

# When to Apply Parallelism in Practice

**Douglas C. Schmidt**

**[d.schmidt@vanderbilt.edu](mailto:d.schmidt@vanderbilt.edu)**

**[www.dre.vanderbilt.edu/~schmidt](http://www.dre.vanderbilt.edu/~schmidt)**

**Professor of Computer Science**

**Institute for Software  
Integrated Systems**

**Vanderbilt University  
Nashville, Tennessee, USA**



# Learning Objectives in this Part of the Lesson

---

- Understand the meaning of key concepts associated with parallel programming
- Know when to apply parallelism in practice
  - i.e., what conditions must apply to choose parallelism as the programming paradigm




---

# When to Apply Parallelism in Practice

# When to Apply Parallelism in Practice

- Parallelism is not a panacea!!

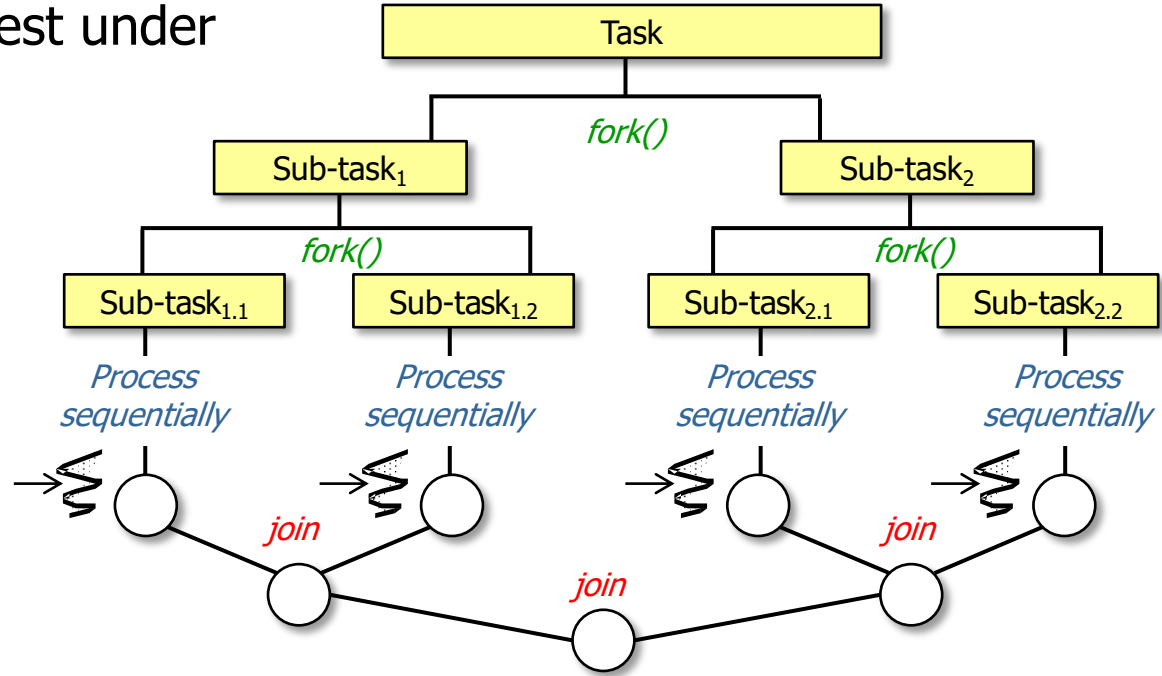


Particularly when there's contention for shared resources

# When to Apply Parallelism in Practice

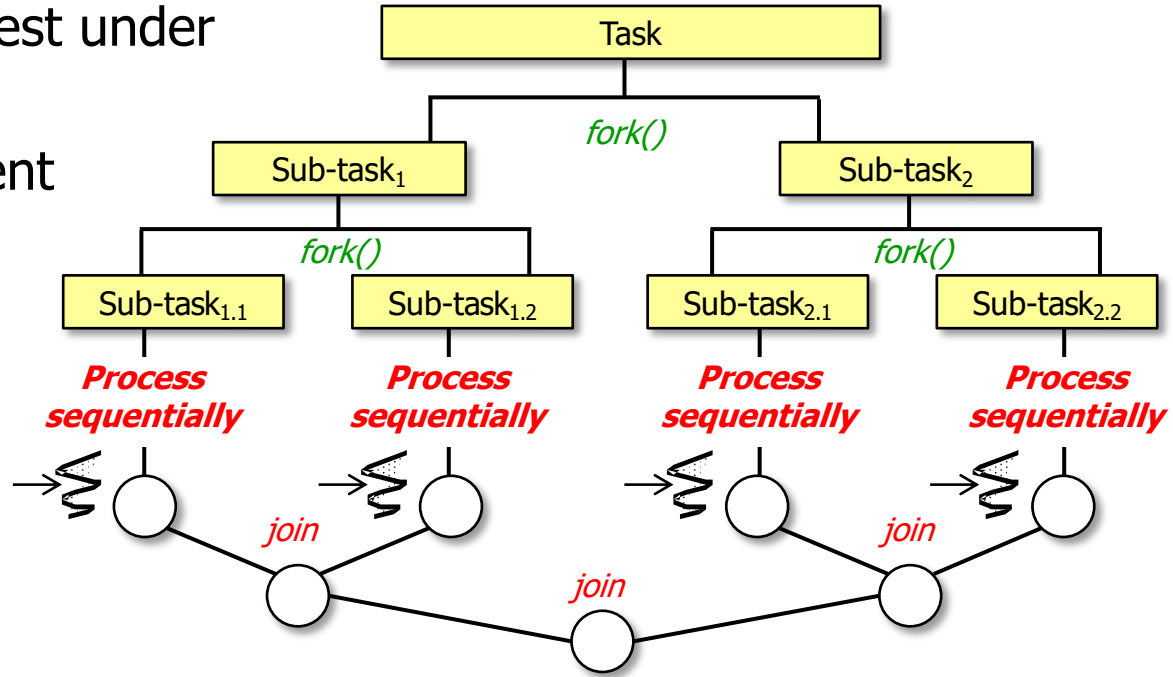
- Instead, parallelism works best under certain conditions

BEST IF USED BY



# When to Apply Parallelism in Practice

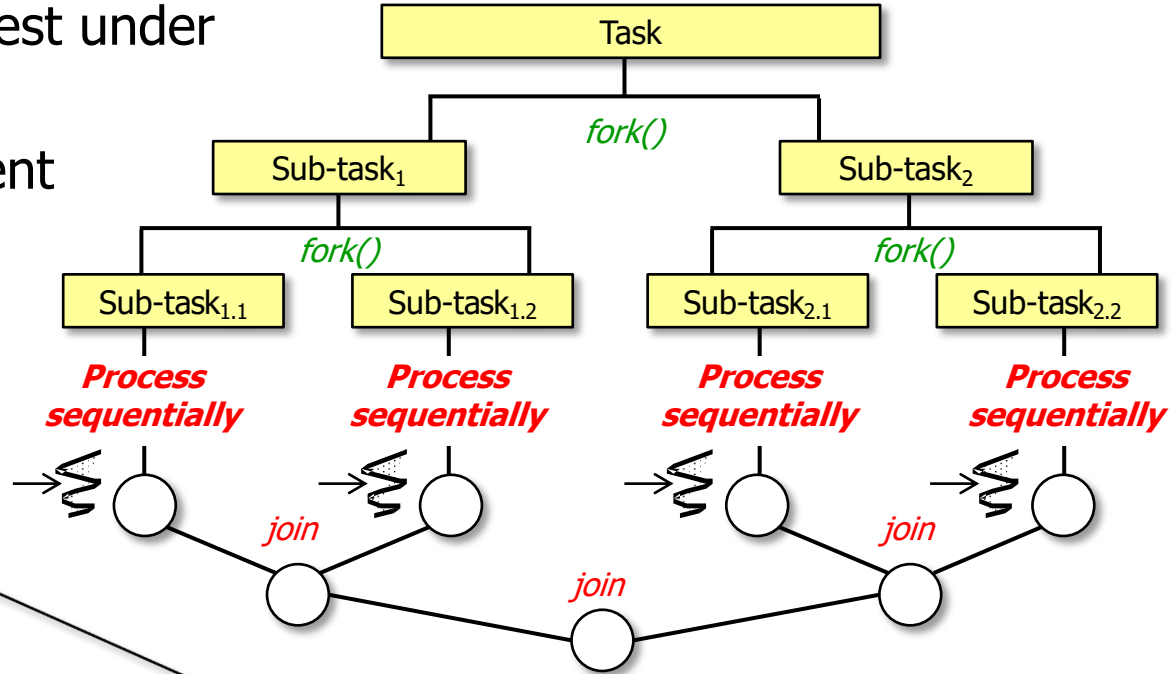
- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent





# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent

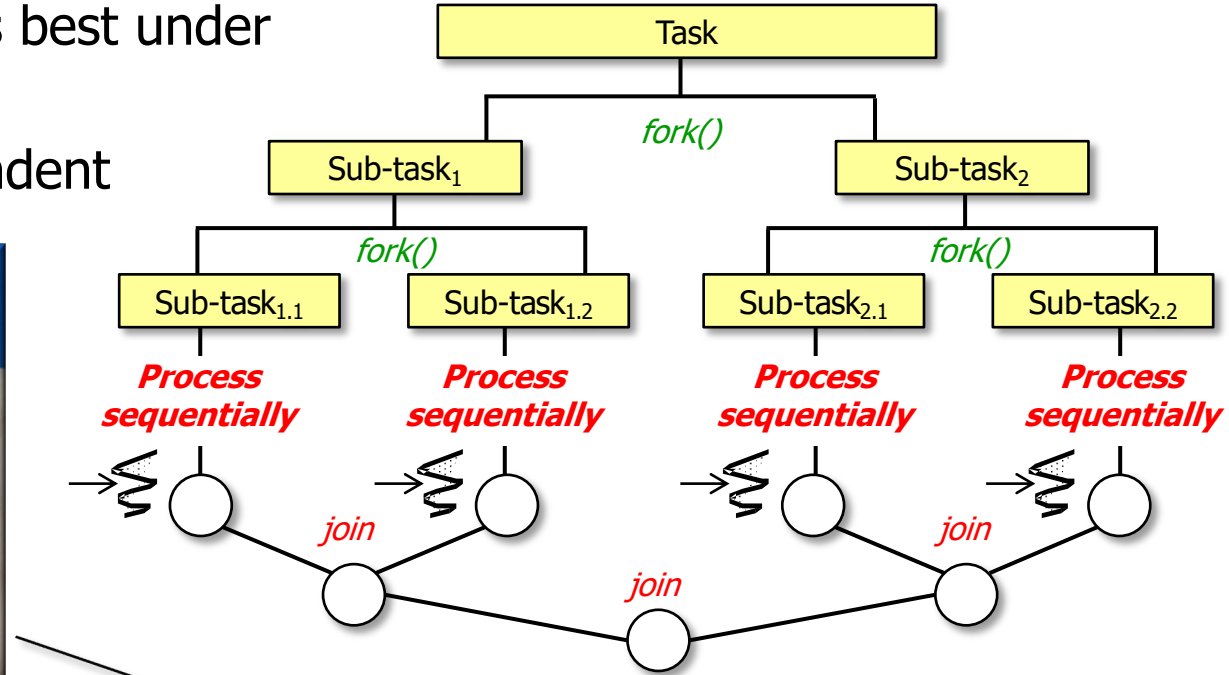



*"Embarrassingly parallel" tasks have little/no dependency or need for communication between tasks or for sharing results between them*

See [en.wikipedia.org/wiki/Embarrassingly\\_parallel](https://en.wikipedia.org/wiki/Embarrassingly_parallel)

