

HPS

zoya.masih@gdwg.de

Zoya Masih

Trace Analysis of ML workflows using Vampir

Table of contents

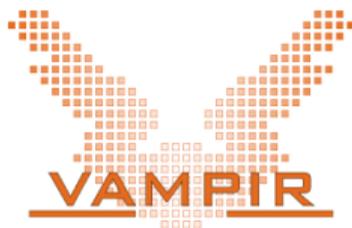
- 1** Introduction
- 2** Vampir/Score-P for Python Codes
- 3** Use case: A simple code
- 4** Use case: LLM-Fine Tuning
- 5** Challenges

Table of Contents

- 1** Introduction
- 2 Vampir/Score-P for Python Codes
- 3 Use case: A simple code
- 4 Use case: LLM-Fine Tuning
- 5 Challenges

What is Vampir

- Vampir is a framework for performance analysis
- Enables developers to study program behavior at a fine level of detail
- The tool is well-proven and widely used in HPC.



Vampir reads OTF2 Files

- Vampir reads Open Trace Format 2 (OTF2) files
- OTF2 is specially designed for massively parallel programs
- Some performance monitoring environments produce OTF2 trace files
 - ▶ like Score-P and TAU

Score-P: The main instrumentor

- compile and run a program, with Score-P prefixed

```
1 scorep gcc -o myApp [-fopenmp] myApp.c  
2 scorep ./myApp  
3 >> output directory: scorep-20240331_1326_33868353060433708
```

Score-P: The main instrumentor

- Enabling the tracing mode generates an OTF2 file
- `export SCOREP_ENABLE_TRACING=true`
- This file is used to capture and store performance data

Table of Contents

- 1 Introduction
- 2 Vampir/Score-P for Python Codes**
- 3 Use case: A simple code
- 4 Use case: LLM-Fine Tuning
- 5 Challenges

The Python-binding

- Score-P only supports Fortran and C/C++ applications by default
- Score-P Python module, supports Python code as well
- The Score-P C-bindings serve as a bridge between Python, C, and Score-P
- Facilitating communication and interaction between these components

Python-binding

Listing: Installing python-binding

```
1 export PATH=/path/to/scorep/bin:$PATH
2
3 git clone https://github.com/score-p/scorep\_binding\_python
4 cd scorep\_python\_bindings/
5 pip3 install .
6 cd ..
7
8 python -m scorep <script.py>
```

Python-binding

Listing: Using Score-P for python-binding

```
1 export SCOREP_ENABLE_TRACING=true
2 export SCOREP_TOTAL_MEMORY=1837MB #estimate with scorep-score
3 python -m scorep --io=posix [--mpp=mpi] --thread=pthread script.py
4 srun -n 2 python -m scorep script.py
```

Table of Contents

- 1 Introduction
- 2 Vampir/Score-P for Python Codes
- 3 Use case: A simple code**
- 4 Use case: LLM-Fine Tuning
- 5 Challenges

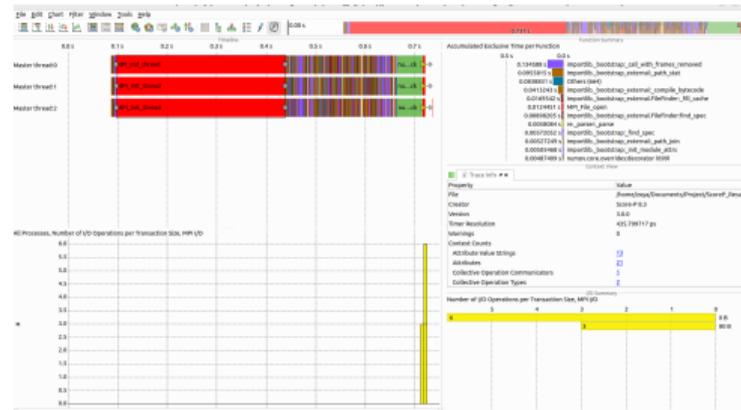
A simple MPI4PY code

■ `srun -n 3 python -m scorep --io=posix --mpp=mpi script.py`

```

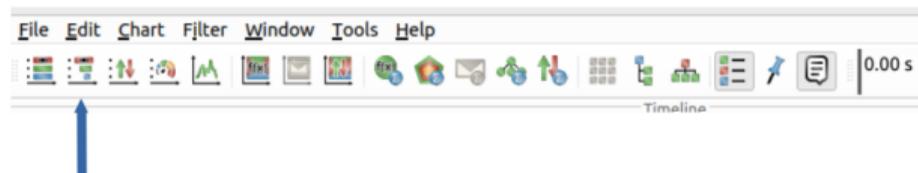
1 from mpi4py import MPI
2 import numpy as np
3 amode =
    ↪ MPI.MODE_WRONLY|MPI.MODE_CREATE
4 comm = MPI.COMM_WORLD
5 fh = MPI.File.Open(comm,
    ↪ "./datafile.contig", amode)
6 buffer = np.empty(10, dtype=np.int)
7 buffer[:] = comm.Get_rank()
8 offset =
    ↪ comm.Get_rank()*buffer.nbytes
9 fh.Write_at_all(offset, buffer)
10 fh.Close()

```



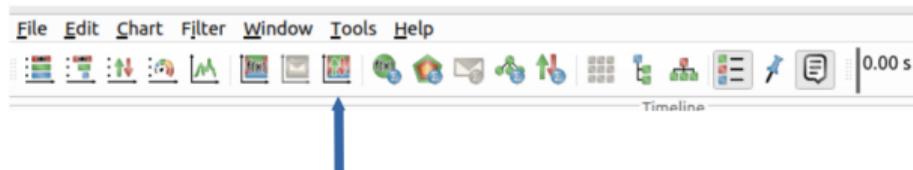
Process Timeline

- Information about different Levels of function calls
- levels show the execution flow and hierarchical relationships

