


B3CC: Concurrency

01: Introduction

Ivo Gabe de Wolff

Hello!

- Ivo Gabe de Wolff
 - i.g.dewolff@uu.nl
- Tom Smeding
 - t.j.smeding@uu.nl
- Most slides and materials were made by Trevor L. McDonell
- Working groups are guided by
 - Ivo Gabe de Wolff
 - Niek Mulleners
 - And 6 TAs



Today

1. Motivation
2. Course formalities

Motivation


Motivation

- **Concurrency:** dealing with lots of things at once
 - Collection of independently executing processes
 - Two or more threads are making progress
- **Parallelism:** doing lots of things at once
 - Simultaneous execution of (possibly related) computations
 - Two or more threads are executing simultaneously

Performance improvements over the years

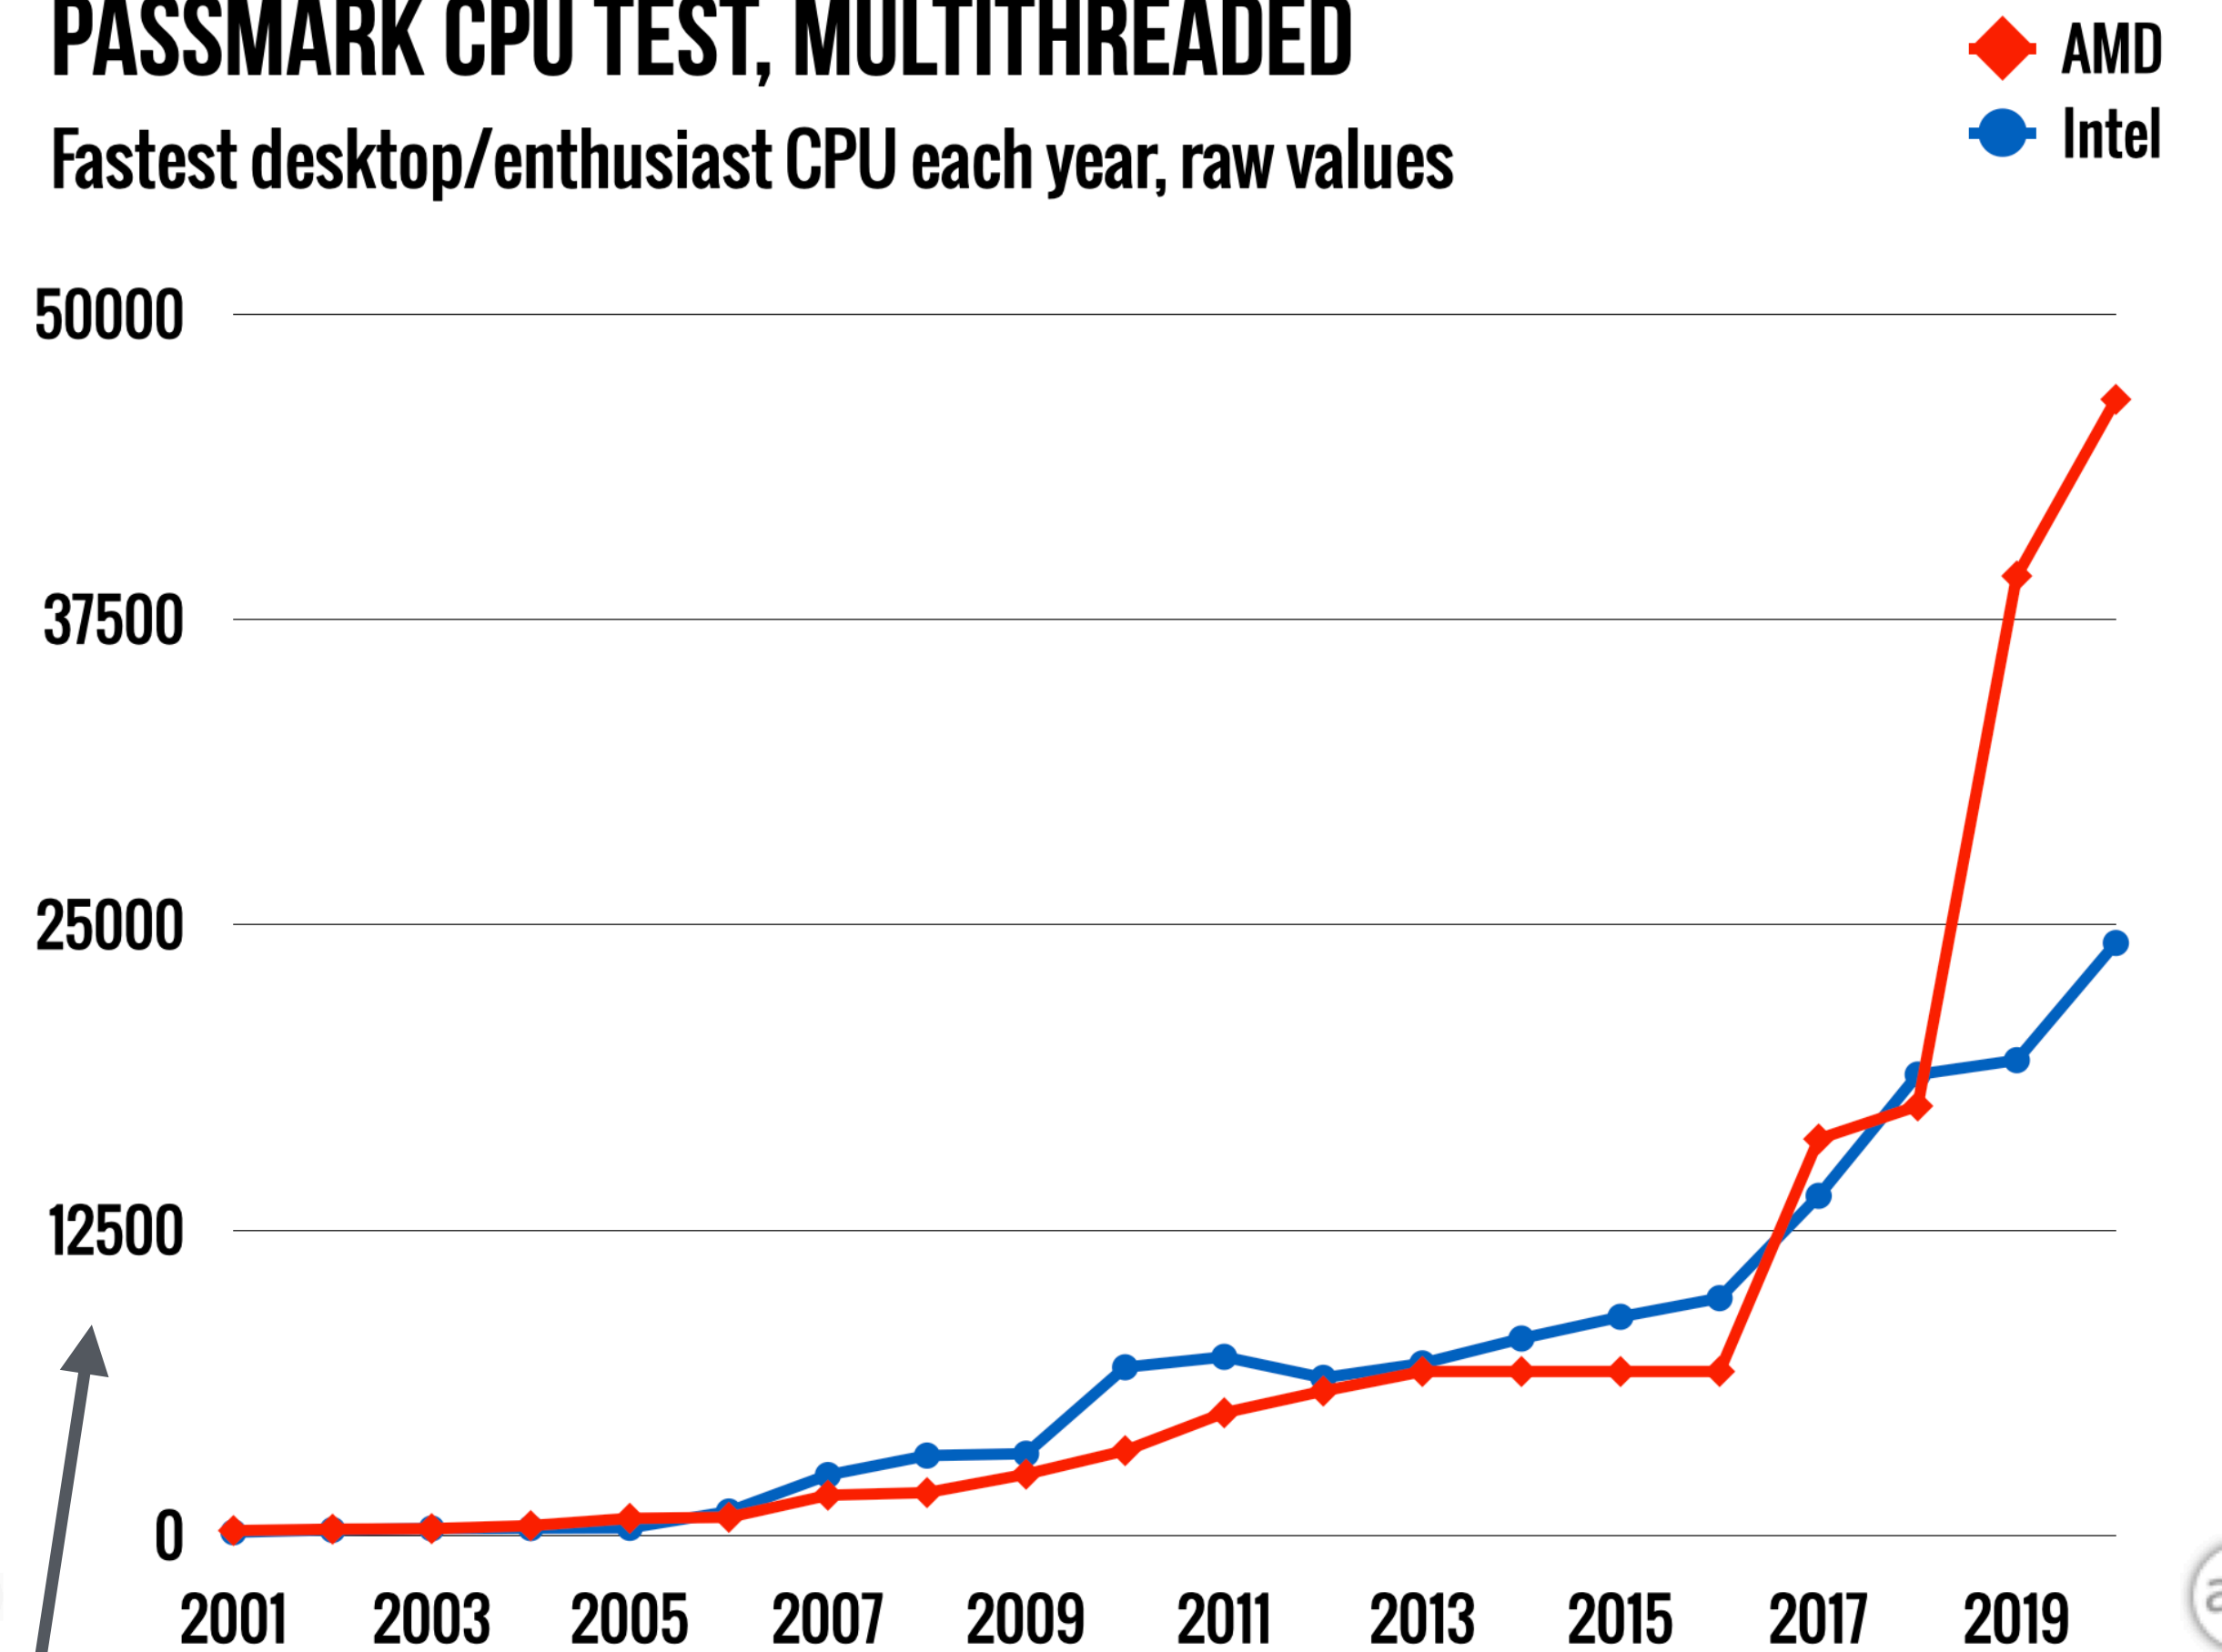
PASSMARK CPU TEST, SINGLETHREADED, 2001-2020

