Orchestrating Tool Chains for Model-based Systems Engineering with RCE

Jan Flink, Robert Mischke, Kathrin Schaffert, Dominik Schneider, and Alexander Weinert[⋆]

Institute for Software Technology
German Aerospace Center (DLR)
51147 Cologne, Germany
`firstname.lastname@dlr.de`

Abstract. When using multiple software tools to analyze, visualize, or optimize models in MBSE, it is often tedious and error-prone to manually coordinate the execution of these tools and to retain their respective input and output data for later analysis. Since such tools often require expertise in their usage as well as diverse run-time environments, it is not straightforward to orchestrate their execution via off-the-shelf software tools. We present RCE, an application developed at the German Aerospace Center that supports engineers in developing and orchestrating the execution of complex tool chains. This application is used in numerous research and development projects in diverse domains and enables and simplifies the creation, analysis, and optimization of models.

1 Introduction

In classical engineering, work processes are organized around textual documents. Such documents contain, e.g., the specification of requirements, preliminary designs of the system under construction, or test reports of prototypes. These documents are written in natural language and are prone to accidental incompleteness, ambiguity, or conflicting definitions.

To alleviate these issues, MBSE (*Model-based Systems Engineering*) is applied in more and more engineering projects. In contrast to classical engineering principles, the processes of MBSE focus on the production and investigation of formal models instead of documents [23].

There exist numerous processes, techniques [8,9], and tools [19] that engineers can use to construct such a model in a reliable, timely, and precise way. Once an initial version of the model is produced, engineers may use the model to evaluate the behavior of the system in specific scenarios. During the development process, the model will go through multiple iterations as the understanding of the system improves and the fidelity of the modeling increases.

The specification language used for describing the model depends strongly on the domain of engineering. One example of such a language is SysML (*Systems Modeling Language*) [15], which describes a system via a set of diagrams. Each diagram provides a different view on the system, e.g., on the structure of its components or on the communication between them in a given usage scenario. There also exist more application-specific modeling languages that allow for greater ease of modeling. One example in the domain of aircraft design is CPACS (*Common Parametric Aircraft Configuration Schema*) [2].

Independently of the modeling language used, most activities of MBSE involve software tools that "understand" the constructed model and perform some operation with it. These operations may involve, e.g., the extension of the model with some characteristics, the evaluation of the model's performance in some scenario, or its optimization with respect to some property.

We present RCE (`https://rcenvironment.de/`), a freely available open-source application for designing, executing, and orchestrating automated distributed tool chains. RCE is being developed at DLR (German Aerospace Center) at the Institute for Software Technology. It satisfies multiple requirements for executing

---

[⋆] Authors are listed in alphabetical order

fully automated tool chains. Among others, RCE allows users to integrate almost arbitrary tools for simulation, analysis, and optimization into tool chains, to publish these integrated tools in a network, and to quickly adapt the design of tool chains to new requirements via a graphical interface.

The remainder of this paper is structured as follows: After defining nomenclature in Section 2, we give an overview over the main features of RCE in Section 3 and describe the design decisions and trade-offs made in the practical implementations of these features in Section 4. We subsequently highlight research and development projects in which RCE has been successfully applied in Section 5 before discussing related work and future directions for development in Section 6 and Section 7, respectively.

2   Nomenclature and Preliminaries

In this section we introduce concepts underlying our work as well as the nomenclature that we use in the remainder of this work.

The major novelty of MBSE in contrast to previous engineering processes is its focus on a common model of the system as the main artifact around which engineering processes are centered. This model is defined in a formal language that is used by all process participants. The availability of such a formal model allows for a greater support by *(software) tools* than for non-model-based engineering processes. These tools support engineers in, e.g., analyzing, optimizing, or visualizing the model. [1, 21]

Parts of each of these activities may be implemented comprehensively in a single tool, such as the visualization of the complete model under construction. In most cases, however, several tools are required to perform an activity, using the output of one tool as the input for one or more succeeding ones. We say that the individual software tools form a *tool chain* in this case.

Consider, e.g., the evaluation of an airplane model in a given scenario. First, the behavior of the model in the scenario needs to be simulated. Afterwards, the collected data needs to be evaluated with respect to at least the performance of the airplane as well as its economic and ecological characteristics. To simplify the comparison of different scenarios, these three evaluations must then be consolidated into a single evaluation. Each step of this process may in practice be implemented in a distinct software tool. We illustrate the individual steps and the relation between them in Figure 1a.

(a) A workflow for evaluating the performance of an airplane in a particular scenario.

