Digitals Twin Thinking for HPC in Weather and Climate

Bryan Lawrence

NCAS &
University of Reading: Departments of Meteorology and Computer Science

19th ECMWF HPC Workshop

Outline

Digital Twins

What is a digital twin? Google says"

- "A digital twin is a digital representation of a physical object, process or service. ... A digital twin is, in essence, a computer program that uses real world data to create simulations that can predict how a product or process will perform."

You would be forgiven for thinking that sounds rather like what we call a model!

Destination Earth defines a digital twins as

- "mirrors of reality, simulators that replicate reality constrained by **real time data**."

You would be forgiven for thinking that sounds like data assimilation.

We have been doing digital twins for a long time!

Information Intervention?

Some[1] have postulated a "leap in information intervention" as part of delivering weather and climate digital twins. What could this mean?

- *"the challenge will be to design a digital twin that allows users to intervene, extract information and influence the system trajectory across time and space, as done — albeit often unwittingly — in the real world."*
- *"tasks like simulations, ..., post-processing ...are ...executed on federated computing infrastructures, feed data into virtual data repositories with **standardized metadata**, and from which a heavily ML-based toolkit extracts information that can be manipulated in any possible way."*
- *"We can literally grab the earth, anyone, and see what the consequences of their actions mean for how things will evolve, what if ..."*
  - (We have to put "their actions" into our model, whether that is an a representation of agricultural policy (a model) or some other process or action.)

[1] Quotes from Bjorn Stevens SC20 talk and Bauer et al "A digital twin of Earth for the green transition".

Information Intervention?

Some[1] have postulated a "leap in information intervention" as part of delivering weather and climate digital twins. Wh

- "the challenge will be to design a digital twin tha intervene, extract information and influence the s time and space, as done — albeit often unwitting ... the real world."
- "tasks like simulations, ..., post-processing ...a federated computing infrastructures, feed data i repositories with **standardized metadata**, and f ML-based toolkit extracts information that can be possible way."
- "We can literally grab the earth, anyone, and see ... the consequences of their actions mean for how things will ... ..."
  - (We have to put "their actions" into our moc representation of agricultural policy (a mode process or action.)

Computational Steering

Data mgmnt. & **workflow engines** to allow (arbitrary?) post-processing.

Embedding and/or **coupling** (arbitrary?) models in workflow.

[1] Quotes from Bjorn Stevens SC20 talk and Bauer et al "A digital twin of Earth for the green trans

Weather and Climate Digital Twins?

"predict how a product or process will perform"
maybe
"project a set (Y) of possible futures and how they interact with X"
($X =$ a **set** of possible societally relevant systems/actions etc)

Weather and Climate Digital Twins?

"predict how a product or process will perform"
maybe
"project a set (Y) of possible futures and how they interact with X"
$(X =$ a **set** of possible societally relevant systems/actions etc)

Leads to Two Key Promises:

Scenario Evaluation

If we do this, then what? Typically being expressed in terms of climate, even if the interest is in high impact weather.

Democratising Access

Going beyond data access, so that those interested in X can exploit the twin in some new ways.

Implementation?

Leads to some questions:

1. What do we mean by scenarios? (Who decides and defines the Y scenarios? How does the X community interact with the decision and definition?)
2. What do we mean by "going beyond" and "data access" to deliver X (and for whom, "democracy" is a very inclusive word?)
3. How is this different from existing practice?

Abstract Practicalities

Assumptions:

1. At scale, scenarios will be hierarchical: from a small set of expensive simulations Y we can investigate X different application scenarios. We might think of X as applications.
2. To be meaningfully different from existing practice the implementation of “going beyond data access” for X must be real.
3. X involves running additional code and or changing parameters/inputs for the code running Y.
4. True computational steering of Y can only be done by a very small number of “users” (possibly as few as one).
5. Recomputation of Y scenarios is not generally affordable to support arbitrary X (next slide)

To compute or re-compute?

- The IPSL contribution to CMIP6 simulated o(50K) "useful" years (and o(150)K years overall), and in doing so consumed o(6) TJ, and stored o(1) PB of data for ongoing analysis.[1]
  - $=>$ making 1 PB of useful data cost 2TJ
- JASMIN currently has o(50)PB of spinning disk and consumes under 250kW[2]
  - $=>$ storing 1 PB of useful data for a year costs 0.16TJ

But

- We might argue that costs for compute will fall faster than for (spinning disk) storage.
- We might then argue that will only happen until we make more use of SSD, then compute and storage will fall in tandem?
- …and then there is tape. Near zero energy cost when at rest!

