

# WORKSWORLD: A Domain for Integrated Numeric Planning and Scheduling of Distributed Pipelined Workflows

Taylor Paul<sup>1</sup>, William Regli<sup>1</sup>

<sup>1</sup>University of Maryland, College Park, MD USA  
thpaul@umd.edu, regli@umd.edu

## Abstract

This work pursues automated planning and scheduling of distributed data pipelines, or workflows. We develop a general workflow and resource graph representation that includes both data processing and sharing components with corresponding network interfaces for scheduling. Leveraging these graphs, we introduce WORKSWORLD, a new domain for numeric domain-independent planners designed for permanently scheduled workflows, like ingest pipelines. Our framework permits users to define data sources, available workflow components, and desired data destinations and formats without explicitly declaring the entire workflow graph as a goal. The planner solves a joint planning and scheduling problem, producing a plan that both builds the workflow graph and schedules its components on the resource graph. We empirically show that a state-of-the-art numeric planner running on commodity hardware with one hour of CPU time and 30 GB of memory can solve linear-chain workflows of up to 14 components across eight sites.

### Domain —

<https://github.com/taylorpaul/WORKSWORLD>

### Benchmarks —

<https://gitlab.com/thpaul/worksworld-benchmarks/>

**Data** — <https://gitlab.com/thpaul/worksworld-benchmarks/-/jobs/13325198684/artifacts/download>

## 1 Introduction

Spending billions of dollars on data platforms and artificial intelligence (AI), most organizations struggle to see any tangible return on investment (ROI) (Challapally et al. 2025). Why? Today’s platforms and algorithms deliver little value without unfettered access to collected, curated, and current data (Heck 2024) from the distributed systems that support daily operations. Establishing effective data pipelines, or *data-intensive workflows* (hereafter, *workflows*), persists as a prerequisite to productive employment of AI tools. Such workflows must process diverse data sources into required formats and feed them into multiple systems and models without excessive cost or prohibitive latency. Optimally planning and scheduling such workflows, especially those

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

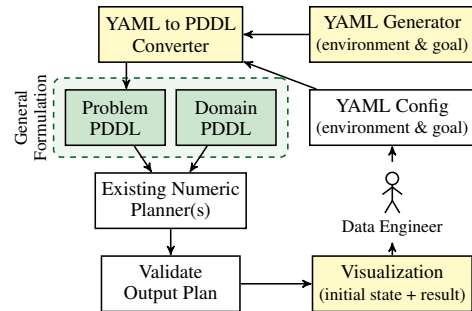


Figure 1: The WORKSWORLD framework. Primary research contributions are green; secondary engineering artifacts provided without experimental evaluation are yellow.

sufficiently complex to meet real-world use cases, most often proves computationally intractable (Versluis and Iosup 2021; Ghallab et al. 1998).

Diverse communities have long pursued efficient placement of computational tasks on distributed heterogeneous resources: high-performance computing (HPC) (Da Silva et al. 2024), distributed computing (Dustdar, Pujol, and Donta 2023), and machine learning (Kreuzberger, Kühl, and Hirschl 2023). These communities share a core problem, yet bring distinct perspectives, objectives, and assumptions that shape both tractability and real-world applicability of solutions. Prior work typically begins with a fully-formed *workflow*—a directed acyclic graph (DAG) of computational tasks connected by data dependency edges—and solves a scheduling problem to place tasks on compute resources while optimizing one or two metrics (e.g., latency, throughput, or cost) (Vivas, Tchernykh, and Castro 2024). We jointly plan and schedule a DAG of both processing and data components, explicitly decoupling compute, storage, and network decisions.

Data-intensive workflows have received increased attention over the past decade (Suter et al. 2023; Heck 2024), inspiring new organizational approaches (e.g. data mesh (Goedegebuure et al. 2024)) and roles (e.g. data engineers). Complexity in planning these workflows increases when data producers (*sources*) and consumers (*sinks*) span multiple networks or *sites*, especially when consumers require the same source data in different formats. A data engi-

neer must decide which data conversions to apply, whether to process data centrally (i.e. the cloud) or simultaneously at multiple sites, and how to place processing tasks across available compute resources while meeting capacity and interoperability constraints. This complexity, combined with the workflow graph’s direct impact on performance and cost, makes such workflows strong candidates for joint planning and scheduling approaches (Long, Dolejsi, and Stolba 2023).

To aid data engineers in their work, we develop the framework depicted in Figure 1 that permits them to define the problem environment, workflow components, and goals (output data formats and locations) as a YAML Ain’t Markup Language (YAML) configuration file. Our framework converts the YAML to a Planning Domain Definition Language (PDDL) problem and solves a numeric planning problem using existing numeric planners and a new PDDL domain we call WORKSWORLD. We validate the output plan<sup>1</sup> (Howey, Long, and Fox 2004) and provide a simple visualization tool depicting both the provided initial state and result of executing the plan. In contrast, existing solutions are largely proprietary and platform-specific (Amado et al. 2025; Sohrabi et al. 2013). To our knowledge, no publicly available system addresses our exact problem formulation. Thus, we demonstrate viability through scalability results rather than a competitive baseline comparison. We find that a state-of-the-art (SOA) numeric planner, the Expressive Numeric Heuristic Search Planner (ENHSP), can plan and schedule satisficing workflows of seven processing and seven data components across eight sites with reasonable CPU time (one hour) and memory (30 GB) on commodity hardware.

## 2 Problem and Motivation

Depicted in Figure 2, we seek to map pipelined workflows (see Definition 1) *permanently*, until deprovisioned, to compute and storage resources across distributed cloud, fog, and edge sites (see Section 3). We do not require the entire workflow graph as input and thus solve an integrated planning and scheduling problem. We seek satisficing mappings of workflows to site resources that minimize compute, storage and network costs while constraining latency.

### Motivating Workflow Examples

We present three workflows spanning distinct cost-latency tradeoff regimes to illustrate the value of decoupling processing, data placement, and routing.

