



# „Why and how to measure greenness of algorithms“

Presentation at Green Software Foundation

Jens Oehlschlägel, Munich, 28.4.2022

# Jens Oehlschägel

## Education

Dipl. Psych. specialized on empirical problem solving  
Dr. hum. biol. in applied medical statistics  
Heterogenous Systems Engineer  
McKinsey Mini-MBA

## Experience

Programming since 1978, S/R-programming since 1996, learning Julia  
Expert Data Scientist at McKinsey 2001-2011  
Expert System Architect focusing OLTP and OLAP Database Clusters at Bertelsmann since 2011

## *greeNsort*® project

virtually save a nuclear power station, the technology is developed

## Aspiration

full-time research and teaching on energy-efficient algorithms





## Private

married, playing piano














greeNsort®

-  **greeNsort®** makes software greener
-  less runtime, energy and CO2
-  less memory hence better hardware amortization
-  saves up to 4 nuclear power stations worldwide

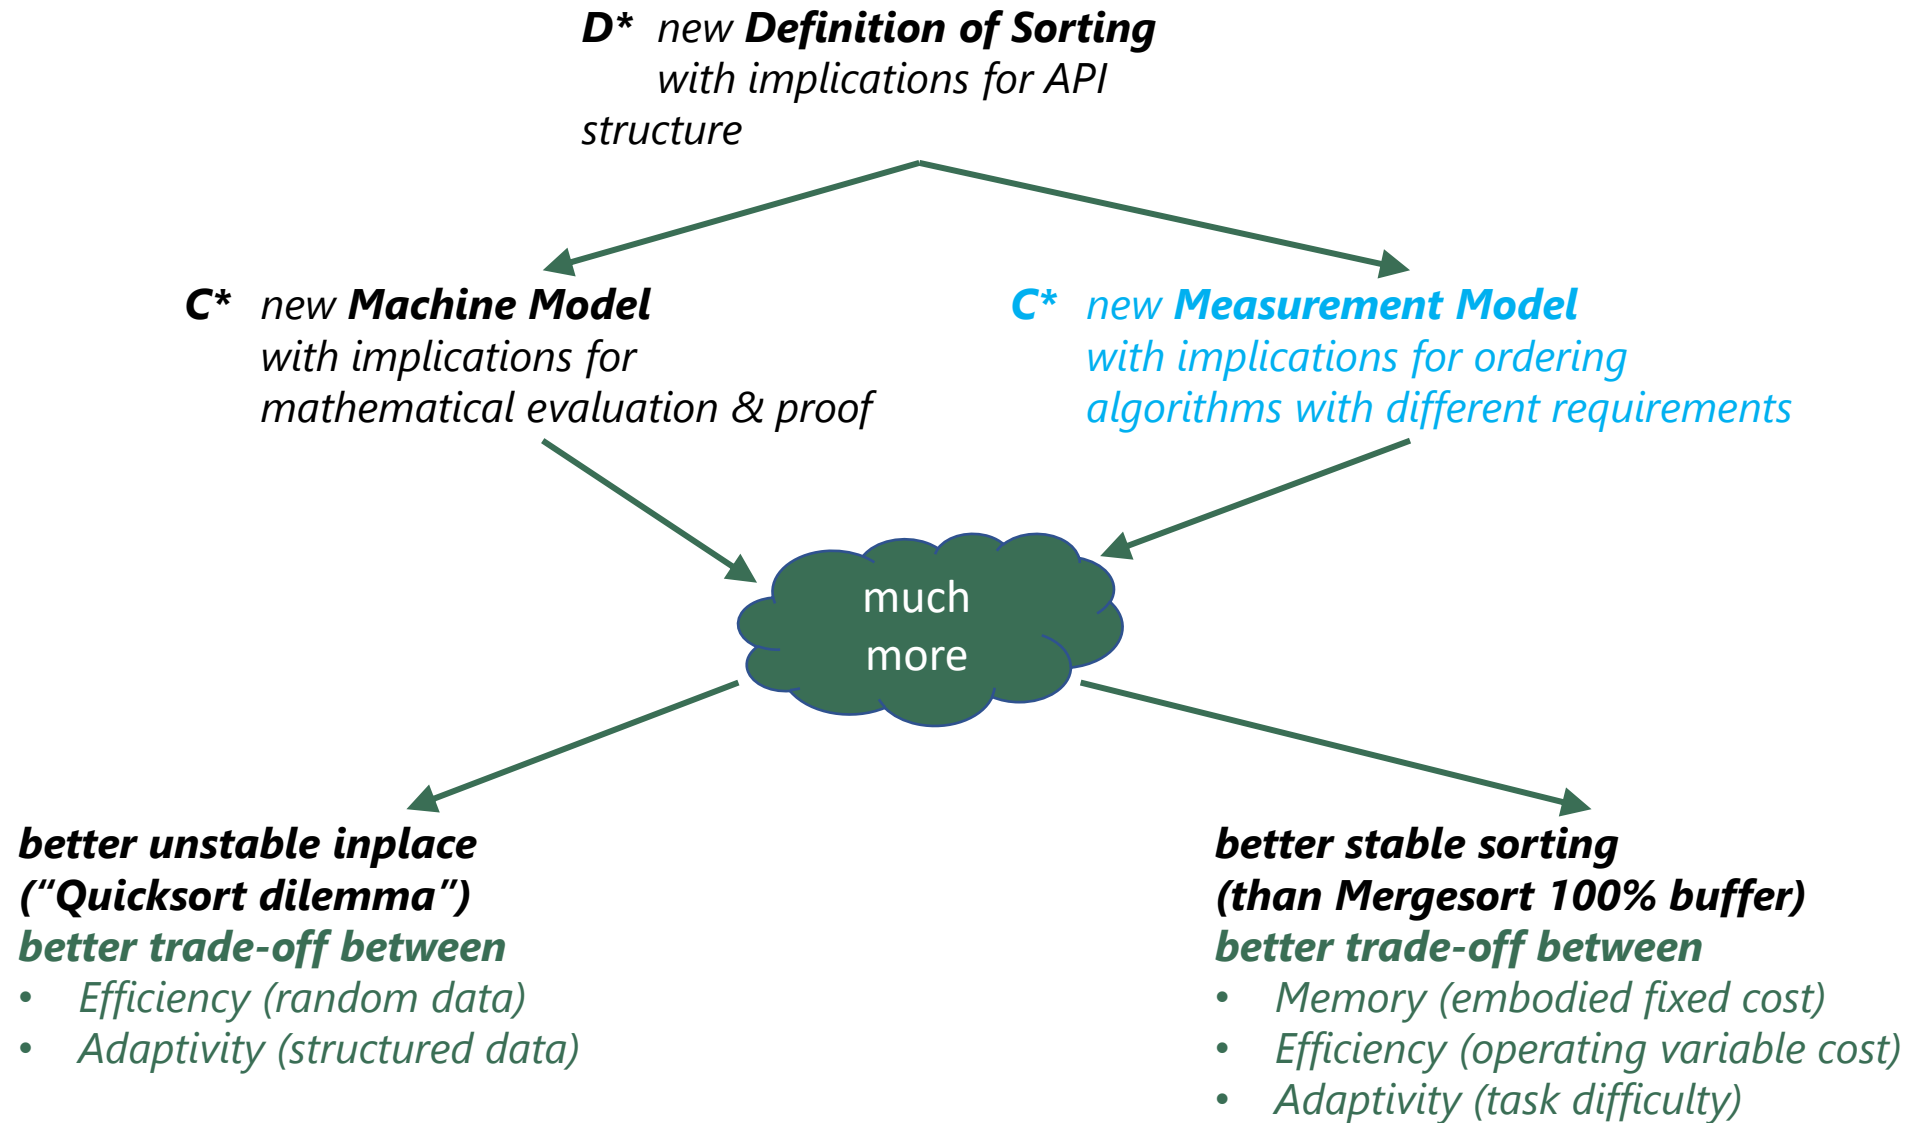
### **How? Better sorting algorithms in central software components**

