Workflow Automation for Cyber Physical System Development Processes

Charles Hartsell Nagabhushan Mahadevan Harmon Nine Institute for Software Integrated Sys. Vanderbilt University charles.a.hartsell@vanderbilt.edu Ted Bapty Abhishek Dubey Gabor Karsai Institute for Software Integrated Sys. Vanderbilt University gabor.karsai@vanderbilt.edu

Abstract—Development of Cyber Physical Systems (CPSs) requires close interaction between developers with expertise in many domains to achieve ever-increasing demands for improved performance, reduced cost, and more system autonomy. Each engineering discipline commonly relies on domain-specific modeling languages, and analysis and execution of these models is often automated with appropriate tooling. However, integration between these heterogeneous models and tools is often lacking, and most of the burden for inter-operation of these tools is placed on system developers. To address this problem, we introduce a workflow modeling language for the automation of complex CPS development processes and implement a platform for execution of these models in the Assurance-based Learning-enabled CPS (ALC) Toolchain. Several illustrative examples are provided which show how these workflow models are able to automate many time-consuming integration tasks previously performed manually by system developers.

Index Terms—cyber physical systems; machine learning; model based design; workflow automation

ACRONYMS

- ALC Assurance-based Learning-enabled CPS
- ANN Artificial Neural Network
- CPS Cyber Physical System
- CNN Convolutional Neural Network
- DSMLDomain Specific Modeling Language
- DAG Directed Acyclic Graph
- LEC Learning Enabled Component
- ML Machine Learning
- UUV Unmanned Underwater Vehicle

I. INTRODUCTION

Modern Cyber Physical System (CPS) design is a complex process requiring multi-disciplinary expertise. Model and component based engineering techniques are proven approaches based on the traditional idea of "separation of concerns". Each engineering domain relies on their own particular modeling languages, often supported by multiple software tools. These tools provide automation for a wide range of common development tasks, and some existing toolchains provide integrated platforms with many such tools. These platforms provide numerous benefits including version control of system models, automatic tracking and maintenance of generated artifacts, and better cross-domain interaction among others. The Assurancebased Learning-enabled CPS (ALC) Toolchain [1] is one such platform which assists in the development of CPSs using data-driven development techniques such as machine learning. Any component which uses these techniques is known as a Learning Enabled Component (LEC).

Increasing system complexity and performance demands require close coordination between disciplines as well as more interaction between the supporting software tools. However, most existing platforms only provide automation for particular processes within their respective domains. Outside of these selected processes, the system developer is left with the burden of performing any required interfacing between tools such as data and model transformations, passing of artifacts from one tool to another, and configuration of tools before execution. To address this problem, we introduce a Domain Specific Modeling Language (DSML) for the purpose of modeling CPS development workflows. This language allows for specification of generic activity graphs where each activity may involve the use of one or more domain-specific tools. The workflow language has executable semantics, hence we implemented both the language and an appropriate execution engine within the ALC Toolchain. The engine, known as the workflow executor, can automate and monitor the execution of these processes, significantly reducing the burden placed on the developer.

The remainder of this paper is organized as follows. Section II presents related research about automation of individual tasks as well as modeling and execution of complete work-flows. Next, Section III provides a brief introduction to the ALC Toolchain, the implementation platform for our workflow modeling language. Section IV provides a description of general workflow automation, explains our workflow modeling language, and details how this language was made executable in the ALC Toolchain. This is followed by a set of illustrative examples in Section V. Finally, we identify areas for future work and give concluding remarks in Sections VI and VII respectively.

II. RELATED RESEARCH

Within each engineering discipline, common domainspecific tasks are often automated with appropriate tool support. For example, TensorFlow-eXtended (TFX) [2] provides an extensible platform for generic Machine Learning (ML) applications. The usual ML training process is divided into a 9step pipeline consisting of smaller individual actions including training data transformation, model learning, and deployment of trained models to production environments. Each step of this pipeline can be customized to fit specific needs before deployment and execution. TFX also automates some common ML processes such as hyper-parameter optimization during model training.

