

MoULDyS: Monitoring of Autonomous Systems in the Presence of Uncertainties

Bineet Ghosh^a, Étienne André^{b,c}

^aThe University of North Carolina at Chapel Hill, NC, USA

^bUniversité de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

^cUniversité Sorbonne Paris Nord, LIPN, CNRS UMR 7030, F-93430 Villetaneuse, France

Abstract

We introduce MoULDyS, that implements efficient offline and online monitoring algorithms of black-box cyber-physical systems w.r.t. safety properties. MoULDyS takes as input an uncertain log (with noisy and missing samples), as well as a bounding model in the form of an uncertain linear system; this latter model plays the role of an over-approximation so as to reduce the number of false alarms. MoULDyS is Python-based and available under the *GNU General Public License v3.0* (gp1-3.0). We further provide easy-to-use scripts to recreate the results of two case studies introduced in an earlier work.

Keywords: energy-aware monitoring, cyber-physical systems, formal methods, monitoring tool.

1. Motivation and Significance

Monitoring consists of analyzing system logs, e.g., for detecting safety violations (see, e.g., [4]). Monitoring has many useful applications such as detecting the cause of a crash of vehicles. As an example, autonomous systems are generally equipped with a device that records their state at periodic or aperiodic time steps—logging the behavior of the system until the time of a failure. A log, comprising of such recorded samples, is then investigated for possible safety violations. Not only the logs can have samples missing at various time steps, but also the recorded samples can have added noise to it, e.g., due to sensor uncertainties. Analyzing such logs to detect possible safety violations, that might have caused a failure, is known as *offline monitoring* when the analysis is done *a posteriori* (see, e.g., [5]). In contrast, it is *online* when performed on-the-fly, when the whole log is not (yet) known (see [6] for a discussion on online verification).

Nr.	Code metadata description	Please fill in this column
C1	Current code version	v1.1
C2	Permanent link to code/repository used for this code version	https://github.com/bineet-coderep/MoULDyS/releases/tag/v1.1
C3	Permanent link to Reproducible Capsule	10.5281/zenodo.7888502
C4	Legal Code License	<i>GNU General Public License v3.0</i> (gpl-3.0) [1]
C5	Code versioning system used	git
C6	Software code languages, tools, and services used	Python, numpy, scipy, mpmath, pandas, Gurobi
C7	Compilation requirements, operating environments and dependencies	Provided in the installation guide [2].
C8	If available, link to developer documentation/manual	Provided in the user guide [3].
C9	Support email for questions	bineet@cs.unc.edu

Table 1: Code metadata




Figure 1: Monitoring at discrete time steps [9]

1 We introduce here MoULDyS¹ (see Table 1 for code metadata), a monitor-
2 ing tool to analyze logs to detect possible safety violations. The specific fea-
3 tures of MoULDyS are twofold: 1) the possibility to monitor aperiodic logs, or
4 periodic logs with missing samples, and with possible noise over the recorded
5 data; and 2) the presence of a bounding model following the formalism of
6 uncertain linear systems.

7 *Uncertain linear systems.* Uncertain linear systems [7, 8] are a special sub-
8 class of non-linear systems that can be used to represent uncertainties and
9 parameters in linear dynamical systems—for example, they can represent un-
10 certain parameters in the cells of the dynamics matrix. Such a formalism is
11 useful in representing an over-approximation of the model when the precise
12 model is unknown or hard to obtain.

¹<https://sites.google.com/view/mouldys>

1 *Monitoring using uncertain linear systems as a bounding model.* A challenge
2 when performing offline monitoring is to “recreate” or guess the samples at
3 the missing time steps. That is, when the system under monitoring is a
4 black-box, with a log in the form of an aperiodic timed sequence of valua-
5 tions of continuous variables (with missing valuations at various time steps):
6 *how to be certain that in between two discrete valuations the specification*
7 *was not violated at another discrete time step at which no logging was per-*
8 *formed?* Consider Fig. 1a, a system for which a logging occurs at every
9 discrete time step. When logging occurs at only *some* time steps (due to
10 some sensor faults, or to save energy with only a sparse, scattered logging),
11 a possible output is in Fig. 1b. In such a setting, *how to be certain that,*
12 *in between two discrete samples, another discrete sample (not recorded) did*
13 *not violate the specification?* That is, in Fig. 1b, there is no way to formally
14 guarantee that the unsafe zone (i.e., above the red, dashed line) was never
15 reached by another discrete sample which was not recorded. In many prac-
16 tical cases, a piecewise-constant or linear approximation (see, e.g., Figs. 1c
17 and 1d, where the large blue dots denote actual samples, while the small
18 green dots denote reconstructed samples using some extrapolation) is arbi-
19 trary and not appropriate—it can yield a “safe” answer, while the actual
20 system could have actually been unsafe at some of the missing time steps.
21 On the contrary, assuming a completely arbitrary dynamics will always yield
22 “potentially unsafe”—thus removing the interest of monitoring. Without any
23 knowledge of the model, one can always assume that the behavior given in
24 Fig. 1e could happen. This behavior shows that the variable x is suddenly
25 crossing the unsafe region (dashed) at some unlogged discrete time step—
26 even though this is unlikely if the dynamics is known to vary “not very fast”.
27 To alleviate such issues, we proposed in [9] an offline monitoring algorithm us-
28 ing a *bounding model*, i.e., a rough overapproximation of the system behavior
29 (originally introduced in [10] in a different context). The proposed method
30 is based on the reachable set computation of uncertain linear systems [11]
31 that can detect safety violations with limited false alarms.