-  Method: analyzing the prior-art and
-  designing and testing many innovative algorithms
-  such that they are sustainable

-  Result? A a thorny 12-year journey
-  identifying poisoned limiting assumptions
-  designing a new theory of sorting with
-  **new sustainability measurements**
-  new methods that are simple and beautiful
-  new algorithms with better trade-offs

# greeNsort® web of innovations – here heavily simplified

already public  
yet confidential



## *Disclaimer: my Perspective*

Manage future

I am not an accountant, I am not interested in reporting the past. At McKinsey I learned to take a **managerial perspective on decisions as being about different choices for the future.**

Scientific ethos

I am a consultant, data scientist and system architect and want sustainable results. Hence I **measure in controlled experiments, unbiased with minimal error variance and try to optimize with impact.**

Pragmatic action

I am not a theoretician, I am not interested in pseudo-precision. I am a practitioner, and in a partial developer role having worked with more than 12 languages I **need pragmatic solutions.**

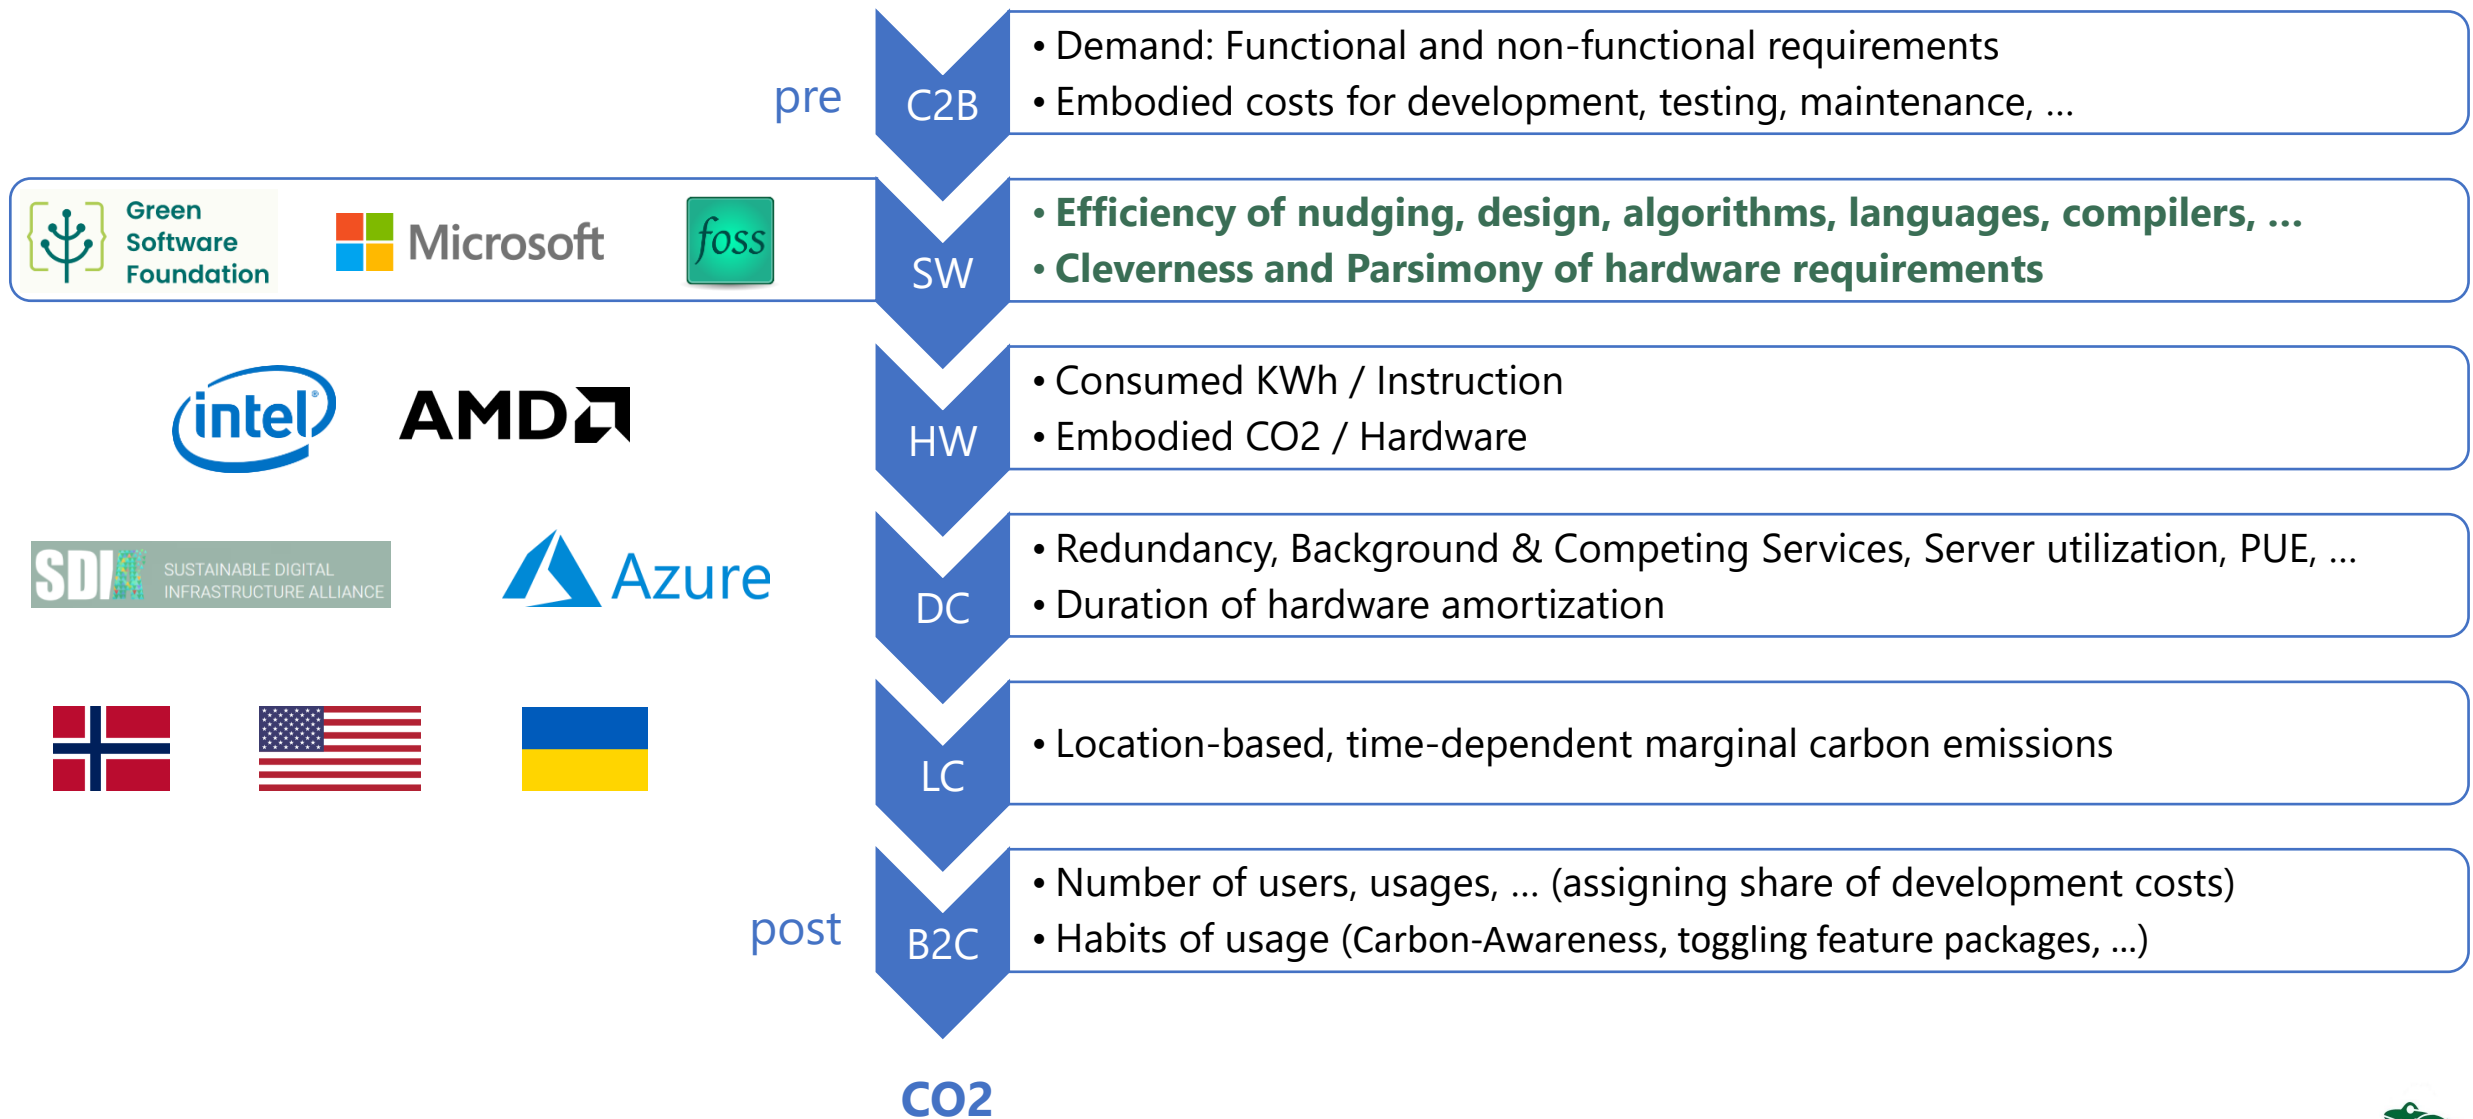
# We have basic agreement on cost measurement

$$\begin{array}{l} \text{assigning fixed costs} \\ M \end{array} = \begin{array}{l} \text{asset costs} \\ TE \end{array} \times \begin{array}{l} \text{share of asset} \\ RS \end{array} \times \begin{array}{l} \text{share of lifetime} \\ TS \end{array}$$

$$\begin{array}{l} \text{measuring variable costs} \\ \text{Energy } E \\ \text{CO2 } O \end{array} = \begin{array}{l} \text{average operational costs} \\ W \end{array} \times \begin{array}{l} \text{runtime} \\ h \end{array}$$
$$O = E \times I$$

$$\begin{array}{l} \text{basis for scaling} \\ SCI \end{array} = \begin{array}{l} \text{cost / task} \\ (O+M) / R \end{array}$$

# Measuring and Optimizing *independent* factors that affect CO2



# Measuring Software Factors that affect CO2 is hard

today out of scope

C2B

- harder to know and harder to influence

**software effect**

SW

- influence of Software Engineering

error variance

HW

- impossible to know and not a feature of software, but ...

error variance

DC

- impossible to know and not a feature of software, but ...

error variance

LC

- impossible to know and not a feature of software, but ...

today out of scope

B2C

- harder to know and harder to influence

strive for  
HW DC LC  
obliviousness

**end-to-end CO2 impossible to know or complicated to measure**



# Do you want to be powered-by renewables? The answer is no.\*



Asim Hussain

## Why?

- Until less than 100% renewables: someone else uses fossile energy, to avoid **greenwashing**, we need to operate with  $I = \text{maximum marginal CO}_2/\text{KWh}$
- Once 100% renewables:  $I = \text{zero CO}_2/\text{KWh}$ , hence an invitation to waste energy

## What follows?

Variable costs: **reducing CO<sub>2</sub>** is not a proper goal for sustainable software

- reducing energy is!


Fixed costs: reducing embodied CO<sub>2</sub>

- by reducing hardware requirements

## Complication

- by choosing architecture with less embodied carbon
- by choosing architecture with less operational energy

} Not independent of HW!  
Which common scale?



\* <https://devblogs.microsoft.com/sustainable-software/do-you-want-to-be-powered-by-renewables-the-answer-is-no/>

# Strategy for software factors that affect CO2

today out of scope

WE

- (try to) minimize requirements
- (try to) **standardize** and **benchmark**

**Task definition**  
/R

**software effect**

SW

- **optimize efficiency**
- **minimize hardware requirements**

**measure**  
**optimize**

zero error variance

HW

- control and standardize KWh / Instruction or Cycle
- control and standardize CO2 / HW

zero error variance

DC

- control and standardize DC inefficiency
- control and standardize duration of hardware amortization