Some platforms take this a step further by automating tasks which have traditionally required multi-disciplinary expertise. The Functional Modeling Compiler (FMC) [3] is one such approach which uses a high-level model specifying the intended function of a CPS to generate multiple potential architectures which implement that functionality. These architecture models are composed of component models from multiple engineering domains (eg. electrical, mechanical, chemical) and can be executed in an appropriate simulator. This allows developers to quickly compare alternative system architectures and analyze how design decisions in one engineering domain will affect system-level performance.

CPS development platforms facilitate multi-disciplinary development and often integrate multiple tools with varying levels of automation. INTO-CPS [4] is one such platform built on top of five particular development tools for system modeling, discrete and continuous model analysis, simulation, and test automation. INTO-CPS is intended for use across many CPS domains and provides examples in several, including automotive, railways, agriculture and building automation. Additionally, INTO-CPS includes recommended development workflows and automates tracking of model provenance and requirements traceability. However, configuration and execution of each operation in the workflow is a manual process. Kulkarni et al. present another framework tailored for smart manufacturing applications [5] which provides tools for the development and refinement of manufacturing models. The authors combine both analytical and data-driven models to facilitate process analysis by domain experts. The framework also integrates and automates execution of several additional libraries including the OpenMDAO Python optimization library [6].

Business Process Model and Notation (BPMN) [7] is an existing modeling standard for representing generic business processes and serves as a general workflow description language. BPMN models typically contain natural language descriptions of each operation in a particular process. Execution engines exist for automating the operations in a BPMN model, but the generic nature of BPMN means very little info required for task execution is contained within the process model. Typically, the individual tools and operations required to complete a process model must be automated

independently, then provide a suitable interface to the BPMN execution engine. Similarly, Maier et al. propose a workflow automation system specifically for CPSs based on a new Project Worker approach [8]. This approach allows developers to package common, repeatable operations into snippets known as Engineering Automation Objects. Additionally, data flowing between workflow operations is identified with Virtual Addresses instead of physical addresses which allows data references to be independent of the underlying storage mediums and tool interfaces. The authors provide conceptual examples and identify potential benefits of this approach, but do not implement the proposed workflow automation system in order to validate these benefits. In [9], Mustafiz et al. introduce the Formalism Transformation Graph (FTG) framework for multiparadigm modeling. The FTG describes the various DSMLs available and how they are related to one another through model transformations which may be automated or manual. This approach also includes a Process Model (PM) which describes development processes as graph of interconnected model artifacts and activities. Artifacts are used as input to or output from activities, and activities may be one of the model transformations described in the FTG or other external operations.

III. ALC TOOLCHAIN

The ALC Toolchain supports multi-disciplinary CPS development by providing an environment with several domainspecific modeling languages and an integrated set of tools. The toolchain is specifically tailored for systems that utilize LECs and supports development workflows like the one shown in Figure 1. In order to promote reproducibility and maintain data provenance, all system models are stored in a versioncontrolled database and management of all relevant data and generated artifacts is automated. The toolchain is built on the WebGME infrastructure [10] which provides a webbased, collaborative modeling environment where changes are automatically and immediately propagated to all active users, not unlike in Google Docs ¹.

The typical development process in the ALC Toolchain, shown in Figure 1, consists of operations divided into two classes: tasks and activities. Tasks are individual actions that require domain-specific knowledge and typically involve the configuration of models or analysis of results. Tasks are shown as grey parallelograms in Figure 1 with examples including creation of a component library, design of an experiment, or evaluation of LEC performance. These actions may be performed manually by a system developer or automated with scripts where possible. Once a model has been configured, it can be deployed as an activity and executed using domainspecific tools. For example, an LEC Training model contains all necessary information to train an Artificial Neural Network using a machine learning library such as TensorFlow [11] or PyTorch [12]. Activities are shown as yellow, rounded rectangles in Figure 1 and include common operations such

¹https://docs.google.com



Fig. 1. ALC Toolchain development workflow.

as data generation, LEC training, and system verification. Additionally, *categories* are used to group related tasks and activities together. System modeling, LEC construction, and verification & assurance are the three primary categories in the ALC Toolchain, shown as blue blocks in Figure 1.

