

# Data Science and Advanced Programming — Lecture 13

## High-Performance Computing with Python

Simon Scheidegger  
Department of Economics, University of Lausanne, Switzerland

December 8th, 2025 | 12:30 - 16:00 | Internef 263

# Lecture Overview

## Schedule:

- ① Why Parallelism? (45 min)
- ② Threading & I/O-Bound (50 min)  
— *Break (15 min)* —
- ③ Multiprocessing & CPU-Bound (55 min)
- ④ Finance Applications (45 min)  
— *Break (10 min)* —
- ⑤ Projects & Best Practices (40 min)

## Format:

- Slides introduce concepts
-  Switch to notebooks for practice
- Back to slides for next concept
- Repeat!

## What You'll Learn:

- Speed up your Python code
- Threading vs. Multiprocessing
- Real finance applications

# Topic 1

## Why Parallelism Matters

Motivation & Core Concepts

01\_motivation.ipynb

# The Problem: Your Code is Too Slow

**Scenario:** You need to price 10,000 exotic options using Monte Carlo simulation

- Each option requires 100,000 simulation paths
- Each simulation takes 0.1 seconds
- Total time:  $10,000 \times 0.1 = 1,000$  seconds  $\approx \mathbf{17 \text{ minutes}}$

## The Reality

Your laptop has 8 CPU cores, but Python is only using **one** of them!

## The Goal

Use all cores  $\rightarrow$  Reduce time to  $\approx \mathbf{2 \text{ minutes}}$

# The End of Free Speed

## Moore's Law (1965-2005):


- Transistor count doubles every 2 years
- Clock speed kept increasing
- Your code got faster *automatically*

## The Wall (~2005):


- Power consumption  $\propto$  frequency<sup>3</sup>
- Clock speeds plateaued at ~4 GHz

## The Solution:

- More cores, not faster cores
- Your laptop: 4-16 cores



# Sequential vs. Parallel Execution



## Key Insight

If tasks are **independent**, we can run them simultaneously on different cores.



## Hands-On Time!

Experience the Speedup

`01_motivation.ipynb`

*Sections 1-2: Sequential vs. Parallel Option Pricing*

# Amdahl's Law: The Limits of Parallelism

## Not everything can be parallelized:


- Reading input data
- Setting up calculations
- Combining final results

## Amdahl's Law:

$$\text{Speedup} = \frac{1}{(1 - p) + \frac{p}{n}}$$

Where:

- $p$  = parallelizable fraction
- $n$  = number of cores



## Takeaway

Even with infinite cores, 50% parallel code gives max 2x speedup!



## Hands-On Time!

Explore Amdahl's Law

`01_motivation.ipynb`

*Sections 3-4: Scaling experiments and visualization*

# Why Finance Needs Parallelism

## CPU-Intensive Tasks:

- Monte Carlo simulations
- Option pricing (exotic derivatives)
- Portfolio optimization
- Risk calculations (VaR, CVaR)
- Backtesting strategies
- Bootstrap inference


## I/O-Intensive Tasks:

- Fetching market data
- Reading multiple data files
- API calls to data providers
- Database queries
- Downloading reports

## The Two Types of Waiting

- **CPU-bound:** Waiting for calculations to finish → **Multiprocessing**
- **I/O-bound:** Waiting for data to arrive → **Threading**

# Python's Parallel Toolkit



Our Focus: `concurrent.futures`

Clean, simple API that works for both threading and multiprocessing.

# First Look: Sequential vs. Parallel

## Sequential (what you're used to):

```
1 results = []
2 for item in data:
3     result = slow_function(item)
4     results.append(result)
5
```

## Parallel (what we'll learn):

```
1 from concurrent.futures import ProcessPoolExecutor
2
3 with ProcessPoolExecutor() as executor:
4     results = list(executor.map(slow_function, data))
5
```

That's It!

Two extra lines of code can give you 4-8x speedup on multi-core machines.