zero error variance

LC

- control and standardize CO2 / KWh as maximum marginal cost

aggregated  
**Standardization**  
factors or  
functions

today out of scope

WE

- (try to) minimize users and usage
- (try to) influence usage patterns

**SW rollout**  
/R

**end-to-end CO2**

# Strategy for Sustainable Algorithms

**CPU-parsimony**  
is key to saving electricity  
(reducing variable cost CO2)

**RAM-parsimony**  
is key to using old hardware longer  
(reducing fixed cost CO2)

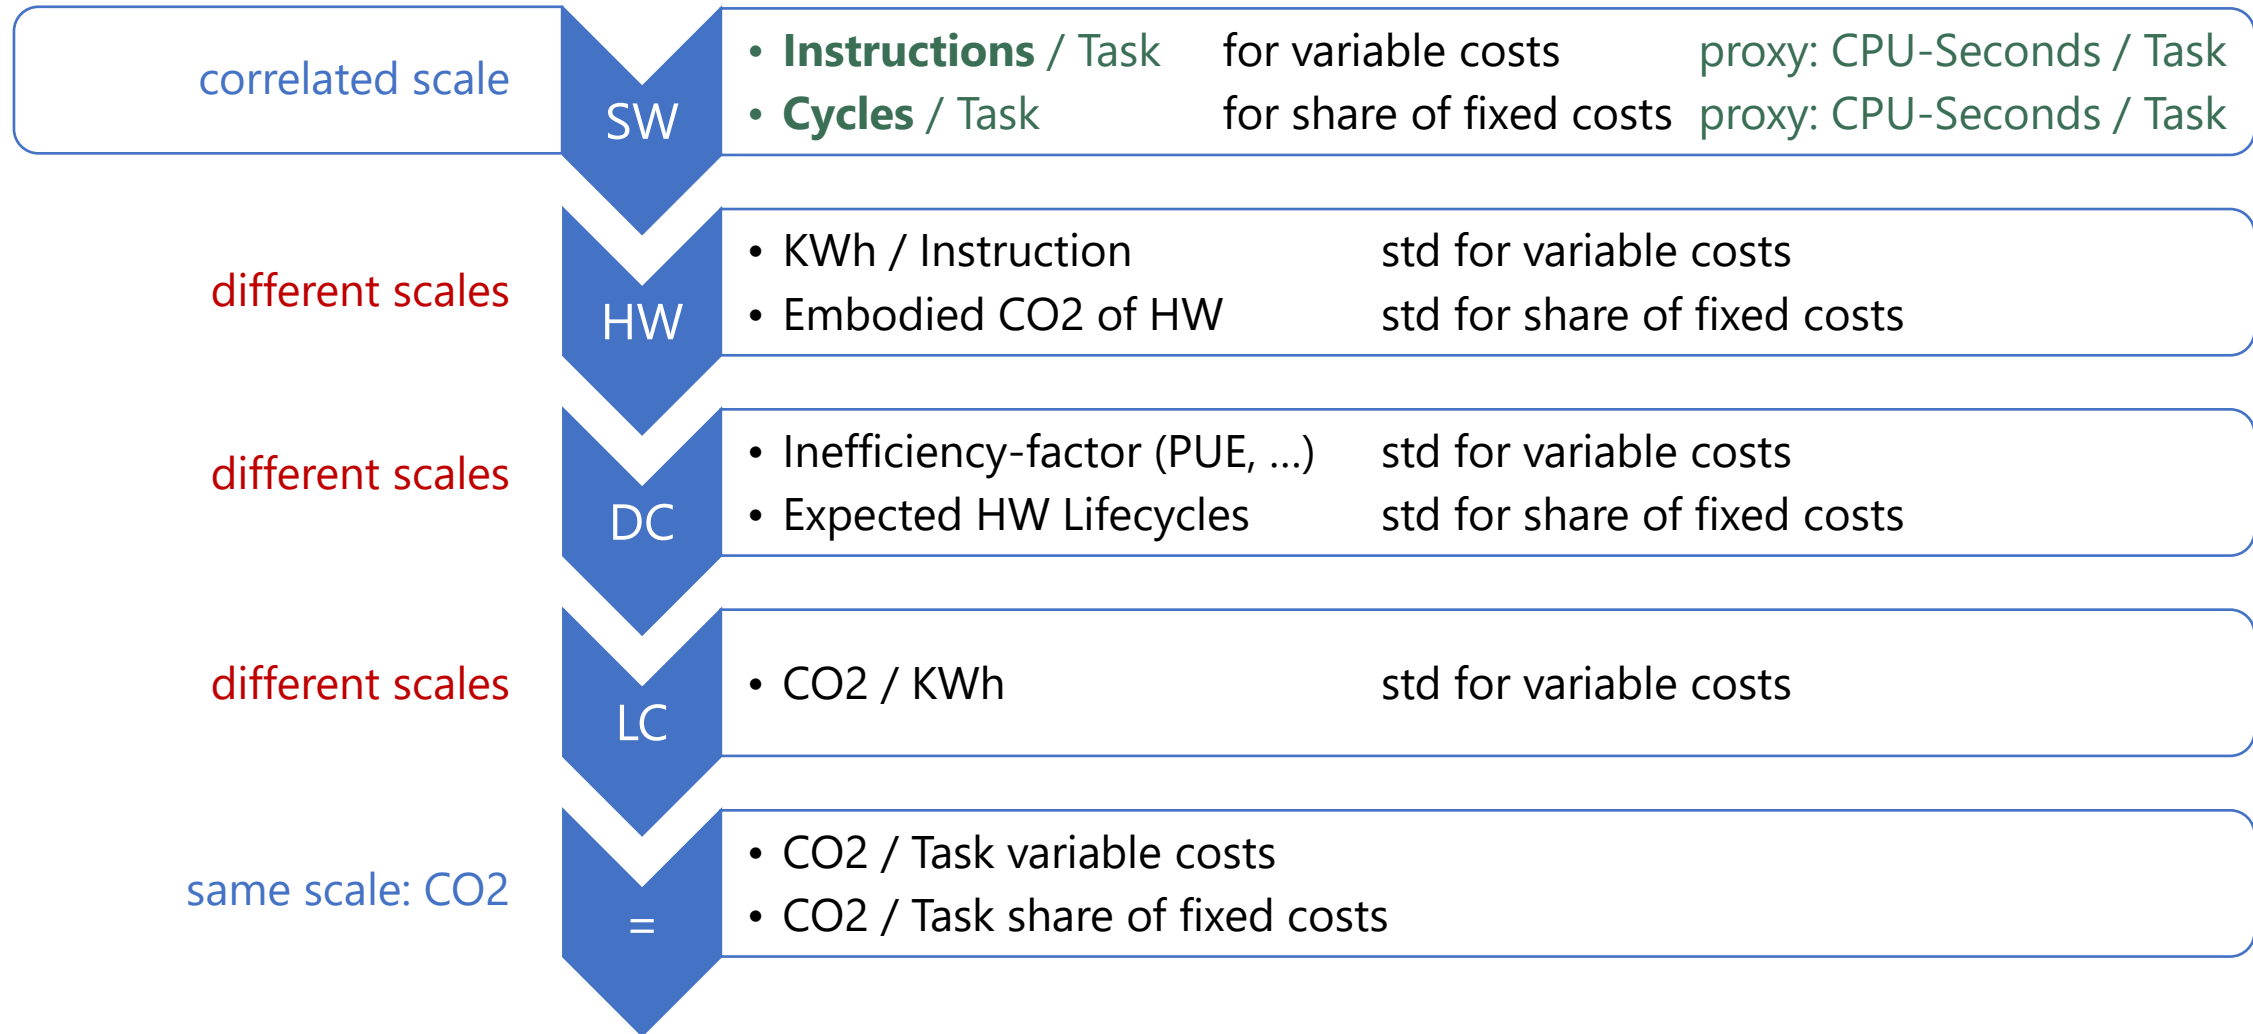


**Optimize for both, RAM and CPU,**  
**try to avoid extremes:**

- minimum buffer RAM (wasting CPU)
- minimum CPU cycles (wasting RAM)