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent

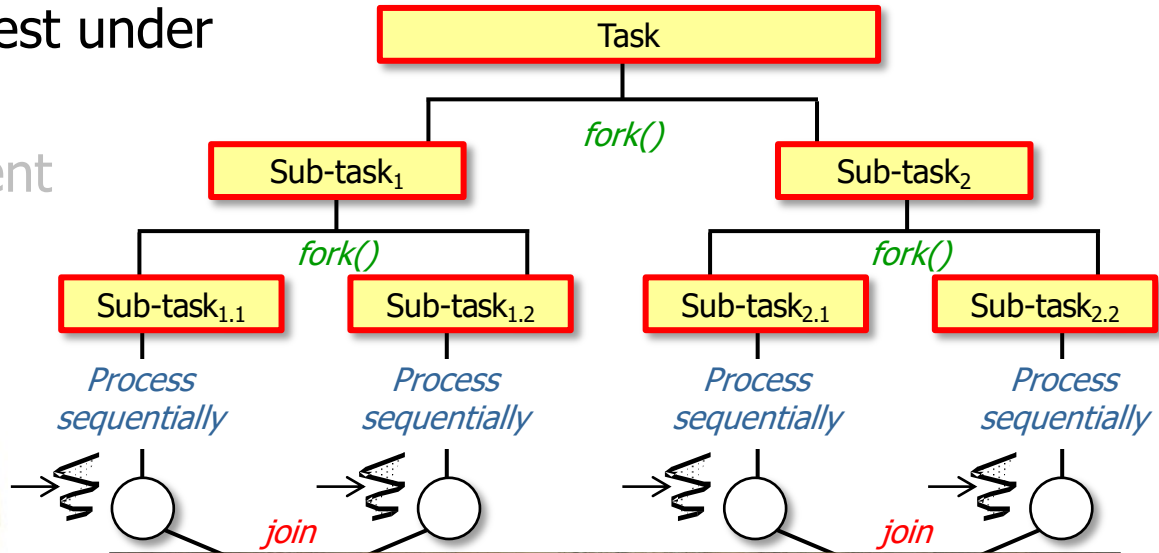


*"Embarrassing" in this context means "overabundance" or "too much of a good thing"!*



# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform

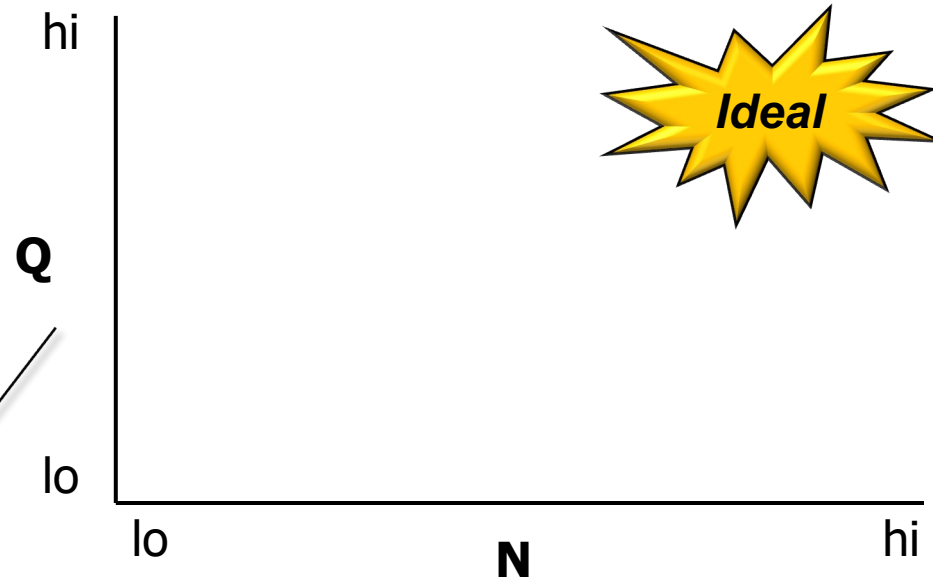


See [en.wikipedia.org/wiki/Terracotta Army](https://en.wikipedia.org/wiki/Terracotta_Army)

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.

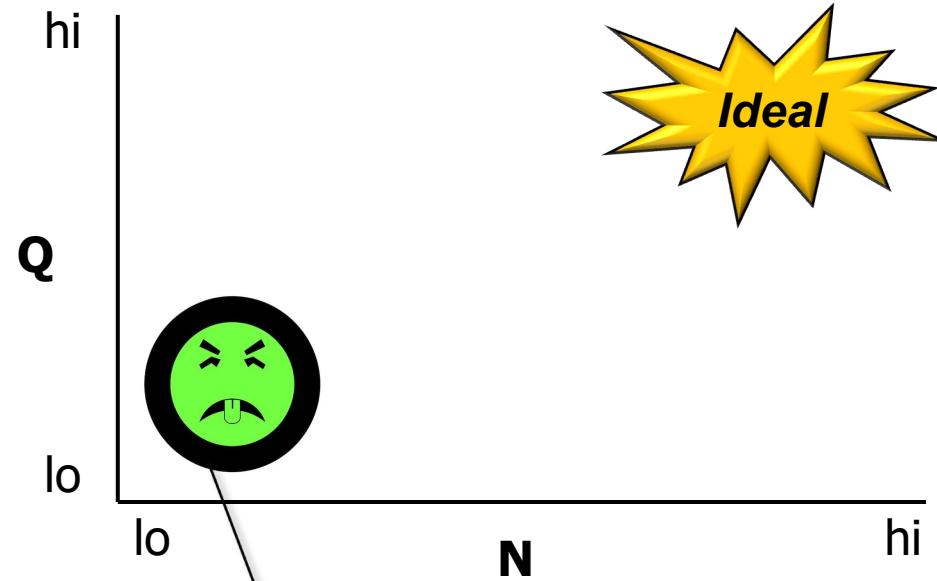
- When tasks are independent
- When there's lots of data & processing to perform
- The "N\*Q" heuristic estimates the benefit of parallelism