ALC provides automation for execution of individual activities through models which describe the desired operations and contain all the information necessary to initialize and deploy the activity. By default, support is included for several common CPS development tasks, particularly those related to ML operations. This automation is designed to be extensible so that new classes of models with different execution semantics can be added in a straightforward manner. Previously, all operations required between execution of activities (eg. passing generated artifacts between models, changing configuration parameters, and starting model execution) were performed manually. These tasks are time consuming to perform and require the user to regularly check if one model execution has completed such that the next execution can be started. The introduction of workflow models into the ALC Toolchain allows developers to automate these tasks and significantly reduce their workload, as discussed in the following sections.

In the remaining sections of this paper, certain modeling and documentation details have been omitted due to space constraints including the ALC and workflow DSML meta-models, certain details from the provided workflow examples, and the workflow API documentation. However, all of these resources are publicly available on the ALC Toolchain website².

IV. WORKFLOW AUTOMATION

Merriam-Webster defines a workflow as "the sequence of steps involved in moving from the beginning to the end of a working process" [13]. In the context of CPS development, workflow steps often include complex tasks such as data generation and analysis, performance evaluation, and systemlevel verification among many others. Typically, operations in a workflow are dependent on results from one or more of the preceding steps. Operations which share the same set of dependencies may be performed in parallel. Additionally, many workflows are iterative processes which are repeated until a predefined threshold is met. To capture these dependencies between operations, workflows are often represented as directed graphs.

Existing CPS development environments usually provide some level of automation for the deployment and execution of individual domain-specific tools, and may enforce certain tool input and output rules to improve the process of interfacing between tools. However, CPS development workflows usually rely on several interconnected tools with various data transformations between tools. This required interfacing between tools is often automated for a few selected workflows with an appropriate toolchain, or automated on a case-by-case basis in an ad-hoc manner with a general purpose scripting

²https://cps-vo.org/group/ALC



Fig. 2. Components of the Workflow modeling language.

language. While these scripting languages are powerful and generic, creating process automation scripts with them is a time-consuming process which often results in scripts that are fragile in use and not easy to generalize. Additionally, objects and activities within these scripts are not linked to system models or artifacts, and the process of manually synchronizing references as development progresses quickly becomes a significant burden.

To address this problem, we introduce a new DSML tailored for the creation and execution of CPS development workflows and provide one implementation of this language in the ALC Toolchain. Within ALC, individual tools are integrated into the toolchain as *activities*. An integrated activity must be able to interpret artifacts used as inputs, often results produced from other activities, and must produce output artifacts which conform to a standard set of rules. The workflow DSML allows developers to quickly chain multiple activities together to model their typical development processes in a structured manner and deploy these processes for execution. This approach is intended to bring the benefits of model-based engineering to CPS workflow automation and is explained in more detail in the following sections.

A. Modeling Language

The Workflow DSML within ALC consists of the eight primary modeling objects shown in Figure 2. Construction of a workflow model starts with the creation of a *job* block, which describes one step in the workflow process and contains exactly one activities block. Within ALC, any executable model can be considered an activity, and execution typically involves the use of domain-specific tools which operate on the input model and produce a result. The activities block within a workflow job specifies which models should be executed during this job, and contains links to one or more existing activity models. Each individual activity model usually needs to be initialized before execution which can be done with an appropriately configured *initializer* block. Activity initialization is often dependent on the results of previous jobs in the workflow, which can be accessed through one or more input ports. The results of each activity contained

within a job can be made available to later jobs through *output ports*. Finally, *iteration* blocks can be used to guide the execution flow of the model by specifying where and how the flow should be interrupted to perform additional iterations. These blocks contain a script which examines any of the outputs from previous jobs, determines if additional iterations of the workflow should be performed, and updates the activity models as needed for the next iteration.

