FACETS: towards whole device modeling with concurrent, tightly coupled components

Presenter: J.R. Cary (Tech-X and U. Colorado)

Location: PPPL

February 24, 2009

The multi-institutional FACETS project was started two years ago to develop new computational capability for multiphysics, whole device modeling that can take advantage of the current and future large parallel hardware. The physics goals include developing an understanding of how a consistent, coupled core-edge-wall plasma evolves, including energy flow, particle recycling, and the variation of power density on divertor plates with plasma under different conditions. FACETS is just now entering the research phase, as it has only recently developed coupled core-edge simulation capability. Hence this talk will concentrate on how we got to this point. This includes a description of the software architecture with issues involved in developing concurrent component parallelism. It also includes what was involved in developing a new core solver (which provides parallelism, with a speedup of more than an order of magnitude). We also discuss some of the issues of verification and validation, as well as some software engineering issues dealing with testing and cross platform build issues. The conclusion will stress lessons learned about how to make a multi-institutional project work (or not), comments on code reuse versus rebuild, and methods for introduction of engineering into computational physics for improved results.

IU
NYU
Lodestar

FACETS has contributors from many institutions

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
PSFC
IU
NYU
Lodestar

- ANL (solvers): McInnes, Zhang, Balay
- ANL (coupling): Larson
- CSU (sensitivity research): Estep, Pham
- General Atomics (GYRO, turbulence understanding, long-range and/or refinement): Candy
- Lehigh (core modeling, SBIR subcontract): Pankin
- LLNL (edge physics): Rognlien, Cohen
- LLNL (interlanguage): Epperly
- ORNL (modeling, evangelism): Cobb
- ORNL (SSGKT): Fahey
- ParaTools (performance analysis): Malony, Morris, Shende
- PPPL (core sources): McCune, Indireshkumar
- UCSD (wall): Pigarov
- Tech-X (framework, core): Cary, Hakim, Kruger, Pletzer, Vadlamani, Miah, Shasharina, Muzsala, VizSchema team

The FACETS project is a multidisciplinary-research project

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT PSFC
BERKELEY LAB
IU
NYU
Lodestar

- Goal is to develop a capability to do computational fusion science research
- Do research with this capability
- BUT also
- How can one do multiphysics computational science?
  - How does one couple?
  - On LCFs?
- How does one get a multi-institutional developer team to work together effectively?
- And other CS/AM research to help FACETS meet its goals

All needed for the FSP

FACETS goal: tight coupling framework for core-edge-wall

- Coupling on short time scales
- Inter-processor and in-memory communication
- Implicit coupling

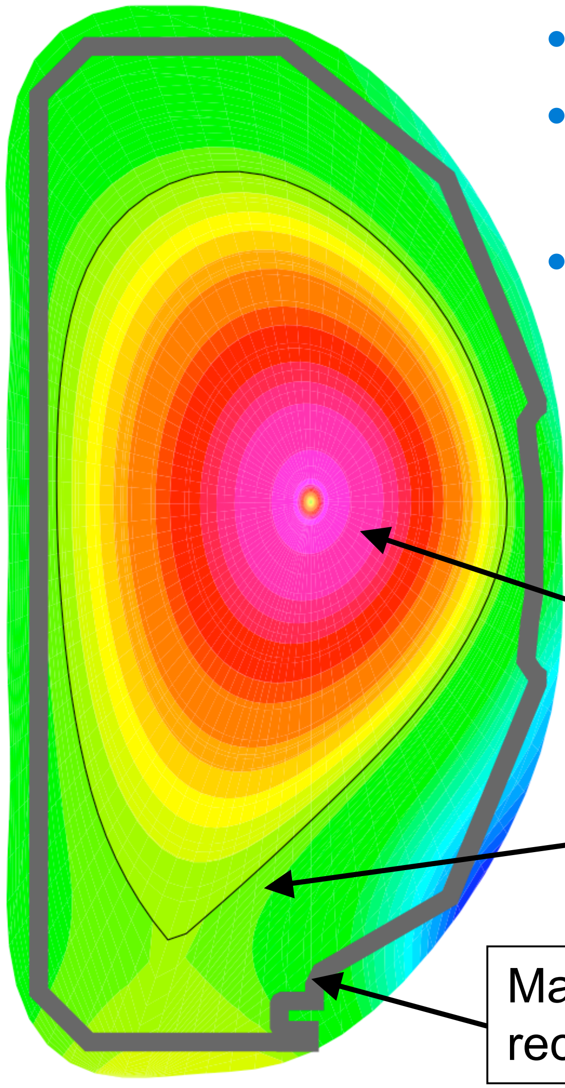

Hot central plasma: nearly completely ionized, magnetic lines lie on flux surfaces, 3D turbulence embedded in **1D** transport

Cooler edge plasma: atomic physics important, magnetic lines terminate on material surfaces, 3D turbulence embedded in **2D** transport

Material walls, embedded hydrogenic species, recycling

IU
NYU
Lodestar

Why do core-edge-wall coupling?

- **A fusion plasma finds a self-consistent, core-edge-wall state**
  - **Energy into the edge determines the pedestal (pressure) height, while the pedestal height is a dominant determiner of interior temperature, and so fusion power**
  - **Particle recycling involves wall loading/discharging, with the particles from the wall determining plasma density which then determines flux into the wall**
- **Coupled simulations vs monolithic codes exploits space and time scale disparities and makes use of proven techniques for incorporating important physics in each region**

Crucial to get structure in place as edge models are being rapidly developed!

Why do reduced-model transport simulation?

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
MIT PSFC
Berkeley Lab
IU
NYU
Lodestar

- Would like to get to 1000 simulation second in one wall-clock hour.
- GYRO simulation: 1 ms on D3D takes 1 wall-clock hour on 128 procs for well resolved.
- Off by 1e6 at the present time. (Improvable with more parallelism, algorithmic improvement?)
- Can be used for refinement (see next talk on SSGKT), study of flattop, and embedded in longer-time-step simulations.

FACETS actually pursuing usage of multiple components with varying levels of fidelity

FACETS only one of proto FSPs (all investigating parallel coupling)

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
MIT PSFC
Berkeley Lab
IU
NYU
Lodestar

- FACETS
  - Physics: core-edge-wall coupling (both surfacial and volumetric couplings)
  - Framework: in-memory, concurrent, implicit coupling
- SWIM
  - Physics: RF-MHD coupling in the core
  - Framework: file-based, sequential coupling
- CPES
  - Physics: Edge transport-MHD coupling
  - Framework: file-based, sequential, diagnostic

All needed for the FSP

Idealized view: surfacial couplings between phase transitions

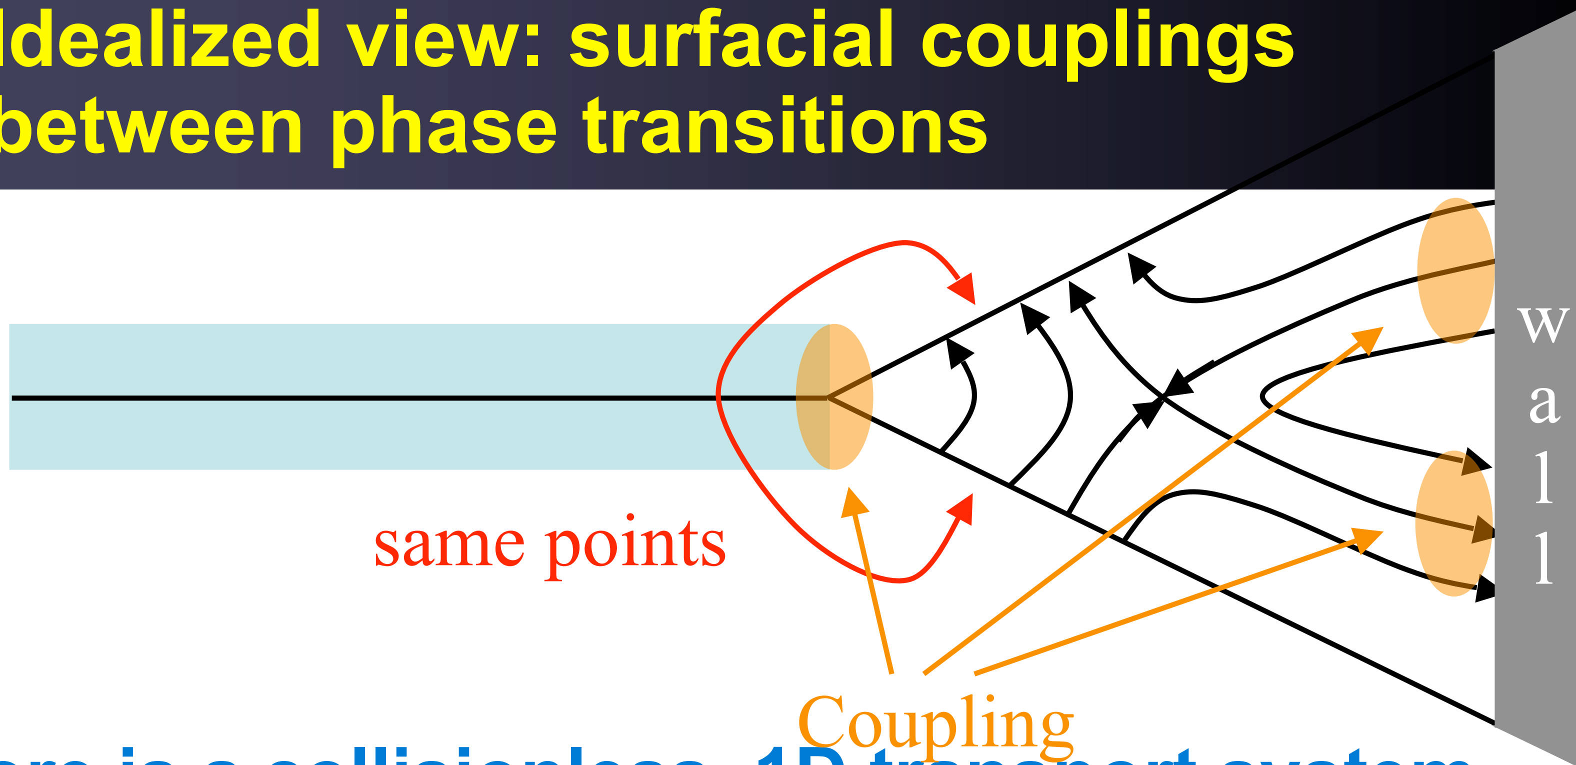

- **Core is a collisionless, 1D transport system with local, only-cross-surface fluxes**
- **Edge is a collisional, 2D transport system**
- **Wall: beginning of a particle trapping matrix**

Surfacial coupling: COUPLING DIMENSIONALITY

volumetric investigated by CPES, SWIM

embedded in out years by FACETS

Idealized view likely okay for edge-wall interaction

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT PSFC
BERKELEY LAB
IU
NYU
Lodestar

- Edge-plasma/wall analogous to atmosphere/ocean
  - Wall acts as a boundary condition for edge plasma
  - Sputtering
  - Secondary electron/ion emission
- Refinement needs wall model to account for internal state
  - Wall has embedded H/D density
  - H/D diffuses in metal, both in and out
  - Impact of electrons, ions, and neutrals can cause release of embedded H/D

Valid basis for independent components

Justification for core-edge coupling needs matching

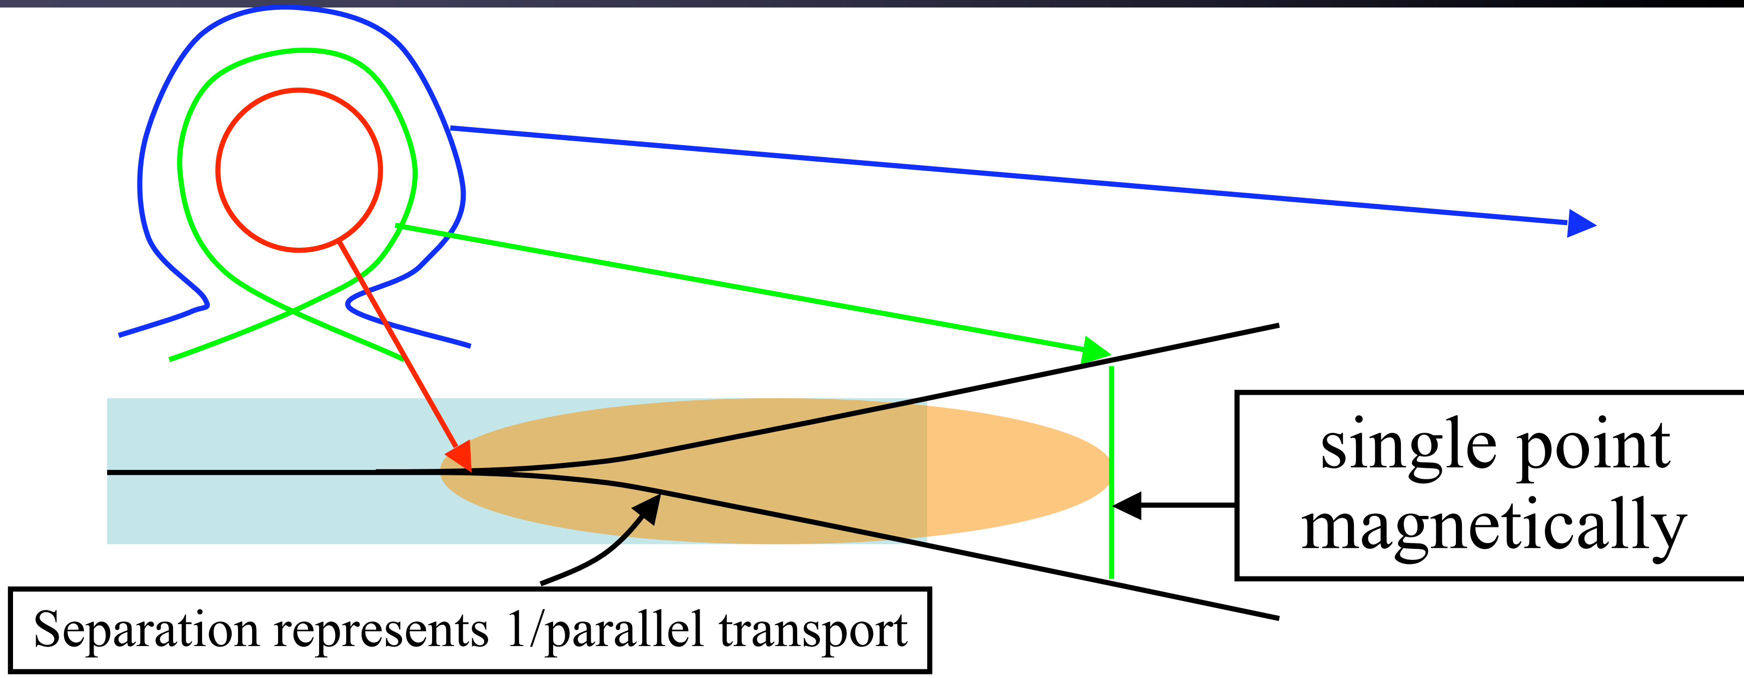

- Sufficiently inside the last closed flux surface, 2D effects are small
- Moving out, plasma has become collisional
- Both approximations exist - allows matching
- Challenge to analytic theorists: provide matching theory

Communication needed for core-edge is very simple

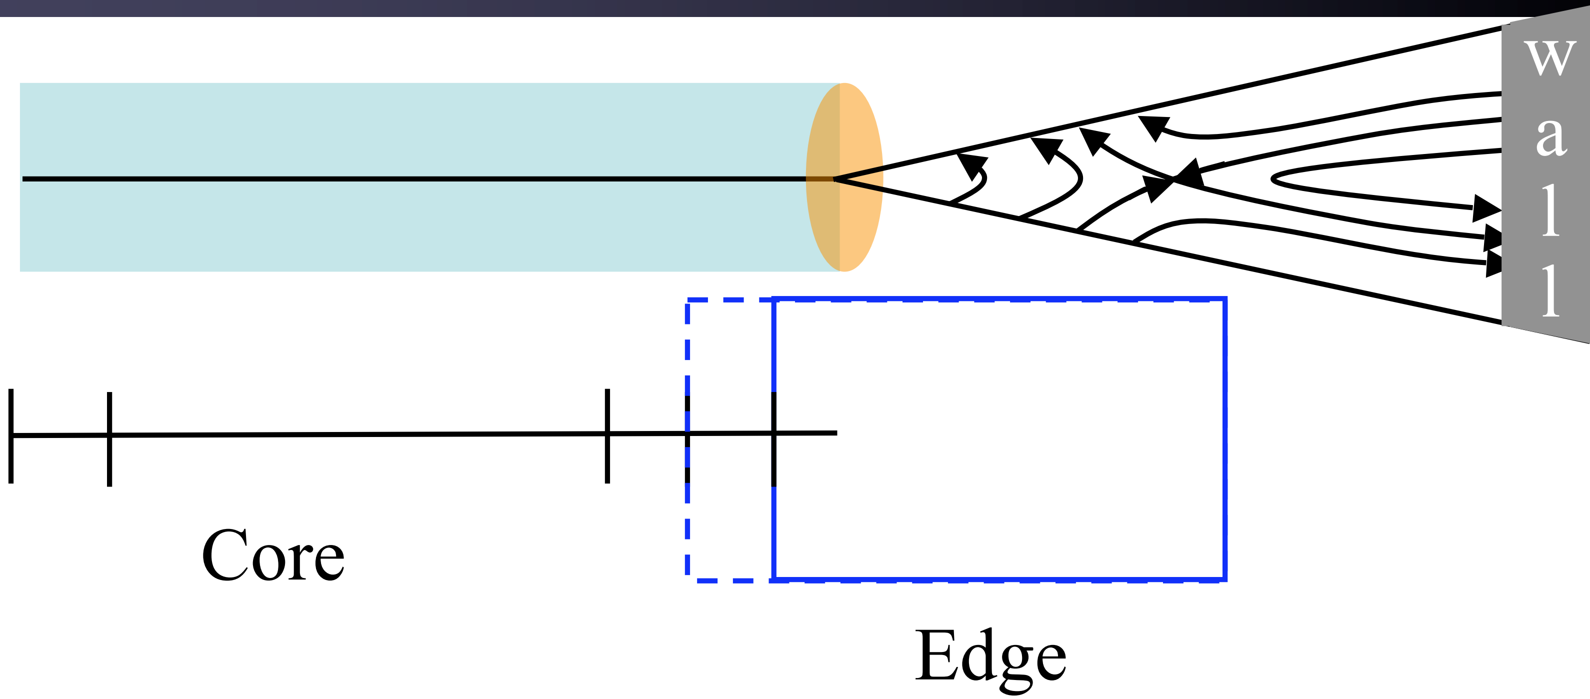

- Establish convention that cell is [lo, hi)
- Edge owns last cell of 1D system
- Edge is given flux of last cell of core
- Core is given density of first cell of edge

(Variation to test coupling math)

Concurrent coupling will be important for multiphysics on LCFs

TECH
Argonne
Colorado State
General Atomics
ParaTools
PPPL
UCSD
IU
NYU
Lodestar

Core solver
UEDGE

Core Solver	UEDGE

Sequential coupling

- Processor use limited by maximum number of procs of weakest link
- Might use fewer processors
- Longer time to solution

Concurrent coupling

- Processor use is the sum of the individual components
- May need more processors
- Shorter time to solution

Components that play well together have matched R_{ec}

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT PSFC
BERKELEY LAB
IU
NYU
Lodestar

$$R_{ec} = \frac{T_{\exp}}{T_{comp}}$$

- Ratio of experimental time to computation
- For physics studies, R_{ec} not restricted (provide computations finish within limits of human patience, lifetime, …)
- For control, need $R_{ec} < 1$

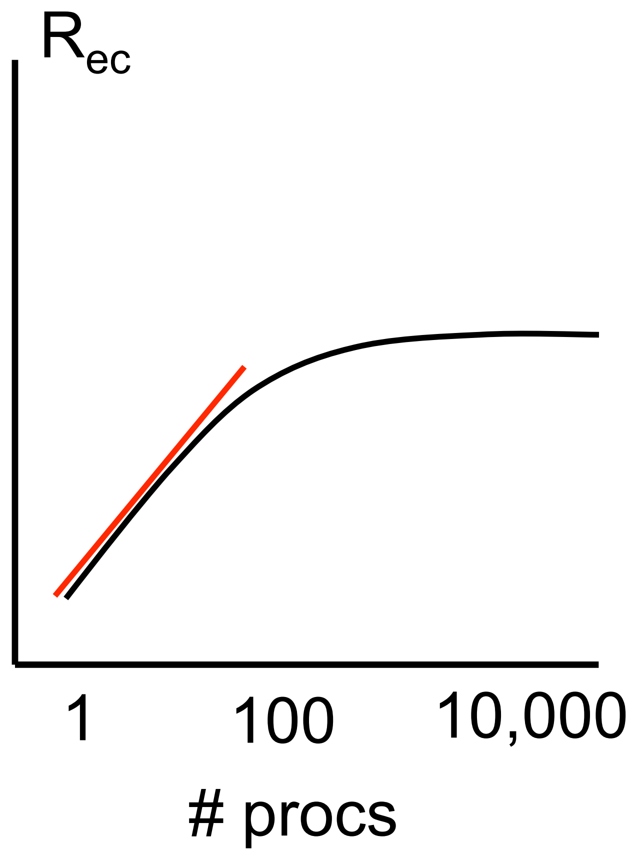

FACETS has determined the components that will play well together

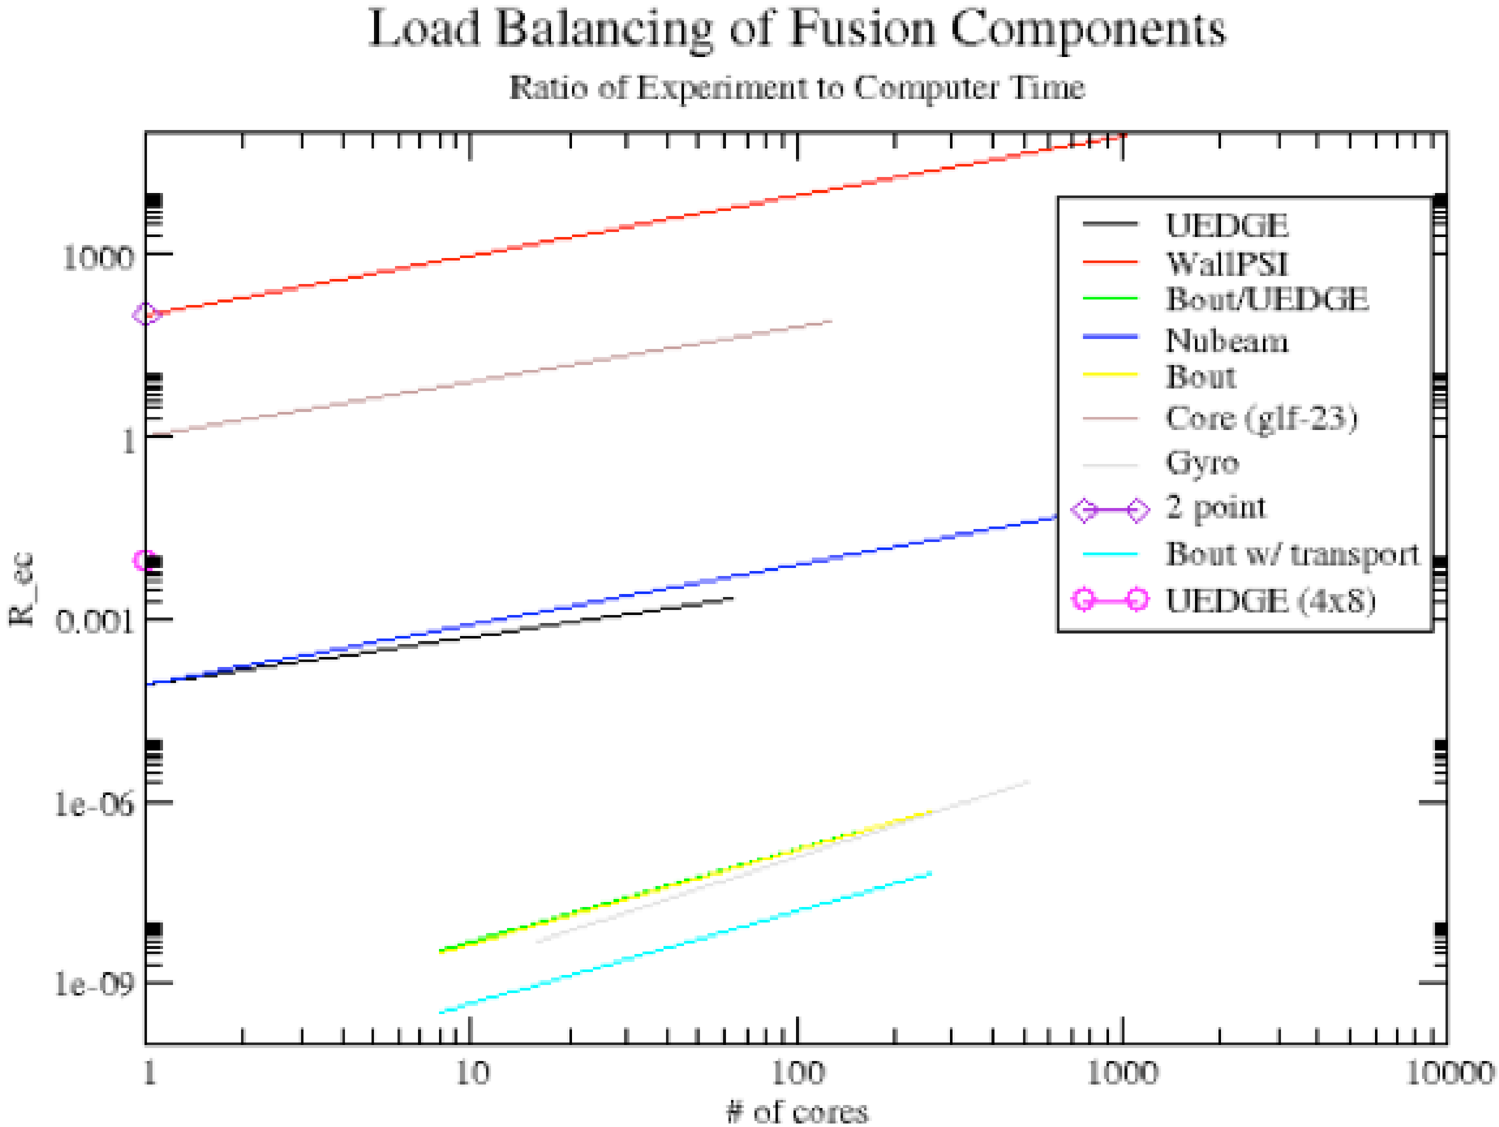

TECH-X, Argonne National Laboratory, Colorado State University, General Atomics, LLNL, ORNL, ParaTools, PPPL, UCSD, Columbia, MIT PSFC, Berkeley Lab, IU, NYU, Lodestar

Computational application development requires multiple efforts

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
MIT PSFC
Berkeley Lab
IU
NYU
Lodestar

- Framework
- Components
  - Core
  - Edge
  - Sources
  - Wall
  - Stability
  - Transport
- Algorithms
- Visualization
- Build systems
- Test harnesses
- Workflow development

Today

- New algorithms for core transport
- Core-transport verification
- Parallel edge
- First core-edge simulations
- Workflow

New core update algorithm: reduced Jacobian evals, quicker convergence, parallelism

- **Development required by lack of parallel core solver in community**
- **1D problems usually considered small, but computations per cell very intensive**
  - **GLF23**
  - **TGLF**
  - **Embedded GYRO**
- **Recent realizations (Jardin et al) that Newton approach to core solver beneficial**

New core solver has dramatically improved convergence.

- Core-Edge for DIII-D shot 118898.03400
- Nested iteration core solver

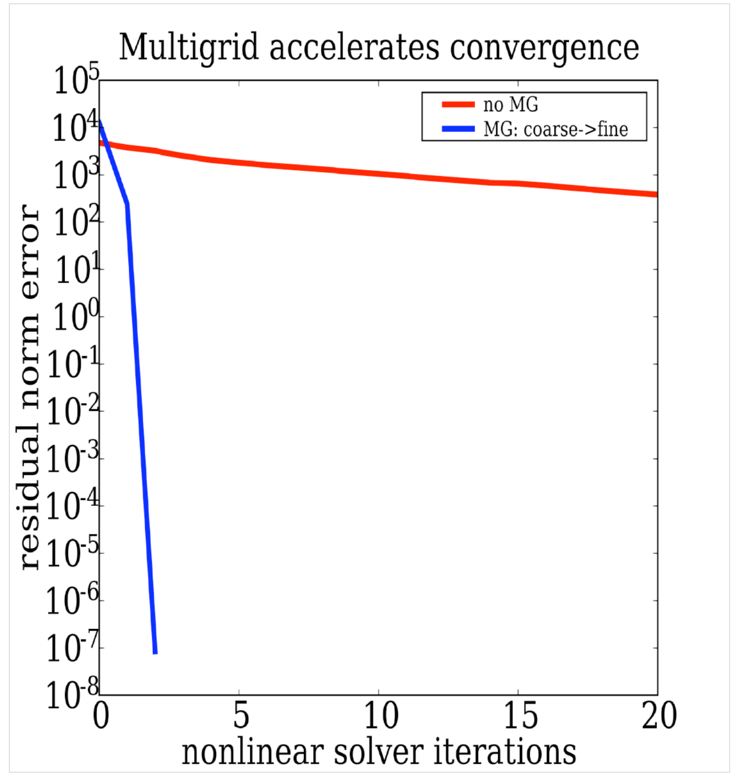

Ref: A. Pletzer, A. Hakim, M. Miah, J. Cary, S. Kruger, S. Vadlamani, and A. Pankin, "Benchmarking the parallel FACETS core solver," Poster presented at the 50th Annual Meeting of the Division of Plasma Physics, Dallas, TX, November 17-21, 2008.

Core solver parallelism getting to 10s or processors

ParaTools

IU
NYU
Lodestar

- **Performance analysis by ParaTools helped find roadblocks**

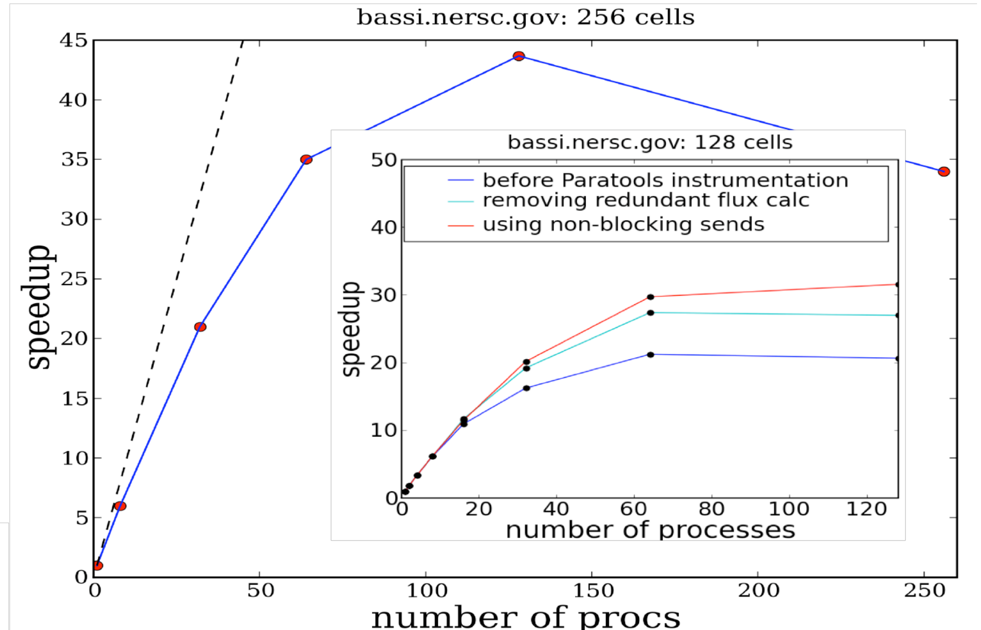

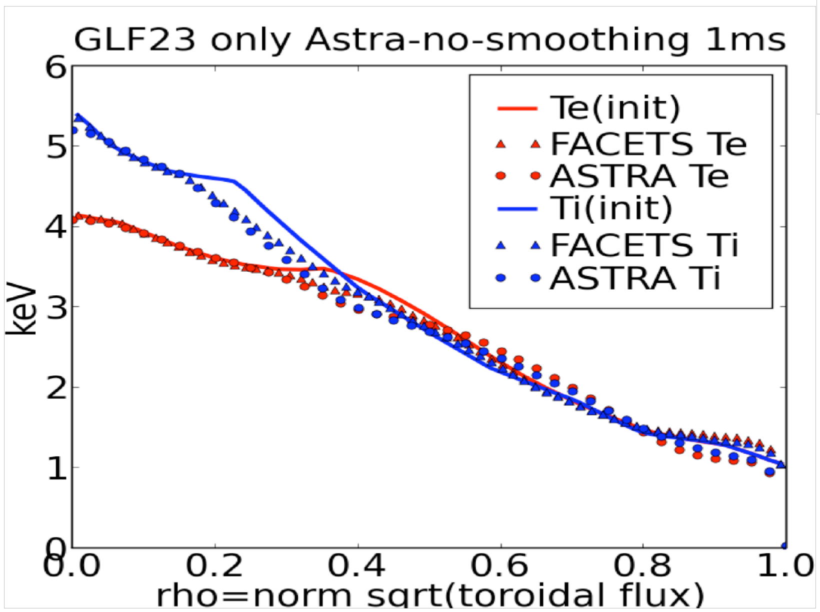

4x larger time step
30x from parallelism

Core verified through comparison with ASTRA

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT PSFC
BERKELEY LAB
IU
NYU
Lodestar

- Well validated case suggested by Alexei Pankin: JET shot 68875
- 10 ms interval [5.0895s, 5.0995s]
- Fixed equilibrium
- Evolve Te and Ti, plasma density kept fixed
- Sources: Neutral beam, radiative from UFILES
- Fluxes:
- Turbulent flux by GLF23
- Neoclassical term
- Equi-partition term (energy channel between ions and electrons)

Te and Ti develop sharp kink at 5s into the discharge, then relax

Electron/ion temperature gradient mode develops

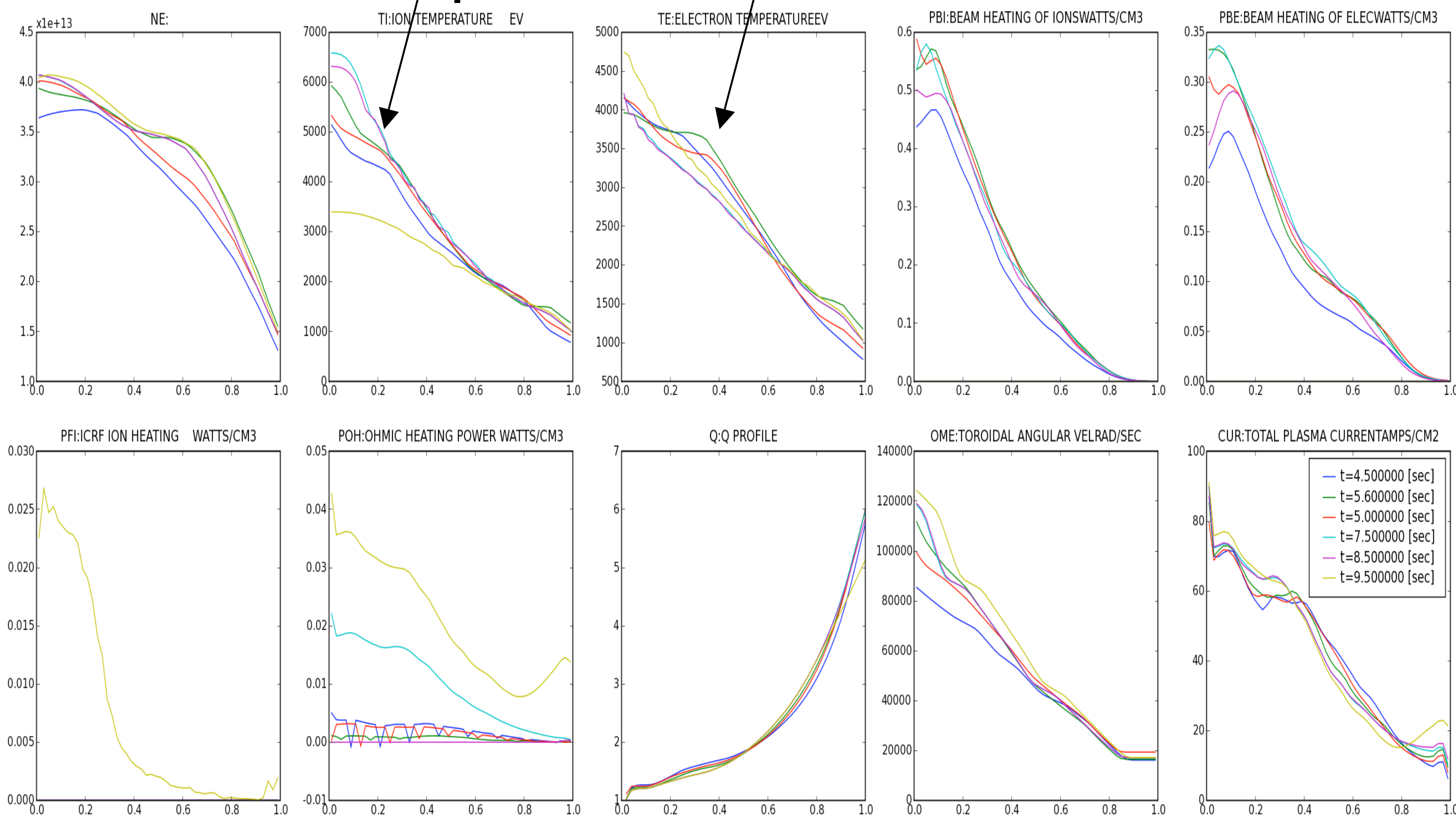

20

ASTRA vs FACETS: All terms in the transport equation are same

- Neoclassical diffusion (Chang-Hinton, 1986) copied “verbatim” from ASTRA to FACETS
- Equi-partition term copied from ASTRA to FACETS
- GLF23 from NTCC library (FACETS uses FMCFM interface)
- Sources prescribed from UFILES
- Some differences in applied boundary conditions (and the way GLF23 is turned off as one approaches the separatrix)

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
MIT PSFC
Berkeley Lab
IU
NYU
Lodestar

FACETS and ASTRA compare well without GLF23

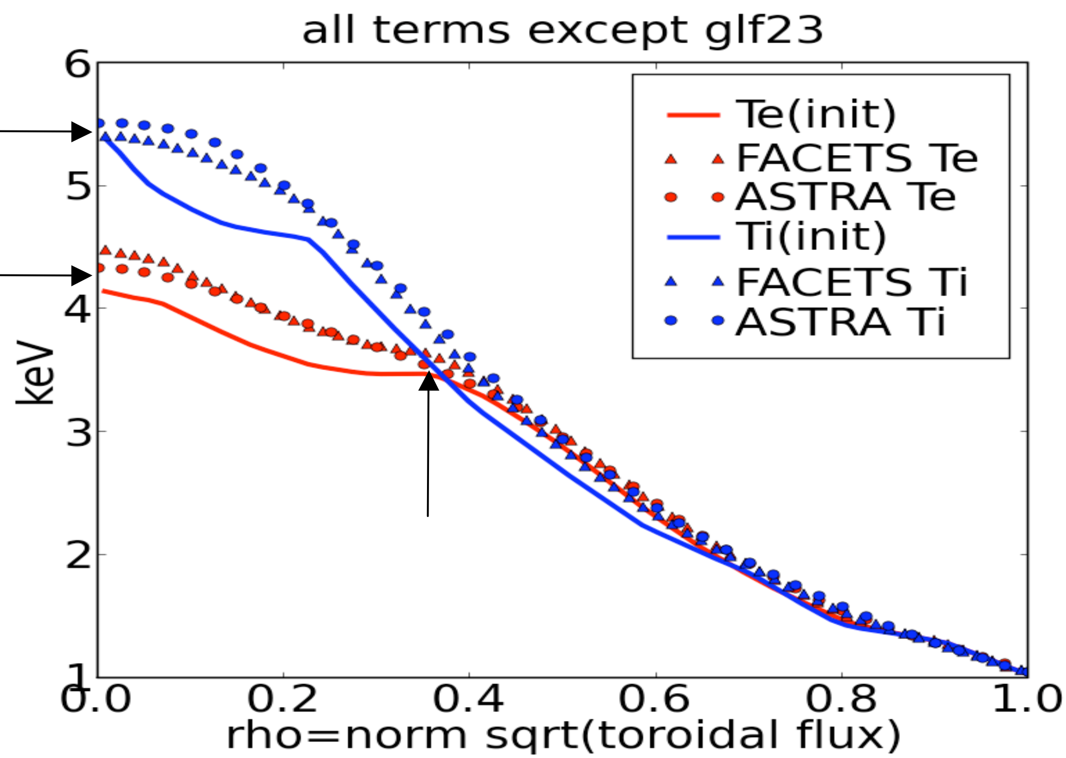

* Neutral beam heating causes Te to increase
* ASTRA has additional diffusivity @ rho=0.4
* Equi-partition term stronger in FACETS @ rho=0

Source terms alone: ASTRA has residual diffusivity on axis

ASTRA smooths profiles

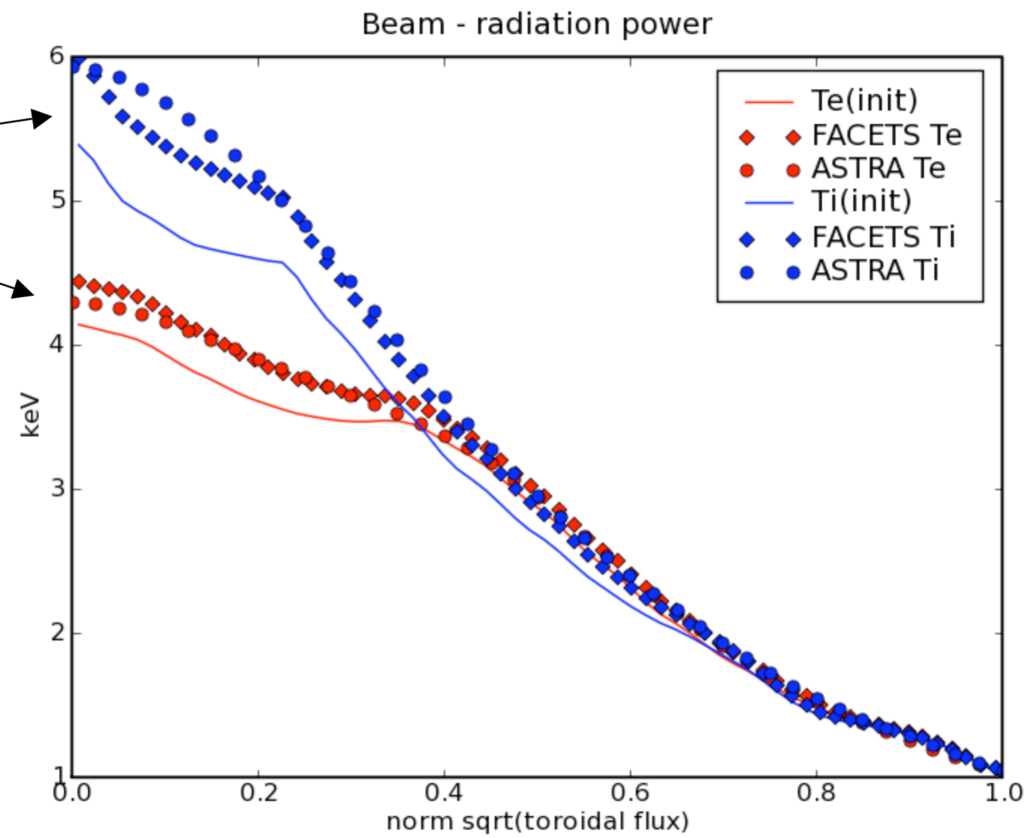

FACETS's Te, Ti shapes are preserved (no diffusivity)

23

Equi-partition solo

- This symmetric term reduces the difference in Ti – Te
- No diffusivity: profile shape is preserved
- No sources: increase in Te must match decrease in Ti

FACETS's profiles preserve initial shape

ASTRA does not quite conserve energy

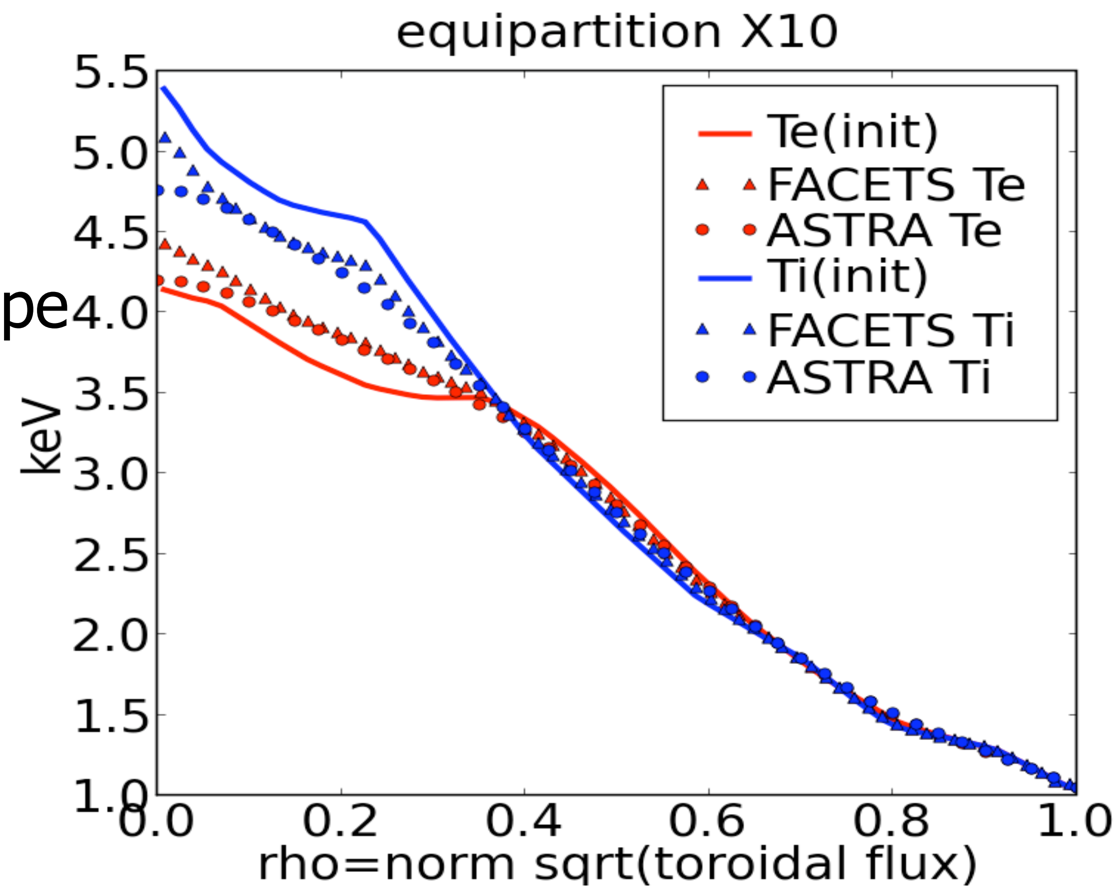

24

FACETS and ASTRA do not compare well when adding GLF23

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT
IU
NYU
Lodestar

All terms included

GLF23 responsible for most of the discrepancy

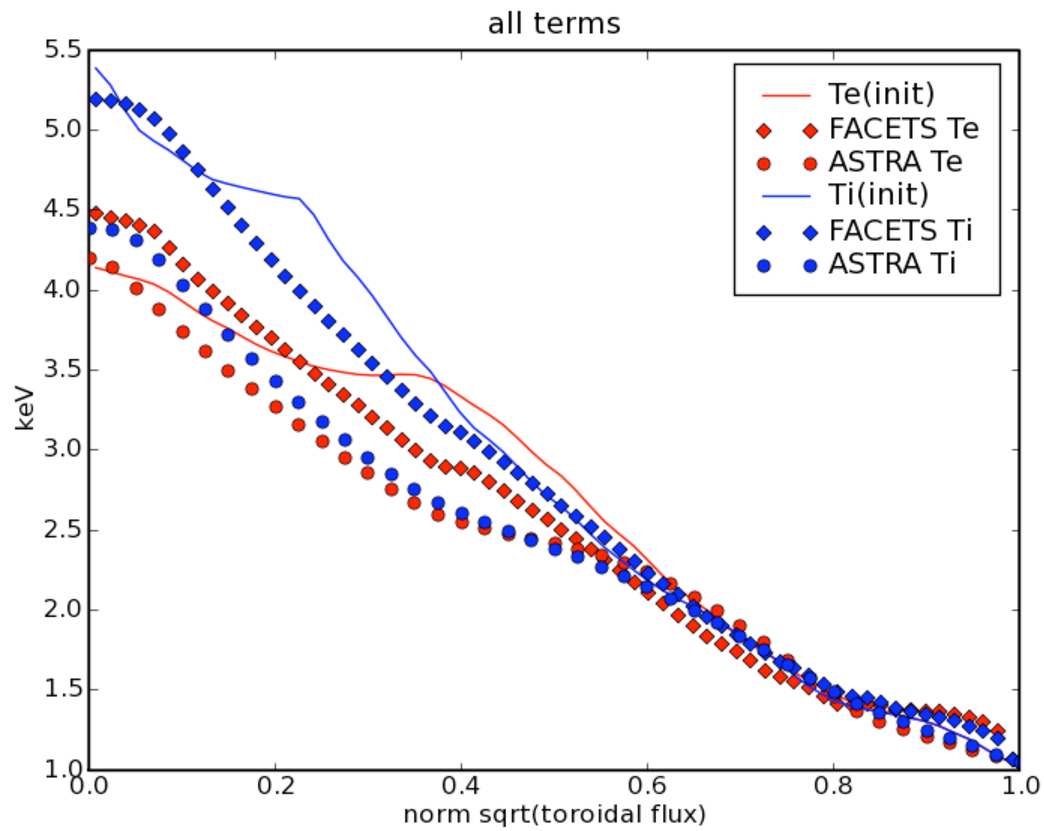

25

Must turn off smoothing of diffusion coefficients in ASTRA

IU
NYU
Lodestar

Now agreement much improved

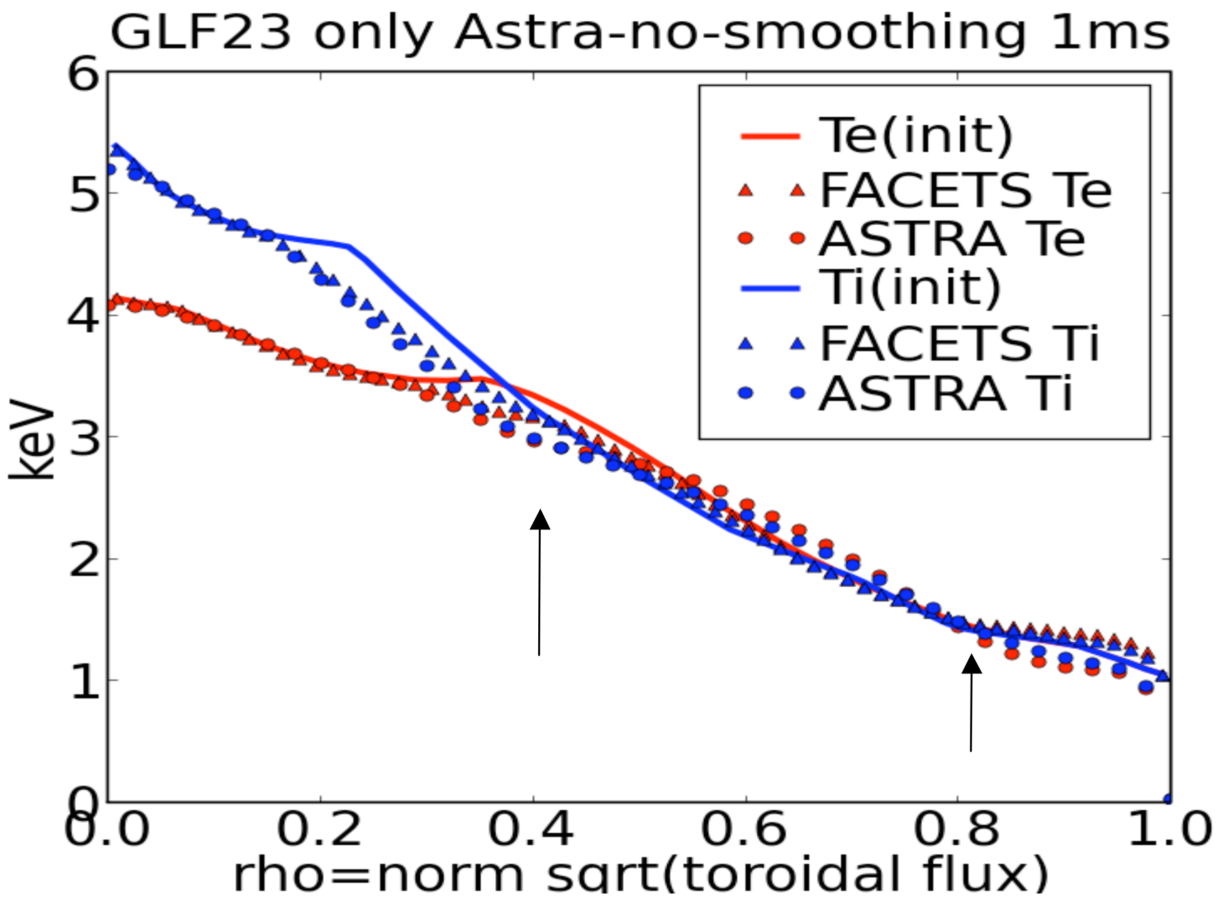

Note: diffusion at rho=0.4 for ASTRA but not FACETS

26

Verification results

- FACETS verification effort shows
  – Good agreement with ASTRA, except for:
    - Highly nonlinear GLF23 term. Smoothing must be turned off in ASTRA
    - Small violation of energy conservation in ASTRA
    - Small extra diffusivity in ASTRA
- ASTRA faster than FACETS for a single time step (4 x)
- FACETS has better numerical stability (so can take 30x larger steps). Potential for 300x wall clock time improvement over ASTRA in parallel

First core-edge simulations now being carried out

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
DSFC
IU
NYU
Lodestar

- Core-Edge for DIII-D shot 118898.03400
- Nested iteration core solver

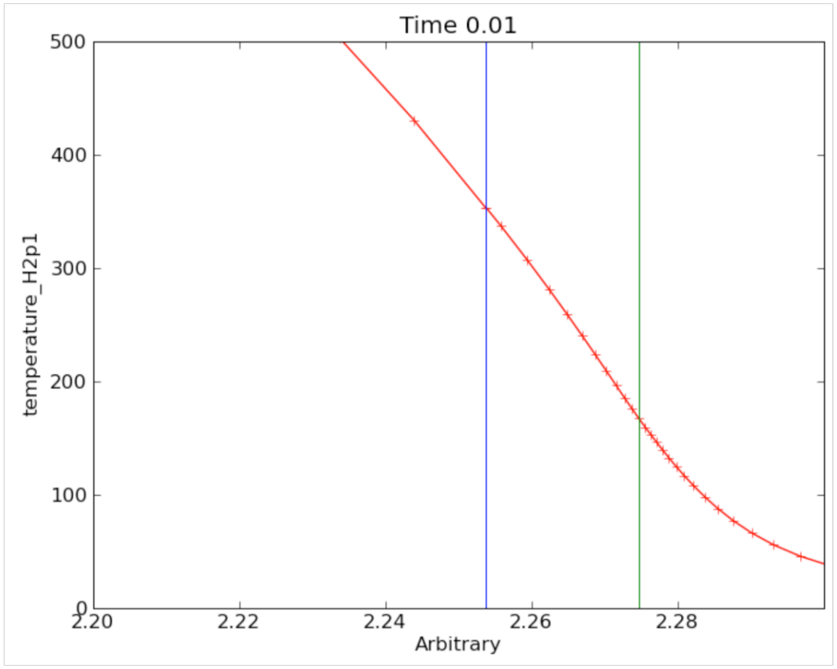

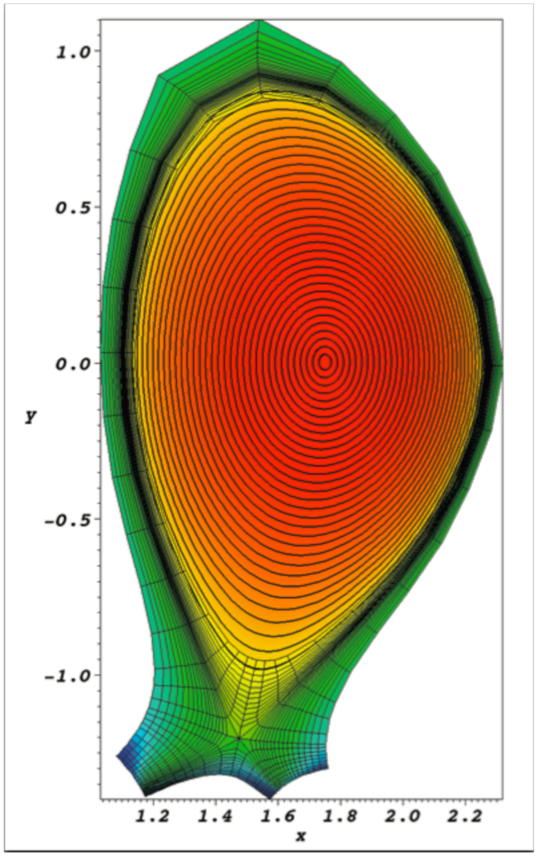

FACETS provided faster, parallel version of UEDGE, new version of wall model

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
MIT PSFC
Berkeley Lab
IU
NYU
Lodestar

- As modeled with WallPSI, static retention is via collisional production of Broken Bond traps followed by population by hydrogen.

- UEDGE parallel speedup using PETSc

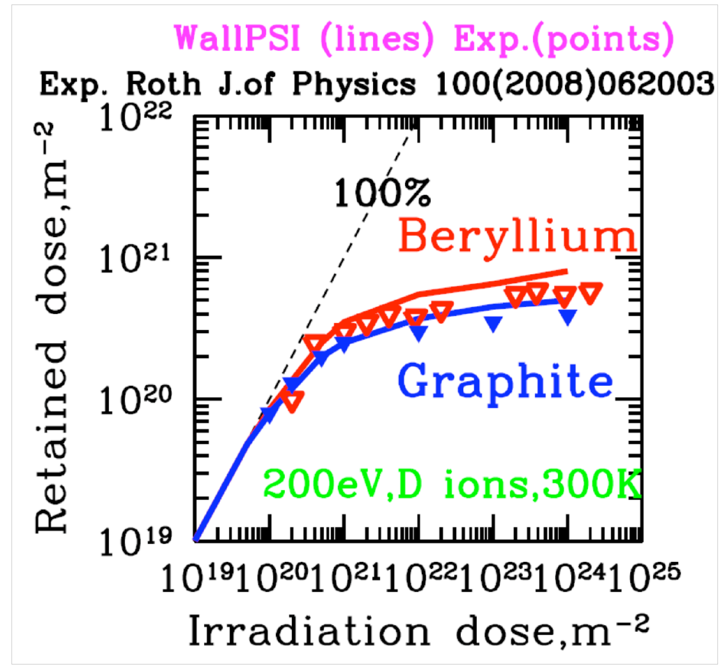

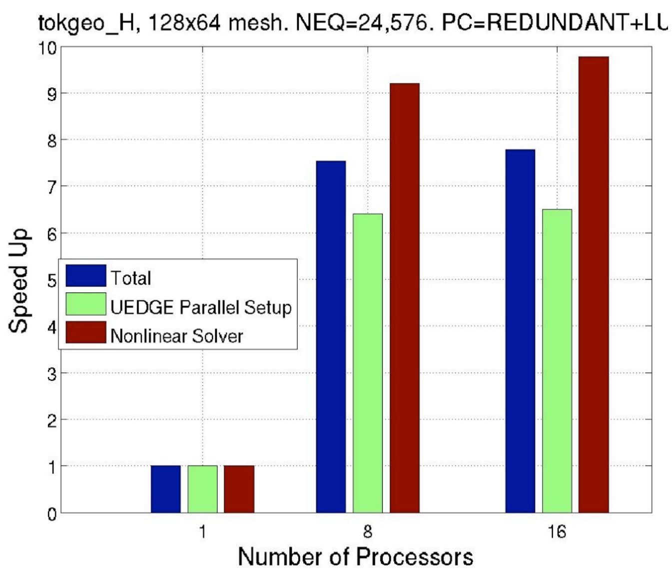

Workflow worked out, first physics

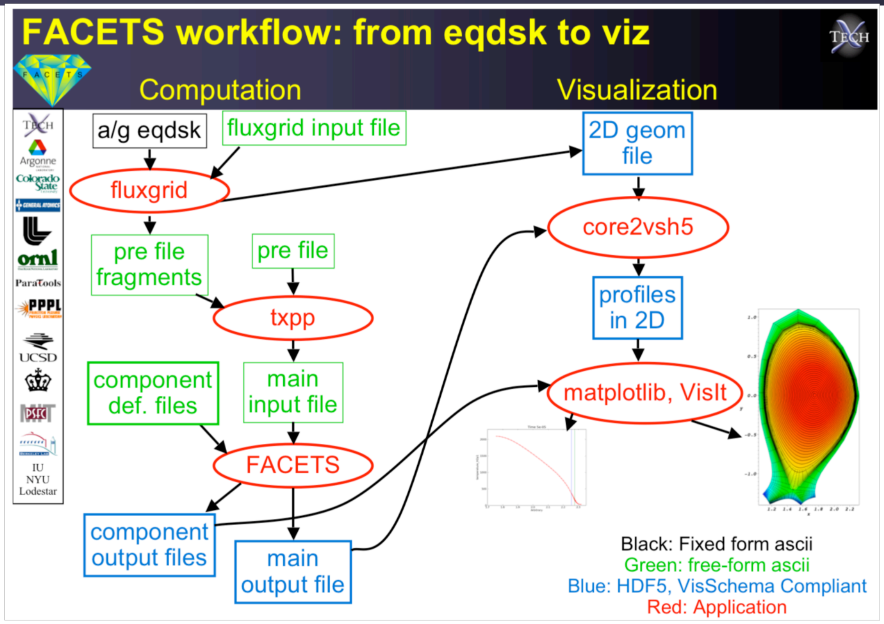

- Comparison of L-mode core build-up profiles vs. expt, with simple UEDGE model (no drifts, const. transport coeffs. from matching expt.) and specified ion beam-heating profiles
- Repeat with fuller UEDGE model
- Power scan, $T_{e,core}$, $T_{e,ped}$ vs expt.
- Fluxes to plasma facing components vs. expt.
- Core buildup profiles vs. expt. With Monte Carlo neutrals (NUBEAM) providing ion heating in core.

Entire system has been built and tested on many systems

- Many Mac Laptops
- Linux workstations - 32/64bit including multi-core
  – General Atomics
  – Lehigh
  – Tech-X
  – ANL
  – LLNL
- Linux clusters
  – PPPL
  – Tech-X
  – U. Oregon
- Bassi (NERSC - AIX)
- In process:
  – Franklin (NERSC - Cray)
  – Intrepic (ALCF - Blue Gene)

IU
NYU
Lodestar

Visualization changing from pretty to useful

- Example: cavity viz showing relevant physics

IU
NYU
Lodestar

FACETS contributing to the development of VizSchema, a markup for scientific data

- https://ice.txcorp.com/trac/vizschema

IU
NYU
Lodestar

Next level of parallelism requires varied communication patterns.

IU
NYU
Lodestar

Do codes have a finite life? parallelism

- Machines have a finite life because construction tools and methods change
- Is there any example of a computational application that transitioned from serial to parallel?
- Serial to parallel requires completely new thining on application organization
  - Start with analysis of communication patterns
- MPI coding is like returning to assembly language

TECH
Argonne
Colorado State
GENERAL ATOMICS
ParaTools
PPPL
UCSD
IU
NYU
Lodestar

With large-code development, many changes

- Separation of builders from users
- Long author lists
- Engineering
  - Design before building
  - Testing
    - Regression
    - Acceptance
  - Performance analysis
  - Upgrades
  - Procedures (workflows)

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
IU
NYU
Lodestar

To make matters worse: many-core!

TECH
Argonne
Colorado State
General Atomics
ornl
ParaTools
PPPL
UCSD
IU
NYU
Lodestar

- Nvidia TESLA
  - XXX cores

- For some applications, software layers may assist (GPULib)
- Should not expect (but can hope for) languages to eliminate the pain
  - We still code to MPI (apologies to Global Arrays, HPC, HPF, …, but that is the truth)
  - We did not predict many-core; what is next?

Lessons learned from initial work that are relevant for FSP

TECH
Argonne
Colorado State
General Atomics
LL
ornl
ParaTools
PPPL
UCSD
MIT PSFC
Berkeley Lab
IU
NYU
Lodestar

- All proto-FSPs face generic coupling issues
  - Interface definitions
  - Mathematical formulation
  - Human communication challenges
  - Generation of "glue code" for componentization
  - Dependency minimization
- In-memory coupling additional challenges beyond file-based coupling
  - Interlanguage interoperability
  - Compiler consistency for fortran
  - Build complexity (link lines, build provenance, …)
  - Symbol collisions
    - commonly included libs such as lsode, blas, …
    - Other: Serial and parallel PETSc

Lessons learned from initial work that are relevant for FSP (continued)

- **Lessons learned**
  - Webex has been crucial
  - Having enough funding for responsible leads (e.g., Livermore leading the edge)
  - Uniformity of code access, builds, testing.
  - Wiki for documentation is helpful
  - Regression testing is critical
  - Maintaining debug capability of individual components (both standalone and within framework)
  - Visualization of coupled components
    - