MPAS-O: 2012 and Beyond

(Where we are going and how we are going to get there.)

Todd Ringler

Theoretical Division
Los Alamos National Laboratory

Climate, Ocean and Sea-Ice Modeling Project
http://public.lanl.gov/ringler/ringler.html

Los Alamos NATIONAL LABORATORY EST. 1943

U.S. DEPARTMENT OF ENERGY | Office of Science

We have a plan ....

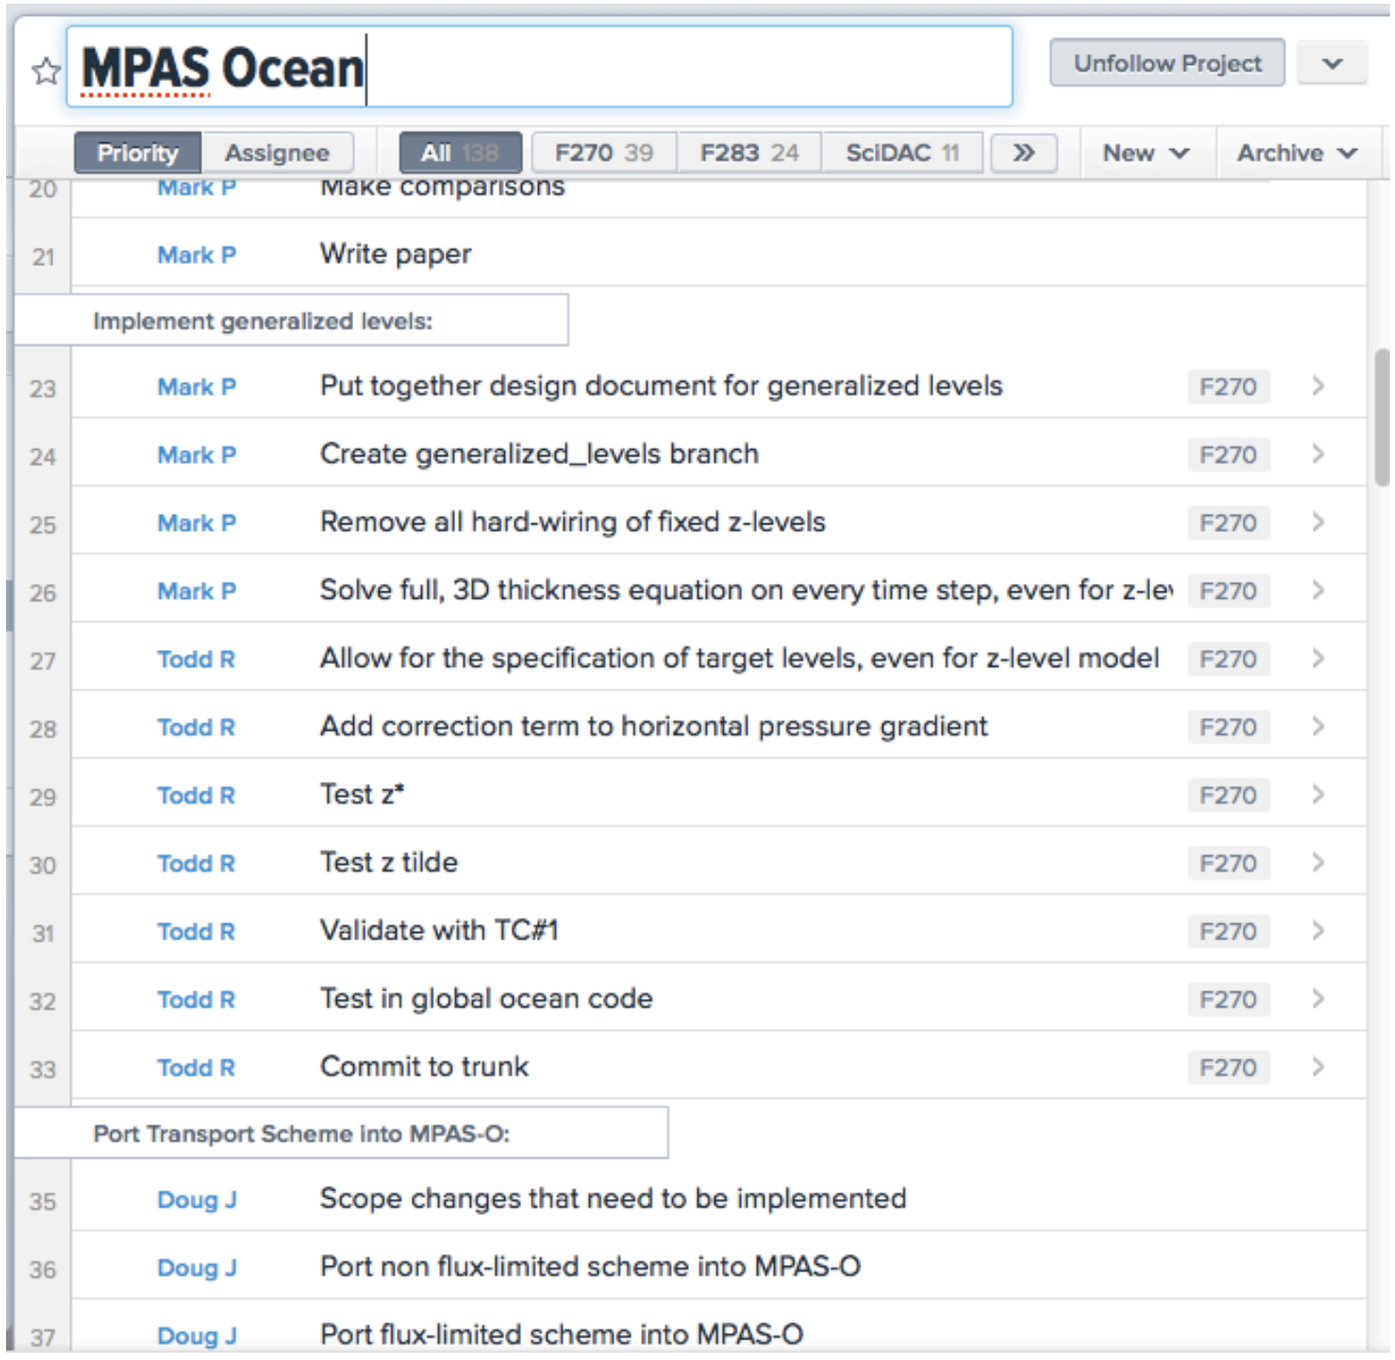

MPAS Ocean

Unfollow Project

Priority | Assignee | All 138 | F270 39 | F283 24 | SciDAC 11 | » | New | Archive

20	Mark P	Make comparisons	
21	Mark P	Write paper	
	Implement generalized levels:		
23	Mark P	Put together design document for generalized levels	F270
24	Mark P	Create generalized_levels branch	F270
25	Mark P	Remove all hard-wiring of fixed z-levels	F270
26	Mark P	Solve full, 3D thickness equation on every time step, even for z-le	F270
27	Todd R	Allow for the specification of target levels, even for z-level model	F270
28	Todd R	Add correction term to horizontal pressure gradient	F270
29	Todd R	Test z*	F270
30	Todd R	Test z tilde	F270
31	Todd R	Validate with TC#1	F270
32	Todd R	Test in global ocean code	F270
33	Todd R	Commit to trunk	F270
	Port Transport Scheme into MPAS-O:		
35	Doug J	Scope changes that need to be implemented	
36	Doug J	Port non flux-limited scheme into MPAS-O	
37	Doug J	Port flux-limited scheme into MPAS-O	

We have currently tracking ~20 development projects, broken into ~110 tasks, to be completed over the next two years.

Each task is linked to a person and includes an estimate of effort required to completed.

The plan was made in the context of available resources ....

At LANL, we are planning to support MPAS-O at a level of approximately 6.0 FTE.

Most of these resources are leveraged from individual projects with specific deliverables, so we are constrained in the extent to which we can reallocate staff work items.

Team:

Ringler (0.5 FTE): management, developer
Petersen (1.0 FTE): primary core developer
Jones (0.0 FTE): COSIM lead
Maltrud (0.0 FTE): applications lead
Lowrie* (0.25 FTE): transport algorithms
Newman (0.25 FTE): implicit time stepping
Chen (1.0 FTE): mesoscale eddy parameterizations
Jacobsen (1.0 FTE): optimization, infrastructure, test-cases
Pratola (1.0 FTE): DART-MPAS coupling
Graham (1.0 FTE): turbulence closures

The plan was also made in view of our assessment of risk:

1. Computational performance: We need to increase our throughput per degree of freedom to be competitive with POP.

2. Simulation quality (Ocean): We need to assess the fidelity of MPAS-O simulations, on a per degree of freedom basis, as compared to POP and other community-accepted global ocean models.

3. Community: Moving this from a LANL-centric activity to a community activity.

Our major goals for 2012:

Our major goals for 2012:

Performance: Implement best-practices profiling tools
Performance: Implement hybrid parallelism (OpenMP/MPI)
Performance: Maximize cache reuse via reordering
Performance: Implement performance tracking infrastructure

Our major goals for 2012:

Performance: Implement best-practices profiling tools
Performance: Implement hybrid parallelism (OpenMP/MPI)
Performance: Maximize cache reuse via reordering
Performance: Implement performance tracking infrastructure

Quality: Codify test-case suite
Quality: Full evaluation of generalized-level vertical coordinate
Quality: Implement "best-in-class" horizontal transport scheme
Quality: Implement mesoscale eddy parameterization
Quality: Couple MPAS-O into DART

Our major goals for 2012:

Performance: Implement best-practices profiling tools
Performance: Implement hybrid parallelism (OpenMP/MPI)
Performance: Maximize cache reuse via reordering
Performance: Implement performance tracking infrastructure

Quality: Codify test-case suite
Quality: Full evaluation of generalized-level vertical coordinate
Quality: Implement "best-in-class" horizontal transport scheme
Quality: Implement mesoscale eddy parameterization
Quality: Couple MPAS-O into DART

Community: Publish the first MPAS-O papers
Community: Start to build external user community
Community: MPAS-O with CORE forcing using CESM drivers(?)

Our major goals for 2013:

Performance: Leverage SUPER autotuning capability
Performance: Port flop-intensive kernels to accelerators

Quality: A full implementation of JFNK time stepping.
Quality: Initial implementation of variable-kappa eddy closure
Quality: Initial implementation of advanced momentum closure

Community: Publish results from MPAS-O test case suite
Community: SourceForge release (along side MPAS-A)
Community: Fully coupled CESM simulations using MPAS-O

Things that we are excited about #1: Generalized Levels

We have incorporated much of our HyPOP infrastructure into MPAS-O.
Specifically, we keep a prognostic equation for thickness at every grid cell.

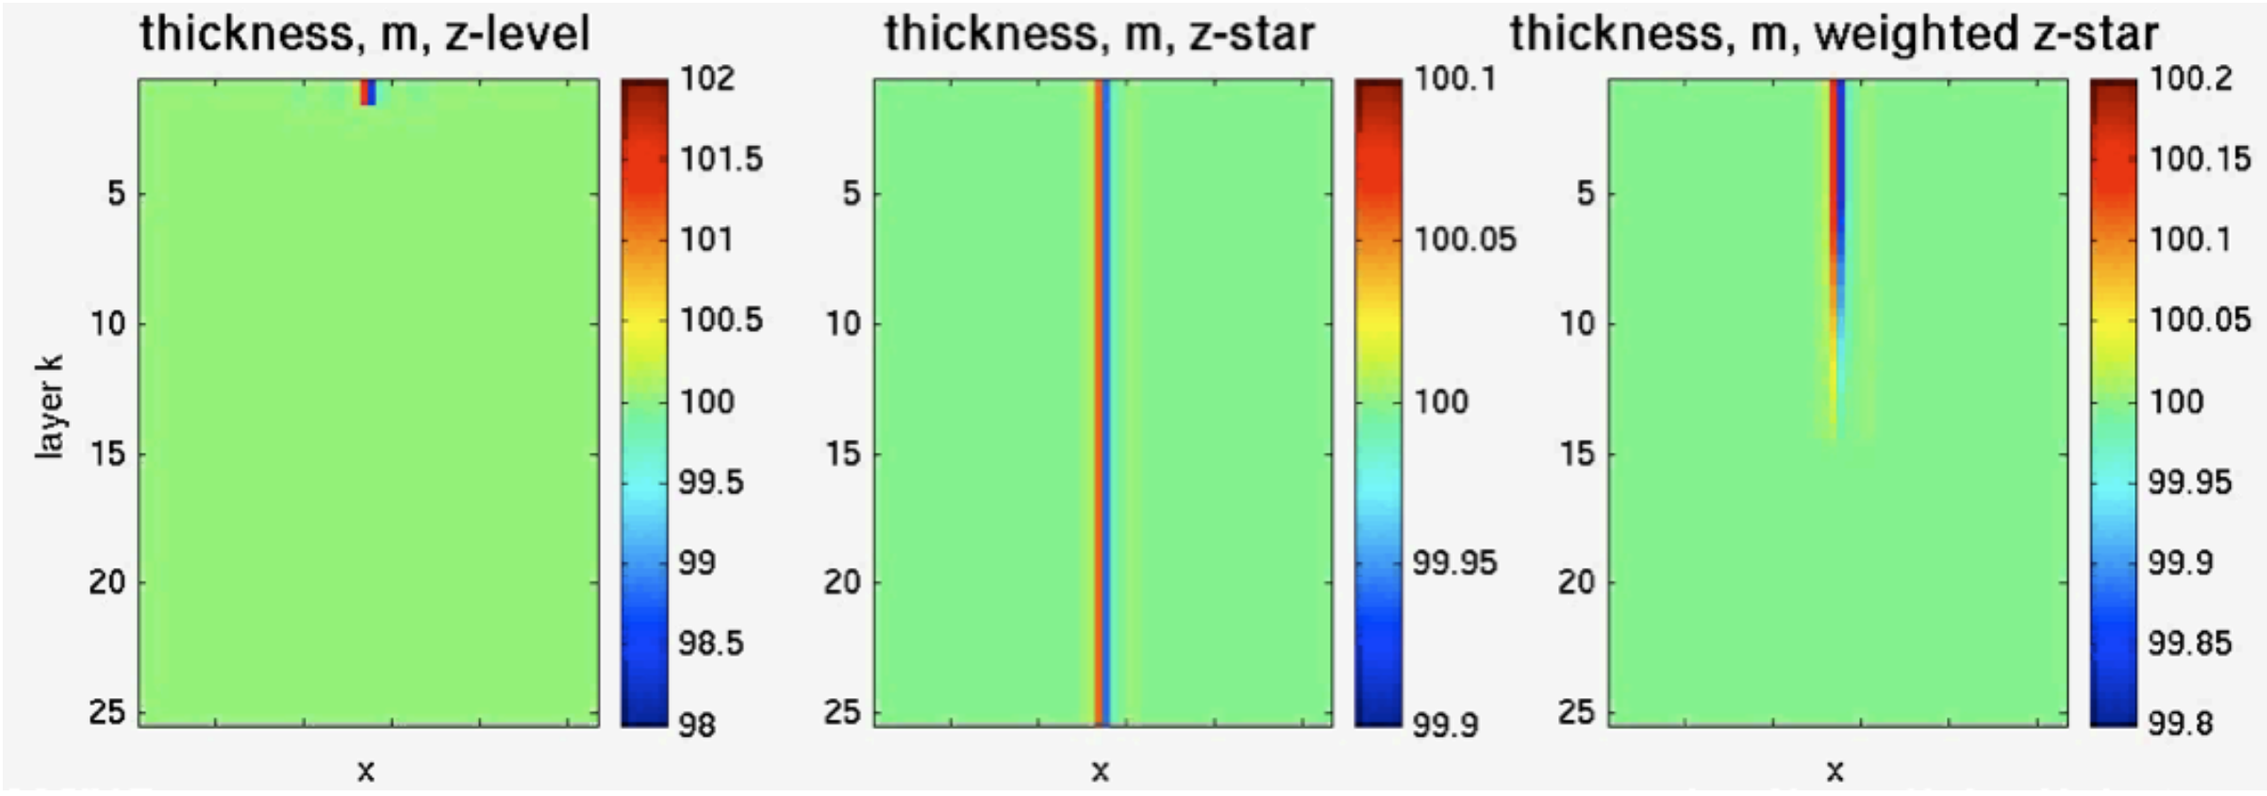

In addition to z-star and z-tilde, we have the opportunity to build more flexible algorithms for the evolution of layer thickness within the construct of a generalized levels approach.

Things that we are excited about #2: Transport

Characteristic Discontinuous Galerkin Transport Scheme

(Prather's method extended to arbitrary order for arbitrary convex polygons without dimensional splitting.)

$$\int_{\Omega_k} \left[\beta_{k,i}(\rho T)^{n+1} - (\phi_{k,i}\rho T)^n\right] \, d\Omega + \sum_{f\in\mathcal{F}(k)} \int_{\Delta\Omega'_f} \left(s_{k,f}\phi_{k,i}\rho T\right)^n \, d\Omega = 0$$

Basis functions that allow "on cell" reconstruction of tracer, T.

Basis function evaluated along characteristics.

Lagrangian pre-image of each face, f

Flux of mass-weighted tracer concentration across each face.

