

# PSI/J: A Portable Interface for Submitting, Monitoring, and Managing Jobs

Mihael Hategan-Marandiu<sup>1,4</sup>, Andre Merzky<sup>3</sup>, Nicholson Collier<sup>1,4</sup>, Ketan Maheshwari<sup>6</sup>, Jonathan Ozik<sup>1,4</sup>, Matteo Turilli<sup>3,5</sup>, Andreas Wilke<sup>1,4</sup>, Justin M. Wozniak<sup>4</sup>, Kyle Chard<sup>1,4</sup>, Ian Foster<sup>1,4</sup>, Rafael Ferreira da Silva<sup>6</sup>, Shantenu Jha<sup>3,5</sup>, Daniel Laney<sup>2</sup>

<sup>1</sup>University of Chicago, Chicago, IL, USA <sup>2</sup>Lawrence Livermore National Laboratory, Livermore, CA, USA

<sup>3</sup>Brookhaven National Laboratory, Upton, NY, USA; <sup>4</sup>Argonne National Laboratory, Lemont, IL, USA

<sup>5</sup>Rutgers, New Brunswick, NJ, USA; <sup>6</sup>Oak Ridge National Laboratory, Oak Ridge, TN, USA

**Abstract**—It is generally desirable for high-performance computing (HPC) applications to be portable between HPC systems, for example to make use of more performant hardware, make effective use of allocations, and to co-locate compute jobs with large datasets. Unfortunately, moving scientific applications between HPC systems is challenging for various reasons, most notably that HPC systems have different HPC schedulers. We introduce PSI/J, a job management abstraction API intended to simplify the construction of software components and applications that are portable over various HPC scheduler implementations. We argue that such a system is both necessary and that no viable alternative currently exists. We analyze similar notable APIs and attempt to determine the factors that influenced their evolution and adoption by the HPC community. We base the design of PSI/J on that analysis. We describe how PSI/J has been integrated in three workflow systems and one application, and also show via experiments that PSI/J imposes minimal overhead.

## I. INTRODUCTION

A portable HPC application, intended to submit and manage jobs across multiple HPC systems, each with a potentially different *local resource manager* (LRM), such as Slurm and PBS, must be able to generate job descriptions for each LRM type and know how to interact with each LRM’s tools (e.g., qsub, qstat) or APIs. The semantics of these operations are, to a large extent, the same across the various LRMs and the differences are to be found in the details of how LRMs expose their functionality. Consequently, reasonable software architecture practices dictate that a portable HPC application be designed such that all the logic that is not specific to a given LRM be handled on an abstract level and the relevant translation to specific LRMs be done as low as possible in the application’s component hierarchy. That is, under reasonable assumptions, every application intended to be portable across multiple LRMs must have a component that translates abstract job information into specific LRM scripts, commands, or API calls. For convenience, we will call such a component a Job Abstraction API (or, in short, JAAPI).

Without a common JAAPI, portable HPC applications have no choice but to implement their own custom solution. This leads to unnecessary redundancy, which could potentially be overlooked if the effort required to implement and maintain a JAAPI was negligible. However, evidence gathered from known past JAAPIs shows no indication of triviality in either design, implementation, or maintenance. On the contrary, it

appears that the complexity of JAAPIs is frequently underestimated. This costly redundancy underlines the need for both reusable JAAPIs as well as an understanding of how their adoption can be facilitated and what prevents it.

The previous statements are abstract and do not do justice to just how prevalent this problem is. To provide some context, nearly every project that is incidental to our team’s work (Balsam [1], Maestro [2], Parsl [3], RADICAL-Pilot [4], Swift/T [5]) has, separately, needed to design and implement a JAAPI, with little or no coordination but with very similar functionality, as shown in Table I. And every JAAPI designed and implemented as part of such a project represents a hidden cost that must be funded but is rarely acknowledged as more than part of the normal cost of writing higher level HPC tools [6].

A number of JAAPIs are known to have emerged over time. Of those that are still actively maintained or developed, none show evidence of widespread adoption. This is in direct contradiction with the clear need for a JAAPI. In this paper, we attempt to examine the causes of this contradiction and use our conclusions to provide insight and possible suggestions for facilitating the adoption of reusable JAAPIs as well as presenting recommendations for aiding in the survivability of JAAPIs. Informed by our analysis, we propose PSI/J: an open, language-agnostic, and minimalistic JAAPI that is meant to satisfy most of the needs of portable HPC applications while also providing pass-through capability for more advanced usage scenarios.

We discuss past (and some present) similar efforts in an attempt to gain some insight into the reasons behind their limited availability and/or adoption in Section II. With the wisdom gained in Section II, we outline the design of the PSI/J API in Section III as well as the relevant details of the Python binding [7] and the testing infrastructure in Section IV. We proceed with a number of examples of applications that are using PSI/J in Section V and report performance numbers meant to illustrate the overhead that our PSI/J Python implementation imposes in Section VI. We conclude with a discussion of the potential opportunities and obstacles that lie ahead in Section VII.

Application	Schedulers Supported	Language	Remote
Balsam [1]	Local, Cobalt, LSF, PBSPro, Slurm	Python	No
Maestro [2]	Local, Slurm, LSF, Flux	Python	No
Parsl [3]	Cobalt, HTCondor, LSF, PBSPro, SGE, Slurm, Torque, AWS, Azure, Google Cloud, Kubernetes	Python	Yes
RADICAL-Pilot [4]	Cobalt, Condor, CPI, LoadLeveler, LSF, PBS, PBSPro, SGE, Slurm, Torque	Python	Yes
Swift/T [5]	Cobalt, AWS, LSF, PBS, SGE, Slurm	Java, C, Tcl	No

TABLE I  
SELECT PORTABLE HPC APPLICATIONS AND CHARACTERISTICS OF THEIR EMBEDDED JAAPIs.