(b) A workflow for determining an airplane that behave optimally in a range of scenarios. The gray boxes denote copies of the workflow illustrated in Figure 1a.
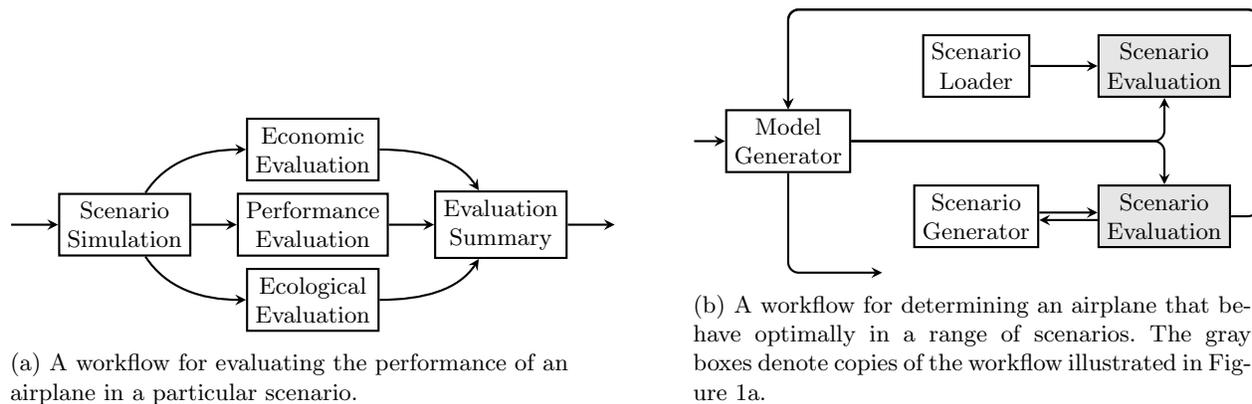
Fig. 1: A workflow for determining an optimal airplane.

While simple tool chains may suffice for many engineering tasks, they may become arbitrarily complex and require iterated or nested execution. Consider, e.g., the task of evaluating the behavior of the same airplane shape with different engines in diverse scenarios. These scenarios may either be based on historical data, or they may be generated according to anticipated operational scenarios. Given an initial airplane shape, for each candidate engine first a model will be constructed that combines the model of the airplane shape with a

model of the engine. The behavior of this combined model can then be simulated in multiple scenarios, where each simulation results in a comprehensive evaluation of the model's performance. By iterating over each candidate engine, the performance of all engines can subsequently be compared. We illustrate the relation between the individual steps in such an evaluation in Figure 1b.

Tools in a workflow may be distributed among individual computers on a network. In a typical engineering process, each tool in a workflow has an operator who is familiar with its operation and responsible for providing the service. To execute a workflow, the operators of the initial tool(s) first execute these tool(s) before passing the results of this execution on to the operators of subsequent tools. Passing of information among operators often happens via, e.g., shared network drives or via email. The subsequent operators then execute their respective tools and again pass their outputs to subsequent operators. The operators may also pre- or post-process the data they receive and produce using additional "helper" tools.

This non-standardized, manual process is time-consuming, tedious, and prone to errors. It also puts the burden of manual bookkeeping on operators, without an effective way to assure accuracy or completeness. For example, accidentally reusing an intermediate file from a previous, structurally similar run can lead to incorrect results while being very hard to detect.

Due to this, it is often cumbersome or even impossible to track down the source of any deviations or inconsistencies seen in the final output of a workflow. Even more concerning is the potential for hidden, non-obvious errors, which can corrupt research results, and may go undetected indefinitely.

In the following sections, we present our software RCE, which allows engineers to construct and automatically execute *workflows* comprising tools that are available via network connections. These workflows comprise the major tools present in the workflow as well as additional "helper" tools that, e.g., control the flow of information or convert between data formats. RCE imposes few requirements on the software tools contained in the workflow and none on their packaging, i.e., it does not require software to be available as a single file or archive. Thus, it greatly simplifies the construction and execution of workflows comprising a wide range of software, both off-the-shelf and custom-made.

During execution of a workflow, the tools remain on their respective machines, only the input and output data is moved around the network. This allows tool operators and tool developers to retain full control over the version of the tool that they provide to other engineers. Moreover, during execution of a workflow, RCE retains both the inputs and outputs of each individual execution of a tool in a central database. This enables engineers to easily trace unexpected result data to its original occurrence and to retain detailed information about the origin of simulation data.

3 RCE

Targeting engineers' needs for integrating MBSE methods into day-to-day-work, RCE provides a framework for a) controlling the data flow of engineer tools' input and output, b) sharing tools across organizational borders in a secure manner, and c) automatically executing distributed workflows. As of January 2022, RCE supports multiple operating systems such as Microsoft Windows and Windows Server as well as various Linux distributions. It is possible to run RCE with a GUI (Graphical User Interface) or as terminal-only application which makes it suitable for use on both workstations and servers. While GUI instances are mainly used to design and execute workflows and to subsequently access the result data, terminal instances are usually used as a tool server or network connection point.

In the remainder of this section we present a step-by-step introduction to RCE from a user's perspective. We start with the integration of a software tool into RCE and conclude with the execution of a distributed workflow. For a more technical description of RCE's features and capabilities, we refer to work by Boden, Flink, Först, et al. [4]

Before executing a tool via RCE, it first has to be *integrated* into RCE, thus turning it into a *workflow component*. To integrate a tool, RCE requires it to be executable as a terminal application without further user interaction during execution. This means that tools can launch a graphical user interface as part of their operation when needed, but must not require any user interaction to complete their execution. A tool can be integrated either manually via configuration files or with the "Tool Integration Wizard", which automatically

creates these configuration files. This allows users to simply integrate a tool with the aid of the wizard and to export the configuration files to a server installation afterwards. As tools may be executed in heterogeneous environments, users can define commands for the tool's execution under Windows and Linux, making the configuration of an integrated tool independent of the operating system. Part of the integration process is defining a set of *inputs* and *outputs* for the respective workflow component. These inputs and outputs then can be accessed by the user-defined *pre-processing script*, *execution script*, and *post-processing scripts*. These scripts adapt incoming and outgoing data for the tool itself and for the following components within the workflow.

