Lifecycle-Based View on Cyber-Physical System Models Using Extended Hidden Markov Models

Matthias Schaffeld
Faculty of Engineering
University of Duisburg-Essen
Duisburg, Germany
0000-0002-5308-7010

Rebecca Bernemann
Faculty of Engineering
University of Duisburg-Essen
Duisburg, Germany
0000-0002-3240-0952

Torben Weis
Faculty of Engineering
University of Duisburg-Essen
Duisburg, Germany
0000-0001-6594-326X

Barbara König
Faculty of Engineering
University of Duisburg-Essen
Duisburg, Germany
0000-0002-4193-2889

Viktor Matkovic
Faculty of Engineering
University of Duisburg-Essen
Duisburg, Germany
0000-0002-6808-471X

***Abstract*—Many components of Cyber-Physical Systems (CPS) are designed based on models that represent the assumed behavior of the CPS at the time of deployment. However, significant or continuous small changes in the CPS, as well as wear and tear reduce the effectiveness of the CPS and its model and may lead to a total failure of the overall system. In this paper, we propose a novel lifecycle-based view of CPS models. First, we define the model's lifespan as the period from the initial conception of the model until it is no longer fit to represent the system behavior. For better differentiation, a lifespan is divided into the initial, operation, and adaptation phases. In the initial phase, a known-good baseline performance metric is established for the model's suitability to reflect the system behavior. In the operation phase, the model is used for CPS analysis, data smoothing, and fault location while its suitability is monitored. The adaptation phase is intended for necessary adaptations to the model and to the CPS itself, which lead to new iterations. To implement these lifecycle augmentations of the CPS, we use formal modeling in the form of Hidden Markov Models extended by unobservable transitions (ϵ-HMMT) to represent the assumed system behavior and compare the data of the observed system behavior with this modeling. In addition, we are testing our proposed formalism by designing a CPS model based on smart home systems and running a simulation for validation. The simulation covers unforeseen system changes and corrupted data.**

***Index Terms*—cyber-physical systems, model lifecycle, formal modeling, model validation**

I. Introduction

Cyber-Physical Systems (CPS) are the real-time digital monitoring, analysis and control of large and complex physical systems. They achieve this by combining networking, storage capacity, and the mutual integration of their digital and physical components. In previous work [1] we discussed the relationship between the physical components of a CPS and models that reflect their expected behavior and are used to calculate the state of the CPS. We argued that phenomena such as simplifications, misunderstandings, measurement errors, parameter drift, or the modification of the system's components cannot be eliminated, especially in the case of long-term considerations. In addition, we concluded the need for robust CPS models and abstractly introduced the formalism of ϵ Hidden Markov Models with transitional emissions (ϵ-HMMT) as a possible solution. However, as robust as a model formalism may be, the occurrence of environmental changes and subsequent system behavior adjustments mean that the original model captures the actual system behavior with decreasing accuracy over time. It can therefore be assumed that this process will ultimately require substantial adaptations to the CPS and its models. In this paper, we aim to concretize the necessary steps to establish the ϵ-HMMT as a suitable formalism for CPS modeling as well as to conceptualize the various phases of the lifespan of a CPS model. For each of these phases, we discuss how the model can be used as a complement to the CPS and how the ϵ-HMMT can be used to support this purpose, including a model's initial validation, its applications during deployment, and recommended adaptations to the model and/or the CPS.

This work has been partially funded by a BMBF grant for the project KI-LiveS.

II. Related Work

As Lee et al. [2] emphasize, CPS are about the intersection of the physical and the cyber, and as such we must understand their interaction. They argue that a comprehensive CPS model must represent both continuous and discrete dynamics and that details of the physical system are inevitably omitted. However, they also suggest that it is possible to use a purely discrete model for properties of interest while maintaining high accuracy. Derler et al. [3] introduce the concept of design contracts to close gaps between control and software engineering. Design contracts are agreements about system properties that specify requirements and assumptions for the system. While these can be used to ensure that digital and physical components are appropriate from a specified perspective, we cannot expect the requirements and assumptions to be comprehensive and flawless. Especially given the complexity

and inherent uncertainty of the underlying physical systems. This can also be seen in the work of Sanislav et al. [4], who named the detection of environmental changes and system behavior adaptation as the main challenges of CPS design. Putnik et al. [5] state the importance of behavior simulation and (predicted) result analysis but emphasize that due to the inherent and intrinsic complexity of emergent CPS scenarios, common modeling patterns such as model-based design are not adequate. They emphasize the importance of feedback loops between the CPS components and that there must be a continuous loop of learning and realignment through all phases of the CPS to adapt to the uncertainties.

III. System Model

We assume that the digital components of a CPS perceive the physical system as a sequence of discrete events and that there are strong sequential dependencies between observable events. Furthermore, we assume that a formalism exists that allows the assumed behavior of the CPS to be modeled and that serves as the basis for the design and operation of the digital components. As such a formalism we propose, concretize and evaluate an adapted form of Hidden Markov Models (HMM) that allows for unobservable transitions, also known as null transitions or ϵ-transitions. These enable HMMs to change state without the observer being aware of it and have been used in the literature, particularly in the context of speech recognition [6], [7]. Furthermore, in previous work [1] we argued that these ϵ-HMMTs seem to be suitable due to their properties for modeling the uncertainties of the underlying physical system as well as human behavior including sensor and measurement errors.

We extend the formalism of time-homogenous first order HMM [8] and give a compact formal definition of ϵ-HMMTs. We formally define an ϵ-HMMT as a 5-tuple $H = (Q, \Sigma, A, B, \pi)$, i.e.,

- $Q = \{s_1, ..., s_N\}$ is the finite set of system states,
- $\Sigma = \{e_1, ..., e_M\}$ is the alphabet of emission symbols,
- $A = (a_{s_i,s_j})_{i,j}$ is the $|Q| \times |Q|$ state transition probability matrix, where $a_{s_i,s_j} = P[s_j|s_i]$ is the probability of going from s_i to s_j,
- $B = (b_{s_i,s_j,e_m})_{i,j,m}$ is the $|Q| \times |Q| \times |(\Sigma \cup \{\epsilon\})|$ observation symbol probability matrix including $\epsilon = e_{M+1}$, where $b_{s_i,s_j,e_m} = P[e_m|s_i, s_j]$ is the probability of emitting e_m provided that one goes from s_i to s_j, and
- π is the $|Q|$-dimensional vector, which specifies the initial probability distribution, i.e., $\pi_k = P[s_k]$ is the probability of starting in state s_k.

Furthermore, A, B and π have to satisfy the following stochastic constraints:

- $a_{s_i,s_j} \geq 0$ and $\sum_{j=1}^{N} a_{s_i,s_j} = 1$, for $1 \leq i, j \leq N = |Q|$
- $b_{s_i,s_j,e_m} \geq 0$ and $\sum_{m=1}^{M+1} b_{s_i,s_j,e_m} = \begin{cases} 1 & \text{if } a_{s_i,s_j} > 0 \\ 0 & \text{otherwise} \end{cases}$, for $1 \leq i, j \leq |Q|$ and $1 \leq m \leq M + 1 = |(\Sigma \cup \{\epsilon\})|$
- $\pi_k \geq 0$ and $\sum_{k=1}^{N} \pi_k = 1$, for $1 \leq k \leq N = |Q|$

The added ϵ-emissions serve a similar purpose as the empty word of length zero in formal languages. In the context of HMM and HMMT, the non-observable ϵ-emission is added to each state's or transition's set of emissions. The ϵ-emissions of transitions between two different states then allow a change of hidden state without emitting a symbol, therefore accounting for the occurrence of missing hidden state transitions and analogously missing observations or false negatives. False positives introduce observations of events that did not occur and that have to be filtered out of the sequence of observations O. They can be accounted for by introducing loop transitions to each state, assuming that states cannot occur multiple times in succession. Therefore, we can account for incorrect observations, resulting from the occurrence of false negatives and/or false positives. Another problem is misclassification of event types, i.e., within observation sequences, the type of an observed event differs from the ground truth. However, by integrating the rate of misclassification of event types into relevant emission probabilities, we account for these as well.

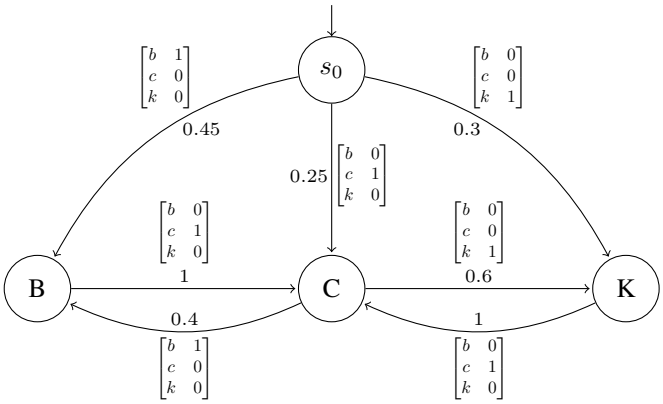

Fig. 1. Exemplary visualization of an HMMT without error consideration.

As an example, consider a bedroom (B) and a kitchen (K) connected by a corridor (C), each equipped with one occupancy sensor of the same name. Figure 1 depicts the modeling of such a system using the HMMT formalism, which does not consider false positives or false negatives. Assuming sensors B and C trigger ($O = bc$), this modeling gives the correct explanation that, states B and C were passed through. We now assume that only sensors B and K trigger, but not C ($O = bk$). The HMMT modeling is not able to explain this scenario. However, since we must pass through the corridor to get from B to K, we can conclude that either a false negative occurred and the correct sequence of states is B, C, K, the events b or c have been misclassified, or at least one of the observations is a false positive.

Figure 2 depicts similar modeling of this scenario using the ϵ-HMMT formalism, that considers false positives, false negatives, and misclassifications. Considering the same example ($O = bk$), there are now several possible explanations. The most probable of these is the state sequence s_0, B, C, K emitting b on the first, ϵ on the second and k on the third

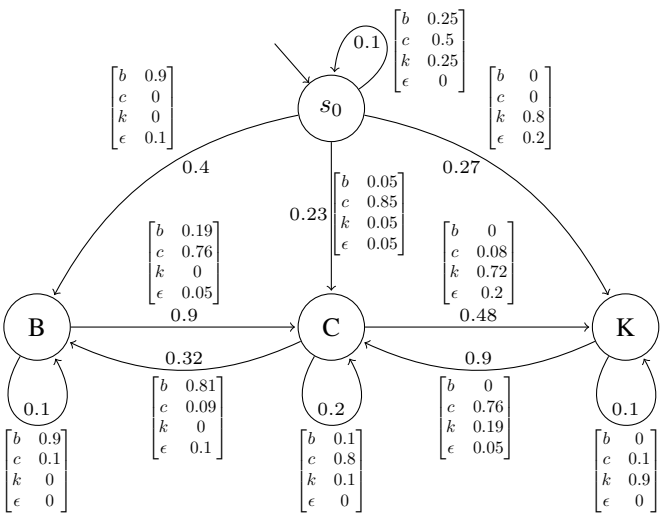

Fig. 2. Exemplary visualization of an ϵ-HMMT, considering false negatives, false positives and misclassifications.

transition. Its probability can be calculated as follows:

$$P(B|s_0)\cdot P(b|s_0,B)\cdot P(C|B)\cdot P(\epsilon|B,C)\cdot P(K|C)\cdot P(k|C,K) \\ = 0.4\cdot 0.9\cdot 0.9\cdot 0.05\cdot 0.48\cdot 0.72 \approx 0.0056$$

The most probable explanation is therefore a false negative for the observation c after b.

For standard HMM there exist two well-known algorithms, the Viterbi [8]–[10] and the Baum-Welch algorithm [11]. Given an observation sequence, the Viterbi algorithm computes the most likely sequence of states and its probability. Given a model structure and initial probabilities the Baum-Welch algorithm adjusts the probabilities based on observed output sequences. Jelinek [7] addressed the introduction of ϵ-transitions into these algorithms and we further proved the correctness of these adapted algorithms, in particular the introduction of ϵ-loops [12].

IV. Approach

To counteract the inherent deviations between a model and the actual system behavior over time, we introduce a lifecycle-based view on CPS models leveraged by the ϵ-HMMT formalism (see Section III). We denote this period from the first conception of the model until it is no longer suitable for representing the system behavior as the lifespan of the model and divide it into the following three phases:

1) Initial
2) Operation
3) Adaptation

The *initial phase* is about quantifying a model's suitability to reflect the actual system behavior to create a known-good baseline performance metric for later references. We consider the following three scenarios.

In the *adequate model* scenario, there is a model that a validator such as a domain expert considers appropriate to represent the actual behavior of the CPS, for example, as the result of a model-driven CPS design process. The model consequently contains all relevant system knowledge and assumptions so that the digital components can adequately carry out their task. However, the recorded CPS data may be corrupted or the system itself may behave unexpectedly.

In the *adequate data* scenario, it is assumed that the collected CPS data is considered adequate by a validator to represent the actual behavior of the CPS. This may, for example, occur if all hardware components of the CPS have recently been maintained. However, the model under consideration may not fully and correctly represent the system behavior. If we can assume the CPS data to be adequate, the modeling formalism should offer the possibility of learning a system model from this data.

In the third *incongruity* scenario, we cannot assume that the model or the data collected adequately represents the behavior of the CPS. This can be the case, for example, if the underlying model contains various assumptions about the system behavior and we simultaneously have to expect significant measurement errors.

All of these scenarios involve varying degrees of confidence in our ability to model correct CPS behavior. However, it is important to quantify the suitability between the CPS and its models so that deviations can be detected later. Therefore, we need a modeling formalism to enable the introduction of a metric that determines the suitability between a model and data of the observed system behavior. In addition, if either the model or the data are deemed adequate, the metric consequently becomes an estimate of the adequateness of the other. Using the ϵ-HMMT as the underlying model formalism then allows us to use the adapted Viterbi algorithm and metrics based on the edit distance to evaluate and establish a known-good baseline performance metric on how well data and model fit together. Furthermore, the adapted Baum-Welch algorithm can be used in the *adequate data* scenario to adjust model probabilities.

The *operation phase* covers the regular operation of the CPS in deployment. The software components perceive the state of the CPS through data that is interpreted based on the underlying model, calculate tasks and possible actions and optionally exercise control through actuators.

First, for this phase, we assume a suitable model formalism to provide robustness to address the inherent uncertainty of the underlying physical system. Using the ϵ-HMMT with its ϵ-emissions greatly enhances the robustness of the classic HMM formalism and is therefore a suitable model formalism for this phase. Second, we require that the formalism provides means to aid in the analysis and maintenance of the CPS. Possible tasks include but are not limited to *monitoring* the suitability assessment, *smoothing data* sequences or *locating faulty system behavior* and identifying their cause. Since CPS are subject to wear, parameter drift, and external influences, we assume that the suitability between the CPS and its models will decrease over time, despite the desired robustness of the model and the continuous maintenance of the CPS. To determine these discrepancies we use the same metric used for the known-good baseline performance in deployment and

monitor its deviation and gradient over time. Furthermore, as long as the model is deemed acceptable, we use the ϵ-HMMT and the adapted Viterbi algorithm for data smoothing purposes, providing further robustness.

The *adaptation phase* assumes that the suitability between the model and the actual system behavior is too poor for the CPS to function properly and for the model to support its operation phase tasks. To return to a point where the system is running safely, significant changes must be made to the model and/or the CPS itself. As major changes are made to the CPS and its models, we refer to the results as new iterations that need to be reevaluated, entering the initial phase again. The purpose of the model in this phase is to locate faulty model or system components and generate recommendations for resolving CPS and model flaws. Since it is difficult to distinguish between the model, the data collection process, or the components of the CPS as the origin of an error, a formalism can only aid to localize the error. This can be done by monitoring the extent to which specific emissions or system states contribute to the overall system error. Ultimately, a domain expert must clarify whether, for example, the problem lies within a model state or whether the respective sensors are causing significant measurement errors. We refer to the iterative sequence of these phases as the life cycle of a model. Using the ϵ-HMMT as the underlying model formalism, enables the adjustment of probabilities according to changed observed system behavior. This can be done by executing the Baum-Welch algorithm iteratively, until we converge to a local (not global) optimum. An open challenge here is how to estimate the probabilities of measurement errors as precisely as possible. It has to be seen if these can be accurately learned by using the Baum-Welch algorithm and by which data and model properties this is influenced. For now, we assume them to be specified by a domain expert or empirically generated.

V. Planned Evaluation

Before using the ϵ-HMMT as a formalism for CPS modeling we first want to establish its soundness. The adapted Viterbi algorithm gives us for an observed sequence (possibly with errors) the most probable sequence, which is compatible with a given ϵ-HMMT model. We quantify the difference between an observed and model-compliant behavior using the edit distance of the two sequences. In the ideal case, (1) the observed event sequences do not contain errors and (2) a model that perfectly represents the system behavior is available. Therefore, the edit distance is always zero, i.e. given any observation the most likely explanation is that this observation is correct. In reality, however, we must consider the possibility of errors in the observed sequence of events and/or the modeling of the CPS. In this case, the magnitude of the difference between observed and model-conform behavior is captured by the edit distance, resulting in non-zero values. To understand how well the ϵ-HMMT as a formalism for modeling CPS works and how it is affected by inducing errors and by different operation patterns, we are currently implementing the following two steps:

***Step 1**:* Generate a multitude of example HMMTs and based on these respective ground truth observation sequences O_α. This portrays the ideal case and should yield an edit distance of 0 in each case, as long as the error rates captured in the model do not interfere.

***Step 2**:* Inject false positives and negatives into the generated sequences O_α resulting in flawed observation sequences O_β and evaluate the edit distance:

Determine Error Base-Line (EBL): Determine the edit distance (ED) between the sequence with errors O_β and the one without the injected errors O_α, resulting in $EBL = ED(O_\alpha, O_\beta)$.

Measure Improvement: Determine the edit distance between the sequence without errors O_α and the sequence O_γ derived from the sequence with errors using the output of the Viterbi algorithm. We expect the result to be $EBL = ED(O_\alpha, O_\beta) > ED(O_\alpha, O_\gamma)$. So we see how well the correction using the Viterbi works.

Diversify and Evaluate Operation Patterns: While keeping the topology of the generated HMMTs we alter the parameter distributions according to different operation patterns and repeat the improvement measurements. By this we want to analyze factors that lead to a deterioration in a model's ability to correct well.

In addition, we plan to learn the probabilities (parameters) of a model using data and the extended Baum-Welch algorithm. Currently, the goal is to collect real-world data and establish the ground truth to properly evaluate the parametric learning of a model in the real world. We aim to achieve this by using known error-prone sensors (e.g. photosensitive sensors) to simulate environmental effects on sensor-based CPS. At this stage, we consider the main challenge to be the establishment of sound ground truth.

VI. Conclusion

The value of domain knowledge and assumptions about system behavior does not end after the components of the CPS have been designed. We propose to capture them through formal modeling and then compare the data of the observed system behavior with this modeling. We expect that such a model will allow various areas of application to complement the CPS throughout its lifespan and propose the ϵ-HMMT as a suitable formalism, while giving steps on how to establish its suitability in this context and how it can be used to support a CPS during its lifespan. This includes quantifying the suitability between a model and the observed system behavior, determining a threshold value for this suitability, locating model and system errors, and recommending model and system adjustments. We will evaluate the capabilities of this approach by testing the ϵ-HMMT formalism we developed with a smart home simulation that can introduce unexpected and changing system behavior as well as data corruption.

References

[1] M. Schaffeld and T. Weis, "Towards robust models of Cyber-Physical Systems," in *UbiComp '21*. New York, NY, USA: Association for Computing Machinery, 2021, p. 104–107. [Online]. Available: https://doi.org/10.1145/3460418.3479314

[2] E. A. Lee and S. A. Seshia, *Introduction to Embedded Systems: A Cyber-Physical Systems approach*, 2nd ed. The MIT Press, 2016.
[3] P. Derler, E. A. Lee, M. Törngren, and S. Tripakis, "Cyber-Physical System design contracts," in *2013 ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*, 2013, pp. 109–118.
[4] T. Sanislav and L. Miclea, "Cyber-physical Systems - concept, challenges and research areas," *Control Engineering and Applied Informatics*, vol. 14, pp. 28–33, 01 2012.
[5] G. Putnik, L. Ferreira, N. Lopes, and Z. Putnik, "What is a Cyber-Physical System: Definitions and models spectrum," *FME Transactions*, vol. 47, pp. 663–674, 01 2019.
[6] L. R. Bahl, F. Jelinek, and R. L. Mercer, "A maximum likelihood approach to continuous speech recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 5, no. 2, pp. 179–190, 1983.
[7] F. Jelinek, *Statistical methods for speech recognition*. MIT press, 1998.
[8] L. R. Rabiner, "A tutorial on Hidden Markov Models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, 1989.
[9] A. Viterbi, "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm," *IEEE Transactions on Information Theory*, vol. 13, no. 2, pp. 260–269, April 1967.
[10] G. F. Jr., "The Viterbi algorithm," *Proceedings of the IEEE*, vol. 61, no. 3, March 1973.
[11] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, "A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains," *The annals of mathematical statistics*, vol. 41, no. 1, pp. 164–171, 1970.
[12] R. Bernemann, B. König, M. Schaffeld, and T. Weis, "Probabilistic systems with hidden state and unobservable transitions," 2022, arXiv:2205.13871. [Online]. Available: https://arxiv.org/abs/2205.13871