

Organizational Strategies that Helped SCREAM Win the Gordon Bell

Luca Bertagna, Hassan Beydoun, Peter Bogenschutz, Andrew Bradley, **Peter Caldwell**, Conrad Clevenger, Aaron Donahue, Chris Eldred, Jim Foucar, Chris Golaz, Oksana Guba, Ben Hillman, Jeff Johnson, Noel Keen, Wuyin Lin, Andy Salinger, Balwinder Singh, Mark Taylor, Chris Terai, Paul Ullrich, and Xingqiu Yuan

acm Association for Computing Machinery

Specific Types of Contributions
ACM Gordon Bell Prize for Climate Modelling
 Innovations in applying high-performance computing to climate modelling applications

"The Simple Cloud-Resolving E3SM Atmosphere Model Running on the Frontier Exascale System"

U.S. DEPARTMENT OF **ENERGY** Office of Science

SC23 Dependable Computing in HPC

SCREAM Lead
 Peter Caldwell
Lead Author
 Noel Keen
White
 James Taylor
Bertagna
 Luca Bertagna
Sreepathi
 Sarat Sreepathi

Conrad Clevenger
Aaron Donahue
James Foucar
Oksana Guba
Benjamin Hillman
Jayesh Krishna
Matthew Norman
Christopher Terai
Andrew Salinger
Dangqing Wu
Renata McCoy
Ruby Leung
David Bader

Let's be honest: we were lucky



“It was an audacious gamble. I’m glad it paid off”

-Congratulations from Robert Pincus

1. The Gordon Bell Climate category came along just in time
2. Timing for Frontier access was perfect
3. Success with Kokkos, OpenMP delays, and no unveiling of better alternatives was not assured

Being humble is an ingredient of success:

- Modeling honesty and humility encourages a culture of openness
- Arrogance breeds blindness and leads to mistakes

Luck comes from noticing opportunities and being flexible enough to take them

“The pessimist sees difficulty in every opportunity. The optimist sees the opportunity in every difficulty” -Churchill

- **Need:** Exascale was the central goal of E3SM, but we weren't making progress towards it
- **Challenge:** Low-res GCMs already exploit all available parallelism so no GPU speedup
- **Opportunity:**
 - High resolution models *can* benefit from GPUs while providing increased accuracy and localized information for impacts planning
 - Sandia had already ported the hydrostatic dycore to C++/Kokkos and demonstrated that it worked



Choosing science goals to fit DOE's computers = "the cart driving the horse"...

but finding the question unlocked by new tools is a better bet than developing tech to fit the question



Foster a Safe and Supportive Culture

“I’d rather become right than appear right and continue being wrong” -me

Finding, fixing, and avoiding problems is supercharged by:

- Asking questions when you don’t understand
- Not being defensive about being proven wrong

If everyone is comfortable making suggestions, the team will have better ideas



Be a role model by messing up!

E3SM is Multi-Lab... SCREAM is Multi-*Disciplinary*

- Combining atm physicists, computer scientists, and applied mathematicians has been important for:
 - Designing accurate parameterizations
 - Quickly fixing bugs
 - Choosing appropriate optimizations
- Melding disciplines is hard. We:
 - Have lots of meetings
 - Tried to learn each other's tasks
 - Communicate with patience



Keep Speculative Projects Speculative

- Knowing that SCREAM could fail without ruining E3SM allowed us to take chances and take the time to do things right (mostly)
- Isolation from the rest of E3SM helped us keep focused (but is making reintegration harder)

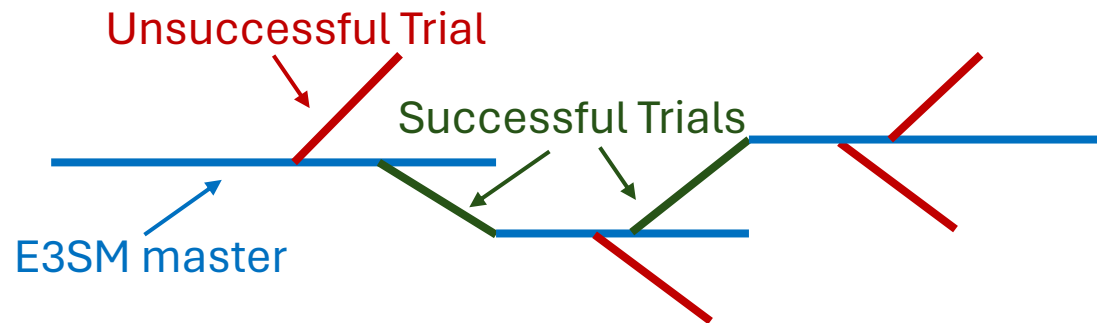


Fig: Trying things and absorbing them into mainline E3SM if/when they're ready maintains high-quality model releases while allowing for exploration

Make Sure Everyone Understands and Buys into the Mission

“If you don’t love what you’re doing, do something else or you won’t be happy and won’t do a good job”
-paraphrased from Bob Rosner (past ANL director)

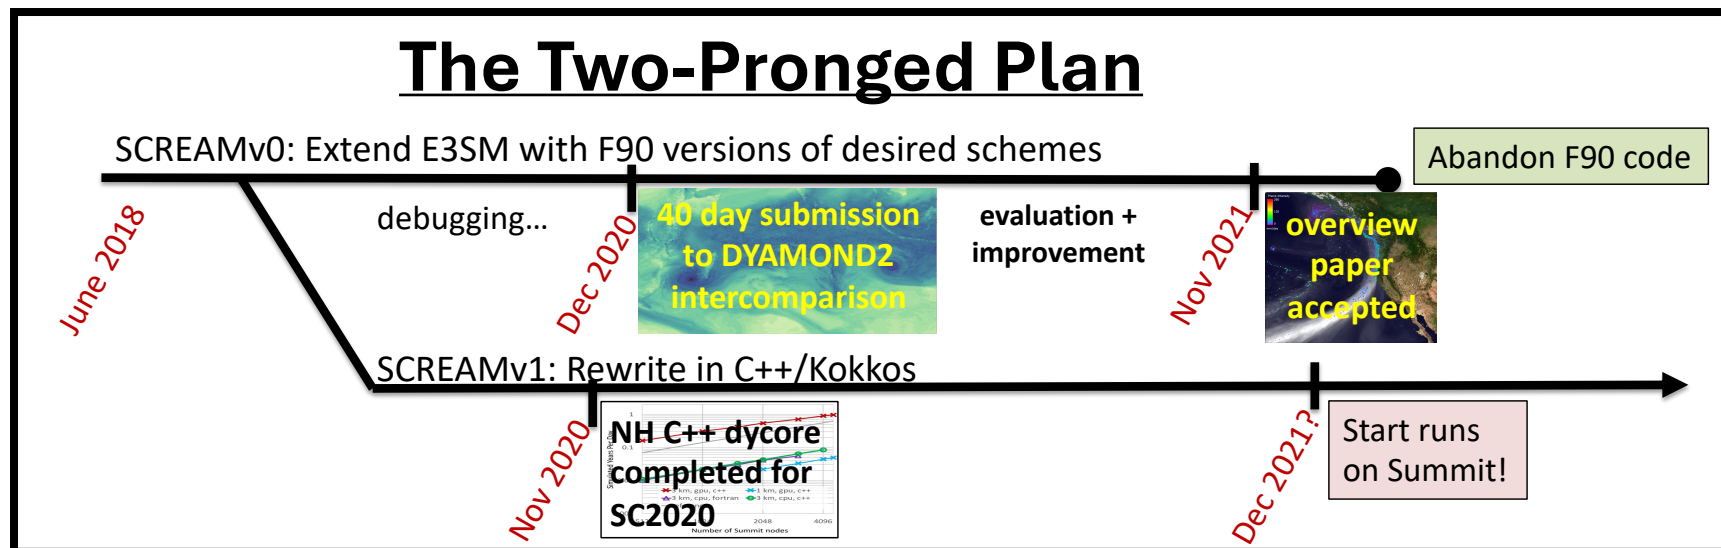
- SCREAM’s mission was particularly easy to articulate: “Make a new global cloud-resolving model that runs on GPUs”
- Frequently summarize key tasks and how they fit into the mission keeps us synced
- Prioritize tasks (and articulate priority)
 - Not everything can be urgent
 - Every year or so we have a sprint to finish something... but we explicitly *don’t* sprint other times
- SCREAM is built by volunteers (for better or worse)



Set Realistic Goals

“Use parameterizations which have already been proven. Developing parameterizations while porting will leave you baffled why your code doesn’t work” -Chris Bretherton (paraphrased)

- Not trying to do too much leads to a tractable plan
- Starting with an F90 version provided a template and gave everyone something to do
- Set “soft” deadlines well in advance of the drop-dead date




Coordinate, Don't Do




- In first aid classes, disaster ensues whenever the leader starts taking on technical tasks rather than maintaining a bird's eye view
 - Without trusting your team to do what's needed, you're sunk
- I shield my team from distractions by taking on bureaucratic tasks



And finally... have fun!



Vision for a SCREAMing future



- 100% correct code (through extensive testing and numerical analysis)
- Development focused on a km-scale physics-based model and ML emulation used to run thousands of years with this physics
- Tune using AI transfer learning from lower-res runs and short forecasts