RCE provides multiple options for sharing integrated tools with other RCE instances. Users can create networks of multiple RCE instances for sharing access to components and distributed workflow execution. A connection between two RCE instances can be either a normal connection in the local network, or it can be an encrypted *Uplink* connection. At least one RCE instance has to be configured as server instance for either incoming local network or Uplink connections. To prevent unauthorized access, workflow components have to be shared explicitly for each component via the "Component Publishing" view. Users can furthermore create arbitrary groups for sharing components. Shared components are then visible for each connected RCE instance only if the respective instances are part of the same authorization group. We describe the authorization and security concept in detail in Section 4.2.

After integrating a tool and thus turning it into a workflow component, users can use shared and local components to construct a workflow via the graphical "Workflow Editor". This editor is part of the GUI and the main view for editing workflows. In addition to integrated components, RCE includes a set of standard components for often-needed basic functionalities. These components easily allow for, e.g., reading and writing files or controlling the data flow within the workflow. Additionally, the users can execute their own Python scripts directly as part of the workflow, handle XML data, or create loops for evaluation or optimization purposes. For the creation of optimization loops, RCE includes the Dakota framework [1]. Users may, however, opt to use other optimization frameworks. It is possible to have multiple instances of the same component within a workflow.

The "Workflow Editor" provides utilities for workflow design, such as connecting the outputs of one component to the inputs of another component or creating labels for marking groups or important calculation steps. Figure 2 shows a workflow within the "Workflow Editor." The yellow boxes represent either an integrated or a standard component whereas the black arrows represent connections from outputs to inputs. The blue and the green rectangle are the labels marking groups of closely related component instances. It is possible to have arbitrarily many connections between a pair of tools. For the sake of readability, connections with the same direction are displayed as a single arrow.

Besides the configuration of component publishing and network connections within their respective views, RCE abstracts the complexity of shared components within the "Workflow Editor": Local and shared integrated components are indistinguishable from one another during the workflow editing process.

After having constructed a workflow using both local and shared workflow components, the user can execute it. When executing a workflow, the user has to determine an RCE instance within the network which serves as *workflow controller*. This workflow controller stores all intermediate execution data and orchestrates each component execution within the workflow. This centralized design allows to disconnect RCE instances during the execution of a workflow if they do not publish components included in the workflow. Consider, e.g., a user who designs a workflow on their mobile computer. They can designate a long-running server as the workflow controller for the execution of the workflow. After starting the execution, the user can disconnect and connect their mobile computer at will without interrupting the workflow's execution.

RCE allows the users to integrate the same tool at multiple locations in the same network. In that case, the users have to choose one of the multiple available integrations for its execution before starting the workflow. After starting a workflow, RCE executes all components which do not need input data and forwards their output data to the subsequent component in the workflow. The workflow controller automatically receives and distributes the data to the respective host instances of each component via the network setup. During the execution of a component representing an integrated tool, all of it's scripts, i.e, the pre-processing, execution,
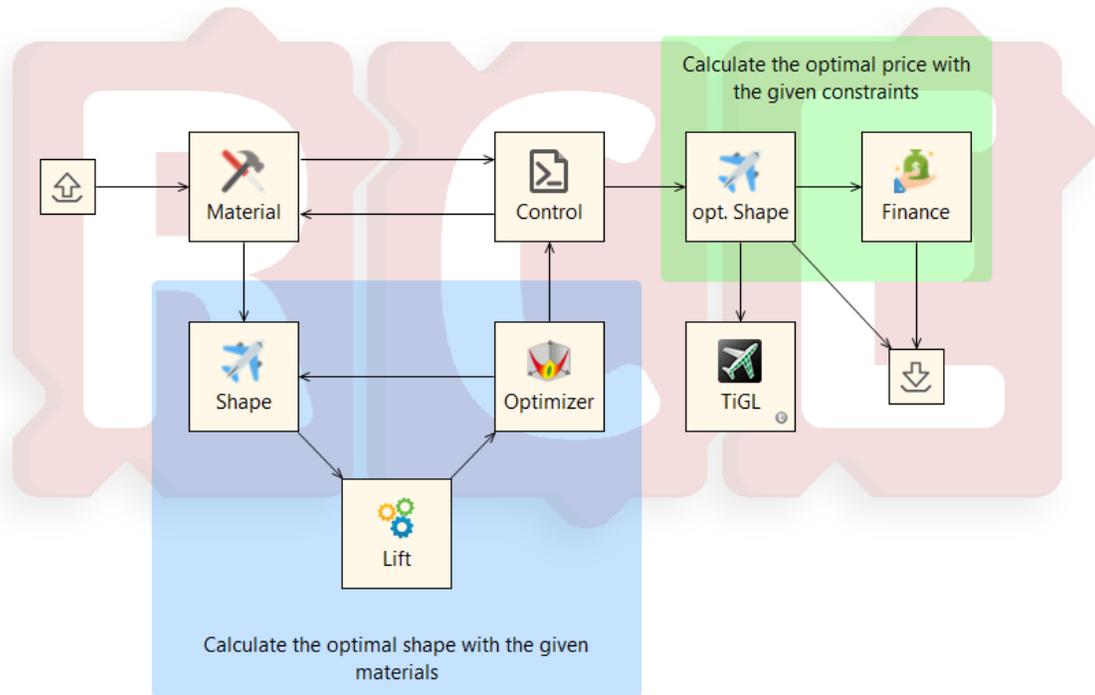
4

Fig. 2: A workflow in RCE. Originally produced by Boden, Flink, Först, et al. [4], used with permission.

and post-processing script, are executed in order. The workflow finishes when no component produces any further output.

All input and output data of all component executions of the workflow run are stored in the "Data Management". Already during the execution of a workflow and after its completion, these data can be accessed by all RCE instances connected to the workflow controller via the RCE "Data Management Browser", which is included in the GUI of RCE. Thus, users can benefit from collaborative access to the result data. Moreover, additional data of the workflow and component executions are stored in the data management and enable the workflow run to be monitored and evaluated. This includes start and end times of executions and log outputs of user integrated tools.

3.1 User Roles Interacting with RCE

As previously described, RCE allows to integrate tools and share the respective components to design workflows to develop and analyze complex systems. Due to the many different tasks RCE covers, users interact with RCE in a number of different roles. Most commonly, we identify the roles of tool integrator, workflow designer, and network administrator in real-life projects. [5] A tool integrator focuses on the integration process for the used tools, e.g., setting up the different execution scripts. A workflow designer, in contrast, can ignore the technical details of the tools and concentrates their work on the semantic correctness of the overall design. Finally, a network administrator does not need to take the workflow design nor the tool integration into account, as their priority is to set-up a secure and preferably simple RCE network structure. In addition to these common roles, users reported the additional roles of workflow executor and data analyst in their respective use-cases.

To support these different modes of interaction, RCE offers specific sets of instruments and the possibility to adapt the GUI to each role's respective requirements. The advantage of this approach is that each role can focus on their task improving the efficiency of the whole workflow development process. Therefore, RCE's GUI consists of multiple views where each view holds an instrument. A set of views is called a

perspective. RCE allows to work with multiple perspectives, enabling seamless switching between tasks and their respective instruments.

3.2 Development Process of RCE

RCE has emerged from earlier, more specialized research projects, and has been under development since 2006. Today, it is an open-source project publicly available for download from `https://rcenvironment.de/` with its source code available from `https://github.com/rcenvironment` on GitHub.

All changes to the code base are tracked with a version control software, allowing to preserve the whole history of the project and to access and rebuild older releases. The development cycle is oriented on semantic versioning [18], meaning there are major, minor, and patch releases on a regular basis. Each release undergoes a quality management procedure containing, among other things, continuous integration, automated tests, and manual testing [10].

Following semantic versioning, all releases within a major cycle are compatible with each other, which is especially important for interoperation within RCE networks. Concerning the loading of existing workflows, we aim to provide backwards compatibility, even for older major release cycles, for as long as possible. Whenever the structure of workflow files needed to change in the past, we have continuously provided automated upgrade capabilities.

4 Practical Considerations

In this section we address practical aspects of constructing, operating, and using distributed tool-based workflows. We connect these aspects to various design decisions of RCE, explain the trade-offs resulting from these decisions, and explain how the operation of RCE networks fits into the technical and administrative environment of real-world organizations.

4.1 Providing Tools as Computational Services

From a purely model-based perspective, the tools within a tool chain can be viewed as black boxes which are executed in an abstract computing environment. In practice, however, the location and environment where tools are installed and executed is highly relevant.

For instance, tools may require specific *operating system and software environments*, e.g., certain operating systems or the presence of specific pre-installed versions of libraries. Further, different locations allow for different degrees of *technical monitoring, access control and IT management*, only some of which may be sufficient to meet organizational policies. Moreover, especially for high-performance computing tools, knowing and selecting the exact *properties of the underlying hardware* is important for optimal performance. Additionally, the *network performance* of an installation location is also relevant for large or frequent data transfers. Also, the *ease of deployment* of tool updates may also vary greatly between locations. Finally, *billing for the usage of underlying resources* (e.g., computing time or licenses) may also be required in some environments. Conversely, the usage of the tool itself may also be subject to *auditing and accounting* needs, especially when used by external partners.

For any combination of these reasons, both developers and administrators typically want to choose the proper locations for the execution of their tools on a case-by-case basis. In contrast, distributing actual tool installations (i.e., binary files or scripts) over the network is generally not desirable, and often not permitted for security reasons. Instead, it is much more useful to publish and consume tools as abstract computational services, while executing them in these explicitly chosen locations. Consequently, RCE's fundamental approach to the distributed execution of multi-tool workflows is to provide the execution of pre-installed tools as remote-callable services.

While this service-based approach offers many advantages, it is not completely without downsides. The most important one is that it is not possible to transparently move the execution of a tool to where the data to be processed resides, or where the generated output data is needed. While mostly irrelevant for small

volumes of data, this can become a performance issue when large amounts of data must be processed or generated.