A complete workflow model, like the examples shown in Figures 4 and 6, consists of one or more jobs connected in a directed graph where each edge is one of two types. Data flow edges are depicted as black lines and represent the flow of generated artifacts from one workflow job into another job. These edges are always directed from an output port to an input port. Process flow edges direct the execution flow of the workflow and may be one of two types. Check edges, shown as blue lines, indicate where the normal flow of the workflow should be interrupted to check if any iteration is needed. If the script contained within the *iteration* block, shown as a grey diamond, indicates that iteration is desired then repeat edges, shown as red lines, indicate where the next iteration should begin. This structured, model-based approach to workflow creation allows developers to quickly automate their usual development processes in a generic, reusable, and extensible manner. The example model in Figure 6 corresponds to one specific instance of the general LEC Construction phase of the ALC development workflow shown in Figure 1 and is discussed in more detail in Section V-C. A meta-model for the workflow DSML described above was created in the ALC Toolchain using the tools provided by the WebGME environment. Due to space constraints, the meta-model has been omitted here, but is available to explore on the ALC Toolchain website. Accessing this model requires registering a user account, clicking the large "LAUNCH" button, then selecting the "ALC_Meta" project.

B. Implementation

We have created a *Workflow Executor* within the ALC Toolchain which handles deployment, execution, and data management for workflow models. Scripts contained in *iteration* blocks which guide looping within the workflow are specified in the Python³ programming language and are invoked on each iteration of their corresponding loop. The user-specified behavioral code within these scripts can be arbitrarily complex. An API is provided for interacting with models in the ALC Toolchain including ability to:

- Access all data produced by previously executed jobs in the workflow.
- Persist selected data across loop iterations.
- Update any model parameters for future activities in the workflow.

The workflow executor is implemented on top of the $Gra-dle^4$ build automation system. Gradle is intended for software

³https://www.python.org/

⁴https://gradle.org/

Execution	Job	Activity	Status
RL_WF			Pending
	RL Exploration		Finished
		RLTrainingLEC1	Finished
	LEC Evaluation		Pending
		Dep_AM_Tracking	Pending

Fig. 3. Workflow status visualization.

build tasks such as compilation, packaging, and deployment. However, it provides a Kotlin⁵ based build script language and is easily extensible for other more specific purposes. In particular, Gradle provides well-defined, easy to use interfaces which are can be integrated with other platforms including the WebGME environment that the ALC Toolchain is built upon. When a workflow model is executed, the executor first extracts all relevant execution information from the model and stores this information in a JavaScript Object Notation (JSON) file. This JSON file is then interpreted by the Gradle build script, which generates the necessary task graph to carry out the workflow dynamically based on the constraints and dependencies specified JSON. The Gradle engine is well-suited to executing the workflow as it provides built-in functionality for:

- dependencies to be expressed between the tasks, ensuring that they are executed in the proper order
- detection, reporting of, and termination of the workflow due to any errors during execution
- determining if a task should be skipped, e.g. in the case that a previous execution of the workflow already executed the task, and its output need not be regenerated.

Gradle is only capable of generating a Directed Acyclic Graph (DAG) of tasks, i.e. no loops can be in the task dependency graph. Therefore, to implement looping in the Workflow DSML, a separate process is executed that contains a simple loop. Each iteration of the loop executes a Gradle process to build and execute a workflow task DAG for that specific iteration. The loop carries information forward from the previous iteration to the next so the appropriate DAG for the workflow is executed on each iteration.

Once a workflow execution begins, the executor tracks the progress of each job in the workflow and provides this information back to the user with a status table like the one shown in Figure 3. If a workflow job encounters an error during execution, the status table will be updated to reflect this. Once an error has been identified and fixed, then the workflow can be resumed from the point where the failure occurred. This way, none of the preceding jobs which finished successfully before the error will be re-executed.