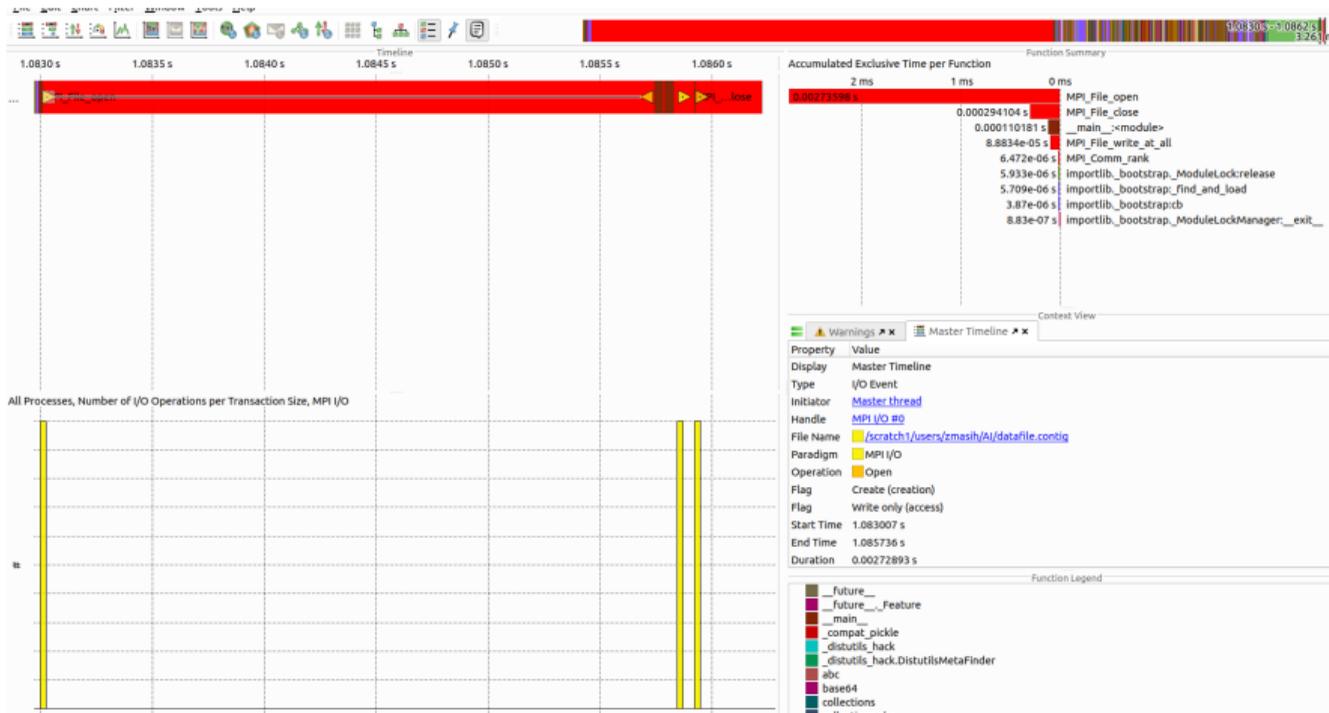


I/O summary

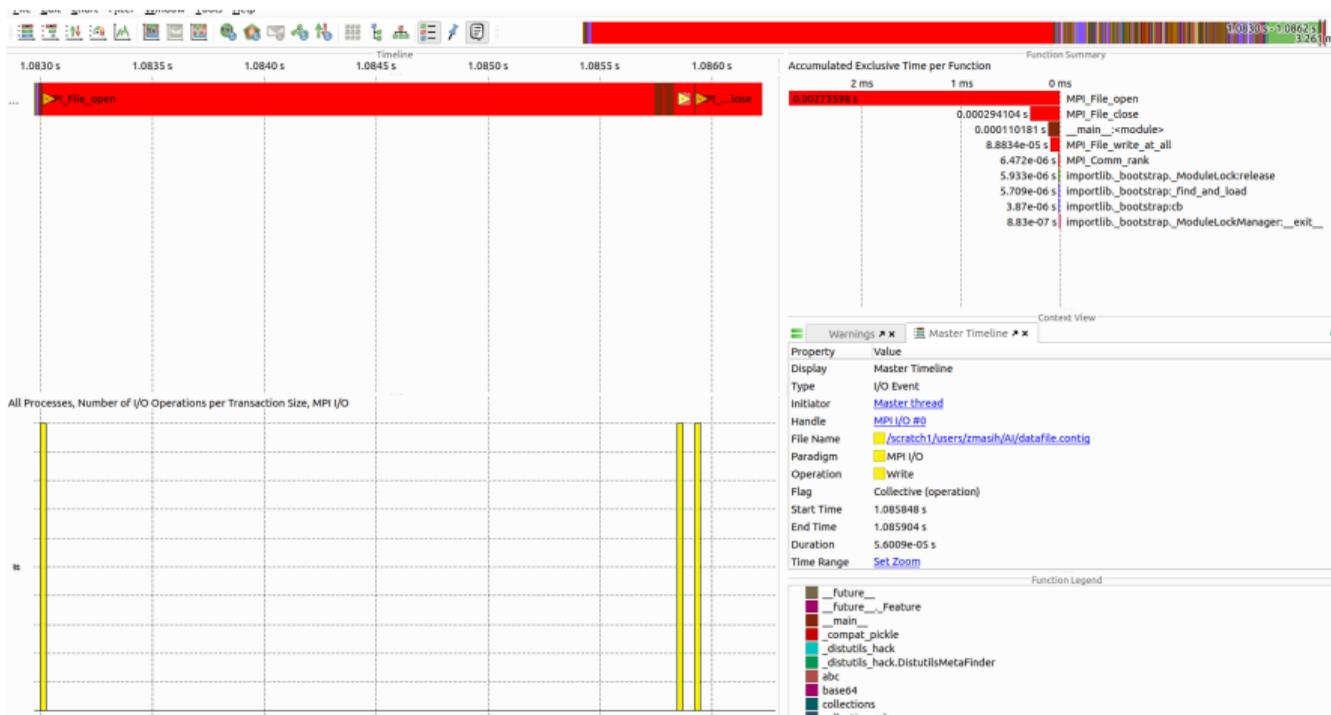
- Vampir highlights I/O operations if I/O performance data has been recorded.
 - ▶ `-io=posix (=runtime:posix)`
- MPI I/O routines are automatically instrumented
 - ▶ if the scorep instrumenter detects the MPI programming paradigm



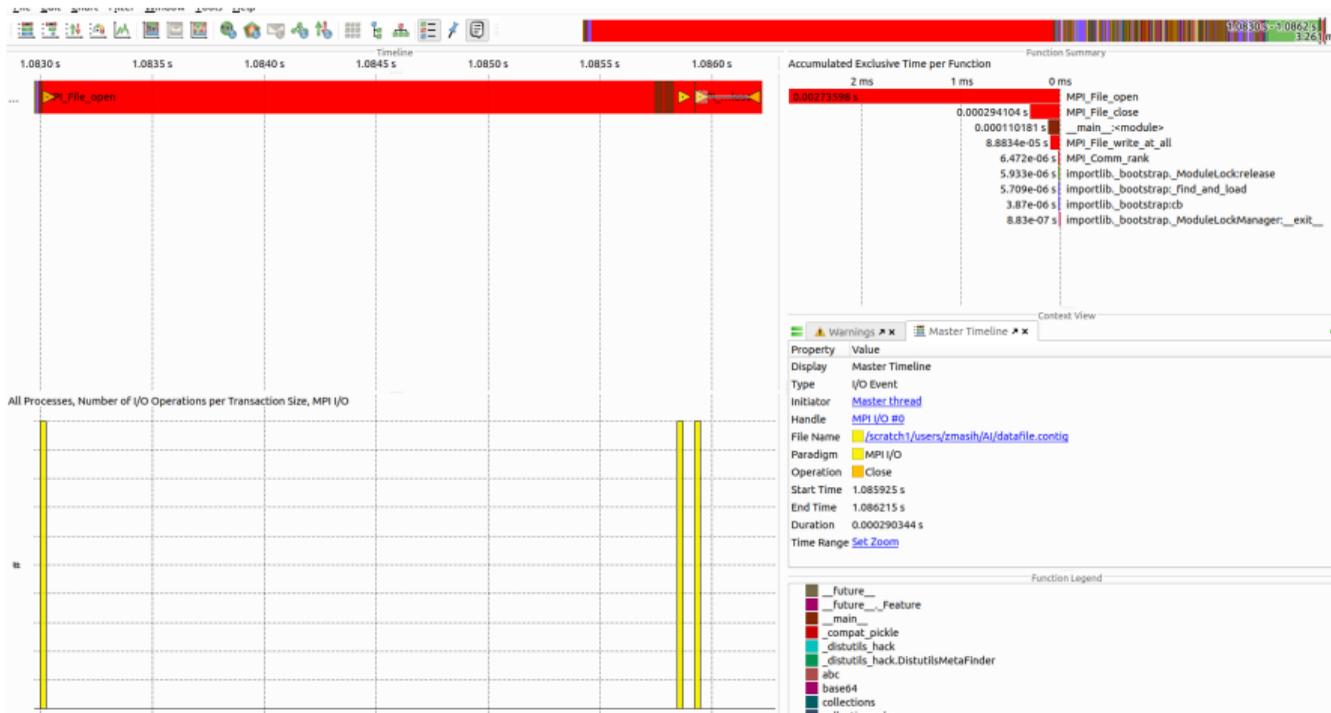
I/O summary



I/O summary



I/O summary



I/O summary

The screenshot shows the 'Context view' window in Vampir, displaying an 'I/O summary' table. The table lists various properties and their corresponding values. The 'Source Code Files' row is highlighted in blue. The table is organized into several sections: Context Counts, Event Counts, and File Substrate.

