

Formalizing and validating properties in Asmeta with Large Language Models (Extended Abstract)

Andrea Bombarda¹[0000-0003-4244-9319], Silvia Bonfanti¹[0000-0001-9679-4551],
Angelo Gargantini¹[0000-0002-4035-0131], and Nico
Pellegrinelli¹[0009-0000-4944-6845]

University of Bergamo, Bergamo, Italy {andrea.bombarda, silvia.bonfanti,
angelo.gargantini, nico.pellegrinelli}@unibg.it

Abstract. Writing temporal logic properties is often a challenging task for users of model-based development frameworks, particularly when translating informal requirements into formal specifications. In this paper, we explore the idea of integrating Large Language Models (LLMs) into the *Asmeta* framework to assist users during the definition, formalization, explanation, and validation of temporal properties. We present a workflow in which an LLM-based agent supports these activities by leveraging the *Asmeta* specification and the feedback produced by the model checker. This work serves as a proof of concept that illustrates the feasibility and potential benefits of such an integration through representative examples.

Keywords: *Asmeta*, Temporal Logic, Large Language Models

1 Introduction

Modern software and system engineering projects are increasingly complex, often involving heterogeneous components, tight safety constraints, and short development cycles. Traditional development approaches, which rely heavily on manual coding and late-stage testing, struggle to ensure early validation, traceability, and consistency between design and implementation. This leads to higher costs and risks, especially in safety-critical domains such as automotive, aerospace, and healthcare. Model-Based Development (MBD) offers an effective alternative. By relying on executable models, it enables early simulation and verification of system behavior, reduces ambiguity in requirements through formal or semi-formal representations, supports automatic code generation to ensure alignment between design and implementation, and enhances traceability and compliance with industry standards. A variety of tools support MBD, and, in this work, we focus on the *Asmeta* framework [1,3]. The framework allows developers to work with Abstract State Machines (ASMs) [4], which extend Finite State Machines (FSMs) by replacing unstructured control states with states that can contain arbitrarily complex data.

2 Objectives

The **Asmeta** framework provides a suite of tools that support developers throughout the software lifecycle, including design, development, and operation. In this work, we concentrate on the design phase, and in particular on verification activities, which are carried out using the model checker integrated into **AsmetaSMV** [2].

Starting from the functional requirements, **Asmeta** enables users to model a system using the **AsmetaL** language. Users can embed Computation Tree Logic (CTL) and Linear Temporal Logic (LTL) properties directly within the ASM model, expressing them over the **Asmeta** signature. The **AsmetaSMV** tool, then, automatically translates the ASM model (including both the system model and the LTL/CTL properties) into a model of the symbolic model checker NuSMV [6], which is used for verification.

Based on feedback collected from **Asmeta** users, particularly students enrolled in courses where **Asmeta** is taught, we observed that many of them encounter difficulties when writing LTL and CTL properties. While formulating properties in natural language is generally straightforward, challenges arise when translating these descriptions into the corresponding **AsmetaL** expressions.

To address this issue, we propose integrating Large Language Models (LLMs) into the workflow. We have identified, after analyzing existing works, at least four ways in which LLMs can assist the designer in the verification of TL properties.

***O1.** Assisting during the definition of the property in natural language:* Given an **Asmeta** specification, LLMs can be leveraged to interpret the domain context and automatically derive properties, expressed in natural language, that are valid within that context. These properties are neither intended to be exhaustive nor universally guaranteed; rather, they serve to improve the user’s understanding of the domain and to potentially reveal constraints or assumptions that were implicit in the original requirements. There are already some experiments proposing Retrieval-Augmented Property Generation in other contexts, like PropertyGPT [7].

***O2.** Helping the designer to write the corresponding LTL/CTL properties:* Given an **Asmeta** specification that includes the system behavior and properties expressed in natural language, the LLM automatically translates these properties into their corresponding LTL or CTL formulations within the **Asmeta** specification. By doing so, we aim to reduce the gap between informal requirement descriptions and formal verification artifacts, ultimately improving usability and lowering the entry barrier for new users.

Furthermore, the generated properties must be formulated with awareness of, and consistency with, the functions defined in the **Asmeta** specification, thereby ensuring that the resulting LTL/CTL formulas are consistent with the operational semantics of the model. This implies that the LLM must be context-aware: it must recognize the roles and types of functions, and incorporate these elements appropriately into temporal operators.

