ETH zürich

Reflecting on the Goal and Baseline of Exascale Computing

Thomas C. Schulthess

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

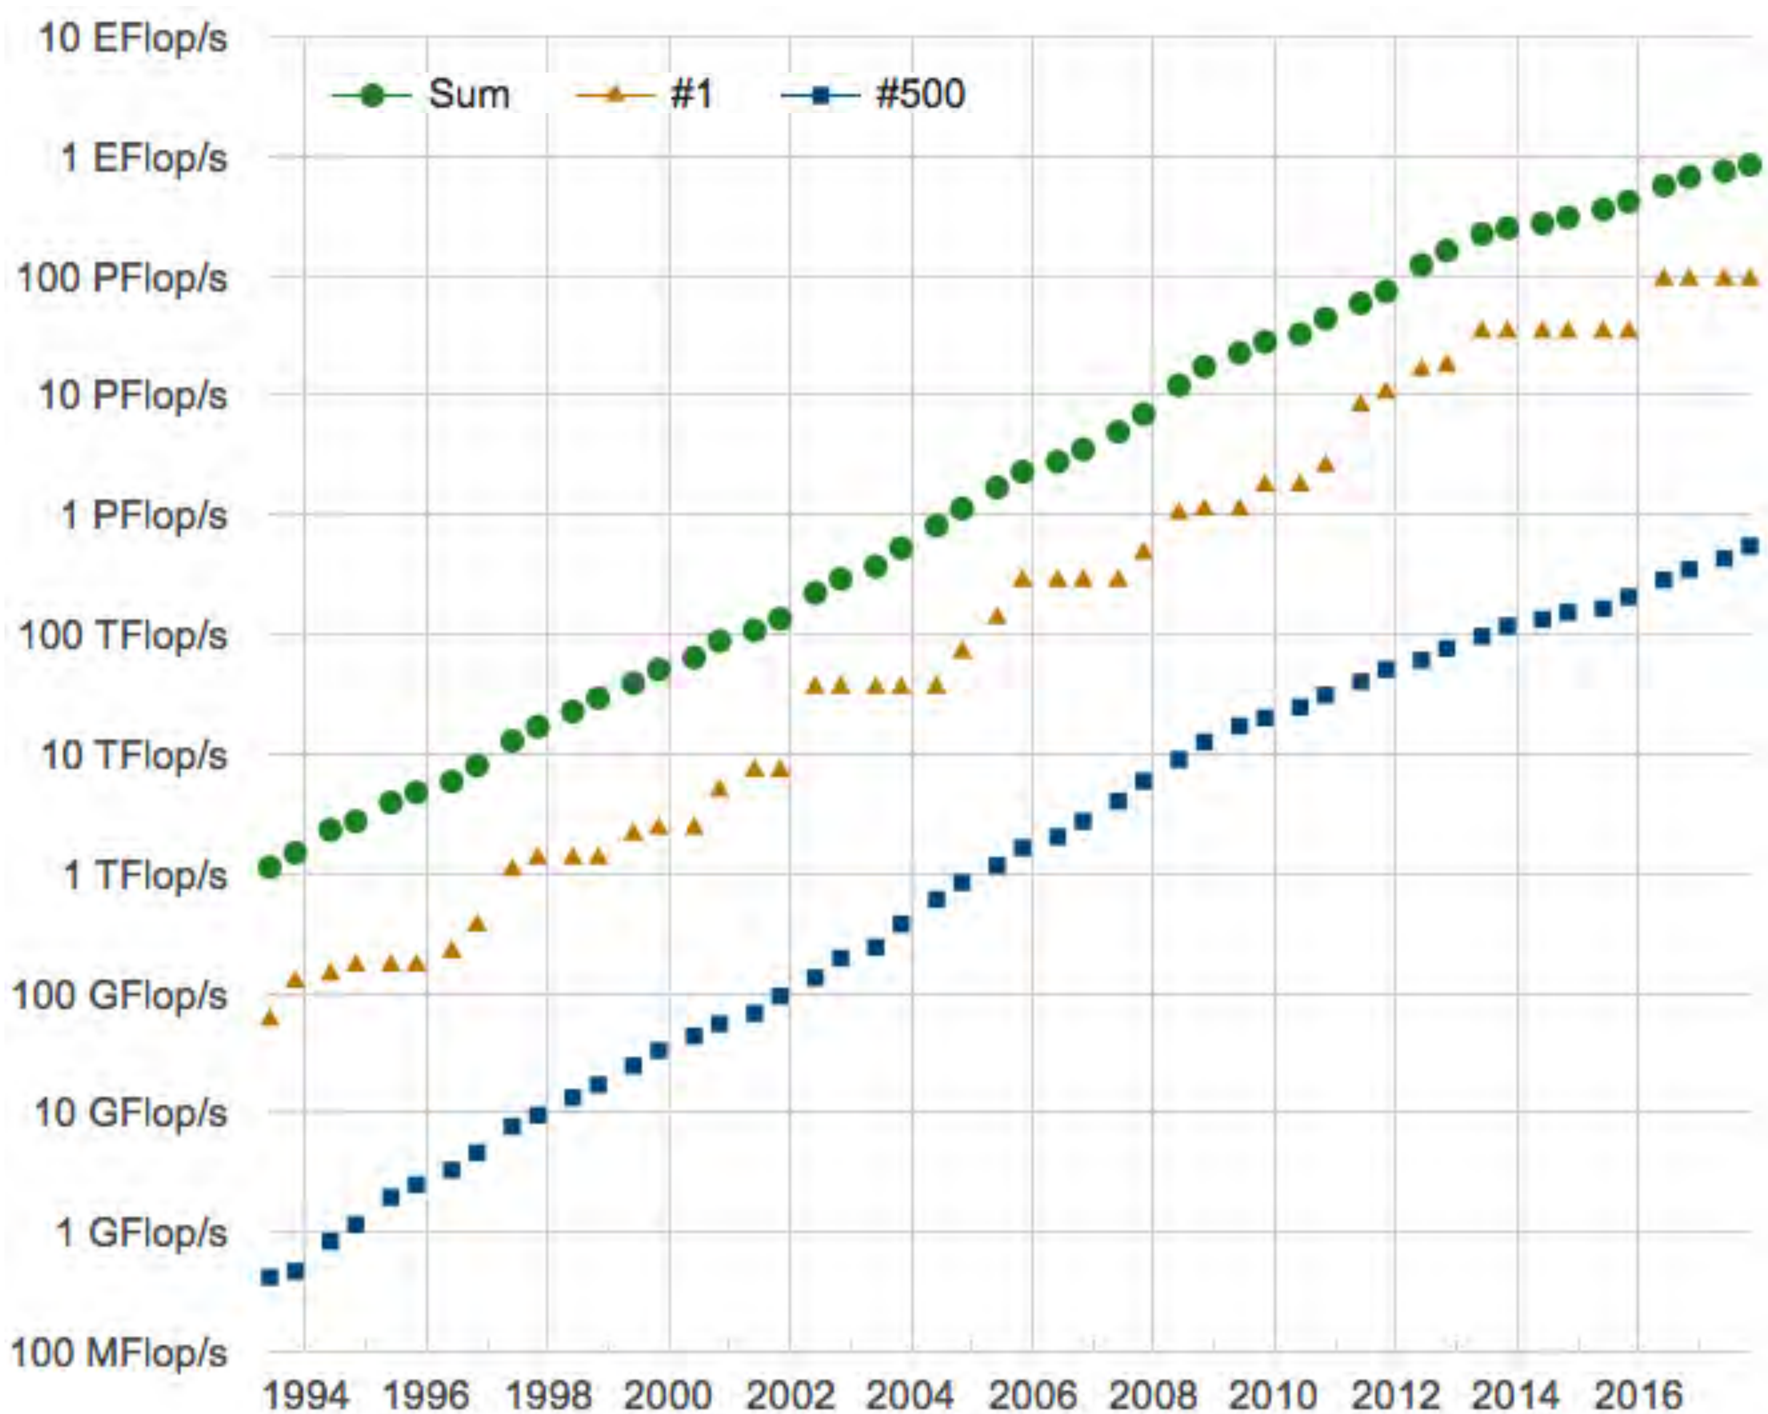

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

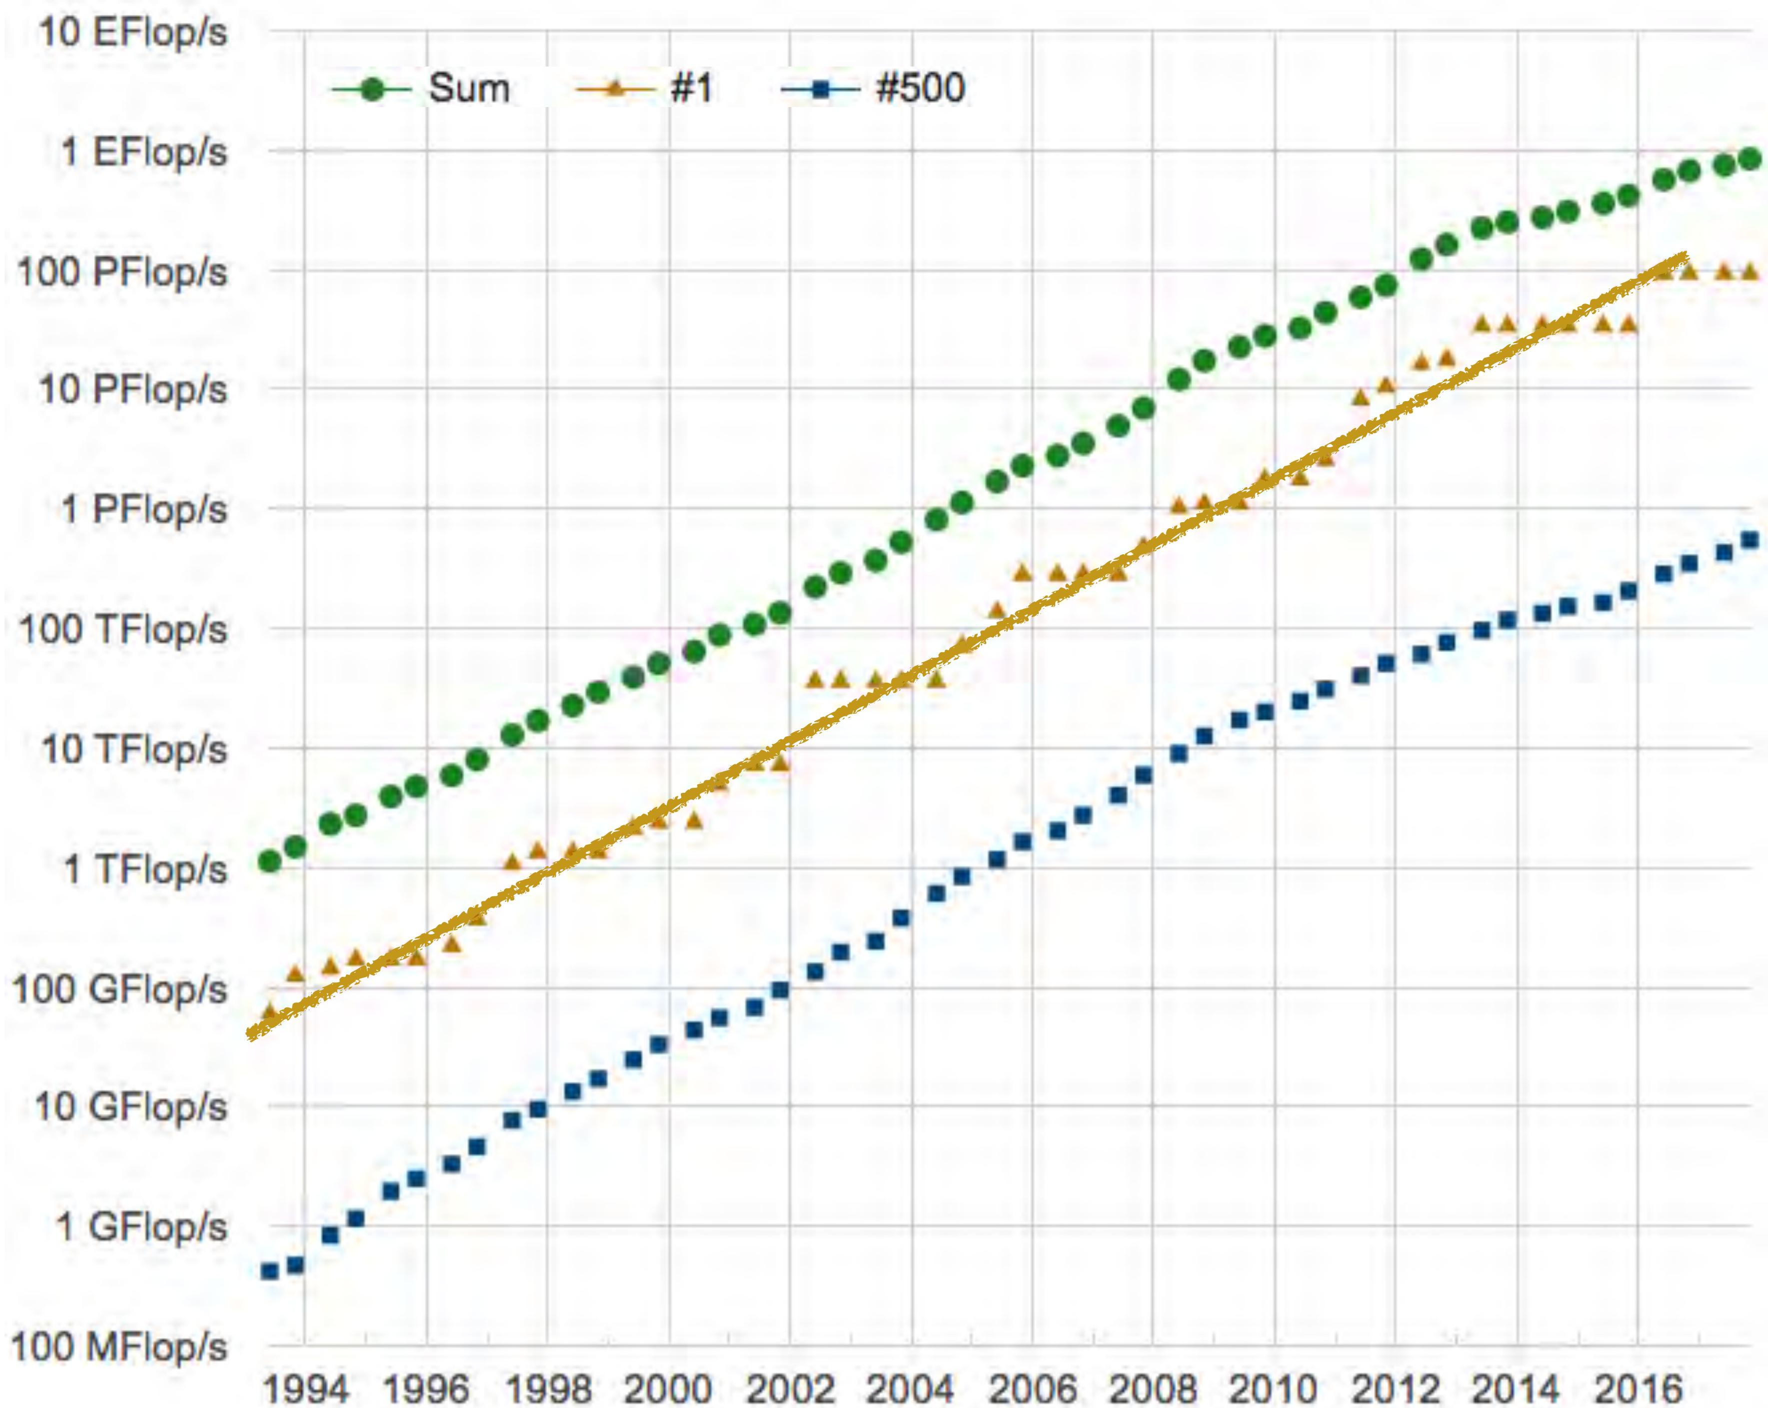

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

ETH zürich

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

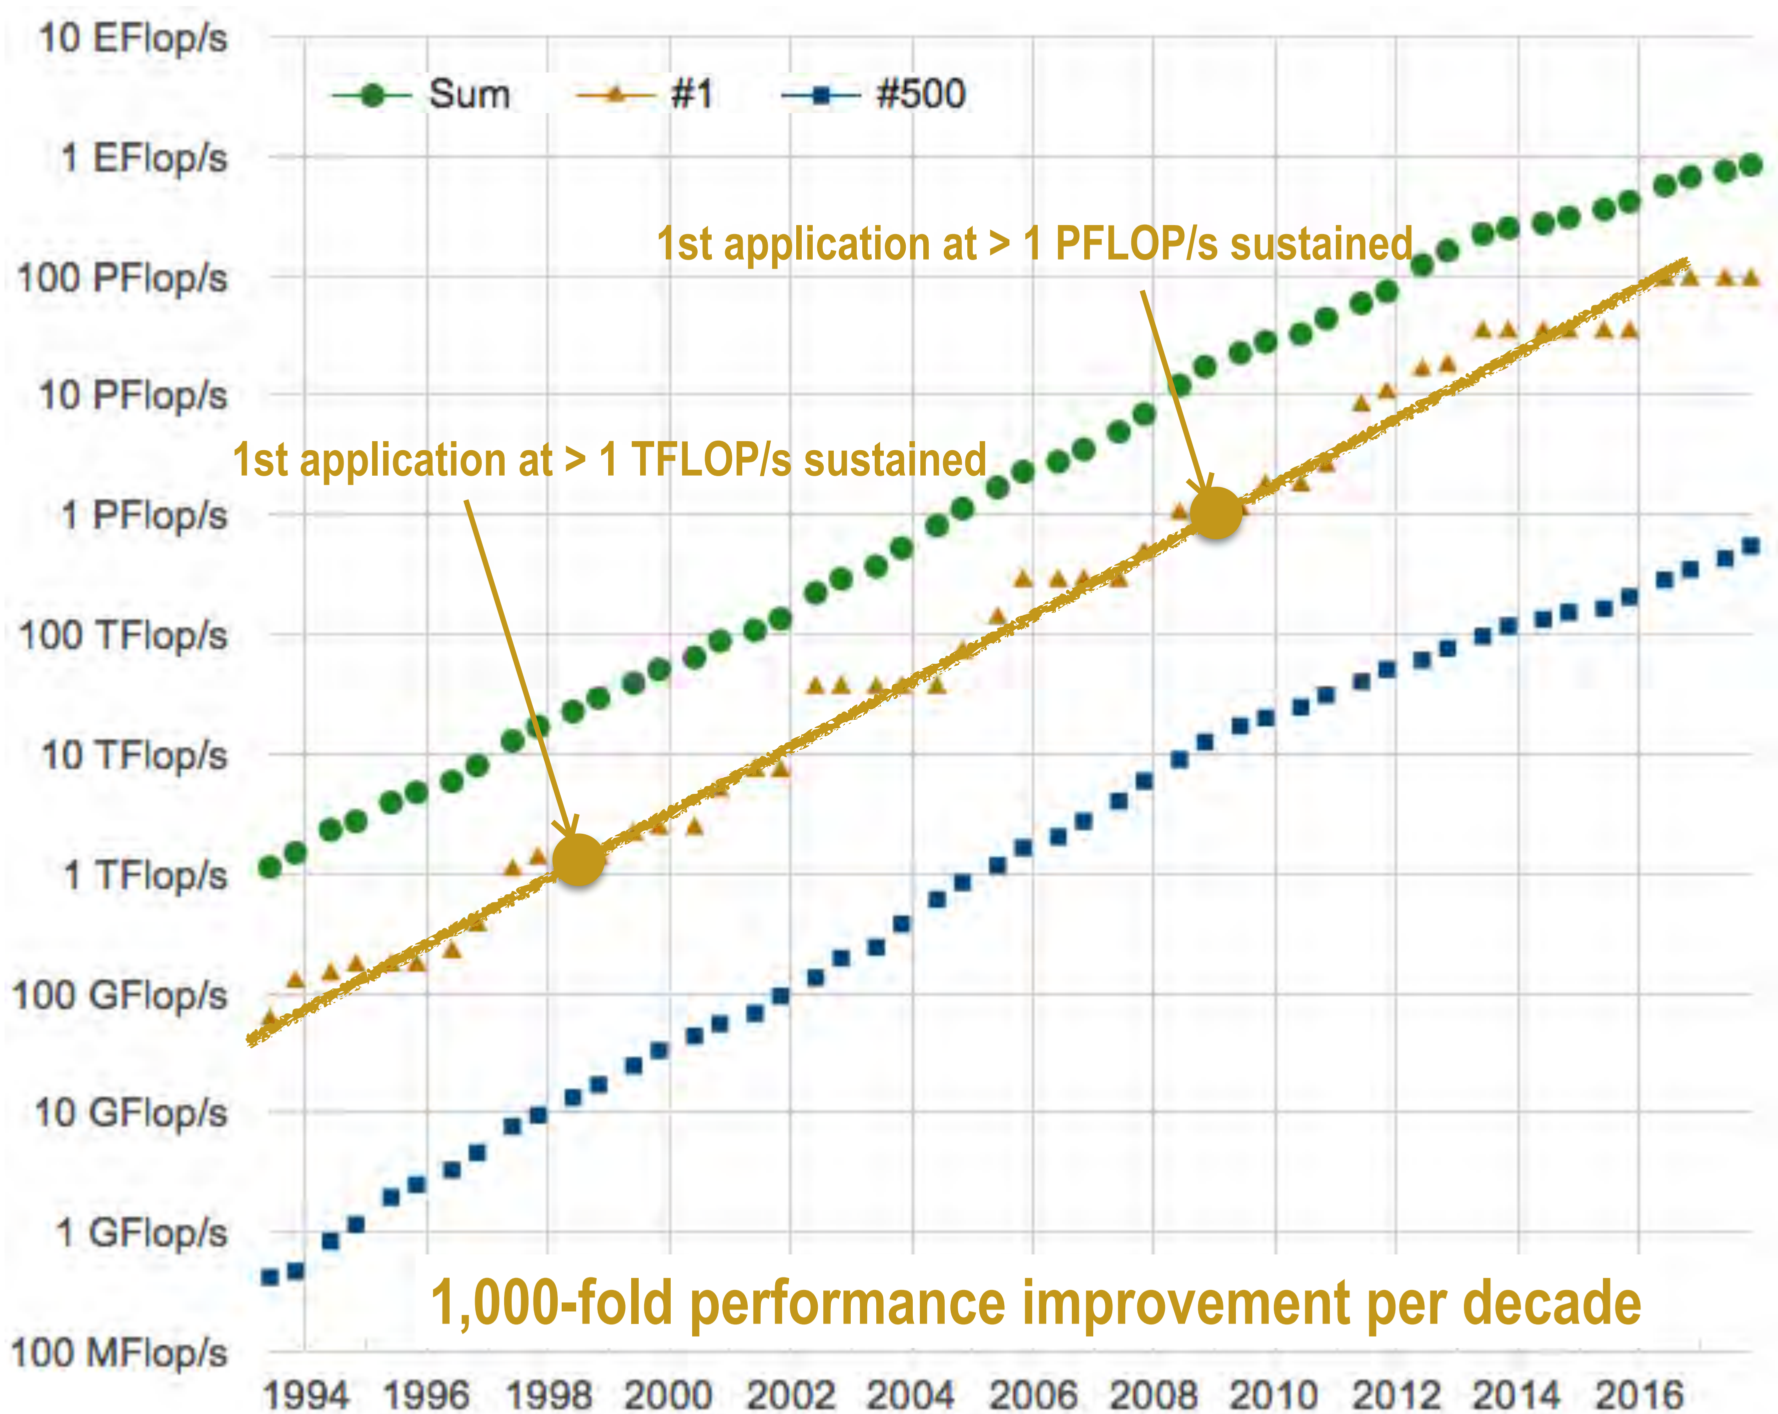

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

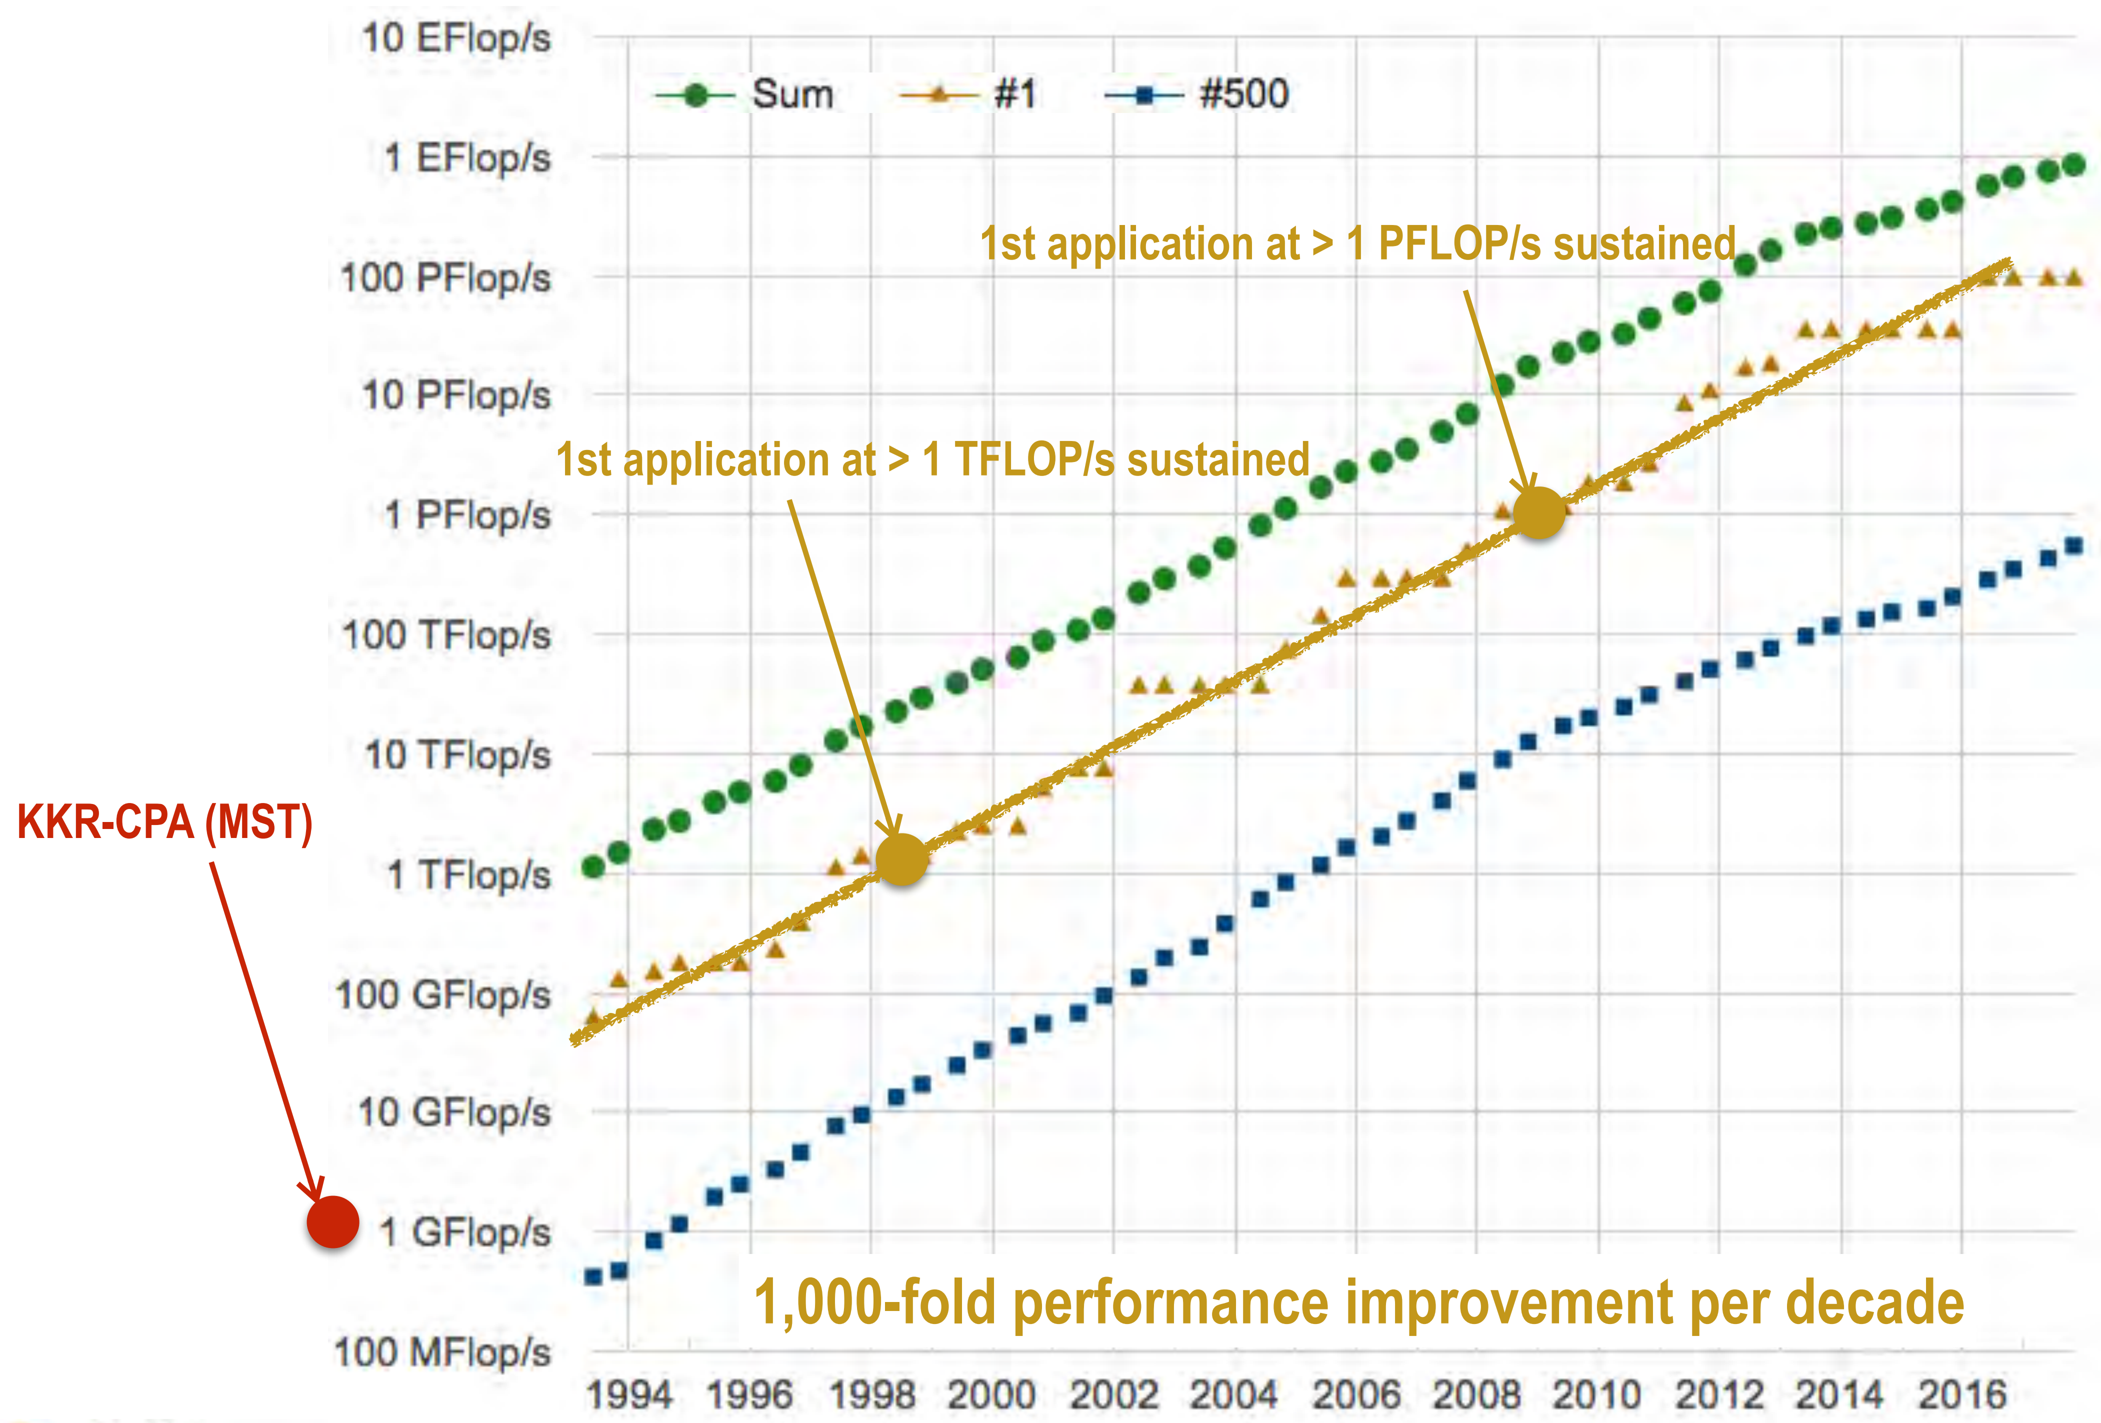

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

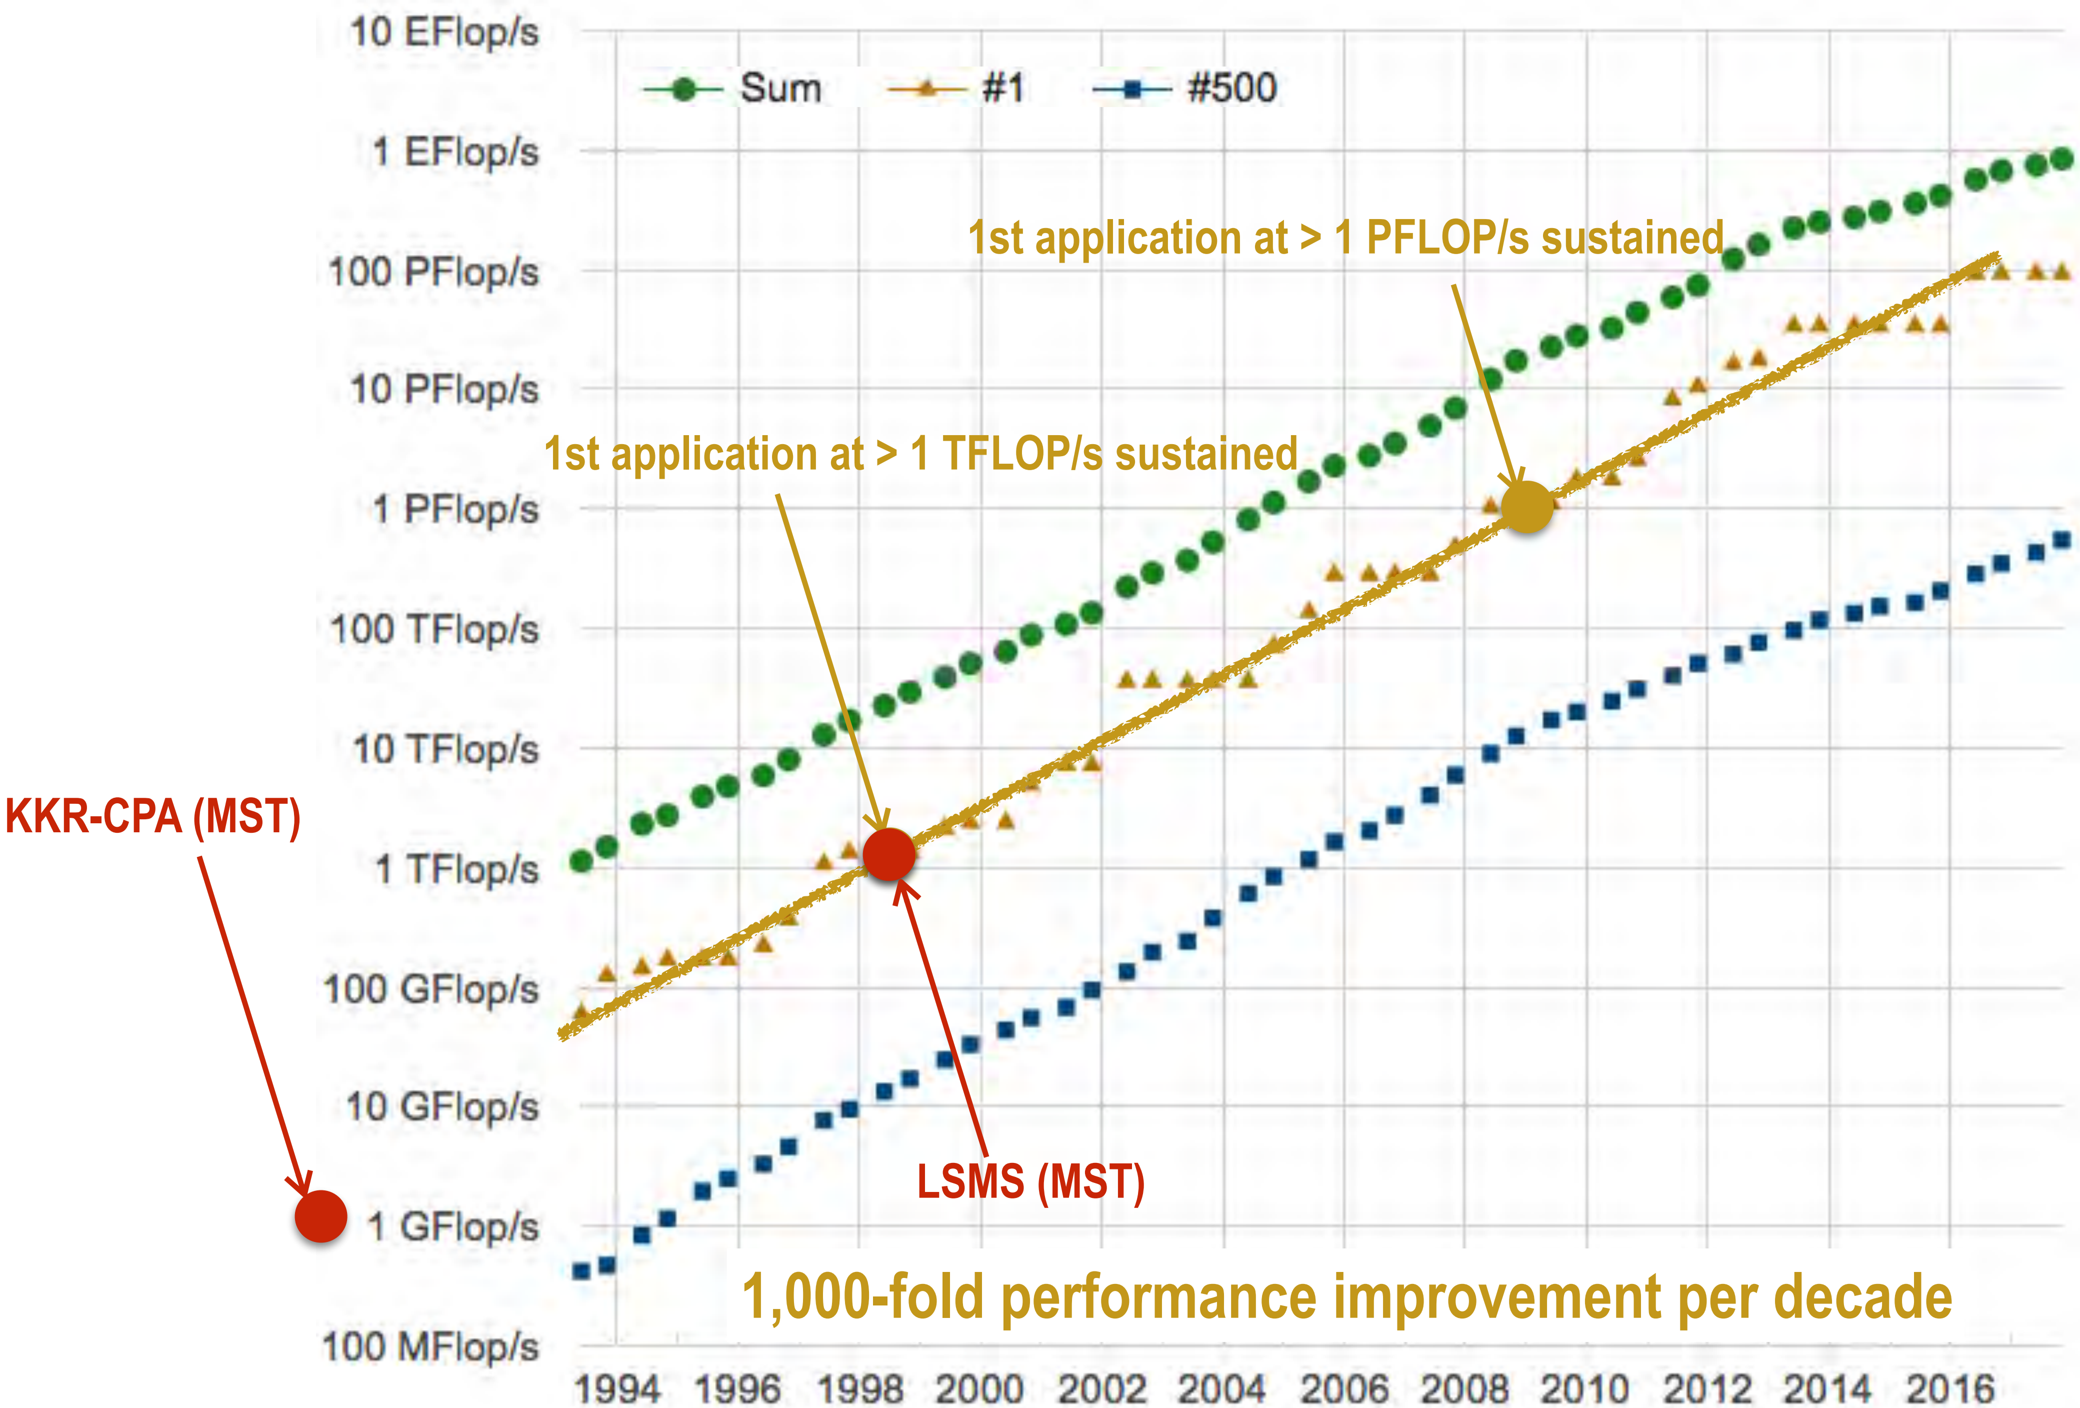