Property	Value
File	/home/zoya/Documents/Project/ScoreP_Results/Python/Trading_io_mpi-io.py/with srun-3processes/scorep-20240426_1317_43386083585181808/traces.otf2
Creator	Score-P 8.3
Version	3.0.0
Timer Resolution	435.799717 ps
Warnings	0
Context Counts	
Attribute Value Strings	13
Attributes	21
Collective Operation Communicators	1
Collective Operation Types	2
Communicators	13
Counter Groups	0
Counters	0
Distinct File Paths	1
Files	1
Function Groups	209
Functions	1056
I/O Handles	11
I/O Paradigms	3
Message Tags	0
Message Types	11
Process Group Paradigms	4
Process Groups	7
Processes	3
Source Code Files	128
Source Code Locations	
Event Counts	
Total	179001
States	178980
Counter	0
Messages	0
Collectives	12
File I/O	9
File Substrate	
Trace Compression	NONE
Trace ID	8f490d36b9ff2e4c

I/O summary

Context View

Trace Info I/O Paradigms I/O Handles Collective Operation Types Attribute Value Strings

Property	Value
Value	--io=posix
Value	--mpp=mpi
Value	-m
Value	/opt/sw/rev/23.12/linux-scientific7-haswell/gcc-11.4.0/anaconda3-2023.09-0-npsf7i/bin/python3.11
Value	/scratch1
Value	Linux
Value	beegfs
Value	beegfs_nodev
Value	machine
Value	mpi-io.py
Value	r31
Value	scorep
Value	shared memory

I/O summary

Context View

 Trace Info  I/O Paradigms  I/O Handles  Collective Operation Types  Attribute Value Strings 

Property	Value
Value	CREATE_HANDLE
Value	DESTROY_HANDLE

I/O summary

Context View

Trace Info ✕ ✕ I/O Paradigms ✕ ✕ I/O Handles ✕ ✕ Collective Operation Types ✕ ✕ Attribute Value Strings ✕ ✕

Property	Value
Handle	MPI I/O #0
Handle	POSIX I/O #0
Handle	POSIX I/O #1
Handle	STDERR_FILENO
Handle	STDIN_FILENO
Handle	STDOUT_FILENO
Handle	all open output streams
Handle	stderr
Handle	stdin
Handle	stdout
Handle	sync - commit buffer cache to disk

I/O summary

Context View

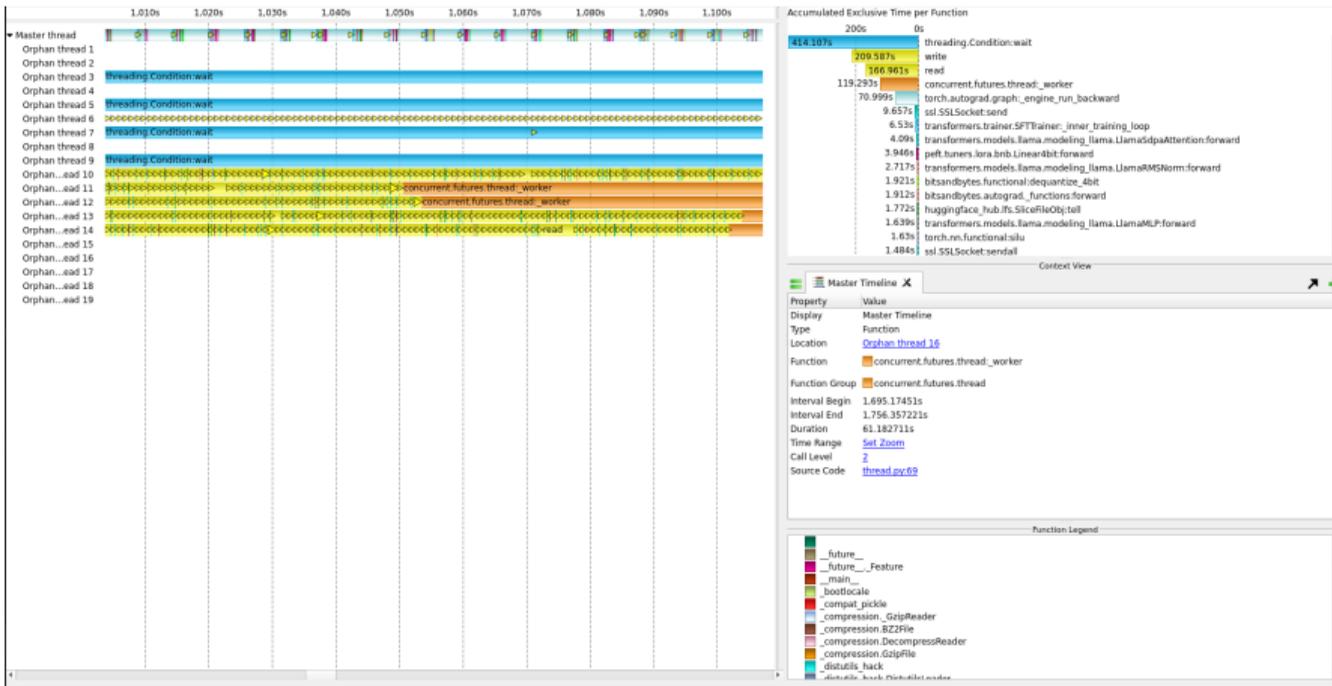
Trace Info ⌵ × I/O Paradigms ⌵ × I/O Handles ⌵ × Collective Operation Types ⌵ × Attribute Value Strings ⌵ ×

Property	Value
Paradigm	ISO C I/O
Paradigm	MPI I/O
Paradigm	POSIX I/O

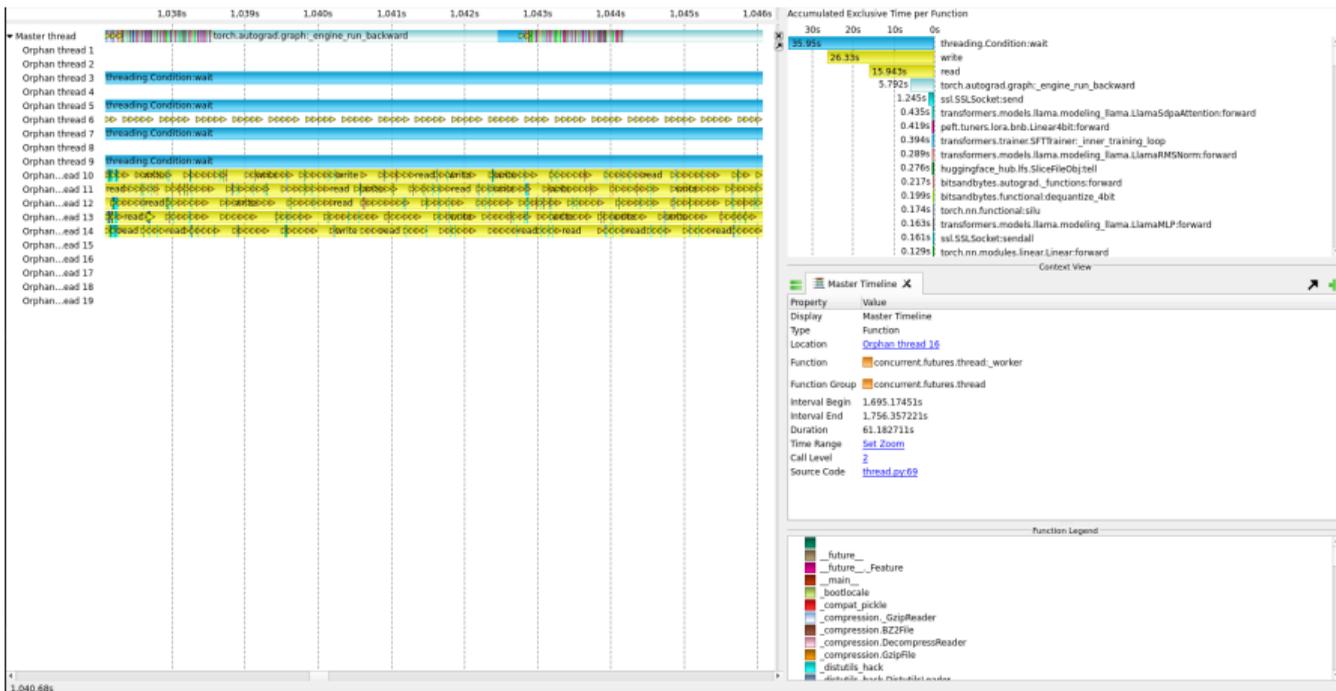
Table of Contents

- 1 Introduction
- 2 Vampir/Score-P for Python Codes
- 3 Use case: A simple code
- 4 Use case: LLM-Fine Tuning**
- 5 Challenges

