Review of "GPU accelerated ice flow solver for unstructured meshes using the Shallow Shelf Approximation (FastIceFlo v1.0)

Dan Martin

July 2023

1 Overview

This article presents algorithm and implementation details to port the SSA solver for the ISSM ice sheet model to GPUs, focusing on the nonlinear momentum solve which takes up the vast majority of computational effort in most ice sheet models. The authors do a good job of explaining their mathematical formulation and present a nice set of real-world test cases to demonstrate performance.

One thing the authors should make clearer is the set of differences (progress) from the previous work by Räss, et al and this work. I think I have a sense of it, but it would be best if it were clarified in the introduction. ("Previous work (Räss 2020, Räss 2019, etc) did xxxx. The work presented here builds upon that by ..." or something like that.)

I do have a set of fairly minor issues which I think should be addressed prior to publication (enumerated below), but once those have been addressed, I support publication in GMD.

2 Specific Points

1. line 30-31: "The performance of CPUs is slowly leveling off" – You could also mention the power-consumption issues pushing us toward GPUs here.

2. line 39: As I mention below, I don't think you should use the word "inertial" here

3. Eqn 3: Do you add any regularization to address the singularity when the strain rate is 0?

4. line 80: You're not really including the full inertial terms here, just a partial time derivative. The "inertial terms" are the nonlinear terms in the material derivative $(u \cdot \nabla u)$. I think you're better off just saying that

you add a pseudo-time to allow iteration to a consistent steady state. (I realize this is what it was called in previous references; if I'd reviewed those, I'd have been equally pedantic for them :-))

5. line 89: Since your pseudo time-step is spatially variable, do you have a sense of how uneven the convergence is?

6. line 98: I'd suggest "the stable explicit CFL time step", or "the explicit CFL-stable time step", etc. to reinforce that stability is driving this choice.

7. Eqn 10: I don't think this is correct – assuming that H is spatially varying (i.e. a non-boring ice sheet), using the chain rule on the RHS of Eqn 9 results in an extra term: $4\mu\frac{\partial v_x}{\partial x}\frac{\partial H}{\partial x}$. Maybe it can be ignored, but I think there needs to be some justification of that.

8. Eqn 11: I'm confused by this, since it seems like you're including an extra $\Delta\tau_W A_x^{old}$ term in the update that would accumulate as it evolves. In a true oscillating wave that would likely get canceled out as it evolves, but this seems like it won't tend toward oscillatory wave solutions. Can you discuss that?

9. line 135 – I think it would be helpful to reference back to Eqn 4 here.

10. Figure 3: Is it possible to include (likely in an additional figure) some sort of plot of norm(residual) (i.e. du/dt) vs. iteration to illustrate how this method performs? Is it a linear convergence, or something better? (in the end, you're comparing against a more-standard iterative method where one would plot residual vs. iteration).

11. line 192: What do you mean by "arithmetic precision"? 32 vs 64? (what number are you using for n_p?)

12. line 200: Can you include a table with the optimal parameters here?

13. line 202: Can you describe what you mean by "optimal solver parameters are unidentifiable"? Are you completely unable to solve the problem? or is it simply that you can't identify optimal parameters (in which case you could still have a result)?

14. line 203: "in the following steps" – Perhaps this phrase is a relic of an edit?

15. line 205: (side note) I've also found that the presence of dynamically important ice shelves (like that in the PIG case) can drastically affect performance in the solution of the momentum balance.

16. line 213: Don't we really care about the integrated power needed to solve this system? (i.e. Watt-hours or Watt-seconds, etc? vs the power, which is a rate...)

17. line 253: "confirms" – I don't think it confirms so much as suggests...

18. line 255: Maybe say "ice sheet evolution" instead of "grounding-line migration"?

19. line 280: I'd probably say "not practical" or "not possible without extreme computational resources" vs. "not possible" (after all, we now have exascale computers (well, one exascale computer, at least)).

20. References: The links to the DOIs are messed up here, with repetitions of "https://doi.org"

3 Typos and grammar fixes

1. line 50: "West" should be capitalized in West Antarctica.

2. line 50: There is an extra space after "read"

3. line 71: "require to"

4. line 82: "allowing us to..."

5. line 103: "process to diverge" → "process diverging"

6. line 113: "allows us to aggressively..."

7. line 114: "method scale to" → "method scale as"

8. line 123: "limiter" → "limit"

9. Eqns 14, 17, 18 are badly formatted

10. line 165: "West Antarctica"

11. line 175: "As a CPU" → "For the CPU comparison" or something like that

12. line 177: missing space after "enabled"

13. line 187: "of PT CUDA: → "of the PT CUDA"

14. line 194: should be a colon after 8, not a semicolon, I think.

15. line 209: "the NVIDIA"

16. line 210:"The Power"

17. line 211: Should it be "sheet: thermal..."? (colon vs. semicolon)

18. line 266: "the 2-D SSA..."