- *N is the # of data elements to process*
- *Q quantifies CPU processing intensity for each data element*

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
  - The "N\*Q" heuristic estimates the benefit of parallelism

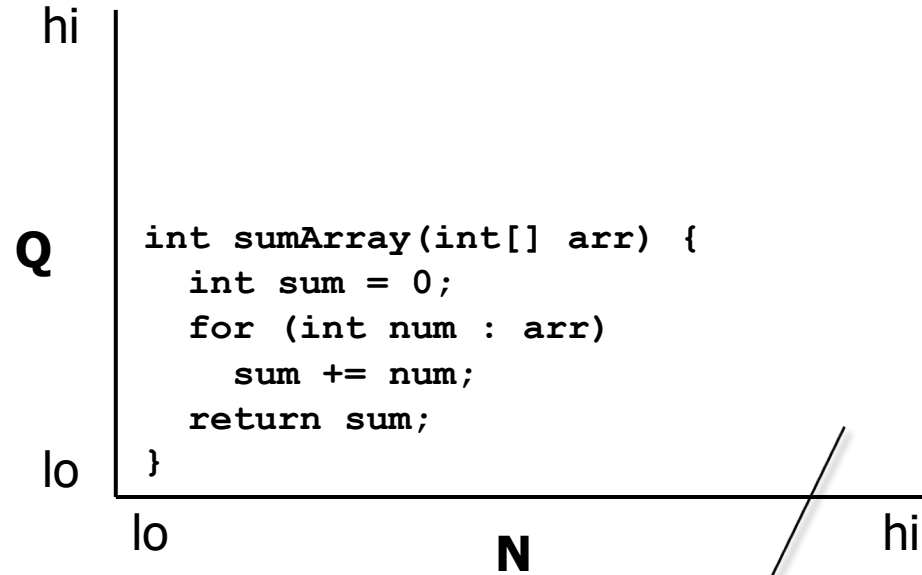


**Low N, Low Q:** *The situation generally does not favor parallelization due to overhead costs incurred*

In this case, it's usually best to stick with sequential programming

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
  - The "N\*Q" heuristic estimates the benefit of parallelism

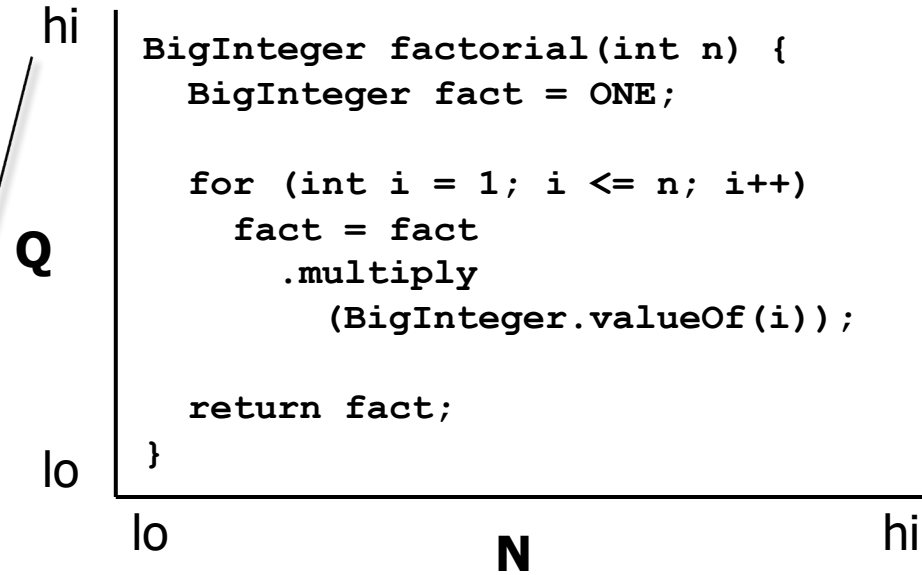


**High N, Low Q:** The overhead of parallelizing may outweigh the benefits, as the computational work per element is trivial

Brian Goetz recommends 'N' be > 10,000

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
  - The "N\*Q" heuristic estimates the benefit of parallelism

