motor. For example, the analysis of the power supply defines that a failure of this component will cause an omission of electrical power. In turn, the analysis of the controller defines that an omission of the start signal will be caused by a controller failure while a commission of the stop signal can be caused by electromagnetic interference or in response to an abnormally high measurement of motor load. The analysis of the load sensor defines that the latter is indeed a plausible failure mode that can be caused in normal conditions of loading if the sensor is positively biased.

An overall view of the global propagation of failure in the system can be automatically captured by traversing the model and by following the causal links specified in the local safety analyses of the components that are progressively encountered during this traversal. The result of this process for the above example is the fault tree that is illustrated in Figure 18. Note that this mechanically synthesised fault tree records the propagation of failure in a very strict and methodical way. It starts from an output failure, the omission of mechanical power, and following dependencies between components in the model it systematically records other component failures that progressively contribute to this event. The logical structure of the tree is determined only by interconnections between the components and the local analyses of those components. This

logical structure is straightforward and can be easily understood, unlike the structure of many manually constructed fault trees, which is often defined by implicit assumptions made by analysts. Note that although in this example the tree only incorporates OR gates, other logical symbols such as AND and priority AND gates would appear in the tree structure if such relationships were originally specified in some of the local analyses.

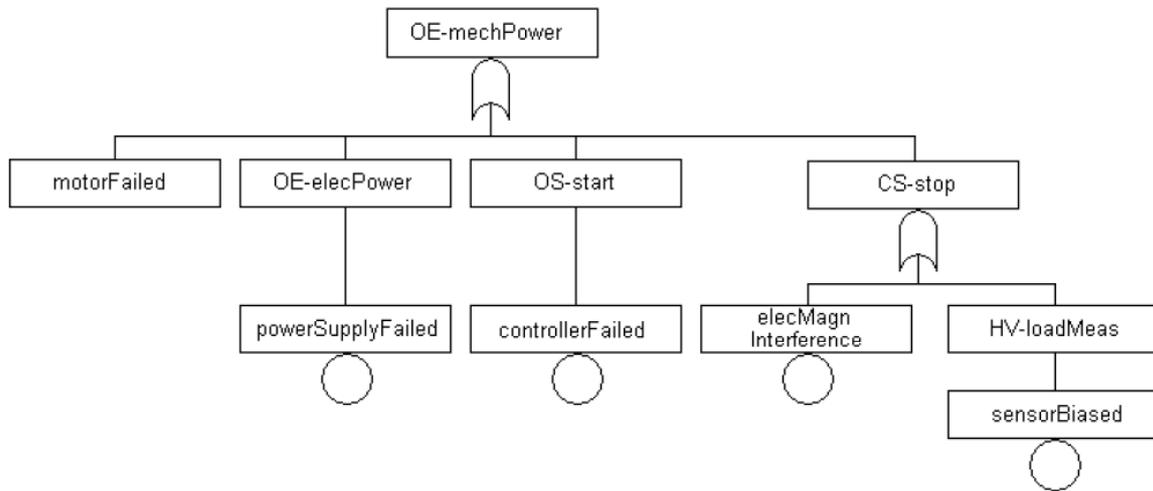


Figure 18 - Mechanically constructed fault tree for the example system

3.3.2.4 Analysis Phase

In the final phase, the synthesised system fault trees are analysed, both qualitatively and quantitatively, and from these results the FMEA is created. Firstly, the fault trees undergo qualitative analysis to obtain their minimal cut sets, which reduces them in size and complexity. This is typically done using a mixture of classical logical reduction techniques, which usually means applying logical rules to reduce complex expressions, and more modern techniques developed specifically for HiP-HOPS. Once the minimal cut sets have been obtained, they are analysed quantitatively, which produces unavailability values for the top events of each fault tree.

The last step is to combine all of the data produced into an FMEA, which is a table that concisely illustrates the results. The FMEA shows the direct relationships between component failures and system failures, and so it is possible to see both how a failure for a given component affects everything else in the system and also how likely that failure is. However, a classical FMEA only shows the *direct effects* of single failure modes on the system, but because of the way this FMEA is generated from a series of fault trees, the HiP-HOPS is not restricted in the same way, and the FMEAs produced also show what the *further effects* of a failure mode are; these are the effects that the failure has on the system when it occurs in conjunction with other failure modes. Figure 19 shows this concept.

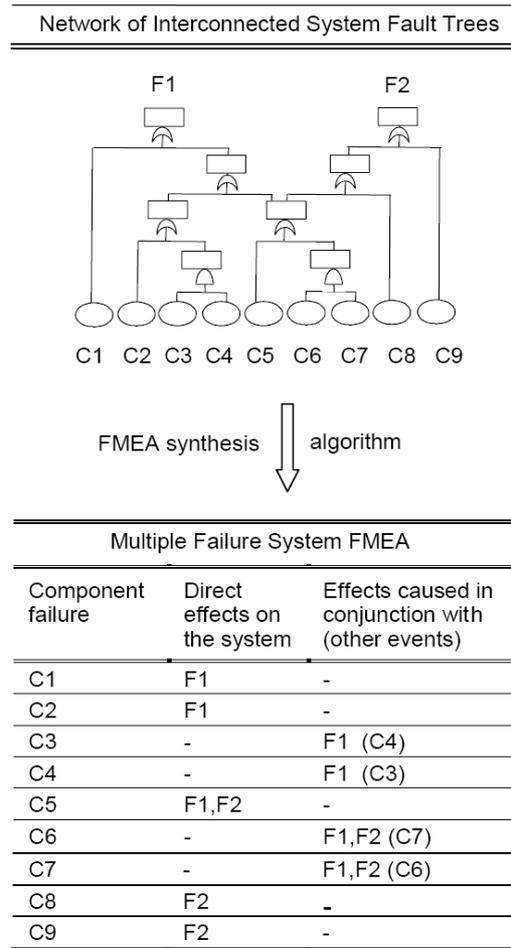


Figure 19 - The conversion of fault trees to FMEA

In Figure 19, F1 and F2 are system failures, and C1 – C9 are component failures. For C3, C4, C6 and C7, there are no direct effects on the system – that is, if only one of these components fail, nothing happens. However, they do have further effects; for example, C3 and C4 both occurring in conjunction will cause F1 to occur. The FMEAs produced, then, show all of the effects on the system, either singly or in combination, of a particular component failure mode. This is especially useful because it allows the designer to identify failure modes that contribute to multiple system failures (e.g. C5 in the example). These common cause failures represent especially vulnerable points in the system, and are prime candidates for redundancy or extra reliable components.

3.3.3 Multi-perspective analysis of EAST-ADL models

EAST-ADL does not limit the modeller to a single architecture, as HiP-HOPS does. Not only are the Error Model and nominal model separate, but EAST-ADL provides multiple layers or levels of potential modelling as well as different views or *perspectives* of a model. HiP-HOPS is expected to be applied to EAST-ADL models at the analysis/artefact and design levels, and at these levels EAST-ADL offers a number of different perspectives. At the analysis level, the primary perspective is the Functional Analysis Architecture (FAA), which is a relatively abstract view of the functions of the system. At the design level, the model is more concrete and is separated into the Functional Design Architecture (FDA) or software perspective and the Hardware Design Architecture (HDA) or hardware perspective. There is also the possibility of a middleware perspective at the design level.

Although HiP-HOPS is designed to be able to support iterative analyses at different levels of abstraction, and thus can cope equally well with an initial qualitative analysis of an abstract functional model and a later quantitative analysis of a more detailed component model, it was only designed to support a single perspective. The concept of separating software and hardware perspectives of the same system is without analogue in HiP-HOPS. This is because HiP-HOPS assumes a compositional model that combines hardware and software, with individual software functions (or subsystems of functions) contained within the hardware processors that run them. Therefore, although HiP-HOPS is capable of analysing both software and hardware, it requires them to be in the same model, with the hardware the primary means of failure propagation: it is not possible for failures in one software function to propagate directly to another remote, physically separated software function; instead software failures must propagate up to their parent hardware component and then propagate along hardware connections to other hardware components that carry the affected functions.

This can be emulated in EAST-ADL by combining everything into a single perspective, but the advantages of having multiple perspectives are then lost. Instead, HiP-HOPS has been extended to provide native support for *multi-perspective analysis*: the analysis of failure behaviour across two or more separate perspectives of the same system (e.g. H/W and S/W). This required the introduction of the concept of a 'perspective' to the HiP-HOPS model hierarchy.

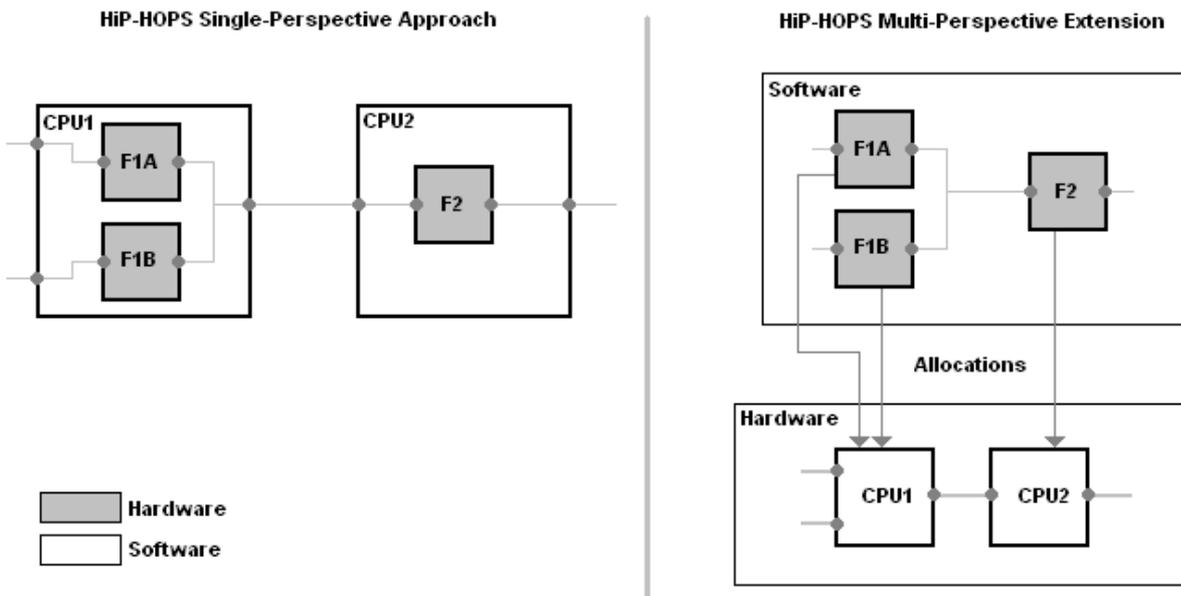


Figure 20 - Multi-Perspective Modelling in HiP-HOPS

In EAST-ADL, hardware and software entities are in separate architectures (HDA and FDA) and communication from one perspective to the other is accomplished by means of *allocation* relationships: software functions are allocated to hardware components that execute them. Failures of the hardware components will propagate to the software functions allocated to them. Failures can also propagate from H/W to H/W and also from S/W to S/W, but this form of connection is restricted to only H/W → S/W propagations. This avoids problems inherent in the multi-perspective modelling paradigm where failures can initiate intractable circular logic loops when propagating through the system (e.g. a S/W fault could cause a H/W failure that in turn is caused by the same S/W fault occurring).

HiP-HOPS has been extended with a similar approach, as shown in Figure 20. Although the original single-perspective approach is still possible, when analysing FDA/HDA models from

EAST-ADL, the new multi-perspective representation can be used. To accomplish this, a new 'Perspective' layer in the model hierarchy was added:

- *Model* (top-level entity representing an entire system under analysis)
- *Perspective* (contains a different perspective of the entire system, e.g. S/W, H/W)
- *Subsystem* (subsystem containing one or more components)
- *Component* (component representing some H/W or S/W entity)
- *Implementation* (the failure behaviour of a component; may also contain a subsystem)

Any number of perspectives can be used, but one perspective must be the 'primary' perspective. Since EAST-ADL assumes that the FDA is the primary perspective, and restricts failures to propagating from H/W to S/W but not from S/W to H/W, it is assumed that an exported EAST-ADL model will use the software perspective as the primary perspective in HiP-HOPS too.

As in EAST-ADL, HiP-HOPS also provides means for failures to propagate from one perspective to another. As in EAST-ADL, it is possible to define 'allocations' of components in one perspective to components in another perspective. Failures can propagate along these allocations by means of 'External Propagation' constructs. In fact, in both EAST-ADL and in HiP-HOPS, it is possible to define *more* than one possible allocation, e.g. using variability constructs in EAST-ADL. This allows scope for potential optimisation of the model on the basis of changing the allocation of software functions to different hardware components, and indeed the optimisation capabilities of HiP-HOPS are also being extended in that direction.

HiP-HOPS provides for certain types of 'local failure data'. Traditionally this has included definitions for Basic Events (representing internal failure modes), Output Deviations (representing the propagation of failures at component outputs), and Common Cause Failures (linked to common model-level failures). Each, once declared, is visible to other definitions in that component, e.g. a basic event can be used in an output deviation of the same component.

Due to multi-perspective analysis, HiP-HOPS now also supports the definition of 'External Propagation' failure data. These are equivalent to Output Deviations, but rather than representing the propagation of failure from the component outputs, they represent the propagation of failure across allocation relationships. Also, unlike other failure data types, External Propagations are visible to allocated components as well as the local component. This means that they can be used in the definition Output Deviations of allocated components.

For example, consider CPU2 and F2 in Figure 20 above. We may define the local failure data of CPU2 to include:

- Basic event "CPU Failure", optionally with appropriate failure rates etc.
- Output Deviation "Omission-output", which is caused by either "Omission-input" or "CPU Failure" (defined above) or "EMI" (defined below).
- Potential Common Cause "EMI", which is caused by electromagnetic interference affecting the entire system.

In addition, because we know that CPU2 can have software functions allocated to it, we define an External Propagation like so:

- External Propagation "Omission", which is caused by "Omission-input", "CPU Failure", or "EMI".

Then when we define the local failure data for F2, which is a software function allocated to CPU2, we can reference this external propagation in F2's output deviation:

- Basic event "Undetected bug"
- Output Deviation "Omission-fout", caused by "Omission-fin" or "Omission". The external propagation is referred to more explicitly as "FromAllocation(Omission)".

Because any component can only be allocated to one component at a time (it has at most one 'current allocation'), we know that "Omission" must refer to an external propagation found in CPU2. HiP-HOPS can therefore link this "Omission" in F2 to the failure logic for "Omission" in CPU2 during its fault tree synthesis process, connecting the perspectives together and allowing for propagation of failure from hardware to software.

Note also that the semantics of common cause failures have been changed too; whereas before common cause failures (CCFs) were globally visible throughout the model, now CCFs are defined per perspective, thus a CCF defined for the hardware perspective (e.g. flooding, fire) would not be accessible from a component in the software perspective, for instance.

HiP-HOPS also allows a more flexible form of connection by means of the 'LocalGoto' and 'GlobalGoto' declarations. These support a more abstract type of connection for situations where there is no form of allocation relationship between two components in different perspectives, but propagation is nevertheless necessary. In these cases, the Gotos act like special forms of input deviations that can connect directly to another output deviation elsewhere in the model. Local gotos only connect to output deviations in the same subsystem (hence 'local') whereas Global gotos can connect to any fully-qualified output deviation in any perspective. For example:

- "O-F2.out = LocalGoto(O-F1A.out)" is valid and states that O-F2.out is caused by O-F1A.out.
- "O-F2.out = LocalGoto(O-CPU1.out)" is invalid as the output deviations are not in the same subsystem.
- "O-F2.out = GlobalGoto(O-CPU1.out)" is valid and states that O-F2.out is caused by O-CPU1.out.