Fastest desktop/enthusiast CPU each year, raw values



PASSMARK CPU TEST, MULTITHREADED

Fastest desktop/enthusiast CPU each year, raw values



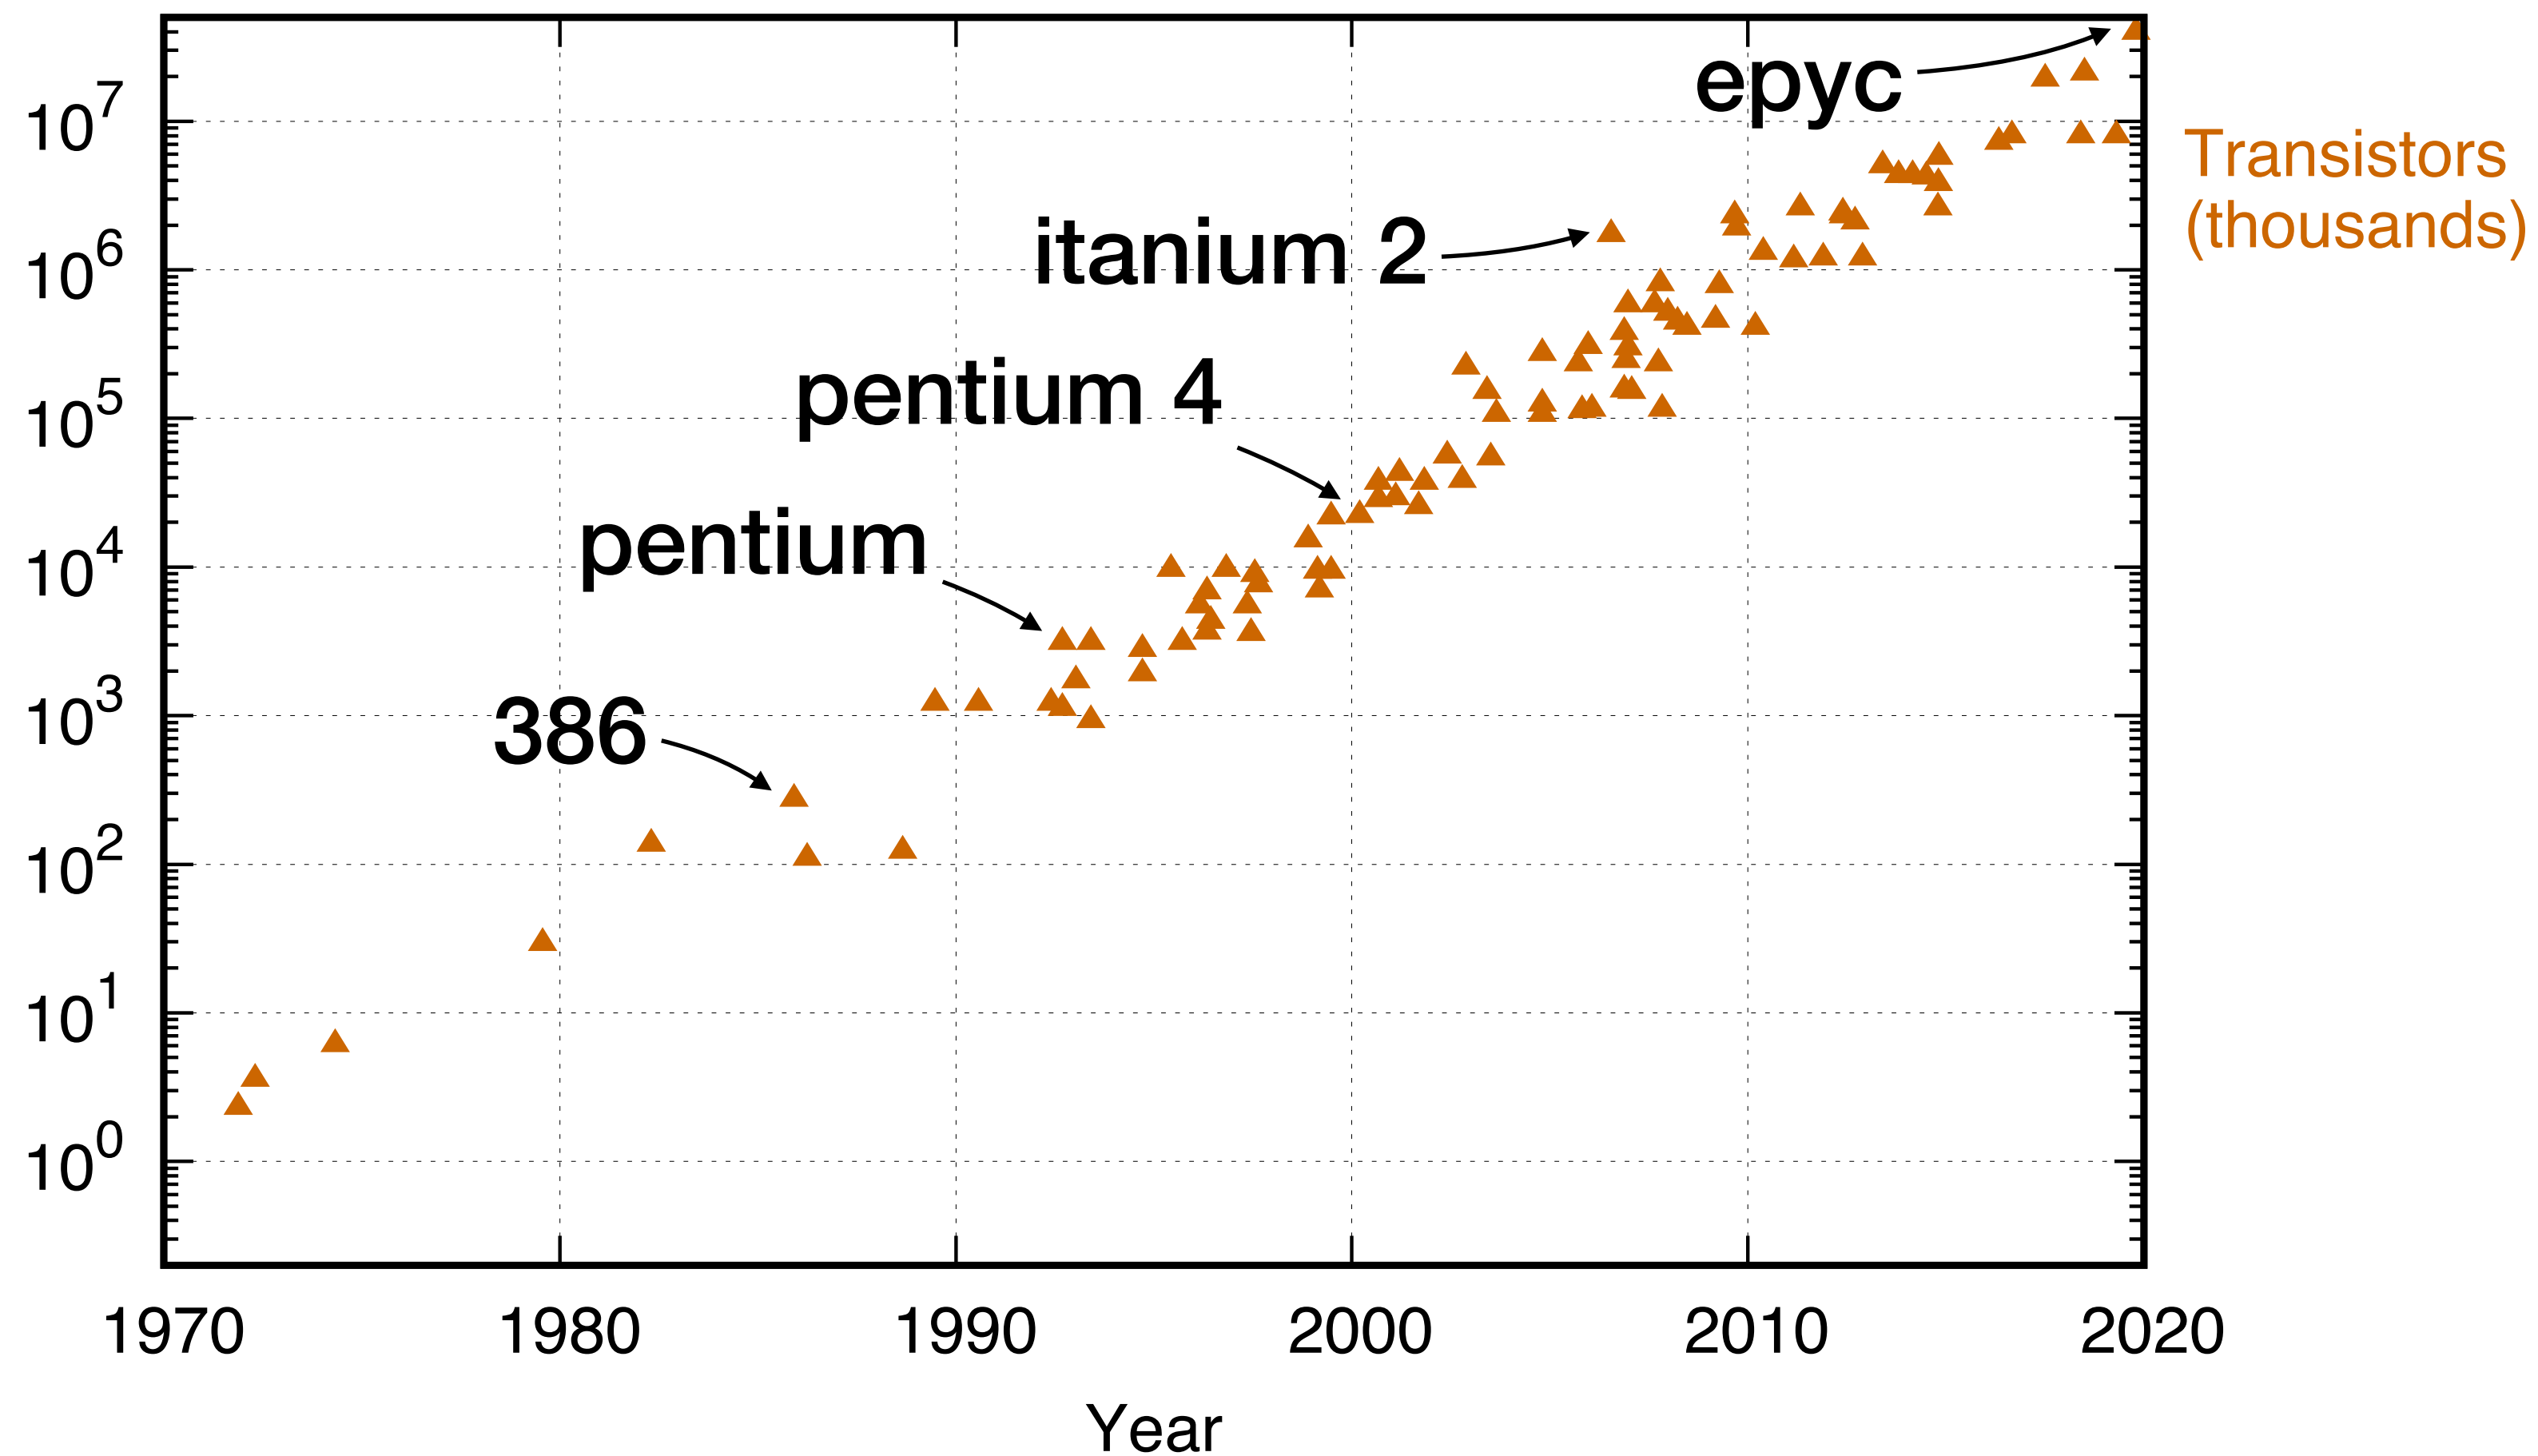
Different scale

The free lunch is over

- “The free lunch is over” (2005)
 - Today virtually all processors include multiple cores/processing elements
 - This has become the primary method for increasing performance
 - This has consequences for the programmer

Why?

48 Years of Microprocessor Trend Data



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2019 by K. Rupp

Why?


- Moore's curve (1965)
 - Observation that the number of transistors in an integrated circuit doubles roughly every two years
 - Based on production cost and yield (success rate of production) of chips
 - Not a law in any sense of the word (don't call it that)

Why?

- Dennard scaling: As transistors get smaller, power density remains constant
 - Voltage and current decrease at same rate as transistor size
- Smaller transistors allowed higher clock frequencies
 - As signal delays are proportional to transistor size

Why?

