

From the Desktop to the Grid and Cloud: Conversion of KNIME Workflows to WS-PGRADE

Luis de la Garza

Center for Bioinformatics

Dept. of Computer Science

University of Tübingen, Germany

delagarza@informatik.uni-tuebingen.de

Fabian Aicheler

Center for Bioinformatics

Dept. of Computer Science

University of Tübingen, Germany

Oliver Kohlbacher

Center for Bioinformatics

Dept. of Computer Science

Faculty of Medicine

Quantitative Biology Center, University of Tübingen

Max Planck Institute for Developmental Biology

Germany

Abstract—Computational analyses for research usually consist of a complicated orchestration of data flows, software libraries, visualization, selection of adequate parameters, etc. Structuring these complex activities into a collaboration of simple, reproducible and well defined tasks brings down complexity and increases reproducibility. This is the basic notion of workflows.

Workflow engines allow users to create and execute workflows, each having unique features. In some cases, certain features offered by platforms are royalty-based, hindering use in the scientific community.

We present our efforts to convert whole workflows created in the Konstanz Information Miner Analytics Platform to the Web Services Parallel Grid Runtime and Developer Environment. We see the former as a great workflow editor due to its considerable user base and user-friendly graphical interface. We deem the latter as a great backend engine able to interact with most major distributed computing interfaces. We introduce work that provides a platform-independent tool representation, thus assisting in the conversion of whole workflows. We also present the challenges inherent to workflow conversion across systems, as well as the ones posed by the conversion between the chosen workflow engines, along with our proposed solution to overcome these challenges.

The combined features of these two platforms (i.e., intuitive workflow design on a desktop computer and execution of workflows on distributed high performance computing interfaces) greatly benefit researchers and minimize time spent in technical chores not directly related to their area of research.

Keywords—WS-PGRADE, KNIME, conversion, fine-grained interoperability, workflow

I. INTRODUCTION

Computers are essential in various scientific fields. Example domains requiring high-performance computing (HPC) include vaccine design, astrophysics, or the multidisciplinary field of bioinformatics. Here, the declining costs of both data generation and storage in the last few years [1] pushed bioinformaticians into using HPC resources such as grids and clouds.

Simultaneously, the scope of research is getting more and more refined and complex. As such, upholding the scientific method increases in difficulty: Being able to reproduce previously observed results when keeping all variables constant, can often be an arduous task. Consequently, journals and news outlets have repeatedly reported cases of published but irreproducible results [2], [3], [4].

Researchers often break down big, complicated analyses into smaller units of work that are easier to manage. These so-called tasks perform one specific function and take an input along with controlling parameters to produce a defined output. Input usually takes the form of files, whereas output could also be for example a set of visualizations. The combination of tasks is often referred to as a workflow. Task outputs can be passed on as inputs to other tasks, defining an order of execution for each step of the comprising workflow. Adoption of workflows not only increases reproducibility but also offers the following benefits:

- Storage of intermediate results (e.g., for troubleshooting, additional analysis, bottleneck identification)
- Simplified substitution of single tasks (e.g., for benchmarking, testing purposes)
- Parallel execution of workflow branches (i.e., *parameter sweep*)
- Reusability of components
- Independent, parallel development of specialized tasks

A. Workflow Interoperability and Conversion

Throughout this work we will use workflow terminology and representation consistent with our previous work [5], [6]. Figures 1 and 2 briefly summarize this.



Fig. 1. **The abstract layer of a workflow.** Vertices represent tasks, edges indicate the execution order. At this point, no implementation or technical details are represented.

Since the *abstract* workflow layer contains solely application domain information, it is independent of the execution requirements. Thus, the *abstract* layer remains unchanged across workflow engines. In contrast, the *concrete* workflow layer, the workflow engine and the executing platform are tightly coupled. This divergence of *concrete* layers across engines makes workflow interoperability challenging. Furthermore,

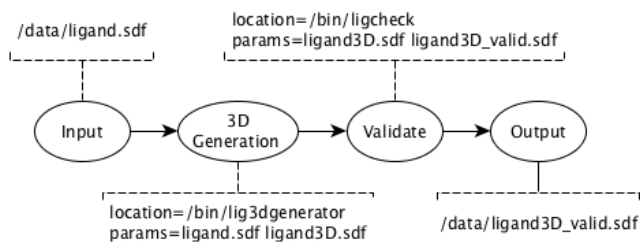


Fig. 2. The *concrete* layer of a workflow. The *concrete* layer contains implicit application domain information. But unlike the *abstract* layer, vertices are annotated with extra attributes. These are the required resources to execute the portrayed tasks.

workflow engines often contain distinct features, complicating conversion across platforms.