1. *Archival Workflow*: Digital media archiving (Hou et al. 2006). Records from three distributed sites must be encoded, compressed, and stored at a central cloud site (hours tolerable). Tradeoff: compress at source (less transfer cost) or centrally to consolidate compute?
2. *Sensor Workflow*: Wildfire detection (Altintas et al. 2022). Edge sensors generate raw data requiring pre-processing before algorithm inference to detect smoke and trigger simulations (minutes tolerable before human

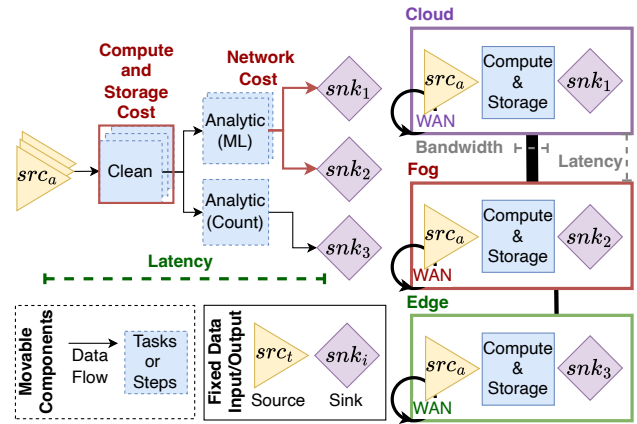


Figure 2: Our scheduling problem consists of mapping movable components in the left-hand DAG to utilize compute, storage and network resources across sites.

review). Tradeoff: transfer raw data to the cloud, or move pre-processing and inference to edge sites?

3. *Automation Workflow*: Cybersecurity. Firewall events undergo normalization and anomaly detection to trigger automatic IP blocking (sub-second latency required). Tradeoff: meeting strict latency bounds while placing compute and data to minimize cost.

These workflows span the cost-latency tradeoff space relevant to our research: workflow purpose drives latency tolerance, while data volume and compute placement decisions drive cost. They motivate our research and guide our experimental setup in Section 5.

## 3 Background

Now we lay down a foundation, built by others, on which our contribution rests. We briefly define the computing environment in which we plan, scope our workflow scheduling problem, and provide the required AI planning prerequisites.

### Compute Environment

The distributed computing research community provide abstractions to describe the computing environment they study. The early 2000’s saw a heavy focus on Grid Computing (Foster, Kesselman, and Tuecke 2001) before cloud-computing (Mell and Grance 2011) stole much of the community’s focus. When latency sensitive workflows did not fit the cloud only approach, the Cloud-Compute Continuum (Moreschini et al. 2022) emerged to define three general categories of compute resources: cloud, fog and edge which we leverage in describing different sites and their compute, storage and network resources.

### Pipelined Workflows Scheduling

Scheduling compute tasks across available resources proves an incredibly broad topic studied by multiple communities across many decades (Ghallab, Nau, and Traverso 2004, Ch. 15). There is no best scheduling approach in general to

<sup>1</sup><https://github.com/KCL-Planning/VAL>

workflows due to the NP-hardness of the problem and the influence of specific scheduling context on the time and space required for finding acceptable solutions (Versluis and Iosup 2021). We find a scoping to pipelined workflows scheduling clarifying in terms of our context; see Definition 1.

**Definition 1** (Pipelined Workflow Scheduling). Pipelined workflow scheduling encompasses executing the same workflow on varying datasets of identical size across parallel machines utilizing task, data, pipelined and/or replicated parallelism. (Benoit et al. 2013)

In our setting, the datasets are individual or groups of source data messages aggregated at our input data components. Benoit et al. 2013 provide a survey and taxonomy to classify pipelined workflow scheduling characteristics and specify complexity for scheduling linear task graphs. With their taxonomy, we classify our problem as scheduling pipelined workflows across related-heterogenous compute and storage resources at distributed sites connected by structured-heterogenous, bandwidth bounded, network links. We employ a multi-objective performance model that minimizes resource utilization (i.e. cost) while constraining latency (corresponding to *makespan* in single dataset DAG scheduling). However, our *permanent* scheduling of components to resources rather than repeatedly scheduling dataset execution tasks removes the temporal sequencing that drives the scheduling complexity in classical pipelined workflows. The hardness of WORKSWORLD instead arises from multi-resource placement under network constraints, which proves structurally similar to the NP-hard virtual network embedding problem (Fischer et al. 2013).

## Classical and Numeric Planning

AI Planning, or automated planning and scheduling, consists of leveraging computers and algorithms to generate plans to execute towards achieving a set of goals; Ghallab, Nau and Traverso (2004) cover in detail topics we discuss briefly.

*Classical planning* possesses the most restrictive assumptions but maintains a general complexity of PSPACE-complete (Bylander 1994); it does not permit reasoning about numeric variables. Planning systems that interact with real-world physical systems often require *numeric planning* (Scala et al. 2016). Numbers increase the complexity in the worst case to semi-decidable and removes bounds on plan length possible in propositional planning (Haslum et al. 2019). Even so, we leverage numeric planning in our work due to our foundational dependency on numeric measurements of compute, storage and network resource availability. Due to the *permanent* nature of our data pipelines, we do not leverage *temporal planning* to reason about time as the ordering of actions (i.e. sequential planning) to schedule and connect components suffices.

In the past three decades, the AI planning community produced several solvers capable of planning with numerics: MetricFF (Hoffmann 2003), LP-td (Gerevini, Saetti, and Serina 2006), SMTPLAN<sup>+</sup> (Cashmore et al. 2016). We exclusively evaluate ENHSP<sup>2</sup> due to the functionality it supports

(i.e. Quantifiers in pre-conditions and goals), its support for parsing large problem files and domains (resp. MetricFF), and its effective pruning of the search space with Additive Interval-Based Relaxation (AIBR) preprocessing (Scala et al. 2016). We leave finding additional compatible planners and testing them for future work (e.g. SMTPLAN<sup>+</sup>, GOOSE (Chen and Thiébaux 2024)).

With our scope established, we turn to defining the general numeric planning problem we solve. We reproduce others’ work (Fox and Long 2003; Scala et al. 2016, 2020) minimally for the readers’ benefit. The state of our system ( $s_i \in \mathcal{S}$ ) amounts to a total assignment of a set of propositional variables ( $\mathcal{F}$ ) and a partial assignment of numeric variables ( $\mathcal{X} := \{x \mid x \in \mathbb{R}\}$ ); let the set of all variables  $\mathcal{V} = \mathcal{F} \cup \mathcal{X}$ . Propositional conditions are positive literals ( $p = \top$ ) or negative literals ( $p = \perp$ ) where  $p \in \mathcal{F}$ . Numeric conditions take the form  $\langle \xi, \triangleright, l \rangle$  where  $\xi$  is a numeric expression over a subset of  $\mathcal{X}$ ,  $\triangleright \in \{\leq, <, =, >, \geq\}$  and  $l \in \mathbb{Q}$ . The operators permitted in these expressions are specific to the planner (or heuristic) leveraged, we currently leverage the standard set of operators  $+, -, \times, \div$  in WORKSWORLD; ENHSP supports more. With the basics covered, Definition 2 describes a numeric action and Definition 3 formulates the numeric planning problem.

