LA-UR-19-23274
Approved for public release; distribution is unlimited.

Title: Calculating Radiation View Factors Using genre: A Case Study

Author(s): Neill, David Howard

Intended for: Report

Issued: 2019-04-11

Disclaimer:
Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Calculating Radiation View Factors Using *genre*:

A Case Study

David Neill-Asanza, CCS-2
Los Alamos National Laboratory

April 9, 2019

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

What is *genre*?

- Computes radiative heat transfer view factors between faces.
  - *View factor:* proportion of field of view covered by a face.
    - Used in radiosity equation to calculate radiation incident on a face.
  - Chaparral library[2] from Sandia used to compute view factors
- *genre* has four main stages:
  1. Read mesh file (.gen) and generate enclosure surface mesh
  2. Calculate view factors (hemicube algorithm)
  3. View factor matrix smoothing
  4. Write radiation enclosure file (.re)
- The view factor computation is expensive.
- **How can we get the best bang for our buck?**

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Hemicube Algorithm

- Approximates exact solution of view factor integral
- Calculates fraction of a face's field of view covered by another face
- Discretizes unit hemisphere as a unit hemicube.
- To reduce bias, hemicube is rotated by random angle about face normal

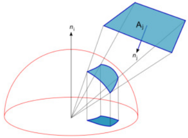

Calculating view factors by projection onto unit hemisphere.[3]

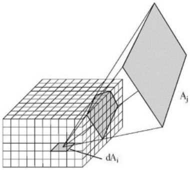

Discretization into a hemicube.[4]

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

View Factor Matrix

- ϕ_{ij} is the **view factor** of face *j* relative to face *i*.
  - Fraction of power radiated from face *i* incident on face *j*.
- $\phi_{i\infty}$ is the view factor between face *i* and the ambient at infinity (when there is a hole in the enclosure).

Properties:

- $0 \leq \phi_{ij} < 1$, 0 if no line-of-sight from face *i* to *j*
- $0 \leq \phi_{i\infty} < 1$, 0 if face *i* doesn't see the ambient
- Reciprocity: $A_i \phi_{ij} = A_j \phi_{ji}$
  - Structurally symmetric matrix: $\phi_{ij} = 0 \Leftrightarrow \phi_{ji} = 0$
  - If face *i* sees face *j*, then the converse is also true
- Unit row sum: $\Sigma_j \phi_{ij} + \phi_{i\infty} = 1$
  - Radiation leaving face *i* is conserved

View Factor Matrix Smoothing

- Reciprocity enforces *structural symmetry*:

$$A_i \phi_{ij} = A_j \phi_{ji} \quad \text{implies} \quad \phi_{ij} = 0 \Leftrightarrow \phi_{ji} = 0$$

- Energy must be conserved:

$$\Sigma_j \phi_{ij} + \phi_{i\infty} = 1$$

- Why is smoothing necessary?
  - Guarantees VF matrix properties.
  - Hemicube resolution is finite
    - Face *A* sees face *B*, but face *B* does not see *A*.
- Smoothing stage runs after the hemicube algorithm. Two steps:
  1. Ensure structural symmetry (free)
  2. CG method used to enforce unit row sums (most expensive)

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

View Factor Matrix Smoothing

- Hemicube resolution is finite
  - Face *A* sees face *B*, but face *B* does not see *A*.

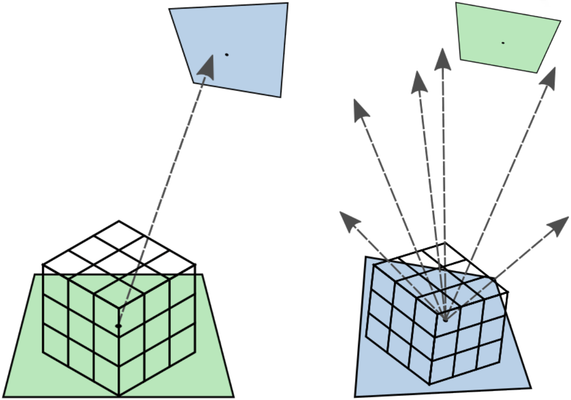

UNCLASSIFIED

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Minimum Separation

- The hemicube algorithm assumes that the distance between faces is *much greater* than the diameter of the faces.[2]
  - Faces may be subdivided to produce sub-faces satisfying the proximity assumption
- How much is ‘much greater’?
  - The **min_separation** parameter defines the minimum ratio of distance to diameter between any two faces.
  - A face f_i is subdivided so that all sub-faces satisfy the condition

$$min\_separation \leq \frac{d_{min}}{subface\_diameter}$$

where d_{min} is the minimum distance between f_i and all other faces.

