GRuM – A Flexible Model-Driven Runtime Monitoring Framework and its Application to Automated Aerial and Ground Vehicles

Michael Vierhauser[a], Antonio Garmendia[b], Marco Stadler[a], Manuel Wimmer[c], Jane Cleland-Huang[d]

[a]*Johannes Kepler University Linz, LIT Secure and Correct Systems Lab, Linz, 4040, Austria*
[b]*Universidad Autónoma de Madrid,, Departamento de Ingeniería Informática, Madrid, 28049, Spain*
[c]*Johannes Kepler University Linz, Institute for Business Informatics - Software Engineering & CDL-MINT, Linz, 4040, Austria*
[d]*University of Notre Dame, Department of Computer Science and Eng., Notre Dame, 46617, IN, United States*

Abstract

Runtime monitoring is critical for ensuring safe operation and for enabling self-adaptive behavior of Cyber-Physical Systems (CPS). Monitors are established by identifying runtime properties of interest, creating probes to instrument the system, and defining constraints to be checked at runtime. For many systems, implementing and setting up a monitoring platform can be tedious and time-consuming, as generic monitoring platforms do not adequately cover domain-specific monitoring requirements. This situation is exacerbated when the System under Monitoring (*SuM*) evolves, requiring changes in the monitoring platform. Most existing approaches lack support for the automated generation and setup of monitors for diverse technologies and do not provide adequate support for dealing with system evolution. In this paper, we present *GRuM* (**G**enerating CPS **Ru**ntime **M**onitors), a framework that combines model-driven techniques and runtime monitoring, to automatically generate a customized monitoring platform for a given *SuM*. Relevant properties are captured in a Domain Model Fragment, and changes to the *SuM* can be easily accommodated by automatically regenerating the platform code. To demonstrate the feasibility and performance we evaluated *GRuM* against two different systems using Turtle-Bot robots and Unmanned Aerial Vehicles. Results show that *GRuM* facilitates the creation and evolution of a runtime monitoring platform with little effort and that the platform can handle a substantial amount of events and data.

Keywords:
Cyber-Physical Systems, Runtime Monitoring, Model-Driven Engineering

Email addresses: michael.vierhauser@jku.at (Michael Vierhauser), antonio.garmendia@jku.at (Antonio Garmendia), marco.stadler@jku.at (Marco Stadler), manuel.wimmer@jku.at (Manuel Wimmer), janeclelandhuang@nd.edu (Jane Cleland-Huang)

1. Introduction

Robotic systems are increasingly used in the context of shop floor automation [1, 2], as autonomous vehicles for medical device delivery [3], and search-and-rescue [4] operations. Such Cyber-Physical Systems (CPS) exhibit tight integration between hardware and software components and frequently interact with humans. This in turn introduces a number of safety concerns that need to be mitigated [5, 6], resulting in the need for customized and system-specific runtime monitoring support. For example, when Unmanned Aerial Vehicles (UAVs) are engaged in search-and-rescue flights in close proximity to humans, or when robots operate on a factory floor, precautionary measures need to be taken to ensure that the CPS adheres to its specified requirements and operates within its predefined safety envelope. Therefore, support is required for monitoring diverse properties at runtime [7]. However, the heterogeneity of hardware and software components within a CPS means that creating and implementing monitors, and subsequently collecting, processing, and checking the required data is often an arduous and time-consuming task.

Runtime information is typically collected from the System under Monitoring (*SuM*), through source code instrumentation [8, 9], using dedicated data-buses [10, 11], or other collection services. It is then processed at runtime to check constraints to ascertain whether the system is behaving according to its specified requirements, or if deviations from expected behavior have occurred. While off-the-shelf monitoring approaches support application performance monitoring [9] or checking Service Level Agreements [12, 13], they often provide inadequate support for instrumenting custom systems, or defining complex temporal or structural constraints. The problem is exacerbated as CPS are typically long-running, with ongoing maintenance and evolution activities that require both, the data collection mechanisms and constraints to be updated and adapted. If the monitoring infrastructure does not co-evolve with the *SuM*, constraints may become stale, and data outdated due to changes in the underlying components. Few approaches have addressed the challenge of automatically generating customized, system-specific, runtime monitoring solutions and their maintenance and evolution support. In closely related work, Model-Driven Engineering (MDE), which automatically generates source code from system models, has previously been used to model entire systems, including their runtime properties [14, 15]. When a new feature or functionality is introduced, the model is updated, and its respective code regenerated. As a result, MDE has been shown to improve productivity by enabling developers to specify a system at a higher level of abstraction. However, the application of MDE has fallen short of enabling the generation and evolution of a complete monitoring platform [16].

To address these shortcomings, we propose and evaluate *GRuM*, a model-driven framework for **G**enerating CPS **Ru**ntime **M**onitors. We aim to enable (i) the generation of a customized monitoring platform with support for data collection and analysis, which (ii) can be readily extended and updated when changes to the monitored system occur. *GRuM* provides a lightweight monitoring solution that only requires modeling

those parts of the system that are relevant for the monitoring infrastructure. As a result, it can be applied to diverse projects, whether or not MDE is adopted as the primary development paradigm. This paper builds upon our earlier work [17, 18], in which we collected requirements for model-driven monitoring and derived an initial architecture. We significantly extend that work by introducing a refined architecture, extending our weaving and meta-model, and providing code generators for the monitoring platform. Furthermore, we have greatly extended the automation support for automatically generating the platform and monitors, and include a thorough evaluation of *GRuM* using two diverse systems. More specifically, we evaluate our approach against a system for deploying UAVs for search-and-rescue and a second system that uses a set of mobile robots to assess indoor air quality.

The contributions of this work are as follows: First, we present a monitoring meta-model that enables us to generate a custom monitoring platform with a runtime model, instead of implementing it manually. This provides a higher automation potential, for more efficient development of monitoring solutions. Second, probes are automatically generated for the target technology. Given this platform, we can leverage different types of (off-the-shelf) constraint engines depending on the monitoring needs, as shown in our evaluation (using a CEP Engine for temporal constraints and event sequences) without the need to reimplement other parts. Finally, once a probe generator for a specific technology has been implemented, evolving the monitoring framework alongside the *SuM* is greatly simplified and only requires a few steps to update the model and automatically regenerate the platform[1].

The remainder of the paper is structured as follows: example describes open challenges. framework presents an overview of the *GRuM* framework, while application describes its application and implementation. In evalmethod we first describe our evaluation objectives and setup, and in Sections 6 to 8 we report on our evaluation results for two real-world systems and discuss results and threats to validity. Finally, relwork describes related work and, conc draws conclusions and proposes future directions.

2. Monitoring Challenges and Motivating Example

Runtime monitoring plays a pivotal role in checking and ensuring that a system operates safely within its predefined operational capabilities and that it adheres to its specified constraints related to both functional and non-functional requirements related to performance, safety, and other quality concerns. It plays an intrinsic role in creating Digital Shadows and Digital Twins [19, 20], and provides support for the MAPE-K loop in self-adaptive systems [21, 22].

[1] All models and source code, as well as the generated monitoring infrastructures, are available on GitHub as an open-source system at `https://github.com/LIT-Rumors/grum-public`

However, the field of runtime monitoring is quite diverse with approaches that have been designed to support various purposes, technologies, and types of system architectures [23, 24]. Examples include service-based systems [13, 25], application performance monitoring [9], and goal-oriented monitoring approaches [26, 27]. The fact that existing approaches are typically customized for a specific type of system, application domain, or support particular types of constraints, has further hampered the broader adoption of runtime monitoring frameworks. As a result, significant upfront investment is required to implement and maintain custom monitoring infrastructures, often without inbuilt support for maintenance and evolution after deployment. In previous work, Rabiser *et al.* [28] investigated several different monitoring frameworks and ultimately derived a reference architecture, identifying three common components: *Monitoring Setup* related to defining monitors, generating probes, and instrumenting a system, *Monitoring Execution* concerning data collection, processing and constraint checking, and finally *Monitoring Support* targeting additional functionality such as data persistence and external applications.

The continuous maintenance, enhancement, and hence evolution, that complex CPS are subject to, has a significant impact on any monitoring infrastructure that collects and analyses data from a system. Runtime monitoring frameworks need to embrace these changes and co-evolve with the *SuM*, providing support for updating, and maintaining the monitoring framework itself as part of the system development process [29, 30].

These challenges are illustrated in *Dronology*, an open-source, multi-UAV system [31] with support for planning and executing UAV missions. Establishing a monitor for Dronology requires in-depth knowledge of the system in order to identify and collect necessary data, establish a monitoring platform that aggregates data from different UAVs and then stores the data for subsequent analysis, and ultimately to select and configure a suitable constraint engine to analyze the data and provide evaluation results. In the case of Dronology, instrumentation is needed to collect UAV data (e.g., GPS coordinates, attitude, mission status), transmit the data to a Ground Control Station which is responsible for aggregating, analyzing, and persisting the data, and developing a dashboard for visualizing data to the UAV operators. In addition, a constraint engine needs to be configured to check constraints, such as whether the UAVs remain within their designated altitude band, avoid no-fly zones, and are flight-worthy.

Ongoing changes to both, the hardware and software can adversely affect the monitoring system. For example, if a new water sampling capability were introduced to Dronology, it would require additions and modifications to the hardware and software components. New downward-facing sensors would be required to ensure that the UAV maintains a stable distance from the water during the collection process. Furthermore, as light reflecting from the water could cause dangerous altitude fluctuations, the monitoring system would need to be updated to continually check that altitude remained within an acceptable range. In order to reduce the cost and effort of creating and maintaining a runtime monitor as illustrated for Dronology, it is necessary to (*i*) provide automated

support for creating a system-specific solution, including the setup and configuration of the monitoring system and specification of its respective constraints, and (*ii*) avoid the additional effort of maintaining the monitoring system itself by enabling the infrastructure to co-evolve alongside the *SuM*.

3. The *GRuM* Framework

To provide support for the different parts of a runtime monitoring architecture and to ease the task of maintaining and co-evolving monitors, *GRuM* leverages MDE techniques to specify relevant monitoring properties and ultimately to generate a complete monitoring infrastructure. It follows Rabiser *et al.*'s reference architecture [28], providing support for data collection, analysis (i.e., checking constraints), and visualization. However, one of the novel characteristics of *GRuM* is that for a *SuM*, both a *Set of Probes*, as well as a fully customized *Monitoring Platform* can be generated based on a model describing the parts to be monitored. Fig. 1 provides an overview of *GRuM*'s architecture. The *Modeling* part relies on a Monitoring Meta-Model (MMM) and a system-specific Domain Model Fragment that is populated for a *SuM*; the *Code Gen-*

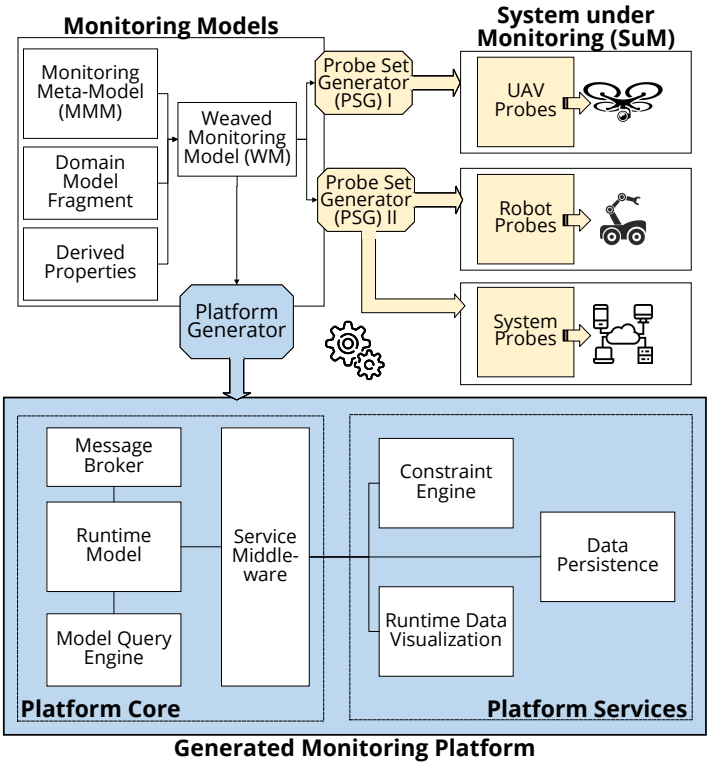

Figure 1: Architectural overview of the *GRuM* components for modeling the domain and relevant constraints, and for generating probes and the core monitoring platform.

erators responsible for generating the Monitoring Platform and the probes for collecting data from the *SuM*; and finally the generated *Monitoring Platform* itself, providing a number of runtime capabilities such as a runtime model, a query engine, and a middleware layer to connect external services. Properties of interest from the *SuM* that are collected and monitored via Probes can be accessed in various different ways. This, to a large extent depends on the code generator, and the resulting Probes that are generated from the model. For example, when generating a simple monitoring API (cf. application), the *SuM* needs to be "accessible", meaning that API calls can be added. If additional means of data collection are required, a code generation that used byte-code instrumentation [23, 32], or aspect orientation [33] might be employed.