However, these gotos are considered potentially harmful to a coherent analysis and should only be used sparingly, because they break the connection between the propagation of failure and the architecture of the system.

Because these new connections (allocations and gotos) connect existing HiP-HOPS constructs (namely, output deviations), they fit relatively seamlessly into the HiP-HOPS fault tree synthesis process. Once the fault trees are generated, then they can be analysed as normal without any further distinction between one perspective and another.

3.3.4 Tool specification and exchange formats

The EAST-ADL Safety Analysis Plugin is a tool providing the means to perform safety analysis on systems modelled using the EAST-ADL language. In EAST-ADL, the user explicitly models the propagation of errors in an 'error model'. The error model can be created easily because it is built

from local error information. For system safety, however, the global effects of the error propagations are of interest. Tools such as HiP-HOPS can calculate the global effects through safety analyses like FMEAs and FTAs by combining the local error information. To provide a view of the global failure effects to EAST-ADL users without additional effort on the users' side, we provide an automated link between EAST-ADL and HiP-HOPS.

3.3.5 Tool Integration

To establish the link between EAST-ADL design tools such as Papyrus and safety analysis tools such as HiP-HOPS, they have to be integrated. Integration is described by Wassermann [5] to have five different aspects:

- *control integration*: programs can interoperate;
- *data integration*: programs can use each others data;
- *presentation integration*: programs have a unified GUI;
- *platform integration*: services provided by platform;
- *process integration*: the software development processes are integrated.

The plugin mainly performs three forms of integration: data, control and presentation integration. All of these forms are realised in the form of a plugin for the Eclipse framework.

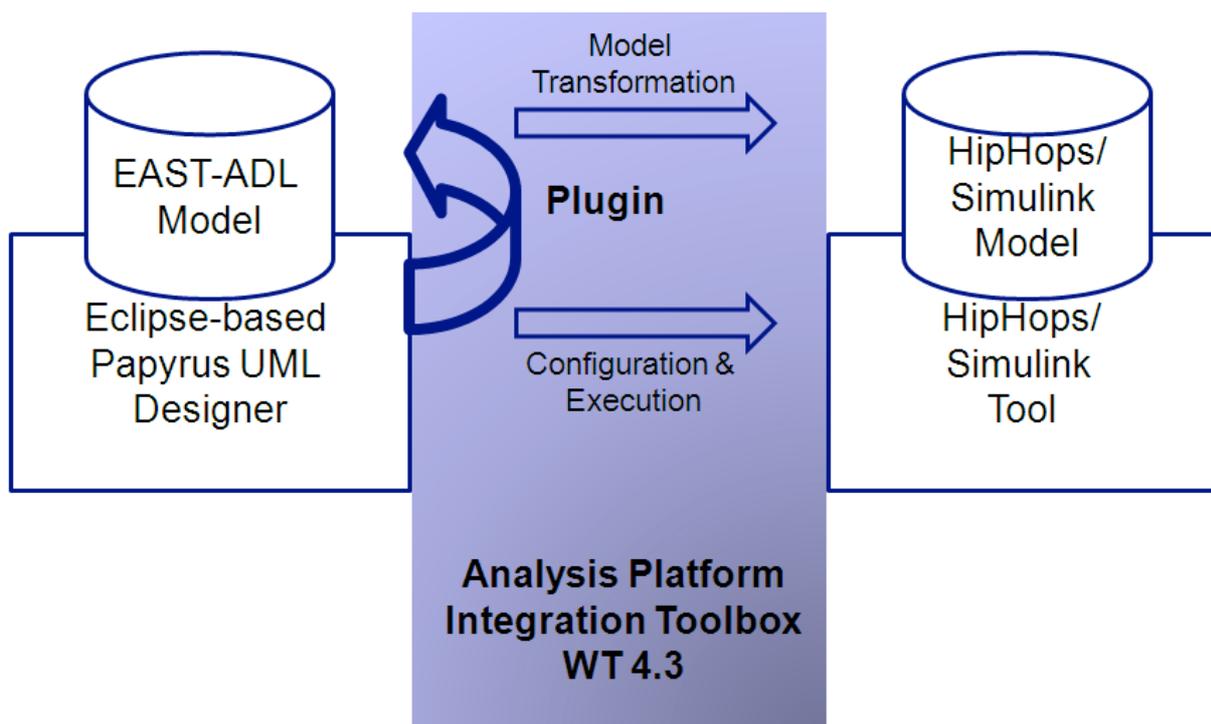


Figure 21 - Integration between EAST-ADL and HiP-HOPS

3.3.5.1 Control and Presentation Integration

This link is provided in the form of a plugin. This ensures seamless integration in the modelling environment (Papyrus/Eclipse) and keeps the safety analysis overhead experienced by the user as low as possible and thus allows for an iterative development process.

Without tool support, safety analysis requires tedious, manual work that is frequently seen as an obstacle by engineers, often reducing the scope of this task or limiting it to a single safety analysis. Automation of safety analysis has several advantages: it makes safety analysis easier, it is readily available, and it allows the engineers to obtain a thorough and quick analysis of their design. This rapid feedback based on analysis results allows engineers to perform more micro iterations in the development process, where each iteration refines and improves the previously built model.

3.3.5.2 Data Integration

Data integration in this context is concerned with the transformation of error modelling data. We transform from an EAST-ADL representation to a HiP-HOPS representation, while preserving the semantics. State of the art data integration for model-based development is supported by powerful model transformation engines and languages. Different transformation languages and engines are available, each of them solving a particular problem especially well. Identifying the right model transformation language/engine for the task at hand is a fundamental part of the solution.

3.3.5.3 Transformation Design

The model transformation process from EAST-ADL to HiP-HOPS is partitioned into two steps. Each step has a separate purpose and concern.

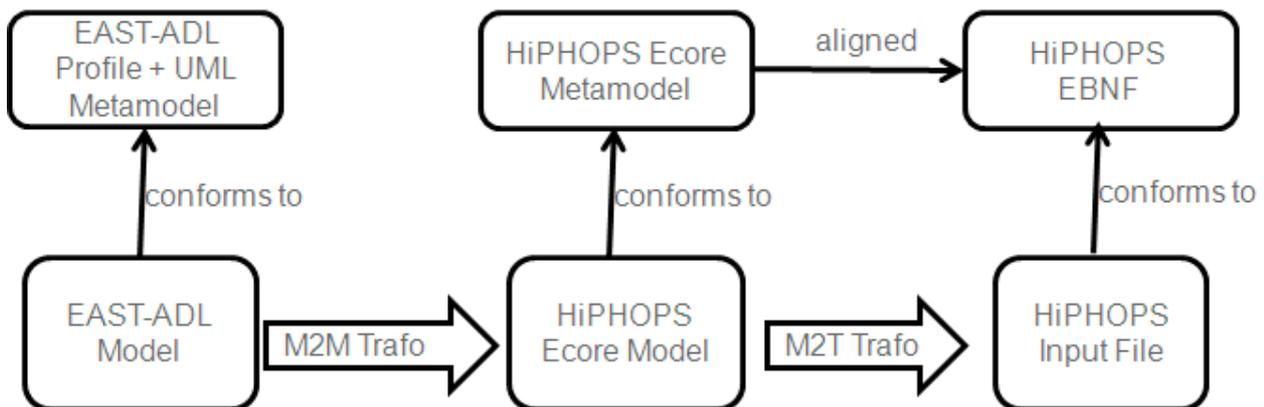


Figure 22 - Model transformation process

(1) *Semantic Mapping Transformation*: The first transformation step transforms an EAST-ADL model that is created in the Papyrus UML modelling tool and creates an intermediate model. The structure of the intermediate model resembles the HiP-HOPS grammar, so it is close to the structure of the desired output. This stage performs the semantic mapping between the domains of EAST-ADL and that of HiP-HOPS. However, this stage is not concerned with the actual representation of the data.

(2) *Representation Transformation*: The second transformation step takes the intermediate model and creates the input file for the HiP-HOPS program. This step is mainly concerned with the representation of the information according to the concrete syntax required by HiP-HOPS.

In the following section we discuss the benefits of this solution.

- Our solution separates two different concerns of the transformation from EAST-ADL to HiP-HOPS: (1) the semantic mapping between the domains of EAST-ADL and that of HiP-HOPS and the (2) details of the concrete syntax of the HiP-HOPS input file.

- Each transformation is a separate, self-contained module, which can be developed, changed and tested independently. This decomposition into two separate transformations allows us to parallelise the work on the two transformations and reduce development time. It also allows the two transformations to evolve independently without affecting each other, e.g. a change in the HiP-HOPS grammar will only affect the representation transformation.

As discussed in the section on data integration, different transformation engines have different strengths that can be played out for different concerns. The solution allows us to select the best tool for each concern.

3.3.6 Semantic Mapping Transformation

The purpose of the Semantic Mapping Transformation is to map concepts from EAST-ADL to HiP-HOPS in a way that preserves the semantics of the original model. The mapping between EAST concepts and HiP-HOPS concepts is explained in the table below.

EAST-ADL	HiPHOPS
ErrorModelType	System
ErrorModelType.errorConnector of type ErrorPropagationLink	System.Lines
ErrorModelType.parts of type ErrorModelPrototype	System.Component
ErrorModelPrototype.type.errorPort of type ErrorPort	System.Component.Ports
ErrorModelPrototype	System.Component.Implementation
ErrorModelPrototype.type.errorBehaviorDescription.internalErrorEvent of type ErrorEvent	System.Component.Implementation.FData.basicEvent
ErrorModelPrototype.type.genericDescription of type String	System.Component.Implementation.FData.outputDeviation
ErrorModelPrototype.type of type ErrorModelType	System.Component.Implementation.System (recursion)

Figure 23 - Mapping of concepts from EAST-ADL to HiP-HOPS

Model to model transformations are well suited for a semantic mapping transformation. Both input and output of a model to model transformation are models themselves. Mapping patterns can be described by relational and declarative transformation languages in a concise manner. We chose the ATLAS Transformation Language (ATL), a language that allows a choice of relational and imperative constructs. It furthermore allows processing of models that have a profiled metamodel, i.e., a metamodel that consists of a metamodel and a profile description. In the case of EAST-ADL the metamodel consists of the UML metamodel and the EAST-ADL profile.

3.3.7 Representation Transformation

The purpose of the Representation Transformation is the generation of a textual description based on the intermediate model. The intermediate model is designed to have structure which is aligned to HiP-HOPS. No structural changes are required in this transformation. The focus is on serialising the model as text.

Textual representations can be generated particularly well with model to text transformation languages. We chose the Xpand language from OpenArchitectureWare. Xpand is a template-based model transformation language, which incorporates the output in the form of templates into the control structure. The intermediate model is explored using a depth-first strategy.

EAST-ADL models created in Papyrus have a metamodel that is a composition of several separate metamodels. In the case of EAST-ADL, the metamodel consists of the UML metamodel and the EAST-ADL profile. These artefacts are composed by the Eclipse Framework to an EAST-ADL metamodel at runtime. The EAST-ADL metamodel corresponding to a Papyrus EAST-ADL model is not a separate artefact, and this complicates the model transformation and limits the choice of model transformation engines.

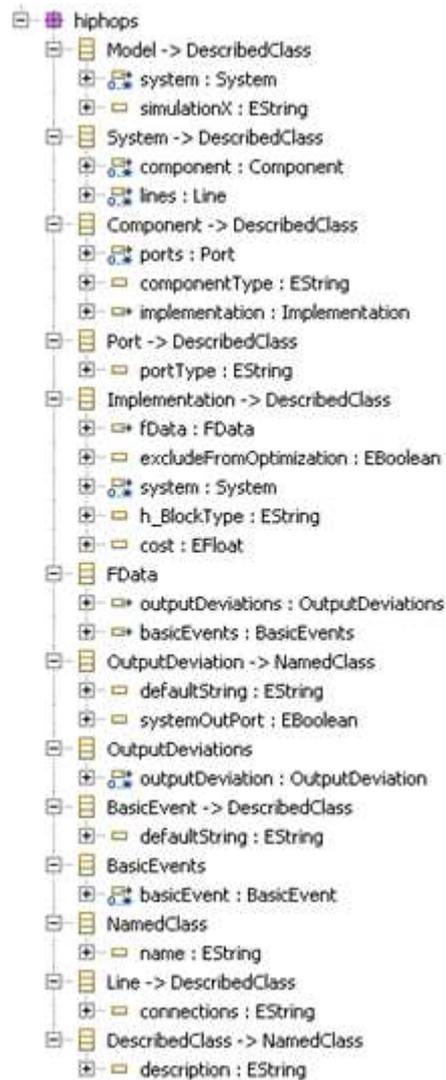


Figure 24 - Intermediate HiP-HOPS model

The intermediate model conforms to the HiP-HOPS ecore metamodel. It is aligned to the HiP-HOPS grammar. It also conforms to Ecore, and thus processable with the Eclipse Modelling Framework (EMF).

3.4 Timing Analysis

Timing modeling in EAST-ADL results from the work done in the TIMMO project, which produced a dedicated language called TADL (see TIMMO deliverable D6 for instance available from <http://www.timmo.org/>). TADL concepts were integrated in the course of the ATESS2 project in the EAST-ADL language.

Essentially three packages structure the timing concepts: *Timing* which defines core elements and their organization, *Events* which lists various kinds of events that can be associated to structural elements, such arrival of data on ports, and *TimingConstraints* which lists all possible constraints one can model – delays, synchronization, etc.

The modeling principle is the following. *TimingConstraints* are associated to an *EventChain*, which defines the scope of the constraint. The *EventChain* in turn relates to various *Events*, such as arrival of a data on a Port. *Events* point to *structural* elements, such as Ports or Functions

Based on this modelling concepts timing analysis can be performed. However TADL does not cover implementation level concepts such as tasks, which are needed to conduct a detailed schedulability analysis for instance. Thus the analysis that one can perform at this level is restricted to feasibility assessment regarding e.g. response times, age, synchronization and resource load balancing and assessing. This provides an interesting insight on how the software implementation could later be defined. To go beyond this, one needs to go to the implementation level (i.e. Autosar architecture) where the allocation of execution to tasks is defined. For this, MARTE provides a good basis. In the experiments done on timing analysis in ATESS2, the extra information needed in EAST-ADL models were added using MARTE constructs. The following section provides a summary of the main elements in both languages, EAST-ADL (and TIMMO) and MARTE.

3.4.1 EAST-ADL/TADL

EAST-ADL divides timing information into timing requirements and timing properties, where the actual timing properties of a solution must satisfy the specified timing requirements. EAST-ADL currently focuses on modeling of timing requirements on the functional abstraction levels of the architecture description language. The implementation level, i.e. AUTOSAR, is currently not explicitly considered, but it is expected that the information can be modeled in a similar way. The same holds for timing properties on both the functional abstraction levels and the implementation level.