JAAPI	Type	Schedulers Supported	A/Synchronous	Languages	Remote	Last Updated
DRMAA [8]	Spec. & Impl.	SGE, HTCondor, PBS, LSF, GridWay, LoadLeveler, Slurm, Unicore	Async.	C, Perl, Python, Ruby, Java, Go, Erlang	Varies	Varies, 2017 (PBS), 2018 (Python)
Globus GRAM [9]	Impl.	PBS, SGE, LSF, Condor	Both	C, Python, Java	Yes	2018
Unicore [10]	Impl.	PBS, LoadLeveler, LSF, Slurm	Both	Java	Yes	2023
HTCondor [11]	Impl.	PBS, LSF, SGE, AWS	Both	C, Python	Yes	2023
SAGA [12]	Spec. & Impl.	SSH, HTCondor, PBS, SGE, Slurm, LSF, LoadLeveler, EC2	Both	Python, C++, Java	Yes	2023 (Python)
Java CoG API	Impl.	PBS, Cobalt, HTCondor, SGE, Slurm, LSF, AWS, GCE, Globus GRAM, SSH	Both	Java	Yes	2017

TABLE II  
A LIST OF WELL KNOWN JAAPIs. UNICORE, HTCONDOR, AND RADICAL-SAGA ARE BEING ACTIVELY MAINTAINED AT THE TIME OF THIS WRITING.

## II. BACKGROUND AND RELATED WORK

In this section we describe both previous and current JAAPIs, discuss factors that influence the adoption and sustainability of JAAPI implementations, and analyze characteristics necessary to create a sustainable JAAPI.

### A. Previous and current efforts

The idea of a unified JAAPI is far from new. During the golden era of Grid Computing, there were a number of high profile JAAPIs. Perhaps one of the most notable was the Globus Toolkit [13] and, specifically, Globus GRAM [9]. While now end-of-life, the Globus Toolkit is, directly and indirectly, a major source of inspiration for PSI/J.

Globus GRAM offers a set of services for submitting and monitoring jobs on diverse LRMs. Implemented as a set of services, and with a custom protocol and resource specification, GRAM offers a complete ecosystem for job management. GRAM is designed for Grid environments, and thus focuses on remote submission via integration of components for grid security, file transfer, and enforcement of authorization and access policies. GRAM is designed for deployment in multi-user environments and thus requires administrator privileges for deployment. At its peak, GRAM was deployed widely both in the US and internationally as an interface to different LRMs on many HPC clusters, including most of the leading supercomputers. Some of the ideas included in Globus con-

tinue today in the cloud-hosted Globus service [14] and in the Globus Compute service [15].

After the first generation of Grid Computing tools, an effort emerged to build community consensus on Grid protocols and tools. This effort, the Open Grid Forum (OGF), was the birthplace of another major JAAPI: the Simple API for Grid Applications (SAGA) [12]. An OGF standard, SAGA was implemented in C++ [16], Java [17] and, lastly, in Python via RADICAL-SAGA [18]. RADICAL-SAGA is still maintained and used by RADICAL-Pilot [4] and a few other distributed computing projects [18] to interface with LRMs. SAGA and RADICAL-SAGA have directly inspired the design of PSI/J. OGF promoted a number of other notable efforts, of which we mention OGSA BES [19] and JSDL [20]. In the European Union, the UNICORE project [10] provided a JAAPI in the form of the UNICORE Abstract Job Object.

At about the same time that SAGA was being designed, the Java CoG Kit [21], which started as a pure Java implementation of the Globus client suite, developed an “abstraction API” [22] for various job execution implementations, including Globus GRAM versions 2 and 3, SSH, LRMs, etc. This abstraction API would later be used by the Swift/K [23] workflow system both directly and indirectly as part of the Coaster pilot job system [24]. Along with SAGA, it remains one of the main sources of inspiration for PSI/J.

Another product of the OGF is the Distributed Resource

Management Application API [8] (DRMAA), which is an API that is meant to be exposed by HPC schedulers such that applications can submit and manage jobs through the DRMAA API rather than by interacting directly with the job scheduler. Applications are expected to link with the local DRMAA shared library and gain access to the local scheduler functionality. Implementations for DRMAA exist in C for Slurm, PBS/Torque, LSF. Wrappers for the C libraries can be found for Python, Perl, Ruby, Erlang, and Go. Since the peak days of the OGF, the HPC community seems to have moved forward and some of the C DRMAA implementations for certain LRMs have not been updated in more than 8 years. Nonetheless, DRMAA appears to be the most widely used JAAPI in existence.

### B. The JAAPI condition

We now consider the factors that can influence the adoption and survivability of a JAAPI. We can classify relevant factors into two categories: *intrinsic* to a particular design and/or implementation or *extrinsic*. Intrinsic factors are factors that are almost entirely under the control of the team or organization enacting the design or implementation of a JAAPI, whereas extrinsic factors are a result of the circumstances under which the enactment takes place and for which the team or organization can at most produce mitigating strategies. This distinction between intrinsic and extrinsic factors can sometimes be blurry, but, in most cases, a dominant direction can be distinguished. In what follows, we will attempt to enumerate and discuss prominent factors. We caution the reader that this discussion is not meant as authoritative, but as a starting point for future discourse and, in the case of PSI/J, as the lens through which we perceive the suitability of a JAAPI solution. We will discuss extrinsic factors first, since they influence the intrinsic factors: complexity of the problem, funding structure, and the research nature of JAAPIs. We follow with the intrinsic factors: complexity of design, design inflexibility, and implementation quality.

1) *Problem complexity*: Designing and implementing a JAAPI is made complex by two characteristics of HPC systems: the heterogeneity of HPC sites and the stringent security associated with accessing most HPC sites. Even when different sites use the same LRM (e.g., Slurm), experience shows that there is significant variability in how LRMs are configured across sites. The assurance that a JAAPI functions properly on an HPC site rarely arises from it working on a similar site and is more often contingent on the JAAPI being tested on the specific HPC site. Consequently, ensuring that a JAAPI works in general requires that it be tested on a significant subset of the target set of HPC sites. But because of the stringent security policies associated with HPC sites, it is virtually impossible for a small development team to gain access to such a significant set of HPC sites.

We arrive at an odd intersection: nearly every portable HPC application that manages jobs needs to have a component that is nearly impossible to test adequately. Furthermore, the inability to test a component that sits at the base of the

component hierarchy of an application has consequences for the entire application. While speculative, it does not stretch the imagination to picture why most HPC workflow systems, written with ad-hoc JAAPIs, fail to get traction outside the organizations in which they were created.