V. EXAMPLES

The following sections present several examples showing how common workflows used in ALC can be automated with appropriate workflow models. All examples are in context of an Unmanned Underwater Vehicle (UUV) tasked with following a pipeline on the seafloor using two LECs along with several non-LEC components. Both LECs are neural network based, but use different training algorithms provided by different machine learning libraries. The Perception LEC performs image segmentation on input sonar images and produces a segmented image where every pixel in the original image has been assigned to an object class (eg. pipeline, seafloor, obstacle). This is accomplished with a Convolutional Neural Network (CNN) based on the SegNet [14] architecture and trained with a supervised learning approach using the PyTorch [12] library. The Control LEC must then use this information, as well as additional vehicle state information, to output appropriate control commands for the UUV. The control LEC is a two-layer, feed-forward network trained using reinforcement learning techniques provided by the Tensorforce [15] library. For a detailed introduction to supervised and reinforcement learning techniques see [16] and [17] respectively. Additionally, each of the example models shown in the following sections can be explored in full detail on the ALC Toolchain website.

A. Reinforcement Learning & Evaluation

The first example workflow involves training the control LEC over many simulated, randomly-generated scenarios known as episodes. During each episode, the LEC explores the possible action space while using the provided reward function to observe which control actions produce desirable results for the UUV system. Network weights are updated accordingly, and this process continues for a set number of episodes. Once the training process is completed, the LEC can be deployed back to the simulator in a non-learning mode (ie. network weights are no longer updated) for system-level evaluation. During this job, the control LEC is evaluated based on system-level performance metrics of the UUV instead of the component-level metric provided by the reward function. Evaluation jobs are often configured as a *campaign* where the same evaluation procedure is repeated in a variety of scenarios which can be provided as a specific set or generated with an appropriate scenario description language such as Scenic [18].

This training and evaluation process can easily be modeled with our workflow modeling language as shown in the top of Figure 4. The training process is captured in the "RL Exploration" block which provides a trained LEC output. This output is fed into the "LEC Evaluation" block and used to initialize the evaluation models. The result of each evaluation activity is also made available, but is not used in this workflow. While the ALC Toolchain provides automation of each individual activity, previously no automation between activities was available and all data passing and configuration between models required manual configuration. This particular workflow model was created by an experienced ALC user in





Fig. 4. Workflow models for example 1 (top) and 2 (bottom). The "HyperParameter Search" block in example 2 contains a script, shown with a dotted outline in the figure.

about 15 minutes, but may take hours or days to execute depending on the number of learning episodes, evaluation scenarios, and available computational resources. Additionally, the status visualizer for this example is shown in Figure 3. In this case, the "RL Exploration" job had completed successfully, indicated by the green "Finished" status, and the "LEC Evaluation" job was pending deployment to an execution server.

B. Hyperparameter Optimization

The next example involves automating hyperparameter optimization for the supervised training of the perception LEC, a common process in the machine learning field. The goal of this process is to find a set of training configuration parameters, known as *hyperparameters*, which results in an LEC with the minimal total loss when evaluated against a particular set of test data. While many common loss functions are available, we use a standard mean-squared-error loss function in this example. Common optimization search algorithms include grid search, random search [19], and Bayesian optimization [20]. For the purpose of this example we automate a simple grid search over two particular hyperparameters, gradient-descent optimizer and number of epochs, but more advanced search algorithms can be automated in a similar fashion.

