

Automated Planning for Optimal Data Pipeline Instantiation

Leonardo Amado^{5,1}, Adriano Vogel^{1,2}, Dalvan Griebler¹, Gabriel Paludo Licks³, Eric Simon⁴ and Felipe Meneguzzi^{5,1}

¹Pontifical Catholic University of Rio Grande do Sul, Brazil

²Johannes Kepler University Linz, Austria

³Sapienza University of Rome, Italy

⁴SAP Labs, France

⁵University of Aberdeen, Scotland, UK

Abstract

Data pipeline frameworks provide abstractions for implementing sequences of data-intensive transformation operators, automating the deployment and execution of such transformations in a cluster. Deploying a data pipeline, however, requires computing resources to be allocated in a data center, ideally minimizing the overhead for communicating data and executing operators in the pipeline while considering each operator's execution requirements. In this paper, we model the problem of optimal data pipeline deployment as planning with action costs, where we propose heuristics aiming to minimize total execution time. Experimental results indicate that the heuristics can outperform the baseline deployment and that a heuristic based on connections outperforms other strategies.

Keywords

Automated Planning, Data Pipelines, Pipeline optimization

1. Introduction

A recent report from the European Commission forecasts an increase to 8.4 million data professionals in the EU27 by 2025, with 1.8 million positions to be added in the period between 2020 and 2025, leading to a shortage of approximately 484,000 positions [1]. The shortage of highly trained human resources calls for data processing platforms accessible to a broader category of users (the so-called citizen data scientists) having strong data expertise, moderate statistics training, and low expertise in programming and distributed computing [2, 3]. Several recent technology trends acknowledge this situation through the development of data pipeline frameworks [4, 5, 6, 7, 8, 9] relying on a flow-based programming model [10], in which the data pipelines are implemented as directed graphs where nodes represent computational operators that transform data and require diverse runtimes, and edges represent the flow of data between these operators.

In these frameworks, the goal is to enable non-technical experts (aka citizen data scientists) to quickly build, debug, and maintain high-quality data pipelines at scale within a cloud cluster by providing high-level abstractions in their programming model. Since users design data pipelines, and are not assumed to comprehend the underlying engine that deploys and executes the pipeline, the resulting data pipelines derive a non-optimal allocation of resources to execute it. Thus, data pipeline engines must provide optimization techniques to assure an optimal usage of computing resources in a distributed system and minimize data-intensive communication costs during graph execution.

The problem of optimizing a pipeline execution is similar to the task of constraint optimization [11], which has been a focus of research on automated planning. Indeed, research in that area has focused on developing algorithms to compute optimized solutions for optimization problems with constraints, leading to fast and reliable algorithms [12]. With such a research background, we develop a model

CAIPI@ECAI'25: ECAI Workshop on AI-based Planning for Complex Real-World Applications, 25th October, 2025, Bologna, Italy

*Corresponding Author. Email: leonardo.amado@abdn.ac.uk.

**Work done while affiliated with the Pontifical Catholic University of Rio Grande do Sul (PUCRS).



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

for automated planning algorithms that use heuristic search and state-dependent action costs to solve the problem of optimally allocating computation tasks to the worker nodes of a distributed system for executing a data pipeline. We chose to formalize the problem in terms of automated planning (instead of, e.g., a discrete optimization problem) partly to showcase recent advances on numeric planning, and partially to assess the limits of current heuristics and planning systems on that space. We evaluate our optimization methods using the pipeline engine of the SAP Data Intelligence commercial product ¹, called VFlow ², which adopts a flow-based programming model, and whose execution model is based on a distributed Kubernetes container orchestration platform.

This work provides two main technical contributions: I) We model several optimization methods that determine how the operators of a data pipeline should be deployed on the worker nodes of Kubernetes clusters; II) We perform an experimental analysis of the proposed approaches with a real-world industry data pipeline processing system, report our findings and indicate directions for future performance improvement.

We organize this article as follows. Section 2 provides an overview of automated planning and data pipelines. In Section 3, we model the elements and constraints of deploying data pipelines and computing plans. Then, in Section 4 we describe the approaches to pipeline grouping for reducing the total execution time. Section 5 explains our experimental methodology that compares the results of the heuristic approaches against solutions from a random baseline approach on various workloads. Section 6 shows our empirical performance results for the different approaches using pipelines with a sequential or parallel topology in which we vary the number of special operators and the computational load. Finally, Section 8 summarizes our key contributions and points toward future work.

2. Background

This section introduces the formal background upon which our application relies. We start with a formalization of Automated Planning and the elements of problems in this formalism in Section 2.1. Section 2.2 then formalize data pipelines and how their definition relates to deployment in an underlying high performance infrastructure. This allows us to formalize the pipeline optimization problem of Section 3.

2.1. Automated Planning

We leverage automated planning techniques to drive the optimization. Automated planning is a sub-area of Artificial Intelligence concerned with finding sequences or policies of actions (*i.e.*, plans) able to transition an agent from a given initial state to a desirable goal state [13, 14, Ch. 1]. A classical planning problem can be described as one of graph search, in which the nodes in the graph represent states, edges represent transitions between states that are caused by applying an action at a state. The solution to a planning problem is a sequence of actions (edges) that forms a path to traverse this graph from a unique initial state to any one of the goal states. Domain experts seldom define such graphs explicitly, but rather use a logic-based description language (which we define later) to induce the actual search graph. This allows graphs with huge state-spaces to be defined compactly. In what follows we adopt the planning terminology from Ghallab *et al.* [15] to represent states and actions in planning domain problems.