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Minimum Separation: Face Subdivision

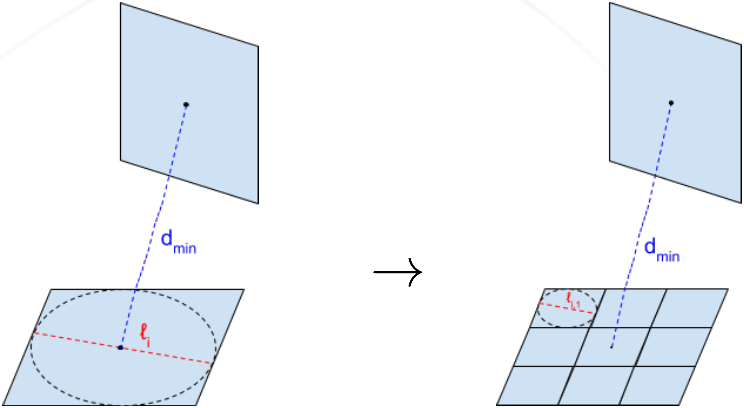

Check that face f_i satisfies
$min\_sep > d_{min}/\ell_i$

Subdivide until
$min\_sep > d_{min}/\ell_{i,k}$ for all k

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

What parameters are important?

Three main factors affect runtime and quality of solution.

1. **hemicube_resolution:** number of subdivisions in one dimension.
   - Setting hc_res=n divides each hemicube face into n^2 regions.
2. **min_separation:** minimum ratio of distance to diameter between any two faces.
3. **max_subdivision:** maximum face subdivisions in one dimension
   - Setting max_subdivision=n allows up to n^2 sub-faces per face.
   - This limit will not be exceeded, regardless of **min_separation**.
   - *genre* prints the maximum number of subdivisions needed to satisfy the given **min_separation**
     - If this limit isn't reached, result may be low quality.

Los Alamos NATIONAL LABORATORY EST. 1943

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Other Parameters (won't be discussed)

- blocking_enclosure
- partial_enclosure
- partial_area
- BSP_max_tree_depth
- BSP_min_leaf_length
- spatial_tolerance
- smoothing_tolerance
- smoothing_max_iter
- smoothing_weight

For more information see the Chaparral User Manual.[2]

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Case Study: Basic-Hemi

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Basic-Hemi: Pure Uranium Cast Simulation

- Test problem for validating Truchas against experiment.[1]
- Pure uranium cast simulation
  - Hemispheric shell graphite mold
    - 170mm outer diameter
    - 10mm thick shell
  - Mold preheat
  - Mold fill (gravity pour)
  - Cooldown and solidification
- View factors used for mold preheat and cooldown stages
  - Re-calculating view factors during pour is prohibitively expensive

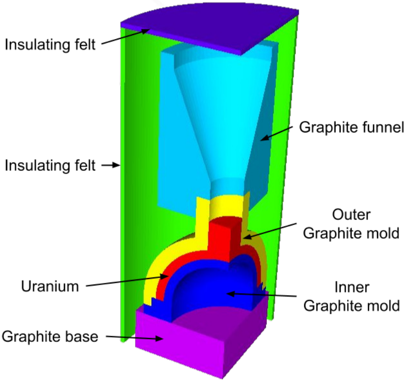

Basic-Hemi Geometry.[1]

UNCLASSIFIED

Basic-Hemi: Preheat Mesh and Enclosures

- Cooldown mesh
  - Unstructured grid
  - Variable resolution

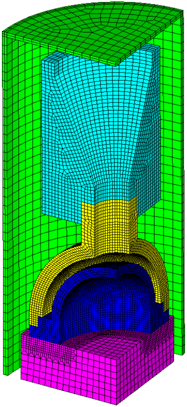

- View factor enclosures
  - Blue surfaces: inner enclosure
  - Green surfaces: outer enclosure

