Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology

Challenges in Convection-Resolving Climate Modeling

Christoph Schär[1], David Leutwyler[1], Nikolina Ban[1], Oliver Fuhrer[2], Michael Keller[1,3], Torsten Hoefler[4], Xavier Lapillonne[2,3], Daniel Lüthi[1], Linda Schlemmer[1], Thomas Schulthess[5], Heini Wernli[1]

[1]Atmospheric and Climate Science, ETH Zürich
[2]MeteoSwiss, Zürich
[3]Center for Climate Systems Modeling (C2SM), ETH Zürich
[4]Computer Science, ETH Zürich
[5]Swiss Center for Scientific Computing (CSCS), Lugano
http://www.iac.ethz.ch/people/schaer

GEWEX Workshop, 6th September 2016, Boulder, CO,USA

Aquaplanet

Cloud fraction [%]

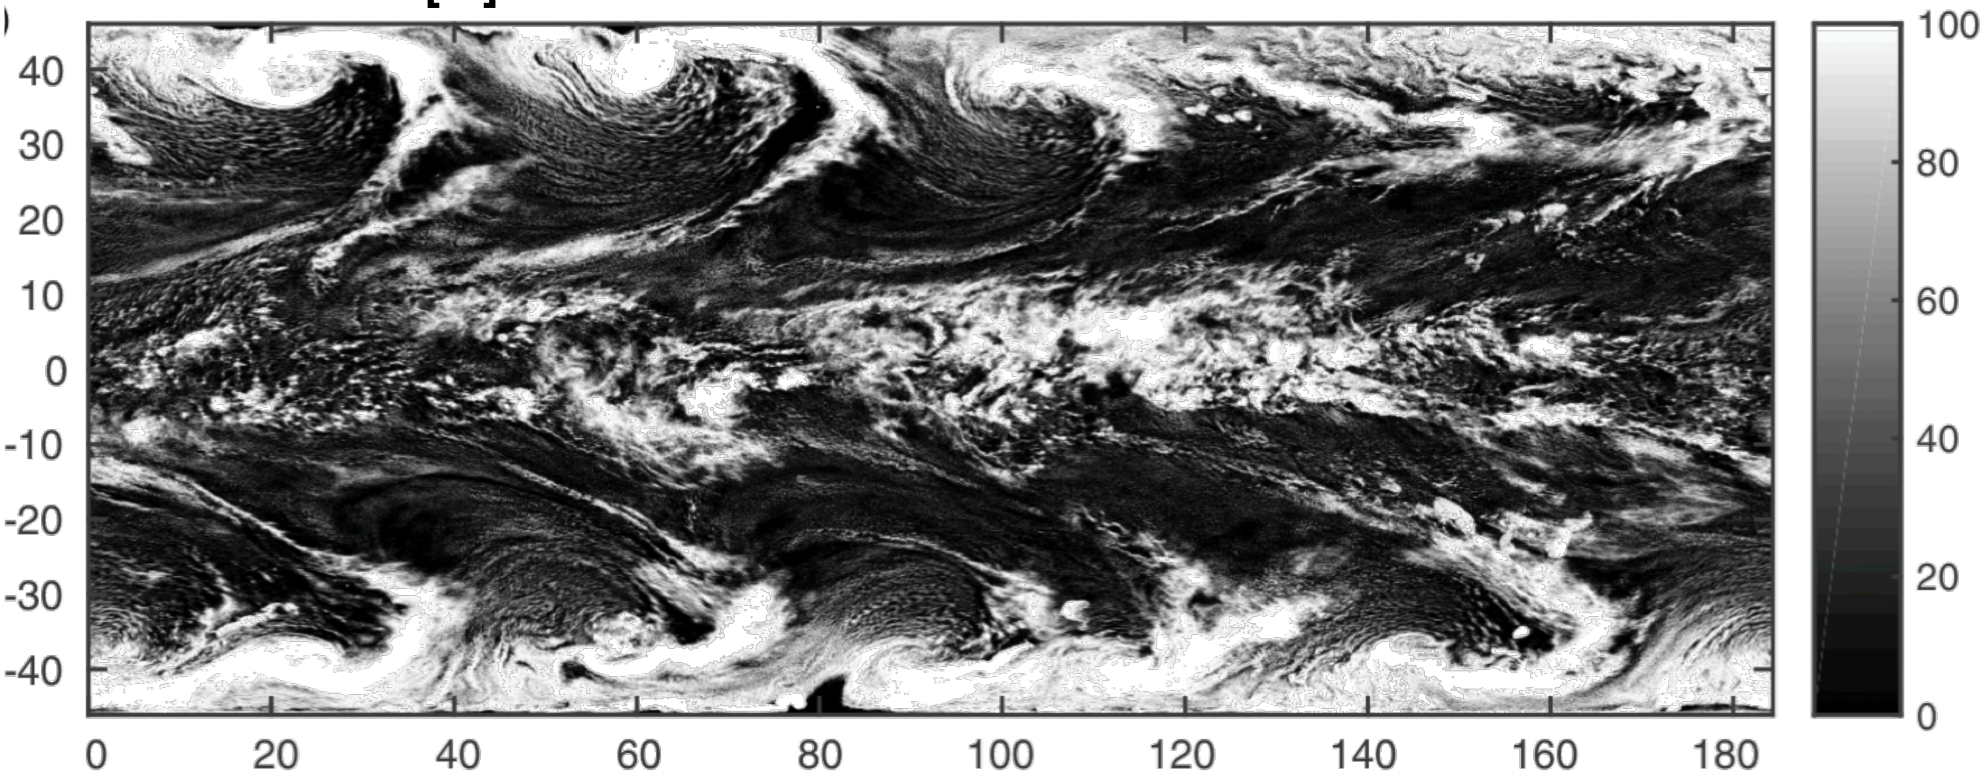

Bretherton and Khairoutdinov (2015)

Δ=4 km, run for O(weeks)
SAM model, anelastic
near global aquaplanet (20,480 x 10,240 km)

LES simulations over Netherlands

Simulated cloud field on July 6, 2004

Δx=100 m (LES)

4000x4000x300 gridpoints

3 cases simulated

Domain decomposition using 256 GPUs

Very limited set of parameterizations
(e.g. no radiation)

Periodic lateral boundary conditions
non-rotating framework (f=0)

Observed

(Schalkwijk et al., 2015, BAMS)

Tropical Atlantic

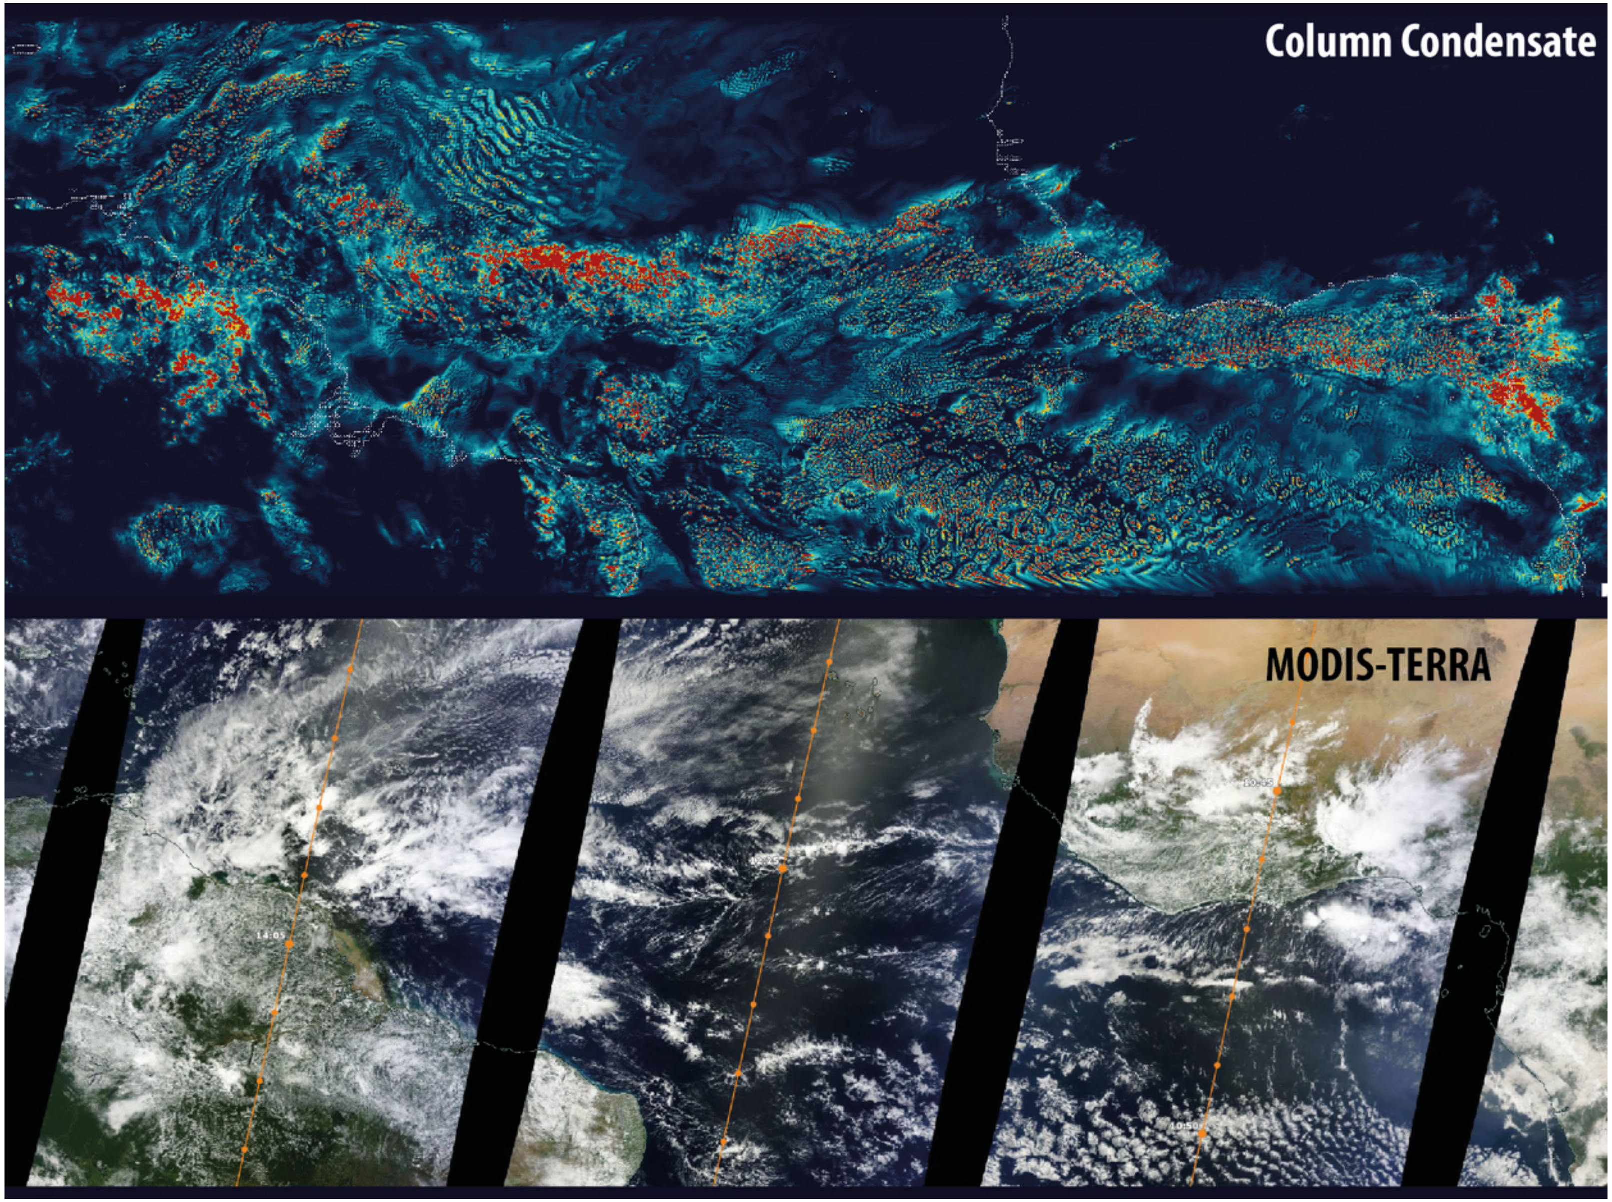

MPI Hamburg
(Stevens et al.)

Δ=2.5 km
Run for O(days)
Tropical Atlantic
ICON Model

European-scale simulations

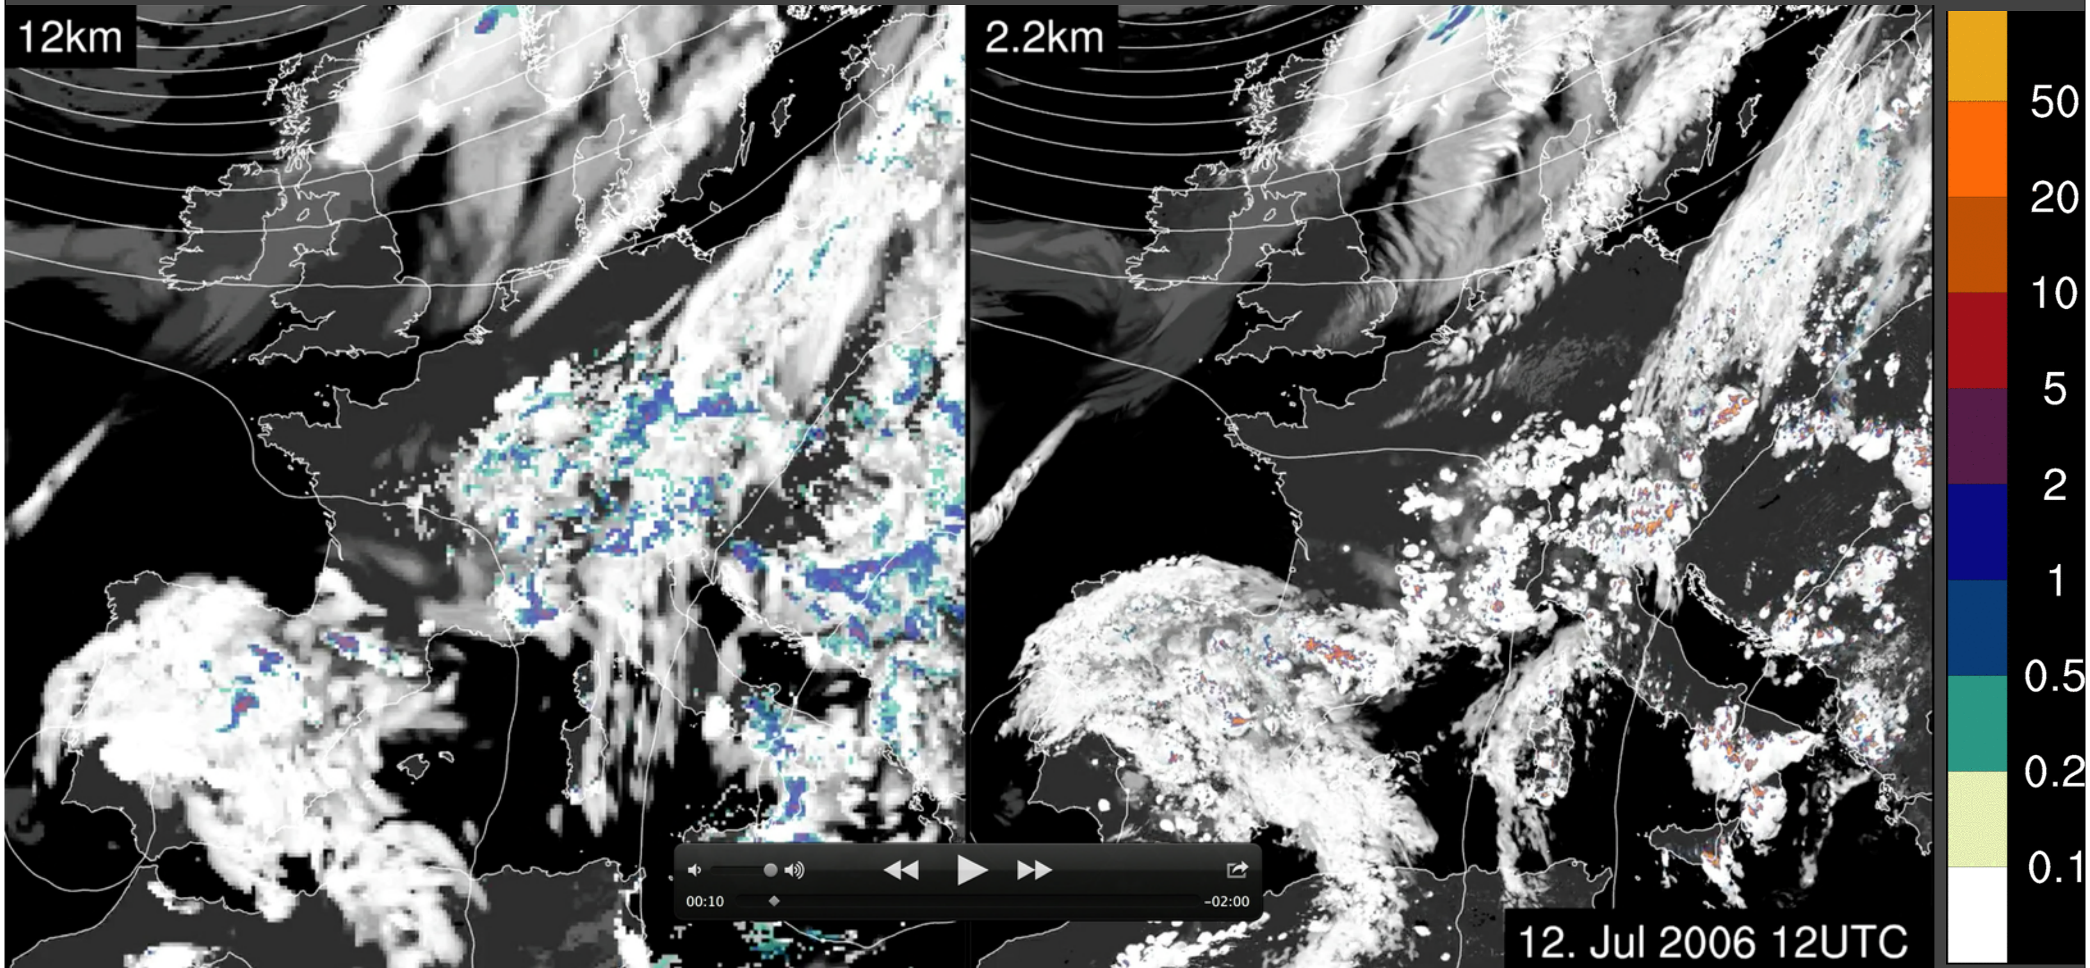

[mm/h]

Schär, ETH Zürich

David Leutwyler, ETH Zurich, animations via crCLIM: http://www.c2sm.ethz.ch/research/crCLIM

Contents

Convection over lake Constance

- Switch off convection parameterization
- More closely based on first principles

April 27, 2016, Lake Constance (Kurt Abderhalden)

Convection and flash-flooding

Flashflood in Wil (Switzerland)
June 15, 2015

Flooding in
St. Gingolphe, Valais
May 1, 2015

Convection and debris flows

Pierre Zufferey, https://youtu.be/0ENe7wDKP6I

Debris flow
July 22, 2016

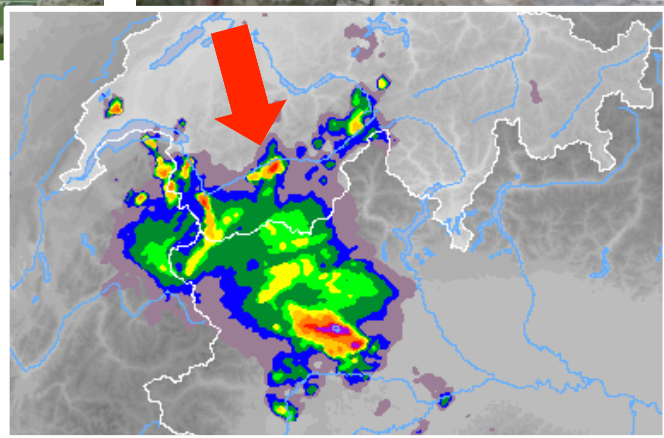

Triggered by a convective cell

Representation of feedback processes

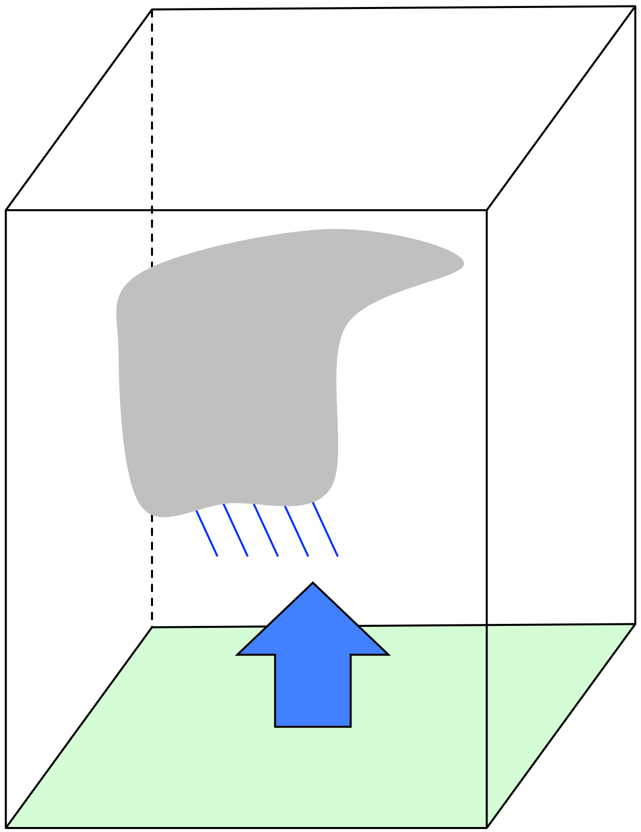

Parameterized convection:
Dominated by
vertical exchange
(column view of feedbacks)

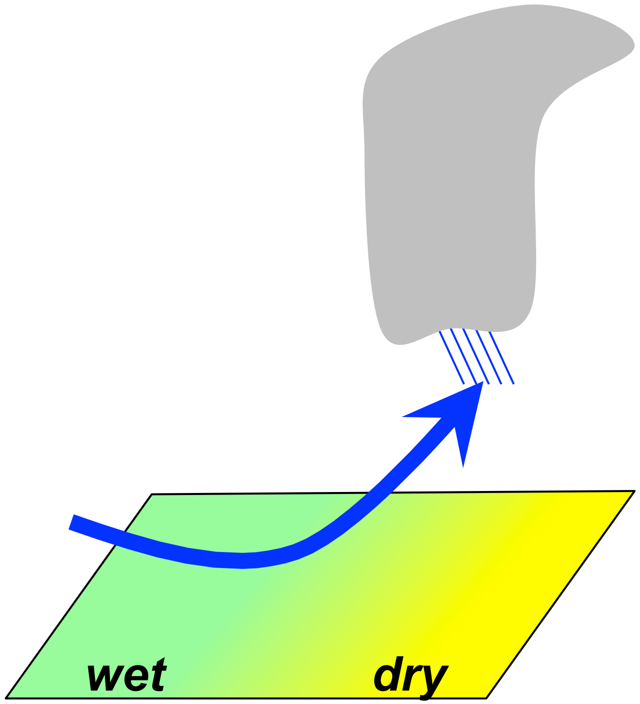

Explicit convection:
Three-dimensional meso-scale
circulations matter
(3D view of feedbacks)

(e.g. SMP feedback; Hohenegger et al. 2009, Taylor et al. 2011, Froidevaux et al. 2013)

Diurnal convection over Europe

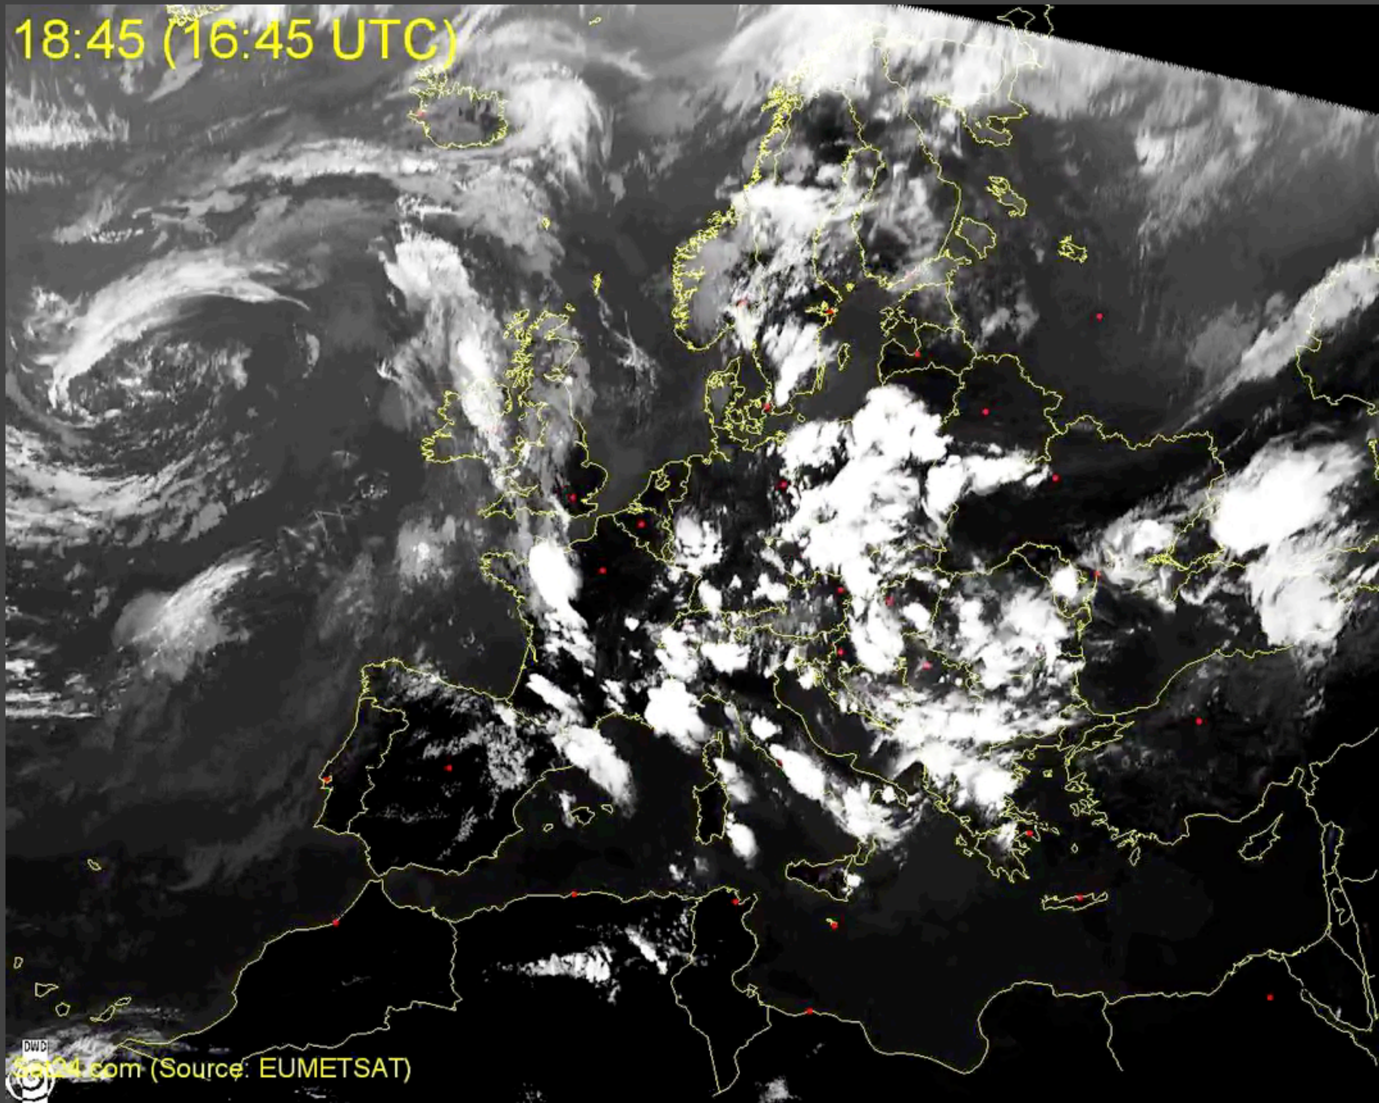

Schär, ETH Zürich

(SEVIRI 10.8µm, June 30 till July 2, 2009; Michael Keller, ETH Zürich)

Contents

Climate simulations at km-scale

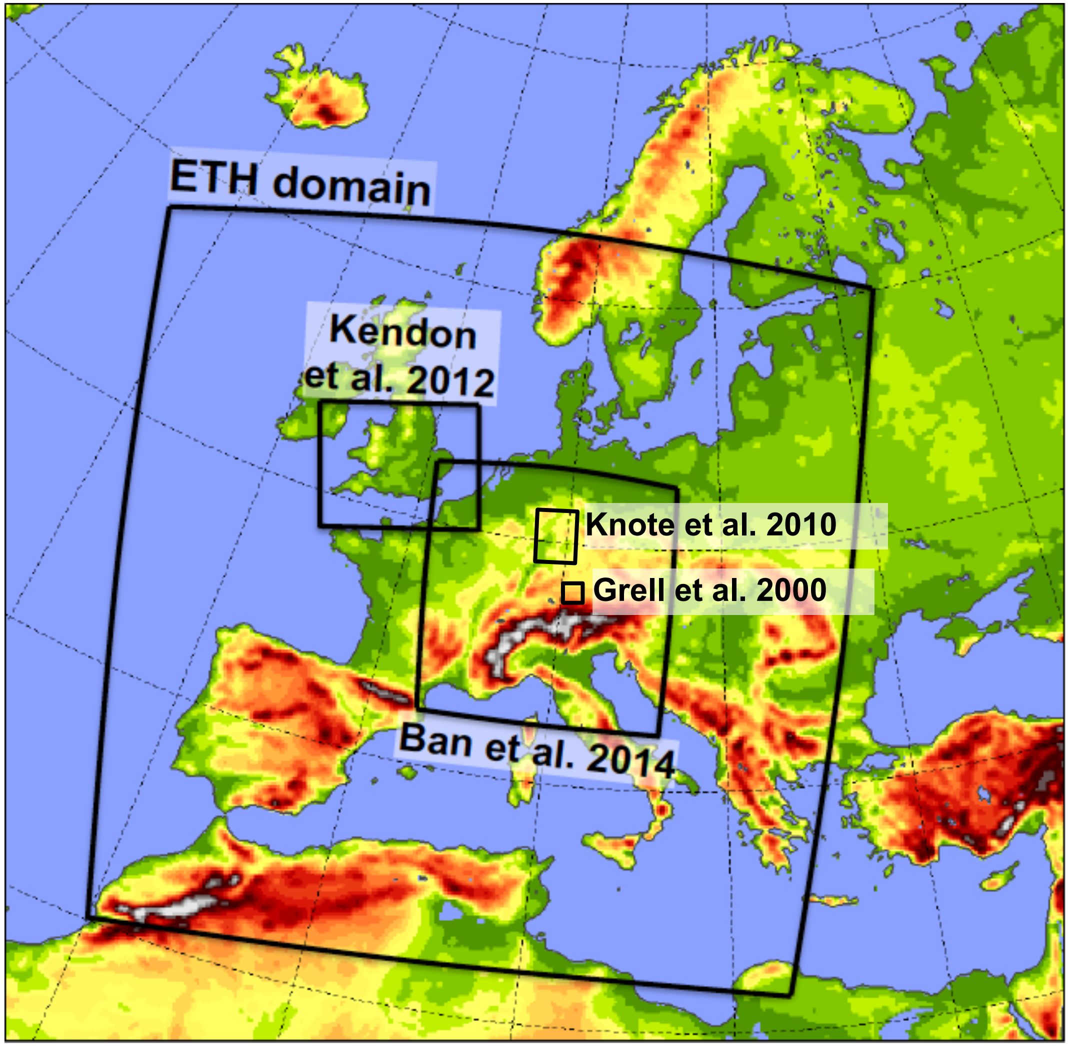

Grell et al. 2000:
46x46 gridpoints at 1 km
14 months

Knote et al. 2010:
ca 200 x 150 gridpoints at 1.3 km
several decades

Kendon et al. 2012:
ca 400 x 300 gridpoints at 1.5 km
several decades

This presentation:

Ban et al. (2014, JGR; 2015 GRL):
500 x 500 gridpoints at 2.2 km
several decades
driven by ERA-Interim and
MPI-ESM-LR (RCP 8.5)

Leutwyler et al. (2016, GMD, in press):
1536 x 1536 x 60 gridpoints at 2.2 km
one decade completed
driven by ERA-interim

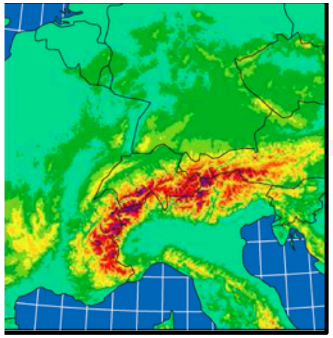

Alpine domain
2.2km (500x500x60)

Validation of diurnal cycle

10-year long simulation driven by ERA-Interim;
Validation against 62 rain-gauge stations in Switzerland (JJA)

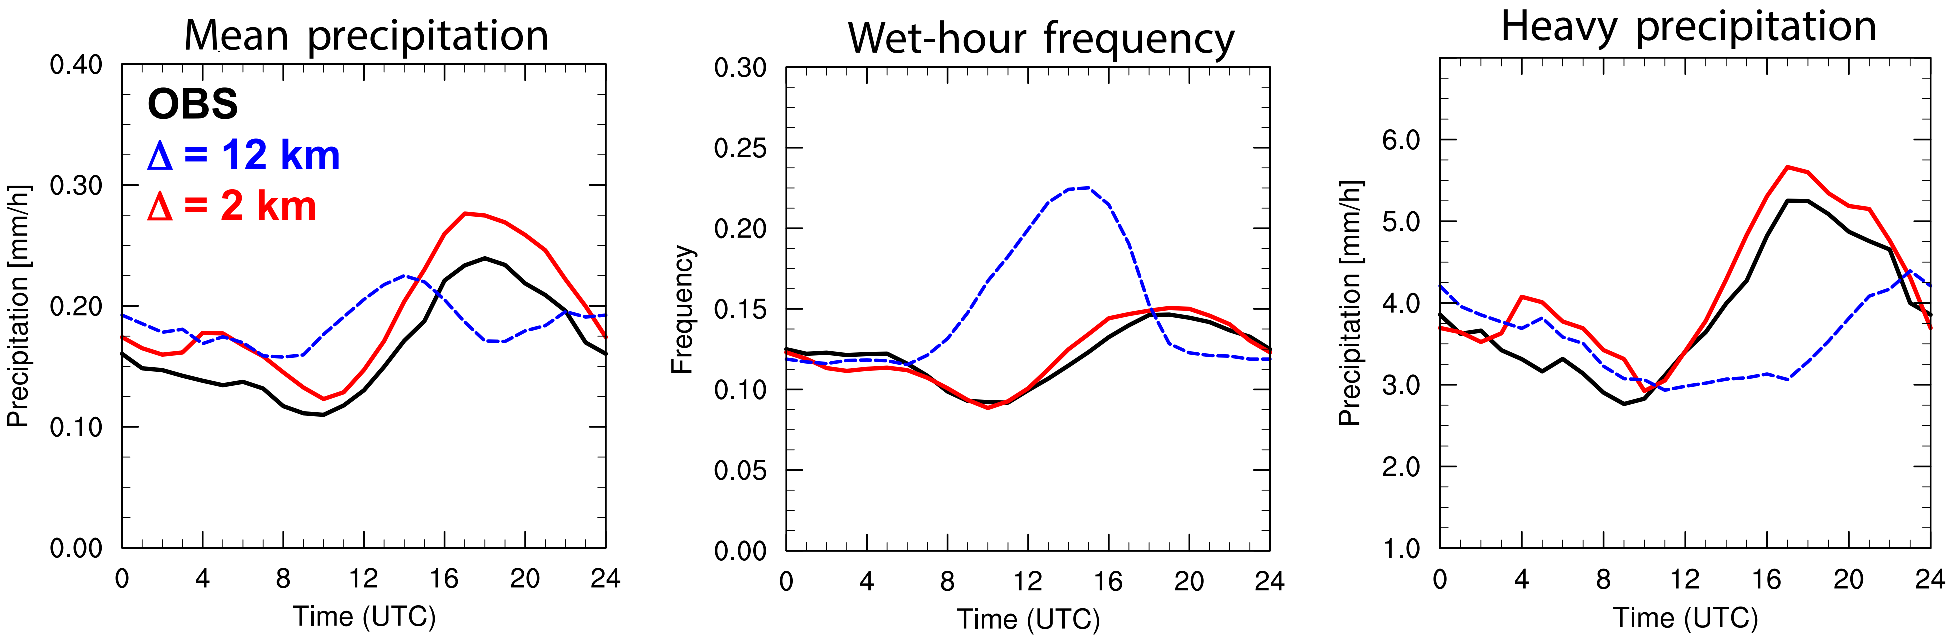

poor representation of diurnal cycle with Δ=12 km
dramatic improvement with Δ=2 km

(Ban et al. 2015, GRL)

Domain size matters

The statistics of convective cell needs to develop within computational domain!

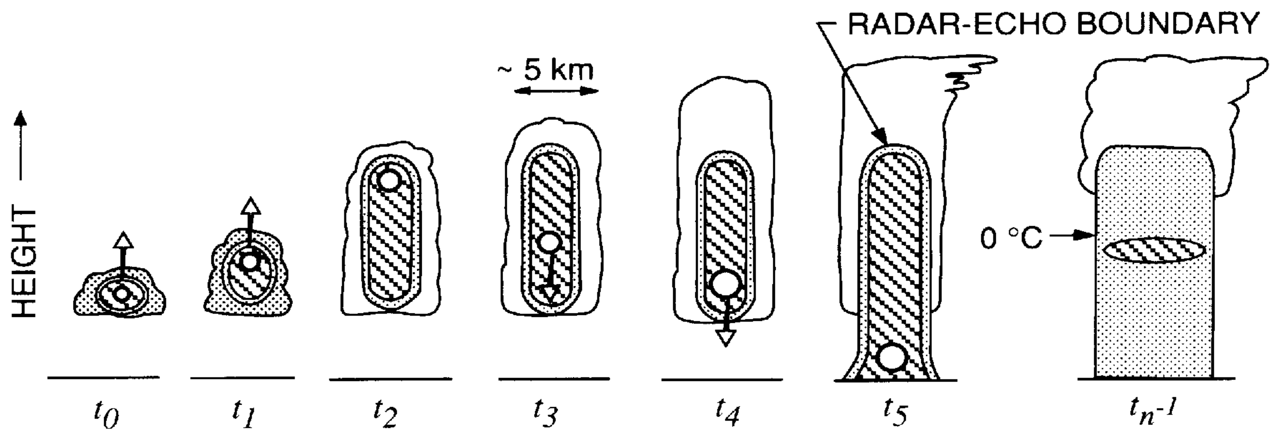

(Houze 2014)

Lifecycle of a convective cell:

Lifetime:	6h
Propagation:	10 m/s
Distance:	200 km

- **A boundary zone of 100-200 km is affected by transition from parameterized to explicit convection.**
- **Very small domains damage the statistics of convection.**
- **Our simulations use wide lateral relaxation zone (50 grid points).**

European-scale simulations

- **GPU-version of COSMO model**
  - Large effort led by O. Fuhrer (MeteoSwiss)
    - runs entirely on GPUs
    - dynamical core rewritten in C++
    - parameterizations use OpenACC
  - Also used for operational NWP (Δ=1 km)
  - Runs on Piz Daint (Cray XC30, CSCS)

Piz Daint: Linpac peak performance: 6×10^{15} Flop/s

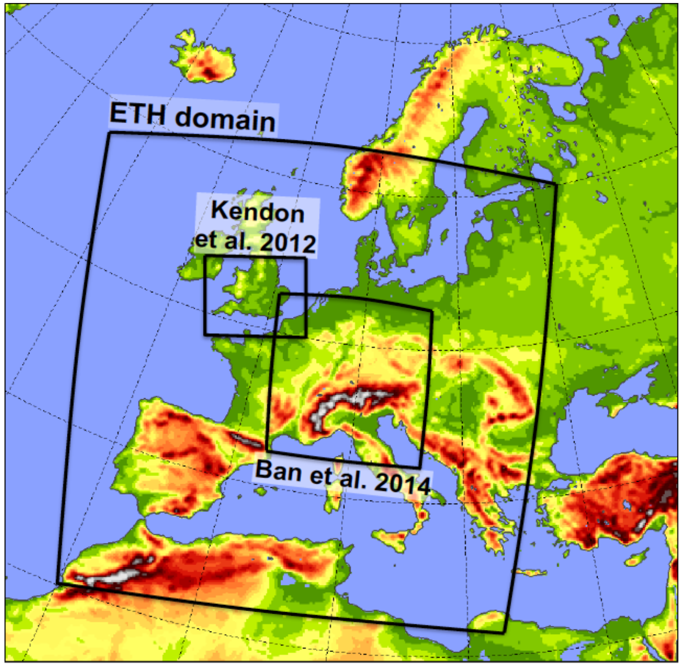

- **European-scale climate simulations**
  - Δ=2.2 km, 1536 x 1536 x 60 grid points
  - Uses intermediate-resolution Δ=12 km simulation
  - Able to run 1 year in 5 days wall-clock time
  - Completed 10 years driven by ERA-Interim

Oliver Fuhrer (MeteoSwiss), Xavier Lapillone (C2SM / ETH), et al.;
Thomas Schulthess (CSCS), et al.;
PhD of David Leutwyler, Leutwyler et al. (2016)

Simulations at 12 and 2 km

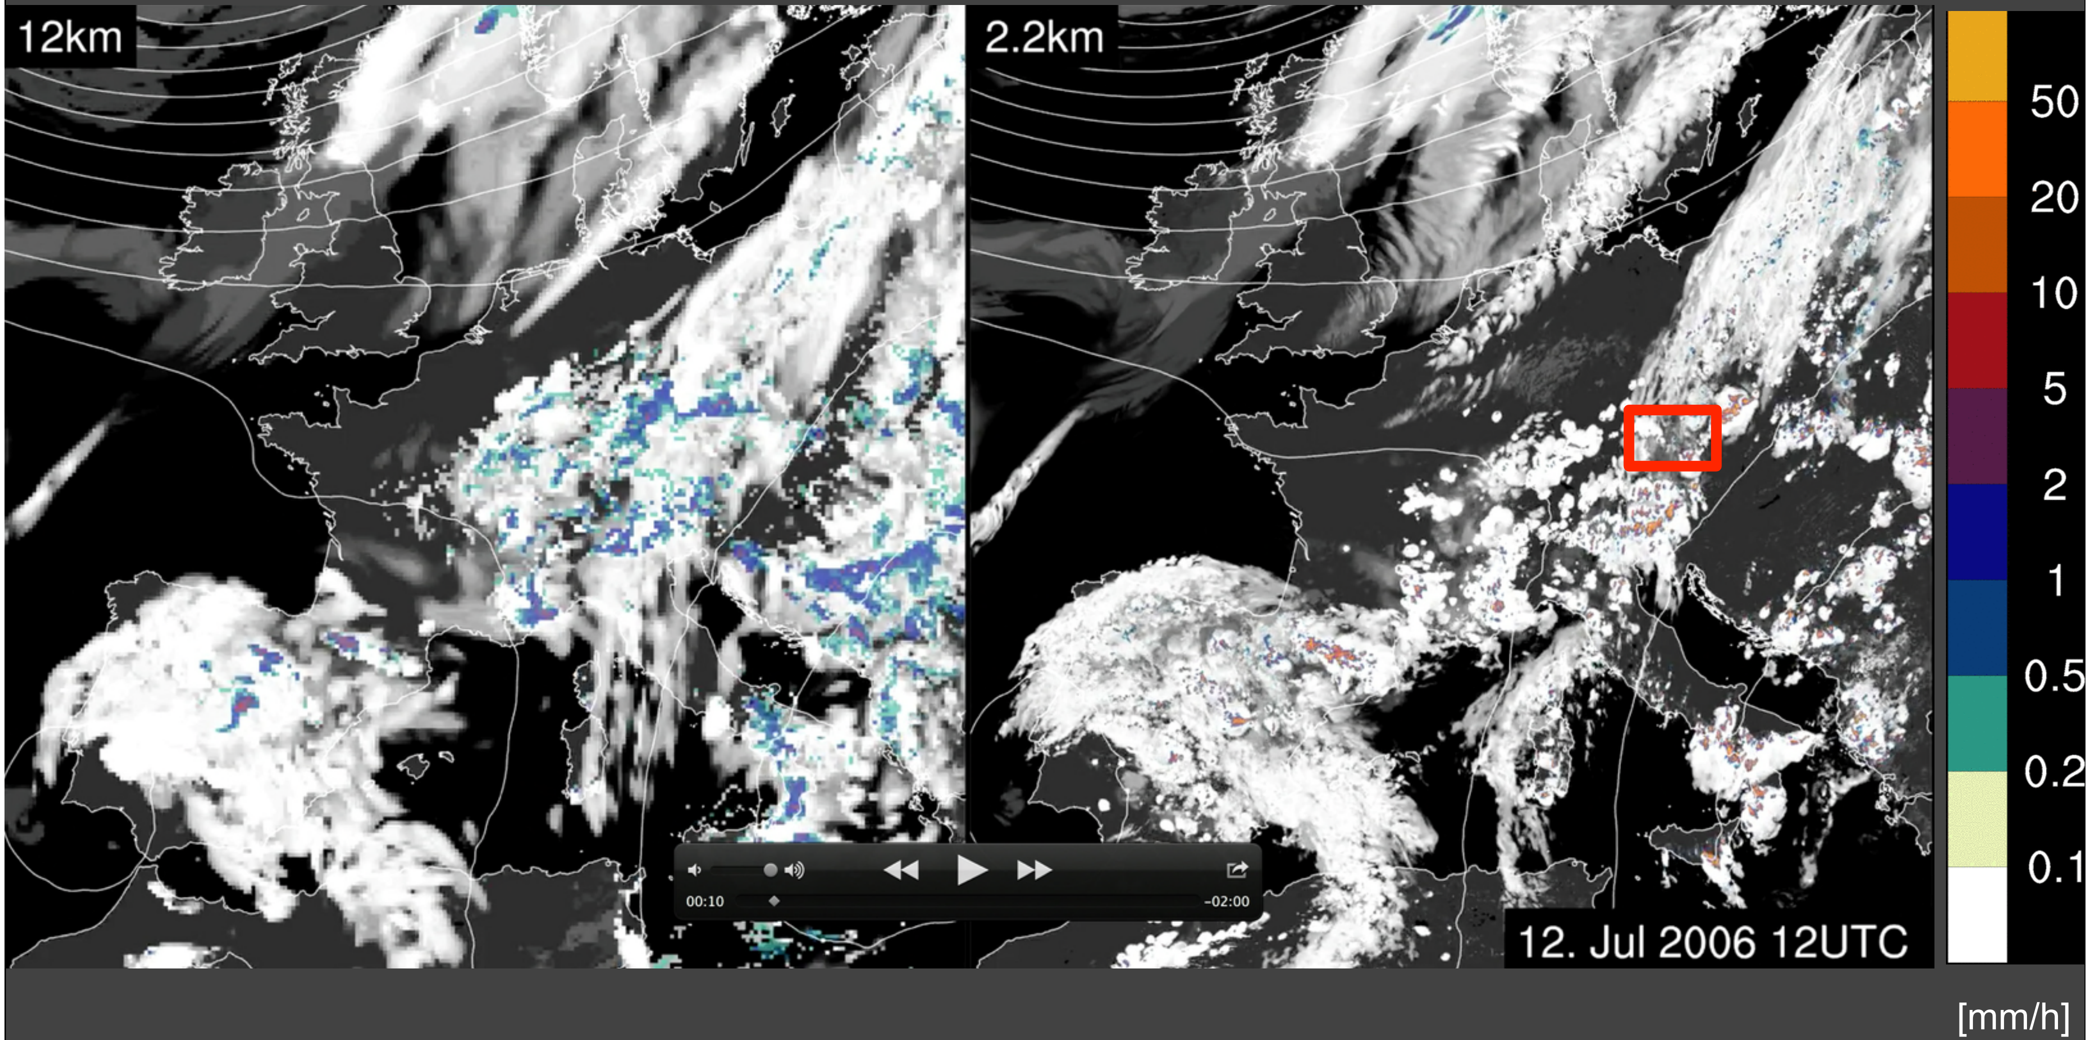

Schär, ETH Zürich

Leutwyler et al. (2016), ETH Zurich, animations via crCLIM: http://www.c2sm.ethz.ch/research/crCLIM

Contents

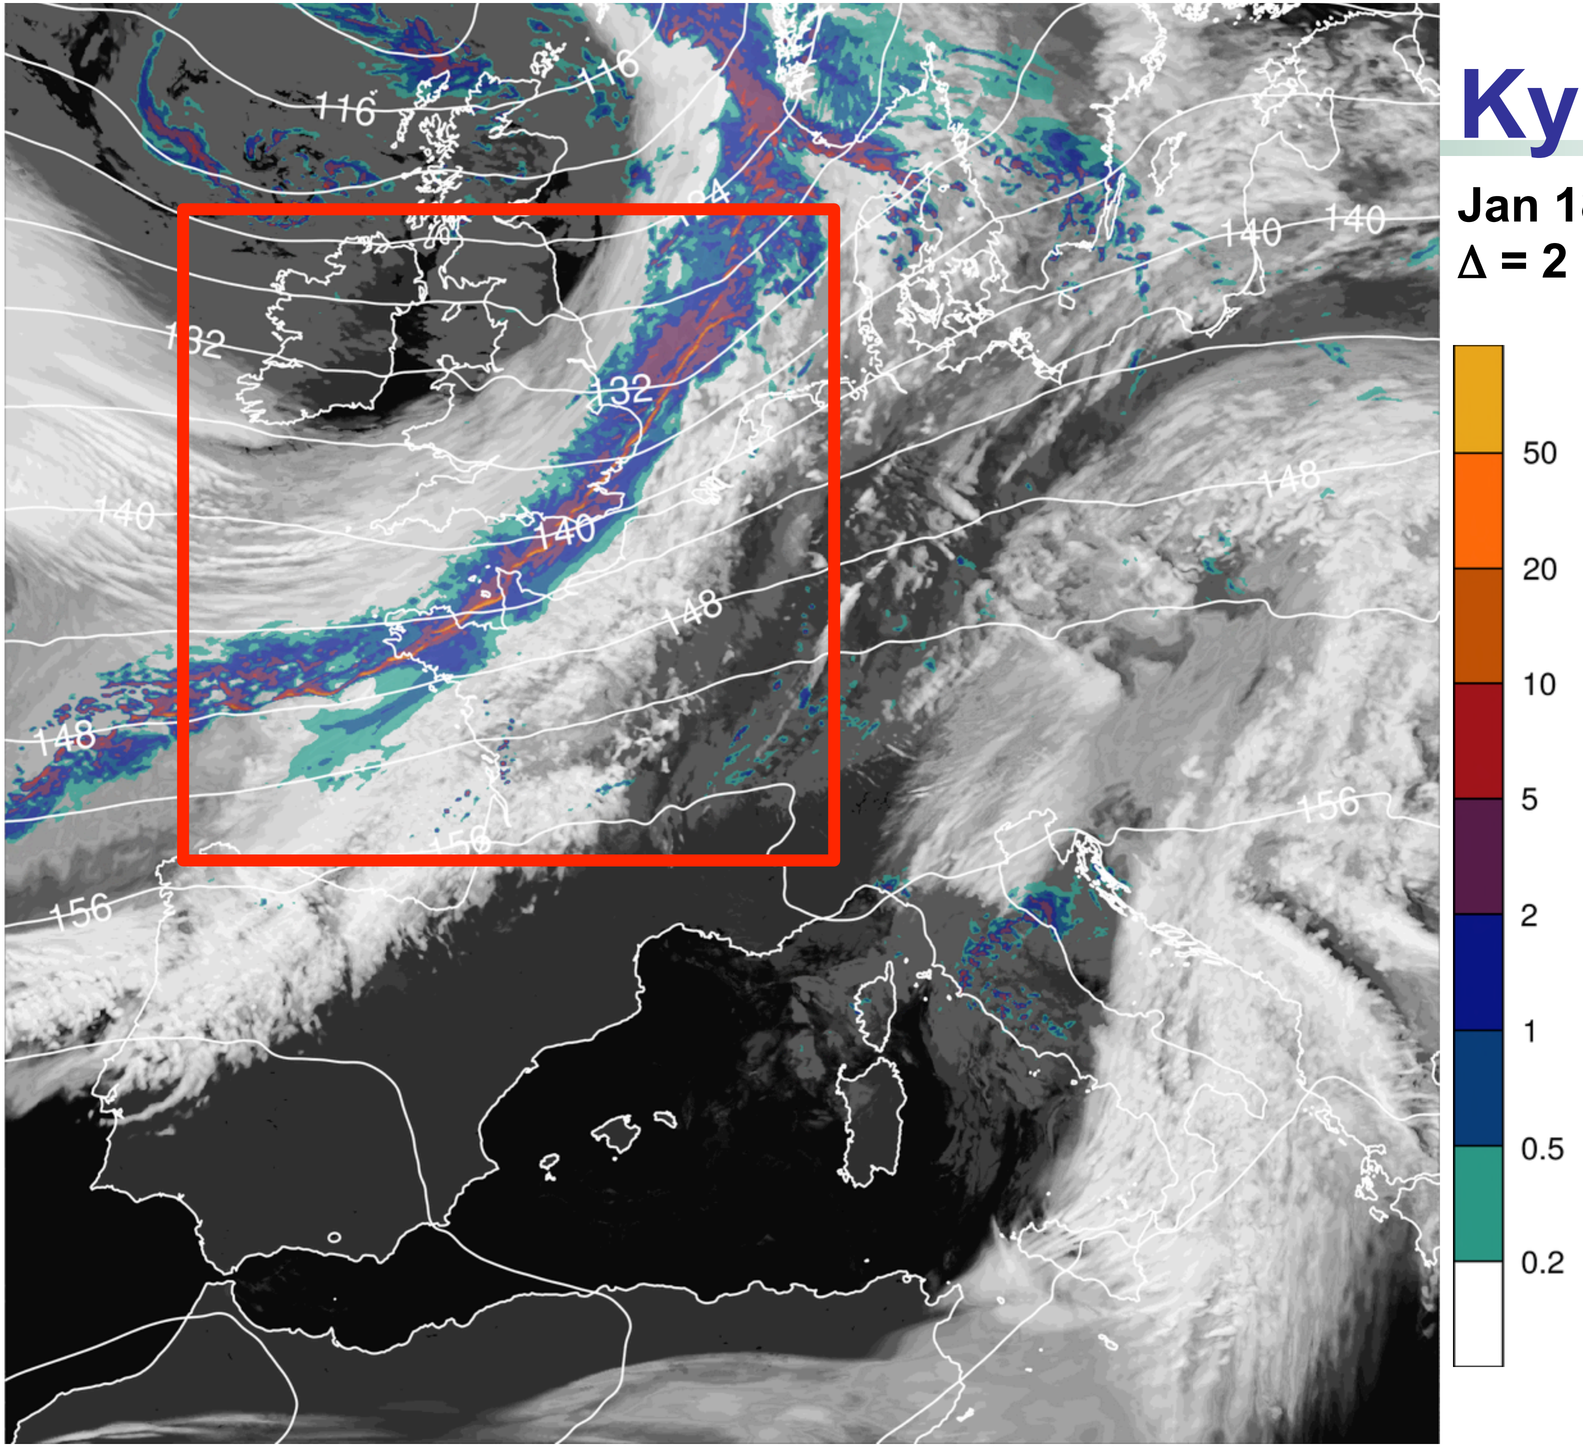

Kyrill

Jan 18, 2007, 12 UTC
Δ = 2 km

Kyrill

Jan 17, 2007, 12 UTC
Δ = 50 km

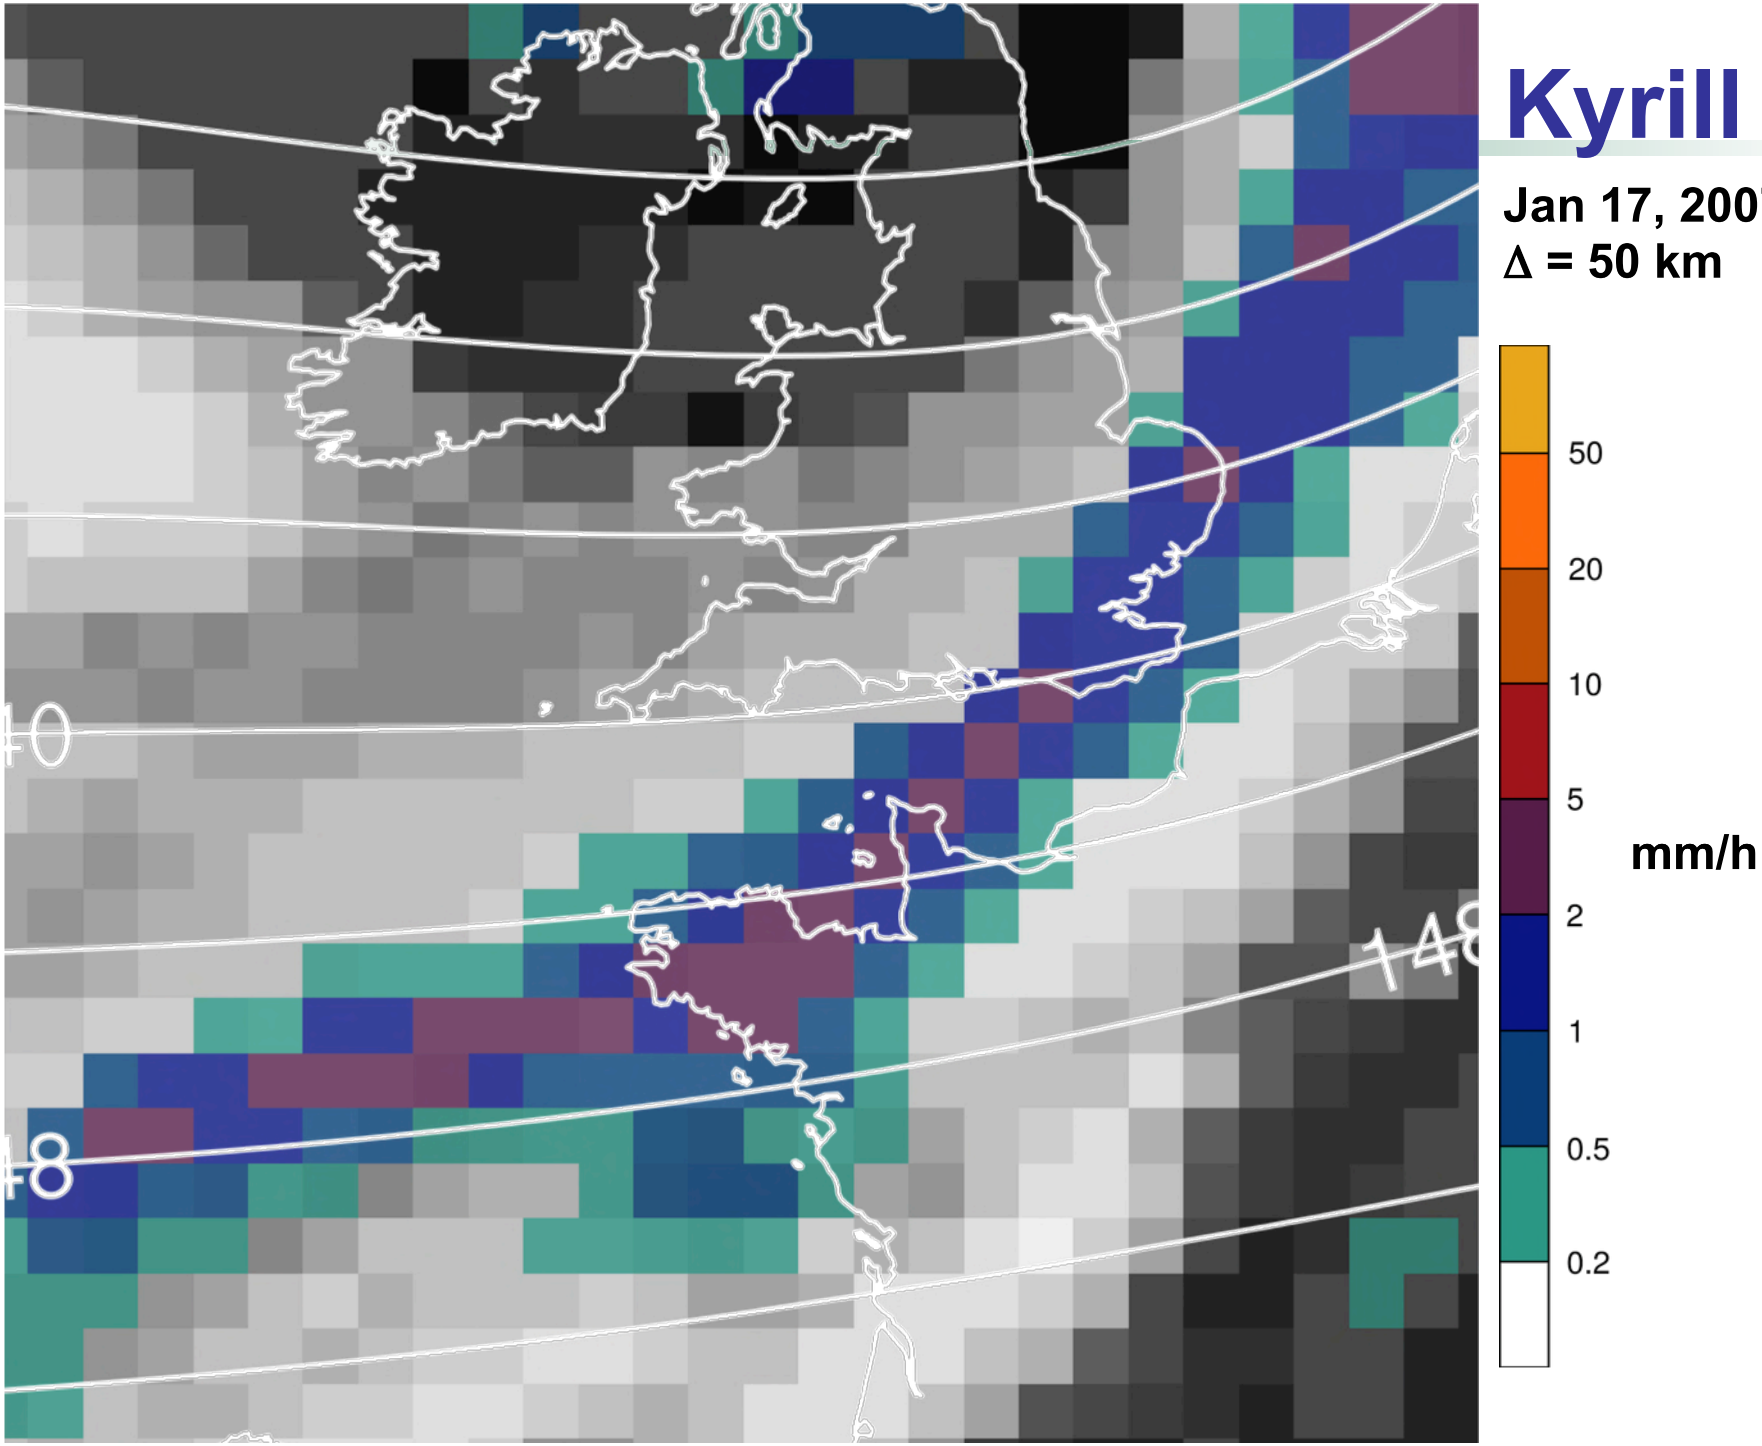

Kyrill

Jan 17, 2007, 12 UTC
Δ = 12 km

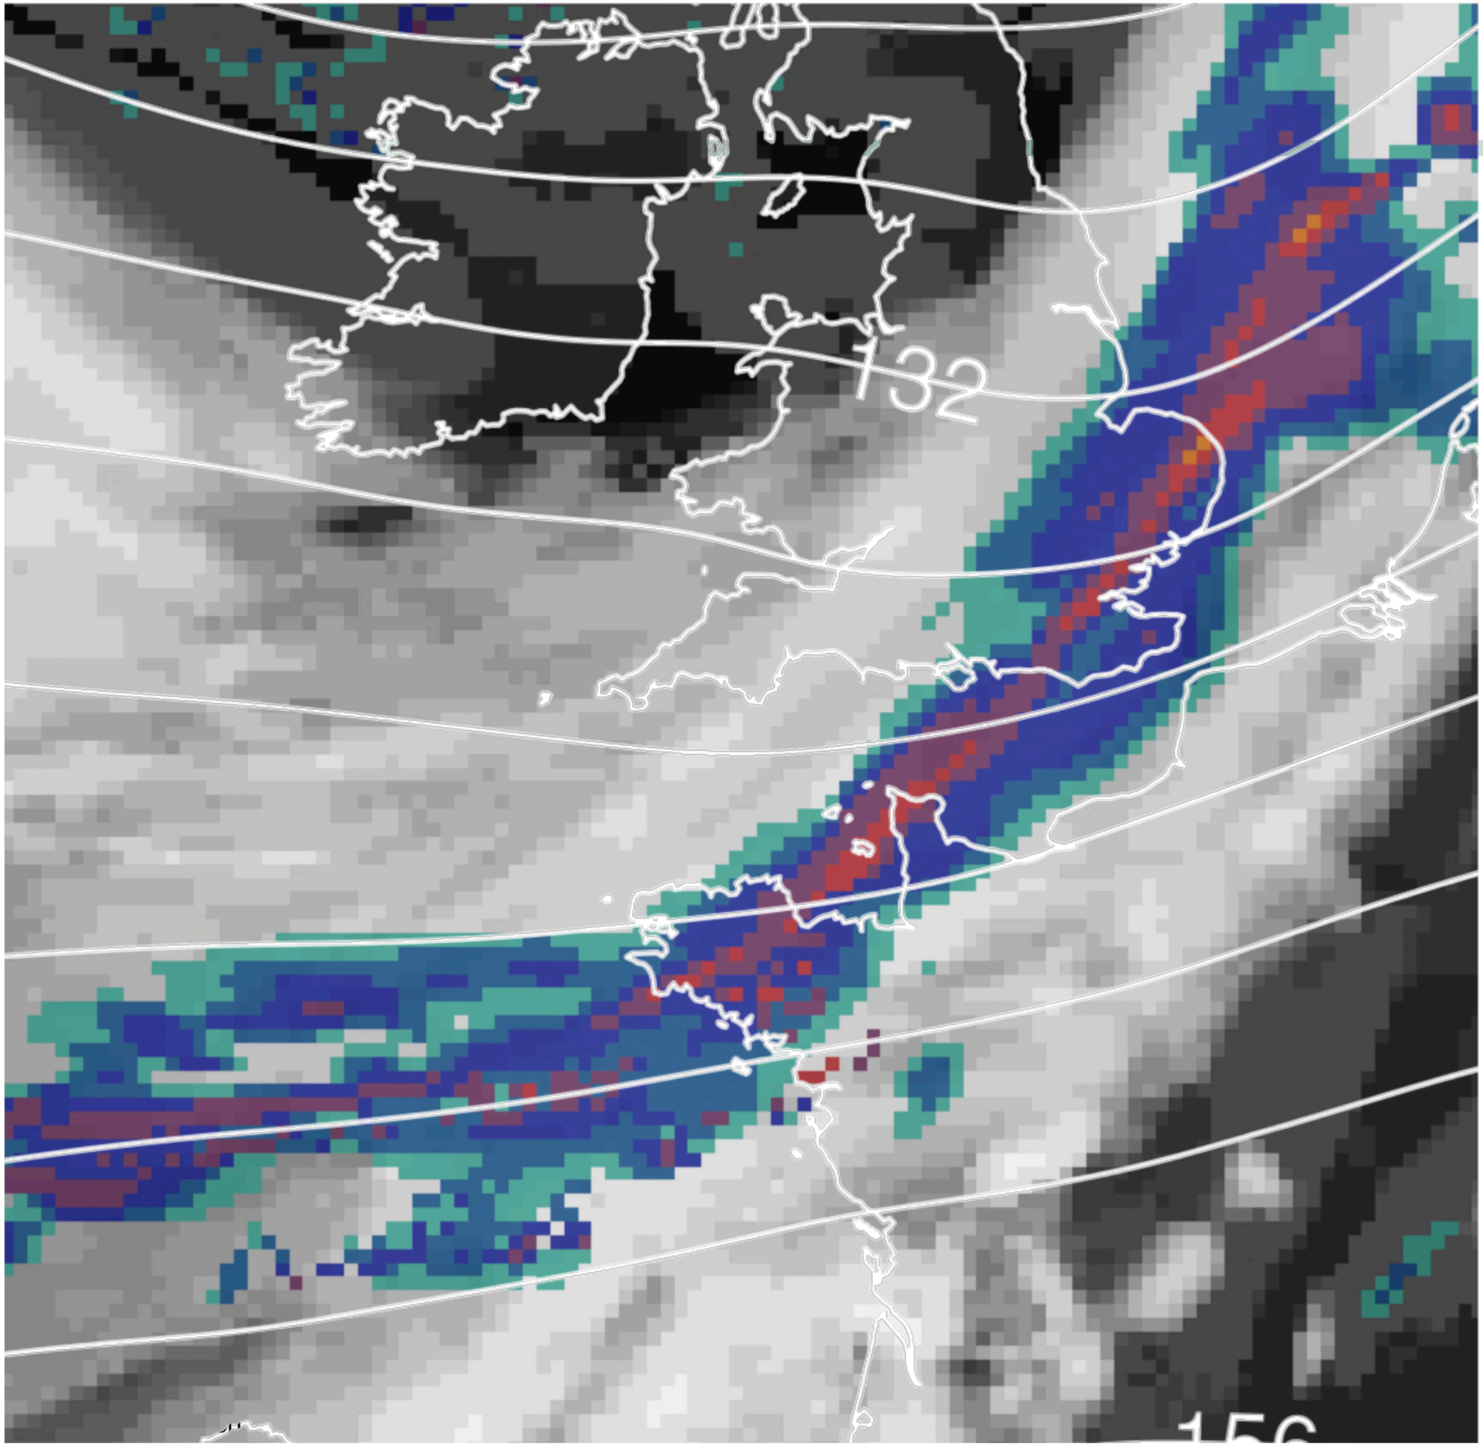

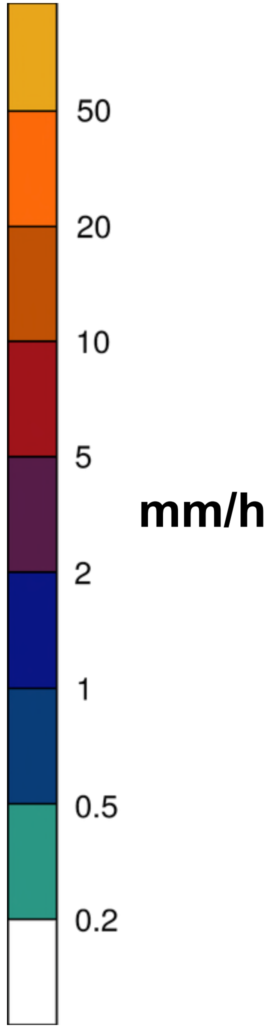

Kyrill

Jan 17, 2007, 12 UTC
Δ = 2 km

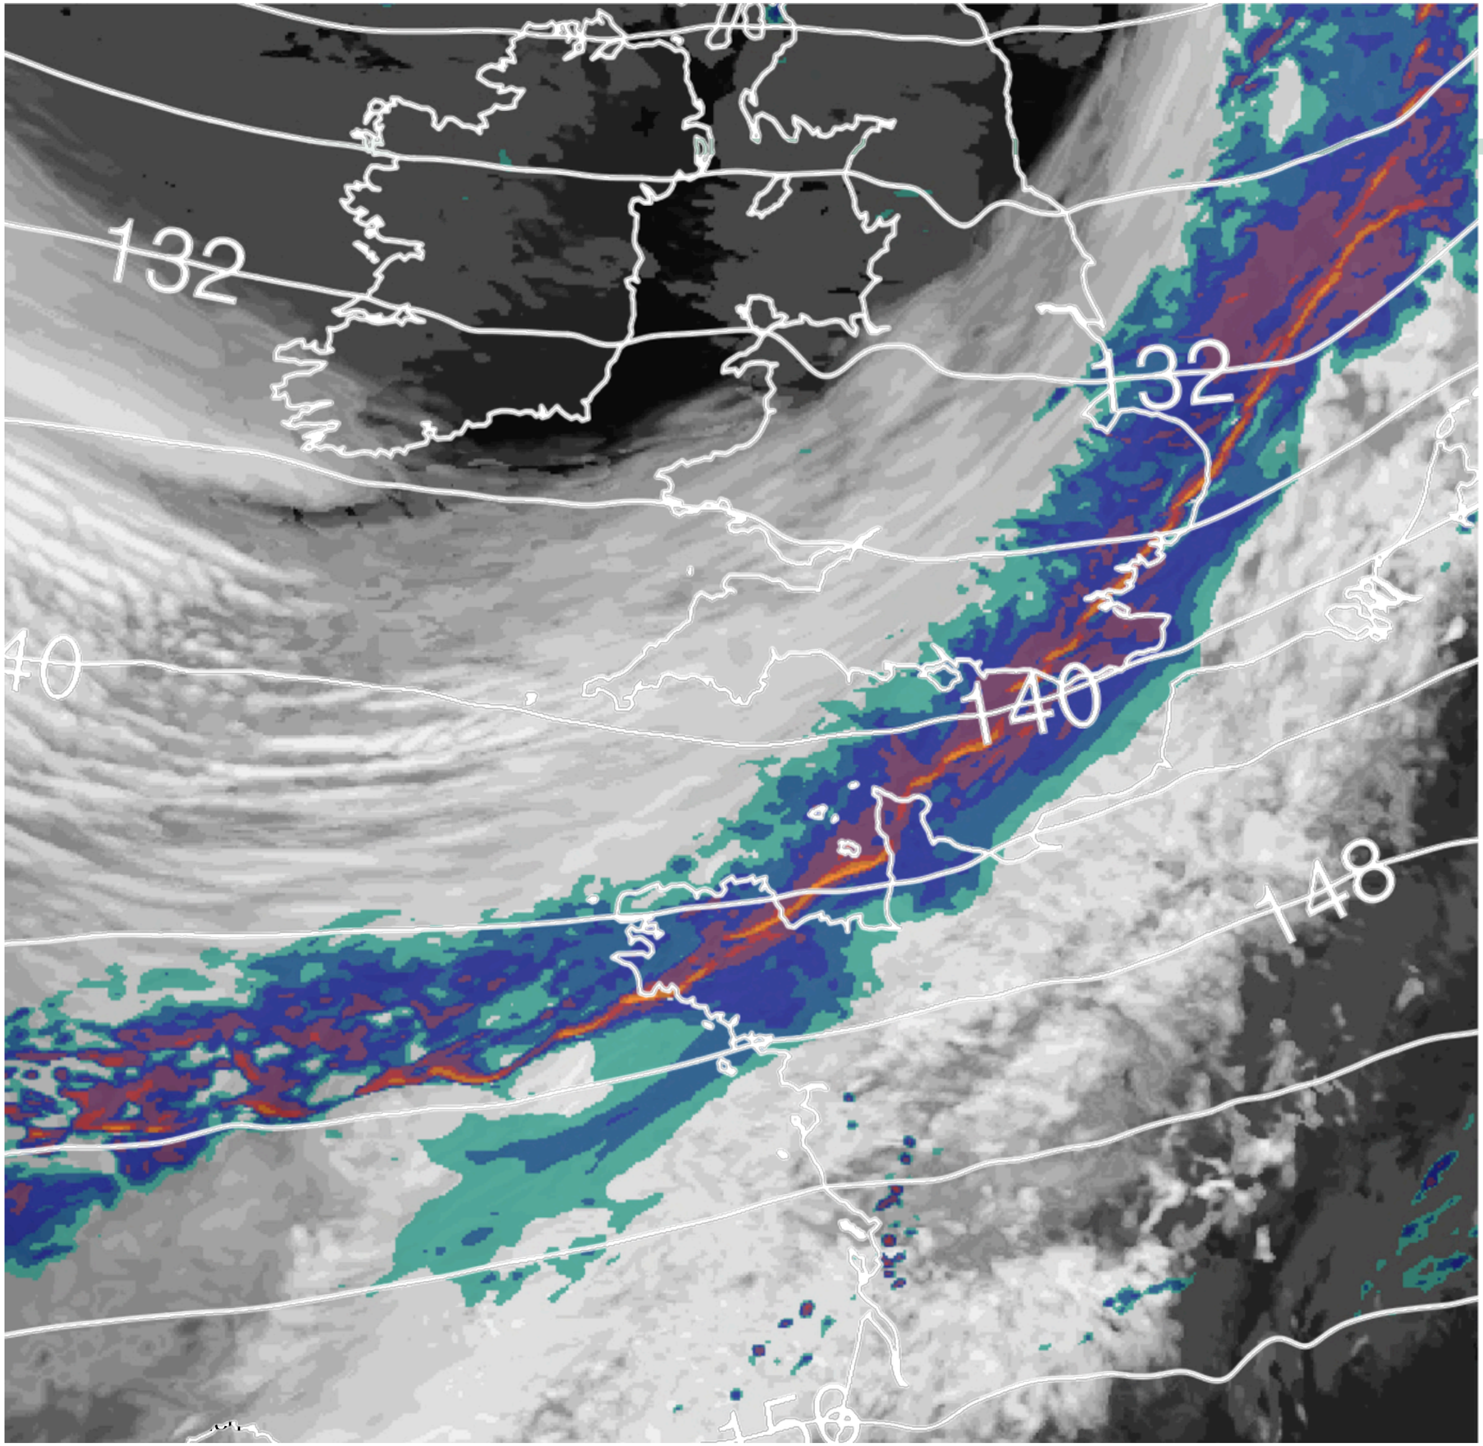

50
20
10
5
2
1
0.5
0.2

mm/h

Kyrill

Break-up of cold-frontal rain band at Δ = 2 km

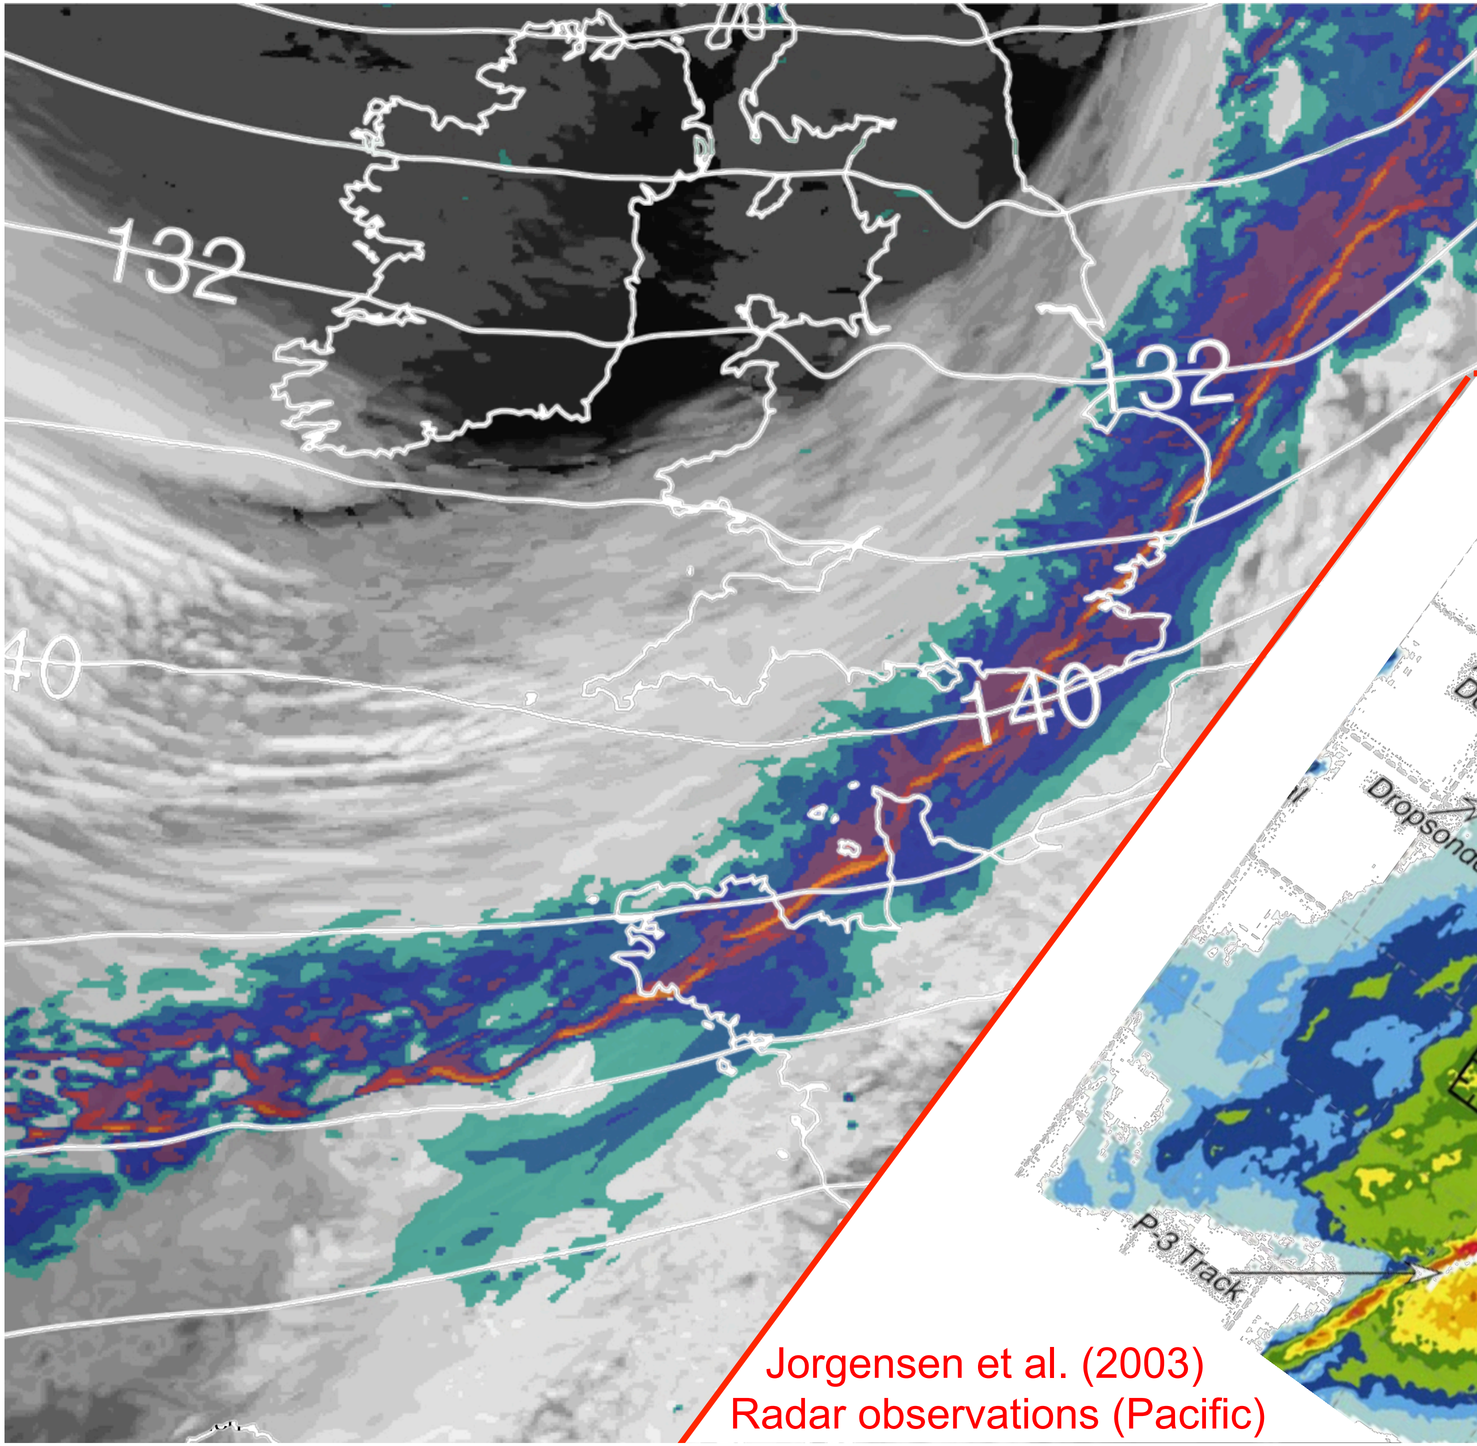

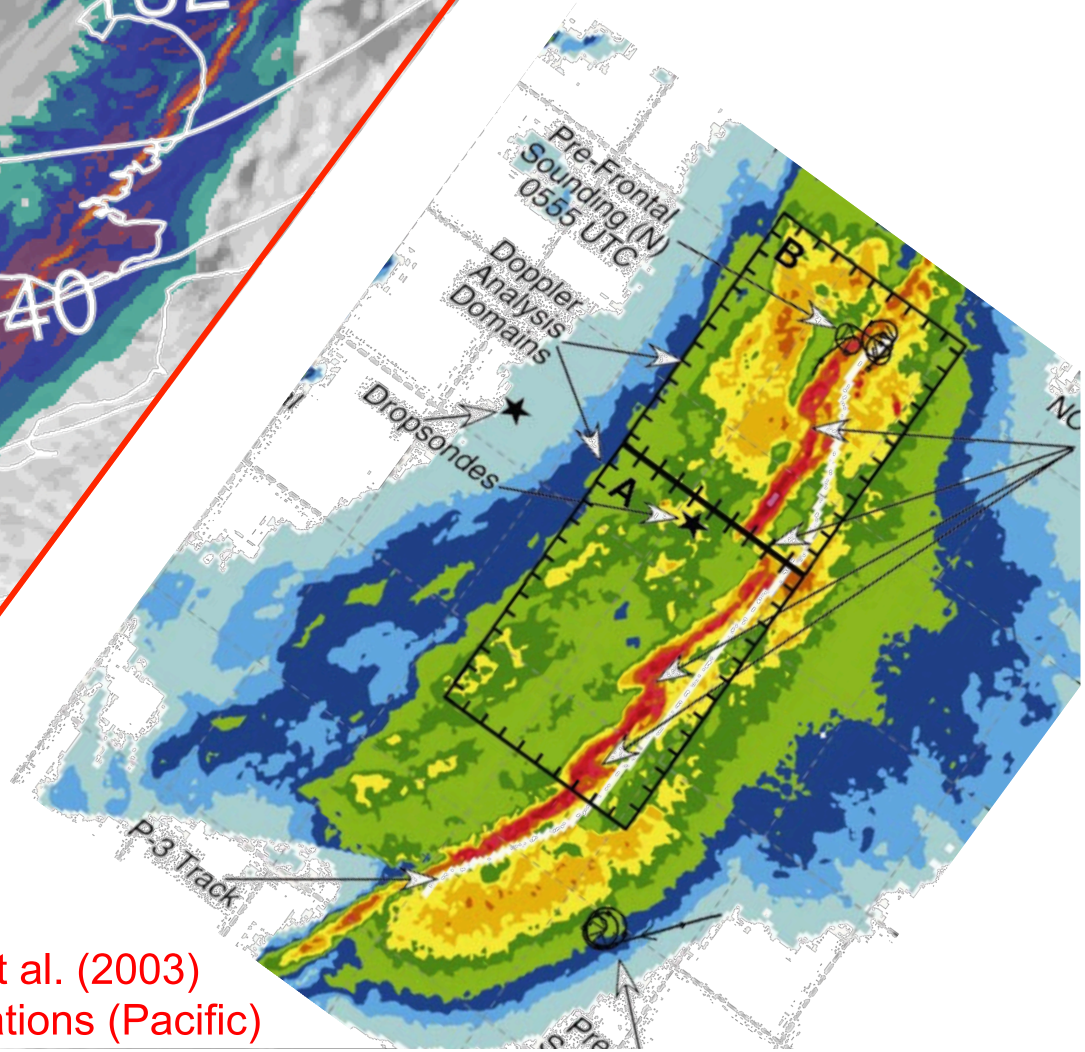

Jorgensen et al. (2003)
Radar observations (Pacific)

Representing convection

Statistics of vertical wind at 500 hPa level, 6h resolution

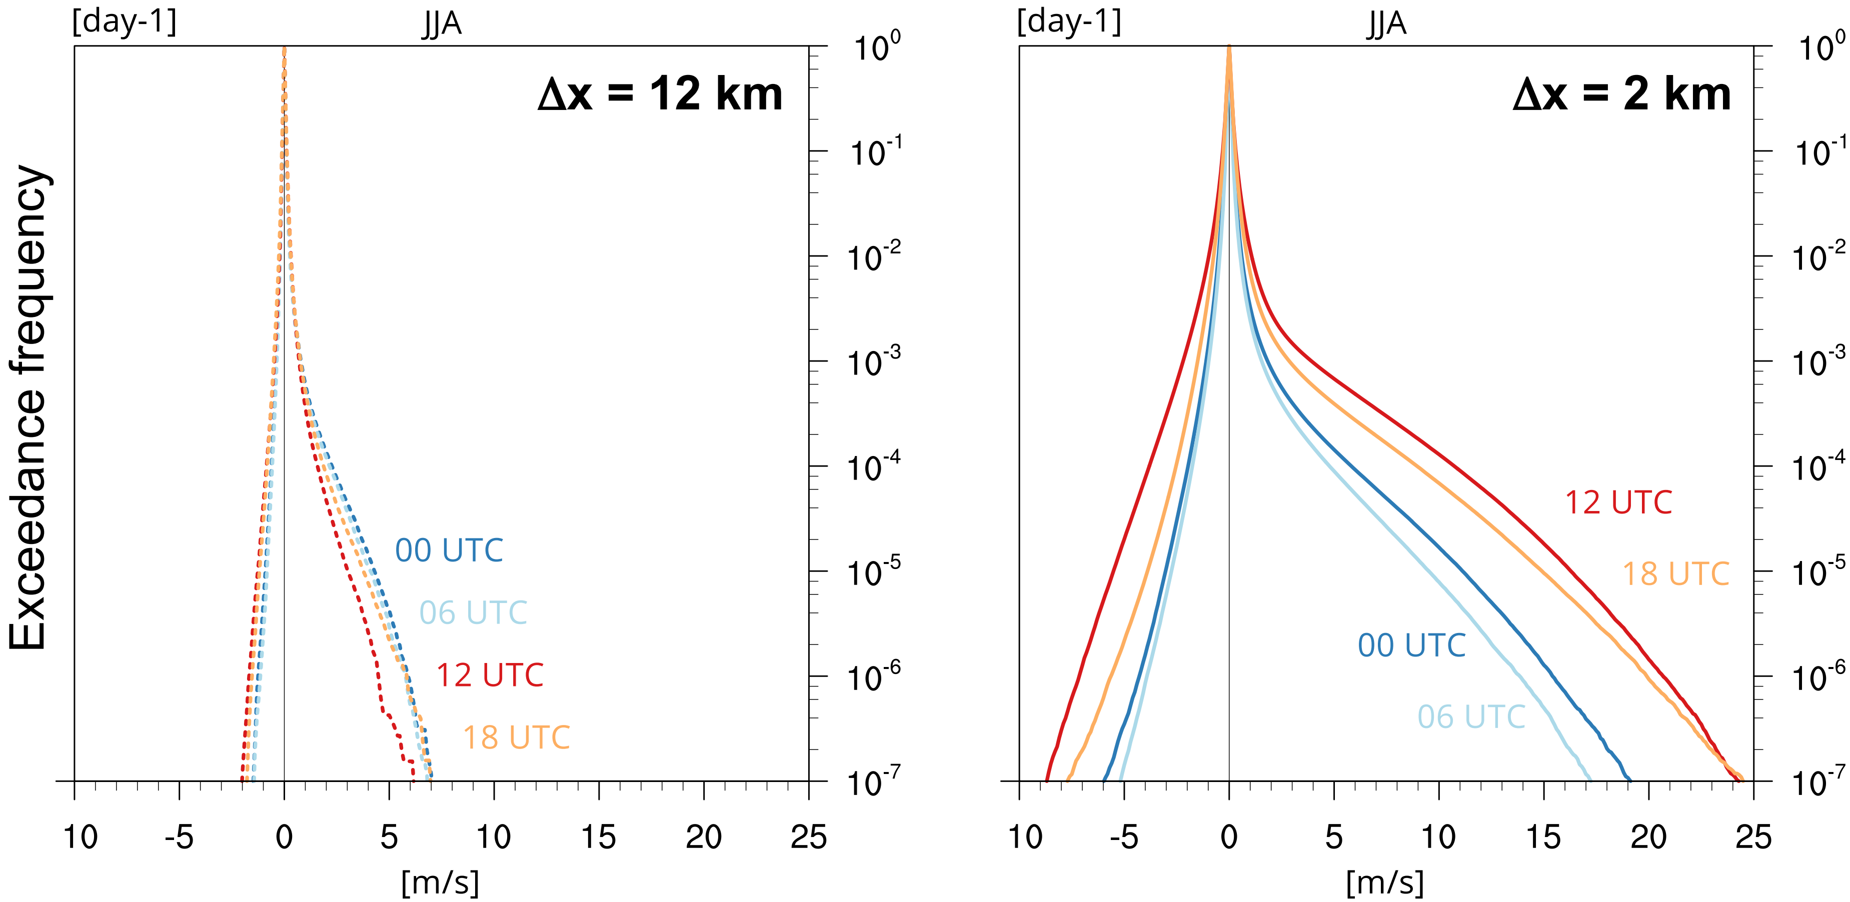

Leutwyler et al. (in prep), ETH Zurich

Gust fronts and cold-air pools

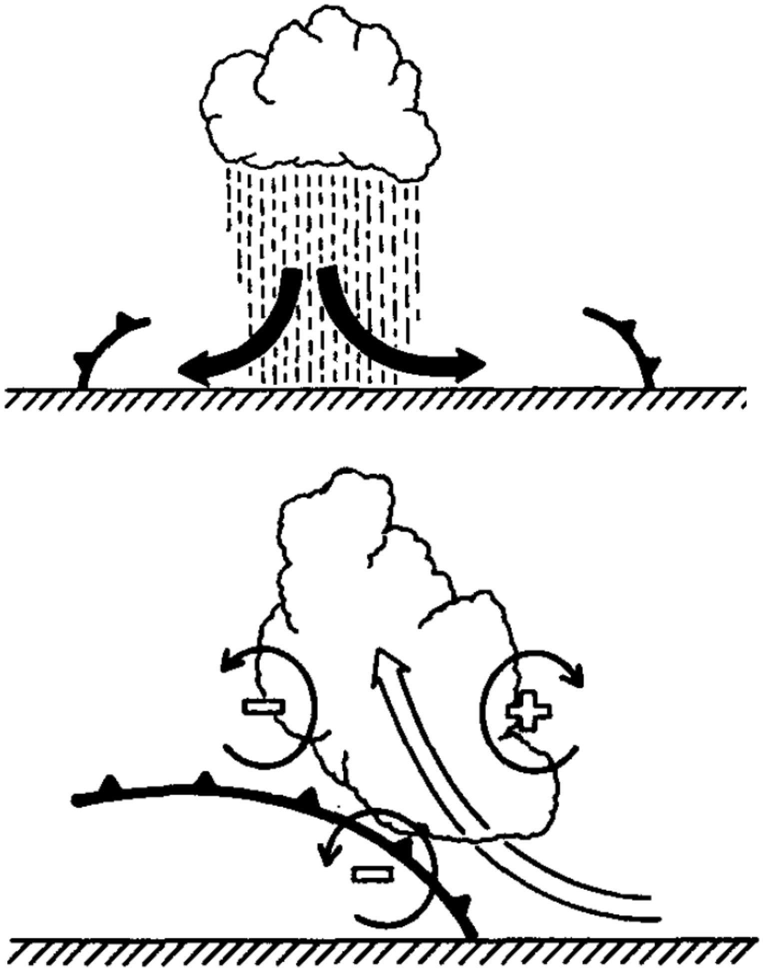

(Rotunno et al. 1988)

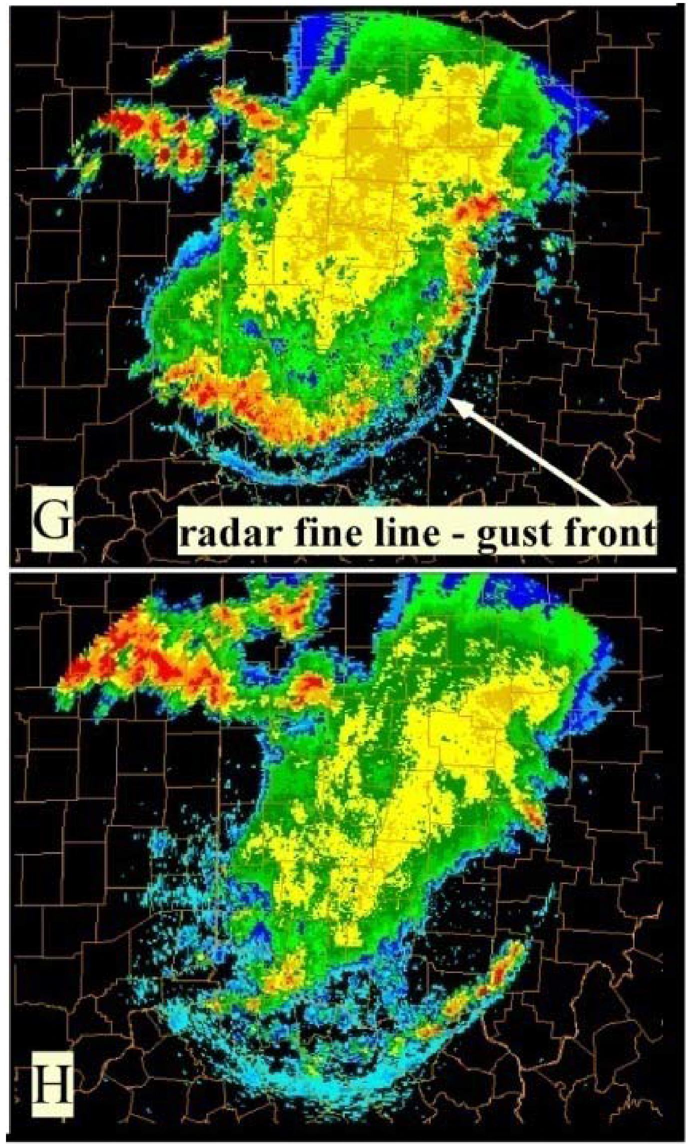

(Jin-Yi Yu, University of California, Irvine)

Simulations at 12 and 2 km

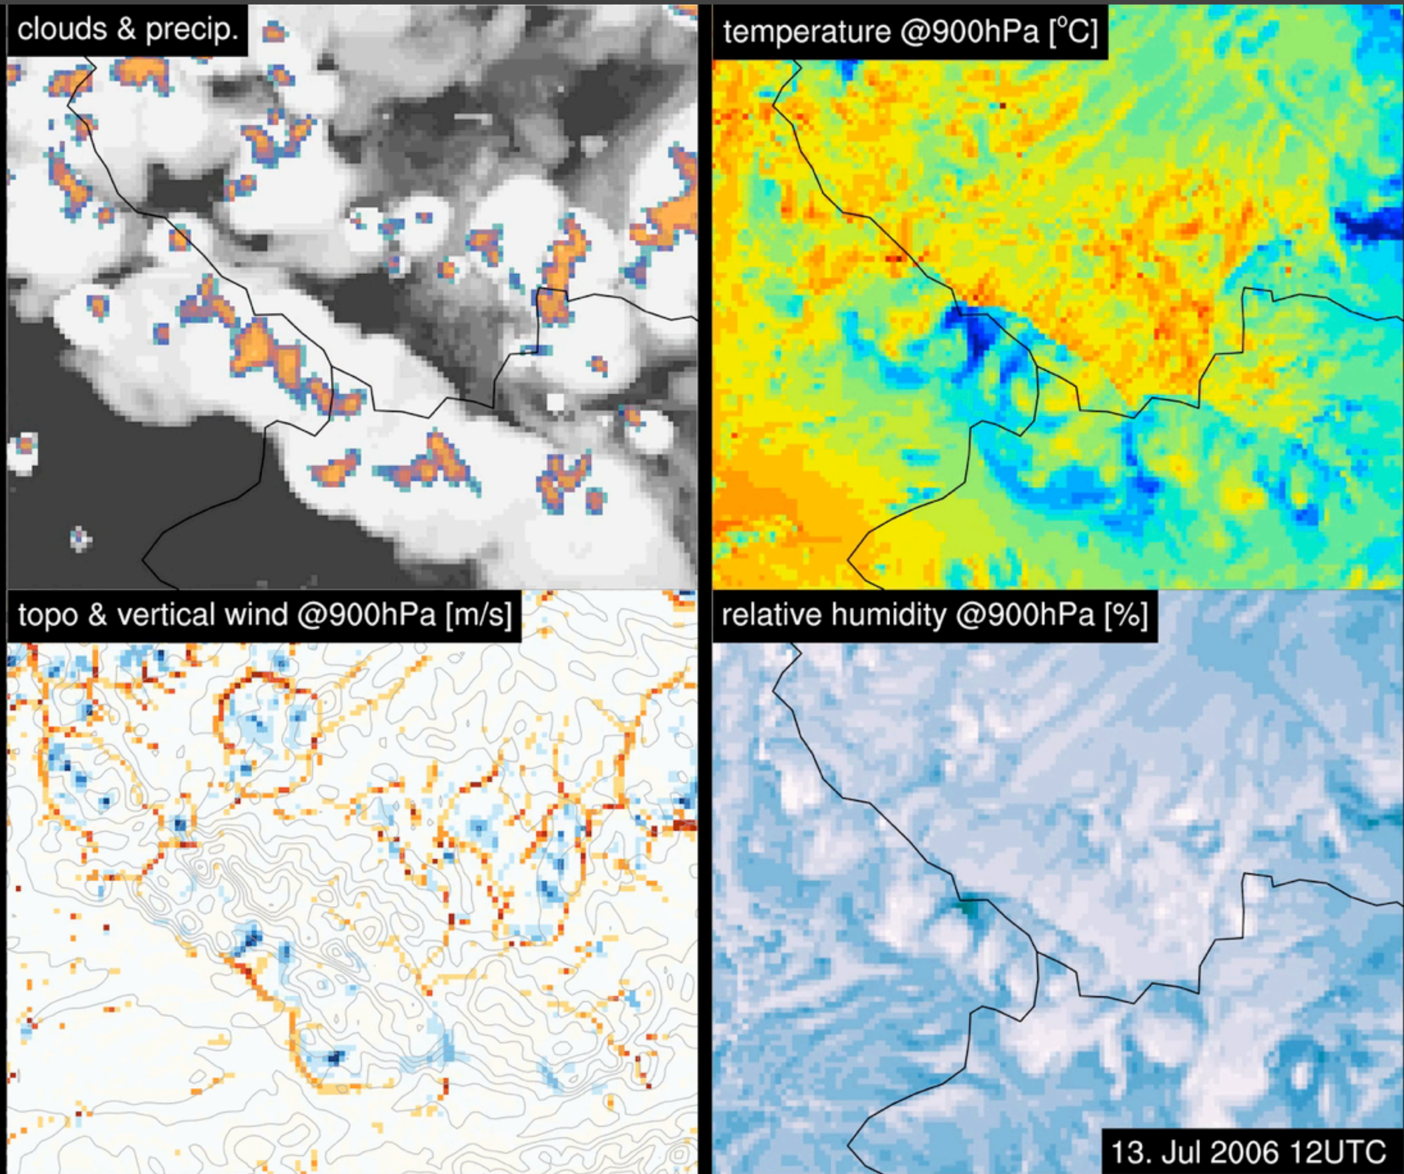

Leutwyler et al. (2016), ETH Zurich, animations via crCLIM: http://www.c2sm.ethz.ch/research/crCLIM

Contents

Parameterization challenges

Main points:

1) **Most models have a history of implicit or explicit tuning and calibration. With explicit convection, balance may change.**

2) **Examples (see also Prein et al. 2015):**
   - Microphysics: role of graupel and hail, cloud-radiative feedbacks, others
   - Turbulence: LES or not?
   - Topography and numerics: new role at high-resolution
   - Soil hydrology: soil-moisture temperature feedback **<= next slides**

Temperature and soil-moisture biases

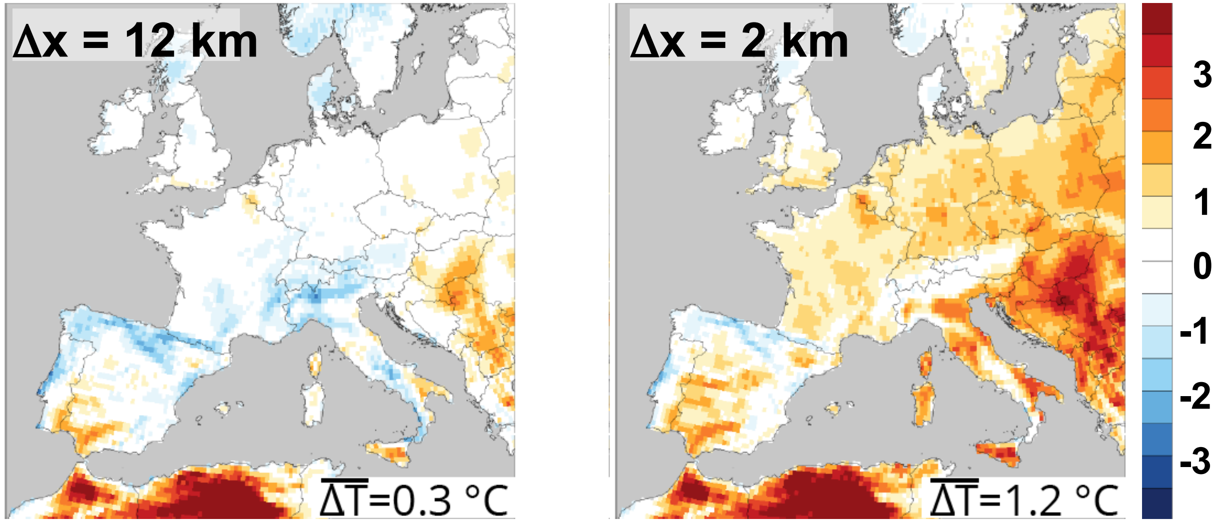

JJA temperature bias

2 km run is warmer by about 1 K

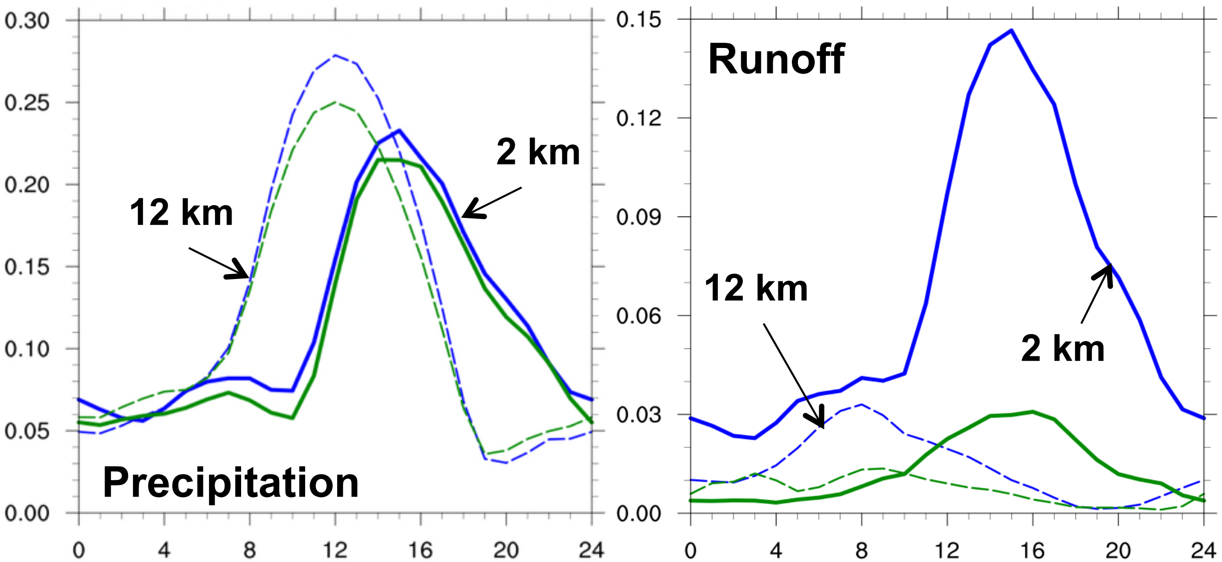

Mean diurnal cycle of precipitation and runoff

12 km has peak runoff in morning

2 km has much larger runoff

Sensitive to infiltration capacity

High precipitation intensity @ 2km
=> large surface runoff
=> dryer soils
=> higher surface temperature

(Leutwyler et al. 2016, Keller et al. 2016, Linda Schlemmer)

Temperature and soil-moisture biases

Lapse-rate derived from surface data

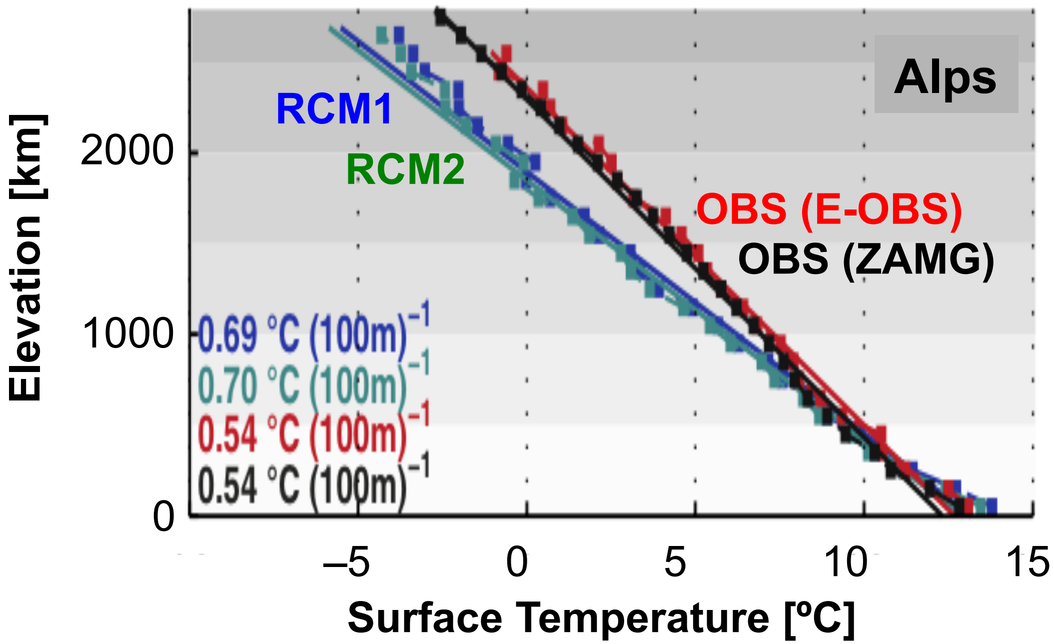

Relative soil moisture (August 2007)

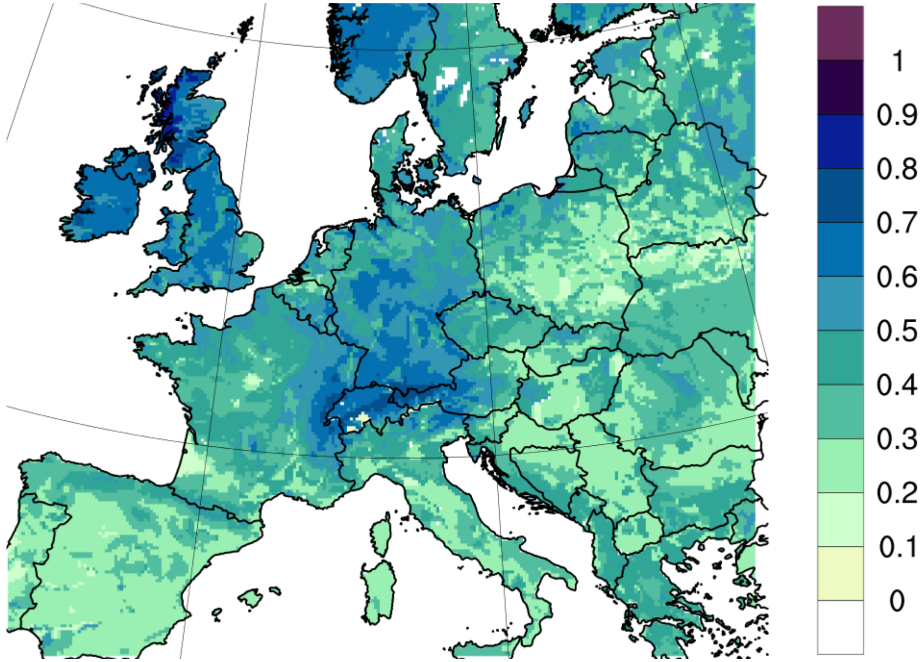

Kotlarski et al. 2012: lapse-rates are systematically overestimated by RCMs

Michael Keller, ETH, PhD 2016: mountains are wetter than valleys

High-altitude enhanced precipitation leads to wet soils and overestimation of lapse-rates

Contents

Compute challenges

Main points:

1) **Modern supercomputers increasingly use a heterogeneous / hierarchical design (e.g. with accelerators and/or GPUs).**

2) **This trend is likely to continue, e.g. driven by considerations of energy consumption** (e.g. Schulthess 2015)

3) **Moving data is as important as compute operations: It becomes imperative to reduce communication as far as feasible.**

4) **Most atmospheric models use double precision, in reduced computational precision would be sufficient** (Düben and Palmer 2014)

Compute Challenge

Emerging hardware architectures
are highly heterogeneous

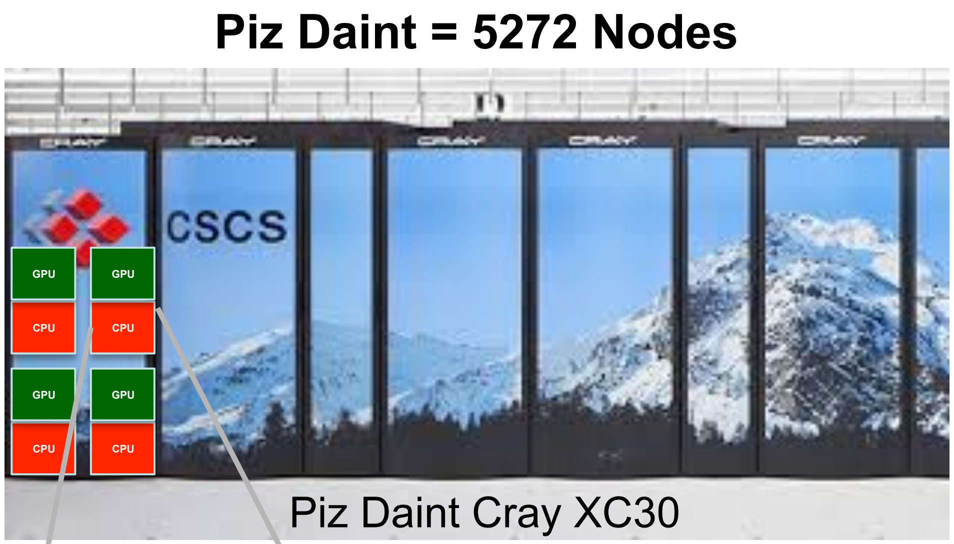

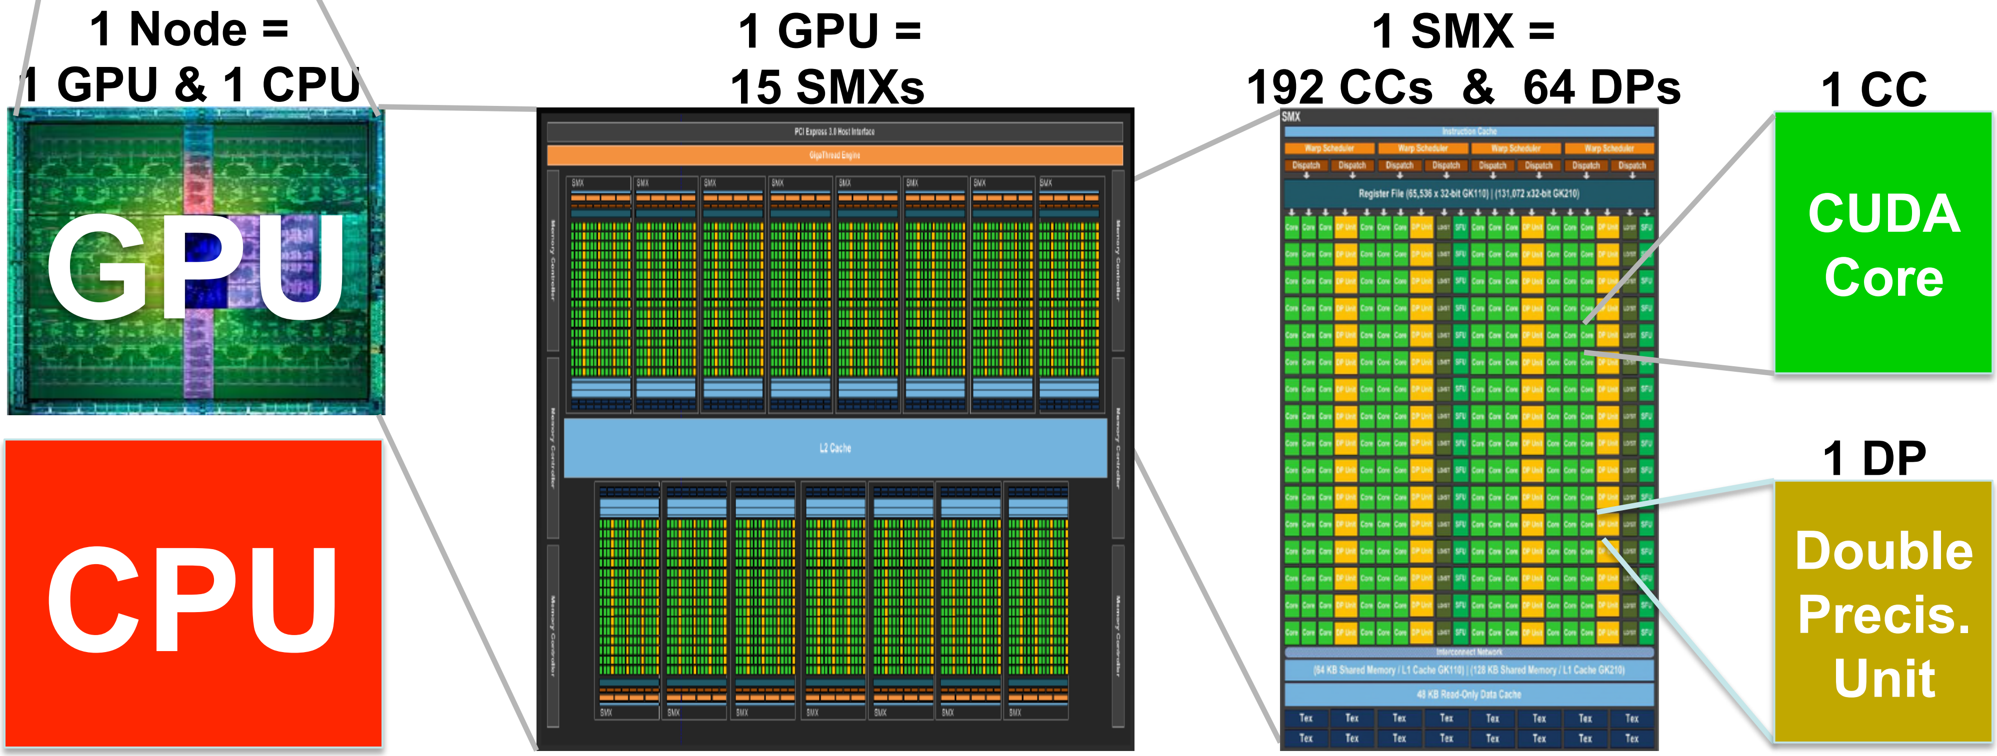

Cray 1 versus Kepler GPU

(released 1975)

Performance: 160 MFlops
Main memory: 8 MByte
Weight: 5.5 tons (Cray-1A)
Cost: about 10 Million $

Kepler GPU (GK110)
GPU = Graphics Processing Unit

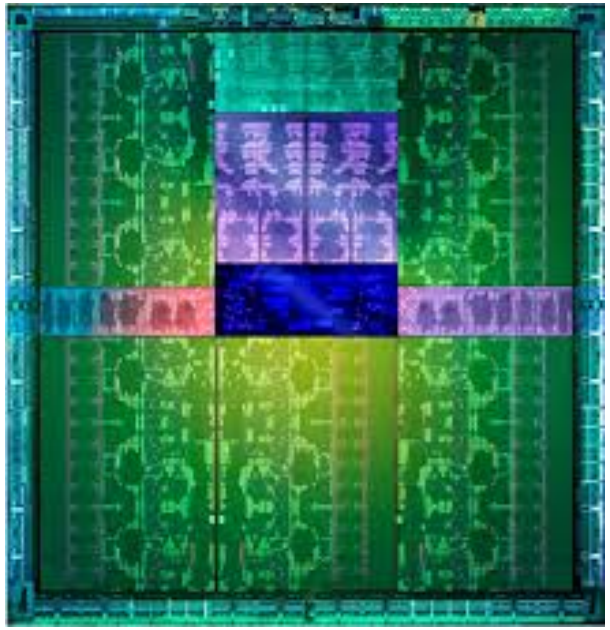

Peak performance:
Double precision: 1311 Gflops
Single Precision: 3.95 TFlops
Memory:
L1/2 Cache: 960 + 1536 kByte
VRAM: 6 GB
Weight: about 0.5 kg
Cost: about 3000 $

Compute challenge

GPU-Version of COSMO model

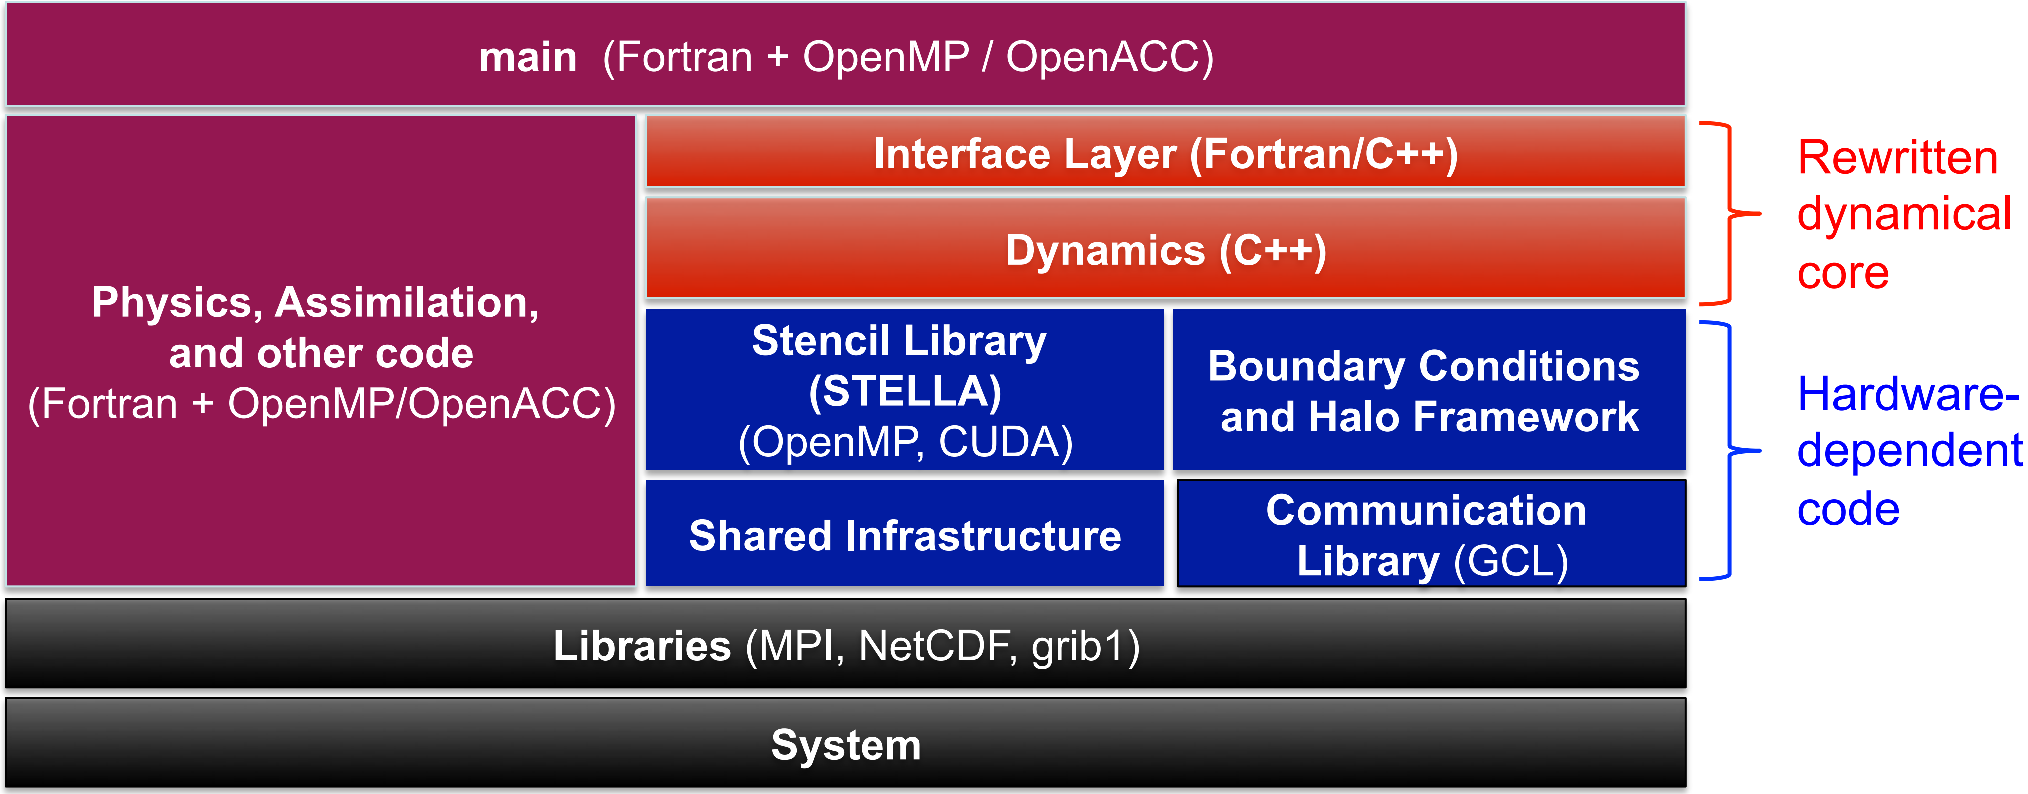

Fuhrer et al. (2014), http://superfri.org/superfri/article/view/17
Gysi et al. (2015), http://sc15.supercomputing.org/schedule/event_detail?evid=pap298
Lapillone and Fuhrer (2015), http://www.worldscientific.com/doi/abs/10.1142/S0129626414500030

Propagation in the atmosphere

Signal after 6 min (30 time steps)

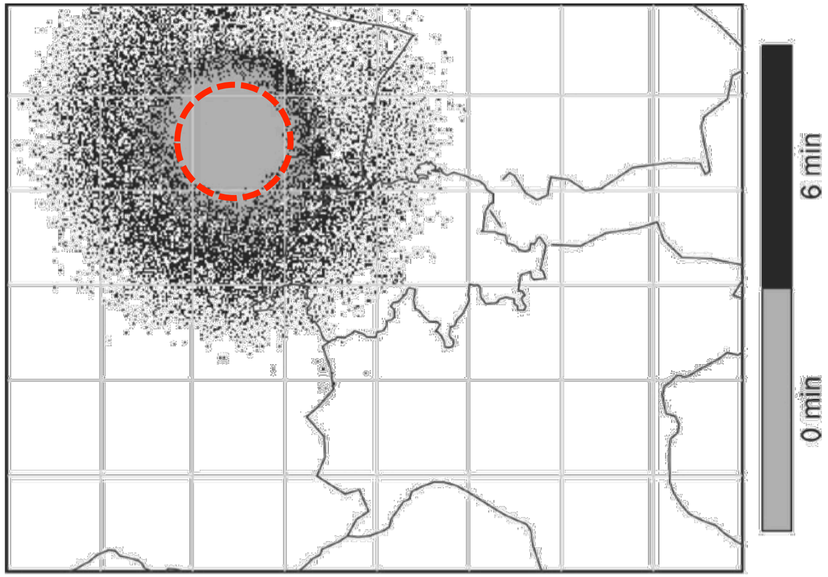

Propagation of initial perturbation
by sound waves (about 310 m/s)

Not physically relevant,
but numerically!

Signal after 6 hours

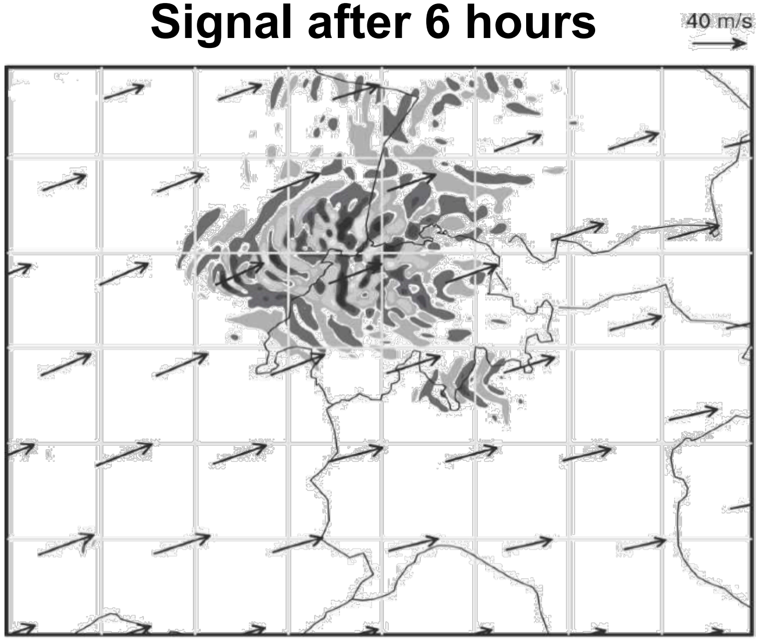

Propagation by and growth of
gravity waves

(Hohenegger and Schär, 2007, JAS)

Propagation in models

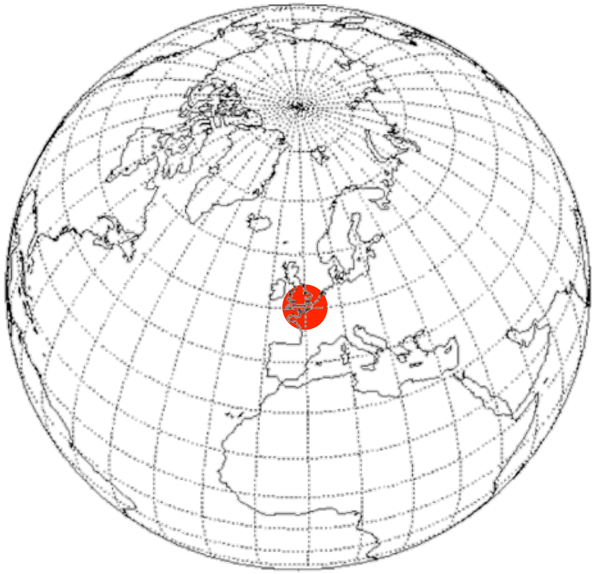

Physical propagation in the atmosphere
(ca 400 m/s)

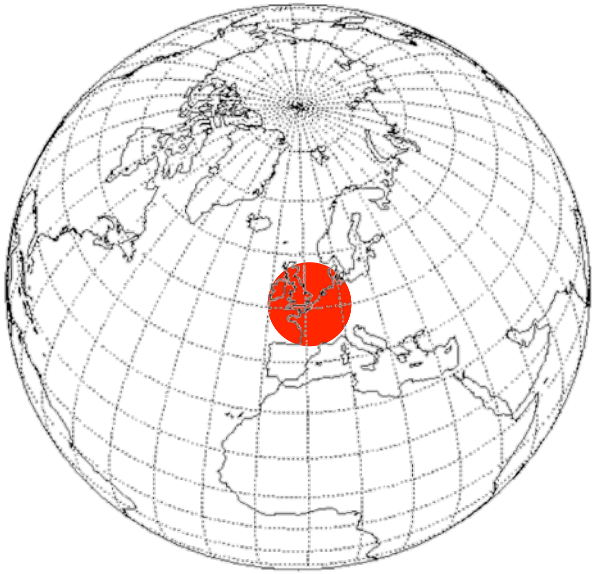

Propagation of data in a split-explicit model
(ca 1000 m/s)

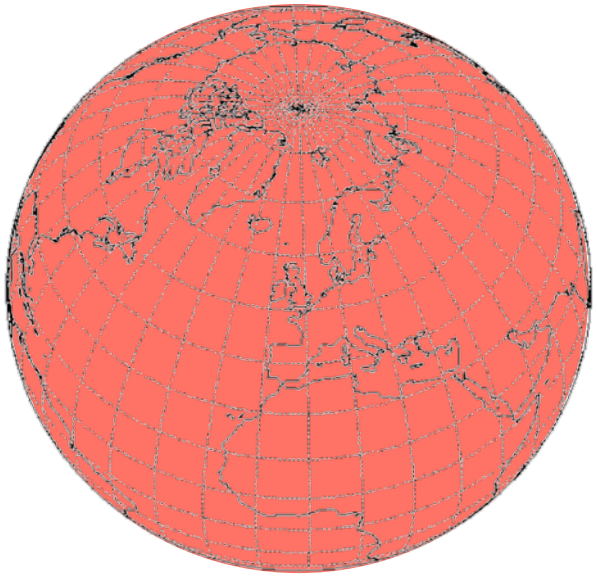

Propagation of data in a global spectral model
(global communication at each time step)

In order to minimize communication of data, numerical methods should reflect principles of physical propagation

Contents

Data challenges

Main points:

1) **Data storage is becoming a critical bottle neck:**
   Mass storage requirements (with a dramatically reduced output list):
   - 10-year simulations of David Leutwyler: 55 Terabyte
   - Global simulations at same resolution: 25 Petabyte

2) **Fundamental limitation is I/O bandwidth**

3) **New strategy explored in project crCLIM at ETH**
   - perform online analysis
   - rerun simulations (or virtualize workflow)
   - make model bit-reproducible across platforms

Output Challenge

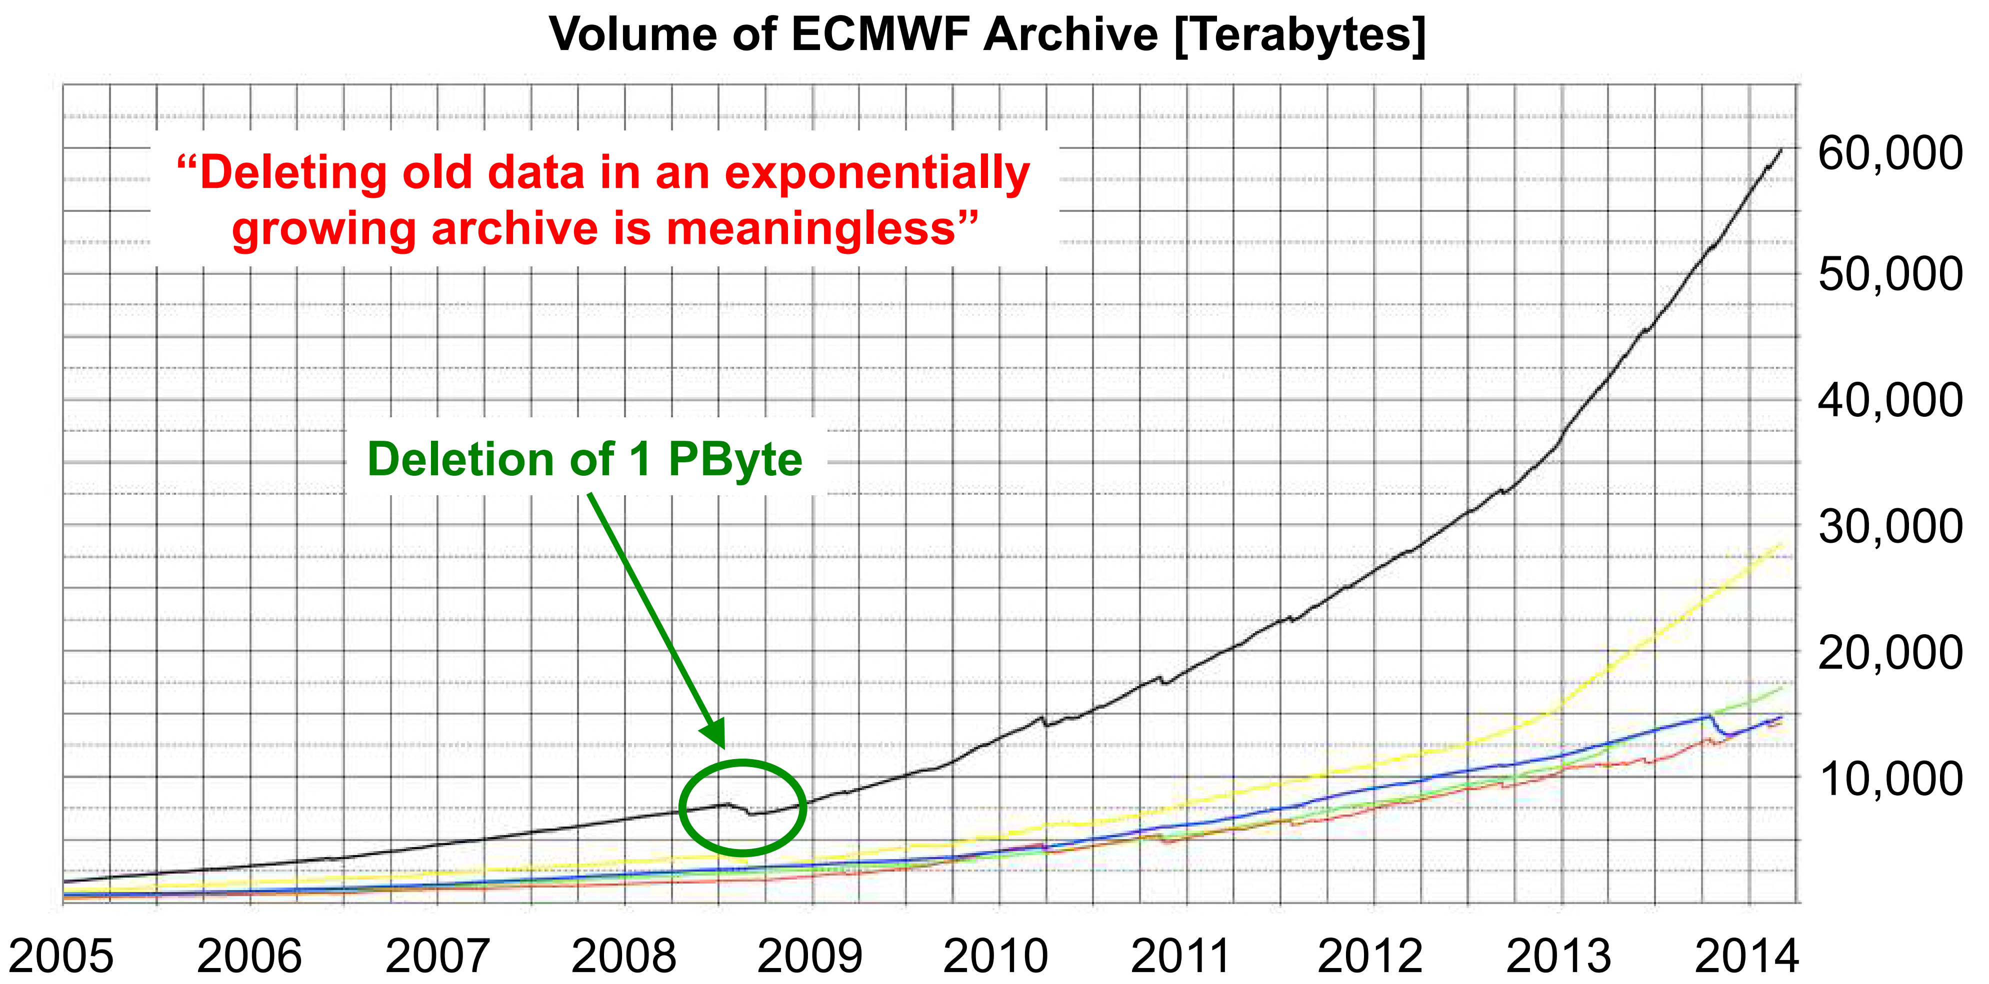

(Baudouin Raoult, ECMWF)

Output challenge

Fundamental limitation is I/O bandwidth

- **Consider use of online analysis rather than mass storage**

- **Recreate model trajectory and use online analysis**

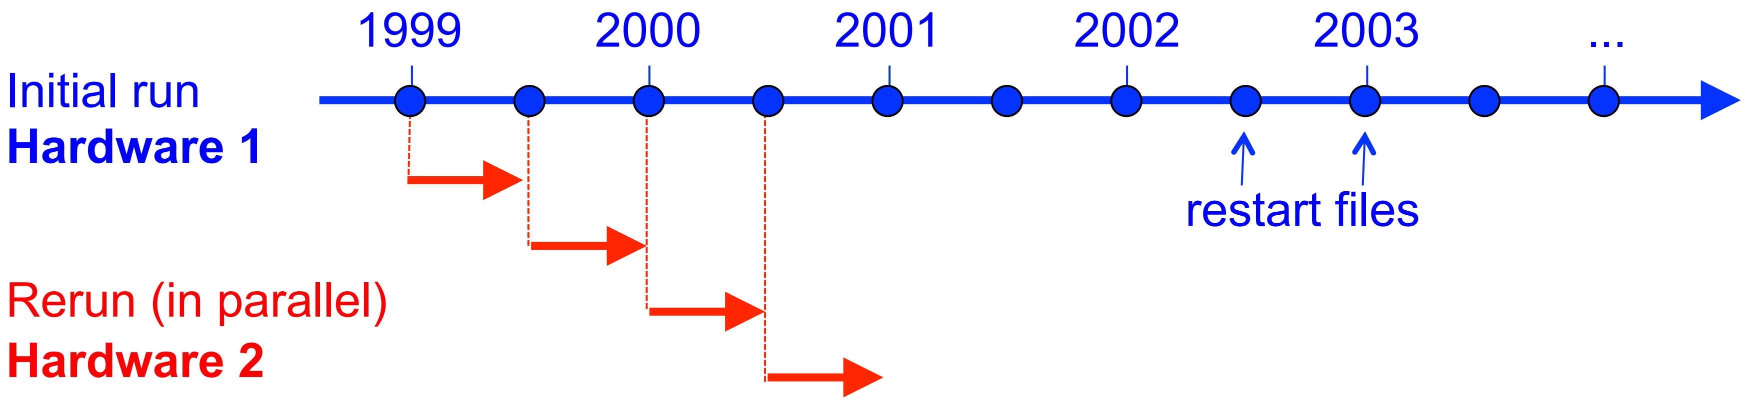

- **Requires bit-reproducible code across hardware platforms**

  Initial work (Andrea Arteaga, Oliver Fuhrer and Torsten Hoefler) suggests this is feasible at reasonable costs and partly guaranteed by IEEE standards

peak performance: $6x10^{15}$ Flop/s

What is feasible today?

- **Decade-long European-scale simulations**
  - Able to run 1 day in 20 minutes (1 year in 5 days)
  - Domain-decomposition with 12 x 12 domains, each running on a GPU/CPU node (144 nodes, 2.8% of PizDaint)

- **CORDEX simulations – Strong scaling: Increase # of nodes for given domain**
  - Resolution of 12 km, 150 year long, domain covering Europe. Able to run 1 year in 18 hours on 10 nodes (poor strong scaling)
  - On dedicated PizDaint (5272 nodes): **Large 500-member ensemble feasible**

- **Global simulations – Weak scaling: Increase domain size with # of nodes**
  - Exploit excellent weak scaling on dedicated Piz Daint (5272 nodes): At a resolution of 2.8 km, whole planet could be covered.
  - **In principle, global convection-resolving AGCM simulations feasible today!**
  - **Would require online analysis (I/O bandwidth becomes critical bottle neck).** See project crCLIM at ETH: http://www.c2sm.ethz.ch/research/crCLIM.html

References

Ban, N., J. Schmidli and C. Schär, 2014: Evaluation of the convection-resolving regional climate modelling approach in decade-long simulations. J. Geophys. Res. Atmos., 119, 7889–7907, http://dx.doi.org/10.1002/2014JD021478

Ban N., J. Schmidli and C. Schär, 2015: Heavy precipitation in a changing climate: Does short-term summer precipitation increase faster? Geophys. Res. Lett., 42 (4), 1165–1172 http://dx.doi.org/10.1002/2014GL062588

Bellprat, O., S. Kotlarski, D. Lüthi, R. De Elìa, A. Frigon, R. Laprise, C. Schär, 2016: Objective calibration of regional climate models: Application over Europe and North America. J. Climate, 29, 819-838, http://dx.doi.org/10.1175/JCLI-D-15-0302.1

Fischer, A.M., D.E. Keller, M.A. Liniger, J. Rajczak, C. Schär, C. Appenzeller, 2015: Projected changes in precipitation intensity and frequency in Switzerland: a multi-model perspective. Int. J. Climatol., 35 (11), 3204-3219, http://dx.doi.org/10.1002/joc.4162

Fuhrer, O., Osuna, C., Lapillonne, X., Gysi, T., Cumming, B., Arteaga, A., and Schulthess, T. C., 2014: Towards a performance portable, architecture agnostic implementation strategy for weather and climate models, Super- comp. Front. Innov., 1, http://superfri.org/superfri/article/view/17

Gysi, T., Osuna, C., Fuhrer, O., Bianco, M., and Schulthess, T. C. 2015: STELLA: A Domain-specific Tool for Structured Grid Methods in Weather and Climate Models, in: Proceedings of the International Con- ference for High Performance Computing, Networking, Storage and Analysis, SC ’15, pp. 41:1–41:12, http://dx.doi.org/10.1145/2807591.2807627

Keller, M., O. Fuhrer, J. Schmidli, M Stengel, R. Stöckli, and C. Schär, 2016: Evaluation of convection-resolving models using satellite data: The diurnal cycle of summer convection over the Alps. Meteorol. Z., published online, http://dx.doi.org/10.1127/metz/2015/0715

Kerkhoff, C., H.R. Künsch and C. Schär, 2014: Assessment of Bias Assumptions for Climate Models. J. Climate, 27, 6799-6818, http://dx.doi.org/10.1175/JCLI-D-13-00716.1

Kröner, N., S. Kotlarski, E. Fischer, D. Lüthi, E. Zubler, C. Schär, 2016: Separating Climate Change Signals into Thermodynamic, Lapse-rate and Circulation Effects: Theory and Application to the European Summer Climate. Climate Dynamics, online, http://dx.doi.org/10.1007/s00382-016-3276-3

Langhans, W., J. Schmidli and C. Schär, 2012: Bulk convergence of kilometer-scale simulations of moist convection over complex terrain, J. Atmos. Sci., 69 (7), 2207-2228, http://dx.doi.org/10.1175/JAS-D-11-0252.1

Leutwyler, D., O. Fuhrer, X. Lapillonne, D. Lüthi, C. Schär, 2016: Towards European-Scale Convection-Resolving Climate Simulations. Geophys. Model Development, in press, http://dx.doi.org/10.5194/gmd-2016-119

Leutwyler, D., et al., 2017: Evaluation of the Convection-Resolving Climate Modeling Approach on Continental Scales, in preparation

Schär, C., N. Ban, E.M. Fischer, J. Rajczak, J. Schmidli, C. Frei, F. Giorgi, T.R. Karl, E.J. Kendon, A.M.G. Klein Tank, P.A. O’Gorman, J. Sillmann, X. Zhang, F.W. Zwiers, 2016: Percentile indices for assessing changes in heavy precipitation events. Climatic Change, published online, http://dx.doi.org/10.1007/s10584-016-1669-2

Animations: can be downloaded via crCLIM website at http://www.c2sm.ethz.ch/research/crCLIM