**Definition 2** (Numeric Action). A numeric action  $a$  consists of two sets:  $\text{pre}(a)$  the set of numeric and propositional pre-conditions permitting a planner to select an action and  $\text{eff}(a)$  the set of propositional and numeric effects that define the resulting state changes after executing the action ( $a = \langle \text{pre}(a), \text{eff}(a) \rangle$ ). Propositional effects set ( $p = \top$ ) or unset ( $p = \perp$ ) a proposition  $p \in \mathcal{F}$ . Numeric effects can assign ( $=$ ), increase ( $+=$ ) or decrease ( $-=$ ) numeric variables  $x \in \mathcal{X}$  by  $\xi$ . An action’s  $\text{eff}(a)$  must only update a variable  $x$  at most once.

**Definition 3** (Numeric Planning Problem). A numeric planning problem ( $\Pi = \langle s_0, \mathcal{A}, \mathcal{G}, \mathcal{V} \rangle$ ) consists of an initial state ( $s_0$ ), which is an initial assignment (complete for  $\mathcal{F}$ ) to the variables in  $\mathcal{V} = \mathcal{F} \cup \mathcal{X}$ , a set of numeric actions ( $\mathcal{A}$ ), and a set of propositional or numeric goal conditions ( $\mathcal{G}$ ).

The act of planning involves selecting a sequence of actions that transitions our system from the initial state  $s_0$  to a goal state  $s_g$  that satisfies ( $\models$ ) all  $g \in \mathcal{G}$  or  $s_g \models \mathcal{G}$ . This sequence of actions is the planner’s output plan  $\pi$ . We reproduce Definition 4 from (Scala et al. 2016) with a slight modification to the description of an optimal plan; see their work for more details on the applicability of numeric actions in a state.

**Definition 4** (Plan). A plan  $\pi$  for  $\Pi = \langle s_0, \mathcal{A}, \mathcal{G}, \mathcal{V} \rangle$  is a sequence of actions  $a_0, \dots, a_{n-1}$  from  $\mathcal{A}$  such that each action in  $\pi$  is applicable in the state resulting from the application of its predecessors, i.e.,  $s_0 \models \text{pre}(a_0)$ ,  $\text{succ}(s_0, a_0) \models \text{pre}(a_1)$ , etc, and  $\text{succ}(s_{n-1}, a_{n-1}) \models \mathcal{G}$ . A plan  $\pi$  is said to be optimal if, among all valid plans, it has a minimal number of actions (by default) or it minimizes or maximizes some specified  $x \in \mathcal{X}$  or  $\xi$  over  $\{x\}$ .

With our problem and desired solution sufficiently specified, we turn to describing the required inputs for domain-independent planners.

<sup>2</sup><https://sites.google.com/view/enhsp/>

## Planning Inputs

The ability to encode a planning problem into a common format permits the creation of domain-independent solvers; a single solver for many types, or domains, of problems. The AI Planning Community produced several versions of Planning Domain Definition Language (PDDL) to standardize inputs across planners (Ghallab et al. 1998; Fox and Long 2003; Edelkamp and Hoffmann 2004; Gerevini and Long 2005) We encode the WORKSWORLD in PDDL 2.1 level 2 (Fox and Long 2003) which permits our use of numeric fluents (or functions); we do not leverage the durative actions of level 3. For an in-depth overview of PDDL see (Haslum et al. 2019).

To solve the planning problem of Definition 3, AI planners require a planning instance ( $\mathcal{I} = \langle \text{Domain}, \text{Problem} \rangle$ ) encoded in PDDL as input. While the domain gets defined once, generating problem instances in PDDL, especially when  $s_0$  or  $\mathcal{V}$  prove large, remains a well-known barrier to users leveraging AI planners. Therefore we permit users (or systems) to more concisely define the problem instance using a popular human-readable data serialization language (YAML<sup>3</sup>) to produce a problem configuration file with common units and minimal requirements. We leverage a templating library (Jinja<sup>4</sup>) and python scripts to produce problem instances in PDDL from the supplied configuration file. The PDDL contains additional values from calculations across input values as well as consistent units to minimize mathematical operations required during planning (i.e. storage in MB and bandwidth in MB/s); see example YAML and PDDL files available online<sup>5</sup>.

## 4 General Problem Formulation

We start by defining our environment, which informs our resource graph, and the workflows we seek to efficiently distribute. Figure 3 depicts our sites and their main resources. The *site*, see Definition 5, proves the foundational abstraction of our environment. We assume that an organization manages specific sites and peers them, bringing them from the set of globally available sites ( $\mathbb{S}$ ) into a subset ( $S_m^{peers}$ ) considered for scheduling; see Equation 1. Peered sites must have at least one link connecting to the set of peers. These logical links (e.g. in a Software-Defined WAN) can be direct, or composite consisting of two hops. These sites possess data *sources* of a specific data type ( $src_i$ ) that serve as inputs to our workflows. They also may have *sinks* which are applications that consume workflow outputs.

$$S_m^{peers} := \{s_n \mid s_n \in \mathbb{S} \text{ and } \exists l(s_m, s_n)\} \quad (1)$$

**Definition 5 (Site).** We define a *site* ( $s_m$ ) that exists in the set of peered sites ( $S_m^{peers}$ ) as a list of three sets: the set of interfaces ( $I_m$ ) located at the site, a set of relations ( $L_m$ ) to describe intrasite and intersite network connections to peered sites, and the set of site-level annotations ( $A_m$ ) to capture

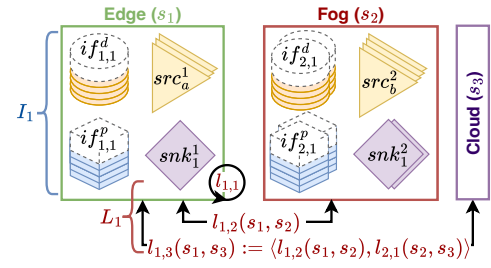


Figure 3: Sites and the relevant annotations and resources for our scheduling problem.

properties of the site where (indices  $m, i, k \in \mathbb{N}$ ):

$$\begin{aligned} s_m &:= \{ I_m, L_m, A_m \} \text{ such that} \\ I_m &:= \{ if_{m,i} \} \\ L_m &:= \{ l_{m,k}(s_m, s_n) \mid s_m, s_n \in S_m^{peers} \} \\ A_m &:= \{ \text{properties of site } m \} \end{aligned} \quad (2)$$

For this work, we present an interface abstraction to expose compute and storage resources for scheduling at sites. Thus, we schedule on a single interface over a clustered set of resources and simplify the number of metrics we must track. We identify two types of interfaces: *data sharing* ( $d$ ) and *data processing* ( $p$ ). We envision *data sharing interfaces* ( $DSIs$ ) to have different performance characteristics and resource costs (i.e. disk, CPU, RAM). There are existing open-source tools that an organization could standardize across sites for streaming, object storage or file sharing (e.g. Apache Kafka, CEPH, SeaweedFS, etc.). Data processing interfaces ( $DPIs$ ) prove most similar to resource managers like Kubernetes, Docker Swarm, or Slurm. Making these interfaces available as network-level resources for applications to leverage requires changes to the current network model (TCP/IP) which many others explore (Zhang et al. 2014; Król et al. 2019; Paul and Regli 2024).