## Hands-On Time!

Complete Topic 1 Exercises

`01_motivation.ipynb`

*Section 5: Exercises (put options, volatility grid)*

# Topic 2

## Threading & I/O-Bound Tasks

When Waiting is the Bottleneck

02\_threading\_io\_bound.ipynb


# What is a Thread?

## Process:

- Independent program, own memory
- Heavy to create, true parallelism

## Thread:

- Lives inside a process, shares memory
- Lightweight, concurrent (not parallel\*)



## \*The Python Catch

Due to the GIL, Python threads don't run truly in parallel for CPU work.

# The Global Interpreter Lock (GIL)

## What is the GIL?

- A mutex in CPython
- Only one thread executes Python bytecode at a time
- Protects memory management



■ Python code ■ I/O wait

## What this means:

- CPU-bound code: threads don't help
- I/O-bound code: threads help a lot!

## Key Insight

Threads take turns using the CPU while others wait for I/O.

## Why I/O works:

- GIL is released during I/O
- While one thread waits for data...
- ...another can do work!

# I/O-Bound vs. CPU-Bound

Characteristic	I/O-Bound	CPU-Bound
Bottleneck	Waiting for data	Calculations
CPU usage	Low (lots of idle)	High (near 100%)
Solution	Threading	Multiprocessing
Finance Examples		
	Fetching stock prices	Monte Carlo simulation
	Reading CSV files	Portfolio optimization
	API calls	VaR calculations
	Database queries	Option pricing

## How to Tell?

Run your code and check CPU usage. If it's low while code runs slowly → I/O-bound.



## Hands-On Time!

### I/O-Bound Demo

02\_threading\_io\_bound.ipynb

*Section 1: See threading in action with simulated data fetching*

# ThreadPoolExecutor: The Simple Way

```
1 from concurrent.futures import ThreadPoolExecutor
2 import time
3
4 def fetch_stock_data(ticker):
5     """Simulate fetching data (I/O operation)"""
6     time.sleep(0.5) # Simulate network delay
7     return {"ticker": ticker, "price": 100.0}
8
9 tickers = ["AAPL", "GOOGL", "MSFT", "AMZN"]
10
11 # Parallel fetching
12 with ThreadPoolExecutor(max_workers=4) as executor:
13     results = list(executor.map(fetch_stock_data, tickers))
14
```

**Sequential:**  $4 \times 0.5\text{s} = 2.0\text{s}$

**Parallel:**  $\approx 0.5\text{s}$  (4x faster!)

# Pattern 1: executor.map()

**Use when:** Same function, many inputs, order matters

```
1 from concurrent.futures import ThreadPoolExecutor
2
3 def process(item):
4     return item * 2
5
6 items = [1, 2, 3, 4, 5]
7
8 with ThreadPoolExecutor(max_workers=4) as executor:
9     results = list(executor.map(process, items)) # [2, 4, 6, 8, 10]
10
```

## Key Properties

Results maintain input order • Simple syntax • Good for homogeneous tasks

## Pattern 2: executor.submit() + as\_completed()

**Use when:** Want results as they finish (not in order)

```
1 from concurrent.futures import ThreadPoolExecutor, as_completed
2
3 tickers = ["AAPL", "GOOGL", "MSFT", "AMZN"]
4
5 with ThreadPoolExecutor() as executor:
6     # Submit tasks
7     futures = {executor.submit(fetch_data, t): t
8                 for t in tickers}
9
10    # Process results as they complete
11    for future in as_completed(futures):
12        ticker = futures[future]
13        result = future.result()
14        print(f"{ticker}: got data!")
15
```

