Automated Configuration Synthesis for Machine Learning Models: A git-Based Requirement and Architecture Management System

Abdullatif AlShriaf
Chalmers and University of Gothenburg
Gothenburg, Sweden
Email: shriaf@chalmers.se

Hans-Martin Heyn
Chalmers and University of Gothenburg
Gothenburg, Sweden
Email: heyn@chalmers.se

Eric Knauss
Chalmers and University of Gothenburg
Gothenburg, Sweden
Email: eric.knauss@cse.gu.se

arXiv:2404.17244v1 [cs.SE] 26 Apr 2024

I. Background

The design of complex distributed systems typically follows a hierarchical process, supported by highly specialized views for decomposing the design task. Requirements and architecture often evolve simultaneously, requiring an architectural framework that supports integrated and collaborative design, including non-functional requirements and quality views. The framework must ensure the traceability of design decisions in order to build safety cases. Integrating requirements into software development is vital for aligning intended functionality with implemented code. However, extracting data from semi-formal requirements and maintaining alignment poses challenges due to its ambiguity and variability making extracting consistent information challenging. Aligning these requirements with other project artifacts can also be difficult due to interpretation differences, often requiring manual effort and leading to complexity and potential inconsistencies in development [1].

II. Introduction

We propose GRAMS[1] as a git-based Requirements and Architecture Management System which we demonstrate can automate configuration extraction, reduce manual effort, enhance configurability, and simplify runtime configuration management. We will show that it improves traceability by establishing clear links between requirements, configurations, and code, and facilitates maintenance through a structured approach. GRAMS is intended for use by i) machine learning engineers and architects tasked with designing the structure of ML-based software systems, and ii) system administrators/MLOps engineers responsible for deploying and managing the runtime environment of software systems. GRAMS builds upon the open-source tool *T-Reqs* for managing textual requirements in git [2]. It extends T-Reqs by implementing a structured architectural view system based on the concept of a compositional architectural framework for the development of AI-enabled software systems (Figure 1, [3]). Using a schema-type template, GRAMS compiles a machine-readable YAML *intermediary* file to communicate AI model specifications. This file can then be forwarded and converted as part of a deployment-toolchain, particularly for example with the purpose of AI model optimization.

[1] https://gitlab.com/latiif/configgen

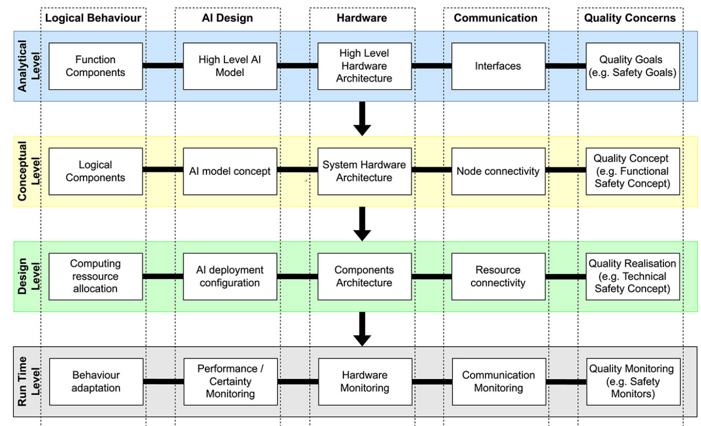

Figure 1. The VEDLIoT Architectural Framework [3].

III. Context: Compositional Architectural Framework

The system design approach supported by GRAMS follows a scalable and compositional architectural framework that provides flexibility to manage requirements in software projects of varying complexity. The architectural framework organizes quality concerns, including ethical considerations such as explainability and security. Supporting middle-out design, GRAMS integrates existing design decisions, enhancing alignment between requirements and architecture [3]. Guiding rules ensure traceability, vital for managing requirements throughout development. The framework's adaptability to diverse abstraction levels enables GRAMS to cater to varied project needs.

IV. T-Reqs: Text-based Requirements Management System

GRAMS extends T-Reqs, an open-source tool for the management of textual requirements in git. T-Reqs is developed at

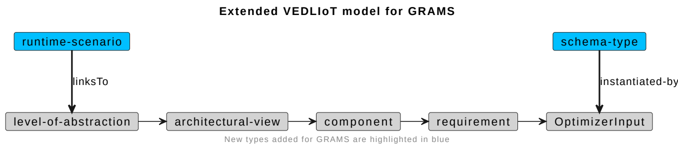

Figure 2. Extended Architectural Framework for GRAMS: New types added for GRAMS are highlighted in blue

```
1: for each level-of-abstraction that linksTo to a runtime-scenario do
2:     paths ← elements start with level-of-abstraction && end with OptimizerInput
3:     valid ← []
4:     for each path in paths do
5:         jsonBody ← JSON object inside the OptimizerInput
6:         jsonSchema ← JSON schema in schema-type linked by instantiatedBy
7:         if jsonBody valid against jsonSchema then
8:             Annotate OptimizerInput with schema-type
9:             valid ← valid + path
10:        end if
11:    end for
12: end for
```

Figure 3. Framework traversal algorithm for GRAMS in pseudo-code.

Chalmers and University of Gothenburg and used in industrial settings, including at Ericsson [4]. T-Reqs utilizes a TTIM, or Type and Trace Information Model, that extends the functionality of the basic Traceability Information Model (TIM) by incorporating type information. While a TIM primarily focuses on representing trace-links and relationships between abstract trace artifacts [5], the TTIM defines types of traced elements and their connections, which can be optional or required [2]. T-Reqs operates as a static analysis tool on the model described by the TTIM.

V. Applying GRAMS

First, GRAMS identifies explicit runtime scenarios, then heuristically traverses the architectural framework, and finally locating configurations/schemes and mapping them to the identified runtime scenarios. In essence, GRAMS executes the TTIM, as illustrated in Figure 2, along with a traversal algorithm which is outlined in Figure 3. Finally, the aggregation capabilities of the system are presented in generic formats such as YAML or PlantUML. Leveraging the architectural structure laid by the framework, we introduce two new elements, as shown in Figure 2:

- *runtime-scenario*: matches particular models to specific contexts and links to a level of abstraction which GRAMS will scan and traverse all the way to an AI-model optimiser input referred to as *OptimizerInput*.
- *schema-type*: Annotates the *OptimizerInput* and contains the actual schema for the JSON object present in the instantiating *OptimizerInput*.

Multiple *runtime-scenario*s can exist in the framework. From a single *runtime-scenario*, multiple *OptimizerInputs* can be traced. The traversal follows the algorithm outlined in Figure 3.

VI. Quality Assurance Checks in GRAMS

To ensure output quality and usability, GRAMS conducts the following checks:

- Consistency is crucial to prevent conflicts and maintain coherence [6]. The initial ”dumb” check ensures alignment with the framework's meta-model (see Section III).
- GRAMS validates the internal correctness of the configuration-relevant requirements specification by checking JSON schema compliance with *schema-type* and *OptimizerInput*.
- GRAMS verifies semantic equivalence between relevant parts of the configuration schema and *OptimizerInput* schemas to confirm configuration relevance to *runtime-scenarios*.

VII. Discussion & Conclusion

Having demonstrated the feasibility of leveraging the architectural framework for specific runtime configurations, the necessity for explicit information about dynamic property schema becomes apparent. Integrating architecture with requirement context facilitates traceable configuration generation, anchored in system design and requirements. GRAMS' scalability and reliance on the architectural framework provide a smoother topology for traversing all possible variable inputs. By specifying the requirements for the system and each model in every runtime scenario, we enable traceability and support safety argumentation. Additionally, generating configurations for the model optimizer from these requirements facilitates frequent changes and iterative development, necessitating the continuous upkeep of requirements to ensure alignment between the optimized models and their specifications. Yet, questions arise: Is the intermediary format sufficient for tracing relevant properties and generating final configurations? Are GRAMS checks overly strict, and what adjustments are needed? Additionally, should GRAMS be extended to generic systems not reliant on the compositional architectural framework? These queries prompt further examination.

References

[1] A. Zaki-Ismail, M. Osama, M. Abdelrazek, J. Grundy, and A. Ibrahim, “Requirements formality levels analysis and transformation of formal notations into semi-formal and informal notations,” ser. Proc. of the Int. Conf. on Software Engineering and Knowledge Engineering, SEKE. KSI Research Inc., 2021, pp. 303–308.

[2] E. Knauss, G. Liebel, J. Horkoff, R. Wohlrab, R. Kasauli, F. Lange, and P. Gildert, “T-reqs: Tool support for managing requirements in large-scale agile system development,” in *2018 IEEE 26th International Requirements Engineering Conference (RE)*, 2018, pp. 502–503.

[3] H.-M. Heyn, E. Knauss, and P. Pelliccione, “A compositional approach to creating architecture frameworks with an application to distributed ai systems,” *Systems and Software (JSS)*, vol. 198, 2023.

[4] N. Theodórsdóttir, Audur Katarína Ojensa, *Evaluation of Text–Based Requirements Engineering Tools*. Master Theses at Chalmers University of Technology, 2022, available online: https://odr.chalmers.se/items/678b7cf4-9d5f-4ba0-a34e-4eda0965ecd8.

[5] O. Gotel, J. Cleland-Huang, J. H. Hayes, A. Zisman, A. Egyed, P. Grünbacher, A. Dekhtyar, G. Antoniol, J. I. Maletic, and P. Mäder, “Traceability fundamentals,” in *Software and Systems Traceability*, J. Cleland-Huang, O. Gotel, and A. Zisman, Eds. Springer, 2012, pp. 3–22.

[6] S. Lauesen, *Software Requirements-Styles and Techniques*. Addison Wesley, 01 2002.

````
 1 <treqs-element id="id2" type="schema-type">
 2 The  optimizer that describes ethernet latency has the following schema type:
 3
 4 ```json
 5 {
 6   "$schema": "http://json-schema.org/draft-07/schema#",
 7   "title": "EthernetLatency",
 8   "type": "object",
 9   "properties": {
10     "value": {
11       "type": "number",
12       "description": "Latency value in milliseconds"
13     },
14     "unit": {
15       "type": "string",
16       "description": "Unit of measurement for latency (e.g., milliseconds)"
17     }
18   },
19   "required": ["value"]
20 }
21 ```
22 <treqs-link type="instantiated-by" target="f3f65cbc9d9711edac23f218984a2e78"/>
23 </treqs-element>
````

Figure 4. A *schema-type* element that includes the JSON schema that describes an *OptimizerInput* property related to *ethernet latency*.

APPENDIX

A. Instantiating the Extended Architectural Framework

An important component of the extended architectural framework outlined in Figure 2, is the *schema-type*. An example of which is shown in Figure 4.

B. Target System Configuration Schema

The JSON schema for the configuration of the target system is given to GRAMS, not only for validation checking, but also to be included in the output of GRAMS. Figure 5 illustrates such a sample schema not only for the parts of the configuration that are to be populated from the requirement model, but for all properties. GRAMS traverses the schema and identifies and performs checks on the relevant properties.

C. Intermediary Output Format

The intermediarey output format is YAML and it has two main sections:

- *config_schema* defines a JSON schema for configuring the target system, specifying schemas for extracting properties from the architectural framework. These properties may be nested, and the schema serves as input to GRAMS. Additionally, it's included in the output, enabling direct use of the intermediary output by the target system or by a tool for specific adjustments.
- *optimizer_inputs* contains a list of elements specifying *OptimizerInput* properties, for each element its *file_name*, *label*, *placement*, *treqs_type*, and *uid*, along with traceability path pointing to requirements going all the way up to a *runtime_scenario*. Additionally, it includes a schema defining the JSON schema for the property and value indicating the specific constraint in milliseconds. This structure repeats for each requirement.

D. GRAMS Output

GRAMS output can be displayed in two formats: PlantUML or YAML, the latter is intended for consumption of the target system (or other intermediary tooling) whereas the PlantUML version is for providing an overview. A sample PlantUML and

```
 1 {
 2   "$schema": "http://json-schema.org/draft-07/schema#",
 3   "type": "object",
 4   "properties": {
 5     "LatencyConstraints": {
 6       "type": "object",
 7       "properties": {
 8         "EthernetLatency": {
 9           "type": "object",
10           "properties": {
11             "value": {
12               "type": "number"
13             },
14             "unit": {
15               "type": "string"
16             }
17           },
18           "required": [
19             "value"
20           ]
21         },
22         "ModelLatency": {
23           "type": "object",
24           "properties": {
25             "value": {
26               "type": "number"
27             },
28             "unit": {
29               "type": "string"
30             }
31           },
32           "required": [
33             "value"
34           ]
35         }
36       },
37       "required": [
38         "ModelLatency",
39         "EthernetLatency"
40       ]
41     }
42   },
43   "required": [
44     "LatencyConstraints"
45   ]
46 }
```

Figure 5. Input to GRAMS: A sample JSON schema describing the configuration of a target system that has two properties: *ethernet latency* and *model latency*.

YAML outputs are hosted as snippets, whereas Figure 6 is the visualization on the PlantUML code. A recording of GRAMS' YAML output is available.

E. GRAMS Checks

As outlined in VI, the following recordings demonstrate how GRAMS detects and reacts to the various checks:

- Meta-model inconsistencies
- Internal schema correctness
- Semantic equivalence between configuration schema and framework's schema.

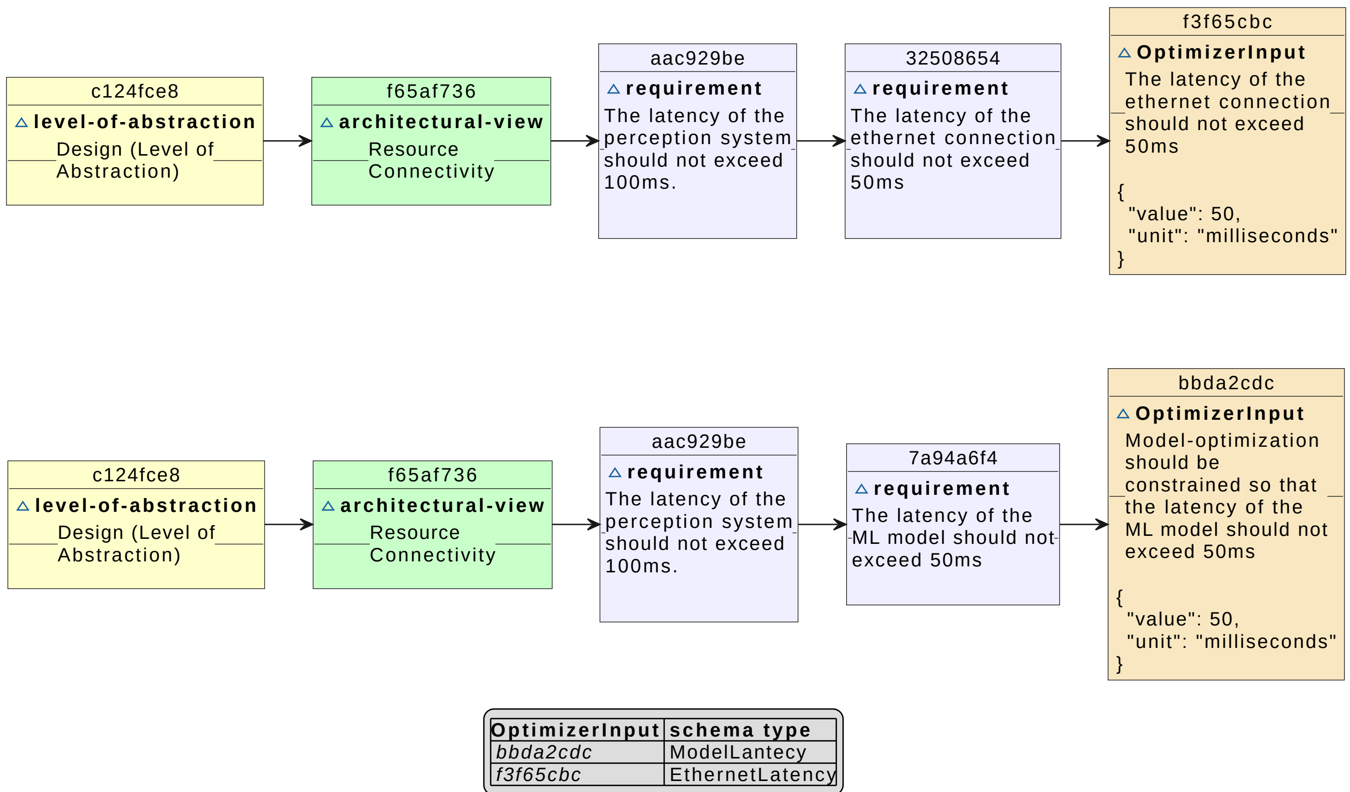

OptimizerInput	schema type
bbda2cdc	ModelLantecy
f3f65cbc	EthernetLatency

Figure 6. Rendering of GRAMS PlantUML output which illustrates two paths stemming from the same *runtime-scenario* down to an *OptimizerInput* property. The legend shows the types of the two identified properties.