48 Years of Microprocessor Trend Data



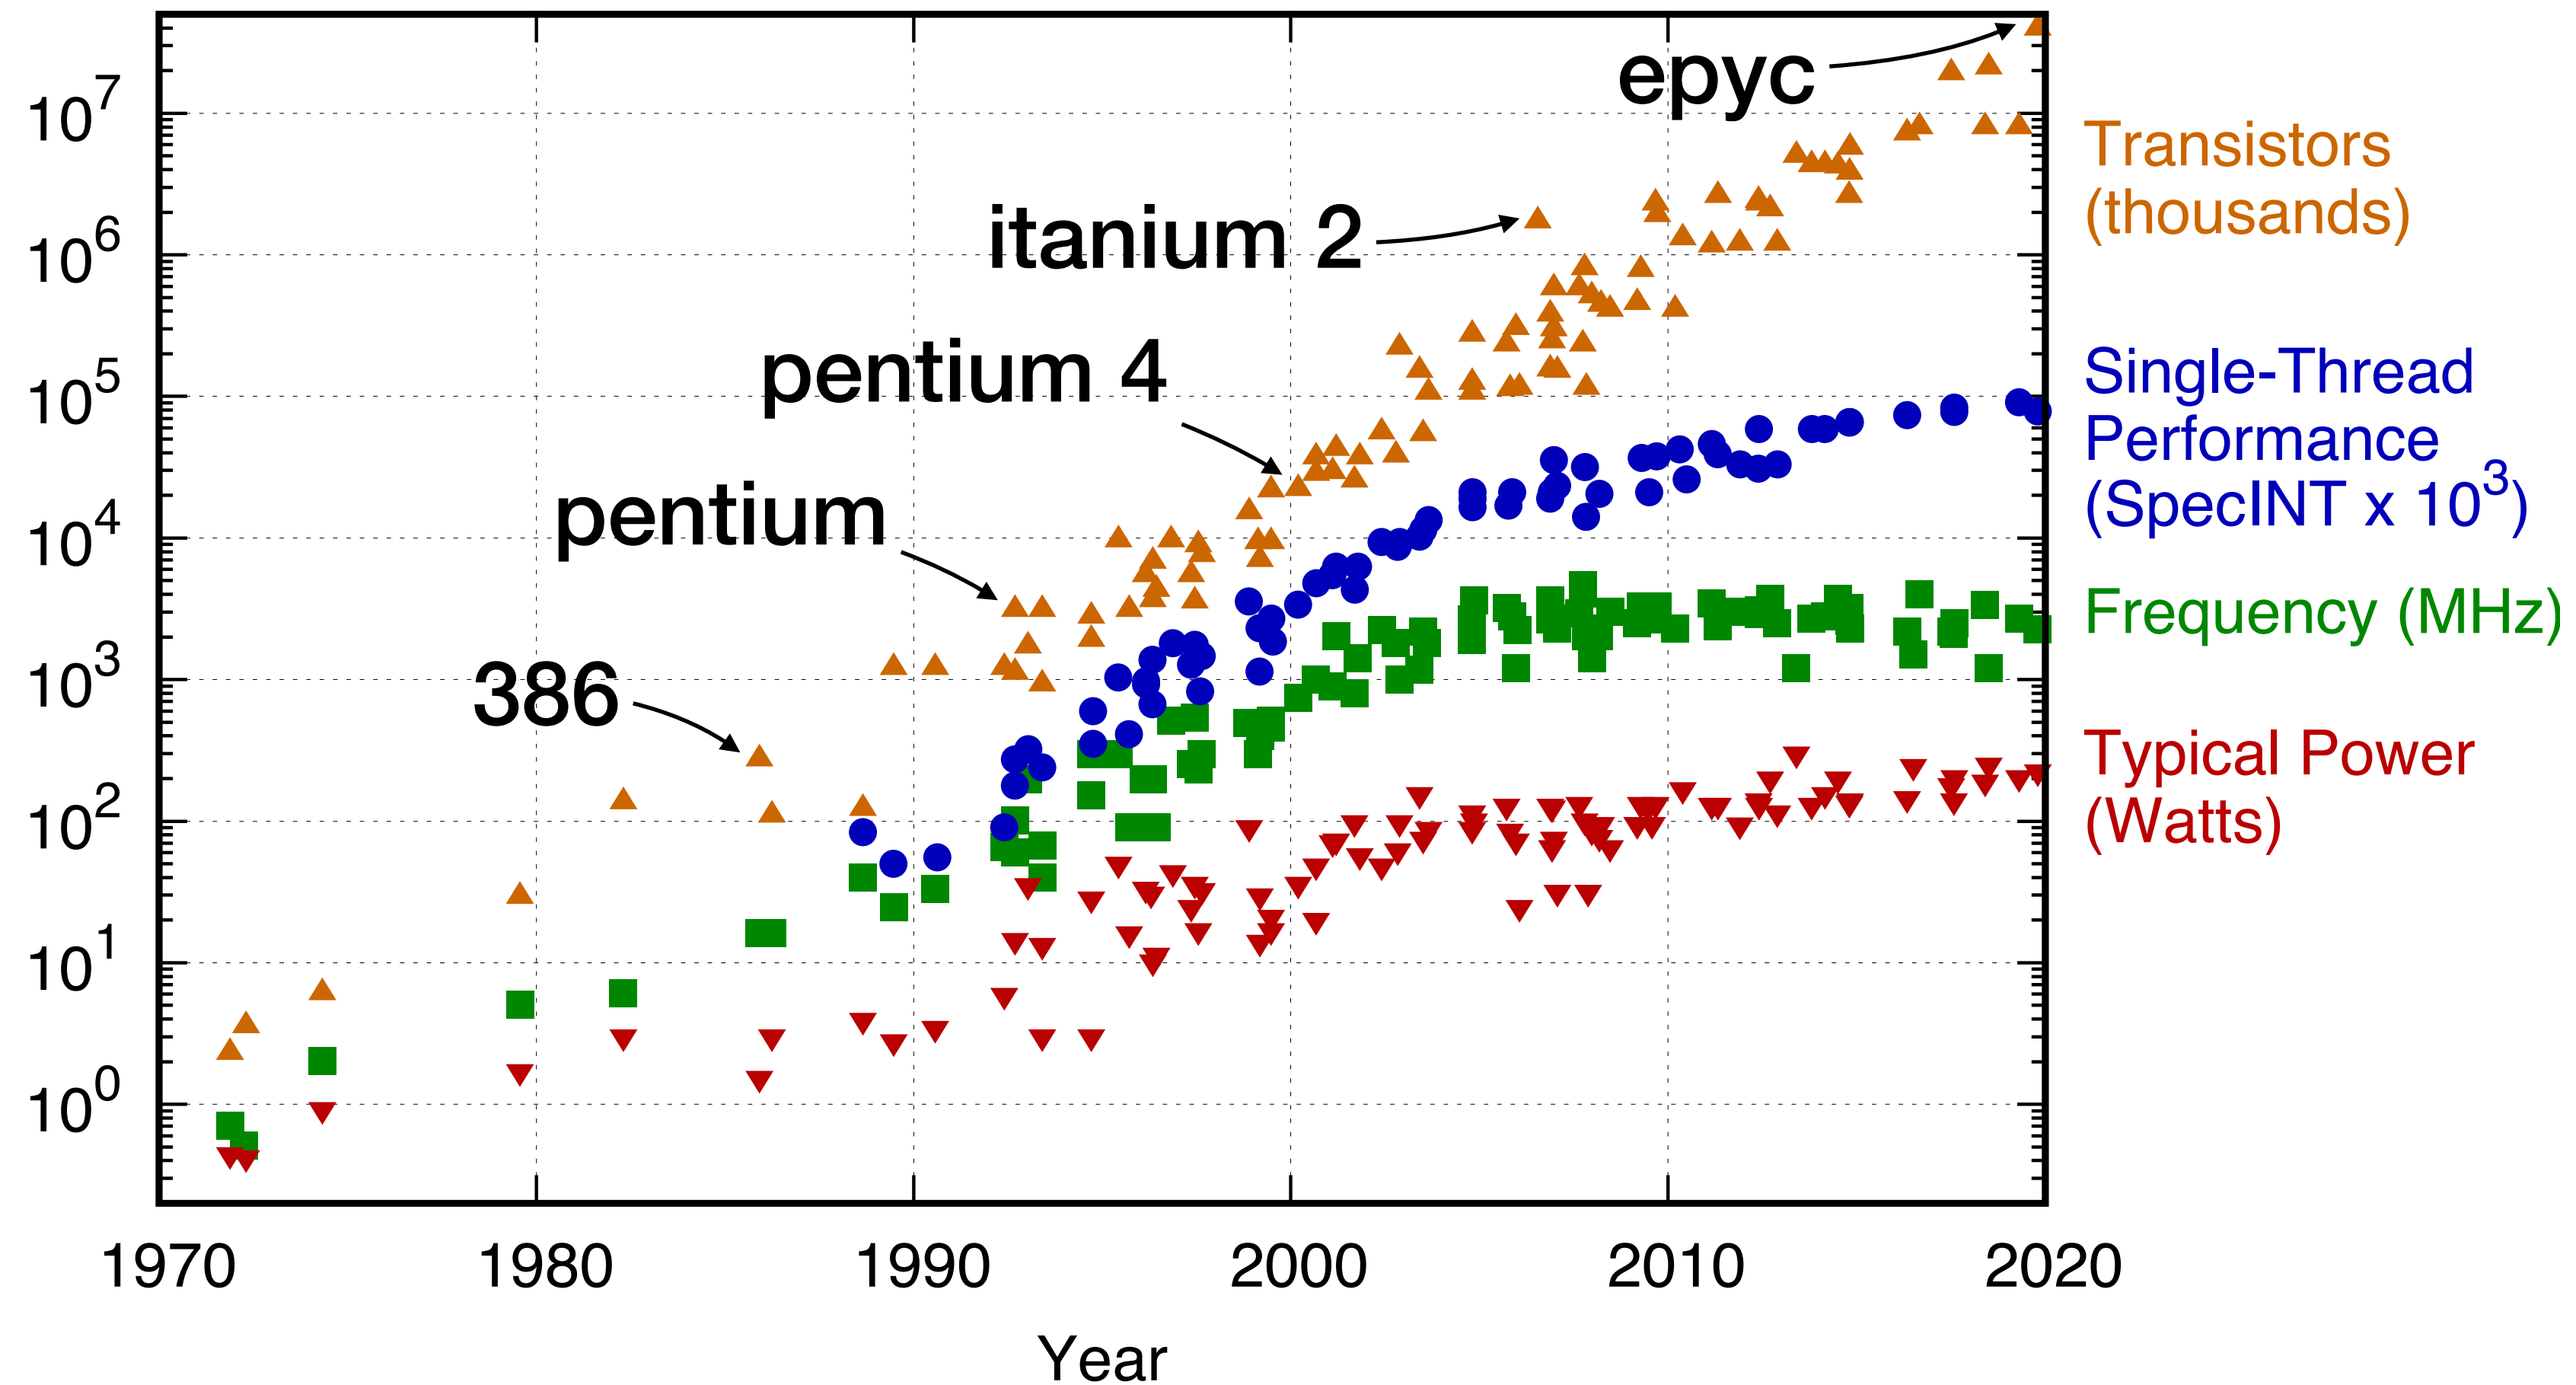
Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2019 by K. Rupp