O3. Explanation of the TL properties using natural language: Formalizing requirements enables systematic reasoning about inconsistencies, the detection of ambiguities, and the identification of critical issues in system models. Temporal logic formulae are a natural choice for specifying requirements related to desired system behaviors. However, understanding and effectively using temporal logic demands a strong formal background. Therefore, there is a need for approaches that make temporal logic formulae more interpretable for engineers, domain experts, and other stakeholders involved in the development process. For this goal, we assume that LLMs can be used to translate TL properties back to natural language. For this scope, we plan to adopt (or adapt, if necessary) the technique presented in [5].

O4. Validation of the TL properties using witnesses and counterexamples: When a property fails verification, **AsmetaSMV** produces a counterexample that can be automatically exported as an **Avalla** scenario. Likewise, successful executions could be exported as scenarios that highlight witnesses. Large Language Models (LLMs) can support developers’ understanding by translating **Avalla** scenarios into natural language, thereby clarifying their meaning and describing the behavior they implement.

We report on preliminary investigations focused on outlining the process to be adopted for integrating an LLM-based assistant into the **Asmeta** framework. The goal at this stage is to define the process required to incorporate an LLM-based assistant into the **Asmeta** framework.

3 Methodology and Walkthrough Demonstration

This work follows a multi-step methodology aimed at designing, integrating, and empirically evaluating an LLM-based assistant for the automatic generation of temporal properties in **Asmeta** specifications. Since this is a preliminary study, our focus is on defining the overall process, identifying key challenges, and conducting initial feasibility analyses rather than delivering a fully operational framework.

The LLM-based assistant, whose workflow is reported in Figure 1, should be able to generate and explain LTL and CTL properties that are not only syntactically correct, but also meaningful with respect to the given **Asmeta** model. This implies that the assistant must interpret the signature of the specification, understand the types of functions, and generate formulas or explanations that consistently refer to elements defined by the model.

The user begins by providing both an **Asmeta** specification and the informal requirements of properties to be formalized, within a specifically designed prompt. The assistant, then, extracts relevant signatures from the **Asmeta** specification and translates the requirements into an LTL or CTL formula embedded into the **Asmeta** specification, or provides explanations complying with the requested tasks and using the same functions as in the **Asmeta** specification.




Fig. 1. Our solution workflow.

<pre> asm Clock import StandardLibrary signature: domain Second subsetof Integer domain Minute subsetof Integer domain Hour subsetof Integer monitored signal: Boolean controlled sec: Second controlled min: Minute controlled h: Hour definitions: domain Second = {0 : 59} </pre>	<pre> domain Minute = {0 : 59} domain Hour = {0 : 23} macro rule r_IncMinHours = par if min = 59 then h := (h + 1) mod 24 endif min := (min + 1) mod 60 endpar main rule r_Main = if signal then par if sec = 59 then r_IncMinHours[] endif sec := (sec + 1) mod 60 endpar endif default init s0: function sec = 0 function min = 0 function h = 0 </pre>
--	---

Listing 1. Asmeta specification of a simple Clock

We expect users to write requirements by following the EARS (Easy Approach to Requirements Syntax) [8], a structured way to write requirements so they are clear, consistent, and less ambiguous.

In the following, we describe how we envision that our approach would help in each of the four objectives described in Section 2. Our demonstration is based on direct interactions with the ChatGPT 5.2 Instant model and aims at checking the feasibility of our approach.

Let us consider the Asmeta specification of a simple clock as in Listing 1. For the objective **O1**, users may ask the LLM-Based Agent to list the *three* most important properties guaranteed by the Asmeta specification. In this case, the LLM replies with the three following properties (each with its own description):

- Time variables always stay within valid clock ranges;
- Correct carry-over from seconds to minutes and hours;
- Time advances only when the signal is true.

Although some of these properties may be straightforward, providing a natural language explanation can help lower the barrier to understanding Asmeta specifications.