32 We also considered in [9] an *online* monitoring algorithm, aiming at en-
33 ergetic efficiency, by recording samples only when required (i.e., when the
34 system may get closer to a violation). MoULDyS implements both our offline
35 and online monitoring algorithms [9]. We also provide here the steps to easily
36 recreate the results of the two case studies in [9].

37 *Experimental setting.* Given an aperiodic (i.e., missing valuations at various
38 time steps) and a noisy log (i.e., the valuations that are present in the log
39 can have an added noise—an overapproximation of the actual state), MoULDyS
40 can perform offline monitoring of the system to detect safety of the system

1 behavior. Further, MoULDyS can be used in an online setting to log only when
2 necessary—thus targeting energy efficiency while logging. MoULDyS can be
3 run on a standard laptop with a Linux operating system (see the installation
4 guide for details [2]). The details of how to use the tool, with illustrative
5 examples, can be found in the user guide [3].

6 *Outline.* Section 2 describes the software; Section 3 describes the architec-
7 ture; Section 4 describes the functionalities; Section 5 describes two illustra-
8 tive examples from medical and automotive domains, with necessary steps
9 to recreate their results; Section 6 discusses the impact of MoULDyS; Sec-
10 tion 7 briefly reviews related works; Section 8 concludes and discusses future
11 research.

12 2. Software Description

13 MoULDyS is an open-source software, implemented in Python, running
14 on Linux platforms. Our experiments [9] suggest that MoULDyS is able to
15 perform monitoring of reasonably large systems in a reasonable time. For
16 example, we are able to monitor a five-dimensional system (i.e., a system
17 with 5 continuous variables monitored at the same time) for 2000 time steps,
18 with only actual 300 samples (note that, the fewer samples, the higher is
19 the monitor computation time—as it is required to “recreate” the missing
20 samples using reachable sets) in under 2.5 minutes on a standard laptop.

21 We believe that MoULDyS will not just be helpful to engineers to analyze
22 logs to detect safety violations in several areas of research (such as in robotics
23 to detect collision and other undesirable behaviors), but also to researchers
24 to further develop monitoring-based approaches—in that case MoULDyS can
25 be used for comparison.

26 3. Software Architecture

27 The architecture of MoULDyS is given in Fig. 2. Each block in Fig. 2 rep-
28 represents a core part (either functional or input) of the tool, and the arrows
29 indicate the flow of data. The dataflow of MoULDyS, for both Figs. 3 and 4
30 (offline and online monitoring respectively), starts from the blocks in the left
31 (e.g., bounding model, unsafe set, etc.) and ends at the extreme right of the
32 figures (outputting the safety status, visualization, and/or synthesized log).
33 Since both the offline (Block 1 of Figs. 2 and 3) and online monitoring algo-
34 rithm (Block 2 in Figs. 2 and 4) uses reachability of uncertain linear systems
35 (Block 3 of Fig. 2), a data exchange occurs between Block 1 and Block 3,
36 as well as Block 2 and Block 3. MoULDyS implements a built-in reachability

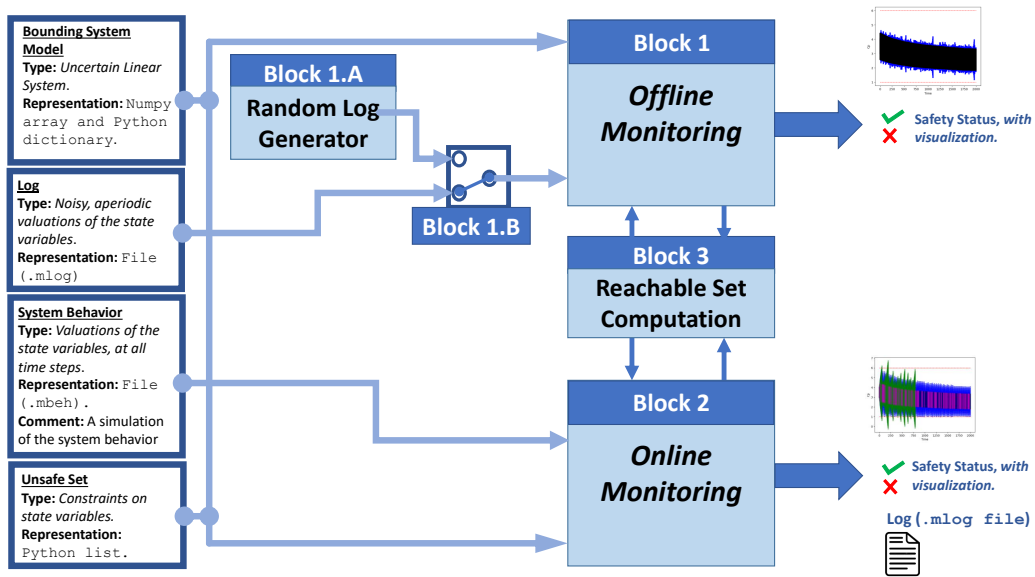


Figure 2: MoULDyS *Architecture*. Each block identifies a core part of the tool, and the arrows indicate the flow of data.