Timing information on the functional abstraction levels is perceived as follows: timing requirements for a function can be captured on logical abstraction levels where no concrete hardware is yet available. This allows the specification of general timing requirements such as end-to-end delays from sensors to actuators regardless of how the final solution is built. This reflects the notion that a purely logical functional specification is not concerned with its technical realization, i.e. how many ECUs or bus systems are ultimately involved. What matters from the functional perspective are the recurring end-to-end delays of a control application, which need to keep pace with the real plant. Specifying timing requirements on the implementation level might be both too late in the development process and rather difficult because of language complexity (e.g. AUTOSAR) and the number of details on this level.

Timing properties are characteristics of a solution, e.g. actual response times, and should be reflected in the functional abstraction levels.

In EAST-ADL, timing requirements are divided into various kinds of delays (or latencies) for single time-consuming modeling entities as well as specific requirements for temporal synchronization of input or output data. The delays are either end-to-end delays, which are subject to segmentation along the functional decomposition track (e.g. end-to-end delay for a top-level function), or the delays form part of an end-to-end timing chain, and thus constitute segments of such an end-to-

end timing chain. Furthermore, delays can be sub-classified as being induced either by data transformation performed by a “FunctionPrototype” or data transmission via a “FunctionConnector”.

More precisely, EAST-ADL Timing concepts is based on

- Events : relate to EAST-ADL and AUTOSAR structural entities and depict observables : e.g. data arriving on a port, triggering of function execution, etc.
- EventChains : bind together events to establish sequences/relations between events e.g. to capture a complete end-to-end flow requirement between data sent by a sensor to output by an actuator.
- Constraints : put temporal constraints on sets of events or on event chains, e.g. deadlines to be met, expected delays, patterns of data arrival, synchronization of outputs on a set of ports, etc.

One example is featured in Figure 25.

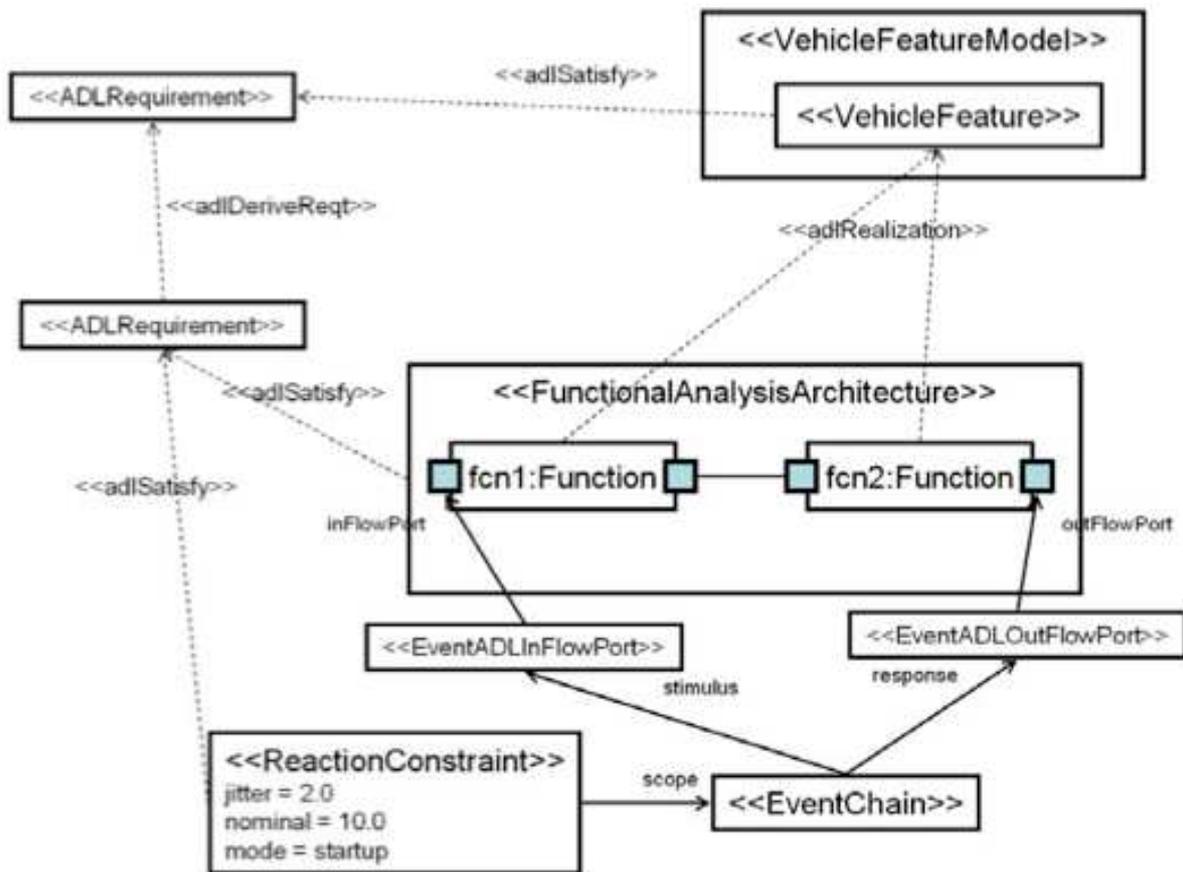


Figure 25 - Timing concepts

Analysis on the basis of analysis-level timing concepts concerns consistency among timing constraints, for example ensuring that the response constraints of parts do not exceed the response constraint of their composition, or that synchronization constraints on a set of events are not tighter than the variation among the response constraints that result in those events. Control performance based on specified timing constraints is also a typical concern on this level.

Analysis on the basis of design-level timing concepts can indicate feasibility of a solution where functions of a functional design architecture are allocated to nodes of a hardware design architecture. Nodes have a processing speed and functions have an execution time and

information about triggering. It is also possible to assign link speeds to take into account communication delays based on the data size of information exchange. In cases where communication buses are not unambiguous, explicit allocation of function connectors to communication buses is possible.

Standard schedulability algorithms can compute response times on this basis, using assumptions regarding function-to-task, data-to-frame, priorities, etc.

3.4.2 MARTE

The domain model for non-functional analysis in MARTE is organized around the notion of Analysis Context (see Figure 26). An analysis context is the root concept used to collect relevant quantitative information for performing a specific analysis scenario. Starting with the analysis context and its elements, a tool can follow the links of the model to extract the information that it needs to perform the model analysis.

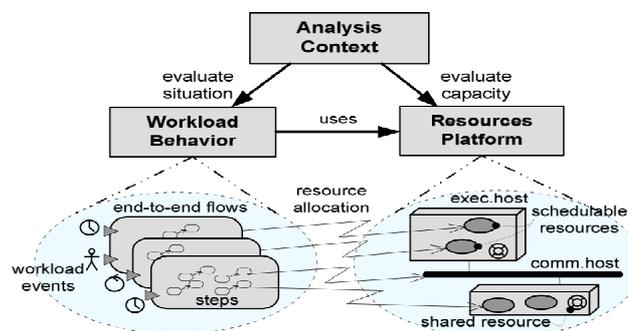


Figure 26: Organization of analysis-specific model elements in MARTE

Within the application model, “workload behavior” describe how often events will arrive and how much data are provided as inputs to the system, and how response events are generated in return by the system. In addition, workload behavior models provide information about the processing behavior that is used to execute the various tasks.

On the other hand, “resource platform” models provide information about the properties of the computing and communication resources that are available within a system, such as processor speed and communication bus bandwidth. The system model thereby captures information about the applications and the available resource platform of the system, and it also defines the mapping of application tasks to computation or communication resources.

The main analysis techniques in the scope of MARTE are scheduling-aware timing analysis. This kind of techniques offers a mathematically-sound way to determine whether a system will meet its timing requirements and how it could be balanced (e.g., task allocation, priority assignment, component deployment) by still respecting timeliness. To this purpose, scheduling analysis tools are a key component in this chain. Implementation of scheduling analysis tools depends a lot on the kind of scheduling analysis supported. In particular, the domain of scheduling analysis for automotive applications has received special attention in recent real-time systems literature. A comprehensive summary of analysis techniques for automotive applications is provided in [17]. Among relevant advances in this field, the holistic approach [18] extends the classical single-processor scheduling theory and applies it for specific combinations of input event models, resource sharing and communication policies. The global view on the system allows taking global dependencies into account (offsets between the activation instants of tasks), thus providing tightly calculated analysis bounds. A different approach for distributed architectures is the compositional scheduling analysis [19] [20]. The basic idea of this approach is to break down the analysis calculation into separate local component analyses (e.g., mono-processor analysis with RMA) and to integrate them for system-level analysis by evaluating the propagation of event stream models.

In our experiment in ATESS2 we worked with the MAST tool from the University of Cantabria, which supports classical RMA analysis and holistic analysis [21][22].

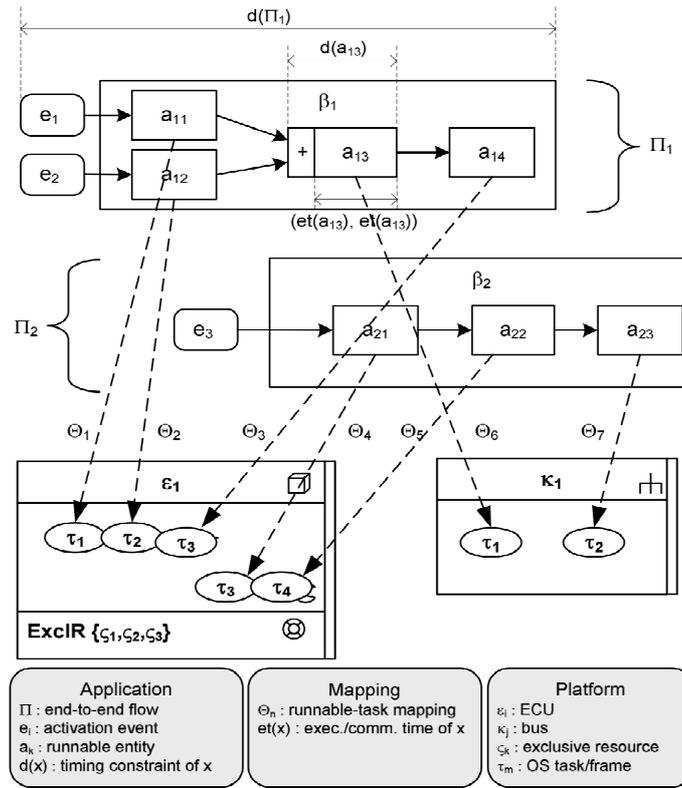


Figure 27: Simplified canonical model for scheduling analysis

Figure 27 shows a simplified canonical model of the modeling features required for scheduling analysis, which is discussed in this section.

Timing constraints

We can distinguish at least three abstraction layers of time constructs:

- Time-related abstraction layer, level 1: time is modeled as a set of fundamental notions such as time instants, durations, time bases, or clocks. They can be specified by relative/absolute durations (maximum time intervals) or instants (occurrence of a timeout event).
- Time-related abstraction layer, level 2: MARTE provides mechanisms to annotate timing requirements and constraints in models. Basic timing constraints include deadlines and maximum jitters. One key-modeling feature is the concept of observation. Observations provide marking points in models to specify real-time assertions. Some typical assertions have been predefined in ready-to-use patterns, such as jitters or conditional time constraints.
- Time-related abstraction layer, level 3: time concepts are defined as part of the behavior, not mere annotations. This set of constructs cover both physical and logical time. Most of the current scheduling analyses use the notion of physical time as measured by a unique time base. However, distributed applications often experience problems for agreement on consistent time reading due to clock synchronization. This means that scheduling them depends on different time bases, and therefore constraints must refer to specific clocks.

End-to-end flows

In MARTE, end-to-end flows describe logical units of processing work in the system, which contend for the use of processing resources (e.g., processors and buses). Let us also notice that firstly data and control can be part of the processing, and secondly, different kind of timing constraints can be attached to end-to-end flows (e.g., deadlines or output jitters).

One important feature in MARTE is that end-to-end flows can be represented in behavioral views (e.g., Sequence or Activity diagrams) complementing component models. This approach allows modelers to specify multiple end-to-end flow configurations that could be likely related to (a) specific operational modes, (b) alternative execution chains, or (c) different quantitative scenarios of activation parameters or other non-functional annotations.

Activation events

Both event-triggered and time-triggered paradigms are often involved in automotive applications. Event-triggered means that tasks are started, or messages are transmitted, following the occurrence of one (or a conjunction of) significant event (e.g., "a door has been opened"). Time-triggered consists of tasks started, or messages transmitted, at predetermined points in time, usually periodically.

In MARTE, activation models are denoted by means of workload events. Workload events can be modeled under different forms: by known patterns (e.g., periodic, aperiodic, sporadic or burst), by irregular time intervals, by trace files, or by workload generator models (e.g., state machine models). Workload events also enable to specify additional parameters for periodic and aperiodic patterns such as jitters, burst parameters, and distribution probabilities.

SW and HW resources

What is needed for scheduling analyses is to take into account the impact of the OS and hardware resources on applications (e.g., overheads due to the OS and the stack of communication layers or throughputs and bandwidths of underlying networks). Among these aspects, access protocols to mutual exclusive resources are of paramount importance in scheduling analysis of modern multiprocessor architectures.

The MARTE analysis model distinguishes two kinds of processing resources: *execution hosts*, which include for example processors, coprocessors and controllers, and *communication hosts*, which include networks and buses.

Processing resources can be characterized by throughput properties as for example processing rate, efficiency properties such as utilization, and overhead properties such as blocking times and clock overheads.

3.4.3 Summary

Both languages are complementary. In general MARTE is more focused on a behavioral description of a system in which timing information is a first-class input, whereas EAST-ADL assumes a fixed execution semantics for the functions, resulting in a more structured and static nature of models. Both approaches rely on end-to-end flow descriptions. The main difference is the level of detail in which both languages regard allocation. On this point MARTE requires more: an initial mapping from functions to tasks is needed, which is explained by the fact that MARTE considers the whole design life-cycle, where EAST-ADL delegates to AUTOSAR the details of the implementation levels.

Tooling for timing analysis is provided as a result from CEA LIST work in the EDONA project. It is for the moment MARTE-centric and takes a MARTE-annotated model as input and interacts with the MAST schedulability analysis tool (from Univ. of Cantabria, Spain).

3.4.4 Timing analysis tooling

A first version of a timing analysis gateway was developed in the context of the EDONA project <http://www.edona.fr> by CEA LIST. It consists of a set of plugins which take a MARTE-annotated model of a system and enable the modeler to run various schedulability analyses provided by a timing analysis tool. The plugins make a bridge to two specific tools:

MAST <http://mast.unican.es/>

RT-Druid <http://www.evidence.eu.com/content/view/28/51/>.

The model taken as input is a MARTE or EAST-ADL design architecture model, enriched with schedulability analysis parameters provided by MARTE annotations; these combine to produce a view of the system suitable for task scheduling and allocation to hardware. Based on this, the plugin transforms the model into the specific analysis tool format, provides an interface to the tool and returns analysis results. A typical scenario allows for the testing of whether a task allocation scheme allows for a system to be schedulable given particular hardware performance, or provide parameter values for task priority assignment.