One way to alleviate these problems is the development of platform-independent workflow representations, e.g., the Interoperable Workflow Intermediate Representation (IWIR) [7] and Yet another Workflow Language (YAWL) [8] to enable fine-grained interoperability (FGI). However, platform-independent workflow representations do not address workflow implementations. The Sharing interoperable Workflows for large-scale scientific Simulation on available distributed computing interfaces project (SHIWA) [9], for instance, provides execution of workflows built on different workflow engines by uploading them to the SHIWA Simulation Platform. Users handling data subject to privacy restrictions (e.g., patient data) might find it an unsuitable solution.

A proper workflow conversion across engines requires that the *abstract* layer remains unchanged (i.e., source and target workflow can be considered logically equivalent). The location of resources, how different engines implement single nodes and logical constructs (e.g., *parameter sweep*) are some of the aspects to be considered. Features unique to one engine represent a complication. Figure 3 shows an example of a simplified workflow conversion.

II. IMPLEMENTATION

The Web Services Parallel Grid Runtime and Developer Environment Portal (WS-PGRADE) [10] is a web-based workflow engine that interacts with a wide array of resource managers (e.g. Moab, LSF) to access distributed computing interfaces (DCIs). This makes it a great *back-end* workflow execution engine. Tasks of the same workflow can be executed on different DCIs. However, workflow creation is a multi-step process, posing problems for users without adequate training.

The Konstanz Information Miner Analytics Platform (KNIME Analytics Platform) [11] is hosted on a personal computer. It features an intuitive interface, contains more than 1,000 pre-loaded tools and hundreds of sample workflows. Addition of new tools requires knowledge of the Java programming language—an aspect that might keep some users away from this feature. A couple of royalty-based variants (i.e., the so-called *KNIME Collaborative Extensions*) are offered to

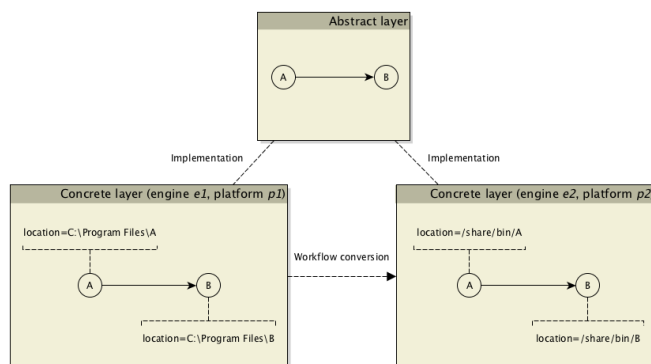


Fig. 3. **Workflow conversion challenges.** Two different engines (i.e., $e1$, $e2$) running on two different platforms (i.e., $p1$, $p2$) contain different *concrete* layers of the same workflow. The *abstract* layer, however, remains unchanged. A successful workflow conversion must take into account not only the differences among the source and target engines, but must also consider the source and target platforms or operating systems.

remotely execute workflows, however, WS-PGRADE offers a wider support for resource managers to access DCIs.

We focus on providing fine-grained interoperability between a great workflow editor such as the KNIME Analytics Platform and a versatile, scalable workflow execution platform such as WS-PGRADE.

The first step to provide interoperability is to represent tasks in a platform-independent manner. Certain attributes of tool execution remain unchanged across platforms (e.g., version and parameters), while some others change (e.g., location of executables, input and output files). Attributes in need of adjustment have to be identified. A platform-independent tool representation facilitates the task conversion across platforms and thus the conversion of full workflows.

One of the first challenges in the conversion between these engines is the maintenance of a database that relates tools on the user's computer with tools on each of the target DCI platforms. The next set of challenges concerns the implementation of nodes and logical workflow constructs. The KNIME Analytics Platform implements *parameter sweep* via node-delimited workflow sections (i.e., using *ZipLoopStart*, *ZipLoopEnd* nodes). WS-PGRADE delimits such sections with *generator* and *collector* ports. Furthermore, WS-PGRADE allows users to assign data files directly to input ports. The KNIME Analytics Platform, however, requires a dedicated node (e.g., *Input File*, *Input Files*), whose output port refers to a file and this reference can be channeled to an input port.