2) *Funding structure*: One of the major barriers to adoption of a JAAPI is the natural desire of adopters to have some confidence that the JAAPI will continue to be supported and maintained for the lifetime of the project that the JAAPI is meant to be integrated into. As it is undesirable to foresee the demise of one's own project, it follows that JAAPIs are expected to be supported for the foreseeable future in order for them to be a compelling target for adoption. Unfortunately, JAAPIs are mostly the products of research funding, which is, by nature, limited to a few years. Even if new funds are received, they must be justified in terms of an active research goal. This situation is problematic and leaves no room to the possibility of funding needed infrastructure. It is akin to requiring that highways be re-built with new materials or receive significant technological improvements every few years in order for them to continue to be in a usable state. Unfortunately, this results in higher overall costs, since lower quality JAAPIs now become a hidden cost of all portable HPC applications. Pushing the highway metaphor a bit further, it is like governments opting to fund multiple parallel roads for use by individual companies in order to avoid the costs of maintaining a single common road.

3) *Research nature of JAAPIs*: Traditionally, the largest customers of HPC resources have been scientific research projects. The biggest supercomputers are overwhelmingly associated with academic or research institutions. The ecosystem surrounding HPC resources and applications is, unsurprisingly, also tightly coupled with the research world. One of the main outputs in the research world are academic publications. Infrastructure software, such as JAAPIs, should have stability as one of their main goals. Stability implies changes that are limited to marginal improvements and the addressing of defects, changes that are rarely interesting enough to align with notable novelty that would result in research output. This leads to a conflict that, simultaneously, pushes JAAPIs towards stability and towards re-inventing themselves, which can add unnecessary complexity over time. Further, it encourages developers to focus on novelty, exploring new, yet untested, approaches necessary for publication.

4) *Complexity of design*: Complexity is perhaps one of the most visible aspects of a JAAPI solution. An overly complex design leads to difficulties in implementation, maintenance, and deployment. A complex design can arise in many ways. An often encountered scenario is that in which a large collaboration attempts to create an all-encompassing design. This leads to the inclusion of legitimate but rarely needed features that a compliant implementation is then forced to deliver. For example, the SAGA specification mandates that certain classes be implemented in versions that allow both synchronous and asynchronous invocation of their methods, even when such methods implement fast operations. Complexity of design is

not limited to a JAAPI specification and can apply equally to implementations.

5) *Design inflexibility*: Design inflexibility is the inability of a JAAPI design to adapt to changes in the software ecosystem used to enact implementations. It was typical of the JAAPIs that were products of the OGF and which included remote job submission to do so using an extensive stack of XML technologies which have since experienced a significant decline in adoption for new implementations. A subcategory of design inflexibility is programming language specificity, which explicitly or indirectly biases a design towards a specific programming language that may be subject to a decline in popularity.

6) *Implementation quality*: All JAAPI implementations are, in the end, software products. And, like other software products, they get consciously or subconsciously evaluated with respect to what potential users expect from a software product. Along with functional metrics, such as lack of errors, weight is also given to non-functional aspects such as code quality, presentation, documentation, ease of use, support, and so on. Ensuring that a well rounded software product is put forward requires team discipline that is often difficult to find in the research world where it is not uncommon for teams to consist of collaborations that are put together as part of a single funded project and where no clear organizational structure exists except PIs and not-PIs.

### C. Discussion

We now turn towards an analysis of known JAAPIs as seen through the prism of the earlier considerations. We argue that a successful JAAPI must exhibit most of the desirable qualities mentioned previously while avoiding or suitably compensating for the difficulties. We, again, caution the reader: this analysis suffers heavily from survival bias as we only discuss JAAPIs that has managed to achieve sufficient prominence.

Perhaps the most successful JAAPI still in existence is DRMAA. The DRMAA specification is focused on a single problem and avoids many of the complexities associated with JAAPIs by only mandating local access to LRMs. While the DRMAA specification allows for implementations in multiple languages, it is primarily geared towards the C language, which is a form of design inflexibility. Libraries exist for all major LRMs, which, in large part, is a consequence of the DRMAA specification being written by representatives from LRM suppliers, some of which provide DRMAA libraries with LRM releases. The existence of the C libraries allow wrappers for other languages to be written without needing a full re-implementation. Nonetheless, the choice of most LRM specific DRMAA libraries to use proprietary APIs to communicate with the scheduler makes community support difficult, which is necessary without a continued source of funding or commitment from LRM suppliers. It also makes it challenging to maintain multiple LRM adapters as part of a single effort, which implies that some internal routines are duplicated across adapters and come with separate and often dissimilar characteristics. This appears to be a consequence

of a design that fails to see beyond the implementation of DRMAA for a single LRM.

Still under active support, RADICAL-SAGA is an implementation of the SAGA specification. The SAGA specification is an extensive document that encompasses job management, file management, data stream management, remote procedure calls, and more. RADICAL-SAGA suffers from the rigidity imposed by having to adhere to this complex standard. For example, compliance requires a relatively complex implementation that limits portability and results in implementation quality issues, such as difficulty in tracing errors. In other words, the likelihood that a third party will implement SAGA and maintain the resulting library is significantly reduced. Indeed, the team maintaining RADICAL-SAGA, who are co-authors to this paper, have contributed significantly to the SAGA specification and maintain its Python implementation. Furthermore, RADICAL-SAGA's Grid Computing origins demanded a vast array of capabilities that are currently available through third party libraries. .

Globus GRAM was an integral part of the Globus Toolkit. It allowed remote job submission when most people still accessed the Internet by dial-up. Implementing a remote-capable JAAPI required solving a number of security and networking problems for which no off-the-shelf solutions were readily available. Consequently, a great deal of complexity went into writing and maintaining it, which had ripple effects for maintenance costs. Additionally, in normal usage scenarios, the Globus GRAM server required a non-trivial amount of effort from HPC cluster administrators to be installed and supported. It was also initially closed-source, which contributed to the improbability of being supported by the community in the event of a funding shortage. Subsequent versions addressed the issue of community openness by releasing specifications through the OGF.

