The design and implementation of the Pilot application for the SPD Online Filter

Проектирование и реализация пилотного приложения для SPD Online Filter

L. Romanychev[a,b,1], *D. Oleynik*[a,2], *N. Greben*[a,3]

Романычев Л. Р.[a,b,1], *Олейник Д. А.*[a,2], *Гребень Н. В.*[a,3]

[a] Joint Institute for Nuclear Research, Dubna, 141980 Russia

[a] Объединенный институт ядерных исследований, Дубна, 141980 Россия

[b] St. Petersburg State University, St. Petersburg, 198504 Russia

[b] Санкт-Петербургский государственный университет, Санкт-Петербург, 198504, Россия

В данной работе представлено адаптивное пилотное программное обеспечение, разработанное для SPD Online Filter. Описывается архитектура, компоненты и рабочий процесс, который использует механизм позднего связывания для эффективного использования ресурсов. Ключевой особенностью является адаптивная стратегия управления рабочей нагрузкой, которая максимизирует утилизацию ресурсов, позволяя пилотам на узлах с GPU динамически выполнять задачи, предназначенные для CPU. Представлены результаты тестирования с использованием эмулятора системы сбора данных.

This paper presents the design and implementation of the Pilot application for the SPD Online Filter. We describe its architecture, components, and operational workflow, which leverages a late-binding mechanism for efficient resource utilization. A key feature is an adaptive workload management strategy that maximizes resource utilization by enabling GPU-node pilots to dynamically execute CPU-bound jobs. The results of testing using a DAQ emulator are presented.

PACS: 89.20.Ff; 07.05.-t

Introduction

The Spin Physics Detector (SPD) at the NICA collider complex [1] will produce a vast amount of data, requiring an advanced high-throughput computing (HTC) infrastructure for its processing. The SPD Online Filter is being developed to handle this challenge through multi-step data processing workflows. This paper describes the design and implementation of a key component of this system: the Pilot agent application.

Pilot-based systems are a proven solution for HTC infrastructure, providing a flexible layer for executing jobs on heterogeneous resources [2–4].

[1] E-mail: romanychev@jinr.ru

[2] E-mail: danila@jinr.ru

[3] E-mail: ngreben@jinr.ru

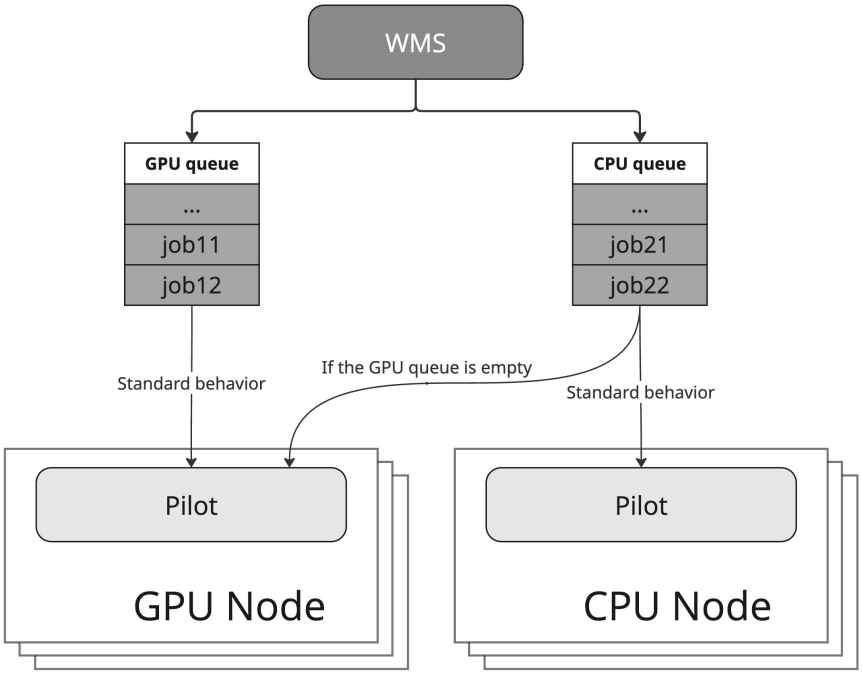


Fig. 1. Adaptive queue management: GPU pilot fetching from CPU queue

In this paradigm, a lightweight "pilot" acts as an intermediary: once active on a compute node, it contacts a Workload Management System (WMS) to fetch a payload. This approach is widely used in major experiments, but many projects develop their own systems, leading to a variety of implementations. Our work builds upon this established concept by introducing an adaptive layer for resource management tailored to the needs of the SPD experiment [5].