Some features present in the KNIME Analytics Platform are not found in WS-PGRADE. The former requires ports to declare which data types they are compatible with and supports file lists as inputs; the latter is more flexible and lacks native support of file lists as inputs (i.e., each input or output port is related to one file). Different to WS-PGRADE, *KNIME Nodes* produce outputs not only via output ports: They can also set *flow variables*, which can be read further down the execution flow.

The KNIME Analytics Platform is a Java program with a graphic interface. *KNIME Nodes* are then instances of Java classes that live inside the process which launched the KNIME Analytics Platform. In other words, they require a running instance of the KNIME Analytics Platform to be executed, making their execution on a DCI a challenge.

The following sections describe our approach to address the mentioned challenges.

A. Conversion of Nodes: Addressing Disparities between Workflow Engines

The KNIME Analytics Platform features a *node repository* in which users can select any of the available nodes (see Figure 4). Creation of workflows in the KNIME Analytics Platform requires a single step, thus the *abstract* and *concrete* layers are merged into the user-friendly workflow editor. Each *KNIME Node* performs a specific task and defines a fixed number of input, output ports. Each port is associated to a *port type*, which is similar to content types (e.g., *csv*, *pdb*). Only ports of compatible types can be interconnected. Furthermore, *KNIME Nodes* rely on the assumption that incoming and outgoing data are arranged in custom in-memory *data tables*. Each *KNIME Node* iterates over the rows of incoming data and is able to modify the contents of the input table, as well as its structure (e.g., by adding columns or rows). File handling is done by using these same data tables, their cells containing uniform resource identifiers (URI) pointing to the needed files.

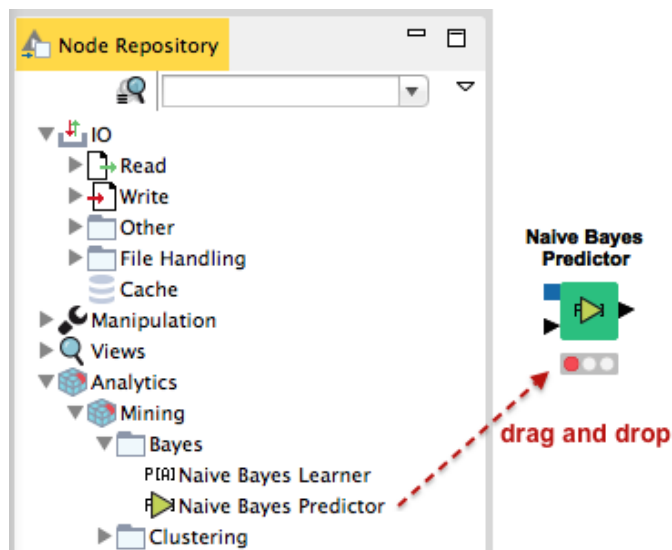


Fig. 4. **The KNIME Analytics Platform Node Repository.** Available nodes can be selected from the node repository by *drag and dropping* them into the current workflow editor.

WS-PGRADE, on the other hand, requires the creation of an *abstract* and a *concrete* workflow in a multi-step process (see Figure 5). During the creation of the *concrete* workflow, users input the required attributes and command line to associate a node to a specific remote binary. In contrast to the KNIME Analytics Platform, WS-PGRADE allows to assign files directly to input ports and it doesn't perform a strict type

checking: Any output port can be connected to any input port. Additionally, the structure of the incoming and outgoing files is arbitrary.

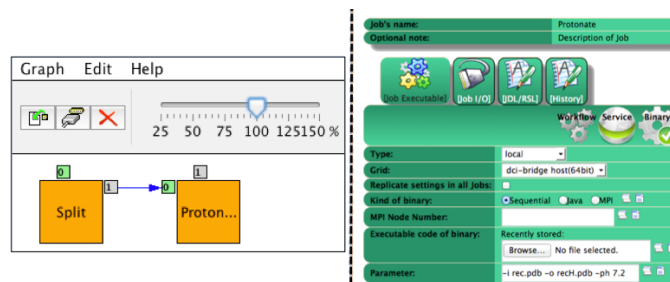


Fig. 5. **The multi-step creation of Workflows in WS-PGRADE.** An *abstract* is first created (left pane), after which a *concrete* can be created and configured (right pane).

