On the Paths to Exascale:

Will We be Hungry?

Presentation by Mike Rezny, Monash University, Australia
michael.rezny@monash.edu

4th ENES Workshop
High Performance Computing for Climate and Weather
Toulouse, 7th April 2016

Credits

In collaboration with:

Reinhard Budich, MPI-M

And from discussions with:

Mick Carter, Met Office

Bryan Lawrence, NCAS

V. Balaji, NOAA/GFDL

Sponsors

ARC Centre of Excellence for Climate System Science

Monash University

Met Office

MPI-M

NERC

IS-ENES2

ENES

ESiWACE

Introduction

Developing, maintaining, executing, and analysing Weather and Climate Models (ESMs) is already an extremely challenging and complex task!

The progress to date has been achieved by teams of very talented, experienced, dedicated, and tenacious scientists and model developers.

Let's consider two types of complexity:

- Science Complexity
- Computational (IT) Complexity

Development of the Hadley Centre Earth System Model

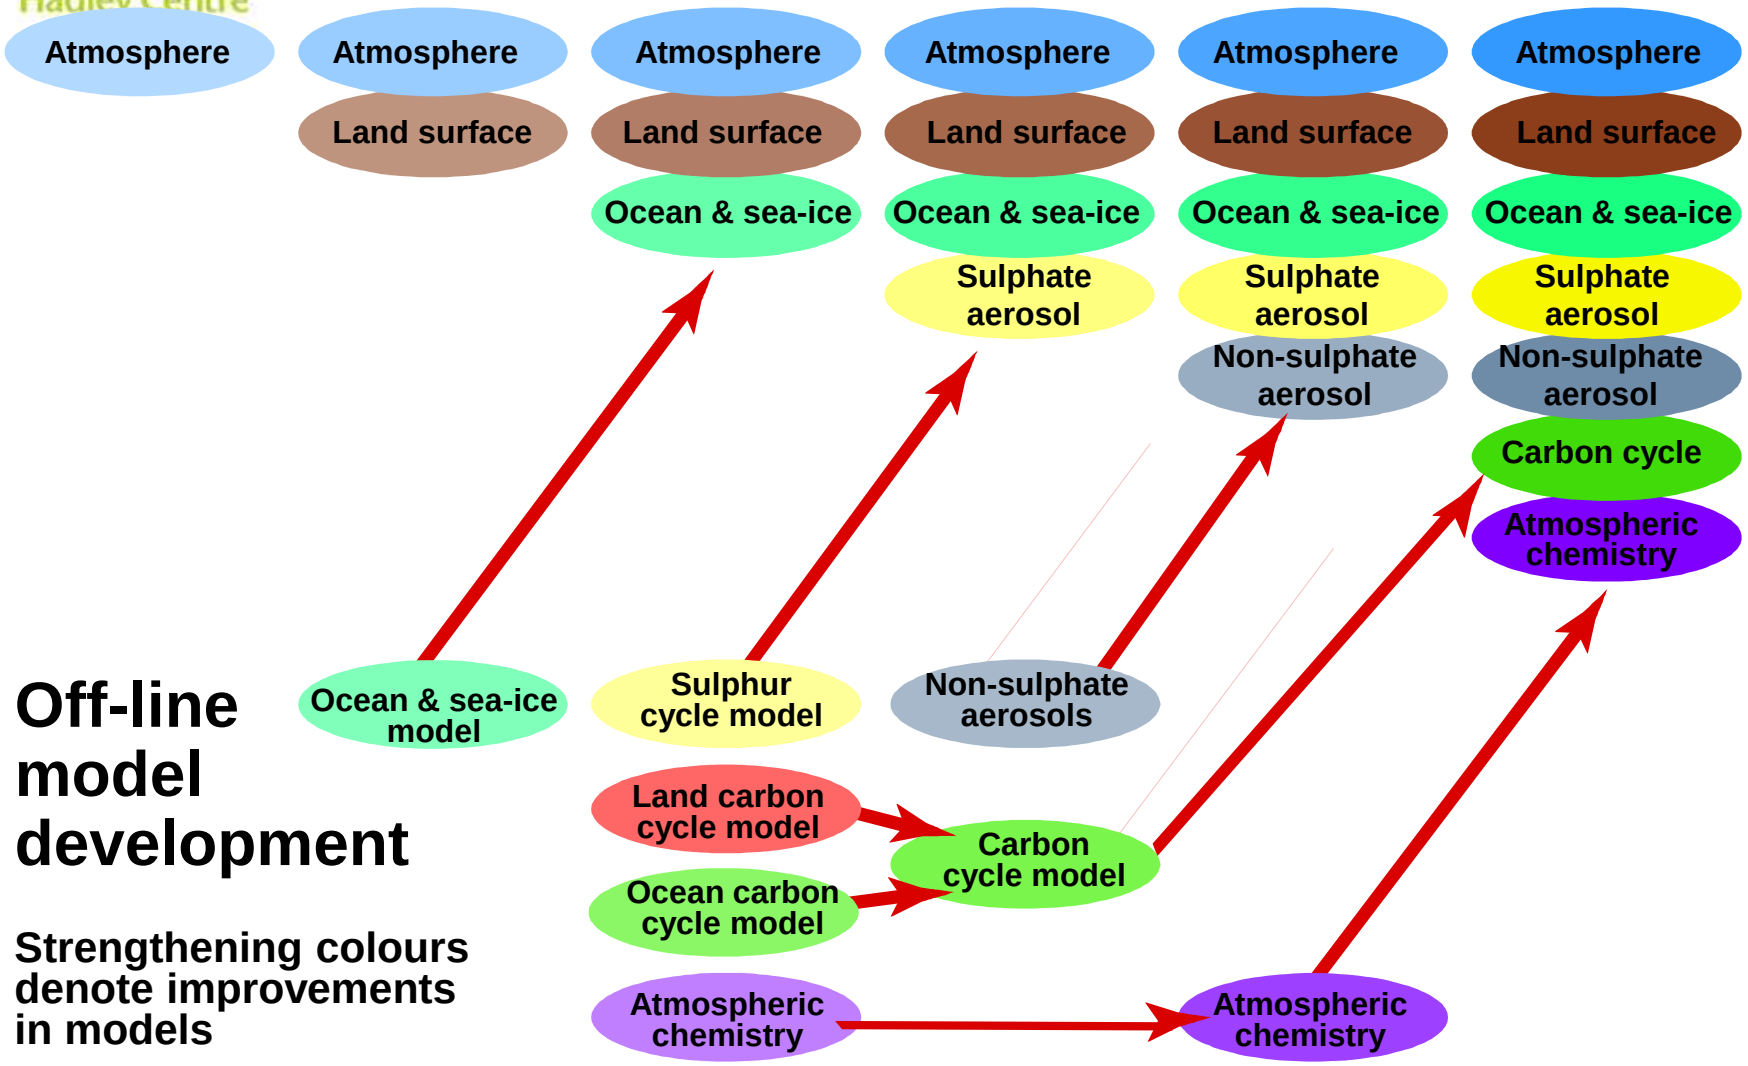

Off-line model development

Strengthening colours denote improvements in models

Science Complexity is increasing:

- More components
- More complex interactions between components
- More derived data
- More complex and much larger ensembles
- More observational data for assimilation
- Higher resolution
- More complexity in the grids / meshes
  - Unstructured, extruded meshes
  - Multiple meshes
  - Multi-grid methods
  - High-order, mixed finite element methods

Boundary Conditions:

- Little or no increase in the performance of each core
- Increased performance is needed to meet research and operational requirements

Computational Complexity is increasing:

- More cores per socket are provided
- Performance improvement is mainly by higher levels of parallelism
- Heterogeneous hardware architectures with accelerators
- Increasingly complex solutions are being applied to achieve the required performance:
  - Numerics
  - Algorithms
  - Software design
  - Workflows

Difficulty is in applying these **complex computational** solutions to existing large **complex science** applications

One Example of Increased Computational Complexity: **Concurrency**

Present and future models may need to incorporate up to **7** levels of concurrency to meet performance requirements with anticipated hardware.

- Vectorisation
- Distributed-memory concurrency across nodes / sockets (MPI)
- Shared-memory concurrency within nodes / sockets (OpenMP)
- Explicit coupling of independent models (Oasis-MCT, Cpl, YAC)
- Asynchronous I/O (XIOS, I/O servers)
- Executing model components concurrently (ESMF, FMS)
- Hardware accelerators (GPU, OpenACC)

The “**FREE**” Lunch is over!

Was there ever a Free Lunch?

But certainly, the menu choices on offer have increased over time

MegaFlop Diner

GigaFlop Café

TeraFlop Bistro

PetaFlop Brasserie

Le Restaurant Grand Exascale

Is the "**FREE Lunch**" over?

because at

Le Restaurant Grand Exascale

We are told:

"**FLOPS** are **FREE!**"

Are YOU Hungry?

What’s the catch?

Yes, the FLOPS might be FREE!

But, the catch is:

it is difficult to reach the Restaurant!

The PATHs to EXASCALE

This is not a PATH to EXASCALE

This is a PATH to EXASCALE

Why the **Grand Canyon**?

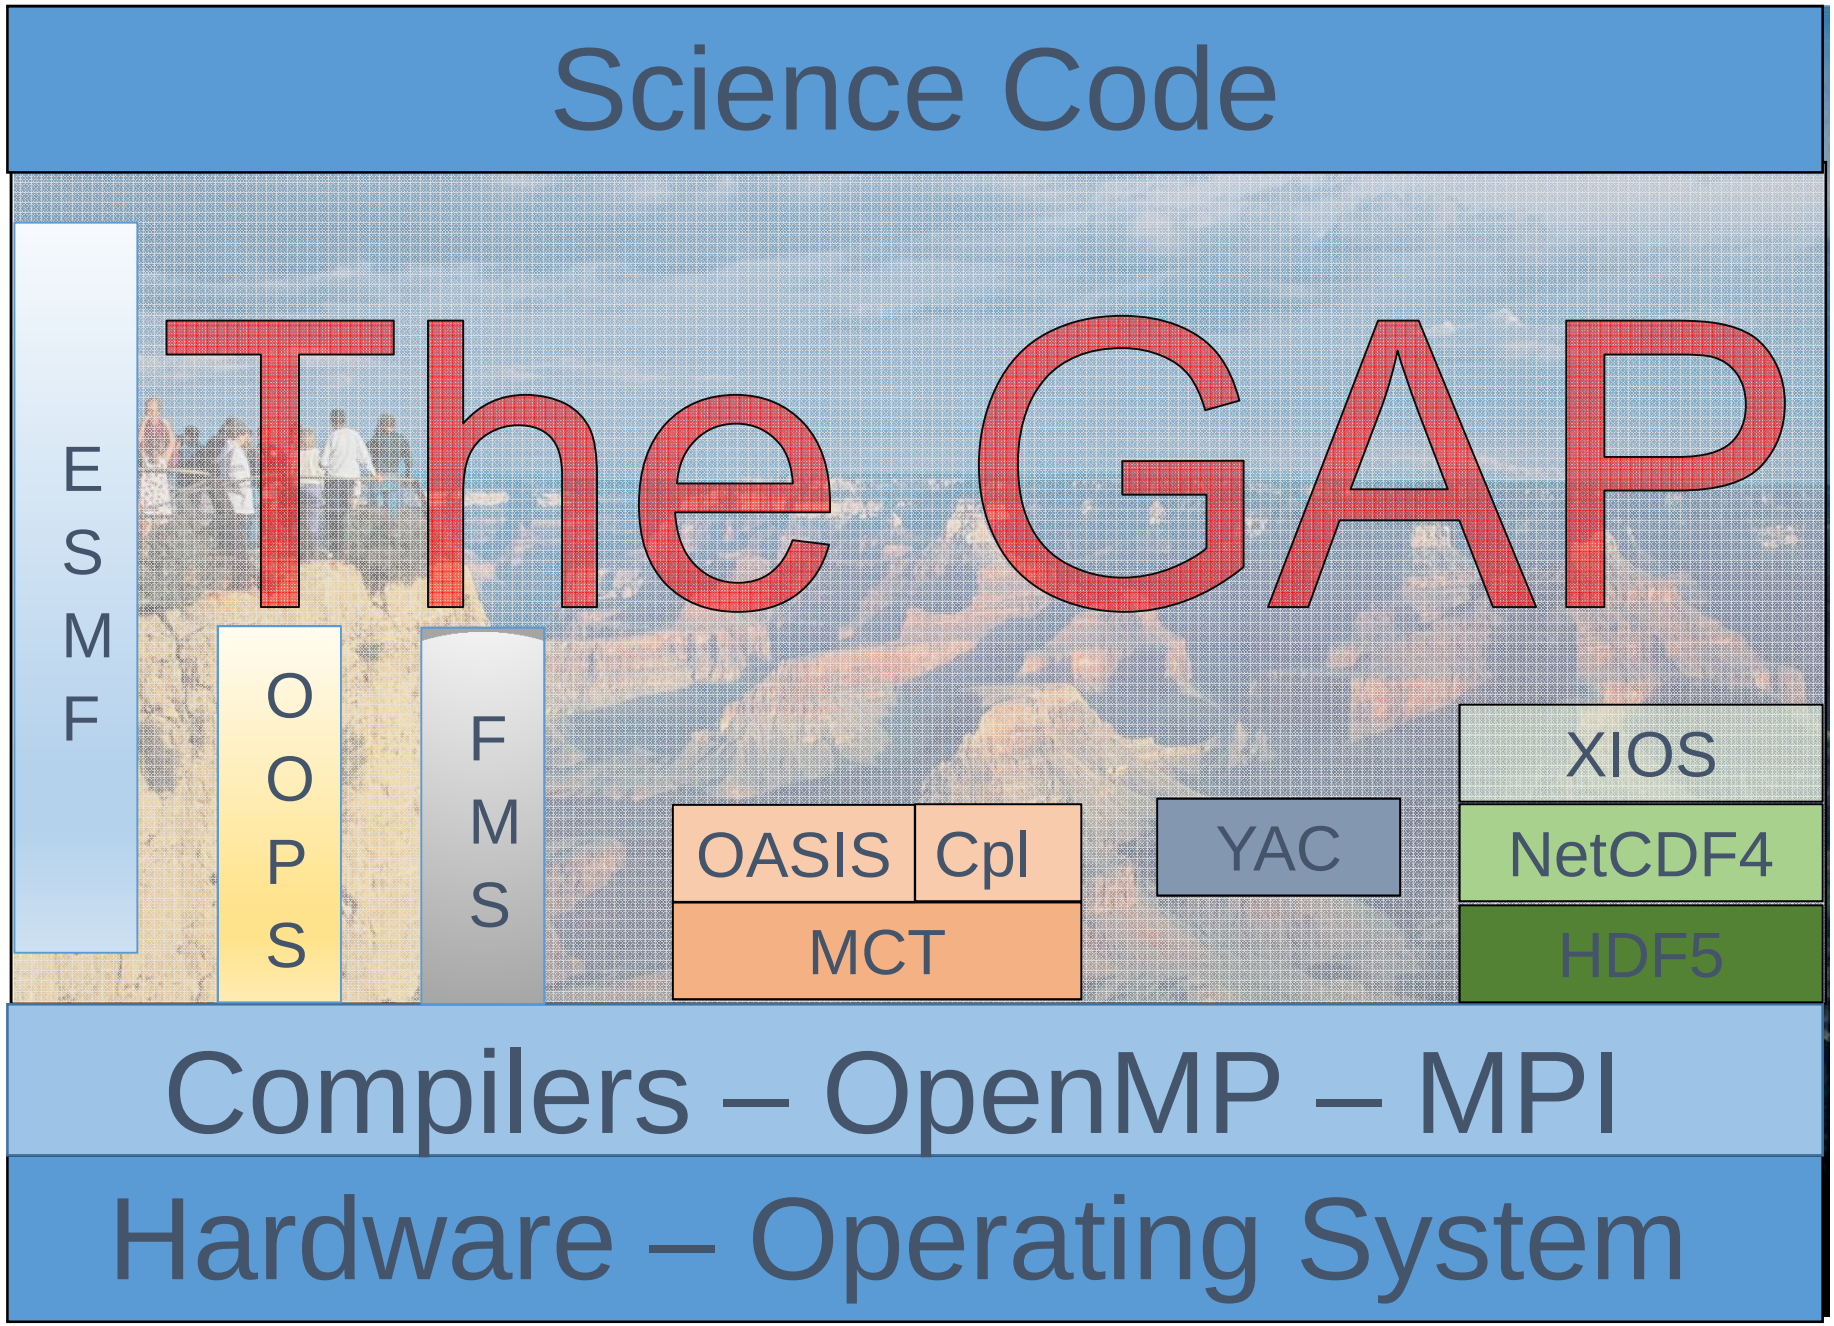
Science Code
The GAP
ESMF
OOPS
FMS
OASIS
Cpl
MCT
YAC
XIOS
NetCDF4
HDF5
Compilers – OpenMP – MPI
Hardware – Operating System

The commonly-used support libraries in ESM Development are just too low-level

Using MPI is similar to Assembly Language Programming

```
ld hl,Result ; load the address of the Result variable in HL
ld (hl),a ;    store value of A into the byte pointed by HL
```

```
MPI_Send(buf, count, type, dest, tag, comm, err)
MPI_Recv(buf, count, type, source, tag, com, status, err)
```

There is not even a commonly-used standard library that provides the most basic functions applicable to ALL ESMs:

Clock, Calendar, and Event handling

MPI has been around for over 20 years!

There is no commonly-used standard library to support halo swapping on rectangular grids.

Higher-level tools do exist that provide various functionalities needed for ESM development.

But they are usually developed to meet the specific needs of only a subset of potential users.

Trying to combine a number of tools and adapt them to another model is often difficult.

How do we intend to cross **The Gap**?


671
U.S.NAVY


www.humancannonball.us
IMERYS
IMERYS

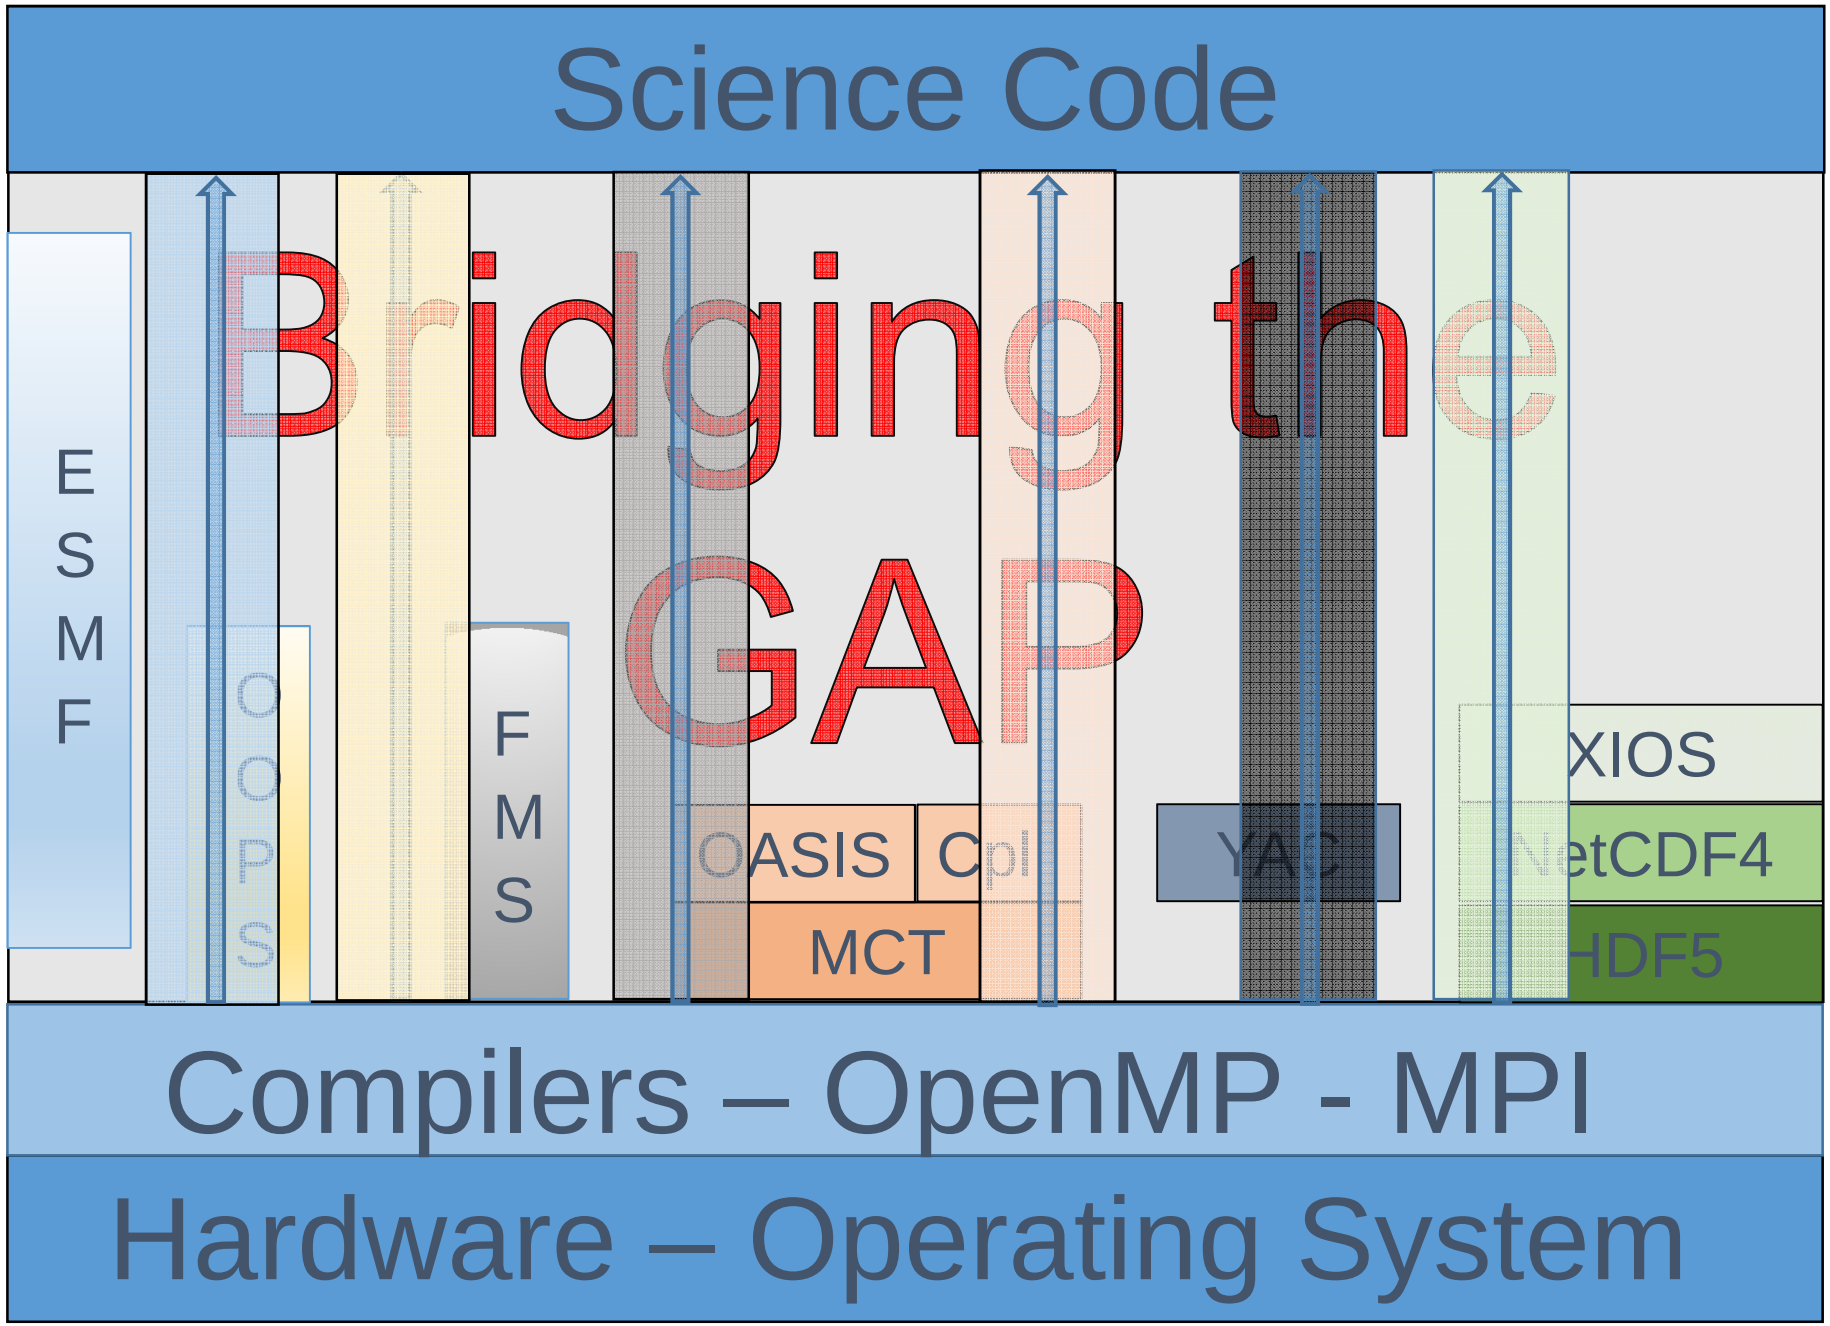
Science Code
Bridging the GAP
ESMF
COUPS
FMS
OASIS
Cpl
MCT
YAXT
XIOS
NetCDF4
HDF5
Compilers – OpenMP - MPI
Hardware – Operating System

Can this be the PATH to EXASCALE?

ESMF is one example of a package with a rather complete set of functionality that attempts to

Bridge the Gap.

Why has it not had more acceptance by major ESM developers?

What requirements does it not fulfil for wider acceptance?

Presentation by V. Balaji, Honolulu, **2004**:

Proposal for an Earth System Modeling Environment (ESME)

We seek to **unite** the data (ESG) and model (ESMF) **communities** with climate scientists (IPCC, CMIP) to develop the model metadata layer, and the relational database of models and data that would be based on it.

This effort would be closely allied with the PRISM / CAPRI efforts in the same domain.

To Bridge the Gap or try to Reduce the Gap:

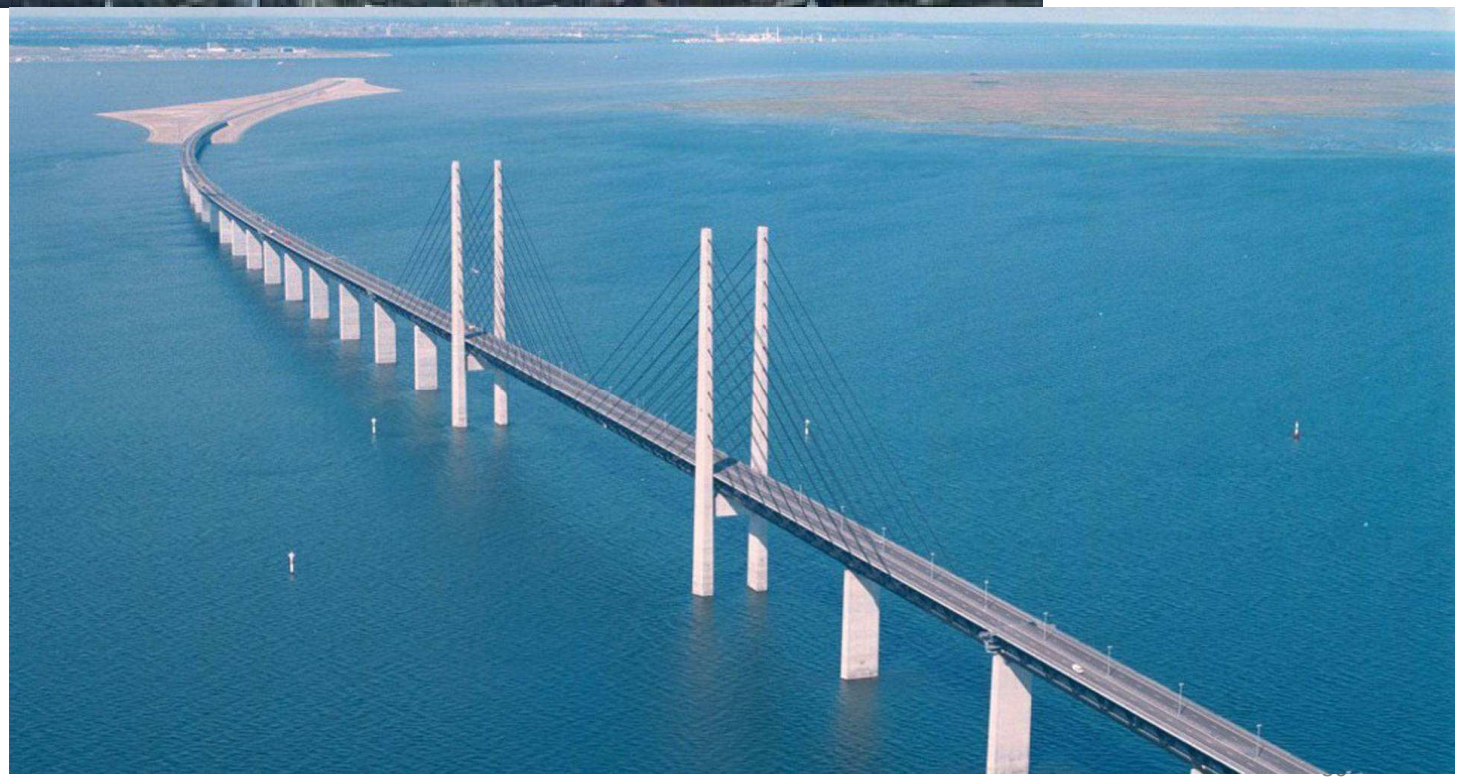

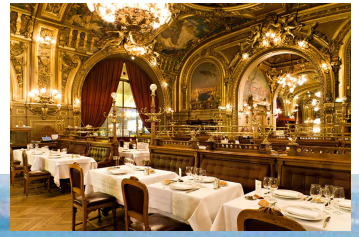

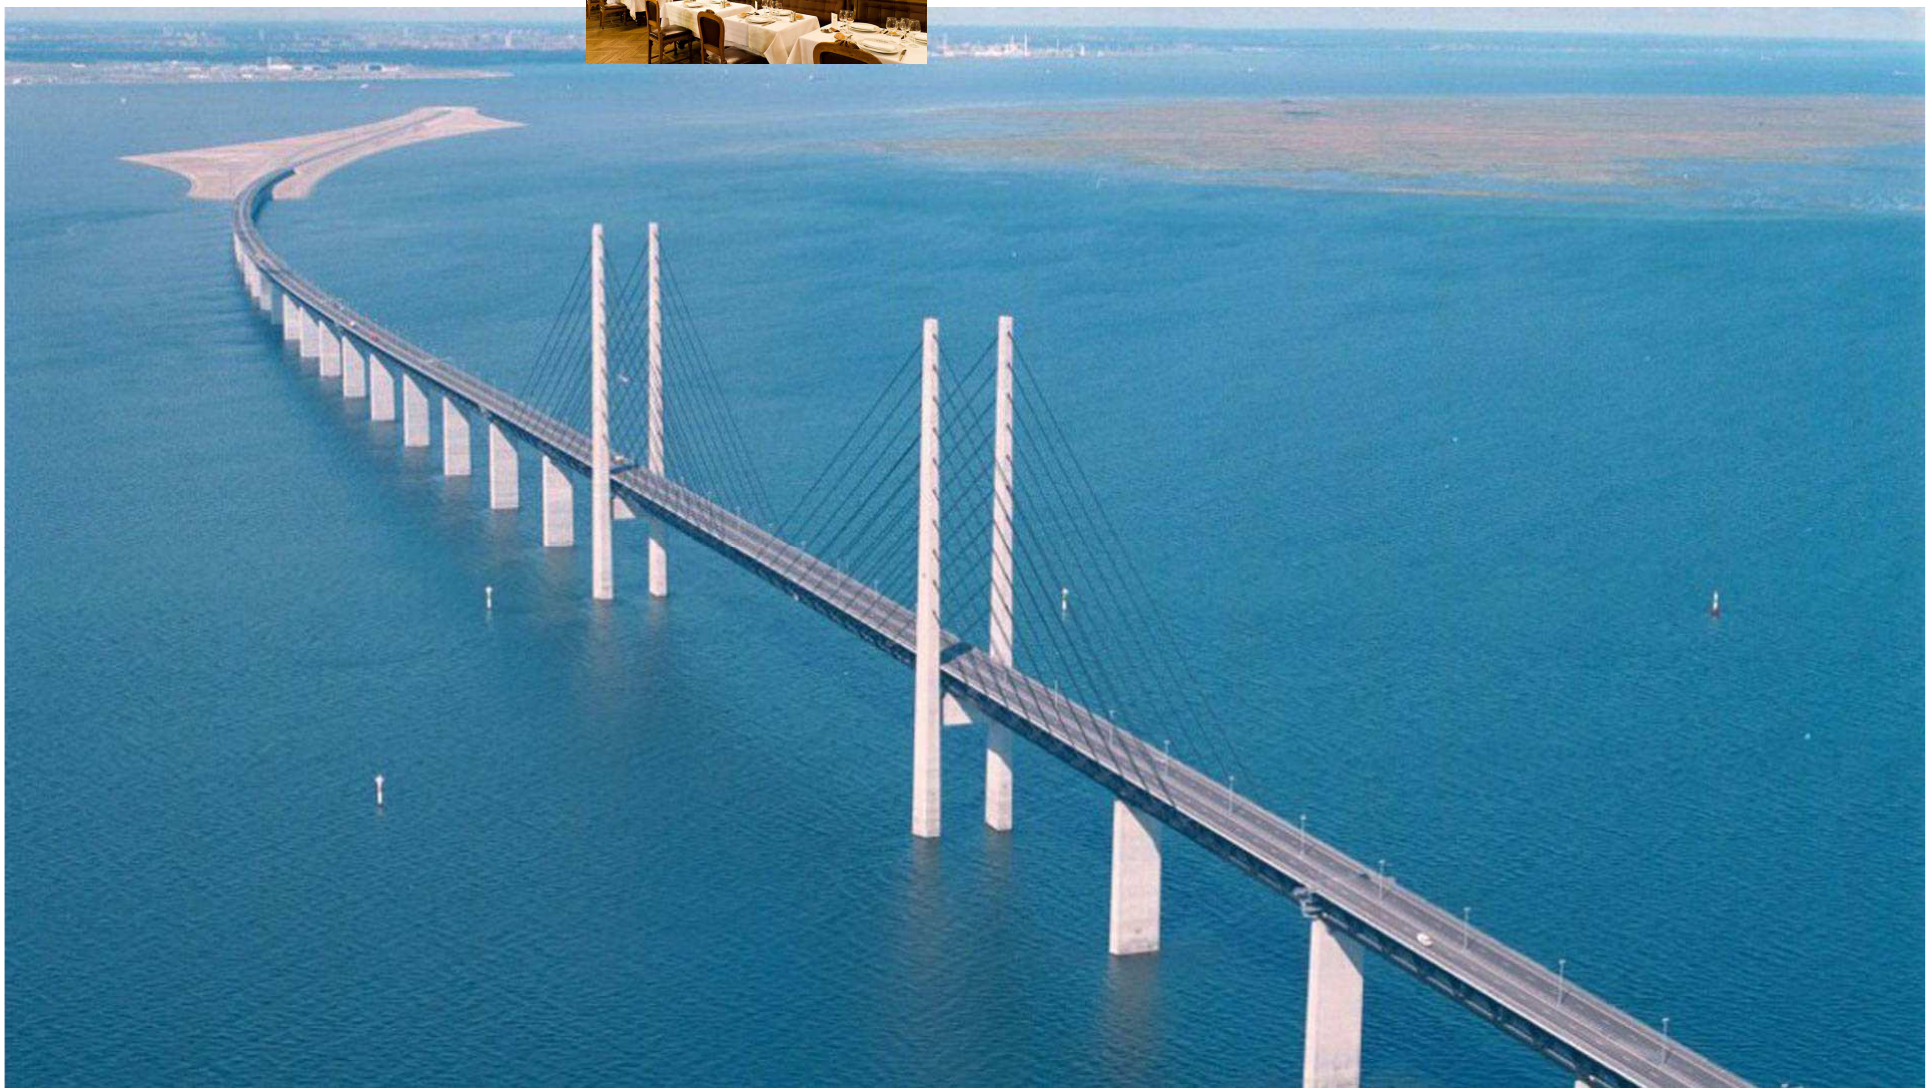

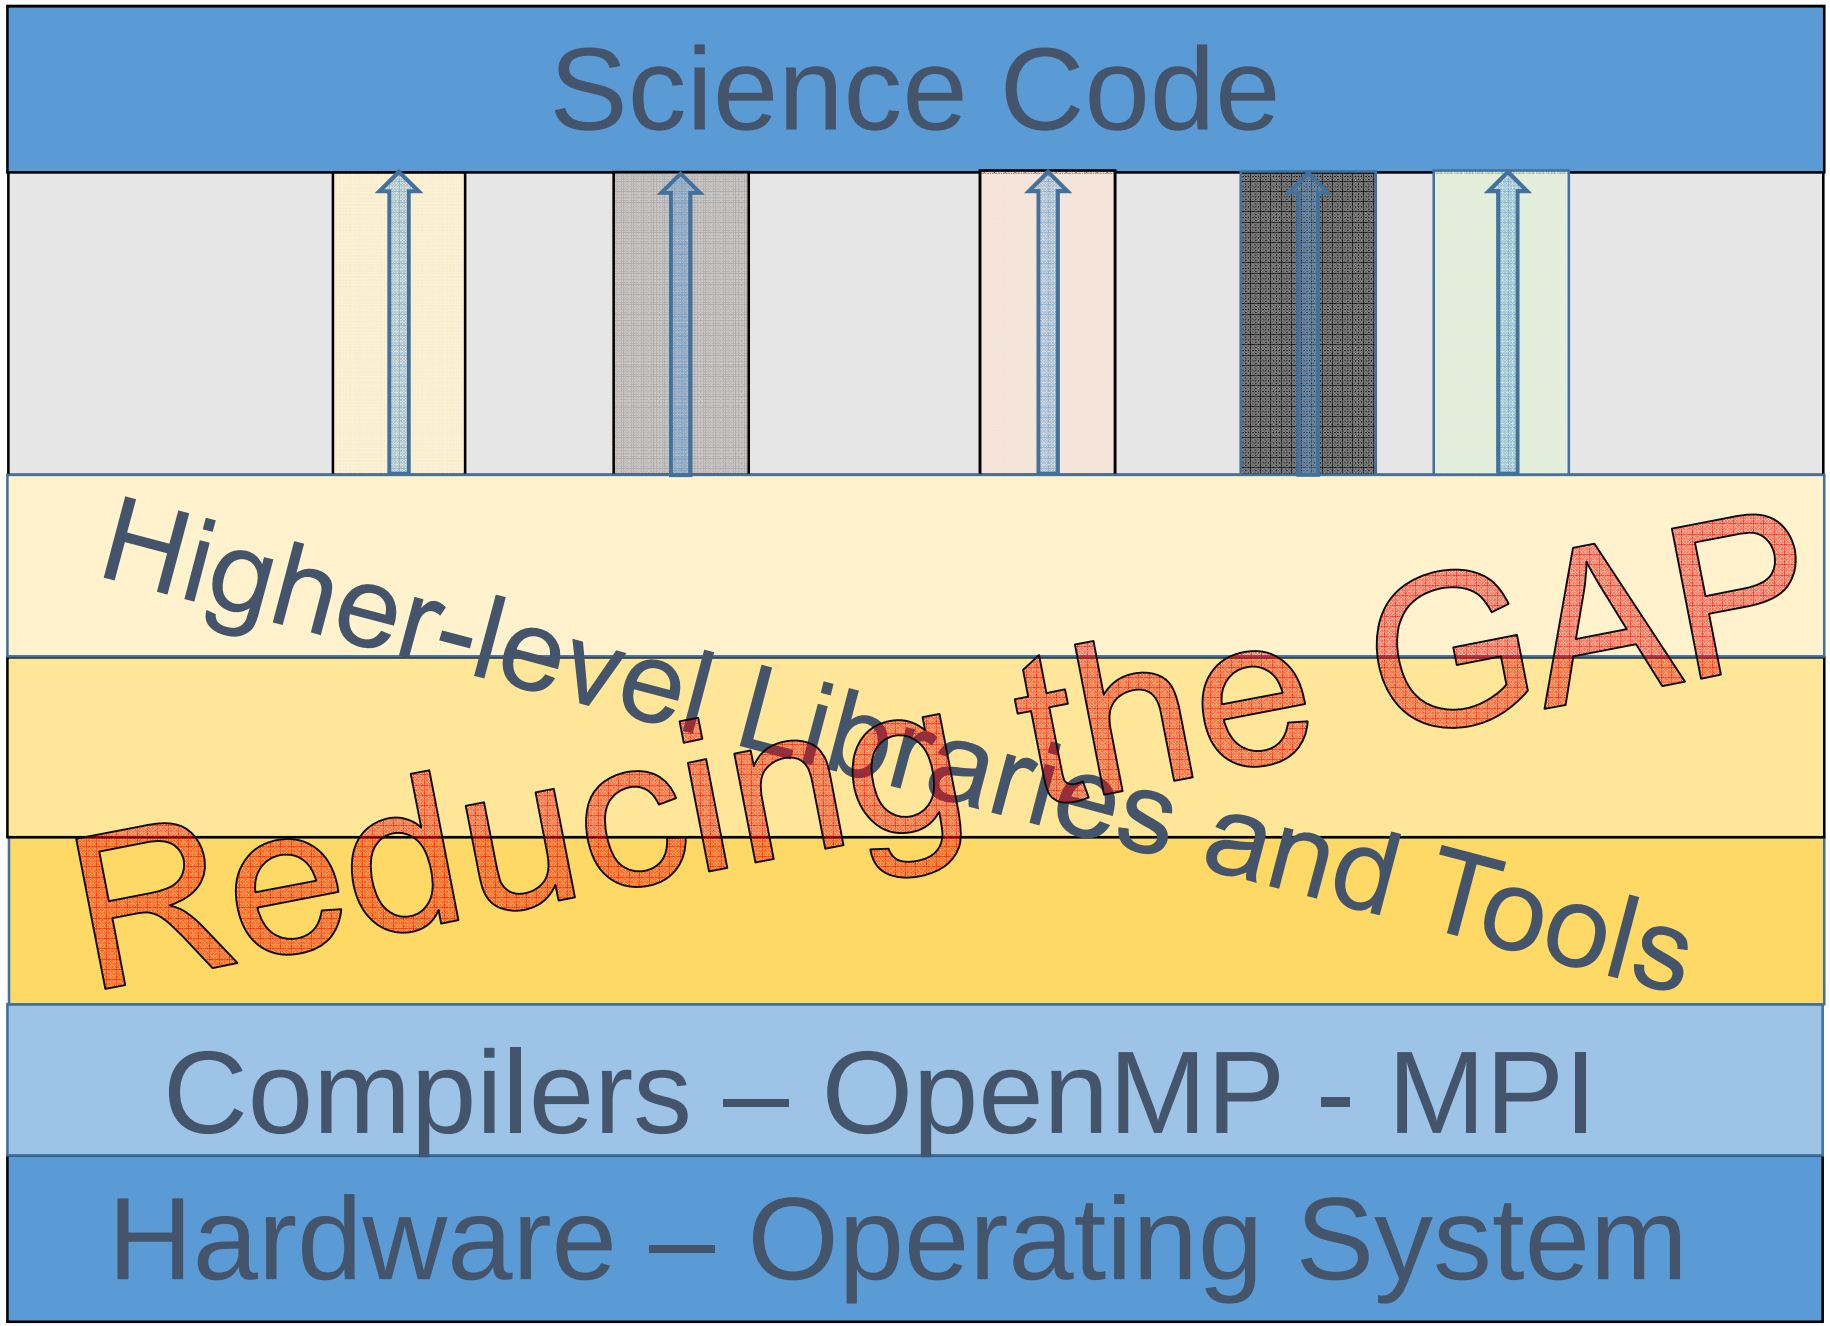
Science Code
Higher-level Libraries and Tools
Reducing the GAP
Compilers – OpenMP - MPI
Hardware – Operating System

Can we, as a community, work collaboratively to:

- Discuss,
- Specify,
- Design,
- Develop,
- Maintain, and
- Document

libraries and tools that will help all of us to

Progressively Reduce the Gap in ESM development?

The motivation for this will come with the realisation, by our community, that the

significant challenges to providing

ESM Infrastructure are too complex and our resources too limited

for it to be developed individually by each modelling centre.

What is necessary is a **Methodology** we all agree upon that will assist us towards achieving the collaborative development of the tools we all need in ESM!

Including Exascale Development Support

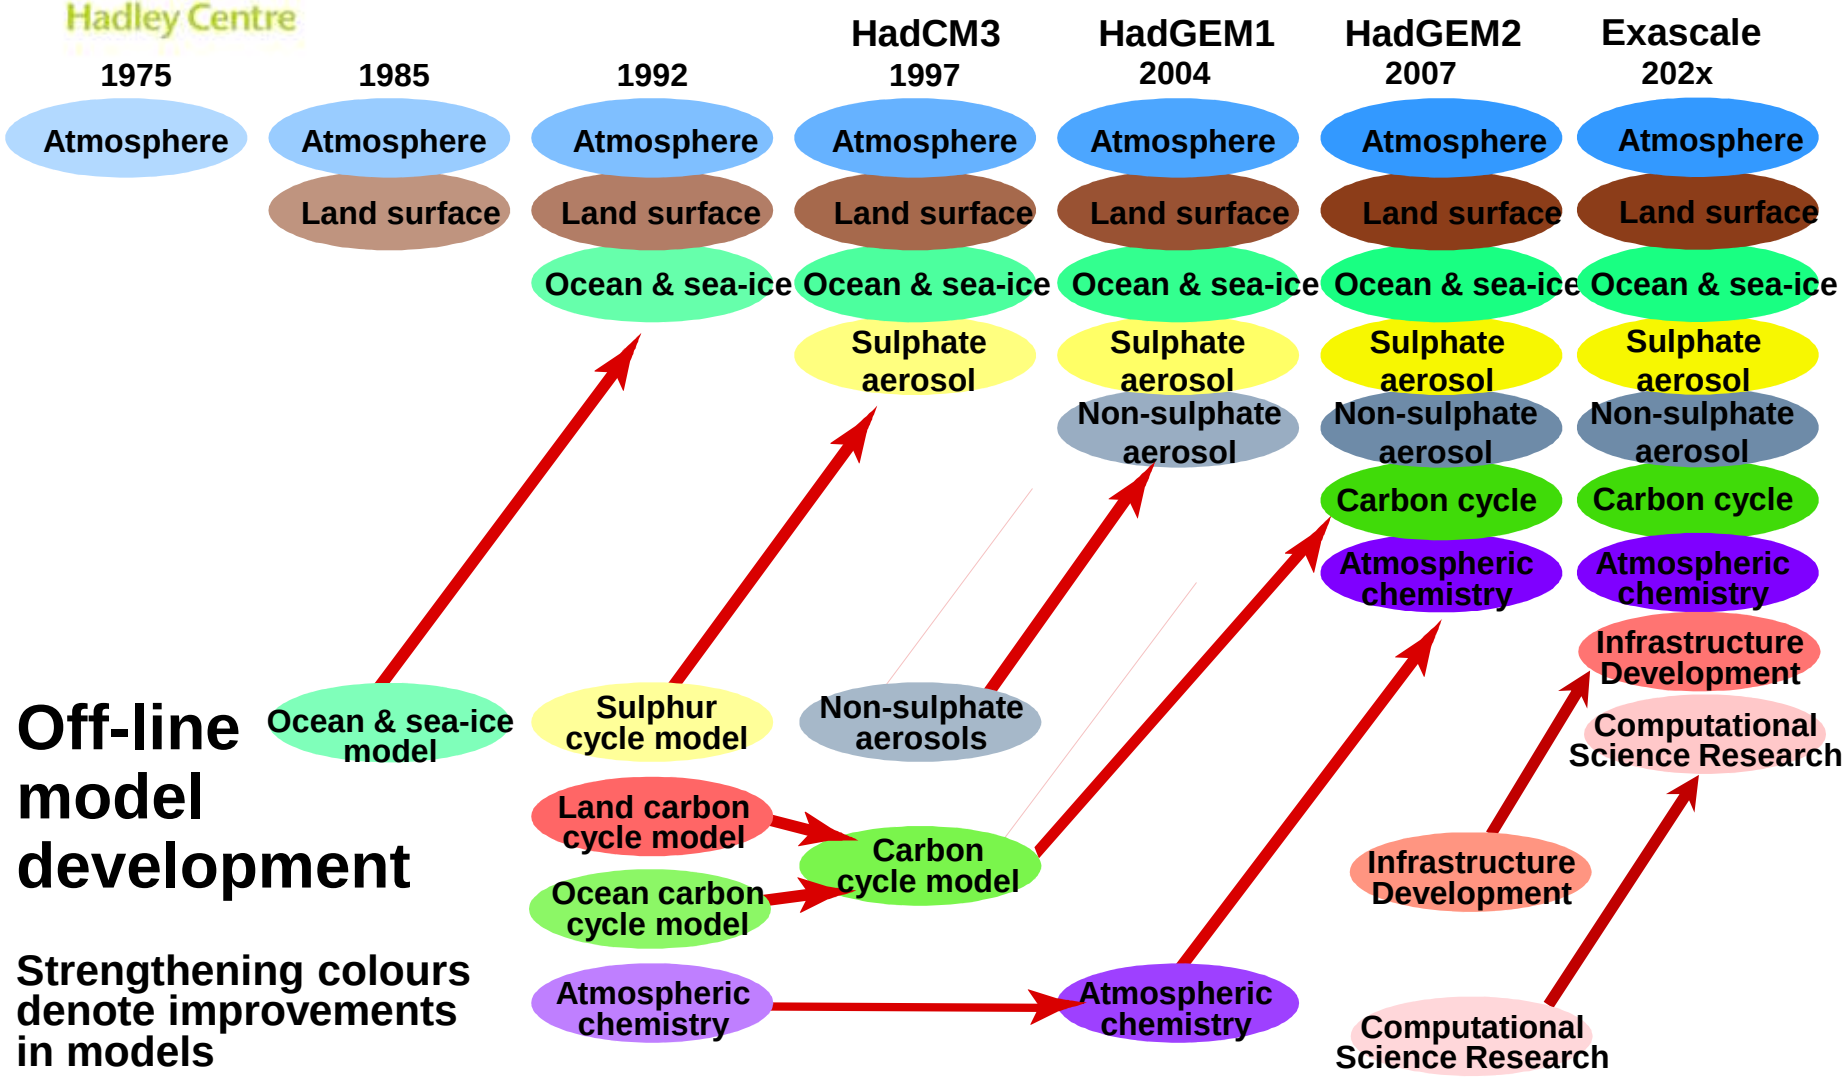

There are significant precedents in other communities:

Large Hadron Collider,
Square Kilometer Array,
Hubble Telescope,
Earth Observation Satellites

Can we do better on the way to Exascale?

Can we harness the collective wisdom, experience, and determination that exists within our community to progressively

Reduce the Gap

And perhaps one day eventually

Bridge the Gap?

Earth

System

Modelling

Standards

ESMS

should NOT be a plan or project to deliver a complete ESM Infrastructure.

ESMF already did something like that and it has not had wide acceptance by the major Climate and Weather Communities.

ESMS

should be a **methodology** to facilitate and encourage collaboration in building key parts of ESM Infrastructure that can be used by the entire ESM community.

ESMS

Examples of **Internal Focus**:

- A common library supporting Clock, Calendar, and Event Handling
- Common tools for mesh generation and partitioning
- Fortran/C/C++ Unit testing framework extended to include Performance Testing for MPI, OpenMP, Hybrid, and Hardware Accelerators
- A common library, such as MCT, used by all couplers.
- File formats for unstructured meshes
  (ugrid, GridSpec?)
- Regridding tools and libraries

ESMS

Examples of **External Focus**:

- Ensure that:
  - Fortran
  - MPI
  - OpenMP
  - OpenAcc
  - NetCDF / HDF5

  develops accounting for ESM community needs.

ESMS

Another example of **External Focus**:

- Work with vendors to provide a unified voice to express ESM hardware needs
- Provide vendors with standard libraries that can be supported and optimized for our entire community, in a similar way as:
  - OpenMP, MPI, BLAS, LAPACK, ScaLAPACK, FFTW, PETSc, etc.

ESMS

Develop an **Ontology** for the ESM community

A formal specification which clearly defines the terms in our domain and the relationships between them

We **WILL** all need to find a
Path to Exascale!

The difficulty in getting there is in our hands.

ESMS

would help us to **Reduce the Gap** and
share some of the Path to Exascale.

And provide some food for us all since we are still hungry!

Thank You!

Any Questions?