While RCE cannot, and is not meant to, solve this installation problem automatically, its flexible network structure allows users to develop solutions for these scenarios. As the same tool can be provided from any number of locations, any tool that regularly processes large amounts of data can be installed close to the source or the receiver of that data. By also publishing this additional installation within the RCE network, this location becomes available as a selectable option when starting a workflow involving that tool. By then choosing the appropriate tool locations at the start of the workflow, users can reduce the performance impact of large data transfers.

4.2 Authorization of Tool Access

At the time of its inception, RCE was built for closely-tied working groups with no need for access control on tools or data. Nowadays, however, it is being deployed in multidisciplinary environments involving different projects, organizational units, and increasingly, even multiple organizations. While these setups are protected from outside access using standard IT security measures, there are still valid reasons for further controlling the access to published tools.

For instance, tools frequently represent novel scientific or engineering approaches, which may not have been published yet. While its author(s) may trust selected colleagues with access to it (e.g., for testing and evaluation), they may not want to provide it to a wider audience yet.

Tools may also be installed in a location where they can use confidential code or data for their computation. For example, an industry partner may allow academic partners to simulate certain scenarios using proprietary technical datasets, without sharing that data directly. Given organization-wide and unlimited access, however, a malicious actor could potentially craft specialized workflows to reverse engineer parts of these datasets. To prevent this, such partners need an option to restrict these simulations to a minimal, well-known, and trusted set of users.

As an additional example, tools may also use limited or costly computational resources. In that case, it is desirable to restrict the tool's usage to authorized members for proper billing, or to avoid unexpected cost overruns.

Starting with RCE 9, RCE offers a limited security mechanism to address these needs. This mechanism consists of a group-based concept, with each group being defined by ownership of a shared secret cryptographical key. From a user's perspective, joining a group is similar to accessing a password-protected area. Every user can publish tools for any number of groups they are a member of, and consume all tools that are published for at least one group they are a member of. If a tool is considered generally useful, and free from the restrictions mentioned above, RCE still offers the option of publishing it to anybody within the same RCE network.

Unlike a standard login-based approach, RCE's access control is decentralized, i.e., there is no central server for checking a user's authentication. Instead, any user can create a new group at any time, and invite other users to join it. These groups are implicitly disposed when they are not used for sharing tools anymore. This approach was chosen to match RCE's focus on flexibility and bottom-up cooperation, and to avoid the overhead of a static authentication infrastructure. We discuss a future evolution of this system in Section 7.4.

4.3 Iterative Development and Updating of Tools

Especially in research, frequent development, iteration, and experimentation regarding tools within a workflow is common. These tool updates may involve new algorithmic approaches being tested, new features added, or issues processing particular input data being fixed.

RCE is well-suited to this iterative work due to several properties. Firstly, any tool that is updated in its execution location is immediately available in the new version to all users within the tool network. Secondly, especially in early development, it is possible for a tool's author to publish a current version directly from their development machine. While execution is usually not as performant as running it on a server or cluster, the

rapid iteration and debugging possibilities this allows are often a great advantage. Lastly, the deterministic and structured approach of RCE's workflow execution makes it easy to re-run existing workflows with new versions of a tool. This can be both used to validate pre-existing behavior when no changes are expected, and also to immediately verify the effect of improvements within the workflow's context.

4.4 Cross-Organizational Tool and Workflow Publication

RCE is widely used for collaboration between organizational units within larger organizations, e.g., multiple institutes of a single research center, each of which is specialized on different parts of the domain space. A typical setup for this is a multidisciplinary workflow comprising tools published by these organizational units, thereby contributing to the overall workflow.

A natural expansion of this scenario is the integration of tools provided by domain experts located in disjoint organizations. For example, industry partners with access to physical test sites or prototype designs and academic partners exploring new algorithms and methods may mutually benefit from combining their tools.

While this change regarding tool locations is virtually irrelevant from the MBSE domain perspective, it has profound technical, administrative, legal, and security implications. From a technical standpoint, one major challenge is having all involved parties agree on a common protocol for exchanging access to the relevant tools. Administratively, one important aspect is explicit control of which tool services should be made available to which partners, and potentially reviewing their usage. Legally, both research and industry organizations are acutely aware of the need to protect intellectual property, especially in the form of tool code or data. Lastly, IT security is of paramount importance, both as a requirement in and by itself, as well as providing the foundation for these legal and administrative needs.

As RCE was not originally designed with cross-organization workflows in mind, adopting its standard network features to this would be difficult. Instead, we opted to design and implement the entirely new, security-first network protocol named *Uplink*. It is publicly available in RCE 10 as an experimental feature and is currently being tested by several project groups.

Instead of attempting to retroactively restrict operations within the standard RCE network, its design philosophy is the opposite. Starting from scratch, only a restricted and safe set of operations is explicitly provided as needed. In its current form, it supports announcing tool services and consuming them, with the auxiliary feature of transferring tool documentation to users. In contrast, accessing the data management of RCE instances in other organizations or starting remote workflows is explicitly forbidden. While this is – by design – more restrictive than a local network, it allows fully automated integration of cross-organization tools into local MBSE workflows.