Adding nodes to the KNIME Analytics Platform requires knowledge of the Java programming language. Generic KNIME Nodes (GKN) [5], [6] was developed to add nodes without programming experience by allowing arbitrary command line tools to behave as *KNIME Nodes* and to seamlessly interact with other nodes inside the KNIME Analytics Platform. The only requirement is the representation of the tools by Common Tool Descriptors (CTDs), which are XML files describing the inputs, outputs and parameters of a tool [5], [6]. Currently, several software suites [12], [13], [14] are able to parse and generate CTDs (i.e., they are *CTD-enabled*). Figure 6 illustrates how CTDs interact with *CTD-enabled* tools.

We introduce KNIME2gUSE, an extension to the KNIME Analytics Platform which converts workflows from the KNIME Analytics Platform to WS-PGRADE, combining the features of both engines and overcoming their disadvantages.

Conversion of *KNIME Nodes* that were imported using GKN is somewhat trivial. Each of these nodes represents an external tool that is independent of the KNIME Analytics Platform. In this case, the matching binary for the represented tool is required on each of the target DCIs.

We identify *native* nodes as those *KNIME Nodes* that were not imported using GKN (i.e., pre-packaged nodes, nodes added as third-party extensions or nodes added by the user via other means). Each *native KNIME Node* is an instance of a Java class managed by the KNIME Analytics Platform. Such nodes exist only in the context of the process that hosts the KNIME Analytics Platform. Execution of a single *KNIME Node* requires a running instance of the KNIME Analytics Platform and converting these nodes is not trivial. Furthermore, a suitable distribution of the KNIME Analytics Platform must be present on each of the target DCIs.

Data between *KNIME Nodes* can only be channeled between ports with compatible data types. Since channeled data are in-memory representations of table-formatted data (i.e., *data tables*), we have devised a solution that allows *native KNIME Nodes* to be executed *as if* they were command line tools: During the export process, *native KNIME Nodes* are individually packed into a small KNIME workflow. Each such

generated workflow contains a copy of the original node, along with any user-established settings. Since inputs and outputs for the exported node won't be channeled *inside* an instance of the KNIME Analytics Platform, extra reader and writer nodes (i.e., *Table Reader* and *Table Writer*) are also included in this small workflow. These nodes allow the serialization and deserialization of the in-memory data format required by *native KNIME Nodes*. The KNIME Analytics Platform can execute workflows in a so-called *batch mode*, without the need of a graphical user interface. A suitable command line is automatically generated during our export process. When the *batch mode* execution of this generated workflow is started, input files will be read into the KNIME *data table* format; upon completion, any output will be serialized from the KNIME *data table* format into a file.

The work previously presented in [5], [6] introduced work we have done in the field and showcased conversion of KNIME workflows composed solely of nodes that were imported via GKN. We have extended KNIME2gUSE in order to convert workflows composed of any kind of nodes. Figures 7 and 8 depict how the conversion of nodes is performed.

```
$ Ligand3DGenerator -i in.sdf -o out3D.sdf -ff MMFF94
```

```
params.ctd
<tool name="Ligand3DGenerator">
  <ITEM name="i" type="input-file" value="in.sdf" />
  <ITEM name="o" type="output-file" value="out3D.sdf" />
  <ITEM name="ff" type="string" value="MMFF94" />
</tool>
```

```
$ Ligand3DGenerator -par params.ctd
```

Fig. 6. A CTD in action. Top section: Parameters needed for the tool *Ligand3DGenerator* [14] to be executed, namely input file, output file and the desired force field. Middle section: A sample CTD snippet representing an execution of the *Ligand3DGenerator* using the shown parameter values. Bottom section: A CTD-enabled tool with the given sample CTD.

B. Conversion of Workflows: Exporting KNIME workflows to WS-PGRADE

The KNIME2gUSE plug-in produces files that can be imported into WS-PGRADE, ready to be executed on any configured DCI with minor modifications.

We have chosen WS-PGRADE as the target engine for the export process due to the fact that it interacts directly with a wide selection of resource and cloud managers (a feature not present in the royalty-based KNIME editions that allow remote execution). It also features workflow submission, control, monitoring and statistics. These are functionalities which resource managers or cloud engines often lack.

