SAND2007-4336C

Performance of AMG based Preconditioners for Fully-coupled Newton-Krylov Solution of Transport/Reaction

John N. Shadid

Computational Science R&D Group
Sandia National Laboratories
Albuquerque, NM

Collaborators:

Paul T. Lin
Ray Tuminaro
Marzio Sala (ETH)

Office of Science
Science for America's Future
MICS Program

ASC
DOE / NNSA

ICIAM 2007, July 16-20, 2007

Sandia National Laboratories

Outline

- **Motivation: Multiple-time-scale Nonlinear Transport-Reaction Systems**
- **Outline of Numerical Solution Algorithms**
- **Why Newton-Krylov Methods?**
    - **Multiple-time-scale Systems**
    - **Characterization of Complex Solution Spaces**
    - **Optimization**
- **Solution Algorithm Performance**
    - **Parallel and Algorithmic Scaling of DD preconditioners**
    - **Multi-level Preconditioners**
        - **1 level & 2 Level Geometric Domain Decomposition**
        - **N-level Aggressive Coarsening Graph-based Block DD/AMG**
- **Conclusions**
- **Extended List of Collaborators**

What's Next? A subjective Computational Science View

⇒Achieving Predictive Simulations of Complex Multi-physics Systems (PDEs)

What are multi-physics systems? (A multiple-time-scale perspective)

These systems are characterized by a myriad of complex, interacting, nonlinear multiple time and length scale physical mechanisms.

These mechanisms can balance to produce:

- **steady-state behavior,**
- **nearly balance to evolve a solution on a dynamical time scale that is long relative to the component time scales,**
- **or can be dominated by one, or a few processes, that drive a short dynamical time scale consistent with these dominating modes.**

e.g. Fusion Reactors (Tokomak -ITER; Pulsed - NIF & Z-pinch); Fission Reactors (GNEP); Astrophysics; Combustion; Chemical Processing; Fuel Cells; etc.

An Example – Fire simulations (Algorithm R&D Code)

Sandia National Laboratories

Multiple-time-scale systems: E.g. Methanol Pool Fire
LES-ksgs and Flamelet Combustion Model (w/ T. Smith – MPSalsa)

2D axisymmetric Simulation

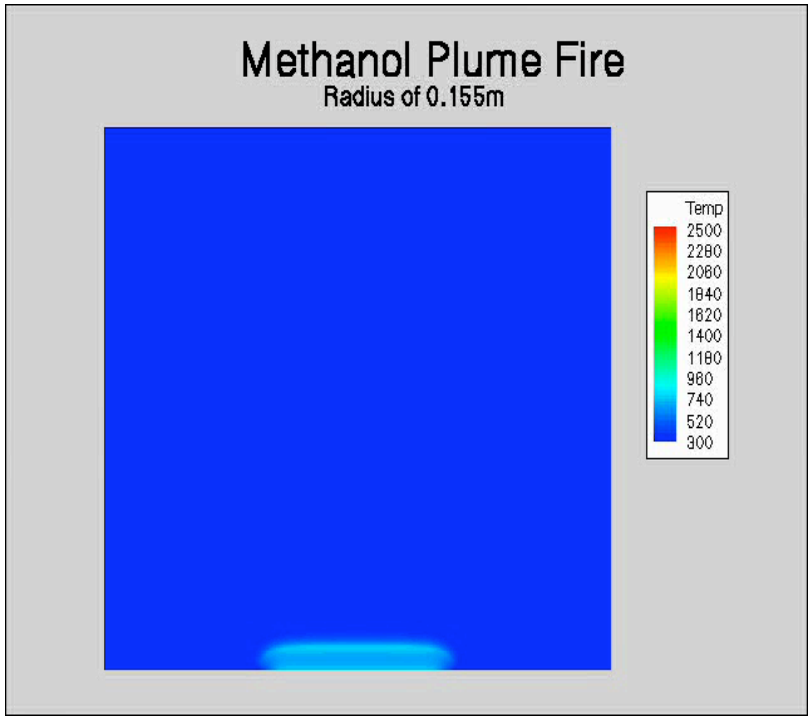

Full 3D Simulation (note: non-axisymmetric mode)

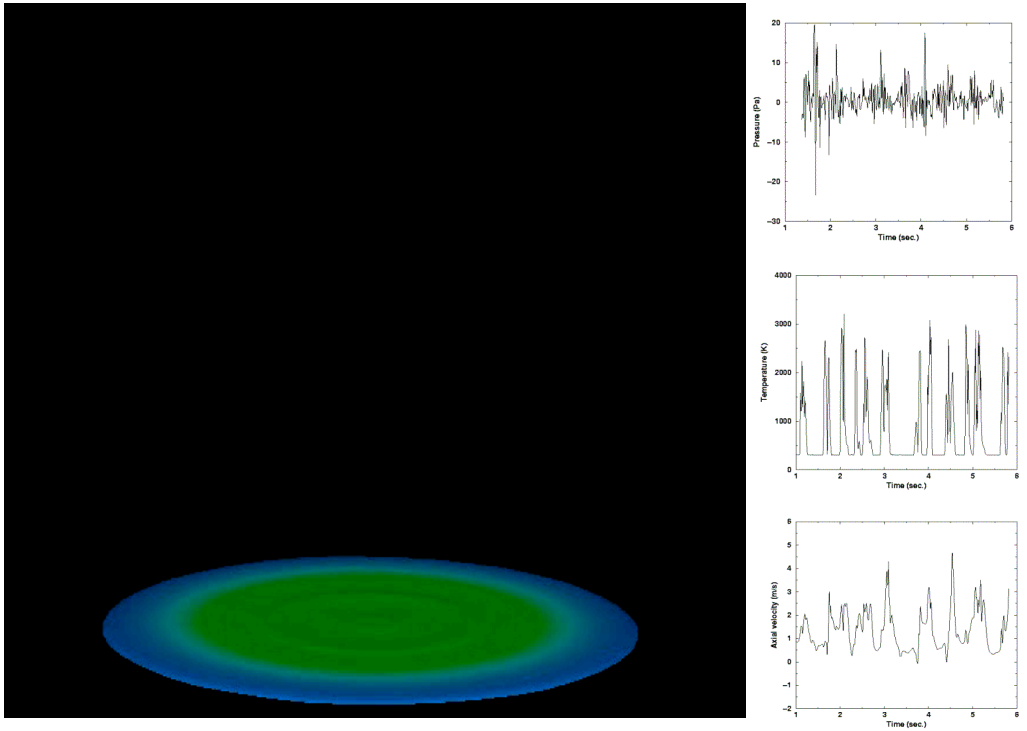

Physical time scales (sec.):

- **Chemical kinetics: 10^{-12} to 10^{-5}**
- **Momentum diffusion: 10^{-6}**
- **Heat conduction: 10^{-6}**
- **Convection: 10^{-3} to 10^{-1}**
- **Buoyancy (puffing freq. = 2.8Hz): 10^{-1} to 10^{0}**
- **Meandering mode: 10^{0}**

Numerical Solution of Multiple-time-scale Multiphysics Systems

A Perspective:

Historically, linearization, semi-implicit, operator split methods and decoupled solution strategies were devised out of necessity in a time when limitations in computer memory and CPU power were acute.

The resulting numerical stability, accuracy and appropriate time step controls are only heuristically understood, in most cases. Solution of these complex systems can be fragile and exhibit non-intuitive instabilities or they can be stable but very inaccurate.

We believe that with the recent significant increase in computing resources and advances in numerical methods our earlier choices should be critically evaluated.

We need to pursue new approaches that include robust, accurate, scalable, efficient and predictive simulation technologies for complex multi-physics systems.

Sandia National Laboratories

<u>Transport/Reaction System:</u> A very broad range of scientific and engineering applications require the detailed computational analysis of strongly coupled continuum transport and non-equilibrium reaction physics with multiple time and length scales.

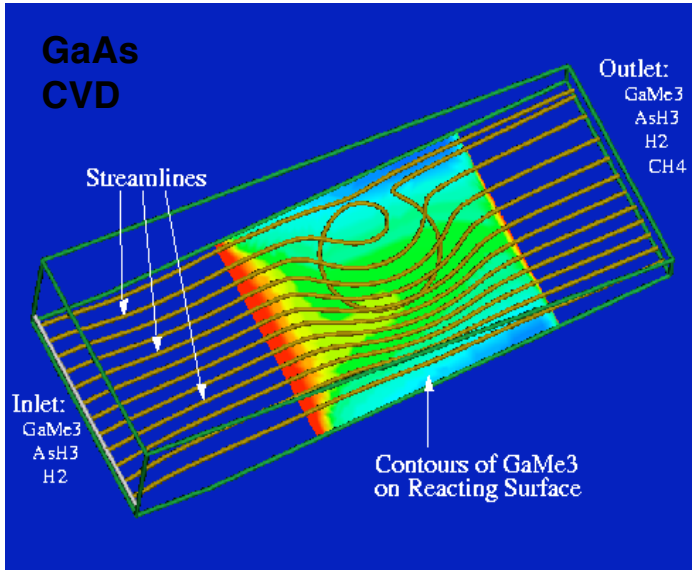

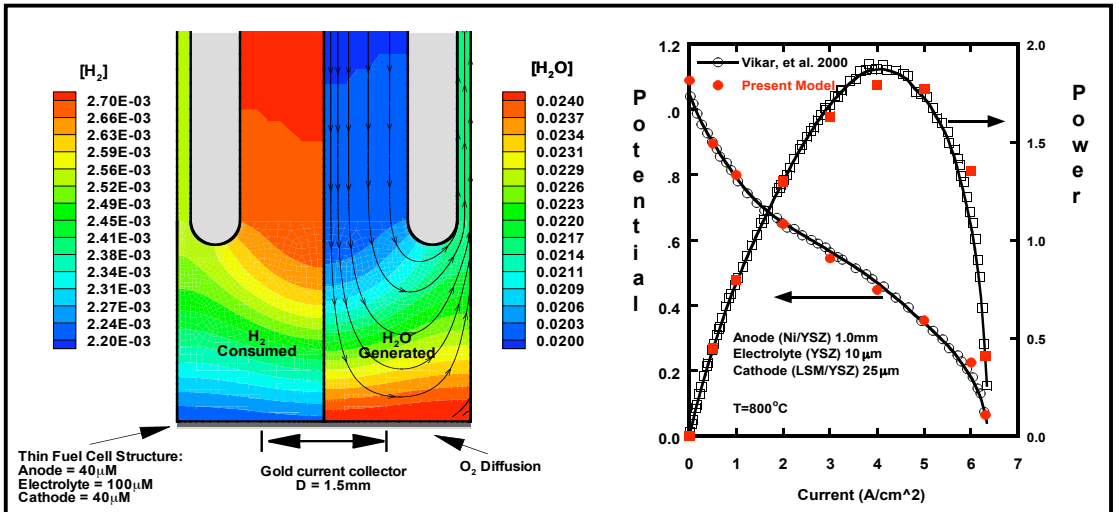

Simulation of Experimental Solid Oxide H_2 Fuel Cell

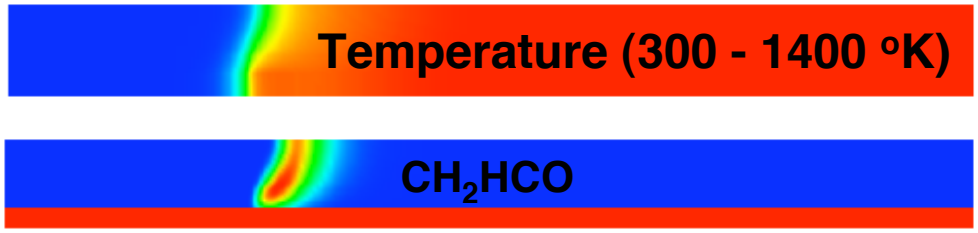

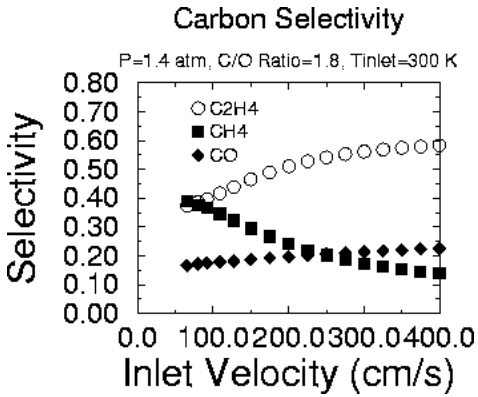

Partial Catalytic Oxidation of Ethane on PT
22 gas-phase species, 77 reactions
17 surface-phase species, 35 reactions

MPSalsa Simulations

Sandia National Laboratories

Summary of PDE Residual Equations for Transport / Reaction Systems; Strong Form Seek $\mathbf{x} \ni \mathbf{R(x)} = \mathbf{0}$

Governing Equation	PDE Residual
Momentum	$\mathbf{R_m} = \frac{\partial(\rho\mathbf{u})}{\partial t} + \nabla\bullet(\rho\mathbf{u}\otimes\mathbf{u} - \mathbf{T}) - \rho\mathbf{g};\quad \mathbf{T} = -\left(P + \frac{2}{3}\nabla\bullet\mathbf{u}\right)\mathbf{I} + \mu\left[\nabla\mathbf{u} + \nabla\mathbf{u}^T\right]$
Total Mass	$R_p = \frac{\partial\rho}{\partial t} + \nabla\bullet(\rho\mathbf{u})$
Thermal Energy	$R_T = \hat{C}_P\left[\frac{\partial(\rho T)}{\partial t} + \nabla\bullet(\rho\mathbf{u}T + \mathbf{q})\right] - \Theta - \dot{Q} + \sum_{k=1}^{N}\mathbf{j_k}\bullet\hat{C}_{P,k}\nabla T - \sum_{k=1}^{N}h_k W_k\dot{\omega}_k$
Species Mass Fraction	$R_{Y_k} = \frac{\partial(\rho Y_k)}{\partial t} + \nabla\bullet(\mathbf{u}Y_k + \mathbf{j}_k) - W_k\dot{\omega}_k;\qquad k = 1,2,\ldots,N-1\ ;\ \sum_{k=1}^{N}Y_k = 1$

General Case a Strongly Coupled, Multiple Time Scale, Nonlinear, Nonsymmetric Indefinite System

Sandia National Laboratories

Summary of the Computational Formulation and Numerical Methods

- Time Integration: 1st and 2nd order fully-implicit; adaptive error control e.g. BDF2
- FE Spatial Discretization: Stabilized FE (Hughes et. al.,) 2D/3D linear and quadratic unstructured FE.
- Parallel Formulation: Nodal based (**Chaco/Zoltan**), distributed fully summed matrices with unstructured communication (**Trilinos/AztecOO**)
- Nonlinear Solver: Inexact Newton method; adaptive convergence criteria and backtracking (**Trilinos/NOX**)
- Bifurcation, Stability Analysis and Optimization: Continuation/bifurcation library (**Trilinos/LOCA**), eigensystem analysis (**ARPACK, Trilinos/Anasazi**), multi-parameter optimization (**Dakota, Moocho**)
- Linear Solvers: Preconditioned DD and multi-level Krylov methods (**Trilinos/AztecOO/ML**)

Solver Software: software.sandia.gov/trilinos

Sandia National Laboratories

Why Newton-Krylov Methods?

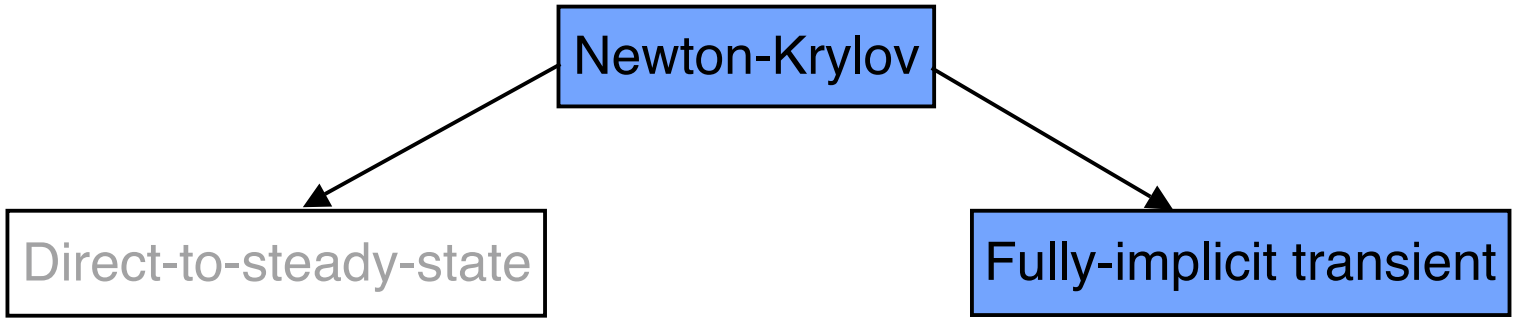

$$\mathbf{F}(\dot{\mathbf{x}}, \mathbf{x}, \lambda_1, \lambda_2, \lambda_3, ..) = \mathbf{0}$$

e.g.

$$\left.\frac{\partial c}{\partial t}\right|^{n+1} + \nabla \bullet \left(\left[\rho c \mathbf{u} \right]^{n+1} \right) - \nabla \bullet \left[D^{n+1} \nabla c^{n+1} \right] + S_c^{n+1} = 0$$

Stability and Accuracy Properties

- Stable (stiff systems)
- High order methods
- Variable order techniques
- Local and global error control possible
- Can be stable and accurate run at the dynamical time-scale of interest in multiple-time-scale systems

Sandia National Laboratories

Operator Splitting Methods can Sometimes Destroy a Critical Balance Present in the Coupled Physics. (Brusselator)

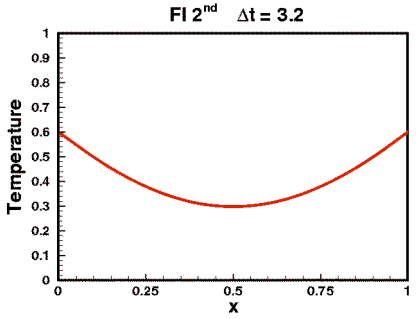

FI

$\Delta t / T \sim 1/3$

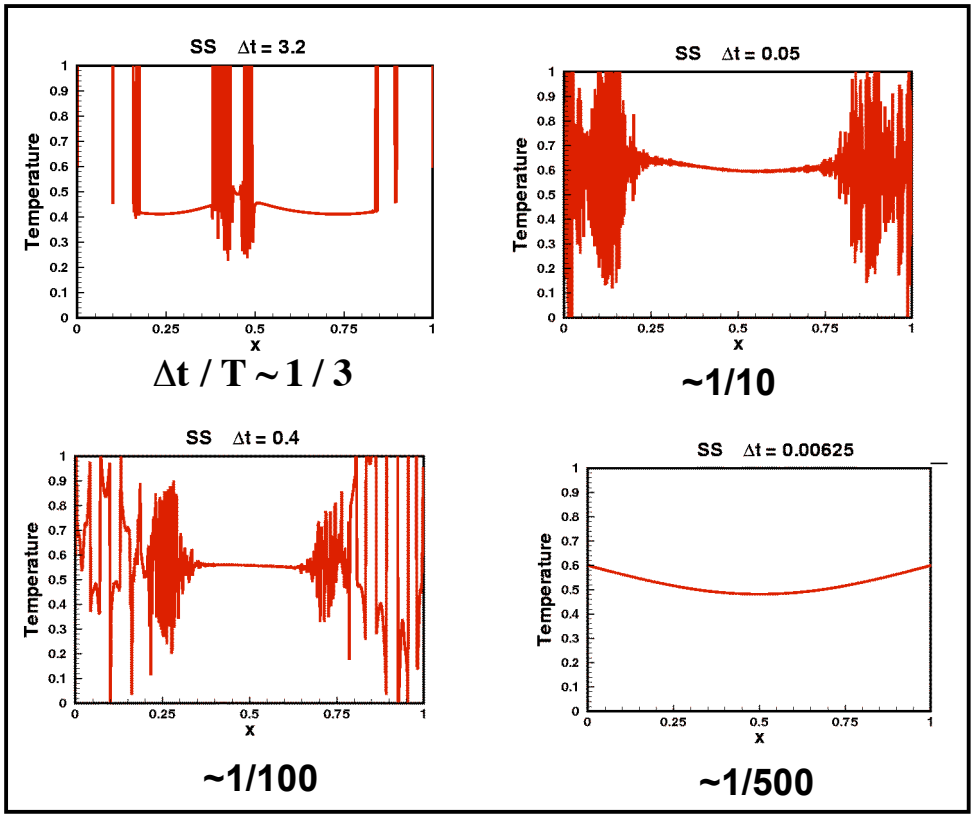

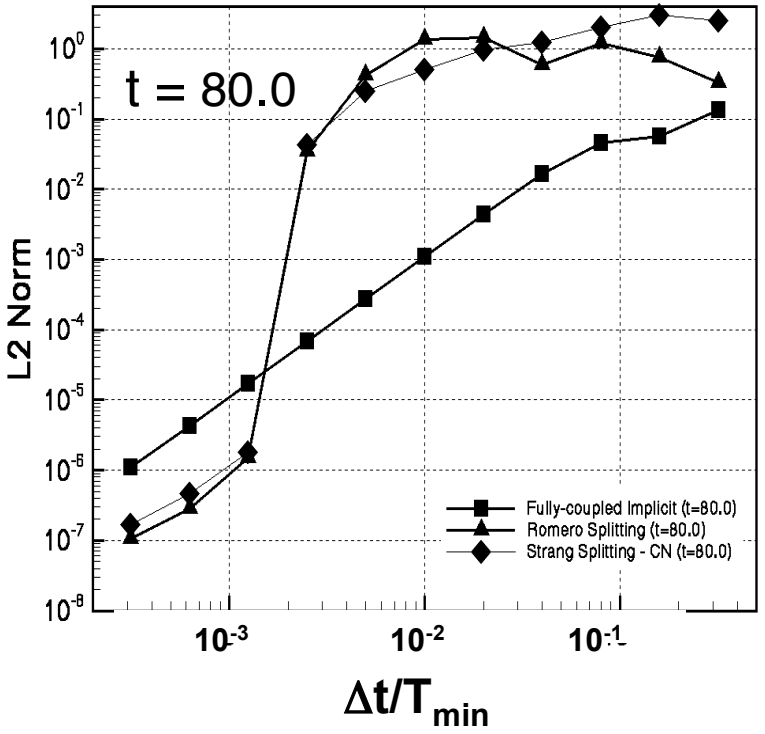

Multiple time scales:
Ropp, Shadid, JCP 2004, 2005
Ober, Shadid JCP 2004

Sandia National Laboratories

Why Newton-Krylov Methods?

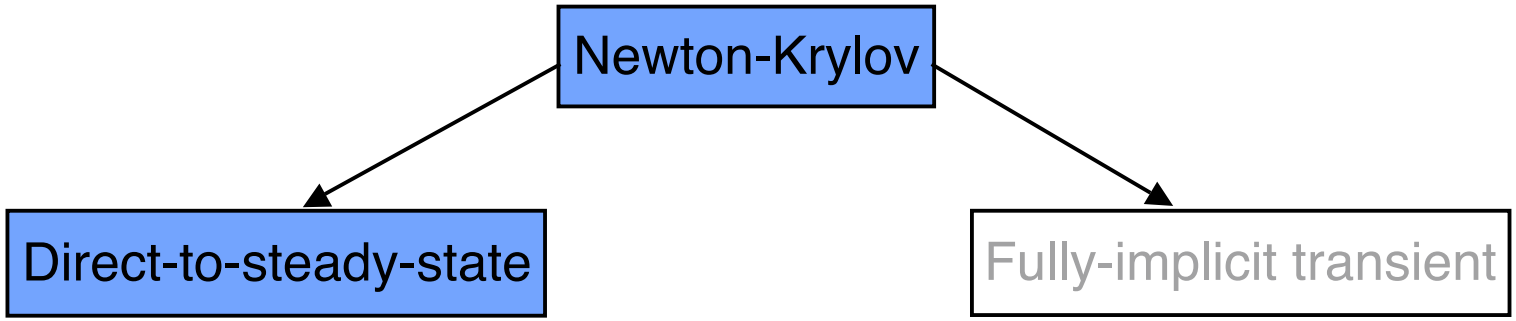

Convergence properties

- Strongly coupled multi-physics often requires a strongly coupled nonlinear solver
- Quadratic convergence near solutions (backtracking, adaptive convergence criteria)
- Often only require a few iterations to converge, if close to solution, independent of problem size

$$\mathbf{F}(\mathbf{x}, \lambda_1, \lambda_2, \lambda_3, ..) = \mathbf{0}$$

Inexact Newton-Krylov

Solve $\mathbf{J}\mathbf{p}_k = -\mathbf{F}(\mathbf{x}_k)$; until $\dfrac{\left\|\mathbf{J}\mathbf{p}_k + \mathbf{F}(\mathbf{x}_k)\right\|}{\left\|\mathbf{F}(\mathbf{x}_k)\right\|} \le \eta_k$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Theta\mathbf{p}_k$$

Jacobian Free N-K Variant

$$\mathbf{M}\mathbf{p}_k = \mathbf{v}$$

$$\mathbf{J}\mathbf{p}_k = \frac{\mathbf{F}(\mathbf{x} + \delta\mathbf{p}_k) - \mathbf{F}(\mathbf{x})}{\delta} \text{ ; or by AD}$$

See e.g. Knoll & Keyes, JCP 2004

Why Newton-Krylov Methods?

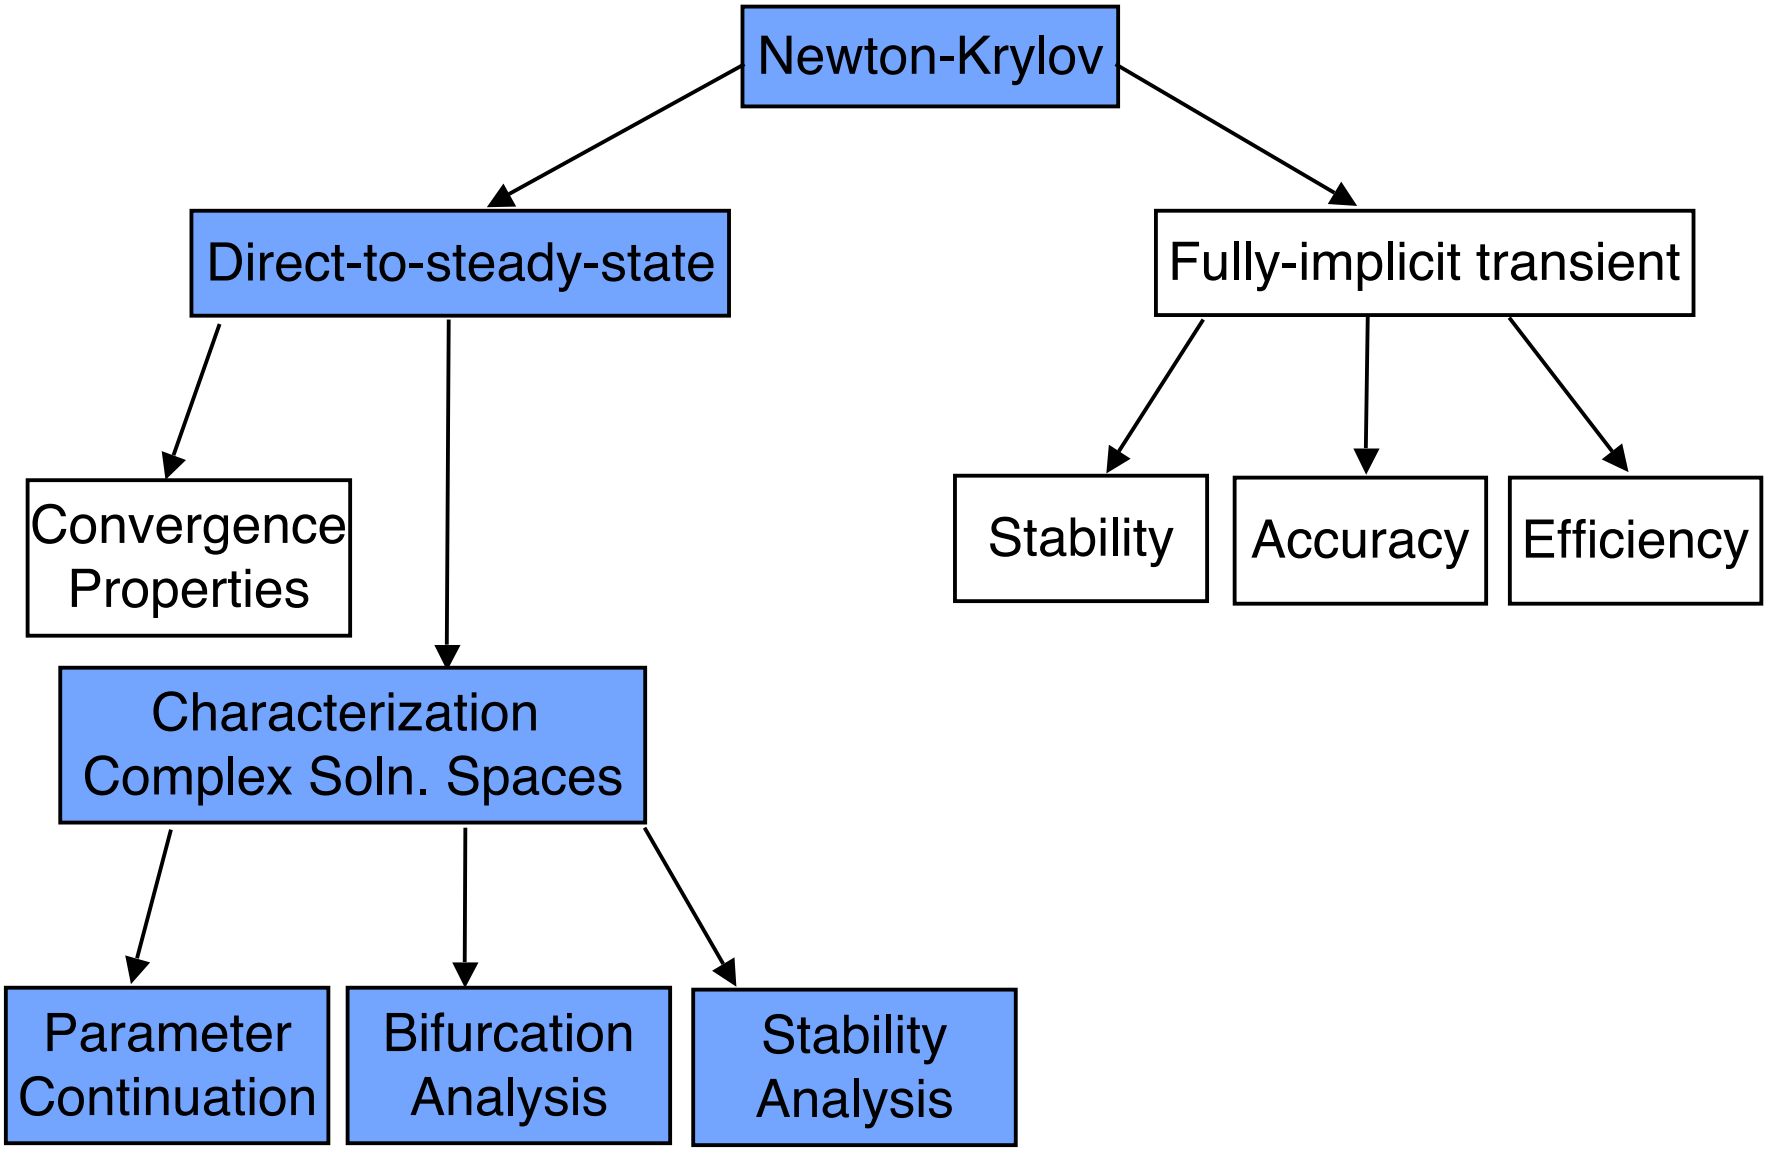

Sandia National Laboratories

Bifurcation Analysis of a Steady Reacting H_2, O_2,, Ar, Opposed Flow Jet Reactor

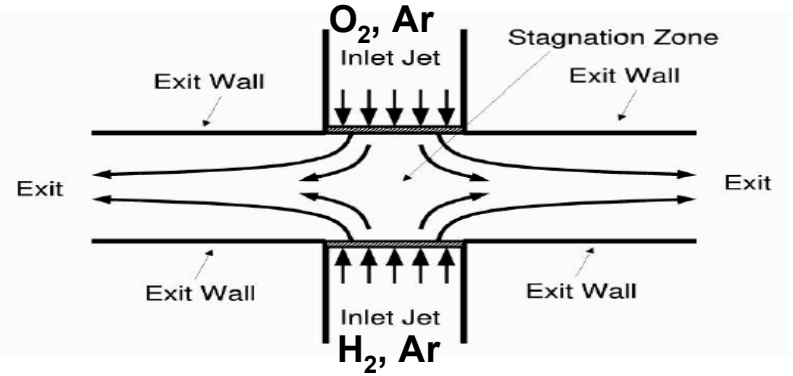

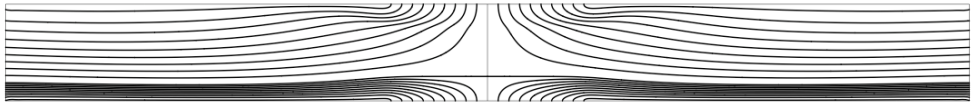

Streamlines

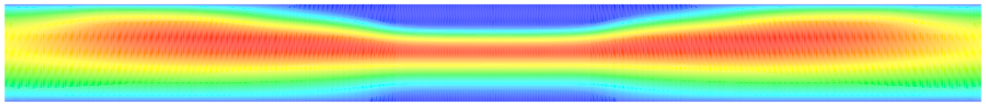

Temperature (Min. 300°K, Max 2727°K)

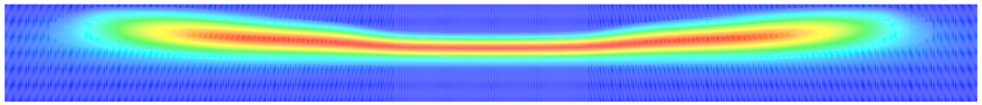

OH (Min. 0.0, Max 0.177)

70 steady state reacting flow solves (10 species, 19 reactions)

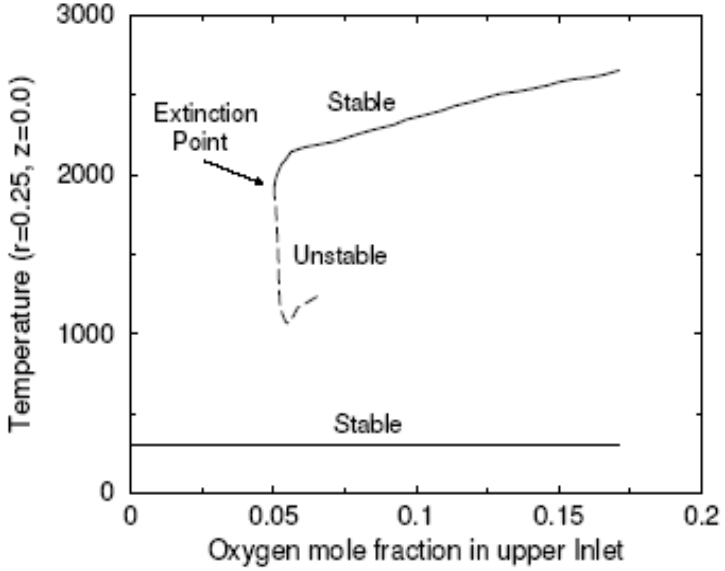

Approx. Time scales (sec.):

- Chemical kinetics: 10^{-12} to 10^{-4}
- Momentum diffusion: 10^{-6}
- Heat conduction: 10^{-6}
- Mass diffusion: 10^{-5} to 10^{-4}
- Convection: 10^{-5} to 10^{-4}
- Diffusion flame dynamics: ∞ (steady)

(w/ Pawlowski, Salinger – MPSalsa)

Sandia National Laboratories

Why Newton-Krylov Methods?

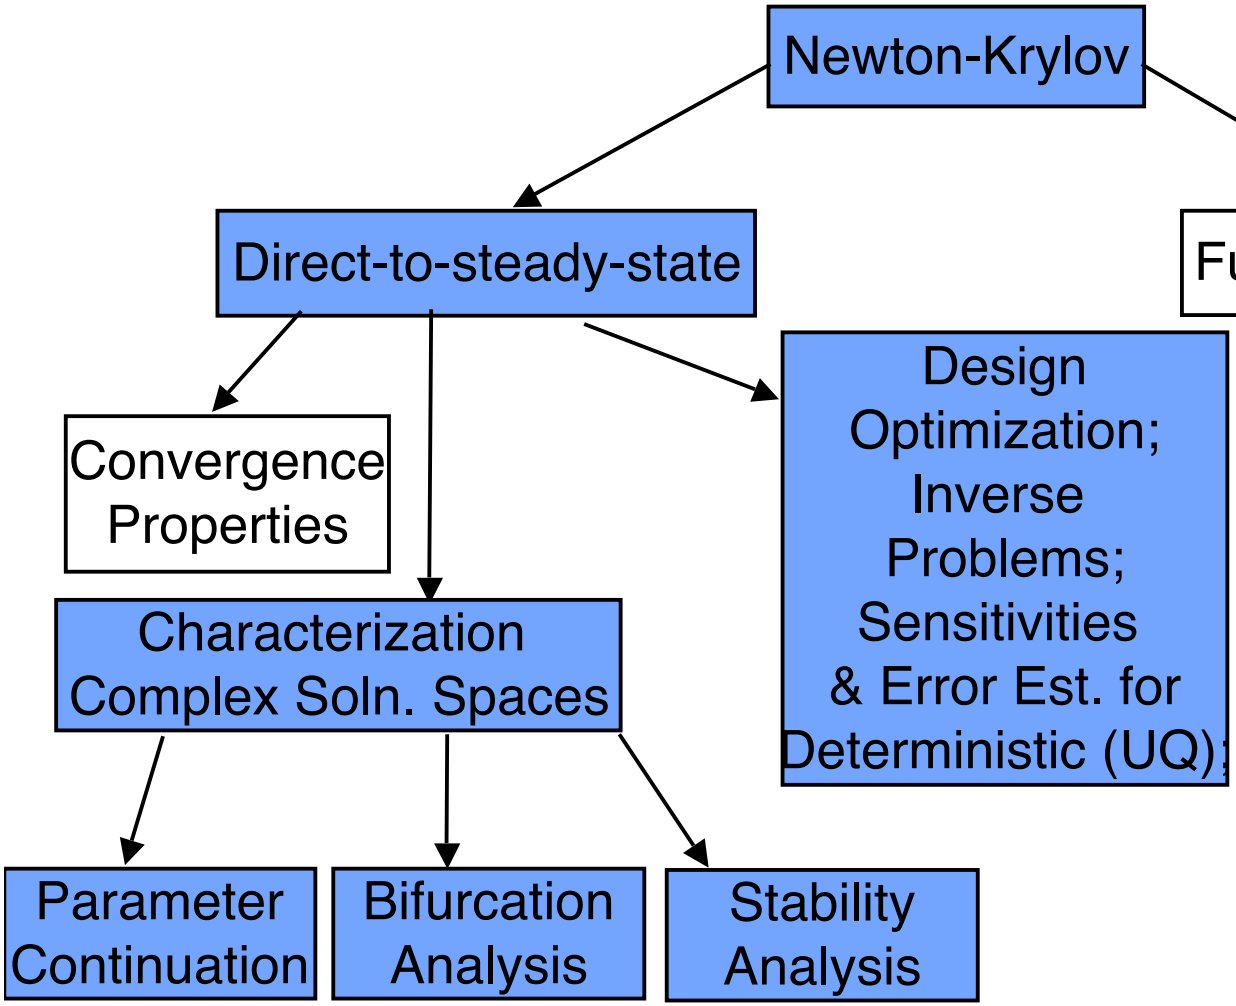

UQ Probabilistic approaches:

- **Sampling (e.g quadrature, MC, etc.),**
- **Direct (e.g Polynomial Chaos)**

Sandia National Laboratories

PDE Constrained Optimization of Poly-Silicon CVD Reactor
Unstructured FE Reacting Flow MPSalsa code

Poly-Silicon Epitaxy from Trichlorosilane in Hydrogen Carrier;

3D (u,v,w,P,T)
3 chemical species
1.2M unknowns

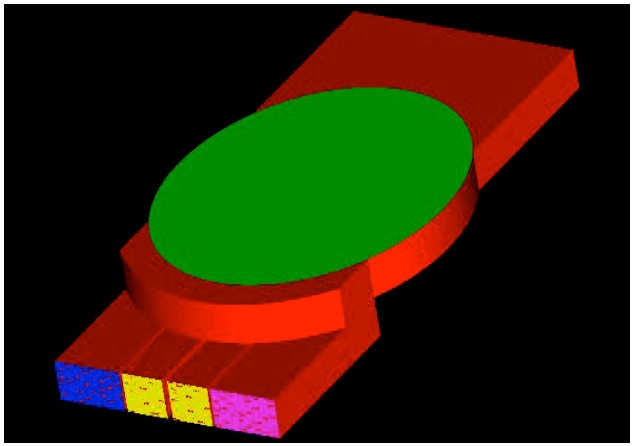

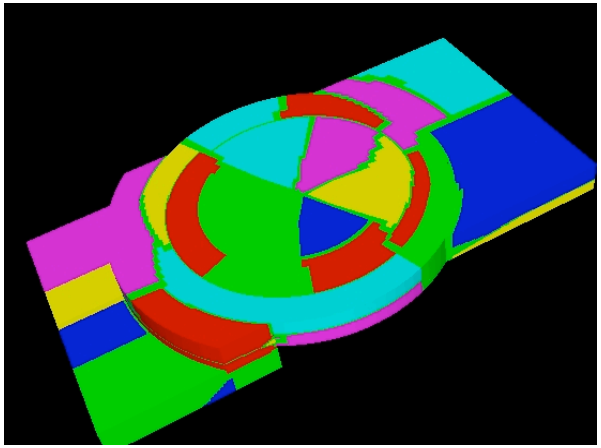

3D

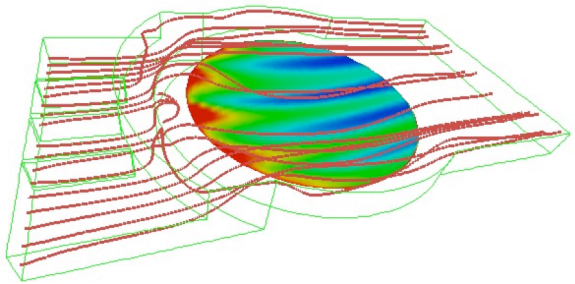

2D

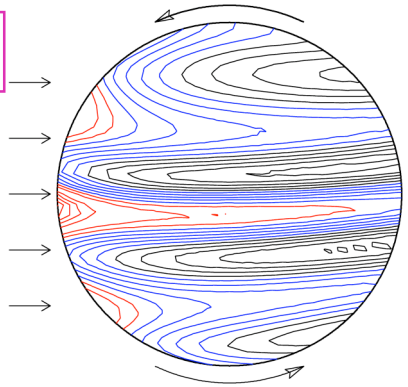

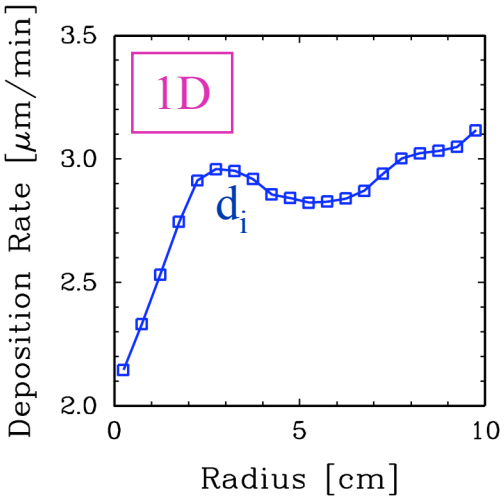

0D Objective Function:

$$f = \frac{1}{2}\sum_{\text{radii}} \left(d_i / d_{ave} - 1 \right)^2$$

Sandia National Laboratories

PDE Constrained Optimization of Poly-Silicon CVD Reactor

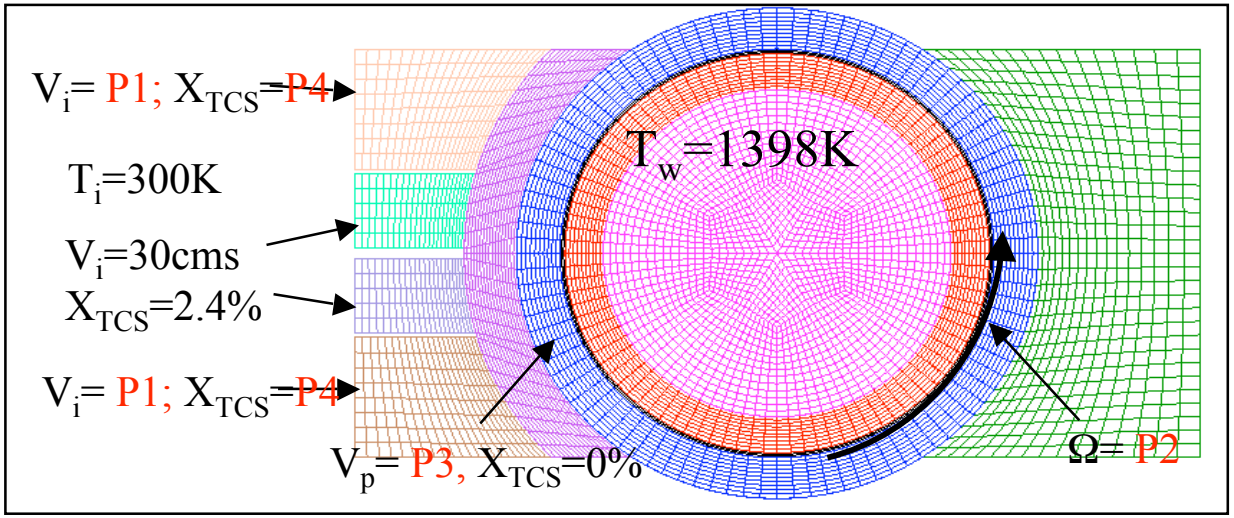

PDE Constrained Optimization:

Minimize: $f(\mathbf{x},\mathbf{p})$
such that: $\mathbf{F}(\mathbf{x},\mathbf{p})=\mathbf{0}$

Use Newton's Method solve KKT system

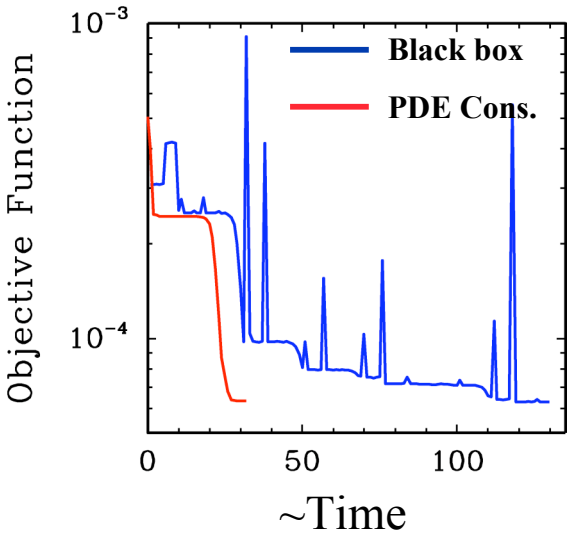

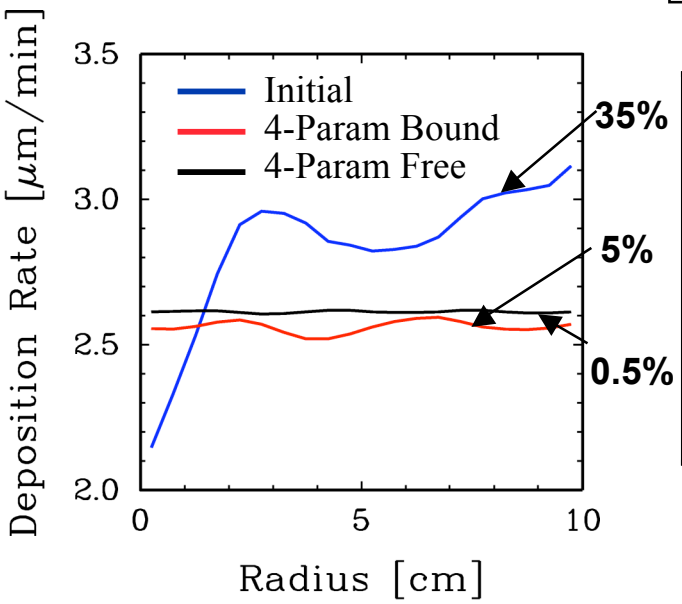

Unks	Procs	Time (hrs.)
1.2	48	6.2 (3GHz Cluster)
4.8M	128	~ 6 (Red Storm: XT3)
38M	1024	~ 7 (Red Storm: XT3)

W/Pawlowski, Salinger, van Bloemen Waanders, Bartlett, Lin - SNL

Sandia National Laboratories

Why Newton-Krylov Methods?

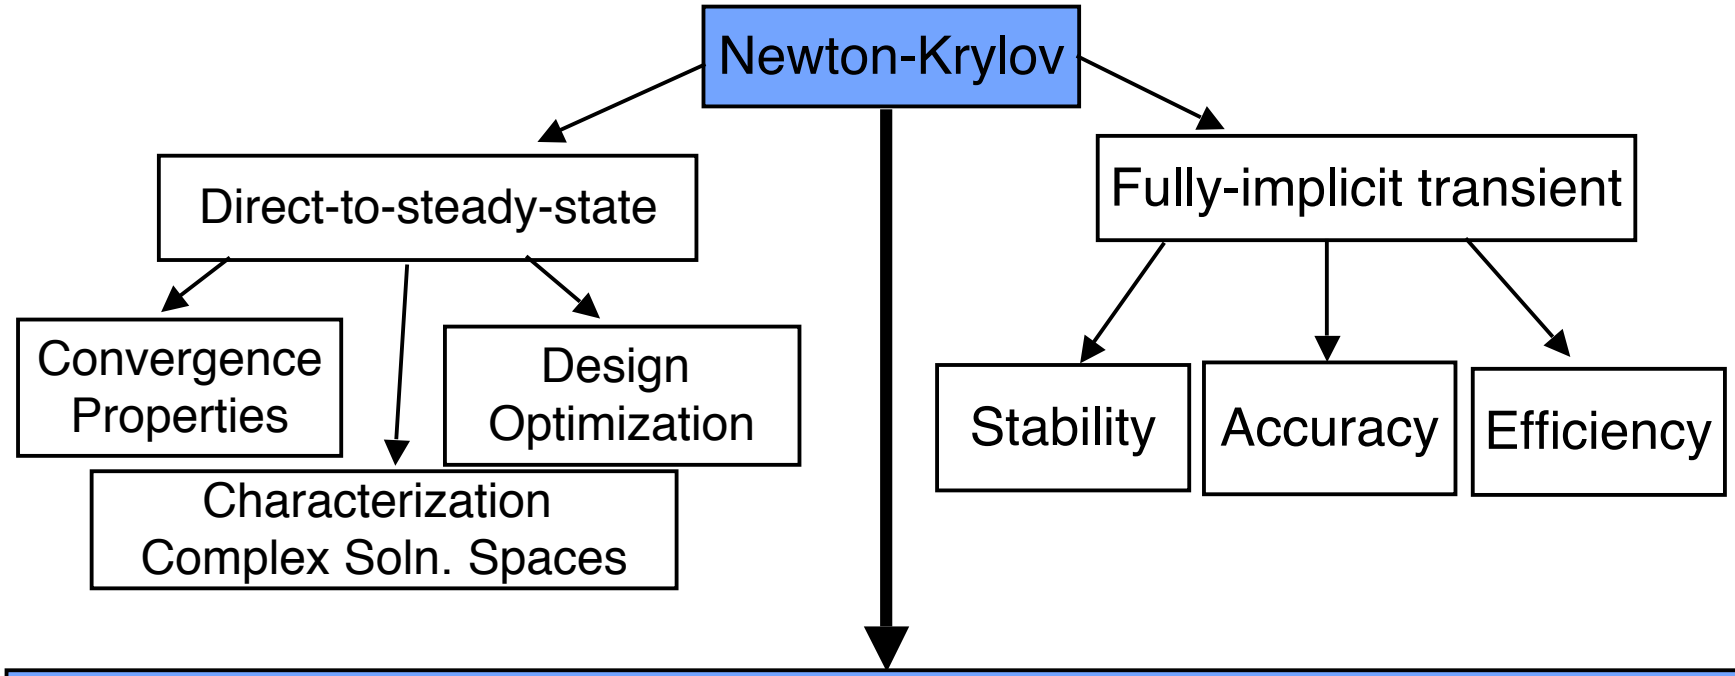

Very Large Problems -> Parallel Iterative Solution of Sub-problems

Krylov Methods - Robust, Scalable and Efficient Parallel Preconditioners

- Approximate Block Factorizations
- Physics-based Preconditioners
- Multi-level solvers for systems and scalar equations

Sandia National Laboratories

Parallel Scaled Efficiency of 1- level DD Preconditioners: Steady Reacting H_2, O_2 Opposed Jet Reactor

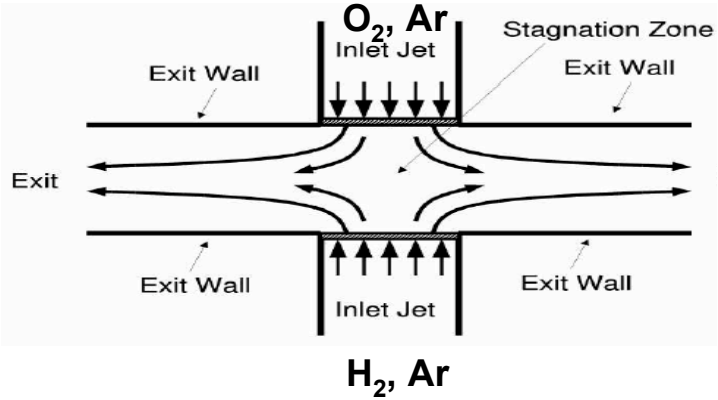

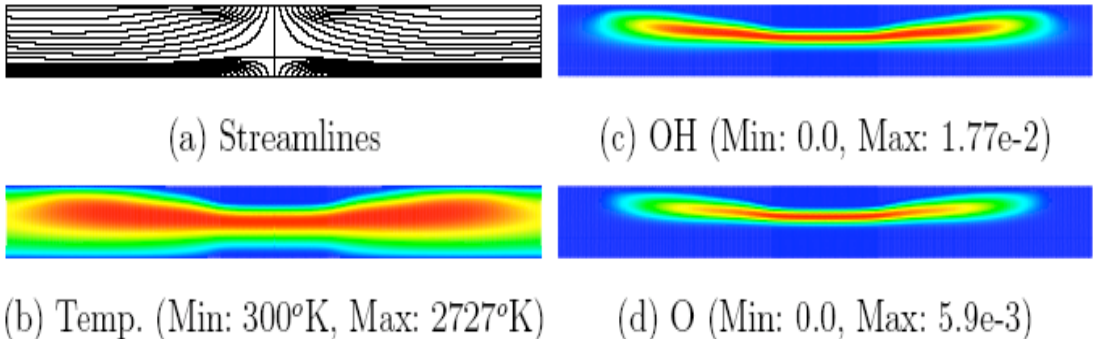

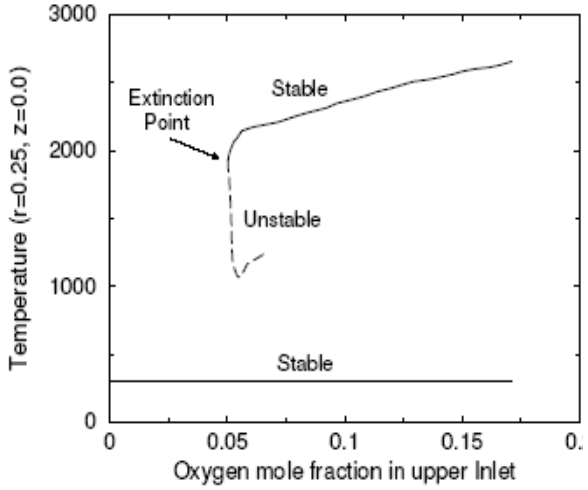

Num Procs	rf	Num. unknowns	avg time /matrix fill (sec)	scaled eff.	avg time /linear iter (sec)	scaled eff.
16	0	199,374	13.94	—	0.5219	—
64	1	790,734	13.86	1.01	0.5286	0.99
256	2	3,149,454	14.06	0.99	0.5363	0.97
1024	3	12,570,894	13.99	1.00	0.5369	0.97

Table 8
Scaled Efficiency of hydrogen/oxygen diffusion flame simulation with DD-ILU TFQMR. **(10 species, 19 reactions)**

Sandia National Laboratories

ML library: Multilevel Preconditioners

(R. Tuminaro, M. Sala)

2-level and N-level Aggressive Coarsening Graph-based Block AMG

- Aggregation is used to produce a coarse operator
  - **Create graph where vertices are block nonzeros in matrix A_k**
  - **Edge between vertices i and j included if block $B_k(i,j)$ contains nonzeros**
  - **Decompose graph into aggregates (subgraphs) [Metis/ParMetis]**
- Construction of simple restriction/interpolation operators (e.g. piecewise constants on agg.)
- Construction of A_{k-1} as $A_{k-1} = R_{k-1} A_k I_{k-1}$

- Nonsmoothed aggregation
- Domain decomposition smoothers (sub-domain GS and ILU)
- Coarse grid solver can use fewer processors than for fine mesh solve (direct/approximate/iterative)

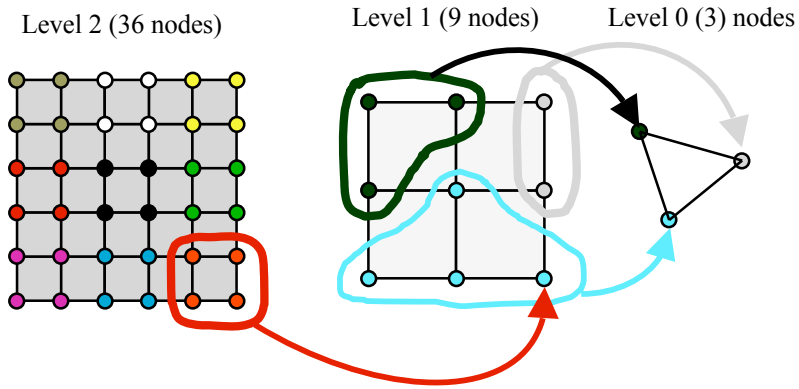

Visualization of effect of partition of matrix graph on mesh

Aggregation based Multigrid:
- **Vanek, Mandel, Brezina, 1996**
- **Vanek, Brezina, Mandel, 2001**

Aggregation used in DD:
- **Paglieri, Scheinine, Formaggia, Quateroni, 1997**
- **Jenkins, Kelley, Miller, Kees, 2000**
- **Toselli, Lasser, 2000**
- **Sala, Formaggia, 2001**

Multilevel Preconditioner Scaling Study: 3D Thermal Buoyancy Driven Convection

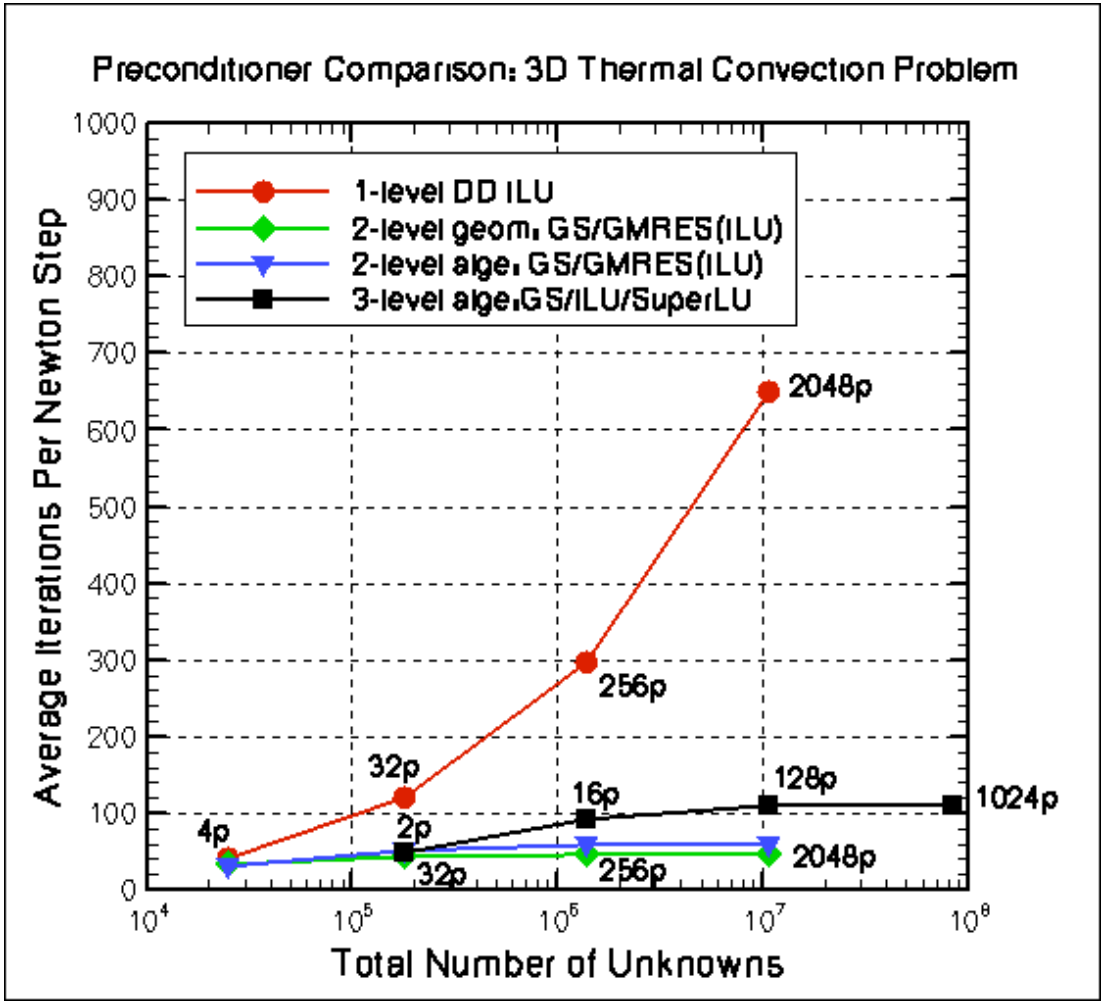

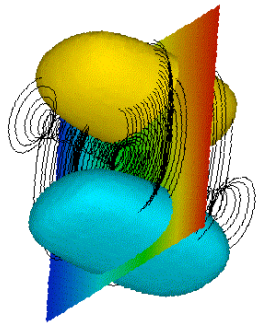

Sandia National Laboratories

Comparison of 1-level with 2-level geometric & algebraic
2D & 3D Thermal Convection Problem

proc	fine grid unknowns	1 - level Method Ilu DD		coarse unknowns		2-level: ilu-superlu			
						geometric		algebraic	
		avg its per Newt step	time (sec)	geometric	algebraic	avg its per Newt step	time (sec)	avg its per Newt step	time (sec)
1	4356	41	23	100	96	29	18	28	20
4	16,900	98	62	324	320	37	25	40	27
16	66,564	251	275	1156	1088	40	34	50	39
64	264,196	603	1,399	4356	4096	38	57	57	69
256	1,052,676	1,478	8,085	16900	16384	37	151	63	191

proc	fine grid unknowns	1 - level Method Ilu DD		coarse unknowns		2-level: gs2-superlu			
						geometric		algebraic	
		avg its per Newt step	time (sec)	geometric	algebraic	avg its per Newt step	time (sec)	avg its per Newt step	time (sec)
4	24.565	40[5]	123	135	120	36[5]	101	30[4]	71
32	179,685	112[5]	282	625	480	44[4]	107	50[4]	109
256	1.373.125	296[5]	863	3,645	2560	47[5]	179	58[4]	152
2048	10,733,445	650[5]	2,915	24,565		47[4]	546	59[4]	681

Analysis: Sala; Math. Modeling and Numer. Anal., 2004
Sala, Shadid, Tuminaro; accepted in SIMAX
Numerical Exp:
Lin, Sala, Shadid, Tuminaro; accepted in IJNME

- Coarse mesh: SuperLU direct solver
- Run on Sandia ASCI Red machine

Sandia National Laboratories

Preliminary Scaling with Reactions: Poly-Silicon CVD

- Deposition of poly-Silicon
  - + **3 species (8 unknowns per node)**
- Steady-state calculation restarted from steady-state solution at lower pressure (one solution in continuation run)
- Run on Cplant

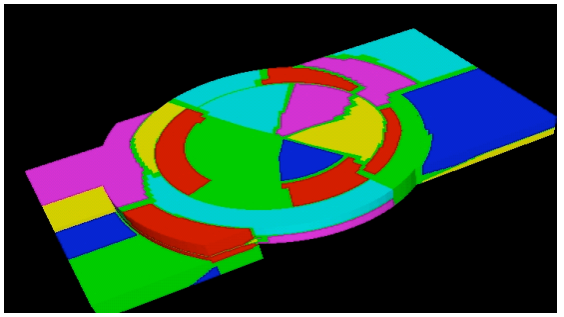

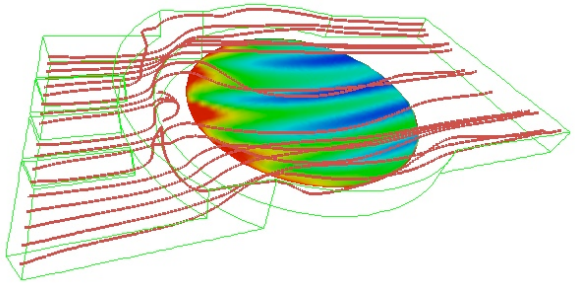

proc	fine unknowns	coarser unknowns	
		no auxiliary matrix	auxiliary matrix
2	87,400	3560/296/24	17K/1240/120/16
16	636K	22K/1816/224/72	120K/9096/880/136
128	4.85M	161K/12K/1440/368	967K/75K/6744/1296

		1-level (ILU)		5-level (GS/ILU/ILU/ILU/KLU) damp 0.67			
				no auxiliary matrix		auxiliary matrix	
proc	unknowns	ave its/ Newt step	time (sec)	ave its/ Newt step	time (sec)	ave its/ Newt step	time (sec)
2	87,400	47 [6]	768	67 [6]	659	45 [6]	690
16	636K	88 [8]	1536	102 [8]	1254	60 [8]	1202
128	4.85M	202 [9]	5765	182 [9]	3099	86 [9]	2417

- Aspect Ratio Effect; Even with auxiliary matrix, still does not scale; more work necessary

Sandia National Laboratories

Drift-Diffusion Equations for Semiconductor Modeling

Electric potential

$$\lambda^2 \nabla \cdot (\epsilon_r \mathbf{E}) = p - n + C \qquad \mathbf{E} = -\nabla \psi$$

Current conservation

$$\nabla \cdot \mathbf{J}_n = \frac{\partial n}{\partial t} + R \qquad \mathbf{J}_n = \mu_n n \mathbf{E} + D_n \nabla n$$

$$-\nabla \cdot \mathbf{J}_p = \frac{\partial p}{\partial t} + R \qquad \mathbf{J}_p = \mu_p p \mathbf{E} - D_p \nabla p$$

"constitutive" relation

- ψ: electric potential
- n: electron density
- p: hole density
- C: doping profile
- R: generation-recombination term

Stabilized FE method (Charon - Hennigan, Hoekstra, Lin, Shadid)

Sandia National Laboratories

Weak Scaling Study: 1-level and 3-level (with and without EMIN)
2x1.5 micron NPN BJT Steady-State Drift-Diffusion

- Charon FEM semiconductor device modeling code
- 2D steady-state drift-diffusion bias 0.3V
- 3-level AMG preconditioner (ML library); EMIN2 with block scaling

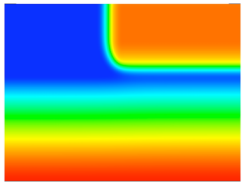

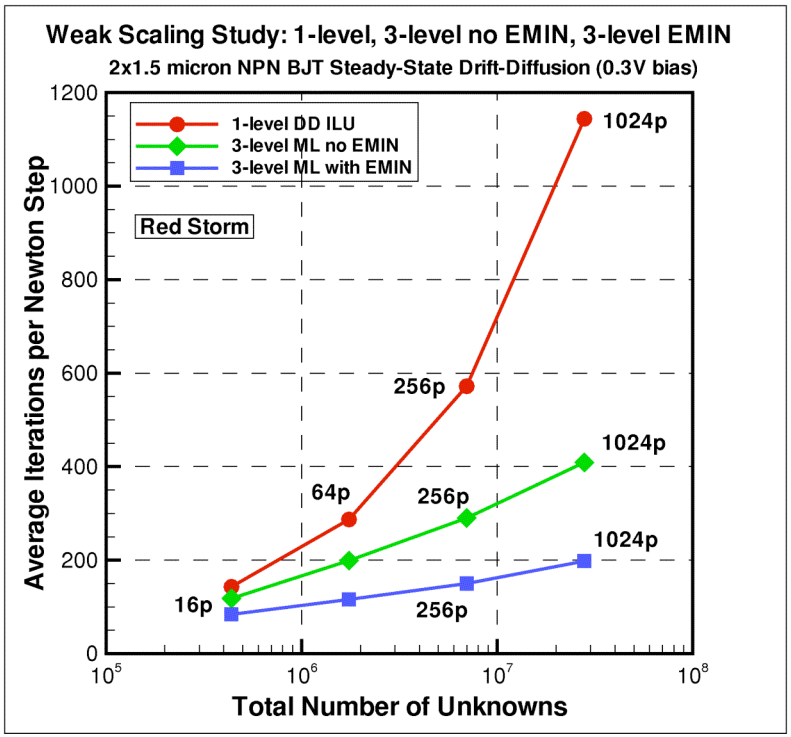

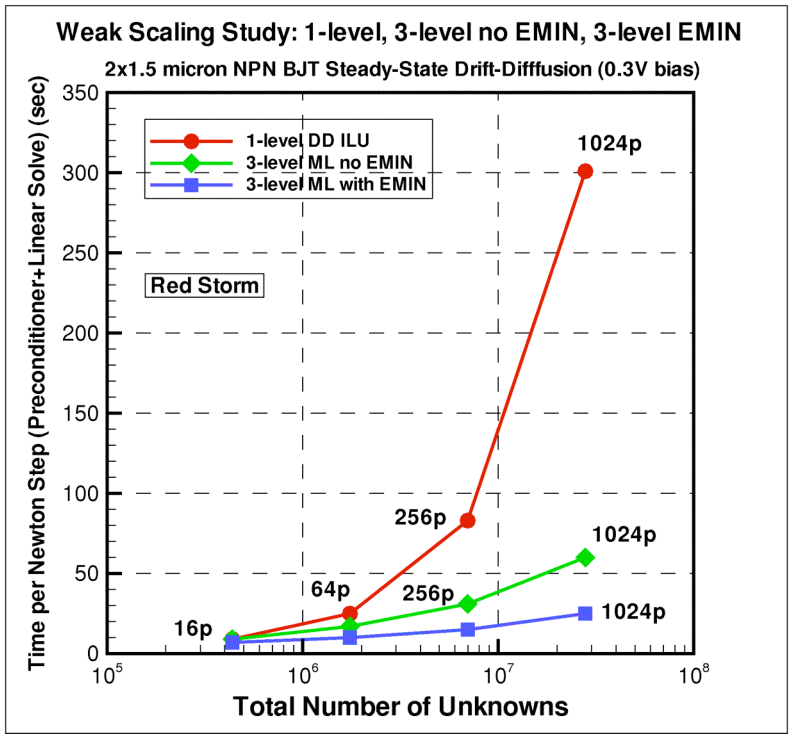

- "Time": construct preconditioner and perform linear solve

Sandia National Laboratories

Currently: Low Mach number Single Fluid Resistive MHD Formulation (Unstructured FE)

Conservation Law System

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \bullet \mathbf{F} + \mathbf{S} = \mathbf{0}$$

Magnetic Flux

$$\mathbf{U} = \begin{bmatrix} \rho \\ \rho\mathbf{v} \\ \Sigma_{tot} \\ \mathbf{B} \end{bmatrix} \quad \mathbf{F} = \begin{bmatrix} \rho\mathbf{v} \\ \rho\mathbf{v}\otimes\mathbf{v} - \frac{1}{\mu_0}\mathbf{B}\otimes\mathbf{B} - \mathbf{T} + \frac{1}{2\mu_0}\|\mathbf{B}\|^2\mathbf{I} \\ \rho E\mathbf{v} - \mathbf{T}\cdot\mathbf{v} + \mathbf{E}\times\mathbf{B} + \mathbf{q} \\ \mathbf{v}\otimes\mathbf{B} - \mathbf{B}\otimes\mathbf{v} - \frac{\eta}{\mu_0}\left(\nabla\mathbf{B} - \nabla\mathbf{B}^T\right) \end{bmatrix} \quad \mathbf{S} = \begin{bmatrix} 0 \\ \mathbf{0} \\ Q^{rad} + Q \\ \mathbf{0} \end{bmatrix}$$

$$\nabla \cdot \mathbf{B} = \mathbf{0}$$

$$\Sigma_{tot} = \rho E + \frac{1}{2\mu_0}\|\mathbf{B}\|^2 \qquad E = e + \frac{1}{2}\|\mathbf{v}\|^2$$

$$\mathbf{E} = -\mathbf{v}\times\mathbf{B} + \eta\mathbf{J} + \underbrace{\frac{1}{en}(\mathbf{J}\times\mathbf{B} - \nabla\mathbf{P_e})}_{\text{Hall}}\cdot$$

Vector Potential

$$\mathbf{U} = \begin{bmatrix} \rho \\ \rho\mathbf{v} \\ \Sigma_{tot} \\ \mathbf{A} \end{bmatrix} \quad \mathbf{F} = \begin{bmatrix} \rho\mathbf{v} \\ \rho\mathbf{v}\otimes\mathbf{v} - \frac{1}{\mu_0}\mathbf{B}\otimes\mathbf{B} - \mathbf{T} + \frac{1}{2\mu_0}\|\mathbf{B}\|^2\mathbf{I} \\ \rho E\mathbf{v} - \mathbf{T}\cdot\mathbf{v} + \mathbf{E}\times\mathbf{B} + \mathbf{q} \\ \mathbf{v}\otimes\mathbf{A} + \mathbf{A}\otimes\mathbf{v} - \frac{\eta}{\mu_0}\left(\nabla\mathbf{A} - \nabla\mathbf{A}^T\right) \end{bmatrix} \quad \mathbf{S} = \begin{bmatrix} 0 \\ \mathbf{0} \\ Q^{rad} + Q \\ \mathbf{S}_A \end{bmatrix}$$

$$\mathbf{B} = \nabla \times \mathbf{A}.$$

$$\mathbf{S}_A = \nabla\left(\frac{\eta}{\mu_0}\right)\cdot\left[\nabla\mathbf{A} - \nabla\mathbf{A}^T\right] + \left[\mathbf{A}\times(\nabla\times\mathbf{v}) - \left(\frac{1}{\mu_0}\nabla\eta + \mathbf{v}\right)(\nabla\cdot\mathbf{A}) - \mathbf{A}(\nabla\cdot\mathbf{v}) - \nabla(\mathbf{A}\cdot\mathbf{v})\right]$$

In 2D we have: $$\frac{\partial \mathcal{A}}{\partial t} + \nabla\cdot\left(\mathcal{A}\mathbf{v} - \frac{\eta}{\mu_0}\nabla\mathcal{A}\right) + \nabla\left(\frac{\eta}{\mu_0}\right)\cdot\nabla\mathcal{A} - \mathcal{A}(\nabla\cdot\mathbf{v}) = 0$$

Sandia National Laboratories

Initial Hydro-magnetic Rayleigh-Bernard Stability Study

Hydromagnetic Rayleigh-Bernard stability study (Rigid boundaries):

Thermal Convection above with no magnetic field (Q = 0)
Ra_{cr} = 1707.5 (Chandrashekhar)

Stable flow with magnetic field (non-zero stable flow / magnetics):

Ra = 2000, Q = 4

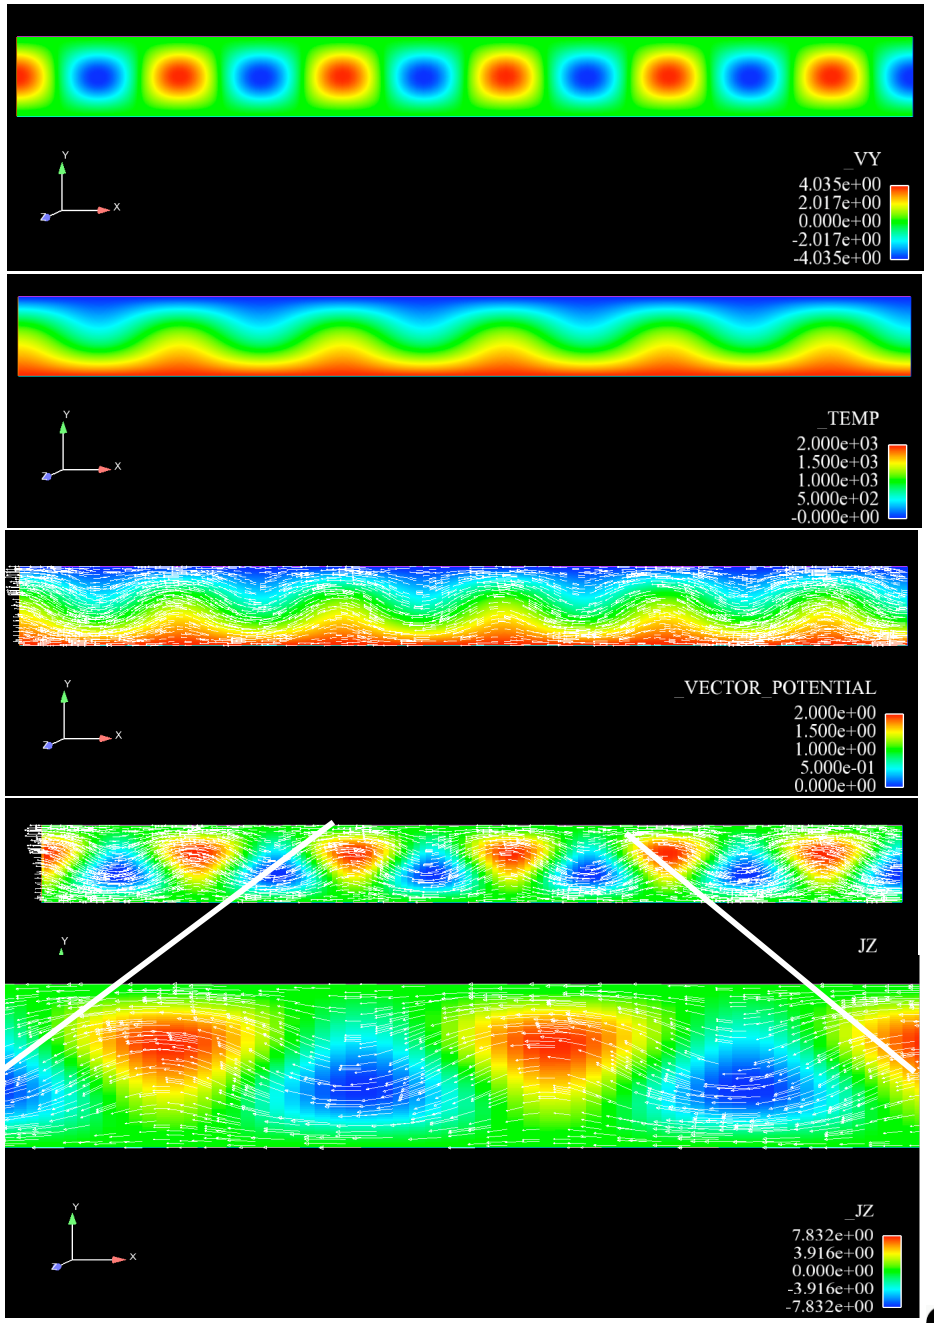

Sandia National Laboratories

Initial ML Results on Hydro-magnetic Rayleigh-Bernard Stability Study (Rigid Boundaries)

Thermal Convection with no magnetic field (Q = 0)

Ra_{cr} = 1707.5 (Chandrashekhar)

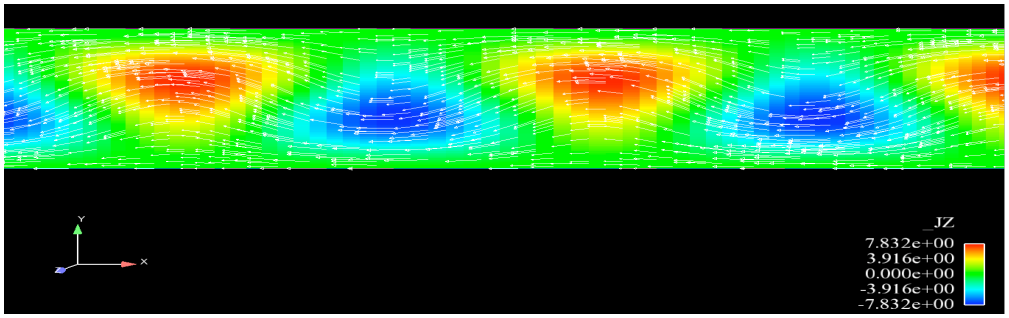

Flow with magnetic field impulsively turned on:

Ra = 2000, Q = 4
(stable state is with flow and magnetics)

Initial Serial Results for aggressive block coarsening AMG

Mesh	Total unknowns	method	Coarse problem unknowns	Linear Its	Construct Precond. (sec.)	Solve Time (sec)
200x80	81,405	1 level	-----	283	4.0	110
200x80	81,405	2 level Agg = 50	1,625	202	4.7	69
200x80	81,405	2 level Agg = 10	8140	178	5.4	52
400x160	322,805	1 level	-------	346	15.6	668
400x160	322,805	2 level Agg = 50	6,455	216	21.8	315

Preconditioners

1 Level solver: ILU fill = 2
2 Level solver: fine: ILU fill = 2, Coarse: direct sparse solve (KLU)

(w/ Pawlowski, Banks, Lin)

Sandia National Laboratories

Incompressible Fluid Mechanics: A Particular View of Block Preconditioning (Kay & Loglin; Wathen & Silvester; Elman.)

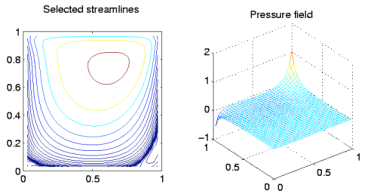

- **For N-S system.** $\begin{bmatrix} F & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathrm{u} \\ p \end{bmatrix} = \begin{bmatrix} \mathrm{f} \\ 0 \end{bmatrix}$ Block LU leads to $(S = BF^{-1}B^T).$

$$\begin{bmatrix} F & B^T \\ B & 0 \end{bmatrix} = \begin{bmatrix} I & 0 \\ BF^{-1} & I \end{bmatrix} \begin{bmatrix} F & B^T \\ 0 & -S \end{bmatrix}$$

Observation: use as Right Preconditioner!
(all eigenvalues 1, Jordon blocks of dimension at most 2)

$$\begin{bmatrix} F & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} F & B^T \\ 0 & -S \end{bmatrix}^{-1} = \begin{bmatrix} I & 0 \\ BF^{-1} & I \end{bmatrix}$$

- **A heuristic motivation: B^T and F contain derivatives. If they can be commuted, then we can change the order of the operators:**

$$P = \begin{bmatrix} F & B^T \\ 0 & -\widehat{S} \end{bmatrix}$$

$$\begin{aligned} S &= \left[BF^{-1}B^T \right] \\ &\approx \left[BB^T F_p^{-1} \right] \\ \widehat{S} &= A_p F_p^{-1}, \end{aligned}$$

$$\begin{aligned} A_p &= \Delta_h \\ F_p &= -\nu\Delta_h + \mathrm{u}^{n-1} \cdot \nabla_h \end{aligned}$$

- **Solve simplified component systems with parallel aggregation (AMG in ML); Convection/diffusion and "pressure Poisson"**

$$F\widehat{u} = rhs \qquad \text{Momentum conv./diff solve}$$
$$A_p\widehat{p} = rhs \qquad \text{Pressure Poisson solve}$$

Sandia National Laboratories

Incompressible Fluid Mechanics: A Comparison of ABF Preconditioning: Fp, DD ILU with MPSalsa

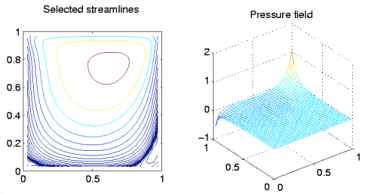

Mesh Independence

Re	Mesh	DD	Fp	Proc
10	64 x 64	79.4	19.4	1
	128 x 128	220.6	21.2	4
	256 x 256	1018.6	23.0	16
100	64 x 64	86.5	34.2	1
	128 x 128	300.3	35.9	4
	256 x 256	1603.9	41.3	16
500	64 x 64	89.7	102.2	1
	128 x 128	334.9	99.5	4
	256 x 256	5433.1	101.3	16

The values in each column represent the average number of outer linear iterations per Newton Step.

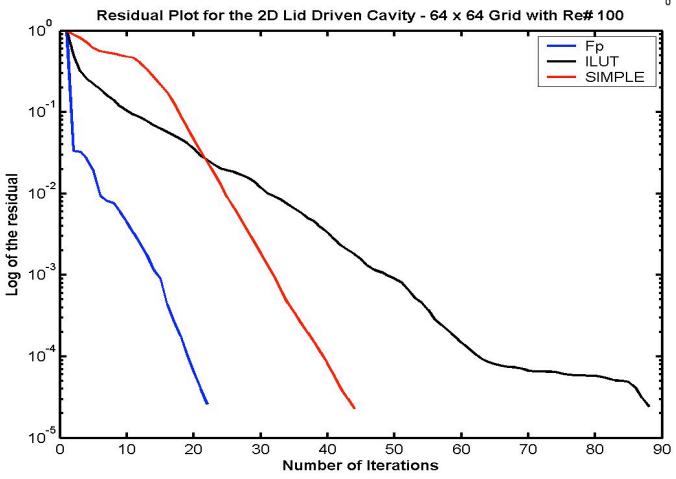

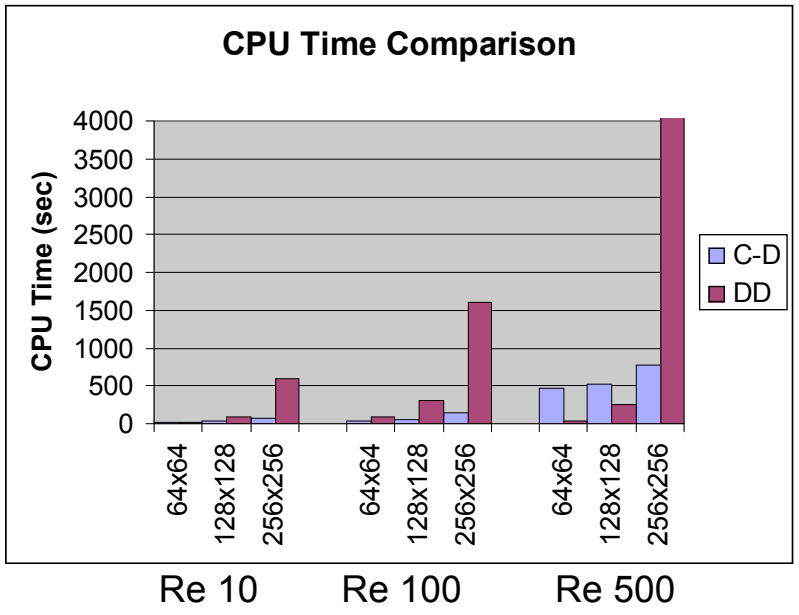

Sandia National Laboratories

MPSalsa Steady Problem Results

2D Flow over a Diamond Obstruction

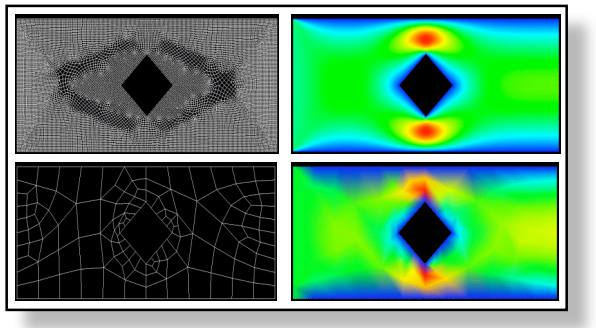

Re #	Unknowns	DD	C-D	Nprocs
10	64K	110.8 (186.6)	20.5 (138.8)	1
	256K	284.6 (1657.4)	22.5 (266.2)	4
	1M	1329.0 (7825.5)	22.9 (501.0)	16
	4M	NC (NC)	29.4 (1841.7)	64
25	64K	101.7 (198.8)	32.9 (248.0)	1
	256K	273.8 (1583.1)	35.9 (480.6)	4
	1M	1104.8 (7631.5)	38.3 (956.9)	16
	4M	NC (NC)	48.0 (4189.8)	64
40	64K	70.4 (267.2)	54.6 (565.8)	1
	256K	203.9 (1420.7)	70.1 (1280.9)	4
	1M	997.1 (8188.2)	65.4 (2011.7)	16
	4M	NC (NC)	79.8 (9387.9)	64

Transient and Pseudo-transient Solvers: 3D Lid Driven Cavity, **N = 64 & Re = 500**

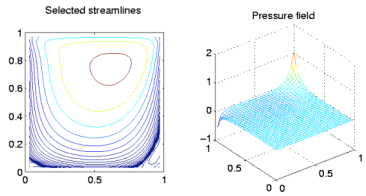

CFL	Time (secs.)	Oseen Steps	Saddle Pt. Solves	Ap (press.)	F (conv. / diff.)
0.1	83.0	2	2	20	2
0.5	85.6	2	2	20	2
1	79.3	2	2	20	2
10	110.0	2	2	20	2
50	92.8	2	2	20	3
100	103.7	3	2	20	3
5,000	238.6	5	5	18	5
10,000	269.6	5	6	18	6
50,000	403.9	6	9	19	8

- **MAC FD scheme used in examples, no upwinding**

100 Procs of ASCI – Red Tflop

Sandia National Laboratories

Multi-level Methods for Coupled Systems of Equations (ML package in Trilinos)

Low Mach No. Flow / Transport and Transport / Reaction - MPSalsa (Nodal FE)

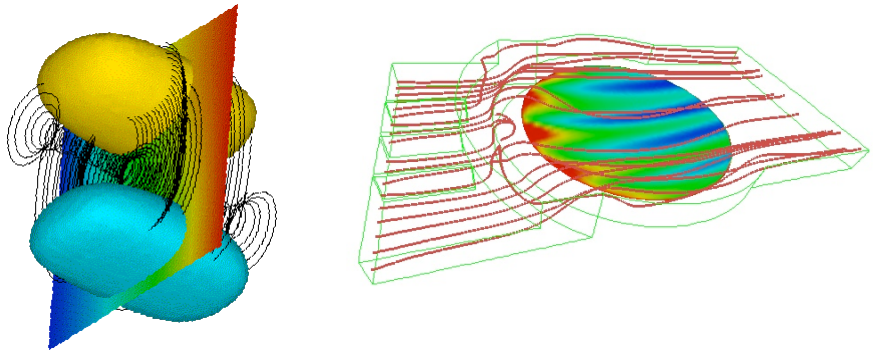

Compressible Euler / Navier-Stokes - Premo (Vertex based FV)

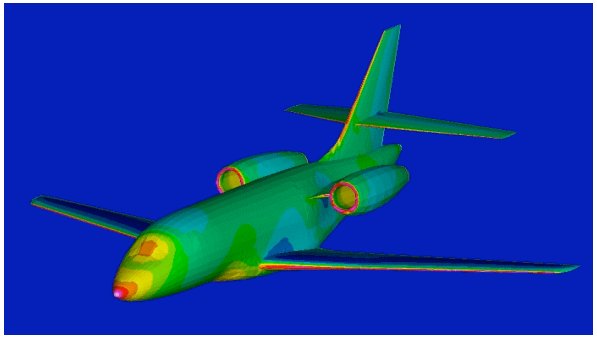

Drift Diffusion and Radiation Damage models with Charged Species in Semiconductors - Charon (Nodal FE)

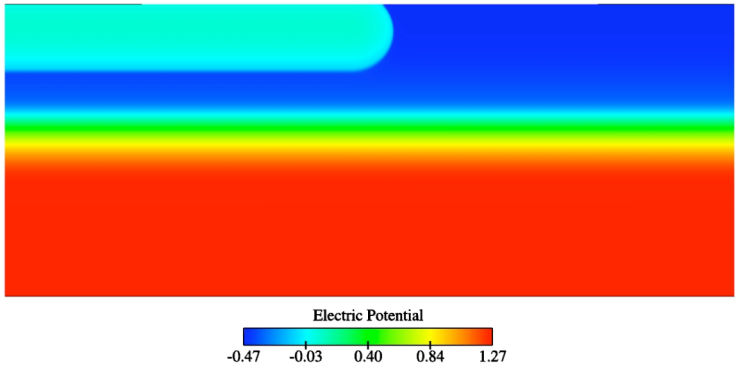

Magnetic Diffusion Solver - Alegra (Edge based FE)

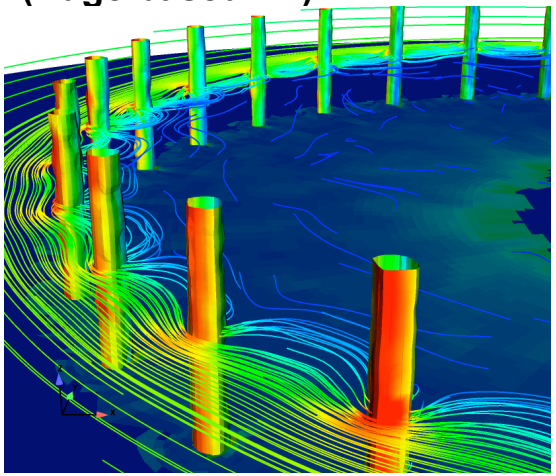

Sandia National Laboratories

Conclusions

- Newton-Krylov methods can provide a very effective, robust and flexible solution technology for analysis and characterization of complex nonlinear solution spaces. For steady state, time dependent and optimization type solutions.