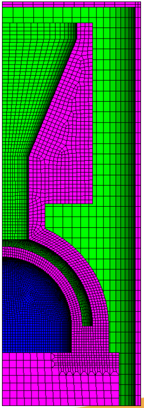

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Methodology

- All tests were run on the Snow cluster.
  - Intel® Xeon® CPUs with 36 cores @ 2.10 GHz
  - 125 GB of main memory
  - Unless otherwise stated, each node was fully subscribed
- Due to QOS limits, total runtime did not exceed 12 hours
- Data collected on four meshes of increasing resolution

Label	Num Faces	VF Matrix Density	RE file size (Mb)	Cell side lengths (mm) median	min	max
OUTER1	9007	∼15.35%	95.49	2.94	1.3	11.
OUTER2	20195	∼15.00%	467.90	1.95	0.70	7.7
OUTER3	36025	∼14.83%	1,470.05	1.47	0.61	5.9
OUTER4	80777	∼14.65%	7,295.27	0.98	0.32	4.0

Table: The four meshes used for the case study

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Los Alamos NATIONAL LABORATORY EST. 1943

Results

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Total Run Time: OUTER3 (single node)

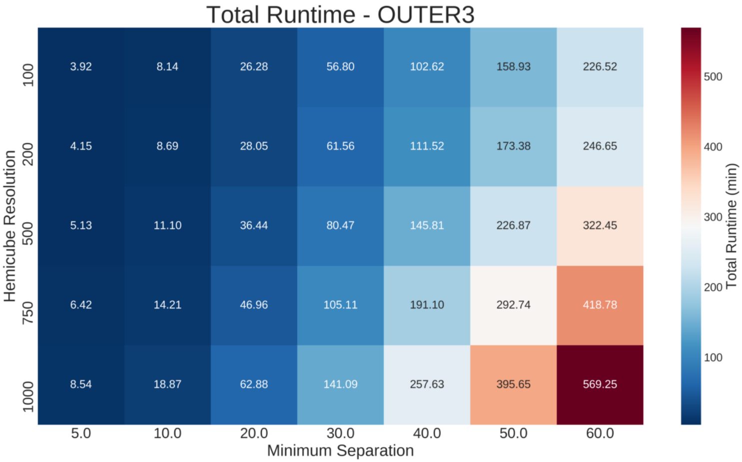

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Total Run Time: OUTER3 (single node)

- Quadratic dependence on hc_res and min_sep
  - Only fitted models of the form:

    $$ax^n + b \quad \text{where} \quad a, b \in \mathbb{R} \text{ and } n \in \mathbb{N}$$

  - Polynomial models tend to over-fit

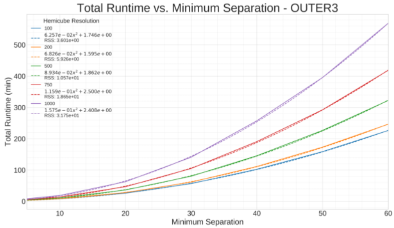

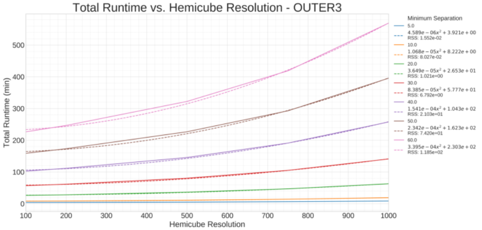

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Run Time Scaling: Stage Breakdown

- All runs on OUTER3, single node, hc_res=500, min_sep=30.0
- Chaparral library accounts for 99% of total *genre* runtime
  - 70% - 95% spent on hemicube algorithm
  - 5% - 30% spent on smoothing
- Poor scaling efficiency.
  - MPI overhead limits smoothing performance.
  - Load balancing is one possible issue