Why?

- Since ~2005 Dennard scaling breaks down
 - Current leakage increased the power usage; power density wasn't constant any more
 - Consequence: can no longer improve performance through frequency scaling alone

Why?

48 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2019 by K. Rupp

Why?

- Traditional approaches to increasing CPU performance:
 - Frequency scaling
 - Caches
 - Micro-architectural improvements
 - Out of order execution (increase utilisation of execution hardware)
 - Branch prediction (guess the outcome of control flow)
 - Speculative execution (do work before knowing if it will be needed)


Why?

- Frequency scaling: The Power Wall
 - Power consumption of transistors does not decrease as fast as density increases
 - Performance limited by power consumption (& dissipation)




Why?

- Caches: The Memory Wall
 - Memory speed does not increase as fast as computing speed
 - Increasingly difficult to hide memory latency




Why?

- Microarchitecture improvements: Instruction Level Parallelism Wall
 - Law of diminishing returns
 - Pollack rule: performance $\propto \sqrt{\text{complexity}}$



Why?

50 Years of Microprocessor Trend Data



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

Performance walls

Limitations for single core performance:

- Power Wall
- Memory Wall
- ILP Wall


Instead, processors became faster by adding more cores

Demo: Counting in parallel

Counting in parallel

Say we need to count the number of persons in this room.

We can do that in parallel!

- Stand up, and remember the number 1
 - Form pairs, if possible
 - Add your two numbers, and now remember that number instead
 - One per pair sits down
 - Again form pairs of two standing persons, and repeat
- 

Concurrency bugs

Therac-25


- Computer controlled medical radiation device (1982)
 - Two operating modes: a low-current electron beam; or high-energy x-rays
- Involved in at least six incidents, resulting in serious injury or death
 - A *race condition* could cause the high-power electron beam to be administered directly to the patient
 - Resulted in radiation doses 100x higher than normal
 - Additional problems related to poor software development practices

Northeast blackout (2003)



- Widespread power outage throughout USA and Canada
 - Second most widespread blackout in history (at the time)
 - Affected an estimated 10M people in Ontario and 45M in 8 US states
 - A *race condition* prevented an alarm from going off
 - Operators were unaware of the need to redistribute power—a minor problem—which cascaded into complete collapse of the electrical grid

Mars Pathfinder



- Launched by NASA in 1996
- *Sojourner* became the first rover to operate outside the Earth-Moon system
- Control computer contained a *priority inversion* bug
 - Triggered under certain high loads causing a system reset
 - Successfully patched remotely

Conclusion?

Theory



Practice



Course formalities

Topics

- Program a multithreaded application
 - Managing threads
 - Synchronise with locks, etc.
 - Software transactional memory
 - Parallelism
- Analyse parallel algorithms with work & span
- Design and implement concurrent algorithms / data structures

Goals

- By the end of the course you should be able to:
 - Design and implement a multithreaded application
 - Understand the difference between concurrency and parallelism
 - Reason about the properties/complexity of parallel algorithms

Homepage

- <https://ics-websites.science.uu.nl/docs/vakken/b3cc/>
 - Feel free to let me know if there are broken links, missing slides, etc.

Teams

- There is an MS Teams Team again this year
 - Join with code: k2st410

Sessions

- Lectures:
 - Mon 17:15 - 19:00, or Mon 15:15 - 17:00, or Tue 13:15 - 15:00
 - Thu 9:00 - 10:45
 - Recorded (at your risk, please remind us if the recording light is off!)
- Working groups:
 - Tue 15:15 - 17:00
 - Thu 11:00 - 12:45
- Participation is expected (please ask questions!)
- **Completing the working group sets is the best way to prepare for both the exams *and* the practicals**

Course components

- Exam (50%)
 - Mid session exam: 17-12-2024 (50%)
 - Final exam: 28-01-2025 (50%)
- Practicals (50%)
 - Assignment 1: 29-11-2024 (individual) (20%)
 - Assignment 2: 20-12-2024 (in pairs) (40%)
 - Assignment 3: 24-01-2025 (in pairs) (40%)
 - If you want to reuse practical grades from last year, send me an email (i.g.dewolff@uu.nl)

Haskell


- General concepts for concurrency are similar across programming languages,
- but some provide better support for concurrency.
- Haskell:
 - Separates pure and impure code, e.g. thread-local work and actual concurrency work
 - Has good support for building and studying abstractions

Haskell

- Haskell allows us to study both low-level and high-level aspects of concurrency.
 - P1: basics of forking threads, locks and communication
 - P3: design and implementation of a quicksort-like algorithm for GPUs
- Tomorrow (15:15): Haskell recap, focusing on the important parts for this course
 - Very boring if you just passed Functional Programming
 - Who is interested?
- Sometimes we use C in examples in the slides, to focus on low-level details

Resources

- Parallel and Concurrent Programming in Haskell
 - <https://simonmar.github.io/pages/pcph.html>
- Many more on the website
 - <https://ics-websites.science.uu.nl/docs/vakken/b3cc/resources.html>



Practicals

- P1: already available, you can start with a sequential (single-threaded) implementation
 - <https://ics-websites.science.uu.nl/docs/vakken/b3cc/assessment.html>

Software installation

- A recent version of GHC (9.4.8)
 - Instructions on our website

Next time...

Tuesday 15:15

- Haskell recap (BBG 161)
- Or: Working groups: start with P1 (BBG 001, 083, 061)

Thursday

- 9:00: Lecture: Threads
- 11:00: Working group: Exercises or P1



Save the date:

On November 14th we protest.

The current Dutch government is planning:

- Caps on the influx of international talent
- Fines for taking longer to complete a degree
- Destructive cuts in academic research funding

Take action to save Dutch higher education and research!

- 1.** Join the protest on November 14th
13:00-15:00 @ Moreelsepark, Utrecht
- 2.** Sign the petition against the budget cuts
- 3.** Help mobilize fellow students & colleagues!
linktr.ee/WOinActie

PROTEST SIGN-UP



PETITION