5 Concrete Use Cases

In the previous section we have given an overview over the design decisions taken when developing RCE and the trade-offs involved in this. In this section we discuss real-life research projects using RCE.

Many users already rely on RCE in their daily work and appreciate its capabilities. [7, 13, 14, 16] In addition to DLR-external projects, there are many research projects at DLR itself in which RCE is used that contain concrete use cases. These use cases are in application domains covering both the domain of aerospace as well as in the domains of traffic and energy.

We now highlight two projects actively using RCE that implement MBSE, namely JETSCREEN and EXACT. Due to its nature as freely distributed open-source software, we omit a complete overview of all projects in which RCE is involved. Beyond that, Boden, Flink, Först, et al. [4] present further projects that use RCE.

5.1 JETSCREEN

The project JETSCREEN (JET Fuel SCREENing and Optimization) aims to support the development of sustainable aviation fuels and reduce climate impact of aviation. To do so, it implements a model and a

simulation-based process to generate key indicators out of a detailed composition of a candidate fuel. These key indicators then support the fuel approval process. The stated objectives of this project are "to develop a screening and optimization platform, which integrates distributed design tools and generic experiments to assess the risks and benefits of alternative fuels, and to optimize alternative fuels for a maximum energy per kilogram of fuel and a reduction of pollutants emissions." [20]

To achieve this goal, project members implemented a cross-organizational RCE network. In this network, the individual disciplinary predictive and simulation software tools of the project partners were integrated into RCE and made accessible to other participants. Based on this infrastructure, RCE workflows were developed which enable the automatic, distributed execution of the entire simulation processes. These workflows include all the competencies of the project partners that are necessary for the evaluation of fuels. To support the possibility of cross-company networking, the Uplink capabilities of RCE were further developed as part of this project.

5.2   EXACT

The currently ongoing DLR-internal project EXACT (Exploration of Electric Aircraft Concepts and Technologies) has the overall aim of evaluating the potential of climate-neutral aircraft on a system level approach. This evaluation includes both operational as well as social aspects. The DLR research divisions of aviation, energy, space and transport are participating in this project. In doing so, they, among other things, provide and extend their domain-specific simulation software tools. As part of the project, various aircraft concepts are designed taking into account various factors such as energy carrier or seat-class. In this use case, RCE is applied to enable the fast and flexible integration of new technologies and disciplines into multidisciplinary dimensioning processes. These iterative processes include several disciplinary, semi-empirical and physics-based models with low to medium fidelity. [22]

For this purpose, a DLR-internal RCE network was set up and the simulation tools were integrated into RCE. This means that all tools and models from the project partners are accessible within the integration framework for setting up the design tool chain. In the context of this project, the ability of RCE to simply integrate an RCE workflow itself as a workflow component including the possibility to make it available in the network for reuse is being extended. In particular, a GUI for this purpose is being developed.

6   Related Work

The application that is closest to RCE in terms of features it offers is ModelCenter, developed by Phoenix Integration [17]. Similarly to RCE, this application allows users to combine their purpose-built software tools into a workflow and to execute that workflow automatically. Model Center is proprietary software and not freely available. RCE, in contrast, is developed as open-source software and freely available under a permissive license.

Moreover, Apache Nifi [3] allows its users to construct data pipelines. Such a pipeline consists of individual processors that process incoming data and pass their output to one or more succeeding processors. These pipelines are conceptually similar to the workflows that can be implemented in RCE. Nifi, however, focuses on indefinitely-running pipelines that continually process incoming data from a variety of sources. The focus of RCE, in contrast, lies on the implementation of long-running workflows that process a batch of data in a single execution.

Another project similar to RCE is the Air Traffic Management (ATM) Test Bed by NASA [6]. This application allows users to simulate and experiment with concepts for the management of air traffic. Similarly to RCE, it features a graphical canvas on which users can arrange and connect individual scenario generation and simulation tools. In contrast to RCE, however, the ATM Test Bed specifically targets the domain of air traffic management. It is, to the best of our knowledge, not well-suited to, e.g., the simulation of individual planes or the application in other domains such as the simulation of spacecraft or energy systems.

7 Future Work

RCE is under active development as part of multiple research and engineering projects. A common theme of most ongoing work is maintaining the stability and further improving the efficiency and user-friendliness of RCE. Additionally, we are continuously monitoring technical and funcional needs, and support users in deploying RCE for new use cases. In this section, we highlight some of the main topics for the future development of RCE.

7.1 Decentralized Data Management

One major strength of RCE is its support for automatically storing detailed input, intermediate, and output data for each executed workflow. Each data transfer between integrated tools is recorded for later inspection.

While the execution of the individual tools is decentralized, this data archiving is centralized on the RCE instance controlling the workflow execution. This setup guarantees that all archived data is consistent and cannot be partitioned due to later unavailability of other instances. The drawback of this approach, however, is the performance overhead of transferring all intermediate artifacts to a central server, which also limits its scalability. Additionally, this centralized data retention can be problematic for data which is affected by special intellectual property regulations.

For future work, we are investigating the usefulness of decentralized data storage mechanisms for improved execution performance and scalability. Depending on project needs, special provisions for more fine-granular access control to proprietary data may also be added.