Architecture of the SPD pilot software

The overall architecture of the Pilot and its integration with the Workload Management System is described in detail in [6]. On each compute node, a persistent background daemon starts and monitors Pilot agents.

A key concept is *Late Binding*, which assigns jobs to pilots only when they are active and ready, reducing queue times. The WMS manages distinct queues for different job types. GPU pilots can fetch jobs from CPU queues when their primary queue is empty, maximizing resource utilization (Fig. 1). This adaptive behavior ensures efficient use of computational resources while maintaining system responsiveness and fault tolerance.

Implementation and workflow

The Pilot implementation uses a multi-level process and thread architecture for reliable job execution. The execution flow is illustrated in Fig. 2.

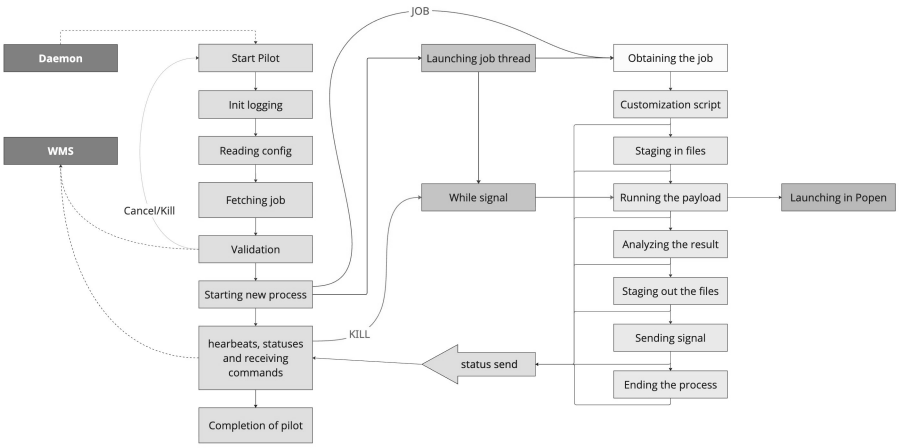


Fig. 2. Pilot execution pipeline and inter-thread communication

1. Process and Thread Architecture

The Pilot employs a three-level hierarchy: the Main Process handles communication with the WMS, a Payload Process executes the job in isolation, and a Job Thread within the Payload Process manages the execution stages. This separation ensures that failures in job execution do not affect the main communication loop.

The Main Process continuously sends heartbeats to the WMS and receives commands. When a job is obtained from RabbitMQ, the Main Process creates a Payload Process using Python's multiprocessing module. The Payload Process spawns a Job Thread that executes the job stages sequentially: Pre-processing, Stage-in, Payload execution, Post-processing, and Stage-out.

2. Thread Synchronization and Cancellation

Communication between processes and threads is implemented using queues and events. The Main Process and Payload Process share two queues: one for passing job data and another for status updates. A `multiprocessing.Event` object provides a synchronization mechanism for cancellation.

The event-based cancellation works as follows: when the WMS sends a "stop" command, the Main Process sets the event by calling `event.set()`. The Payload Process monitors this event in two places: (1) in the monitoring loop that checks log file size, and (2) within the payload execution method, where it periodically checks `event.is_set()` while the subprocess is running.

When `event.is_set()` returns True, the Payload Process terminates the job subprocess gracefully and sets the job status to Failed. Similarly, the Job Thread checks the event after each stage completion. If the event is set, the thread stops processing remaining stages and exits. This ensures that all child threads and processes terminate promptly when cancellation is requested, preventing resource leaks and zombie processes.

3. Job State Model

The job lifecycle is tracked through states: *Registered*, *Extracted* (managed by WMS), *Obtained*, *Pre-processing*, *Stage-in*, *Running*, *Post-processing*, *Stage-out*, *Completed* (managed by Pilot). Any error moves the job to *Failed*

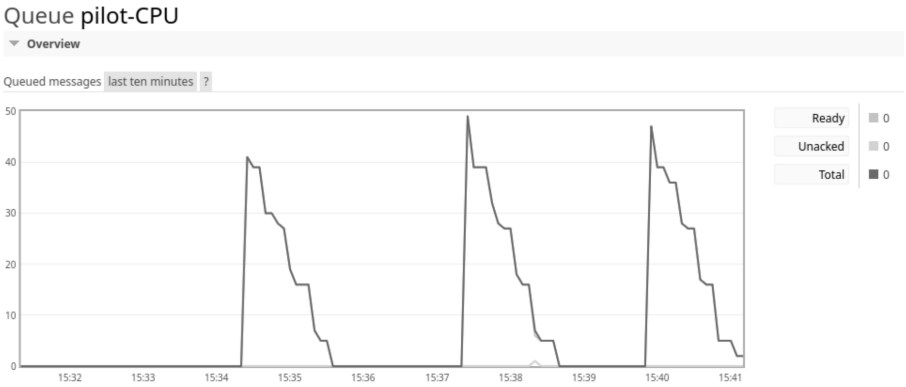


Fig. 3. Queue monitoring: three successful processing cycles

state. Each stage method is wrapped with error handling that automatically uploads logs on failure.

4. Testing and Validation

Testing was conducted on a JINR Cloud virtual machine [7]. Due to limited hardware, 10 daemons were launched on a single node, each spawning pilots. At the current stage, the processor type is configured through a configuration file. A `RAND` processor type option was added for flexible testing without dedicated GPU hardware, allowing the pilot to randomly assign itself either `CPU` or `GPU` type. Fig. 3 shows the queue monitoring dashboard demonstrating three successful processing cycles, where messages are queued, processed, and cleared.

Conclusion

We have designed, implemented, and successfully tested the Pilot application for the SPD Online Filter. Our approach provides a unified, adaptable, and low-overhead interface for job execution. The key features are:

- An adaptive workload management system that maximizes resource utilization by allowing GPU pilots to execute CPU jobs.
- Reduced overhead and high efficiency thanks to the lightweight pilot and late-binding mechanism.
- Robust inter-process and inter-thread communication with event-based cancellation mechanism ensuring proper termination of all child processes and threads when jobs are cancelled or errors occur.

Future work will focus on comprehensive joint testing of SPD Online Filter subsystems, where the pilot software serves as one component of the integrated system.

REFERENCES

1. *Abazov V., others., [SPD Collaboration].* Technical Design Report of the Spin Physics Detector at NICA // Nat. Sci. Rev. — 2024. — V. 1, no. 1. — doi:10.54546/NaturalSciRev.100101.

2. *Sadiku M., Eze K., Musa S.* High-throughput computing // International Journal of Trend in Research and Development. — 2018. — V. 5, no. 4.

3. *Turilli M., Santcroos M., Jha S.* A Comprehensive Perspective on Pilot-Job Systems. — 2015. — URL: `https://arxiv.org/pdf/1508.04180`.

4. *Bal H., others.* Using Pilot Systems to Execute Many Task Workloads on Supercomputers. — 2018. — URL: `https://jsspp.org/papers18/JSSPP_2018_paper_11.pdf`.

5. *Greben N., Romanychev L., Oleynik D., Degtyarev A.* SPD on-line filter: Workload management system and pilot agent // Phys. Part. Nucl. — 2024. — V. 55. — P. 612–614.

6. *Greben N., Oleynik D., Romanychev L.* Workload Management System for SPD Online Filter // Phys. Part. Nucl. Lett. — 2025. — V. 22, no. 5. — P. 1044–1047.

7. *Balashov N., Kuprikov I., Kutovskiy N., others.* Changes and challenges at the JINR and its Member States cloud infrastructures // Phys. Part. Nucl. — 2024. — V. 55. — P. 366–370.