The KNIME Analytics Platform natively supports the association of single input/output ports to a file list determined at *runtime*, a functionality not present in WS-PGRADE. To overcome this, a wrapper script is automatically generated by KNIME2gUSE that zips corresponding files into a single archive. To translate *parameter sweep* sections, conversion removes KNIME Analytics Platform *ZipLoopStart* and *ZipLoop-*

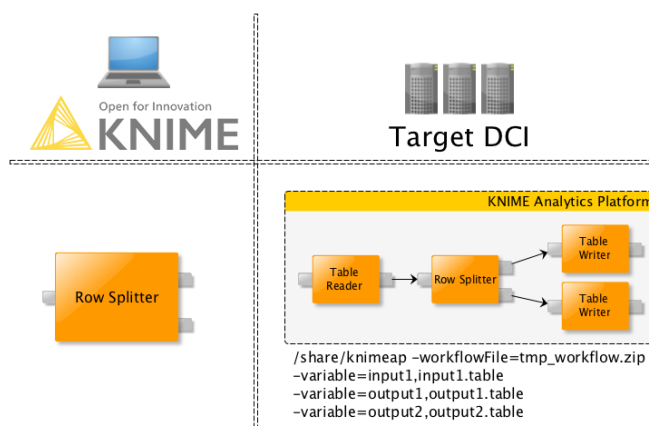


Fig. 7. Conversion of *native KNIME nodes*. Since *Row Splitter* exists only in the context of the KNIME Analytics Platform, it requires an instance of a KNIME Analytics Platform for its execution. KNIME2gUSE generates a workflow containing the required input/output nodes and a copy of *Row Splitter* with the same configuration settings as its origin node. The generated command line invokes a KNIME Analytics Platform in the target DCI in *batch mode*, reads input from a file and writes outputs to files.

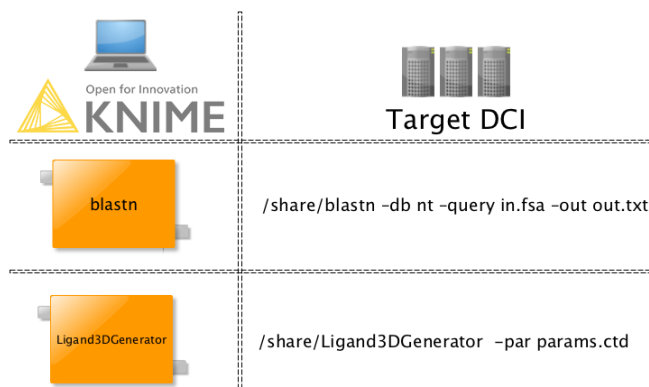


Fig. 8. Conversion of *GKN-imported nodes*. The nodes depicted on the left side directly interact with binaries located on the user's desktop computer. The right column shows the mapped binaries and an equivalent execution on a target DCI. Since *Ligand3DGenerator* is CTD-enabled, a suitable CTD file can be generated; this is not the case for the *blastn* tool.

End nodes and substitutes suitable WS-PGRADE *generator* and *collector* ports.

C. Example Application: Biomarker Discovery in Metabolomics

Metabolomics is a mass spectrometry-based approach aimed to evaluate the entirety of a metabolite sample. Applications include the tracking of chemicals and their transformation products in waste water [15], identification of cancer types via biomarkers [16], [17] and elucidation of disease-underlying mechanisms [18]. Compared to complementary *omics* technologies (e.g., transcriptomics, proteomics), metabolomics is closer to the actual biochemical processes that occur, making it attractive for biomarker development.