HTCondor started out as a cycle-stealing batch scheduling system, but has since grown into a system supporting the routing of jobs to other LRMs as part of a federated HTCondor pool. Various large Condor pools exist, including the Open Science Grid [25], but it is rare for leading HPC resources to be part of a HTCondor pool. HTCondor is a large and complex project with over 1.5k C++ source files and with many contributors. As a consequence, installing and configuring a local HTCondor pool that can act as a JAAPI is rarely the primary way in which portable HPC applications implement LRM abstraction.

The Java CoG Kit abstraction API was an implementation-only JAAPI that included support for job, file transfers, and remote filesystem access. While the Java GoG Kit abstraction API was, at first, a standalone product under the umbrella of the Java GoG Kit, it reached maturity almost exclusively as part of the Swift/K system at a time when Grid Computing was being slowly phased out. Combined with the fact that it was a Java-only implementation and that it suffered from implementation quality issues, such as poor documentation and support, it failed to gain widespread adoption.

The UNICORE (Uniform Interface to Computing Re-

sources) project developed an end-to-end middleware to enable users to execute many-task applications on Grid computing infrastructures. Evolved into a solution supporting general-purpose distributed computing, it requires root privileges on a cluster login nodes in order to deploy and run its server components on HPC platforms. Further, integration with each HPC cluster requires complex and specific customization. This imposes rigid constraints in terms of security, portability, and resource requirements and is an example of design complexity.

### III. DESIGN

In this section we describe the goals that motivated PSI/J, the resulting specification, and the reference implementation.

#### A. Design Goals

The PSI/J API [26] is designed as a minimal interface to submitting and monitoring jobs and their execution state. The need for minimalism is informed, in part, by the observations in Section II that complexity is unlikely to lead to a successful JAAPI solution in the long term. That is coupled with the observation that many custom JAAPIs (e.g., Parsl [3], Swift/K [23], Swift/T [5], Balsam, RADICAL-Pilot) are focused almost exclusively on submitting and monitoring jobs with no further adornments. Therefore we believe that maintaining an orthogonal and therefore simple API is fundamental in facilitating both use and implementation.

The PSI/J API is also designed to allow scalable implementations where scalability is targeted both in the number of handled jobs and in the rate of job submission. Suggestions for implementers relevant to scalable implementations are provided in the specification document. Our reference Python PSI/J implementation makes full use of the scalability features of the API. Specifically, the API is asynchronous in order to support threadless use. The choice of asynchronous API is also motivated by the fact that one can easily transform an asynchronous API into a synchronous one, without incurring a performance penalty. However, the opposite is not true: adding an asynchronous layer on top of a synchronous API does not remove the one-to-one correspondence between threads and jobs at the synchronous API level.

In PSI/J, the simplicity of the API is favored over that of the implementation if a reasonable implementation choice exists for a given design goal [27]. This is motivated by the fact that implementation complexity can be addressed with reusable solutions that multiple implementations can share, whereas API complexity translates into a pervasive overhead for users of the API. For example, bulk operations can, in certain cases, improve scalability of an implementation. Bulk operations are versions of API operations that act on multiple items as opposed to single ones. A bulk *submit* method would take a list of jobs as arguments and submit them all in one call to the underlying LRM under the assumption that the alternative of making multiple calls to the LRM introduces a large combined overhead. However, a simple time windowing function can be used to aggregate individual job submissions

and transparently route the resulting list of jobs to a bulk LRM submit call, if such a call is available. This would retain the simplicity of the API while also allowing the implementation to be efficient and is the favored choice. Certain caveats of this approach should, however, be noted. Time-based clustering of jobs is insufficient in determining jobs that are also related in their various properties and may be unsuitable in creating job arrays as supported by various LRMs. Such support may be added to PSI/J in the future.

A similar problem to bulk submission is that of bulk job status querying. Invoking *qstat* commands individually for each job can result in poor scaling as the number of actively managed jobs increase. The PSI/J specification [26] mandates that implementations query the status of jobs in bulk in order to avoid this issue.

Last but not least, the PSI/J specification is geared towards allowing implementations to live in user space as well as concurrently being installed as system libraries. This can push the burden of supporting a PSI/J implementation deployment away from system administrators and give users the flexibility to address deployment problems as needed.

#### B. The Structure of the API

The PSI/J specification is organized into three basic layers, depending on the level of functionality that it describes. The *local layer* (see Figure 1) defines the API needed to interact with job schedulers locally. That is, the location of the job scheduler is implicit and assumed to be on the same machine as the one on which the client application is running. The *remote layer* (see Figure 2) defines additional API elements needed to submit jobs to a remote scheduler, running on a different machine than the one on which client code is running. The *nested layer* (Layer 2) is meant to add API elements needed to interact with pilot job implementations. At the time of the writing of this paper, only the local layer of the PSI/J specification is available publicly and work is underway in drafting the remote layer. Nonetheless, the bulk of the specification rests with the local layer, with the other two layers containing only incremental updates needed to describe endpoints, credentials, and service configuration.




Fig. 1. Illustration of the local layer of PSI/J.

The PSI/J layers are related by an inclusive relationship, in the sense that a nested layer implementation is expected to also contain all the API elements of the remote and local layers. However, one can also identify a functional relationship between layers. For example, the remote layer functionality can be implemented by adding remote invocation capabilities (see Figure 2) to a local layer implementation, whereas a

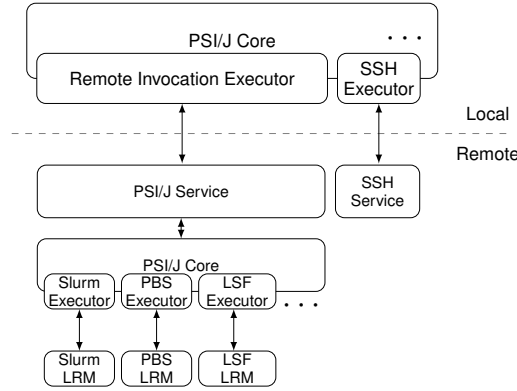


Fig. 2. Intended usage scenario for the remote layer of PSI/J.

nested layer implementation requires a mechanism to submit the pilot jobs to a HPC scheduler, which can be achieved with either a local or remote layer. The remote layer is also suitable for interfacing with libraries that natively implement remote job execution, such as SSH.