7.2 Provenance Recording

Furthermore, while the data management of RCE stores all input and output data for each execution of the components of a workflow, it only retains limited information about the environment in which each component was executed. The data management also does not currently store the explicit relation between the output of one component and the input of the subsequent component. Instead, it only retains a copy of the executed workflow from which the user can reconstruct the input-output relations between component executions.

Explicitly storing such implicit information about the provenance of data would greatly assist users in developing and debugging workflows. It would additionally allow for easier auditability and replicability of workflow execution. We are currently working on a prototypical implementation of provenance recording in RCE using the W3C PROV standard and provenance templates [11, 12].

7.3 Footprint Reduction, APIs, and Ecosystem Integration

Currently, RCE is packaged as a single software application that can be used in all contexts discussed in this work. This includes interactive work via its GUI, deployment as a tool publishing and workflow coordination platform, as an intermediate communication relay, or for batch execution by external scripts. While this single release package is flexible enough to be used in all of these roles, it could be optimized by providing more specialized packages for each. Especially for usage in containerized or cloud environments, reducing its deployment footprint by omitting client-side features would be beneficial.

Furthermore, while RCE instances are designed for interconnecting with each other, and for actively managing other tools, it provides relatively few options for integrating itself into existing software ecosystems. To address this, we intend to provide explicit APIs that can then be invoked by other software, e.g. for starting and querying workflow executions, accessing the data management, or managing network connections.

7.4 User Management and Access Control

Currently, RCE only provides a basic mechanism for controlling access to published tools, the "authorization groups" described in Section 4.2. The common security boundary at the moment is the access to RCE networks themselves; internally, there are few restrictions so far. In particular, RCE does not feature detailed mechanisms for restricting the access to workflow data, more fine-granular permissions for tool execution, or the monitoring of workflows yet. Moreover, the access groups for controlling tool access are managed by individual RCE users, i.e., the management of user groups is decentralized.

From an IT administrator's perspective, a centralized approach would fit better into the IT structures of most organizations. This particularly involves central control and revocation of access keys. It is, e.g., critical for IT security that members leaving an organization are automatically prevented from accessing future work within the organization's network. From a user's perspective, the approach currently implemented in RCE has the benefit of allowing flexible and low-overhead creation of new groups. A centralized approach, in contrast, would require organizational structures, responsibilities, and procedures for managing access control.

Permissible and practical methods for user authentication are largely determined by the policies of the companies using RCE. The variety of existing authentication schemes and approaches within different organizations complicates designing a common and generally applicable approach. Thus, it is infeasible to develop a solution for user management and user authentication that serves all use cases encountered in practice at once. As the networks and workflows built with RCE continue to grow larger, however, a more centralized and sophisticated framework for access control and authentication will be indispensable.

We are looking to design and implement a framework that allows for lightweight user management, that can be sustainably integrated with existing solutions, and that allows for sufficiently granular control over information in coordination with future project partners.

7.5 Cloud Computing

The increasing adoption of cloud computing in science and engineering offers new options and challenges for the deployment of scientific tools. While RCE does not yet have explicit support for cloud computing, its focus on distributed workflow management would be a natural fit for it. With individual tools deployed in a cloud environment using standardized interfaces, RCE could provide the infrastructure for orchestrating the execution of such tools in workflows.

However, the practical aspects regarding tool installations listed in Section 4.1 still apply. Thus, significant research and design work is still needed for safely moving real-world tools and workflows into cloud environments.

Finally, in addition to the topics mentioned in this section, we are working closely with known users of RCE to determine and anticipate current and future requirements towards MBSE support in general, and particular needs regarding the software itself. This work helps us prioritize and steer future development.

8 Conclusion

In this work we have presented RCE, an application developed at DLR that allows its users to construct and automatically execute tool chains comprised of both off-the-shelf and tailor-made software tools. This aids engineers in constructing, analyzing, and optimizing the model constructed as part of MBSE. We have subsequently given an overview over RCE's features and discussed the tradeoffs made in their practical implementation. Moreover, we have highlighted projects in which RCE has been used and given an outlook on future areas of development.

RCE has enabled, supported, and simplified numerous research and development projects in a wide variety of fields over the years. It is used in ongoing development projects, its adoption has significantly increased over the years, and there is great interest in a plethora of directions for future development on the side of users and project partners. We are thus looking forward to continue development of RCE to adapt it to new use cases and to evolve it in tandem with the state of the art in the wider field of application.

References

1. Adams, B.M., Bohnhoff, W.J., Dalbey, K.R., Ebeida, M.S., Eddy, J.P., Eldred, M.S., Hooper, R.W., Hough, P.D., Hu, K.T., Jakeman, J.D., Khalil, M., Maupin, K.A., Monschke, J.A., Ridgway, E.M., Rushdi, A.A., Seidl, D.T., Stephens, J.A., Swiler, L.P., Winokur, J.G.: Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.14 User's Manual. Tech. rep., Sandia National Laboratories (2021)

2. Alder, M., Moerland, E., Jepsen, J., Nagel, B.: Recent Advances in Establishing a Common Language for Aircraft Design with CPACS. In: Aerospace Europe Conference (2020)

3. Apache Software Foundation, Cloudera, Hortonworks: Apache Nifi, `https://www.w3.org/TR/prov-dm/`

4. Boden, B., Flink, J., Först, N., Mischke, R., Schaffert, K., Weinert, A., Wohlan, A., Schreiber, A.: RCE: An Integration Environment for Engineering and Science. SoftwareX **15**, 100759 (Jul 2021). https://doi.org/10.1016/j.softx.2021.100759

5. Boden, B., Mischke, R., Weinert, A., Schreiber, A.: Supporting the composition of domain-specific software via task-specific roles. In: Interconnecting Code Workshop (ICW). ACM (Mar 2020). https://doi.org/10.1145/3397537.3399576

6. Chan, W.N., Barmore, B.E., Kibler, J., Lee, P., O'Connor, N., Palopo, K., Thipphavong, D., Zelinski, S.: Overview of NASA's Air Traffic Management-eXploration (ATM-X) Project. In: AIAA Aviation Technology, Integration, and Operations Conference. AIAA (2018)

7. Cruz, A.A.R., Fregnani, J.A., Mattos, B.S., Murca, M.C.R.: Identification of the Actual Mission Profiles and Their Impact on the Integrated Aircraft and Airline Network Optimization. In: AIAA Aviation 2021 Forum. AIAA. https://doi.org/10.2514/6.2021-2453

8. Estefan, J.A.: Survey of Model-Based Systems Engineering (MBSE) Methodologies. Tech. rep., Incose MBSE Focus Group (2008)

9. Madni, A.M., Sievers, M.: Model-based systems engineering: Motivation, current status, and research opportunities. Systems Engineering **21**(3), 172–190 (May 2018). https://doi.org/10.1002/sys.21438

10. Mischke, R., Schaffert, K., Schneider, D., Weinert, A.: Automated and Manual Testing in the Development of the Research Software RCE. In: SE4Science Workshop 2022, pp. 531–544. Springer. https://doi.org/10.1007/978-3-031-08760-8_44

11. Moreau, L., Batlajery, B.V., Huynh, T.D., Michaelides, D., Packer, H.: A Templating System to Generate Provenance. IEEE Transactions on Software Engineering **44**(2), 103–121 (Feb 2018). https://doi.org/10.1109/tse.2017.2659745

12. Moreau, L., Missier, P., Belhajjame, K., B'Far, R., Cheney, J., Coppens, S., Cresswell, S., Gil, Y., Groth, P., Klyne, G., Lebo, T., McCusker, J., Miles, S., Myers, J., Sahoo, S., Tilmes, C.: PROV-DM: The PROV Data Model. Online (2013), `https://www.w3.org/TR/prov-dm/`, accessed October 11, 2021

13. Niklaß, M., Dzikus, N., Swaid, M., Berling, J., Lührs, B., Lau, A., Terekhov, I., Gollnick, V.: A Collaborative Approach for an Integrated Modeling of Urban Air Transportation Systems. Aerospace **7**(5), 50 (Apr 2020). https://doi.org/10.3390/aerospace7050050

14. Nöding, M., Bertsch, L.: Application of Noise Certification Regulations within Conceptual Aircraft Design. Aerospace **8**(8), 210 (Aug 2021). https://doi.org/10.3390/aerospace8080210

15. Object Management Group: OMG System Modeling Language. Tech. rep., Object Management Group (2019)

16. Papanikolaou, A., Harries, S., Hooijmans, P., Marzi, J., Néna, R.L., Torben, S., Yrjänäinen, A., Boden, B.: A Holistic Approach to Ship Design: Tools and Applications. Journal of Ship Research pp. 1–29 (Dec 2020). https://doi.org/10.5957/josr.12190070

17. Phoenix Integration: Model Center, `https://www.phoenix-int.com/`

18. Preston-Werner, T.: Semantic Versioning 2.0.0 (2013), `https://semver.org/`, accessed October 15, 2021

19. Rashid, M., Anwar, M.W., Khan, A.M.: Toward the tools selection in model based system engineering for embedded systems—A systematic literature review. Journal of Systems and Software **106**, 150–163 (Aug 2015). https://doi.org/10.1016/j.jss.2015.04.089

20. Rauch, B.: JETSCREEN: JET fuel SCREENing and optimization (2020), `https://cordis.europa.eu/project/id/723525`, accessed October 12, 2021

21. Siggel, M., Kleinert, J., Stollenwerk, T., Maierl, R.: TiGL: An Open Source Computational Geometry Library for Parametric Aircraft Design. Mathematics in Computer Science **13**(3), 367–389 (Jul 2019). https://doi.org/10.1007/s11786-019-00401-y

22. Silberhorn, D., Hartmann, J., Dzikus, N.M., Atanasov, G., Zill, T., Brand, U., Trillos, J.C.G., Oswald, M., Vogt, T., Wilken, D., Grimme, W.: The Air-Vehicle as a Complex System of Air Transport Energy Systems. In: AIAA Aviation 2020 Forum. AIAA. https://doi.org/10.2514/6.2020-2622

23. Technical Operations International Council on Systems Engineering (INCOSE): Systems Engineering Vision 2020 (2007), available at `https://sdincose.org/wp-content/uploads/2011/12/SEVision2020_20071003_v2_03.pdf`, accesssed October 15, 2021