When to use this pattern

Progress feedback, early termination, heterogeneous task times



## Hands-On Time!

Practice Both Patterns

02\_threading\_io\_bound.ipynb

*Section 2-3: executor.map() vs submit() + as\_completed()*

# Handling Exceptions in Threads

```
1 from concurrent.futures import ThreadPoolExecutor, as_completed
2
3 def risky_fetch(ticker):
4     if ticker == "BAD":
5         raise ValueError(f"Invalid ticker: {ticker}")
6     return {"ticker": ticker, "price": 100.0}
7
8 tickers = ["AAPL", "BAD", "MSFT"]
9
10 with ThreadPoolExecutor() as executor:
11     futures = {executor.submit(risky_fetch, t): t for t in tickers}
12
13 for future in as_completed(futures):
14     ticker = futures[future]
15     try:
16         result = future.result()
17         print(f"{ticker}: {result}")
18     except Exception as e:
19         print(f"{ticker}: ERROR - {e}")
20
```

# Threading: Key Takeaways

## When to Use Threading

- Fetching data from multiple sources
- Reading/writing multiple files
- Any task where you're waiting for external resources

## When NOT to Use Threading

- Heavy computations (Monte Carlo, optimization)
- Number crunching → Use multiprocessing instead!

## Best Practices

- Use ThreadPoolExecutor (not raw threads)
- Always handle exceptions
- Use context managers (with statement)



## Hands-On Time!

Complete Threading Exercises

02\_threading\_io\_bound.ipynb

*Section 4-5: File processing and exercises*



## **Break Time**

15 minutes

Next up: Multiprocessing for CPU-bound tasks

# Topic 3

## Multiprocessing & CPU-Bound Tasks

True Parallelism for Heavy Computation

03\_multiprocessing\_cpu\_bound.ipynb


# Why Multiprocessing?

## The GIL Problem:

- Threads share one GIL
- Only one runs Python at a time
- CPU-bound code doesn't speed up

## The Solution:

- Use separate processes
- Each process has its own GIL
- Each process has its own Python interpreter
- True parallel execution!



Result

3 processes = 3 cores working simultaneously

# Threads vs. Processes: Trade-offs

Aspect	Threads	Processes
Memory	Shared	Separate (copied)
Creation	Fast	Slower
Communication	Easy (shared vars)	Harder (serialization)
GIL	Blocked	Bypassed
Best for	I/O-bound	CPU-bound

---

concurrent.futures	ThreadPoolExecutor	ProcessPoolExecutor
--------------------	--------------------	---------------------

## Process Overhead

Creating processes is slower and uses more memory. The task must be substantial enough to overcome this overhead.



## Hands-On Time!

Threading vs. Multiprocessing

03\_multiprocessing\_cpu\_bound.ipynb

*Section 1: See why threads fail for CPU-bound tasks*

# ProcessPoolExecutor: Same API, True Parallelism

```
1 from concurrent.futures import ProcessPoolExecutor
2 import numpy as np
3
4 def monte_carlo_pi(n_samples):
5     x = np.random.random(n_samples)
6     y = np.random.random(n_samples)
7     inside = np.sum(x**2 + y**2 <= 1)
8     return 4 * inside / n_samples
9
10 # Split work across 4 processes
11 with ProcessPoolExecutor(max_workers=4) as executor:
12     estimates = list(executor.map(monte_carlo_pi, [1_000_000] * 4))
13
14 pi_estimate = np.mean(estimates)
15
```

# Finance Example: Monte Carlo Option Pricing

```
1 import numpy as np
2 from concurrent.futures import ProcessPoolExecutor
3
4 def price_european_call(args):
5     S0, K, T, r, sigma, n_paths = args
6     Z = np.random.standard_normal(n_paths)
7     ST = S0 * np.exp((r - 0.5*sigma**2)*T + sigma*np.sqrt(T)*Z)
8     payoffs = np.maximum(ST - K, 0)
9     return np.exp(-r * T) * np.mean(payoffs)
10
11 params = (100, 100, 1.0, 0.05, 0.2, 250_000) # S0,K,T,r,sigma,paths
12
13 with ProcessPoolExecutor(max_workers=4) as executor:
14     prices = list(executor.map(price_european_call, [params]*8))
15 option_price = np.mean(prices)
16
```