In the remote layer, multiple remote schedulers on multiple HPC schedulers are meant to be accessible concurrently, from the same client process. Similarly, in the local layer, it can be desirable to be able to submit simple test jobs that can be run locally, using a forked process. To support this scenario, PSI/J-Python adopts a multiple-dispatch mechanism, in which bindings to underlying job execution mechanisms can coexist and be used concurrently. Such bindings are called “executors”. In this sense, it differs fundamentally from most DRMAA implementations, which require that the client executable be dynamically or statically linked in order to switch to an alternate LRM.

### C. Implementations

As previously stated, the implementations of PSI/J are meant to be usable as user space components. This is intended to overcome issues that stem from delays in updates of system software on HPC clusters. Cluster systems software, especially on HPC systems that are part of larger organizations, are known to go through lengthy processes of approvals for updates on software packages that are deployed system-wide. This can result in large delays between the reporting of an issue with such software and the availability of a fix on a given HPC cluster.

We provide a reference Python implementation of PSI/J. The Python language was chosen due to its widespread use among higher level tools within the scientific community. However, in order to facilitate portability, most LRM specific functionality is described as textual templates using a templating library with binding in a large number of languages. The PSI/J specification and implementations can be logically split into two main parts: the *core classes*, which are independent of the underlying execution mechanism (e.g., LRM), and the *executors and launchers*, which implement LRM and cluster specific functionality. Executors are entities that know

how to communicate with specific LRMs, whereas launchers describe the command used to launch multi-node jobs once the job resources are allocated. The core classes are used to build and manage jobs on an abstract level, independent of the underlying executor. Once a job, with all the necessary information, is built, one obtains an instance of an executor and uses it to submit the job. We show a simple example of this process for the Python binding of PSI/J in Figure 3.

```
job = Job(
    spec=JobSpec(
        executable='/opt/cps/bin/NOARCH.x',
        arguments=['-qmp-geom', '8', '4', '4', '4',
                  'do_arg.vml', 'evo_arg.vml',
                  'eig_arg.vml', '0.00', 'Overlap'],
        stdout_path=cwd + '/eig.out',
        stderr_path=cwd + '/eig.err',
        resources=ResourceSpecV1(process_count=512),
        launcher='srun'
    )
)
ex = JobExecutor.get_instance('slurm')
ex.submit(job)
```

Fig. 3. Simple example for submitting a job in PSI/J Python with boilerplate removed.

The Python implementation of PSI/J uses a dynamic plugin discovery mechanism which allows a PSI/J core to detect executor and launcher implementations that are installed in different places from the PSI/J core. This allows for various scenarios, such as a stable system provided core using a user customized executor implementation or a user installed core using a system provided executor implementation (a scenario somewhat similar to that of DRMAA). Executors for Slurm, PBSPro, LSF, Cobalt, and Flux [28] are provided with the current version of the reference implementation. A “local” executor that runs jobs using a simple fork mechanism is also provided. Launcher implementations are provided for all LRM specific launchers (*srun*, *aprun*, *jsrun*, etc.) as well as for generic launchers, such as *mpirun*. The dynamic plugin system allows PSI/J Python to be extensible with new executors and launchers without necessarily requiring that the additional executors or launchers be part of the PSI/J Python code base.

We have opted for the executors provided by the reference PSI/J Python implementation to use publicly available LRM interfaces, which, in most cases, consist of well known commands, such as *qsub* and *qstat*. This is a deliberate choice that assumes that established public LRM interfaces are less likely to change and lead to incompatibilities than proprietary ones. This is in contrast to many DRMAA implementations which use proprietary APIs to communicate with LRMs. The choice made by the PSI/J Python reference implementation does not, however, preclude one from writing executors using proprietary APIs.

## IV. DEPLOYMENT

One of the major problems that PSI/J is meant to address is the difficulty involved in adequately testing JAAPIs on a

widespread set of HPC resources. We can assert that it is virtually impossible for a single research team to gain and maintain access to a sufficient number of resources in order to cover a satisfactory range of configurations. Traditionally, testing on resources that are not accessible to the JAAPI developers tends to be done by their users in the process of using the higher level system that employs the JAAPI. We propose to improve on this scenario by allowing any user with access to a HPC cluster (or any computing resource), when authorized to do so, to set up daily runs of a PSI/J implementation test suite and automatically report the results to a centralized test results aggregation site. The aggregation site can then be consulted by PSI/J developers who can react to issues detected on various systems.

The Python implementation of PSI/J allows its test suite to be configured to run daily on a given machine, typically using the standard *Cron* tool, and upload detailed results of the tests to a central location. By default, the system is set up to test the main branch as well as branches belonging to GitHub pull requests that are made by the core team. While not currently implemented, we envision adding a tagging mechanism that allows the test suite to also be run on external pull requests, but only if tagged by a member of the core team. This allows both continuous testing of the main branch of the Python implementation of PSI/J, as well as that of relevant proposed changes. A *minimal uploads* mode, which strips all non numeric and non boolean information from the results but retains the detailed information locally, is provided for use on sensitive and/or classified systems. When errors are detected on such a system and if desired, PSI/J developers have an option of emailing the test maintainers to ask for detailed information which can then be manually inspected and edited for sensitive information by the test maintainers.

The results of the tests are presented through a dashboard, which can display the data in multiple ways. One of the display modes shows aggregate results for each site over the past few days. Maintainers of PSI/J implementations can quickly identify sites on which one or more tests are failing. One can then select a particular site and inspect the actual *pytest* output, and example of which is shown in Figure 4. This effectively allows maintainers and developers of PSI/J to run tests and get detailed results on resources that they would not normally have access to. It also unburdens them from setting up, maintaining, and prioritizing test sites while allowing users to gain confidence that PSI/J implementations function correctly on their sites.

## V. EXAMPLES

In this section we describe the use of PSI/J by three workflow programming frameworks, Parsl, RADICAL-Pilot, and Swift/T, and one application framework, OSPREY.

### A. Parsl

Parsl is a parallel programming library for Python that allows users to write programs in Python that orchestrate parallel execution of Python functions and other external application.