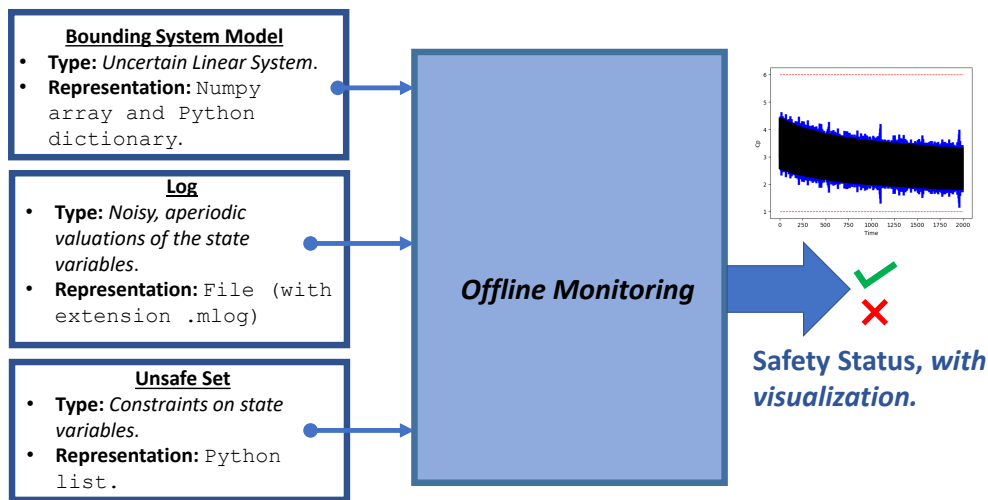


Figure 3: Dataflow diagram of the offline monitoring functionality of MoULDyS.

- 1 algorithm. The native reachable set computation support facilitates faster
- 2 computing (i.e., no data exchange with third party tool is required, which
- 3 would potentially require additional data reformatting) with no additional
- 4 tool installation required. MoULDyS employs an algorithm proposed in [11] to

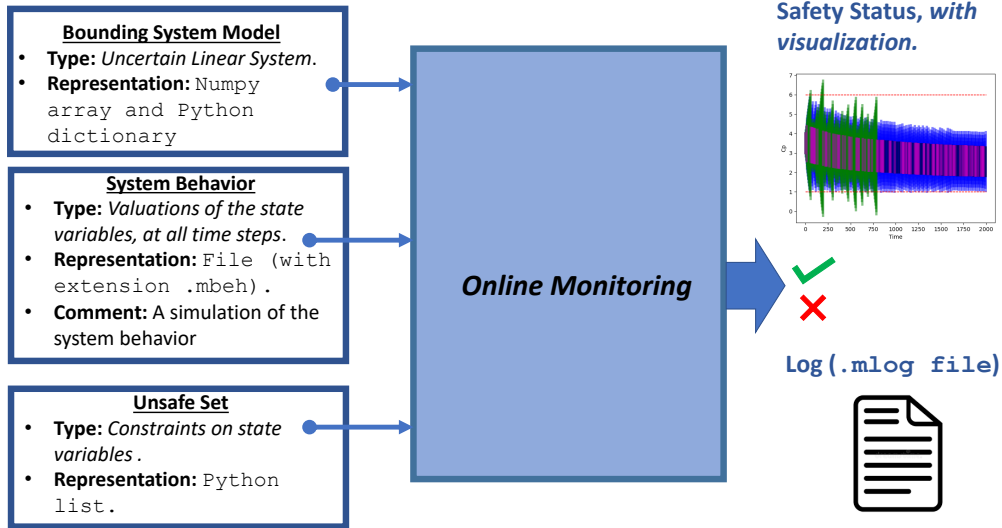


Figure 4: Dataflow diagram of the online monitoring functionality of MoULDyS.

1 compute the reachable set of uncertain linear systems. The algorithm first
 2 computes the reachable set of the nominal dynamics (which excludes uncer-
 3 tainties) and then computes the reachable set related to the uncertainties
 4 in the dynamics. These two sets are then combined using the Minkowski
 5 sum to obtain the reachable set of the entire dynamics. Although computing
 6 the reachable set of the nominal dynamics is straightforward, the reachable
 7 set related to uncertainties is challenging to compute. After obtaining the
 8 reachable sets, MoULDyS verifies the safety of these sets by comparing them
 9 against provided safety specifications. These safety specifications are con-
 10 straints on state variables, which can be complex and involve multiple state
 11 variables (e.g., linear inequalities involving several state variables). To rep-
 12 resent such safety specifications, MoULDyS uses a special type of polytope
 13 called *zonotopes*, which can be expressed as an affine transformation of a
 14 unit box. MoULDyS can check multiple safety specifications, each represented
 15 as a zonotope and involving several state variables.

16 3.1. Implementation

17 MoULDyS is implemented using Python 3.7.x, and runs in a Linux envi-
 18 ronment. The architecture is given in Fig. 2.

19 The tool can be used in the following two ways. On the one hand, users

1 can use it through the provided virtual machine², which already contains all
2 the necessary dependencies and has the path variable set. Nevertheless, users
3 are still required to obtain and install the Gurobi license themselves, since
4 Gurobi only grants free academic licenses to individuals. On the other hand,
5 the tool can also be downloaded and setup from its public GitHub repository³.
6 If users aim to recreate the results in the paper or simply employ it for basic
7 monitoring purposes, using the provided virtual machine is recommended.
8 However, if the tool will be used for research and development purposes, it
9 is recommended to download and set it up on a local machine. The detailed
10 installation instructions are provided in [2]. The system model (represented
11 as an uncertain linear system), in MoULDyS, is represented with a `numpy`
12 array and a dictionary. The unsafe set is represented with a Python list. An
13 example code of how to encode the system model and its unsafe specification
14 can be found in the public repository⁴. The log and the system behavior,
15 on the other hand, is given as a file to MoULDyS (MoULDyS can also generate
16 random logs; discussed later). Logs can either be represented as zonotopes
17 or intervals (an example of the required file can be found in MoULDyS public
18 repository⁵).