**Definition 6 (Interface).** We define an *interface* ( $if$ ) as a network-level resource to schedule workflow components with a type ( $t$ ): data sharing ( $d$ ) or processing ( $p$ ). A set of metrics ( $IF^R$ ) describes the total and available resources. ( $IF^C$ ) is the set of components instantiated on the  $if$  and a set of annotations ( $IF^A$ ) describes hardware, software and performance characteristics. We define ( $m, i \in \mathbb{N}$ ):

$$if_{m,i}^t := \{ IF_{m,i}^R, IF_{m,i}^C, IF_{m,i}^A \} \text{ such that } t \in \{d, p\} \quad (3)$$

We establish our workflow components in Definition 7 that we will schedule on our interfaces and our rules for connecting them into a workflow; see Definition 8.

**Definition 7 (Workflow Component).** We define a *workflow component* ( $wfc^{c,t}$ ) as a logical boundary of work required by a workflow, delineating two classes of work: data processing ( $p$ ) or data sharing ( $d$ ) of a specific type ( $t$ ) from the

<sup>3</sup><https://yaml.org/>

<sup>4</sup><https://jinja.palletsprojects.com/en/stable/>

<sup>5</sup><https://github.com/taylorpaul/WORKSWORLD>

set of known component types ( $\mathbb{T}_c$ ), such that  $(x, y \in \mathbb{N})$ :

$$\begin{aligned}
wfc_x^{c,t} &:= \{C_{c,t}^R, C_x^{IO}, C_{t,x}^A \mid c \in \{p, d\}, t \in \mathbb{T}_c\} \\
C_{c,t}^R &:= \{amt_{c,t}^r \mid r \in \{\text{comp, store, mem, bw}\}\} \\
C_x^{IO} &:= \{C_{y \rightarrow x}^{in}, C_{x \rightarrow y}^{out} \mid y \neq x, l \in L_m \text{ at } s_m\} \\
C_{t,x}^A &:= \{\text{properties of } t \text{ or } x\}
\end{aligned} \tag{4}$$

**Definition 8 (Workflow).** We define a *workflow* as a DAG of *workflow components* with components *only* connecting via a *link* to one or more components of another class;  $d \rightarrow p$  or  $p \rightarrow d$  but *never*  $d \rightarrow d$ . A workflow *must* start and end with one or more data components.

With these definitions in hand, we prove prepared to pursue a planning and scheduling approach that can implement our model and produce actions required to map our workflows onto our connected site resources.

## 5 Methodology

We document the WORKSWORLD domain and describe its complexity. Then we describe our research questions and corresponding experiments.

### WORKSWORLD

Our main contribution is a new numeric planning domain, WORKSWORLD, that permits automated planning of pipelined workflow graphs to process and deliver data in the correct format to the desired site. We describe the domain’s primary objects and actions leveraging a single-site representation and discuss the added complexity when growing the resource and workflow graph. We developed and tested several iterations of the domain to achieve what we believe represents the right balance of expressivity, realism and complexity.

**Object Hierarchy** We introduce 17 object types to implement our model from Section 4 in our PDDL domain; see Figure 4. For our resource graph we define types *sites*, connected via *links* (direct and composite), *interfaces* (DSI/data processing interface (DPI)) and *resource*. We leverage constants of type resource in our domain: *storage*, *compute*, *network* and *config*. These constants describe resources available on (links and interfaces) or required (workflow components) by *instance* subtypes. We employ the config resource to reason about whether or not we must copy a workflow component configuration to a site or if it already exists there. Thus we can account for the cost of copying configurations like container images or model weights that may prove prohibitively large. We also provide the domain constants *input* and *output* of type *direction*, always relative to a processing component. Figure 4 also contains the abbreviations for our object types we employ throughout the paper.

**Predicates ( $\mathcal{F}$ )** Table 1 lists our 12 predicates; we discuss the most relevant here and point the reader to comments in the PDDL domain for details. The *available\_at* predicate declares what links or interfaces exist at a site. We also use the predicate to describe if a configuration for a workflow component is available at a site. The *fixed* predicate

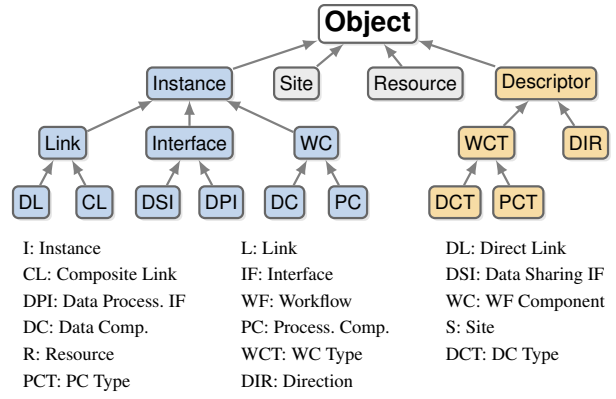


Figure 4: PDDL type hierarchy for WORKSWORLD and abbreviations utilized throughout paper.

Predicates ( $\mathcal{F}$ )	Numeric Fluents ( $\mathcal{X}$ )
(available_at I R S)	(resource_total I R)
(scheduled_on WC IF)	(resource_available I R)
(type_of WC WCT)	(work_amount WC R)
(fixed WC)	(msg_max_rate WC)
(linked L S S)	(msg_actual_rate WC IF)
(link_uses CL DL)	(msg_size DCT)
(connected L PC DPI DC DSI DIR)	(network_latency L)
(has_input WC IF)	(work_cost_weight R)
(input_format PC DCT)	(total_cost)
(output_format PC DCT)	(absolute_latency)
(processed_by DC DSI PC DPI)	
(has_data DC DSI DC DSI)	

Table 1: Available predicates and numeric fluents in WORKSWORLD.

marks workflow components that the planner cannot schedule. We use this for source and sink data components fixed to specific sites, while  $\neg$ fixed components can be scheduled at multiple sites. The *has\_data* predicate tracks data provenance through a planned workflow graph. In our goals, it describes what source data component we want processed into the format required by the destination sink data component.