Modeling the hyperparameter optimization process requires introducing the *iteration* concept as shown in the bottom left of Figure 4. In this workflow, there is only one job named "SL Model Training" which initializes the training activity model with labeled data sets before executing the training process for the perception LEC. The trained LEC is produced as an output, but is unused in this workflow. This job is then connected to a "Hyperparameter Search" iteration block, shown in an exploded view, containing the script which implements the iteration logic. In this case, each iteration will repeat the training process with one of three learning optimizers: Stochastic Gradient Descent (SGD) [21], Adam [22], or RMSprop [23]. Additionally, each optimizer will be run for 2, 4, 6, and 8 training epochs to evaluate how quickly each optimizer converges to the minimal loss. This results in a 3 by 4 search grid for a total of 12 training iterations, and the final loss value of each combination of hyperparameters is shown in Figure 5. In this case, both the RMSprop and Adam optimizers converge much faster than the standard SGD technique, but the RMSprop optimizer has a large jump in error rate when trained for 4 epochs. This could be due to other training parameters which were not considered in this test, such as a learning rate that is too large. Increasing the number of learning parameters used in our workflow grid search could help confirm this hypothesis, but was out of scope for this paper. Once all 12 iterations are completed, the workflow will terminate automatically. Similarly to the previous model, this workflow model required less than an hour of developer time to construct, but may require many computational hours to execute all 12 LEC training cycles. Without the workflow DSML and executor, the developer would have to regularly monitor for completion of each training execution, then manually update the training activity model with the next set of network hyperparameters and launch the next iteration of the grid search. Manual monitoring for job completion is tedious and typically results in available computing hardware sitting idle until the developer notices job completion. Additionally,

Learning	Number of Epochs				
Optimizer	2	4	6	8	
SGD	1.057	0.996	0.951	0.904	
Adam	0.057	0.032	0.022	0.011	
RMSProp	0.147	2.485	0.010	0.010	

Fig. 5. Perception LEC loss after training with various hyperparameters. Large loss values are colored red while small loss values are colored green.

manually updating model parameters and references for this relatively simple example is a quick, straight-forward process, but quickly becomes difficult as the complexity of the workflow increases.

C. ALC Development Workflow

The last example demonstrates automation for the full "LEC Construction" portion of the ALC development workflow shown in Figure 1. This process consists of generating labeled training data through simulation of the UUV system in a variety of scenarios. For this system, variations between scenarios include different configurations of the pipeline on the seafloor, randomized starting location of the UUV, and randomly generated obstacles in the path of the vehicle among others. The collected data can then be passed to a supervised learning model in order to train the perception LEC. At this point, the developer may chose to perform a hyperparameter search to optimize the performance of the LEC against an evaluation data set. Once an optimal training configuration has been identified, both the training data set and the trained LEC are sent to an assurance monitor training job. Assurance monitors are software components which act in parallel to an LEC and provide a confidence metric for each output produced by that LEC. This metric is based on the similarity between a new, previously unseen input and all inputs contained in the data set used to train the particular LEC. This technique is known as Inductive Conformal Prediction [24], and there are a number of algorithms for efficiently computing confidence metrics. In this example, we use one particular algorithm [25] based on Support Vector Data Description [26]. Finally, an evaluation model can be configured to use the trained perception LEC and corresponding assurance monitor as part of the complete UUV system. If the LEC shows unsatisfactory performance during a particular evaluation scenario, then the developer may chose to add this scenario to the original data generation step and repeat the process.

Figure 6 shows how this full LEC construction workflow can be modeled. As with the previous two examples, each type of activity is contained in a workflow job block with appropriate input and output ports to control the flow of generated artifacts. The last job in the workflow, LEC Evaluation, is depicted in an exploded view showing how individual activity models can be selected for this job. In this case, only a single evaluation model named "Dep_AM_Tracking" has been selected to run. However, this model is configured as a campaign so that several simulated scenarios will be run to evaluate the trained LEC. Additionally, this workflow model contains two loops: an inner loop for hyperparameter optimization and an outer loop for iterating the entire process with additional training scenarios. The inner loop follows the same procedure outlined in Section V-B to optimize the LEC training procedure. The outer loop contains a script which identifies any evaluation scenarios which resulted in unsatisfactory performance and updates the first workflow job, "Data Collection", to include these specific scenarios in the training data set. Also note that the assurance monitor training job is dependent upon the trained LEC produced by the hyperparameter search loop. Because of this dependency, the assurance monitor job will not run until after the entire hyperparameter search has completed. Since the hyperparameter search is repeated in full on every iteration of the outer loop, this workflow typically executes many computationally intensive LEC training activities. As with the previous examples, this workflow model was constructed and deployed in less than an hour of developer time, but required several computational hours to complete.