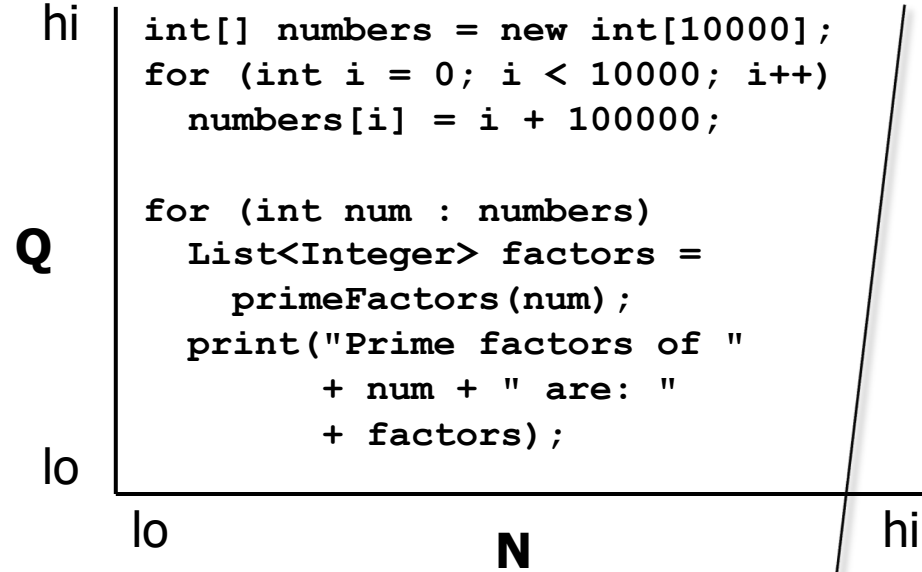


**Low N, High Q:** The computational workload for each data element is high, so even a small N can benefit from parallelization because work can be partitioned across multiple cores, thereby reducing the total time for computation

Often seen in simulations, complex math computations, or graphics rendering

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
  - The "N\*Q" heuristic estimates the benefit of parallelism