19 Both the online and offline monitoring algorithm (Block 1 and Block 2
20 of Fig. 2) have been implemented in Python using standard libraries, namely
21 `numpy`, `scipy` and `mpmath`. The reachable set computation (Block 3 of
22 Fig. 2) module implements the algorithm proposed in [11] in Python. Both
23 the online and the offline module require performing intersection checking of
24 zonotopes [9], which has been implemented as an optimization formulation
25 using Gurobi. Gurobi has been further used to visualize the reachable sets.

26 *Installation.* While the virtual machine comes preinstalled with the required
27 dependencies, setting up MoULDyS on a local machine requires installing the
28 following dependencies: `numpy`, `scipy`, `mpmath`, `pandas`, `Gurobi` along with
29 `gurobipy` (Gurobi requires an *ad-hoc* install but is free to use for academic
30 purposes). The detailed steps for installing MoULDyS are given in the instal-
31 lation guide [2].

32 *User Guide.* A tutorial on how to use several functionalities of MoULDyS,
33 along with sample codes (encoding a toy dynamics), is given in the user

²[10.5281/zenodo.7888502](https://doi.org/10.5281/zenodo.7888502)

³<https://github.com/bineet-coderep/MoULDyS/releases/tag/v1.1>

⁴<https://www.github.com/bineet-coderep/MoULDyS/blob/main/src/tutorial/TutorialOfflineMonitoring.py>

⁵https://www.github.com/bineet-coderep/MoULDyS/blob/main/data/toyEg_5_interval.mlog

1 guide [3].

2 4. Software Functionalities

3 In this section, we discuss the core functionalities of MoULDyS:

4 **Offline Monitoring** The offline monitoring requires the bounding model of
5 the system (represented as an uncertain linear system) to be given as
6 input. Further, a log of the system behavior is required, which can be
7 achieved by the following ways: *i*) If the user already has a log to be
8 monitored, it can be simply passed as an input to Block 1 of Fig. 2.
9 *ii*) Alternatively, MoULDyS can also generate a random (noisy and aperi-
10 odic) log of the system, from a given initial set, using Block 1.A. The
11 selection between the two possible choices is facilitated by Block 1.B.
12 Analyzing the log, the final output of the offline monitoring is either
13 **safe** (indicating the system behavior is certainly safe at all time steps),
14 or **possibly-unsafe** (indicating the system might have shown unsafe
15 behavior). An example code snippet to perform offline monitoring, on
16 a toy example, can be found in its public repository.⁶

17 **Online Monitoring** The online monitoring requires the bounding model of
18 the system (represented as an uncertain linear system) to be given as
19 input, as well as the actual behavior of the system. The actual behav-
20 ior of the system is given as a file (with extension `.mbeh`) representing
21 the values of the state variables at every time step—an example of
22 the expected file (with extension `.mbeh`), containing valuation of the
23 state variables at every time step, can be found in its public reposi-
24 tory.⁷ The final output of this feature is the safety status of the system
25 (**safe/possibly-unsafe**), and a synthesized log. An example code
26 snippet to perform online monitoring, on a toy example, can be found
27 in its public repository.⁸

28 As a side functionality, MoULDyS also allows to generate random logs
29 (using Block 1.A). While this is not strictly speaking part of monitoring,
30 it greatly helps to perform experiments using MoULDyS. Basically, given a

⁶<https://www.github.com/bineet-coderep/MoULDyS/blob/main/src/tutorial/TutorialOfflineMonitoring.py>

⁷https://www.github.com/bineet-coderep/MoULDyS/blob/main/data/toyEg_5_interval.mbeh

⁸<https://www.github.com/bineet-coderep/MoULDyS/blob/main/src/tutorial/TutorialOnlineMonitoring.py>

1 bounding model, MoULDyS can generate a random log following the bounding
2 model.

3 5. Illustrative examples

4 In this section, we briefly recall the two case studies presented in [9]
5 that use a prototype version of MoULDyS. The two case studies, automated
6 anesthesia delivery and adaptive cruise control, demonstrate the applicability
7 and usability of MoULDyS. Further, we provide detailed steps to recreate the
8 results presented in [9] using MoULDyS.

9 **Anesthesia** [12] presents an automated anaesthesia delivery model, with
10 the drug propofol. The system models the metabolization of the drug
11 by the body, and the depth of hypnosis. The state variables encode
12 the various concentration levels—that must be within a certain limit
13 at all times—modeling the metabolization of the drug and the depth of
14 hypnosis. Note that a higher concentration level would mean that the
15 patient remains unconscious for a longer period of time, while a lower
16 concentration level would mean that patient remains conscious during
17 the surgery—which can be traumatic. MoULDyS can help performing an
18 automated monitoring of the patients without compromising on safety.

19 **Adaptive Cruise Control (ACC)** [13] presents a model of ACC with
20 state variables as velocity, distance between two vehicles, and the ve-
21 locity of the lead vehicle. Offline monitoring provides an automated
22 way to detect the *cause of the crash* and *who was at fault*. Similarly,
23 consider a vehicle driving on a highway with a vehicle in its sight. The
24 ACC unit will have to continuously read sensor values to track sev-
25 eral parameters, such as acceleration of the lead vehicle, braking force,
26 etc.—causing a waste of energy. In these cases, deploying online mon-
27 itoring on the vehicle ACC will ensure that the sensor values are only
28 read when there is a potential unsafe behavior—thus saving energy.
29 [9] provides several such practical cases where monitoring is useful—in
30 [9, Figs. 5 and 6] (also recalled in the appendix in Figs. B.8 and B.9
31 respectively).

32 The case studies mentioned above study the following aspects with re-
33 gards to monitoring: *i*) Impact of number of samples in the log. *ii*) Impact
34 of uncertainties in the samples of the log. *iii*) Demonstrating online moni-
35 toring. *iv*) Comparing offline and online monitoring. In the following, we
36 provide the detailed steps to recreate the results in [9, Section 5]. In par-
37 ticular, the results that we wish to recreate, from [9], are given in Figs. B.6
38 to B.9.

1 5.1. Recreating Results

2 The results of the Anesthesia case study and the ACC case study can
3 be recreated by using the scripts provided in the GitHub repository⁹. The
4 detailed steps to recreate the results from both the case studies are given
5 in [14].

6 Note that, in [9, Section 5], logs were randomly generated *during our*
7 *experiments*, and therefore the results from [9, Section 5] cannot be *stricto*
8 *sensu* be recreated, as their reproducibility script re-generates a random log,
9 which may differ from the one actually shown in [9, Section 5].

10 Therefore, we modified our scripts so that the log is given as an input,
11 on which monitoring is performed: we thus generated the logs statically, and
12 embedded them in the reproducibility capsule, in order for users to reproduce
13 exactly the result we present here. These logs were generated with the same
14 logging probabilities (and initial sets) as in [9, Section 5]. In the rest of the
15 section, we discuss the steps for replicating the new Figs. B.6 to B.9.

16 *Anesthesia*. The main results of the Anesthesia case study are pro-
17 vided in Figs. B.6 and B.7 (variant of [9, Figures 3 and 4]). We
18 use `python Anesthesia.py -offline i`, with $i \in \{1, 2, 3, 4\}$ to recre-
19 ate Figs. B.6a to B.6d respectively. To recreate Fig. B.7a, we
20 use `python Anesthesia.py -online`. To recreate Fig. B.7b, we use
21 `python Anesthesia.py -compare`. The `Anesthesia.py` script is provided
22 in its GitHub repository.¹⁰

23 *ACC*. The main results of the ACC case study are provided in Figs. B.8
24 and B.9 (variant of [9, Figures 5 and 6]). The results of the ACC case study
25 can be recreated in a similar manner to the Anesthesia case study, by simply
26 using the `ACC.py`¹¹ script instead of the `Anesthesia.py` script.

27 6. Impact

28 While MoULDyS remains a prototype based on a recent algorithm [9], and
29 is by no means a widely used software in an industrial context for the time be-
30 ing, we believe it has an interesting potential to gain up a user base interested
31 in monitoring black-box cyber-physical systems against safety properties. To
32 the best of our knowledge, it is the first software allowing monitoring logs

⁹https://www.github.com/bineet-coderep/MoULDyS/tree/main/src/recreate_results_from_paper

¹⁰https://www.github.com/bineet-coderep/MoULDyS/blob/main/src/recreate_results_from_paper/Anesthesia.py

¹¹https://github.com/bineet-coderep/MoULDyS/blob/main/src/recreate_results_from_paper/ACC.py

1 featuring not only uncertainty (due to sensor’s imperfect behavior) but also
2 potentially missing samples, together with a bounding model going beyond
3 the class of linear systems. This is in contrast with, e.g., [10] in which our
4 bounding model was restricted to linear models.

5 In addition, MoULDyS can perform *online* monitoring, with a focus on
6 energetic efficiency: by triggering a sample recording only when necessary
7 (i.e., when MoULDyS informs the system that it may get close to an unsafe
8 behavior according to its online algorithm), the system saves energy, i.e., only
9 records samples (which needs network bandwidth usage, as well as processor
10 and memory usage) when necessary instead of at every time unit.

11 Our applications recalled in Section 5 show that MoULDyS can be applied
12 to challenging domains such as health and autonomous driving, giving inter-
13 esting results (e.g., limited number of false alarms) in a reasonable execution
14 time making it suitable to real-time applications.

15 7. Related works

16 Monitoring complex systems, and notably cyber-physical systems, drew
17 a lot of attention in the last decades [4]. We briefly review close works in the
18 following.

19 MONPOLY [15] is a monitoring tool taking as specification formulas ex-
20 pressed using MFOTL (metric first-order temporal logic). It is entirely
21 black-box: the only input beyond the formula is the *log*, i.e., a sequence of
22 timestamped system events, potentially with numeric arguments (e.g., “@10
23 `withdraw (Alice,6000)`”, expressing that a withdrawal occurs at times-
24 tamp 10).

25 In [16], the focus is on online monitoring over real-valued signals, using
26 MTL as the specification formalism. Again, the system is black-box.

27 In [17], parametric timed pattern matching is made, on an entirely black-
28 box system, i.e., without any prior knowledge of the system; the tools used
29 are IMITATOR [18] and a prototypal tool ParamMONAA. The output is a set
30 of intervals where a property is valid/violated, possibly with a set of timing
31 parameter valuations.

32 In [10], we proposed *model-bounded monitoring*: instead of monitoring a
33 black-box system against a sole specification, we use in addition a (limited,
34 over-approximated) knowledge of the system, to eliminate false positives.
35 This over-approximated knowledge is given in [10] in the form of a *linear hy-*
36 *brid automaton* (LHA) [19]. We use in [10] both an *ad-hoc* implementation,
37 and another one based on PHAVerLite [20]. In this work, we share with [10]
38 the principle of using an over-approximation of the model to rule out some

1 violation of the specification, which comes in contrast with the aforemen-
2 tioned works. However, we consider here a different formalism, and we work
3 on discrete samples. In terms of expressiveness of the over-approximated
4 model, while our approach can be seen as less expressive than [10], in the
5 sense that we have a single (uncertain) dynamics (as opposed to LHAs, where
6 a different dynamics can be defined in each mode), our dynamics is also sig-
7 nificantly *more expressive* than the LHA dynamics of [10]; we consider not
8 only the class of linear dynamical systems, but even fit into a special case of
9 non-linear systems, by allowing *uncertainty* in the model dynamics.

10 In [21, 22], a monitor is constructed from a system model in differential
11 dynamic logic [23]. The main difference between [21, 22] and our approach
12 relies in the system model: in [21, 22], the compliance between the model
13 and the behavior is checked at runtime, while our model is assumed to be an
14 over-approximation of the behavior—which is by assumption compliant with
15 the model.

16 8. Conclusion and Future Work

17 Monitoring black-box complex cyber-physical systems can be delicate,
18 and may lead to false alarms. MoULDyS is a Python-based tool implementing
19 offline and online monitoring algorithms. A first crux of MoULDyS is to be able
20 to manage logs with uncertainty over the logged state variables, as well as
21 missing samples. A second crux is the use of a bounding model in the form
22 of uncertain linear systems, helping to reduce the number of false alarms.
23 MoULDyS can analyze logs efficiently to detect possible safety violations that
24 might have caused an unsafe behavior. Further, MoULDyS can also be used
25 in an online setting where the system is sampled only when there is a risk
26 of safety violation. As a result, the online monitoring is able to decrease
27 the number of samples, therefore reducing energy consumption at runtime.
28 MoULDyS is available under the GNU General Public License.

29 In future, we wish to extend MoULDyS to support uncertainty not only
30 in the log valuations (the value of a sensor at a given timestamp), but also
31 uncertainty in the log timestamps themselves: this makes sense when some
32 sensors are distributed with drifting clocks, or when network delays make the
33 exact recording timestamp imprecise.

34 Acknowledgements (optional)

35 Bineet Ghosh was supported by the National Science Foundation (NSF)
36 of the United States of America under grant number 2038960. This work is
37 partially supported by the ANR-NRF French-Singaporean research program

1 ProMiS (ANR-19-CE25-0015 / 2019 ANR NRF 0092) and by ANR BisoUS
2 (ANR-22-CE48-0012).

3 References

- 4 [1] GNU General Public License v3.0, [https://www.gnu.org/licenses/
5 gpl-3.0.en.html](https://www.gnu.org/licenses/gpl-3.0.en.html).
- 6 [2] MoULDyS installation guide, [https://www.github.com/
7 bineet-coderep/MoULDyS/blob/main/documentation/
8 installation_guide.md](https://www.github.com/bineet-coderep/MoULDyS/blob/main/documentation/installation_guide.md) (2022).
- 9 [3] MoULDyS user guide, [https://www.github.com/bineet-coderep/
10 MoULDyS/blob/main/documentation/user_guide.pdf](https://www.github.com/bineet-coderep/MoULDyS/blob/main/documentation/user_guide.pdf) (2022).
- 11 [4] E. Bartocci, J. V. Deshmukh, A. Donzé, G. E. Fainekos, O. Maler,
12 D. Ničković, S. Sankaranarayanan, Specification-based monitoring of
13 cyber-physical systems: A survey on theory, tools and applications,
14 in: E. Bartocci, Y. Falcone (Eds.), Lectures on Runtime Verification
15 – Introductory and Advanced Topics, Vol. 10457 of Lecture Notes
16 in Computer Science, Springer, 2018, pp. 135–175. [doi:10.1007/
17 978-3-319-75632-5_5](https://doi.org/10.1007/978-3-319-75632-5_5).
- 18 [5] D. A. Basin, G. Caronni, S. Ereth, M. Harvan, F. Klaedtke, H. Mantel,
19 Scalable offline monitoring of temporal specifications, Formal
20 Methods in System Design 49 (1-2) (2016) 75–108. [doi:10.1007/
21 s10703-016-0242-y](https://doi.org/10.1007/s10703-016-0242-y).
- 22 [6] O. Maler, Some thoughts on runtime verification, in: Y. Falcone,
23 C. Sánchez (Eds.), RV, Vol. 10012 of Lecture Notes in Computer Science,
24 Springer, 2016, pp. 3–14. [doi:10.1007/978-3-319-46982-9_1](https://doi.org/10.1007/978-3-319-46982-9_1).
- 25 [7] R. Lal, P. Prabhakar, Bounded error flowpipe computation of parameterized
26 linear systems, in: A. Girault, N. Guan (Eds.), EMSOFT, IEEE,
27 2015, pp. 237–246. [doi:10.1109/EMSOFT.2015.7318279](https://doi.org/10.1109/EMSOFT.2015.7318279).
- 28 [8] B. Ghosh, P. S. Duggirala, Robust reachable set: Accounting for uncertainties
29 in linear dynamical systems, ACM Transactions on Embedded
30 Computing Systems 18 (5s) (2019) 97:1–97:22. [doi:10.1145/3358229](https://doi.org/10.1145/3358229).
- 31 [9] B. Ghosh, É. André, Monitoring of scattered uncertain logs using uncertain
32 linear dynamical systems, in: M. Mousavi, A. Philippou (Eds.),
33 FORTE, Vol. 13273 of Lecture Notes in Computer Science, Springer,
34 2022, pp. 67–87. [doi:10.1007/978-3-031-08679-3_5](https://doi.org/10.1007/978-3-031-08679-3_5).