The most fundamental part of a classical planning task is the planning domain. A planning domain is a formal description of the dynamics of an environment in which an agent acts. Definition 1 formalizes a planning domain.

Definition 1 (Planning Domain). *A planning domain Ξ consists of a pair $\langle \mathcal{F}, \mathcal{A} \rangle$, which specifies the knowledge of the domain, and comprises: a finite set of facts \mathcal{F} , *i.e.*, a set of ground instantiated predicates,*

¹<https://www.sap.com/products/technology-platform/data-intelligence.html>

²<https://api.sap.com/api/vflow/overview>

defining the environment state properties; and a finite set of actions \mathcal{A} , which is technically a set of ground instantiated operators, representing the actions that can be performed in the environment.

States $s \subseteq \mathcal{F}$ consist of facts from a planning domain indicating properties that are true at any moment in time and follow the closed world assumption so that any fact not included in a state is false. Conditions or formulas in our formalism comprise positive and negated facts ($f, \neg f$) representing an implicit conjunctive formula indicating what must be true (alternatively, false) in a state. The positive part of a condition c ($\text{pos}(c)$) comprises the positive facts, and the negated part of a condition $\text{neg}(c)$ comprises the negated facts. We say a state s supports a condition c , $s \models c$ (alternatively, c is valid in s), if and only if all positive facts are present in s , and all negated facts are absent in s , i.e., $s \models c$ if and only if $(s \cup \text{pos}(c) = s) \wedge (s \cap \text{neg}(c) = \emptyset)$. An action $a \in \mathcal{A}$ is represented by a tuple $o = \langle \text{pre}(a), \text{eff}(a), \text{cost}(a) \rangle$ containing the preconditions $\text{pre}(a)$, the effects $\text{eff}(a)$, and a non-negative cost $\text{cost}(a)$, indicating possible transitions between states.

The transition of a state s into a new state s' using an action a is represented as the transition function $s' = \gamma(s, a)$. The transition is valid iff $s \models \text{pre}(a)$, and $s' = (s \cup \text{pos}(\text{eff}(a))) - \text{neg}(\text{eff}(a))$.

A planning task (or planning instance) is the formal description of a task to be solved in a given planning domain. Planning tasks comprise a planning domain Ξ and a planning problem. A planning problem describes the finite set of objects of the environment, the initial state from which the planning problem starts, and the goal state which an agent desires to achieve.

Definition 2 (Planning Task). *A planning task is a tuple $\Pi = \langle \Xi, \mathcal{I}, G \rangle$, in which \mathcal{I} is an initial state, G is a goal condition, and Ξ is a planning domain.*

The solution for a planning task, which we call a *plan*, is a sequence of actions that an agent can perform in a planning domain to achieve the goal state G from the initial state \mathcal{I} . Definition 3 formalizes the notion of a plan and an optimal plan.

Definition 3 (Plan). *A plan π for a planning task Π is a sequence of actions $\pi = \langle a_1, \dots, a_n \rangle$ that induces a sequence of states $\langle s_0, s_1, \dots, s_n \rangle$ such that $\mathcal{I} = s_0 \models \text{pre}(a_1)$, $s_n \models G$ and that every state $s_i \in \pi$ is such that $s_{i-1} \models \text{pre}(a_i)$ and $s_i = \gamma(s_{i-1}, a_i)$. The cost of a plan is the sum of the cost all of its actions such that $\text{cost}(\pi) = \sum_{i=1}^n \text{cost}(a_i)$. We say a plan is optimal, denoted π^* , if its cost is minimal, that is, if no other plan π has $\text{cost}(\pi) < \text{cost}(\pi^*)$.*

Note that Definition 3 allows for multiple optimal plans, which may have the same cost, and are thus equivalently optimal. Modern planners use the Planning Domain Definition Language (PDDL) as a standardized domain and problem representation medium [16]. PDDL encodes a number of expressive planning formalisms, such as STRIPS-style [17] planning tasks. For simplicity, most planning formalizations assume that every action in \mathcal{A} has cost 1, so that the optimal plan is the plan with the smallest number of actions. Here, we define actions with varying costs that depend on the state of the environment, where the environment considered in this work relates to the configurations of a platform to run data pipelines.

2.2. Designing and deploying data pipelines

In the flow-based programming model [10], a data pipeline corresponds to a type of application that processes large amounts of data and produces results promptly. We refer to processing as the capacity of data pipelines to ingest, transform, and output data with high throughput and low latency [18]. At the design level, the data pipelines considered in this article are graphs in which nodes are operators and edges represent the flow of data, following the flow-based programming paradigm [10]. Operators³ are responsible for performing operations with data. Operators can be written using different programming

³Note that we explicitly make a distinction between *operators* (i.e., elements of the data pipeline formalism) and *actions* (i.e., state transformation functions in the planning formalism)

languages using a Software Development Kit (SDK) provided by the data pipeline framework, formalized as Definition 4.