3.1. Monitoring Models

MDE has been used extensively to generate applications and system components across a wide variety of domains, including automotive, railroad systems, business process engineering, and embedded systems [34, 35, 36]. *GRuM* uses model-driven techniques to generate a customized monitoring platform for the *SuM*. It only requires the parts of the *SuM* that need monitoring to be modeled, instead of the entire *SuM*. This represents a significant time-saving as it is often the case that only a small subset of properties and data created by a system are of interest at runtime. For example, a UAV's PX4 flight controller [37] has over 1,200 configurable properties and sensor values. Of these, a few values are regularly checked during the preflight arming process, whilst other values, such as GPS coordinates, altitude, attitude, velocity, satellite links, temperature, and camera/gimbal settings are typically monitored during flight; however, many hundreds of properties, such as internal flight controller states are unlikely to be monitored at runtime. The *GRuM* Monitoring Setup specifies exactly what parts of the system should be monitored and provides the mechanisms for performing the monitoring.

The challenge of creating a *Probe* [38], to collect runtime information from the *SuM* and making it available to the monitoring infrastructure, is that each *SuM* uses different technologies, diverse architectural styles, and consists of various software and hardware components. Therefore, in practice, monitors are typically defined and developed individually for each type of system or technology, such as byte-code instrumentation or dedicated service busses in service-based systems [9, 11].

GRuM addresses this challenge by providing a generic approach for describing monitoring properties and their resulting monitors, so that instead of building a custom set of probes, we specify the setup in a standardized way and generate a custom monitor directly from the specification. *GRuM* leverages MDE to describe the components of the *SuM* that need to be monitored, and then subsequently generates the monitoring infrastructure. To support this process we have created a dedicated MMM that can be used to define the relevant parts of the system for the resulting probes and the Monitoring Platform.

3.1.1. Monitoring Meta-Model (MMM)

Our MMM, shown in Fig. 2 was derived from analyzing the monitoring reference architecture from Rabiser *et al.* and existing runtime monitoring frameworks [28]. The `MoConfig` element defines the monitoring configuration for a specific system and is subsequently used to generate the code for a target system (cf. code$_{g}eneration$).$Itcontainsinformationregardingthen$ $processinganalysissuchasdataaggregationorqueryingthedatamodel(cf.runtimemodel)$.

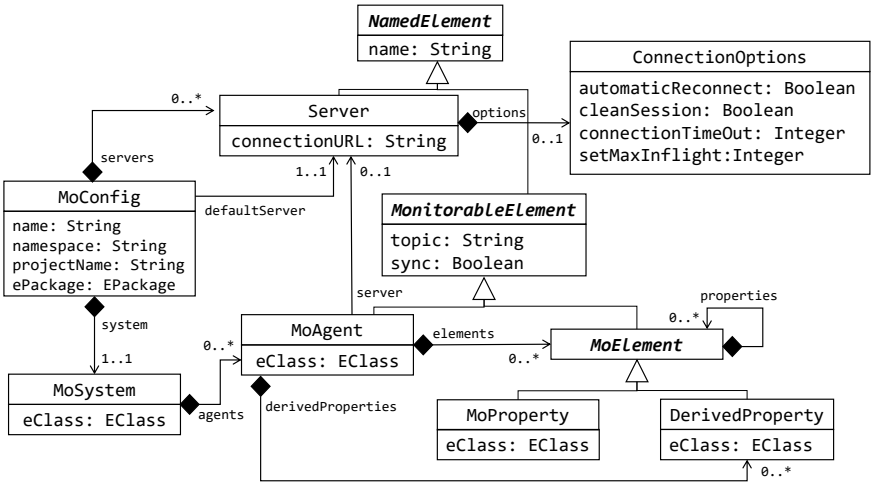

Figure 2: *GRuM*'s Monitoring Meta-model (MMM)

3.1.2. Domain Model Fragment

The MMM provides the foundation for all systems monitored with *GRuM*, requiring only those properties of the *SuM* that must be collected and analyzed at runtime to be specified in a *Domain Model Fragment*. Fig. 3 contains a partial Domain Model Fragment for the Dronology system, showing the `Drone` elements, each of which consists of different properties, such as the state (`DroneState`), a dedicated flightplan (`FlightPlan`), and safety checks at startup (`StartUpCheck`). The state, in turn, includes sub-properties such as battery and location information. While a `Drone` has many other properties that are either set during startup or at runtime, the Domain Model Fragment describes the subset of properties to be collected and analyzed by the monitoring framework. Properties can represent primitive data types (e.g., numbers, strings) or more complex aggregated elements.

3.1.3. Weaved Monitoring Model

In order to support the generation of monitoring code, the different elements specified in the Domain Model Fragment need to be linked to the concepts specified on the MMM.

These links are specified in the *Weaved Monitoring Model (WM)* and are subsequently used to determine which code fragments are generated for which element in the

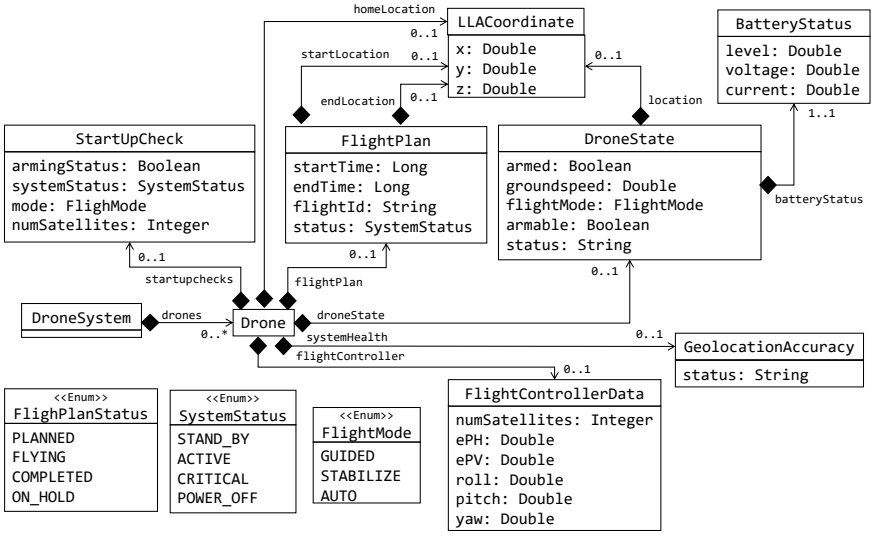

Figure 3: Excerpt of Dronology's Domain Model Fragment

Domain Model Fragment. Model weaving has been used in the past to connect different models to establish dedicated links between them [39, 40, 41, 42, 43]. As described by Del Fabro *et al.* [41], the weaving model is not self-contained but is meant to be used in conjunction with its related models and to provide the respective link semantics for the application scenario – in our case, the creation of the runtime monitoring platform.

In the WM, for example, a `Drone` element from the Dronology Domain Model Fragment is linked to the *MonitorableAgent* from the MMM, whilst `DroneState` and `FlightPlan` are children of the `Drone` element and linked to a *MonitorableProperty* as they can be directly collected from the system. The resulting WM (cf. Listing 1) is a hierarchical representation of the monitored system. Additionally, the WM also contains information about the topic mappings as well as the server configuration (represented by the `MoConfig` element) which is used to automatically generate topic subscriptions and connection code.

This clear separation of concerns, the monitoring information on the one hand, and the system configuration on the other, facilitates easy updates of the Domain Model Fragment when new properties or agents are added. It also supports reuse of existing models in cases where new *SuM*s also use model-driven techniques to generate system code.

3.2. Code Generation

A runtime monitoring platform relies upon several different components to collect, aggregate, and analyze runtime data. Based on the information encoded in the WM and the Domain Model Fragment, *GRuM* automatically generates (i) a set of probes for collecting information from the system, and (ii) an instance of the Monitoring Platform for the *SuM*.

Listing 1: Excerpt of the WM for the Dronology system

```
MoConfig DroneConf
    projectName DroneProject
    defaultServer localserver
    ePackage Dronology
    servers{
        Server localServer
        connectionURL "tcp://192.168.0.55"
    }
    system MoSystem{
        agents{
            MoAgent "Drone"{
             eClass "Dronology.Drone"
                elements{
                    MoProperty "DroneState"
                    topic "state"
                    sync true
                    eClass "Dronology.Dronestate"
   ... }
```

3.2.1. SuM Probes

Existing monitoring approaches use a variety of techniques to collect information from the running system. In order to reduce the overhead of manually implementing system-specific probes and to create data collection mechanisms for each system individually, *GRuM* generates a technology and language-specific set of probes. The Probe Set Generator (PSG) is independent of the data that is actually collected from the system but is dependent upon the underlying technology. For example, a PSG for a Java-based system that generates target Java code can be reused for other Java-based systems, and code can be re-generated when new properties or types of agents are added to the Domain Model Fragment. New PSGs can be easily added to *GRuM* when probes for a new technology or type of system are required without any need to redesign or re-implement the underlying monitoring platform.

3.2.2. Monitoring Platform

Runtime data provided via probes subsequently has to be collected and aggregated, so that the data can be analyzed, constraints checked, and results visualized. A topic-based message broker sends runtime data from the *SuM* (via the probes) to the generated infrastructure. The respective topics and topic subscriptions are generated automatically based on the agents and properties defined in the WM. Each `MoAgent` represents a root topic (e.g., "Drone"), and each `MoProperty` a respective sub-topic, e.g., "State". At runtime, each activated drone sends information through these topics (e.g., "Drone/-Drone1/State" for an update of the State property). This data is then used to populate and update the runtime model.

Every change in the model triggers a query against the *Model Query Engine* where constraints, i.e., checks on the model can be defined. Constraints are then evaluated on the runtime model and violations are generated when an evaluation fails. For example,

for the drone system, a constraint in the query engine could check whether the reported FlightPlan received from the probe contains a valid flight id and a set of valid coordinates the UAV should fly to. In case the constraint check fails, a violation is generated which can then be displayed in the user interface and the user can cancel the flight, or if desired, the flight could be aborted automatically.

3.2.3. Monitoring Middleware

In addition to the monitoring platform, *GRuM* also generates a middleware component that provides a *SuM*-specific interface to, for example, connect external services or applications. This allows applications to register to specific properties, and receive notifications when a change in the model occurs, for example, when a new agent has been added, or when properties are updated. Each agent and property is thereby accessible via a dedicated API method that corresponds to information encoded in the WM (LISTING 1). The mapping between the updated property/agent and the topic where the information is published in the middleware is established via the topic name specified in the WM. Examples of external applications that can be attached include additional constraint engines to support specific types of constraints (e.g., temporal constraints or event patterns), attaching a database to store runtime data, or adding a new UI. Separating the platform's core capabilities from external services decouples the data-receiving parts of the *SuM* from application services that build on top of it. The interface provides access to all information collected and aggregated in the runtime model from the *SuM*.

4. Applying *GRuM*

The process of using *GRuM* to specify and generate a *monitoring platform* and then deploying the platform to monitor a system at runtime is described in Fig. 4. The upper part involves creating the respective models and generating the platform code, whilst the lower part uses the resulting platform to define constraints and to connect external services, before deploying the platform for the *SuM*.