- 1 [10] M. Waga, É. André, I. Hasuo, Model-bounded monitoring of hybrid
2 systems, *ACM Transactions on Cyber-Physical Systems* 6 (4) (2022)
3 30:1–30:26. doi:10.1145/3529095.
- 4 [11] B. Ghosh, P. S. Duggirala, Robustness of safety for linear dynamical
5 systems: Symbolic and numerical approaches, Tech. Rep. 2109.07632,
6 arXiv (2021).
- 7 [12] V. Gan, G. A. Dumont, I. Mitchell, Benchmark problem: A PK/PD
8 model and safety constraints for anesthesia delivery, in: G. Frehse,
9 M. Althoff (Eds.), ARCH@CPSWeek, Vol. 34 of EPiC Series in Com-
10 puting, EasyChair, 2014, pp. 1–8. doi:10.29007/8drm.
- 11 [13] P. Nilsson, O. Hussien, A. Balkan, Y. Chen, A. D. Ames, J. W. Grizzle,
12 N. Ozay, H. Peng, P. Tabuada, Correct-by-construction adaptive cruise
13 control: Two approaches, *IEEE Transactions on Control Systems Tech-*
14 *nology* 24 (4) (2016) 1294–1307. doi:10.1109/TCST.2015.2501351.
- 15 [14] MoULDyS: Recreating results, [https://www.github.com/
16 bineet-coderep/MoULDyS/blob/main/documentation/recreate_
17 results.md](https://www.github.com/bineet-coderep/MoULDyS/blob/main/documentation/recreate_results.md) (2022).
- 18 [15] D. A. Basin, F. Klaedtke, E. Zalinescu, The MonPoly monitoring tool,
19 in: G. Reger, K. Havelund (Eds.), RV-CuBES, Vol. 3 of Kalpa Publica-
20 tions in Computing, EasyChair, 2017, pp. 19–28.
- 21 [16] K. Mamouras, A. Chattopadhyay, Z. Wang, A compositional frame-
22 work for quantitative online monitoring over continuous-time signals,
23 in: L. Feng, D. Fisman (Eds.), RV, Vol. 12974 of Lecture Notes
24 in Computer Science, Springer, 2021, pp. 142–163. doi:10.1007/
25 978-3-030-88494-9_8.
- 26 [17] M. Waga, É. André, I. Hasuo, Parametric timed pattern matching, *ACM*
27 *Transactions on Software Engineering and Methodology* 32 (1) (2022)
28 10:1–10:35. doi:10.1145/3517194.
- 29 [18] É. André, IMITATOR 3: Synthesis of timing parameters beyond de-
30 cidability, in: R. Leino, A. Silva (Eds.), CAV, Vol. 12759 of Lecture
31 Notes in Computer Science, Springer, 2021, pp. 1–14. doi:10.1007/
32 978-3-030-81685-8_26.
- 33 [19] N. Halbwachs, Y.-É. Proy, P. Raymond, Verification of linear hybrid
34 systems by means of convex approximations, in: B. Le Charlier (Ed.),



Figure A.5: MoULDyS Logo.

- 1 SAS, Vol. 864 of Lecture Notes in Computer Science, Springer, 1994,
2 pp. 223–237. [doi:10.1007/3-540-58485-4_43](https://doi.org/10.1007/3-540-58485-4_43).
- 3 [20] A. Becchi, E. Zaffanella, Revisiting polyhedral analysis for hybrid
4 systems, in: B. E. Chang (Ed.), SAS, Vol. 11822 of Lecture Notes
5 in Computer Science, Springer, 2019, pp. 183–202. [doi:10.1007/
6 978-3-030-32304-2_10](https://doi.org/10.1007/978-3-030-32304-2_10).
- 7 [21] S. Mitsch, A. Platzer, ModelPlex: verified runtime validation of verified
8 cyber-physical system models, *Formal Methods in System Design* 49 (1-
9 2) (2016) 33–74. [doi:10.1007/s10703-016-0241-z](https://doi.org/10.1007/s10703-016-0241-z).
- 10 [22] S. Mitsch, A. Platzer, [Verified runtime validation for partially observable
11 hybrid systems](https://arxiv.org/abs/1811.06502), Tech. rep. (2018). [arXiv:1811.06502](https://arxiv.org/abs/1811.06502).
12 URL <http://arxiv.org/abs/1811.06502>
- 13 [23] A. Platzer, The complete proof theory of hybrid systems, in: LICS, IEEE
14 Computer Society, 2012, pp. 541–550. [doi:10.1109/LICS.2012.64](https://doi.org/10.1109/LICS.2012.64).

15 **Appendix A. MoULDyS: A Monitoring Tool for Autonomous Systems**

16 The various details of MoULDyS are given as follows:

17 **Logo** The tool logo is given in [Fig. A.5](#).

18 **Webpage** The tool webpage can be found here: [https://www.sites.
19 google.com/view/mouldys](https://www.sites.google.com/view/mouldys).

1 **Code** MoULDyS is an open-source tool under the `gpl-3.0` license. The
2 code can be found in a public GitHub repository: [https://github.com/
3 bineet-coderep/MoULDyS/releases/tag/v1.1](https://github.com/bineet-coderep/MoULDyS/releases/tag/v1.1).

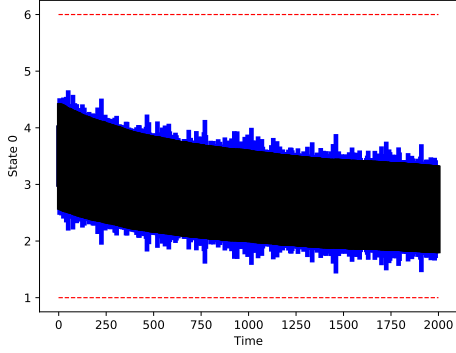
4 **Installation Guide** The installation guide is available in [2].

5 **User Guide** The user guide is available in [3].

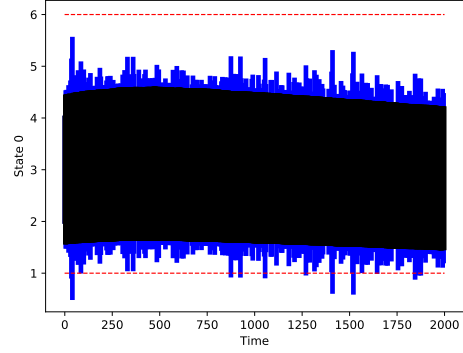
6 **Result Recreation Guide Guide** A prototype version of MoULDyS was
7 used to perform the experiments in [9]. The detailed steps to recre-
8 ate the results in [9], through easy-to-use scripts, are available in [14].

9 **Appendix B. Recreating Experimental Results**

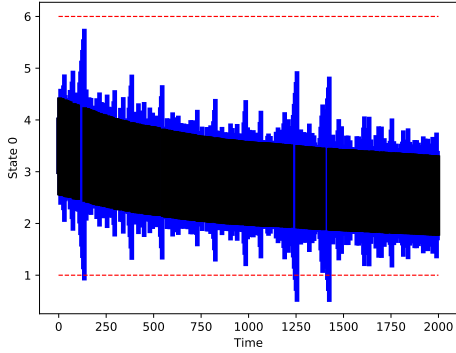
10 The results to be recreated for the Anesthesia case study are given in
11 [Figs. B.6](#) and [B.7](#). The results to be recreated for the ACC case study are
12 given in [Figs. B.8](#) and [B.9](#).




(a) Monitoring with frequent samples, and low uncertainty



(b) Monitoring with frequent samples, and high uncertainty




(c) Monitoring with sporadic samples, and low uncertainty




(d) Monitoring with sporadic samples, and high uncertainty

Figure B.6: *Offline Monitoring (Anesthesia)*. We plot the change in concentration level of c_p with time. The volume of the samples increases from left to right, and the probability of logging increases from bottom to top. The blue regions are the reachable sets showing the over-approximate reachable sets as computed by the offline monitoring, the black regions are the samples from the log given to the offline monitoring algorithm, and the red dotted line represents safe distance level. Note that although **Figure 1** and **Figure 4** (Figs. B.6b and B.6c) reachable sets' seem to intersect with the red line (unsafe set), the refinement module infers them to be *unreachable*, therefore concluding the system behavior as *safe*—unlike Fig. B.6d. These plots are a stochastic recreation of [9, Figure 3].




(a) **Online Monitoring**




(b) **Compare Online and Offline Monitoring**


Figure B.7: *Online Monitoring (Anesthesia)*. We plot the change in concentration level of c_p with time. The blue regions are the reachable sets showing the over-approximate reachable sets as computed by the online monitoring, the black regions are the samples generated when the logging system was triggered by the online monitoring algorithm, and the red dotted line represents safe concentration levels. *Online Monitoring (Fig. B.7a)*: We apply our online monitoring to the anesthesia model. *Compare (Fig. B.7b)*: We compare our online and offline algorithms. The green regions are the reachable sets showing the over-approximate reachable sets between two consecutive samples from the offline logs, the magenta regions are the offline logs, given as an input to the offline monitoring algorithm, generated by the logging system, and the red dotted line represents safe concentration levels. The blue regions are the reachable sets showing the over-approximate reachable sets as computed by the online monitoring, the black regions are the samples generated when the logging system was triggered by the online monitoring algorithm, and the red dotted line represents safe concentration levels. These plots are a stochastic recreation of [9, Figure 4].




(a) Monitoring with frequent samples, and low uncertainty



(b) Monitoring with frequent samples, and high uncertainty




(c) Monitoring with sporadic samples, and low uncertainty




(d) Monitoring with sporadic samples, and high uncertainty

Figure B.8: *Offline Monitoring (ACC)*. We plot the change in distance h between the vehicles with time. The volume of the samples increases from left to right, and the probability of logging increases from bottom to top. These plots are a stochastic recreation of [9, Figure 5]



(a) Online Monitoring



(b) Compare Online and Offline Monitoring

Figure B.9: *Online Monitoring (ACC)*. We plot the change in distance between two vehicle h with time. The color coding is same as Fig. B.7. *Online Monitoring (Fig. B.9a)*: We apply our online monitoring to the ACC model. *Compare (Fig. B.9b)*: We compare our online and offline algorithms. These plots are a stochastic recreation of [9, Figure 6]