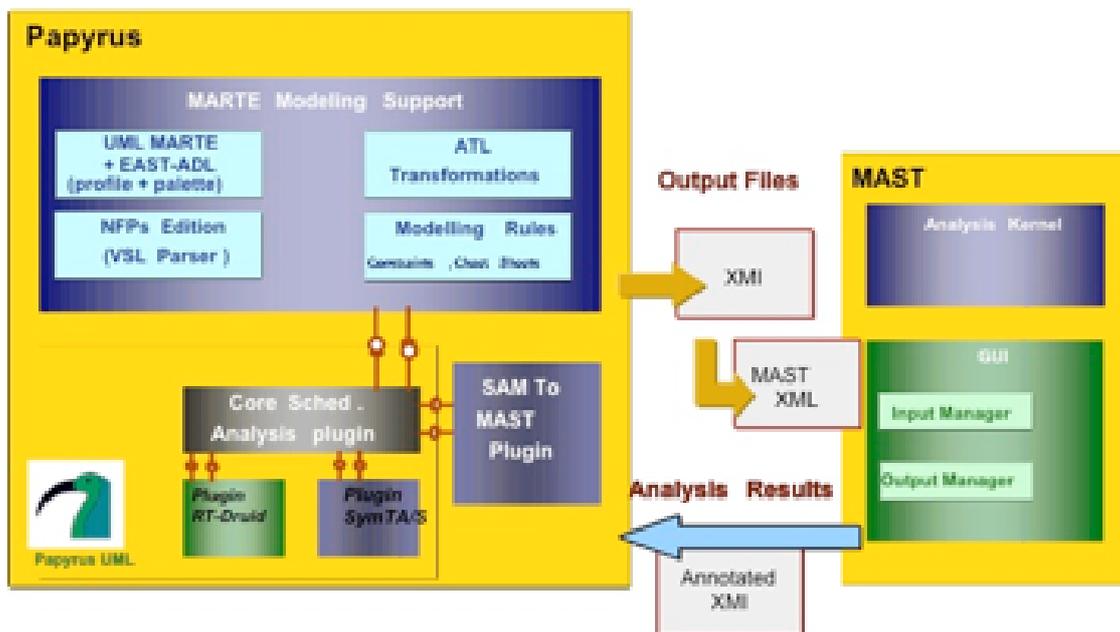


Figure 28 - Timing analysis framework

This plugin, now called Optimum, has been improved since and now features its own timing analysis algorithm, which no longer needs a connection to an external engine. Besides, part of the information needed, e.g. the task allocation is now performed using a wizard, which helps modeler point to the correct structural elements and fill in information.

Documentation and tutorial are in preparation.

In future versions Optimum will allow to perform a “traditional” schedulability analysis on implementation-level models (typically EAST-ADL models + tasks allocation) but also a simple resource load estimation based on the functional architecture and allocation to hardware. This is seen as being a good candidate for an analysis criteria to be taken into account in an optimization chain. A complete schedulability analysis requires additional structural elements not taken into account by the optimization engine – i.e. typically the tasks and task allocation are not considered by the optimization because it is not part of EAST-ADL models.

3.5 Analysis and Synthesis concepts to support Electrical Vehicles

This section describes analysis and synthesis concepts relevant for electrical vehicles and how we intend to support them in MAENAD.

3.5.1 Identified Needs

The engineering of an electrical vehicle is similar to that of a conventional vehicle in many ways. The particular needs we want to address in this section concern the analysis and synthesis of electrical power consumption, current analysis, power architecture analysis, cable length analysis, component cost analysis and component count analysis. This needs to be supported both statically and dynamically. Other possible FEV-related analyses include:

- Insulation analysis, to examine overall resistance and voltage compliance
- RESS short circuit analysis
- Charging inlet voltage analysis
- Transient voltage analysis, e.g. for accidental discharge of capacitors into accessible parts of the EV
- Range analysis, to determine maximum range of the EV given the battery capacity and electricity consumption of auxiliary components
- Power analysis in key-off mode, to prevent excessive discharge due to vehicle equipment operating in standby mode
- Powered braking analysis, to ensure correct braking operation in depleted-power scenarios

Constraints can be linked to EAST-ADL and AUTOSAR elements. The *mode* element defines where the constraint is applicable, and it is possible to analyse the total set of constraint in each mode. To find an average or total measure, the result can be considered in steady state where the modes change according to the behavioral definition. Below we will discuss what this means for each type of constraint studied.

3.5.2 Component count and cost analysis

The number of components produced and their combined cost is a critical parameter for the assessment of a solution, e.g. during optimisation or manual exploration of alternatives. The total cost is defined as the development cost for each component type plus the sum of piece cost multiplied by piece count summed over all component types:

$$\text{Total Cost} = \text{Sum of Development Costs} + (\text{Piece Cost} \times \text{Piece Count}) \text{ for all components}$$

To know the piece count, it is necessary to define the absolute or relative number of elements in the feature tree or in the artefact model. A constraint solver can compute the number of elements of any given component based on such constraints, and it can also warn if the model is underspecified or inconsistent.

In the example in Figure 30, 20000 vehicles are produced. Since 16% of all vehicles go to the US market, and 50% of those have trim level Prime, it is possible to deduce that 1600 such vehicles will be produced. 10% of all vehicles will be Plus, so these represent another 2000 vehicles. Both trim levels have ABSPlus which means 3800 PlusECUs will be produced and thus 16200 Standard ECUs.

Having established the number of components produced, it is possible to take this into account during optimisation: solutions with good price and performance for high volume products are favored before solutions that only benefit low volume products.

For example, if a high-volume vehicle needs an advanced component, while the low volume vehicle can do with a simpler component, it may well be better to equip all vehicles with the advanced component. This may also open up for after-sales revenue by selling upgraded functionality that relies on the better ECU. Figure 29 shows that the eliminated fixed cost for the low-end ECU compensates for the higher piece cost. This is consistent with Figure 30, as there is no configuration decision that states whether standard or plusECU shall be used. In the takerate model, it is undefined whether standardECU or PlusECU is used, unless the ABSPlus is selected (in which case, PlusECU is mandatory).

$50*3800+300000 + 45*16200+300000=1519000$ $50*3800+300000 + 50*16200+0=1300000$
--

Figure 29 - Total Product Line cost for same and different ECU

The assumption is that this is the stated rate; the constraints are part of the same GenericConstraintSet and thus intended to be consistent and simultaneously applicable. Further, the semantic assumption is that the root element of each product feature tree represents all vehicles.

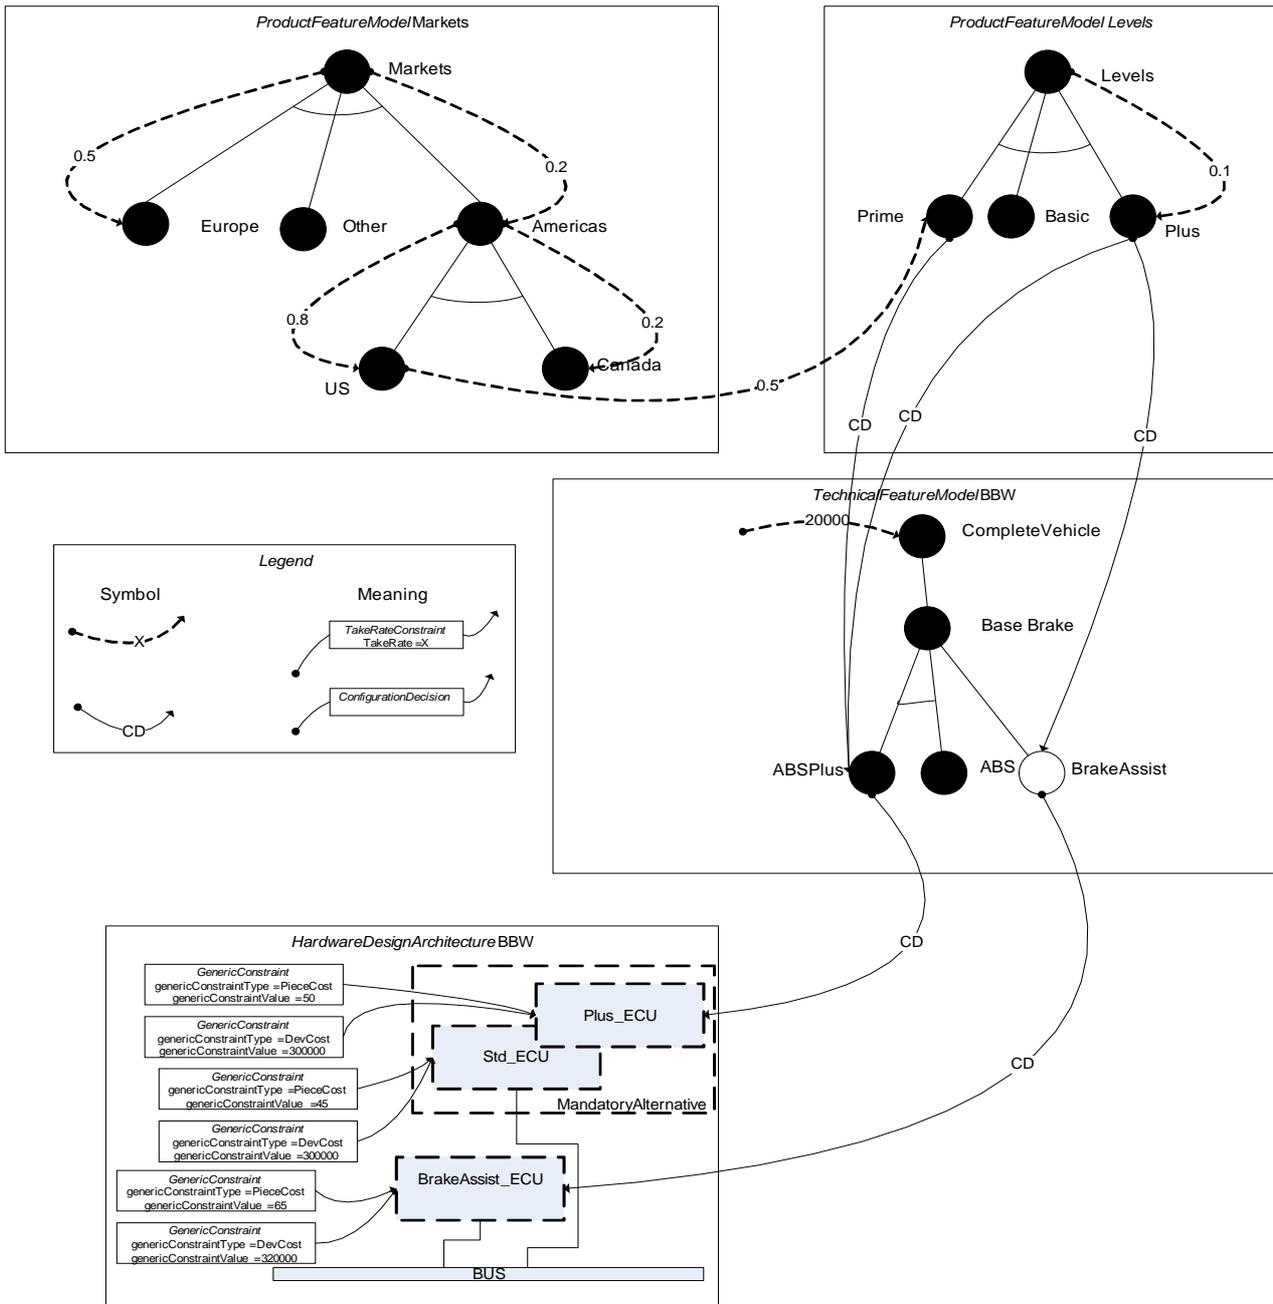


Figure 30 - Feature models and constraints to consider in optimization

One way to resolve the take rate constraints is to translate it to a constraint programming problem and use a suitable solver. Figure 31 shows how a constraint programme in Prolog may look. It also contains a minimisation criteria for cost.

```

% Prolog CLP(FD) pseudo-code
:- use_module(library(clpfd)).
% The feature tree as constraints
feature_tree(TotVehicles, Cost, Variables):-
    % All variables should be greater than zero
    TotVehicles#>=0,
    VehiclesOther#>=0,
    VehiclesEurope#>=0,
    VehiclesAmerica#>=0,
    VehiclesUS#>=0,
    VehiclesCanada#>=0,
    Prime#>=0,
    Basic#>=0,
    Plus#>=0,
    BaseBrake#>=0,
    ABSPlus#>=0,
    ABS#>=0,
    BrakeAssist#>=0,
    PlusECU#>=0,
    StdECU#>=0,
    BrakeAssistECU#>=0,
    PlusECUYesNo in 0..1,
    StdECUYesNo in 0..1,
    BrakeAssistECUYesNo in 0..1,
    Cost#>=0,

    % Constraints from Markets
    VehiclesEurope#=TotVehicles*0.5,
    VehiclesAmericas#=TotVehicles*0.2,
    TotVehicles#=VehiclesEurope+VehiclesAmericas+VehiclesOther,
    VehiclesUS#=VehiclesAmerica*0.8,
    VehiclesCanada#=VehiclesAmerica*0.2,
    VehiclesAmerica#=VehiclesUS+VehiclesCanada, % redundant constraint

    % Constraints from Levels
    TotVehicles#=Prime+Basic+Plus,
    Prime#=0.5*VehiclesUS,
    Plus#=0.1*TotVehicles,

    % Constraints from BBW feature tree
    BaseBrake#=TotVehicles,
    BaseBrake#=ABSPlus+ABS,
    ABSPlus#=Prime+Plus,
    BrakeAssist#=Plus,

    % Constraints from BBW architecture design
    StdECU#<=ABS
    PlusECU#>=ABSPlus
    PlusECU#<=ABSPlus+ABS
    StdECU#+PlusECU#=ABS+ABSPlus
    BrakeAssisECU#=BrakeAssist,

    % Objective function (single objective)
    PlusECU#>0 #<=> PlusECUYesNo,
    StdECU#>0 #<=> StdECUYesNo,
    BrakeAssistECU#>0 #<=> BrakeAssistECUYesNo,

    Cost#=PlusECUYesNo*300000+PlusECU*50+StdECUYesNo*300000+StdECU*45+BrakeAssistECUYesNo*320000+Brak
eAssistECU*65,

    Variables=[VehiclesEurope, VehiclesAmerica, VehiclesUS, VehiclesCanada, Prime, Basic, Plus, BaseBrake, ABSPlus, ABS, Bra
keAssist, PlusECU, StdECU, BrakeAssistECU].
% Find minimum cost solution for given number of vehicles
find_cost(Cost, Variables):-
    feature_tree(20000, Cost, Variables),
    labeling([minimize(Cost)], Variables).
% Find maximum number of vehicles for maximum given cost
find_vehicles(TotVehicles, Variables):-
    Cost#=<1600000,
    feature_tree(TotVehicles, Cost, Variables),
    labeling([maximize(TotVehicles)], Variables).

```

Figure 31 - Prolog code corresponding to the TakeRate constraints

3.5.3 Cable Length

Cable Length is a static measure and is established by considering all cableLength constraints allocated to elements of the HardwareDesignArchitecture, typically the HWconnectors. In case such fine granularity is not needed it is also possible to define cable lengths for components or component compositions that use a set of cables. In such case the cable length is specified without defining which cables it refers to.

To find the total cable length, the *genericConstraintValue* over all *GenericConstraints* with the *genericConstraintType* cableLength is added up.

If there is variability, all variation groups of the architecture need to be resolved first.

If an entire product line is considered, it is necessary to either sum the cable lengths over each fully resolved architecture and iterate over all fully resolved vehicles while summing up the total cable lengths, or calculate component counts over the entire product line (see TakeRate constraint) so that the component's cable length is multiplied with the component count and summed.