For objective **O2**, users may ask the LLM-Based Agent to encode in CTL a property verifying that when the min function reaches the value 59, it is set to 0 in the next state. With this prompt, the LLM-Based Agent replies with the following CTL property:

<pre>scenario ClockScenario load Clock.asm set signal := true; step;</pre>	<pre>check sec = 1 and min = 0 and h = 0; set signal := false; step; check sec = 1 and min = 0 and h = 0;</pre>
--	---

Listing 2. Avalla scenario for the Clock specification

AG (min = 59 implies AX (min = 0))

A property such as the one generated by our LLM-Based Agent must contain only valid operators and functions. While LLMs can be fed with an **Asmeta** specification, it may happen that the LLM generates a syntactically wrong property. For this reason, as shown in Figure 1, we integrate our solution with **AsmetaSMV**. Once the property is generated, the model checker’s feedback is leveraged. If errors are detected, the model checker’s output is provided together with the prompt to correct the property formulation.

For objective **O3**, suppose that the considered **Asmeta** specification contains the property given before. Users may leverage our LLM-Based Agent to explain in natural language what the property means. In this example, the output of our Agent is the following:

In every reachable state of the system, if the value of min is 59, then in all possible next states the value of min will be 0.

Finally, for the objective **O4**, if the Avalla scenario in Listing 2 is available (e.g., as the output of the counterexample generation), users may be interested in having its high-level to better understand the reasons behind a possible failure. In this case, the LLM-Based Agent provides the following explanation:

The scenario verifies that the clock advances by one second when the signal is true, and stops advancing when the signal is false, preserving the current time.

4 Conclusion

In this paper, we explored the idea of integrating Large Language Models into the **Asmeta** framework to support users in writing, understanding, and validating temporal logic properties. Motivated by the difficulties encountered by practitioners when translating informal requirements into LTL and CTL specifications, we outlined a conceptual workflow in which an LLM-based agent assists the designer across multiple verification-related tasks, ranging from property elicitation in natural language to the interpretation of witnesses and counterexamples.

Through simple examples, we showed that LLMs can plausibly bridge the gap between informal requirements and formal verification artifacts, while remaining consistent with the semantics of **Asmeta** specifications and the feedback provided by the model checker. Future work will focus on turning this conceptual workflow into a concrete implementation within **Asmeta**.

References

1. Arcaini, P., Bombarda, A., Bonfanti, S., Gargantini, A., Riccobene, E., Scandurra, P.: The ASMETA Approach to Safety Assurance of Software Systems. In: *Logic, Computation and Rigorous Methods*, pp. 215–238. Springer (2021). https://doi.org/10.1007/978-3-030-76020-5_13
2. Arcaini, P., Gargantini, A., Riccobene, E.: *AsmetaSMV: A Way to Link High-Level ASM Models to Low-Level NuSMV Specifications*, p. 61–74. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-11811-1_6
3. Bombarda, A., Bonfanti, S., Gargantini, A., Riccobene, E., Scandurra, P.: *ASMETA Tool Set for Rigorous System Design*, p. 492–517. Springer Nature Switzerland (Sep 2024). https://doi.org/10.1007/978-3-031-71177-0_28
4. Börger, E., Stärk, R.: *Abstract State Machines*. Springer Berlin Heidelberg (2003). <https://doi.org/10.1007/978-3-642-18216-7>
5. Cherukuri, H., Ferrari, A., Spoletini, P.: *Towards Explainable Formal Methods: From LTL to Natural Language with Neural Machine Translation*. In: *Requirements Engineering: Foundation for Software Quality*. pp. 79–86. Springer International Publishing, Cham (2022)
6. Cimatti, A., Clarke, E., Giunchiglia, E., et al.: *NuSMV 2: An OpenSource Tool for Symbolic Model Checking*, p. 359–364. Springer Berlin Heidelberg (2002). https://doi.org/10.1007/3-540-45657-0_29
7. Liu, Y., Xue, Y., Wu, D., et al.: *PropertyGPT: LLM-driven Formal Verification of Smart Contracts through Retrieval-Augmented Property Generation*. In: *Proceedings 2025 Network and Distributed System Security Symposium*. Internet Society (2025). <https://doi.org/10.14722/ndss.2025.241357>
8. Mavin, A., Wilkinson, P., Harwood, A., Novak, M.: *Easy Approach to Requirements Syntax (EARS)*. In: *2009 17th IEEE International Requirements Engineering Conference*. pp. 317–322 (2009). <https://doi.org/10.1109/RE.2009.9>