4.1. Creating the Models

In the first step of the process, relevant system information needs to be identified. This could include hardware and software components, such as sensors and their respective data, that provide valuable or critical runtime information. Various approaches can be used to identify and document relevant runtime requirements, safety properties, or QoS attributes [16, 44, 45, 46, 47]; however, the identified properties must be (1) specified in the *SuM*'s *Domain Model Fragment*. In case the *SuM* itself is model-driven, existing models of the system can be leveraged. Once the Domain Model Fragment is populated with an initial set of properties and attributes, these are (2) linked to the different monitoring concepts via the *Weaved Monitoring Model*. The WM captures monitored data, structured by the different agents and their constituent monitorable properties. This

step is vital for the subsequent code generation, so that the platform code generator can generate the specific runtime model and the hierarchical topic structure.

4.2. Platform and Probe Generation

Based on the WM, (3) two different types of monitoring code is generated for the *SuM*. First, a set of probes is generated using a technology-specific PSG which is either selected from an existing library or created from scratch. *GRuM* provides templates to ease the implementation of these generators, and new PSGs can be stored in the library and then reused for any other *SuM* using the same technology. The second code generator is technology-independent and is shared across all *SuM* applications. It (4) generates the monitoring platform containing the runtime model, data aggregation, and analysis component.

4.3. Platform Usage & Runtime Monitoring

Generated probes can be (5) directly deployed to the *SuM*. Depending upon their type (e.g., byte-code instrumentation probes, or data interceptors), they are integrated directly by the developers, or as part of the continuous integration process, into the source code or the binary files of the system. The generated platform then allows (6) users to define constraints and runtime checks. *GRuM* deliberately does not store constraint information in the Domain Model Fragment, and hence does not generate the constraints, since these are typically added incrementally, and subject to modifications at runtime [24]. *GRuM* supports constraint checks in two different ways. First, it provides an integrated model query engine that can query `MoProperties` and second, it is used to define checks on the model, for example, to ensure that a property stays within a certain range. The monitoring middleware generates notifications when a constraint is violated. Additional constraint engines can be connected via the middleware, providing the connectivity needed for constraint violations to be detected and reported. All

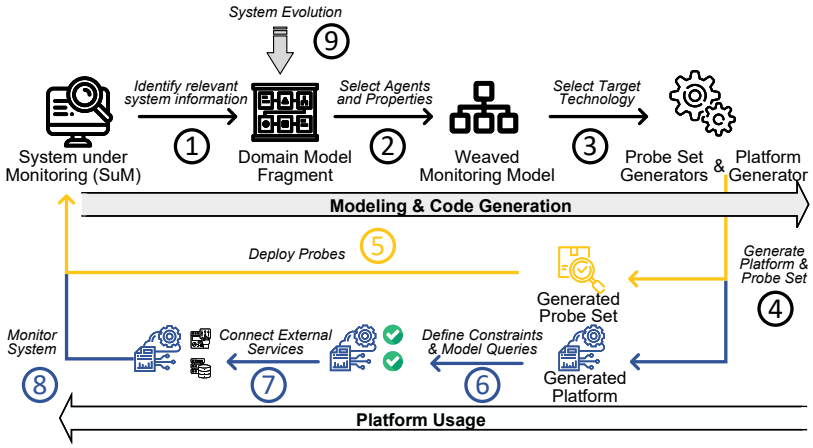

Figure 4: *GRuM* process for specifying, generating, and utilizing a monitoring platform

additional services (7) are connected via the middleware layer, and once connected, the platform (8) can be deployed to collect and analyze runtime data sent from *SuM* via the generated probes.

4.4. System Evolution

Since all major components of the platform and the probes are generated by *GRuM*, changes made to the *SuM* can be easily incorporated into the monitoring platform as well. This significantly reduces the effort of co-evolving the monitoring platform, as well as the probes with the *SuM*. Changes in the *SuM* can be reflected in the Domain Model Fragment and WM, and the code for the probes and platform can then be automatically re-generated. Previously defined constraints as well as external services remain unaffected and can be reused or deactivated (in case a constraint is checking a specific property that no longer exists, or is not monitored anymore). For example, if an additional property of an existing agent needs to be added to the set of monitored properties, this change can be directly performed at the model level. The respective element is first added to the Domain Model Fragment and then linked in the WM. Once this change has been made, probes and platform code can be completely regenerated without the need to manually adapt the code. In case a property is removed for which a model query has been defined the query engine provides an error message in the query description so that stale and invalid constraints can be ignored or removed.

4.5. GRuM Implementation

We implemented a fully operational prototype[2] of *GRuM* supporting the aforementioned parts based on the Eclipse Modeling Framework (EMF) [48]. EMF provides capabilities to describe meta-models by using its Ecore language which we used for the MMM and the Domain Model Fragment. For the automatic code generation, we use Roaster [49] to parse and create Java source files. Additionally, we used XText [50] to define a textual domain-specific language (DSL) for the WM (see example in Listing 1) which allows easy creation and modification of the WM for a specific *SuM*.

Probe Set Generator: As part of *GRuM* we provide a template and utility functions (e.g., for packaging and deployment) to create an implementation of a PSG for a target language/technology. A concrete implementation needs to iterate over the elements defined in the WM and then generate topic mappings for the respective agents and their constituent properties, as well as the code to send the data to the message broker. As part of our evaluation, we have implemented two distinct code generators to support two different target languages and technologies. The first generates a Java-based Probe Set, whilst the second targets ROS-based systems and generates a Python component with ROS-node subscriptions (cf. eval).

[2]Further implementation details and documentation is provided in our GitHub repository: `https://github.com/LIT-Rumors/grum-public`

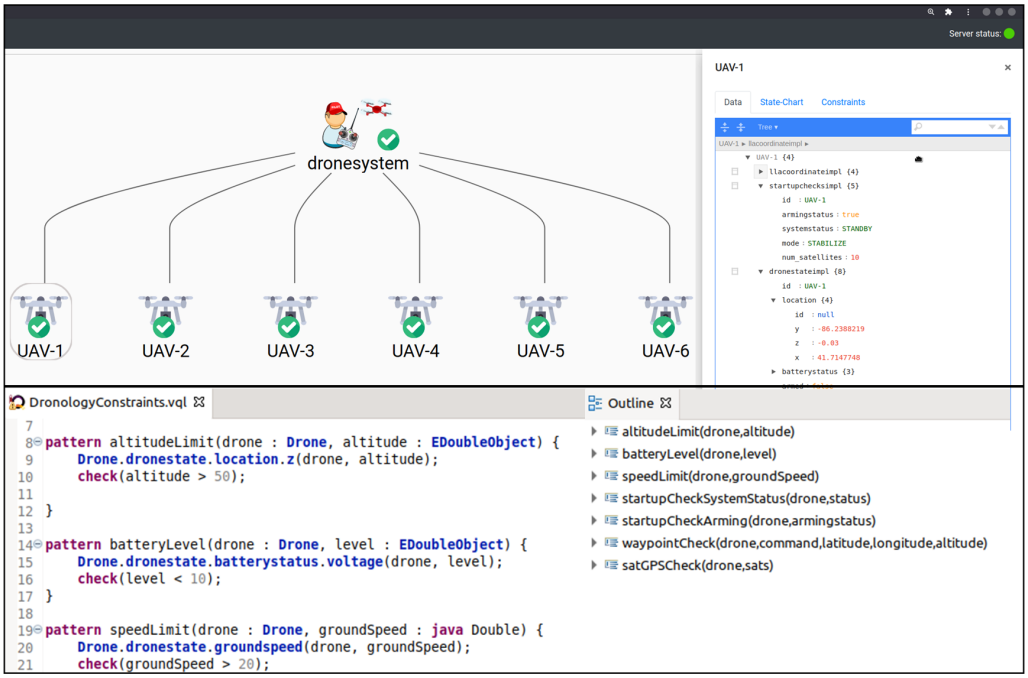

Figure 5: User Interface Prototype (upper half) and examples of Viatra model queries (lower half)

Monitoring Platform Generator: As EMF natively generates Java code from Ecore models, we leverage this feature, and our runtime model and the platform core components are implemented in Java (i.e., Java code is generated by the platform generator for each *SuM*) regardless of the *SuM*'s native programming language. At runtime, an instance of the Domain Model Fragment is created and used as the runtime model. All model code, code for instantiating the model, as glue code for setting up the platform is generated automatically. For data transmission, we use MQTT [51], which provides a scalable, topic-based publish and subscribe mechanism which is also automatically configured based on the WM.

The generated middleware layer is also Java-based and for each property and agent specified in the model, a respective notification method is generated. This allows external applications to register to be notified about changes in the runtime model. Furthermore, data is published on predefined topics (as specified in the WM), which can directly be accessed by any service or application subscribed to the respective topic via the MQTT broker. In case multiple instances of an agent are active (e.g., three UAVs are monitored at the same time) the messages are augmented with the respective "agent id" so that property updates received via the middleware can be assigned to the proper agent. This separation between the incoming information from the probes and outgoing information to the external services ensures that all information is handled by the runtime model which serves as the single point of information for all connected services.

Constraint Evaluation & Visualization: Integrated within the monitoring platform core, our implementation uses the Viatra model query engine, which is tightly integrated into EMF and directly operates on the model for specifying checks on runtime properties. However, while Viatra provides tools and incremental evaluation [52, 53], it does not support more complex constraints, such as temporal ones. We, therefore, connected the Esper CEP engine [54] via the middleware to support Complex Event Processing (CEP) as an additional service for checking temporal sequences and occurrences of events.

5. Evaluation Objectives and Setup

When evaluating our *GRuM* framework we address three main aspects. First, we address the general applicability of our approach to diverse systems that should be monitored. Second, we focus on the evolution aspect where we assess the effort that is in fact required when the *SuM* evolves compared to other monitoring approaches. Finally, we target efficiency and evaluate the performance of the generated monitoring platform to ensure its suitability in real-world application scenarios. For this purpose we used our *GRuM* implementation and investigate the following research questions:

RQ1: Can GRuM be used to create a runtime monitoring platform for a CPS with reasonable effort?
While representing a binary question, it is an important one for validating that *GRuM* works as specified. We, therefore, address the question by using *GRuM* to generate a monitoring infrastructure, collect events from the system and evaluate constraints at runtime.

RQ2: What additional effort is required to update the platform and monitors when the SuM evolves?
With the second research question we start with the generated platform (RQ1) and then analyze the effort needed to co-evolve *GRuM*'s models and monitoring platform when the *SuM* is modified. The goal hereby is to assess how different change scenarios impact *GRuM* and the resulting changes required.

RQ3: To what extent can GRuM's generated code be used to efficiently monitor time-sensitive runtime data in a CPS system?
The last RQ directly addresses the pertinent question of whether the code generated by *GRuM* is efficient. This is an important question, as generated generic code is often perceived to be less efficient than manually constructed code, and efficiency is particularly important in a runtime monitoring system. We address this question by measuring the model-set-latency of the monitoring platform, i.e., by measuring the time it takes to update the model at runtime when different property values change.

Two different researchers, both co-authors of this paper, performed the tasks associated with research questions RQ1 and RQ2. One was assigned to Dronology and one

to the TurtleBot system. Both researchers were familiar with the respective *SuM*, had previous experience in Java development (for the Dronology use case) and ROS and Python (for the ROS TurtleBot use case), as well as basic experience with Eclipse and the Eclipse Modelling Framework for creating Ecore models. Specific knowledge about code generation with Xtend was only necessary for implementing the initial Probe Generator prototypes, but not for subsequently generating a new *GRuM* instance based on the existing generators.

5.1. Use Cases

We explored the three research questions in two diverse robotic systems.

Case 1 – Dronology: The first use case was the previously discussed Dronology UAV system which we developed from 2016 to 2020 using Java and DroneKit [31] and is now available in the open domain. Dronology is fully compatible with both physical UAVs and simulated ones and has been used in extensive field deployments. The only difference between simulations and physical deployments is a single string used to establish communication with a real or simulated UAV. For our experiments, we utilized the high-fidelity Ardupilot simulation. All parts of the Dronology architecture pertaining to UAV management (such as route creation, mission planning, user interfaces, etc.,) are Java-based, while only the Groundcontrol Station responsible for forwarding messages to the UAV's flight controllers is Python-based. In our evaluation we, therefore, focus on the Java system and generate the Java monitoring Probes for relevant parts[3].