The screenshot shows a web interface for test results. At the top, there's a breadcrumb trail: 'test\_multiple\_jobs: test\_getting\_started\_multiple\_jobs[slurm:slurm]'. Below this is a code block containing a Python script snippet. The script defines a function `test_getting_started_multiple_jobs` which interacts with a batch scheduler. The error message at the bottom of the code block reads: `psij.exceptions.SubmitException: sbatch: error: Batch job submission failed: Job violates`. Below the code, the dashboard displays the site ID 'behop.lrcr.anl.gov' and the test start time 'Fri, Apr 21, 2023 - 17:19:42'.

Fig. 4. Screenshot of the test results view of the PSI/J testing dashboard showing a failed test and the Pytest report.

Parsl exposes a straightforward programming model in which developers identify opportunities for concurrent execution by adding decorators to Python functions, called Parsl *apps*. Any call to a Parsl app creates a new task that executes concurrently with the main program and any other task(s) that are currently executing. Parsl returns a *future* for each task, allowing the developer to compose dynamic workflows by passing futures between apps. When the program is executed, Parsl manages task dependencies, data transfers, and execution of tasks on connected resources.

Parsl implements a flexible runtime model via which tasks can be executed on different resources, from laptops to supercomputers. The runtime model builds on two extensible abstractions: the *provider*, a JAAPI that manages the provisioning and management of compute resources (e.g., from batch schedulers or cloud providers); the *executor* (not to be confused with PSI/J executors) manages the execution of tasks on those resources. Parsl, like other workflow systems, includes implementations for various providers (e.g., Slurm, PBS, Cobalt, LSF). Parsl has been extended to leverage PSI/J as a community-developed JAAPI. Integration was straightforward, as PSI/J implements a compatible model and required making only minor adjustments to certain library calls.

### B. RADICAL-Pilot

RADICAL-Pilot (RP) enables the execution of one or more workloads comprised of heterogeneous tasks on one or more HPC platforms. RP offers five unique features: (1) concurrent execution of tasks with five types of heterogeneity; (2) concurrent execution of multiple workloads on a single pilot, across multiple pilots and across multiple HPC platforms; (3) support of all major DoE and NSF HPC platforms, offering



reliable scaling behavior of  $O(10^5)$  concurrent tasks on up to  $O(6 * 10^5)$  cores and  $O(3.6 * 10^4)$  GPUs; (4) support for seventeen methods to launch tasks; and (5) integration with third party middleware like workflow and runtime systems. The five types of task heterogeneity supported by RP are: (1) type of task (executable, function or method); (2) parallelism (scalar, MPI, OpenMP, or multiprocess/thread); (3) compute support (CPU and GPU); (4) size (1 hardware thread to 8000 compute nodes); and duration (0 seconds to 48 hours).

As a pilot system, RP schedules tasks concurrently and sequentially, depending on available resources, and defines scheduling policies for executing tasks on all the acquired resources. As such, RP requires scheduling a job on an HPC machine via its batch system to acquire resources, which makes supporting diverse platforms with the same code base challenging. So far, RP has been using RADICAL-SAGA to support all the major batch systems (see Table II) but, due to limitations listed in §II, RP has integrated PSI/J to offer the same interoperability but via leaner, more extensible, and easier to integrate interfaces.

### C. Swift/T

Swift/T is a dataflow language with automatic parallelization capabilities [29], with a runtime based on MPI messaging [30]. Swift/T is designed to run large numbers of small tasks at the CPU core-level granularity at a scale of  $O(10^5 - 10^6)$  processors for tasks that run for  $O(10)$  seconds, although larger-granularity tasks, including tasks that are MPI codes, are typical. Swift/T originally used a collection of hand-coded shell scripts to launch itself on scheduler systems including PBS, SLURM, LSF, Cobalt, and plain `mpirun` execution. Using PSI/J provided the opportunity to reduce the maintenance burden of supporting the various schedulers and their often customized installations on exotic large-scale computing systems.

At run time, Swift/T needs to launch its runtime on a certain MPI environment. The Swift/T runtime, Turbine, is packaged as a Tcl library that links to MPI, ADLB [31], and optional application libraries including Python, R, and the JVM library [32], all at the C level. Running a Swift/T program essentially consists of invoking a command of the form `mpirun tclsh workflow.tic`, where `workflow.tic` is a Tcl program translated from the high-level user workflow specification. The complexities involved include setting the job specification and environment as requested by the user with an appropriate request for hardware resources and environment settings for all of the libraries.

A PSI/J module was implemented for Swift/T as done for its other scheduler scripts. However, the PSI/J module is much simpler than those for specific LRMs since all of the complexity is pushed into the PSI/J layer. A small stub of Python code gathers the settings as encoded by Turbine and presents them to PSI/J, which constructs and launches the job. Thus, simply by specifying PSI/J as the scheduler, Swift/T users are able to benefit from the portability and feature set of PSI/J without additional work from the Swift/T maintainers.

### D. OSPREY: Open Science Platform for Robust Epidemic Analysis

The COVID-19 pandemic has highlighted both the utility and the difficulties associated with applying computational epidemiology to decision-making in times of crisis and uncertainty. The Open Science Platform for Robust Epidemic Analysis (OSPREY) [33] was developed while taking the lessons learned from supporting public health stakeholders during the pandemic. The platform seeks to create a readily deployable capability for producing scenario analyses and forecasts of epidemiological quantities of interest, such as cases, resource needs, and disease outcomes [34], [35], utilizing distributed HPC resources.

One characteristic of epidemic analyses is that their computational demand can vary dramatically over time. Furthermore, computational availability can fluctuate due to demand and resource priorities. Also, for epidemic analyses to be useful as decision support tools, they need to provide actionable insights quickly. To support these requirements, OSPREY uses scalable worker pools, elastic sets of pilot jobs tuned to run specific tasks within the epidemic analyses, such as agent-based disease models or ML/AI computations. These worker pools are dynamically provisioned based on computational demand. The client code that triggers the analyses remotely, e.g., from a researcher laptop, is able to actively monitor the worker pools and terminate them as needed, both of which are done through PSI/J's ability to query and interact with LRMs. The remote authentication and interaction is enabled by Globus Compute [15]. The worker pool is submitted for execution to a LRM using the Swift/T [5] framework, which returns a job identifier. Using the PSI/J Python API, a `psij.Job` object is created and attach to a native job using an instance of a PSI/J `JobExecutor`. The `Job` object's methods and attributes can then be used to query the worker pool job's status or cancel it.

## VI. PERFORMANCE

An important characteristic of abstraction libraries is the amount of overhead that they introduce on top of the underlying components that they abstract, since a sufficiently large overhead may be a deterrent in adopting an abstraction library. We present a number of measurements that are meant to quantify the overhead introduced by PSI/J-Python.

We first consider the scaling of the time required to run dummy jobs with the number of jobs using the *local* executor. This measures overheads introduced by the executor itself as well as the overhead of constructing `Job` objects and querying their status. We run all jobs synchronously and sequentially in loops. We run the experiments on a modern laptop (AMD Ryzen 7 PRO 4750U) running Linux. Results collected for PSI/J-Python, a number of other Python methods of invoking an executable, as well as a Perl version of the same are shown in Figure 5. The results show that the local PSI/J executor introduces an overhead of about 10 milliseconds per job.

All PSI/J executors invoke applications indirectly, using a launcher script, which is specific to the launching mechanism



desired by users (e.g., `mpirun`, `srun`) but also allows for more complex usage scenarios, such as pre- and post-launch scripts. We first collect results from all methods while uniformly invoking the application through a PSI/J-Python launcher script. In order to ascertain the typical overhead of PSI/J-Python launcher scripts, we compare timing results of the default PSI/J-Python launcher script, with those of a minimal wrapper script which simply runs the job as a subprocess, and, finally, to those without the launcher script. The corresponding results are shown in Figure 6. The results show an average overhead of about 2 milliseconds per job attributable to the use of a launcher script. There is no material difference between the default PSI/J-Python launcher script and a minimal wrapper script.




Fig. 5. Time required to run  $N_{jobs}$  dummy jobs locally using various methods. The “Popen() + poll()” method starts the job using the Popen() call and then repeatedly calls “poll()” on the Popen object in a busy loop.




Fig. 6. Time required to run  $N_{jobs}$  dummy jobs with and without a PSI/J-Python launcher script.

Finally, we explore if PSI/J-Python imposes an observable load on a queuing system. We do so by submitting and monitoring a number of jobs while measuring the latency of the queuing system by repeatedly timing a simple *qstat* command that queries the status of a known job. We run experiments on three supercomputers at ORNL: Crusher, Frontier, and Summit. These supercomputers use Slurm and LSF schedulers. The results, presented in Figure 7, show that, up to 100 jobs, PSI/J does not appear to affect queue performance in a significant way. Nonetheless, a large variance in the time taken to run the *qstat* command is observed.

The numbers shown here suggest that the overheads imposed by the PSI/J Python implementation add an insignificant overhead to the typical jobs and LRMs that the PSI/J Python implementation is expected to manage.

## VII. DISCUSSION AND NEXT STEPS

We presented PSI/J, a minimal, composable, and extensible JAAPI specification whose goal is to address many of the




Fig. 7. Dependence of the time taken to run a *qstat* command on the number of jobs actively managed by PSI/J.

hardships associated with achieving a durable JAAPI solution. A reference Python implementation which is meant to encourage community participation is also provided. Additionally, we provide an important tool that allows testing of the reference implementation in ways that would otherwise be significantly challenging.

With most of the JAAPIs described in this paper, it is apparent that the teams responsible for the JAAPIs that ended up getting considerable traction with the community were aware of most challenges and problems involved in designing, implementing, and supporting their solutions successfully. Globus GRAM came at a time when C was the favored language and few modern networking or security libraries or protocols existed. It did what it had to do in order to provide a remote JAAPI solution. The SAGA specification came at a time when the Grid appeared destined for long term domination of the HPC world. It attempted to introduce a standard that addressed the many needs that bridged reality with what was necessary. This gives weight to the idea that extrinsic factors are largely responsible for what the lack of a cohesive JAAPI solution.

We aim to refine both the specification and reference implementation to allow a transition to community-directed support and development in order to mitigate against the likely funding difficulties associated with JAAPIs. We believe that this is a viable solution that is predicated on the assumption that, if provided with a sufficiently well designed JAAPI and a sufficiently well executed reference implementation, the community will stand to gain from supporting a common solution rather than many disparate ones.

## ACKNOWLEDGEMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-CONF-826133), Argonne National Laboratory under Contract DE-AC02-06CH11357, and Brookhaven National Laboratory under Contract DESC0012704. This research used resources of the OLCF at ORNL, which is

supported by the Office of Science of the U.S. DOE under Contract No. DE-AC05-00OR22725. OSPREY work was supported by the National Science Foundation under Grant No. 2200234, the National Institutes of Health under grant R01DA055502, and the DOE Office of Science through the Bio-preparedness Research Virtual Environment (BRaVE) initiative. This manuscript has been authored in part by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid up, irrevocable, worldwide license to publish or reproduce the published form of the manuscript, or allow others to do so, for U.S. Government purposes. The DOE will provide public access to these results in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