A common analysis approach for studies interested in com-

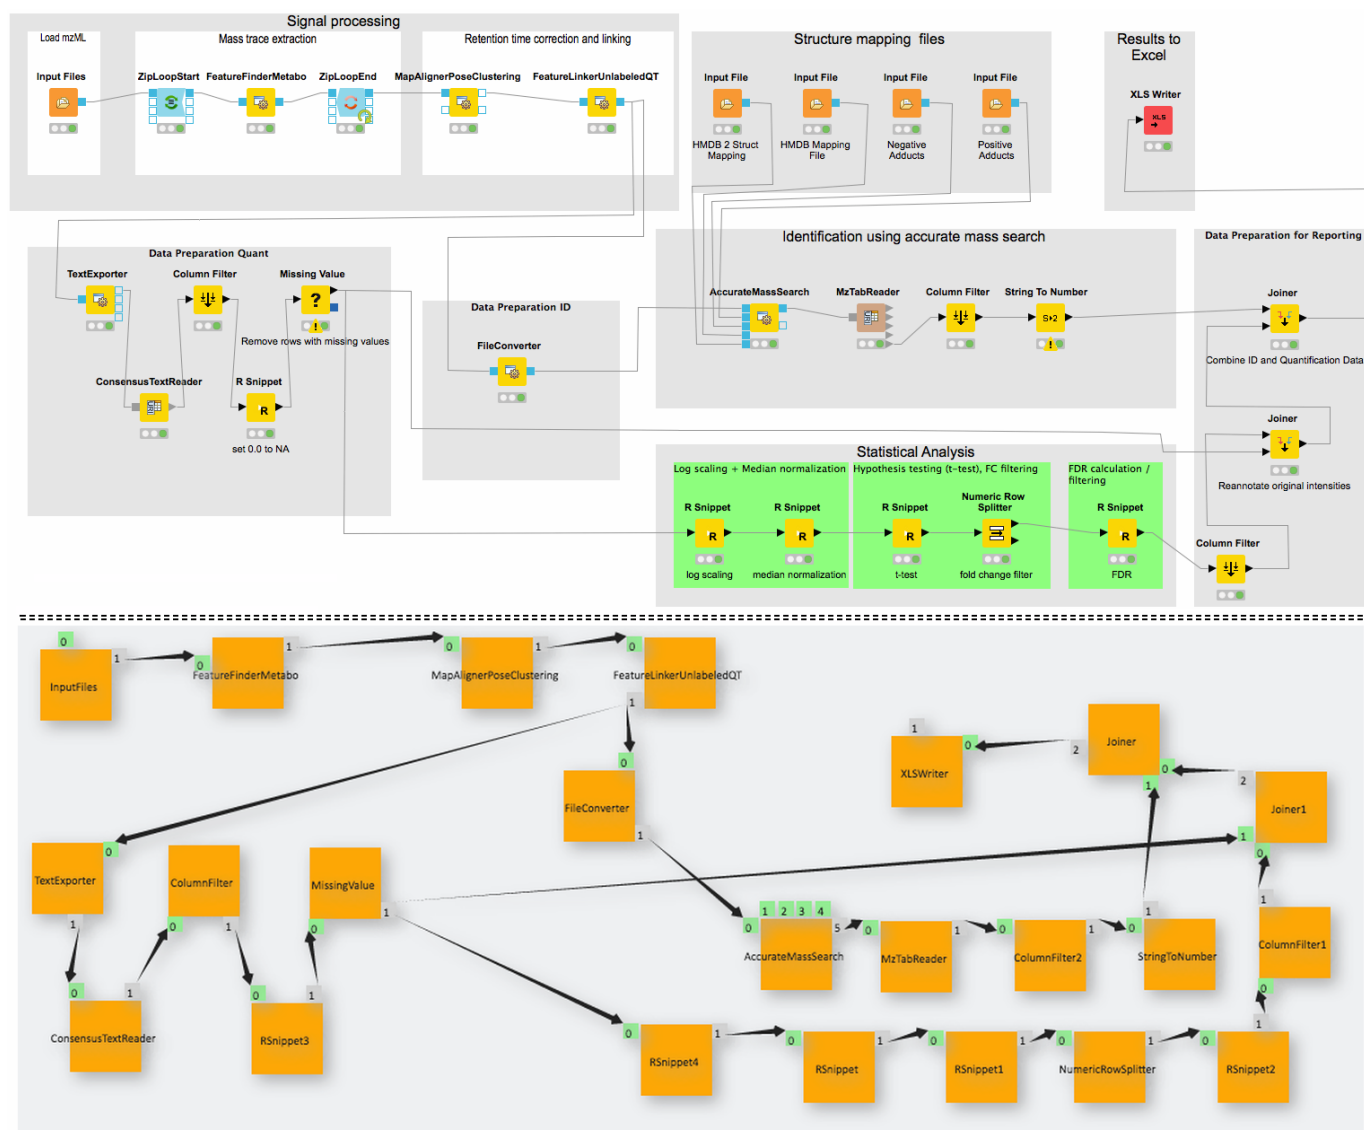


Fig. 9. **Metabolomics biomarker discovery workflow both in the KNIME Analytics Platform and WS-PGRADE.** Top: Workflow implemented in the KNIME Analytics Platform, OpenMS [12] was imported using GKN. Bottom: Workflow generated by KNIME2gUSE imported into WS-PGRADE (workflow slightly edited for visual clarity). Even if some elements are missing after conversion (e.g., *ZipLoopStart*, *ZipLoopEnd*, *Input File* nodes), both versions have the same *abstract* layer. This is due to the difference in implementation of logically equivalent constructs in WS-PGRADE and the KNIME Analytics Platform, such as *parameter sweep*.

parative metabolite concentrations is label-free quantification. The independence from chemical labels allows the direct comparison of small molecules across an arbitrary number of samples. As a consequence, the need to evaluate hundreds of gigabyte-sized samples in concert is already common. Numbers and sizes of concurrently evaluated samples are steadily increasing, emphasizing the necessity for distributed computing.

We provide an example workflow for metabolomics biomarker discovery using OpenMS [12] for mass spectrometry algorithms as well as various *native KNIME Nodes* (including nodes for the R scripting language). The KNIME workflow and its converted WS-PGRADE version are shown in Figure

9. We assume some initial preparations were performed prior to the execution of the workflow, namely, conversion from closed mass spectrometer vendor formats to the open *mzML* format and data reduction by means of peak picking, which could also be implemented in KNIME via OpenMS [12] tools.

Using a detection method for so-called small molecules [19], we adapted a label-free quantification pipeline [20]. The quantification part of our biomarker discovery workflow consists of sample specific feature detection (i.e., finding the convex hulls and respective centroids of analyte mass traces) followed by temporal alignment of samples and the quantification of corresponding features across samples.

Downstream small molecule identification was done via mass-based search in the Human Metabolome Database. Included sample normalization allows for comparison of analyte abundances across samples. Analytes whose abundances vary significantly after false discovery rate correction are annotated with the mass-based identifications and exported to a Microsoft Excel Spreadsheet (XLS format).

III. FUTURE WORK

The KNIME Analytics Platform features *Metanodes* encapsulating complete workflows. We would like to extend KNIME2gUSE to support their conversion. Furthermore, seeing that considerable effort has been put into creating platform-independent workflow representation formats, we would like to add IWIR and YAWL file generation to KNIME2gUSE. We would also like to extend our converter to support other workflow engines, such as Galaxy.

IV. CONCLUSION

Workflows assist reproducibility and minimize time spent validating research by reducing analysis complexity. There are currently several workflow engines with user-friendly interfaces that support remote execution of workflows. However, we feel that their scalability and support of major resource managers is still lacking. In contrast, HPC infrastructures and their resource managers rarely support the execution and control of workflows. As a consequence, HPC users often require programming skills to handle the channeling of data as well as to submit, monitor and control the respective computing jobs.

We present our efforts to support workflow export from the KNIME Analytics Platform to WS-PGRADE, identified challenges for both node and workflow conversion and detailed our solutions. KNIME offers remote workflow execution, but it is a royalty-based solution and support of DCIs is limited—an aspect in which WS-PGRADE excels. KNIME2gUSE brings together a user-friendly and intuitive workflow engine for personal computers together with a scalable HPC workflow platform that interacts with several DCIs.

We thus provide the individual advantages of both engines without any of their shortcomings. Overall, our methods decrease time spent designing workflows and troubleshooting conversion for different workflow engines.

ACKNOWLEDGMENT

The authors would like to thank Bernd Wiswedel, Thorsten Meinel, Patrick Winter and Michael Berthold for their support, patience and help in developing the KNIME2gUSE extension.

This work was supported by the German Network for Bioinformatics Infrastructure (Deutsches Netzwerk für Bioinformatik-Infrastruktur, de.NBI).