Case 2 – ROS TurtleBot Robots The second system used TurtleBot3 robots [55] equipped with sensors to collect $C0_2$ measurements. A TurtleBot is a small mobile robot using the Robot Operating System (ROS) [56] as a platform which has emerged as a de-facto standard for both research and industry. The TurtleBot platform is also used frequently for prototyping in research and education in the area of CPS in general and robotics in particular. [57, 58, 59]. ROS is an open-source platform for robotics software development using a component-based architecture with nodes that interact via a publish-subscribe pattern. ROS supports both hardware and high-fidelity simulation using, for example, the Gazebo simulation environment [60]. As ROS applications are commonly implemented in Python, we used this second case to demonstrate the applicability of *GRuM* with different target languages and technologies. As part of the prototype we implemented an application for controlling the TurtleBots and a number of control scripts for performing actions. Additionally, to augment the original Hardware of the TurtleBot3 robot, we attached an MQ-135 $C0_2$ sensor to facilitate mobile sensor collection with the robots. All required parts for this case were implemented in Python, and further details can be found in the open-source GitHub repository.

[3]Furhter details about the architecture, requirements, etc. can also be found on the project website: `https://dronology.info`

6. Evaluation

In the following we address our three research questions for creating (RQ1), evolving (RQ2) and evaluating the performance (RQ3) of *GRuM*.

6.1. RQ1 – Creating a GRuM instance

In the first part of our evaluation, we used *GRuM* to model MoAgents and MoProperties, created the WM, implemented the reusable PSGs needed for each system, generated the probes and monitoring platform for each system, and specified constraints against the runtime model. Two different researchers, both co-authors of this paper, performed this task. One was assigned to Dronology and one to the TurtleBot system. We then recorded the time taken to complete each activity.

Case 1 – Dronology: As we had not previously created a PSG for supporting probes in Java systems, we started by creating such a generator. Based on *GRuM*'s template for generating deployment packages we used a Java code generation framework (JBoss Roaster) to specify source code for the Java probes to be generated. We implemented a push-based probe generator that used the Domain Model Fragment and the WM (parts shown in Listing 1) to generate dedicated publish methods for each property, as well as utility classes for setting up the connection to the MQTT broker. The implementation effort for creating this reusable code generator was approximately 10 hours.

Modeling & Code Generation: We then developed runtime monitors for the Dronology system. This was guided by our inherent knowledge of the UAV domain and the Dronology system, including its properties, and the constraints that should be checked at runtime. From this, we (1) specified an initial Domain Model Fragment using the Eclipse EMF Ecore Editor. The model included multiple UAVs, representing monitorable agents, controlled by the `DroneSystem`, and several drone-related properties, such as `DroneState` and, `DroneCommands`. In total, we modeled 8 properties with 33 attributes. We also modeled `GeolocationAccuracy` as a derived property computed from each UAV's GPS status and its reported bias. We then (2) created the WM specifying the server configuration and the hierarchical structure of the system. These two steps took around 1-2 hours of effort. As the process was informed by our in-depth knowledge of the system, the reported effort focuses purely on the task of establishing the monitors rather than on "thinking-time" which is expected to be similar whether the monitor is built from scratch or using *GRuM*. We then (3) selected the Java code generator that we had created and (4) generated an initial version of the Dronology monitoring platform and its associated probes, composed of a set of executable Eclipse Plugins (for the platform) and a maven project (for the probe) that were directly integrated into Dronology (5). The only modifications made directly to the Dronology code base were the insertion of one method call to the probes for each monitorable property. All other functionality that was needed to collect data, establish a connection to the platform, and publish data, was fully covered by generated code.

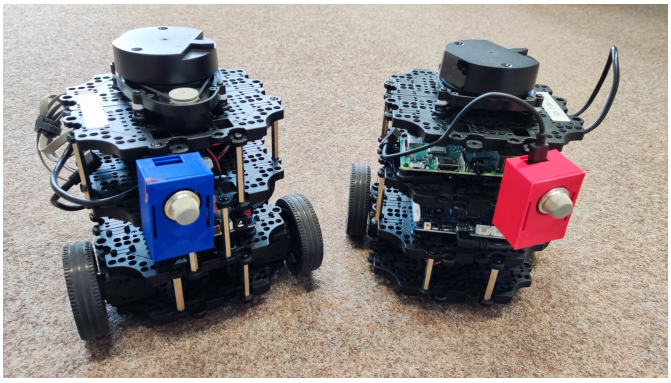

Figure 6: Two of the TurtleBot robots equipped with an additional $C0_2$ sensor.

Platform Usage: We then derived a set of constraints (6). Once again, these were guided by our own knowledge of the Dronology system [61] and its deployment in real-world tests. In total, we created 10 constraints, 7 of which were implemented using the Viatra model query engine [53], with three additional temporal constraints implemented using the Esper CEP Engine [54], as listed in Table 3. Additionally, we (7) developed a generic UI for visualizing agents and runtime data (cf. Fig. 5). Both the Esper CEP engine and our UI were connected via the middleware. Finally, we (8) deployed the platform for the subsequent evaluation. Altogether, the implementation of constraints and extensions was completed within about 20 hours.

Case 2 – ROS TurtleBot: As with the Dronology example, we started by creating a new PSG. For the TurtleBots, the generated code targeted Python to collect data from appropriate ROS topics and forward it to our monitoring platform. We modeled it in the same way as the Java-based generator by implementing a push-based approach. The Probes were implemented as a custom ROS package, consisting of the *ROS API* which provides update-functions for every monitored property of the ROS system, and an *MQTT-Forwarder* for sending data to the monitoring platform. Generated probes were implemented using ROS Noetic Ninjemys distribution packages in combination with the rospy client library for ROS.

Modeling & Code Generation: In contrast to the Dronology system, we developed this application from scratch without prior experience with the system or technology. Therefore, as a starting point, one researcher reviewed documentation for the TurtleBot platform [55], and identified relevant properties and their associated ROS topics over which the data was published. The resulting Domain Model Fragment (1) for the TurtleBot consisted of 10 properties with 23 attributes. We then (2) again created the WM for the Bot system specifying the server configuration and the hierarchical structure of the system. It took approximately 15 hours to investigate ROS properties, identify relevant ROS messages and create the Domain Model Fragment and WM. Of this, we estimate that 5 hours were spent implementing the *GRuM* monitor after properties had been selected.

Platform Usage: After (4) selecting the Python probe generator, the resulting probes were (5) easily integrated into the ROS workspace and executed without further customization. For the constraints (6) we again created 7 constraints for the Viatra query engine and 3 additional temporal constraints with Esper. In this case, constraints included a maximum carbon dioxide concentration and an enforced speed limit when low battery is detected to lengthen battery life and allow the bot to return to its starting position. A complete list of constraints is provided in Table 3. We (7) reused the same generic user interface as developed for the Dronology system, and (8) deployed the platform. In order to evaluate whether *GRuM* can easily cope with changes in the *SuM*, (9) we additionally mounted an external air quality sensor (MQ 135) and connected it to the control board of the TurtleBot. We incorporated this change in the Domain Model Fragment by adding an additional monitorable property and then regenerated the probes and the monitoring platform. The two constraints CST-T06 and CST-T10 specifically use the data reported by the $C0_2$ sensor to check that the value does not exceed a critical threshold and that measurements are collected only when the robot is stationary.

Analysis of Implementation Process: Creating the Domain Model Fragment in both cases was a straightforward task. After analyzing the two systems and selecting properties to be monitored, EMF provided an easy way to create the model, and was similar to using any other UML modeling tool. The implementation of each PSG required an initial one-time effort but the PSG was then reused to support subsequent changes in the system and Domain Model Fragment. Furthermore, once a generator for a certain language and/or technology is made available, it can be reused for any other system (in our case for any other Java-based or ROS-based Python application). To answer RQ1, for both cases we were able to model the monitoring properties with minimal effort. The vast majority of this time was used to specify (and test) temporal constraints with Esper, as no immediate tool support was provided, in contrast to the Viatra constraints (cf. discussion in discussion). For the Dronology system, where we had in-depth knowledge about the system, its properties, and structure considerably less effort was required, whereas, for ROS, we had to familiarize ourselves with the technology, but still were able to create a monitoring platform in less than 20 hours, including the process of selecting relevant properties, creating the model, and the constraints. In the Dronology case, for example, this resulted in 8 lines of code that were added manually to the *SuM*, supported by approximately 200 lines of generated Java probe code, about 700 lines of generated platform code, and an additional 2500 LoC generated by EMF for the runtime model. More importantly, when the *SuM* is changed, only minor implementation effort in the *SuM* is required to invoke *GRuM*-generated functions. When properties are added to the model, the entire platform and probes are regenerated.

6.2. RQ2 – System Evolution Support

In the second part of our evaluation, we specifically focused on the support provided by *GRuM* with regards to monitoring (co-)evolution, and what actions and efforts were

	MoProperty	NPR [#]	SML [ms]
Dronology	Command	470	5.79 (4.42, 6.41)
	FlightControllerData	2,943	5.91 (4.85, 6.52)
	FlightPlan	66	7.43 (5.33, 8.01)
	GPSHealth	6	4.94 (4.32, 7.19)
	HomeLocation	6	11.77 (7.91, 15.46)
	OperationMode	96	5.24 (3.91, 5.74)
	StartupChecks	6	12.86 (6.82, 18.42)
	DroneState	14,702	6.32 (4.99, 6.96)
TurtleBot	AirQuality	1,406	6.60 (5.50, 7.35)
	BatteryStatus	368	6.97 (6.32, 7.63)
	Diagnostic	965	6.66 (5.88, 7.28)
	JointState	1,828	6.35 (5.46, 6.93)
	LaserScan	1,828	6.39 (5.49, 6.96)
	MagneticField	1,828	6.35 (5.40, 6.89)
	Odometry	1,790	6.44 (5.74, 7.06)
	SensorState	1,813	6.50 (5.76, 7.12)
	Velocity	1,829	6.47 (5.62, 7.12)
	VersionInfo	1,822	6.25 (5.55, 6.84)

Table 1: Model-Set-Latency results of the simulation runs (median values). NPR: number of property values collected at runtime / run, SML: time required from event received until set in the runtime model, median, (1st/3rd quartile)

required to add monitorable properties to a system, perform constraint checks using the newly collected data, update an existing element (e.g., change the information and/or type of a property) and remove an existing element from the model. We established a set of tasks that impacted all three primary parts of the reference architecture [28, 62, 63, 64], focusing on activities related to *Monitoring Setup* and *Execution*). Based on these tasks (cf. Table 2), for each of our two systems, we recorded the different steps necessary for adding, updating, and deleting monitoring properties in the system, modifying the respective constraints, and regenerating the platform.

We specifically examined three changes: (*i*) a new property needs to be monitored and therefore added (A), (*ii*) an existing property is modified (e.g., the property name, MQTT topic, or attributes change) (U), and finally, (*iii*) a property is removed and should no longer be monitored (D). For each change, we documented the steps necessary, the tasks that needed to be performed, the elements/artifacts (e.g., models) that needed to be touched, the tools used, and the degree of manual work/automation provided (cf. Table 2).

Monitoring Setup: All changes performed related to the monitoring setup only require editing one of the two models - the Domain Model Fragment or WM, and all code is subsequently generated. For activities related to the monitoring setup, all changes are performed at the model level, including updating the Domain Model Fragment and the WM. Using the Eclipse Ecore framework [65], this is supported by a graphical editor and can be easily achieved with a UML-like diagram. All model code is subsequently generated automatically. Once these changes are applied, the probe API and the monitoring

platform are generated automatically. For creating an executable monitoring platform (and the respective runtime model), no changes are required at the code level for either system.

Monitoring Execution: Once respective probe APIs have been generated, they need to be linked with the *SuM*. If agents and/or properties are added or modified, the new API methods need to be called accordingly. When applying our change scenarios to the Dronology system, this required adding/modifying 1 line of code (1 method call) in the system, calling the respective API method when a value is updated. For the ROS/Python example, the effort is similar. First, the ROS node subscriber needs to be started and the corresponding function call in the probe API executed. Besides this, no additional code is required. All connections, topic publishing, and message forwarding is performed automatically by the generated probe and platform code. In the platform itself, no changes or manual steps are required and the runtime model is executed automatically once the platform has been generated.

Based on analyzing the different change scenarios for two different systems, we can conclude that *GRuM* provides a high degree of automation, reducing the manual effort that is required to adopt changes in the *SuM* in the monitoring platform. More importantly, the majority of steps performed are supported by tools and editors, and only the actual instrumentation, i.e. linking the *SuM* to the respective probes, requires actual changes in the source code of the system. None of the changes performed required a manual change in the code of the monitoring platform, and all necessary code could be generated automatically.