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

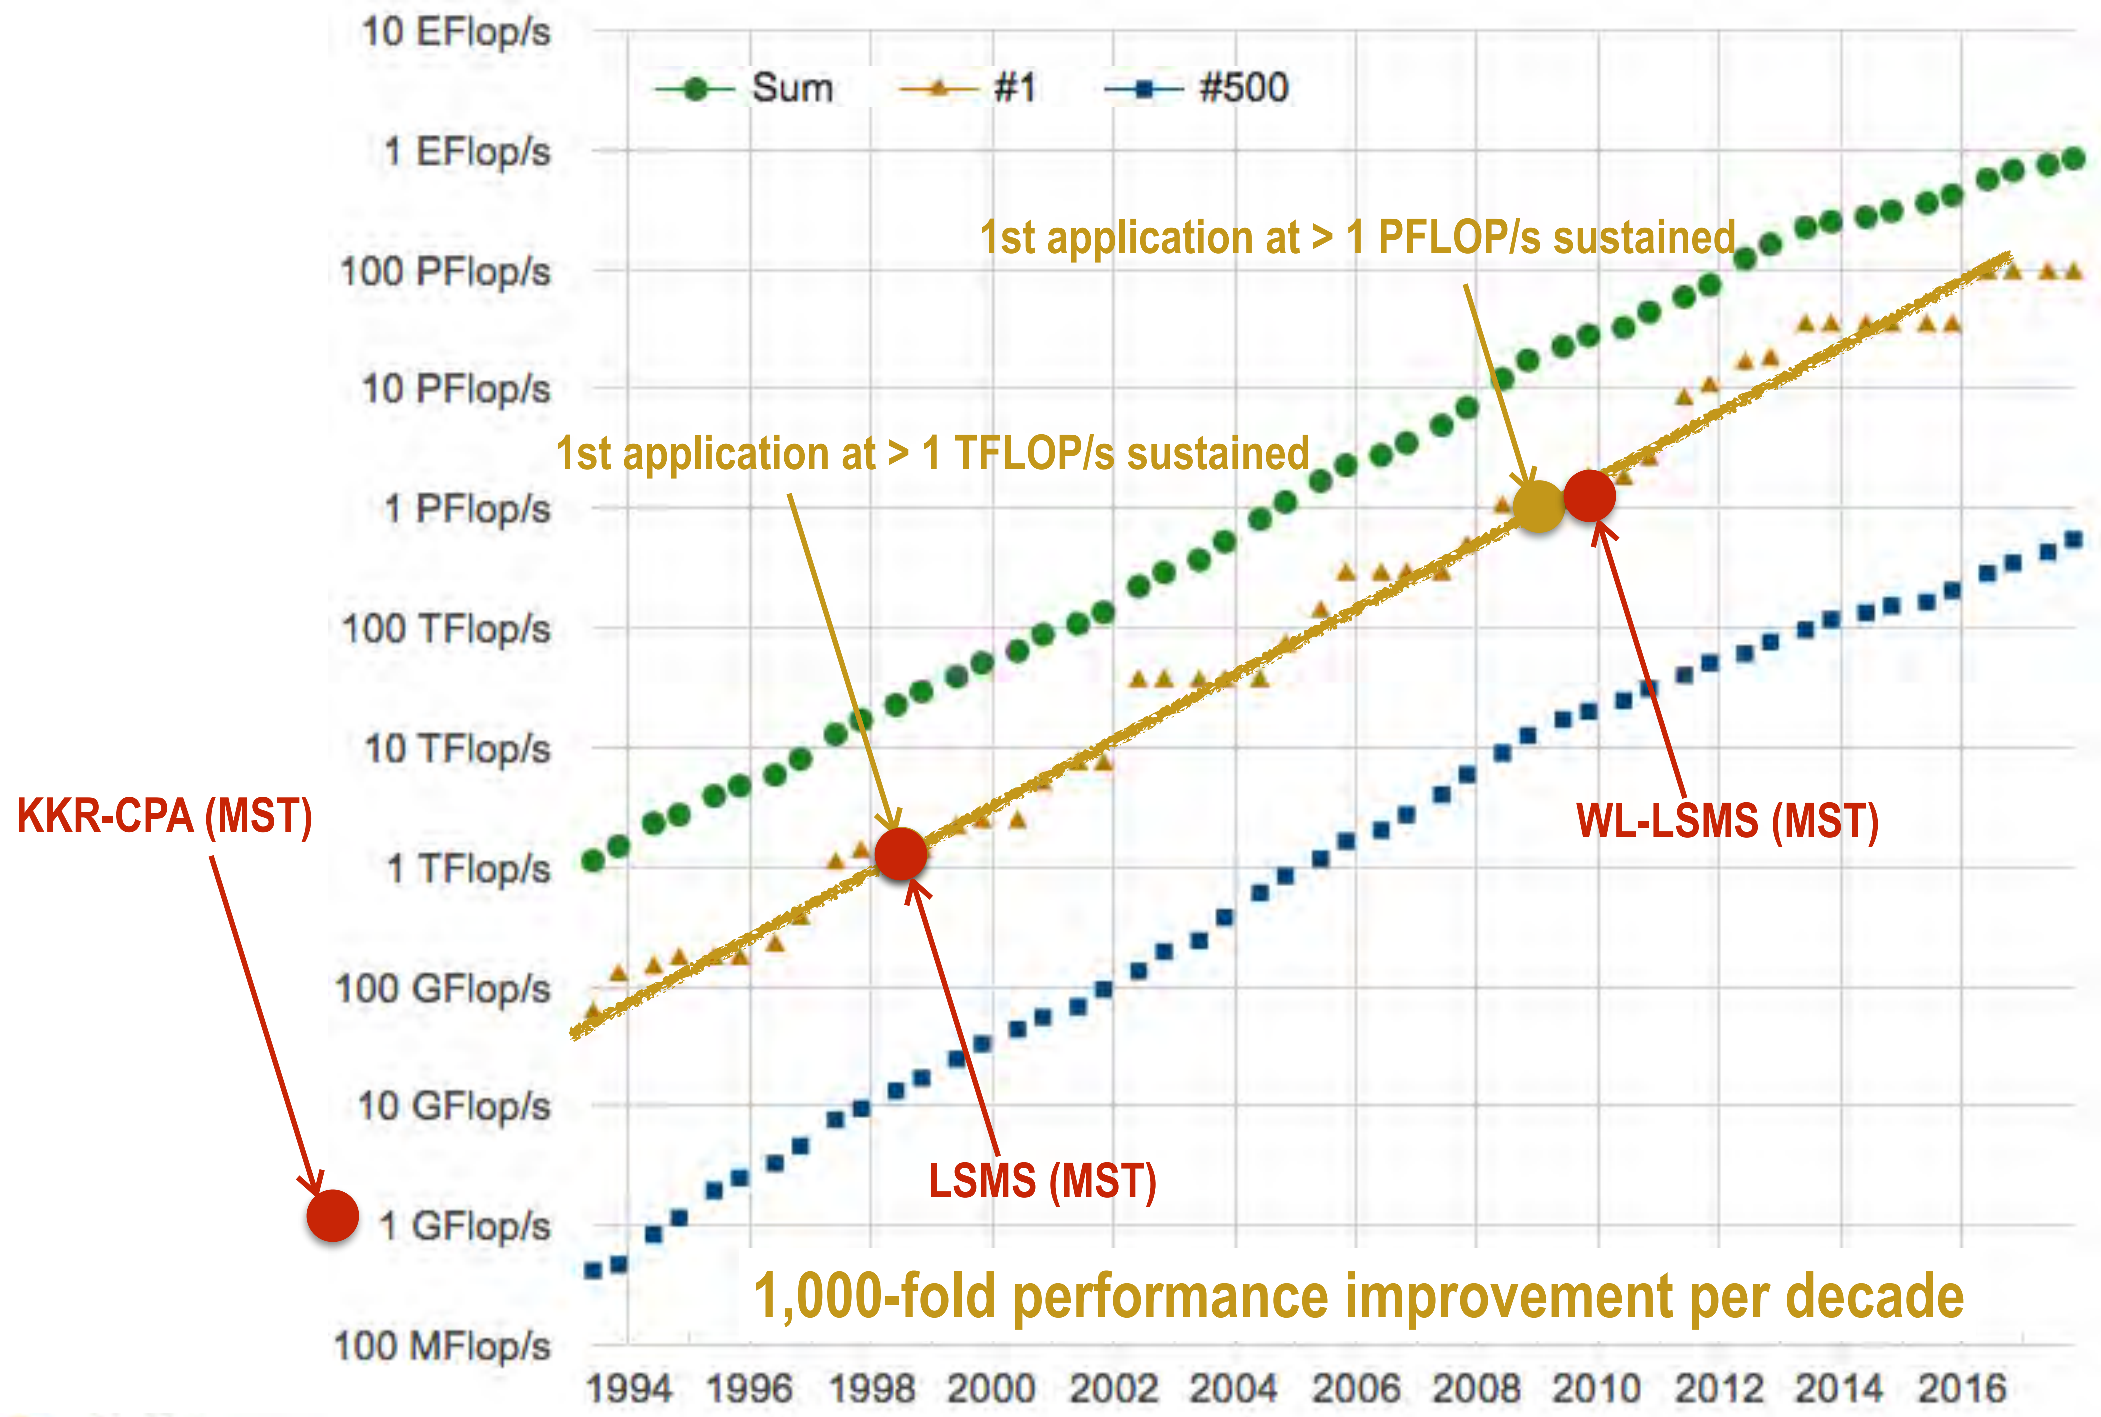

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

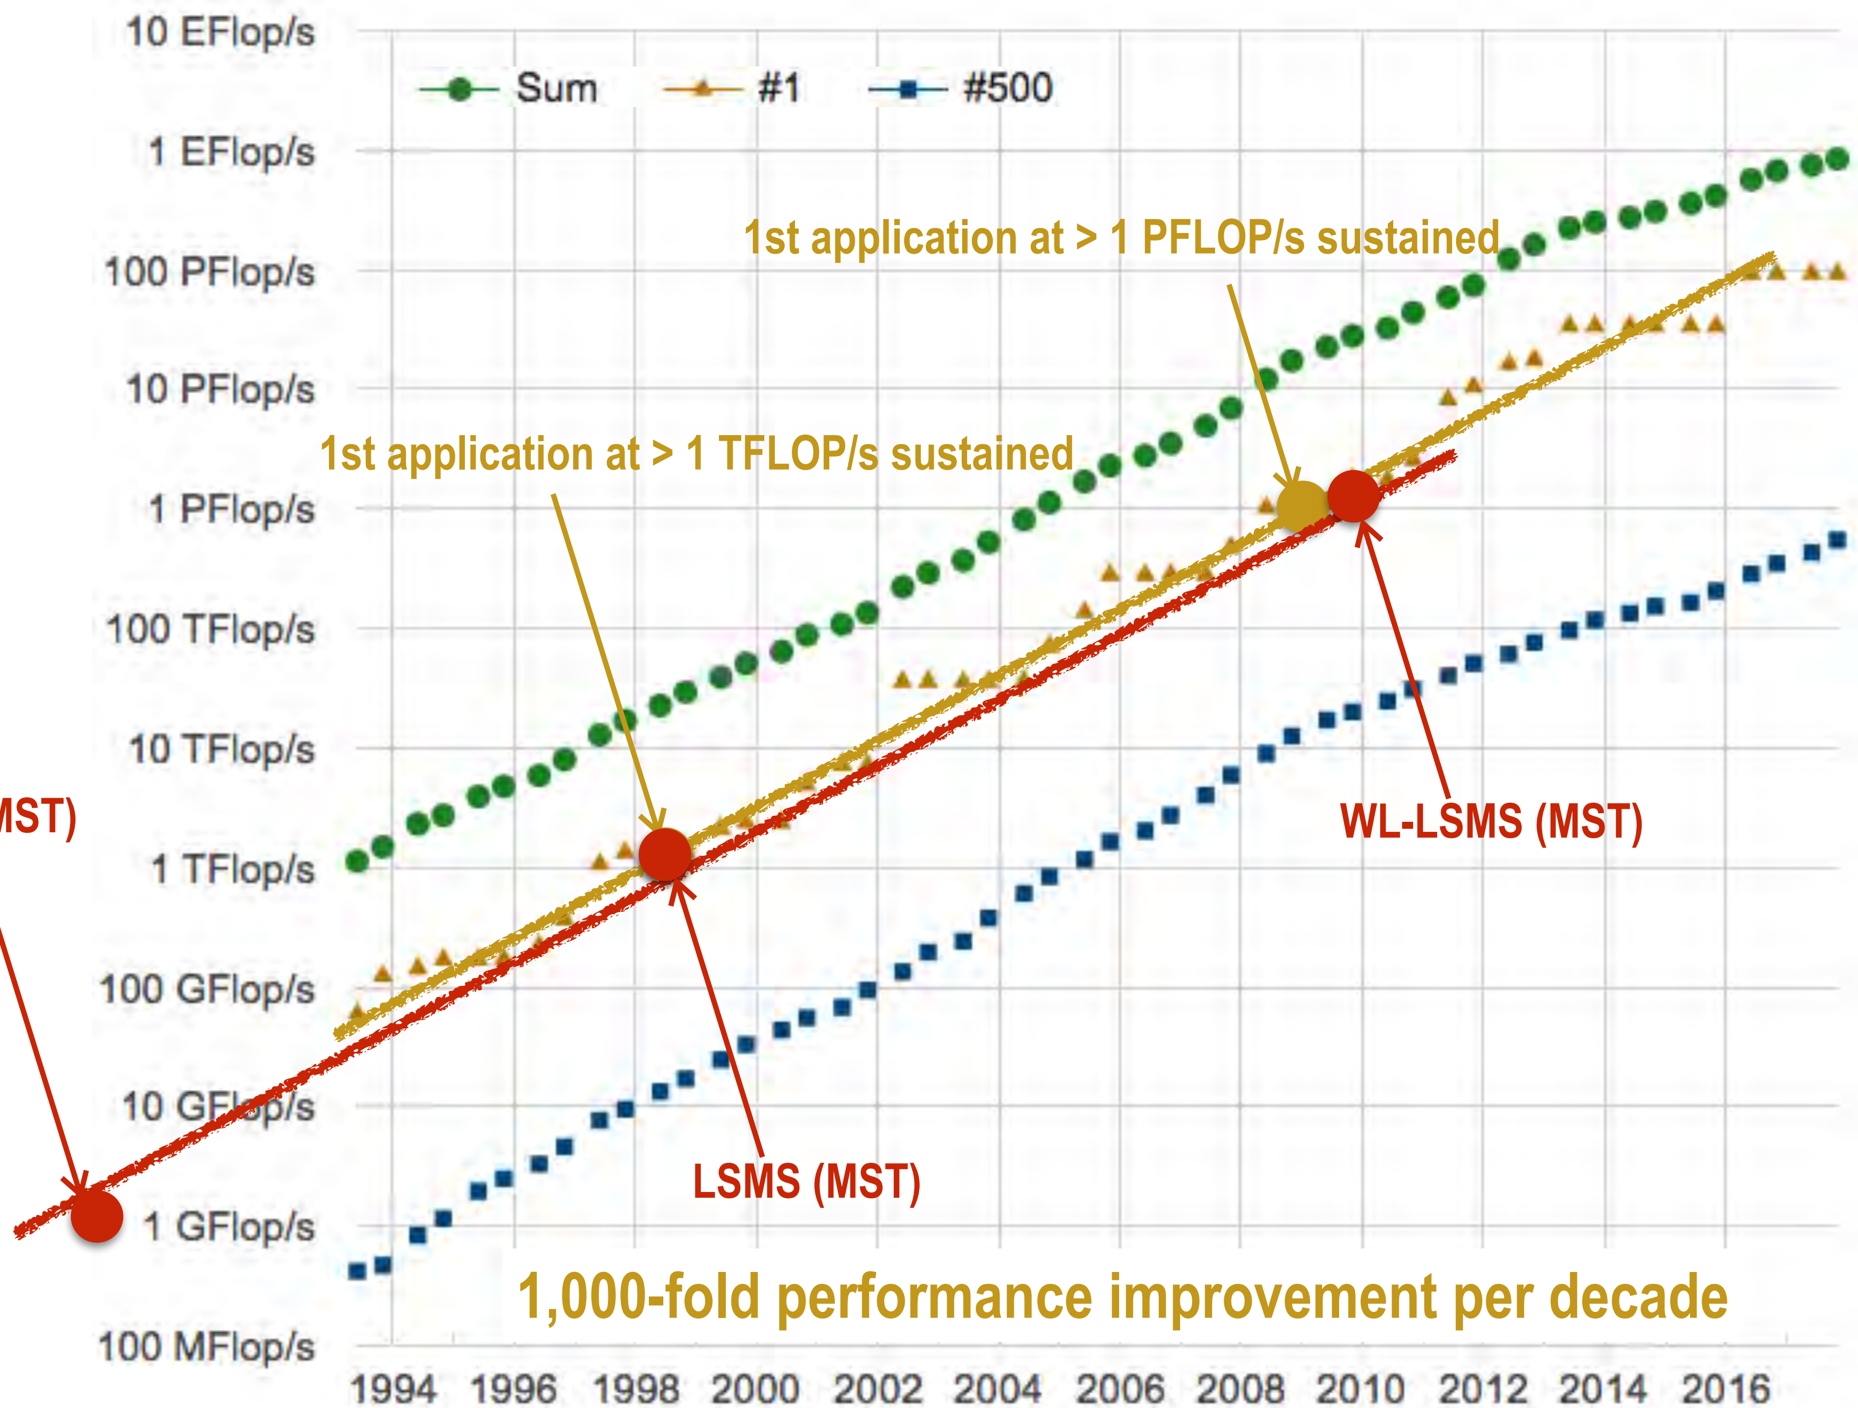

ETH zürich

Tracking supercomputer performance over time?

Linpack benchmark solves: $Ax = b$

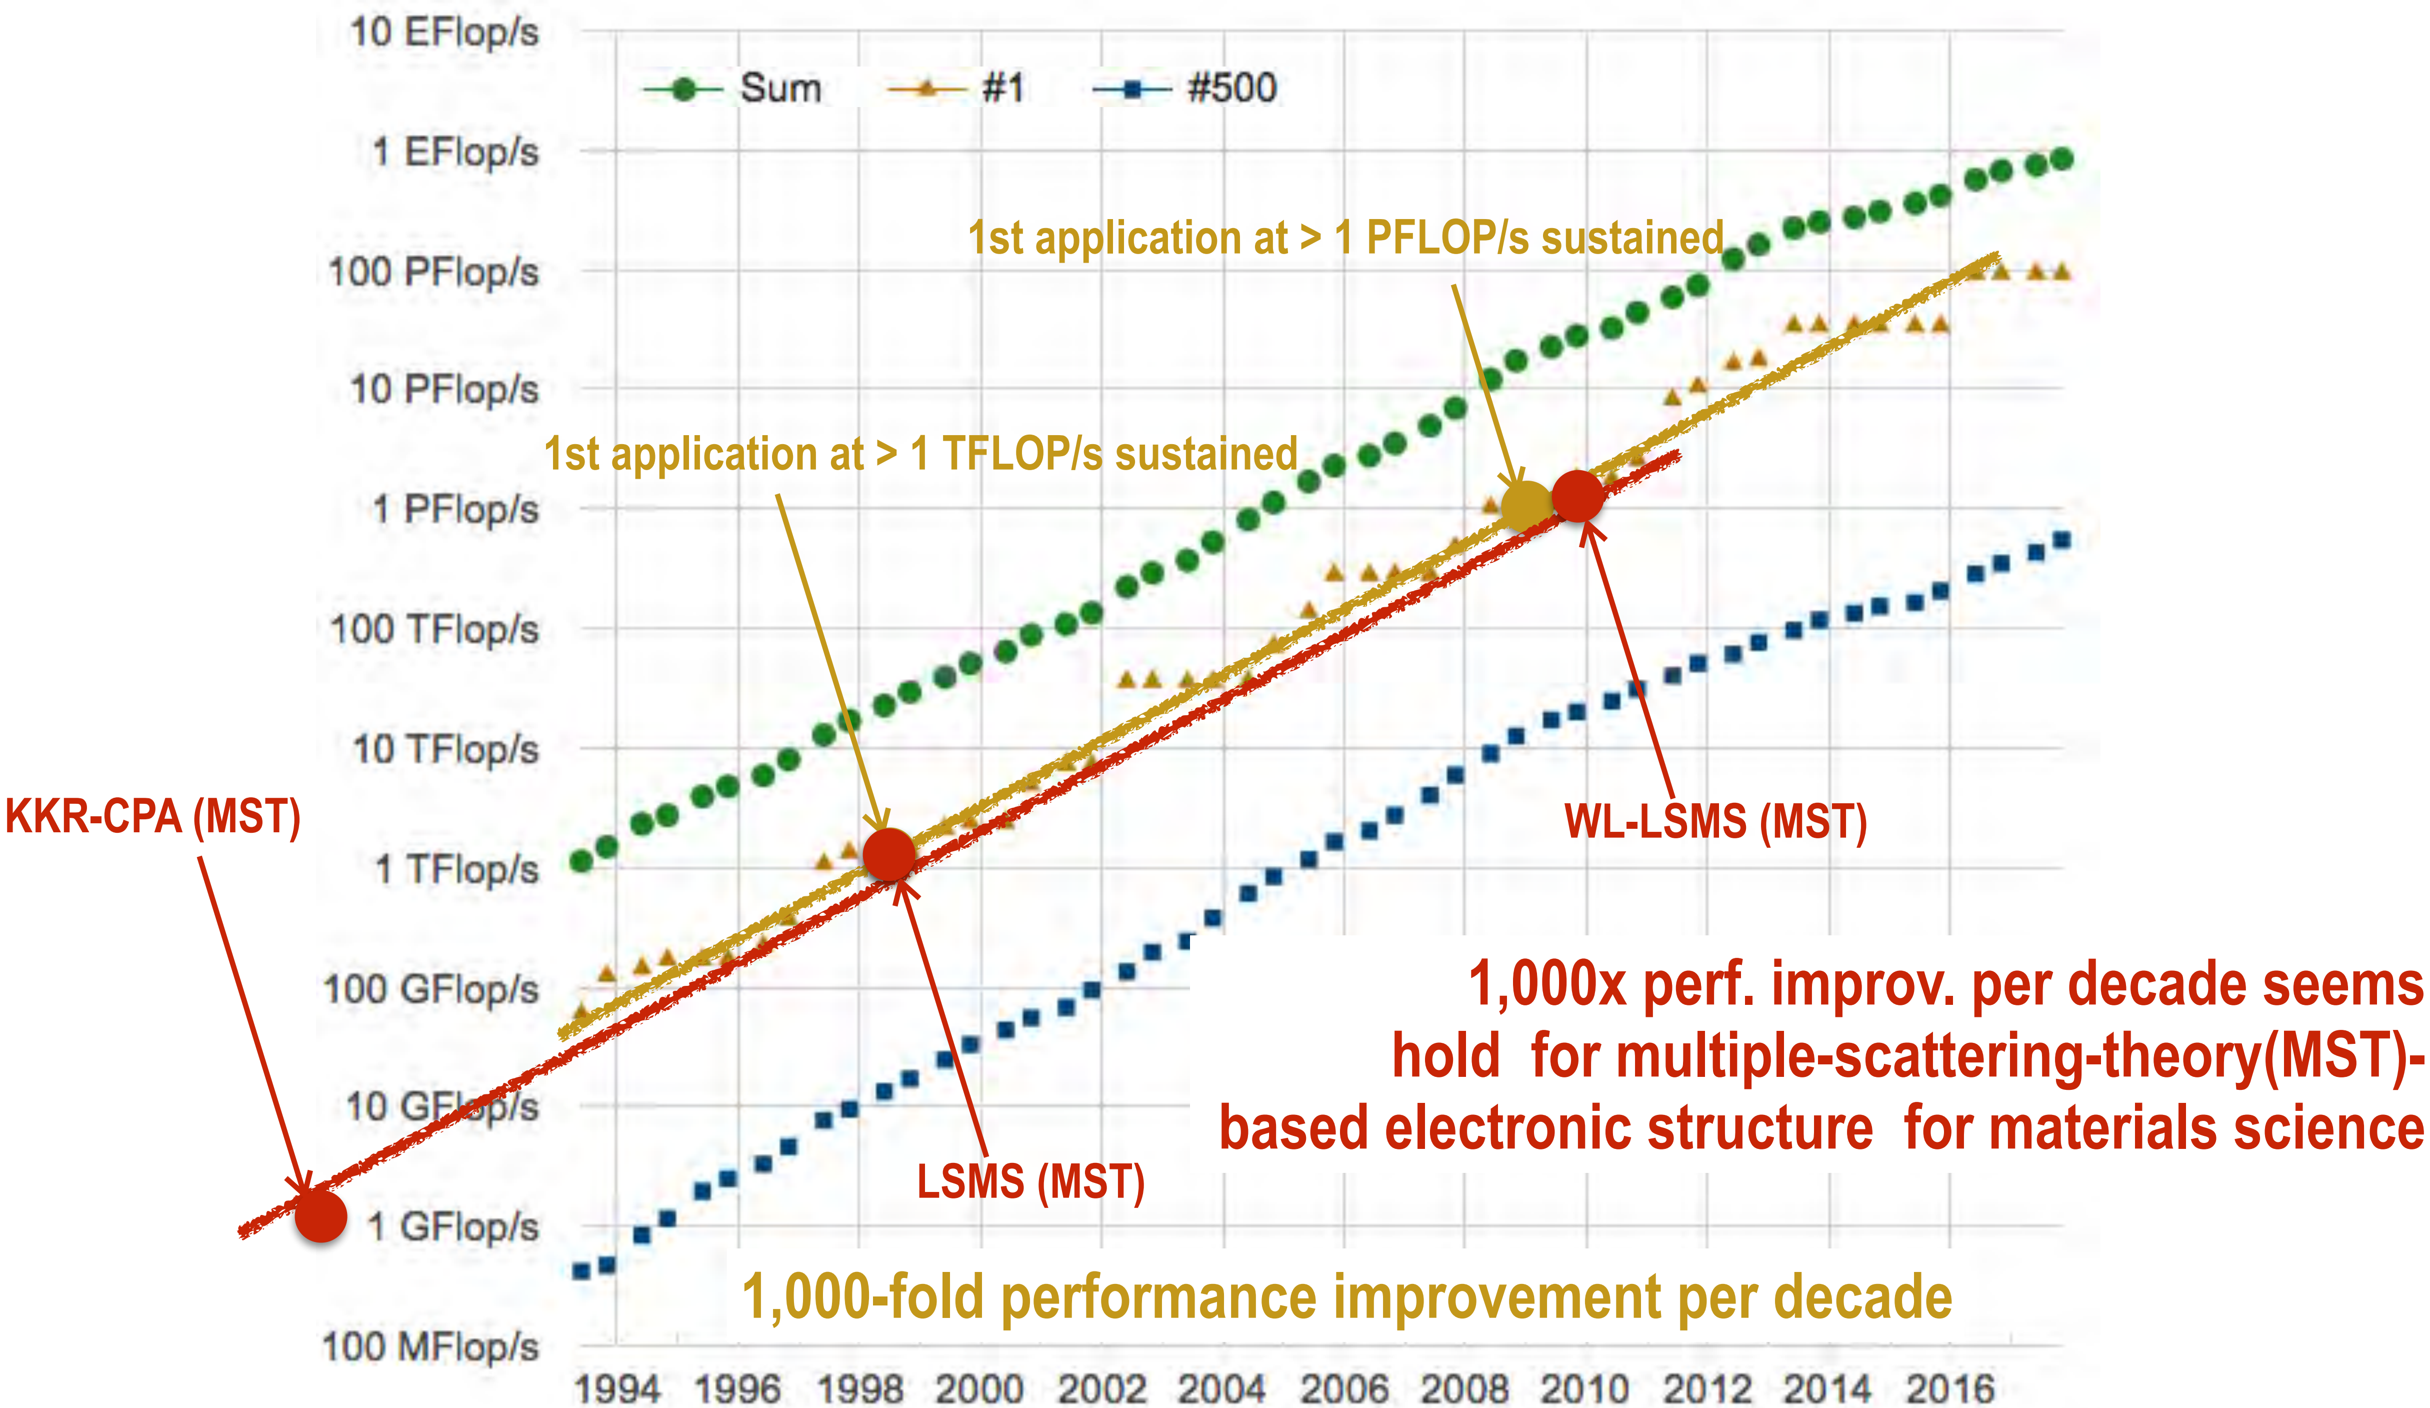

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

“Only” 100-fold performance improvement in climate codes

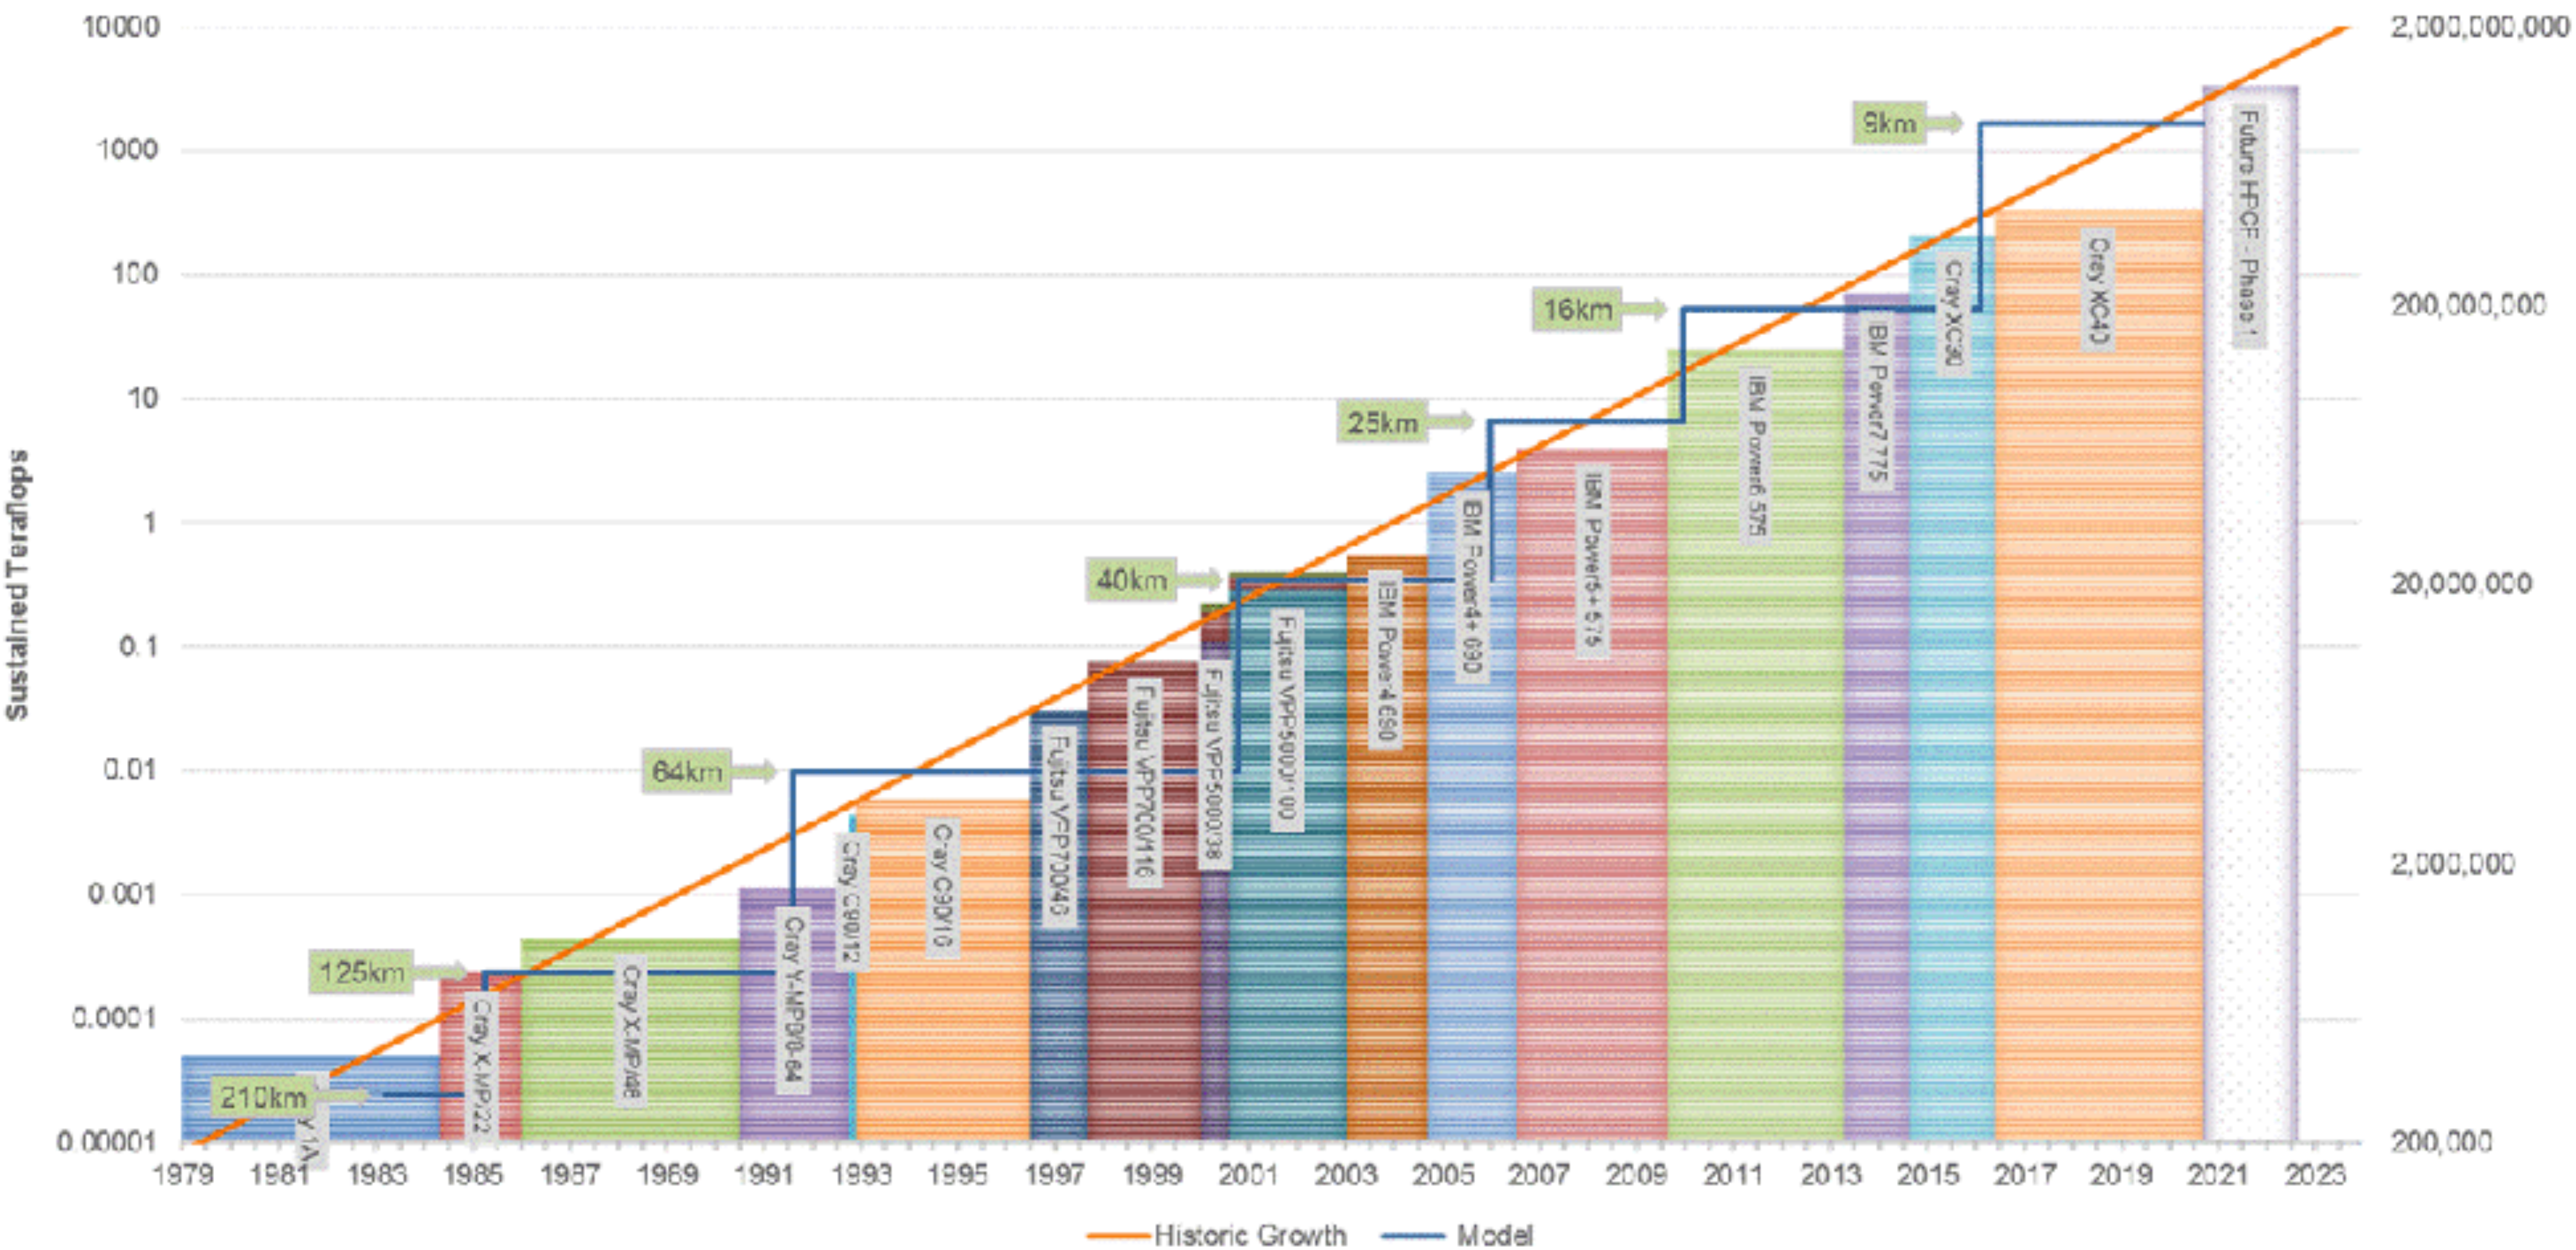

Source: Peter Bauer, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Has the efficiency of weather & climate codes dropped 10-fold every decade?

ETH zürich

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

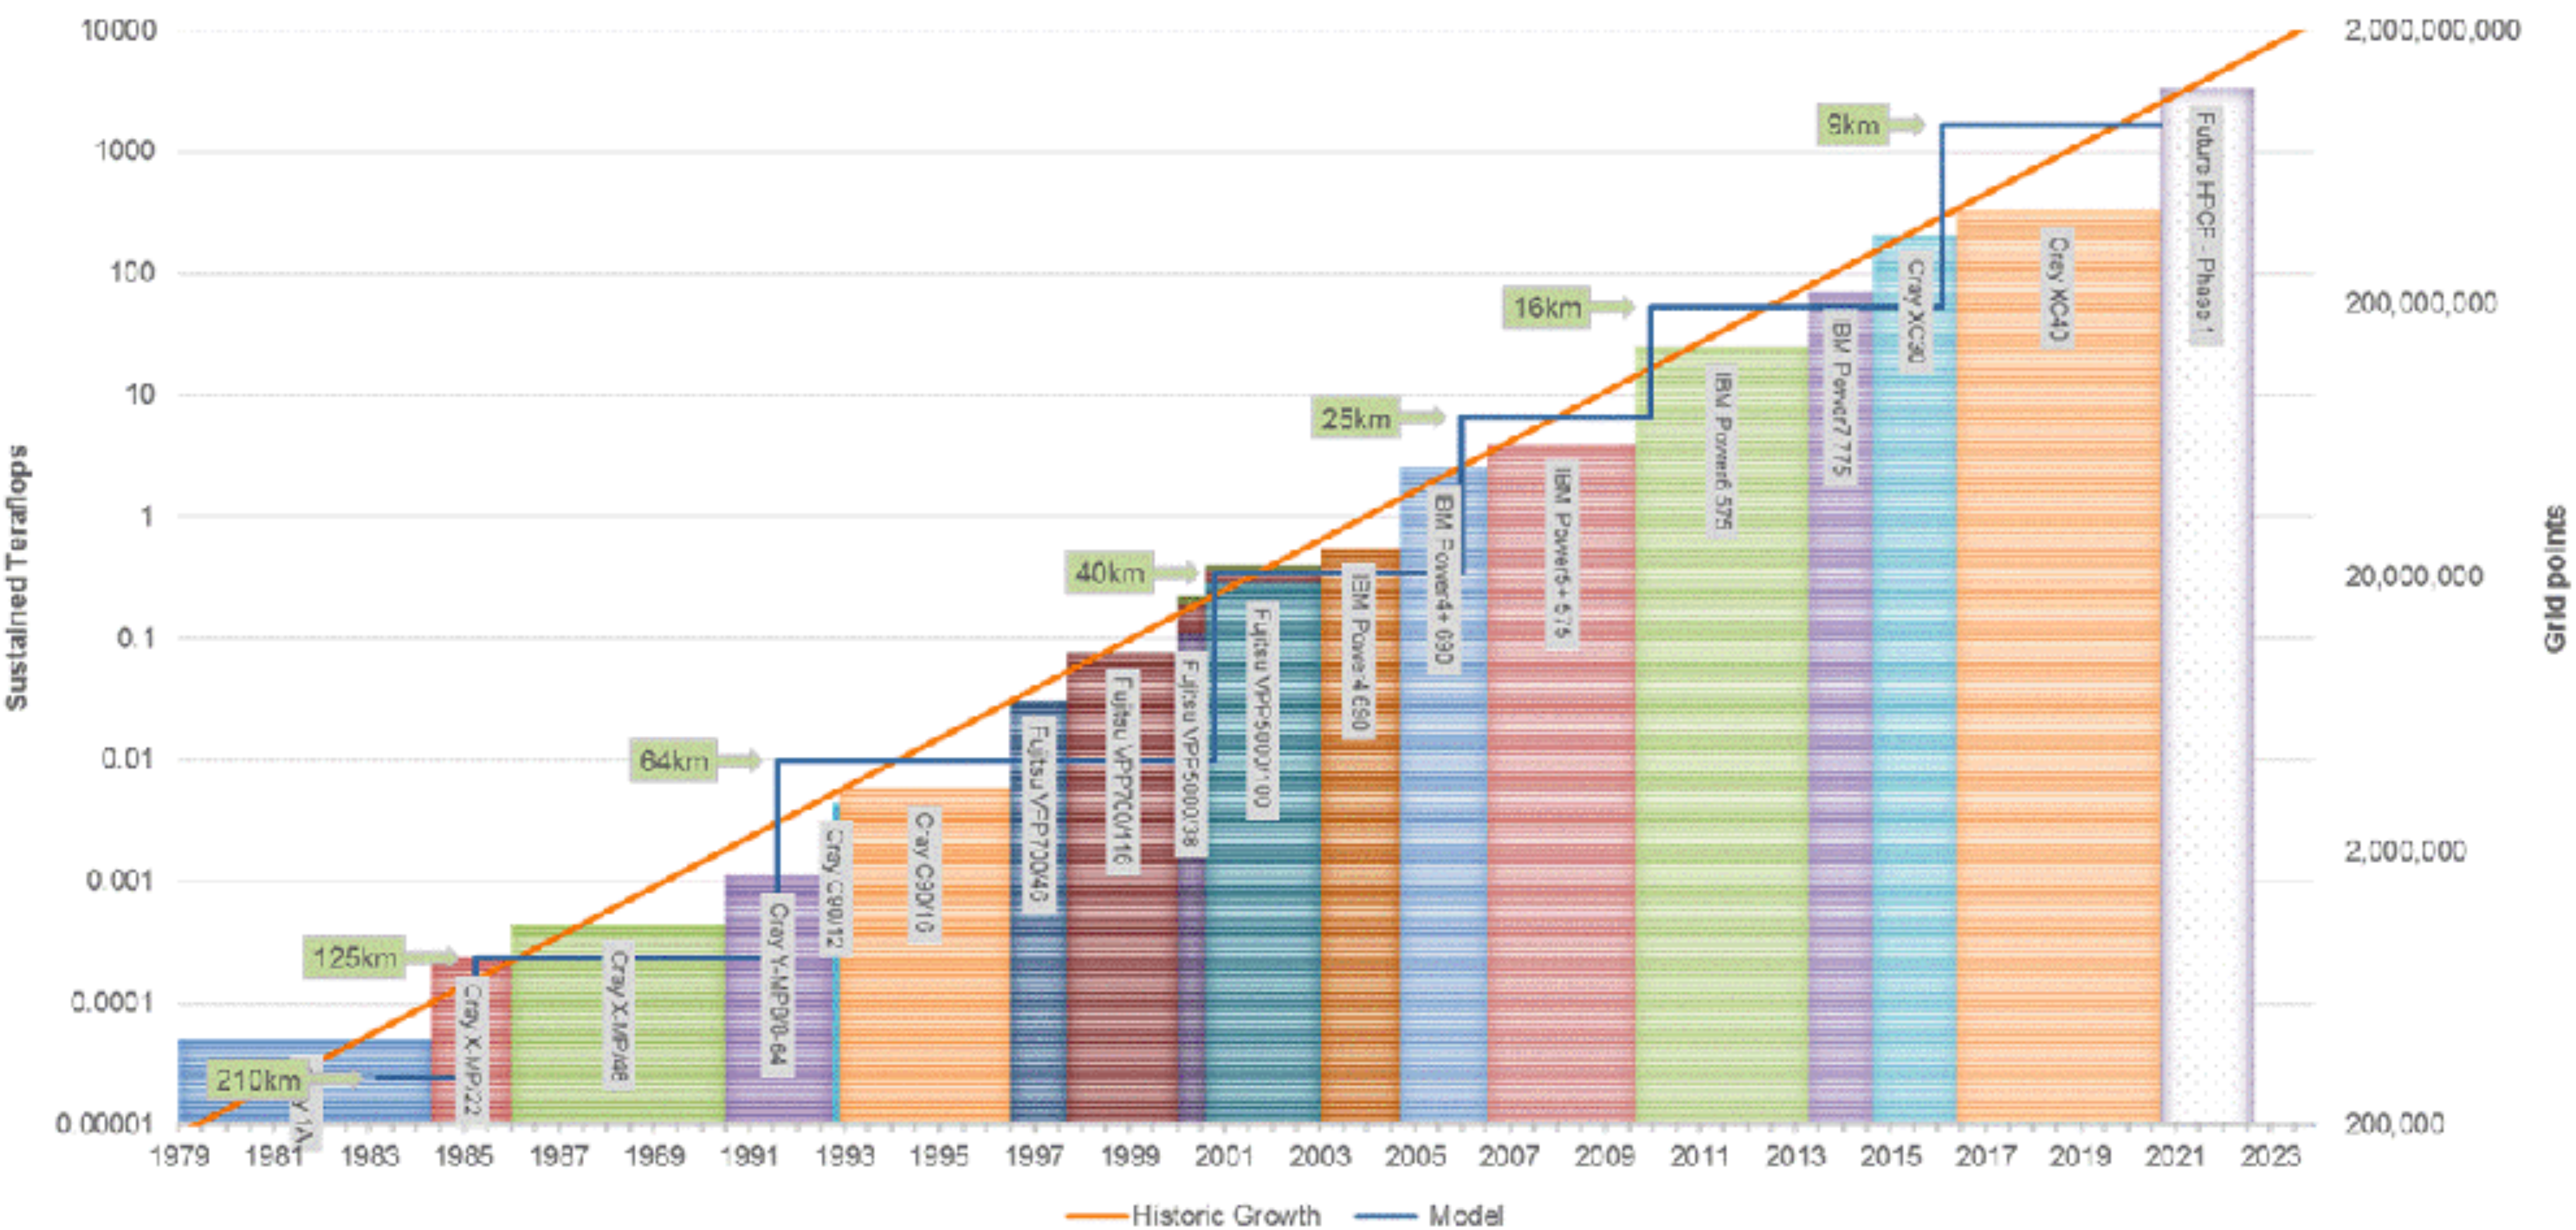

Source: Peter Bauer, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

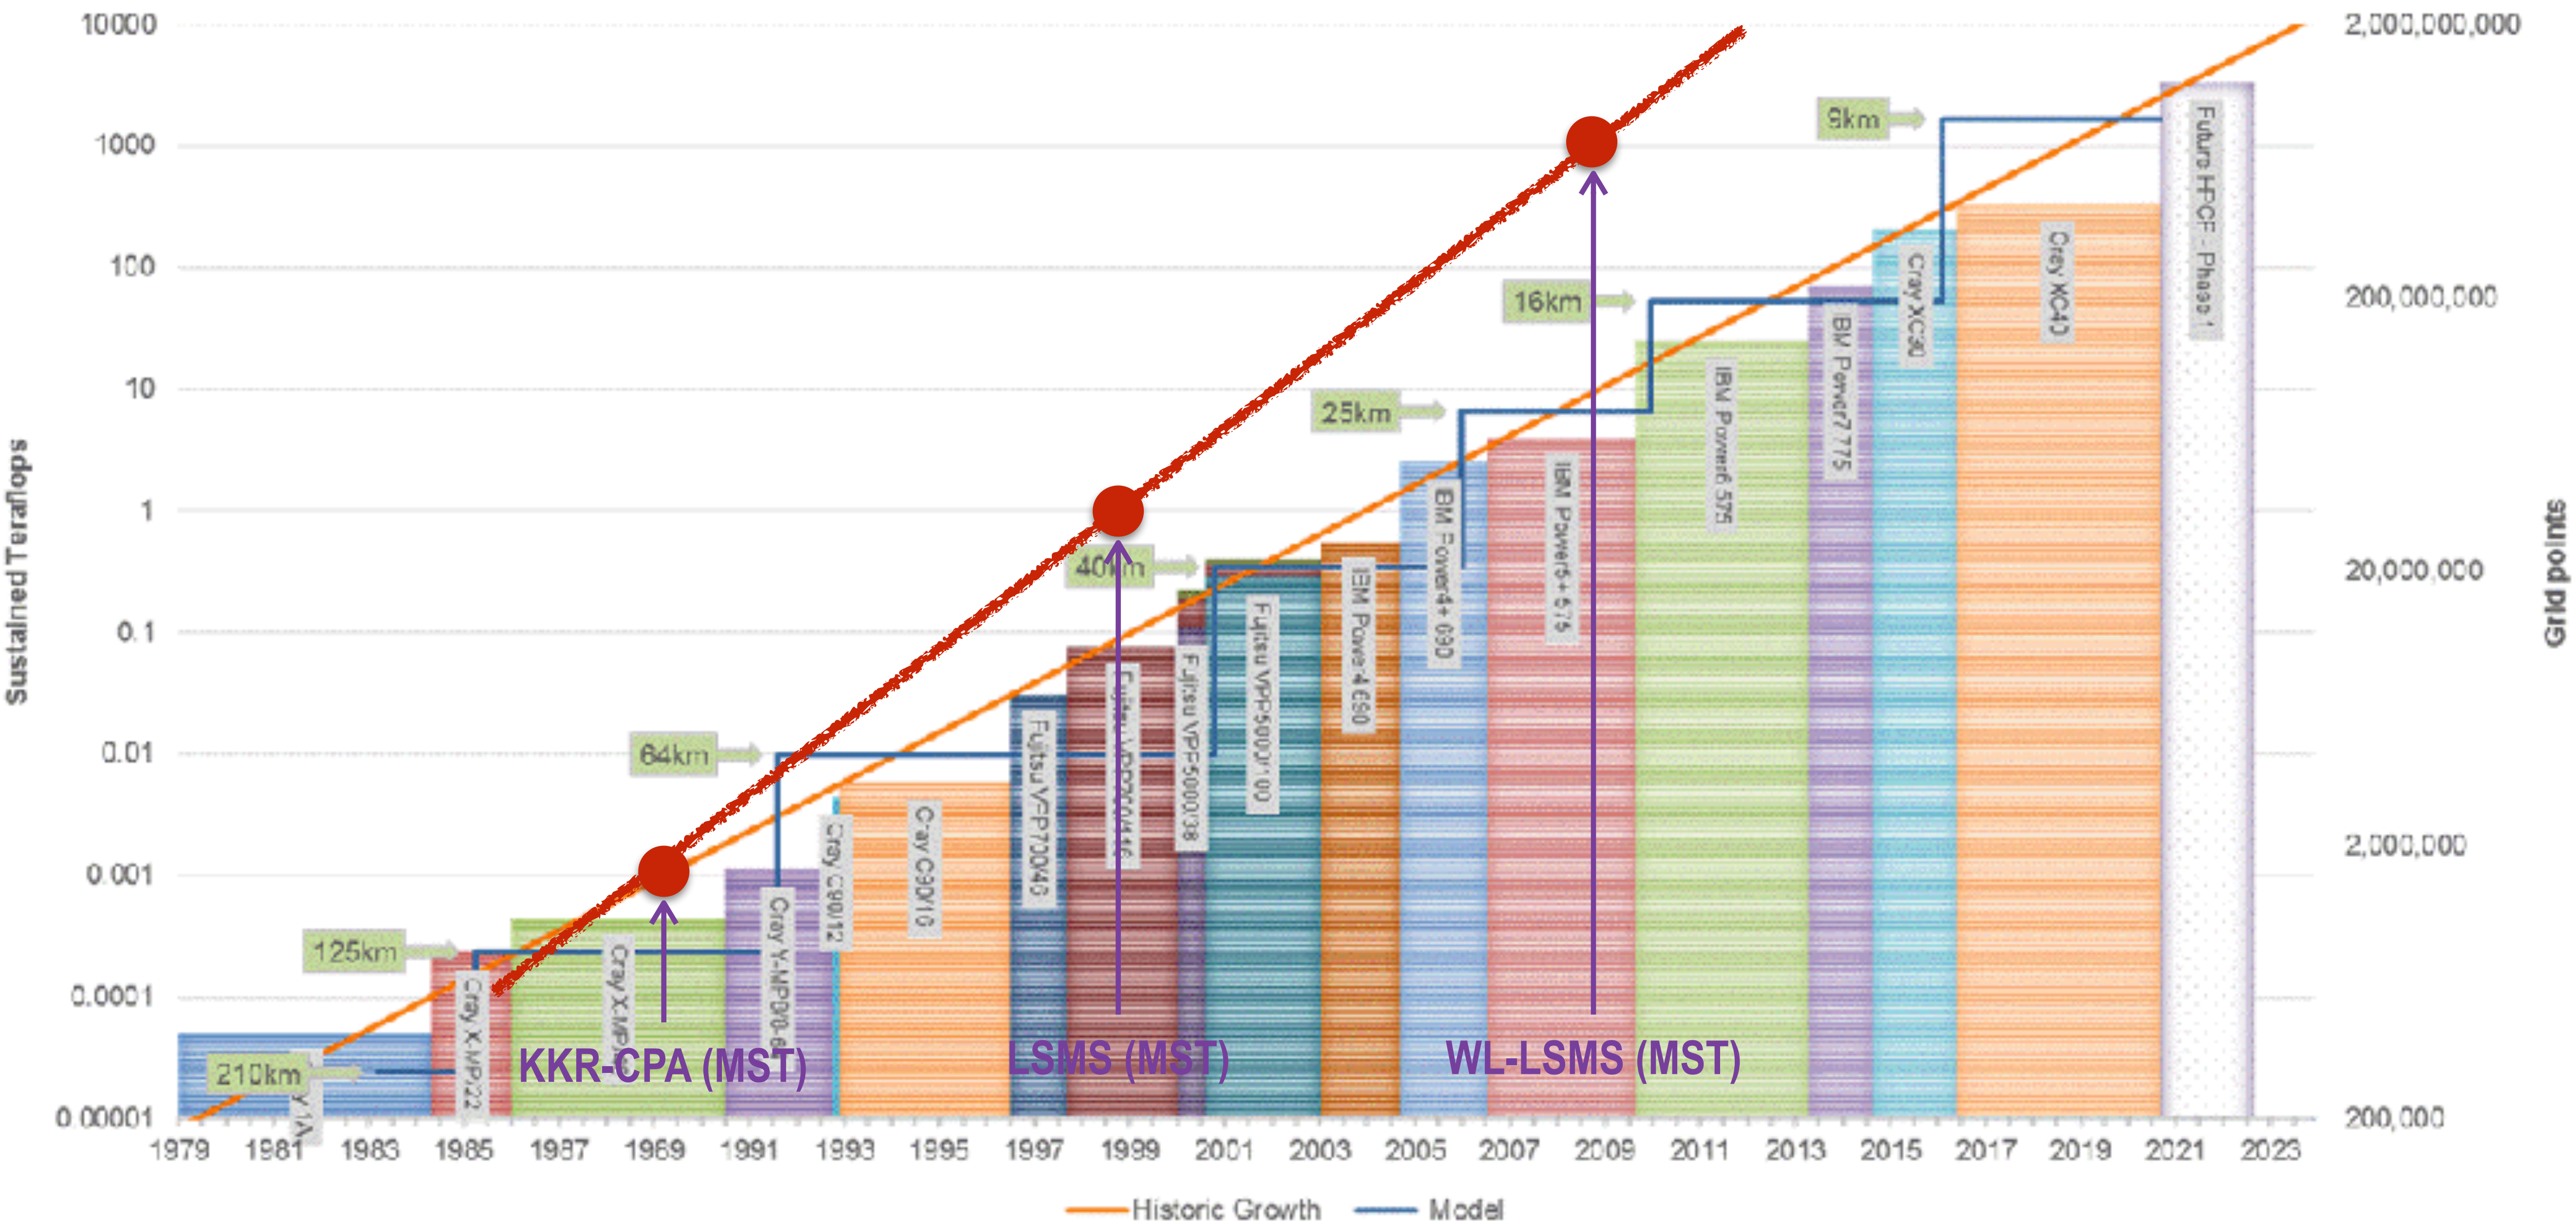

Source: Peter Bauer, ECMWF

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

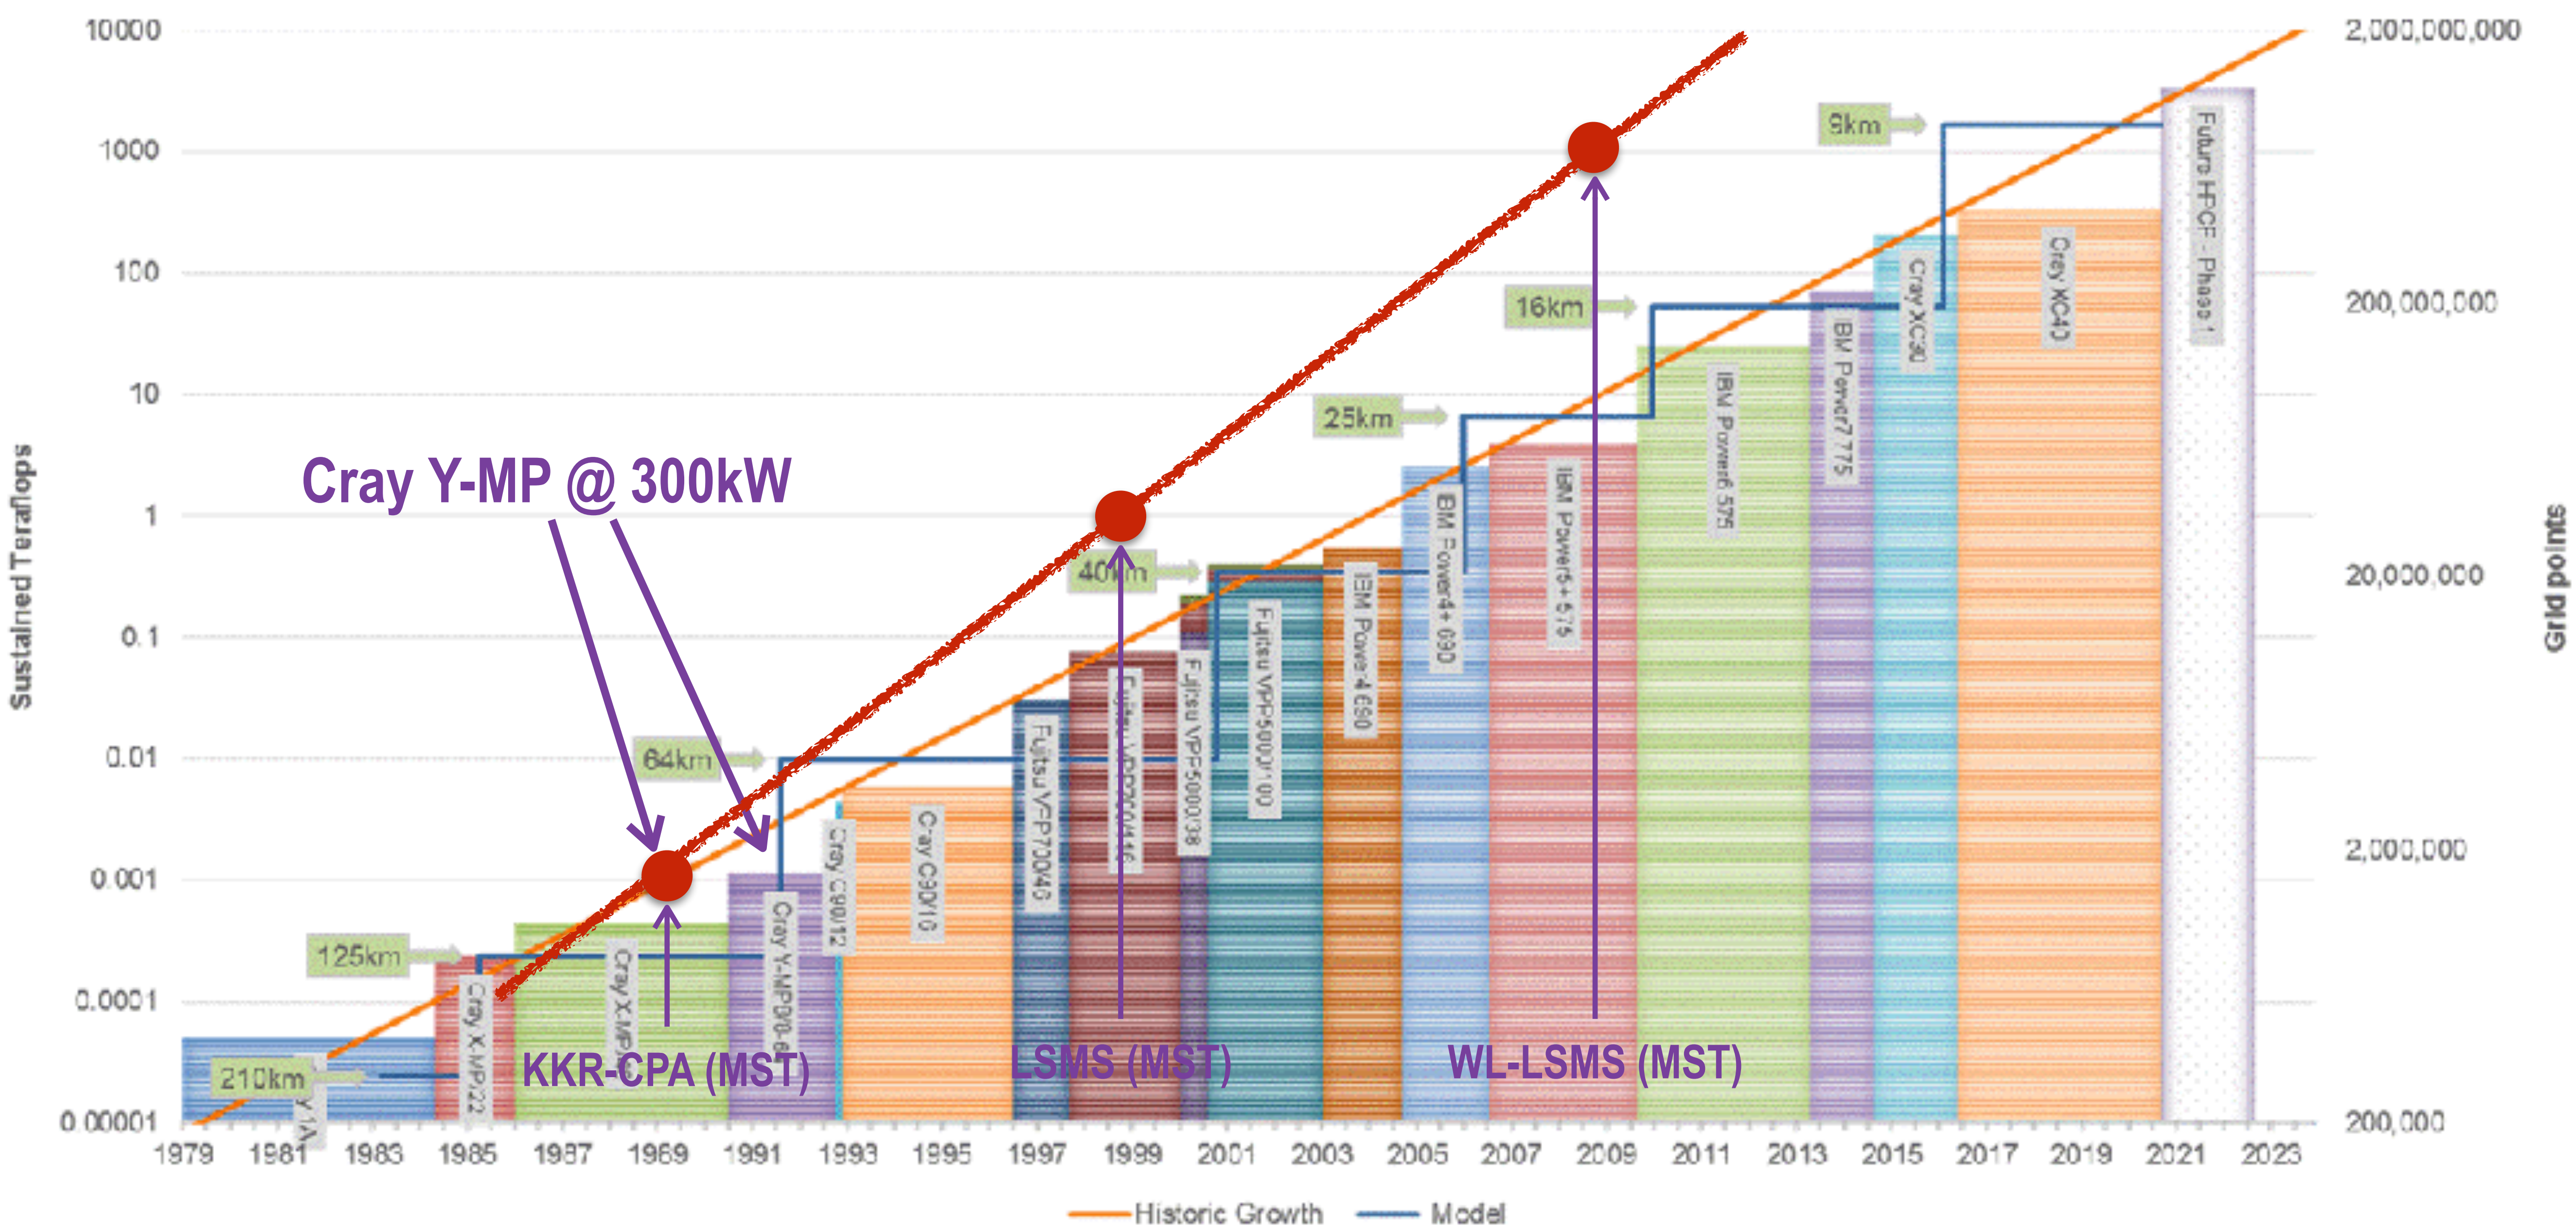

Source: Peter Bauer, ECMWF

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

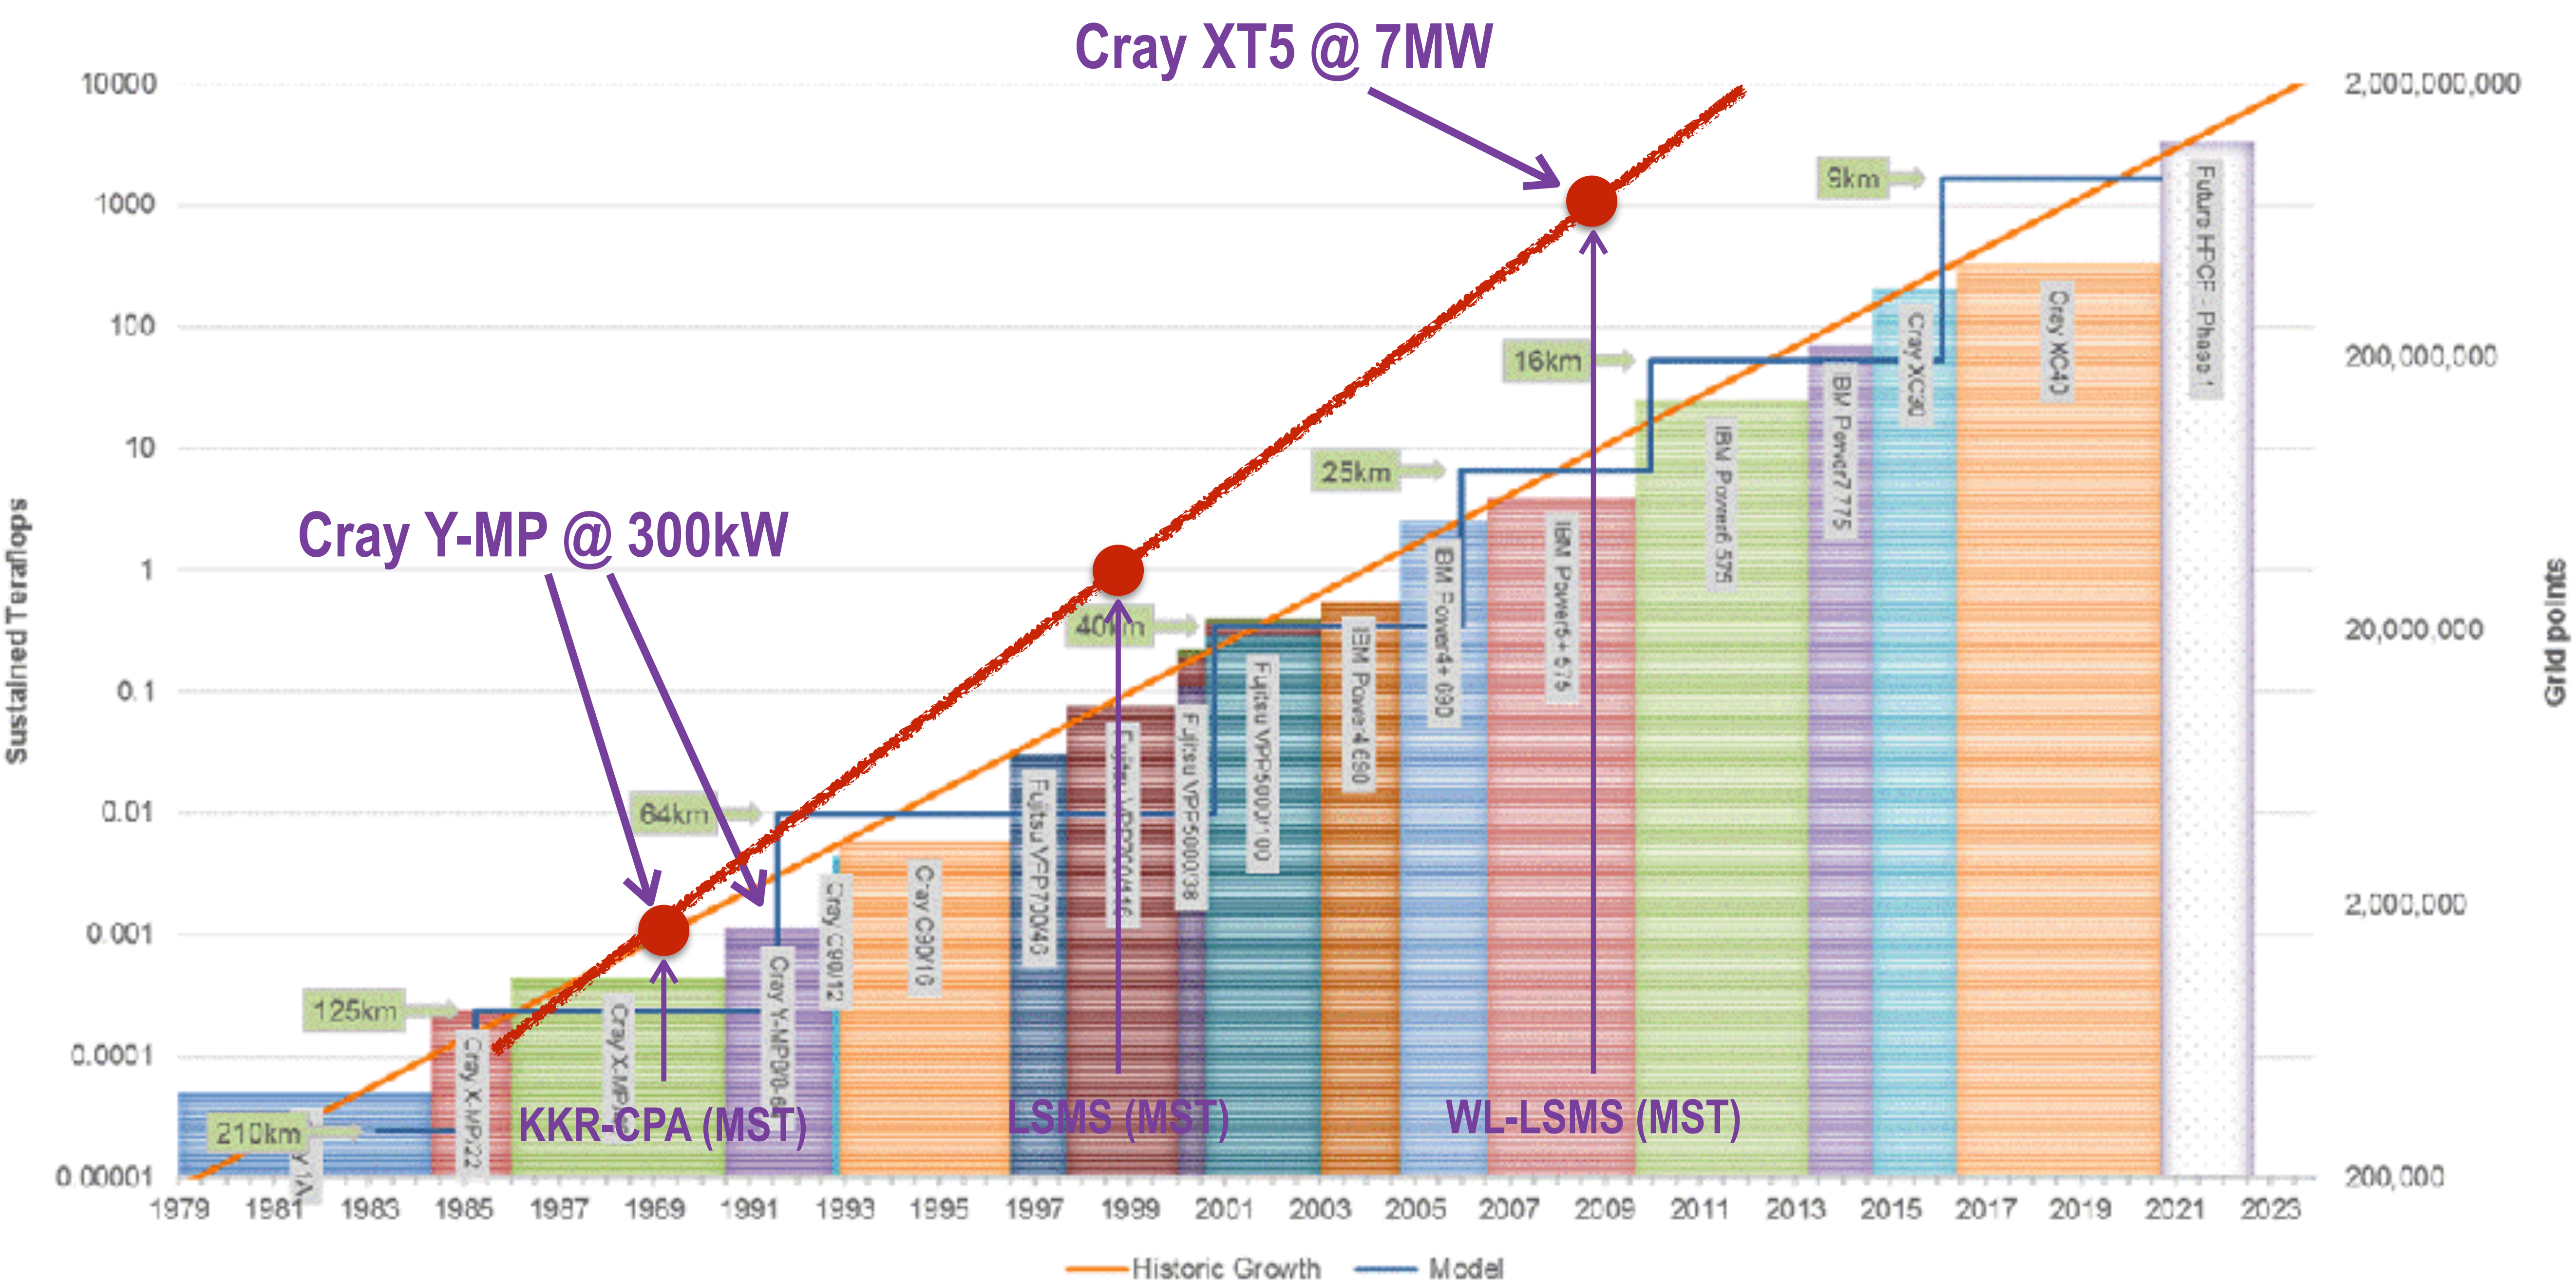

Source: Peter Bauer, ECMWF

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

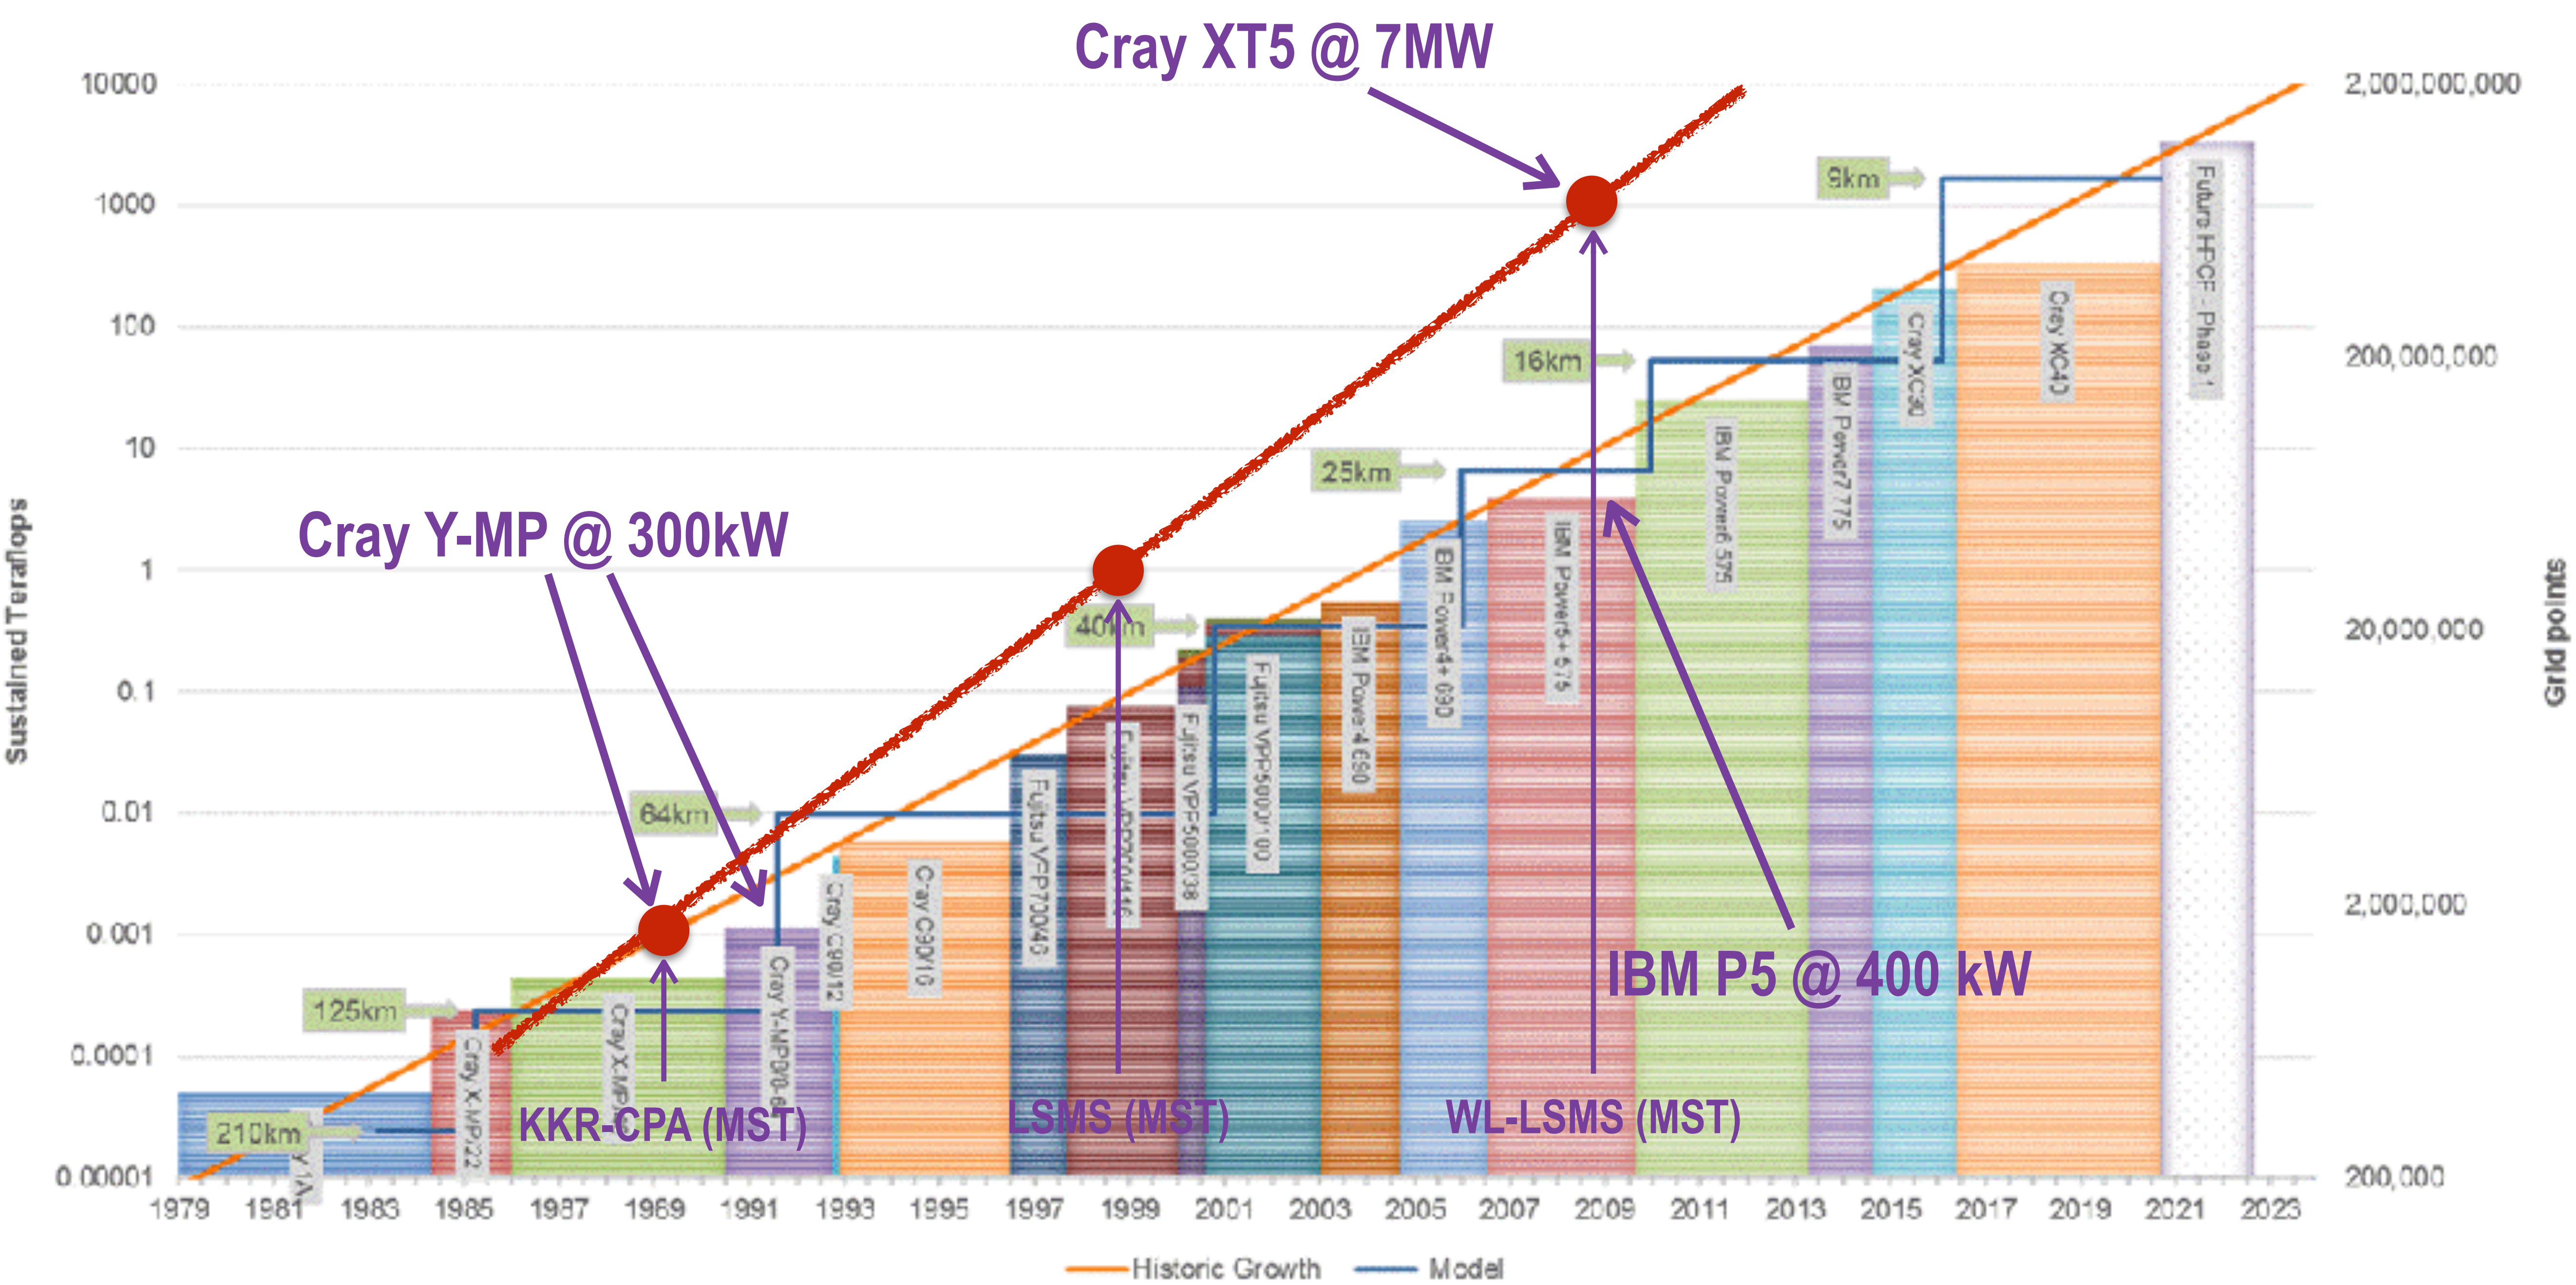

Source: Peter Bauer, ECMWF

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

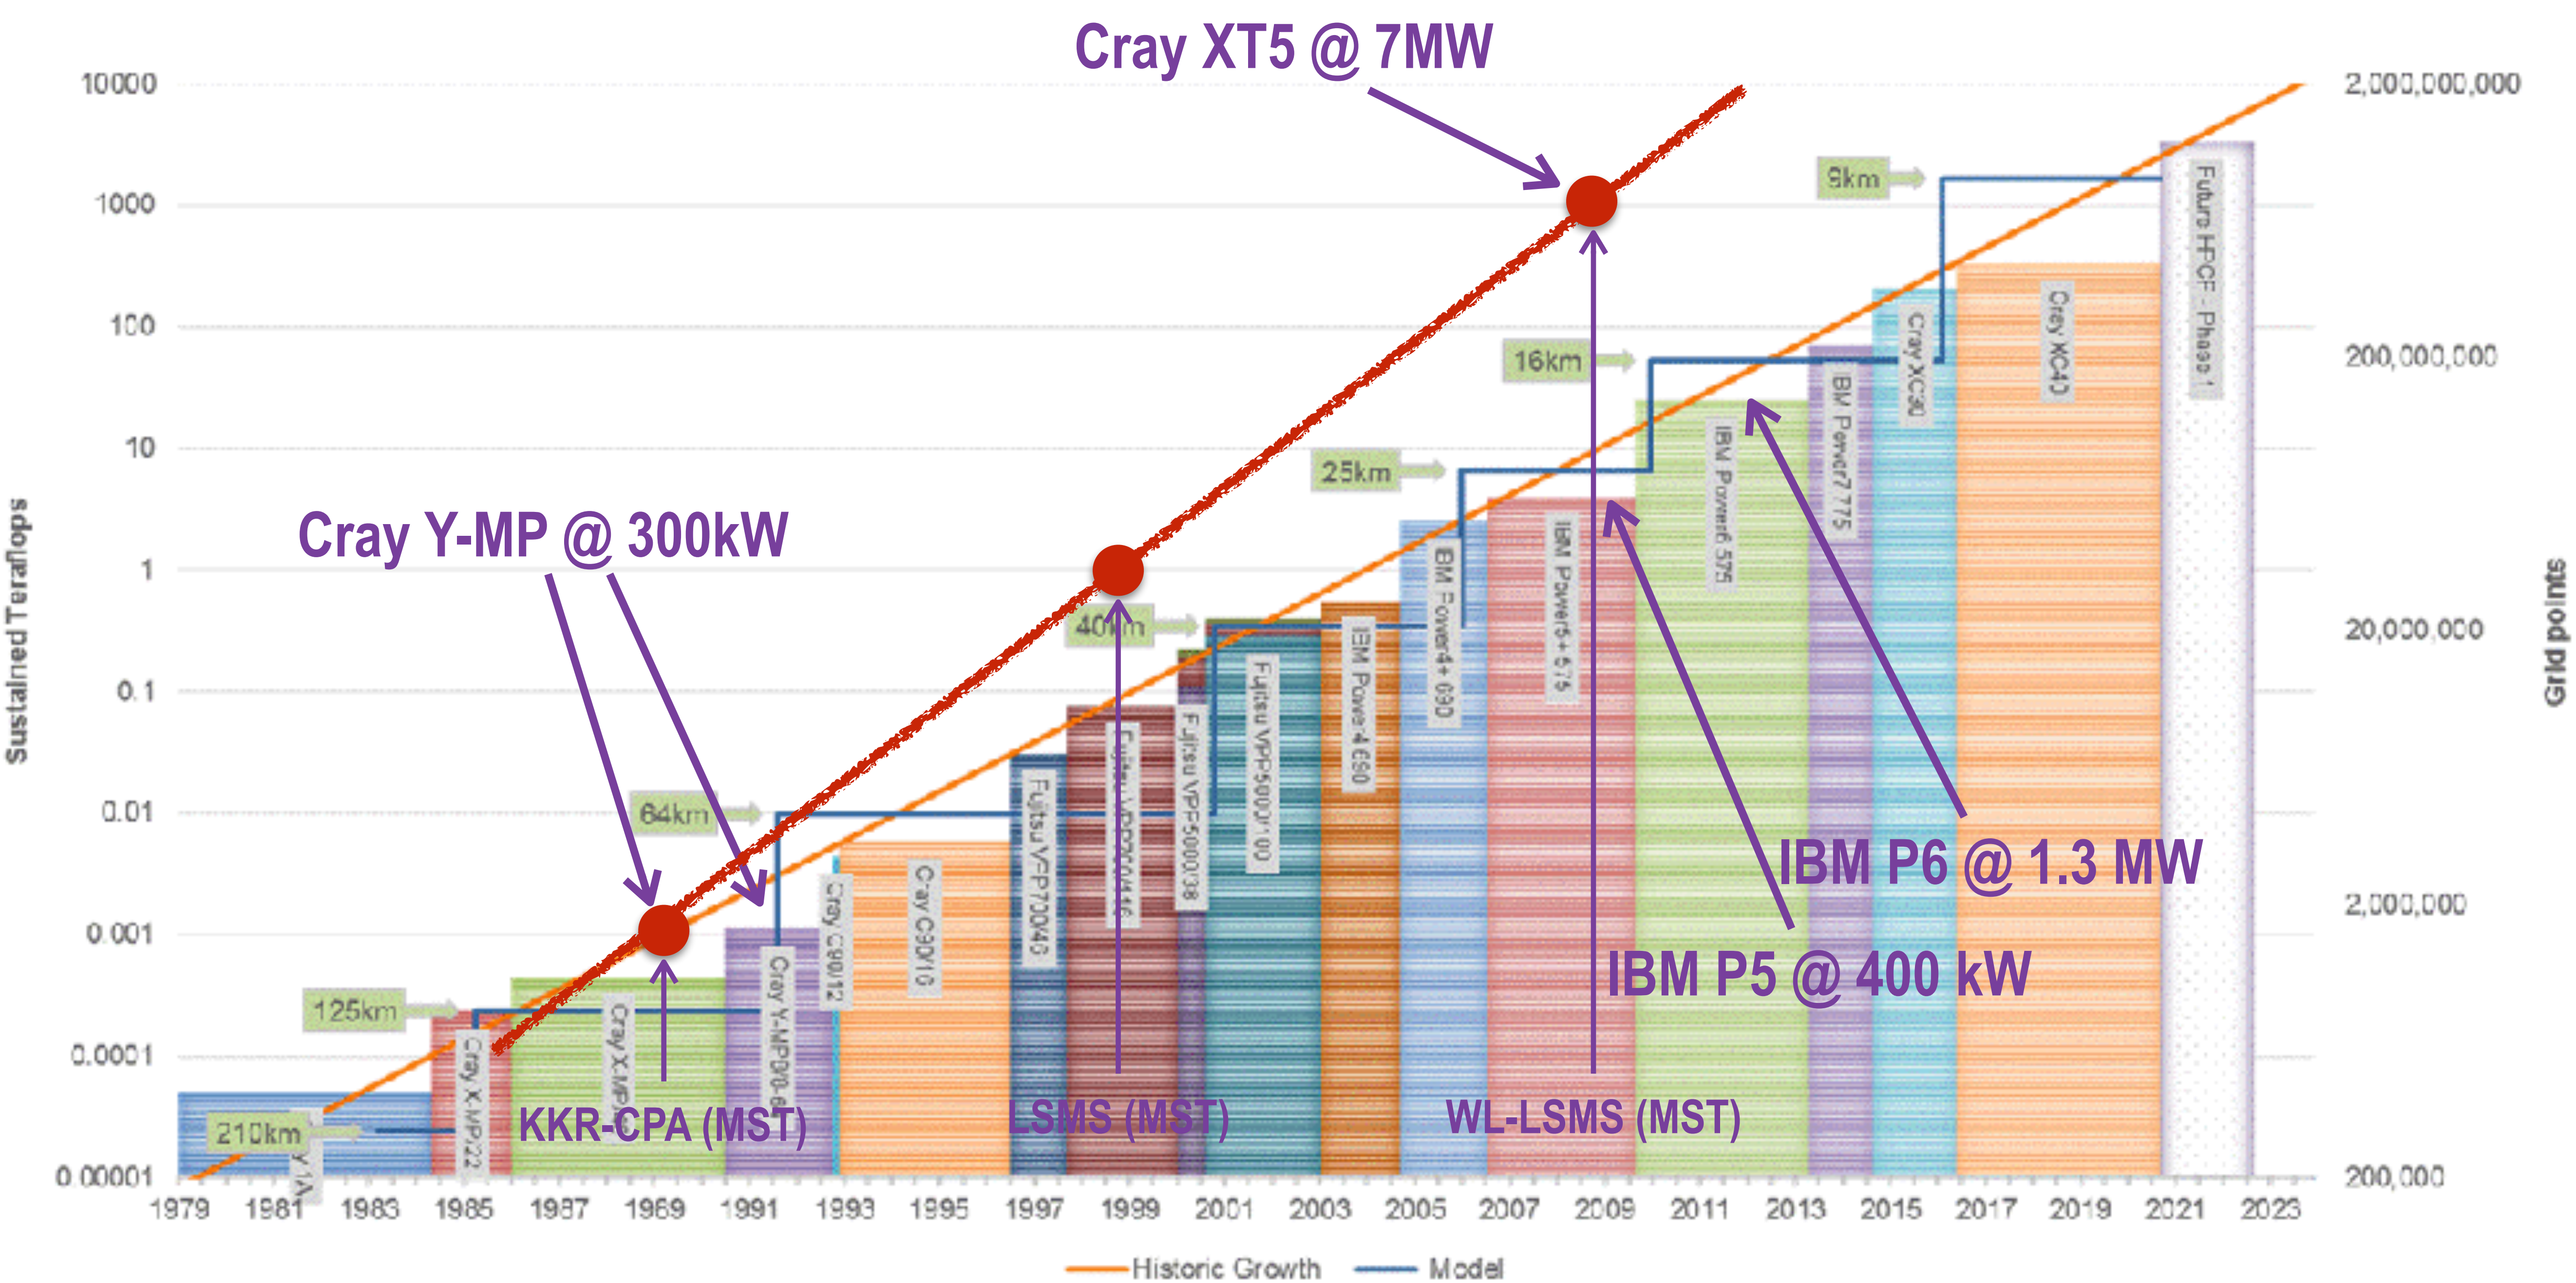

Source: Peter Bauer, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

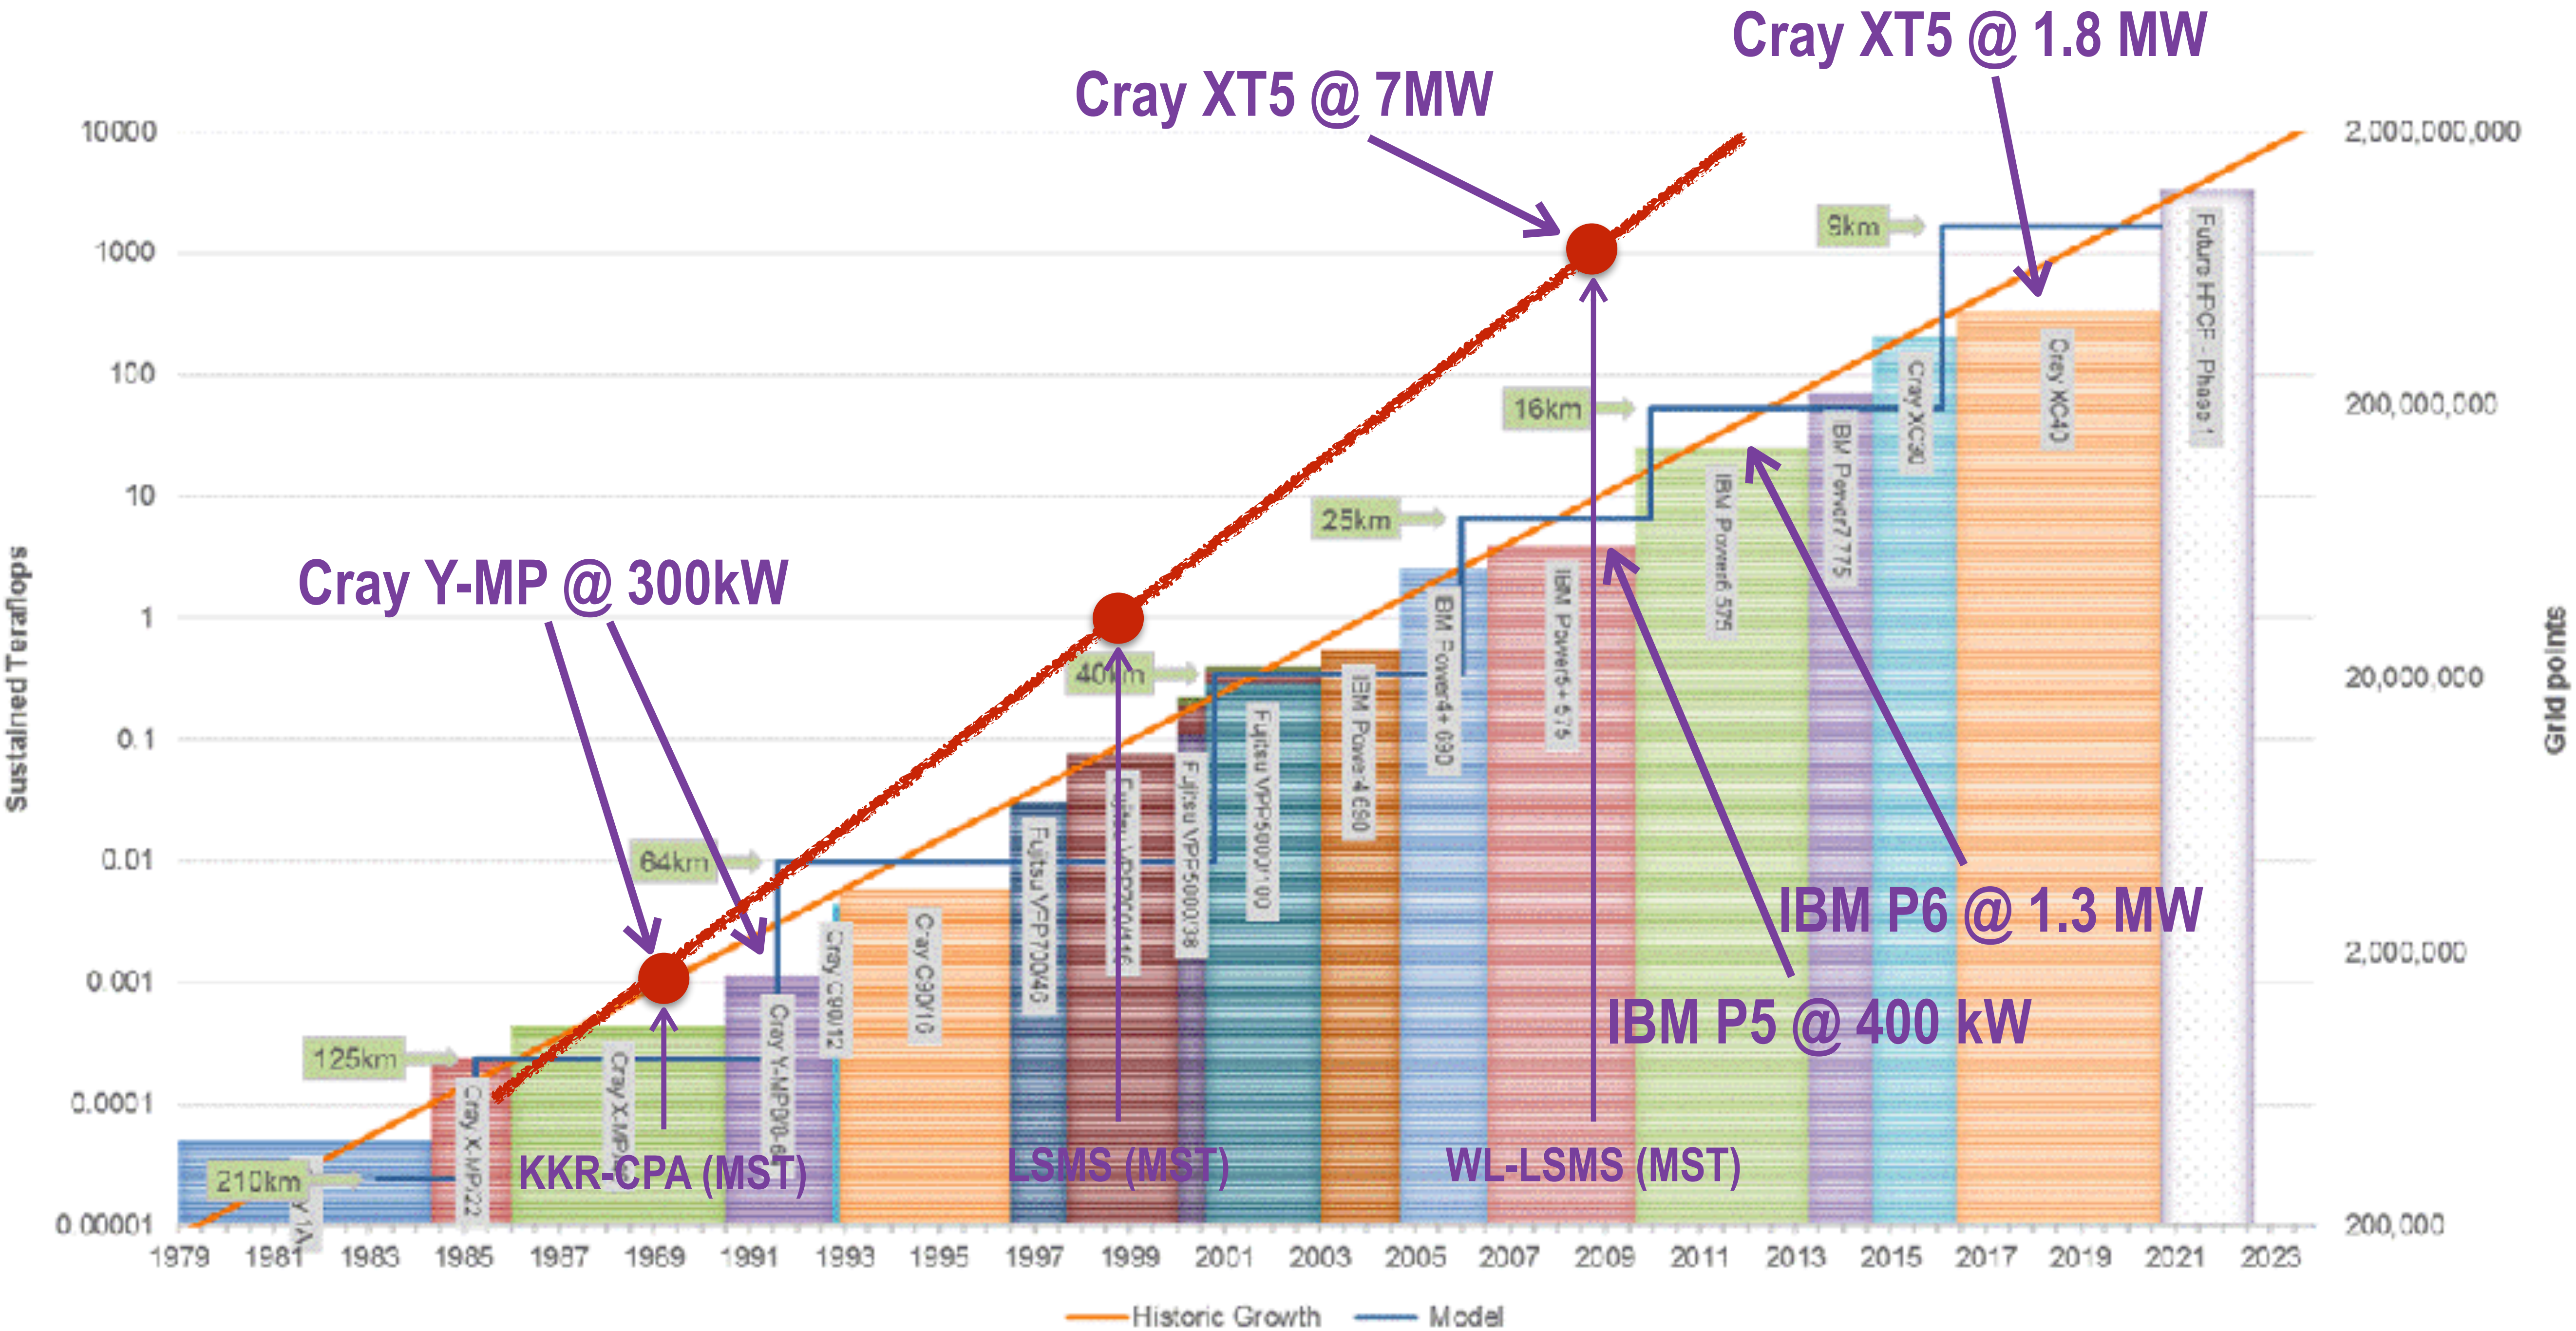

Source: Peter Bauer, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Floating points efficiency dropped from 50% on Cray Y-MP to 5% on today's Cray XC (10x in 2 decades)

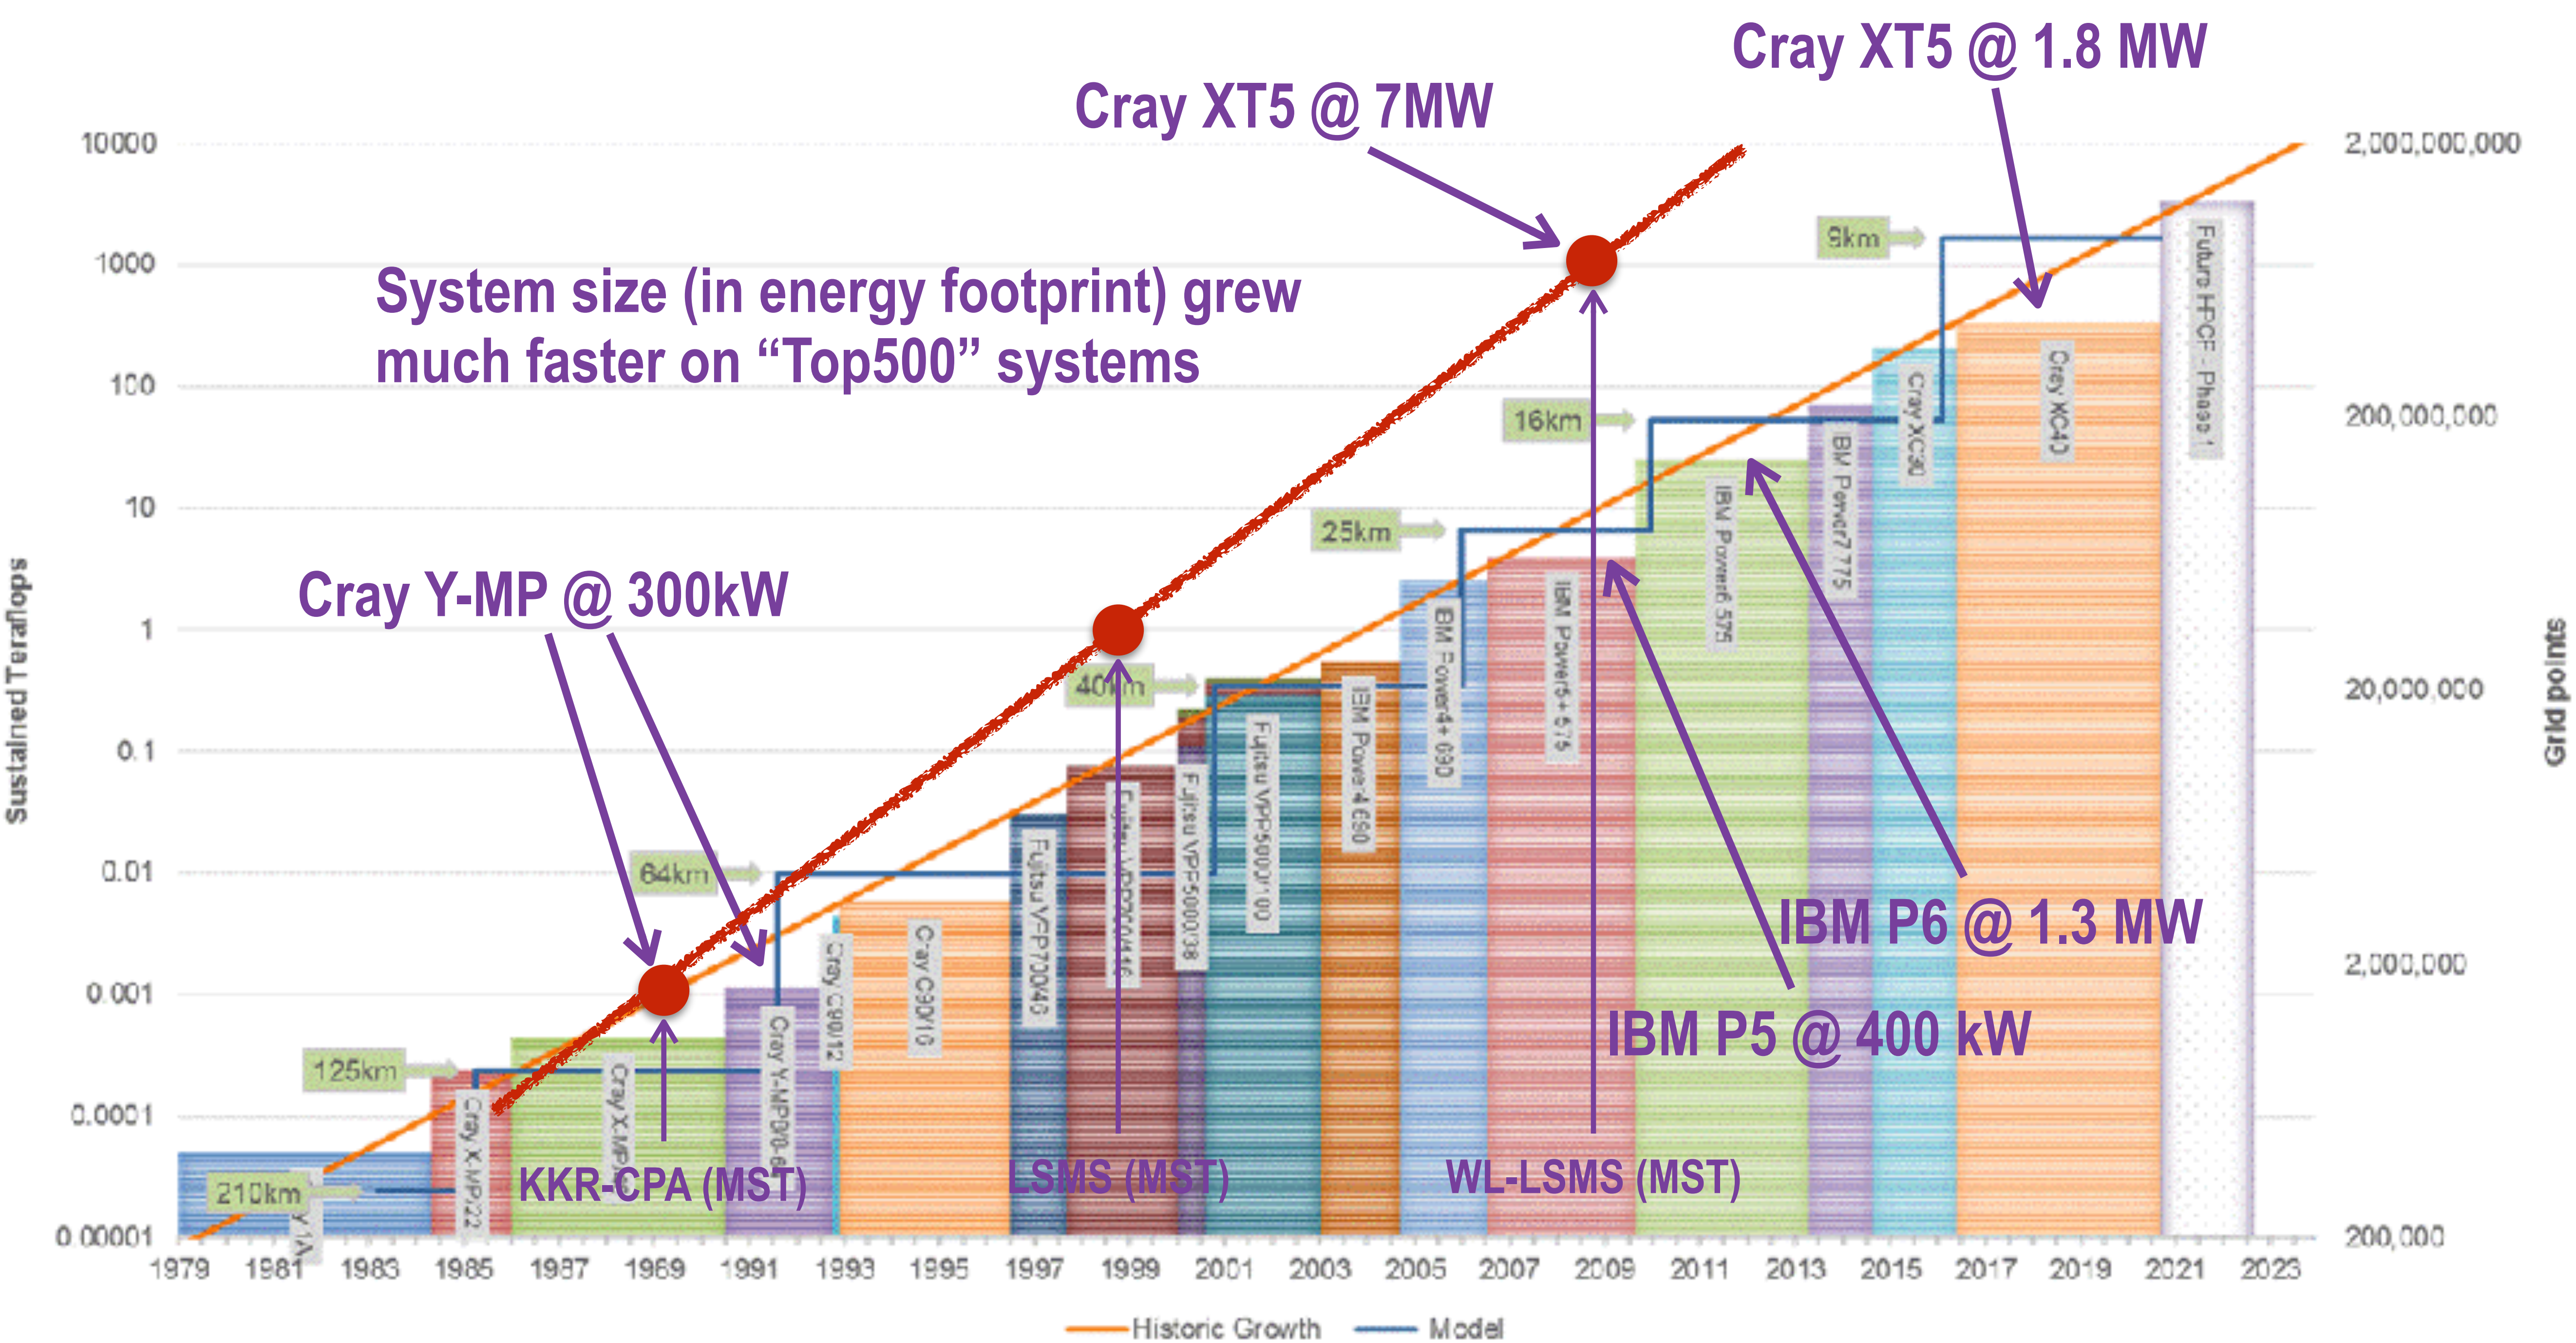

Source: Peter Bauer, ECMWF

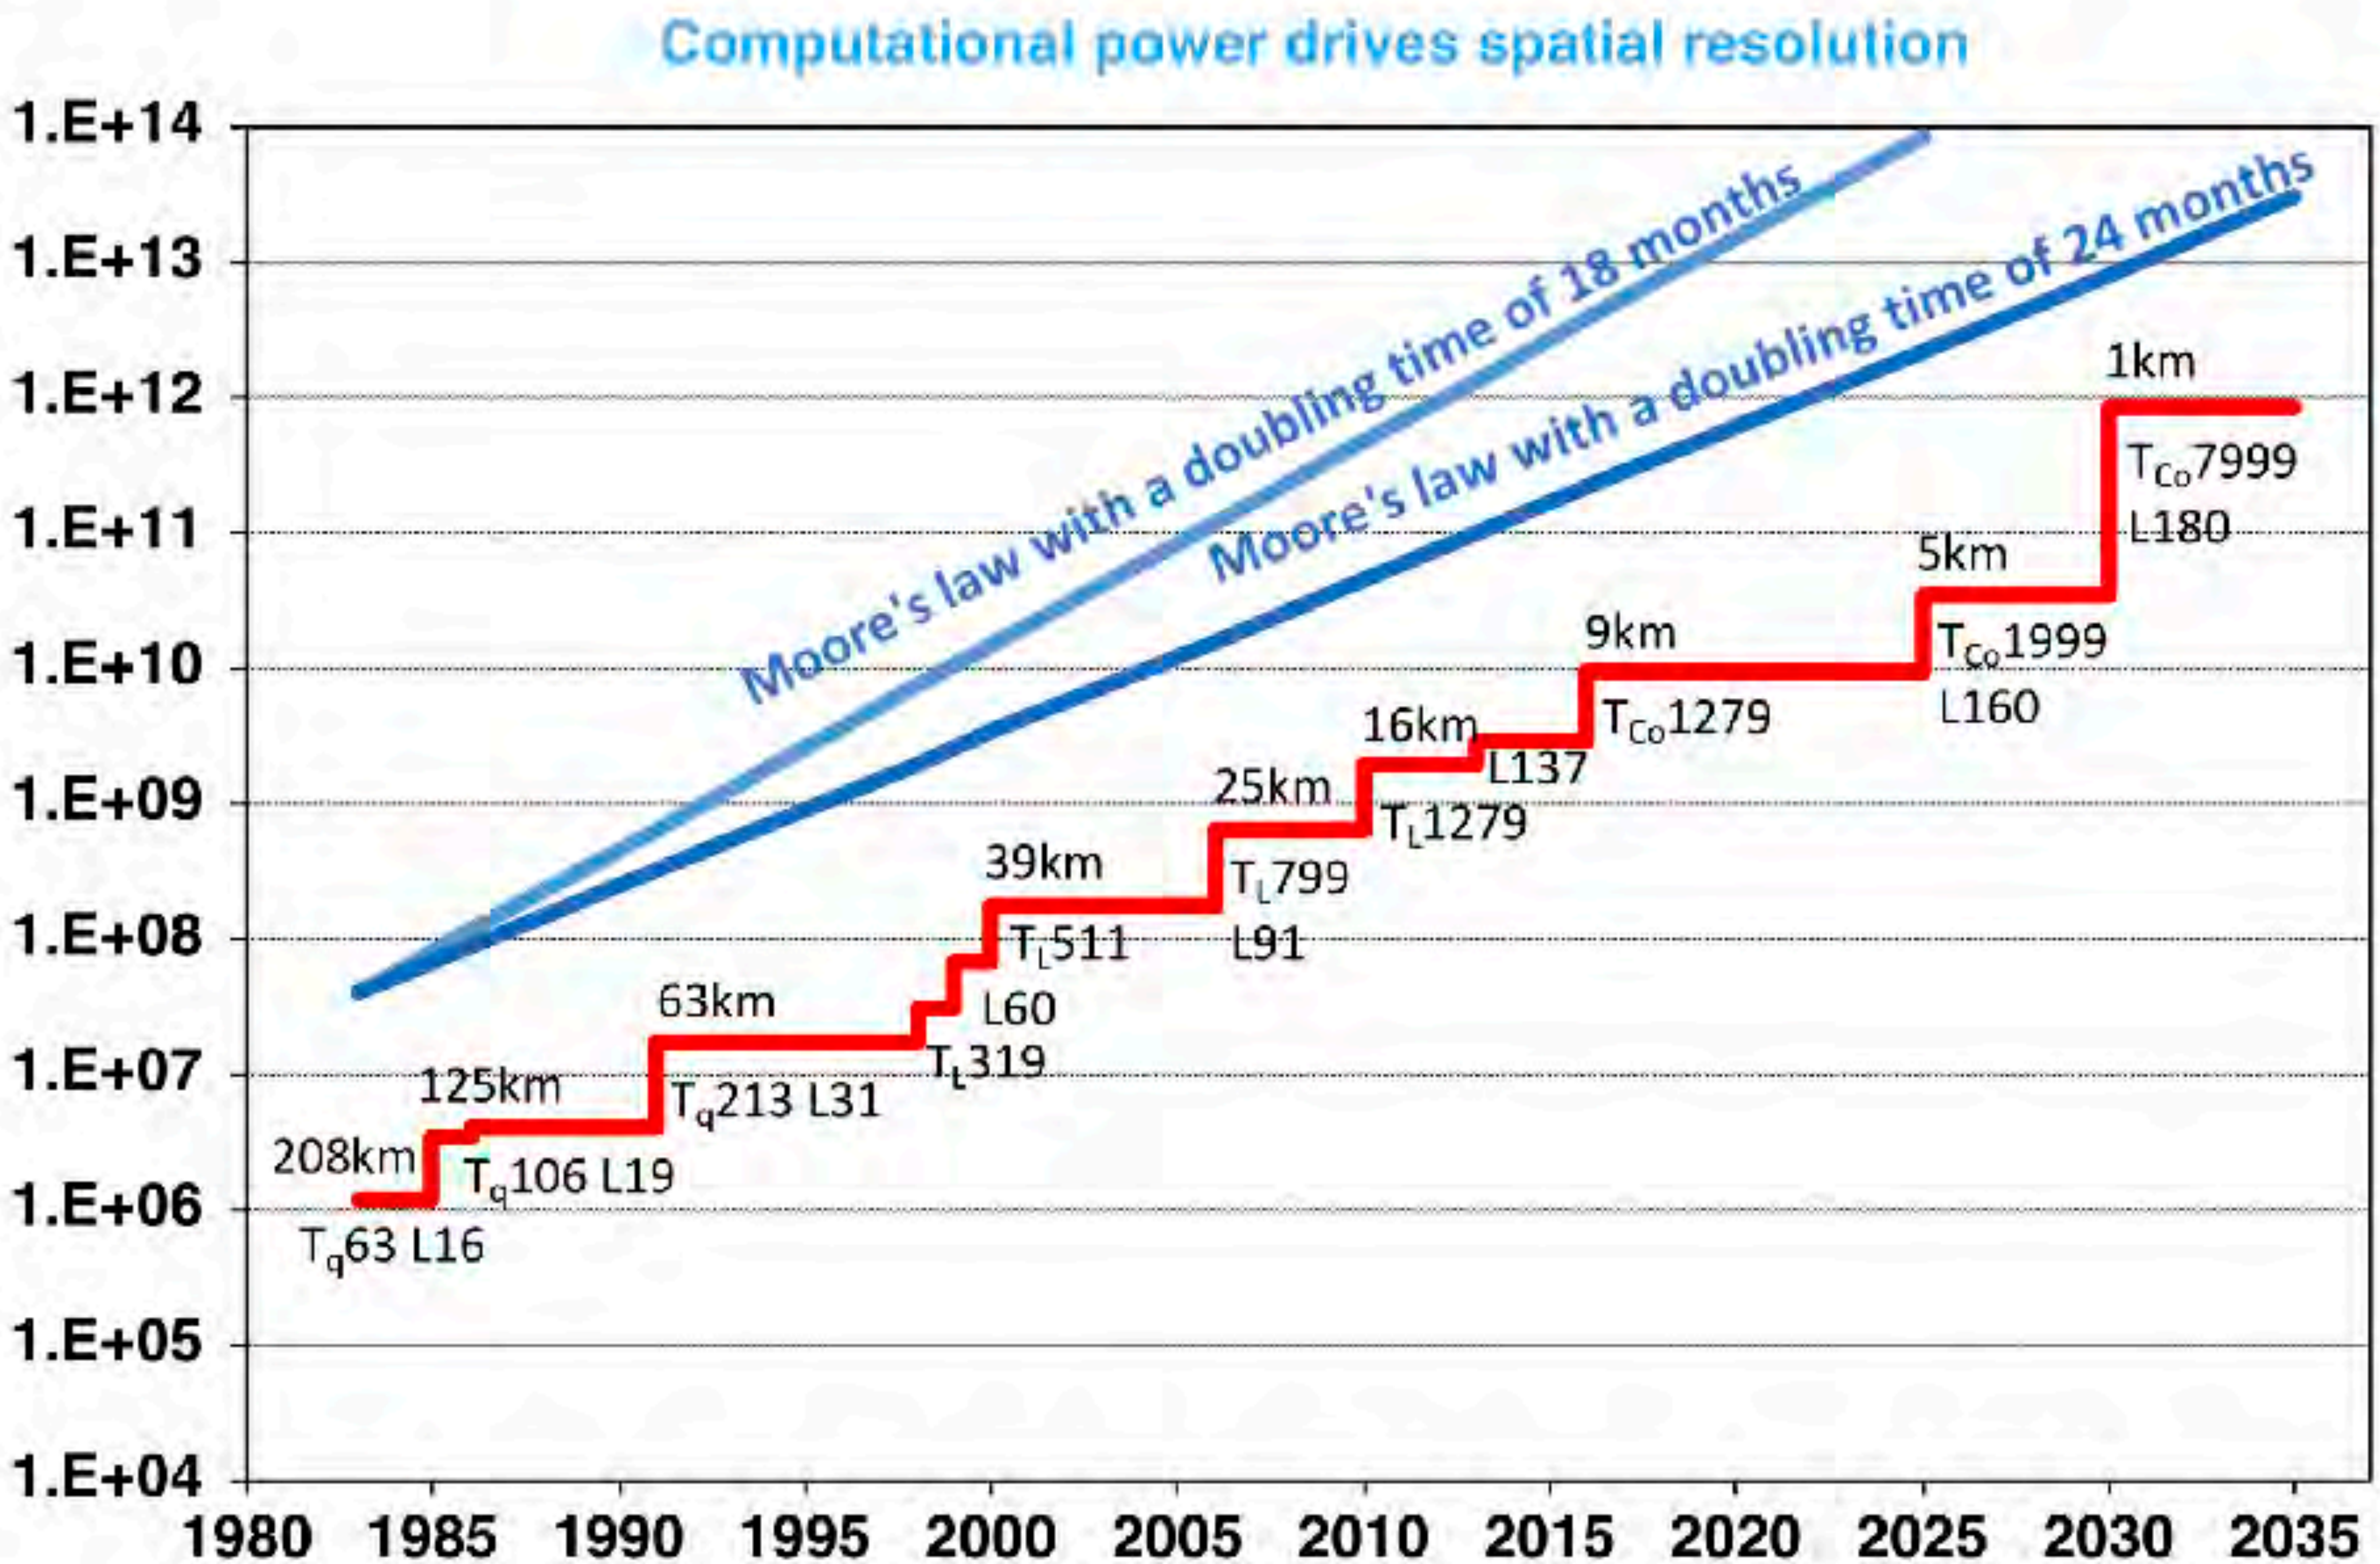

Source: Christoph Schär, ETH Zurich, & Nils Wedi, ECMWF

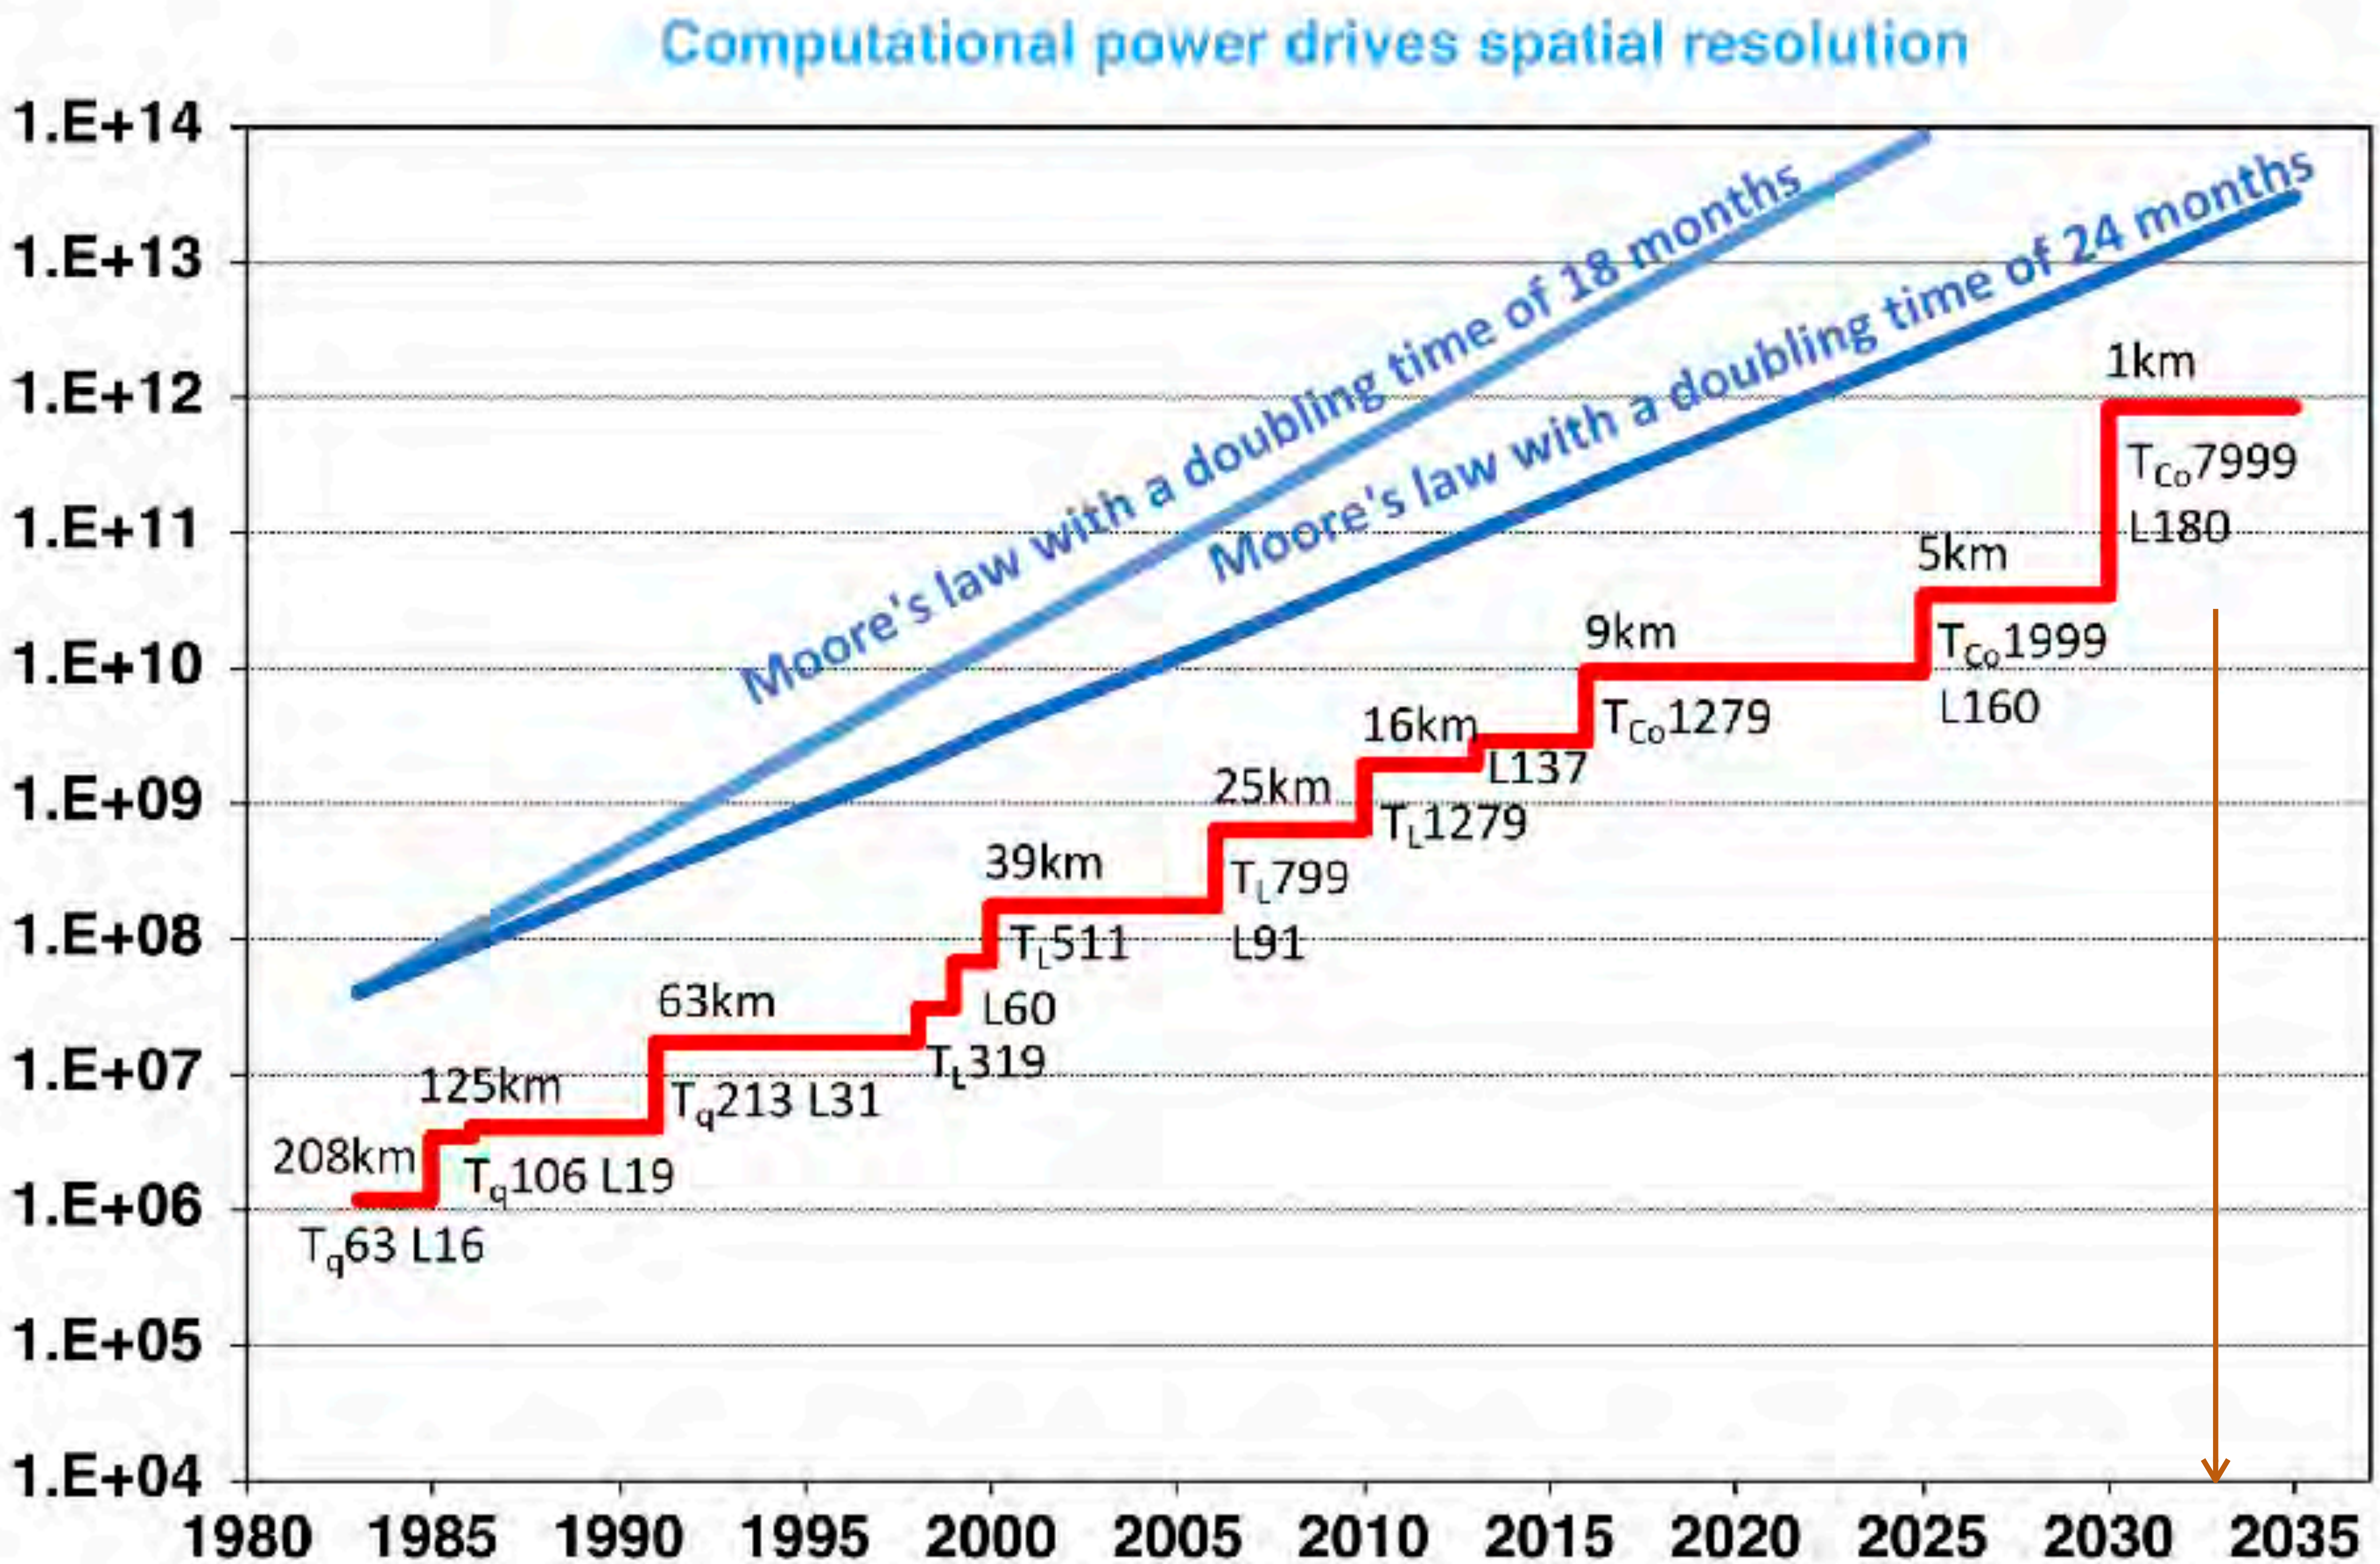

Source: Christoph Schär, ETH Zurich, & Nils Wedi, ECMWF

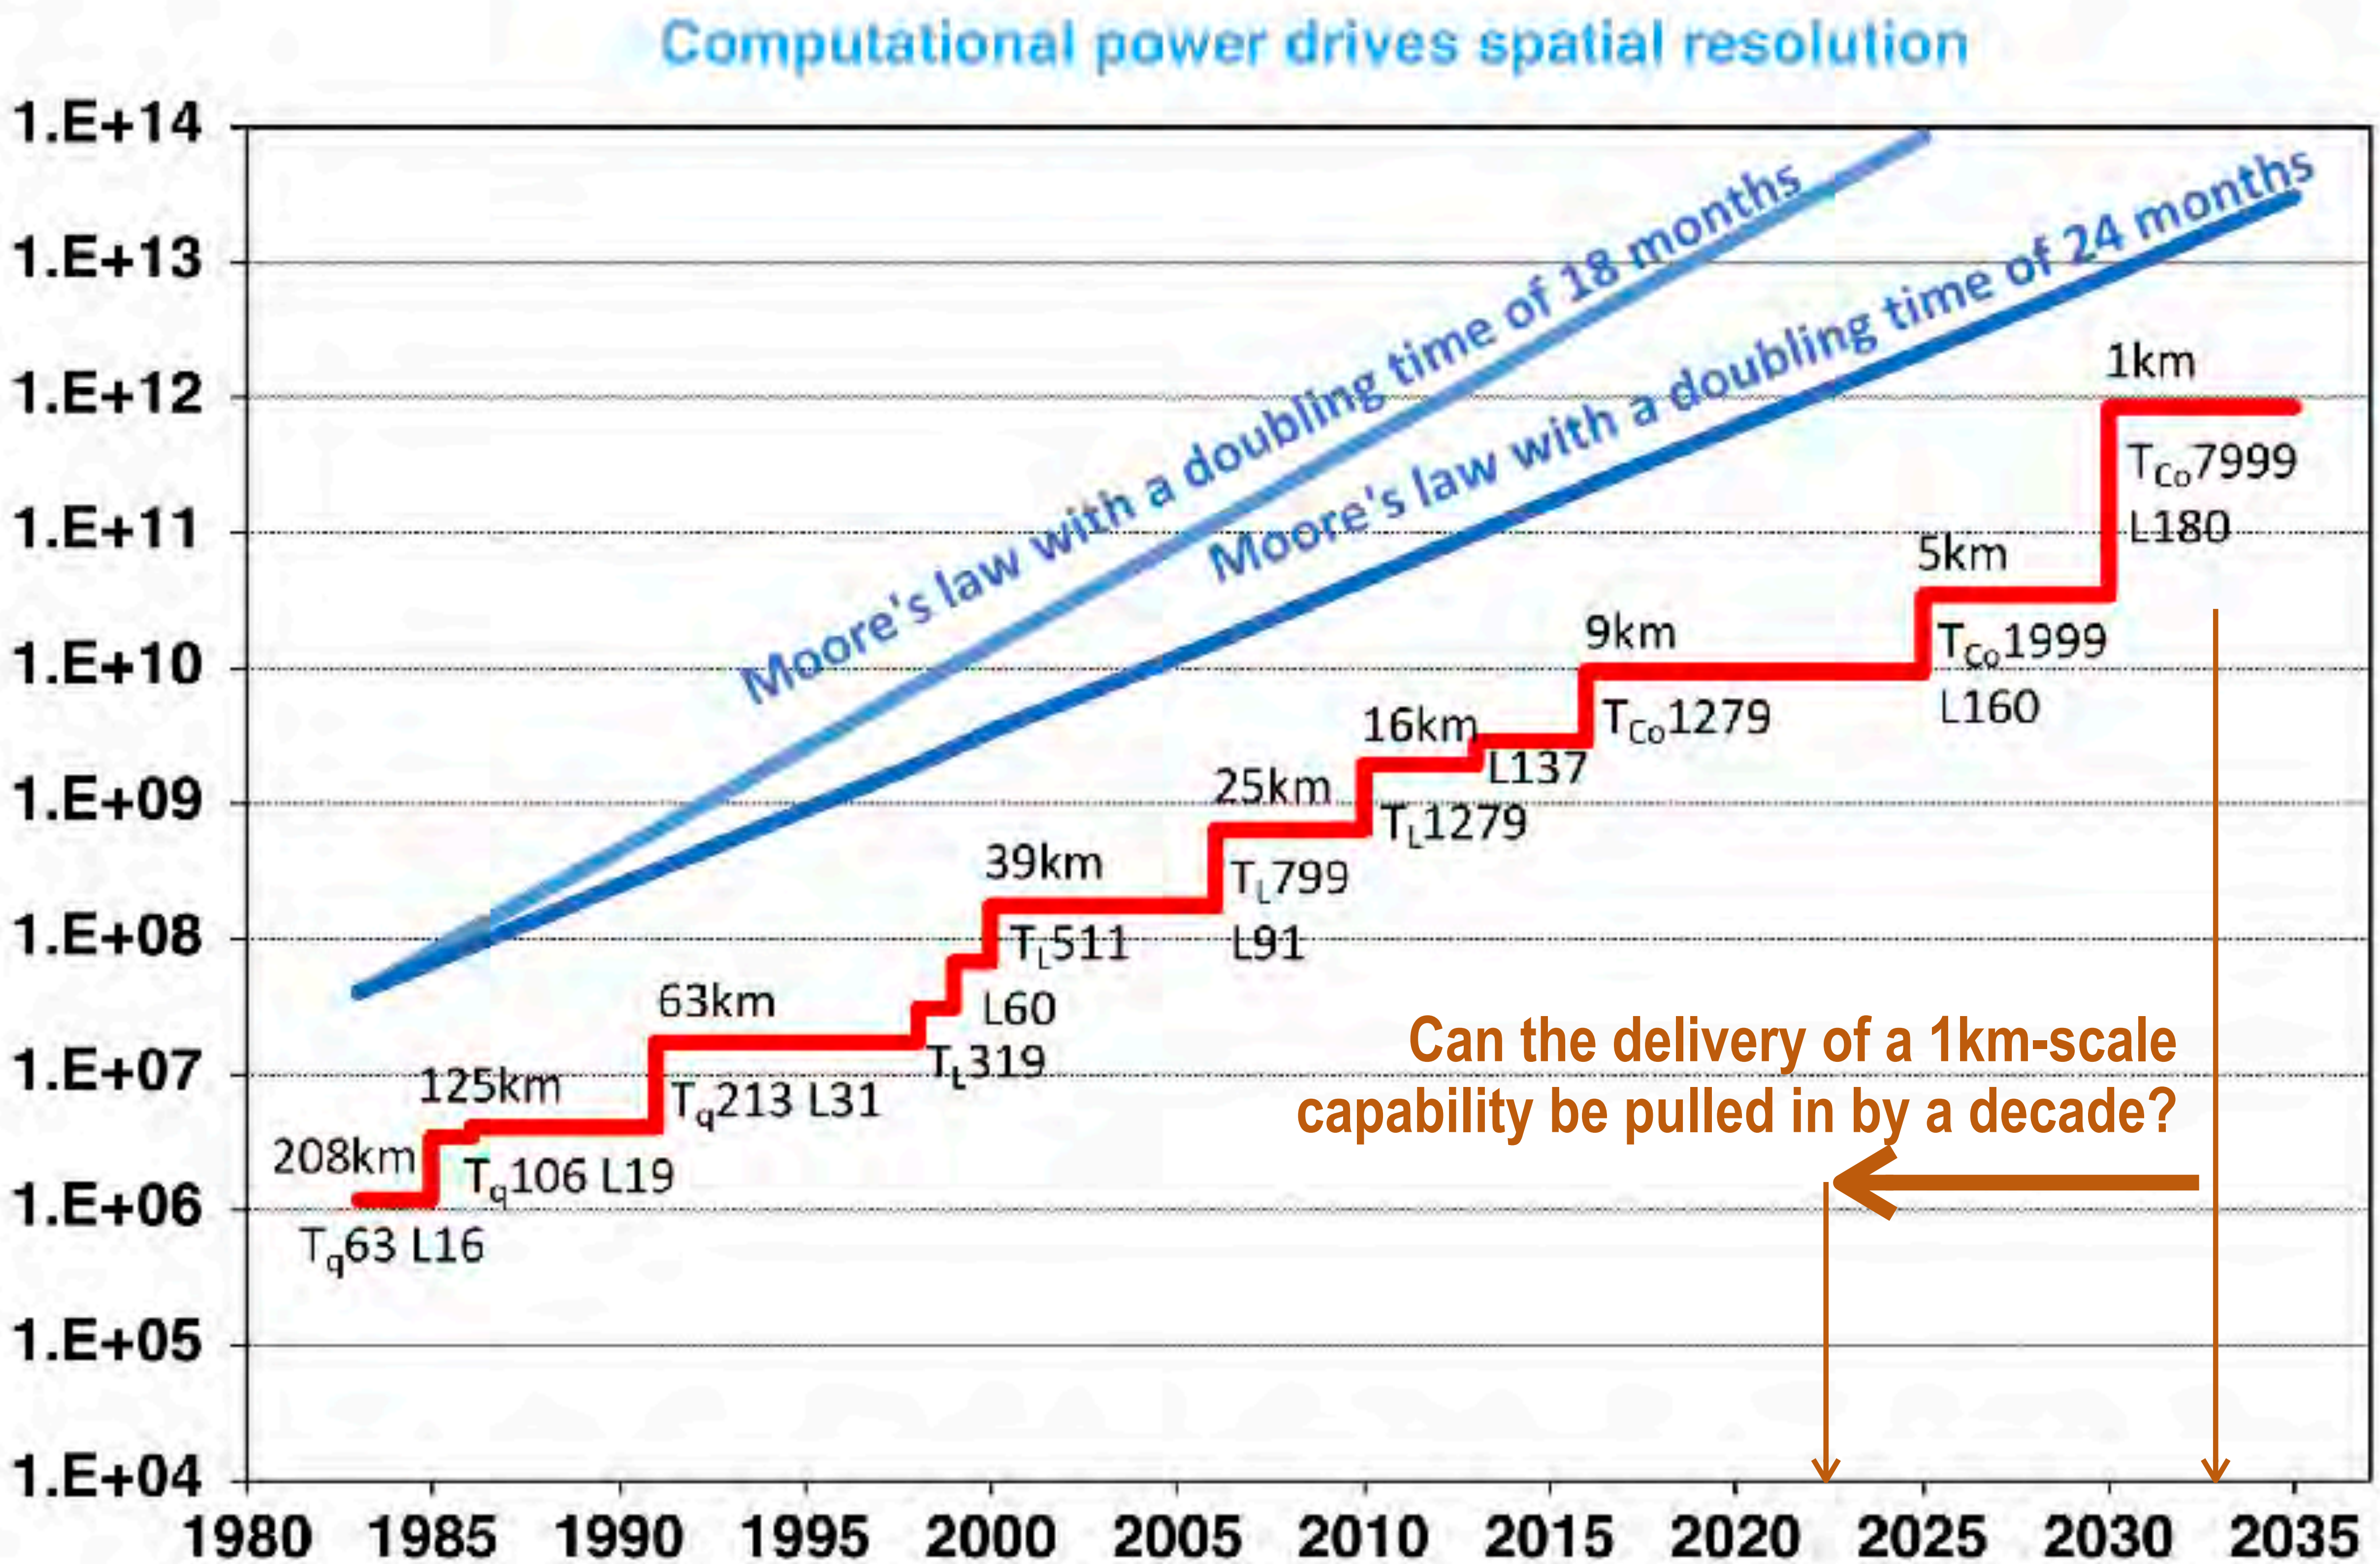

Source: Christoph Schär, ETH Zurich, & Nils Wedi, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

EPECC

Leadership in weather and climate

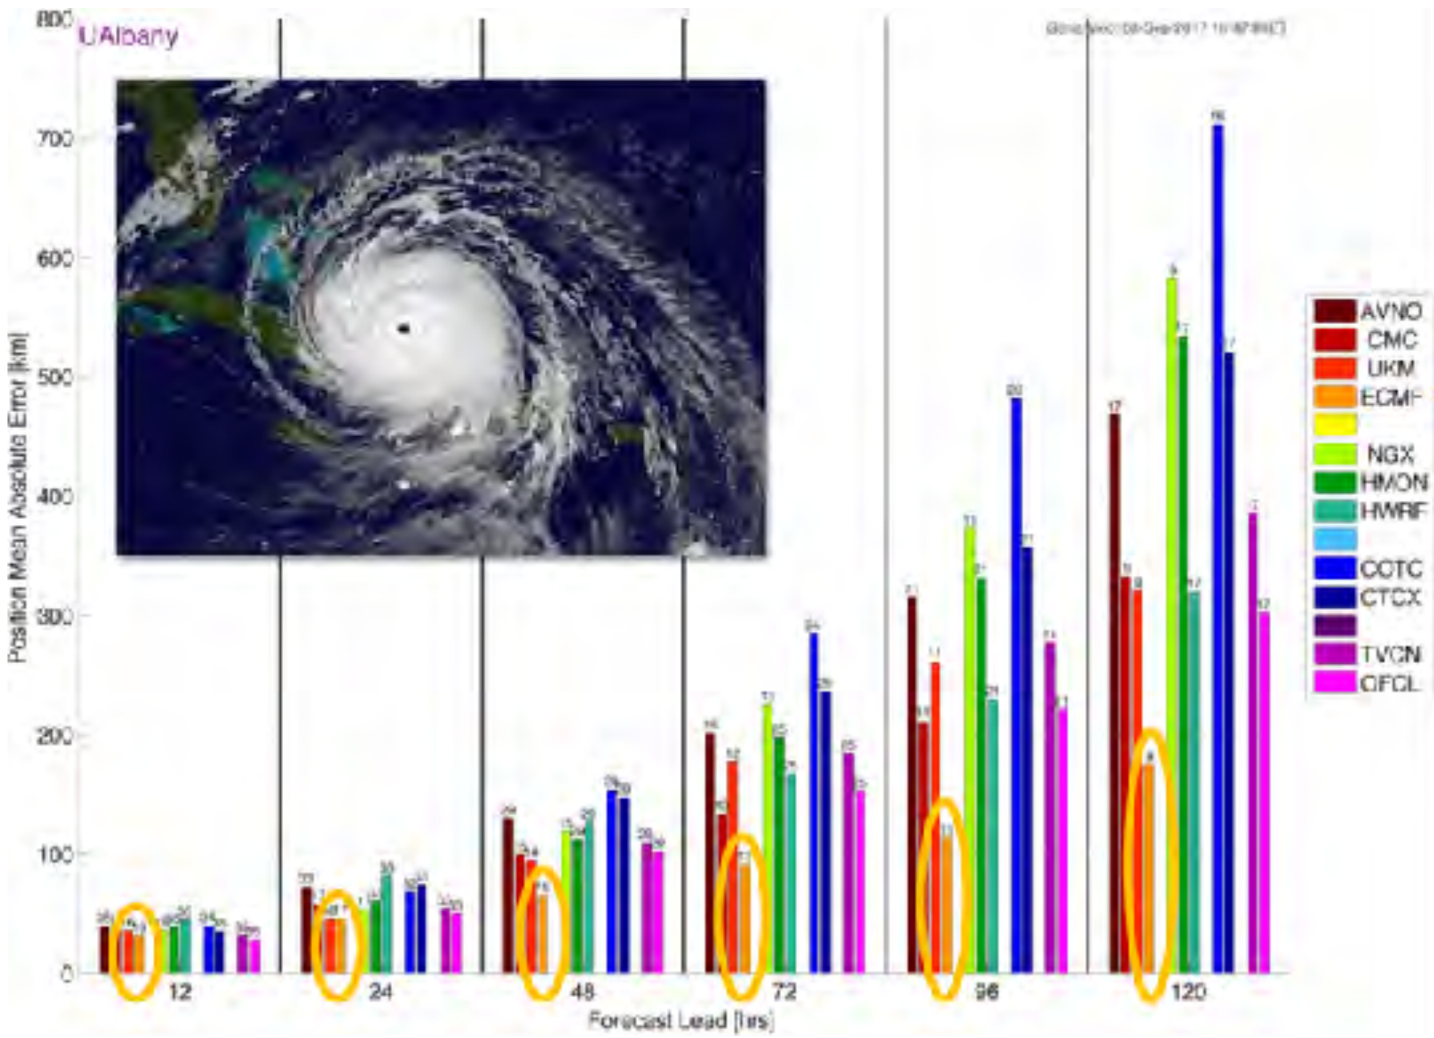

SCIENCE —

The European weather forecast model already kicking America's butt just improved

Better resolution will allow the world's best model to improve local forecasts.

ERIC BERGER (US) - 12/9/2016, 08:15

NATIONAL GEOGRAPHIC

Why Are Europeans Better at Predicting Weather?

ars TECHNICA

SCIENCE —

At times during Harvey, the European model outperformed humans

NOAA's new hurricane model, the HMON, performed terribly.

SCIENTIFIC AMERICAN

Are Europeans Better Than Americans at Forecasting Storms?

Peter Bauer, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

EPECC

Leadership in weather and climate

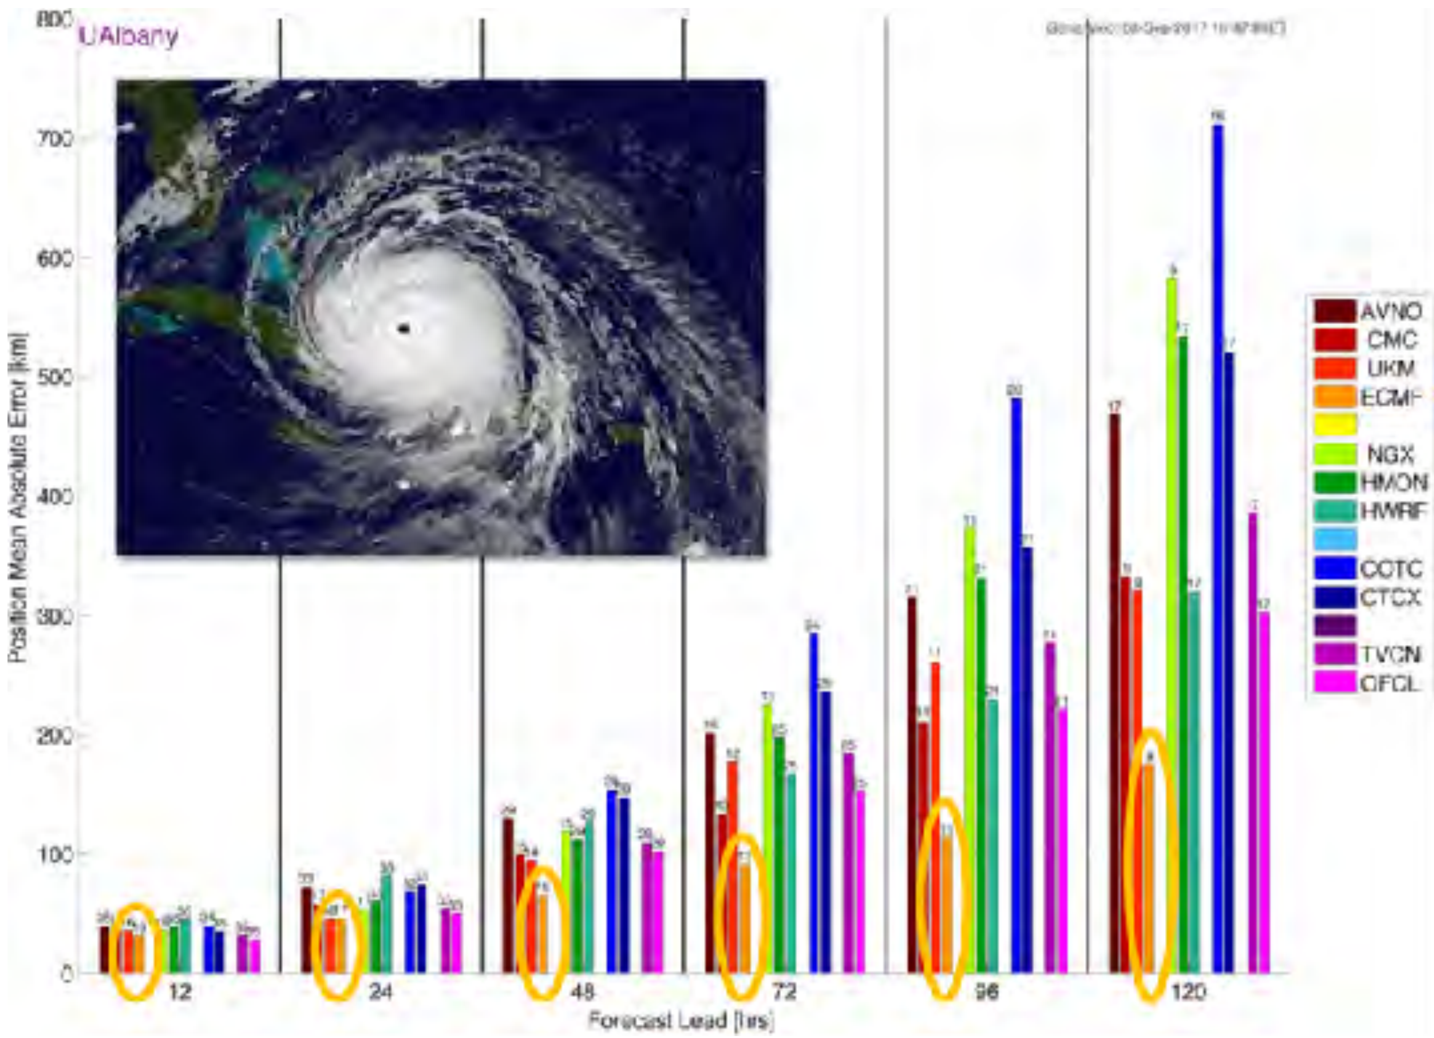

SCIENCE —

The European weather forecast model already kicking America's butt just improved

Better resolution will allow the world's best model to improve local forecasts.

ERIC BERGER (US) - 12/8/2016, 08:15

NATIONAL GEOGRAPHIC

Why Are Europeans Better at Predicting Weather?

ars TECHNICA

SCIENCE —

At times during Harvey, the European model outperformed humans

NOAA's new hurricane model, the HMON, performed terribly.

European model may be the best - but far away from sufficient accuracy and reliability!

SCIENTIFIC AMERICAN

Are Europeans Better Than Americans at Forecasting Storms?

Peter Bauer, ECMWF

The impact of resolution: simulated tropical cyclones

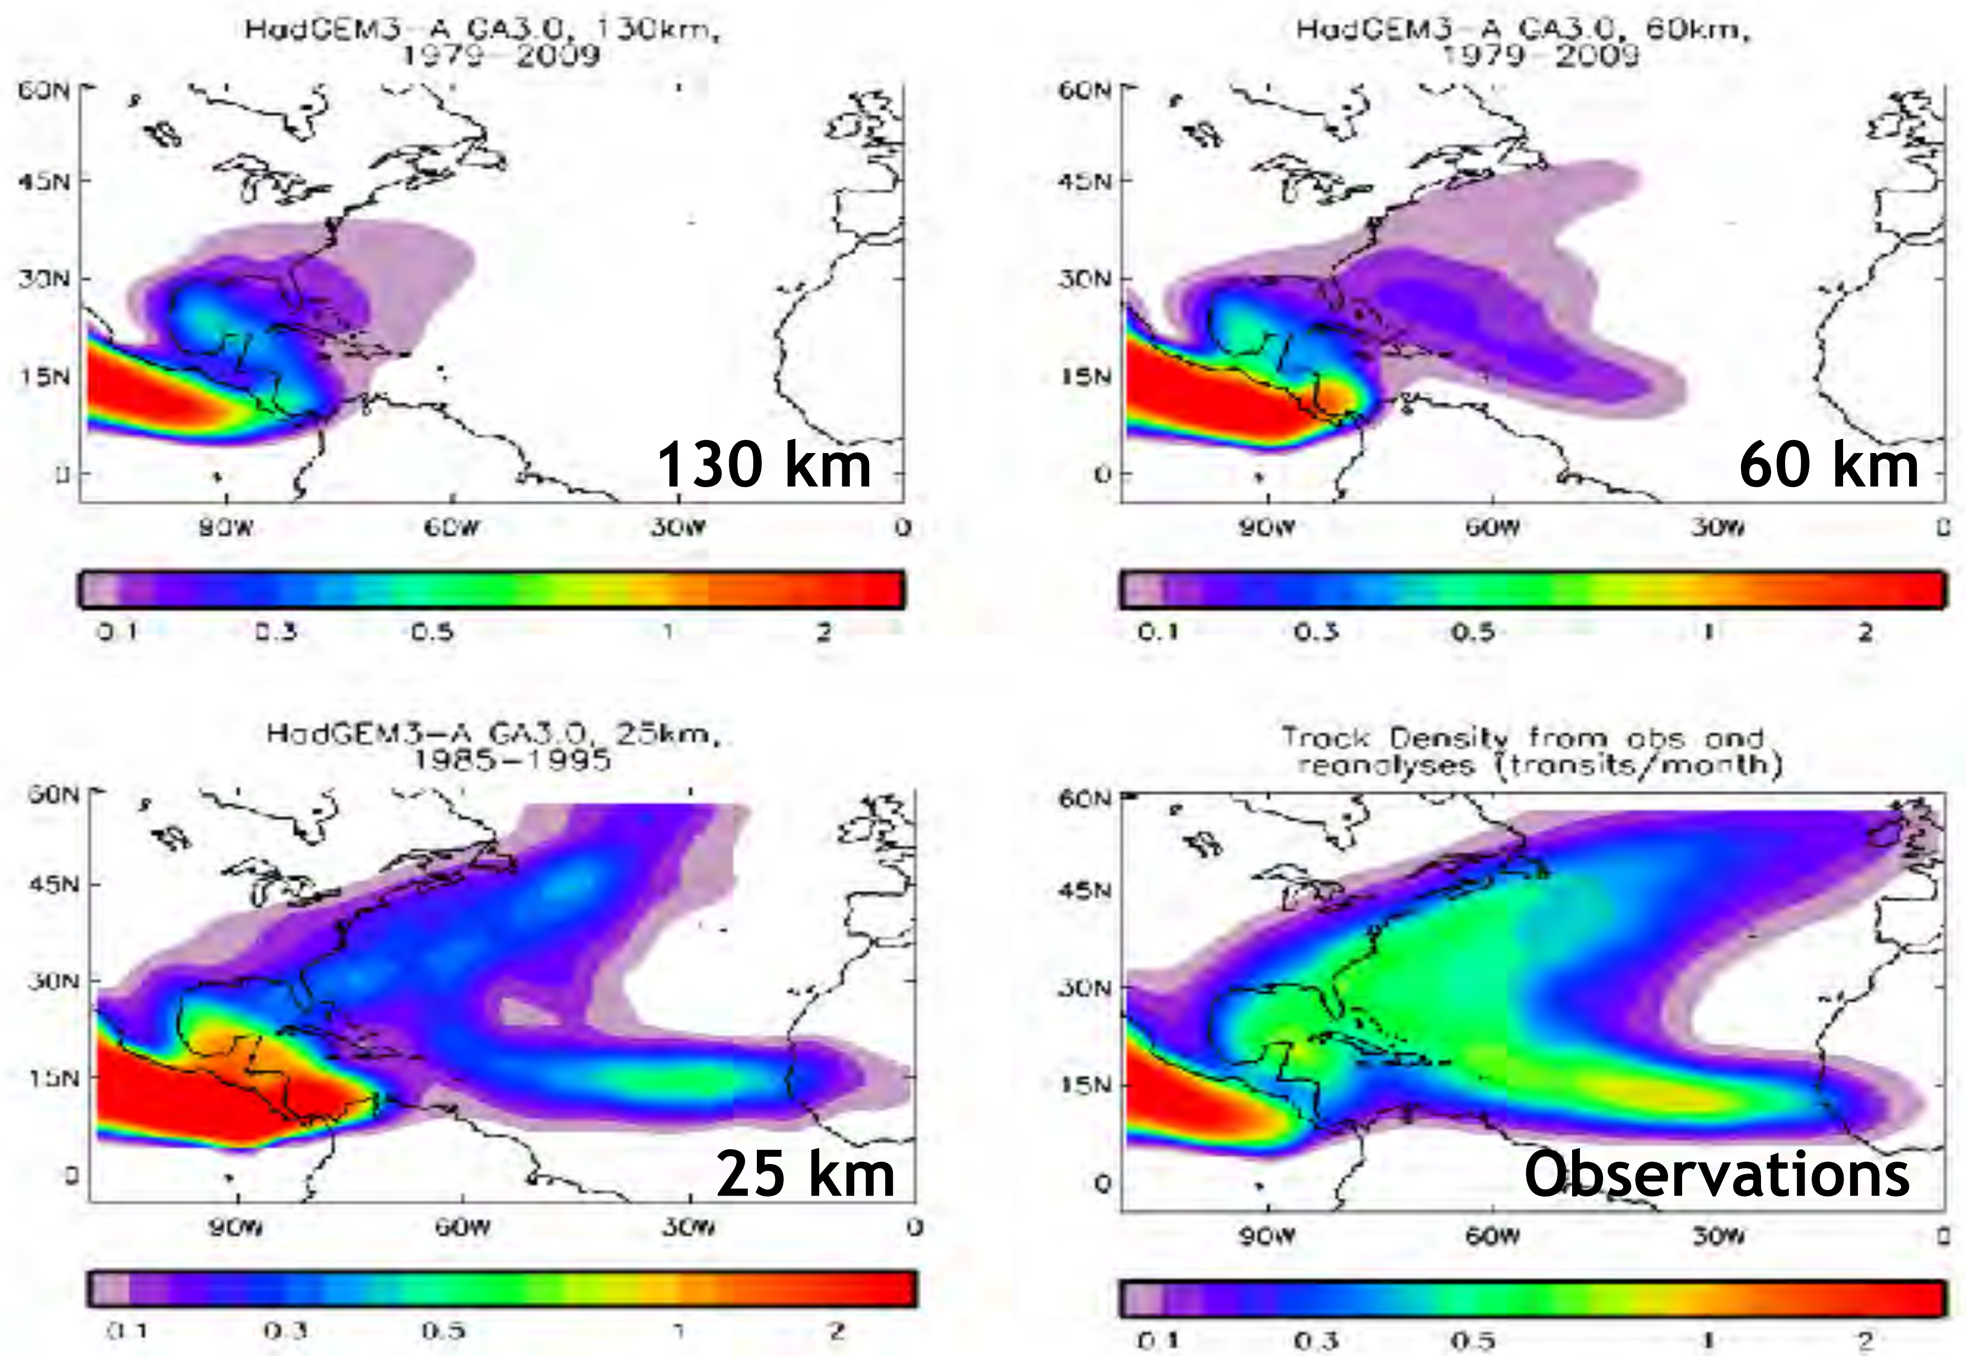

HADGEM3 PRACE UPSCALE, P.L. Vidale (NCAS) and M. Roberts (MO/HC)

Resolving convective clouds (convergence?)

Bulk convergence

Area-averaged bulk effects upon ambient flow:
E.g., heating and moistening of cloud layer

Structural convergence

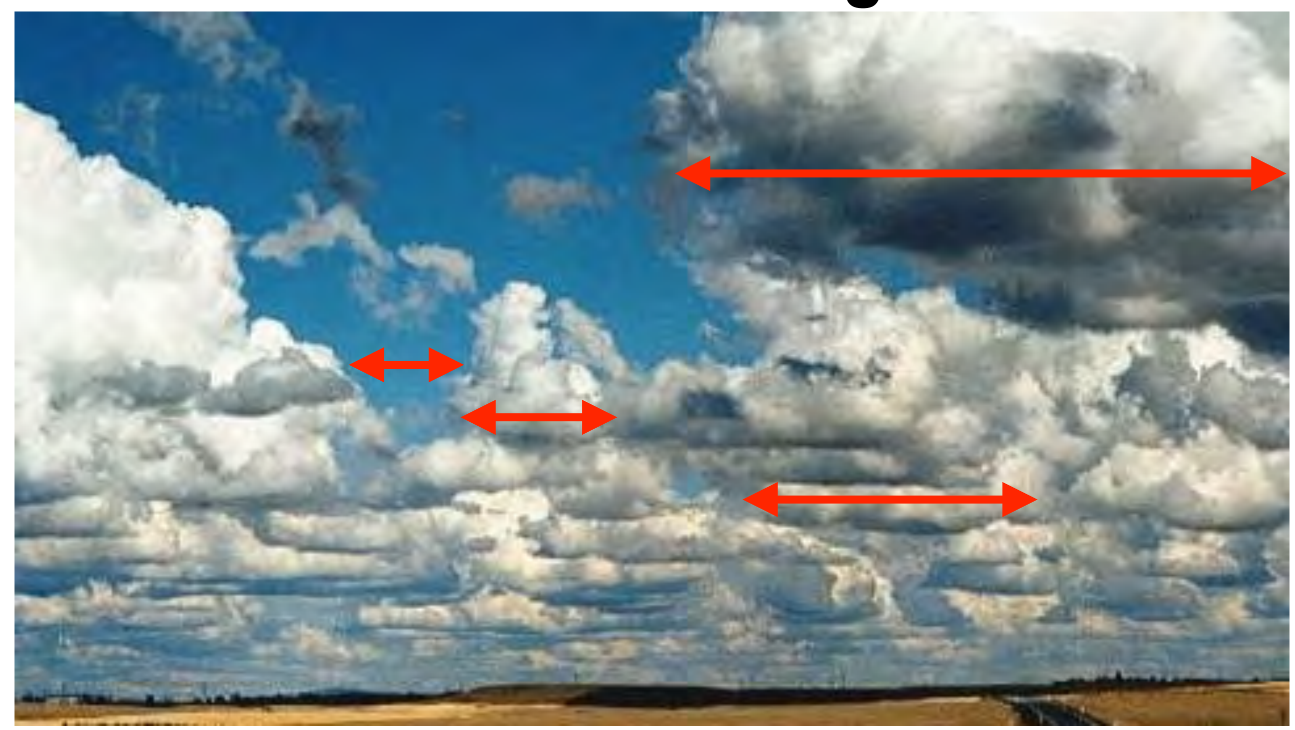

Statistics of cloud ensemble:
E.g., spacing and size of convective clouds

Source: Christoph Schär, ETH Zurich

Structural and bulk convergence

(Panosetti et al. 2018)

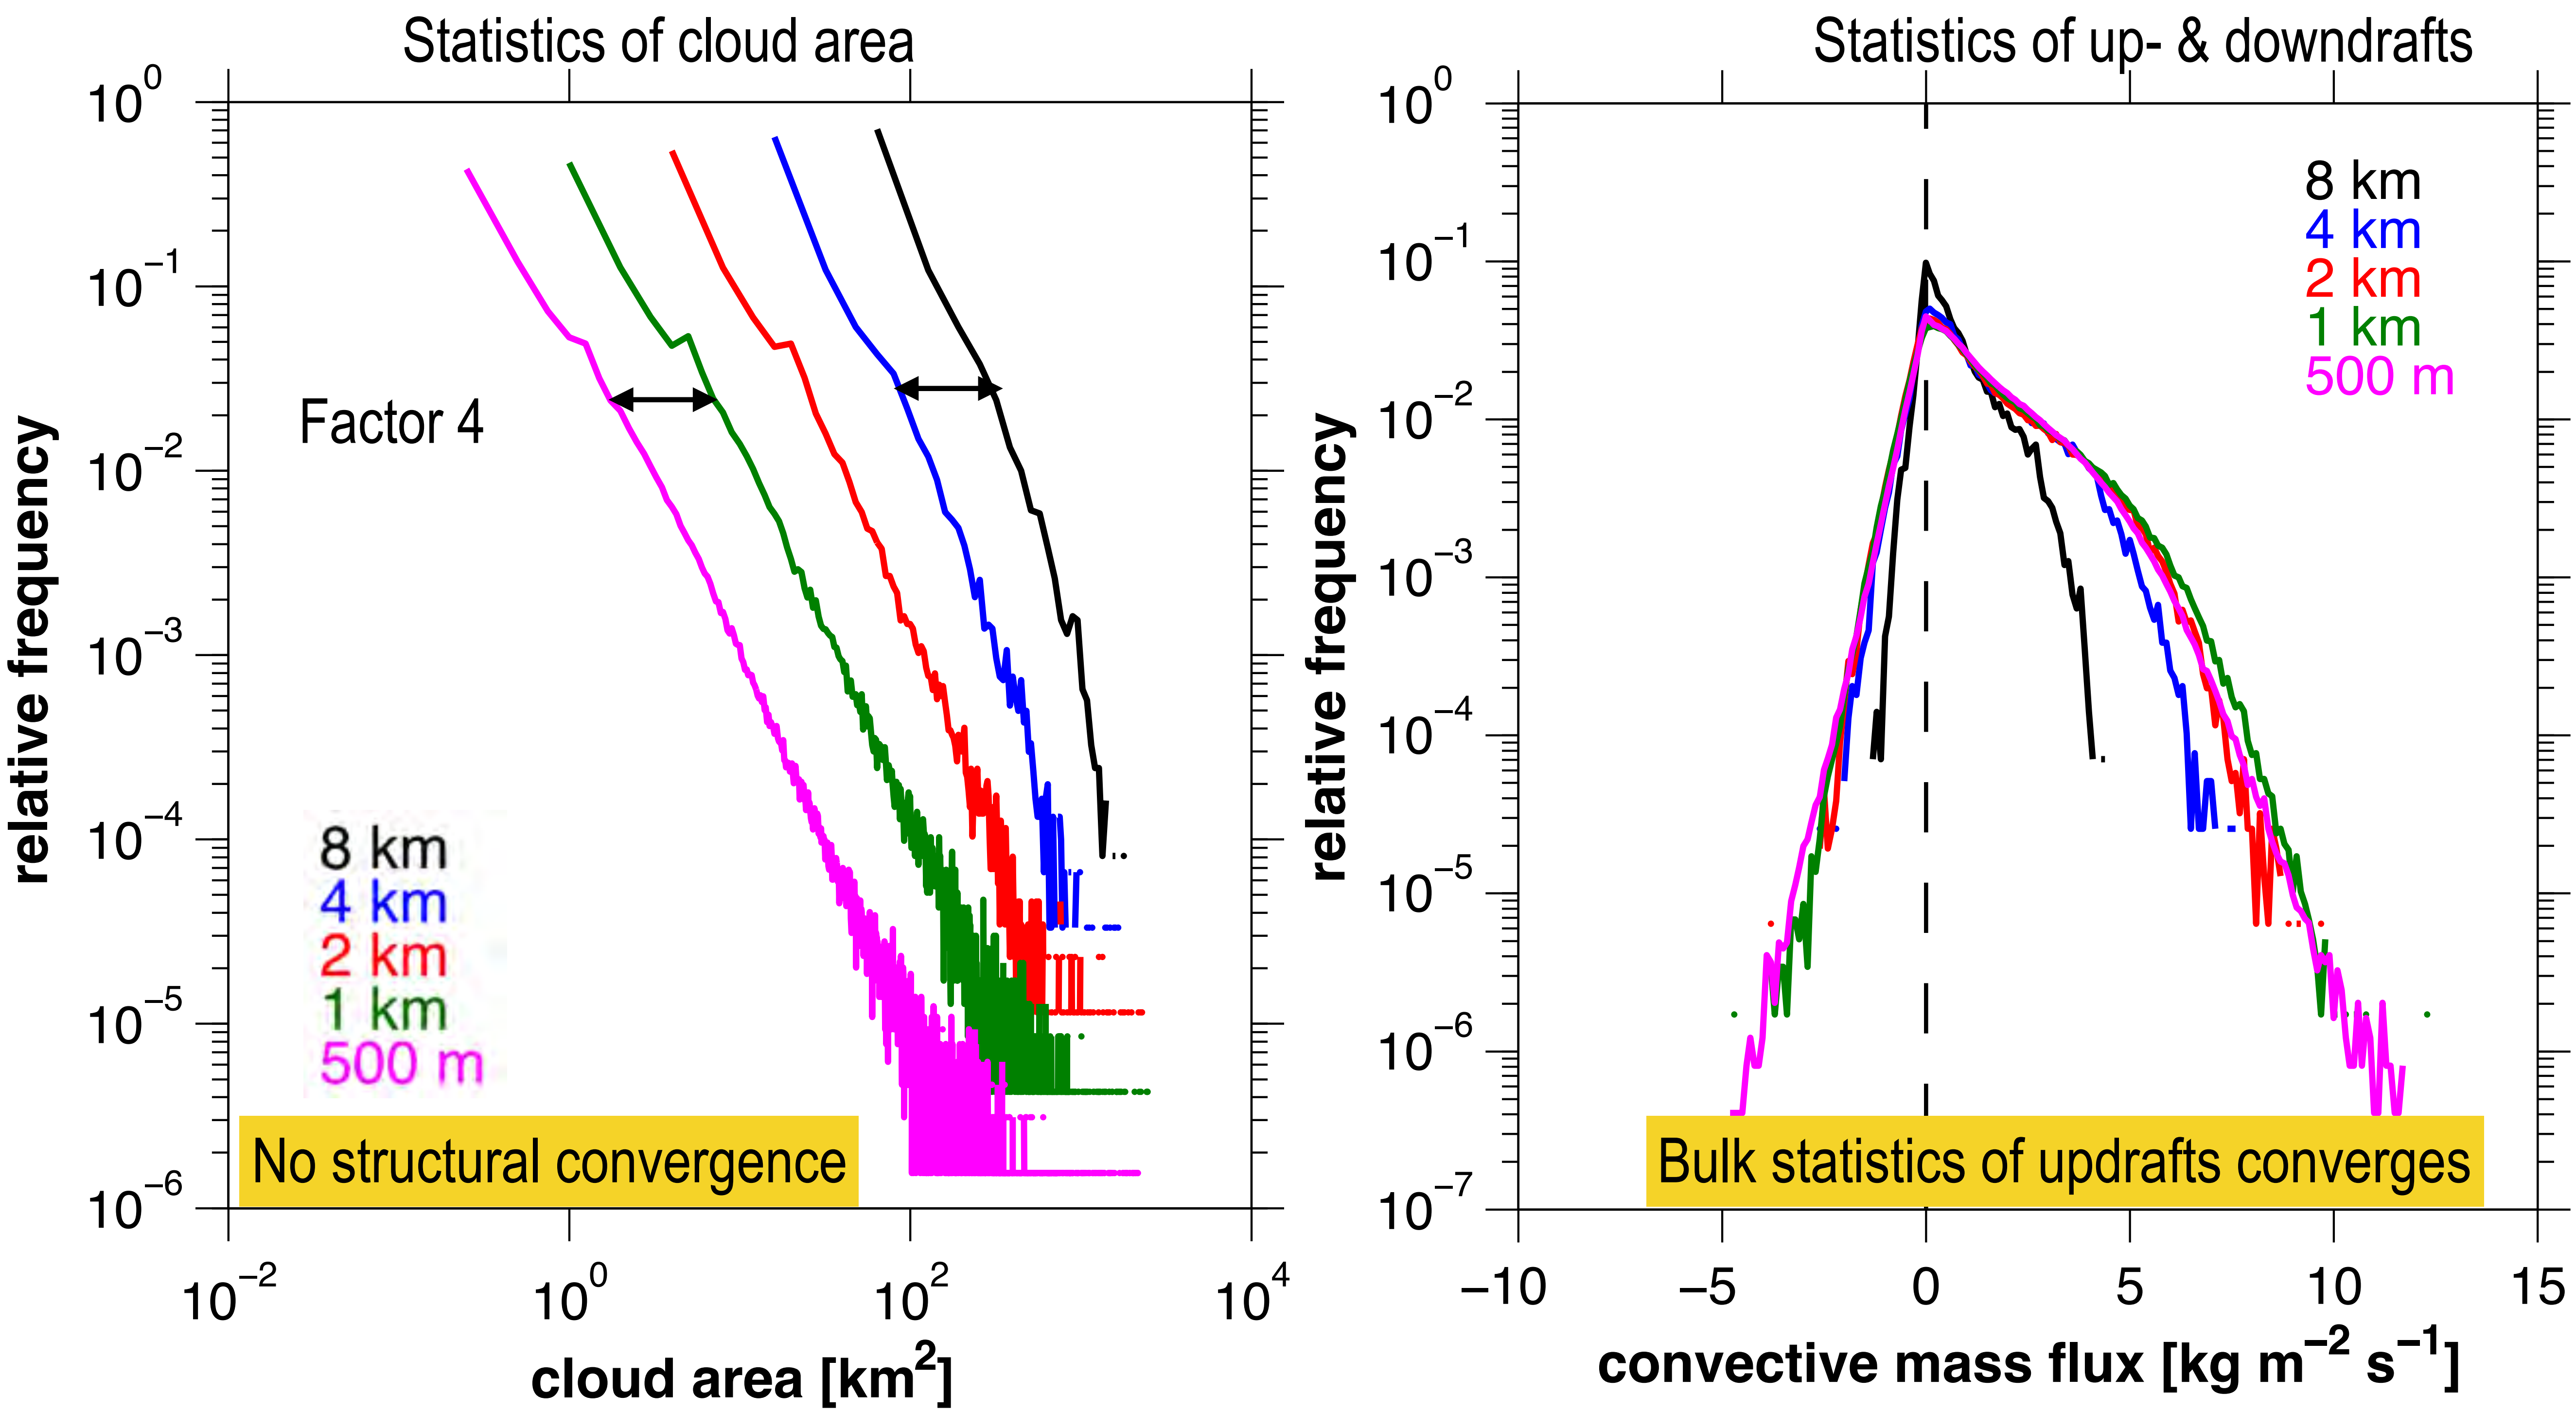

Source: Christoph Schär, ETH Zurich

ETHzürich

What resolution is needed?

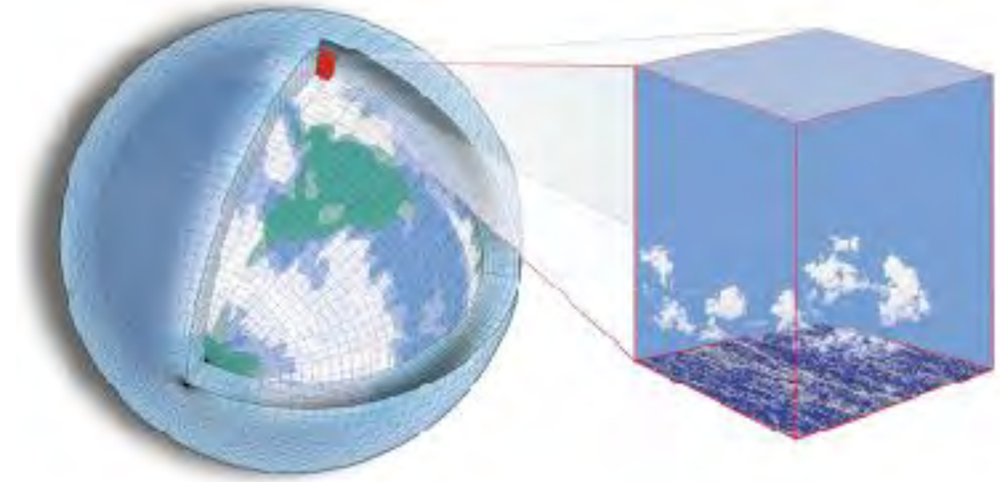

Bjorn Stevens, MPI-M

- **There are threshold scales in the atmosphere and ocean**: going from 100 km to 10 km is incremental, 10 km to 1 km is a leap. At 1km
  - it is no longer necessary to parametrise precipitating convection, ocean eddies, or orographic wave drag and its effect on extratropical storms;
  - ocean bathymetry, overflows and mixing, as well as regional orographic circulation in the atmosphere become resolved;
  - the connection between the remaining parametrisation are now on a physical footing.
- **We spend the last five decades in a paradigm of incremental advances.** Here we incrementally improved the resolution of models from 200 to 20km
- **Exascale allows us to make the leap to 1 km.** This fundamentally changes the structure of our models. We move from crude parametric presentations to an explicit, physics based, description of essential processes.
- **The last such step change was fifty years ago.** This was when, in the late 1960s, climate scientists first introduced global climate models, which were distinguished by their ability to explicitly represent extra-tropical storms, ocean gyres and boundary current.

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Simulation throughput: Simulate Years Per Day (SPYD)

	NWP	Climate in production	Climate spinup
Simulation	10 d	100 y	5’000 y
Desired wall clock time	0.1 d	0.1 y	0.5 y
ratio	100	1'000	10'000
SYPD	0.27	2.7	27

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Simulation throughput: Simulate Years Per Day (SPYD)

	NWP	Climate in production	Climate spinup
Simulation	10 d	100 y	5’000 y
Desired wall clock time	0.1 d	0.1 y	0.5 y
ratio	100	1'000	10'000
SYPD	0.27	2.7	27

Minimal throughout **1 SYPD**, preferred **5 SYPD**

Summary of intermediate goal (reach by 2021?)

Horizontal resolution	1 km (globally quasi-uniform)
Vertical resolution	180 levels (surface to ~100 km)
Time resolution	Less than 1 minute
Coupled	Land-surface/ocean/ocean-waves/sea-ice
Atmosphere	Non-hydrostatic
Precision	Single (32bit) or mixed precision
Compute rate	1 SYPD (simulated year wall-clock day)

Running COSMO 5.0 at global scale on Piz Daint

Scaling to full system size: ~5300 GPU accelerate nodes available

Running a near-global (±80º covering 97% of Earths surface) COSMO 5.0 simulation & IFS

> Either on the hosts processors: Intel Xeon E5 2690v3 (Haswell 12c).

> Or on the GPU accelerator: PCIe version ofNVIDIA GP100 (Pascal) GPU

September 15, 2015

Today's Outlook: GPU-accelerated Weather Forecasting

John Russell

MeteoSwiss New Weather Supercomputer

World's First GPU-Accelerated Weather Forecasting System

2x Racks

48 CPUs

192 Tesla K80 GPUs

> 90% of FLOPS from GPUs

Operational in 2016

5 NVIDIA

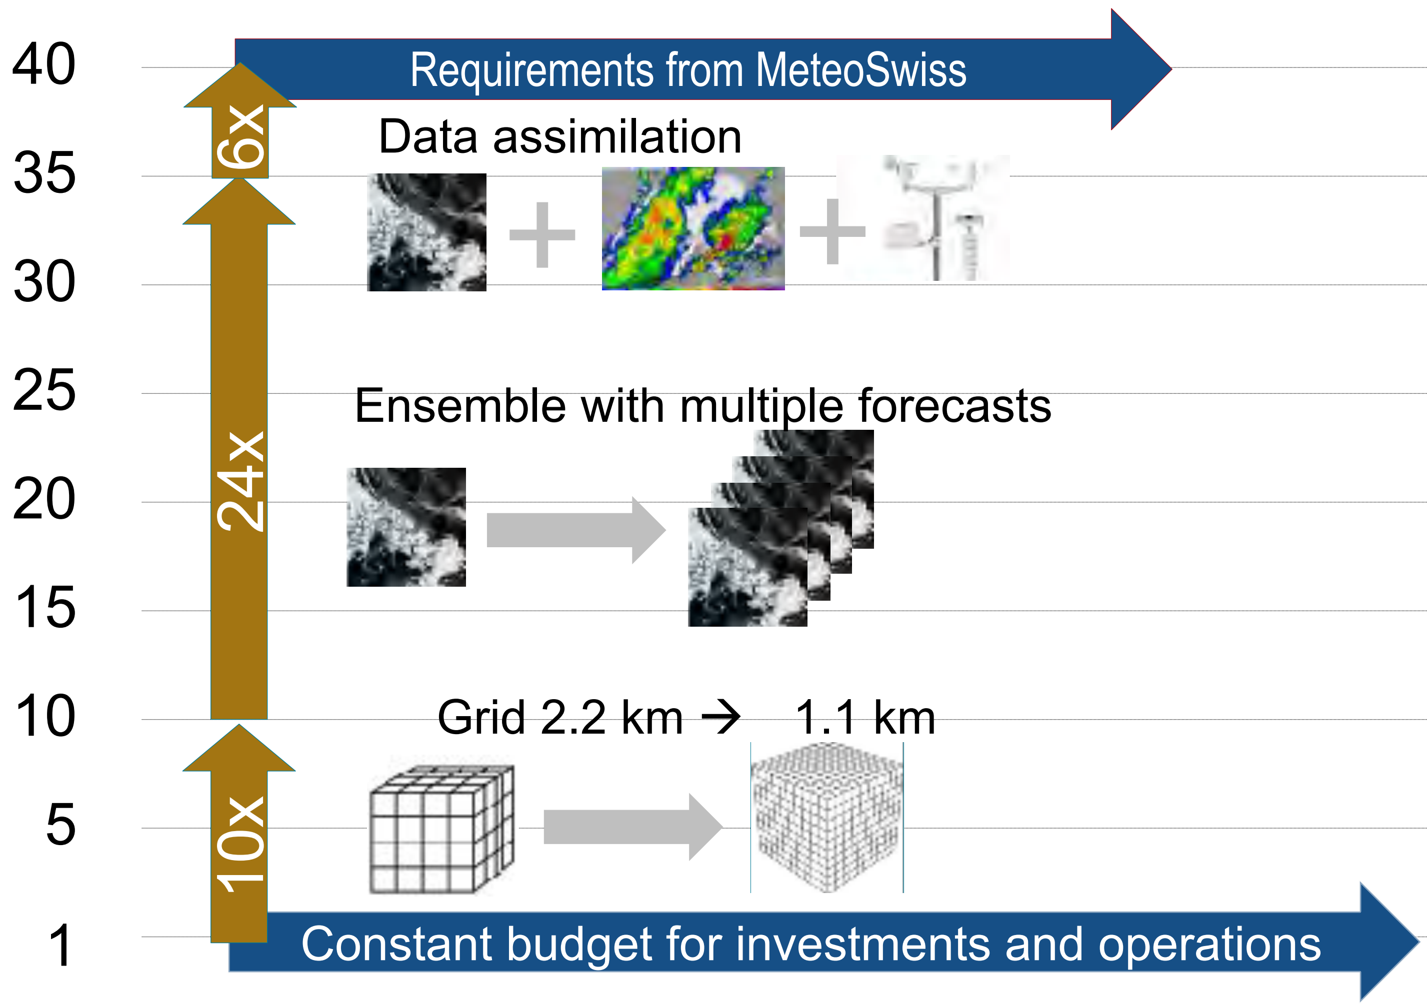
40
35
30
25
20
15
10
5
1
Requirements from MeteoSwiss
6x
24x
10x
Data assimilation
Ensemble with multiple forecasts
Grid 2.2 km → 1.1 km
Constant budget for investments and operations

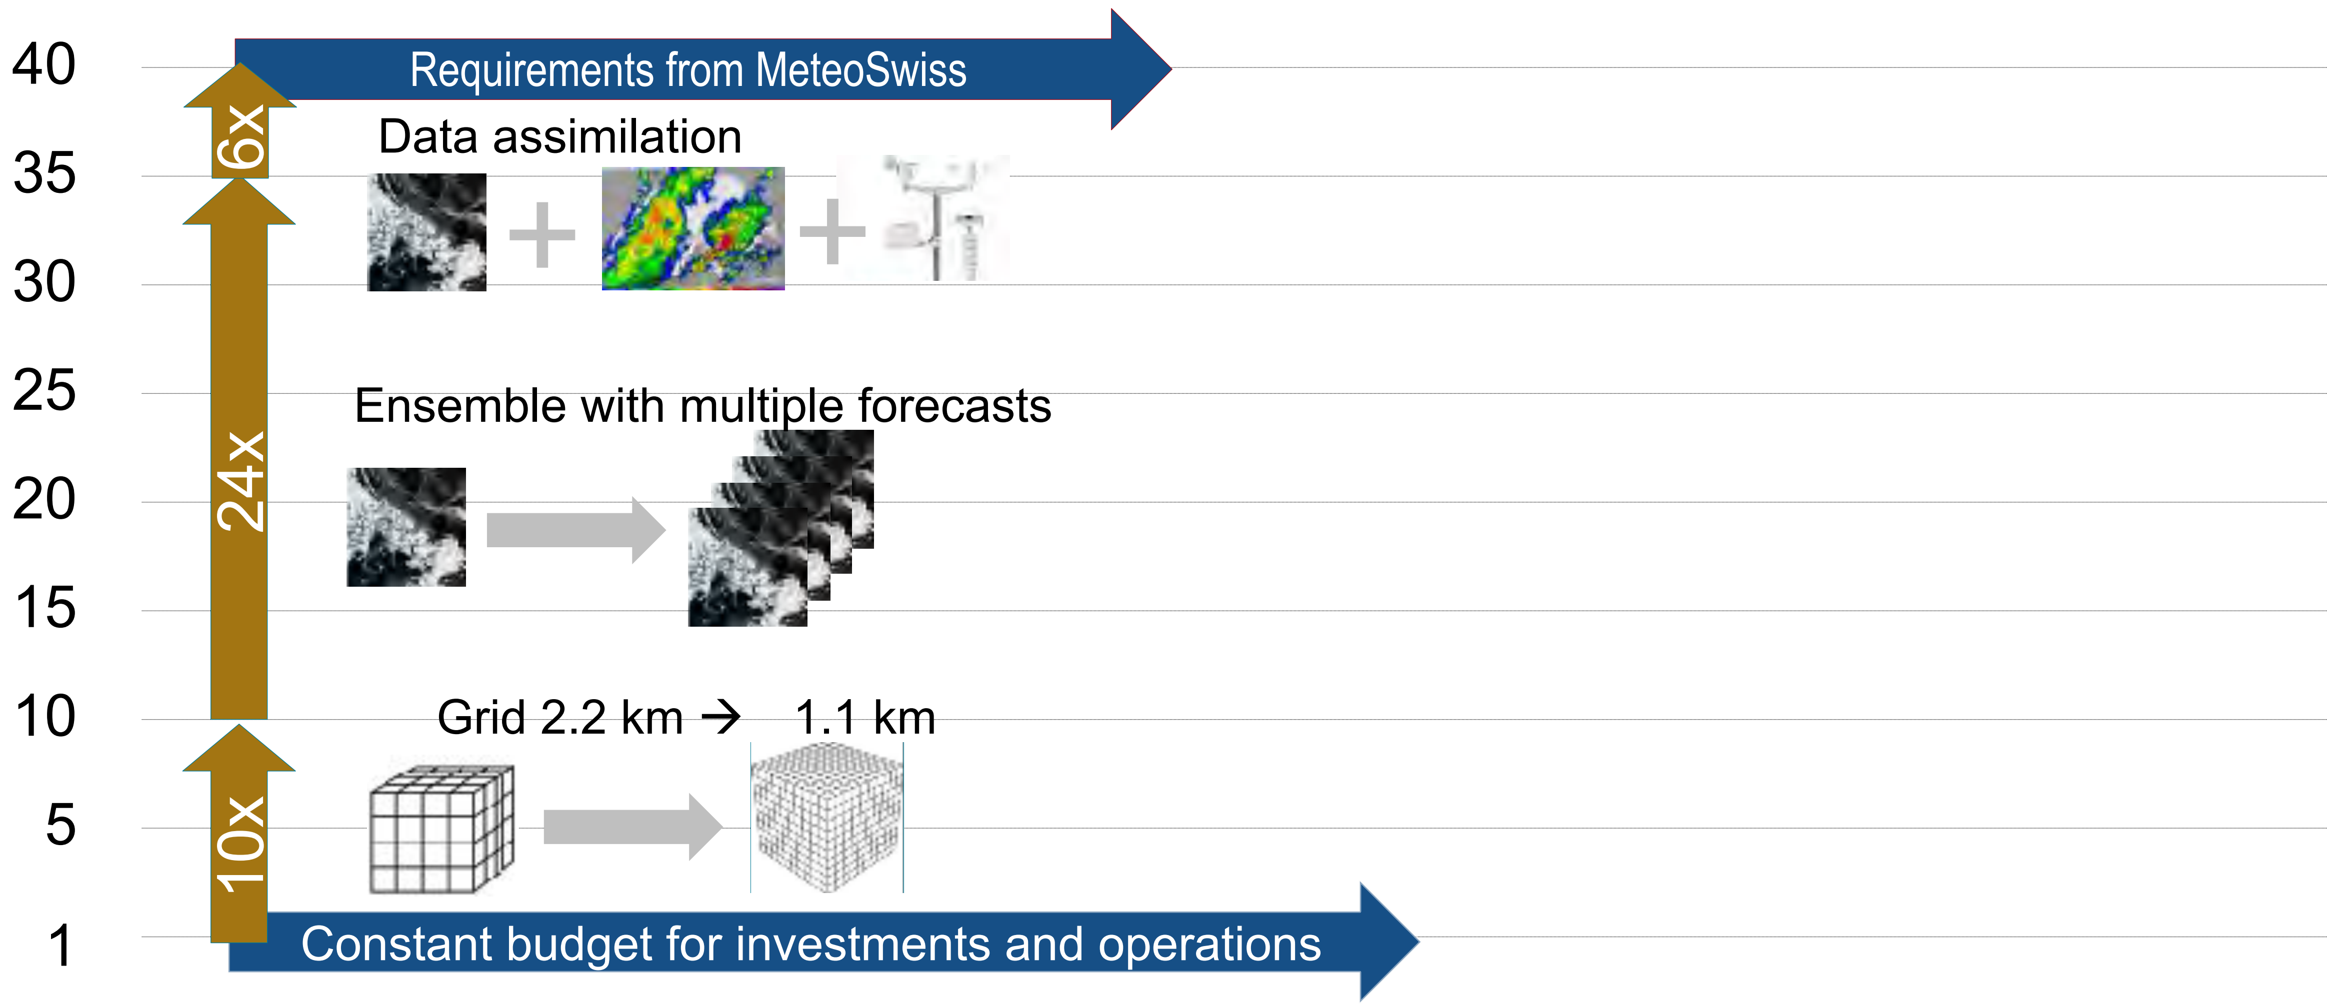

We need a 40x improvement between 2012 and 2015 at constant cost

ETH zürich

Where the factor 40 improvement came from

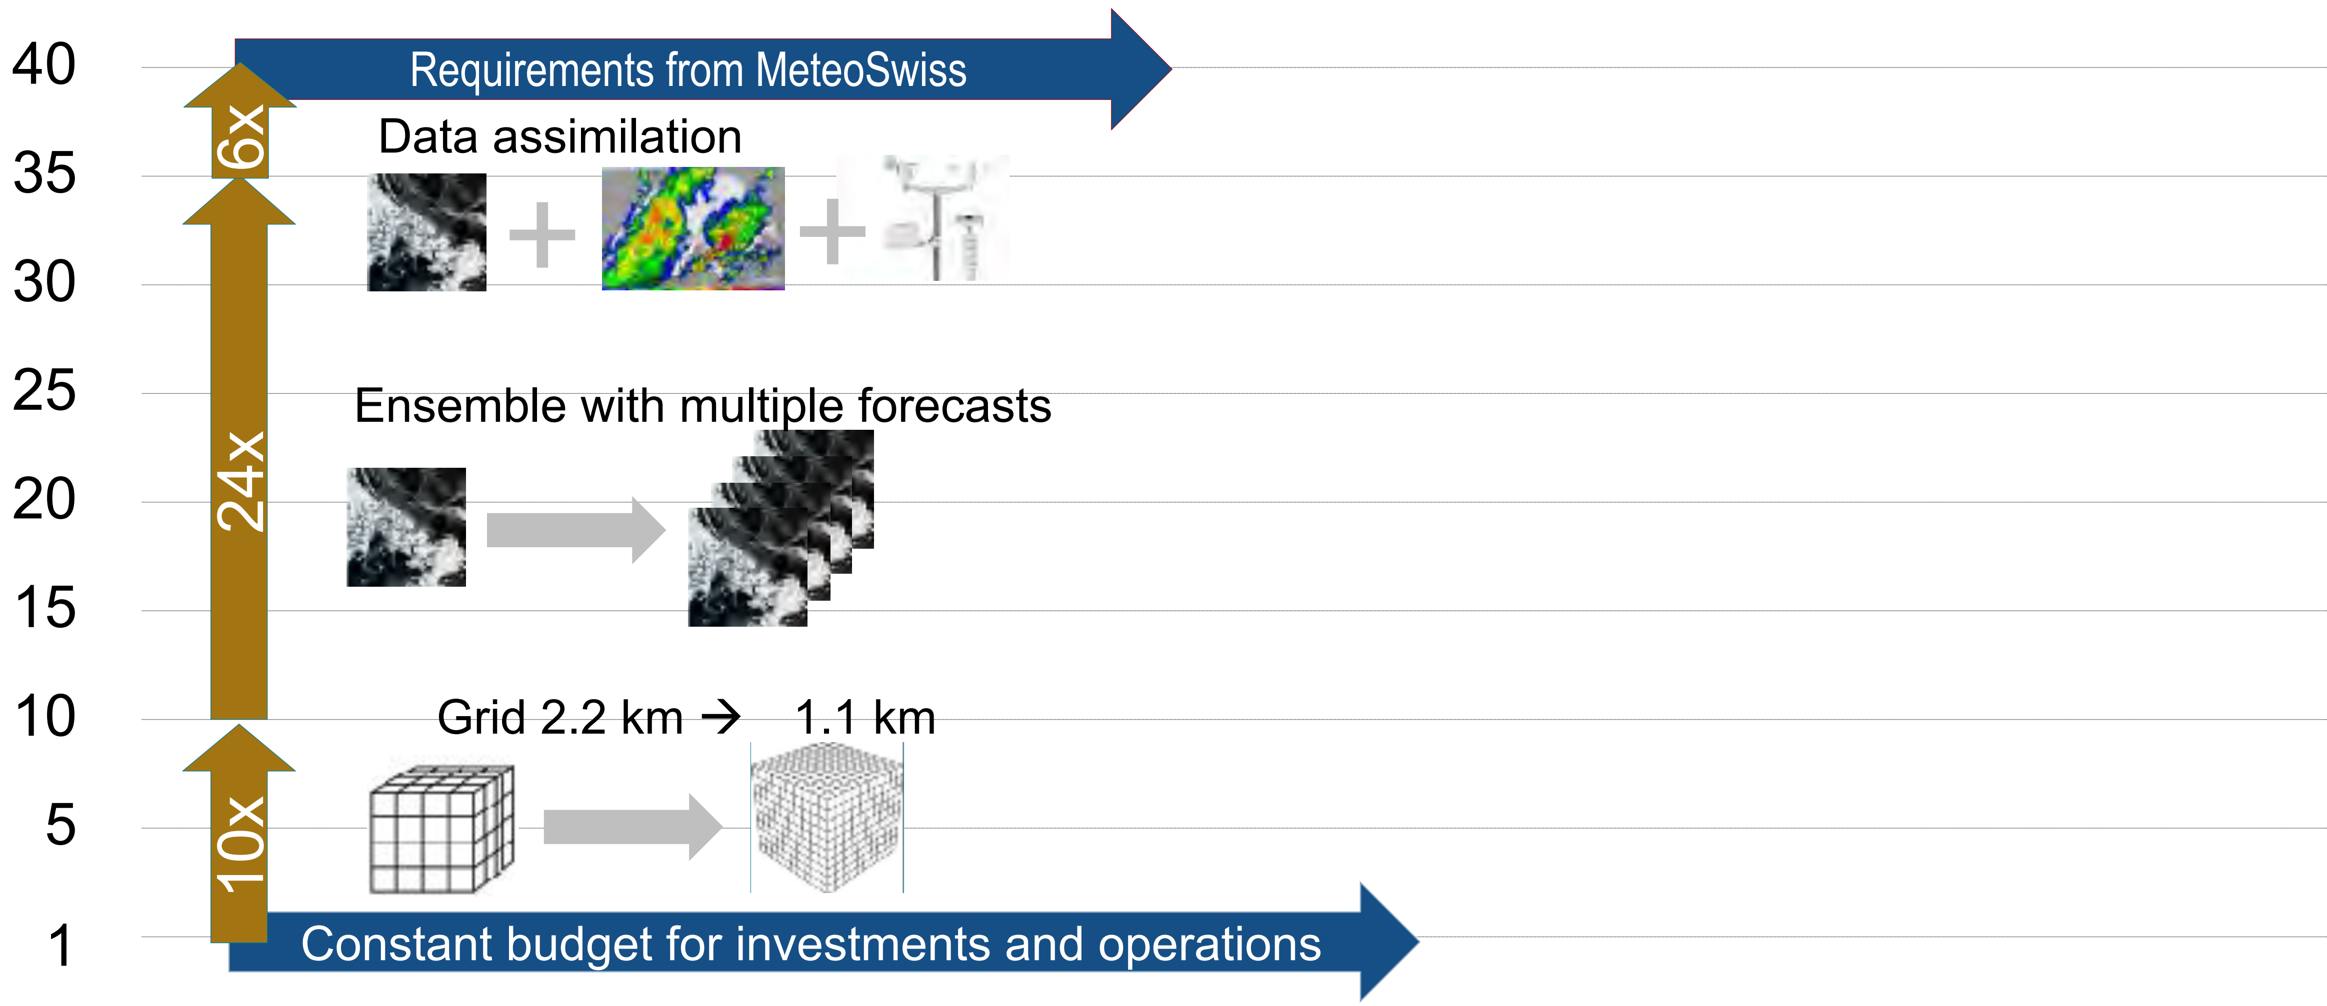

We need a 40x improvement between 2012 and 2015 at constant cost

ETH zürich

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

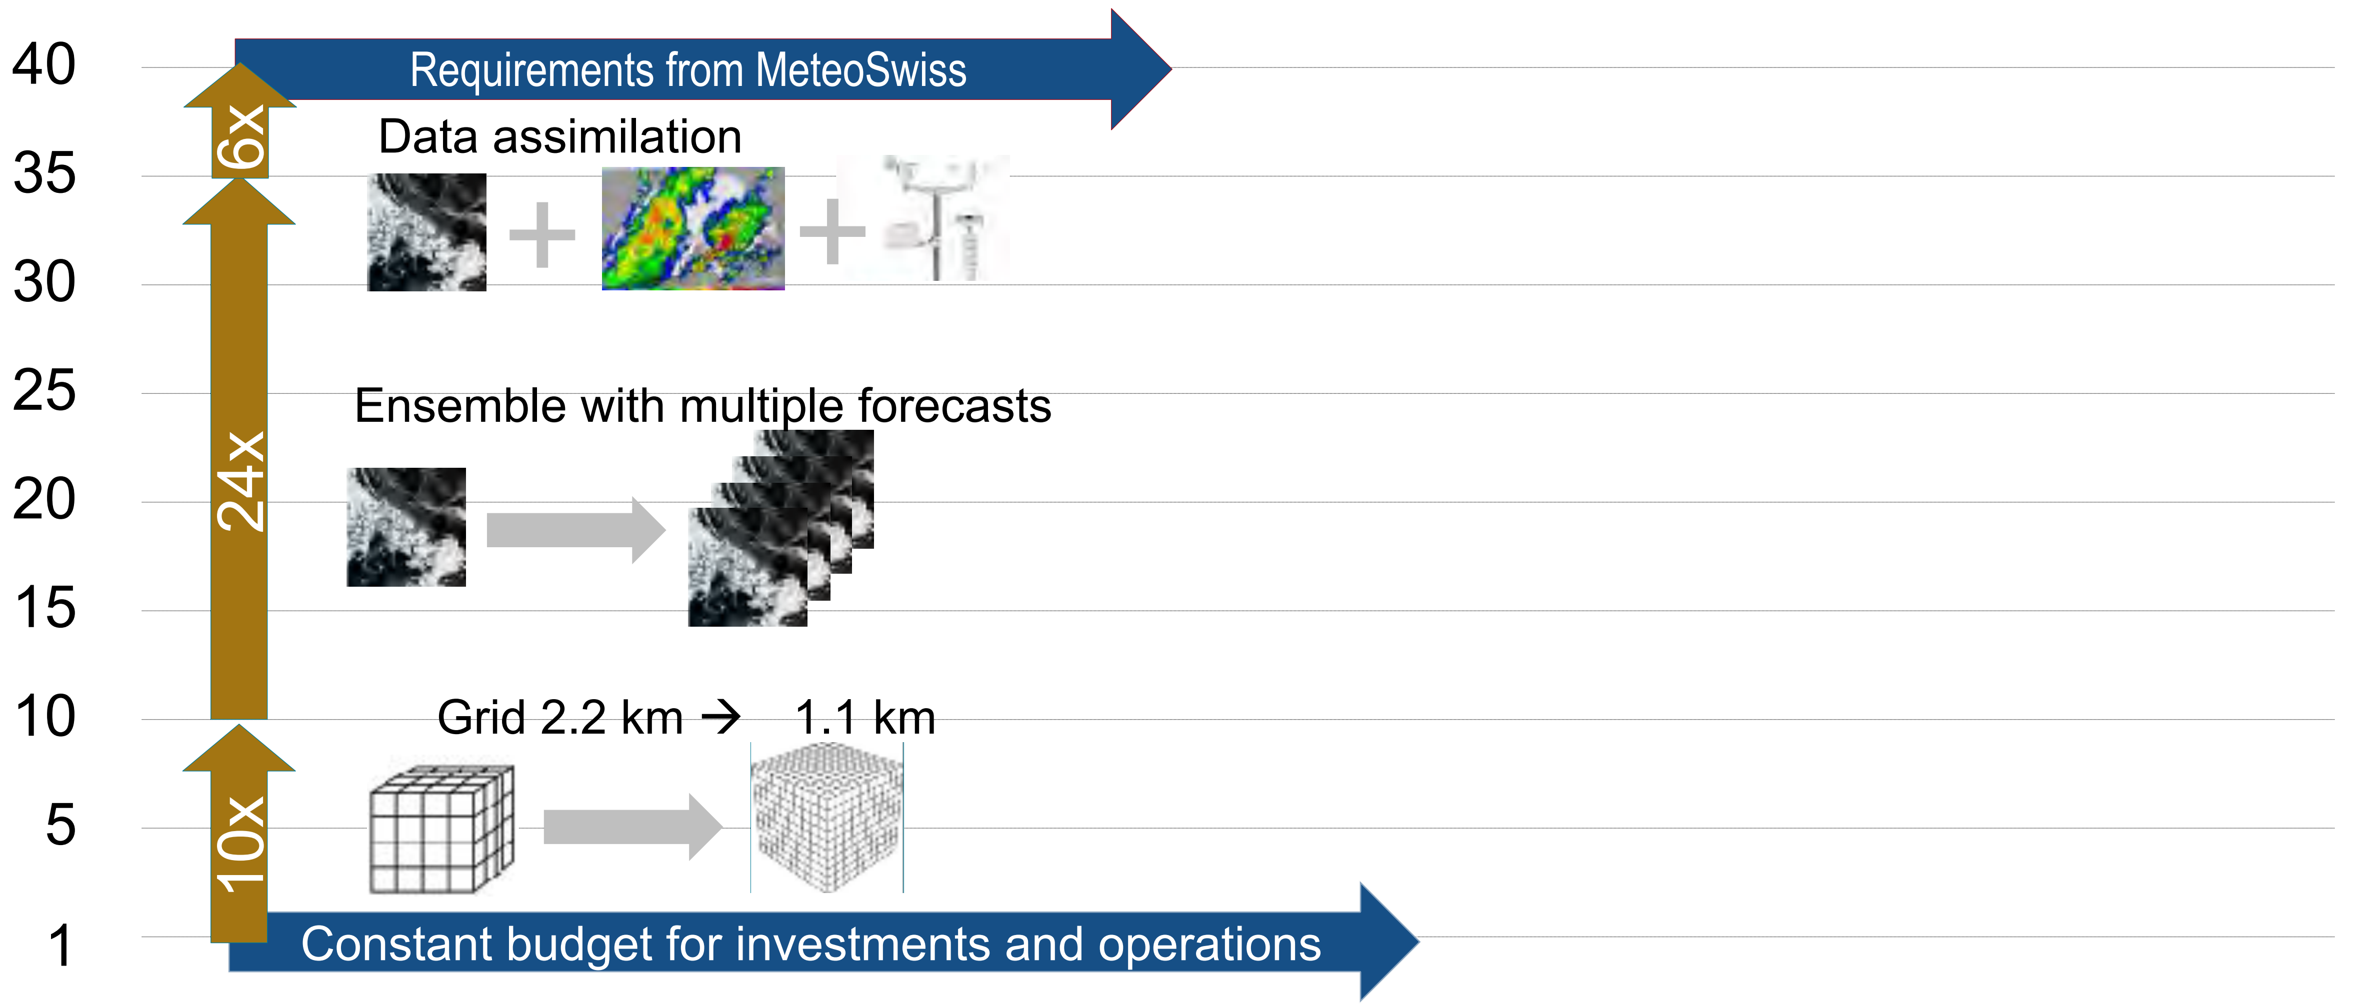

We need a 40x improvement between 2012 and 2015 at constant cost

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

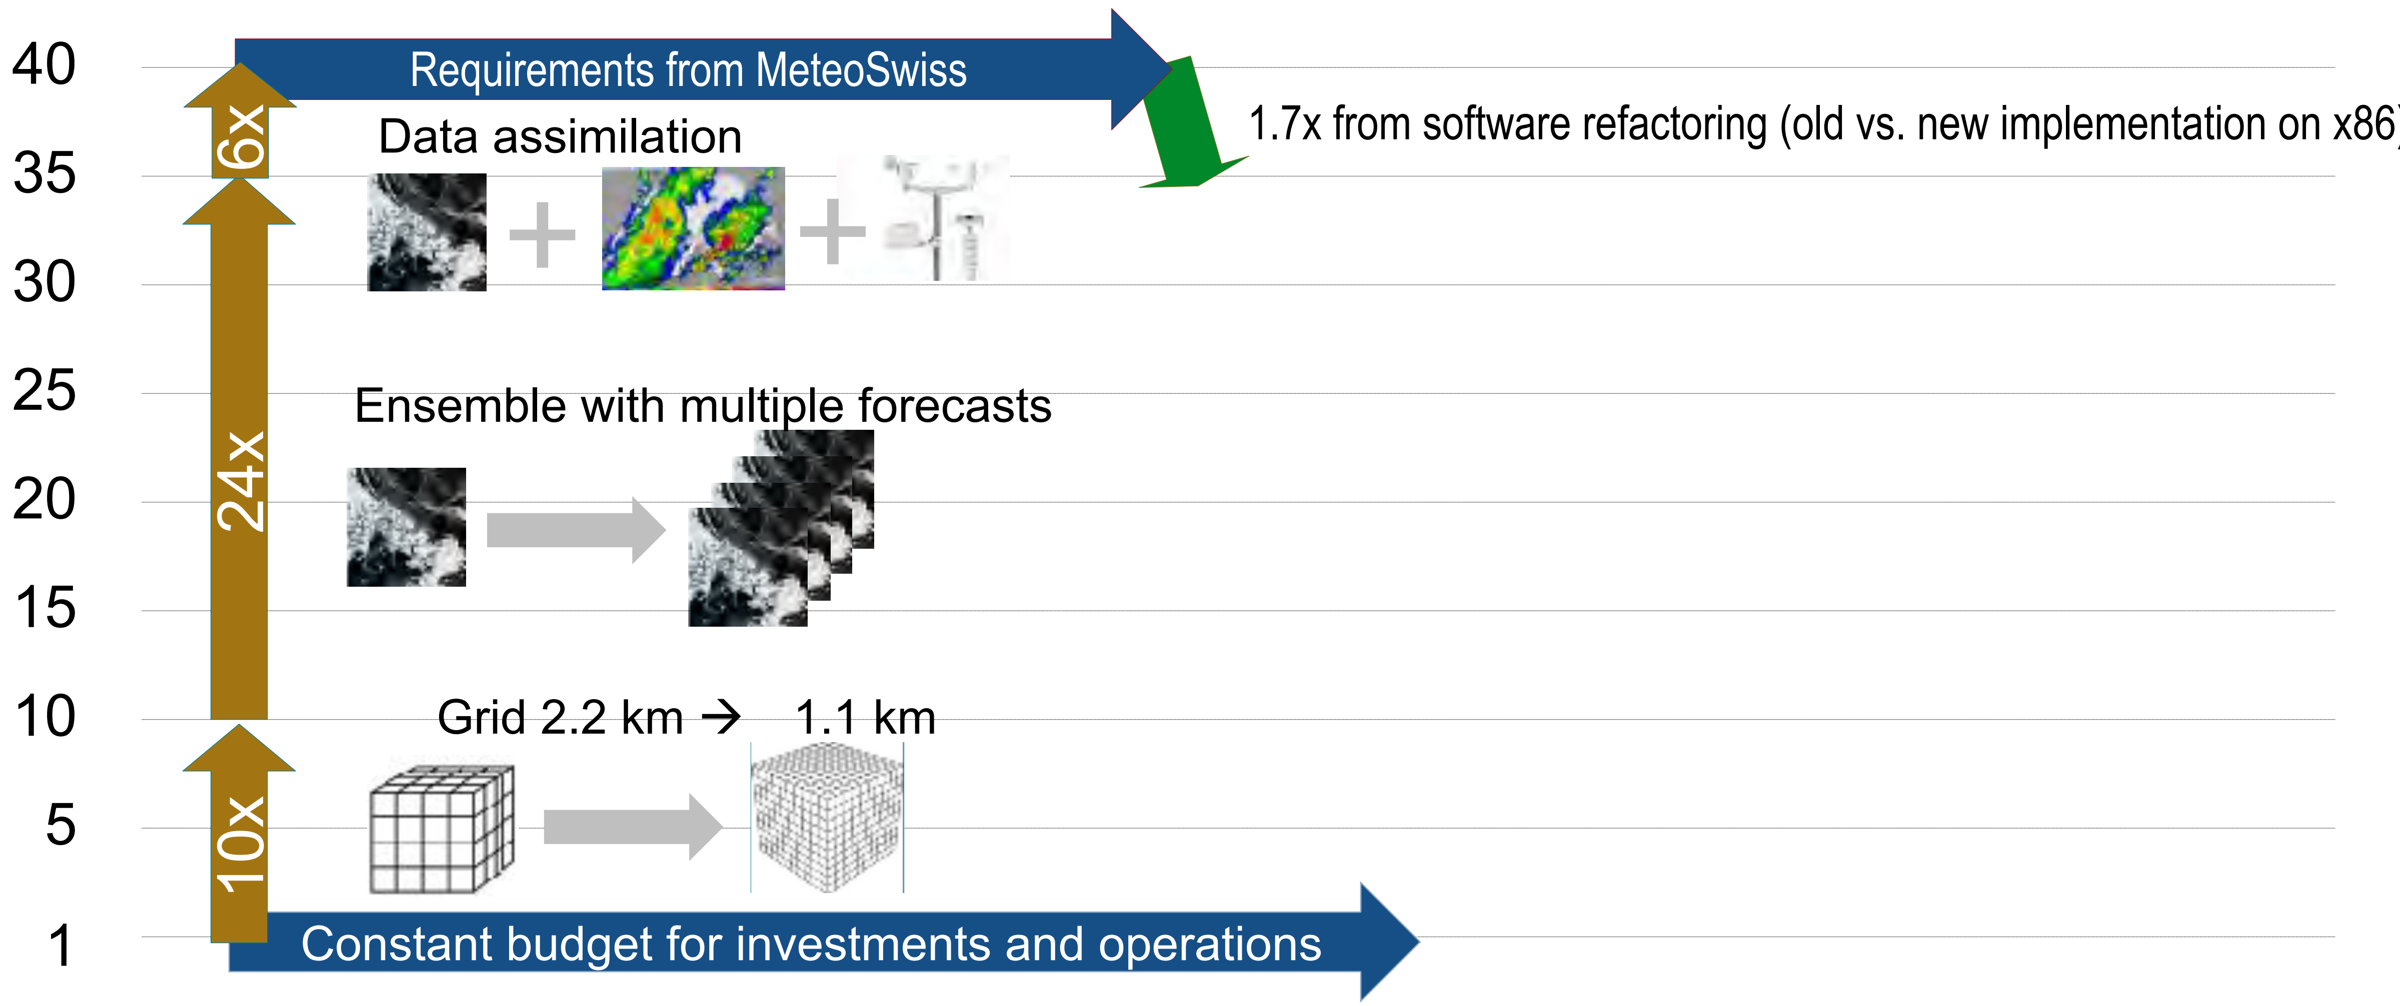

We need a 40x improvement between 2012 and 2015 at constant cost

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

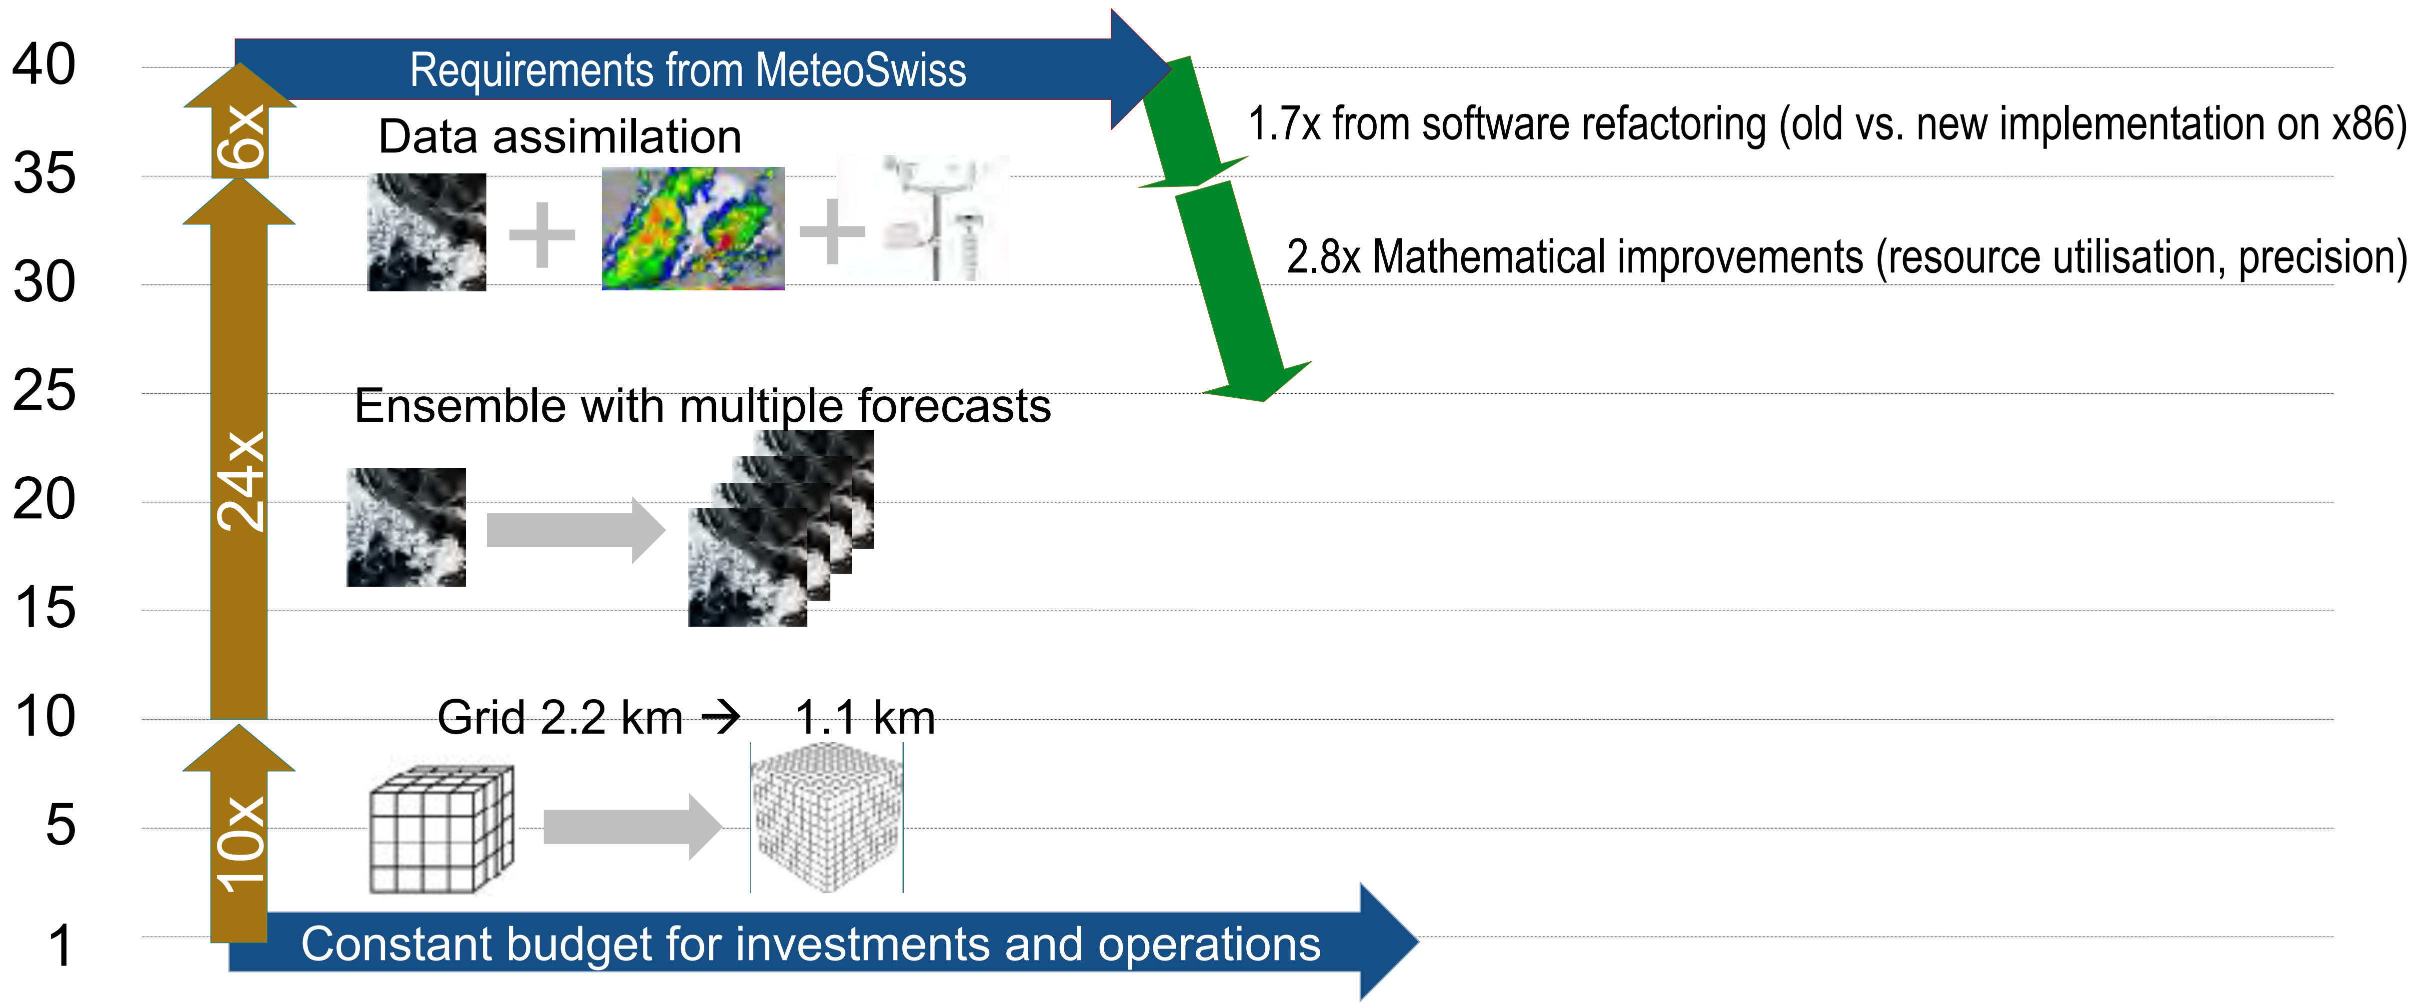

We need a 40x improvement between 2012 and 2015 at constant cost

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

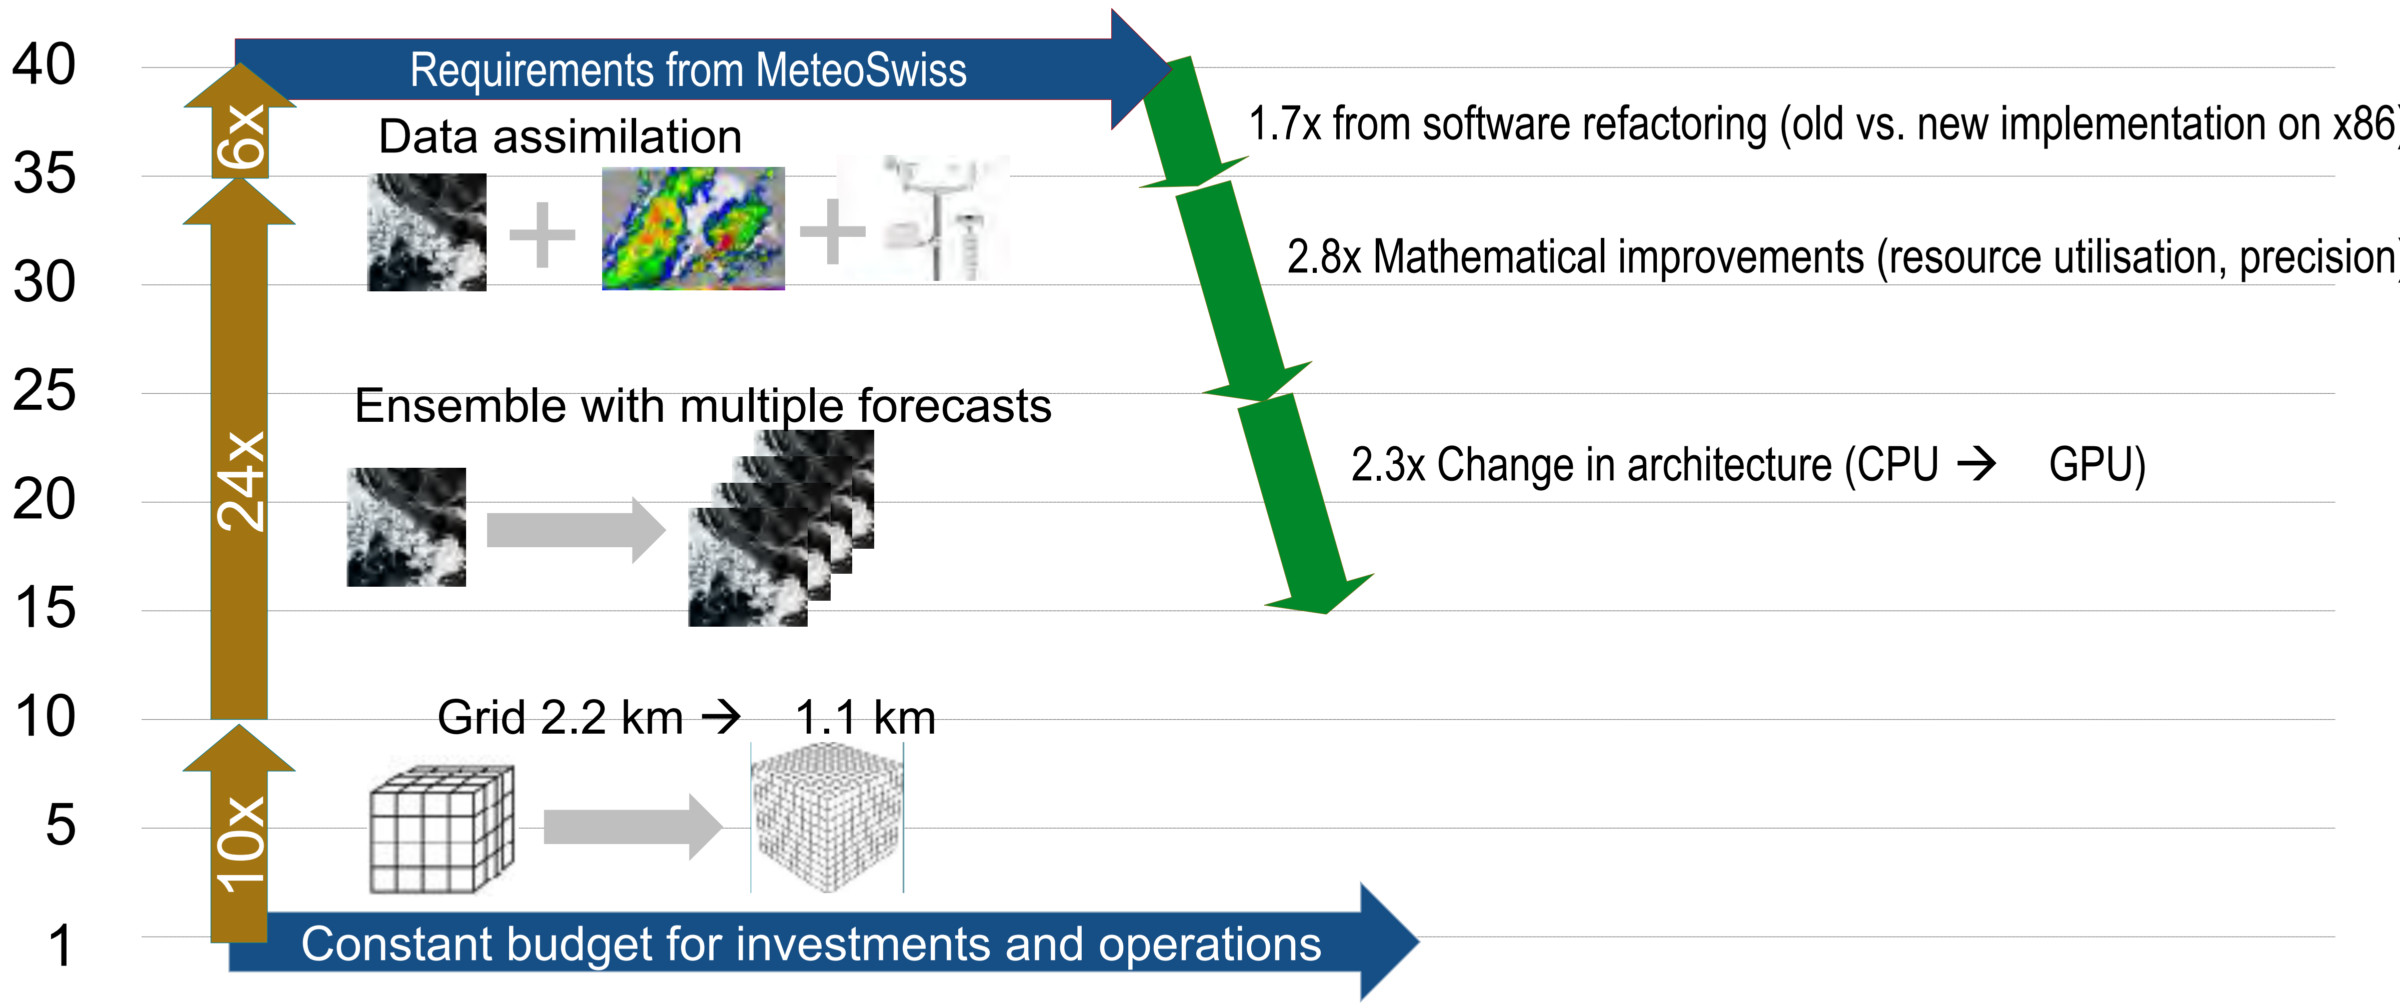

We need a 40x improvement between 2012 and 2015 at constant cost

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

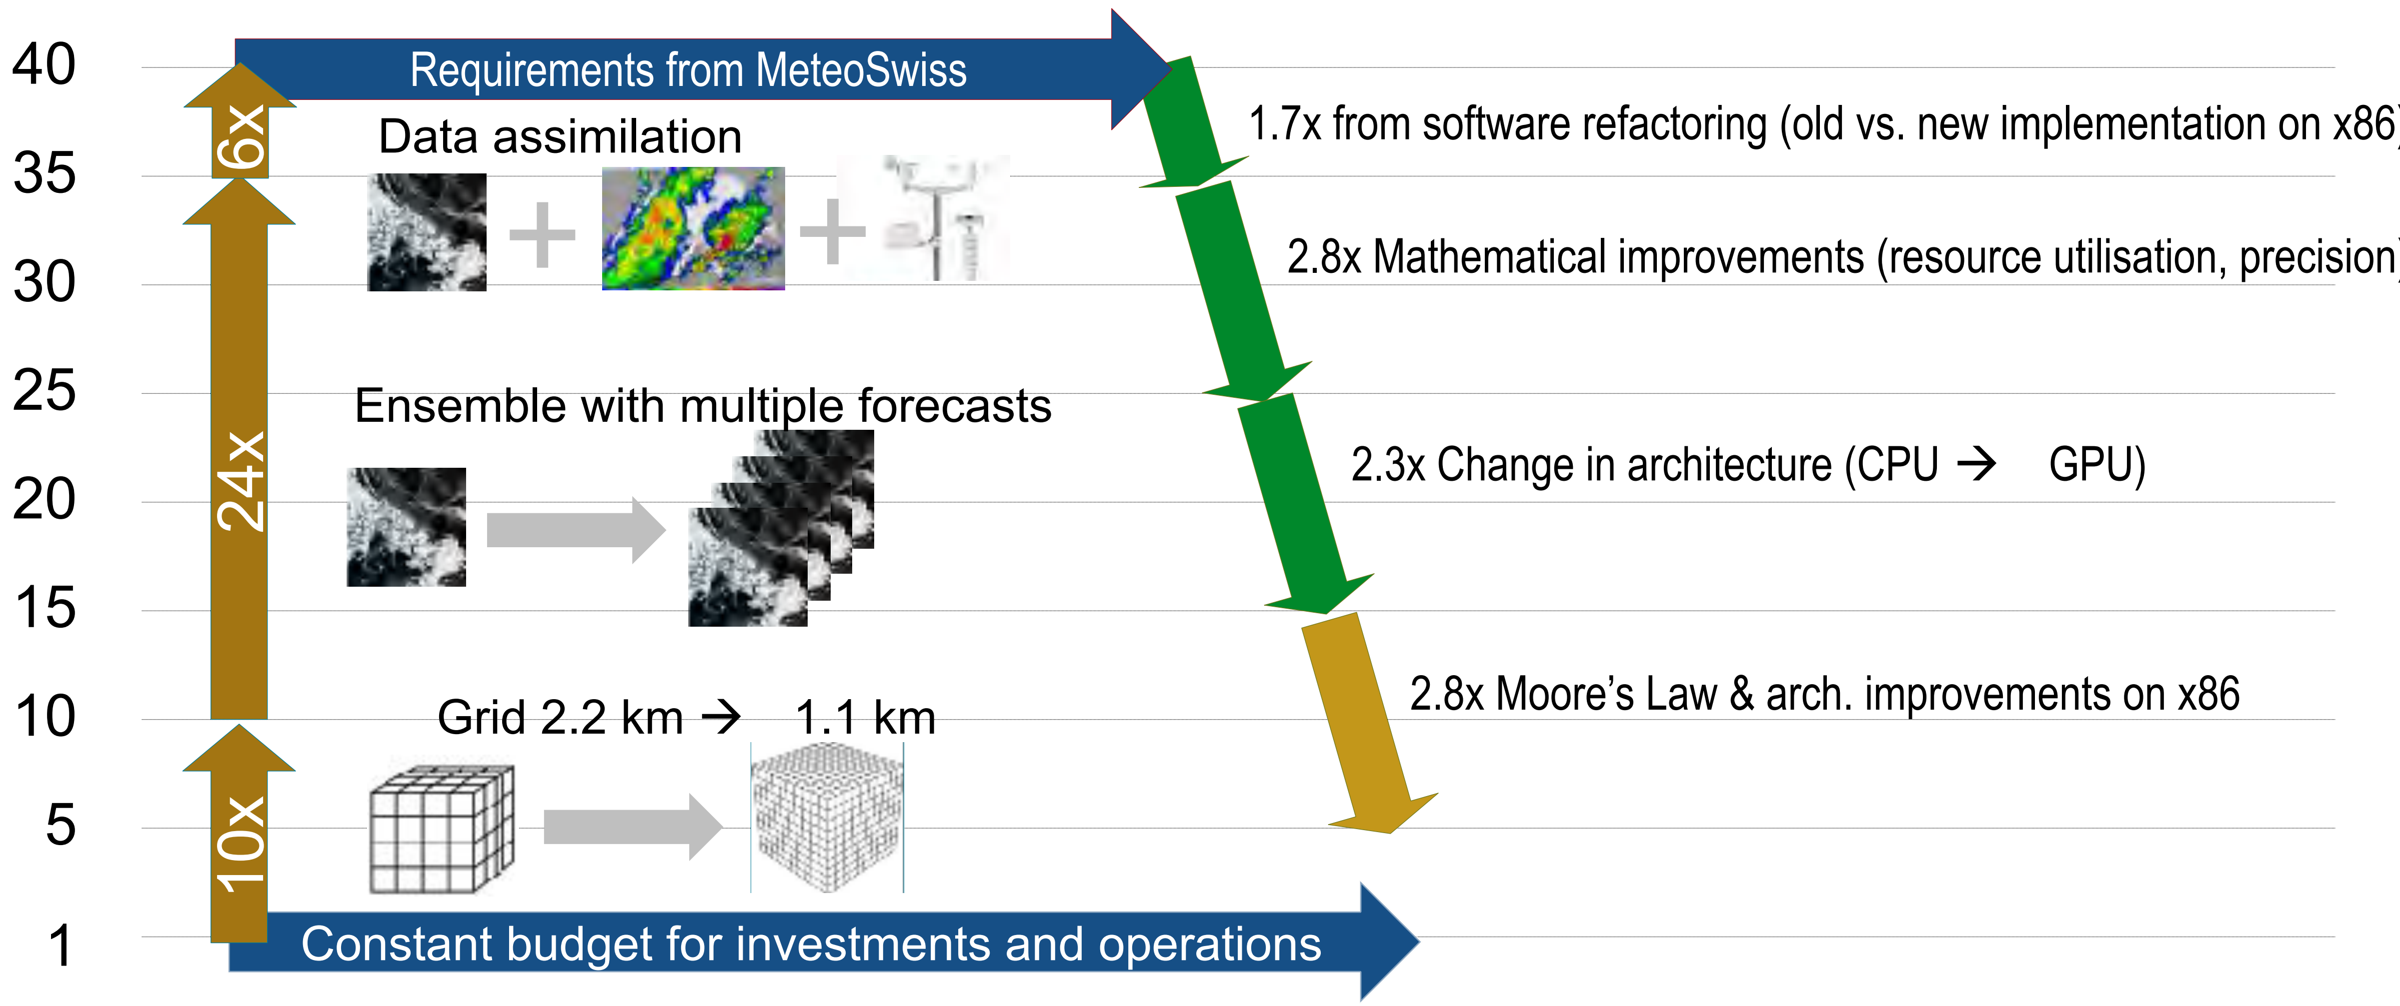

We need a 40x improvement between 2012 and 2015 at constant cost

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

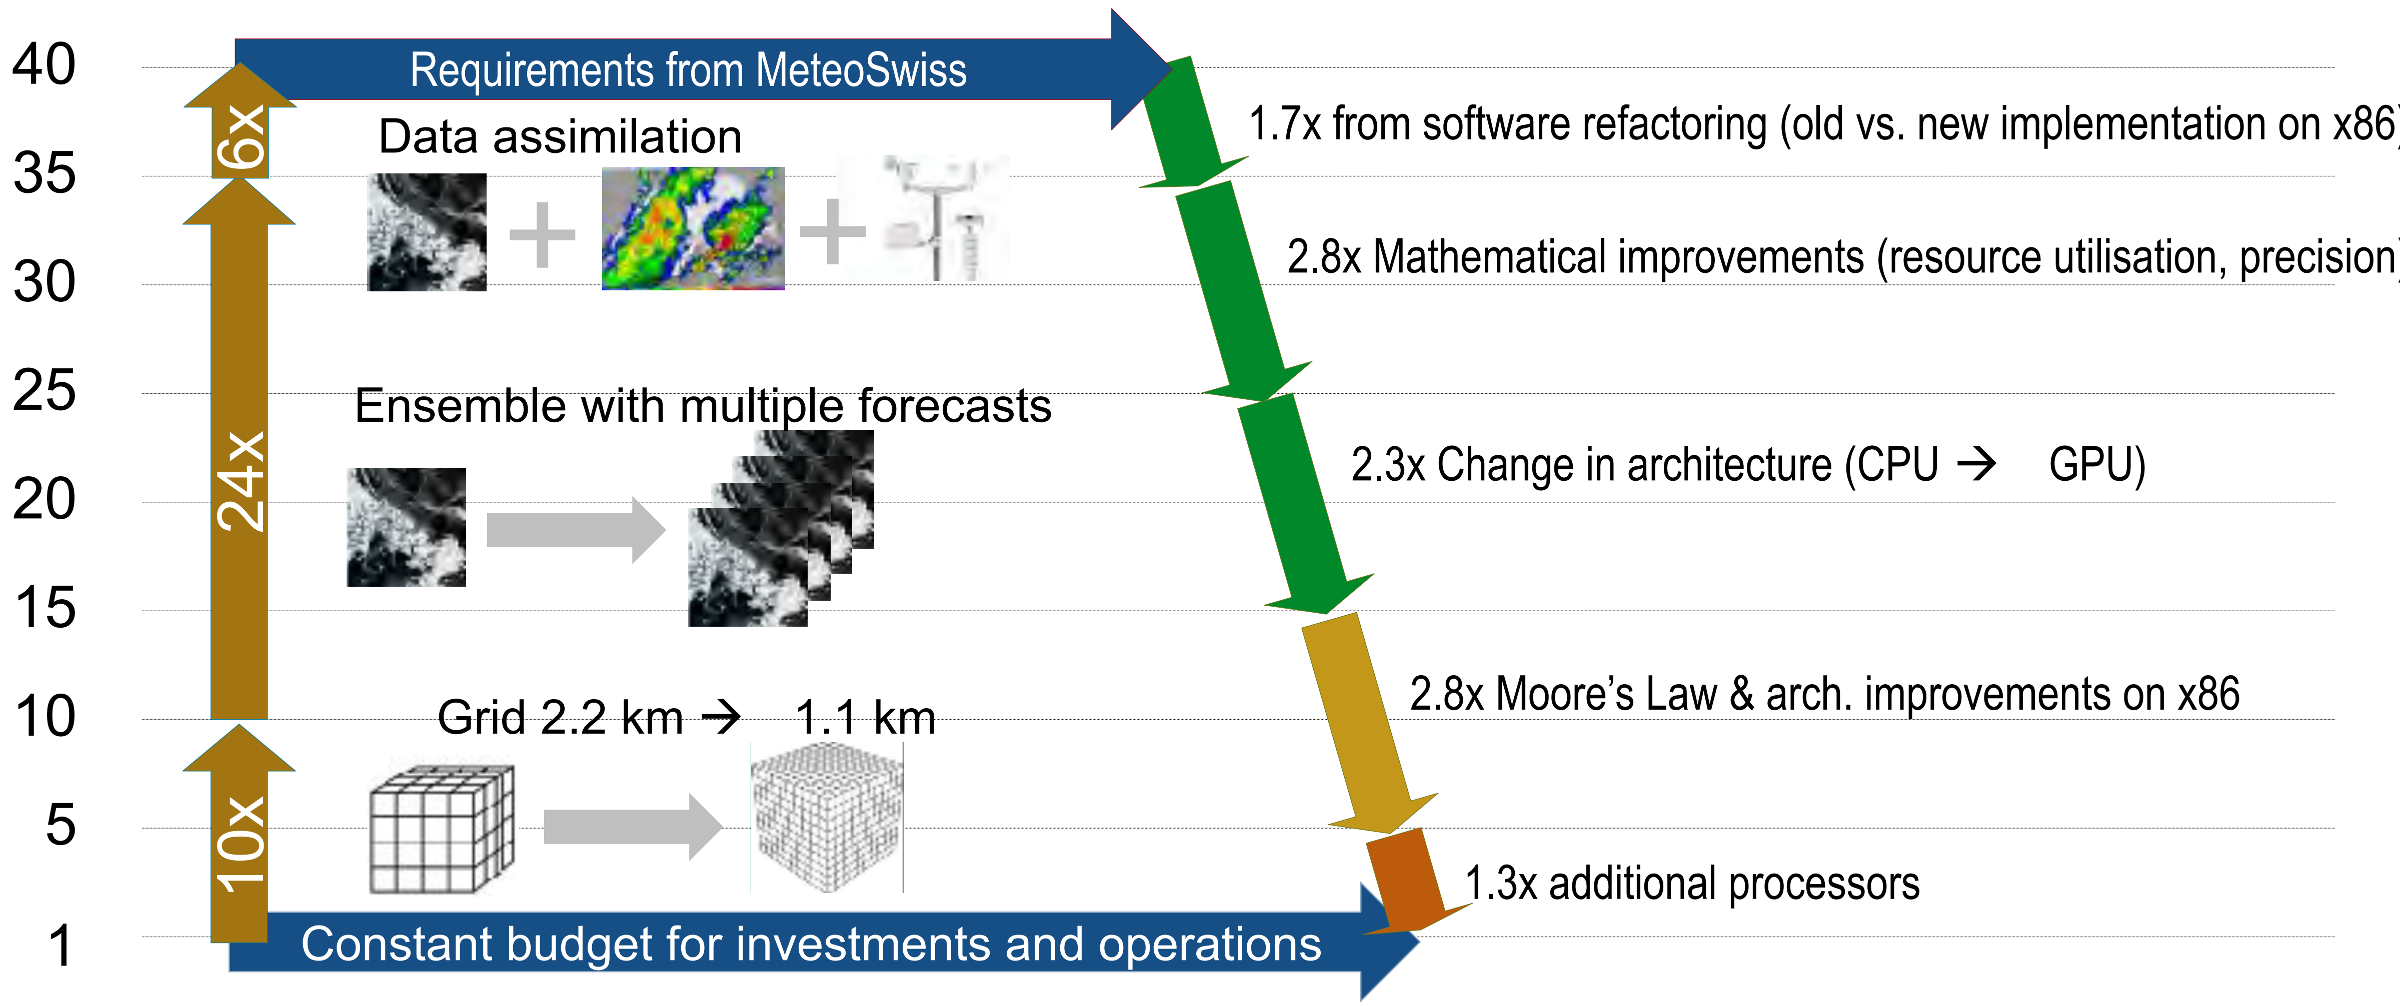

We need a 40x improvement between 2012 and 2015 at constant cost

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

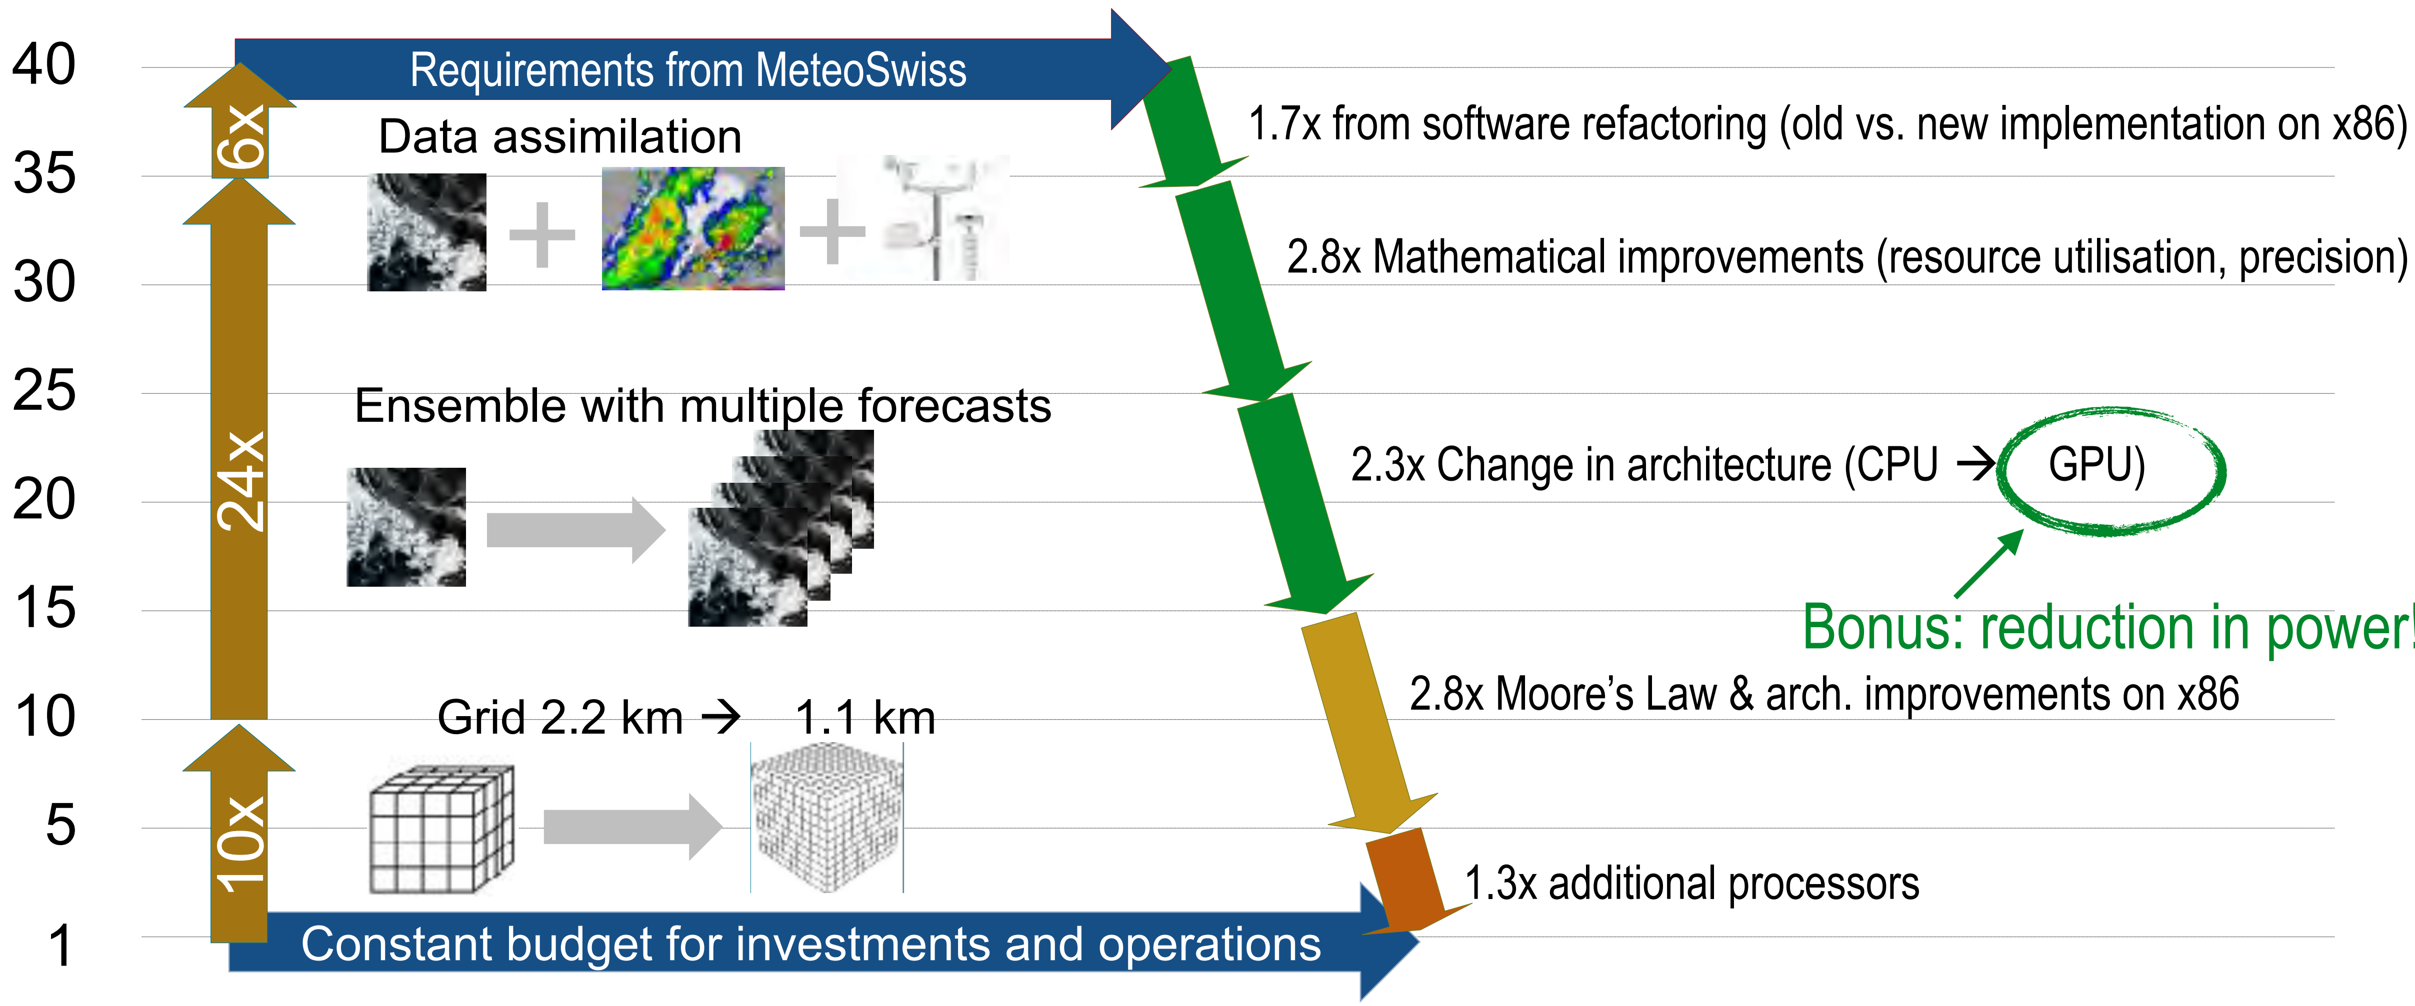

We need a 40x improvement between 2012 and 2015 at constant cost

Where the factor 40 improvement came from

Investment in software allowed mathematical improvements and change in architecture

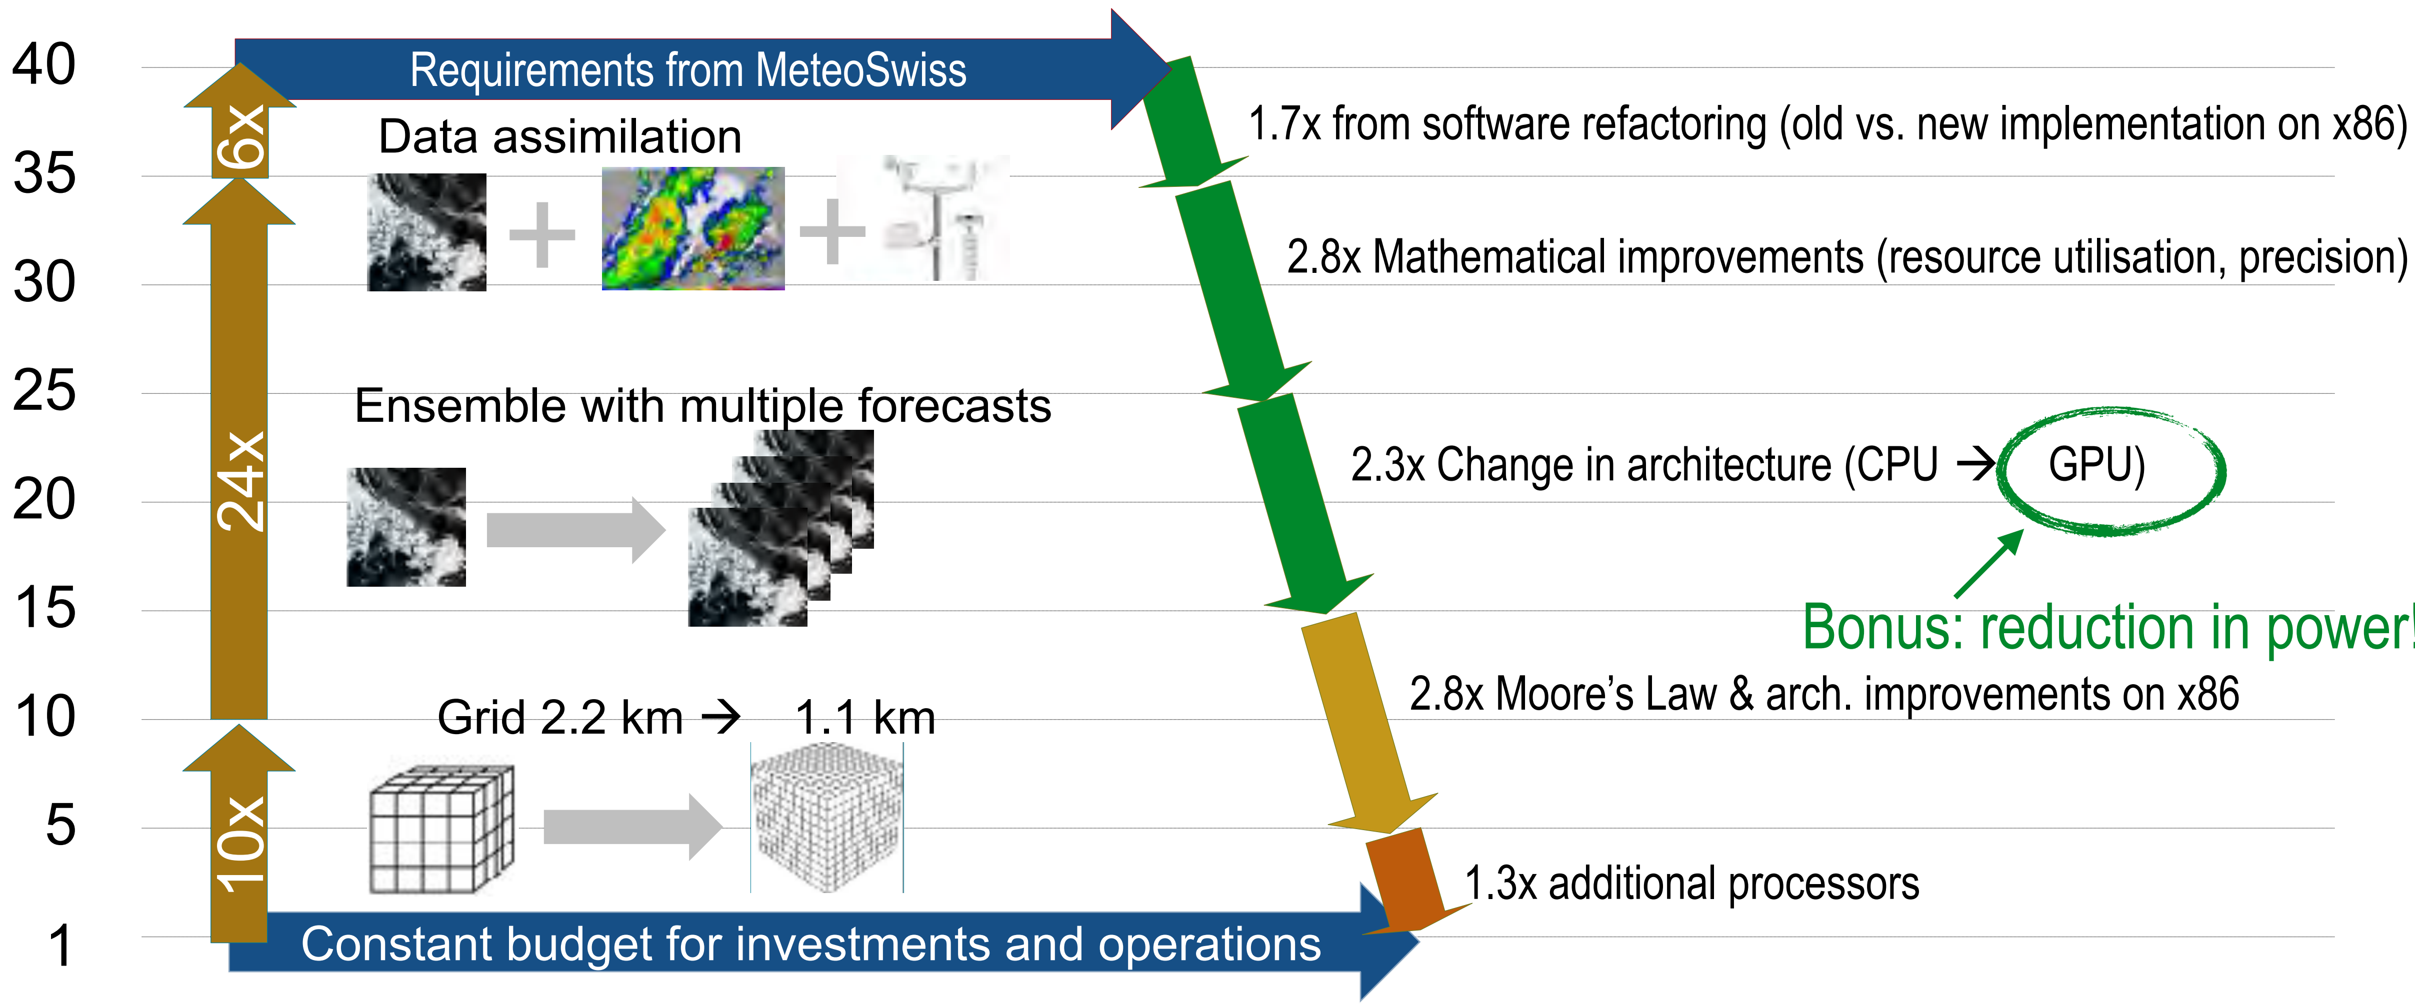

We need a 40x improvement between 2012 and 2015 at constant cost

There is no silver bullet!

Near-global climate simulation at 1km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

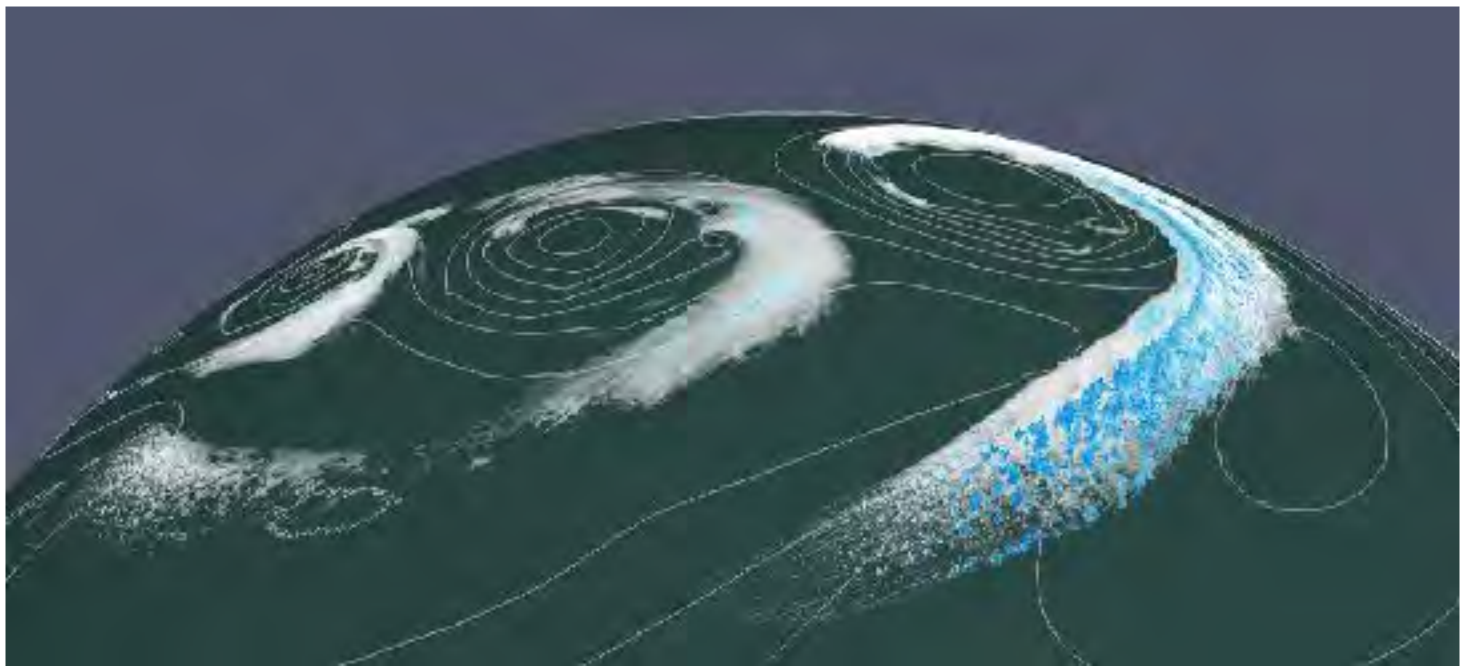

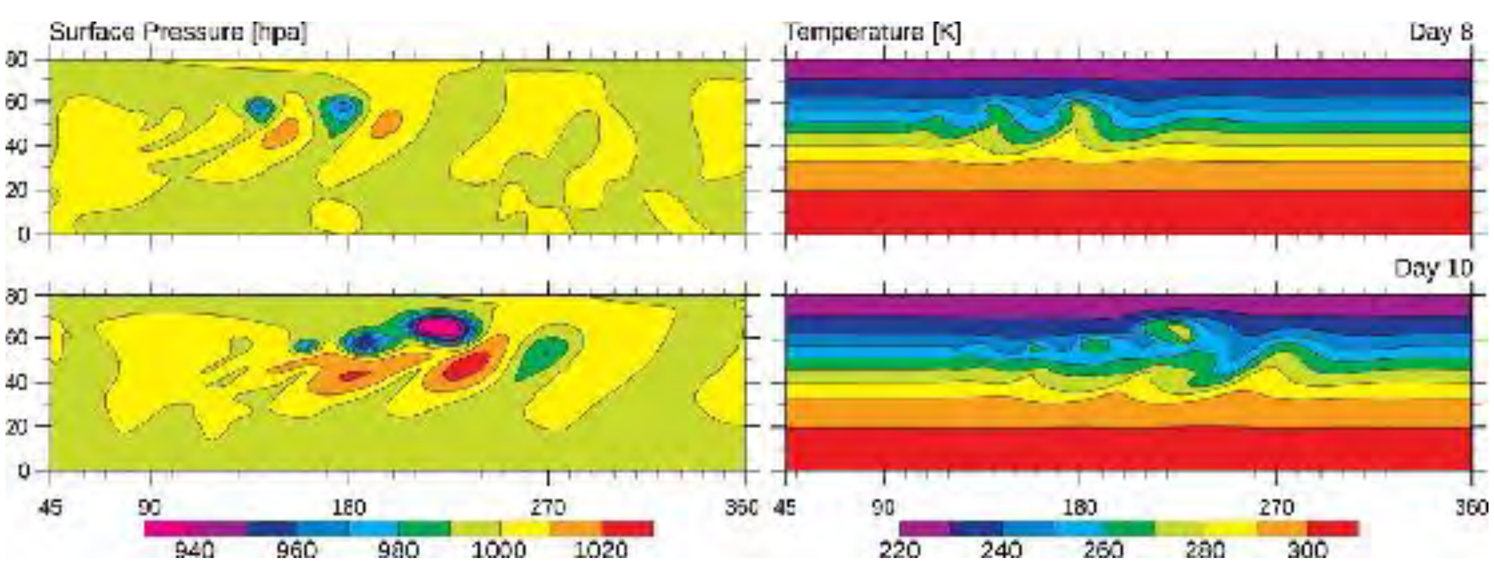

Near-global climate simulation at 1km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

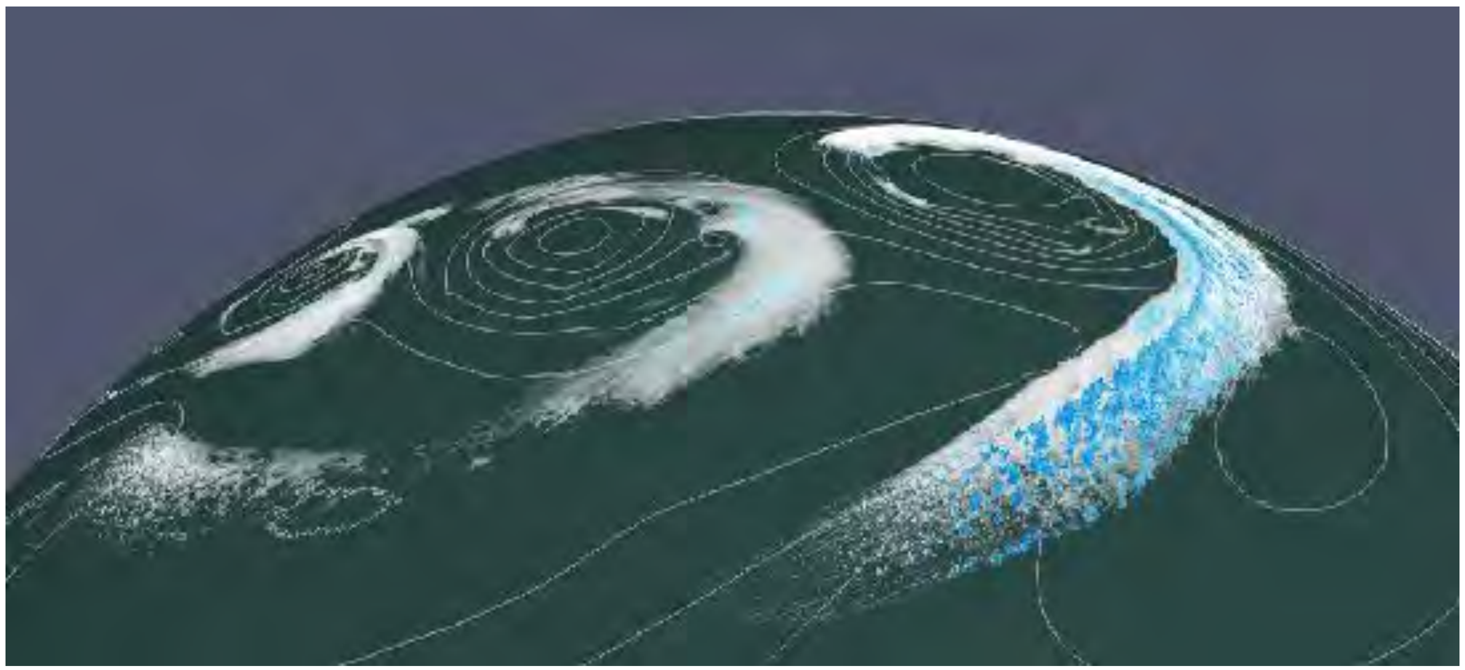

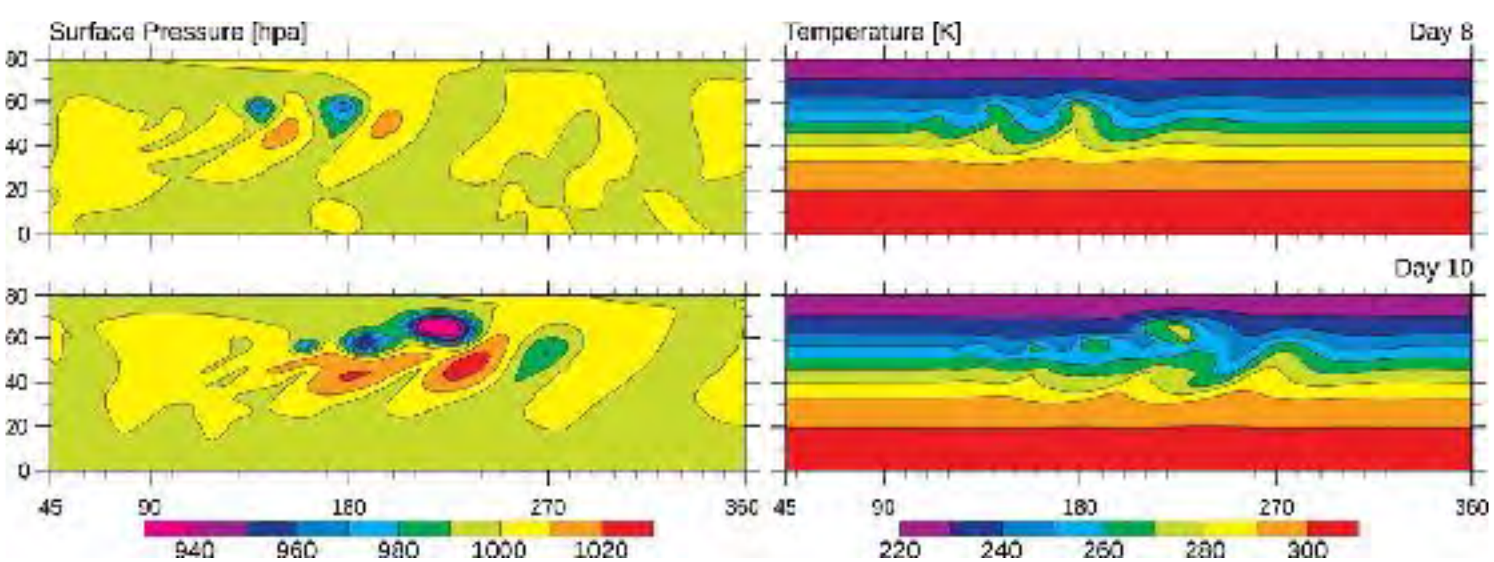

Metric: simulated years per wall-clock day

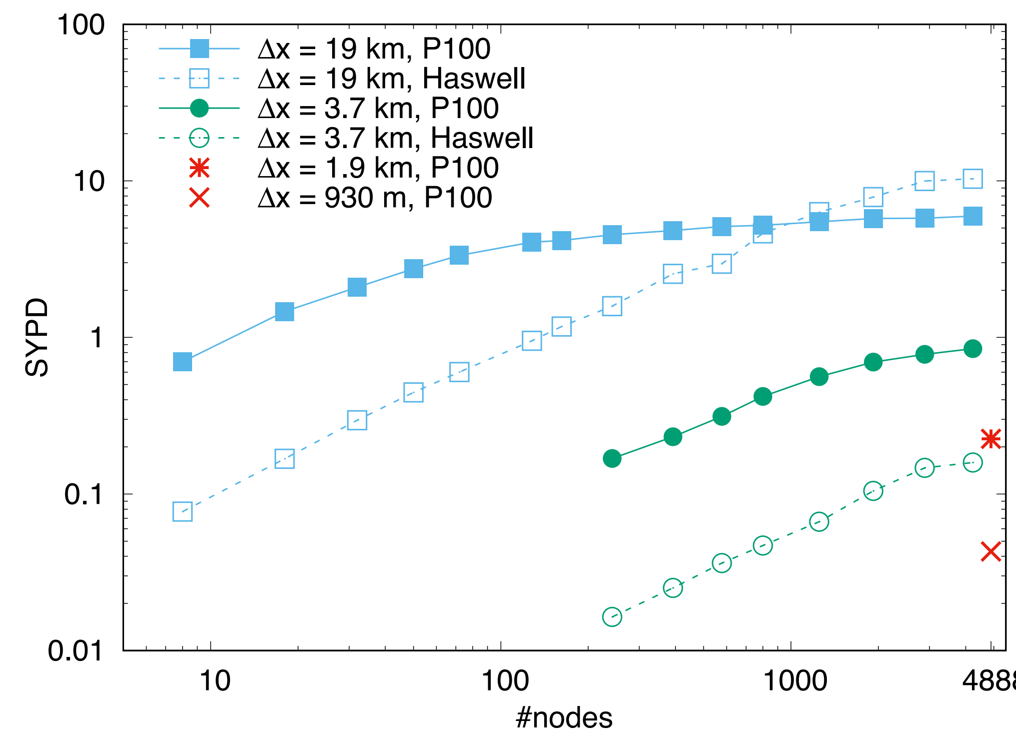

$\langle \Delta x \rangle$	#nodes	Δt [s]	SYPD	MWh/SY	gridpoints
930 m	4,888	6	0.043	596	3.46×10^{10}
1.9 km	4,888	12	0.23	97.8	8.64×10^{9}
47 km	18	300	9.6	0.099	1.39×10^{7}

(c) Time compression (SYPD) and energy cost (MWh/SY) for three moist simulations. At 930 m grid spacing obtained with a full 10d simulation, at 1.9 km from 1,000 steps, and at 47 km from 100 steps

Near-global climate simulation at 1km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

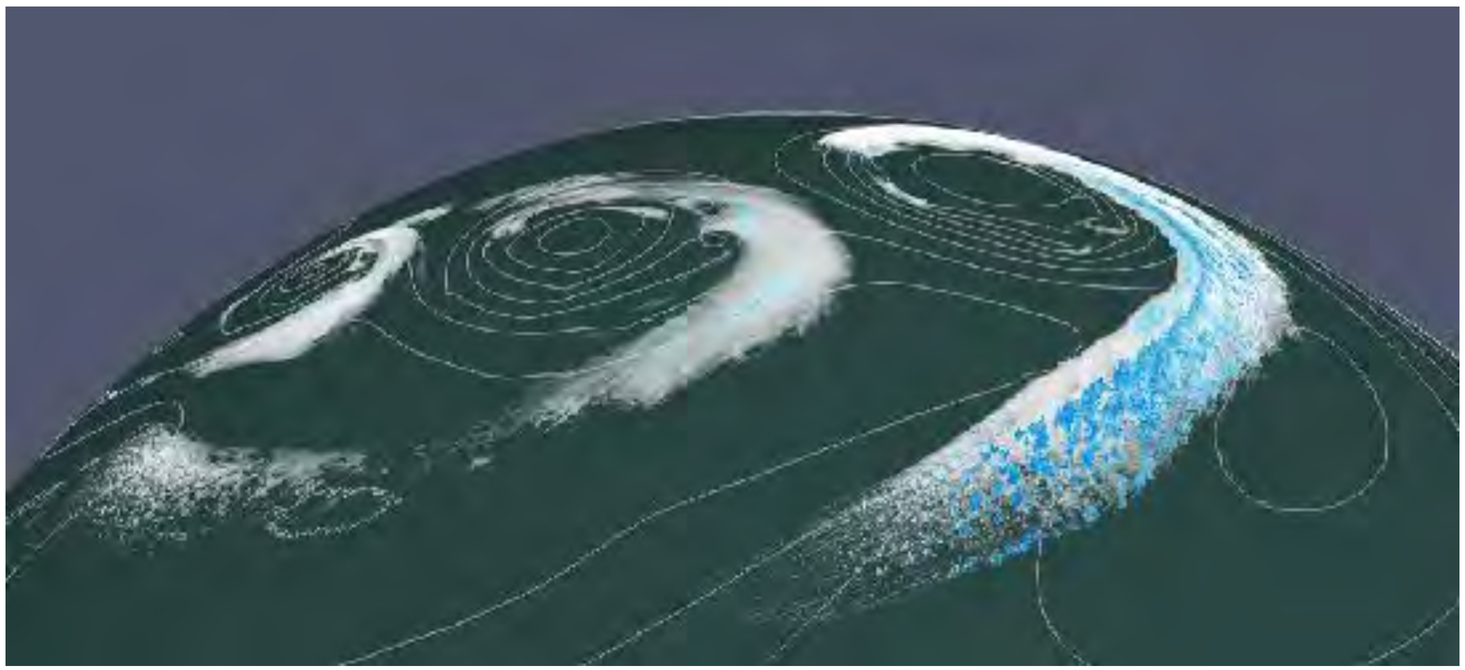

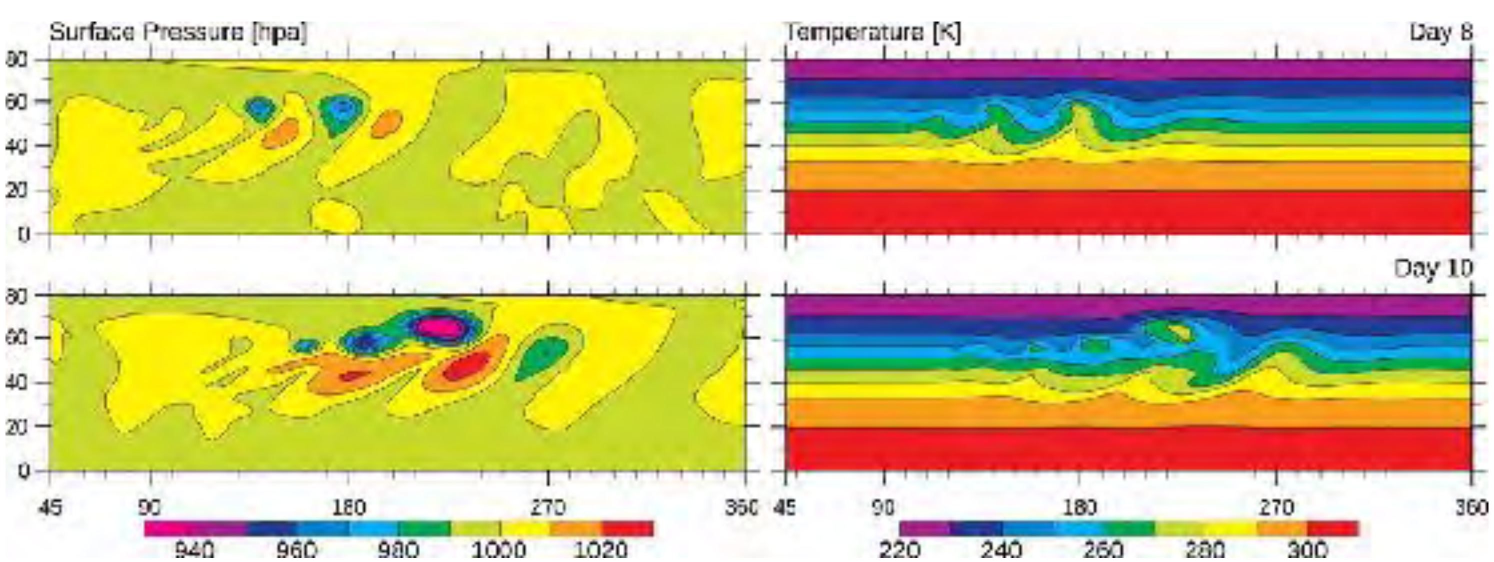

Metric: simulated years per wall-clock day

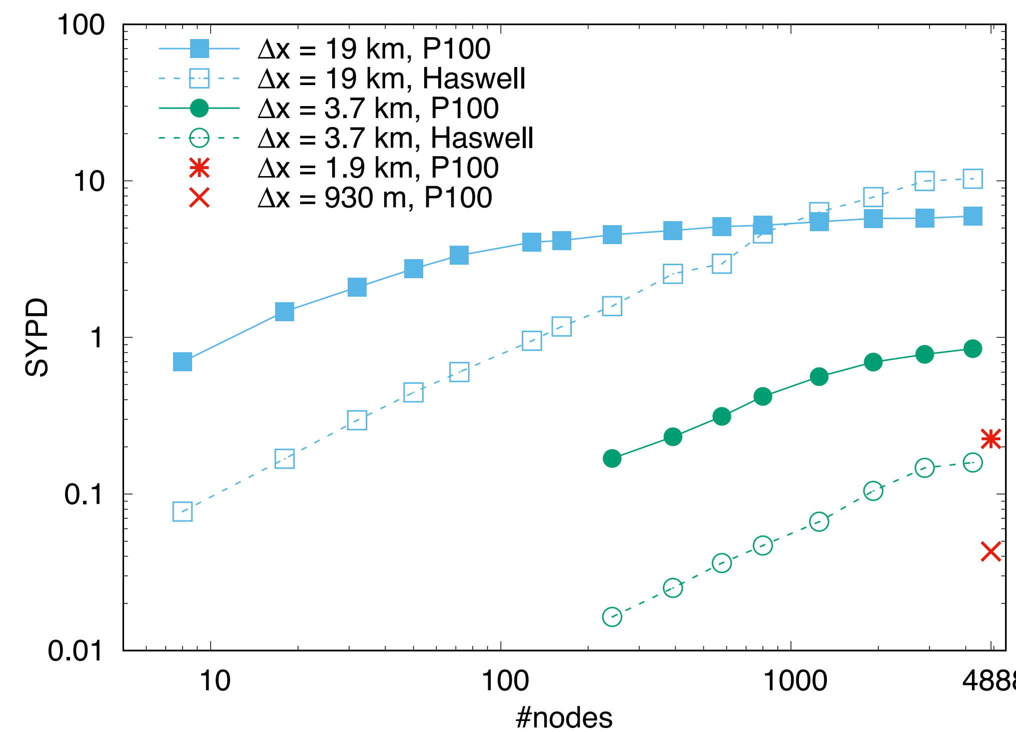

$\langle \Delta x \rangle$	#nodes	Δt [s]	SYPD	MWh/SY	gridpoints
930 m	4,888	6	0.043	596	3.46×10^{10}
1.9 km	4,888	12	0.23	97.8	8.64×10^{9}
47 km	18	300	9.6	0.099	1.39×10^{7}

(c) Time compression (SYPD) and energy cost (MWh/SY) for three moist simulations. At 930 m grid spacing obtained with a full 10d simulation, at 1.9 km from 1,000 steps, and at 47 km from 100 steps

2.5x faster than Yang et al.'s 2016 Gordon Bell winner run on TaihuLight!

The baseline for COSMO-global and IFS

	Near-global COSMO [Fuh2018]		Global IFS [Wed2009]	
	Value	Shortfall	Value	Shortfall
Horizontal resolution	0.93 km (non-uniform)	0.9x	1.25 km	1.25x
Vertical resolution	60 levels (surface to 25 km)	3x	62 levels (surface to 40 km)	3x
Time resolution	6 s (split-explicit with sub-stepping)	-	120 s (semi-implicit)	4x
Coupled	No	1.2x	No	1.2x
Atmosphere	Non-hydrostatic	-	Non-hydrostatic	-
Precision	Double	0.6x	Single	-
Compute rate	0.043 SYPD	23x	0.088 SYPD	11x
Other (I/O, full physics, …)	Limited I/O Only microphysics	1.5x	Full physics, no I/O	-
Total shortfall		**65x**		**198x**

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

Necessary data transfers

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

Necessary data transfers

$$MUE = \mathrm{I/O\ efficiency} \cdot \mathrm{BW\ efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Actual data transfers

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

0.88

Necessary data transfers

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Actual data transfers

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

0.88

Necessary data transfers

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Actual data transfers

Max achievable BW

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

0.88

Necessary data transfers

Achieved BW

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D} \frac{B}{\hat{B}}$$

Actual data transfers

Max achievable BW

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

0.88

Necessary data transfers

Achieved BW

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Actual data transfers

Max achievable BW

0.76

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

0.88

Necessary data transfers

Achieved BW

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\frac{B}{\hat{B}}$$

Actual data transfers

Max achievable BW

0.76

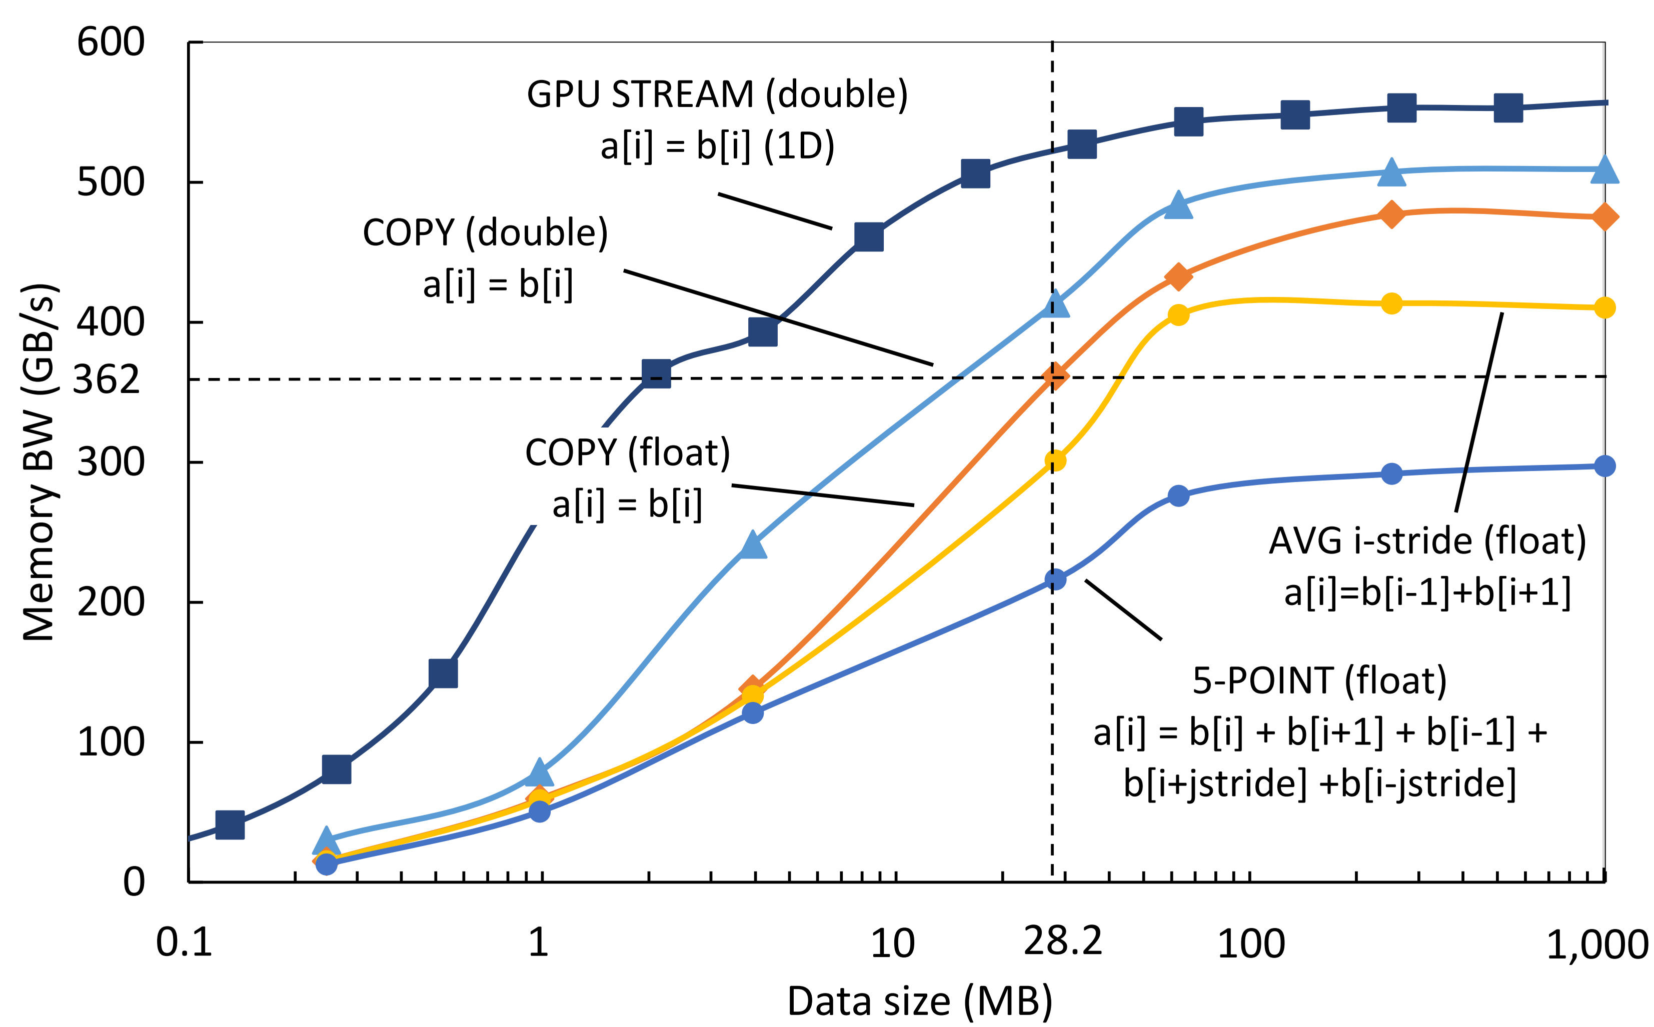

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

Necessary data transfers — 0.88 — Achieved BW

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \frac{Q}{D}\,\frac{B}{\hat{B}}$$

Actual data transfers — 0.76 — Max achievable BW

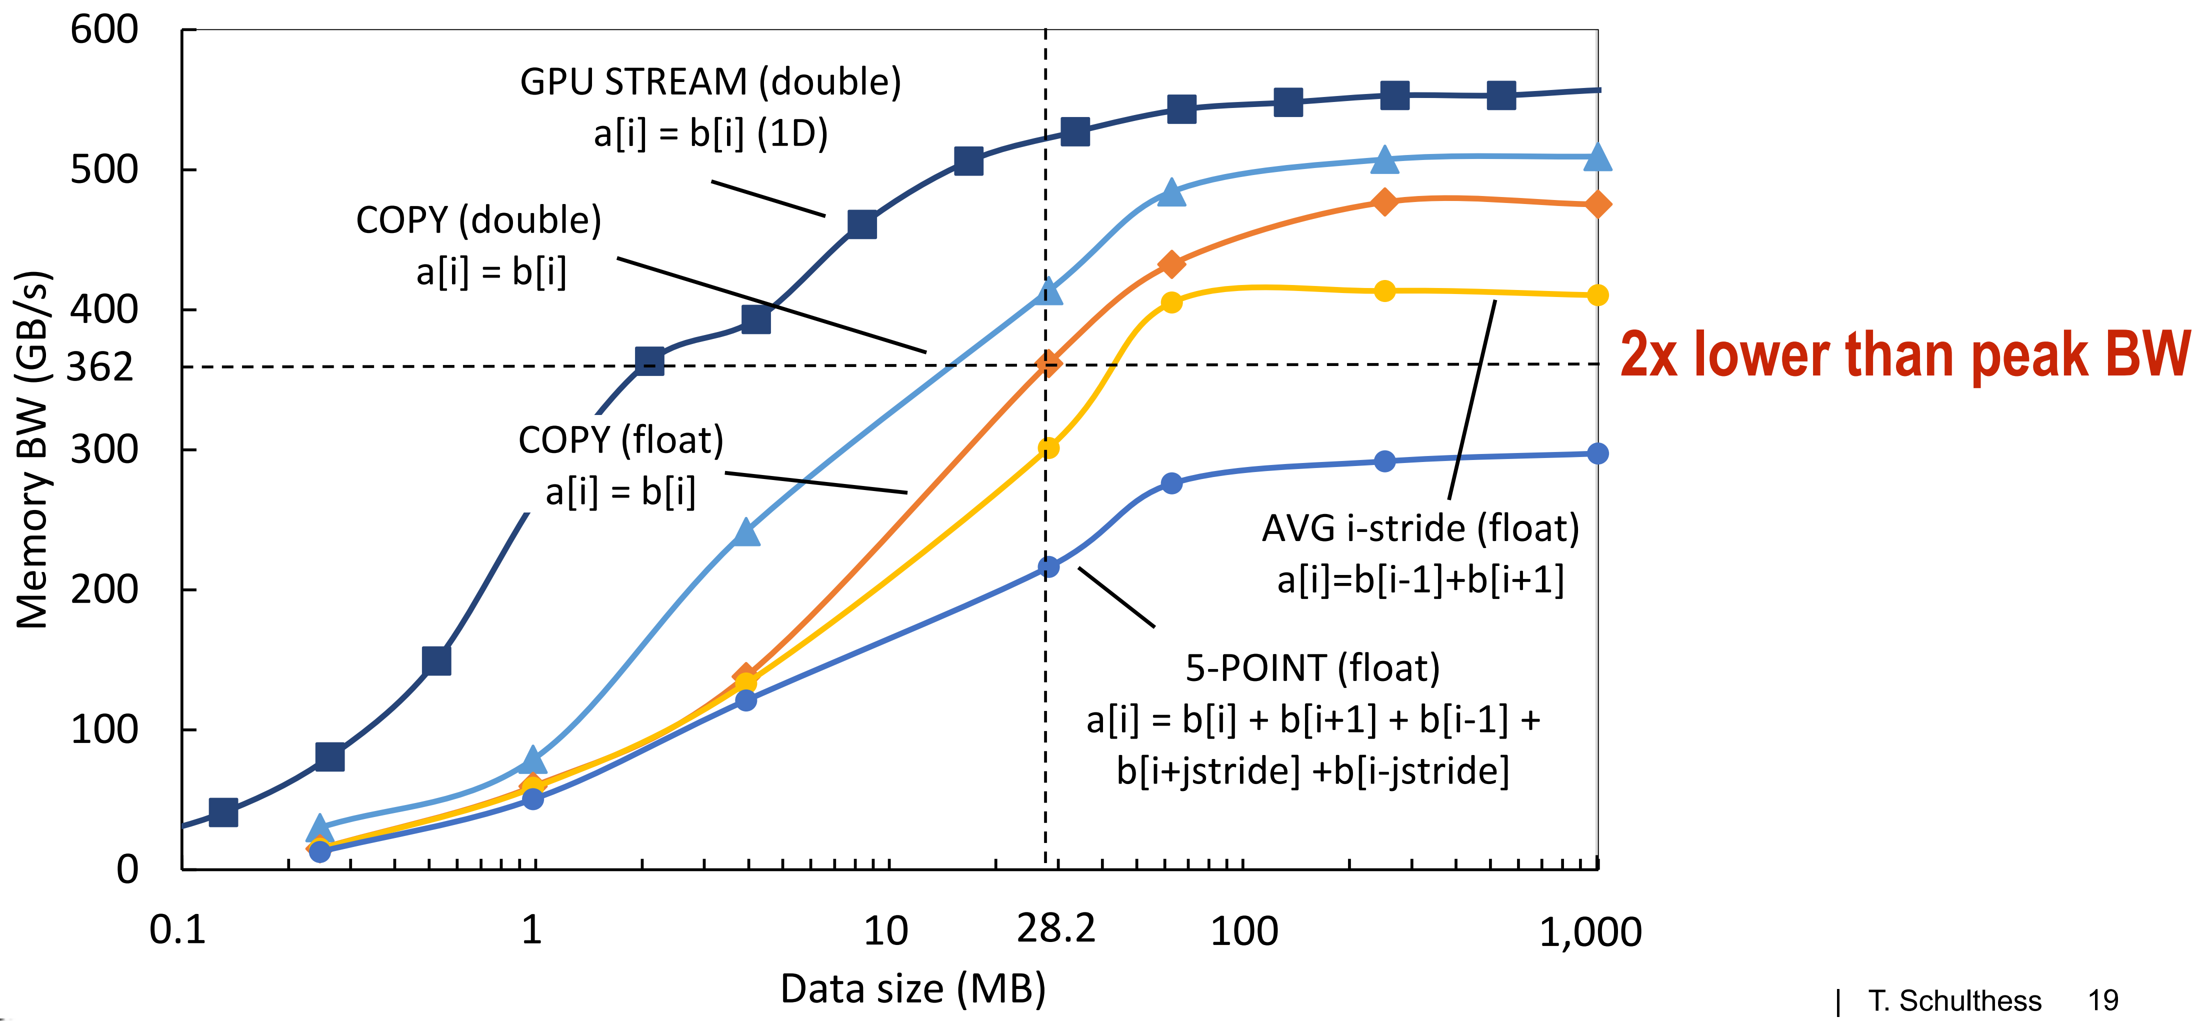

Memory use efficiency

Fuhrer et al., Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-230, published 2018

Necessary data transfers

Achieved BW

$$MUE = \text{I/O efficiency} \cdot \text{BW efficiency} = \overbrace{\frac{Q}{D}}^{0.88}\underbrace{\frac{B}{\hat{B}}}_{0.76} = 0.67$$

Actual data transfers

Max achievable BW

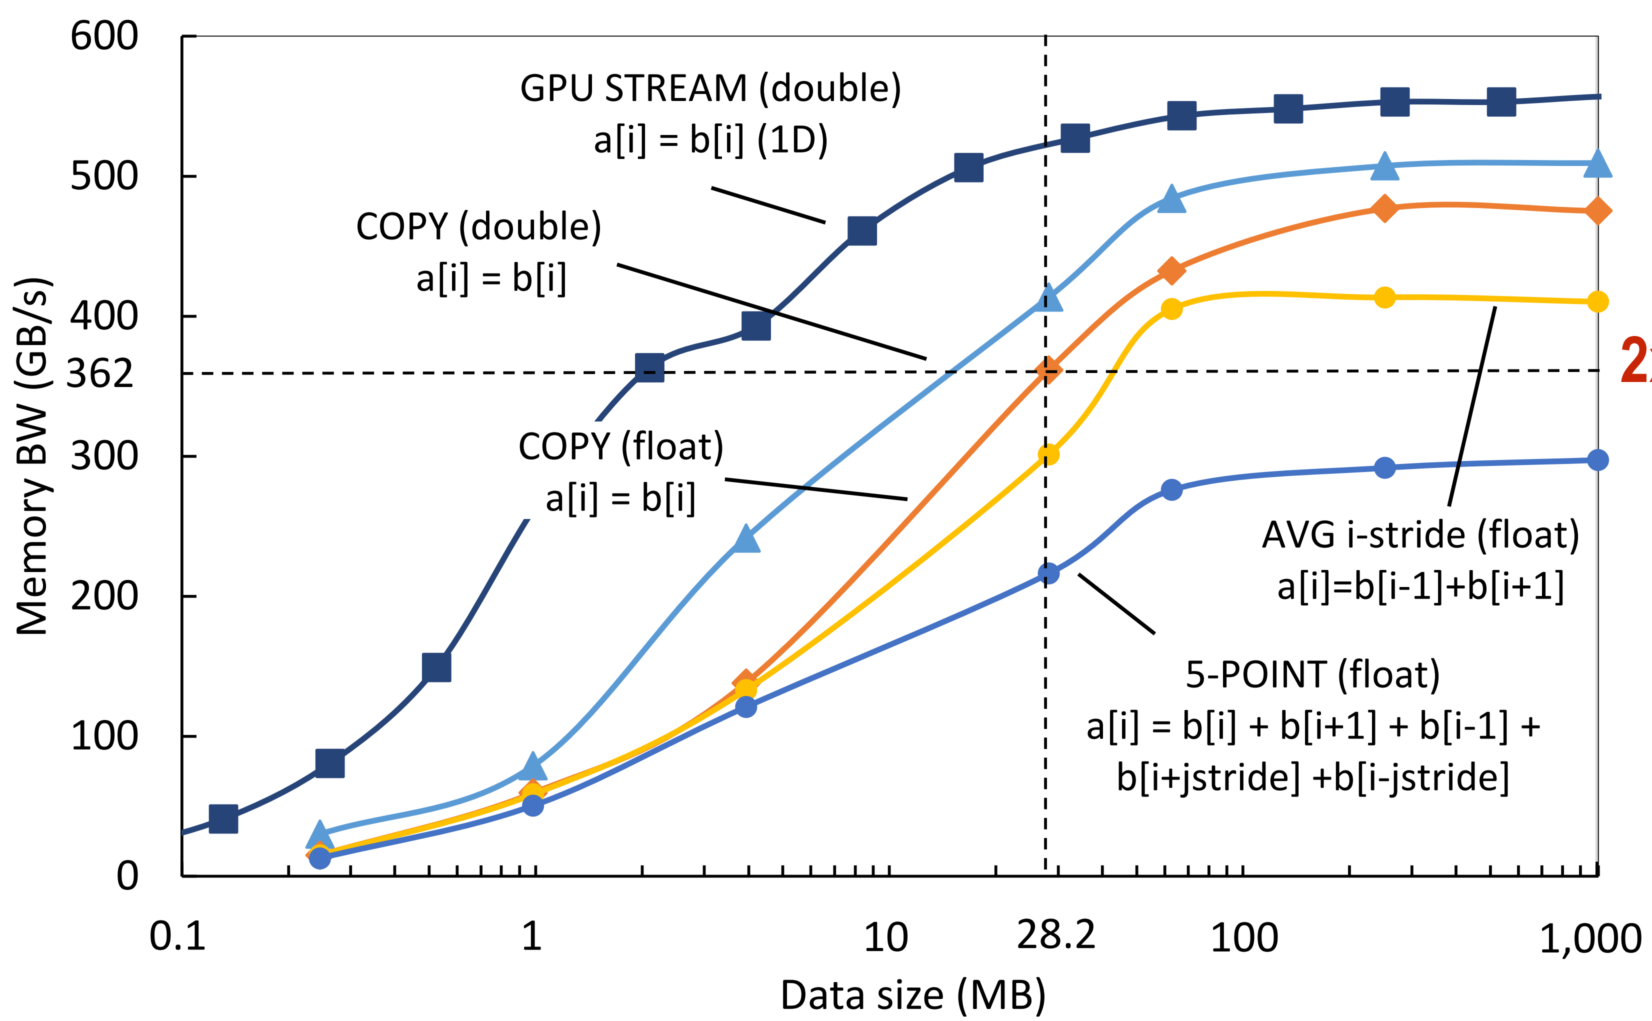

2x lower than peak BW

ETH zürich

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)

   2x fewer grid-columns

   Time step of 10 ms instead of 5 ms

4x

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)

   2x fewer grid-columns

   Time step of 10 ms instead of 5 ms

4x

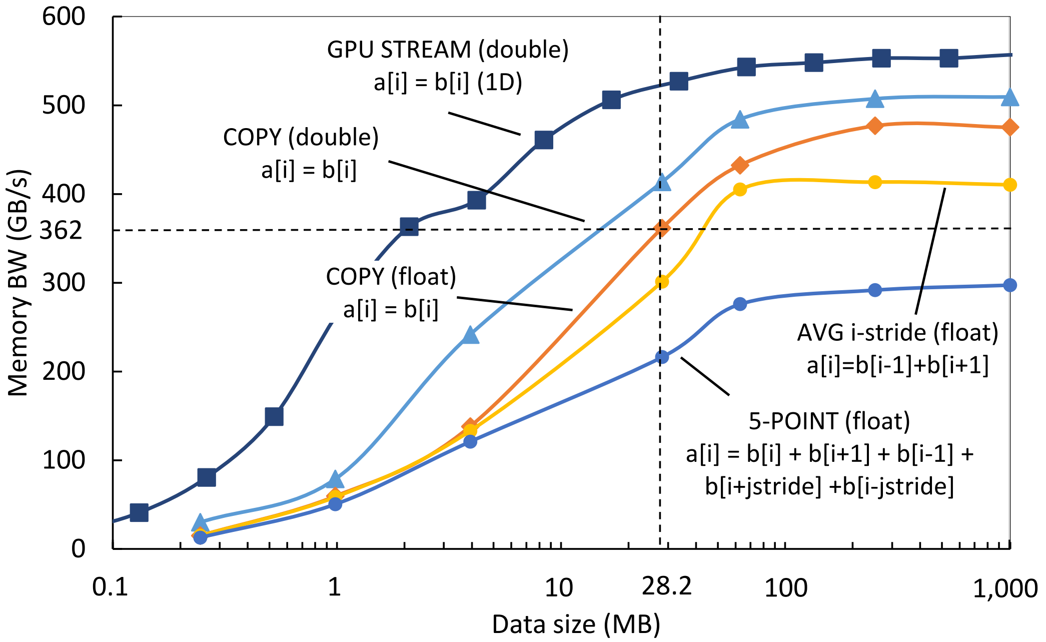

ETH zürich

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)

   2x fewer grid-columns

   Time step of 10 ms instead of 5 ms

   4x

2. Improving BW efficiency

   Improve BW efficiency and peak BW

   (results on Volta show this is realistic)

   2x

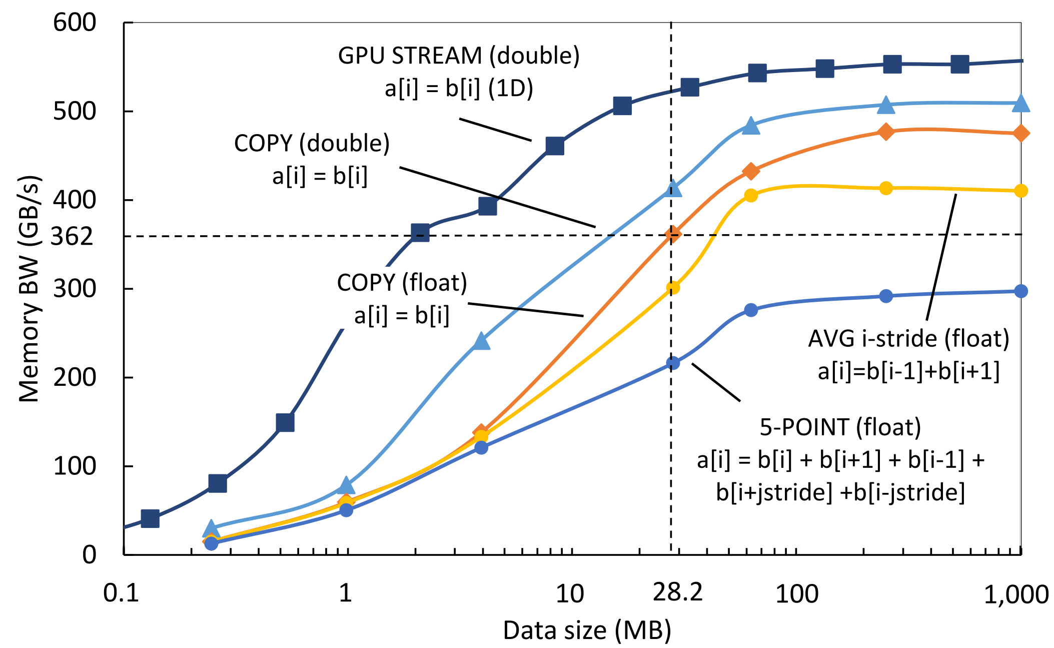

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)

   2x fewer grid-columns

   Time step of 10 ms instead of 5 ms

   4x

2. Improving BW efficiency

   Improve BW efficiency and peak BW

   (results on Volta show this is realistic)

   2x

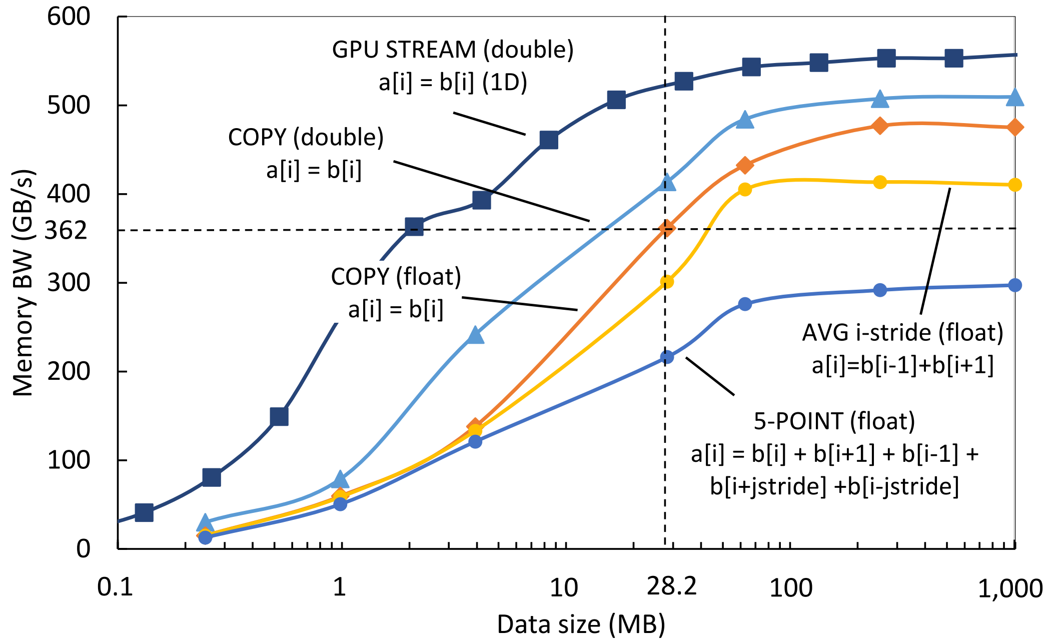

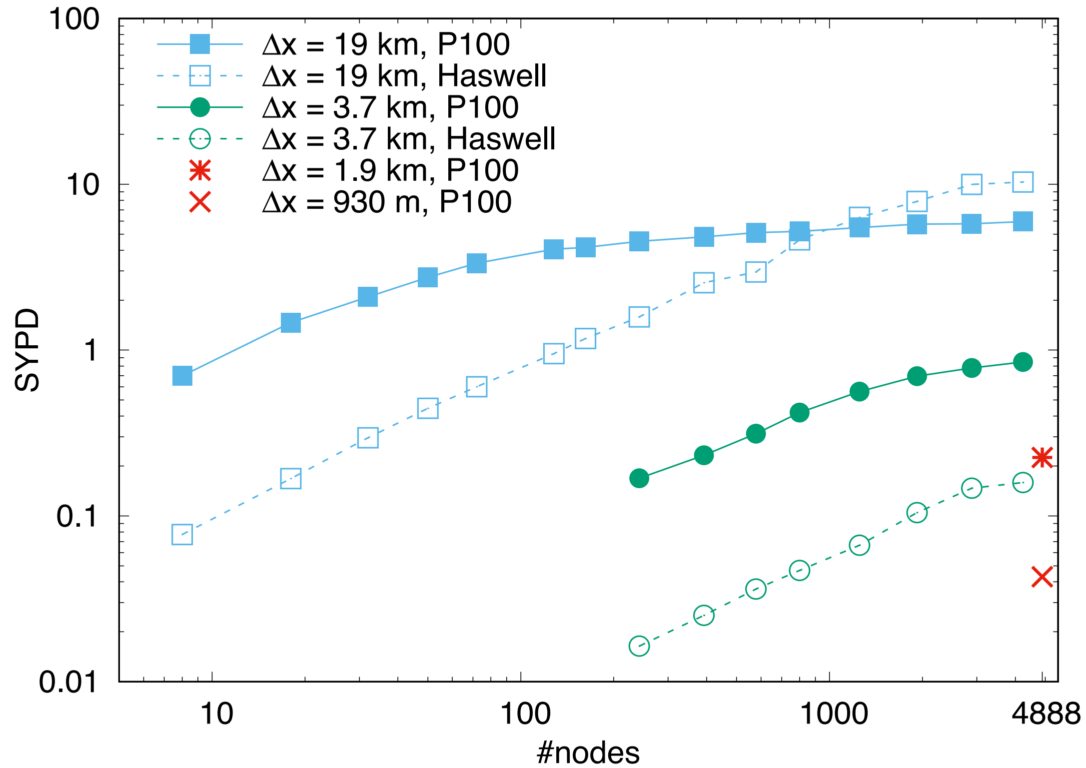

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)

   2x fewer grid-columns

   Time step of 10 ms instead of 5 ms

   4x

2. Improving BW efficiency

   Improve BW efficiency and peak BW

   (results on Volta show this is realistic)

   2x

3. Weak scaling

   4x possible in COSMO, but we reduced available parallelism by factor 2

   2x

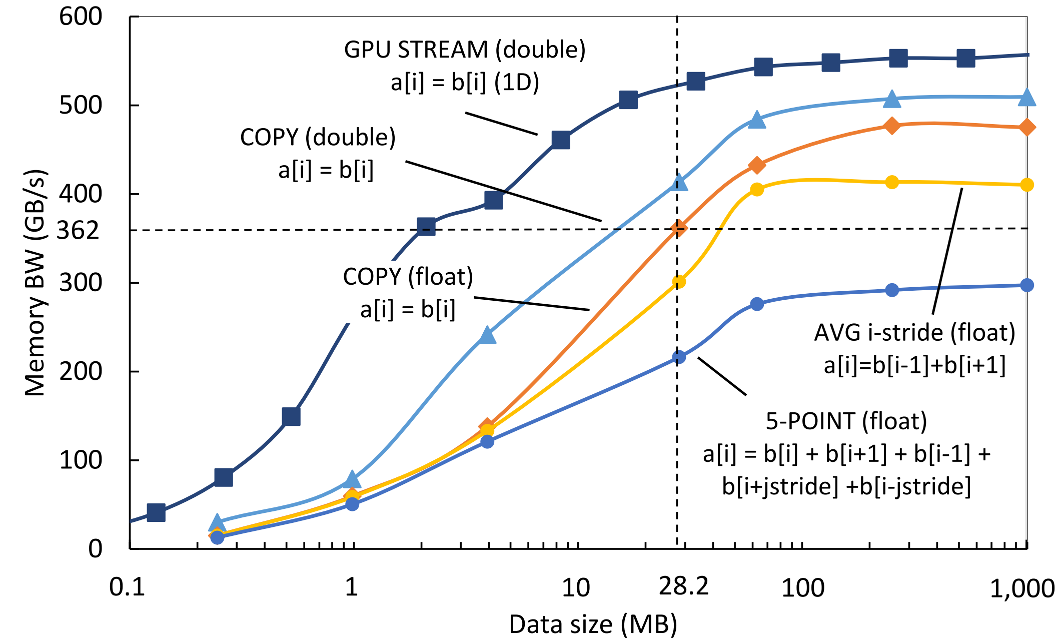

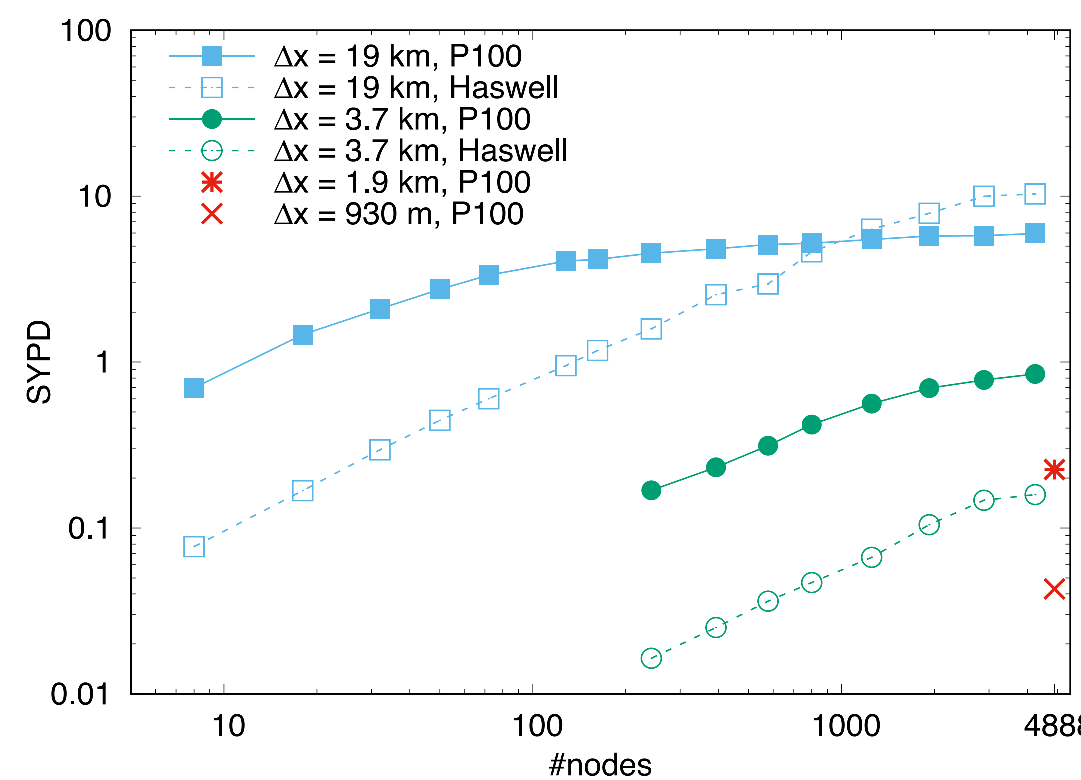

ETH zürich

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)

   2x fewer grid-columns

   Time step of 10 ms instead of 5 ms — **4x**

2. Improving BW efficiency

   Improve BW efficiency and peak BW (results on Volta show this is realistic) — **2x**

3. Weak scaling

   4x possible in COSMO, but we reduced available parallelism by factor 2 — **2x**

4. Remaining reduction in shortfall — **4x**

   Numerical algorithms (larger time steps)

   Further improved processors / memory

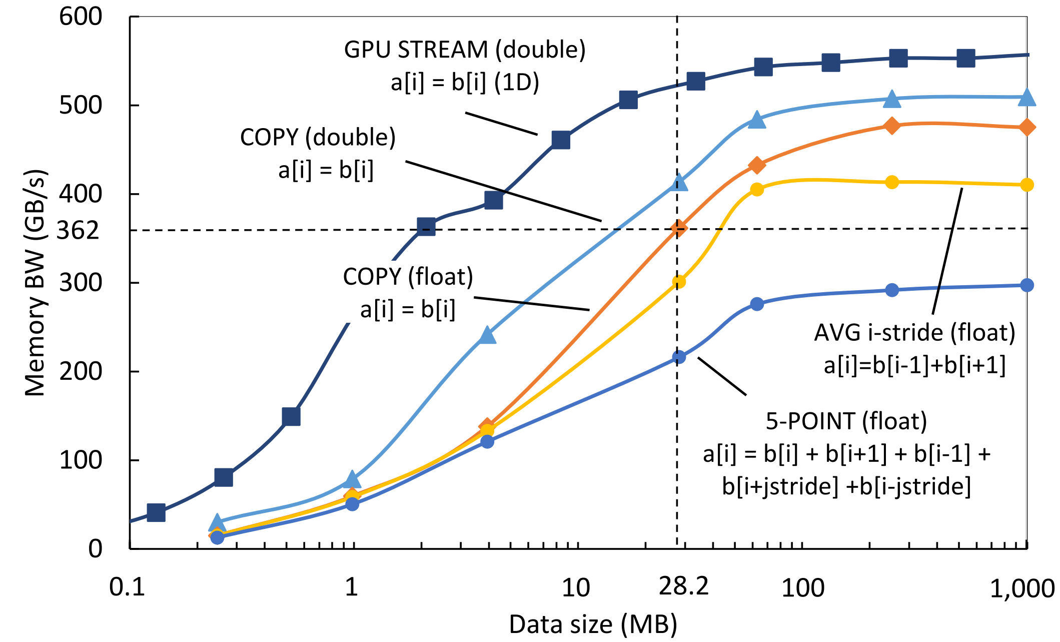

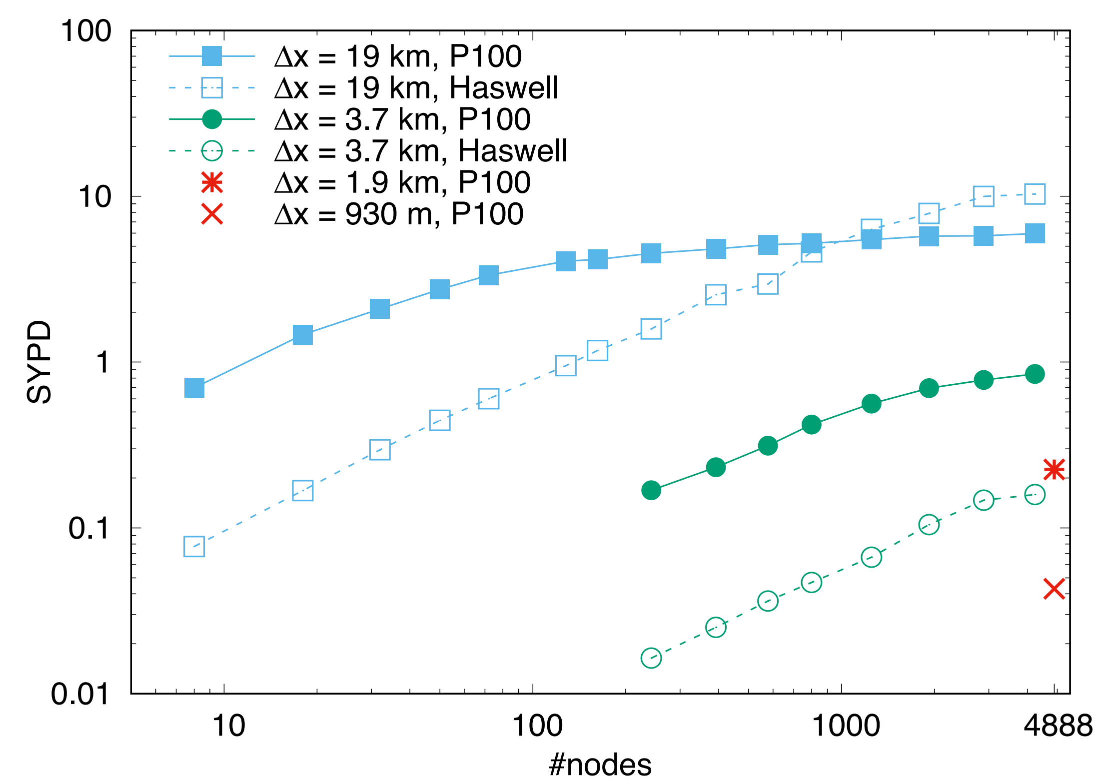

How realistic is it to overcome 65-fold shortfall of a grid-based implementation like COSMO-global?

1. Icosahedral grid (ICON) vs. Lat-long/Cartesian grid (COSMO)
   - 2x fewer grid-columns
   - Time step of 10 ms instead of 5 ms

   4x

2. Improving BW efficiency
   - Improve BW efficiency and peak BW (results on Volta show this is realistic)

   2x

3. Weak scaling
   - 4x possible in COSMO, but we reduced available parallelism by factor 2

   2x

4. Remaining reduction in shortfall

   4x

   - Numerical algorithms (larger time steps)
   - Further improved processors / memory

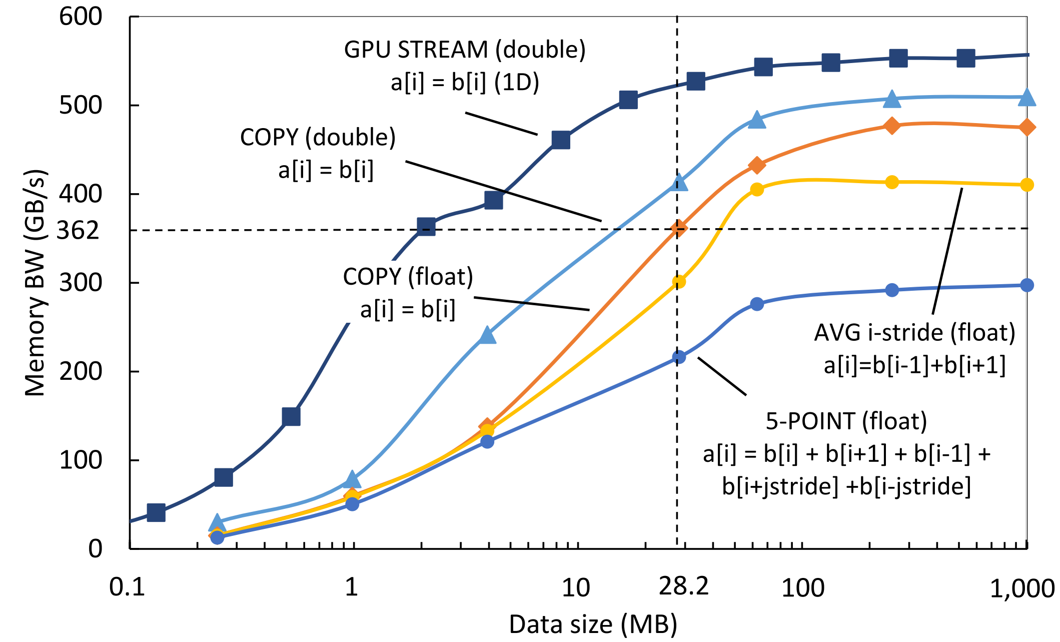

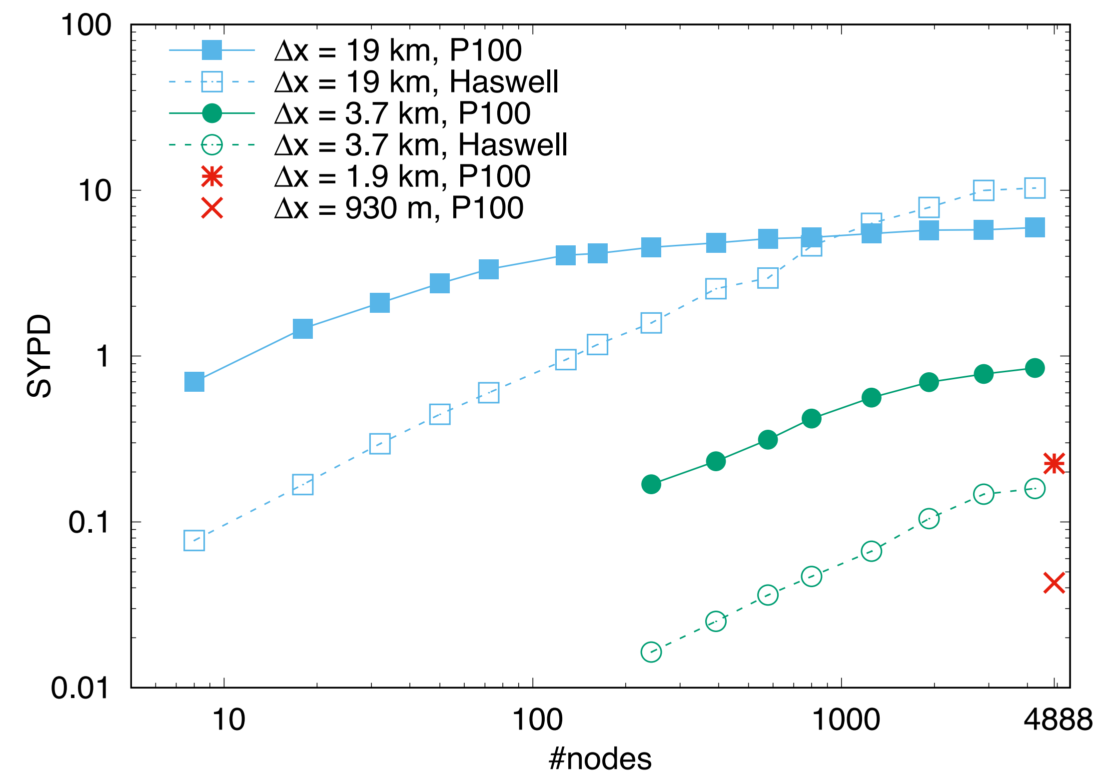

But we don't want to increase the footprint of the 2021 system beyond "Piz Daint"

ETH zürich

The importance of ensembles

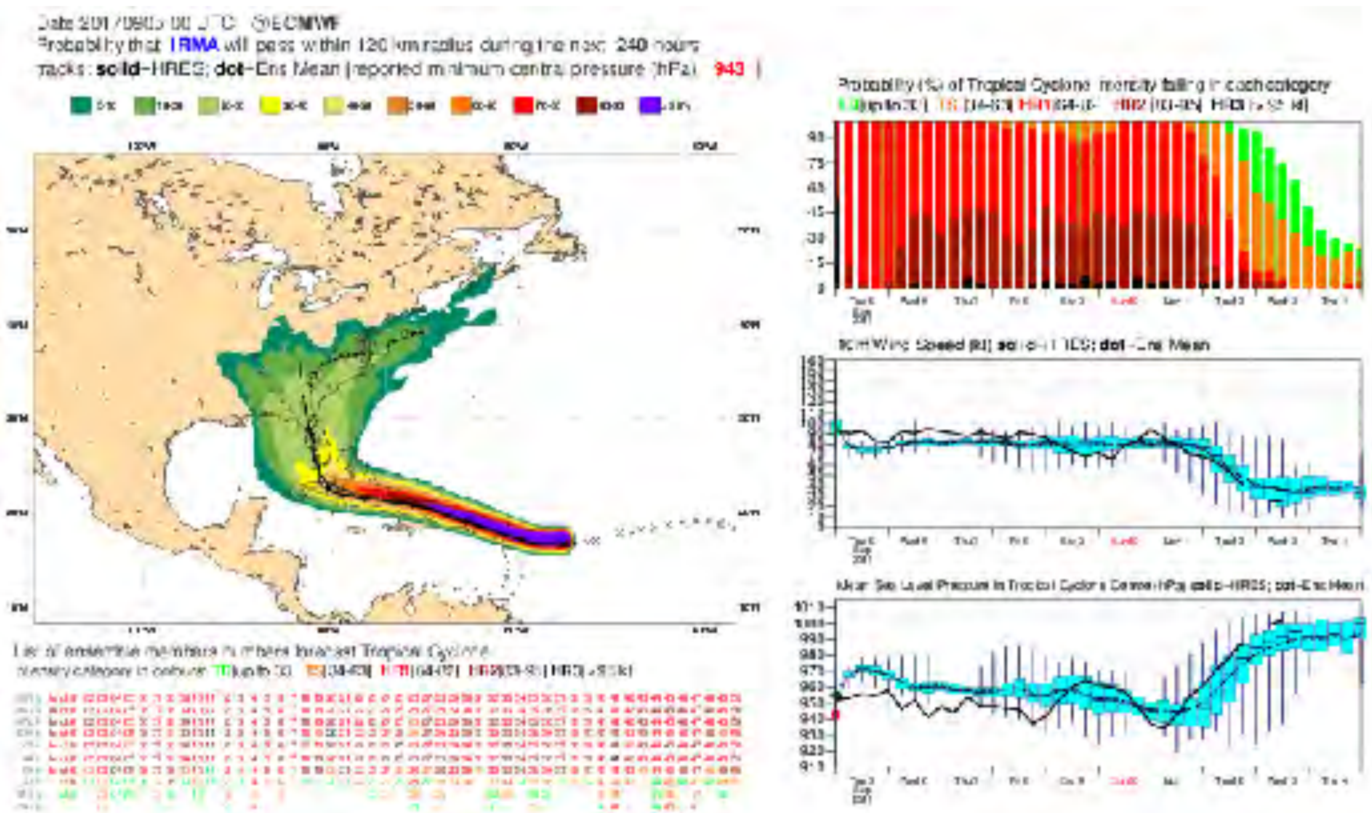

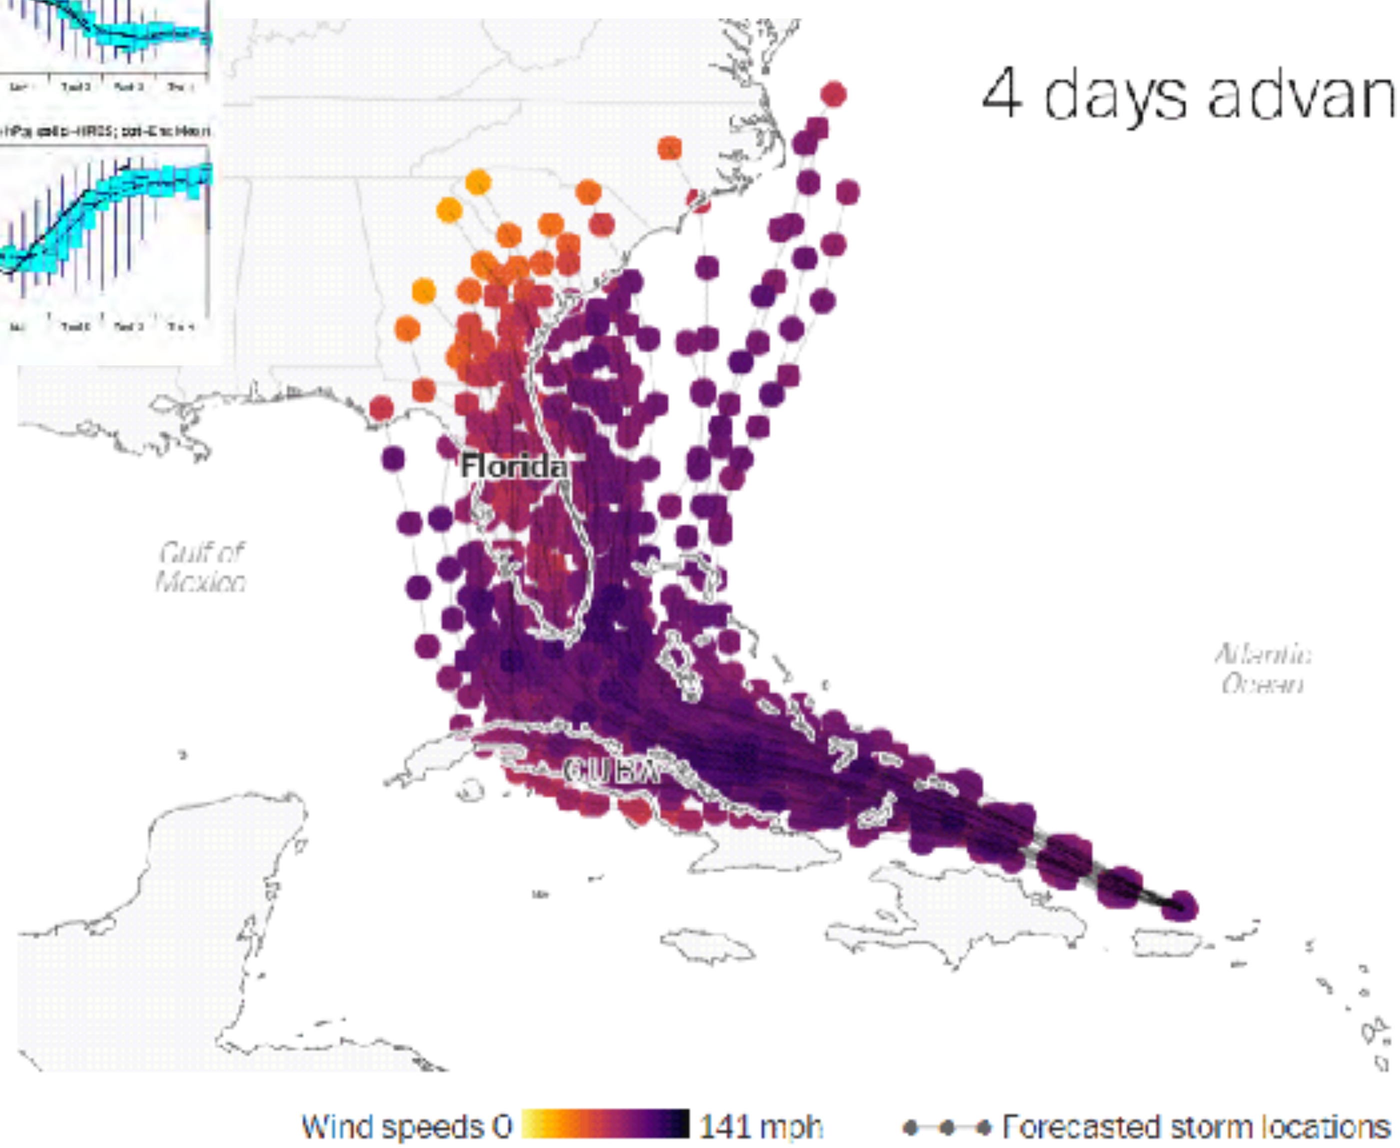

4 days advance

Peter Bauer, ECMWF

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

The importance of ensembles

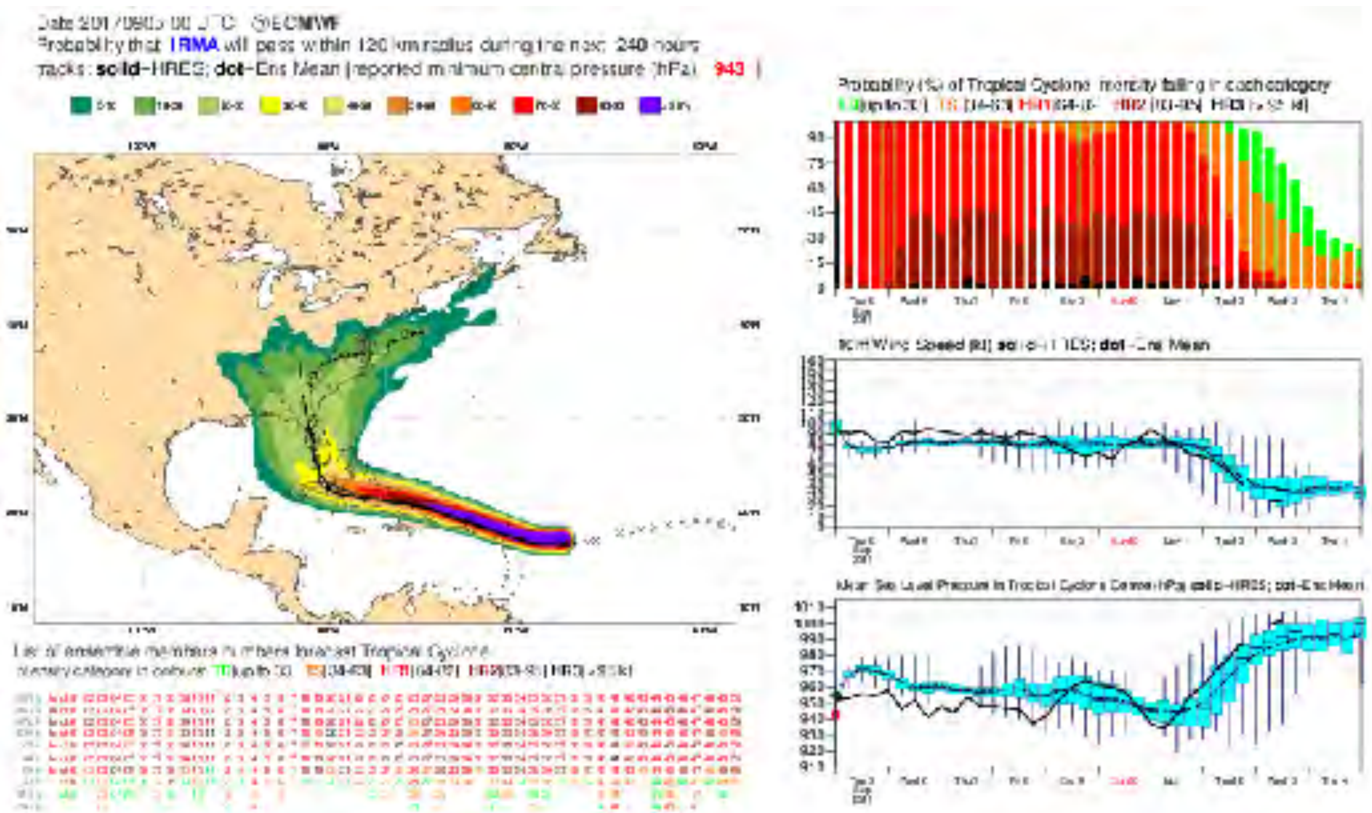

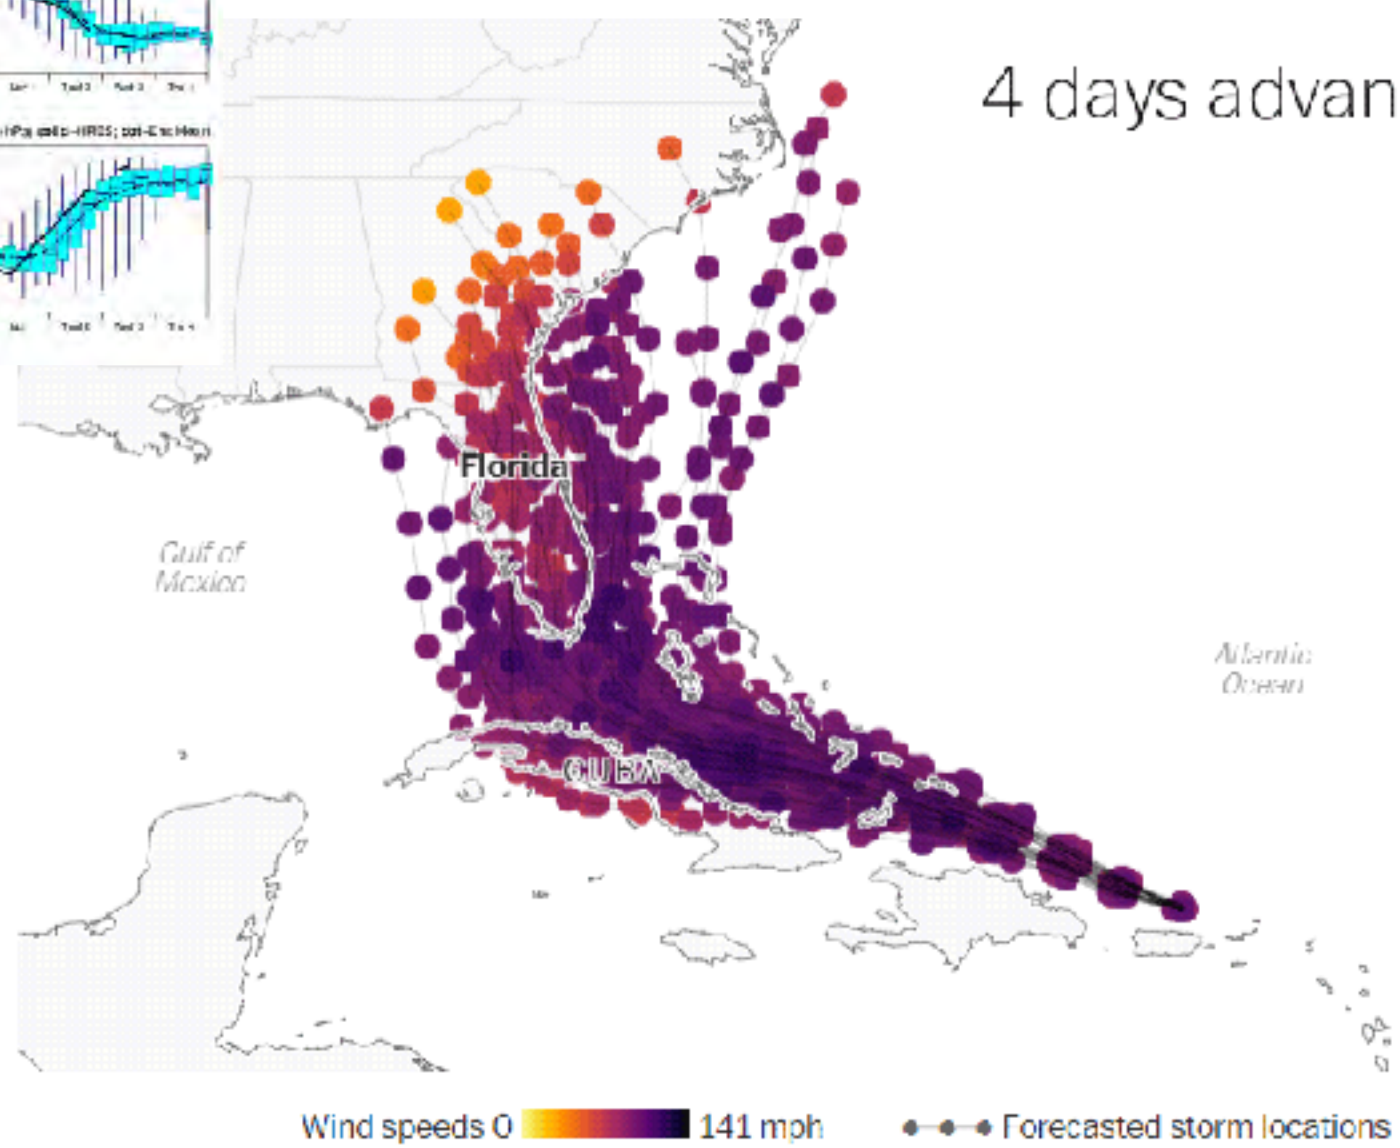

4 days advance

Peter Bauer, ECMWF

Remaining goals beyond 2021 (by 2024?)

1. Improve the throughput to 5 SYPD

2. Reduce the footprint of a single simulation by up to factor 10

Data challenge: what is all models run at 1 km scale?

With current workflow used by the climate community,
IPPC at 1 km scale avg. would produce 50 exabytes of data!

The only way out would be to analyse the model while the simulation is running

New workflow:

1. first set of runs generate model trajectories and checkpoints
2. Reconstruct model trajectories from checkpoints and analyse (as often as necessary)

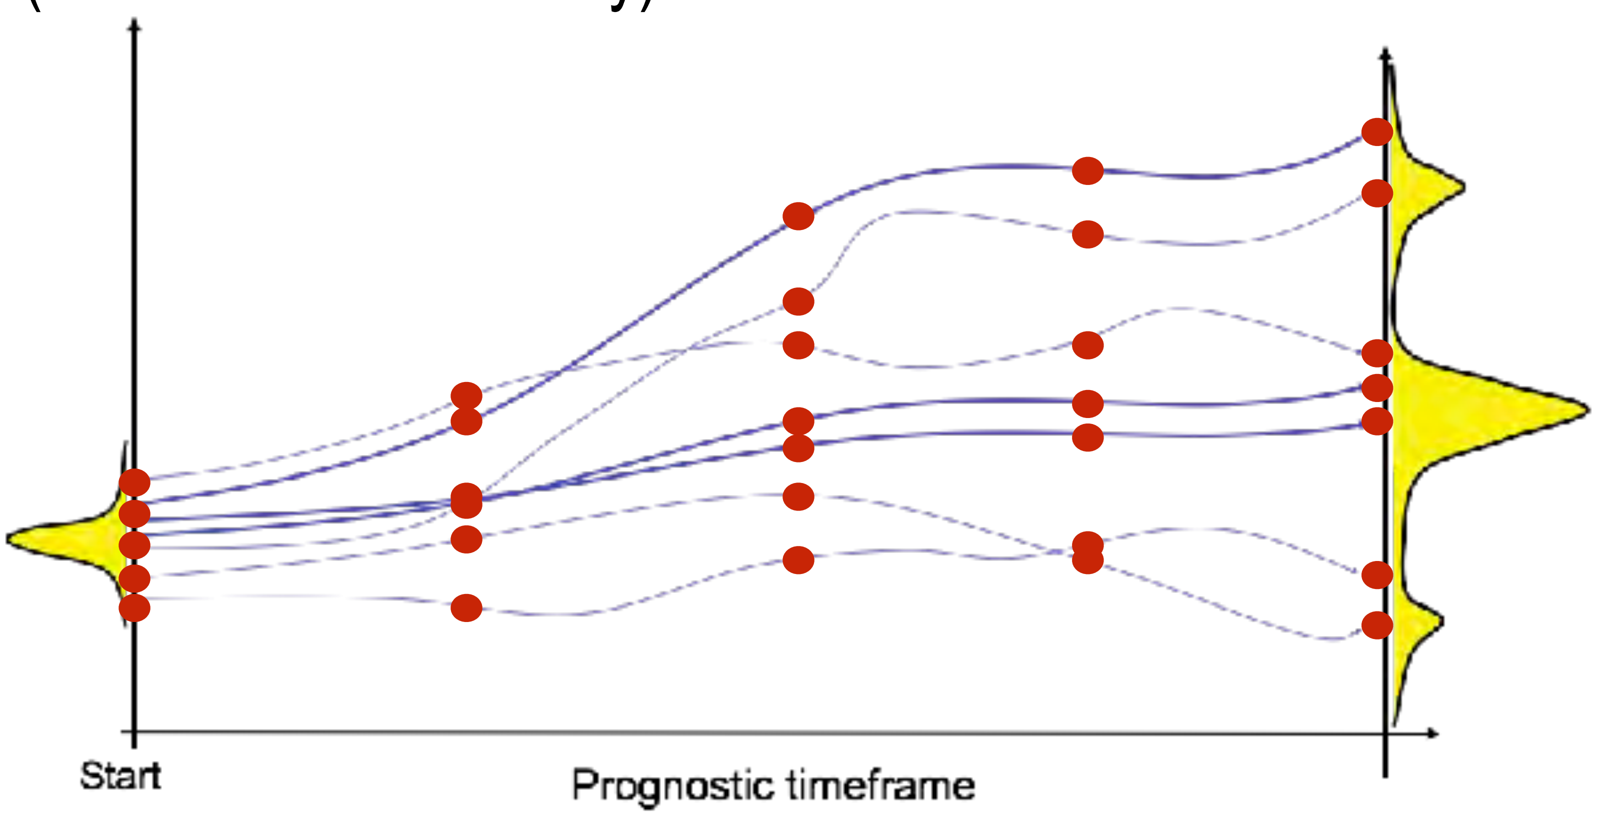

Summary and Conclusions

Summary and Conclusions

- While flop/s may be good to compare performance in history of computing, it is not a good metric to design the systems of the future

Summary and Conclusions

- While flop/s may be good to compare performance in history of computing, it is not a good metric to design the systems of the future
- Given todays challenges use Memory Use Efficiency (MUE) instead

MUE = “I/O Efficiency” X “Bandwidth Efficiency”

Summary and Conclusions

- While flop/s may be good to compare performance in history of computing, it is not a good metric to design the systems of the future
- Given todays challenges use Memory Use Efficiency (MUE) instead

  MUE = “I/O Efficiency” X “Bandwidth Efficiency”

- Convection resolving weather and climate simulations @ 1 km horizontal resolution
  - Represents a big leap for quality of weather and climate simulations
  - Aim is to pull in the milestone by a decade, in the early 2020s rather than 2030s
  - Desired throughput of 1 SYPD by 2021 and 5 SYPD by mid 2020s
  - Need a 50-200x performance improvement to reach 2021 goal of 1SYPD

Summary and Conclusions

- While flop/s may be good to compare performance in history of computing, it is not a good metric to design the systems of the future
- Given todays challenges use Memory Use Efficiency (MUE) instead

  MUE = “I/O Efficiency” X “Bandwidth Efficiency”
- Convection resolving weather and climate simulations @ 1 km horizontal resolution
  - Represents a big leap for quality of weather and climate simulations
  - Aim is to pull in the milestone by a decade, in the early 2020s rather than 2030s
  - Desired throughput of 1 SYPD by 2021 and 5 SYPD by mid 2020s
  - Need a 50-200x performance improvement to reach 2021 goal of 1SYPD
- System improvement needed
  - Improve “Bandwidth Efficiency” for regular but non-stride-1 memory access
  - Improve application software
  - Further improvement of scalability
  - Improve algorithms (time integration has potential)

Collaborators

Tim Palmer (U. of Oxford)

Bjorn Stevens (MPI-M)

Peter Bauer (ECMWF)

Oliver Fuhrer (MeteoSwiss)

Nils Wedi (ECMWF)

Torsten Hoefler (ETH Zurich)

Christoph Schar (ETH Zurich)