3.5.4 Power Consumption

The fidelity in modelling power consumption may vary. The average power regardless of mode could be represented (1), or the different power levels in each mode (2), or a varying power within each mode (3). Establishing the total power or energy consumption over an interval of time would then be treated differently in each case.

Alternative 1) would mean that the *genericConstraintValue* over all *GenericConstraints* with *genericConstraintType* powerConsumption is added up.

Alternative 2) would mean that each unique mode assignment for the integrated architecture is traversed, and the *genericConstraintValue* over all *GenericConstraints* with *genericConstraintType* powerConsumption are multiplied by the time spent in this mode and added.

Alternative 3) would mean that simulation or analytical means are used to establish the instantaneous power in each component. Because power is allowed to vary, the component power must be defined by a function. Typically, power is related to the nominal functional behavior of the component. The *genericConstraintValue* for a powerConstraint could thus be an expression with nominal signals as operands. This is however not supported by EAST-ADL.

The instant view of the power consumption can be used for integrating to energy consumption, averaging power levels for certain intervals, finding peak values, etc.

3.5.5 Current analysis

Current analysis is useful to establish cable and connector dimensions. It can be assumed to be the same procedure as for power, only that the results are divided by voltage.

3.5.6 Power Distribution Analysis

The power distribution is a critical service in an all-electrical vehicle. An analysis of how each hardware component and each SW component, function or feature relies on the power distribution network is thus important. This analysis can be done through the error modelling and error propagation analysis described elsewhere.

Analysis can also be conducted by considering the nominal architecture: each feature is realised by functions allocated to nodes in the HardwareDesignArchitecture. Nodes have HardwarePins that may be used for power. There is a PowerSupply component.

Power distribution is typically organised as independent power networks. Using the elements above, it is possible to analyse and visualise:

- which nodes, sensors and actuators depend on which PowerSupply. In case multiple power supplies are connected to a component, it is assumed that one is sufficient for operation (OR-logic);
- which functions depend on which PowerSupply (based on allocation to nodes, sensors and actuators);
- which features depend on which PowerSupply (based on realisation by functions) In case realising functions relies on different power supplies, it is assumed that all of these are required (AND logic).

The distinction between OR and AND dependency on the power networks can thus be established and visualised.

4 Automatic multi-objective optimisation of system models

This section describes the algorithms necessary to support automatic optimisation of EAST-ADL models. For a fuller introduction to the concepts involved, see D3.1.1, chapter 4.

4.1 Brief description of optimisation definitions and concepts

As systems grow more complex, they require increasingly complex analyses and designs to ensure that they meet all their requirements and do so efficiently, with a minimum of waste and excess cost. In many cases, the complexity and time involved means that this cannot be done manually, at least for more than a handful of possible alternative designs. However, automatic optimisation - particularly the type of model-based, multi-objective optimisation discussed in this document - presents one possible solution.

The goal is to rapidly evaluate many different potential designs to find the optimal (or at least more optimal) candidate. However, even though computers can do this automatically many times quicker than it could ever be done manually, the number of possibilities means that it is still prohibitively expensive - if not impossible - to check *all* of the different solutions. This means that we need some kind of algorithm to guide the process and select which solutions to investigate and which to ignore.

In MAENAD, one of the goals of the project is to allow EAST-ADL models to be optimised according to a variety of objectives, such as safety and timing, to enable better designs to be rapidly evaluated and evolved.

To aid in understanding during the following section, here is a brief summary of the concepts involved:

- **Design/Search Space** - The search space is the name given to the set of all possible solutions that can be derived from modifying the original starting point. In our case, this can also be called the design space, because it is the set of all possible designs that can be obtained via the variability inherent in the original model.
- **Optimisation Objective** - An objective is the goal of the optimisation process and typically takes the form of an evaluation criteria and a direction, e.g. to maximise safety, or minimise cost. Some optimisation approaches combine multiple objectives into a single objective via some kind of weighted formula, but the goal is still to minimise or maximise the result of the formula.
- **Multi-objective Optimisation** is the optimisation of more than one objective simultaneously. This means that multiple evaluation methods are required (to evaluate each objective) and it also requires some kind of heuristic for resolving the inevitable trade-offs between multiple, possibly conflicting objectives.
- **Genetic Algorithms** are one of the more popular types of optimisation algorithms. Some types of genetic algorithms also support multiple objectives. They are so called because they echo the evolutionary process found in nature, where organisms with more adaptable genes tend to survive to pass those genes on to later generations where other less adaptable and less successful genetic lines die out.
- **Evaluation & Fitness Functions** - To be able to judge whether one potential solution is better than another, it has to be evaluated against one (single-objective) or more criteria (multiple-objective). In genetic algorithms, these evaluation functions are often called "fitness functions", because they determine the fitness of an individual and therefore its likelihood of survival into the next generation.

- **Populations & Individuals** - In a genetic algorithm, a population is a collection of individuals, where an individual represents a particular solution (or in this case, design model). Each individual has a different genetic encoding.
- **Encoding** - In a genetic algorithm, this is the "DNA" of the individual, representing its genes. There are different types of encodings, ranging in complexity depending on the complexity of the model variability they are required to capture. In this case an encoding represents a particular configuration of the variability in the original model.
- **Dominance** - In a multi-objective optimisation, there is no global "best" solution, because everything is a trade-off between the various objectives. However, it is possible for one solution to *dominate* another if it is evaluated to be at least as good in all objectives. Thus the dominating solution would always be chosen over the dominated solution (it is the optimal design for those particular objective values).
- **Pareto Frontier** - The Pareto frontier is the set of all dominating solutions that have been found (none of the Pareto solutions dominate each other). Essentially, it is the set of current "optimal" solutions, representing the best trade-offs between the various objectives that have been found so far.

4.2 Comparison of optimisation algorithms

Automatic optimisation is still a relatively new field, but various approaches have been put forward - often inspired by natural processes - to serve as optimisation heuristics and algorithms. Some of the more prominent examples will be discussed below, with their advantages and disadvantages.

4.2.1 Genetic Algorithms (GAs)

Genetic algorithms, as briefly explained above, are optimisation algorithms inspired by the evolutionary processes found in nature. They have been identified as a particularly effective method for solving combinatorial optimisation problems (such as those faced in MAENAD) as they are capable of navigating large, complex search spaces [6]. The general process is as follows:

- A population of candidate solutions is randomly generated. Each individual is represented by a different encoding, which encapsulates the way they can vary from each other.
- Genetic operators - crossover and mutation - are applied to the population to obtain a new set of child solutions, which are added to the population (or, in some genetic algorithms, become the sole new population).
- This process of reproduction continues until the population grows too large, at which point those with the least successful genomes are discarded. This is established by evaluating the solutions in the population according to some criteria and retaining those with the best results - in other words, survival of the fittest.
- After a set number of generations, the optimisation ceases, and the current population should contain the best (i.e., optimal) solutions discovered so far.

The nature of the domain and the context involved can heavily influence the nature of the genetic algorithm and there are many types and variants of GA, and each can be tweaked and modified further according to various parameters. However, they all have the same general features in common: a population of individuals represented by some form of encoding, crossover and mutation operators that can be applied to the encodings to produce new individuals, and some kind of fitness function to evaluate the individuals.

Encodings represent the DNA of an individual. One of the most popular forms of encoding is to represent the individual as a string of integers - the value of each position in the string represents

a different variability point. Take for example a simple model consisting of three components connected in series. Each component can be replicated (i.e., duplicated and connected in parallel to achieve a measure of redundancy). The encoding string for such a system could consist of three digits, each of which refers to one of the three components and indicates how many times it is replicated. Thus the default system would be 111 - indicating no replication - and 121 would indicate that the second component is replicated once, and so on.

This form of encoding is easy to manipulate, but is quite limited in flexibility (for example, it cannot represent hierarchical systems). Other types of encoding are also possible, including tree-based methods, which are more suitable for modelling hierarchical and compositional models. In tree-based approaches, encodings typically consist of a tree of integer strings; some integers open up a new branch (indicating a subsystem) whereas others do not. So for example, the top level encoding of the default system may still be 111, but 121 implies that the second component is replaced by a subsystem, which will in turn have its own sub-encoding, and so forth.

Encodings are important because they make it possible for new individuals to be created in a mechanistic way from the existing population. This is done via two mechanisms, or genetic operators: crossover and mutation. Crossover is essentially reproduction - it takes two individuals and produces a third individual whose encoding is a random mix of the encodings of its parents. So for example, if individual A has an encoding 111, and individual B has an encoding 222, then the child individual C may have an encoding 122. Crossover is designed to persist good gene lines by allowing successful individuals to pass on their DNA to the next generation - the resulting child has characteristics of both parents. Mutation, by contrast, takes an existing individual and randomly alters some part of the encoding to produce a new one, e.g. individual D may be created by taking individual C and mutating it to have an encoding 123 instead. Mutation is meant to ensure that there is an element of randomness to the process, to allow entirely new areas of the design space to be explored and to help avoid the algorithm becoming stuck in a local optimum (a situation where it appears as though an optimum has been found, but in reality there are better solutions elsewhere in the design space, just apparently inaccessible from the current position; mutation enables the algorithm to 'jump' to other areas).

Genetic algorithms can be applied to multiple objectives. Many "multiple objective" algorithms deal with the multiple constraints imposed by the various objectives by combining them into a single objective/fitness function, typically in the form of a weighted formula; in effect, converting a multiple objective problem into a single objective problem. So for example, if the objectives were safety, performance, and cost, each would be normalised and weighted - e.g. $0.4 \times \text{safety}$, $0.3 \times \text{performance}$, $0.3 \times (\text{inverse of cost})$ - and the result of the formula used as the fitness for that individual. However, this is not always desirable and the fitness formula can be very difficult to derive, so better ways of dealing with multiple-objective problems were sought.

One such approach is the penalty-based GA approach. In this approach, one objective is treated as the main objective and maximised - e.g. safety - and other objectives are treated as scaled, weighted constraints and applied as a penalty to the main objective. So for example, a maximum cost limit might be set, and the more the design candidate exceeds the limit, the more harshly the fitness is penalised. Thus even if two solutions both achieve the same reliability, the one that does not violate any constraints - or which violates them by the smallest degree - is preferred by the algorithm. However, successful exploration of the design space often requires infeasible solutions (i.e. those that break the constraints) to be explored too, so the constraints are not hard constraints; instead, a dynamic penalty function is used that factors in not just the degree of constraint violation but also how long the algorithm has been running. By doing this, it allows the algorithm to explore more infeasible solutions early on, when the penalties are still small, and then increasingly focus on more feasible solutions (that do not violate the constraints) as the algorithm progresses. This form of dynamic penalty-based GA has been found to have superior performance compared to either static versions or versions that only allow feasible solutions to be part of the population [7].

Another type of GA is the multi-objective evolutionary algorithm. These are true multi-objective approaches that evaluate against multiple objectives at once and maintain a Pareto front of optimal, non-dominated solutions. This allows a broader set of solutions and generally results in a better coverage of the search space [8], albeit at the cost of higher complexity. These types of algorithms are generally better suited to multi-objective problems because they allow the various trade-offs between the objectives to be better represented, rather than being biased towards one objective and using the others as constraints or weighted penalties etc.

There are many different multi-objective GAs, including:

- PESA-II (Pareto Envelope-based Selection Algorithm 2) [9], which seeks to maximise evenly distributed spread in the Pareto set by including crowding in the selection criteria; solutions that are found in less crowded regions of the search space are preferred for selection to encourage the algorithm to look at under-explored areas of the search space.
- SPEA-II (Strength Pareto Evolutionary Algorithm 2) [10]. Unlike PESA-II, which adopts a "pure elitist" strategy that allows no dominated solutions to be retained, SPEA-II does keep some dominated solutions. The selection criteria are altered to include a dominance strength (how many solutions it dominates) and a density value (how close an individual is to other solutions). Thus although dominated, crowded solutions are possible, non-dominated, non-crowded solutions are preferred.
- NSGA-II (Non-Dominated Sorting Genetic Algorithm 2) [11]. In NSGA-II, both non-dominated and dominated solutions are kept in the same population, ranked according to dominance (based on how many other solutions they dominate). More dominating solutions are preferred during selection, and crowding is used as a tie-breaker if two solutions have equal dominance.

Although a full comparison of all the various approaches is difficult, [12] found that PESA was easy to implement and was efficient, NSGA-II was also efficient and widely used, while SPEA-II was criticised for being computationally expensive.

4.2.2 Tabu Search

Tabu Search is another type of heuristic algorithm. Whereas a GA performs a global search, a tabu search algorithm makes more use of the local region of the design space [13]. It works as follows:

- An initial solution is randomly generated and evaluated.
- Next, all neighbouring solutions - i.e., those that can be obtained by making a single modification to the initial solution - are generated and evaluated.
- The neighbour with the best fitness is chosen as the next solution, and the process repeats.

In this way, Tabu search deterministically moves through the design space. To prevent it from looping or getting stuck, it also maintains a list of forbidden or 'taboo' moves, which is populated by recently tried solutions. This helps force the algorithm to search in new, unexplored regions, although the list is finite and so it is possible to loop back some time later. The whole process is repeated until either a set number of moves have been made or a solution is found for which no better solutions are available. Modifications to the algorithm are also possible, e.g. by using an adaptive penalty function similar to that found in penalty-based GAs.

Tabu search can be more efficient than GAs in some circumstances [13], requiring less evaluation, but they can also be more prone to becoming stuck in local optima. Variations on the algorithm exist, e.g. to offer better support for multiple objectives; one such approach randomly selects one of multiple objectives to be active at each stage [12], giving a more equal balance to the multiple objectives and exploring more of the search space. In this approach, a set of good solutions is remembered; at each stage, the best newly discovered solutions are added, and any dominated solutions in the resultant set are discarded. This produces a Pareto front. To ensure an even

spread of the Pareto front over the search space, a diversification feature can be added wherein the algorithm can randomly select another Pareto solution to be used as the starting point if no new solutions have been discovered for a set number of iterations.

4.2.3 Ant Colony

The Ant Colony optimisation algorithm is another nature-inspired approach, this time taking its source from the way ants navigate. In nature, ant colonies use pheromone trails to establish and navigate efficient routes between their nests and food sites. Although the pheromones evaporate over time, good routes are reinforced and retained as more ants pass through it, while bad routes are left to decay and eventually disappear [14].

The algorithm generally works as follows:

- An initial ant colony is generated.
- The ants conduct a local search and establish initial pheromone trails.
- The better trails (e.g. those leading to better solutions) are reinforced by more ants over a series of iterations.
- The strongest trail is followed and a new ant colony established; the process then repeats.

The Ant Colony approach can have similar performance to Tabu search [14] but can also be more computationally expensive than other approaches, e.g. penalty-based GA. As with the other algorithms, modifications and variants exist, e.g. ones which allow the use of *a priori* knowledge or which store and rank a number of best individuals rather than only a single individual. Using multiple individuals allows a larger portion of the design space to be explored (essentially, by multiple ant colonies) at the expense of performance.