## REFERENCES

- [1] M. A. Salim, T. D. Uram *et al.*, “Balsam: Automated scheduling and execution of dynamic, data-intensive HPC workflows,” *ArXiv*, vol. abs/1909.08704, 2019.
- [2] F. Di Natale, H. Bhatia *et al.*, “A massively parallel infrastructure for adaptive multiscale simulations: Modeling RAS initiation pathway for cancer,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019.
- [3] Y. Babuji, A. Woodard *et al.*, “Parsl: Pervasive parallel programming in Python,” in *28th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, 2019.
- [4] A. Merzky, M. Turilli *et al.*, “Design and performance characterization of Radical-Pilot on leadership-class platforms,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 4, pp. 818–829, 2021.
- [5] J. M. Wozniak, T. G. Armstrong *et al.*, “Swift/T: Large-scale application composition via distributed-memory dataflow processing,” in *2013 13th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing*. IEEE, 2013, pp. 95–102.
- [6] R. Ferreira da Silva, R. M. Badia *et al.*, “Workflows Community Summit 2022: A Roadmap Revolution,” Oak Ridge National Laboratory, Tech. Rep. ORNL/TM-2023/2885, Mar. 2023.
- [7] The Exaworks Team, “PSI/J Python Reference Implementation,” <https://purl.org/psij.io/psij-python>.
- [8] P. Troger, H. Rajic *et al.*, “Standardization of an API for distributed resource management systems,” in *7th IEEE International Symposium on Cluster Computing and the Grid*, 2007, pp. 619–626.
- [9] K. Czajkowski, I. Foster *et al.*, “A resource management architecture for metacomputing systems,” in *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 62–82.
- [10] M. Romberg, “The UNICORE grid infrastructure,” *Scientific Programming*, vol. 10, no. 2, pp. 149–157, 2002.
- [11] D. Thain, T. Tannenbaum, and M. Livny, “Distributed computing in practice: The condor experience: Research articles,” *Concurrency and Computation: Practice and Experience*, vol. 17, no. 2–4, p. 323–356, feb 2005.
- [12] T. Goodale, S. Jha *et al.*, “SAGA: A simple API for grid applications. high-level application programming on the grid,” *Computational Methods in Science and Technology*, vol. 12, no. 1, pp. 7–20, 2006.
- [13] I. Foster and C. Kesselman, “The Globus Toolkit,” in *The Grid: Blueprint for a New Computing Infrastructure*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1998, p. 259–278.
- [14] K. Chard, S. Tuecke, and I. Foster, “Efficient and secure transfer, synchronization, and sharing of big data,” *IEEE Cloud Computing*, vol. 1, no. 3, pp. 46–55, 2014.
- [15] R. Chard, Y. Babuji *et al.*, “funcX: A federated function serving fabric for science,” in *29th International Symposium on High-Performance Parallel and Distributed Computing*. Stockholm Sweden: ACM, 6 2020, pp. 65–76.
- [16] H. Kaiser, A. Merzky *et al.*, “The saga c++ reference implementation: a milestone toward new high-level grid applications,” in *Proceedings of the 2006 ACM/IEEE conference on Supercomputing*, 2006, pp. 184–es.
- [17] M. Russell, P. Dziubecki *et al.*, “The vine toolkit: A java framework for developing grid applications,” in *Parallel Processing and Applied Mathematics: 7th International Conference, PPAM 2007, Gdansk, Poland, September 9-12, 2007 Revised Selected Papers 7*. Springer, 2008, pp. 331–340.
- [18] A. Merzky, O. Weidner, and S. Jha, “Saga: A standardized access layer to heterogeneous distributed computing infrastructure,” *SoftwareX*, vol. 1, pp. 3–8, 2015.
- [19] I. Foster, A. Grimshaw *et al.*, “OGSA Basic Execution Service version 1.0,” 2007, oGF GFD-R.108.
- [20] A. Anjomshoa, F. Brisard *et al.*, “Job Submission Description Language (JSDL) specification, version 1.0,” 2005, open Grid Forum, GFD-56.
- [21] G. Von Laszewski, I. Foster *et al.*, “A Java commodity grid kit,” *Concurrency and Computation: practice and experience*, vol. 13, no. 8-9, pp. 645–662, 2001.
- [22] K. Arnin, “Abstracting the grid,” in *12th Euromicro Conference on Parallel, Distributed and Network-Based Processing*. IEEE, 2004, pp. 250–257.
- [23] Y. Zhao, M. Hategan *et al.*, “Swift: Fast, reliable, loosely coupled parallel computation,” in *IEEE Congress on Services*. IEEE, 2007, pp. 199–206.
- [24] M. Hategan, J. Wozniak, and K. Maheshwari, “Coasters: Uniform resource provisioning and access for clouds and grids,” in *4th IEEE International Conference on Utility and Cloud Computing*. IEEE, 2011, pp. 114–121.
- [25] The Open Science Grid Executive Board, “The open science grid,” *Journal of Physics: Conference Series*, vol. 78, no. 1, p. 012057, jul 2007.
- [26] The Exaworks Team, “A Portable Submission Interface for Jobs (PSI/J),” <https://purl.org/psij.io/spec>.
- [27] M. den Burger, C. Jacobs *et al.*, “What is the price of simplicity? a cross-platform evaluation of the saga api,” in *Euro-Par 2010-Parallel Processing: 16th International Euro-Par Conference, Ischia, Italy, August 31-September 3, 2010, Proceedings, Part 1 16*. Springer, 2010, pp. 392–404.
- [28] D. H. Ahn, J. Garlick *et al.*, “Flux: A next-generation resource management framework for large hpc centers,” in *2014 43rd International Conference on Parallel Processing Workshops*, 2014, pp. 9–17.
- [29] T. G. Armstrong, J. M. Wozniak *et al.*, “Compiler techniques for massively scalable implicit task parallelism,” in *Proc. SC*, 2014.
- [30] J. M. Wozniak, T. G. Armstrong *et al.*, “Swift/T: Scalable data flow programming for distributed-memory task-parallel applications,” in *Proc. CCGrid*, 2013.
- [31] E. L. Lusk, S. C. Pieper, and R. M. Butler, “More scalability, less pain: A simple programming model and its implementation for extreme computing,” *SciDAC Review*, vol. 17, 2010.
- [32] J. M. Wozniak, T. G. Armstrong *et al.*, “Interlanguage parallel scripting for distributed-memory scientific computing,” in *Proc. WORKS @ SC*, 2015.
- [33] N. Collier, J. M. Wozniak *et al.*, “Developing distributed high-performance computing capabilities of an open science platform for robust epidemic analysis,” in *2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2023, pp. 868–877.
- [34] J. Ozik, J. M. Wozniak *et al.*, “A population data-driven workflow for COVID-19 modeling and learning,” *The International Journal of High Performance Computing Applications*, vol. 35, no. 5, pp. 483–499, 9 2021.
- [35] A. L. Hotton, J. Ozik *et al.*, “Impact of changes in protective behaviors and out-of-household activities by age on COVID-19 transmission and hospitalization in Chicago, Illinois,” *Annals of Epidemiology*, p. S1047279722001053, 6 2022.