# Break-down for CPU-cost

if serial CPU-bound:  
just Seconds / Task



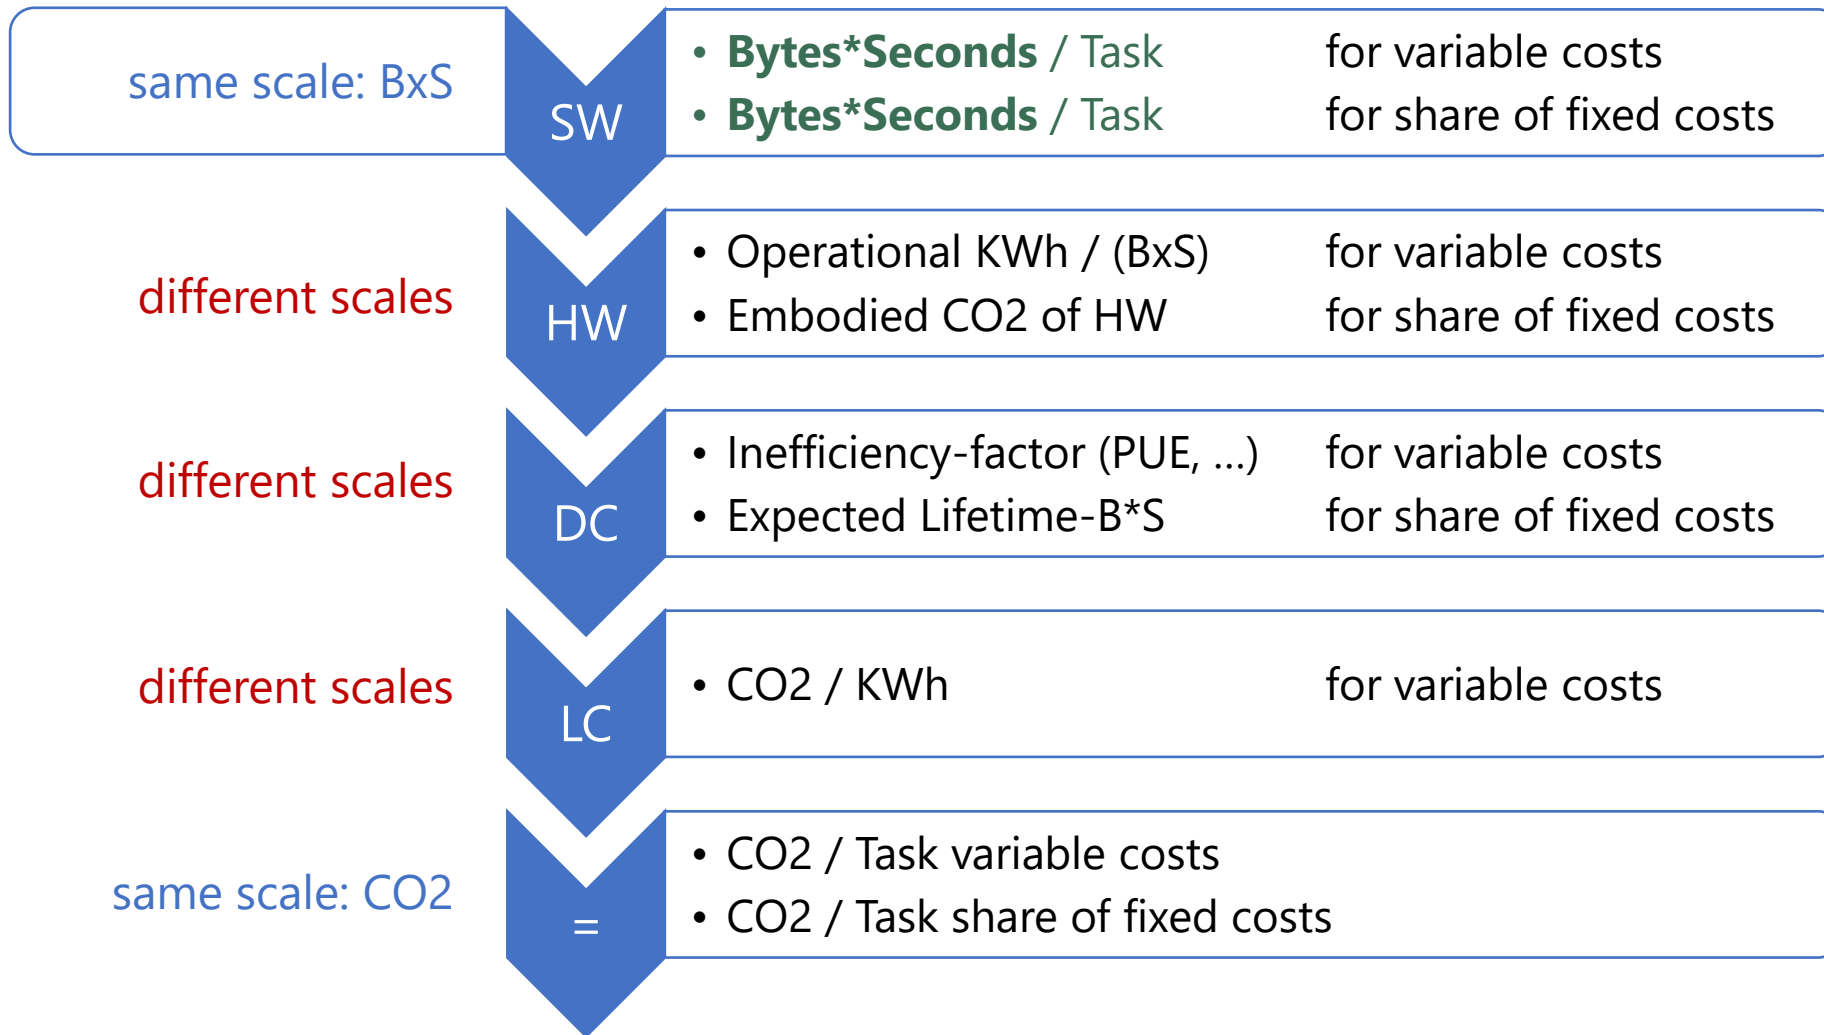
$$= \text{CPU-Seconds} * \text{stdCO2/Instruction(CPU)} + \text{CPU-Seconds} / \text{stdLifetime(CPU)} * \text{stdCO2Embodied(CPU)}$$

$$= \text{CPU-Seconds} * ( \text{stdCO2/Instruction(CPU)} + \text{stdCO2Embodied (CPU)} / \text{stdLifetime(CPU)} )$$

$$= \text{CPU-Seconds} * \text{stdCO2(CPU)} = \text{CPU-Seconds} * \text{const} = \text{CPU-Seconds are a good proxy for benchmarking ratios of CPU!}$$



# Break-down for RAM-cost is different



$$= \text{BxS} * \text{stdCO2perBxS(RAM)} + \text{BxS} * (\text{stdCO2EmbodiedperB(RAM)} / \text{stdLifeseconds(RAM)})$$

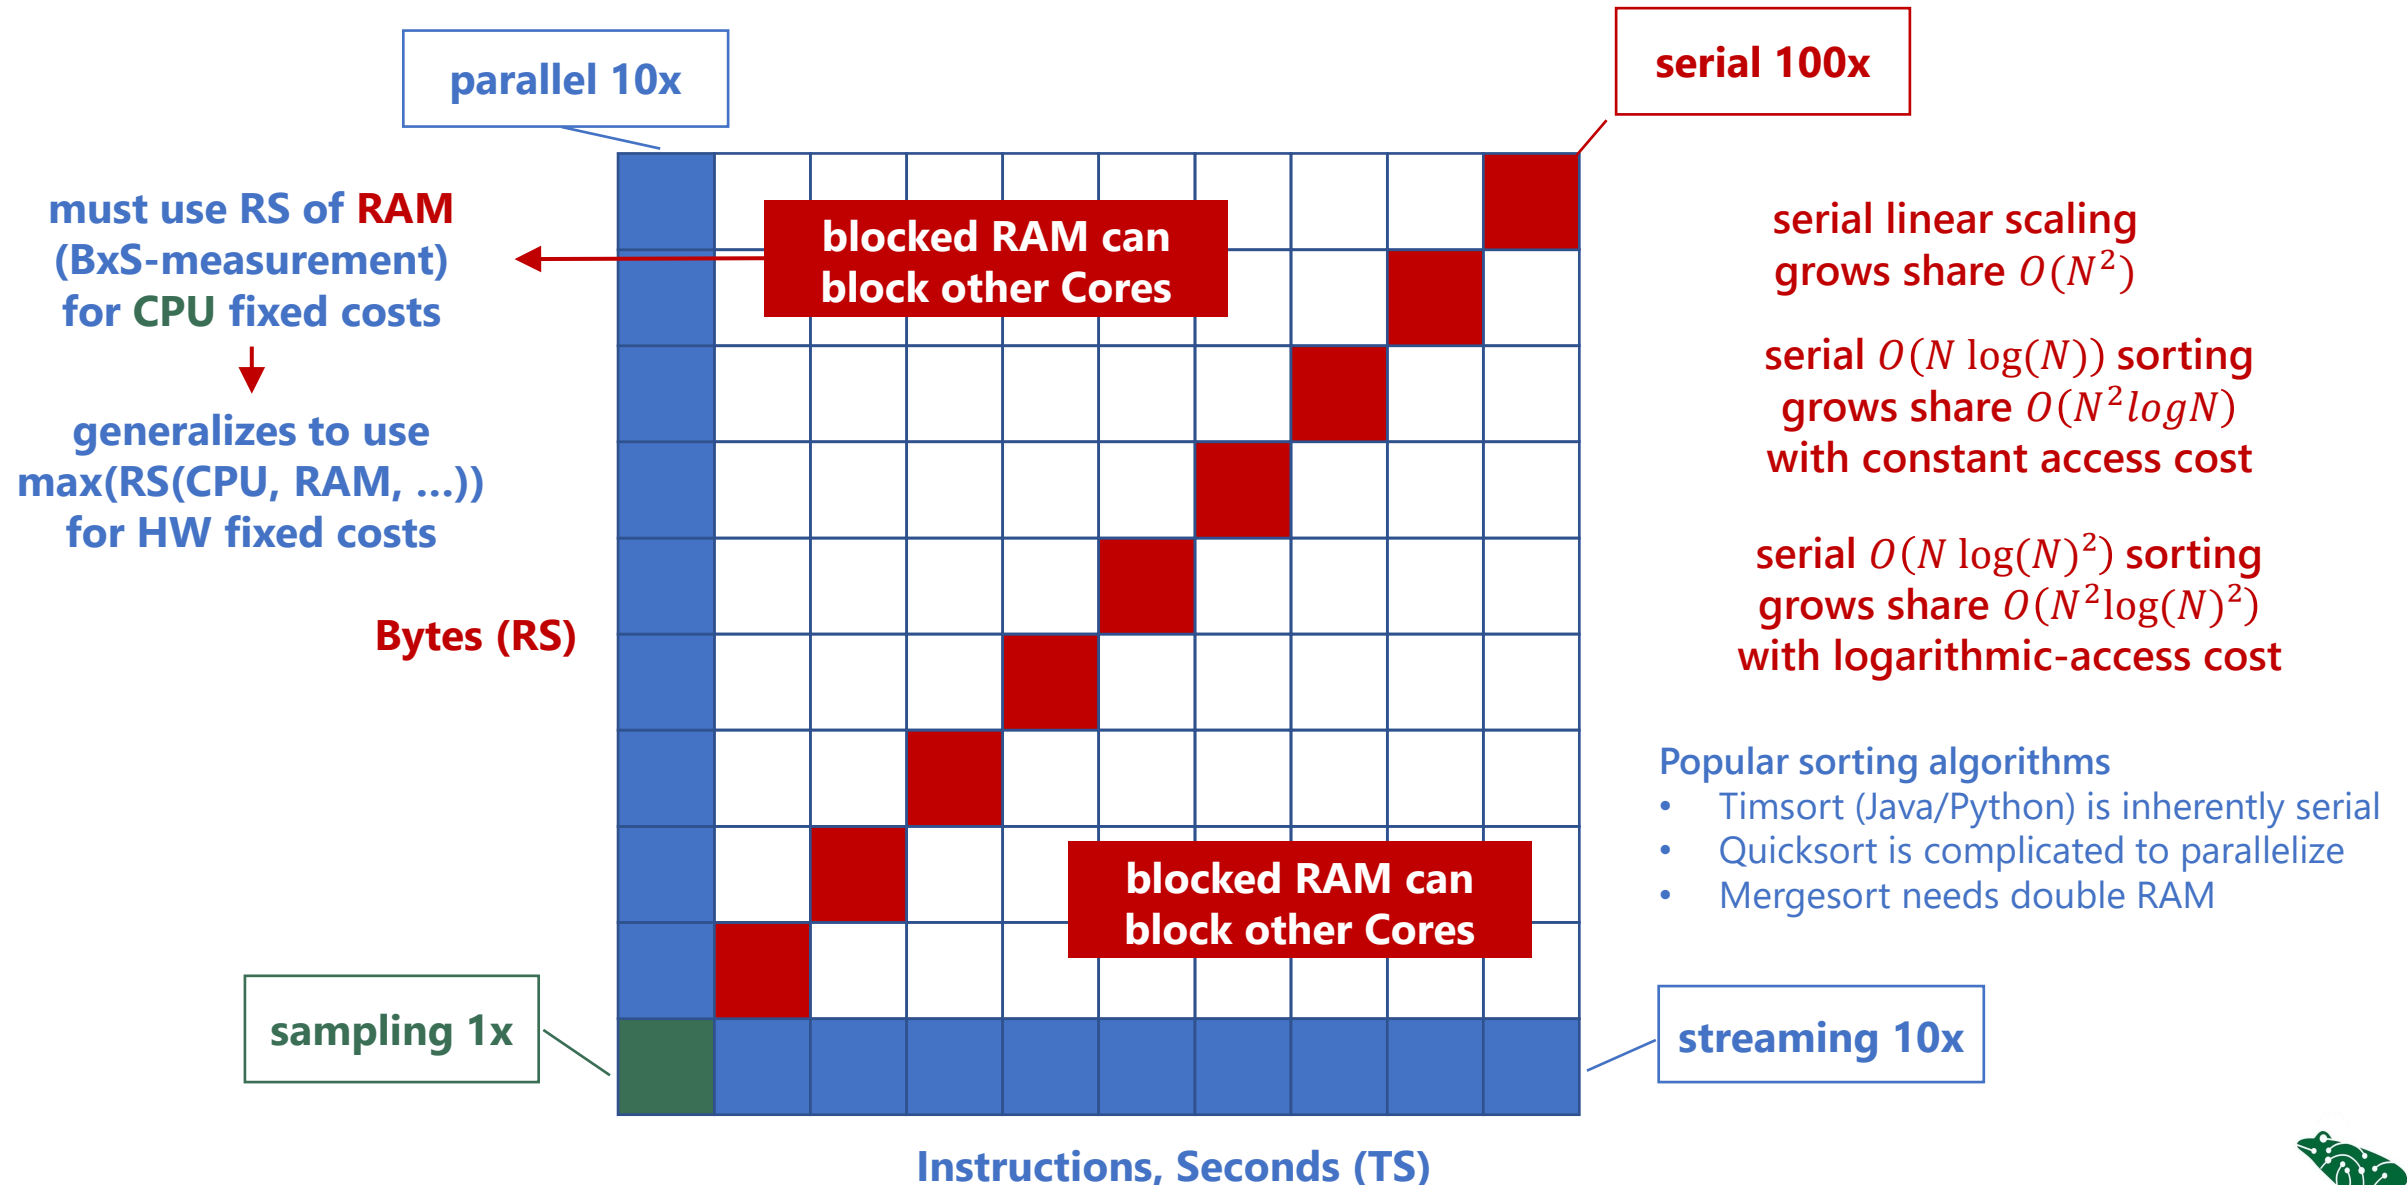
$$= \text{BxS} * (\text{stdCO2perBxS(RAM)} + (\text{stdCO2EmbodiedperB(RAM)} / \text{stdLifeseconds(RAM)}))$$

$$= \text{BxS} * \text{stdCO2(RAM)} = \text{BxS} * \text{const} = \mathbf{\text{BxS is a good measurement for benchmarking ratios of RAM!}}$$

? BxS = N<sup>2</sup> ?



# Multiplying Bytes x Seconds – Scaling 10x bigger data



# Optimize for what? CPU? RAM? Variable? Fixed? End-to-End CO2? There is no perfect solution – but we need a single scale!

	CPU	RAM	
<b>Fixed costs</b>	$\text{Max}(\text{Cycles share}, \text{Bytes} * \text{Seconds share})$	$\text{Bytes} * \text{Seconds share}$	→ correlated or identical
<b>Variable costs</b>	Cycles	$\text{Bytes} * \text{Seconds}$	→ correlated



correlated or identical



identical

How to penalize apps  
with unpredictable CPU usage?

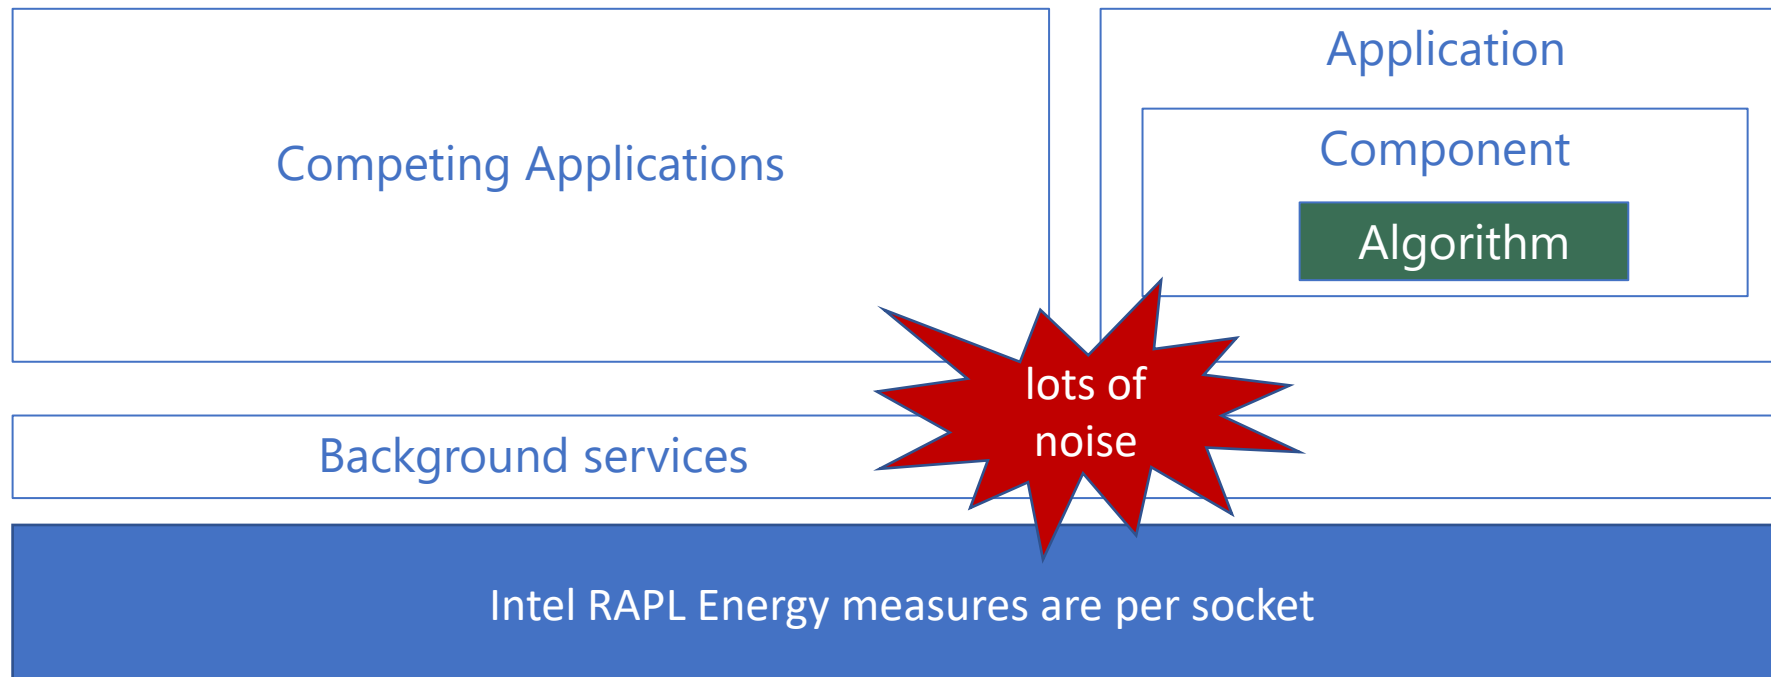
How to penalize apps with  
unpredictable RAM usage?

How to penalize apps  
with random access?

# Measuring and attributing Energy is difficult. End-to-end CO2 even more. Needed are simple pragmatic measurements!

**PAPI** instrumentation is complicated and instrumentation might influence performance


**perf** is simpler but measures only complete applications





# Even external measurement with perf is complicated and expensive

- 1 measure algo duration
  - 2 calculate enclosing window
  - 3 baseline window sleep
  - 4 spawn perf for window
  - 5 wait until perf is likely up
  - 6 measure algorithm
  - 7 refine window
  - 8 repeat from 3
- 
- B estimate baseline
  - O mark out-of-window
  - I subtract idle-baseline
  - C core-share of baseline
  - A statistical analysis



multi-seconds for  
each run to be  
measurable

$$O(R \cdot A \cdot D \cdot N \cdot \log N)$$

where

*R = 50 replications*

*A = 100 algorithms*

*D = 30 data patterns*

*N = 2<sup>24</sup> data size*

~ weeks

*C = #compiler versions*

*H = #HW configurations*

~ crazy

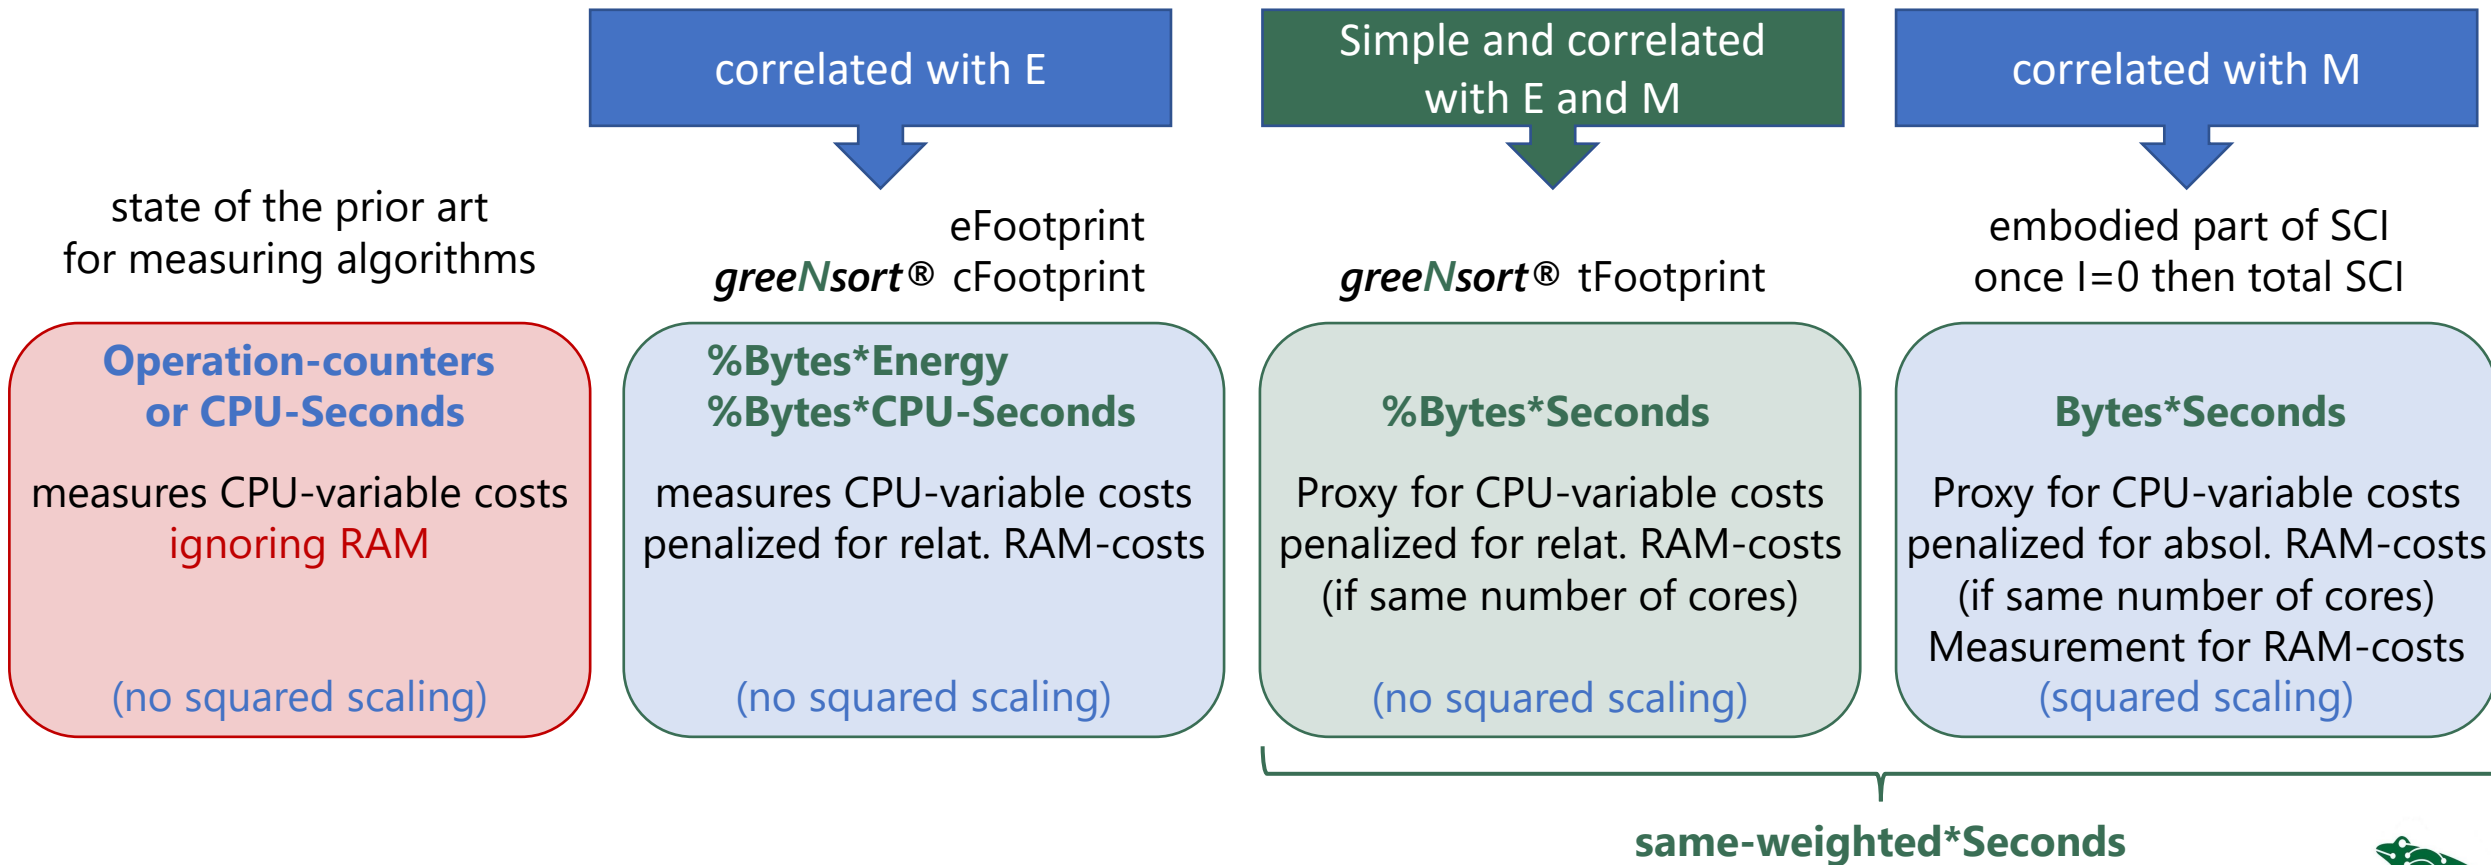


# Carbon proxies: Measuring the greenness of your application\*

*"If you can somehow connect what you are measuring to carbon, then it's a proxy.  
And optimizing for that metric is optimizing for carbon."*



Asim Hussain



\* <https://devblogs.microsoft.com/sustainable-software/carbon-proxies-measuring-the-greenness-of-your-application/>



# Example for the power of the Footprint approach

Rui Pereira, Marco Couto, Francisco Ribeiro, Rui Rua, Jácome Cunha, João Paulo Fernandes, and João Saraiva  
"Energy Efficiency across Programming Languages: How does Energy, Time and Memory Relate?"

## Not all Programming Languages contribute equally to the Climate-Crisis



Now consider  
all the people  
doing  
*Big<sup>2</sup> Data*  
*Auto<sup>3</sup> ML*  
with  
*200 · Python !*

Develop with efficient Algorithms, Libraries and Languages and reduce  
Users, Devices, Applications, Features, Data, AI, VR, Traffic and Usage

possible cooperation  
on draft paper about  
12 software factors

# Result of *greeNsort*<sup>®</sup>

popular Algorithms	RAM	Parallel?
Quicksort (everywhere)	100% (unstable)	complicated (practically big tasks remain serial)
Timsort (Java & Python)	150%	no (extremely adaptive but even serially inefficient)
Mergesort (everywhere)	200%	possible (little adaptive)

<i>greeNsort</i> Algorithms	RAM	Advantage
Ducksort	100% (unstable)	more efficient and adaptive, very simple
Frogsort	105% .. 150%	quite .. very parallel, adaptive, very simple
Squidsort	105% .. 150%	quite .. very parallel, very adaptive, simpler
Octosort	200%	fully parallel, extremely adaptive, very simple

# A final thought: Carbon-Awareness – Solution or Illusion?

indeed a  
feature of software

good old batch



can delay a lot

can collect for macro-batch

can shift to optimal time,  
e.g. solar daytime

neartime



can delay a bit

can collect for micro-batch

cannot shift to optimal time

realtime



cannot delay at all

cannot collect for micro-batch

cannot shift to optimal time

but: if not using hardware during nighttime  
this halves lifetime and doubles fixed costs

but: often batch is needed in the night

## Sorry for leaving little time for final discussion – I still give a conclusion from my side

The most accepted common scale is absolute CO2. For evaluation of a minimal (in-memory) software we need four measurements for fixed cost and variable cost of CPU and RAM, if necessary more:

MEASURE	CPU	RAM	GPU, SSD, ...
<b>Fixed costs</b>	CxS Cores*Seconds	BxS Bytes*Seconds	
<b>Variable costs</b>	Wh Energy	BxS Bytes*Seconds	

Then we need for each of the four measurements a GSF standardized default conversion-factor to CO2:

CO2-FACTOR	CPU (core)	RAM (byte)	GPU, SSD, ...
<b>Fixed costs</b>	gCO2/CxS	gCO2/BxS (for fixed)	
<b>Variable costs</b>	gCO2/Wh	gCO2/BxS (for variable)	

Since the Energy-measurement is complicated and expensive, software-engineers might benefit from quick and easy measurement of **proxies** which replace above measurements, namely  $\%Bytes*(CPU)Seconds$ , as a validation-bridge  $\%Bytes*Wh$  and for final validation the full CO2 calculation above.

Any helper-tools for the above measurements would be highly welcome. They should be easy to use, scale to large code-bases and many users and ideally be applicable for confidential code. Thank you for the invitation.