**Numeric Fluents ( $\mathcal{X}$ )** We leverage ten numeric fluents in WORKSWORLD depicted in Table 1. Links and interfaces must have *resource\_total* and *resource\_available* fluents assigned in the initial state; *work\_amount* fluents describe how many resources (compute cores or storage bytes) scheduling a component on an interface consumes. We multiply *msg\_max\_rate* times *msg\_size* to consume permanent bandwidth resources when connecting components across links. We will discuss the other relevant numeric fluents when we discuss WORKSWORLD’s cost and performance calculations.

**Actions ( $\mathcal{A}$ )** Our action definitions prove the most influential on the compute and memory requirements during planning. The number of objects input into our actions greatly

### Actions (A) & Input Object Types

```
(replicate_code WC S S L)
(schedule_component WC R IF R S)
(connect_direct_link
  PC DPI S DC DCT DSI S DL DIR)
(connect_composite_link
  PC DPI S DC DCT DSI S CL DL DL DIR)
(propagate_input PC DPI DC DSI DC DSI)
(propagate_output PC DPI DC DSI DC DSI)
```

Table 2: Available actions in WORKSWORLD to build and schedule pipelined workflows.

affects the grounding time; the number of actions affects ENHSP AIBR preprocessing step as it loops over each action (Scala et al. 2016) to prune the state space. As we show later, these times dominate our planning time as problem sizes grow due to the size of the state space. Table 2 presents our six actions and the objects they require as inputs. We briefly walkthrough how a planner leverages our actions during planning before discussing how the actions affect cost and performance calculations.

Starting from  $s_0$ , the planner builds a workflow graph satisfying  $\mathcal{G}$  by iteratively scheduling and connecting components. To `schedule_component`, the component’s configuration must exist at the target site; `replicate_code` copies configurations from peered sites as needed. After scheduling required data and processing components, the planner connects them via `connect_direct_link` or `connect_composite_link` based on available links, bandwidth, and latency. It then selects `propagate_input` or `propagate_output` depending on the connection *direction* relative to the processing component, recording that a downstream component `has_data` from an upstream source. This process repeats until the planner achieves all goal conditions or the planner exhausts time limits, memory, or the entire pruned state space.

Separating the propagate effects from our connect actions provides two benefits. First, we avoid conditional effects for input or output in the connect actions, which the specific ENHSP heuristics (i.e. multi-repetition relaxed plan (MRP) (Scala et al. 2020)) we require do not support<sup>6</sup>. Second, we avoid requiring input and output versions of the more complex connect actions which reduces the required object combinations in actions during grounding; see object inputs in Table 2.

**Cost and Performance** In WORKSWORLD we minimize on cost (see Definition 9) and constrain on a message latency metric (see Definition 10) in  $\mathcal{G}$ . Our use of a domain-independent planner permits adapting these metrics with ease compared to domain specific optimization algorithms. We increase the cost metric for every action in Table 2 except the propagate actions. Importantly, we avoid utilizing `resource_available` to avoid state-dependent action costs; which ENHSP only experimentally supports. We instead ensure links and interfaces do

<sup>6</sup><https://sites.google.com/view/enhsp/home/how-to-use-it>

Object Type	F	X	A
Site	0	0	$\mathcal{O}(n^2)$
Direct Link	2	3	$\mathcal{O}(n^2)$
Composite Link	4	3	$\mathcal{O}(n)$
Interface (DSI/DPI)	1	2	$\mathcal{O}(n^2)$
Data Component	2	3	$\mathcal{O}(n^2)$
Processing Component	4	3	$\mathcal{O}(n^2)$
Resource	0	1	$\mathcal{O}(n^2)$

Table 3: Additional predicates (F), numeric fluents (X), and grounded actions (A) required by additional object types.

not get over provisioned via `pre(a)` checking for sufficient `resource_available`.

**Definition 9 (Total Cost).** We define *total-cost* as a weighted (*work-weight R*) resource utilization (*work\_amount WC R*) normalized by resource total (*resource\_total I R*) on a link or interface.

For our latency measure, due to an inability to track the bottleneck path and report the worst case latency, we instead define `absolute_latency` in Definition 10. We increase this numeric fluent in every action that establishes a permanent part of the workflow: the schedule and connect actions. Again, we apply a similar strategy to avoid state-dependent calculations by calculating latencies with a constant numeric fluent (`msg_max_rate`). We avoid more detailed definitions of cost and latency here due to space but point a curious reader to the domain PDDL for details.

**Definition 10 (Absolute Latency).** We define *absolute\_latency* as the sum of the time (seconds) it takes a single message to traverse every edge of the workflow graph.

**WORKSWORLD Expressivity** WORKSWORLD generalizes to any workflow expressible as a DAG of alternating data sharing and processing steps where components remain online until deprovisioned. This covers data pipelines common in sensing, scientific computing, and machine learning. It does *not* cover short-lived batch jobs that maximize compute utilization or minimize makespan across temporal tasks (e.g. some HPC workloads).

**State Space Complexity** We present Table 3 on the grounded state-space complexity increase as our problem sizes grow. We declare objects types that have multiple inputs to the same action (see Table 2) have the potential to impact the state space most significantly. While this table helps understand object-level complexity, there are combinatorial effects when we add sites and/or workflow components. Table 4 presents the state-space increase when adding sites, with their interfaces and links, and also when adding an extra data and processing component. Fortunately, ENHSP smartly prunes this state space before initializing search to permit sufficiently large problem sizes. We explore this in our experiments next.

### Experiments Overview

We possess two research questions we seek to answer about WORKSWORLD:

Problem Size	F (AIBR)	X (AIBR)	A (AIBR)
2WFC, 1S	22 (8)	7 (6)	14 (6)
2WFC, 2S	96 (24)	14 (12)	90 (22)
2WFC, 3S	306 (39)	22 (19)	284 (51)
4WFC, 3S	843 (81)	28 (25)	1146 (109)

Table 4: Grounding comparison: Predicates (F), Numeric Fluents (X), and Grounded Actions (A) before and after AIBR preprocessing.

- (RQ1) How does increasing the size of the workflow and resource graph affect the state space ENHSP searches?
- (RQ2) Does the domain definition permit scaling to large enough problems to plan and schedule workflows centrally across an enterprise?