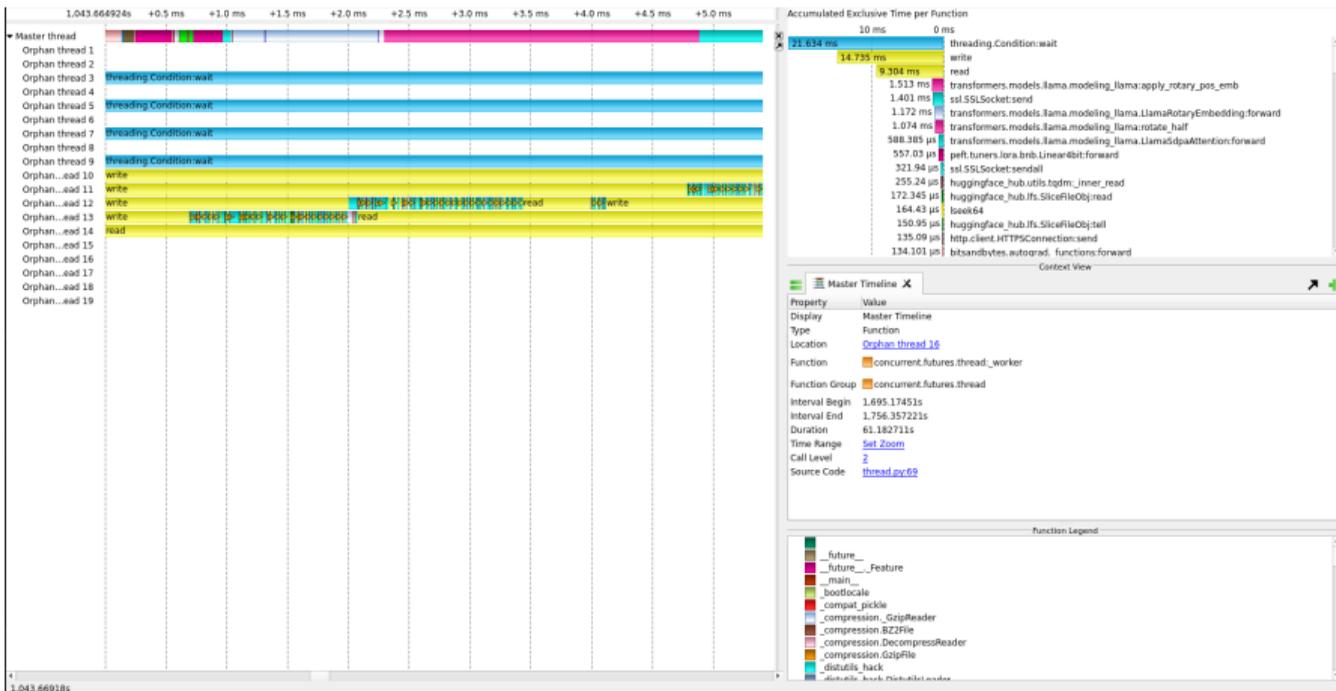
LLM-Fine Tuning



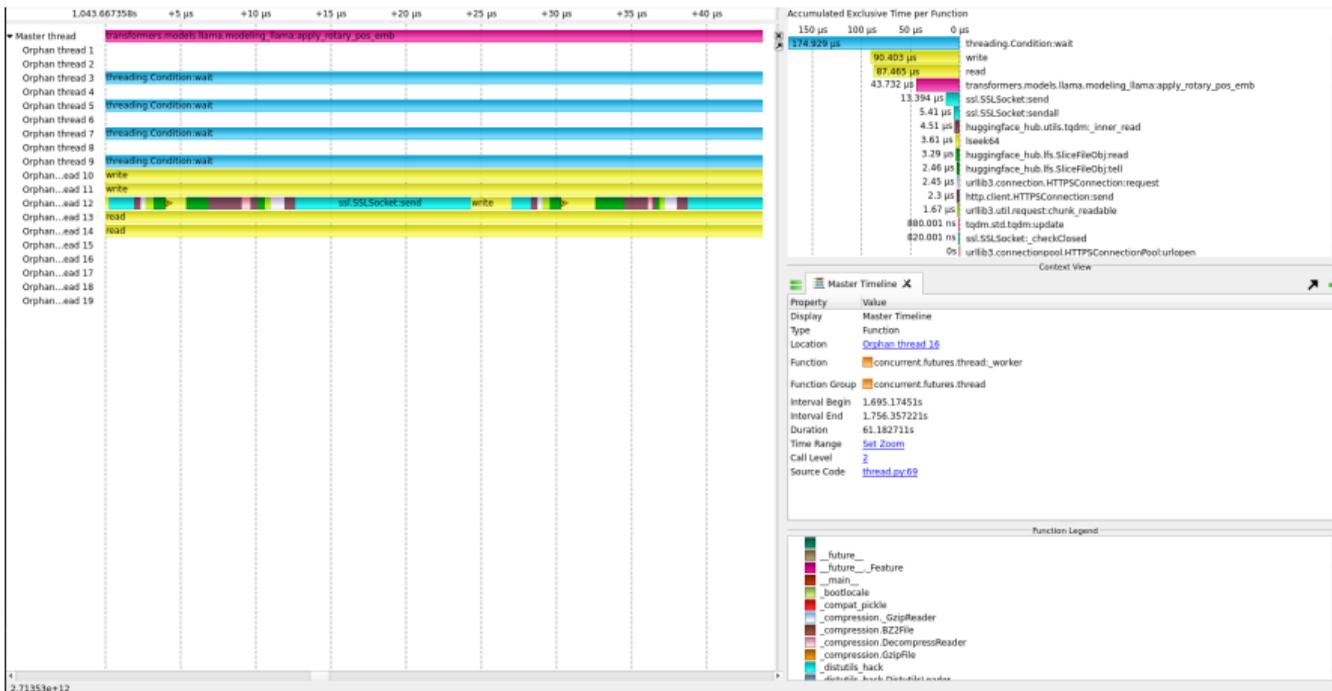
LLM-Fine Tuning



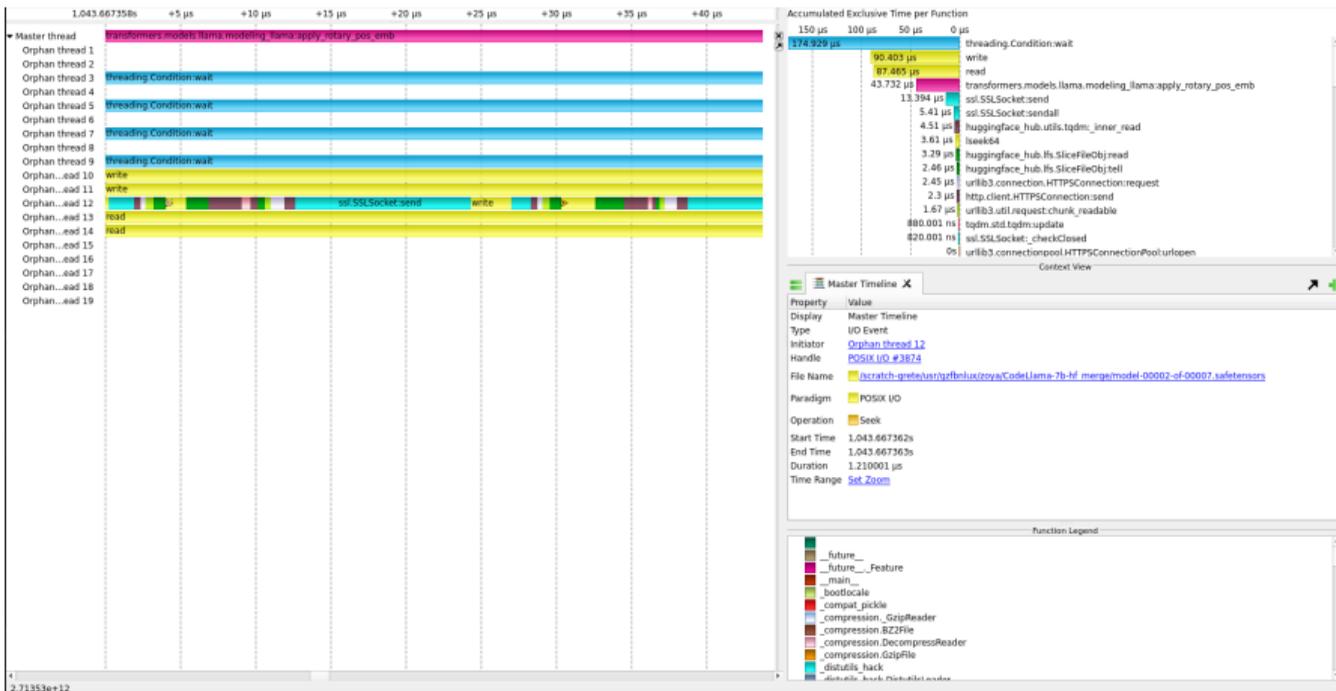
LLM-Fine Tuning



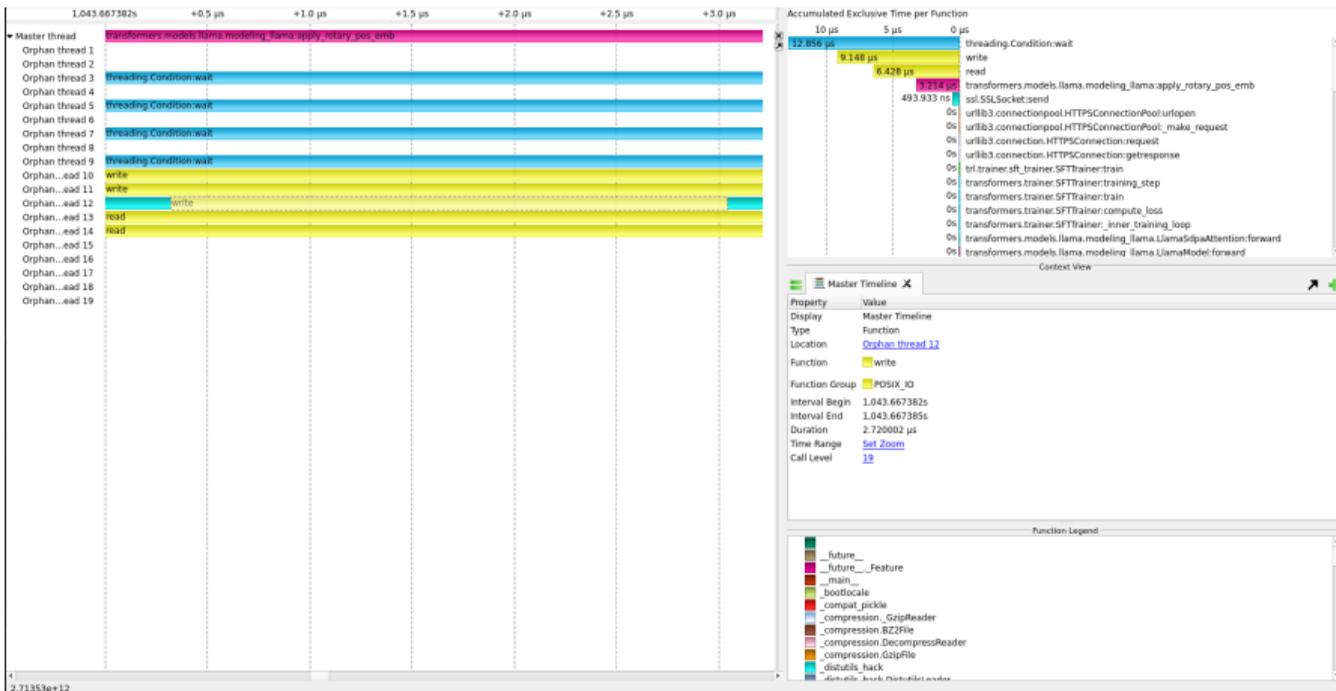
LLM-Fine Tuning



LLM-Fine Tuning