Graphical representation of CDG transport for arbitrary convex polygons

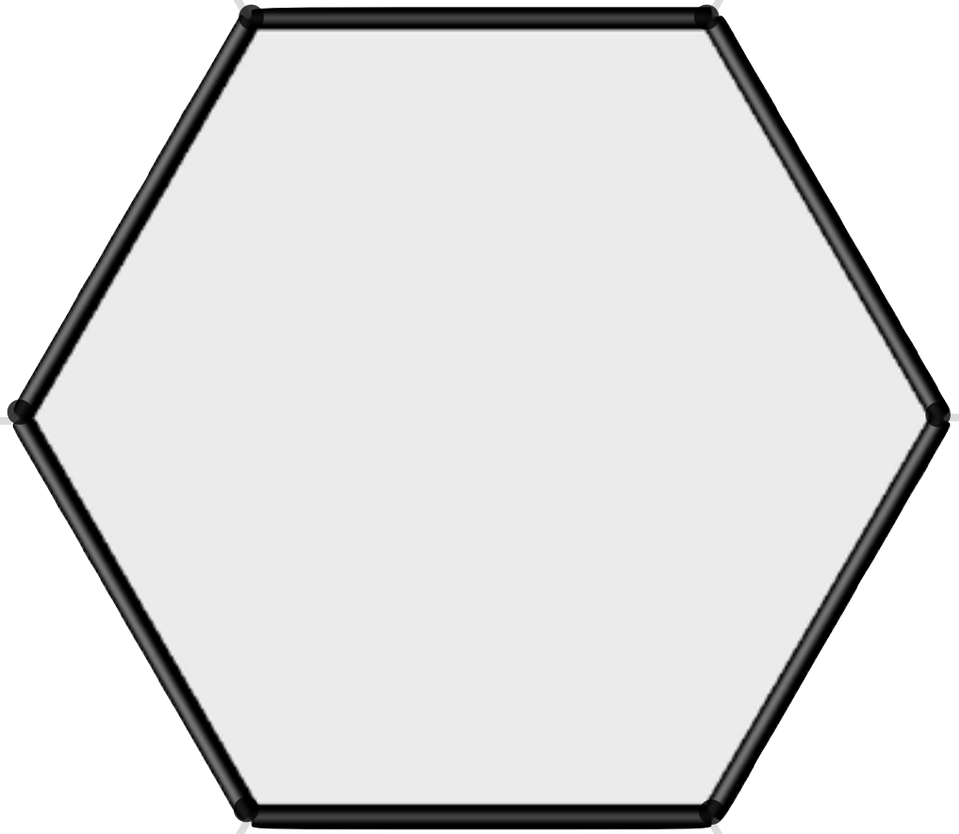

Graphical representation of CDG transport for arbitrary convex polygons

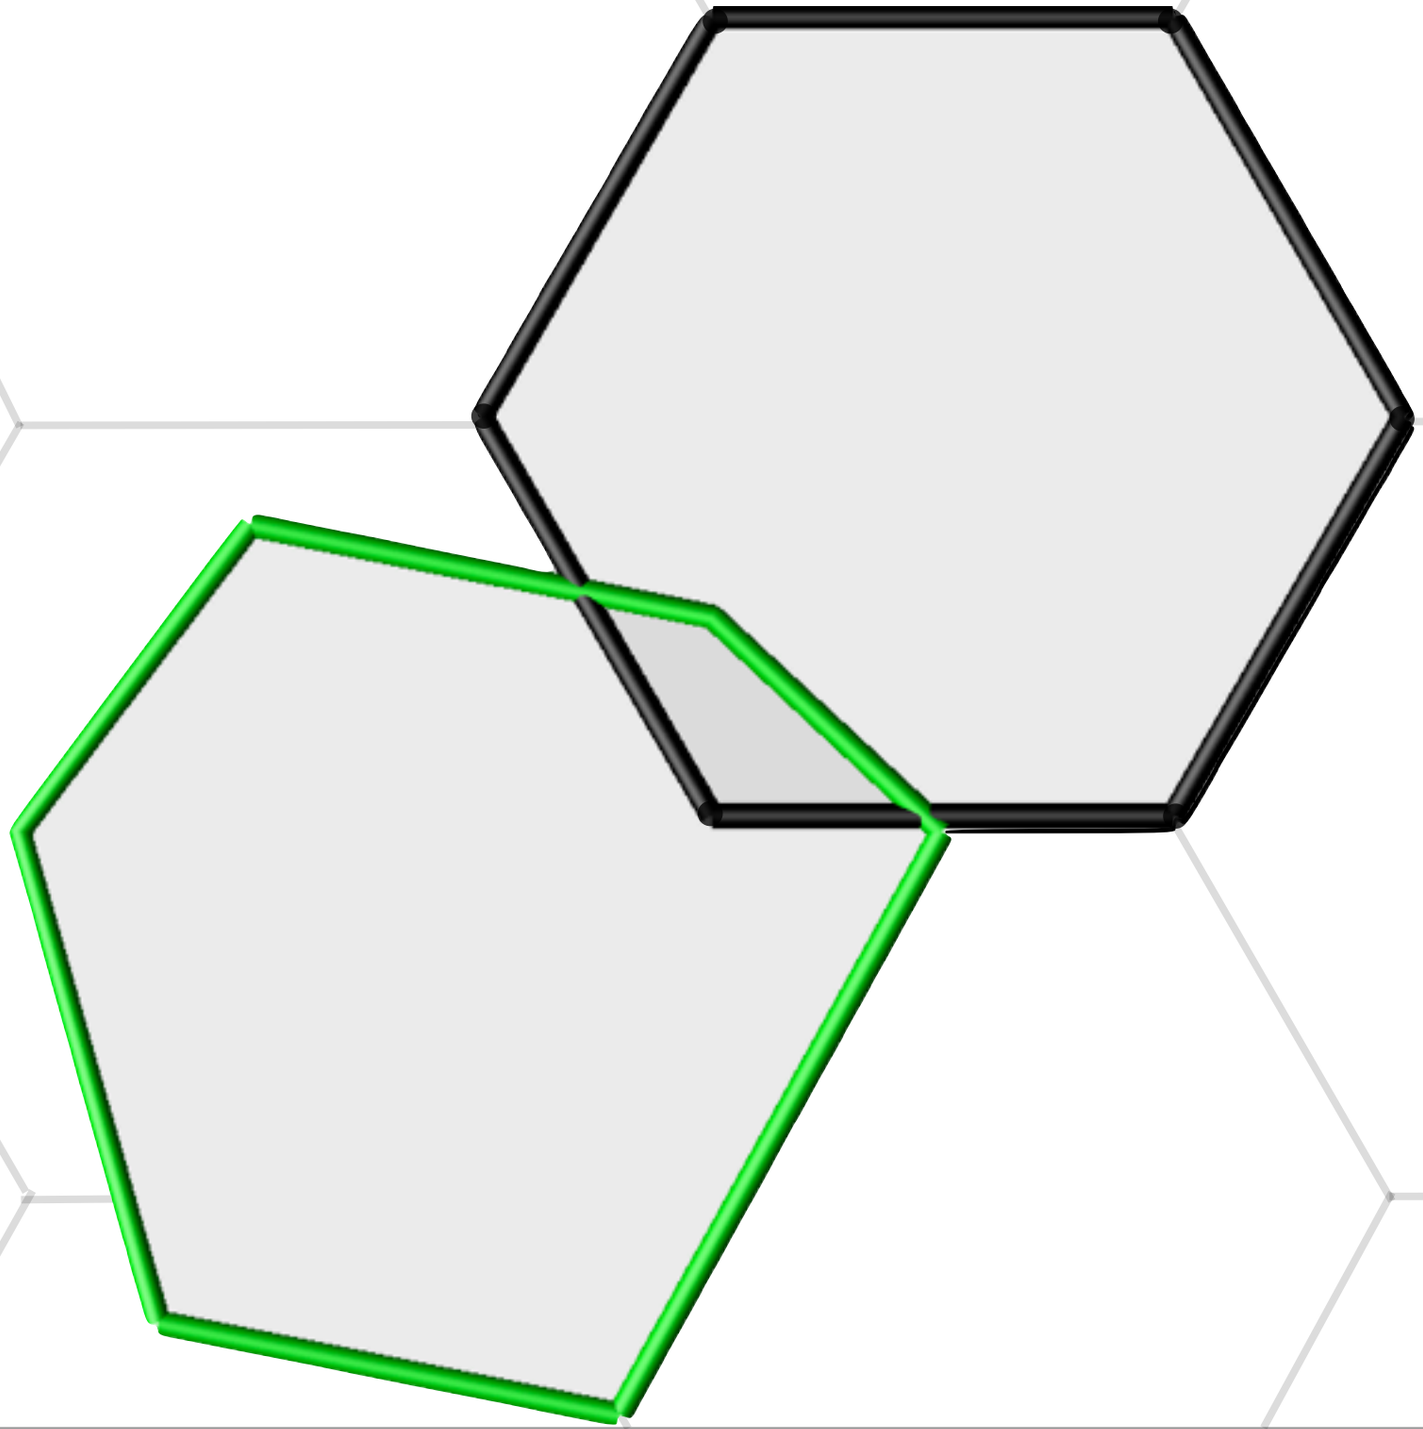

Graphical representation of CDG transport for arbitrary convex polygons

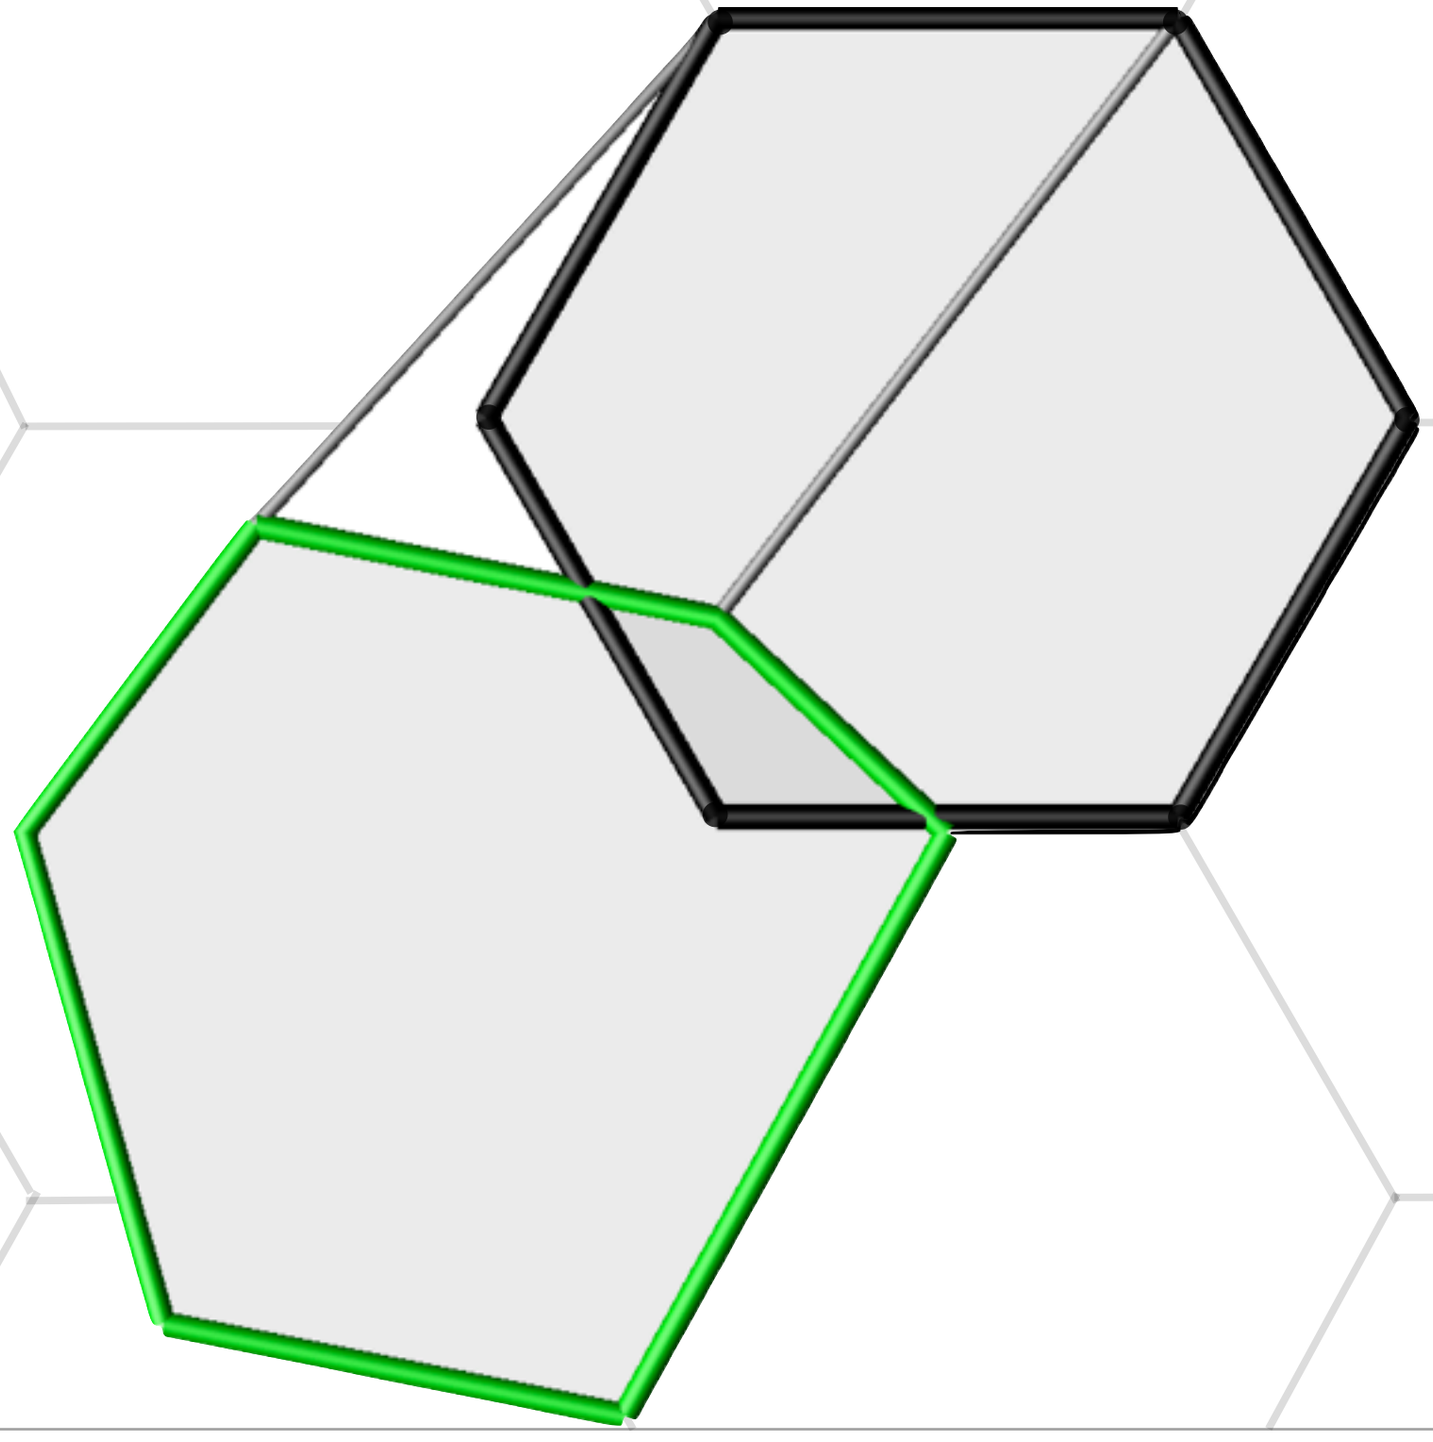

Graphical representation of CDG transport for arbitrary convex polygons

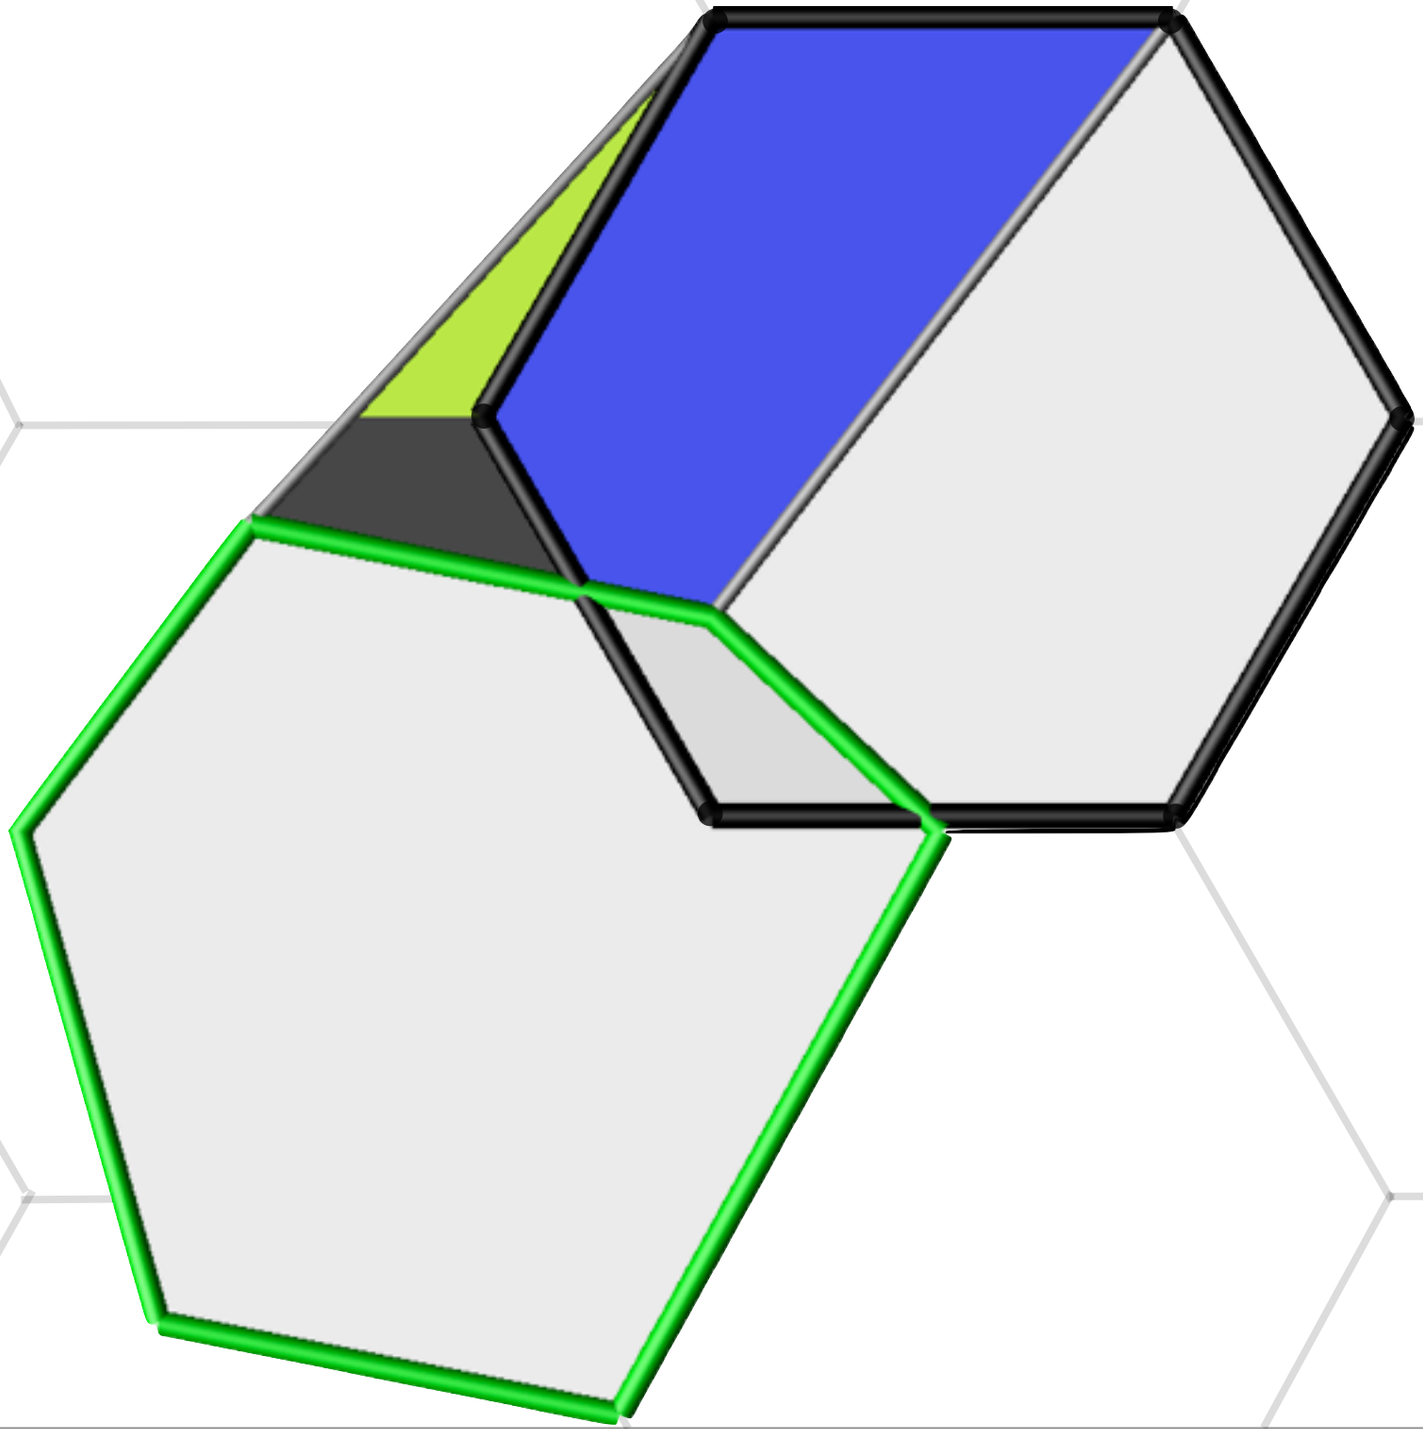

Graphical representation of CDG transport for arbitrary convex polygons

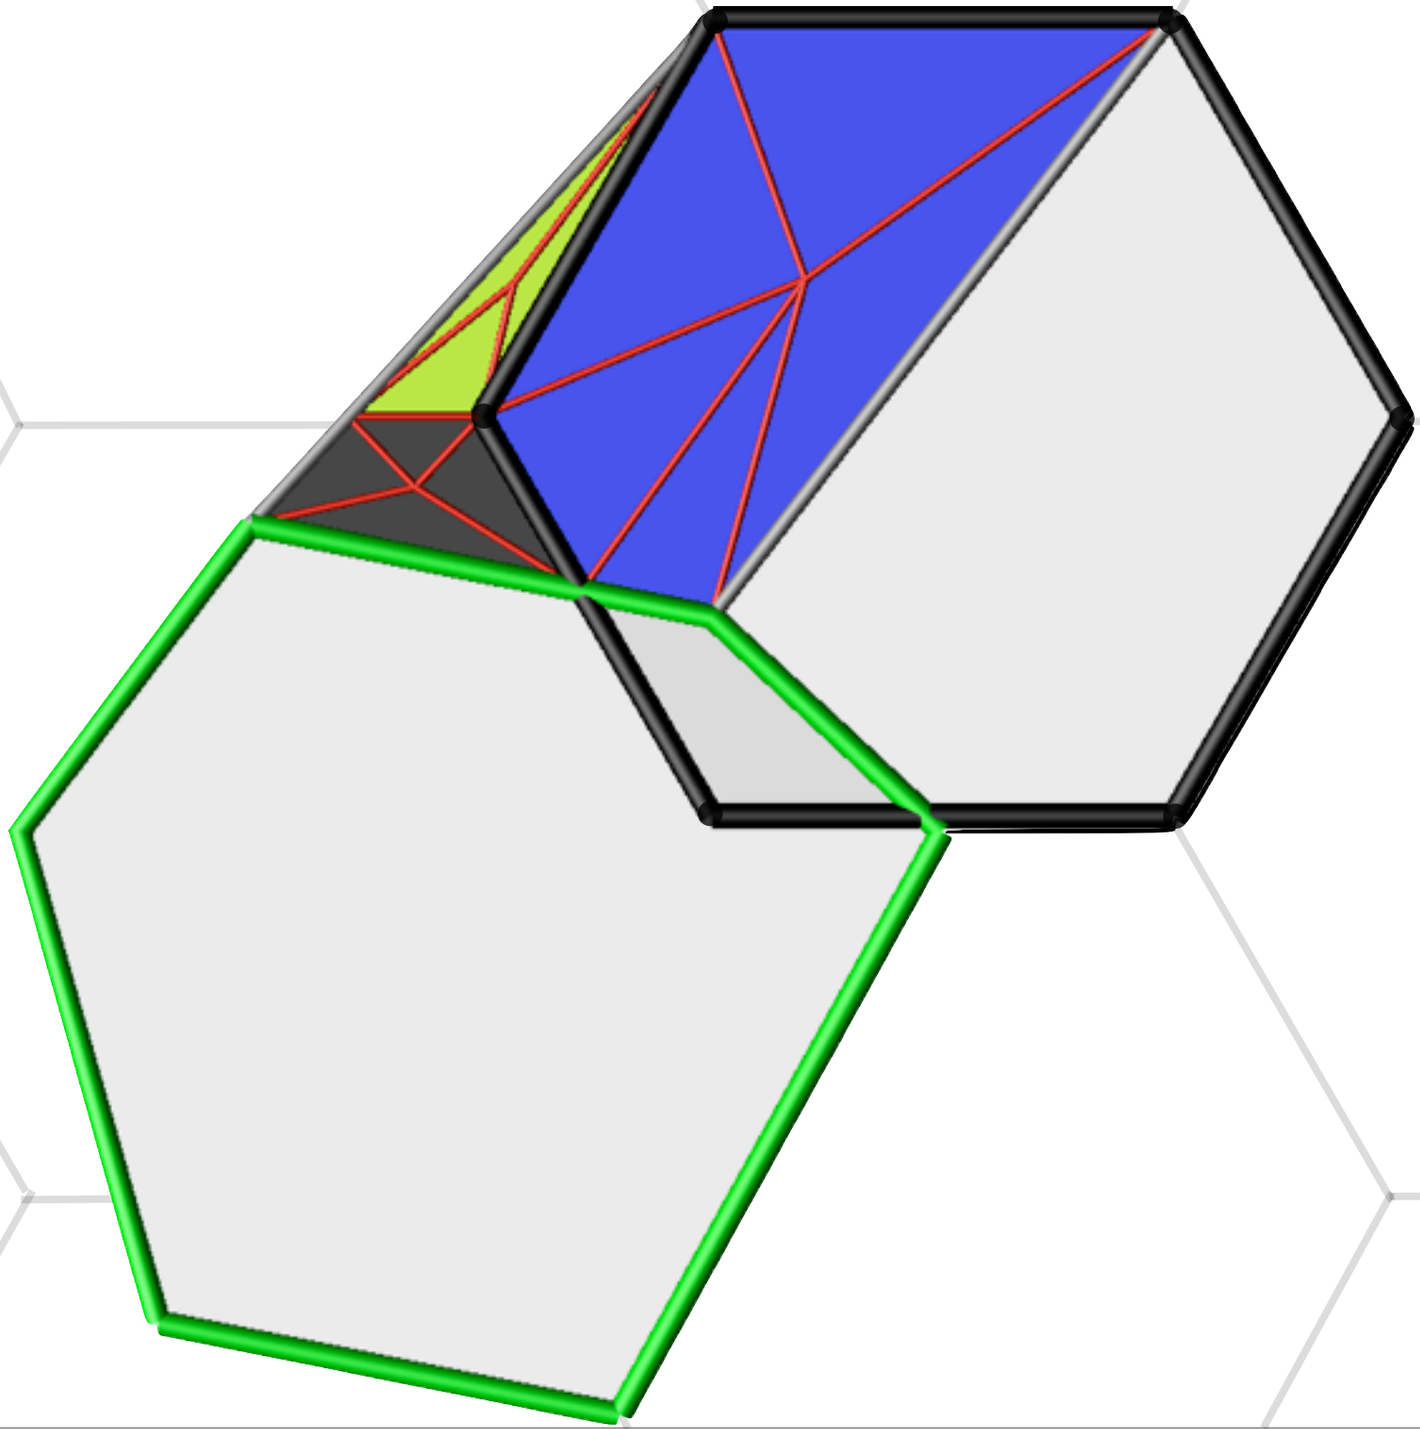

Graphical representation of CDG transport for arbitrary convex polygons

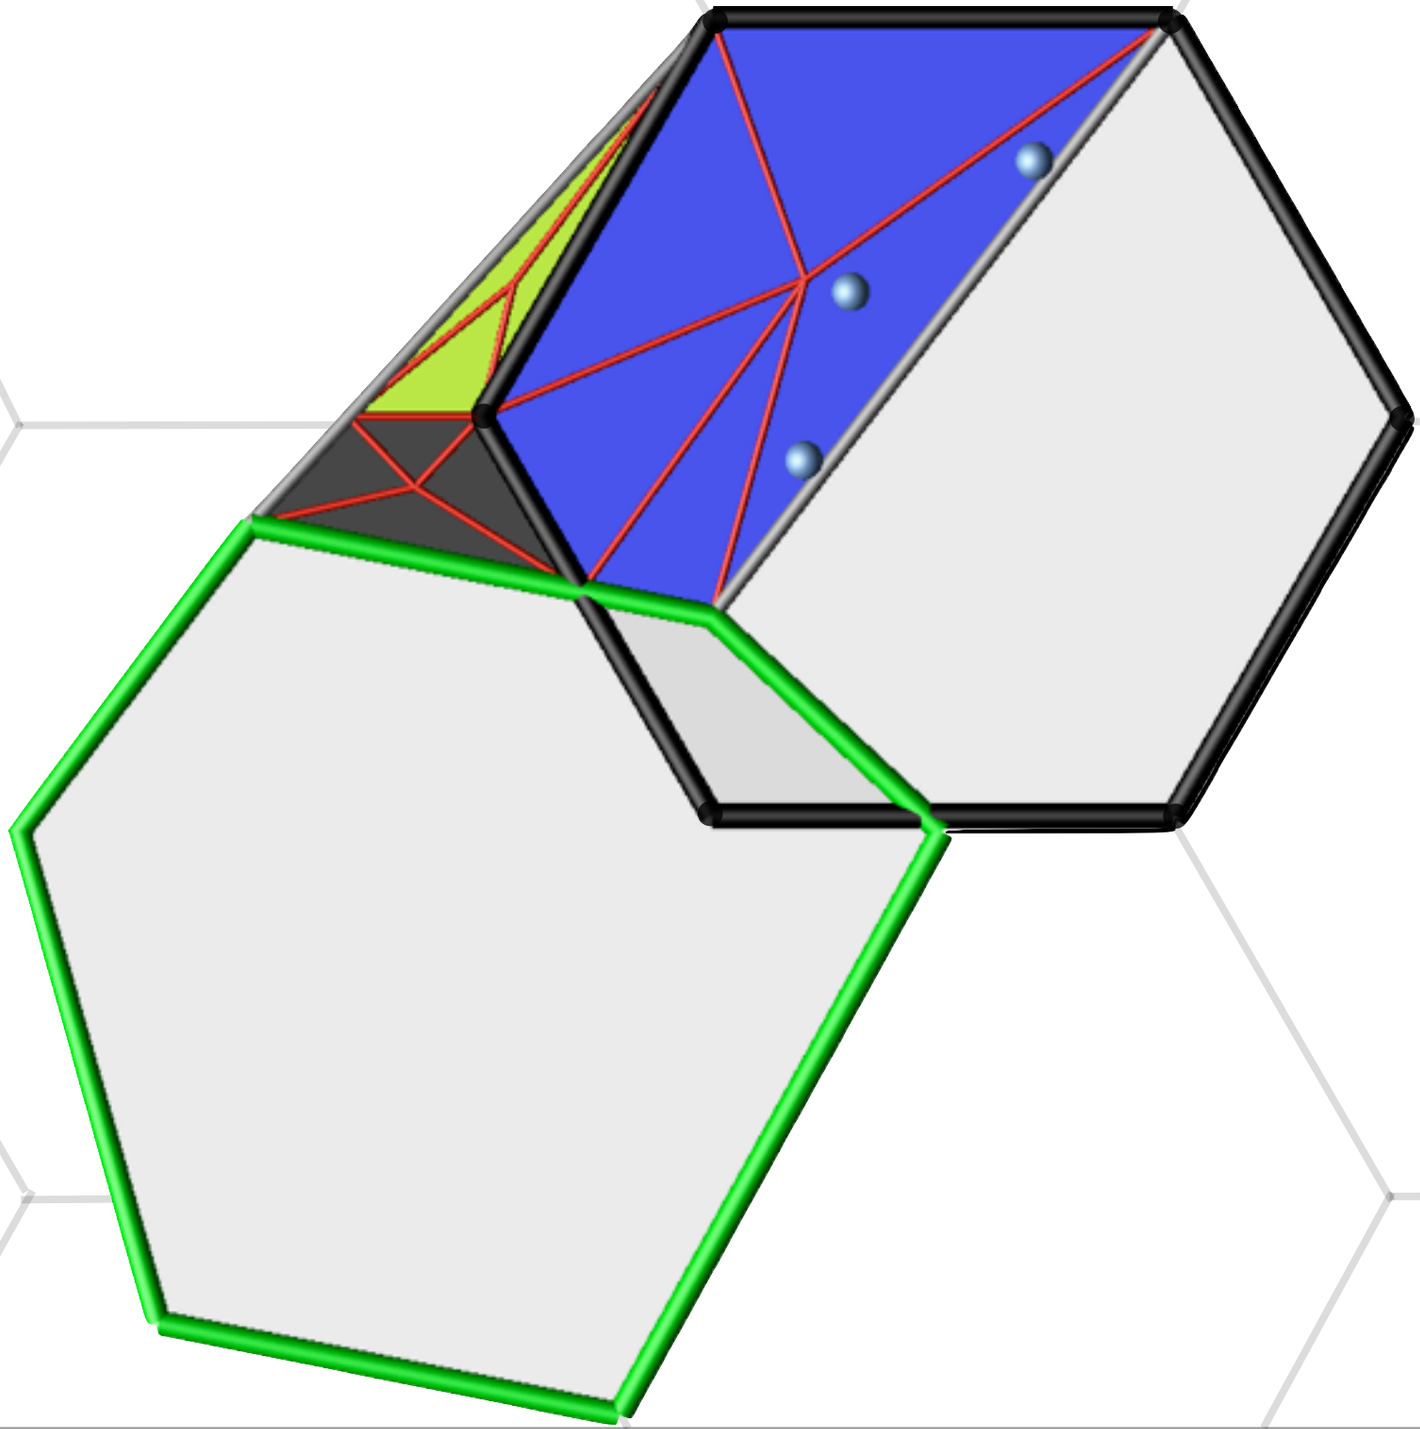

Building a Community ....

Building a Community ....

From the start, our intention has been to distribute MPAS components (atmosphere, ocean, etc.) via 3rd-party, open-source channels (SourceForge). Ultimately, we expect this to be the hub for community engagement.

Building a Community ....

From the start, our intention has been to distribute MPAS components (atmosphere, ocean, etc.) via 3rd-party, open-source channels (SourceForge). Ultimately, we expect this to be the hub for community engagement.

Note: CESM-specific interfaces will be maintained on SourceForge, so coupling MPAS-O into the CESM is a trivial procedure.

Building a Community ....

From the start, our intention has been to distribute MPAS components (atmosphere, ocean, etc.) via 3rd-party, open-source channels (SourceForge). Ultimately, we expect this to be the hub for community engagement.

Note: CESM-specific interfaces will be maintained on SourceForge, so coupling MPAS-O into the CESM is a trivial procedure.

Yet, in the near term (~18 months) we need to engage a relatively small group of scientists to begin using MPAS-O. We expect to build up this small group one person at a time.

Ideas for engaging scientists in this process ....

Ideas for engaging scientists in this process ....

The CESM will now accept an ocean model on an unstructured grid. We could use the milestone of a CESM-driven, MPAS-O simulation with CORE forcing as a way to frame engagement.

Ideas for engaging scientists in this process ....

The CESM will now accept an ocean model on an unstructured grid. We could use the milestone of a CESM-driven, MPAS-O simulation with CORE forcing as a way to frame engagement.

We have the idea to refactor the ocean driver to increase modularity of the ocean component (dynamics, physics, BGC). We could push this idea from a document to software while enlarging the MPAS-O developer team.

Ideas for engaging scientists in this process ....

The CESM will now accept an ocean model on an unstructured grid. We could use the milestone of a CESM-driven, MPAS-O simulation with CORE forcing as a way to frame engagement.

We have the idea to refactor the ocean driver to increase modularity of the ocean component (dynamics, physics, BGC). We could push this idea from a document to software while enlarging the MPAS-O developer team.

We could engage scientists on the topic of local-mesh refinement in the context of idealized configurations.

Thanks!