VI. FUTURE WORK

We have identified several opportunities for future improvements in both the workflow language itself as well as the implementation of the language within the ALC Toolchain. First, the workflow language does not allow general branching operations, which limits their expressiveness. While looping is supported, this can be viewed as a special case of more generic branching. Additionally, the language does not allow for hierarchical construction of workflows. That is, one workflow model cannot contain another workflow model. Hierarchical models can be useful in many scenarios and can promote model reuse. For instance, the hyperparameter optimization loop in the final example presented in Section V-C could have been replaced by the workflow model shown in the bottom of Figure 4. Due to this current limitation, the optimization loop model had to be recreated for both examples. We intend to add both general branching and hierarchical construction features to the workflow DSML.

Within the ALC Toolchain, some support is provided for the debugging of activities including the ability to quickly view all execution logs and perform interactive executions with Jupyter notebooks⁶. However, debugging capabilities for the workflow executor are limited and require understanding of the executor internals. In particular, there is no straightforward way of debugging the user-provided *iteration* scripts which guide the execution of the workflow. We plan to add additional debugging capabilities to allow typical developers to troubleshoot models without detailed understanding of the executor internals. Next, our current implementation of the Workflow Executor does not support any parallel execution of job iterations within loops. For example, each training session of the hyperparameter optimization described in Section V-B

⁶https://jupyter.org/



Fig. 6. Workflow model for example 3. The "LEC Evaluation" workflow job includes an exploded view of the internal model.

could be executed in parallel since the execution of one iteration is not dependent on the results from the previous iteration. However, this type of parallelization would not apply to loops with dependencies between each iteration such as the outer loop of the development cycle described in Section V-C. Additionally, a primary goal of the ALC Toolchain is to assist developers in providing system-level assurance for CPSs. With this in mind, we intend to integrate our workflow models more tightly with existing tools for construction and evaluation of assurance cases. We expect this would reduce both the burden of manually maintaining these assurance cases and the time required to identify system-level properties after design iterations. Finally, the current workflow executor is able to catch error conditions and report them back to the user. However, these errors must be corrected manually as the executor does not have any fault mitigation or recovery capabilities. It is often the case that a developer will not be able to address such issues until hours or days after the error occurred, resulting in wasted computer time. We plan to explore this functionality in future work to alleviate this problem.

VII. CONCLUSION

Modern CPS design requires cross-discipline expertise and involves many heterogeneous system models. These models often include domain-specific tooling for various tasks such as performance analysis, code generation, and system deployment among others. Automation of these tools has typically been limited to execution of individual tasks or to a limited set of common discipline processes. To address this problem, we have introduced a DSML for describing CPS workflows and provided an executive engine for automated execution of these workflows. This, along with the existing automation capabilities provided by the ALC toolchain, helps to bring the full benefits of model-based engineering to workflow automation.

ACKNOWLEDGEMENT

This work was supported by DARPA and the Air Force Research Laboratory. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA or AFRL.