[1] Rounded from figures compiled by IS-ENES3 and reported by M. ACosta at the IS-ENES3 General Assembly, March 2020

[2] Not including cooling, but including a 10K core CPU batch cluster.

To compute or re-compute?

- The IPSL contributi… to CMIP6 simulated o(50K) "useful" years (and o(1… …n doing so consumed o(6) TJ, and sto… …going analysis.[1]
  - => … cost 2TJ
- JASMIN … spinning disk and consumes under 25…
  - => … …r a year costs 0.16TJ

> More caveats: We can do better by looking at specific runs not big bulk averages. We need to understand usage patterns.

But

- We migh… …pute will fall faster than for (spinning…
- We might … …happen until we make more use of SS…, then compute and storage will fall in tandem?
- …and then there is tape. Near zero energy cost when at rest!

Currently vastly cheaper to store for re-use than re-compute and likely to remain so.

[1] Rounded from figures compiled by IS-ENES3 and reported by M. ACosta at the IS-ENES3 General Assembly, March 2020

[2] Not including cooling, but including a 10K core CPU batch cluster.

Phase 1: Individuals and Groups.

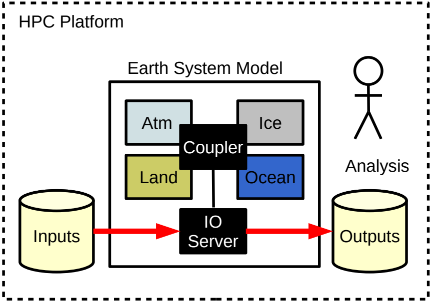

- All parties are well known to each other and communicate using personal contact.
- Often relatively inexpensive. Much re-running of simulations.
- Experiment definitions may or may not be well understood by all parties. Simulator are authors.

Phase 2: Collaborating with your mates. Multiple groups running same simulations

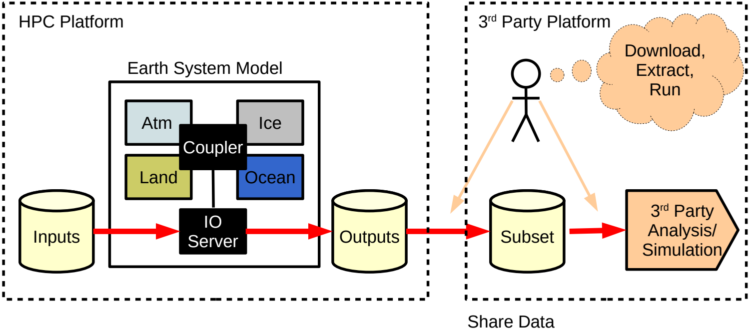

- Sharing files “as is”, with ad-hoc information systems (e.g wikis)
- Considerable overheads for data consumers (unfamiliar formats, documentation etc).
- Experiment definitions may or may not be well understood by all parties. Simulators are authors.

Phase 3: Collaboration at a distance (in time and space). Multiple groups …

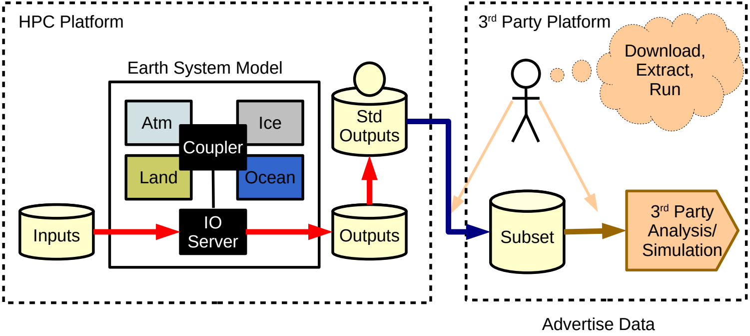

- Data producers and data consumers groups are not (well) known to each other.
- Standardised data files are described using machine readable catalogs and shared using standard protocols. Significant overheads in production to reduce friction in consumption.
- Experiment definitions may or may not be well understood by all parties, but simulators generally not authors.

From an ESM to a Digital Twin - A simple vision

A vision for Digital Twin Model Systems

- Important distinction between **coupled** components and **forced** components.
- Coupled components exchange data in memory.
- Forced components change data via (fast) disk.
- Forced components run in parallel via scheduler using **ephemeral** data.

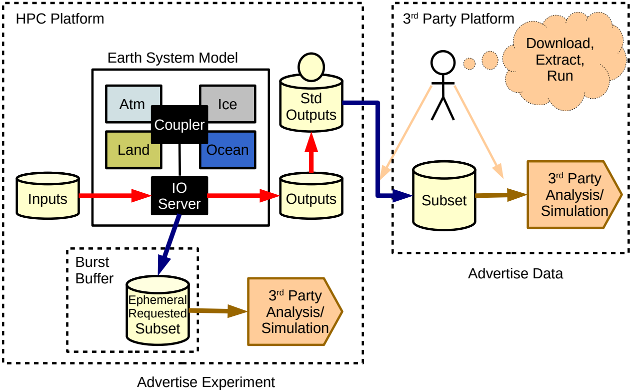

- Note importance of Data Standards for both routes!

- Community have to agree on the experiment that is run so that not only those running the ESM itself, but those exploiting **both** the ephemeral and standard data have what they need and understand the simulation context.
- Standard data goes to data lakes.

Implications - Part 1

- There is a technical route to split notions of X and Y around X being forced by Y.
- Definitions of Y need to be well promulgated ("Advertise Experiments").
  - But X community may still want to interact around definitions of Y.
- Coupling models is "business as usual", we worry about performance, portability, productivity as we always have.
- Higher frequency and/or higher resolution data can be available for a shorter time than standard outputs on high performance storage.
  - The group of users will likely still need standardised data interfaces, but are more likely to be like the phase 1/2 community.
  - The more democratic this is, the more they may need their own analysis environments.
  - Community of third party simulators using forcing data; will need HPC access and resources.

Implications - Part 2

- Interaction between scheduler, burst buffer and third party simulations needs thinking about.
  - Existing HPC concepts all about “job control”, but we might need to think a bit more about pub/sub (both interfaces and “opportunity advertisement”).
  - Those running Y may not be in a position to consider resourcing all X (or even desire to do so, nor in a position to evaluate the cost or value of X …if X is related to human behaviour it might involve very large numbers of small simulations.)
- Increasing role for virtualisation on HPC platform, to support third party computational environments. Containers not a luxury, a core service!
- Third party analysis environments can be very different technically than core simulation environments (a la https://jasmin.ac.uk).

On Advertising Experiments

Simulations are complex:

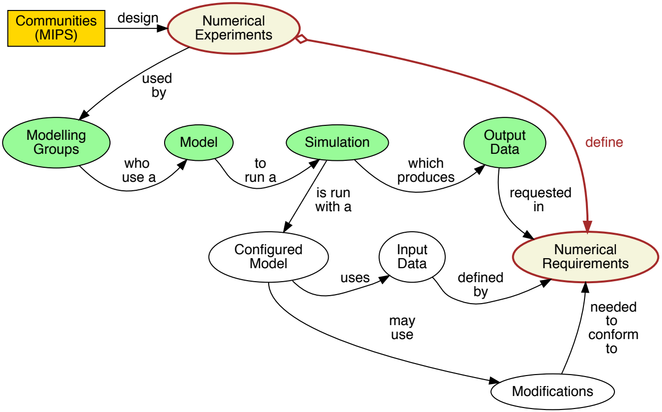

Communities need to agree on requirements

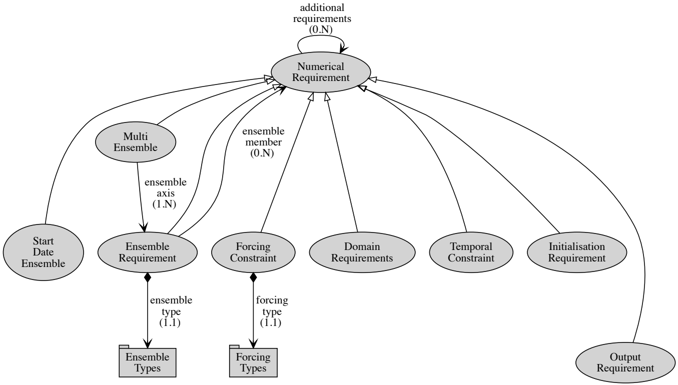

Some of these requirements will be generated by "privileged" X communities, but all X communities need to understand what has been done and why.

See Pascoe et.al (2020):https://doi.org/10.5194/gmd-13-2149-2020.

Summary

- Weather and Climate has a lot of experience with Digital Twinning.
- …but we have made promises about increased potential for "user intervention".
- We cannot afford to recompute or save everything.
- A practical route to increasing user intervention involves providing access to ephemeral data as well as more traditional data sharing.
- Ephemeral data can be used for more detailed analysis and chaining further simulation, but
- We may need to think quite hard about making that possible for others "beyond your mates". Implications for scheduling, new ideas (for HPC) about pub/sub job connection etc.