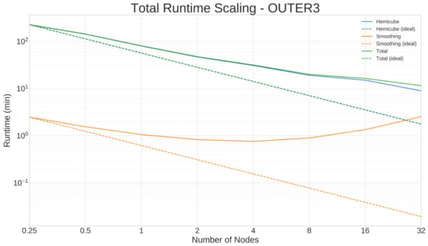

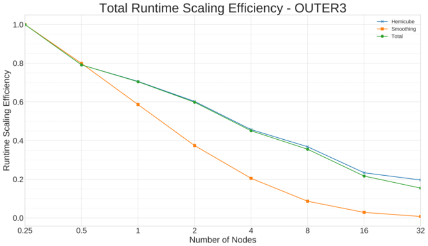

Run Time Scaling: Mesh sizes

- All runs with hc_res=500, min_sep=30.0
- Better scaling on finer meshes.
  - More work to go around, better latency hiding
- Superlinear speed-up on OUTER2
  - Hard to pin down: more cache per rank is one likely reason

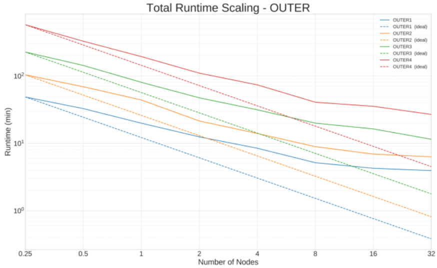

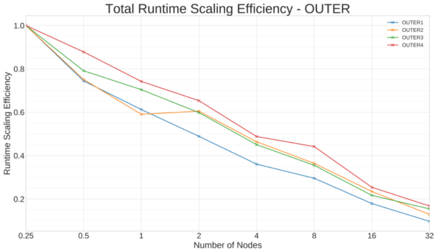

Memory Usage

- Disk space usage
  - Mesh size and geometry are main factors.
  - Parameters don't really affect output file size.
- Main memory (RAM) usage
  - Each rank has a copy of the entire mesh
  - Total memory usage increases linearly with number of ranks

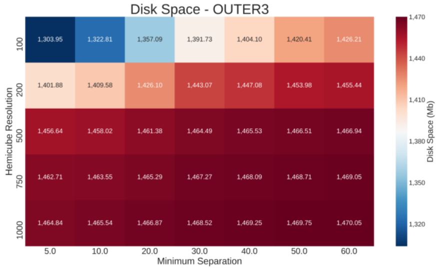

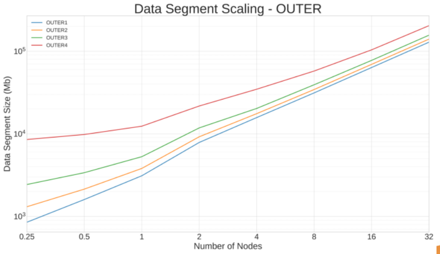

Quality of Solution

Definitions:

- Φ: the view factor matrix
- A: diagonal matrix where a_{ii} is the area of face i
- q: incident power per unit area on each face
- e: exitant power per unit area on each face

These quantities are related by

$$Aq = A\Phi e$$

Comparing solutions:

Given a VF matrix Φ, we compare it against a known high-quality solution Φ_{best}. Taking their difference $\delta\Phi$ yields:

$$A\delta q = A\delta\Phi e$$

Quality of Solution

Three measures of accuracy:

1. ℓ_1 **matrix norm:** max absolute column sum of $\delta\Phi$

$$\underbrace{\|A\delta q\|_\infty}_{\text{max incident energy error on any face}} \leq \|\delta\Phi\|_1 \underbrace{\|Ae\|_\infty}_{\text{max exitant energy of any one face}}$$

2. ℓ_∞ **matrix norm:** max absolute row sum of $\delta\Phi$

$$\underbrace{\|A\delta q\|_1}_{\text{total absolute diff. in incident energy}} \leq \|\delta\Phi\|_\infty \underbrace{\|Ae\|_1}_{\text{total exitant energy}}$$

3. **NZ delta:** percent of total non-zeros added in smoothing step

Los Alamos NATIONAL LABORATORY EST. 1943

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

ℓ_1 norm for OUTER3 (single node)

- Each run compared against most accurate test.
  - hemicube_resolution=1000, min_separation=60.0

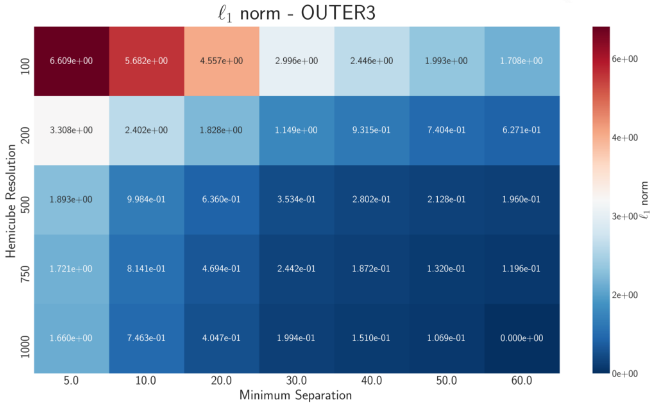

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

ℓ_∞ norm for OUTER3 (single node)

- Each run compared against most accurate test.
  - hemicube_resolution=1000, min_separation=60.0

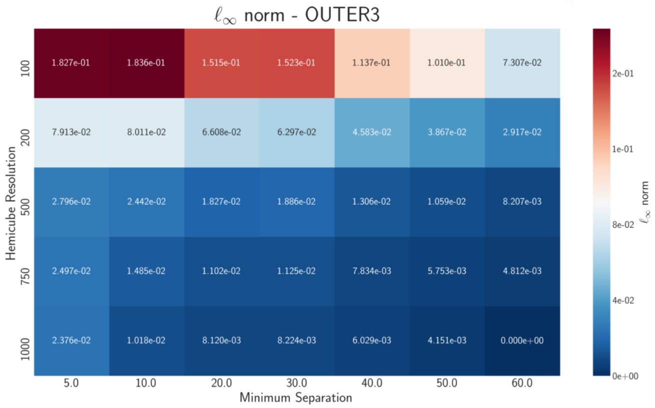

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Non-zero delta for OUTER3 (single node)

- *genre* prints this information after smoothing step
- Good measure of accuracy when high-quality solution is not available

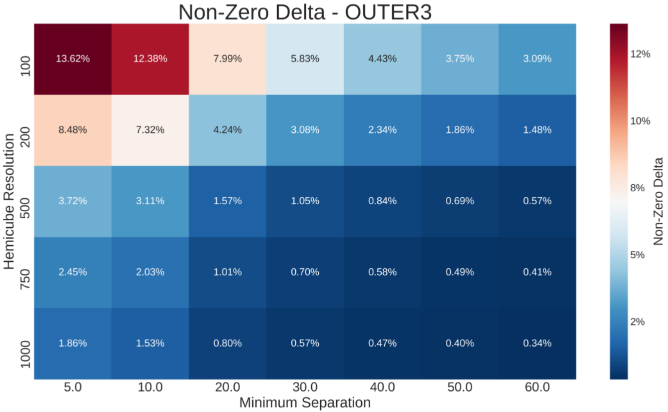

VF Matrix Visualization

- Visualize each column of the VF matrix as a field on the enclosure.
  - If visualizing column *j*, then each face *i* is rendered with ϕ_{ij}.
  - Intuitively, face *j* illuminates the other faces as the only light source.
  - Result should have smooth gradients with no speckling.

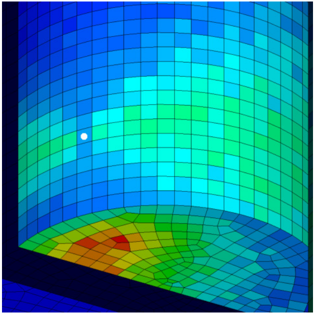

Poorer VF matrix

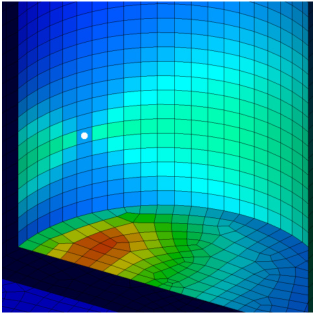

Better VF matrix

Los Alamos NATIONAL LABORATORY EST. 1943
Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

VF Matrix Visualization for OUTER3

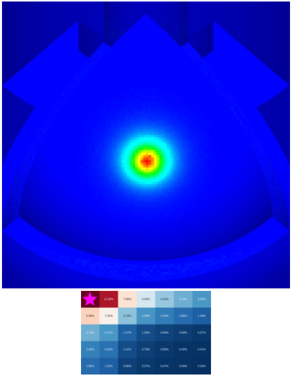

(100, 5.0)

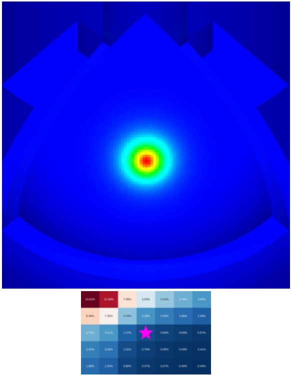

(500, 30.0)

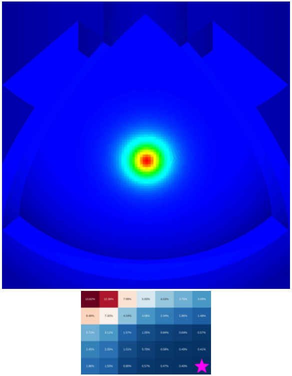

(1000, 60.0)

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

VF Matrix Visualization for OUTER3

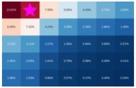

(100, 10.0)

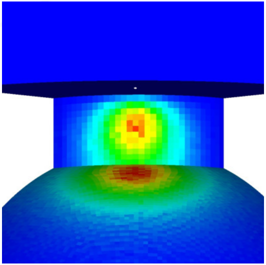

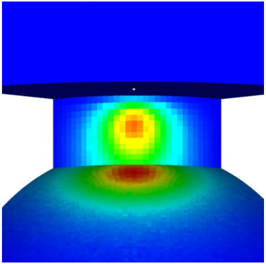

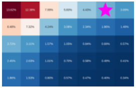

(100, 50.0)

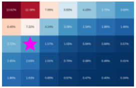

(500, 10.0)

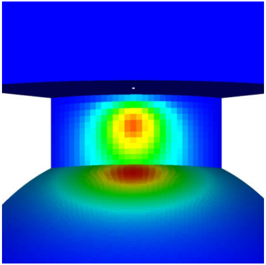

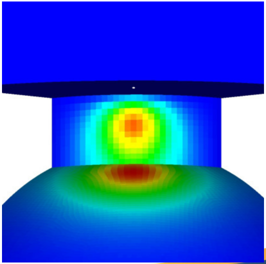

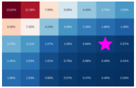

(500, 50.0)

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. Accuracy

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. Accuracy

- What parameter combination gives the best bang for your buck?
- Depends on your priority:
  - Highest quality solution for given runtime
  - Shortest runtime for given solution accuracy
- We can quantify the quality of a parameter combination with respect to these two axes
  - Weigh runtime and accuracy equally.
  - Quality metric: normalized euclidean distance from origin.

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. ℓ_1 norm for OUTER3 (single node)

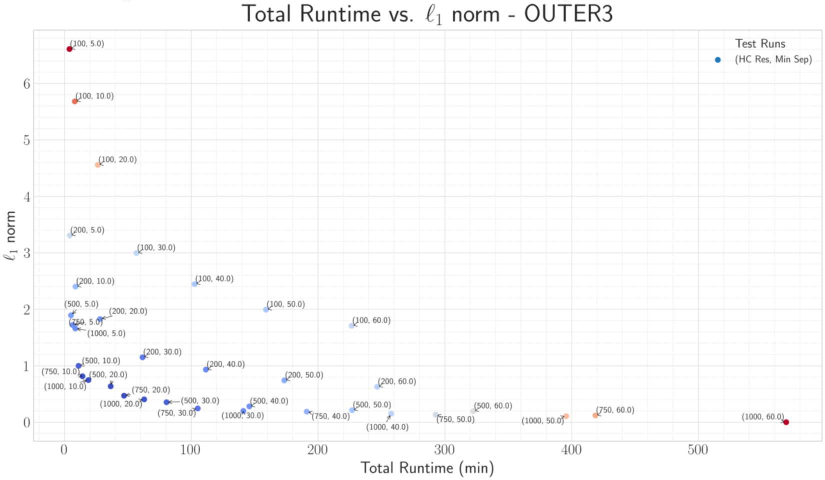

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. ℓ_∞ norm for OUTER3 (single node)

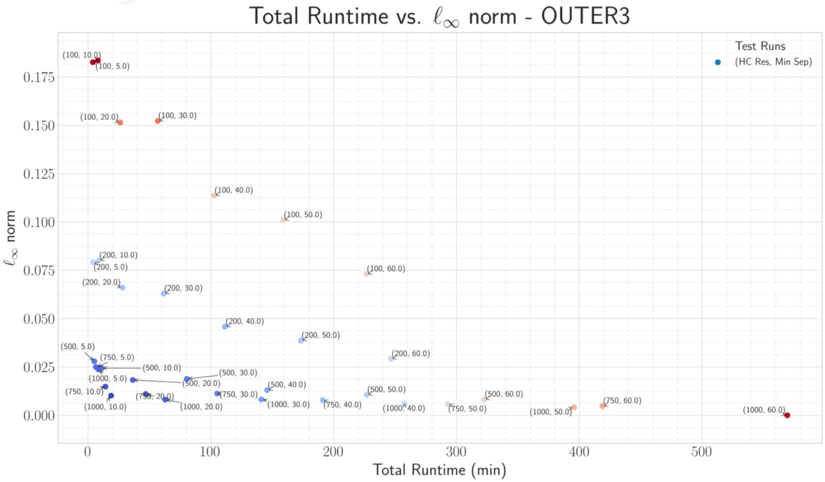

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. NZ delta for OUTER3 (single node)

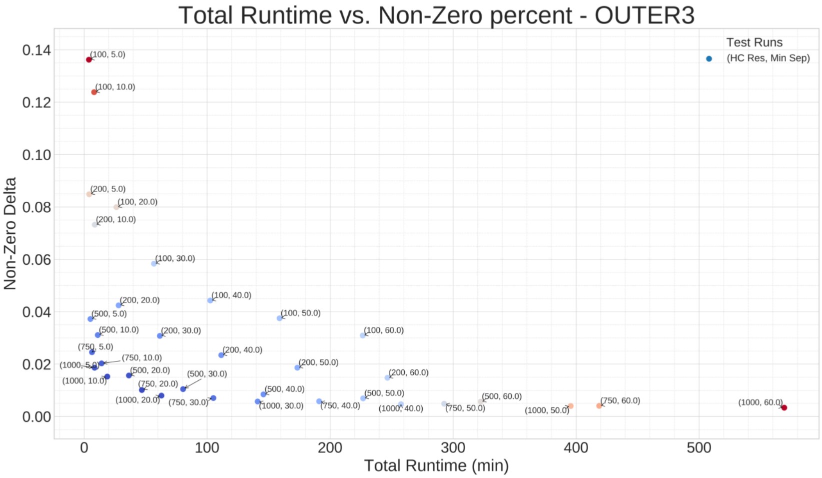

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. NZ delta for OUTER3 (single node)

- Optimal parameters along red curve, depending on constraints

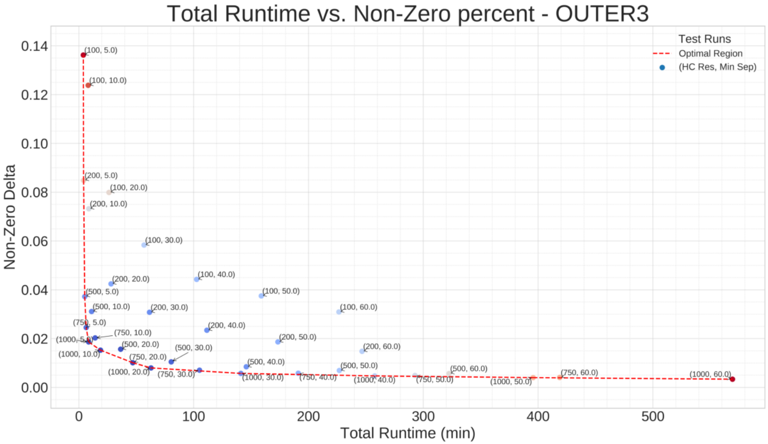

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Runtime vs. NZ delta for OUTER3 (single node)

- hc_res improves solution with less impact on runtime than min_sep
- After a point increasing either costs much more time for only a small improvement in solution accuracy

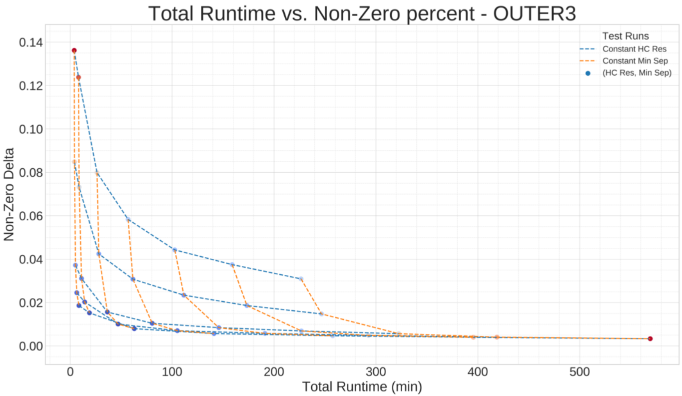

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Quality metric comparison

- As mesh size increases, hc_res becomes more important than min_sep

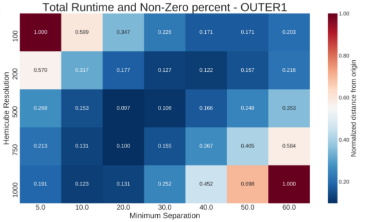

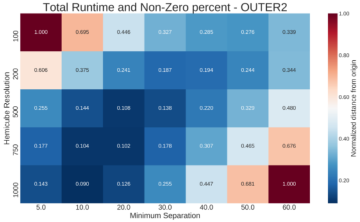

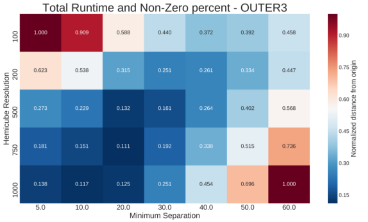

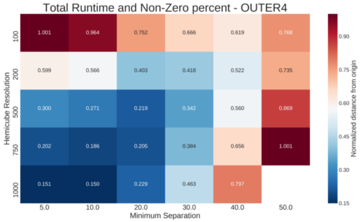

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Conclusions

- Runtime depends quadratically on hc_res and min_sep
- NZ delta is a good measure of accuracy without ‘gold standard’
- As resolution increases, hc_res becomes more important than min_sep
  - At higher resolutions faces are already smaller, so there is less need to subdivide.
  - At lower resolutions, low resolution hemicubes tend to hit most cells. Faces are larger, so greater error from violating proximity assumption.
- Parameter tuning heuristics:
  - A small min_sep is sufficient (10-30). Large values take too long for little benefit
  - Better to increase hc_res (up to a point). 200-750 is sufficient, depending on mesh resolution.
  - These are just suggestions based on basic-hemi geometry.
  - Always visualize VF matrix to ensure quality solutions.

Los Alamos NATIONAL LABORATORY EST. 1943

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

References

BARKER, E. I., CUMMINS, S. J., AND KORZEKWA, D. A.
A 3D uranium casting simulation from preheat to cooldown.
Tech. Rep. LA-UR-08-00472, Los Alamos National Laboratory, 2008.
http://permalink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-UR-08-00472.

GLASS, M. W.
CHAPARRAL: A library for solving large enclosure radiation heat transfer problems.
Tech. Rep. SAND95-2049, Sandia National Laboratory, 1995.
https://digital.library.unt.edu/ark:/67531/metadc618866/.

HEALD, J.
Nusselt analog.
https://commons.wikimedia.org/w/index.php?title=File:Nusselt_analog.svg&oldid=224385674, 2011. [Online; accessed 4-March-2019].

VUEGHS, P., AND BECKERS, P.
Presentation of the hemisphere method.
In *Advanced Computational Methods in Heat Transfer IX* (06 2006), pp. 121–130. doi:10.2495/HT060121.

Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA