High-Performance Computing and Uncertainty Quantification for Particle Transport Problems: Beyond Nuclear Energy

Ryan McClarren
Texas A&M University
Dept. of Nuclear Engineering

6 Apr 2016

TEXAS A&M UNIVERSITY

Acknowledgements

- Graduate Students:
  - Vincent Laboure  , Weixiong Zheng 
  - Pablo Vaquer  , Jonathan Madsen 
  - James Tompkins  , Lance Merchant 
  - Kelli Humbird  , Jacob Landman 
  - Yuriy Ayzmann  , Daniel Holladay, Can Pu, Arnulfo Gonzalez

Acknowledgements

- Collaborators:
  - Prof. Marvin Adams, Prof. Jim Morel, and Prof. Jean Ragus (Texas A&M)
  - Cory Hauck (Oak Ridge National Lab)
  - Martin Frank (RWTH Aachen)
  - Paul Laiu, André Tits, and Dianne O'Leary (U. Maryland)
- Sponsors:
  - US National Science Foundation, Division of Mathematical Sciences
  - US Dept. of Energy, US National Labs: Los Alamos National Laboratory, Lawrence Livermore National Laboratory, Idaho National Laboratory

Section 1

Introduction

TEXAS A&M UNIVERSITY

The model of linear radiation transport describe many important phenomenon.

The linear Boltzmann equation very accurately describes the behavior of a variety of transport processes

- Neutrons in a nuclear reactor,
- Radiation in oil well monitoring, medical imaging, and therapies
- X-rays in high energy density situations: inertial confinement fusion, astrophysical radiating shocks
- Atmospheric radiative transfer
- Neutrinos in core-collapse supernovae
- Electron/ion/photon transport in radiotherapy, space weather, electronics

Nuclear Fission

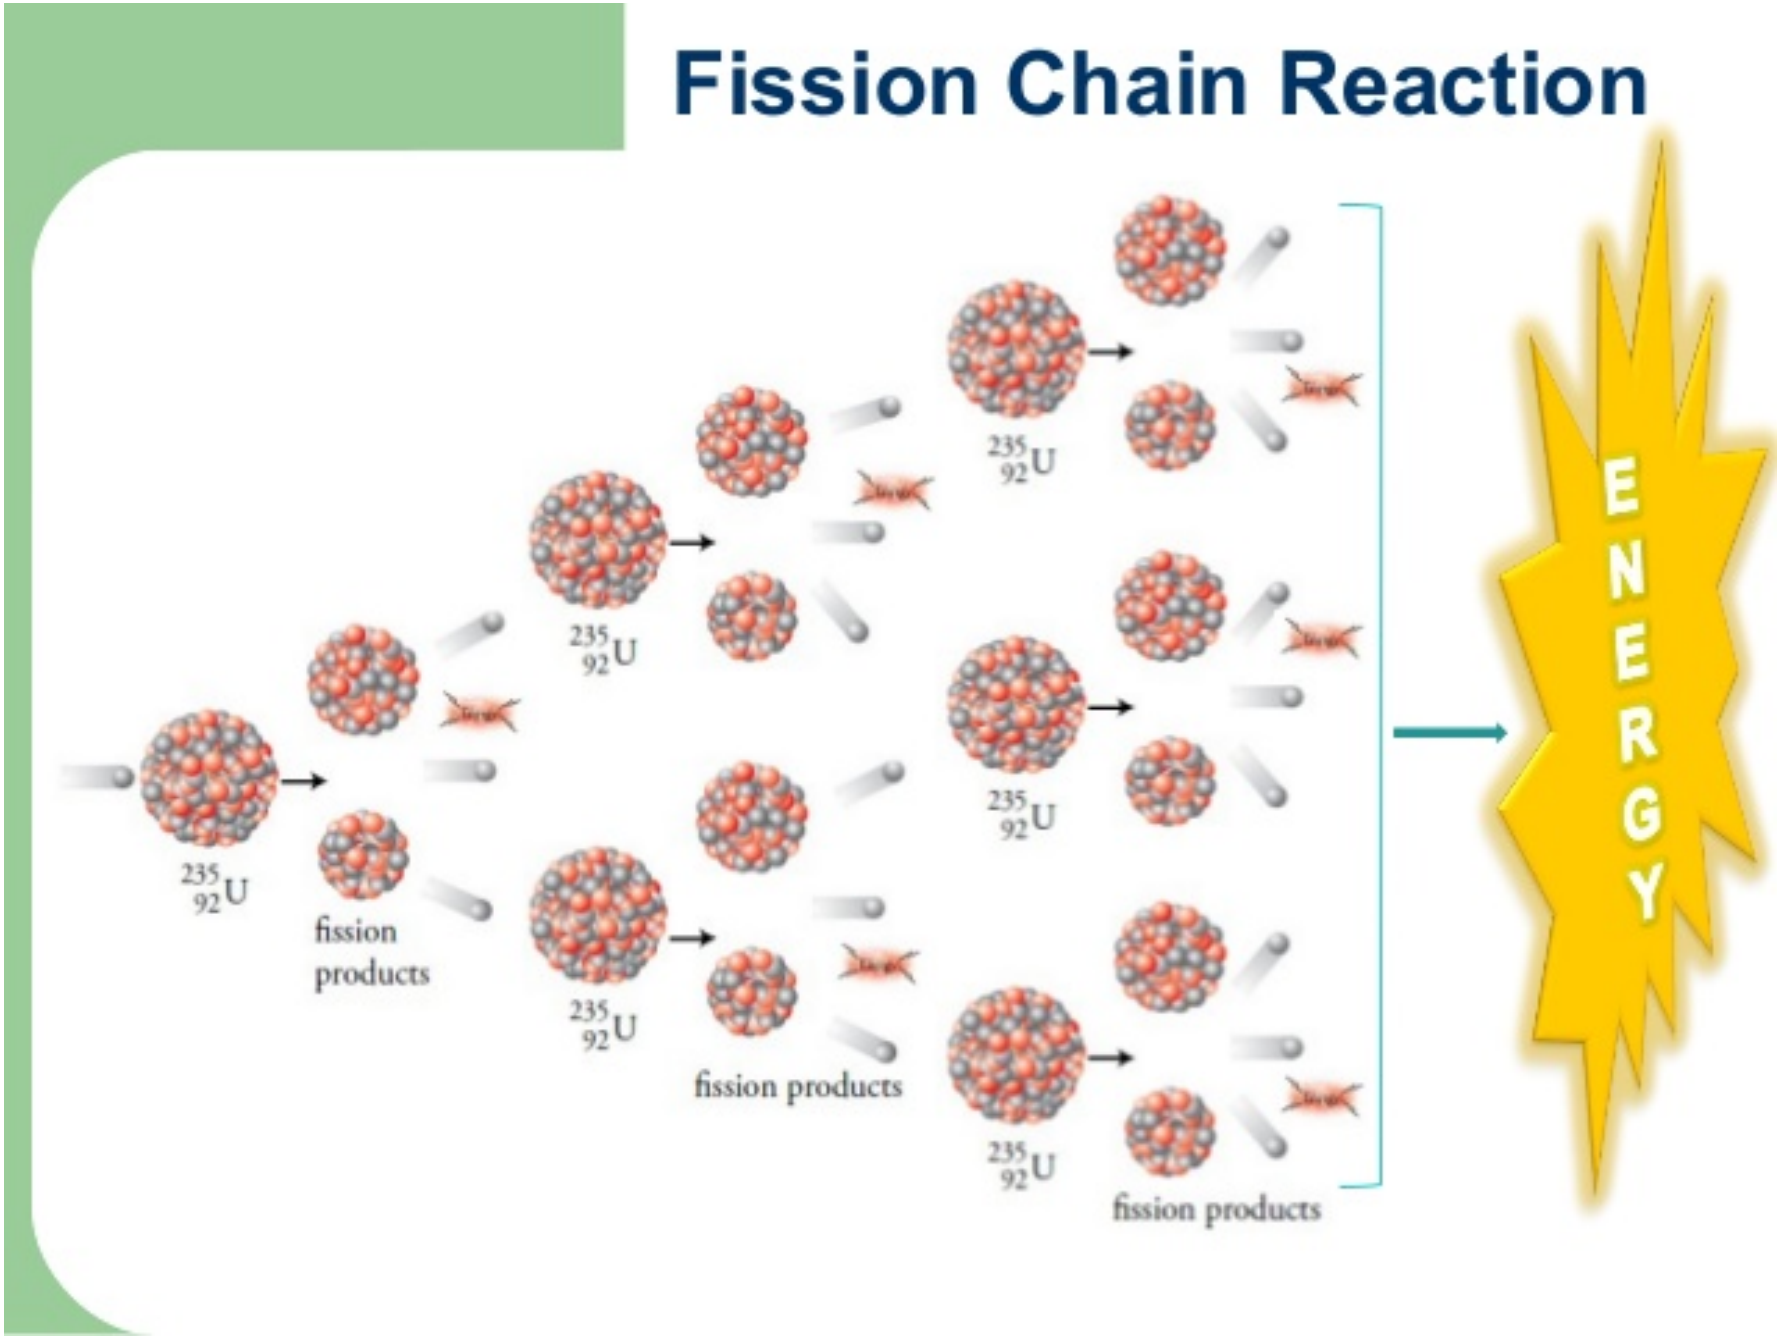

Image slidesharecdn.com/nuclearreactor-150817051416-lva1-app6891/95/nuclear-reactor-16-638.jpg

Nuclear Plant

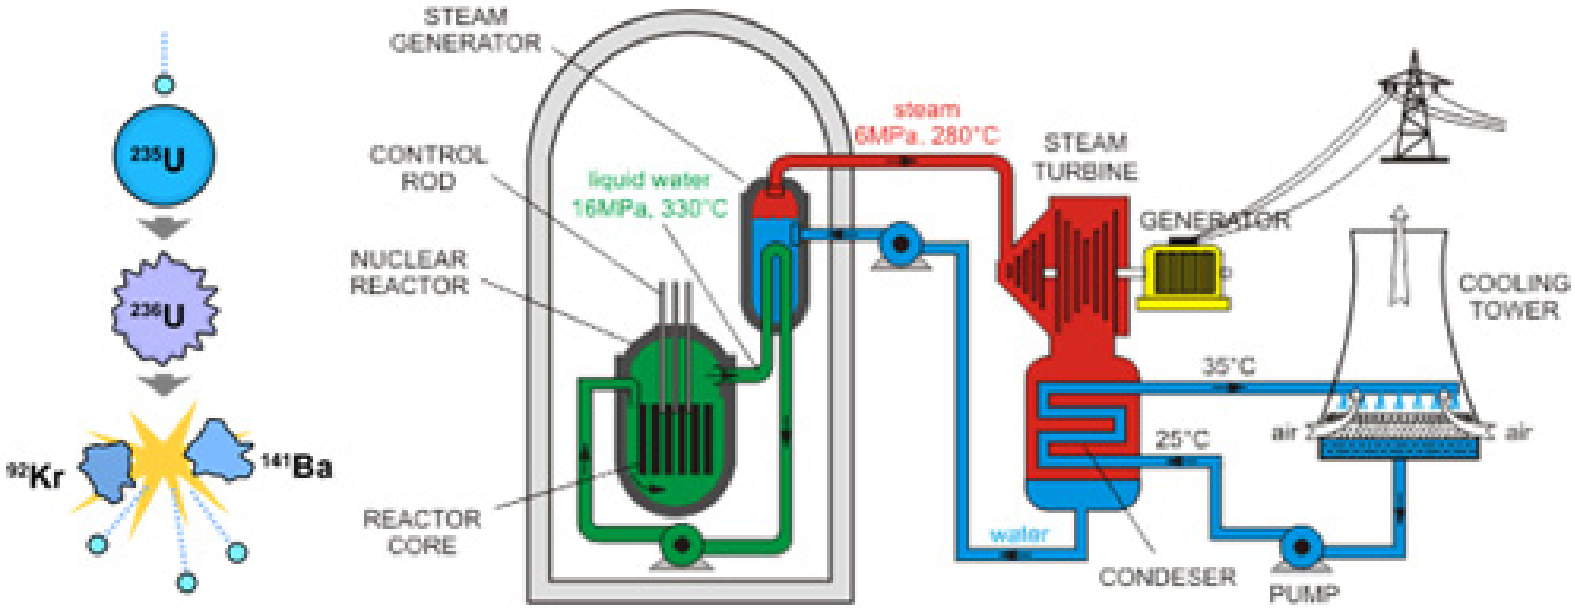

yale.edu/ynhti/curriculum/images/2011/4/11.04.05.02.jpg

Nuclear Plant

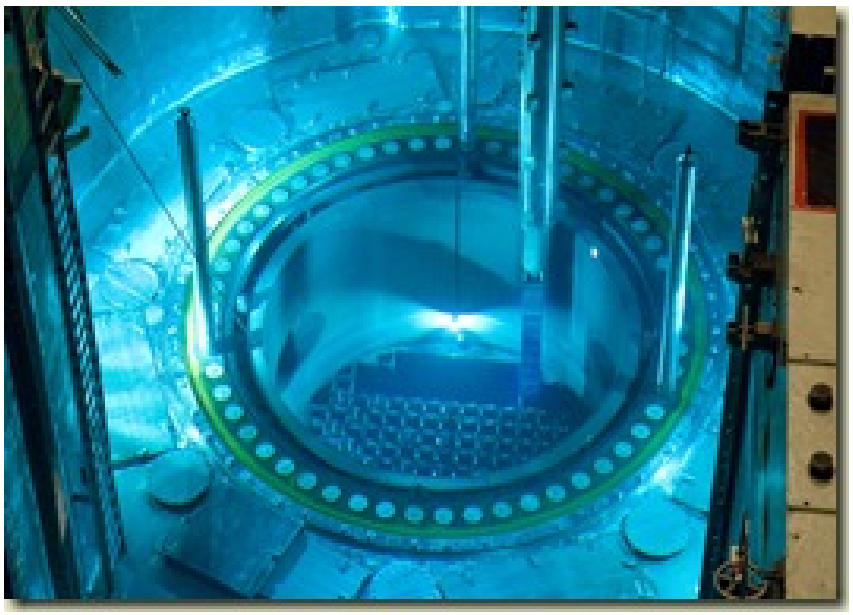

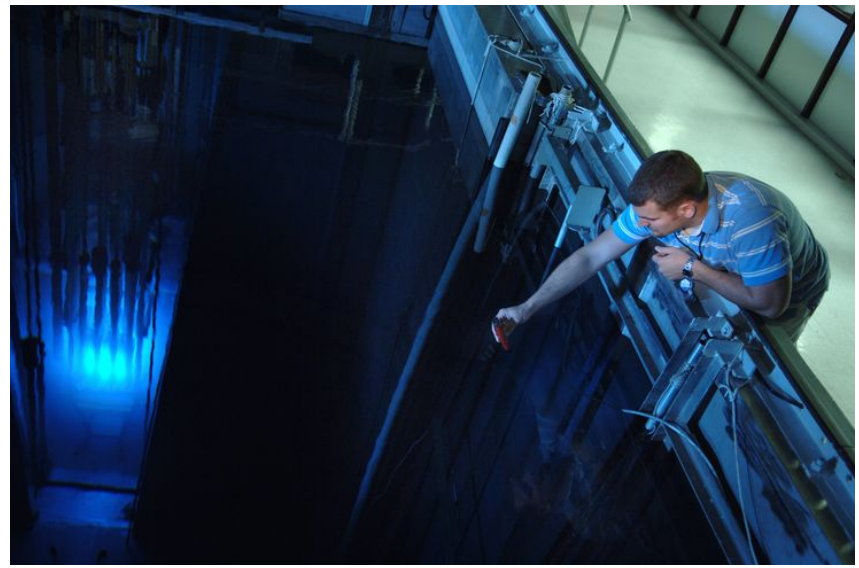

Oil Well Logging

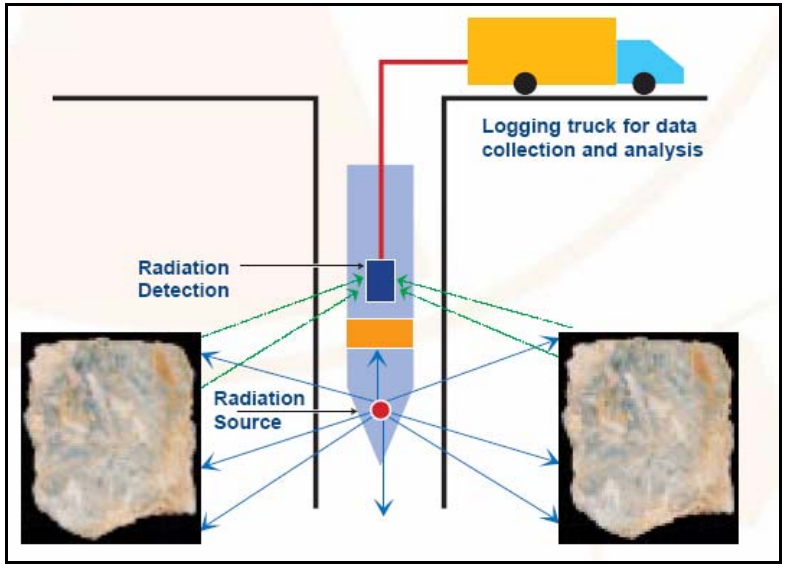

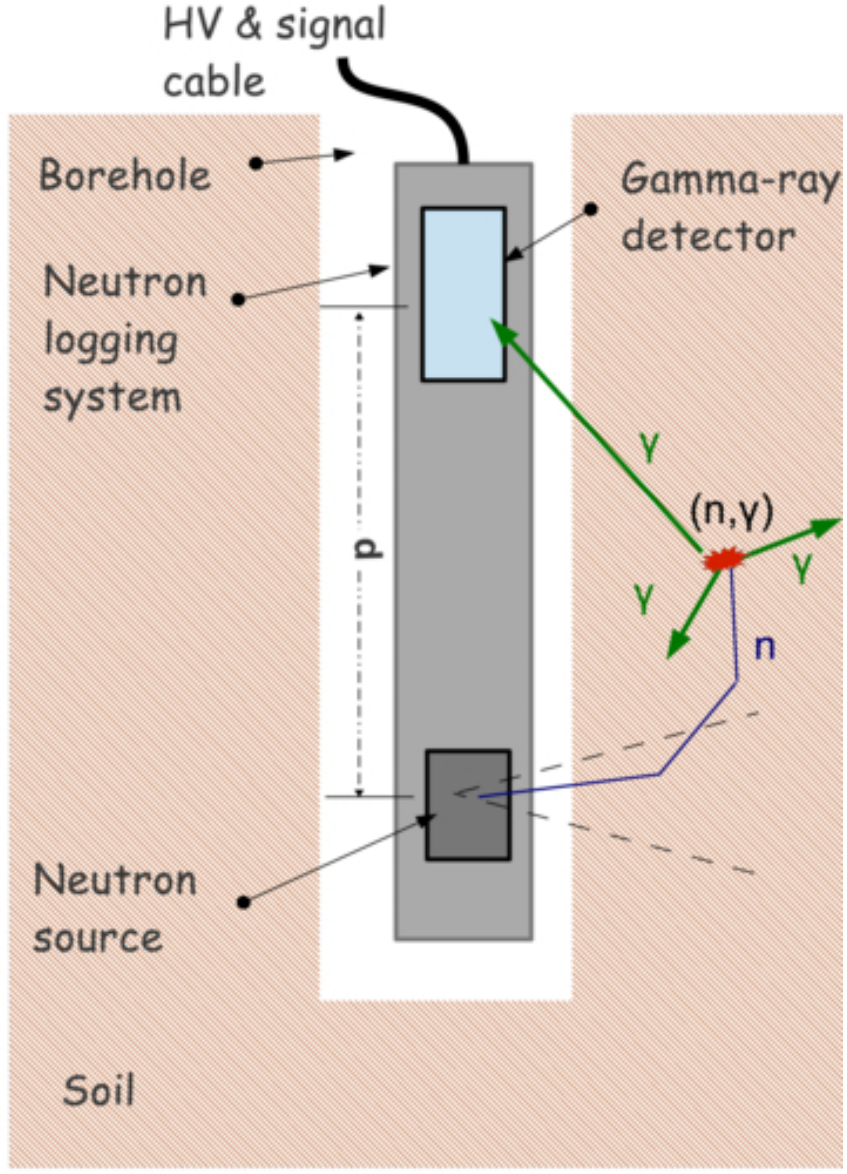

www.kau.edu.sa/Files/320/Researches/47387_18847.pdf, http://wand.images.worldnow.com/images/10020730_G.jpg, http://iopscience.iop.org/0957-0233/24/12/125903/downloadFigure/figure/mst487195f1

Inertial Confinement Fusion

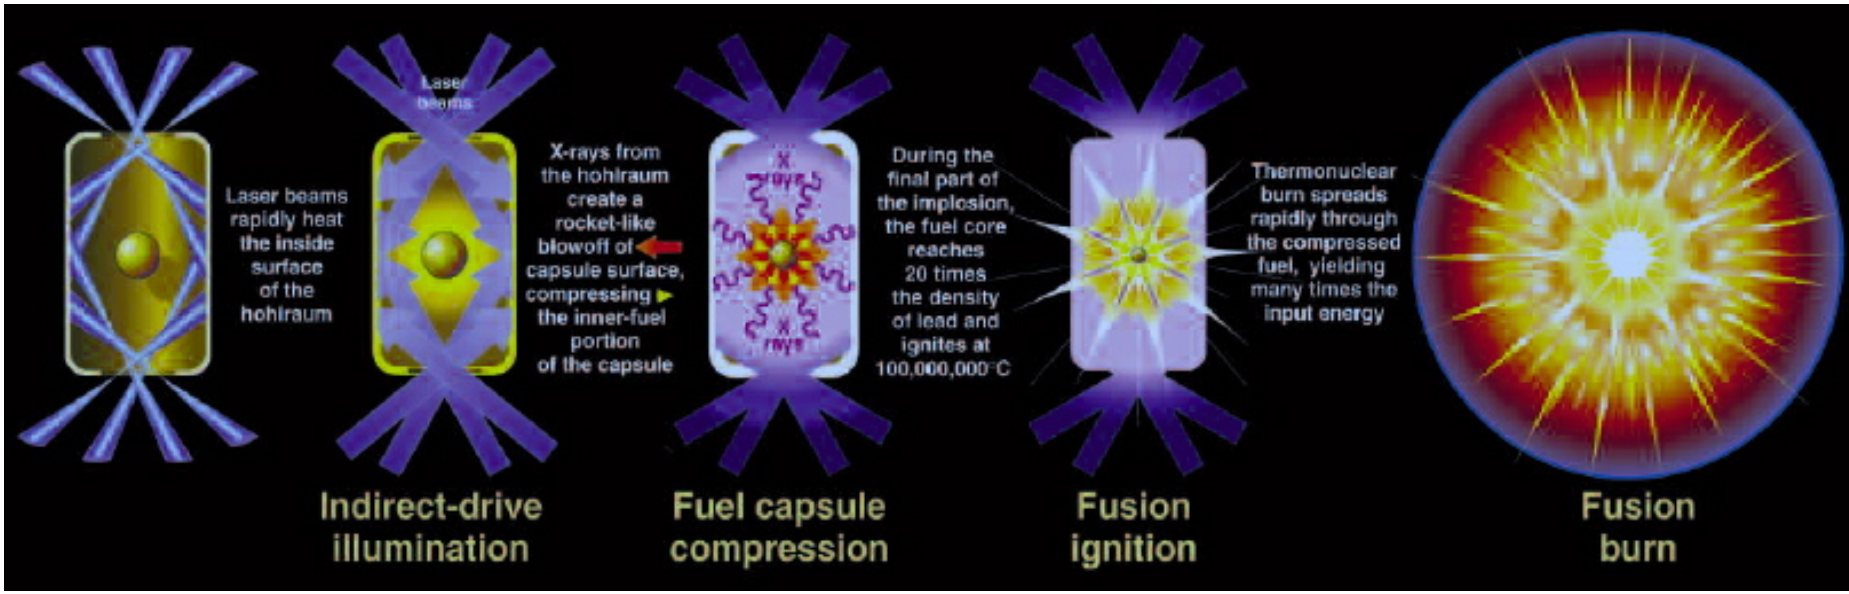

physics.ox.ac.uk/~Gregori/icf/assets/icf-principle.jpg

National Ignition Facility, Livermore, CA

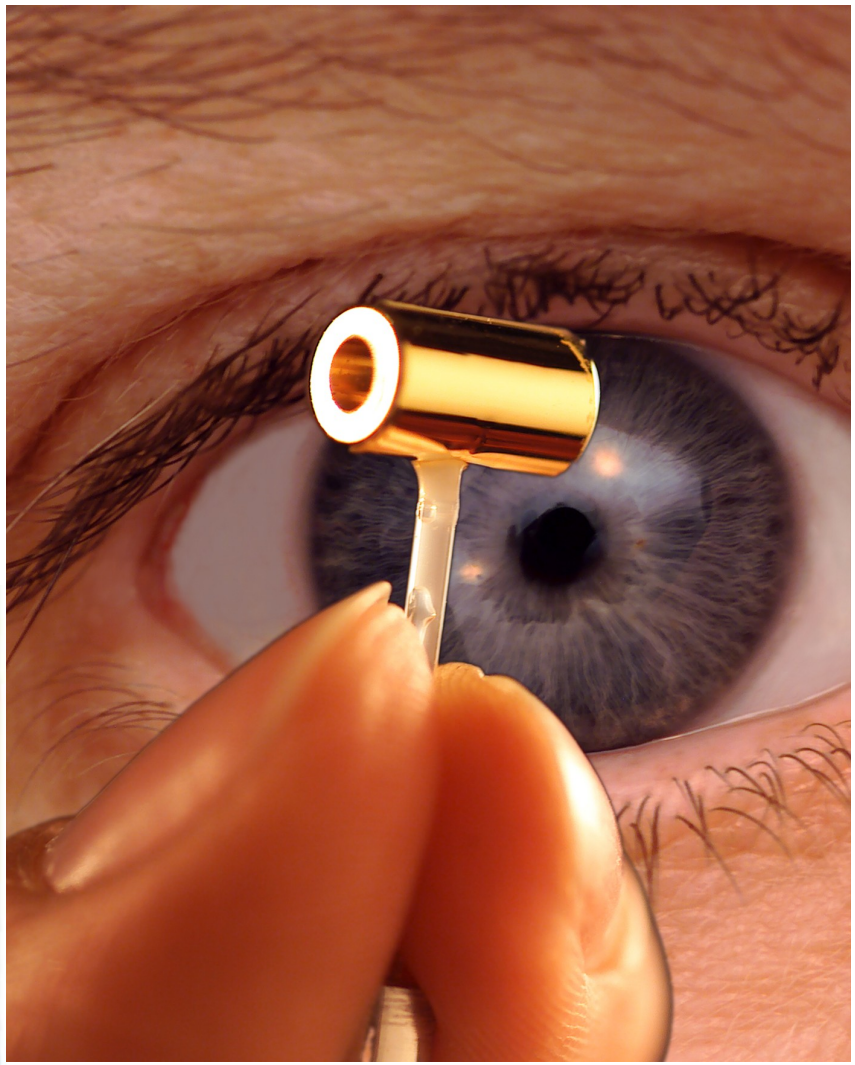

Core-Collapse Supernova

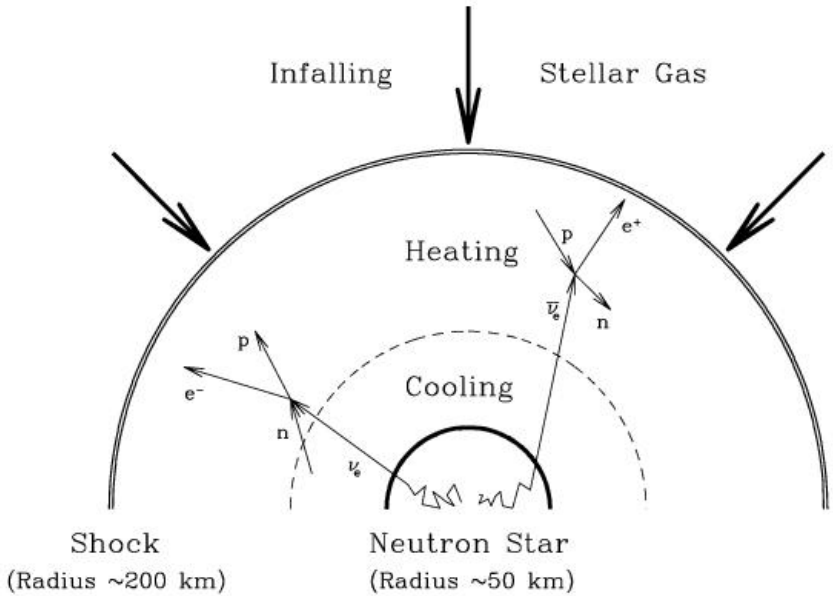

www.mppmu.mpg.de/supernova/shadowing/sn1.jpg

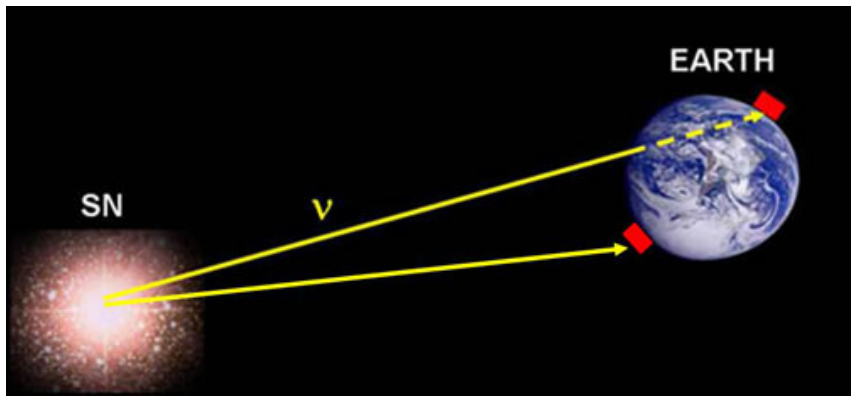

Core-Collapse Supernova

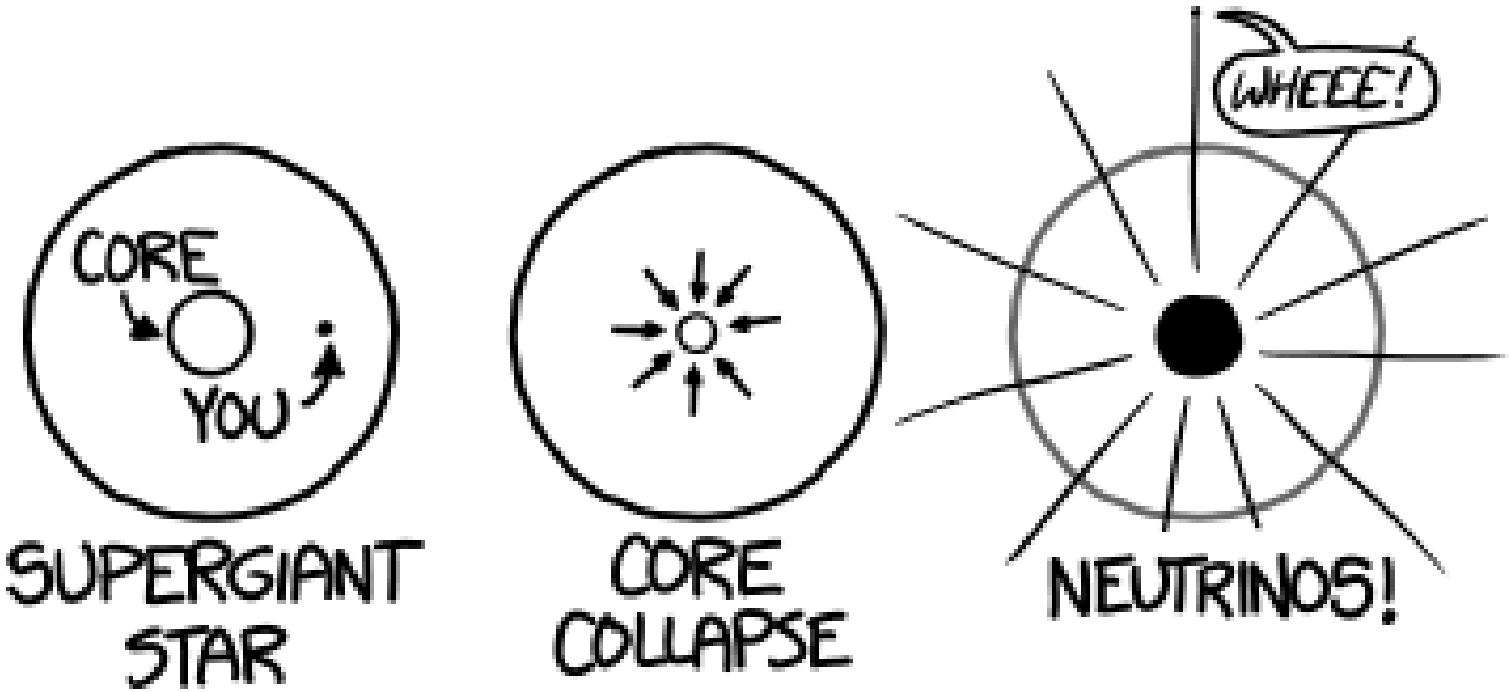

xkcd

Radiation Therapy

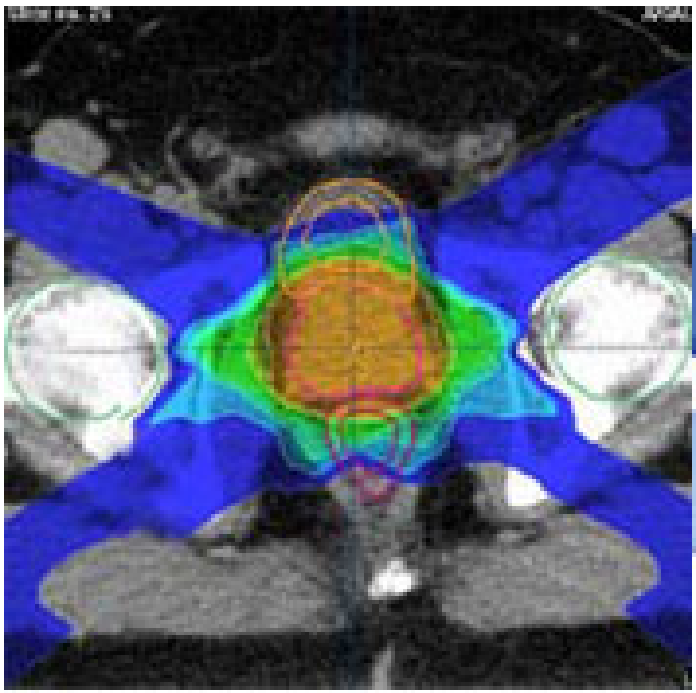

precisionradiotherapy.com/images/PE_Prostateplan.jpg

Atmospheric Radiative Transfer

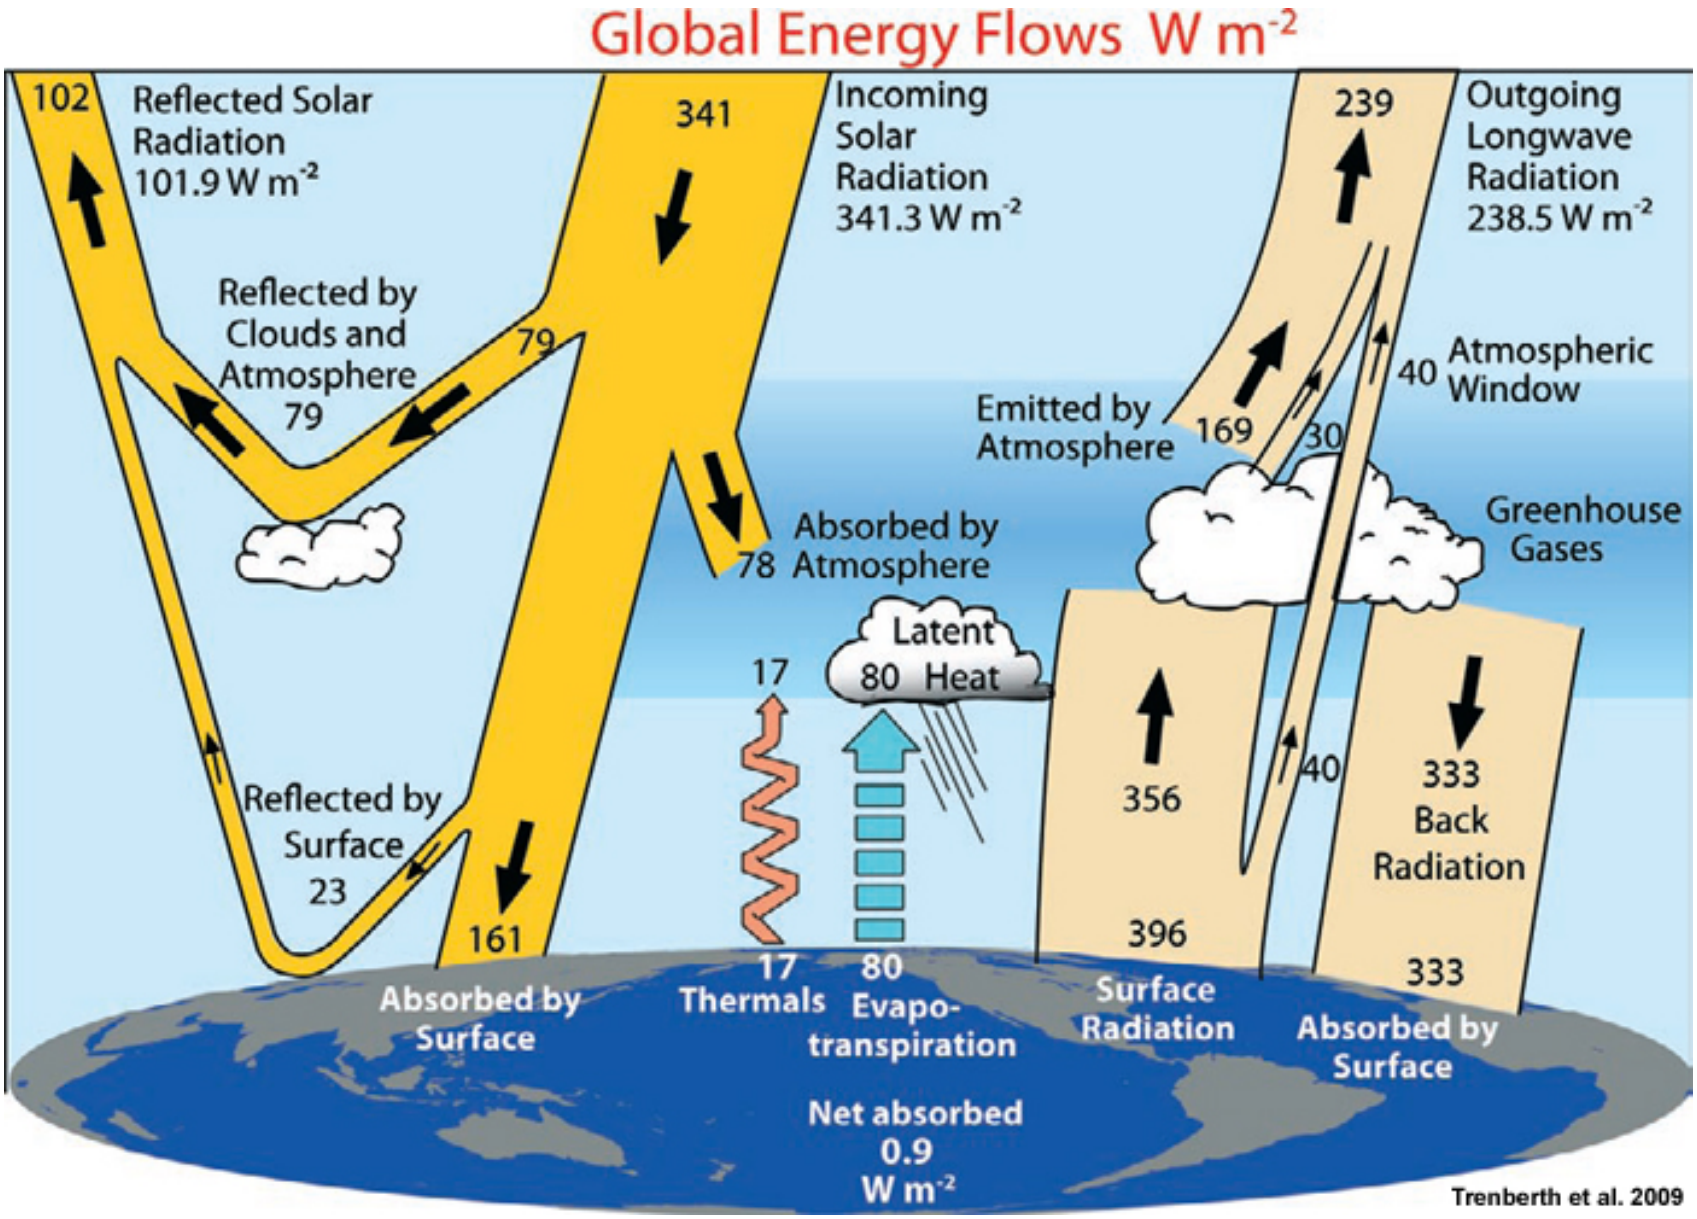

www.goes-r.gov/users/comet/tropical/textbook_2nd_edition/media/graphics/global_energy_flows.jpg

Section 2

Numerical Methods for the Linear Boltzmann Equation

TEXAS A&M UNIVERSITY

The Linear Boltzmann Equation describes the transport of particles that travel in straight lines between collisions.

- We are interested in the phase-space density of particles, N, that travel in straight-lines between collisions. The equation that describes this is the linear-Boltzmann equation:

$$(\partial_t + v\Omega \cdot \nabla + v\sigma_{\mathrm{t}}(x,t))\, N(x,\Omega,v,t) =$$
$$\int_{\mathbb{S}_2} d\Omega' \int_0^\infty dv'\, v'\sigma_s(x,t,\Omega' \to \Omega, v' \to v) N(x,\Omega,v,t) + Q(x,\Omega,v,t)$$

- $\Omega \in \mathbb{S}_2$ is the direction of the particle's flight (angular variable), v is the particle speed.
- The interaction probabilities (cross-sections) are the total cross-section σ_{t} which is the average number of collisions a particle undergoes with the material medium per unit distance travelled, and
- The double-differential scattering cross-section, $\sigma_s(x,t,\Omega' \to \Omega, v' \to v)$ is the mean number of particles that scatter to direction Ω and speed v per particle traveling in the differential phase space element.

TEXAS A&M UNIVERSITY

Simplifications

- For notational simplicity we will make the assumption that the discretization in speed (energy) is a solved problem, and we only need to consider a single speed equation.
- Additionally, we will assume that the scattering is isotropic.
- Both of these are simplifications for real systems. These will contribute to uncertainties later.

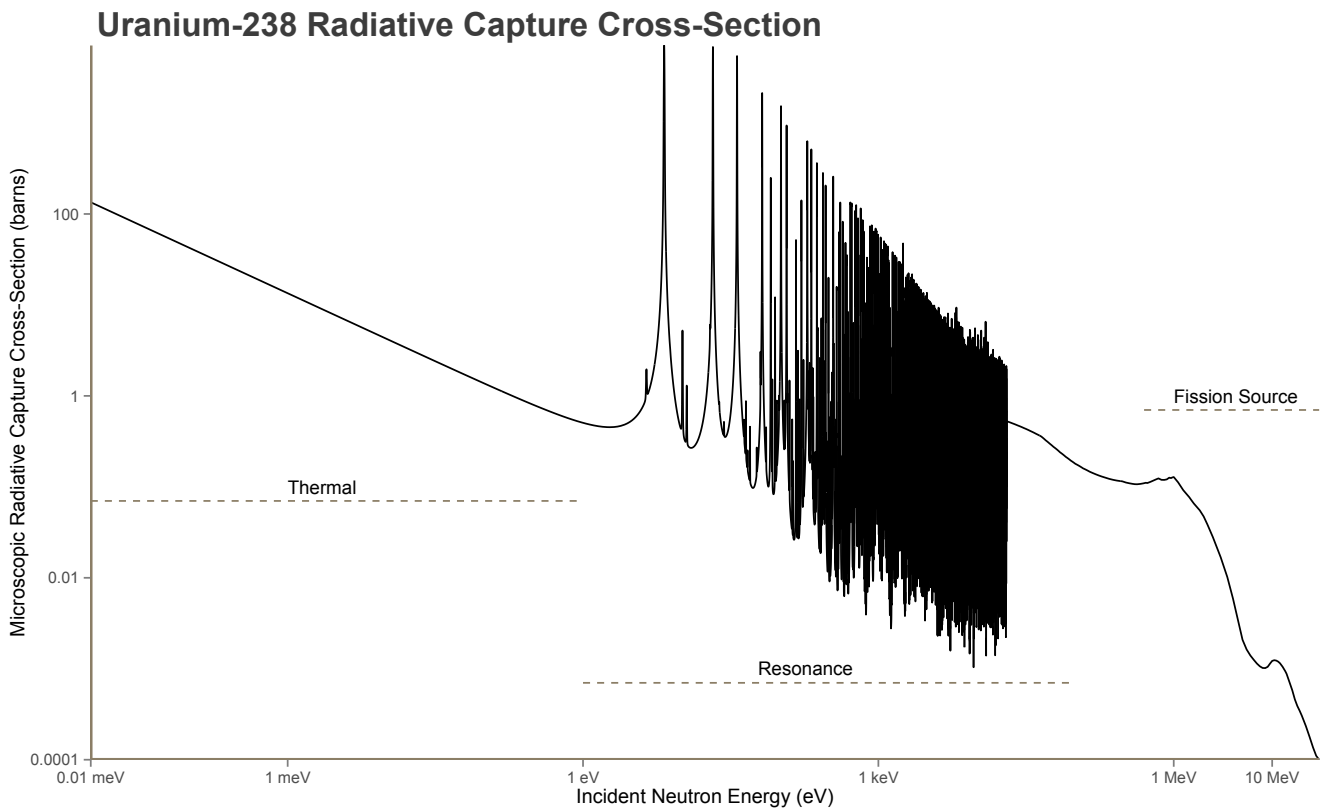

Simplified Equations

- After these simplifications we can write the resulting equation as

$$\left(v^{-1}\partial_t + \Omega \cdot \nabla + \sigma_{\mathrm{t}}(x,t)\right)\psi(x,\Omega,t) = \frac{\sigma_{\mathrm{s}}(x,t)}{4\pi}\phi(x,t) + Q(x,\Omega,t),$$

  where $\psi = vN$ and

$$\phi(x,t) = \int_{\mathbb{S}_2} d\Omega\, \psi(x,\Omega,t) = \langle\psi\rangle.$$

- We will also assume that $v = 1$. This is the same as scaling the time variable.
- Initial condition: $\psi(x,\Omega,0) = F(x,\Omega)$, and boundary conditions are inflow conditions:

$$\psi(x,\Omega,t) = \Gamma(x,\Omega,t) \quad \text{for } x \in \partial V,\ \hat{n}\cdot\Omega < 0.$$

TEXAS A&M UNIVERSITY

Numerical Challenges

- Phase space complexity
  - Need for thousands of unknowns per spatial degree of freedom
- Multiscale phenomenon
  - In problems where the scattering is large, the transport equation asymptotically limits to a diffusion equation for the particles
  - Need numerical methods that preserve this fact when the mesh does not resolve the collision length scales.
- Coupling to other physics (fluid flow, etc.)

Discrete Ordinates (S_n) method

- The discrete ordinates method is a collocation method in angle that solves the transport equation along a particular directions (Ω_j) and uses a quadrature rule, $\{w_j, \Omega_j\}$ to estimate the collision terms. (Chandrasekhar)
- Leads to a simple, triangular system of discrete equations for each direction when the backward Euler method is used in time and a simple iteration strategy is used

$$(\Omega_j \cdot \nabla + \sigma_{\mathrm{t}}^*)\, \psi_j^{\ell+1}(x, t^{n+1}) = \frac{\sigma_{\mathrm{s}}(x,t)}{4\pi} \sum_{j'} w_{j'} \psi_{j'}^{\ell}(x, t^{n+1}) + Q_j^*,$$

  $\sigma_{\mathrm{t}}^* = \sigma_{\mathrm{t}} + \Delta t^{-1}$ and $Q_j^* = Q + \psi_j(x, t^n)$.
- This iteration procedure is called a transport sweep.
- As a result when, $\sigma_{\mathrm{s}}/\sigma_{\mathrm{t}}$ is small this iteration convergences quickly, otherwise need to include the solution of a diffusion equation in the iteration.
- This is the best understood method for deterministic particle transport.

Transport Sweeps and Parallel Computing

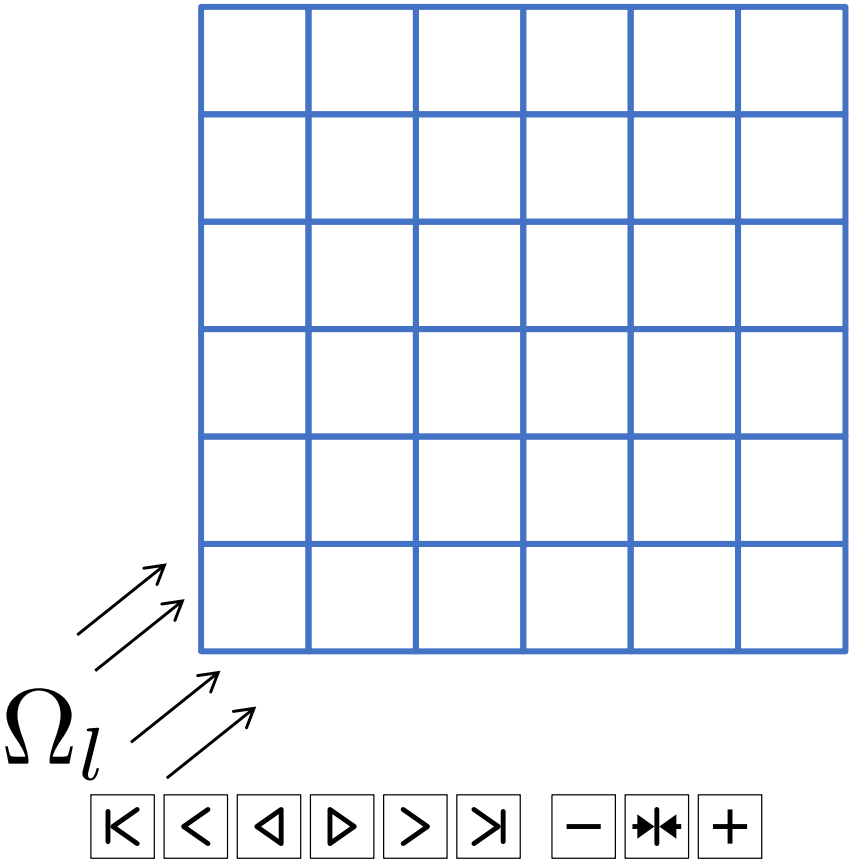

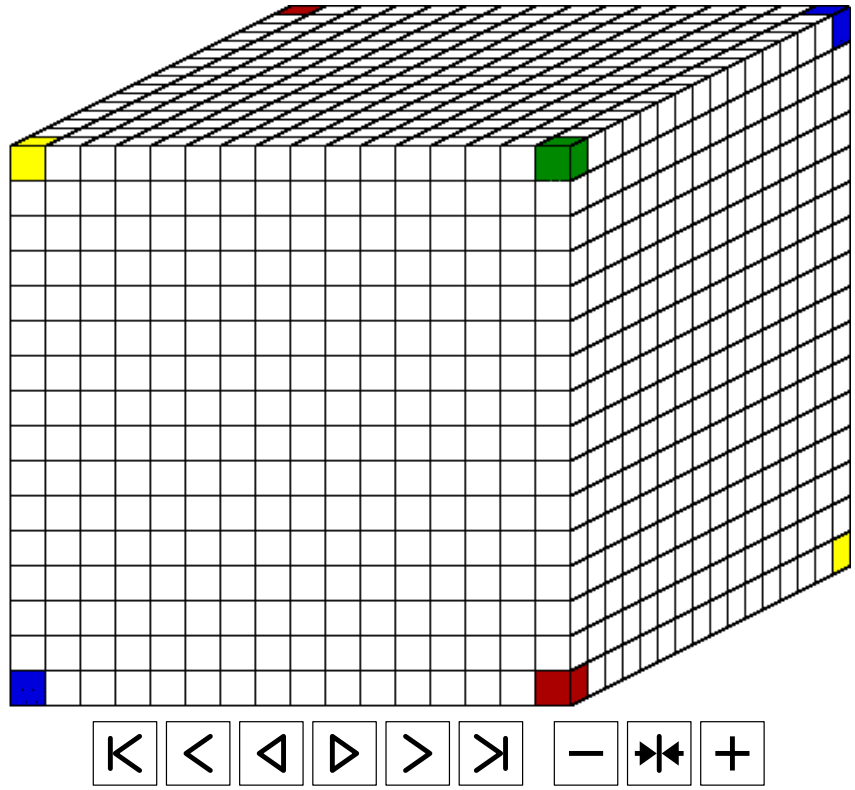

Transport Sweeps and Parallel Computing

- Despite the inherent bottlenecks in a transport sweep, there is a lot of work to group together to gain efficiencies.
- Can scale reasonably well to millions of cores.

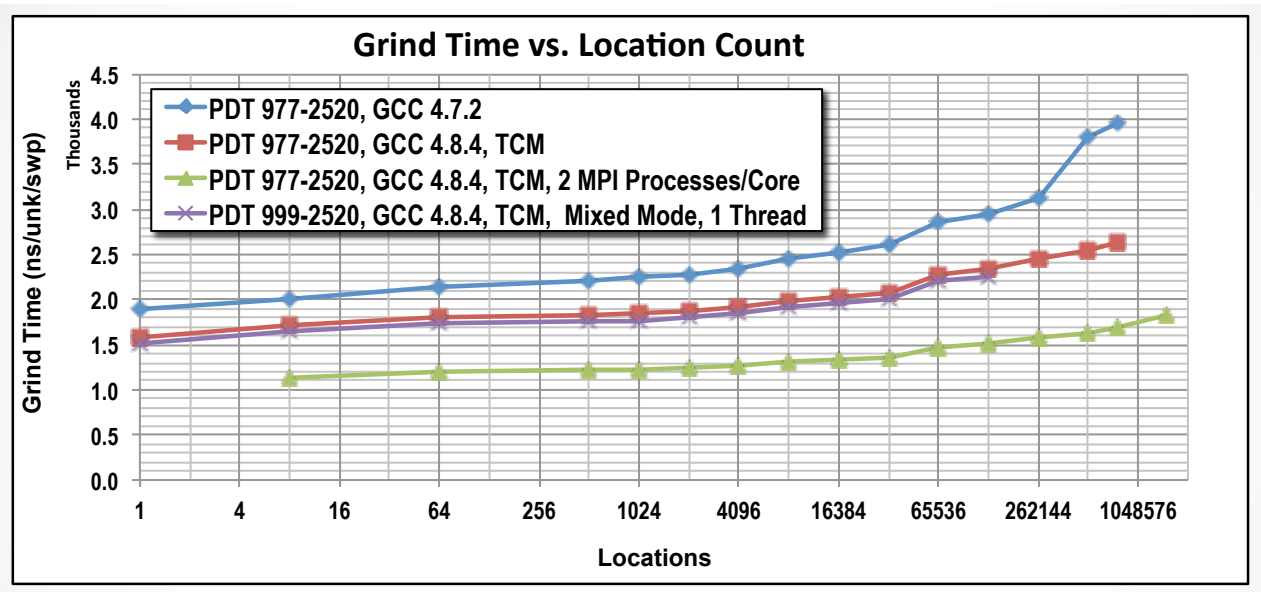

Monte Carlo

- Rather than discretize phase space directly we sample particles and advect them based on stochastic collision processes.
- Can be very accurate and operate on general domains in space and energy.
- Slow convergence $N^{-1/2}$ typically limits applicability.
- For steady-state problems it is considered the gold standard, if you can afford the simulation.

Spherical Harmonic Functions

- Decompose the angle Ω into components

$$\Omega = (\Omega_1, \Omega_2, \Omega_3)^T = (\sin\vartheta\cos(\varphi), \sin\vartheta\sin(\varphi), \cos\vartheta)^T$$

- The normalized, complex spherical harmonic of degree ℓ and order k are

$$Y_\ell^k(\Omega) = \sqrt{\frac{2\ell+1}{4\pi}\frac{(\ell-k)!}{(\ell+k)!}}e^{ik\varphi}P_\ell^k(\cos\vartheta)\,,$$

  where P_ℓ^k is an associated Legendre function.

- For convenience, we use normalized, real-valued spherical harmonics m_ℓ^k and for each degree ℓ. For given $N > 0$, set

$$\mathbf{m}_\ell = (m_\ell^{-\ell}, m_\ell^{-\ell+1}, \ldots, , m_\ell^{\ell-1}, m_\ell^{\ell})^T \quad \text{and} \quad \mathbf{m} = (\mathbf{m}_0^T, \mathbf{m}_1^T, \ldots, \mathbf{m}_N^T)^T$$

- The components of $\mathbf{m}$ form an orthonormal basis for the polynomial space

$$\mathbb{P}_N = \left\{ \sum_{\ell=0}^{N}\sum_{k=-\ell}^{\ell} c_\ell^k m_\ell^k : c_\ell^k \in \mathbb{R} \text{ for } 0 \le \ell \le N,\ |k| \le \ell \right\}. \qquad (1)$$

TEXAS A&M UNIVERSITY

Spherical Harmonics (P_N) Equations

- Spectral approximation in Ω

$$\psi \approx \psi_{P_N} \equiv \mathbf{m}^T \mathbf{u}_{P_N}$$

where $\mathbf{u}_{P_N} = \mathbf{u}_{P_N}(t, x)$ solves the P_N equations

$$\begin{cases} \partial_t \mathbf{u}_{P_N} + \mathbf{A} \cdot \nabla_x \mathbf{u}_{P_N} + \sigma_a \mathbf{u}_{P_N} + \sigma_s \mathbf{G} \mathbf{u}_{P_N} = \mathbf{s}, & (t, x) \in (0, \infty) \times \mathbb{R}^3 \\ \mathbf{u}_{P_N}(0, x) = \langle \mathbf{m} \psi_0(x, \cdot) \rangle, & x \in \mathbb{R}^3 \end{cases}$$

with

- $\mathbf{s} := \langle \mathbf{m} S \rangle$
- $\mathbf{A} \cdot \nabla_x \equiv \sum_{i=1}^3 \mathbf{A}_i \partial_{x_i}$ and each $\mathbf{A}_i = \langle \Omega_i \mathbf{m} \mathbf{m}^T \rangle$ is symmetric
- $\mathbf{G} \geq 0$ is diagonal

- Angle brackets denote integration over $\mathbb{S}^2$: $\langle \cdot \rangle := \int_{\mathbb{S}^2} (\cdot) d\Omega$

TEXAS A&M UNIVERSITY

Properties of the P_N Equations

- Good Stuff
  - Fast convergence for smooth solutions
  - Preserve rotational invariance of the transport operator
  - Harmonics are eigenfunctions of the scattering operator
- Bad Stuff
  - Gibbs phenomena near wave fronts
  - Negative values for the concentration $\langle \psi \rangle$ in multi-D
  - May be ill-posed in steady-state (A_i can have zero eigenvalues)
  - Challenging boundary conditions

The Line Source Problem: All Methods have issues[1]

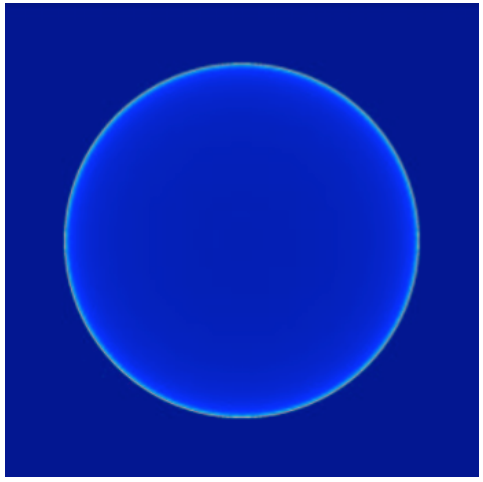

(a) analytic

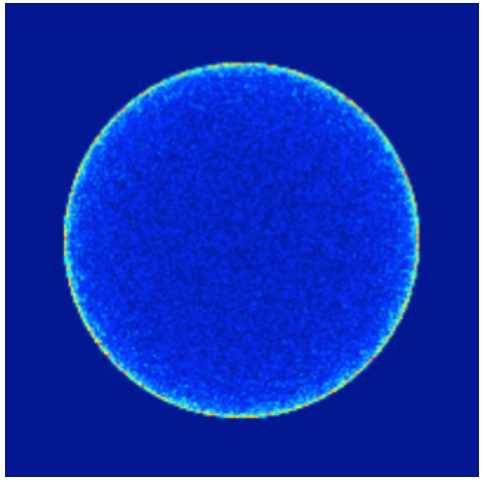

(b) Monte-Carlo

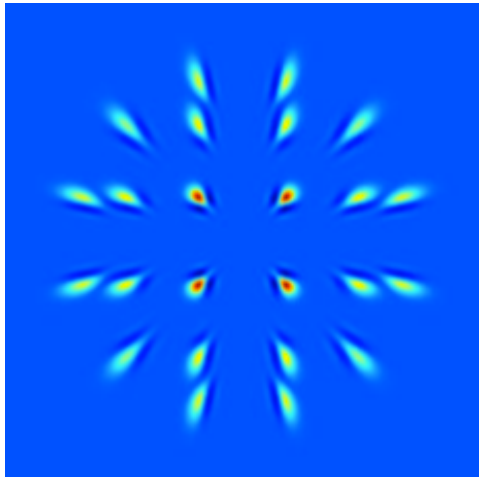

(c) S_6

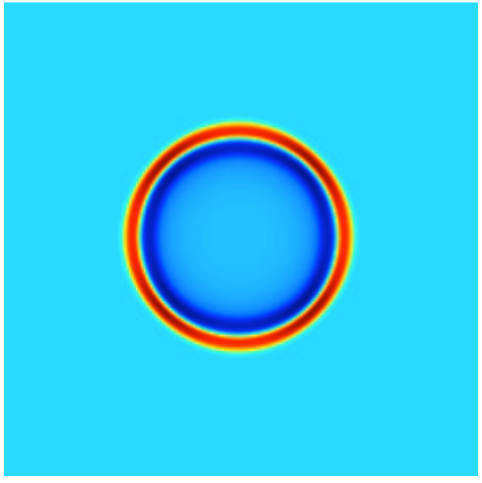

(d) P_1

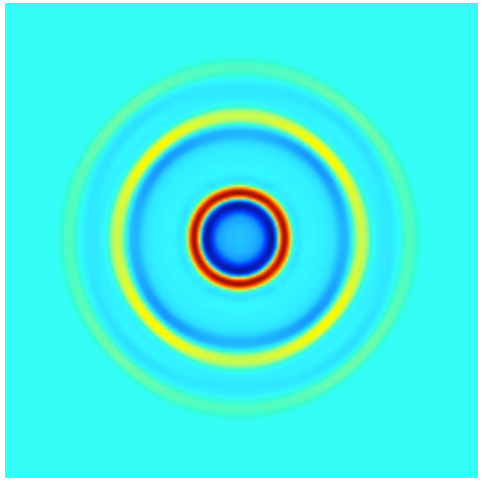

(e) P_5

[1] T. A. Brunner. "Forms of Approximate Radiation Transport", Tech. Rep. SAND2002-1778 Sandia National Laboratories, Jul 2002.

The issue is the closure

- The standard P_N closure simply truncates the expansion for $l > N$.
- The Gibbs oscillations are a result.
- The negative densities are problematic for coupled simulations: what does a negative absorption rate density mean?
- Other methods have been proposed to alleviate this issue
  - The M_N methods use the ansatz

$$\psi \approx e^{\mathbf{p}^T \mathbf{c}}$$

    to close the system.
  - Solve an optimization problem to assure that the ansatz is positive.
- Idea: Apply filters to the expansion to damp oscillations.

Back to the Linesource

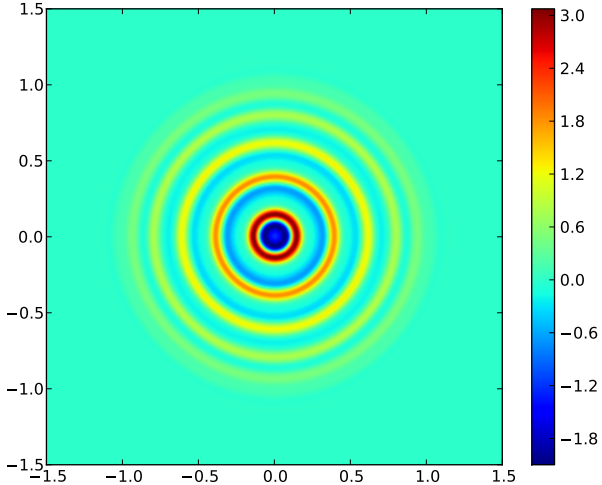

(a) P_{11}

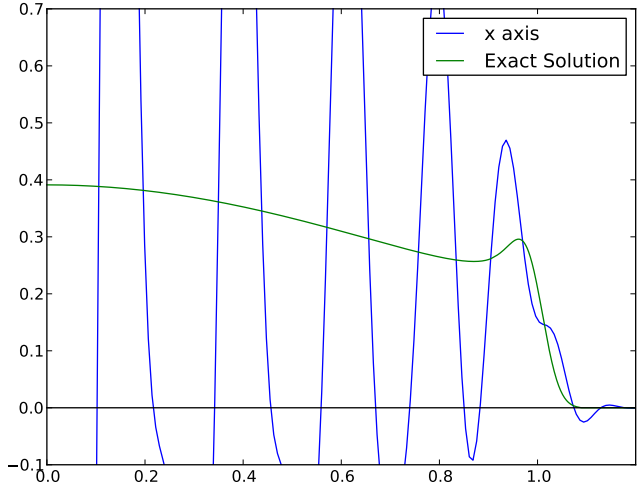

(b) P_{11}-Lineout

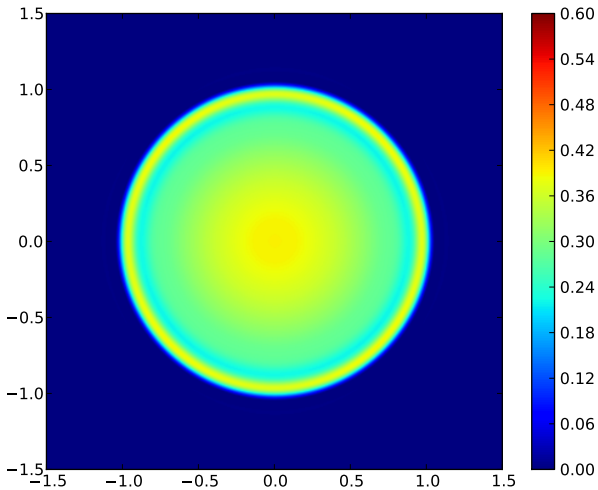

(c) FP_{11}

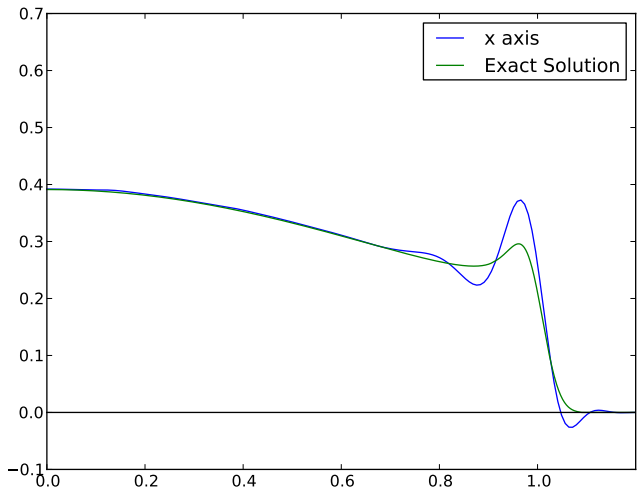

(d) FP_{11}-Lineout

TEXAS A&M UNIVERSITY

Test problem: Crooked Pipe

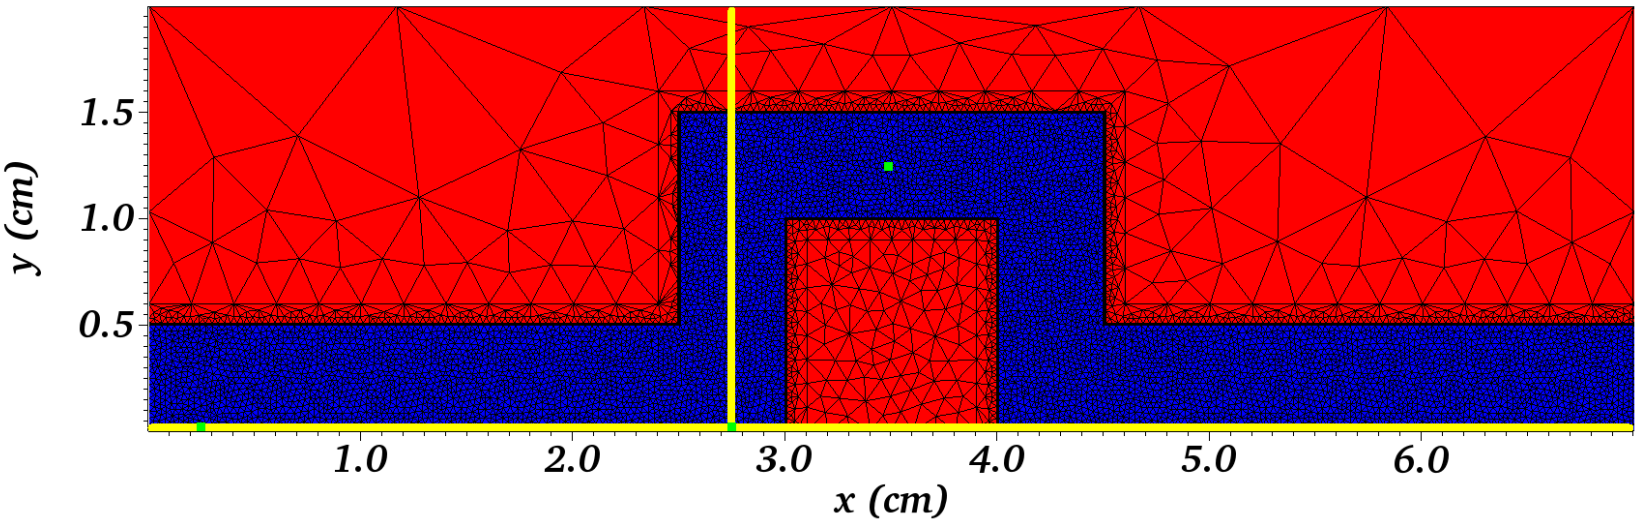

- This is a standard high energy density radiative transfer test problem.
- Blue region is optically thin (little interaction between radiation and material)
- Red regions are optically thick (strong collisions between radiation and material)
- Radiation source at left entrance.

TEXAS A&M UNIVERSITY

Unfiltered calculation

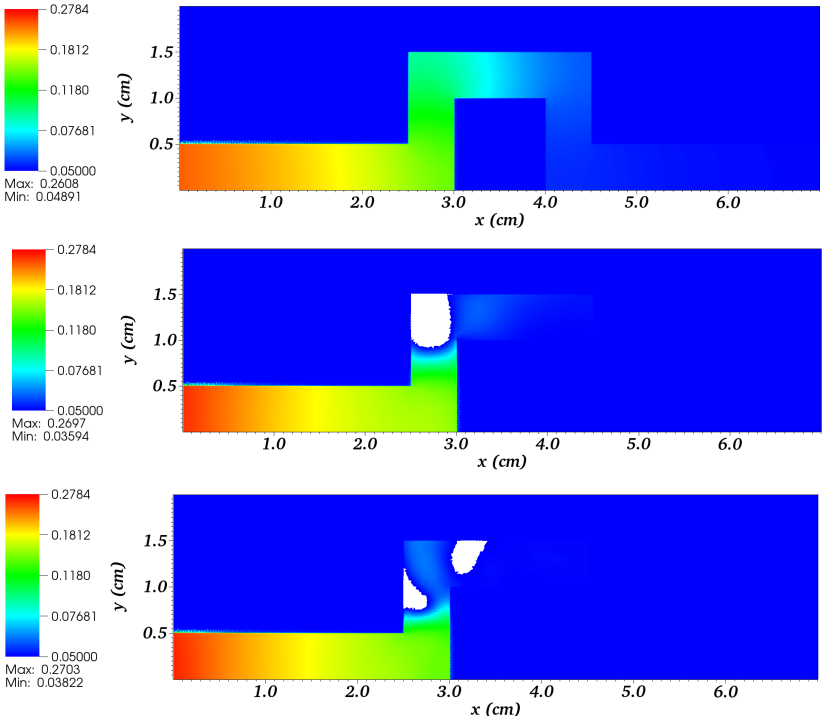

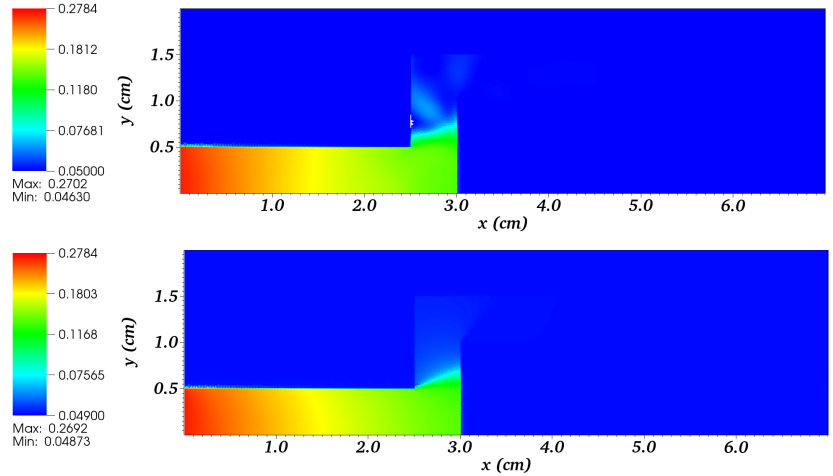

Material temperature (in keV) at $t = 3.5$ ns (from top to bottom)
left: P_1, P_3, P_5
right: P_7, P_{39}

The white regions show where the radiation density is negative.

TEXAS A&M UNIVERSITY

Locally filtered calculation

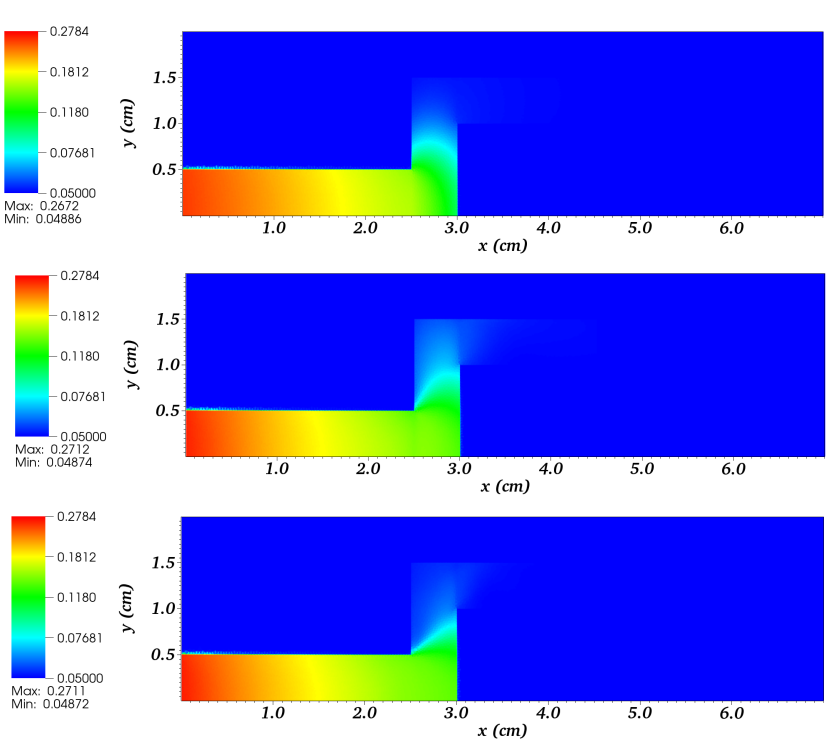

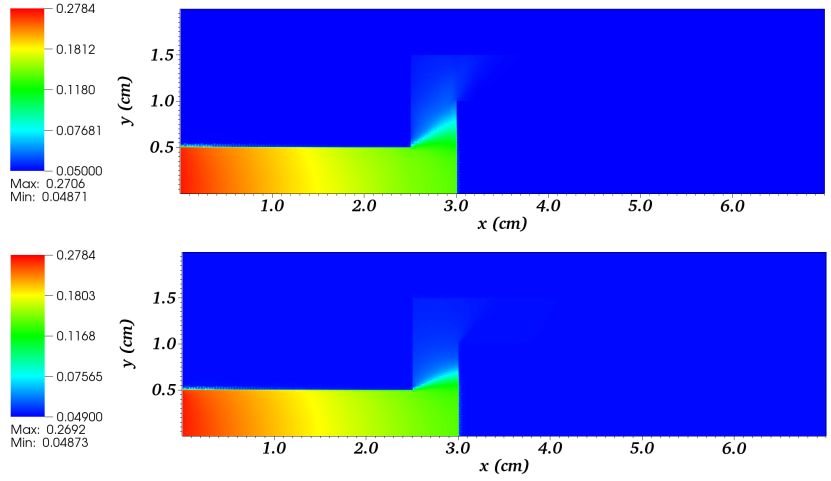

Material temperature (in keV) at $t = 3.5$ ns (from top to bottom)
left: P_1, P_3, P_5
right: P_7, P_{39}

TEXAS A&M UNIVERSITY

Section 3

Uncertainty Quantification

TEXAS A&M UNIVERSITY

The Parametric Uncertainty Quantification Problem

- Consider an output quantity from a transport calculation of the form

$$y = \int dV\, w(x, \Omega, E, t)\psi(x, \Omega, E, t).$$

- Such a quantity could be the total fission energy produced in the system or the number of neutrons leaking out of the reactor.
- We are interested in the estimation of y as a function of the input parameters: $\sigma_t(E)$, $\sigma_s(E)$.
- Because these inputs are functions that depend on energy and the particular material, there can easily be hundreds/thousands of input uncertainties
- Other possibilities are the shape of objects in the system, densities, etc.
- In may be the case that we are primarily interested in first-order or second-order sensitivities to each of the parameters, p_i:

$$y = y_0 + \sum_{i=1}^{I} \alpha_i \frac{\partial y}{\partial p_i} p_i + \sum_{i=1}^{I} \sum_{i'=1}^{I} \alpha_{ii'} \frac{\partial^2 y}{\partial p_i \partial p_i'} p_i p_{i'}.$$

TEXAS A&M UNIVERSITY

Efficient Estimation via regularization

- One could estimate these sensitivities via finite differences or via an adjoint approach. Both have downsides.
- We could interpret the equation above as a regression model and run the code with various values of p_i and estimate the coefficients.
- A greenhorn statistics student will tell you that estimating the coefficients in the regression model

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_p x_p + \epsilon,$$

  is impossible if the number of observations m is less than the number of parameters N, and not likely to be accurate until $m \gg N$.
- This a common problem in many data mining analyses, e.g. my grocery store has 1000s of potential variables that could explain what item I will buy.
- Therefore, unless we want to run a very large number of simulations, the m above, we cannot estimate all the β's.

Efficient Estimation via regularization

- There are approaches that help this issue, but don't exactly fix the problem.
  - Variable selection based on judgment is a key example,
  - Adjoint-based approaches can also help, but are tricky in non-linear, time-dependent, or multi-physics situations. These also are best for single quantity of interest (QoI) situations.
- It turns out we can get robust estimates of sensitivities when the number of simulations is smaller than the number of parameters we want to estimate.
- The reason that this could work is that in most problems many of the sensitivities are effectively zero, i.e. $\beta_i \approx 0$.
- What we need is a technique that determines which of these is zero, *based on the data* and not based on an assumption.
- The issue is that this is clearly an ill-posed problem and we need to constrain the space in which we look for a solution. This is done through regularization of the problem.

The magic of the L1 norm

- We will cast the problem in terms of an optimization problem. For the regression formulation one possibility is the problem

$$\text{Find the } \boldsymbol{\beta} \text{ that minimizes } \|\mathbf{e}\|_2 + \lambda \sum_i |\beta_i|.$$

- This approach is a regularized regression problem called lasso regression because in practice it sets some β_i's to zero and "lassos" the important variables.
- Like ordinary least squares regression it attempts to minimize the sum of the squares of the error, but it also tries to minimize the magnitude of the coefficients (the L1 norm of the vector $\boldsymbol{\beta}$).
- The L1 norm is the reason that certain β_i's are set to zero.

The magic of the L1 norm

- While there is rich literature on why these regularized optimization problems work well in the L1 norm (see for instance the work of Candes and Tao), here is a yeoman's justification of why this might be so.
- Consider the problem of estimating the coefficients in the problem

$$y = a + bx + \epsilon,$$

by minimizing

$$\sum_i \epsilon_i^2 + \left(|a|^p + |b|^p\right)^{1/p}.$$

- The curve of equal value of $\left(|a|^p + |b|^p\right)^{1/p}$ is a circle for $p = 2$ and a diamond for $p = 1$.
- The curves of equal value for $\|\epsilon\|^2$ are ellipses.
- One can show that where the diamond intersects the ellipse of minimum size will be closer to one of the axes.

The magic of the L1 norm

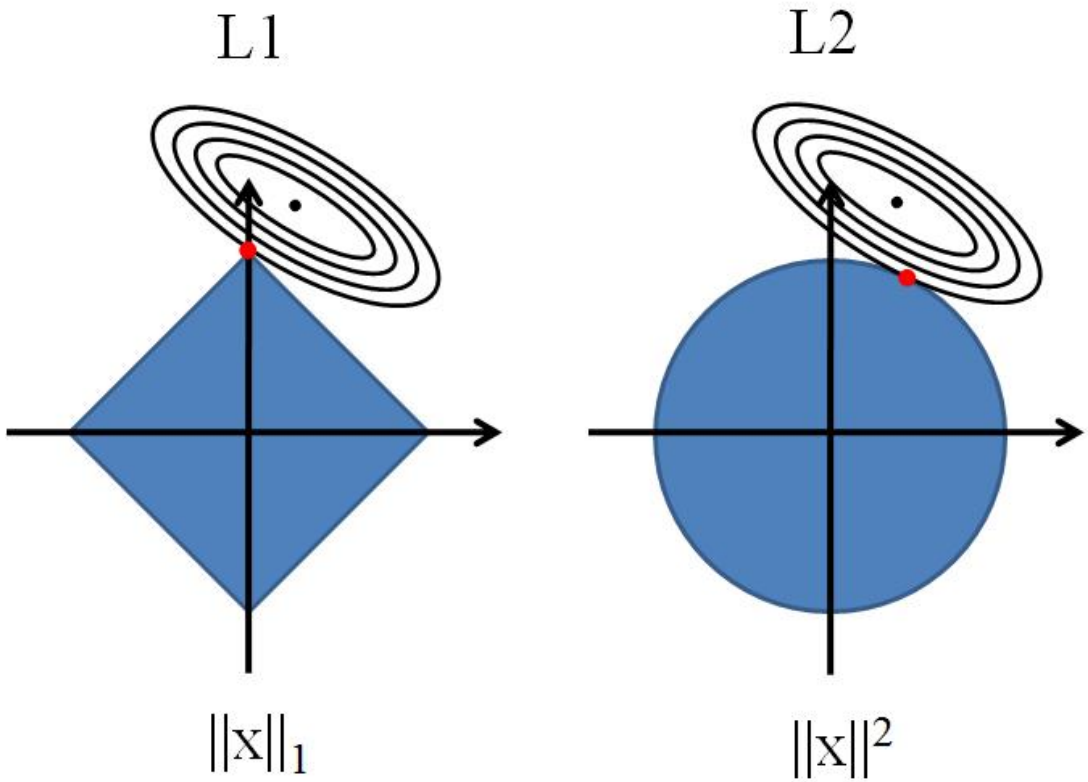

from https://tianyizhou.wordpress.com/2010/08/23/compressed-sensing-review-1-reconstruction-algorithms/

Previous UQ Work

- In the nuclear field, Watanabe et al. used L1 minimization to estimate first-order sensitivity coefficients for a pincell burnup problem with 5000 parameters. They needed 500 simulations to estimate the parameters efficiently. These results did not leverage a regression framework, which could lead to improvement.
- For climate uncertainty analysis, LLNL researchers have used lasso-type approaches to estimate polynomial chaos expansion coefficients.
- In this presentation I'll present the results of a bake-off to compare different approaches to estimate second-order sensitivity coefficients, i.e., the quadratic and interaction terms neglected in a first-order sensitivity analysis.

Non-Bayesian Regularization Regression Approaches

In these methods we explicitly change the minimization problem.

- Lasso regression (OLS plus an ℓ_1 penalty based on size of β's):

$$\beta = \underset{\beta}{\text{argmin}}\{\|\mathbf{Y} - \mathbf{X}\beta\|_2^2 + \lambda_1\|\beta\|_1\} \tag{2}$$

- Ridge regression (OLS plus an ℓ_2 penalty based on size of β's):

$$\beta = \underset{\beta}{\text{argmin}}\{\|\mathbf{Y} - \mathbf{X}\beta\|_2^2 + \lambda_2\|\beta\|_2^2\} \tag{3}$$

- Elastic net regression (Combination of Lasso and Ridge):

$$\beta = \underset{\beta}{\text{argmin}}\left\{\|\mathbf{Y} - \mathbf{X}\beta\|_2^2 + \alpha\lambda_1\|\beta\|_1 + (1-\alpha)\lambda_2\|\beta\|_2^2\right\} \tag{4}$$

- Dantzig selector (Minimize ℓ_∞ error in fit with ℓ_1 penalty on β's):

$$\beta = \underset{\beta}{\text{argmin}}\{\|\beta^T(\mathbf{Y} - \mathbf{X}\beta)\|_\infty + \lambda_1\|\beta\|_1\} \tag{5}$$

Section 4

Model description

TEXAS A&M UNIVERSITY

Problem settings

Lattice of TRIGA fuels pin modeled with MCNP

- QoI: k_{eff}

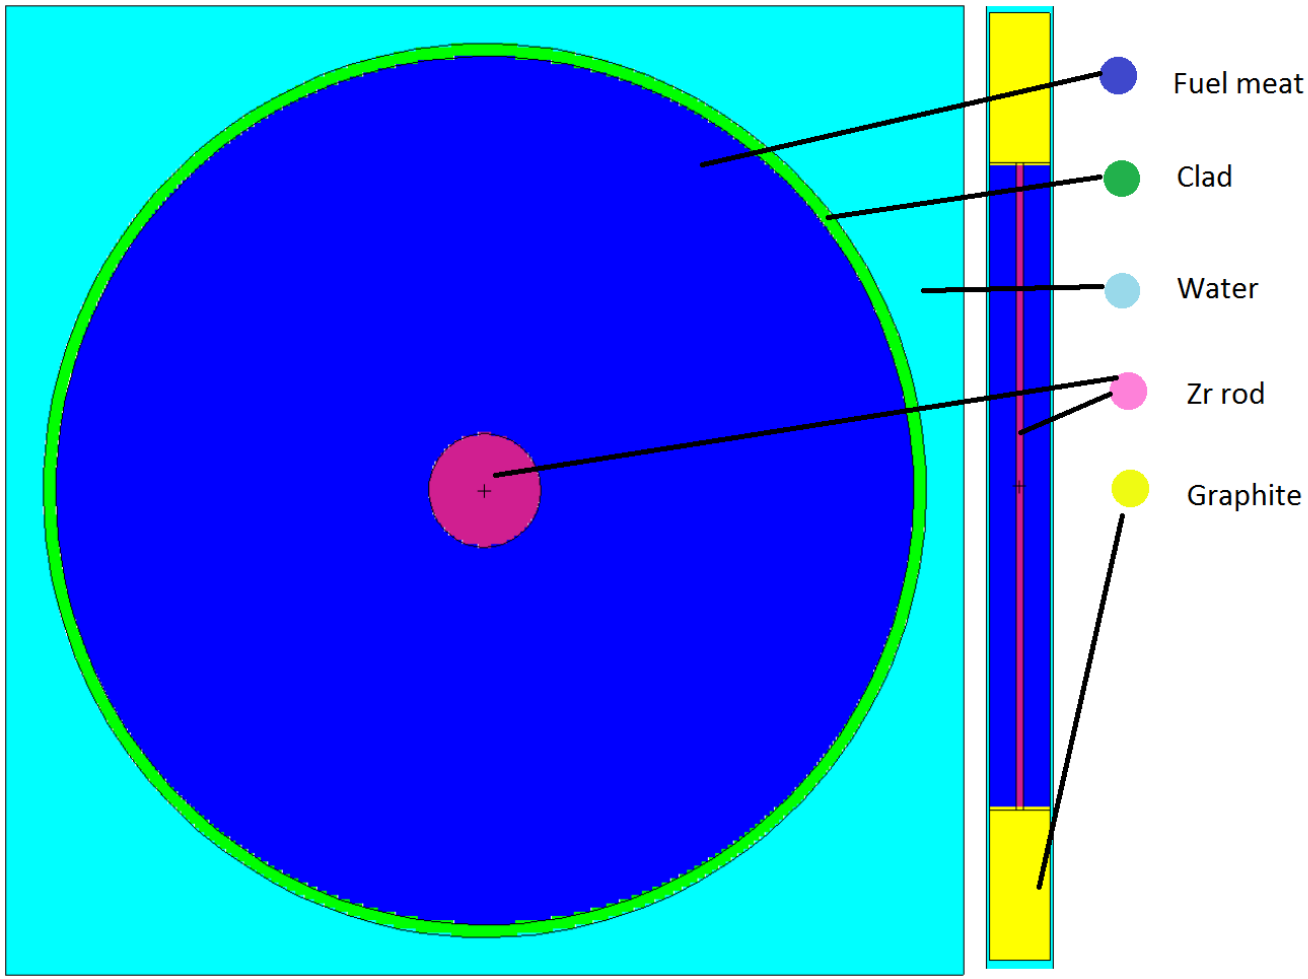

Problem descriptions

There are 299 sensitivity coefficients taken into account in this problem:

- 23 input parameters:
  - 6 geometric parameters: e.g. r-fuel (fuel radius)
  - 17 material parameters: e.g. $\rho-$Zr (Zr rod mass density)
- 253 pairwise interactions (23 choose 2)
- 23 quadratic terms

The aim is to investigate the sensitivity of the criticality to the parameters, especially the second order terms. The model is:

$$\frac{\delta k}{k} \approx \sum_{i=1}^{23} c_i \left(\frac{\delta x_i}{x_i}\right) + \sum_{i=1}^{22} \sum_{j=i+1}^{23} c_{ij} \left(\frac{\delta x_i}{x_i}\right)\left(\frac{\delta x_j}{x_j}\right) + \sum_{i=1}^{23} c_{ii} \left(\frac{\delta x_i}{x_i}\right)^2 \qquad (6)$$

where c_i, c_{ij} and $c_{ii}, i = 1, \cdots, 23, j \neq i$, are the first order, interactive and quadratic sensitivity coefficients, respectively.

Reference data

- We are going to compare reference sensitivity coefficients to the coefficients computed by various regularized regression techniques using many few code runs (cases).
- The reference coefficients are computed using 1058 cases.
  - We need 46 total simulations for the linear and quadratic parameters
  - 1012 simulations are needed for the 253 interactions (4 simulations for each)
- The goal of this research is to see if regularized regression techniques can give coefficient estimates close to the references using many fewer simulation runs than the 1058 cases.

Coefficient Estimation: Interactions (299 samples)

Blue dots are regression estimations, red lines are reference

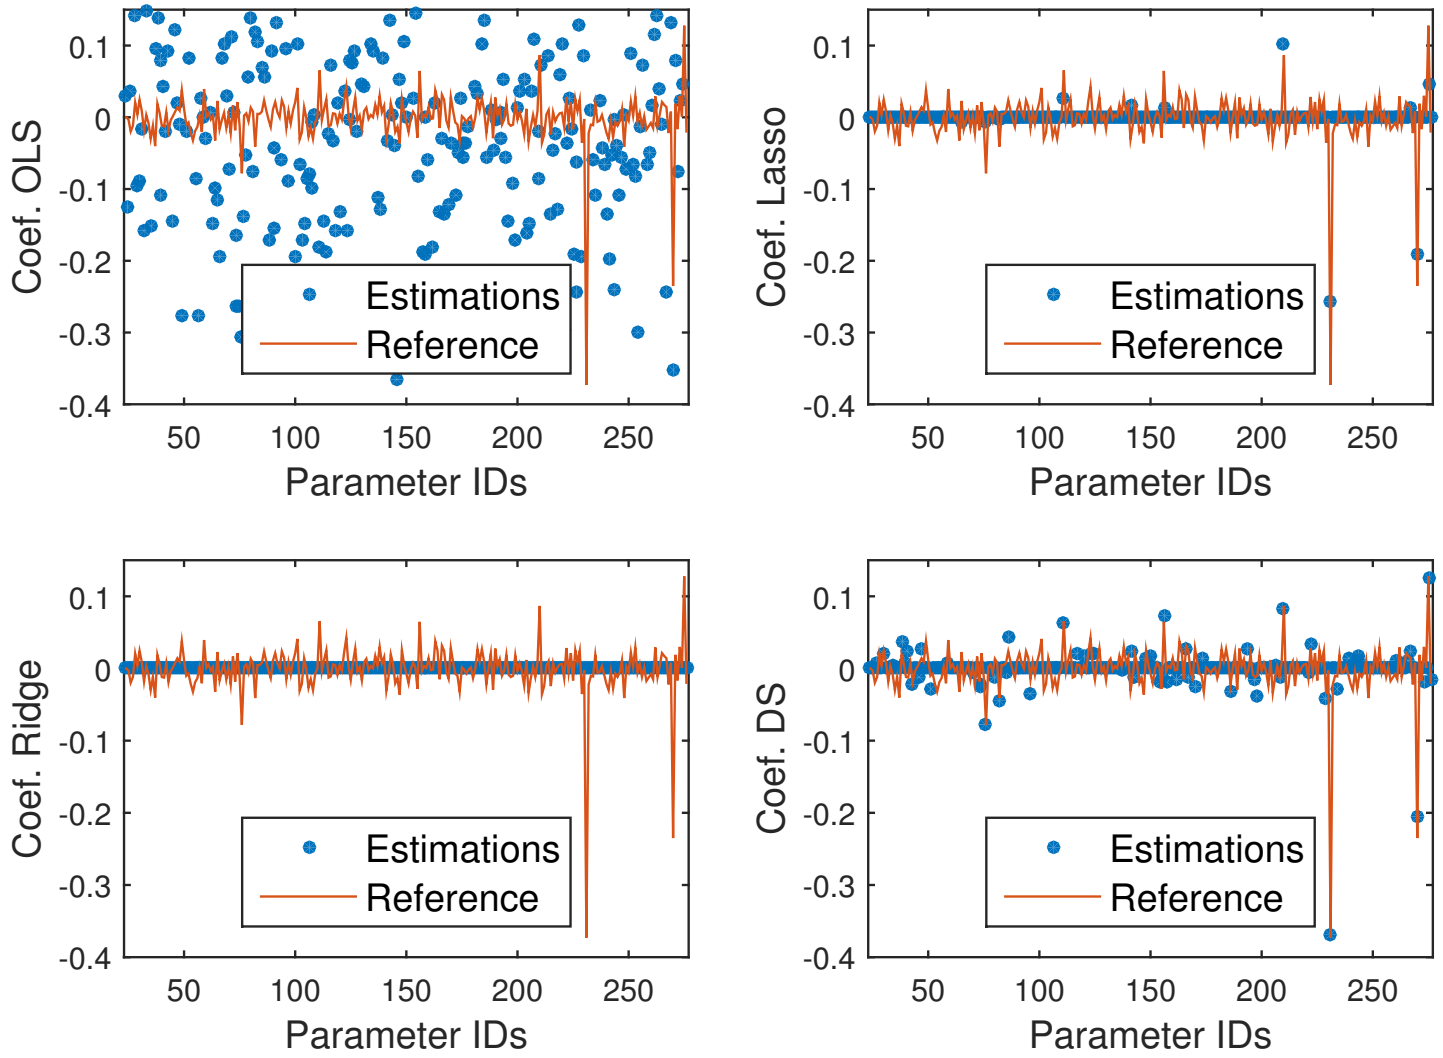

TEXAS A&M UNIVERSITY

Coefficient Estimation: Quadratic (299 samples)

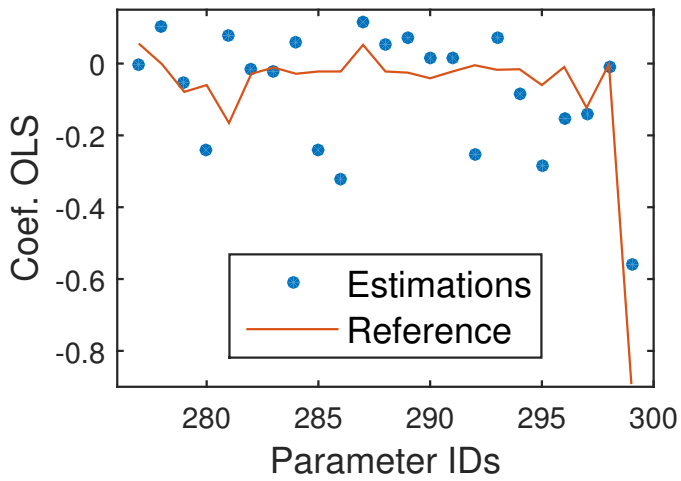

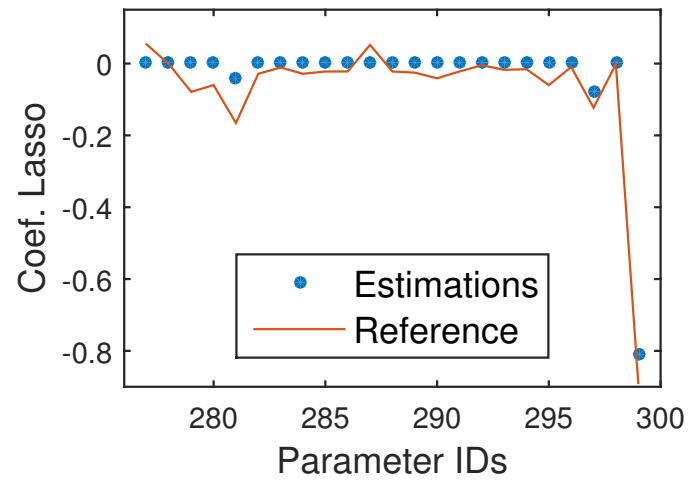

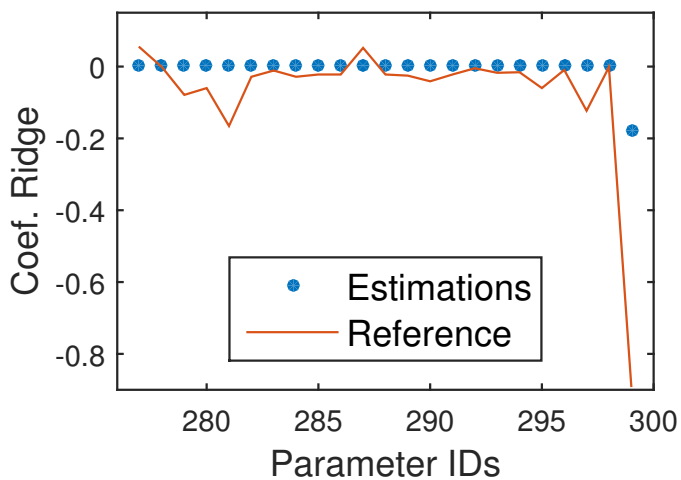

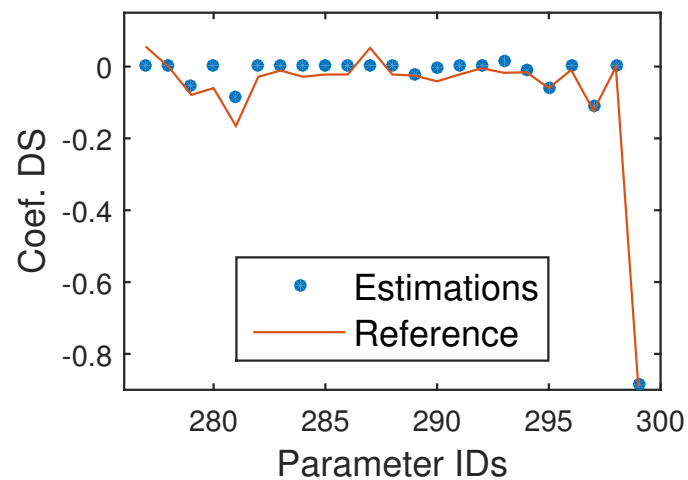

TEXAS A&M UNIVERSITY

Conclusions

- Particle transport problems are important in many fields.
- The efficient solution of these problems is important due to the rich phase space.
- Uncertainty quantification is challenging because one needs many simulations for a given system.
- Much more work to do.