REFERENCES

- [1] C. Hartsell, N. Mahadevan, S. Ramakrishna, A. Dubey, T. Bapty, T. Johnson, X. Koutsoukos, J. Sztipanovits, and G. Karsai, "Model-based design for CPS with Learning-Enabled Components," in *Proceedings of the Workshop on Design Automation for CPS and IoT*. ACM, 2019, pp. 1–9.
- [2] D. Baylor, E. Breck, H.-T. Cheng, N. Fiedel, C. Y. Foo, Z. Haque, S. Haykal, M. Ispir, V. Jain, L. Koc *et al.*, "Tfx: A tensorflow-based production-scale machine learning platform," in *Proceedings of the 23rd* ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, 2017, pp. 1387–1395.
- [3] A. Canedo and J. H. Richter, "Architectural design space exploration of cyber-physical systems using the functional modeling compiler," *Procedia CIRP*, vol. 21, pp. 46–51, 2014.
- [4] P. Larsen, C. Thule, K. Lausdahl, V. Bandur, C. Gamble, E. Brosse, A. Sadovykh, A. Bagnato, and L. Couto, "Integrated tool chain for model-based design of cyber-physical systems," in *The 14th Overture Workshop: Towards Analytical Tool Chains*, ser. Technical report ECE - TR - 28, P. Larsen, N. Plat, and N. Battle, Eds., vol. 4/28. Aarhus University, Department of Engineering, 11 2016, pp. 63–79.
- [5] A. Kulkarni, D. Balasubramanian, G. Karsai, and P. Denno, "An analytical framework for smart manufacturing," in *MATEC Web of Conferences*, vol. 249. EDP Sciences, 2018, p. 03010.
- [6] J. S. Gray, J. T. Hwang, J. R. R. A. Martins, K. T. Moore, and B. A. Naylor, "OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization," *Structural and Multidisciplinary Optimization*, vol. 59, no. 4, pp. 1075–1104, April 2019.
- [7] OMG, "Business Process Model and Notation (BPMN), Version 2.0," https://www.omg.org/spec/BPMN/2.0/PDF, Object Management Group, 2011.
- [8] R. Maier, S. Unverdorben, and M. Gepp, "Efficient implementation of task automation to support multidisciplinary engineering of cps," in 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE). IEEE, 2018, pp. 1388–1393.
- [9] S. Mustafiz, J. Denil, L. Lúcio, and H. Vangheluwe, "The ftg+ pm framework for multi-paradigm modelling: An automotive case study," in *Proceedings of the 6th International Workshop on Multi-Paradigm Modeling*, 2012, pp. 13–18.
- [10] M. Maróti, T. Kecskés, R. Kereskényi, B. Broll, P. Völgyesi, L. Jurácz, T. Levendovszky, and Á. Lédeczi, "Next generation (meta) modeling: Web-and cloud-based collaborative tool infrastructure." *MPM@ MoD-ELS*, vol. 1237, pp. 41–60, 2014.

- [11] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/
- [12] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Systems 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: http://papers.neurips.cc/paper/ 9015-pytorch-an-imperative-style-high-performance-deep-learning-library. pdf
- [13] M.-W. Inc., "Workflow," https://www.merriam-webster.com/dictionary/ workflow, accessed Jan 2020.
- [14] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *CoRR*, vol. abs/1511.00561, 2015. [Online]. Available: http://arxiv.org/ abs/1511.00561
- [15] A. Kuhnle, M. Schaarschmidt, and K. Fricke, "Tensorforce: a tensorflow library for applied reinforcement learning," Web page, 2017. [Online]. Available: https://github.com/tensorforce/tensorforce
- [16] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press,

2016, http://www.deeplearningbook.org.

- [17] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [18] D. Fremont, X. Yue, T. Dreossi, S. Ghosh, A. L. Sangiovanni-Vincentelli, and S. A. Seshia, "Scenic: Language-based scene generation," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2018-8, Apr 2018. [Online]. Available: http://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-8.html
- [19] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *Journal of machine learning research*, vol. 13, no. Feb, pp. 281–305, 2012.
- [20] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in neural information* processing systems, 2012, pp. 2951–2959.
- [21] H. Robbins and S. Monro, "A stochastic approximation method," *The annals of mathematical statistics*, pp. 400–407, 1951.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [23] G. Hinton, N. Srivastava, and K. Swersky, "Lecture 6e rmsprop: Divide the gradient by a running-average of its recent magnitude," http://www.cs.toronto.edu/ tijmen/csc321/slides/lecture_slides_lec6.pdf, accessed Jan 2020.
- [24] H. Papadopoulos, "Inductive conformal prediction: Theory and application to neural networks," in *Tools in artificial intelligence*. InTech, 2008.
- [25] F. Cai and X. Koutsoukos, "Real-time Out-of-distribution Detection in Learning-Enabled Cyber-Physical Systems," in 11th IEEE/ACM Conference on Cyber-Physical Systems (ICCPS 2020), 2020.
- [26] D. M. Tax and R. P. Duin, "Support vector data description," *Machine learning*, vol. 54, no. 1, pp. 45–66, 2004.