- High parallel efficiencies for fully-implicit fully coupled Newton-Krylov iterative solvers for a wide range of problems are possible.

- Parallel multilevel aggressive coarsening block AMG preconditioners for systems have shown promising results for algorithmic scalability and CPU time performance of transport solutions.

(Issues: Strong convection, reaction and FE aspect ratios for multilevel methods)

- Cray XT3 very capable parallel computing platform. Very good scaling results on CFD type of simulations.

Collaborators & Sponsors

Transport / Reaction Solver	Solution Methods				
	Stabilized FE Methods	Non-linear/ stability/ Bifurcation Solvers	Linear Solvers	Eigen Solvers	Optimization
Karen Devine Gary Hennigan Scott Hutchinson Paul Lin Harry Moffat Roger Pawlowski Curt Ober Andy Salinger	Pavel Bochev	Roger Pawlowski Andy Salinger	Ray Tuminaro Vicki Howle Paul Lin Marzio Sala	David Day Rich Lehoucq	Roscoe Bartlett Roger Pawlowski Andy Salinger Bart van Bloemen Waanders
Applications: CVD: SUNY- Buffalo, Sumitomo Sitix CVD&Stability: U. Manchester Fuel Cells: Ford Resc. Biological Cells: UNM, UConn. Stability: Princeton, UNM Optimization: Carnegie-Mellon, Rice Univ.	Theory: Max Gunzburger (Fl. St.) Tom Hughes (U. TX)	Algorithms: Homer Walker (WPI) Joe Simonis (WPI)	Algorithms: Charles Tong (LLNL) Howard Elman (U. Maryland) Bob Shuttleworth (U. MD)		L. Biegler (CMU)

Office of Science
Science for America's Future

MICS Program

ASC™

DOE / NNSA

Homeland Security

ASC Science & Tech.

Sandia National Laboratories

Trilinos: Full Vertical Solver Coverage (Part of DOE: TOPS SciDAC Effort)

Optimization Unconstrained: Constrained:	Find $u \in \Re^n$ that minimizes $g(u)$ Find $x \in \Re^m$ and $u \in \Re^n$ that minimizes $g(x,u)$ s.t. $f(x,u) = 0$	**MOOCHO**
Bifurcation Analysis	Given nonlinear operator $F(x,u) \in \Re^{n+m} \rightarrow \Re^n$ For $F(x,u) = 0$ find space $u \in U \ni \frac{\partial F}{\partial x}$ singular	**LOCA**
Transient Problems DAEs/ODEs:	Solve $f(\dot{x}(t), x(t), t) = 0$ $t \in [0,T], x(0) = x_0, \dot{x}(0) = x_0^{'}$ for $x(t) \in \Re^n, t \in [0,T]$	**Rhythmos**
Nonlinear Problems	Given nonlinear operator $F(x,u) \in \Re^{n+m} \rightarrow \Re^n$ Solve $F(x) = 0$ $x \in \Re^n$	**NOX**
Linear Problems Linear Equations: Eigen Problems:	Given Linear Ops (Matrices) $A, B \in \Re^{m \times n}$ Solve $Ax = b$ for $x \in \Re^n$ Solve $A\nu = \lambda B\nu$ for (all) $\nu \in \Re^n$, $\lambda \in \Re$	**AztecOO** **Belos** **Ifpack, ML, etc...** **Anasazi**
Distributed Linear Algebra Matrix/Graph Equations: Vector Problems:	Compute $y = Ax; A = A(G); A \in \Re^{m \times n}, G \in \Im^{m \times n}$ Compute $y = \alpha x + \beta w; \alpha = \langle x, y \rangle; x, y \in \Re^n$	**Epetra** **Tpetra**

THE END OF TALK

Sandia National Laboratories

Phase II: tentative prolongator

- **For the Laplace equation, the simplest prolongator is:**

$$P_t(i,j) = \begin{cases} 1 & \text{if } i \in \text{aggregate} \quad \text{j} \\ 0 & \text{otherwise} \end{cases}$$

- **The basis functions for the coarse space reads**

$$\Phi_i(x) = \sum_{j \in agg_i} \phi_j(x)$$

1D: starting grid and ϕ_j

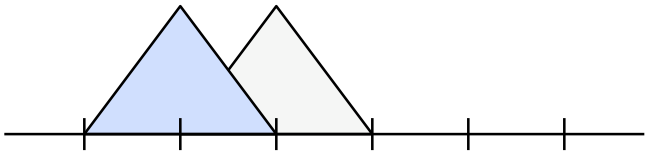

1D: aggregates and φ_i

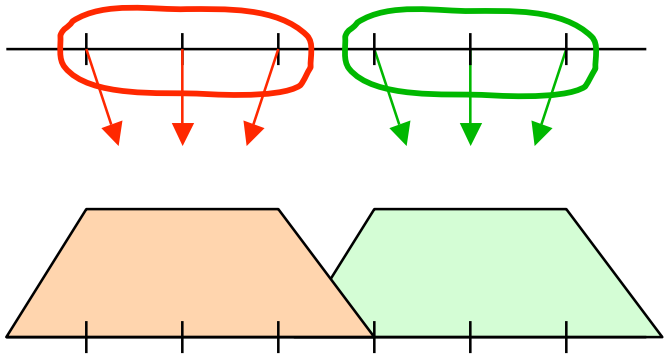

$\Phi_{i's}$ have high energy (1/h for linear functions): need to smooth them to improve the preconditioner

$$P_t =$$

$$\begin{matrix} 1 & & & & & & \\ 1 & & & & & & \\ 1 & & & & & & \\ & 1 & & & & & \\ & 1 & & & & & \\ & 1 & & & & & \\ & & \ddots & & & & \\ & & & & & 1 & \\ & & & & & 1 & \\ & & & & & 1 & \end{matrix}$$

Sandia National Laboratories

Algebraic: Smoothed Aggregation

+ Easily implemented on unstructured grids, 2D/3D
+ Simple representation of restriction/prologator operators
+ (Almost) black-box (for matrices arising from FE discretization of PDE equations)

Aggregation is performed in 4 steps:

I. Generation of the aggregates;
II. Construction of simple restrictions/prolongators
III. Improvement of restrictions/prolongators using (simple) smoothing
IV. Construction A_{k-1} as

$$A_{k-1} = R_k A_k P_k, \quad A_L = A$$

Phase I: Construction of the Aggregates

- **For each level k, create a graph G_k:**
  - # vertices (G_k) = n_k / num PDE eqns
  - An edge between vertices i and j is added iff $a_k(i,j) \neq 0$
  - May ignore `weak' matrix coupling, e.g.

$$a_k(i,j) \ll \max |a_k(i,i), a_k(j,j)|$$

- **Decompose G_k into subgraphs (called aggregates). Each aggregate is composed by a set of contiguous vertices.**

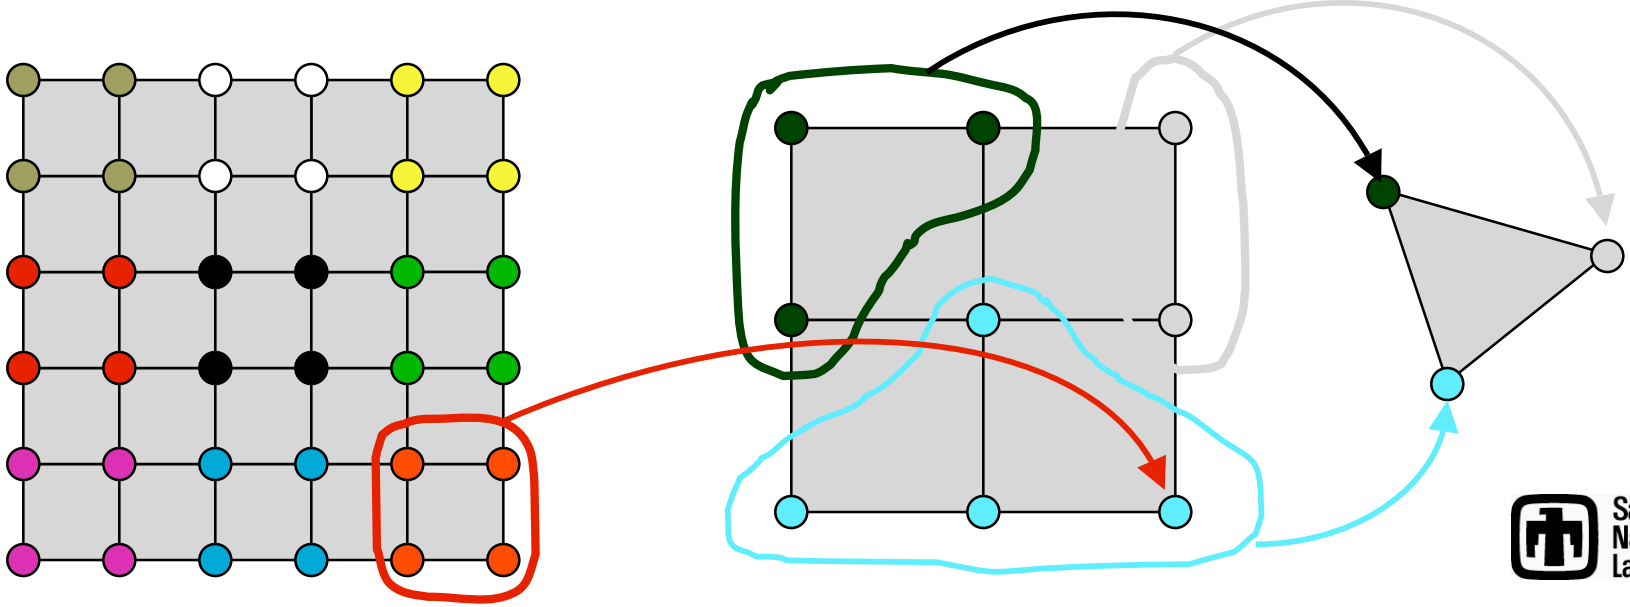

Sandia National Laboratories

Definition of the Aggregates (3)

- **Using only few levels is very close to two-level domain decomposition (DD) preconditioners**
- **Aggregation has been used in DD literature:**
  - + Shallow water equations for 2D (Paglieri, Scheinine, Formaggia, Quarteroni,1997);
  - + 3D potential flows (Formaggia, Scheinine, Quarteroni,1994);
  - + Groundwater flows (Jenkins, Kelley, Miller, Kees,2000);
  - + Discontinuous Galerkin for advection diffusion (Toselli,Lasser, 2000);
  - + Compressible Euler Equations (S.,Formaggia,2001);
  - + …
- **Abstract Schwarz theory can be used to estimate theoretical properties for (additive) preconditioners.**

Two-level Schwarz preconditioners

- If the coarse space is defined using a coarse grid (Dryja and Widlund, 1989, Smith Bjorstad and Gropp, 1996):

$$\kappa\left(P_{add}^{-1}A\right) \leq C_1\left(1+\frac{H}{\delta}\right)$$

- If the coarse space is defined using aggregation, with one aggregate per domain (Lasser and Toselli, 1999):

$$\kappa\left(P_{add}^{-1}A\right) \leq C_2\left(1+\frac{H}{\delta}\right)$$

- If more than 1 aggregate per domain, and H_0 is aggregate diameter (M.S., PhD Thesis, 2003):

$$\kappa\left(P_{add}^{-1}A\right) \leq C_3\left(1+\frac{H}{\delta}\right)\left(1+\frac{H_0}{\delta_0}\right)$$

Sandia National Laboratories

Conclusions (contd.)

- Issues with our formulation:
  - Some difficulties with un-stabilized oscillations for steady state simulations.(pressure/velocity coupling – pursue consistency restoring methods for low order FE , large gradients – non monotonicity preserving method). Pavel's talk to follow.
  - Experience with some counter intuitive behavior of consistent stabilized method for small time step sizes.
    - Loss of stabilization for small time step sizes
    - Loss of second order convergence in time for small time steps

    Current work by Harari, SNL