We design and execute two experiments. First, we leverage our configuration and problem generator to observe state space growth from Table 3 by individually increasing object types. Second, we explore a larger combination of components, interfaces, sites and links that ENHSP can solve in reasonable time (one hour) and space (30 GB of RAM).

**Environment and Configurations** We use Cloud-Lab (Duplyakin et al. 2019) with a reproducible profile<sup>7</sup> that automatically configures our benchmarking environment. We deploy three *c6420*<sup>8</sup> nodes (dual sixteen-core processors at 2.6 GHz, 384GB DDR4 memory) running Kubernetes with open-source tools for metrics collection, storage, and visualization. Container images build and deploy to our Kubernetes cluster via a GitLab pipeline<sup>9</sup>. We constrain solver containers to five cores and 30 GB RAM with a one-hour CPU time limit per run. Each problem runs 20 independent times, and we report the 95th percentile with five-percent confidence intervals for measures with variation. ENHSP runs with a 30GB Java heap using the  $h_{\max}^{\text{MRP}} + H$  heuristic (Scala et al. 2020) (`sat-hmrph flag`<sup>10</sup>).

## 6 Results

Figure 5 shows results from our first experiment solving linear workflow chains while varying a single object type. As a baseline,  $s_0$  contains one site with one intrasite link, one DSI and DPI, and one scheduled source data component. For the top-left plot, we exponentially increase workflow components to schedule from two until failure at 128. For interfaces, we fix two components (to schedule) while exponentially increasing DSIs/DPI from two each (4 total) to failure at 64 each (128 total). For direct links, we scale intrasite links exponentially to 64 without failure. Finally, we add sites exponentially (each with one DSI, one DPI, one intrasite link, and one direct link to a central site), succeeding at 32 sites but failing at 64.

<sup>7</sup><https://gitlab.com/thpaul/k8s-promstack-profile>

<sup>8</sup><https://docs.cloudlab.us/hardware.html>

<sup>9</sup><https://gitlab.com/thpaul/worksworld-benchmarks>

<sup>10</sup><https://sites.google.com/view/enhsp/home/how-to-use-it>

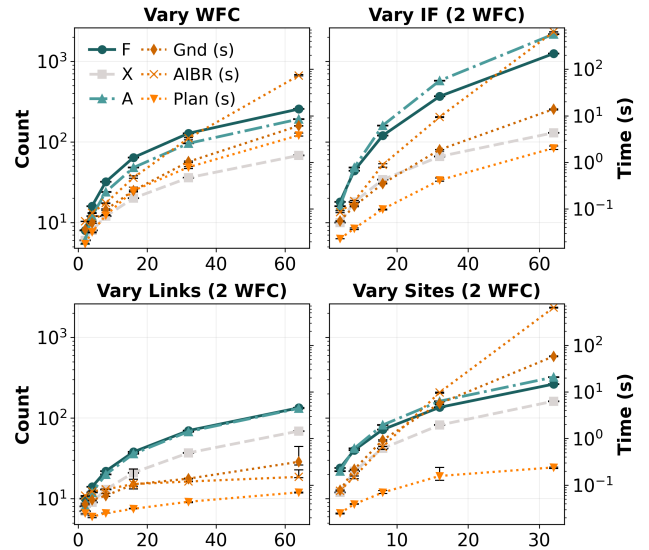


Figure 5: Planning and scheduling linear-chain workflows varying one aspect: workflow components, interfaces, direct links and sites.

Figure 6 presents our results from experiment two. From our first experiment, we selected a resource graph that produces reasonable planning times but is also representative of a mid-sized organization. We include eight sites in our initial state, each with a DSI/DPI with uniform `resource_total` across sites but randomly varied `resource_available` values. We connect them via seven direct links (with eight intrasite links) in a hub-spoke structure with the first site connected to all other sites. We then establish 14 composite links to permit over half of the possible two-hop connections ( $\binom{7}{2} = 21$ ). We increase the size of the workflow scheduled by two's from two to 16 components; 14 components just meet our memory and time limits, 16 components fail. We include total solver container energy measures from Kepler (Amaral et al. 2023), generated by reading the intel RAPL<sup>11</sup> metrics on each node, to report sustainability results and showcase the extensibility of our Kubernetes benchmark environment.

## 7 Discussion

Figure 5 addresses RQ1. In the *Vary WFC* subplot, most metrics transition from exponential to sub-exponential growth between 16 and 32 components, suggesting effective pruning, while AIBR preprocessing remains exponential throughout. For direct links, ENHSP handles added predicates and fluents efficiently. For interfaces, the planner prunes effectively but preprocessing grows exponentially, becoming the planning bottleneck. This effect intensifies when scaling sites (each adding two interfaces and two links). In sum, AIBR preprocessing proves effective yet emerges as the primary scalability bottleneck; accelerating

<sup>11</sup><https://sustainable-computing.io/archive/design/kepler-energy-sources/>

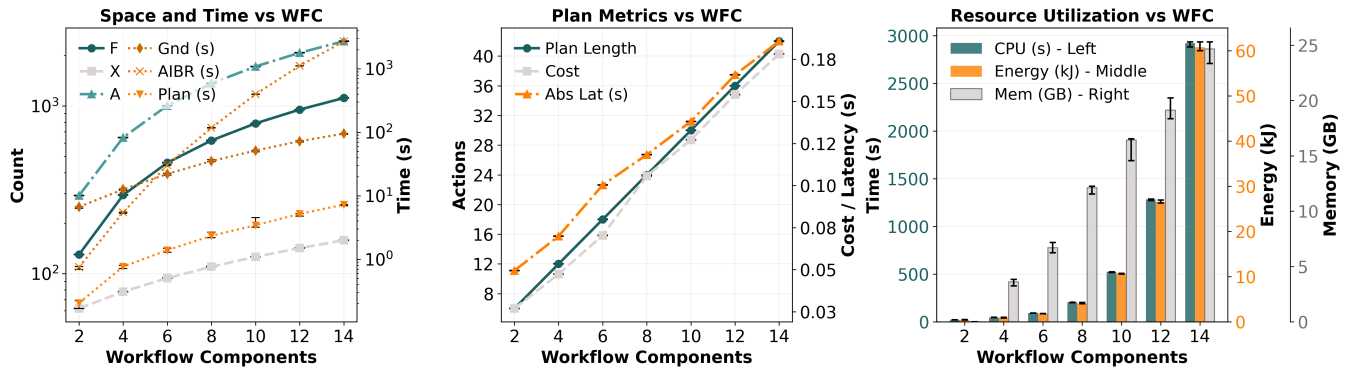


Figure 6: Planning and scheduling variable length linear-chain workflows on a complex resource graph (8S, 16IF, 15DL, 14CL).

this step is a natural future research direction.