## Hands-On Time!

Monte Carlo in Parallel

03\_multiprocessing\_cpu\_bound.ipynb

*Section 2: Option pricing with multiprocessing*

# The Pickling Requirement

Processes have separate memory → data must be serialized (pickled)

## What CAN'T be pickled

Lambda functions, nested functions, file handles, DB connections

```
1 # This will FAIL
2 with ProcessPoolExecutor() as executor:
3     results = executor.map(lambda x: x**2, [1,2,3]) # Error!
4
5 # This WORKS
6 def square(x):
7     return x ** 2
8
9 with ProcessPoolExecutor() as executor:
10    results = executor.map(square, [1,2,3]) # OK!
11
```

# When Parallelization Hurts: The Overhead Trap

```
1 # BAD: Task is too small
2 def add_one(x):
3     return x + 1
4
5 # Overhead of creating processes >> computation time
6 with ProcessPoolExecutor() as executor:
7     results = list(executor.map(add_one, range(100)))
8 # This is SLOWER than sequential!
9
```

## Rule of Thumb

Each task should take at least **10-100ms** to justify process overhead.

**Solution:** Chunk your work into larger batches.



## Hands-On Time!

Chunking and Overhead

03\_multiprocessing\_cpu\_bound.ipynb

*Section 3: Learn when parallelization helps vs. hurts*

# Finance Application: Portfolio VaR

**Value at Risk (VaR):** Maximum expected loss at a confidence level

**Monte Carlo VaR requires:**

- ① Simulate many portfolio return scenarios
- ② Sort returns
- ③ Find the percentile cutoff

**Parallelization strategy:**

- Split simulations across processes
- Each process generates subset of scenarios
- Combine and calculate VaR at the end

Typical Speedup

4-core machine: 3-4x faster

8-core machine: 6-7x faster



## Hands-On Time!

VaR Calculation

03\_multiprocessing\_cpu\_bound.ipynb

*Section 4: Parallel Value-at-Risk*

# Multiprocessing: Key Takeaways

## When to Use

- Monte Carlo simulations
- Parameter grid searches
- Backtesting strategies
- Heavy numerical computation

## Watch Out For

- Overhead for small tasks
- Memory usage (data copied)
- No lambdas (pickling)

## Best Practices

Use `ProcessPoolExecutor` • Chunk small tasks • Profile before/after

# Topic 4

## Real-World Finance Applications

Putting It All Together

04\_finance\_applications.ipynb

# Application 1: Parallel Backtesting

## The Problem:

- Test a trading strategy with different parameters
- 100 parameter combinations  $\times$  10 years of data
- Sequential: hours of waiting

## The Solution:

- Each parameter combination is independent
- Perfect for ProcessPoolExecutor
- Distribute across all CPU cores

## Example Strategy

Moving average crossover: test all combinations of short (5-50 days) and long (20-200 days) windows



## Hands-On Time!

### Parallel Backtesting

04\_finance\_applications.ipynb

*Section 1: Build a parallel strategy backtester*

# Application 2: Bootstrap Confidence Intervals

## The Problem:

- Estimate uncertainty in Sharpe ratio
- Need 10,000+ bootstrap samples
- Each sample: resample data, calculate statistic

## Why it's parallel-friendly:

- Each bootstrap sample is independent
- CPU-bound (resampling + calculations)
- Easy to split: 10,000 samples → 2,500 per core

## Statistical Rigor

Bootstrap gives you confidence intervals without assuming normality — essential for fat-tailed financial returns.



## Hands-On Time!

### Bootstrap Analysis

04\_finance\_applications.ipynb

*Section 2: Parallel bootstrap for Sharpe ratio CI*

# Application 3: Correlation Matrix Computation

## The Problem:

- 500 assets → 124,750 pairwise correlations
- Need rolling correlations over time
- Sequential calculation is slow

## Parallelization Approach:

- Split asset pairs across processes
- Or: parallelize across time windows
- Combine results at the end

## Note

NumPy already parallelizes some operations internally. Profile first!



## Hands-On Time!

Complete Finance Applications

`04_finance_applications.ipynb`

*Section 3: Correlation analysis and wrap-up*



## **Break Time**

10 minutes

Final session: Best practices and projects


# Topic 5

## Best Practices & Projects

Writing Robust Parallel Code

05\_project\_exercises.ipynb

# Decision Tree: Thread or Process?



# Common Pitfall 1: Shared State

## The Problem

Multiple processes modifying the same variable causes unexpected behavior.

```
1 # WRONG - This doesn't work as expected!
2 counter = 0
3
4 def increment():
5     global counter
6     counter += 1 # Each process has its OWN copy!
7
8 with ProcessPoolExecutor() as executor:
9     executor.map(increment, range(100))
10
11 print(counter) # Still 0!
12
```

## Solution

Return values instead of modifying global state. Let the main process aggregate.

# Common Pitfall 2: Too Many Workers

```
1 # WRONG - More workers than cores
2 with ProcessPoolExecutor(max_workers=100) as executor:
3     results = executor.map(cpu_task, data)
4 # Context switching overhead kills performance!
5
6 # RIGHT - Match workers to cores
7 import os
8 n_cores = os.cpu_count()
9 with ProcessPoolExecutor(max_workers=n_cores) as executor:
10     results = executor.map(cpu_task, data)
11
```

## Guidelines

- CPU-bound: workers  $\leq$  number of cores
- I/O-bound: workers can exceed cores (2-4x)
- Memory-heavy: reduce workers to avoid swapping

# Progress Bars with tqdm

```
1 from concurrent.futures import ProcessPoolExecutor, as_completed
2 from tqdm import tqdm
3
4 def slow_task(x):
5     # ... some computation
6     return x ** 2
7
8 items = range(100)
9
10 with ProcessPoolExecutor() as executor:
11     futures = [executor.submit(slow_task, x) for x in items]
12
13 results = []
14 for future in tqdm(as_completed(futures), total=len(items)):
15     results.append(future.result())
16
```

## Output

```
100%|=====| 100/100 [00:05<00:00, 18.32it/s]
```

# Beyond concurrent.futures

When you need more power:

## joblib      Dask      Ray

<ul style="list-style-type: none"><li>• Simple API</li><li>• Memory mapping</li><li>• Good for NumPy</li><li>• scikit-learn uses it</li></ul>	<ul style="list-style-type: none"><li>• Parallel DataFrames</li><li>• Larger-than-RAM data</li><li>• Lazy evaluation</li><li>• Scales to clusters</li></ul>	<ul style="list-style-type: none"><li>• Distributed computing</li><li>• Actor model</li><li>• ML focused</li><li>• Production ready</li></ul>
---	---	---

Start Simple

concurrent.futures handles 90% of use cases. Only reach for specialized tools when you hit its limits.

# Debugging Parallel Code

## Parallel bugs are hard to find:

- Non-deterministic behavior
- Errors in worker processes
- Hard to reproduce

## Strategies:

- ① **Start sequential:** Make sure code works with `max_workers=1`
- ② **Catch exceptions:** Always wrap `future.result()` in `try/except`
- ③ **Test with small data:** Faster iteration, easier to spot issues
- ④ **Use logging:** Print statements get mixed up

## Golden Rule

If it works with `max_workers=1`, it should work with more. If not, you have a parallelism bug.




## Hands-On Time!

Mini-Project Time!

`05_project_exercises.ipynb`

*Choose your project and apply what you've learned*

# Summary: Your Parallel Python Toolkit



## Key Message

You now have the tools to make your finance code run 4-8x faster. Use them wisely!

# Thank You!

## Questions?

Key takeaways:

- I/O-bound → ThreadPoolExecutor
- CPU-bound → ProcessPoolExecutor
- Always profile before and after
- Chunk small tasks to reduce overhead

Happy parallel programming!