Definition 4 (SDK). *Let \mathcal{S} be the set of all SDKs made available by the pipeline engine. An SDK $S \in \mathcal{S}$ provides a programming environment and runtime specific to a programming language, with tools such as an API for the user to program operators and interact with the pipeline framework. Thus, the data pipeline framework can provide SDKs for developing operators in Java, Python, Go, etc.*

Implementations of operators can differ in requirements that need to be satisfied by the runtime system in order to be executable, and *tags* associated with an operator encode these requirements. Definition 5 formalizes the notion of a tag.

Definition 5 (Tag). *Let \mathcal{T} be the set of all possible tags defined by the user. A tag $t \in \mathcal{T}$ denotes a requirement that has to be satisfied for an operator to be executable, i.e. a specific library or package (possibly a specific version of each).*

Definition 6 brings together Definitions 4 and 5 into the formal operator within a pipeline.

Definition 6 (Operator). *A pipeline operator $O = \langle id, S, T \rangle$ comprises an identification id , the SDK $S \in \mathcal{S}$ it is developed with, and a set of tags $T \subseteq \mathcal{T}$ that it requires in order to execute.*

The runtime environment in which an operator must be executed is set by creating a Docker image (later called OS image) [19], which is a lightweight way to package code, run-time and all dependencies, and by associating tags to the OS image that reflect the characteristics of the environment. Definition 7 formalizes the notion of an image in terms of the tag requirements from Definition 5.

Definition 7 (Image). *An image $I \subseteq \mathcal{T}$ is a set of tags that the image specification supports, representing a possible grouping for operators' tags satisfied by the image.*

A designer of a data pipeline can then tag an operator to indicate the runtime environment required for its execution. By default, when deploying a data pipeline for execution, the execution framework searches for an OS image whose tags match all the required tags of the operators in the data pipeline. In this case, all operators become part of a default “group”. Operators in a group are deployed as a container running in a Kubernetes *Pod* [20]: they share the libraries and packages available in the environment defined by the containerized OS image. Definition 8 formalizes such matching of tags to images as a group.

Definition 8 (Group). *Let \mathcal{G} be the set of groups in a data pipeline and \mathcal{O} be the set of operators in a data pipeline. A group $G \in \mathcal{G}$ then is a tuple $\langle O_G, I \rangle$, where $O_G \subseteq \mathcal{O}$ is the set of operators in the group and $I \in \mathcal{I}$ is the image associated to the group. Since all operators have to be grouped, explicitly or implicitly, the intersection between operators in each group has to be empty, i.e. $O_{G_i} \cap O_{G_j} = \emptyset$ for all $G_i, G_j \in \mathcal{G}$ with $i \neq j$.*

However, when the execution framework cannot find a single matching OS image, the user must manually group operators according to their required tags until the remaining non grouped operators can be matched by a single OS image. Alternatively, in order to have control over which requirements a user needs for specific tasks, users can manually group operators and create an OS image specific to the grouped operators. Finally, we bring together the notion of an undeployed Data Pipeline incorporating the elements from Definitions 4–6 in Definition 9.

Definition 9 (Data Pipeline). *Let $E = \langle O_i, O_j \rangle$ be an edge tuple representing the directed connection between two operators O_i and O_j . A data pipeline then is a tuple $\mathcal{D} = \langle \mathcal{G}, \mathcal{E} \rangle$, containing the specified groups and edges between all operators.*

Figure 1 exemplifies a data pipeline with operators grouped using different requirements, specified by tags. Here, the data pipeline is executable if the runtime framework has at least two images I_1 and I_2 , such that $\langle t_1, t_2, t_3, t_5 \rangle \subseteq I_1$ and $\langle t_1, t_3, t_4 \rangle \subseteq I_2$.

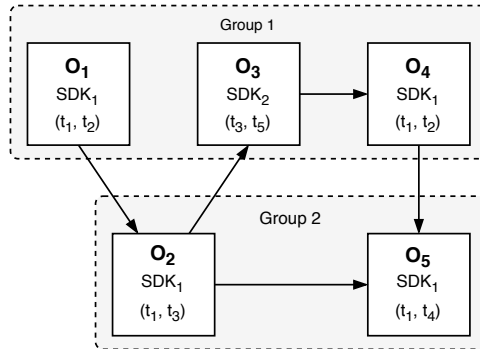


Figure 1: Data pipeline illustration at design level

3. The pipeline optimization problem

At design level, data pipelines are graphs in which nodes correspond to operators and edges represent the flow of data. There may be multiple concrete implementations of data pipelines⁴ and many open source systems do, including Flink [21], Storm and others. In this work, we focus specifically on the implementation provided in the VFlow engine where operators perform computing operations with data. Operators in VFlow can be written using different programming languages using SDKs provided by the data pipeline framework.

Given a data pipeline \mathcal{D} from Definition 9, with a set of connected operators, we must find the optimal grouping configuration, in terms of total execution time, which satisfies the constraints imposed by the operator tags T . It is therefore an optimization problem under constraints, where the constraints are the tags of operators, which must match the tags of an OS image to belong to the same group.

The runtime environment in which operators execute can be set by creating groups of operators. These runtime environments are executing on container images (e.g., Docker images) where each group ideally has the same image associated to it. This allows the operators of the same group share resources, libraries, and packages available in that environment. In order to have control over which requirements a user needs for specific tasks, Users can manually group operators and create an image specific to the grouped operators. In the default group, VFlow associates available images and stock images that are tailored to the basic requirements of each SDK. In case operators are not explicitly grouped by users they belong to an implicit group such that the pipeline framework automatically associates an image to it using the available images and stock images that are tailored to basic requirements of each SDK.

4. Automating operator grouping

In this section, we detail three optimization methods: two greedy approaches and one approach based on deriving specific costs. As a baseline, we define a random approach that optimizes the pipeline by selecting random images and adding random operators. We use PDDL to create descriptions of state-space search problems involving a variety of domains independently of the underlying AI search engine, enabling us to solve the problem of Definition 2. Such flexibility, coupled with the number of efficient planning software available [22, 23], allows us to prototype a number of optimization strategies for grouping in order to achieve desirable deployment properties.

In order to generate an optimized graph, we provide two such specifications to a PDDL planner such as ENHSP [24, 25], and the resulting plan can then be translated into an optimized graph. We illustrate this workflow in Figure 2.

⁴We refer to data pipelines as graphs and use the term interchangeably.

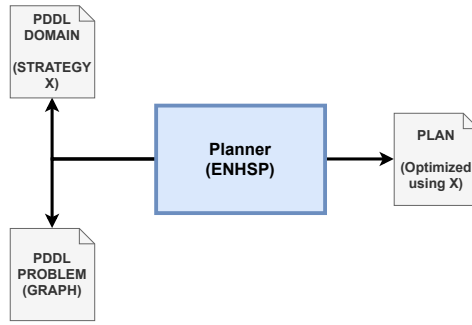


Figure 2: Planner workflow.

4.1. Connection heuristic

The *connection* greedy approach prioritizes the creation of groups with a minimum number of intergroup communication links. To encode this approach using PDDL, we use different weights for intragroup and intergroup communication, respectively 5 and 20. We assign a medium cost for instantiating groups, since creating many groups will lead to more intergroup connections. Figure 3a exemplifies a solution using this approach, where the number of intergroup communication links is only one and two operator groups are created.

4.2. Node heuristic

The *node* approach prioritizes the creation of groups with as many nodes as possible, greedily searching for groups with a larger number of operators. To encode this approach using PDDL, we use the same weights for intragroup and intergroup communication, respectively 5. We assign a very high cost for instantiating groups, which forces the planner to search for configurations with a very small number of groups. Figure 3b illustrates an example of a solution using this heuristic. As we can see, this solution ignores the number of intergroup connections, leading to two groups with many intergroup connections.

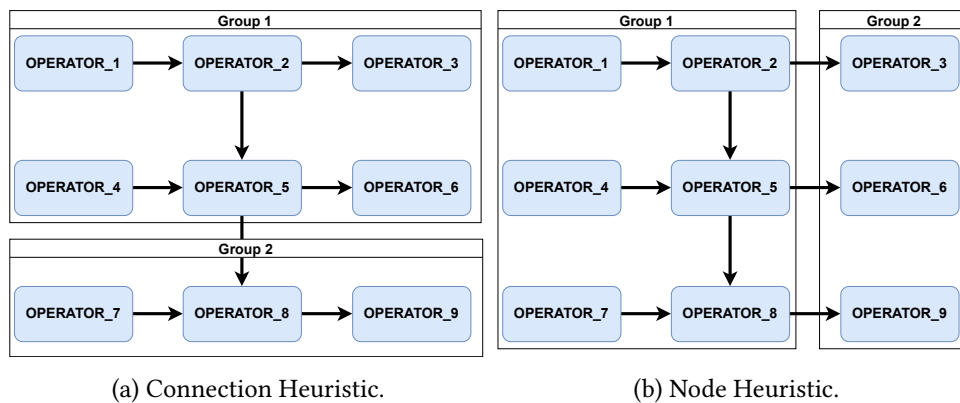


Figure 3: Proposed heuristics.

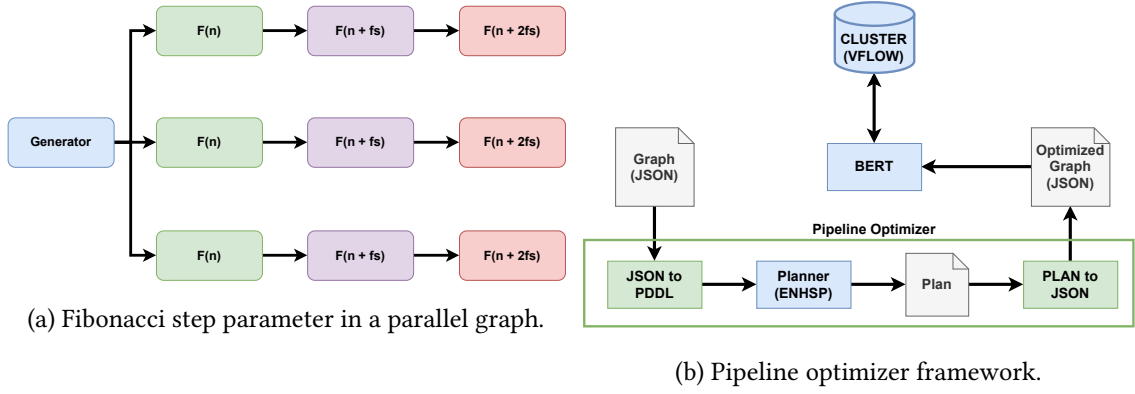


Figure 4: Fibonacci and pipeline optimizer.

4.3. Random baseline

The random baseline approach groups operators by randomly selecting an image from the list of available images and then grouping all operators satisfied by the selected image, if there are any. This algorithm randomly selects images and groups operators until no operators are left to be grouped. This process of randomly selecting groups works as follows. First, we randomly select an available image capable of executing operators based on the available tags of this image. Second, we select all n operators that can run using this image. Third, if the number of operators is greater than 0, we randomly compute a number from the range $[1, n]$, which defines how many of these possible operators our new group will contain. Finally, we generate a new group with the selected operators and repeat this process until all operators are assigned to a group.

5. Evaluation methodology

The evaluation of our solution requires that we design different examples of data pipelines and generate a synthetic workload that incur substantive computing. We instantiate data pipelines following the two most common types of patterns that users design: sequential and parallel executions. As a workload choice, we choose the computation of Fibonacci sequences, which allows us to parameterize the amount of work that needs to be computed by an operator by defining the size of the sequence. We detail these in the following sections.

5.1. Workload configuration

In order to evaluate the grouping methods, we use two patterns of synthetic data pipelines built from the same operator which computes the n^{th} number of the Fibonacci sequence and sends it to the next operator. This operator allows us to control the processing time of a data pipeline. The two patterns of data pipelines, illustrated in Figure 5, are as follows:

1. a topology representing sequential computation steps;
2. a topology representing independent and concurrent paths of sequential computation steps.

We define two parameters to generate pipeline optimization problems with varying degrees of complexity. These parameters are: **number of special tags** (later called *special_ops*) and **Fibonacci step**.

The **special tag** represents operators with specific library requirements beyond what is included in the standard images associated with operators in each sub-engine. Those necessitate specific images to

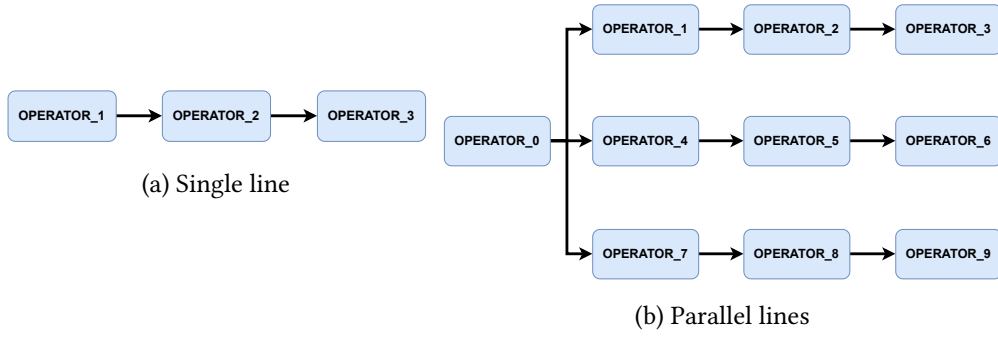


Figure 5: Graph topologies used in evaluation.

be run in VFlow. By default, all operators have default "golang" tags, and all images have the default tag. In addition to the default tag "golang", we create three distinct special tags (*spt-1*, *spt-2*, and *spt-3*), where each special tag belongs to a single specific image. When the parameter "special_ops" has the value n , then n randomly selected operators will receive a special tag. If $n \leq 3$, then all the operators with special tags will have distinct tags between (*spt-1*, *spt-2*, *spt-3*). Otherwise, two operators can have the same special tag. For example, if $n = 4$, two operators will have the special tag *spt-1*.

The goal of the **special tag** is to increase the complexity of the optimization problem, mirroring real deployment scenarios where operators require specific libraries (e.g., machine learning frameworks, specialized data processing tools) that cannot coexist in a single image due to dependency conflicts or size constraints. By default, all generated pipelines already contain one special tag to ensure the problem cannot be trivially solved by including all operators in a single group. As the number of special tags increases, the problem of grouping the operator becomes more difficult.

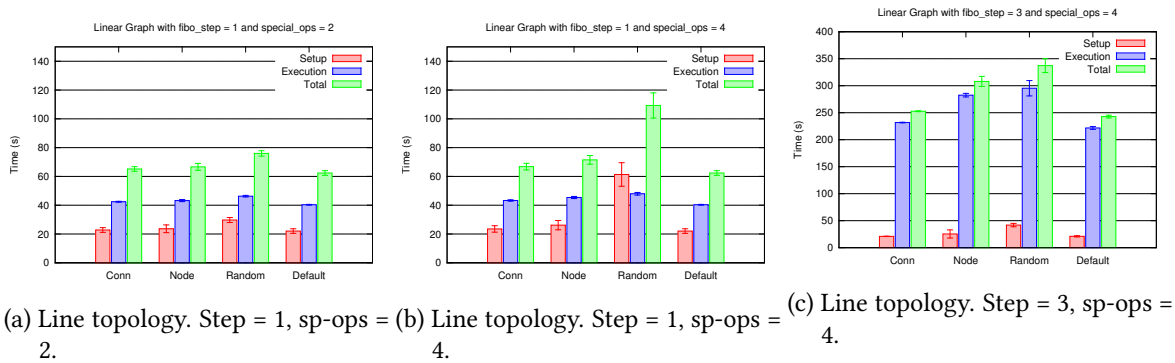


Figure 6: Experiment results for Line topology.

The **Fibonacci step** is an application parameter related to increasing the n values. The n value starts as 5 (the fifth element of the Fibonacci sequence) and the **Fibonacci step** ranges from 2 to 4. As this parameter increases, the overall execution time of the graph increases steeply. Figure 4a illustrates the propagation of the Fibonacci step parameter in a parallel topology, where $F(n)$ is the algorithm that computes the n th element of the sequence, and fs is the *Fibonacci step* parameter. The same logic of an individual line of the parallel topology is applied to the linear topology pipeline.

The pipelines consist of 14 operators, where 12 are computing operators, one is a generator, and one is a terminator. The first operator is a *generator*, which sets up the workload for the other operators. The generator sends a message to all connected operators, sending the n (5 as default) value for the




Figure 7: Parallel topology experiments

Fibonacci operator. This message is sent 5 times, to ensure continuous communication throughout the execution of the operator. The generator is connected with a *terminator* operator, which signals the end of the execution of the pipeline once all other operators finish. Finally, the other twelve operators are Fibonacci operators, which receive a message with the n value and compute the n th element of the Fibonacci sequence. After computing the sequence, each operator sends a message to the next one with $n + \text{FibonacciStep}$, which will set up the execution of the next operator.

5.2. Experiment Execution

To optimize data pipelines, we set up a framework to measure our experiments, which comprises the following processes:

- a parser to convert vFlow graphs into a planning task in PDDL;
- a **numeric planner** to compute a grouping strategy to optimize the parsed pipeline, yielding a PDDL-based plan; and
- a converter to interpret this PDDL plan into a vFlow graph, which is executed in an installation of VFlow and extracts performance metrics, e.g., execution time.

VFlow includes a profiling tool to evaluate the performance of the designed approaches, which allows us to collect data on three metrics: *setup time*, *execution time*, and *total time*. The *setup time* measures how long it took to prepare the environment; *execution time* tracks how long the pipeline took to execute after the setup phase. The *total time* is the setup time plus the execution time.

Figure 4b illustrates the architecture of our approach. VFlow exports data pipeline graphs in JSON format, which serves as the medium of communication between VFlow and our optimizer. A graph described in JSON is sent to a converter, which converts JSON to a planning description language, PDDL. The PDDL has two files, the domain, and problem. The domain is where we define the optimization strategy (the ones we detailed in Section 5.1), dictating how we optimize the given pipeline. The problem comprises the details of the given graph, detailing the operators (as black boxes) that are going to run, how they are connected forming the topology of the given graph, and the necessary tags to run each operator. With these files we execute the planner, ENHSP [23], which outputs a plan. The plan is the optimization steps, which define the number of groups and which specific images are going to be used to optimized the graph. We then use a converter to convert the plan (sequence of optimization steps) to a JSON description of the optimized graph, which runs using the VFlow profiler.

6. Experimental results

In the following, we show the quantitative results we obtained from running our experimental evaluation. We first describe the experimental environment in which we deployed the data pipelines, and follow by

a description of the results obtained for each topology of graphs that we generated, i.e. line and parallel.

6.1. Experimental setup

The experiments are executed in a virtualized cloud cluster composed of two worker nodes. Docker containers using the Centos-72 image run within the worker nodes, where the containers are organized in pods managed by Kubernetes (version 1.19.13).

To evaluate our approaches, we run each graph five times to account for variations in the cluster's performance. We distinguish the measurements of the first execution (cold start) from the other executions (warm starts). As metrics, we compute the mean setup time, execution time, and total time discussed in the previous section.

We compare our approaches (connection, node, and random) to the *default* baseline that has no grouping assuming that all operators are run in a single group. For line topology experiments, it represents an optimum grouping, as all operators are sequential and in the same image. For parallel topology experiments, the default baseline runs all operators in a single group, which prevents them to benefit from the parallelism of the topology. Note that it would be impossible to run the graphs with only one image due if the special tags were real software constraints, so the default baseline is not a viable strategy outside of this evaluation scenario. Considering this, we include it as a method for comparison, highlighting the effects of having no groups in different scenarios.

6.2. Line topology

Our first batch of experiments uses the line topology with sequential operators. We vary the parameters described in Section 5.1, running nine distinct baseline graphs, and optimizing each graph using all four approaches, resulting in 36 graphs. Figures 6a, 6b, 6c show the average results in milliseconds for all approaches, dividing these results into *Setup*, *Execution*, and *Total* time. Results from these experiments show that the heuristics provide performance gains, mainly in terms of setup time. Figures 6b and 6c show that the connection based heuristics also achieved performance gains in execution time. Such results indicate that planning heuristics are effective at optimizing stream processing applications.

6.3. Parallel topology

The experiments with parallel lines topology use three lines of four operators. We vary the parameters described in Section 5.1, running nine distinct baseline graphs with the same procedure we did for the line topology. Figures 7a, 7b, and 7c show the results obtained for three distinct combinations of parameters. Figure 7b shows the random approach performing well, even outperforming some heuristics in execution time. The random approach creates more groups, which increases the graph's setup time. Hence, the approaches' contrast in the total time is insignificant if we consider the standard deviation. However, as the execution time is often longer with higher workloads, in some cases, the execution speed-up can compensate the higher time for setup.

Figures 8a and 8b show the results for the first execution of two parameters configuration. As we can see, the setup time for the first execution is much higher, which impacts the performance of the random approach when the workload is small (Fibonacci step = 1). As the workload increases, execution time gains outweigh setup time losses.

7. Related Work

Researchers have applied automated planning to various stages of software engineering, including application synthesis, deployment, and optimisation of distributed systems. In stream processing and cloud-native environments, planning methods help manage component composition, resource allocation, and runtime adaptation. These approaches differ in formalism (classical planning, HTN models, preference-based search, and cost-sensitive optimisation) and target distinct problems such as




Figure 8: Parallel topology experiments, only first execution.

flow graph generation, operator placement, and service orchestration. In what follows, we illustrate these directions and provide context for our contribution.

Arshad et al. [26] develop Planit, a system that uses temporal planning to automate the initial deployment and dynamic reconfiguration of distributed software systems. Their planner, LPG, generates plans that minimise total execution time, using durative actions and explicit time annotations. While their work addresses deployment, it targets general-purpose component-based systems, focusing on scheduling actions under temporal constraints. In contrast, our work assumes a fixed data pipeline and optimises its deployment over a distributed infrastructure using numeric planning with state-dependent action costs. We do not model reconfiguration or time-based scheduling, and they do not optimise for execution efficiency of the deployed system beyond plan duration. The two approaches operate at different levels of abstraction and solve distinct optimisation problems.

Early work by Riabov and Liu [27] introduces SPPL, a domain-specific extension of classical planning tailored to stream processing composition. Their planner constructs workflows by selecting and connecting components to satisfy goal predicates, modelling streams as immutable entities with structured input/output ports and predicate propagation. This approach exposes structural regularities in stream processing domains that traditional PDDL encodings have struggled to exploit efficiently. Building on this foundation, Sohrabi et al. [28] develop a preference-based HTN planning framework that translates Cascade flow patterns into HTN domains, enabling more expressive composition and optimisation. Their work inherits the compositional focus of SPPL but replaces its custom formalism with a general-purpose HTN planner, allowing integration of user preferences and improved scalability. Unlike prior work on stream processing composition using Classical and HTN planning [27, 28], which focuses on synthesising pipeline structures from abstract flow patterns, our work assumes a fixed pipeline and addresses the orthogonal problem of optimally deploying it on a distributed infrastructure. While all three approaches leverage AI planning, we focus on numeric planning with state-dependent action costs to minimise execution time, rather than preference-based HTN planning for goal satisfaction.

Georgievski et al. [29] and Quenum and Josua [30] also focus on composition, applying HTN planning to generate cloud-ready applications from abstract component models. The former handles instance creation, duplication, and configuration processes using a formal model (Aeolus). By contrast, the latter generates hierarchical plans, matches actions to serverless functions across providers, and refines compositions using constraint satisfaction over QoS attributes. Unlike [28], they target cloud services rather than stream processing. Like the other work in this section, it applies planning to optimise structure or configuration. In contrast, the contribution of this paper optimises deployment.

8. Results discussion and closing remarks

The connection heuristic approach overall had better results. However, although with a high standard deviation, the random strategy with more groups has the best results for the parallel topology. We believe that this happens because the random approach created many groups that are capable of executing the processing steps in parallel. Moreover, the results show that more groups for the line topology do not improve the total execution time. The connection strategy, although consistently good in the line topology, does not consider possible parallel executions.

Our results show that the usage of AI techniques can improve the setup and execution time of pipelines by optimizing the grouping of operators. In this work, our optimization goal was the total execution time. In the future, other desirable evaluation metrics can be included, such as the infrastructure cost of running a graph (which includes the number of pods used) or the resilience of the deployed graph.

A hybrid strategy could also be explored to optimize each line of the parallel topology individually. In future work, we could evaluate grouping strategies based on TCO metrics that measure the total performance time per dollar (on a given cloud platform).

Acknowledgements

We are grateful for the financial support and computing resources from SAP Labs. We also thank the reviewers who pointed us to related work we were not familiar with at the outset of our project.

Declaration on Generative AI

During the preparation of this work, the author(s) used no Generative AI tools for any part of the writing, editing, or revision.

References

- [1] G. Cattaneo, G. M. Tochtermann, M. Glennon, C. La Croce, C. Mitta, The European Data Market Monitoring Tool, Luxembourg: Publications Office of the European Union, 2020.
- [2] B. Burns, Designing Distributed Systems: Patterns and Paradigms for Scalable, Reliable Services, O'Reilly Media, Inc., 2018.
- [3] M. Pi Puig, J. M. Paniago, S. Medina, S. Rodríguez Eguren, L. Libutti, J. Lanciotti, J. D. Antueno, C. Estrebou, F. Chichizola, L. D. Giusti, Intelligent distributed system for energy efficient control, in: Conference on Cloud Computing and Big Data, Springer, 2019, pp. 51–60.
- [4] A. Paleyes, C. Cabrera, L. Lawrence, Towards better data discovery and collection with flow-based programming, in: Proceedings of the 35th International Conference on Neural Information Processing Systems, Sydney, Australia, 2021.
- [5] T. Mahapatra, C. Prehofer, Graphical Flow-based Spark Programming, Journal of Big Data (2020).
- [6] H. Bergius, NoFlo: Flow-based programming for JavaScript, 2022. URL: <https://noflojs.org/>.
- [7] TouK, Nussnacker, 2022. URL: <https://nussknacker.io/documentation/about/Overview>.
- [8] P. Schweizer, Flowpipe, 2022. URL: <https://github.com/PaulSchweizer/flowpipe>.
- [9] OpenJS Foundation, Node-RED, 2022. URL: <https://nodered.org/>.
- [10] M. J.P, Flow-based programming: A New Approach to Application Development, Van Nostrand Reinhold, New York, NY, USA, 1994.
- [11] M. Hirzel, R. Soulé, S. Schneider, B. Gedik, R. Grimm, A Catalog of Stream Processing Optimizations, ACM Computing Surveys 46 (2014) 46.
- [12] P. Haslum, N. Lipovetzky, D. Magazzeni, C. Muise, An Introduction to the Planning Domain Definition Language, volume 13, 3 ed., Morgan & Claypool Publishers, 2019.
- [13] H. Geffner, B. Bonet, A Concise Introduction to Models and Methods for Automated Planning, volume 7, 3 ed., A Concise Introduction to Models and Methods for Automated Planning, 2013.

- [14] K. Chowdhary, *Fundamentals of artificial intelligence*, Springer, 2020.
- [15] M. Ghallab, D. S. Nau, P. Traverso, *Automated Planning and Acting*, 1st ed., Elsevier, 2016. 368p.
- [16] M. Fox, D. Long, PDDL2. 1: An extension to PDDL for expressing temporal planning domains., *Journal of Artificial Intelligence Research* (2003).
- [17] R. Fikes, N. Nilsson, STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving, *Artificial Intelligence* 2 (1971) 189–208.
- [18] F. Hueske, V. Kalavri, *Stream processing with Apache Flink: fundamentals, implementation, and operation of streaming applications*, O'Reilly Media, 2019.
- [19] J. Turnbull, *The Docker Book: Containerization Is the New Virtualization*, James Turnbull, 2014.
- [20] D. Bernstein, Containers and cloud: From LXC to Docker to Kubernetes, *IEEE Cloud Computing* 1 (2014) 81–84.
- [21] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, K. Tzoumas, *Apache Flink™ - Stream and Batch Processing in a Single Engine.*, *IEEE Data Engineering Bulletin* (2015).
- [22] E. Scala, P. Haslum, S. Thiébaux, *Heuristics for Numeric Planning via Subgoalng.*, in: *International Joint Conference on Artificial Intelligence*, 2016.
- [23] E. Scala, *Expressive numeric heuristic search planner (enhsp-20)*, 2020. URL: <https://gitlab.com/enricos83/ENHSP-Public/-/tree/enhsp-20>.
- [24] E. Scala, P. Haslum, D. Magazzeni, S. Thiébaux, *Landmarks for Numeric Planning Problems.*, in: *International Joint Conference on Artificial Intelligence*, 2017, pp. 4384–4390.
- [25] P. Haslum, F. Ivankovic, M. Ramírez, D. Gordon, S. Thiébaux, V. Shivashankar, D. S. Nau, *Extending Classical Planning with State Constraints: Heuristics and Search for Optimal Planning*, *Journal of Artificial Intelligence Research* 62 (2018) 373–431.
- [26] N. Arshad, D. Heimbigner, A. L. Wolf, *Deployment and dynamic reconfiguration planning for distributed software systems*, in: *15th IEEE International Conference on Tools with Artificial Intelligence*, IEEE Computer Society, 2003, pp. 39–46.
- [27] A. Riabov, Z. Liu, *Planning for stream processing systems*, in: M. M. Veloso, S. Kambhampati (Eds.), *Proceedings, The Twentieth AAAI Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference*, AAAI Press / The MIT Press, 2005, pp. 1205–1210.
- [28] S. Sohrabi, O. Udrea, A. Ranganathan, A. Riabov, *HTN planning for the composition of stream processing applications*, in: D. Borrajo, S. Kambhampati, A. Oddi, S. Fratini (Eds.), *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling*, AAAI, 2013.
- [29] I. Georgievski, F. Nizamic, A. Lazovik, M. Aiello, *Cloud ready applications composed via HTN planning*, in: *10th IEEE Conference on Service-Oriented Computing and Applications*, IEEE Computer Society, 2017, pp. 81–89.
- [30] J. G. Quenum, J. Josua, *Multi-cloud serverless function composition*, in: I. Brandic, R. Sakellariou, J. Spillner (Eds.), *IEEE/ACM 14th International Conference on Utility and Cloud Computing*, ACM, 2021, pp. 9:1–9:10.