Figure 6 addresses RQ2. ENHSP with WORKSWORLD can serve as a centralized planner for *linear-chain* workflows at a realistic resource graph size. The sub-exponential growth across all metrics, with the exception of AIBR preprocessing, supports this conclusion. The resource graph (8S, 16IF, 15DL, 14CL) represents a realistic mid-sized enterprise with two cloud, two fog, and four edge sites. However, the linear-chain assumption limits realism; we leave non-linear series-parallel workflows for future work.

## 8 Related Works

A broad body of literature addresses workflow scheduling and planning. We focus on work most similar to our pipelined-workflow approach using AI planning techniques.

First, we discuss available planning domains for compute resource planning and scheduling. Seipp et al. (2022) provide an online repository of existing domains. The *Data-network* domain proves most similar to WORKSWORLD, modeling servers with different RAM capacities that receive, process and send files across a network. Like WORKSWORLD, goals declare desired data formats on specific servers. However, this domain uses PDDL 2.2 level 1 (non-numeric planning) with PDDL 3.1 action costs, while WORKSWORLD fully supports numeric planning for more realistic resource modeling. Numeric planning domains also exist from the International Planning Competition numeric track<sup>12</sup>, though none closely resemble WORKSWORLD.

Second, we discuss works leveraging AI planning for similar problems. Arshad et al. 2003 develop Planit, which uses the LPG planner (Agnew, Hofmeister, and Purtilo 1994) to dynamically reconfigure software systems by scheduling components and point-to-point connectors, whereas we provide DSIs for multi-component data sharing. Kichkaylo et al. 2003 use node and link-level abstractions to place components for applications (secure mail, webcast) by compiling placement into planning problems. Riabov et al. 2005 solve workflow planning by declaring input and output streams but define a PDDL extension (Stream Processing Planning Language) that prevents testing with domain-independent

planners. Shorabi et al. 2013 use hierarchical task network (HTN) planning for workflows in vendor-specific platforms, exploring optimal *planning* of graphs with 30–170 components; we explore satisficing planning and *scheduling* on resources in a vendor-agnostic manner. Most recently, Amado et al. 2025 group data pipeline tasks into Kubernetes pods to minimize execution time for a vendor-specific application. Compared to these works, we operate at a higher abstraction level, reasoning over sites containing resources and generating plans in a vendor and platform-agnostic manner.

## 9 Conclusion

We introduced WORKSWORLD, a numeric domain for integrated planning and scheduling of permanent workflows. We showed empirically that a state-of-the-art planner can solve the joint planning and scheduling problem we defined in reasonable time and space for simple workflows on complex resource graphs. The domain and our other engineering artifacts are available publicly to benefit the AI planning community. Beyond future work mentioned earlier, we intend to explore leveraging state-dependent action costs in the domain to permit greater cost penalty when scheduling on high-utilization instances. This will require normalizing work scheduled with our state-dependent `resource_available` fluent. Supporting state dependent action costs will also permit a more conservative cost estimate based on actual rates instead of max rates for workflow components. Second, we posit that increasing the types of DSIs modelled in WORKSWORLD could permit interesting tradeoffs for a planner. For example, the planner may schedule latency-sensitive components on memory-based interfaces at higher cost, and latency-tolerant components on disk-backed interfaces at lower cost. Finally, our workflow and resource graph facilitate easy division of work that may enable a hierarchy of planners to collaboratively solve the scheduling portion of the problem. We envision a high-level planner that produces the concrete workflow graph and delegates to each site’s low-level planner to schedule and connect its assigned components across available interfaces. This approach may prove more efficient for problems with complex resource and/or workflow graphs.