6.3. RQ3 – Performance of the generated monitoring platforms

As explained earlier, one question that often arises with code-generated solutions is whether they perform efficiently. This part of the evaluation uses the two deployed monitoring solutions from RQ1 to assess *GRuM*'s performance for monitoring a system, collecting runtime information, and evaluating constraints. In both cases, we collected the following metrics: The number of messages, i.e., monitorable properties sent to the platform (NPR), the latency of the platform (set-model-latency), i.e., time required from setting a value to the runtime model (excluding network delay from the system to *GRuM*) (SML), the number of constraint checks/violations reported (NCST) and time required for performing a constraint check after an element in the model changed (TCST).

The goal of this experiment is to assess whether the generated monitoring platform can handle realistic data from either a high-fidelity simulator (Dronology case) or actual physical hardware (TurtleBot case). We excluded communication latency from our experiments, as this would equally affect any monitoring solution that requires central data analysis. Instead, we focused on the latency when data is available and propagated in the runtime model, and constraint evaluation times, two critical characteristics of the monitoring platform itself.

	Task	Description	Artifacts & Tools	Action required in *GRuM*
Monitoring Setup	Monitoring Definition	What should be monitored?	Domain Model Frag., (Eclipse Ecore Model and Generator Model) STEPS: 2	**[A|U|D]** MoProperties and MoAgents can be changed in the Domain Model Fragment using a graphical modeling editor. AUT.: PARTIALLY AUTOMATED When the Domain Model Frag. is updated, Eclipse automatically generates required code – no manual code changes required.
	Monitoring Instrumentation	How are actual Probes defined/created?	Weaving Model (DSL) Eclipse XText editor STEPS: 1-2	**[A|D]** MoProperties and MoAgents from the Domain Model Frag. need to be linked in the WM and topic bindings need to be established using the DSL. **[U]** If only attributes of MoProperties are updated, no change in the WM. AUT.: PARTIALLY AUTOMATED When WM is updated, the respective code generator can be selected in Eclipse, and the platform and Probes are generated automatically.
Monitoring Execution	Monitoring Information Collection	What is required to collect data from the SuM?	Generated API library (Java: JAR file, Python: module for ROS) STEPS: 1	**[A|U|D]** No extra steps necessary for framework. For instrumented system (Java, Python), the generated library needs to be added and the respective API method calls need to be added. AUT.: FULLY AUTOMATED (PLATF.)/MANUAL (SUM)
	Monitoring Information Processing	What is required to manage and distribute data in the mon. framework?	-	**[A|U|D]** No extra steps necessary, code for the monitoring platform is generated automatically and a runtime model is instantiated. AUT.: FULLY AUTOMATED
	Monitoring Information Checking	What is required to define and perform constraint checks?	Viatra editor, Esper rule file STEPS: 1-2	**[A|U|D]** Constraints need to be added/modified either in the Viatra query engine or as new Esper rules. AUT.: MANUAL WITH TOOL SUPPORT.

Table 2: Key tasks in a monitoring framework based on the monitoring reference architecture by Rabiser *et al.* [28]. *Aut.* describes the degree of automation support provided by *GRuM*, and *Steps* refers to the number of artifacts that need to be modified for performing a certain action for Adding (A), Updating (U), and Deleting (D) elements from the monitoring platform and SuM.

Case 1 – Dronology: Using Dronology's Software-in-the-loop (SITL) simulator we created a UAV search mission scenario where multiple UAVs search for a missing person in a predefined area. We deployed Dronology with our generated probes and the sim-

ulation environment to six Raspberry Pis (RPi 4 with 4GB of RAM, running Raspian Buster), each representing an individual UAV and connected via Wifi (similar to companion computers on real UAVs). The monitoring infrastructure was set up on a standard Desktop Computer, with 16GB of RAM, running Ubuntu 20, using Java 11, and Ecore 2.23. We randomly assigned 5 flight routes to each UAV, and then collected runtime data and performed constraint checks. We first performed a standard simulation scenario with no anticipated constraint violations, and then in a second simulation, we seeded errors every 30 seconds in order to evaluate whether violations were reported correctly and in a timely fashion. Both scenarios (with and without errors), were executed 5 times and median values are reported over all 5 runs. An overview of the number of events captured and constraint checks performed is provided in Table 1 and Table 3.

Overall, each simulation run lasted about 41 minutes representing a typical upper bound of a typical battery-powered UAV flight. During this time, a total of 18,295 events were sent to the monitoring platform. The median time to set a property in the model was between 4.9 and 12.9 ms. We only observed higher latencies (> 7.5 ms) for the properties that were collected a single time during UAV startup (12.9 ms) (the startup checks and homelocation) with all other properties well below this limit. The median time to evaluate a constraint for the Viatra query engine was between 0.02 ms and 0.9 ms. As Esper is a stream processing engine that only matches a certain pattern, i.e., a constraint violation, we were only able to count and/or measure constraint violations for the seeded errors and not actual executions. In total 4,168 violations, for both Esper and Viatra constraints were reported. As Esper constraints are temporal, the time from the point an event is added to the event stream to the time a violation is detected varies depending on the type of constraint. For CST-D08 (cf. Table 1), this means that a constraint violation is triggered after 5 minutes (300 seconds), for CST-D08 after 2 seconds, when the second `DroneState` property is not received in time, and after 30 seconds for CST-D10 respectively. In both cases our platform handled diverse constraints, some of which are executed very frequently, some of which are only executed once, providing an evaluation result of fewer than 0.9 ms for the Viatra constraint and 71 ms for a timing constraint evaluated with Esper.

Case 2 – ROS TurtleBot: For the second case, we performed a series of runs with two TurtleBot robots (TurtleBot3 Burger, equipped with a RPi 3, running ROS noetic), tasked with measuring $C0_2$ levels in various offices and hallways to detect and report high concentrations. We used the ROS SLAM (Simultaneous Localization and Mapping) node to create a map of the office space and sent the TurtleBots to different offices multiple times using ROS' 2D Navigation Stack. The monitoring infrastructure was set up on a standard Desktop Computer, with 16GB of RAM, running Ubuntu 20, using Java 11, and Ecore 2.23. Again, we first executed a scenario without any anticipated constraint violations, followed by a second scenario seeded with random errors. Both scenarios (with and without seeded errors), were executed 3 times and we report the median values over the 3 runs. Each run lasted approx. 16 minutes with 15,477 events

CST	Description	Type	NCST [#]	TCST [ms]
D01	*Altitude Restrictions:* The UAV must not exceed a maximum altitude of 50m	Q	492	0.03
D02	*Speed Restrictions:* The UAV must not exceed a maximum speed of 15ms/s	Q	492	0.07
D03	*Minimum Battery Level:* The UAV has to maintain a minimum battery voltage of 10.5 V	Q	492	0.02
D04	*Valid Goto-Commands:* A waypoint sent to the UAV must have valid latitude, longitude and altitude values	Q	286	0.09
D05	*Startup Arming Check:* Before active, the UAV needs to compute its arming checks	Q	6	0.30
D06	*Flight Controller Startup:* Before active, the UAV flight controller needs to properly startup	Q	6	0.17
D07	*GPS:* A minimum number of 10 satellites must be available	Q	1458	0.06
D08	*Flight Completion:* An assigned route for a UAV has to be completed within 5 minutes	T	42	59.00
D09	*Update Frequency:* State data from the UAV, updating its speed and position must be sent at least every 2 sec	T	541	58.70
D10	*Geolocation Accuracy:* When GPS accuracy of the UAV is low for more than 30 sec., an active mission can not be continued and manual control has to be assumed	Q/T	353	71.75
T01	*Movement Speed Limit:* To maintain accuracy during navigation, the bot must not move faster than 2.5m/s	Q	65	0.09
T02	*Minimum Battery Status:* The bot has to maintain a minimum battery level of 5%	Q	65	0.09
T03	*Minimum Battery Voltage:* The bot has to maintain a voltage between 10.5 and 12.5 volts	Q	65	0.06
T04	*Power Supply Health:* The bot's power supply must remain in a healthy condition, e.g., no overheating	Q	65	0.05
T05	*Diagnostics Error:* The operating level of the hardware components (actuator etc.) must not be in an error state	Q	49	0.07
T06	*$C0_2$ Limit:* The $C0_2$ measurements from the air quality sensor must not exceed a threshold of 800ppm	Q	76	0.08
T07	*Obstacle:* The bot must maintain a minimum distance of 5cm to an object, as detected by the Lidar unit	Q	65	0.08
T08	*Speed Reduction:* When the bot operates below 25% battery level, the speed must not exceed 2.0m/s anymore	T	88	52.10
T09	*Diagnostics temporal check:* A stale state of the actuator must change within 10 sec. to another state	T	206	32.58
T10	*Measurements Accuracy*: To ensure accurate measurements, an alert should be raised when measurements are transmitted while moving	T	191	84.00

Table 3: Constraints used in the evaluation. NCST describes the median violations reported in our seeded runs, and TCST shows the median evaluation time from when the value was received until the constraint violation was reported in the platform

collected per run, which were sent to the monitoring platform, resulting in 980 events per minute. The latency to set values in the model ranged between 6.3 ms and 7 ms (cf. Fig. 7).

For the constraint evaluation runs, we obtained very similar results to the previous

experiment. The median time to evaluate a constraint for the Viatra query engine was between 0.05 and 0.09 ms and for Esper constraints between 33 ms and 84 ms, with a median of 935 violations reported per run. All results are listed in Table 1 and Table 3.

6.3.1. Scalability

To validate the scalability of the generated platform, we evaluated its capability to deal with a large number of properties and agents. For the Dronology case, we virtualized the setup used on the Raspberry Pis in a Docker container and executed a simulation scenario with 25 UAVs flying and reporting data at the same time. For the second case with the TurtleBots, we increased the frequency with which data was sent to the monitoring platform. Initially, properties were sent every second and every five seconds for the battery measurements which were reduced to 0.25 seconds resulting in a much higher rate of events and constraint checks. For the approx. 46-minute Dronology run with 25 UAVs, we received 71,567 property updates, i.e., more than 1,530 property changes per minute, and observed a constant number of messages being received by the platform throughout the run without any backup with comparably low latencies between 4.3 and 6.3 ms. During this run, 20,719 constraint checks were performed with evaluation times between 0.08 and 0.13 ms for the Viatra queries and 46 and 61 ms for the Esper temporal checks.

For the TurtleBot run, we received 64,000 updates during the 16 minutes run, i.e., approximately 4,076 properties per minute. 1,173 constraint checks were performed and evaluation times range between 0.07 and 0.1 ms for the Viatra queries and between 63 to 180 ms for the Esper temporal checks.

Analysis of Performance Evaluations: For both cases, we were able to monitor a substantial amount of properties and the platform performed well for updating the runtime model and performing constraint checks. For the Dronology system, we deliberately selected properties with a diverse update frequency, demonstrating that the framework

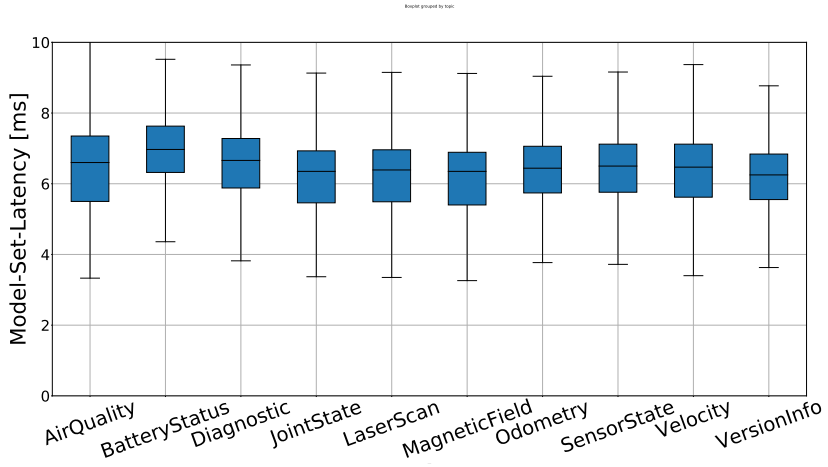

Figure 7: Model-Set-Latency for Runtime properties for the TurtleBot evaluation runs.

works well with data collected only a few times a minute, or several times a second as for the TurtleBots. Additionally, even when scaling up the number of agents in the Dronology system, and when increasing the frequency at which events are sent to the platform for the TurtleBot robots, we achieved almost constant evaluation and latency times. Furthermore, as all experiments were run on a standard desktop machine, both the latency to update the model and time to perform constraint checks can be further improved when deploying the platform on server hardware. With regards to RQ2, we can conclude that the generated platform is capable of handling a substantial amount of diverse properties from *SuM* using different architecture styles and technologies.

6.4. Comparison to related techniques and Trade-Offs

In order to contextualize the results of RQ1 and RQ2 and to draw conclusions about the extent to which *GRuM* can reduce required effort, we provide a qualitative comparison against a selection of related approaches. We do not provide quantitative comparisons – as fully recreating each of the other systems and frameworks is out of scope of this paper. For this purpose, we leverage our previous work in the area of analyzing Runtime Monitoring Frameworks [16] by selecting five [13, 66, 67, 68, 69] approaches that (a) cover a broad spectrum of different application domains and application contexts of runtime monitoring; and additionally, (b) provide sufficient information in the paper to allow a comparison with respect to the capabilities provided for Runtime Monitoring and trade-off ease of use and convenience. We structure our discussion along the previously discussed key tasks of "Monitoring Setup", "Monitoring Execution", and 'Monitoring Support" from the monitoring reference architecture [28]. Results from this analysis are summarized in Table 4.

Monitoring Setup: Our *GRuM* framework relies on a strict separation between the description of the *SuM* (and its constituent properties) in the Domain Model Fragment and any information relevant for monitoring in the Weaved Monitoring model. With regards to monitoring definition, similar to many other approaches, we use a textual description based on a dedicated DSL to specify monitoring-related information. Similar to our approach, Java-MaC [67] and Inzinger *et al.* [69] use their own custom DSLs, and Kieker [68] uses OCL for monitoring definition. In contrast, SPASS-meter and WLSA [13, 66] rely on XML files for defining monitoring scopes. Alternatively, SPASS-meter also supports direct annotation of source code, thereby avoiding the need for a separate file for defining monitors. One negative trade-off that we are aware of is that as a precursor to applying *GRuM* the creation of the Domain Model Fragment requires some additional effort. However, as only relevant properties need to be modeled, the perceived effort is relatively low, and the resulting benefits are the tool-assisted creation of the Weaved Monitoring Model (auto-completion, selection of properties, etc.) which greatly eases its creation and maintenance, and the entire framework generated from it. Furthermore, if models are readily available for the system, they can be easily reused and integrated, thus further reducing the effort of this extra step. With regards to the creation

of probes, we also follow the paradigm of automatically generating probe code for the target *SuM*. Depending on the approach, either Java bytecode instrumentation [66, 67] or aspect-oriented programming [68] techniques are employed. With the concept of dedicated *Probe Set Generators*, we introduced one additional layer. This increases flexibility (similar to, for example, Kieker [68] where templates for Python and Java-based systems are provided). PSGs thus facilitate easy reuse, and once a generator for a specific technology/target system has been created it can be easily integrated and readily reused.

Monitoring Execution: Once created, *GRuM* provides a fully operational monitoring framework that can be readily executed. This includes event collection as well as distribution at runtime. While other approaches combine the definition of event monitoring and analysis (e.g., as part of the Event Definition DSL [67] or SLA metrics [13]) we again focus on separation of concerns between these two parts. The MaC [67] architecture, for its Java implementation, for example, relies on a dedicated instrumentor and compiler to manipulate bytecode and a run-time checker. While this allows for some flexibility, and its concepts can also be applied to other languages, it requires a specific implementation of these components. We avoid this by decoupling probes and the actual framework implementation. Kieker [68] on the other hand follows a similar approach as *GRuM*, with a common and extensible monitoring record model for data collection, and a dedicated record consumer model that facilitates runtime checks. As part of the core capabilities in *GRuM*, we leverage the Viatra query engine that again use the generated Ecore code to provide sophisticated tool support for generating model queries/constraints on the monitored properties. More complex constraints can be added via external services connected through the middleware layer.

Monitoring Support: While the main purpose of *GRuM* is the generation of the "base" monitoring framework, the automatically generated Monitoring Middleware enables easy integration of additional services and tools. We have demonstrated this by connecting a CEP engine, which is an approach also used in the framework proposed by Inzinger *et al.* [69]. The model-to-code transformation, using EMF, *GRuM* generates ready-to-use interfaces and classes and a template Java Maven project, that provides *SuM* specific interfaces and wrapper classes that could significantly ease the task of connecting external applications. While introducing model-based technologies adds to the complexity of *GRuM*, our goal is to "hide" this complexity from the user to the fullest extent possible. Similar to the other monitoring frameworks, we employ a domain-specific language for defining monitoring configurations, ready-to-use code generators, and templates for connecting to external services.

Co-Evolution Support: The aspect of monitoring co-evolution is only lightly addressed by a few other approaches. SPASS-meter briefly addresses the co-evolution aspect, acknowledging that depending on how the scope configurations are defined they may become outdated as the system evolves. When using their external configuration option, specified as XML files, synchronization is required, whereas their inline con-

figurations are directly added to the source code, hence avoiding the synchronization issue. However, the latter requires direct modifications to the source code of the SuM. The results of RQ3, indicate clear benefits when relying on model-based concepts, not only with regards to easy re-generation of the entire framework, but also in terms of the provided support structure. For example, removal or updates of properties from the system/model are immediately reflected in the updated code that is generated (or removed), and error messages and warnings are triggered in the tool-supported editors (e.g., for Viatra model queries that reference a property that no longer exists).

Phase \ Framework	GRuM	[66]	[67]	[68]	[69]	[13]
Monitoring Setup	●	◑	●	◑	◑	◑
Monitoring Execution	●	○	◑	◑	○	○
Monitoring Support	◑*	◑	○	◑	○	○
Monitoring Co-Evolution	●	◑	○	○	○	○

Table 4: Comparision of *GRuM* with regards to support and ease of use of key monitoring tasks [28]. ●= fully supported/available; ◑= partially supported ; ○= not supported/mentioned; (*=supported via the *GRuM* Monitoring Middleware).

7. Threats to Validity

As with any experiment, our evaluation is subject to a number of threats to validity.

Internal validity is related to the rigor of the experimental design. Our evaluation is based on the creation of monitoring models, and the generation of the monitoring platform for two different application scenarios. While these were created by the authors of this paper, the selected properties and resulting constraints were selected from real-world use case scenarios and applications. With regards to the selected systems, for Dronology, we possess in-depth knowledge about its internal design, structure, and functionality which also informed the creation of the Domain Model Fragment and selection of constraints. To avoid bias, we performed the same steps on a second system, the TurtleBot robots where we first had to get familiar with the technology and system characteristics. Additionally, having built a custom, system-specific monitoring solution for one of the systems allowed us to discuss and assess the value of a generic monitoring platform that can easily be applied to different systems and updated when the *SuM* evolves. So far, as part of the evaluation, *GRuM* was used by several authors to create models and generate the Monitoring Platform demonstrating the applicability and evolution efforts. In order to properly assess the usability, and to show broader applicability, we need to involve additional users, e.g., by conducting a dedicated usability study of the framework and provided tools.

External validity refers to the generalizability of results and findings. Our study focused on only two different systems, and therefore, we cannot claim full generalizability of our approach to diverse types of CPS. However, the Dronology system and the TurtleBot systems exhibit significant differences in terms of the technology used, their architectural styles, and properties that are monitored, demonstrating the flexibility of *GRuM* across different types of systems. Further, to minimize the threat of invalid data measurements due to external factors such as OS tasks or interference with other applications, we performed multiple runs of each scenario and calculated the medial values; however, our performance measurements focused on the runtime model and the constraint evaluation and did not include network delays or delays in the underlying *SuM* which likely account for certain latencies such as retrieving the home coordinates of a UAV in Dronology. Most notably, while the model-set-latency and constraint evaluation times are acceptable for the Dronology and TurtleBot systems, we can not currently guarantee strict real-time deadlines for evaluating constraints. Nevertheless, the modular nature of *GRuM* allows key platform components to be replaced, for example, replacing MQTT with a DDS [70] middleware implementation providing real-time capabilities.

Construct validity refers to the extent to which a study measures what it claims to be measuring. RQ1 investigated the ability of *GRuM* to describe relevant properties and to generate a customized monitoring platform. We applied *GRuM* to two different systems with real-world applications and physical hardware using properties derived from existing documentation. We were able to represent all identified properties and constraints with our proposed approach and generated a fully functioning, customized monitoring platform, and showed that in both cases, *GRuM*'s generated platform was able to handle realistic amounts of data and different types of properties. While we have been using a simulator for the Dronology use case, Ardupilot is a high-fidelity simulator providing a realistic execution environment, and our previous experience has confirmed that this closely resembles real-world scenarios.

8. Discussion

Results from our evaluation have shown that our model-based *GRuM* framework can be successfully used to (1) describe relevant properties and (2) collect and check runtime data via the automatically generated monitoring platform. The platform was easy to customize through the use of the middleware – which allowed us to seamlessly add the additional temporal constraint engine. Clear separation of concerns makes the platform highly flexible and maintainable.

The distinction between the platform's core capabilities and add-on services allows the core to be regenerated in response to changes in the *SuM*, while the separation between the generation of the platform and the probe set for collecting runtime data enables easy adaptation of *GRuM* for new technologies and types of systems being monitored. In our prototype implementation, we currently use direct instrumentation for the gener-

ated Probes (i.e., providing probes that expose an API for submitting new data). This, on the one hand, allows for easy use (similar to a logging API), but on the other hand, requires some manual effort to connect the *SuM* to our framework, and requires source code access. If this is not possible and/or not desired, this can be easily changed by simply replacing the respective Probe Generator with a different one. For example, a new Probe Generator for Java-based systems could use byte-code instrumentation [71], or aspect orientation [33] to collect runtime data. The main advantage however is that the platform can be easily tailored to a new system with no programming effort if a probe Set Generator can be reused from the library or with considerably low effort as shown in RQ1. In comparison, as part of our Dronology system, we have previously implemented a custom monitoring solution specifically tailored to the UAV use case. With a similar technology stack (MQTT, and Drools as a Constraint Rule engine), this required approx. a person-month of effort and roughly 1500 lines of code for a platform that is highly specific and interwoven with the Dronology system. In contrast, *GRuM* required some initial modeling effort, but only minimal implementation effort to call the probe API to connect the system to the platform.

With regards to monitoring co-evolution support, we were able to demonstrate that *GRuM* provides a high degree of automation and reduces the manual effort of modifying code in the monitoring platform, when changes to the *SuM* occur. Certain manual tasks are still required (e.g., adding/modifying properties in the models), but these largely take place at the model level, where tools and support for the users are provided. One aspect we are currently investigating as part of our ongoing work is providing additional support for checking the consistency between the WM and the actual source code of the system, by e.g., annotating Java Classes and fields in the system as monitorable, and subsequently leveraging AST analysis to synchronize, or even automatically generate the WM. Additionally, we are planning on extending our evaluation to include a full user study, providing participants with a system and the requirement to monitor and check certain properties, with the goal of assessing the usability of the platform and modeling approach.

Eclipse and the EMF platform provide a convenient way of defining the Domain Model Fragment and the WM. For the WM, a textual representation via a DSL and the provided EMF editor (including code completion and syntax highlighting) enables easy updates and adaptation of the WM and regeneration of the monitoring platform. The evaluation showed that both Domain Model Fragment and WM can be easily created and updated with minimal effort. However, to further ease the task of preparing the generated platform for a specific *SuM*, we plan to provide additional tool support. While creating Viatra model queries is relatively easy and straightforward, defining temporal constraints as Esper queries is more challenging and time-consuming. We, therefore, plan to further investigate the suitability of other constraint and rule engines and to provide a dedicated DSL with constraint templates that get automatically translated to either Viatra queries or Esper constraints as appropriate.

9. Related Work

Related work is primarily in runtime monitoring for CPS, model-based runtime monitoring, and models@runtime.

Runtime Monitoring for CPS: A recent survey by Rabiser *et al.* [28] highlights the runtime monitoring area as an active field of research, presenting a reference architecture for monitoring platforms. We based the architecture of *GRuM* on the architecture described in this work. However, the approach of Rabiser *et al.* is not model-based and does not provide capabilities to generate monitoring support for different types of systems.

Several approaches have been proposed with regards to monitoring CPS and large-scale systems. For example, Doherty *et al.* [72] presented a task planning and execution monitoring framework using temporal action logic to specify the behavior of the system. Vierhauser *et al.* [73] presented a case study on monitoring UAVs using their ReMinDs framework. In the domain of self-adaptive systems, Machin *et al.* [74] proposed an approach for synthesizing monitored autonomous systems. Kieker [68] provided capabilities for inserting probes for intercepting the system execution and monitoring various runtime aspects of the system. The framework also supports adaptation of monitoring rules for activating and deactivating probes relevant to the current monitoring task. HiFi [75] used programmable agents and filters to manually configure the infrastructure at runtime and adapt different agents. However, while some of these approaches provide support for instrumentation, for example via byte-code instrumentation or predefined monitors [76, 9], or customizing the monitoring platform, none of them provide extensive support for automatically generating code or provide an independent-language, applicable across different types of systems and thus lack of a more abstract solution.

Model-based Monitoring & Models@Runtime: In MDE and industrial CPS domain, the term "Digital Twin" has become synonymous with a model of the system instantiated at runtime [20, 77]. Specifically, in Models@Runtime research [78, 79], the models are instantiated at runtime, and then used to check properties and support self-adaption. However, very few approaches employ model-based concepts to facilitate the generation of a complete monitoring platform. Hili *et al.* [80] proposed a model-based architecture for interactive runtime monitoring using model-based techniques. However, while they supported model-to-model transformation as well as automated code generation, their focus was on monitoring real-time and embedded systems. In the domain of robotic applications, MROS by Corbato *et al.* [81] supported runtime adaptation of ROS-based systems using a model-based framework. While their approach provides potentially interesting extensions to our ROS PSG, with *GRuM* we aim to provide a more diverse framework that can generate and provide monitoring support for different types of systems. Brand and Giese [15, 82] have extensively worked in the area of model-based adaptive monitoring. They proposed a generic adaptive monitoring approach, based on analyzing queries on a runtime model and adjusting periodic or event-driven monitoring tasks. While their work also revolves around runtime models and MDE, their focus

is on adaptive monitoring and query support on runtime models, without support for automatically generating the monitoring platform and probes to perform these queries.

Sakizloglou *et al.* [83] used runtime models supported by Viatra to check constraints as part of adaptation rules. Bur *et al.* [84] provide support for querying CPS runtime properties using distributed graph queries. While we plan to explore distributed monitoring in future work, *GRuM* focuses on the specification, creation, and maintenance of a flexible, and extensible monitoring platform.

Formal approaches for Monitoring & Verification: An entire research area is dedicated to formal approaches and runtime verification with several frameworks that specify monitors using formal models, such as temporal logic or Event Calculus. For example, MOP [76] uses a specification formalism either part of a Java class file or specified independently. It consists of a header containing meta-information, a body specifying the desired property to be checked, and a handler performing certain actions in case the property is violated. Chan *et al.* [85] developed a framework for runtime verification of timed and probabilistic non-functional properties in component-based systems. At runtime, the specified assertions can be verified and the framework provides additional capabilities for linking domain-specific models with the system interception code. Drusinsky [86] have presented Temporal Rover supporting Linear Time Temporal Logic (LTL) including real-time, as well as time series constraints. This facilitates monitoring properties such as stability and calculating sum, average, and temporal minima and maxima. While these approaches provide an important foundation for runtime verification and formal constraint checking, with *GRuM* we further focus on providing support for the entire lifecycle, reducing the burden of creating monitors, and allowing for easy adaptation of constraints, monitors, and instrumentation when the underlying system evolves.

10. Conclusion

In this paper, we present *GRuM*, a novel framework for automatically generating runtime monitors. *GRuM* leverages model-driven technologies for defining a Domain Model Fragment and WM, describing relevant properties and data. Depending on the *SuM*'s technology and architectural style, different types of probes collecting runtime data can be generated. Additionally, *GRuM* automatically generates a customized monitoring platform consisting of a runtime model, a model query engine for defining constraint checks, and a dedicated middleware for connecting external services and applications. We evaluated our approach on two different systems, a UAV management and control system, and a set of TurtleBot 3 robots. The results of our experiments have shown that our implementation of a model-driven framework for automatically generating a runtime monitoring platform is well suited for describing the monitoring needs of non-real-time systems for collecting and checking large amounts of runtime data.

As part of our ongoing future work, we plan to explore ways to support dynamic

reconfiguration of the monitoring platform and to provide additional evolution support by synchronizing the models with the *SuM*. Furthermore, we are exploring adding self-adaptation capabilities [87] to our *GRuM* framework. Probes then, instead of just sending data to the runtime model, provide an additional back channel to receive and execute adaptation instructions, for example, triggered by the constraint engine when a specific constraint check fails.

References

[1] C. Mayr-Dorn, M. Winterer, C. Salomon, D. Hohensinger, R. Ramler, Considerations for using block-based languages for industrial robot programming-a case study, in: Proc. of the 2021 IEEE/ACM 3rd Int'l WS on Robotics Software Engineering, IEEE, 2021, pp. 5–12.

[2] D. Sykes, G. Keighren, Industrial-scale environments with bounded uncertainty: a productivity maximisation challenge, in: Proc. of the 2018 IEEE/ACM 1st Int'l WS on Robotics Software Engineering, IEEE, 2018, pp. 29–32.

[3] M. Eichleay, E. Evens, K. Stankevitz, C. Parker, Using the unmanned aerial vehicle delivery decision tool to consider transporting medical supplies via drone, Global Health: Science and Practice 7 (4) (2019) 500–506.

[4] M. Schörner, C. Wanninger, A. Hoffmann, O. Kosak, W. Reif, Architecture for Emergency Control of Autonomous UAV Ensembles, in: Proc. of the 2021 IEEE/ACM 3rd International Workshop on Robotics Software Engineering (RoSE), 2021, pp. 41–46.

[5] N. Nikolakis, V. Maratos, S. Makris, A cyber physical system (cps) approach for safe human-robot collaboration in a shared workplace, Robotics and Computer-Integrated Manufacturing 56 (2019) 233–243.

[6] M. Vierhauser, A. I. Md Nafee, A. Agrawal, J. Cleland-Huang, J. Mason, Hazard analysis for human-on-the-loop interactions in suas systems, in: Proc. of the 29th ACM Joint European Software Engineering Conf. and Symp. on the Foundations of Software Engineering, 2021, pp. 8–19.

[7] H. Shakhatreh, A. H. Sawalmeh, A. Al-Fuqaha, Z. Dou, E. Almaita, I. Khalil, N. S. Othman, A. Khreishah, M. Guizani, Unmanned aerial vehicles (UAVs): A survey on civil applications and key research challenges, IEEE Access 7 (2019) 48572–48634.

[8] B. Li, S. Ji, L. Liao, D. Qiu, M. Sun, Monitoring web services for conformance, in: Proc. of the 7th Int'l Symp. on Service Oriented System Engineering, IEEE, 2013, pp. 92–102.

[9] H. Eichelberger, K. Schmid, Flexible resource monitoring of java programs, Journal of Systems and Software 93 (2014) 163–186.

[10] A. B. Hamida, A. Bertolino, A. Calabrò, G. De Angelis, N. Lago, J. Lesbegueries, Monitoring service choreographies from multiple sources, in: Proc. of the 4th Int'l WS on Software Engineering for Resilient Systems, Springer, 2012, pp. 134–149.

[11] R. Popescu, A. Staikopoulos, S. Clarke, An extensible monitoring and adaptation framework, in: Proc. of the 2009 ICSOC/ServiceWave Workshops, Springer, 2010, pp. 314–324.

[12] C. Ruz, F. Baude, B. Sauvan, Enabling sla monitoring for component-based soa applications–a component-based approach, Proc. of the Work in Progress Session SEAA (2010) 41–42.

[13] A. Keller, H. Ludwig, The wsla framework: Specifying and monitoring service level agreements for web services, Journal of Network and Systems Management (2003) 57–81.

[14] O. Reynolds, A. García-Domínguez, N. Bencomo, Towards automated provenance collection for runtime models to record system history, in: Proc. of the 12th System Analysis and Modelling Conf., 2020, pp. 12–21.

[15] T. Brand, H. Giese, Generic adaptive monitoring based on executed architecture runtime model queries and events, in: Proc. of the 13th Int'l Conf. on Self-Adaptive and Self-Organizing Systems, IEEE, 2019, pp. 17–22.

[16] R. Rabiser, S. Guinea, M. Vierhauser, L. Baresi, P. Grünbacher, A comparison framework for runtime monitoring approaches, Journal of Systems and Software 125 (2017) 309–321.
[17] M. Vierhauser, H. Marah, A. Garmendia, J. Cleland-Huang, M. Wimmer, Towards a model-integrated runtime monitoring infrastructure for Cyber-Physical Systems, in: Proc. of the 2021 IEEE/ACM 43rd Int'l Conf. on Software Engineering: New Ideas and Emerging Results, IEEE, 2021, pp. 96–100.
[18] M. Stadler, M. Vierhauser, A. Garmendia, M. Wimmer, J. Cleland-Huang, Flexible model-driven runtime monitoring support for cyber-physical systems, in: 2022 IEEE/ACM 44th International Conference on Software Engineering: Companion Proceedings (ICSE-Companion), 2022, pp. 350–351.
[19] W. Kritzinger, M. Karner, G. Traar, J. Henjes, W. Sihn, Digital twin in manufacturing: A categorical literature review and classification, IFAC-PapersOnLine 51 (11) (2018) 1016–1022.
[20] T. H.-J. Uhlemann, C. Lehmann, R. Steinhilper, The digital twin: Realizing the cyber-physical production system for industry 4.0, Procedia Cirp 61 (2017) 335–340.
[21] D. Weyns, B. Schmerl, V. Grassi, S. Malek, R. Mirandola, C. Prehofer, J. Wuttke, J. Andersson, H. Giese, K. M. Göschka, On patterns for decentralized control in self-adaptive systems, in: Software Engineering for Self-Adaptive Systems II, Springer, 2013, pp. 76–107.
[22] I. Elgendi, M. F. Hossain, A. Jamalipour, K. S. Munasinghe, Protecting cyber physical systems using a learned MAPE-K model, IEEE Access 7 (2019) 90954–90963.
[23] M. Vierhauser, R. Rabiser, P. Grünbacher, Requirements monitoring frameworks: A systematic review, Information and Software Technology 80 (2016) 89–109.
[24] E. Zavala, X. Franch, J. Marco, Adaptive monitoring: A systematic mapping, Information and software technology 105 (2019) 161–189.
[25] L. Baresi, S. Guinea, Towards dynamic monitoring of ws-bpel processes, in: Proc. of the Int'l Conf. on Service-Oriented Computing, Springer, 2005, pp. 269–282.
[26] Y. Wang, S. A. McIlraith, Y. Yu, J. Mylopoulos, Monitoring and diagnosing software requirements, Automated Software Engineering (2009) 3–35.
[27] L. Baresi, L. Pasquale, P. Spoletini, Fuzzy goals for requirements-driven adaptation, in: Proc. of the Int'l Requirements Engineering Conf., IEEE, 2010, pp. 125–134.
[28] R. Rabiser, K. Schmid, H. Eichelberger, M. Vierhauser, S. Guinea, P. Grünbacher, A domain analysis of resource and requirements monitoring: Towards a comprehensive model of the software monitoring domain, Information and Software Technology 111 (2019) 86–109.
[29] S. Martínez-Fernández, A. M. Vollmer, A. Jedlitschka, X. Franch, L. López, P. Ram, P. Rodríguez, S. Aaramaa, A. Bagnato, M. Choraś, et al., Continuously assessing and improving software quality with software analytics tools: a case study, IEEE access 7 (2019) 68219–68239.
[30] R. France, B. Rumpe, Model-driven development of complex software: A research roadmap, in: Future of Software Engineering (FOSE'07), IEEE, 2007, pp. 37–54.
[31] J. Cleland-Huang, M. Vierhauser, S. Bayley, Dronology: an incubator for Cyber-Physical Systems research, in: Proc. of the 40th Int'l Conf. on Software Engineering: New Ideas and Emerging Results, 2018, pp. 109–112.
[32] A. Goldberg, K. Haveland, Instrumentation of java bytecode for runtime analysis, in: Formal Techniques for Java-like Programs, no. NAS2-00065, 2003.
[33] I. Cassar, A. Francalanza, L. Aceto, A. Ingólfsdóttir, et al., A survey of runtime monitoring instrumentation techniques, in: Proceedings Second International Workshop on Pre-and Post-Deployment Verification Techniques, PrePost@ iFM 2017, 2017, pp. 15–28.
[34] A. Bucchiarone, J. Cabot, R. F. Paige, A. Pierantonio, Grand challenges in model-driven engineering: an analysis of the state of the research, SoSyM 19 (1) (2020) 5–13.
[35] E. Quiñones, S. Royuela, C. Scordino, P. Gai, L. M. Pinho, L. Nogueira, J. Rollo, T. Cucinotta, A. Biondi, A. Hamann, D. Ziegenbein, H. Saoud, R. Soulat, B. Forsberg, L. Benini, G. Mando, L. Rucher, The ampere project: A model-driven development framework for highly parallel and energy-efficient computation supporting multi-criteria optimization, in: Proc. of the 23rd Int'l Symp. on Real-Time Dist. Comp., 2020, pp. 201–206. doi:10.1109/ISORC49007.2020.00042.
[36] T. Z. Asici, B. Karaduman, R. Eslampanah, M. Challenger, J. Denil, H. Vangheluwe, Applying model

driven engineering techniques to the development of contiki-based iot systems, in: Proc. of the 1st Int'l WS on Software Engineering Research Practices for the Internet of Things, 2019, pp. 25–32.
[37] PX4, Open Source Flight Controller, https://px4.io, [Last accessed 01-06-2021] (2021).
[38] M. Mansouri-Samani, M. Sloman, Monitoring distributed systems, IEEE network 7 (6) (1993) 20–30.
[39] M. D. Del Fabro, J. Bézivin, F. Jouault, E. Breton, G. Gueltas, Amw: a generic model weaver, in: 1 ere Journées sur l'Ingénierie Dirigée par les Modèles (IDM05), 2005, pp. 105–114.
[40] M. D. Del Fabro, J. Bézivin, F. Jouault, P. Valduriez, et al., Applying generic model management to data mapping, in: BDA, 2005.
[41] M. D. Del Fabro, J. Bézivin, P. Valduriez, Weaving models with the eclipse amw plugin, in: Eclipse Modeling Symposium, Eclipse Summit Europe, Vol. 2006, Citeseer, 2006, pp. 37–44.
[42] F. Jouault, B. Vanhooff, H. Bruneliere, G. Doux, Y. Berbers, J. Bézivin, Inter-dsl coordination support by combining megamodeling and model weaving, in: Proceedings of the 2010 ACM Symposium on Applied Computing, 2010, pp. 2011–2018.
[43] R. Hawkins, I. Habli, D. Kolovos, R. Paige, T. Kelly, Weaving an assurance case from design: a model-based approach, in: 2015 IEEE 16th International Symposium on High Assurance Systems Engineering, IEEE, 2015, pp. 110–117.
[44] K. Aizawa, K. Tei, S. Honiden, Identifying safety properties guaranteed in changed environment at runtime, in: Proc. of the 2018 IEEE Int'l Conf. on Agents, IEEE, 2018, pp. 75–80.
[45] N. Maiden, Monitoring our requirements, IEEE software 30 (1) (2013) 16–17.
[46] V. R. B. G. Caldiera, H. D. Rombach, The goal question metric approach, Encyclopedia of software engineering (1994) 528–532.
[47] M. Galster, E. Bucherer, A taxonomy for identifying and specifying non-functional requirements in service-oriented development, in: Proc. of the 2008 IEEE Congress on Services-Part I, IEEE, 2008, pp. 345–352.
[48] E. Foundaition, Eclipse modeling framework (emf), https://www.eclipse.org/modeling/emf, [Last accessed 01-06-2021] (2021).
[49] JBoss, Roaster java source parser library, https://github.com/forge/roaster, [Last accessed 01-06-2021] (2021).
[50] Eclipse Foundation, Xtext language engineering framework, https://www.eclipse.org/Xtext, [Last accessed 01-06-2021] (2021).
[51] MQTT, Mqtt message broker, https://mqtt.org, [Last accessed 01-06-2021] (2023).
[52] G. Bergmann, Á. Horváth, I. Ráth, D. Varró, A. Balogh, Z. Balogh, A. Ökrös, Incremental evaluation of model queries over emf models, in: International Conference on Model Driven Engineering Languages and Systems, Springer, 2010, pp. 76–90.
[53] G. Bergmann, I. Dávid, Á. Hegedüs, Á. Horváth, I. Ráth, Z. Ujhelyi, D. Varró, Viatra 3: A reactive model transformation platform, in: Proc. of the Int'l Conf. on Theory and Practice of Model Transformations, Springer, 2015, pp. 101–110.
[54] Espertech, Esper complex event processing, https://www.espertech.com/esper, [Last accessed 01-06-2021] (2021).
[55] ROBOTIS, ROBOTIS e-Manual for TurtleBot3, [Last accessed 01-06-2021] (2021).
URL https://emanual.robotis.com/docs/en/platform/turtlebot3/overview
[56] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, et al., Ros: an open-source robot operating system, in: ICRA workshop on open source software, Vol. 3, Kobe, Japan, 2009, p. 5.
[57] H. Aagela, M. Al-Nesf, V. Holmes, An asus_xtion_probased indoor mapping using a raspberry pi with turtlebot robot turtlebot robot, in: Proc. of the 23rd Int'l Conf. on Automation and Computing, IEEE, 2017, pp. 1–5.
[58] A. Koubâa, M.-F. Sriti, Y. Javed, M. Alajlan, B. Qureshi, F. Ellouze, A. Mahmoud, Turtlebot at office: A service-oriented software architecture for personal assistant robots using ros, in: Proc. of the 2016 Int'l Conf. on Autonomous Robot Systems and Competitions, IEEE, 2016, pp. 270–276.
[59] R. Amsters, P. Slaets, Turtlebot 3 as a robotics education platform, in: Proc. of the Int'l Conf. on

Robotics in Education, Springer, 2019, pp. 170–181.
[60] Open Robotics, Gazebo, https://gazebosim.org/home, [Last accessed 01-01-2023] (2023).
[61] Software Engineering Datasets for Small Unmanned Aerial Systems, `https://github.com/SAREC-Lab/sUAS-UseCases`, [Last accessed 01-06-2021] (2021).
[62] J. Pimentel, J. Castro, E. Santos, A. Finkelstein, Towards requirements and architecture co-evolution, in: International Conference on Advanced Information Systems Engineering, Springer, 2012, pp. 159–170.
[63] L. Cianciaruso, F. Di Forenza, E. Di Nitto, M. Miglierina, N. Ferry, A. Solberg, Using models at runtime to support adaptable monitoring of multi-clouds applications, in: 2014 16th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, IEEE, 2014, pp. 401–408.
[64] R. Morrison, D. Balasubramaniam, F. Oquendo, B. Warboys, R. M. Greenwood, An active architecture approach to dynamic systems co-evolution, in: European Conference on Software Architecture, Springer, 2007, pp. 2–10.
[65] Eclipse, Ecore tools - graphical modeling for ecore, `https://www.eclipse.org/ecoretools`, [Last accessed 01-06-2022] (2021).
[66] H. Eichelberger, K. Schmid, Flexible resource monitoring of java programs, Journal of Systems and Software (2014) 163–186.
[67] M. Kim, S. Kannan, I. Lee, O. Sokolsky, M. Viswanathan, Java-mac: a run-time assurance tool for java programs, Electronic Notes in Theoretical Computer Science 55 (2) (2001) 218–235.
[68] J. Ehlers, W. Hasselbring, A self-adaptive monitoring framework for component-based software systems, in: Proc. of the 5th Europ. Conf. on Software Architecture, Springer, 2011, pp. 278–286.
[69] C. Inzinger, B. Satzger, W. Hummer, S. Dustdar, Specification and deployment of distributed monitoring and adaptation infrastructures, in: Proc. of the Service-Oriented Computing-ICSOC 2012 Workshops: ICSOC 2012, International Workshops ASC, DISA, PAASC, SCEB, SeMaPS, WESOA, and Satellite Events, Shanghai, China, November 12-15, 2012, Revised Selected Papers 10, Springer, 2013, pp. 167–178.
[70] DDS Foundation, Data distribution service middleware, `https://www.dds-foundation.org`, [Last accessed 01-06-2021] (2021).
[71] K. Havelund, G. Roşu, Monitoring java programs with java pathexplorer, Electronic Notes in Theoretical Computer Science 55 (2) (2001) 200–217.
[72] P. Doherty, J. Kvarnstròm, F. Heintz, A temporal logic-based planning and execution monitoring framework for unmanned aircraft systems, Autonomous Agents and Multi-Agent Systems 19 (3) (2009) 332–377.
[73] M. Vierhauser, R. Rabiser, P. Grünbacher, K. Seyerlehner, S. Wallner, H. Zeisel, Reminds: A flexible runtime monitoring framework for systems of systems, Journal of Systems and Software 112 (2016) 123–136.
[74] M. Machin, F. Dufossé, J.-P. Blanquart, J. Guiochet, D. Powell, H. Waeselynck, Specifying safety monitors for autonomous systems using model-checking, in: Proc. of the 33rd Int'l Conf. on Computer Safety, Springer, 2014, pp. 262–277.
[75] E. Al-Shaer, H. Abdel-Wahab, K. Maly, Hifi: A new monitoring architecture for distributed systems management, in: Proc. of the 19th IEEE Int'l Conf. on Distributed Computing Systems, IEEE, 1999, pp. 171–178.
[76] F. Chen, G. Rocsu, Mop: An efficient and generic runtime verification framework, in: Proc. of the 22nd Annual ACM SIGPLAN Conf. on Object-oriented Programming Systems and Applications, ACM, 2007, pp. 569–588.
[77] J. C. Kirchhof, J. Michael, B. Rumpe, S. Varga, A. Wortmann, Model-driven digital twin construction: synthesizing the integration of Cyber-Physical Systems with their information systems, in: Proc. of the 23rd ACM/IEEE Int'l Conf. on Model Driven Engineering Languages and Systems, 2020, pp. 90–101.
[78] G. Blair, N. Bencomo, R. B. France, Models@ run.time, Computer 42 (10) (2009) 22–27.

[79] N. Bencomo, S. Götz, H. Song, Models@run.time: a guided tour of the state of the art and research challenges, Software & Systems Modeling 18 (5) (2019) 3049–3082.
[80] N. Hili, M. Bagherzadeh, K. Jahed, J. Dingel, A model-based architecture for interactive run-time monitoring, Software and Systems Modeling (2020) 1–23.
[81] C. H. Corbato, D. Bozhinoski, M. G. Oviedo, G. van der Hoorn, N. H. Garcia, H. Deshpande, J. Tjerngren, A. Wasowski, Mros: Runtime adaptation for robot control architectures, arXiv preprint arXiv:2010.09145 (2020).
[82] T. Brand, H. Giese, Towards software architecture runtime models for continuous adaptive monitoring, in: Proc. of the 13th Int'l WS on Models@run.time, 2018, pp. 72–77.
[83] L. Sakizloglou, S. Ghahremani, T. Brand, M. Barkowsky, H. Giese, Towards highly scalable runtime models with history, in: Proc. of the IEEE/ACM 15th Int'l Symp. on Software Engineering for Adaptive and Self-Managing Systems, ACM, 2020, p. 188–194.
[84] M. Búr, G. Szilágyi, A. Vörös, D. Varró, Distributed graph queries for runtime monitoring of Cyber-Physical Systems, in: International Conference on Fundamental Approaches to Software Engineering, Springer, Cham, 2018, pp. 111–128.
[85] K. Chan, I. Poernomo, H. Schmidt, J. Jayaputera, A model-oriented framework for runtime monitoring of nonfunctional properties, in: Quality of Software Architectures and Software Quality, Springer, 2005, pp. 38–52.
[86] D. Drusinsky, Monitoring temporal rules combined with time series, in: International Conference on Computer Aided Verification, Springer, 2003, pp. 114–117.
[87] J. O. Kephart, D. M. Chess, The vision of autonomic computing, Computer 36 (1) (2003) 41–50.