4.2.4 Simulated Annealing

Simulated Annealing is inspired by the process of metallurgical annealing: when a metal is heated, the atoms are free to change their position randomly in the high-energy environment. By controlling how quickly the metal cools, it is possible to encourage the atoms in the metal to form in stronger, more beneficial arrangements, thus strengthening the metal. Simulated annealing works in a similar way:

- An initial solution is generated and a high 'global temperature' is set.
- The solution is randomly mutated to another solution by making a single change.
- Acceptance of the solution depends on the temperature; if it is high, then highly random changes are possible and acceptable; if it is low, then only minor changes are possible.
- The process repeats but the global temperature is steadily reduced.

The result of the process is that large areas of the design space are explored early on via highly random changes, and later on only minor improvements are made. Simulated annealing follows a greedy method - it seeks to achieve a global optimum by making decisions that are locally optimum at each stage [15].

Simulated annealing's primary advantage is that it is very fast, but as a result it is not always as thorough as other algorithms.

4.3 Tool support for automatic optimisation

The main tool support for automatic optimisation in MAENAD is the HiP-HOPS optimisation engine, developed by the University of Hull as part of the HiP-HOPS analysis tool. HiP-HOPS implements a multi-objective genetic algorithm to allow for the automatic optimisation of system models (in HiP-HOPS format). Currently, only three objectives are natively supported - unavailability (as determined by the HiP-HOPS safety analysis engine, which is based on FTA) cost, and weight. The latter two are calculated as a simple sum of component costs/weights.

4.3.1 Variability in HiP-HOPS

Variability is the means by which the design space is defined. In a HiP-HOPS model, variability is represented by **implementations**. Each component in the model can have one or more possible implementations, each of which will have different characteristics (cost, safety etc). It is also possible for implementations to have their own subsystems, meaning that the variability is compositional. This allows several different types of variability to be supported, including:

- **In-place Substitution** - one implementation can be substituted for another that performs the same purpose and has the same interface, but with different cost and reliability attributes.
- **Compositional Substitution** - an implementation can also be swapped for a subsystem that performs the same purpose and has the same interface, but consists of multiple subcomponents (and possibly more subsystems). This is also known as *architectural substitution*.
- **Replication** - by means of compositional substitution, an implementation can be replaced by a subsystem which contains the same implementation(s), but replicated to achieve redundancy.

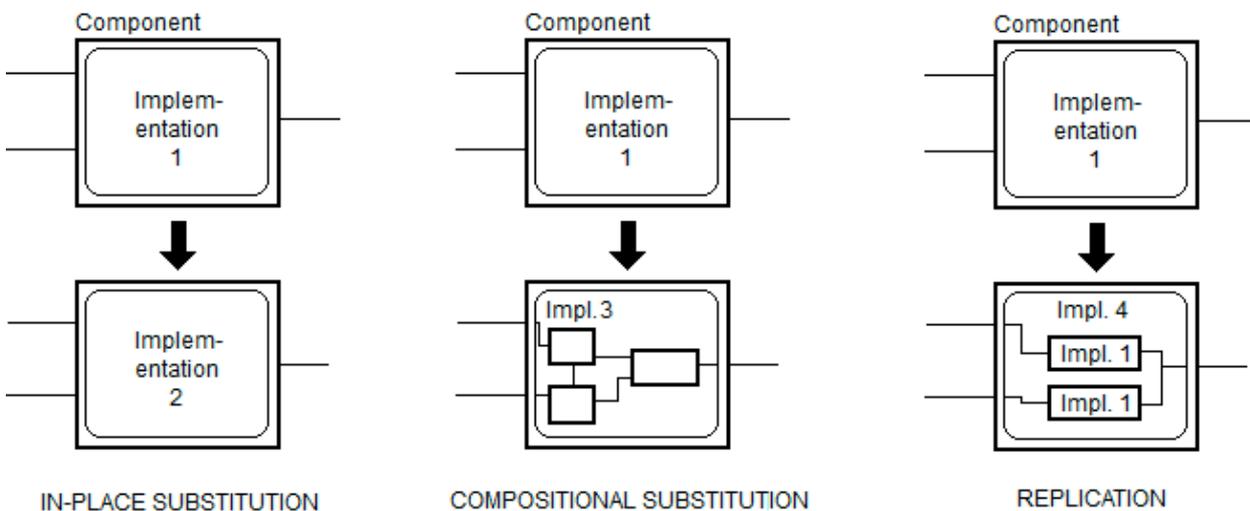


Figure 32 - HiP-HOPS Variability

All of these approaches require that the original model is annotated with the additional implementations by the designer. The designer can also add disable or enable certain implementations for a given optimisation.

Currently, each HiP-HOPS Component defines its own interface in terms of Ports and connections to Lines, but the failure behaviour and characteristics of that Component - including any subsystems - is defined solely by the Implementations. This means that the Component interface

is fixed, but the behaviour (and sub-architecture) can be varied by means of the Implementations. Each Implementation has:

- a name
- an optional description
- a flag to exclude it from the current optimisation (e.g. to turn off an particular implementation)
- an optional cost figure
- an optional weight figure
- a setting to indicate whether failure logic is hierarchical or flat
- a set of failure data, including basic events, input & output deviations and their failure logic, any propagations to other perspectives, and any common cause failures
- an optional subsystem

Currently, HiP-HOPS variability is not harmonised with the more extensive variability support found in EAST-ADL. One of the goals in MAENAD is to look at how these can be brought closer together and to identify any EAST-ADL variability features that HiP-HOPS should be extended to support.

4.3.2 HiP-HOPS Optimisation Algorithm

As stated above, HiP-HOPS employs a multi-objective optimisation algorithm. In particular, it implements a heavily modified version of the NSGA-II algorithm [16].

4.3.2.1 Encoding

HiP-HOPS originally used a traditional fixed-length integer string encoding, but to support proper compositionality in the model, it was upgraded to use a tree-based encoding that echoes the hierarchical structure of the model itself:

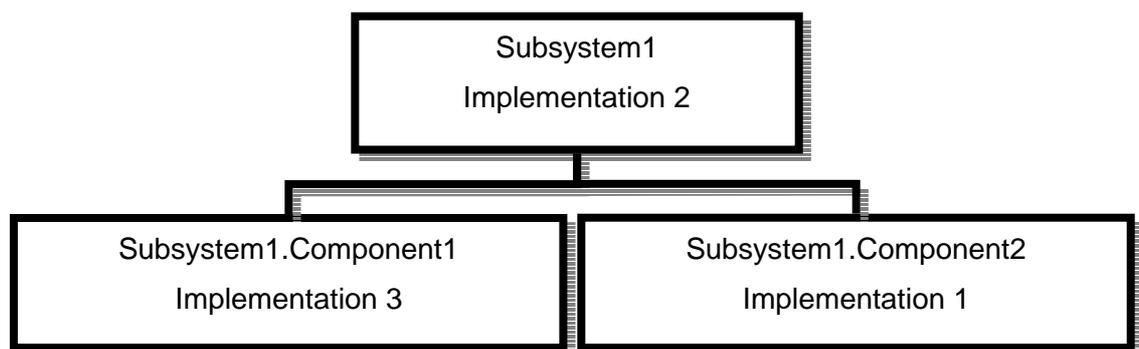


Figure 33 - Tree Encoding in HiP-HOPS

In the tree encoding, a certain implementation may have a subsystem of its own, the components of which may in turn have their own subsystems too. A 'default' implementation is always set as this defines the structure of the original design model and facilitates normal safety analysis.

Mutation is handled by randomly changing the implementation of one or more of the nodes in the tree. In some cases this can result in a subsystem being removed, in others it results in a subsystem being added; in the latter case, subsystem's implementations are randomly selected.

Crossover is achieved by performing a simultaneous depth-first traversal of the tree encodings of both parents and the child (the child being constructed in the process). There is a 50% chance of a node from either parent being inherited. Where a subsystem node is inherited from one parent that the other parent does not possess, the entire subsystem is inherited from that parent; where the subsystem is present in both parents, the traversal continues and the sub-nodes are randomly inherited from one or the other parent.

4.3.2.2 Algorithm

In pseudo-code, the HiP-HOPS algorithm is as follows:

```
Randomly initialise population
While (currentGeneration < maxGeneration)
    Generate child population
    Calculate total population crowding
    Crop population that exceeds sizeLimit
End While
```

The sub-functions are defined below:

Generate child population

```
While (childPopulationSize < maxChildPopulationSize)
    Select parent1 using Binary Tournament Selection
    Select parent2 using Binary Tournament Selection
    Create child using Crossover of parent1 and parent2
    Mutate child
    Add child to childPopulation
End While
```

Crossover

```
For Each Node in Tree Encoding
    If parent1 node not equal to parent2 node
        Flip a coin to choose parent1 or parent2
        Child node is set to selected parent node
    Else
        Child node is set to (equal) parent node
    End If
End For Each
```

Mutation

```
For Each Node in Tree Encoding
  If within mutationRate
    Randomise node and subNodes
  End If
End For Each
```

4.3.3 Achieving full multi-objective optimisation in MAENAD

To achieve the optimisation objectives in the MAENAD project, there are three main steps that must be fulfilled:

- **It must be possible to define the design space.** This is primarily a language issue and means ensuring that EAST-ADL and its variability mechanisms can be used to define optimisation-based variability as well as normal product line variability. Although the language definitions are largely beyond the scope of this document, in practical terms these mechanisms must also be harmonised with HiP-HOPS and any plugins, so that tool support exists for them.
- **It must be possible to evaluate the design candidates.** This will require a close link with both existing and developing analysis tools. HiP-HOPS already has a direct connection with its safety analysis elements but to allow other forms of analysis, e.g. timing, it will likely require some substantial modifications to the optimisation engine, particularly if the engine is to use a modelling tool such as Papyrus as its base.
- **The multi-objective optimisation heuristics must allow for an efficient exploration of the design space.** In practice, the demands of accomplishing the MAENAD objectives are likely to require extensions, modifications, and tweaks to the optimisation algorithm implemented in HiP-HOPS.

The first step is the variability support. If it is not possible to define the design space, then no optimisation can take place. The central principle here is **substitutability**: the language has to allow the designer to specify alternative versions of model elements while ensuring that the alternatives are sufficiently constrained such that they can readily be substituted for one another by the optimisation algorithm. If swapping one alternative for another results in an invalid model, then they are not substitutable. Of particular difficulty here is ensuring a consistent interface; some variants may have more ports or connections than other variants, and so it has to be possible to ensure that those connections remain valid regardless of which variant is chosen. For example, if a variant with more connections to other components is chosen, then failure propagation should be possible along those connections.

This leads on to the second step - that of evaluation of design variants. The notion of substitutability is encapsulated in practice by the process of **variability resolution** wherein a model containing variability is *resolved* to produce a new model in which all the variability has been fixed to choose a particular configuration with a concrete set of objective attributes. Until this problem is solved, and it becomes possible to produce models with resolved variability that can be subsequently evaluated, true multi-objective optimisation of EAST-ADL models will not be possible due to the heterogeneity of analysis tools involved.

Although the tool architecture is still to be fully defined, the overall work process can be divided into four main elements (as shown by the rectangles in Figure 34):

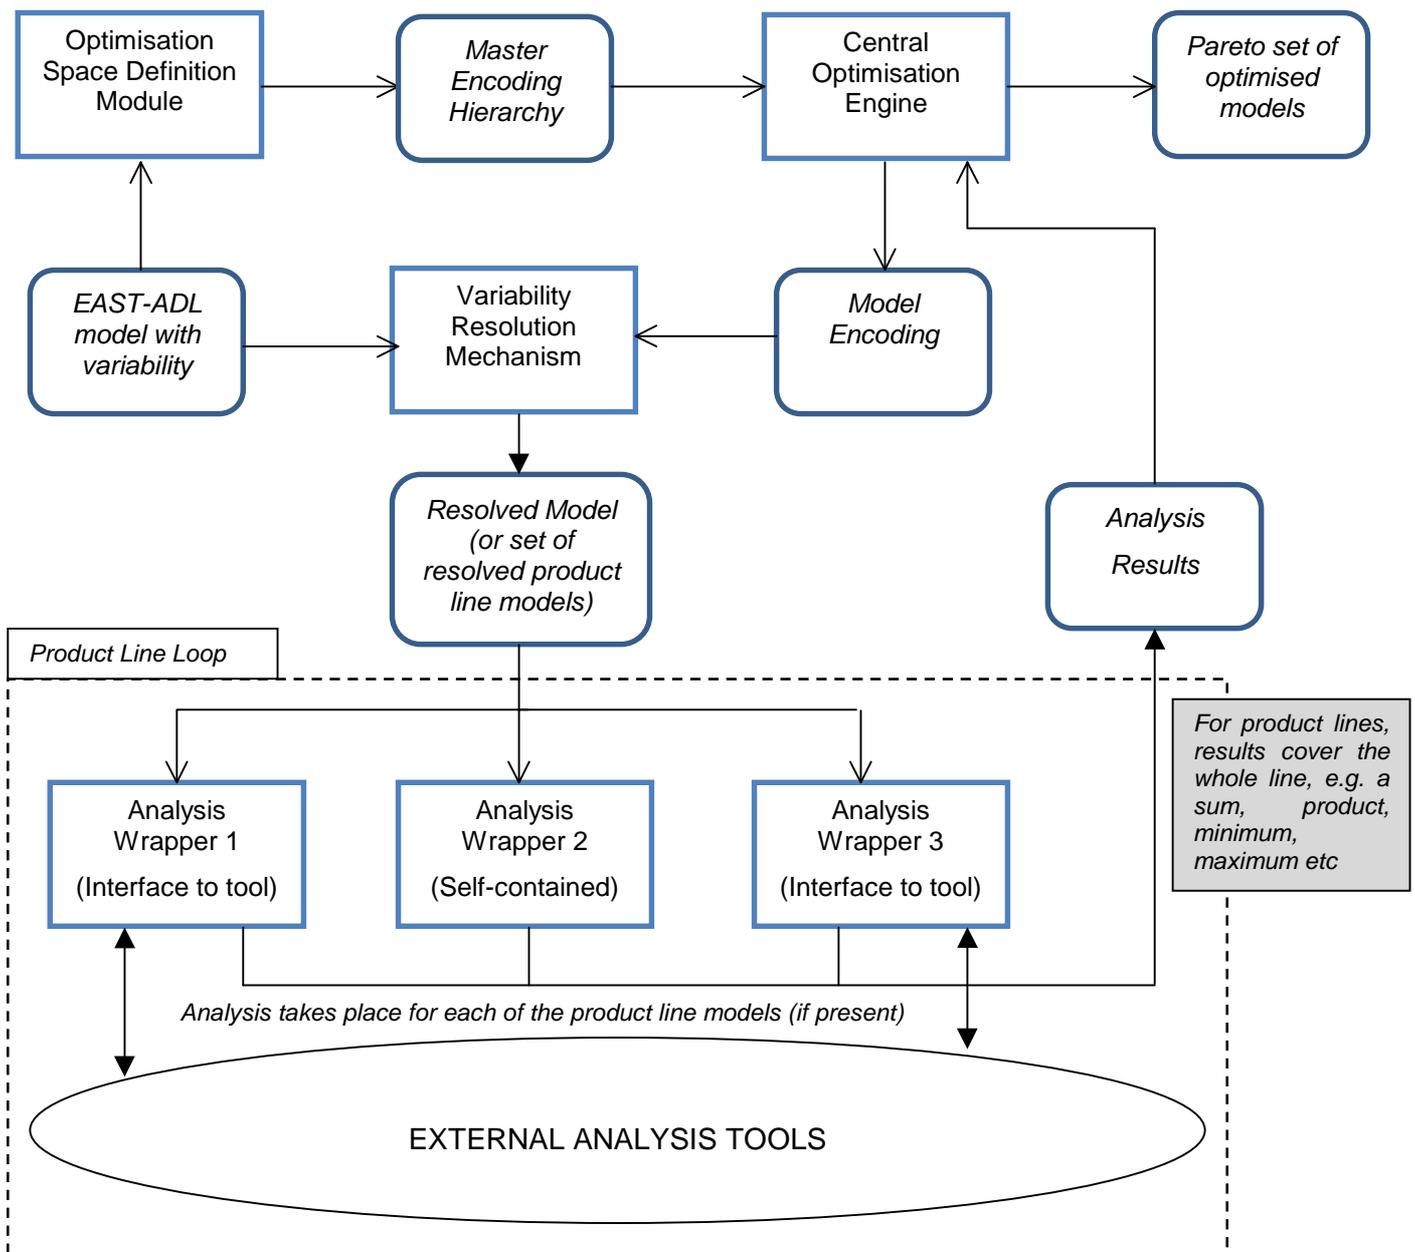


Figure 34 - Optimisation process

This diagram shows the flow of operations in the optimisation process. The initial model (with variability) is passed to the OSDM, which generates a Master Encoding Hierarchy. This defines the search space for the Central Optimisation Engine, which generates different model encodings that are passed to the VRM to be resolved into analysable models. The analysis wrappers then evaluate them according to different objectives and return the results. Eventually the optimisation settles on a set of Pareto optimal solutions, which constitute the final results of the process.

The architecture itself will be further explained below.

4.3.4 Optimisation Architecture Elements

4.3.4.1 Optimisation Space Definition Module (OSDM)

Function

The OSDM provides the input for the overall optimisation process by taking the original model, which contains variability elements to define the variation points of the design, and derives from this input the set of possible encodings of the optimisation search space – the *Master Encoding Hierarchy*, which is a structure that can be manipulated to obtain the set of all valid design candidates (also called population). This allows the Central Optimisation Engine (below) to explore the design space automatically by using an abstract structure rather than needing semantic knowledge of the design model itself. This is “abstract” in the sense that it does not include any information on the structure, behaviour or other properties of the individual candidates; it just provides a means to unambiguously denote each candidate. Unlike the other elements, which are used in each optimisation iteration, the OSDM is only required once, at the beginning of the optimisation, not iteratively in each optimisation cycle as the other elements.

Inputs

- A variant-rich model, i.e. one containing variability in the form of variation points. Provided as an external input. May contain both optimisation variability and product line variability (the precise distinction between the two types is still to be determined; it might be binding time or it might be something more explicit).

Outputs

- Master Encoding Hierarchy (MEH) of the optimisation space. This will be used by the optimisation engine to select and denote candidates to be analysed in each optimisation cycle. This notion of optimisation space corresponds exactly to what is called a configuration space in feature modeling terminology. Therefore, a feature model would be a natural fit for representing the optimisation space. However, for pragmatic implementation reasons a different form of encoding might be required to simplify the implementation of the OSDM and/or the optimisation engine or to enable the reuse of existing implementations (to be investigated). The format of the MEH is yet to be finalised but is likely to be based on XML somehow.

Discussion

The optimisation space definition module (OSDM) takes a model with variability and derives from that an abstract definition of the optimisation space, i.e. a structure which can be manipulated to obtain the set of all valid design candidates – the Master Encoding Hierarchy (MEH). This definition is “abstract” in the sense that it does not include any information on the structure, behaviour or other properties of the individual candidates; it just provides a means to unambiguously denote each candidate.

It takes the form of a tree structure which shows each point of variability in the model; therefore it is analogous (and very similar to) an ordinary feature model. The precise mechanics are yet to be determined, but the principle is that each node represents a variability point and the nature of the node describes the type of variability. For example, assume we have a simple model with one optional component (C1) and another component (C2) that can have one of three implementations (A, B, or C). Implementation C itself is a subsystem that can be replicated upto 3 times for redundancy purposes. This model may produce a MEH like this:

MASTER ENCODING

```

+
|
+----> C1: 0..1   +----> Implementation 1
|
+----> C2: 1      +----> Implementation A
|
|
+----> Implementation B
|
|
+----> Implementation C: 1..3

```

Note that the above is purely illustrative and not meant to be a definitive example of what the MEH may ultimately look like. In particular, the ‘types’ of variability we wish to include are not yet finalised. The example above uses optionality, multiple implementations, and replication, but we may decide to use other types (like k-out-of-n selections) or fewer types. Also to be determined is whether or not it is necessary (or possible) to ‘link’ variability points so that e.g. selecting one point automatically enables another variability point too, so we can model dependencies between them.

Regarding product line variability, the exact handling of product line variability vs optimisation variability is yet to be determined, both in terms of how they are distinguished in the EAST-ADL model and also in terms of how they appear in the Master Encoding Hierarchy. It is likely that only ‘true’ optimisation variability will be present in the MEH, and the product line variability remains unresolved in the model; in this scenario, the VRM would then be responsible for generating all possible product line examples (resolving both optimisation and product line variability) for analysis purposes. The OSDM could also have a mode in which it strips out or hides any irrelevant variability, e.g. allowing optimisation to take place on only one product line variant. In its initial incarnation, this is likely to be the case, with future extensions to allow full product line functionality (thus we plan for the full functionality but only implement a subset initially).

- **Optimisation Space Definition Module (OSDM):** this element provides the input for the overall optimisation process by taking the original model, which contains variability elements to define the variation points of the design, and derives from this input the possible encodings of the optimisation search space. This allows the optimisation engine (below) to explore the design space automatically by using an abstract structure rather than needing semantic knowledge of the design model itself. Unlike the other elements, which are used in each optimisation iteration, the OSDM is only required once, at the beginning of the optimisation.
- **Optimisation engine:** this is the driver of the process, responsible for exploring the design space (by choosing which encodings to evaluate) and for collecting the best solutions so far. The likely algorithm to use for the optimisation is a genetic algorithm, but others could be used (including options to use a simple random selection or enumeration of a small design space, for example).
- **Variability Resolution Mechanism (VRM):** this element is responsible for taking the original model (with variability) together with an encoding and then producing a new model from them, with all variability resolved. This model can then be subjected to analysis for the purposes of evaluation.
- **Analysis:** this element is responsible for analysing the model according to the objectives given to the optimisation engine. Typically there would be a separate analysis for each objective, although in practical terms these analyses could be carried out by the same tool or in the same

process. The results of the analysis/analyses are then fed back to the optimisation engine, which uses them to decide whether or not to retain that candidate model.

Apart from the OSDM, the optimisation process therefore consists of a loop - from engine to VRM, to analysis tools, and back to engine - which is iterated continuously until the optimisation is complete.

The OSDM is necessary to ensure that the optimisation engine itself does not need to know anything about the semantics of the models it deals with other than what variability is present in the original model and how it can be encoded. It is then job of the VRM to convert a particular encoding - i.e., a particular configuration of the variability in the model - to produce a new model in which all the variability is resolved. This can then be analysed by the tools which feed their results back to the optimisation engine, which determines whether to keep or discard that particular design candidate on the basis of those results.

4.3.4.2 Central Optimisation Engine (COE)

Function

The Central Optimisation Engine is the driver of the optimisation process, responsible for exploring the design space (by choosing which encodings to evaluate) and for collecting the best solutions so far. The COE will use genetic algorithms based on HiP-HOPS technology, but HiP-HOPS itself will only be used as an analysis tool. Instead the intention is to create a new implementation that can more tightly integrate with the other elements of the optimisation architecture (e.g. in the form of a Papyrus plugin).

When optimisation is required, the optimisation engine receives from the OSDM a tree containing the different variation points in the model together with a set of optimisation objectives. This is the *master encoding*, the base DNA of the model to be optimised. To perform optimisation, the optimisation engine chooses a particular encoding (by selecting which variation points to use and which not to use according to its internal heuristics) and passes this on to the VRM. The VRM returns a version of the original model with the variability resolved (although product line variability may still be unresolved). This resolved model is then passed to the plugin interfaces to the various analysis engines (timing, safety, cost etc), according to the optimisation objectives. Once complete, the analysis results are returned so the optimisation algorithm can evaluate them against its objectives and then repeat the cycle.

Inputs

For starting optimisation:

- A set of optimisation parameters. These should specify which analyses are to take place (e.g. safety, timing) and with which tools, which particular attributes are to be optimised (e.g. cost, reliability, schedulability etc), and whether to maximise or minimise those objectives. Any constraints may also be included here (e.g. a minimum level of safety, a maximum cost etc). The set of optimisation parameters comes from the user (requiring some kind of user interface).
- The Master Encoding Hierarchy, which describes all the variability points in the model. The master encoding tree should be provided by the OSDM. It is likely to be in XML format and provided in-memory (but there are other alternatives).

During optimisation, in each iteration the engine receives:

- A resolved model; after passing the current encoding to the VRM, the optimisation engine receives a model with the variability resolved accordingly. It can then pass this to the analysis plugins.

- The set of analysis results for each objective; this is essentially a single floating point number for each of the objectives (e.g. total cost, system reliability etc). It receives this information from each of the analysis plugins, which are interfaces to the relevant analysis tools. It is possible for multiple objectives to be analysed by the same tool, as this is transparent to the optimisation engine (which only needs to know which plugin to call for which objective).

Outputs

During optimisation, in each iteration the engine provides:

- A configuration encoding, which is passed to the VRM. The VRM can then resolve the variability in the model according to that encoding (equivalent to configuring the feature model). This is likely to be in the same format as the Master Encoding Hierarchy.
- Execution instructions to the analysis plugins, together with the resolved model to be analysed. Note that the model may still contain *product line* variability; this will be discussed later.

In terms of optimisation results:

- When the optimisation process finishes (e.g. because it meets its time limit, or because no progress has been made for X number of iterations, or because the user manually stops the process etc), it will output the Pareto set of all optimal encodings it has found so far, together with the evaluation results for each objective. Ultimately this would probably be presented to the user in the form of a selection dialogue; when the user selects a particular encoding (and set of evaluation results), the VRM can take that encoding and provide the user with the resolved model in question. This avoids the need to store all of the optimal models themselves. Thus the Pareto set will probably be stored as a set of XML-based model encodings, together with their analysis results.

Discussion

The Central Optimisation Engine is the driver of the whole process, but paradoxically is also the element that needs to know the least about the model being optimised. The engine will likely be implemented using genetic algorithms as a Papyrus plugin (i.e. Java), to allow a closer integration of the different optimisation modules.

Because the COE would (probably) be responsible for dispatching the models to be analysed by the various analysis tools, it will also need to know when those analyses are complete. There is therefore the need for some kind of event or callback mechanism so that the COE knows to carry on with the next iteration of the optimisation. Another question is to decide on the parameters for finishing the optimisation: should it run for a set number of iterations, for a given time period, can it be interrupted etc? However, these are technical details and if we choose the simplest option now (e.g. number of iterations), we can expand the solution with other methods over time.

4.3.4.3 Variability Resolution Mechanism (VRM)

Function

The Variability Resolution Mechanism (VRM) is responsible for taking the original model (with variability) together with an encoding and then producing a new model from them, with all variability resolved. This model can then be subjected to analysis for the purposes of evaluation.

Inputs

- A variant-rich model, i.e. one containing variability in the form of variation points. Provided as an external input.
- A configuration encoding (from the optimisation engine). The precise format of this encoding has to be clarified. One option is to use feature configurations but it is likely to be in a similar format as the Master Encoding Hierarchy. It will likely be XML-based and passed in-memory from the COE.

Outputs

- A fully resolved model, i.e. one that does no longer contain any (optimisation) variability. No special representation has to be devised for this, because it will just be a normal model with variation points removed. Will be forwarded by the optimisation engine to the analysis tool(s), together with the analysis execution instructions.

Discussion

As the name suggests, the main responsibility of the variability resolution mechanism (VRM) is to take a model with variability (provided as an external input to the overall optimization process) plus the encoding of a particular configuration (provided by the optimisation engine) and produce from that a fully resolved model in which all variation points have been replaced by the correct variant according to the configuration defined by the encoding. In principle, it is also conceivable to accept an encoding that represents an only partial configuration and then produce a partially resolved model, i.e. one with some variation points being resolved but other remaining in the model unchanged. In the context of optimisation, however, this is probably not useful because most analyses will require a fully resolved model.

One complication is with product line optimisation. The mostly likely solution is that the VRM will first resolve all optimisation variability according to the received encoding (resulting in a partially resolved model), and then generate multiple fully resolved models, one for each piece of the product line. It could then pass this set of models (and the production numbers of each one) to the COE so that analyses can take place across the whole product line. However, the precise details are yet to be fully established.

Additional care may be necessary to ensure that optimisation variability, once resolved, still produces a viable, meaningful model, in which the component/function interfaces etc align correctly. We may need to establish limits or constraints on what is possible to ensure this is the case, but this may require testing of many examples before potential problems can be identified.

4.3.4.4 Analysis Wrappers

The analysis tools are responsible for analysing the model according to the objectives given to the optimisation engine. Typically there would be a separate analysis for each objective, although in practical terms these analyses could be carried out by the same tool or in the same process. The results of the analysis/analyses are then fed back to the optimisation engine, which uses them to decide whether or not to retain that candidate model.

The interface to each analysis tool is provided by a particular analysis wrapper / plugin. Although they should (ideally) present a common interface to the COE, the interfaces to the analysis tools can vary, as each tool has different requirements. In some cases the wrapper may be the analysis tool, e.g. in the case of a custom cost analysis plugin. In other cases, it performs model transformations and passes it to an external tool (e.g. the HiP-HOPS plugin). It should not matter to the rest of the architecture whether an analysis plugin is self-contained or only a gateway to an external tool.

Inputs

- Each analysis takes place on a model (or set of models) with all variability resolved, i.e., each analysis looks at only one design candidate and product line variant at a time. It may also take extra parameters as needed per each analysis tool (e.g. to control the type or goal of the analysis).

Outputs

- Each analysis tool outputs its own results; ideally we would need some way of standardising the output to make it more easily read in by the optimisation engine, which uses the results to evaluate the design candidate. Ultimately we only need a single floating point value back.

Discussion

The integration of the analysis tools is perhaps the most challenging part of the whole process. Clear candidates for optimisation are safety/dependability analysis (with HiP-HOPS), timing analysis, and cost analysis (either simple summation with HiP-HOPS or more advanced cost analysis with a dedicated plug in). Other possible analyses include things like EMI, weight, temperature, and perhaps some kind of behavioural analysis.

Each of these presents its own particular set of problems. For example, safety analysis requires an error model be present in the EAST-ADL model, but it is not fully clear yet how the error model may be annotated with variability and stay synchronised with the nominal model. Cost analysis may work as a simple summation for a single model, but will not work across a product line as large volumes of components are cheaper than small volumes, and so the cost of a given component cannot necessarily be determined until after the whole product line has been analysed and the number of needed components of that type calculated. Thus for product line optimisation, a dedicated cost analysis plugin may be necessary. With timing analysis, performance may be an issue and so a subset of the full timing analysis may be necessary, e.g. only schedulability analysis, concentrating on information present in the EAST-ADL model (without need for tasks, runnable allocations etc).

In general, analysis should take place on a fully resolved model, i.e. one with no optimisation or product line variability left unconfigured. Thus to analyse a product line, we need to iteratively analyse the complete set of product line models for a given optimisation encoding. Then a loop takes place, cycling through all of the models in the set and analysing them. The VRM is responsible for generating this set, but the analysis wrappers/interfaces are responsible for handling the analysis and calculating the final result (e.g. a sum, product, minimum, maximum etc). In some cases this could be done sequentially, e.g. if the final result was just a total or a product; in other cases, as with cost analysis, it may be necessary to analyse them all as a set and then calculate the result at the end. Possible performance enhancements may include running multiple analyses simultaneously or perhaps finding some way of 'sharing' analysis results (e.g. if two models are substantially the same, perhaps some results could be cached), though this latter idea is likely to present substantial technical challenges.

4.3.5 Product Line Optimisation

One of the additional things we hope to include is the notion of product line optimisation. Variability in EAST-ADL was originally included to help cater for different product lines that may have different sets of features but share an overall commonality of architecture. It is hoped that the optimisation process can be extended to work with not just a single model of a product line, but all models of a product line simultaneously.

The dashed area in Figure 34 represents this 'inner optimisation loop'. The goal is that the optimisation variability will be separate to the product line variability, enabling the former to be resolved by the VRM but not the latter. This would result in a model being sent to the analysis wrappers that has had all of the optimisation variability resolved, to settle on a particular architecture to evaluate, but retains the product line variability showing which products have which features. The 'inner loop' would then iterate exhaustively through the different variations in the product line, evaluating each of them separately, and then combining them in an analysis-specific manner to return an overall heuristic evaluation value for that particular optimisation candidate. This process would therefore allow for different product line characteristics to be included in the consideration and evaluation of different architectural design candidates.

In most cases, the VRM would be responsible for iterating through the product line variability, but the analysis wrappers themselves would have to be able to combine the results for each variant in a meaningful fashion. For example, cost may be determined by the sum of each piece cost multiplied by the piece count for that piece, whereas for something like safety, the wrapper may use a maximum unavailability value as its result instead. This would have to be determined on a case-by-case basis, as each analysis would have different semantics and would require different calculations, but it may be possible to create a set of generic operators (e.g. 'sum', 'product', 'minimum', 'maximum') that can be applied to each objective.

Product line optimisation is still being investigated and is intended to be undertaken only after the general optimisation process has been successfully achieved. However, the architecture and tools are being planned in such a way that they can support the future extension to product line optimisation with minimum impact.

4.3.6 The Plan

The plan for developing the MAENAD Optimisation Architecture is as follows:

1. The first step is to construct an example model. This should feature only a subset of the full optimisation problem, so e.g. only dependability and cost information, perhaps some limited types of optimisation variability (e.g. only optionality and alternative implementations), and no product line variability. It should be modelled in Papyrus / EAST-ADL and include all configuration data for the optimisation. Its purposes is to serve as a testbed and proof of concept for an initial optimisation prototype, and to help us identify (and solve) and gaps or problems, whether in the theory or in the implementation.
2. In parallel (and afterwards), development of a prototype optimisation architecture should take place, i.e. the OSDM, VRM, COE, and a few analysis wrappers (all implemented as Papyrus-compatible plugins in Java). The initial goal would be to target dependability and cost analyses. Since there is no product line variability at this stage, HiP-HOPS can be used for both analyses, which is beneficial as we have full control over HiP-HOPS and can more easily make any necessary modifications than we could with a third-party tool. This prototype architecture would serve as a proof of concept.
3. The goal for M24 (end of second year) would be to use the new prototype implementation to fully analyse the example model. If successful, we can begin extending both with new features and functionality.
4. Moving into the third year, the goal would be to add remaining analysis tools (e.g. timing, other FEV-specific analyses), incorporate functionality to handle product line variability, and improve/streamline the UI and interfaces (e.g. providing a results GUI to automatically select an optimal model and configure it for viewing). In parallel we would want a full-size case study with enough variability and information present to test out all the new features we have implemented.

5 Implementation support for variability management

This section provides an overview of variability-related plugins used in MAENAD and their relevance for the analysis and synthesis concepts described above.

The tool support for variability is provided in three major parts:

1. **CVM Variability Management Framework:** core CVM with Model Editor, textual VSL Editor, etc.
2. **EAST-ADL Bridge for CVM Organizer:** EAST-ADL related functionality
3. **EPM Function Editor**

No. 1 is a stand-alone version of CVM covering the fundamental variability management concepts of EAST-ADL but without relying on the rest of EAST-ADL. No. 2 then integrates CVM into the EAST-ADL language and framework, using the EAST-ADL UML2 profile from CEA as a basis. This also makes it possible to use CVM together with Papyrus. No. 3 is an experimental domain specific editor for editing the core of EAST-ADL elements, especially *FunctionTypes* and their internal structure. It was initially implemented to experiment with various editing related aspects of variability management in EAST-ADL; for example, to see how the editing of a *FunctionType*'s public feature model and internal binding can best be supported with editing and visualisation functionality in a DSL editor.

The main intention of the above tools is to define system variability from a product line management perspective. This means the defined variations represent differences between distinct products being offered to the end-customer, for example a medium-class versus a luxury car or a passenger vehicle versus a truck. We refer to this kind of variations as *product line variability*.

When using EAST-ADL variability concepts and the corresponding tooling in the context of optimisation, however, there are several important differences to note. In this case, the variations of interest do not represent different variants of the complete system that are actually being built and offered to the end-customer but they represent different solutions for realising and implementing the system that seem worth considering during development. This kind of variation is referred to as *design space variability*.

The key defining difference between product line and design space variability can be summarised as follows: in the case of design space variability, ultimately only a single – the optimal – variant is of interest and will be built and delivered; in the case of product line variability, all variants are intended to be built and shipped.

To illustrate the relevance of this distinction let us consider a brief example. If we were to conduct a cost optimisation across system variants that represent product line variability, e.g. defining the differences between our model range from compact cars via medium-class models up to high-end luxury vehicles, we would always end up with the minimum configuration of the cheapest low-end model. Such a cost optimisation isn't of much use. This clearly shows that not all forms of variability are suited for optimisation. (Note that this is not to say that product line variability does not play any role in design space optimisation; for example, take rates in product line variability may provide relevant information for design space optimisations.)

The distinction between the above two forms of variability is extremely important for the future work in MAENAD, esp. WT 3.2, mainly because:

1. The difference and relation between product line and design space variability is not yet well understood in science and most often only one of these two forms of variation is investigated at a time. For example, it is not yet clear if these two classes of variations can always be clearly separated, or if there are cases where a particular variation in a system model is partly product line as well as design space oriented (and if so, how to deal with these cases).

2. The EAST-ADL language concepts for variability management were mainly targeted at product line variability. The same applies to the related tooling. Therefore, amendments to the concepts and tools might be required in order to provide a comprehensive solution for defining both product line and design space optimisation within a single system model.

For the MAENAD objective of utilising EAST-ADL's variability handling mechanisms in the context of optimisation, the above two items constitute several demanding research challenges. If solved, however, this will allow to deal with product line and design space variability in a concise and consistent form within a single framework, thus enabling a product-line aware optimisation of system models.

6 Summary and conclusions

This document describes some of the techniques, approaches, and algorithms being looked at and developed as part of the MAENAD project. They range from techniques to support safety analysis under the ISO 26262 standard, such as automated ASIL decomposition, to various analysis approaches to take advantage of the capabilities of the EAST-ADL language, to optimisation algorithms that can be used - with the aid of variability mechanisms defining the design space - to automatically optimise design models.

In terms of the project objectives:

- O1-1: Support in EAST-ADL for the safety process of ISO 26262 and representation of safety requirements.

Being able to support ISO 26262 is important for any automotive language or toolset. The initial safety process is briefly defined in Section 2. ISO 26262's approach to safety revolves to a significant degree around the use of ASILs - automotive safety integrity levels - which are used to define required levels of safety and which are defined as part of a hazard analysis of the system. EAST-ADL provides support for this process and for the definition of ASILs.

- O1-2: Automatic allocation of safety requirements (ASILs).

The standard also describes how ASILs can be decomposed across a system; this can be a difficult task if performed manually, and as a result, in MAENAD we have developed a concept for the automatic decomposition and allocation of safety requirements, as described in Section 2.4. It involves the use of the HiP-HOPS analysis tool and its FTA capabilities, and functions by assigning ASIL constraints to the results of the FTA (minimal cut sets, which are sets of basic events that can cause system failures). Once these constraints (known as Derived Safety Requirements) have been applied, the algorithm looks at how they interact to determine what possible assignments of ASILs to basic events will meet the overall requirements. This set of possible ASILs - known as Safety Requirement Allocations - can then be used to determine ASIL assignments for other system elements, such as input/output failures and whole components.

One disadvantage is that the algorithm is subject to combinatorial explosion, and thus as part of MAENAD, we are also investigating the idea of reformulating the ASIL decomposition process as an optimisation problem.

- O2-1: Dependability analysis of EAST-ADL models (with capabilities for multi-mode and temporal analysis of failures, together with integrated assessment of HW-SW design perspectives).

To support ISO 26262, general safety analysis of system models - such as FTA and FMEA - has to be possible. In MAENAD, these capabilities are primarily provided by the HiP-HOPS analysis tool, which is described in Section 3. As well as providing automatic analysis capabilities, it has also been extended to support the assessment of multiple perspectives (such as separated hardware and software concerns, as found e.g. in EAST-ADL). Section 3 also describes other tools and approaches that can be used for the purposes of dependability analysis.

- O2-2: Behavioural Simulation of EAST-ADL models.

Section 3 also describes other analysis approaches, including behavioural simulation approaches. These tools and techniques provide valuable capabilities when applied to EAST-ADL models,

enabling the modeller to form a dynamic view of the system design in addition to the primarily static view typically given by most dependability analyses.

- O2-3: Timing Analysis of EAST-ADL models.

The third type of model-based analysis included MAENAD is timing and performance analysis and this is also briefly described in Section 3. Further integration between EAST-ADL and MARTE is being investigated in MAENAD.

- O3-1: Extension of EAST-ADL with semantics to support multi-objective optimisation for product lines.

Section 4 describes the initial concepts towards the implementation of a multi-objective optimisation capability for EAST-ADL models. This is centred on the use of the HiP-HOPS optimisation engine, although the precise tool architecture has yet to be determined. Feasibility of the approach relies upon the use of variability mechanisms to define the design space and then the resolution of this variability - to achieve substitution of one element of a design for another - to enable evaluation of the design candidates. Variability implementation support is summarised in Section 5.

At present, the optimisation capabilities are relatively immature and are expected to be developed further later in the MAENAD project.

7 **References**

- [1] INTERNATIONAL ORGANIZATION FOR STANDARDIZATION: *Final Draft International Standard ISO/FDIS 26262*, 2010.
- [2] Bengtsson, J. and Wang, Y. 2004. *Timed Automata: Semantics, Algorithms and Tools*, Lecture Notes in Computer Science, Volume 3098/2004, 87-124
- [3] Zeigler, B.P., Praehofer, H., and Kim T.G. 2000. *Theory of Modelling and Simulation*, 2nd Edition, Academic Press, 2000.
- [4] Y.I. Papadopoulos and J.A. McDermid. (1999) "Hierarchically performed hazard origin and propagation studies," in Computer Safety, Reliability and Security, 18th International Conference, SAFECOMP'99, Toulouse, France, September, Proceedings, ser. LNCS, M. Felici, K. Kanoun, and A. Pasquini, Eds., vol. 1698. Springer, 1999, pp.139–152
- [5] Wasserman, A.I. 1990. *Tool integration in software engineering environments*. Proceedings of the International Workshop on Software Engineering Environments, 1990, Chinon, France. pp 137-149.
- [6] Coit, D. W. and Smith A.E. 1996a. Penalty guided genetic search for reliability design optimisation. *Computers and Industrial Engineering*. **30**(4), pp.895-904.
- [7] Coit, D W and Smith A E. 1996b. Reliability optimisation of series-parallel systems using a genetic algorithm. *IEEE Transactions on Reliability*. **45**(2), pp.254-260.
- [8] Salazar, D, Rocco C M, and Galvan B J. 2006. Optimisation of constrained multiple-objective reliability problems using evolutionary algorithms. *Reliability Engineering and System Safety*. **91**, pp.1057-1070.
- [9] Corne, D W, Jerram N R, Knowles J D, and Oates M J. 2001. PESA-II: Region-based Selection in Evolutionary Multiobjective Optimisation. *In: Proceedings of the Genetic and Evolutionary Computation Conference GECCO*. Morgan Kaufmann Publishers, pp.283-290.
- [10] Zitzler, E, Laumanns M, and Thiele L. 2001. SPEA2: Improving the strength Pareto evolutionary algorithm. *In: Proceedings EUROGEN 2001 Evolutionary methods for Design, Optimisation and Control with Applications to Industrial Problems*. Athens, Greece.
- [11] Deb, K, Pratap A, Agarwal S, and Meyarivan T. 2002. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*. **6**(2), pp.182-197.
- [12] Konak, A, Coit D.,W, and Smith A E. 2006. Multi-objective optimisation using genetic algorithms: A tutorial. *Reliability Engineering and System Safety*. **91**, pp.992-1007.
- [13] Kulturel-Konak S., Smith A.E, Coit D.W. 2003. Efficiently solving the redundancy allocation problem using tabu search. *IIE Transactions*. **35**, pp.515-526.
- [14] Liang, Y C and Smith A E. 2004. An ant colony optimisation algorithm for the redundancy allocation problem (RAP). *IEEE Transactions on Reliability*. **53**(3), pp.417-423.
- [15] Kim, H, Bae C., and Park S. 2004. Simulated annealing algorithm for redundancy optimization with multiple component choices. *In: Advanced Reliability Modelling, Proceedings of the 2004 Asian Internationals Workshop*. World Scientific, pp.237-244.
- [16] Parker, D.J. 2010. Multi-Objective Optimisation of Safety-Critical Hierarchical Systems. PhD Thesis, University of Hull, UK.
- [17] E. Frank, Reinhard Wilhelm, Rolf Ernst, Alberto L. Sangiovanni-Vincentelli, Marco Di Natale: Methods, Tools and Standards for the Analysis, Evaluation and Design of Modern Automotive Architectures. DATE 2008: 659-663

- [18] Palencia, J. C. and González Harbour, M. 1998. Schedulability Analysis for Tasks with Static and Dynamic Offsets. In Proceedings of the IEEE Real-Time Systems Symposium (December 02 - 04, 1998). RTSS. IEEE Computer Society, Washington, DC, 26.
- [19] Richter, K., Ziegenbein, D., Jersak, M., and Ernst, R. 2002. Model composition for scheduling analysis in platform design. In Proceedings of the 39th Annual Design Automation Conference (New Orleans, Louisiana, USA, June 10 - 14, 2002). DAC '02. ACM, New York, NY, 287-292
- [20] R. Henia, A. Hamann, M. Jersak, R. Racu, K. Richter, and R. Ernst. System level performance analysis - the SymTA/S approach. IEE Proceedings Computers and Digital Techniques, 152(2):148–166, March 2005.
- [21] UML-MAST project web site. <http://mast.unican.es/umlmast/>
- [22] J. Drake, M. Harbour, J. Gutierrez, P. Lopez, J. Medina, J. Palencia. Modeling and Analysis suite for Real-time Applications (MAST 1.3.7). <http://mast.unican.es/umlmast/>