<sup>12</sup><https://github.com/ipc2023-numeric/ipc2023-dataset>

## References

- Agnew, B.; Hofmeister, C.; and Purtilo, J. 1994. Planning for change: A reconfiguration language for distributed systems. *Distributed Systems Engineering*, 1(5): 313.
- Altintas, I.; Perez, I.; Mishin, D.; Trouillaud, A.; Irving, C.; Graham, J.; Tatineni, M.; Defanti, T.; Strande, S.; Smarr, L.; and Norman, M. 2022. Towards a Dynamic Composability Approach for using Heterogeneous Systems in Remote Sensing. In *2022 IEEE 18th International Conference on e-Science (e-Science)*. IEEE.
- Amado, L. R.; Vogel, A.; Griebler, D.; Licks, G. P.; Simon, E.; and Meneguzzi, F. 2025. Automated Planning for Optimal Data Pipeline Instantiation. *CoRR*, abs/2503.12626.
- Amaral, M.; Chen, H.; Chiba, T.; Nakazawa, R.; Choochotkaew, S.; Lee, E. K.; and Eilam, T. 2023. Kepler: A Framework to Calculate the Energy Consumption of Containerized Applications. In *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*, 69–71.
- Arshad, N.; Heimigner, D.; and Wolf, A. L. 2003. Deployment and dynamic reconfiguration planning for distributed software systems. In *Proceedings. 15th IEEE International Conference on Tools with Artificial Intelligence*, 39–46. IEEE.
- Benoit, A.; Çatalyürek, Ü. V.; Robert, Y.; and Saule, E. 2013. A survey of pipelined workflow scheduling: Models and algorithms. *ACM Computing Surveys (CSUR)*, 45(4): 1–36.
- Bylander, T. 1994. The computational complexity of propositional STRIPS planning. *Artificial Intelligence*, 69(1-2): 165–204.
- Cashmore, M.; Fox, M.; Long, D.; and Magazzeni, D. 2016. A compilation of the full PDDL+ language into SMT. In *Proceedings of the international conference on automated planning and scheduling*, volume 26, 79–87.
- Challapally, A.; Pease, C.; Raskar, R.; and Chari, P. 2025. The GenAI Divide: State of AI in Business 2025. Technical report, MIT NANDA. Preliminary findings from Project NANDA.
- Chen, D.; and Thiébaux, S. 2024. Graph learning for numeric planning. *Advances in Neural Information Processing Systems*, 37: 91156–91183.
- Da Silva, R. F.; Bard, D.; Chard, K.; De Witt, S.; Foster, I. T.; Gibbs, T.; Goble, C.; Godoy, W.; Gustafsson, J.; Haus, U.-U.; et al. 2024. Workflows community summit 2024: Future trends and challenges in scientific workflows. *arXiv preprint arXiv:2410.14943*.
- Duplyakin, D.; Ricci, R.; Maricq, A.; Wong, G.; Duerig, J.; Eide, E.; Stoller, L.; Hibler, M.; Johnson, D.; Webb, K.; Akella, A.; Wang, K.; Ricart, G.; Landweber, L.; Elliott, C.; Zink, M.; Cecchet, E.; Kar, S.; and Mishra, P. 2019. The Design and Operation of CloudLab. In *Proceedings of the USENIX Annual Technical Conference (ATC)*, 1–14.
- Dustdar, S.; Pujol, V. C.; and Donta, P. K. 2023. On Distributed Computing Continuum Systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(4): 4092–4105.
- Edelkamp, S.; and Hoffmann, J. 2004. PDDL2. 2: The language for the classical part of the 4th international planning competition. Technical report, Technical Report 195, University of Freiburg.
- Fischer, A.; Botero, J. F.; Beck, M. T.; De Meer, H.; and Hesselbach, X. 2013. Virtual network embedding: A survey. *IEEE Communications Surveys & Tutorials*, 15(4): 1888–1906.
- Foster, I.; Kesselman, C.; and Tuecke, S. 2001. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *The International Journal of High Performance Computing Applications*, 15(3): 200–222.
- Fox, M.; and Long, D. 2003. PDDL2. 1: An extension to PDDL for expressing temporal planning domains. *Journal of artificial intelligence research*, 20: 61–124.
- Gerevini, A.; and Long, D. 2005. Plan constraints and preferences in PDDL3. Technical report, Technical Report 2005-08-07, Department of Electronics for Automation . . . .
- Gerevini, A.; Saetti, A.; and Serina, I. 2006. An approach to temporal planning and scheduling in domains with predictable exogenous events. *Journal of Artificial Intelligence Research*, 25: 187–231.
- Ghallab, M.; Howe, A.; Knoblock, C.; McDermott, I. D.; Ram, A.; Veloso, M.; Weld, D.; Sri, D. W.; Barrett, A.; Christianson, D.; et al. 1998. PDDL — The Planning Domain Definition Language Version 1.2. *Technical Report, Tech. Rep.*
- Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: theory and practice*. Elsevier.
- Goedegebuure, A.; Kumara, I.; Driessen, S.; Van Den Heuvel, W.-J.; Monsieur, G.; Tamburri, D. A.; and Nucci, D. D. 2024. Data mesh: a systematic gray literature review. *ACM Computing Surveys*, 57(1): 1–36.
- Haslum, P.; Lipovetzky, N.; Magazzeni, D.; Muise, C.; Brachman, R.; Rossi, F.; and Stone, P. 2019. *An introduction to the planning domain definition language*, volume 13. Springer.
- Heck, P. 2024. What About the Data? A Mapping Study on Data Engineering for AI Systems. In *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering - Software Engineering for AI, CAIN 2024*, 43–52. ACM.
- Hoffmann, J. 2003. The Metric-FF Planning System: Translating “Ignoring Delete Lists” to Numeric State Variables. *Journal of artificial intelligence research*, 20: 291–341.
- Hou, C.-Y.; Altintas, I.; Jaeger-Frank, E.; Gilbert, L.; Moore, R.; Rajasekar, A.; and Marciano, R. 2006. A scientific workflow solution to the archiving of digital media. In *2006 Workshop on Workflows in Support of Large-Scale Science*, 1–10.
- Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence*, 294–301. IEEE.
- Kichkaylo, T.; Ivan, A.; and Karamcheti, V. 2003. Constrained component deployment in wide-area networks using AI planning techniques. In *Proceedings International*

*Parallel and Distributed Processing Symposium*, 10–pp. IEEE.

Kreuzberger, D.; Kühl, N.; and Hirschl, S. 2023. Machine learning operations (mlops): Overview, definition, and architecture. *IEEE access*, 11: 31866–31879.

Król, M.; Mastorakis, S.; Oran, D.; and Kutscher, D. 2019. Compute First Networking: Distributed Computing meets ICN. In *Proceedings of the 6th ACM Conference on Information-Centric Networking, ICN '19*, 67–77. ACM.

Long, D.; Dolejsi, J.; and Stolba, M. 2023. Scheduling problems in PDDL. In *Workshop on Knowledge Engineering for Planning and Scheduling*.

Mell, P. M.; and Grance, T. 2011. *The NIST definition of cloud computing*. National Institute of Standards and Technology.

Moreschini, S.; Pecorelli, F.; Li, X.; Naz, S.; Hastbacka, D.; and Taibi, D. 2022. Cloud Continuum: The Definition. *IEEE Access*, 10: 131876–131886.

Paul, T.; and Regli, W. 2024. Accelerating the Operation of Complex Workflows through Standard Data Interfaces. *arXiv preprint arXiv:2412.13339*.

Riabov, A.; and Liu, Z. 2005. Planning for stream processing systems. In *AAAI*, 1205–1210.

Scala, E.; Haslum, P.; Thiébaux, S.; and Ramirez, M. 2016. Interval-based relaxation for general numeric planning. In *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, 655–663.

Scala, E.; Saetti, A.; Serina, I.; and Gerevini, A. E. 2020. Search-Guidance Mechanisms for Numeric Planning Through Subgoal Relaxation. In *Proceedings of the 30th International Conference on Automated Planning and Scheduling (ICAPS)*, 226–234. AAAI Press.

Seipp, J.; Torralba, Á.; and Hoffmann, J. 2022. PDDL Generators. <https://doi.org/10.5281/zenodo.6382173>.

Sohrabi, S.; Udrea, O.; Ranganathan, A.; and Riabov, A. 2013. HTN planning for the composition of stream processing applications. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 23, 443–451.

Suter, F.; Da Silva, R. F.; Gainaru, A.; and Klasky, S. 2023. Driving Next-Generation Workflows from the Data Plane. In *2023 IEEE 19th International Conference on e-Science (e-Science)*, 1–10.

Versluis, L.; and Iosup, A. 2021. A survey of domains in workflow scheduling in computing infrastructures: Community and keyword analysis, emerging trends, and taxonomies. *Future Generation Computer Systems*, 123: 156–177.

Vivas, A.; Tchernykh, A.; and Castro, H. 2024. Trends, Approaches, and Gaps in Scientific Workflow Scheduling: A Systematic Review. *IEEE Access*, 12: 182203–182231.

Zhang, L.; Afanasyev, A.; Burke, J.; Jacobson, V.; Claffy, K.; Crowley, P.; Papadopoulos, C.; Wang, L.; and Zhang, B. 2014. Named data networking. *ACM SIGCOMM Computer Communication Review*, 44(3): 66–73.