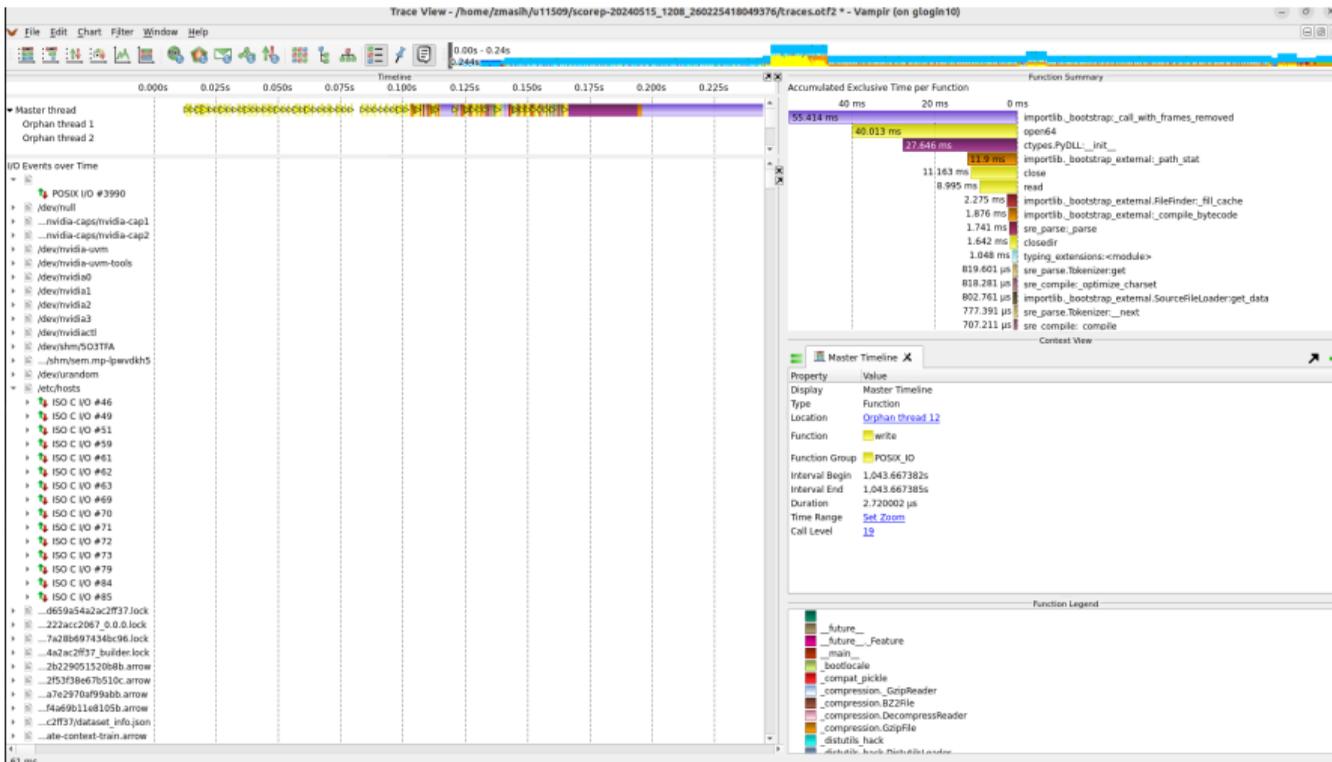
LLM-Fine Tuning



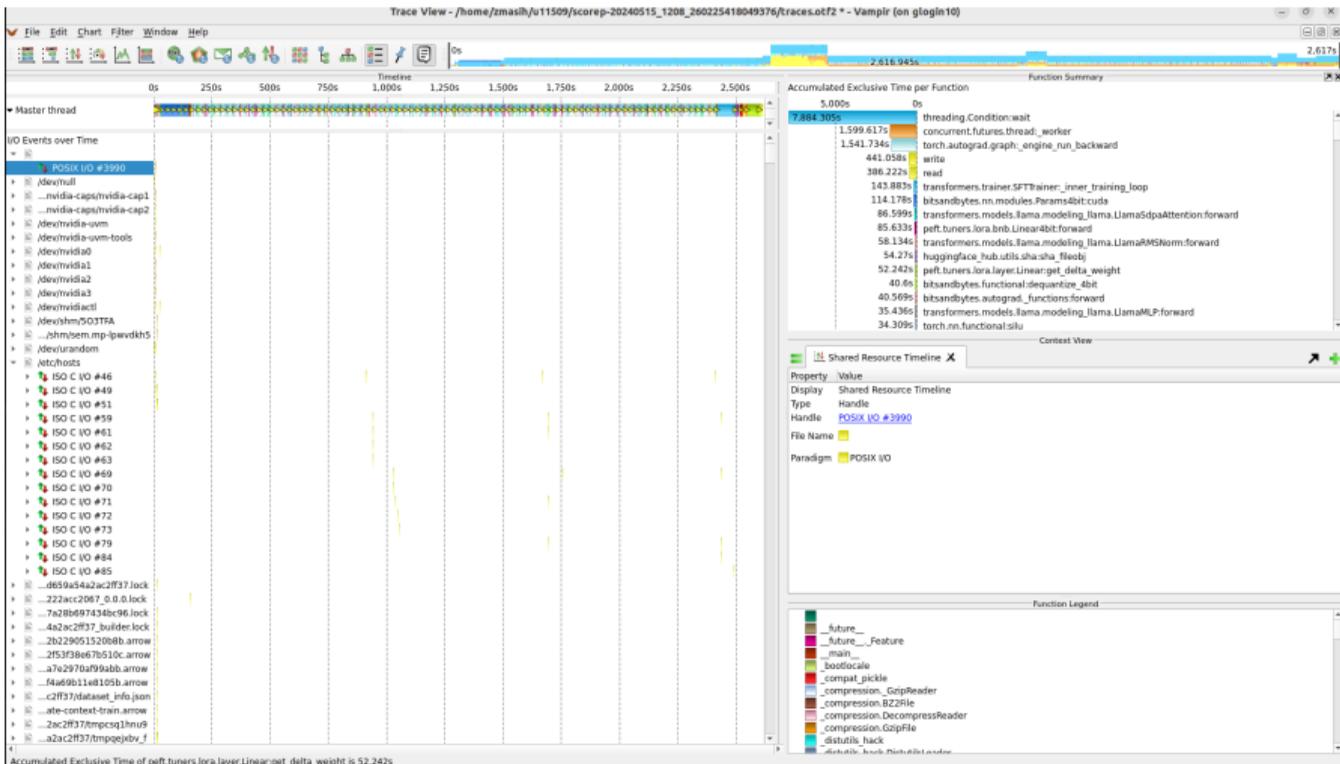
LLM-Fine Tuning



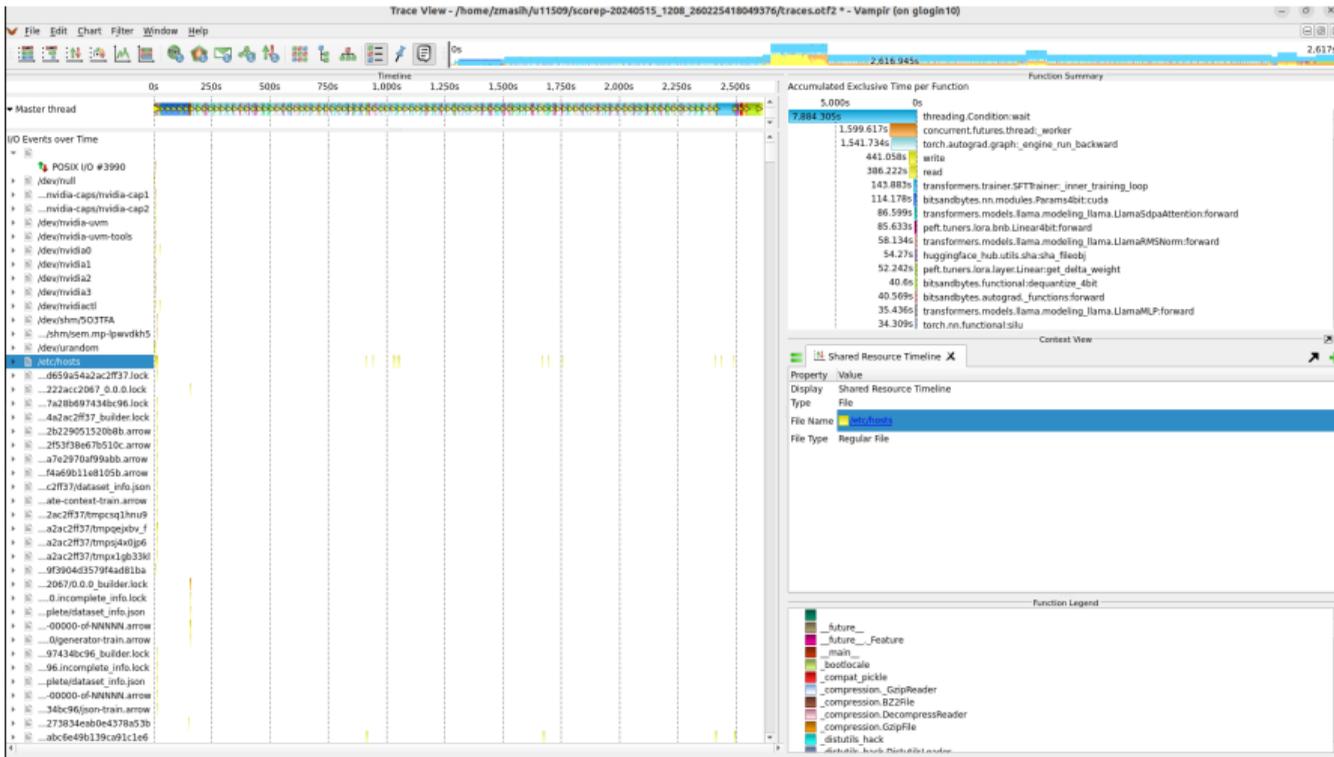
LLM-Fine Tuning



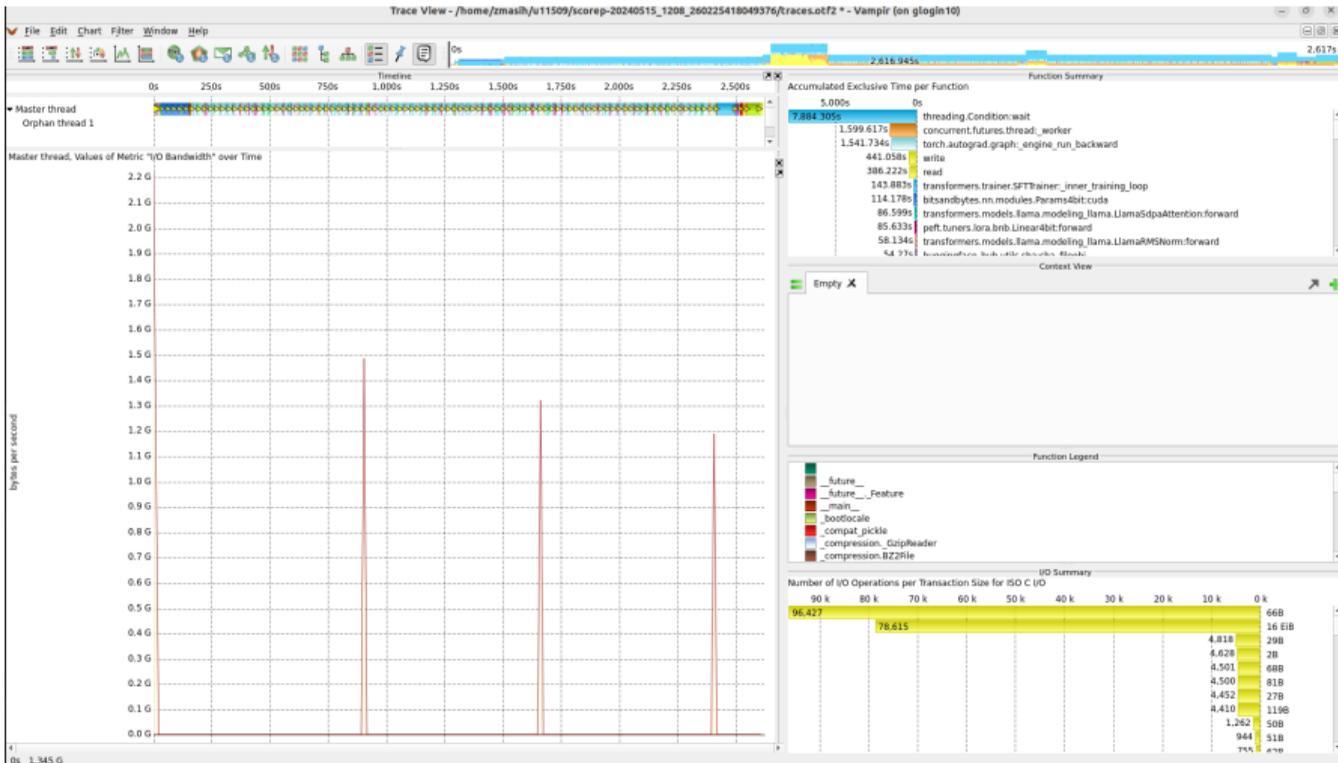
LLM-Fine Tuning



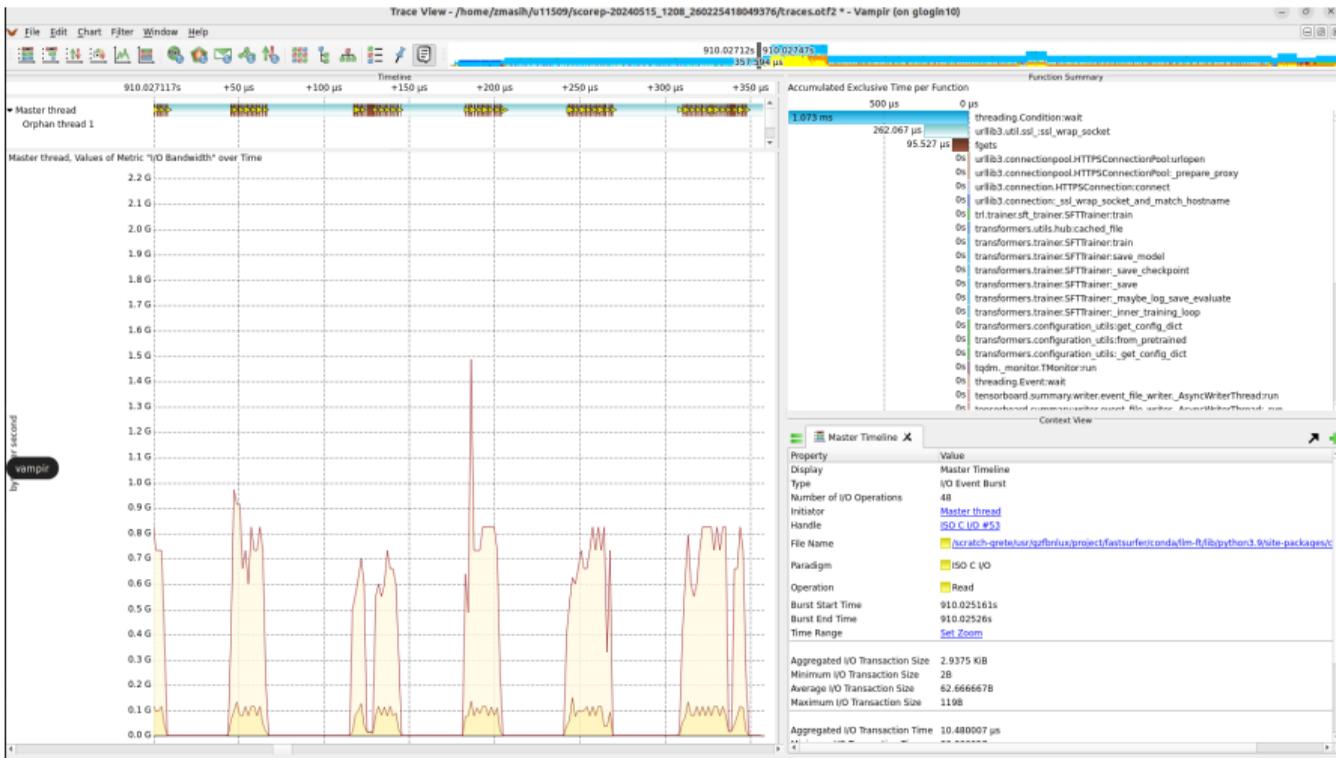
LLM-Fine Tuning



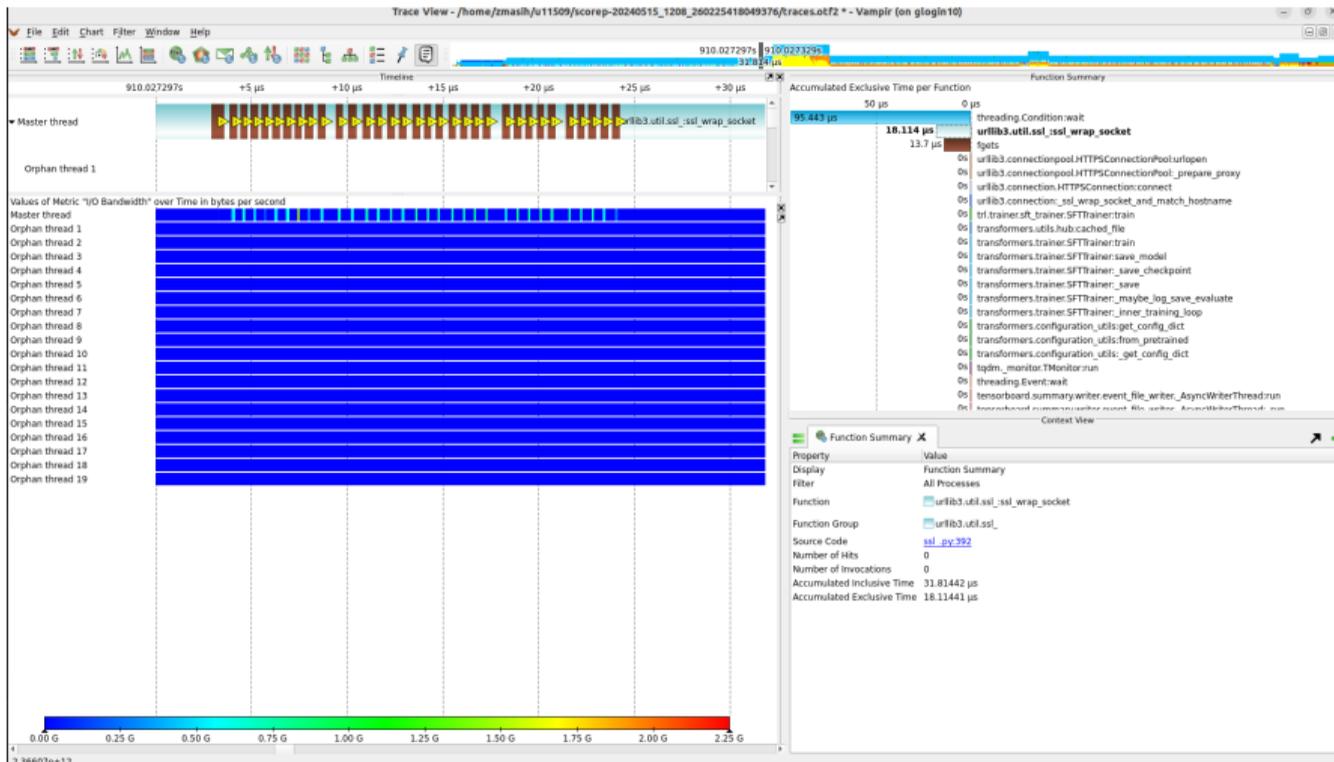
LLM-Fine Tuning



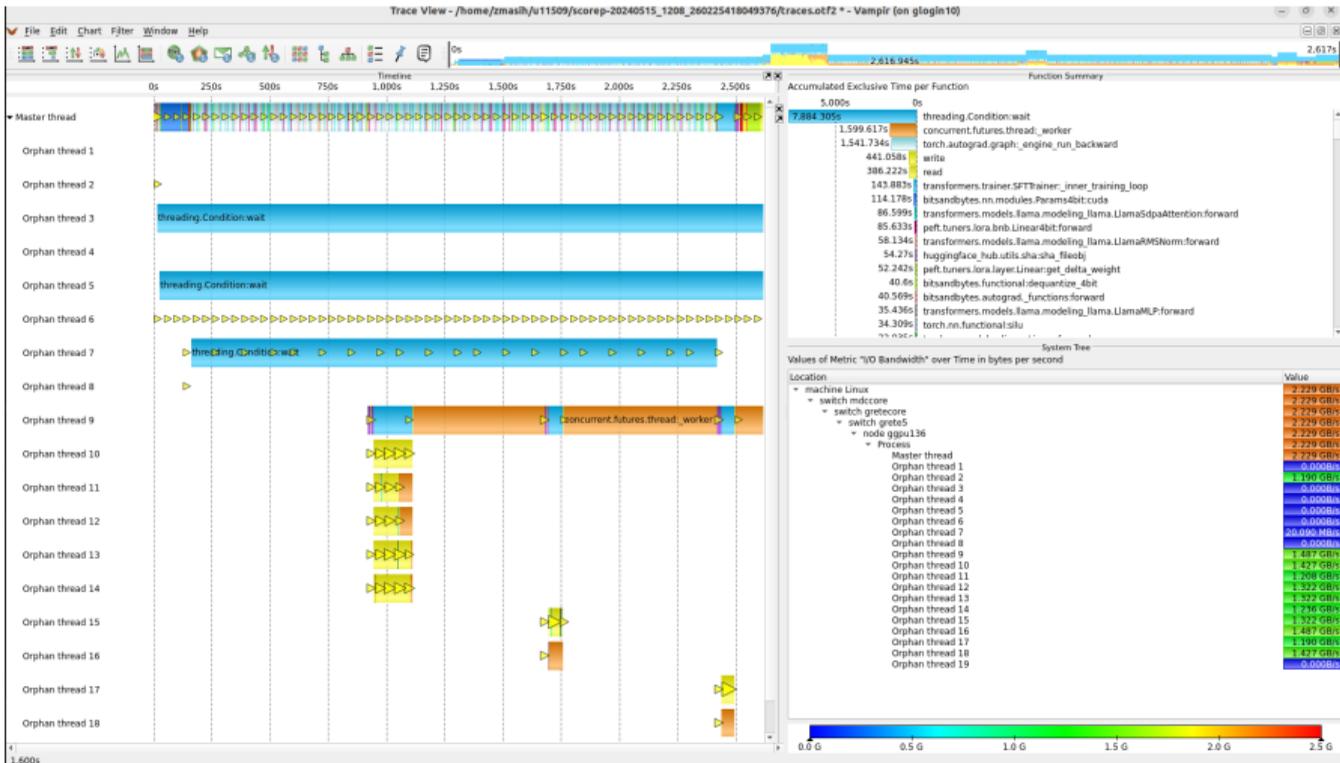
LLM-Fine Tuning



LLM-Fine Tuning



LLM-Fine Tuning



LLM-Fine Tuning

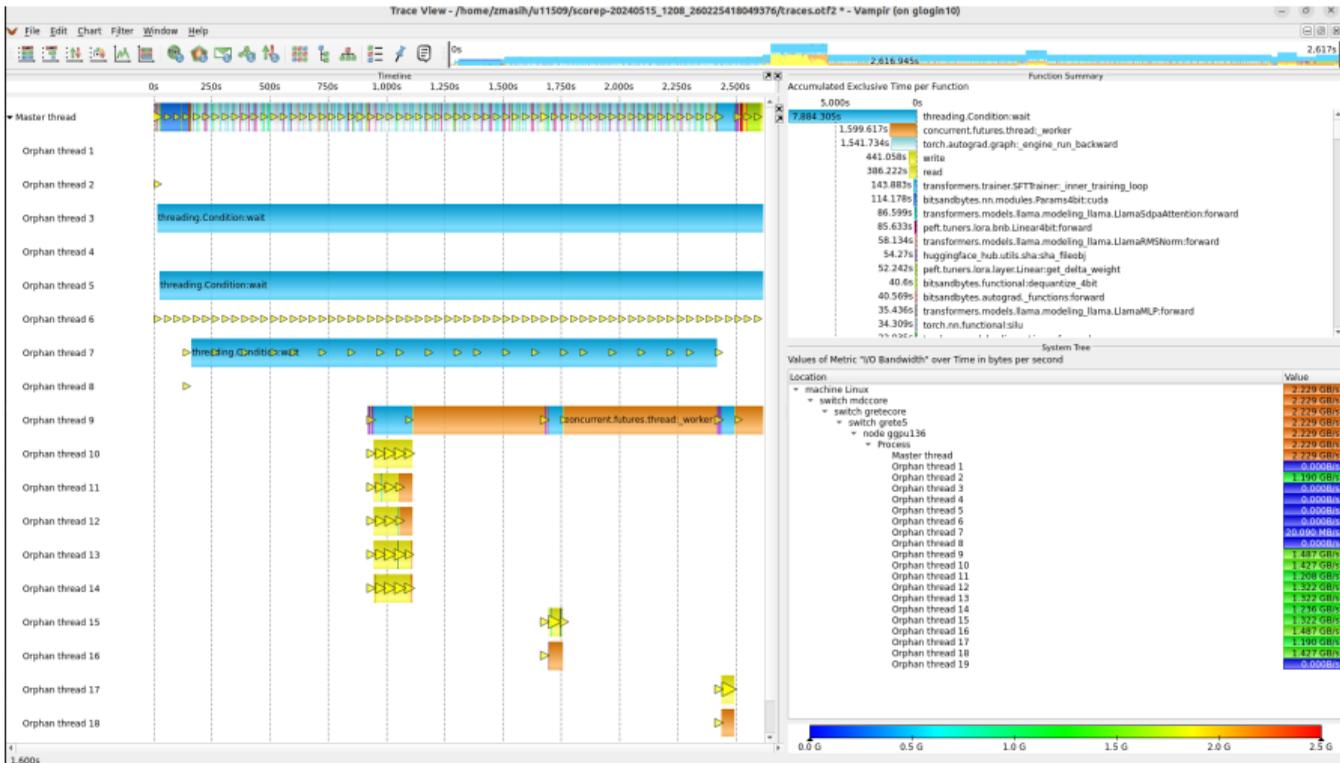


Table of Contents

- 1 Introduction
- 2 Vampir/Score-P for Python Codes
- 3 Use case: A simple code
- 4 Use case: LLM-Fine Tuning
- 5 Challenges**

Challenges

- Generally, X11 forwarding can be slow and inefficient
 - ▶ The basic underlying design not well suited to the rising demands of GUI's?
 - ▶ RFB (remote frame buffer) protocol used in VNC May help
- ScoreP python-binding currently supports only MPI and SHMEM
- Instrumenting applications is not always straightforward

End

Wish it was helpful

Happy Analysis