REFERENCES

- [1] C. S. Greene, J. Tan, M. Ung, J. H. Moore, and C. Cheng, "Big data bioinformatics." *Journal of cellular physiology*, vol. 229, no. 12, pp. 1896–900, Dec. 2014.
- [2] M. McNutt, "Reproducibility." *Science (New York, N.Y.)*, vol. 343, no. 6168, p. 229, 2014.
- [3] "Trouble at the lab," *The Economist*, oct 2013.
- [4] M. Baker, "Over half of psychology studies fail reproducibility test," *Nature*, Aug. 2015.
- [5] L. de la Garza, J. Veit, A. Szolek, M. Röttig, S. Aiche, S. Gesing, K. Reinert, and O. Kohlbacher, "From the desktop to the grid: scalable bioinformatics via workflow conversion," *BMC Bioinformatics*, vol. 17, no. 1, pp. 1–12, 2016.
- [6] L. de la Garza, J. Krüger, C. Schärfe, M. Röttig, S. Aiche, K. Reinert, and O. Kohlbacher, "From the desktop to the grid: conversion of knime workflows to guse." in *IWSG*, 2013.
- [7] K. Plankensteiner, J. Montagnat, and R. Prodan, "IWIR: A Language Enabling Portability Across Grid Workflow Systems," in *SIGMOD Rec.*, vol. 34, no. 3, 2011, pp. 97–106.
- [8] W. van der Aalst and A. ter Hofstede, "YAWL: yet another workflow language," *Information Systems*, vol. 30, no. 4, pp. 245–275, Jun. 2005.
- [9] G. Terstyanszky, T. Kukla, T. Kiss, P. Kacsuk, A. Balasko, and Z. Farkas, "Enabling scientific workflow sharing through coarse-grained interoperability," *Future Generation Computer Systems*, vol. 37, pp. 46–59, 2014.
- [10] P. Kacsuk, Z. Farkas, M. Kozlovsky, G. Hermann, A. Balasko, K. Karoczkai, and I. Marton, "WS-PGRADE/gUSE Generic DCI Gateway Framework for a Large Variety of User Communities," *Journal of Grid Computing*, vol. 10, no. 4, pp. 601–630, 2012.
- [11] M. R. Berthold, N. Cebren, F. Dill, T. R. Gabriel, T. Kötter, T. Meinel, P. Ohl, K. Thiel, and B. Wiswedel, "KNIME - the Konstanz information miner," *ACM SIGKDD Explorations Newsletter*, vol. 11, no. 1, p. 26, Nov. 2009.
- [12] M. Sturm, A. Bertsch, C. Gröpl, A. Hildebrandt, R. Hussong, E. Lange, N. Pfeifer, O. Schulz-Trieglaff, A. Zerck, K. Reinert, and O. Kohlbacher, "OpenMS - an open-source software framework for mass spectrometry." *BMC Bioinformatics*, vol. 9, p. 163, 2008.
- [13] A. Döring, D. Weese, T. Rausch, and K. Reinert, "SeqAn an efficient, generic C++ library for sequence analysis." *BMC bioinformatics*, vol. 9, no. 1, p. 11, Jan. 2008.
- [14] A. Hildebrandt, A. K. Dehof, A. Rurainski, A. Bertsch, M. Schumann, N. C. Toussaint, A. Moll, D. Stöckel, S. Nickels, S. C. Mueller, H.-P. Lenhof, and O. Kohlbacher, "BALL—biochemical algorithms library 1.3." *BMC bioinformatics*, vol. 11, p. 531, 2010.
- [15] E. L. Schymanski, H. P. Singer, P. Longrée, M. Loos, M. Ruff, M. A. Stravs, C. Ripollés Vidal, and J. Hollender, "Strategies to characterize polar organic contamination in wastewater: exploring the capability of high resolution mass spectrometry." *Environmental science & technology*, vol. 48, no. 3, pp. 1811–8, Jan. 2014.
- [16] M. Sugimoto, D. T. Wong, A. Hirayama, T. Soga, and M. Tomita, "Capillary electrophoresis mass spectrometry-based saliva metabolomics identified oral, breast and pancreatic cancer-specific profiles." *Metabolomics : Official journal of the Metabolomic Society*, vol. 6, no. 1, pp. 78–95, mar 2010.
- [17] C. Denkert, J. Budczies, T. Kind, W. Weichert, P. Tablack, J. Sehoul, S. Niesporek, D. Könsen, M. Dietel, and O. Fiehn, "Mass spectrometry-based metabolic profiling reveals different metabolite patterns in invasive ovarian carcinomas and ovarian borderline tumors." *Cancer research*, vol. 66, no. 22, pp. 10795–804, Nov. 2006.
- [18] J. S. Hansen, X. Zhao, M. Irmeler, X. Liu, M. Hoene, M. Scheler, Y. Li, J. Beckers, M. Hrab? de Angelis, H.-U. Häring, B. K. Pedersen, R. Lehmann, G. Xu, P. Plomgaard, and C. Weigert, "Type 2 diabetes alters metabolic and transcriptional signatures of glucose and amino acid metabolism during exercise and recovery." *Diabetologia*, vol. 58, no. 8, pp. 1845–54, Aug. 2015.
- [19] E. Kenar, H. Franken, S. Forcisi, K. Wörmann, H.-U. Häring, R. Lehmann, P. Schmitt-Kopplin, A. Zell, and O. Kohlbacher, "Automated label-free quantification of metabolites from liquid chromatography-mass spectrometry data." *Molecular & cellular proteomics : MCP*, vol. 13, no. 1, pp. 348–59, jan 2014.
- [20] H. Weisser, S. Nahnsen, J. Grossmann, L. Nilse, A. Quandt, H. Brauer, M. Sturm, E. Kenar, O. Kohlbacher, R. Aebersold, and L. Malmström, "An automated pipeline for high-throughput label-free quantitative proteomics." *Journal of proteome research*, vol. 12, no. 4, pp. 1628–44, Apr. 2013.