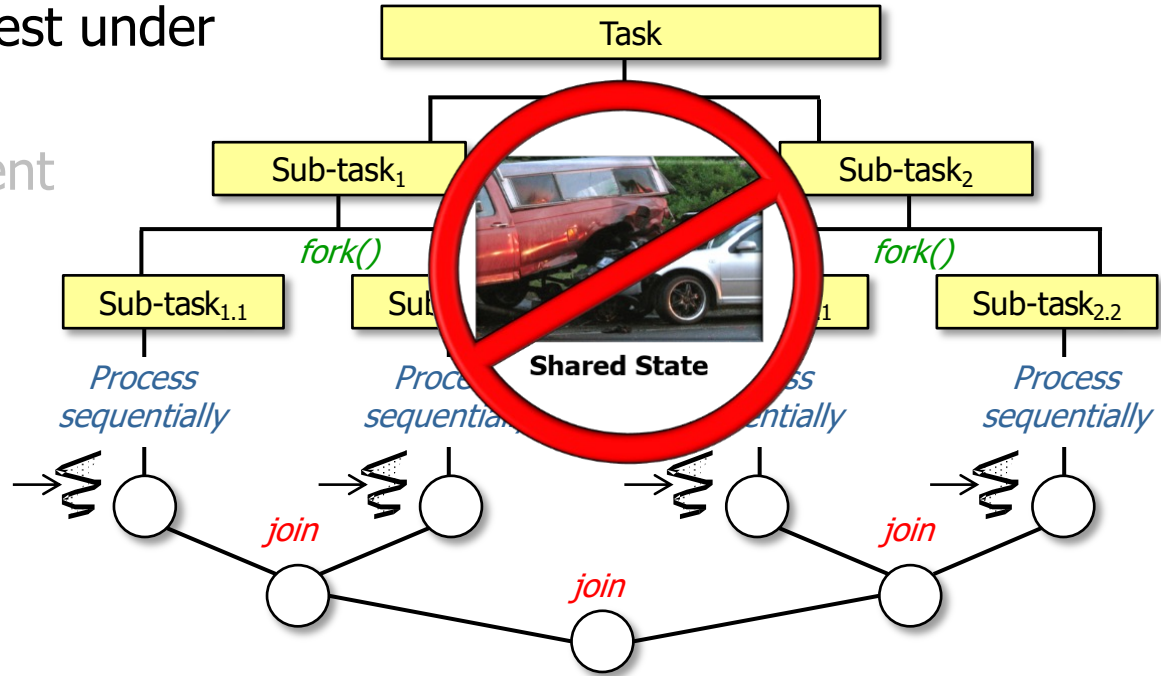


**High N, High Q:** The potential for parallel speedup is significant



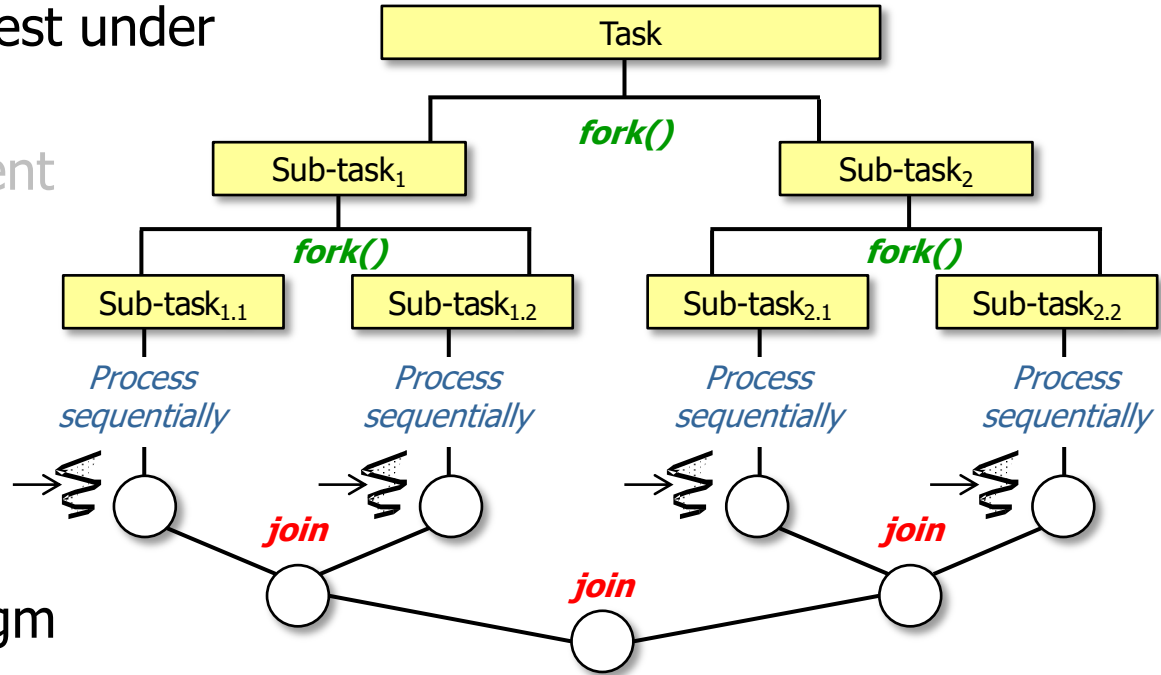
# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
  - When tasks neither block nor share mutable state



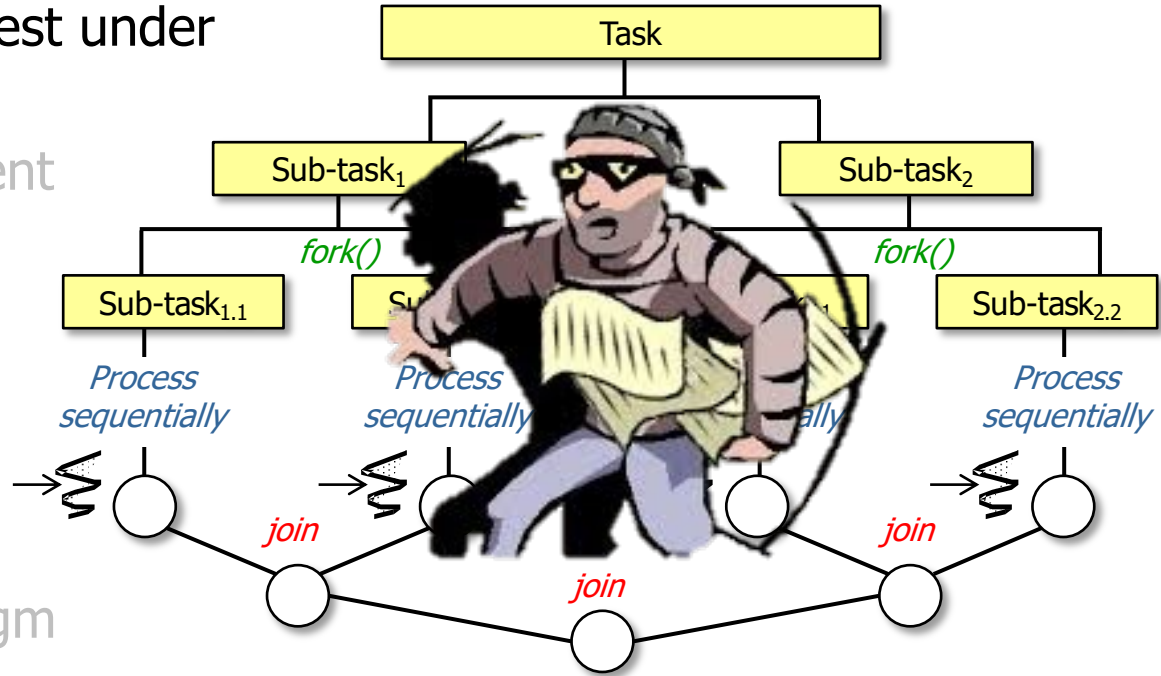
# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
- When tasks neither block nor share mutable state
  - Hence Java's focus on
    - The "fork-join" paradigm
    - To avoid sharing mutable state



# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.
  - When tasks are independent
  - When there's lots of data & processing to perform
  - When tasks neither block nor share mutable state
    - Hence Java's focus on
      - The "fork-join" paradigm
      - "Work-stealing"
      - To avoid blocking

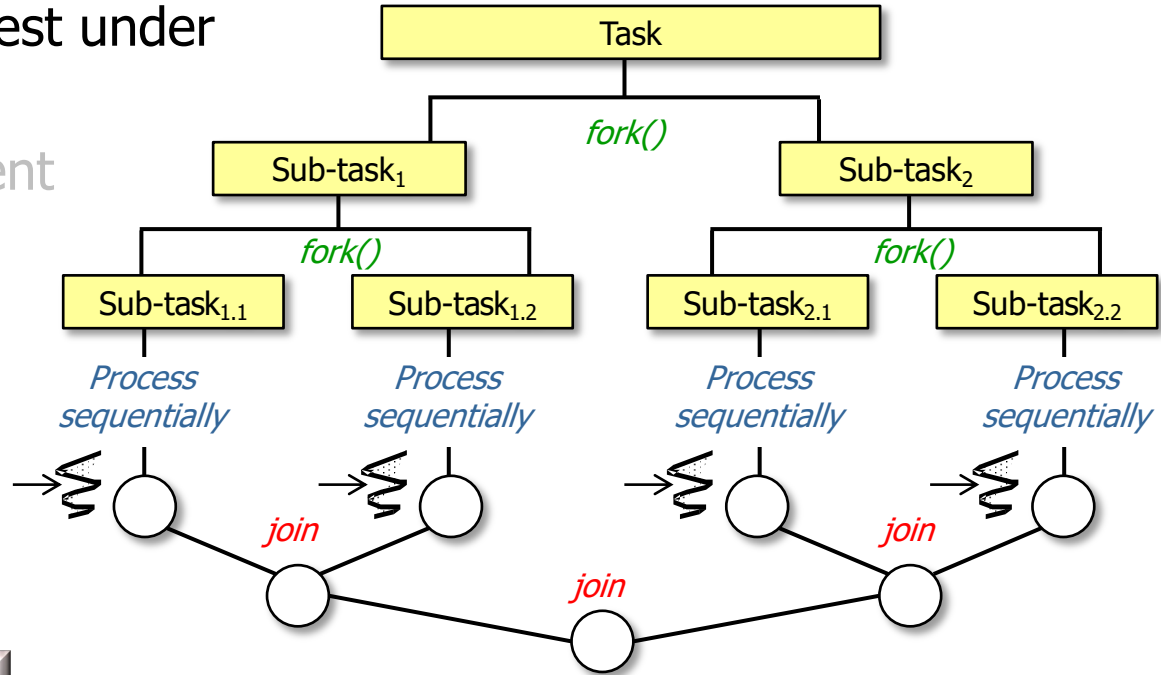


See [en.wikipedia.org/wiki/Work\\_stealing](https://en.wikipedia.org/wiki/Work_stealing)

# When to Apply Parallelism in Practice

- Instead, parallelism works best under certain conditions, e.g.

- When tasks are independent
- When there's lots of data & processing to perform
- When tasks neither block nor share mutable state
- When there are many cores and/or processors



*The more, the merrier*



---

# End of When to Apply Parallel Programming in Practice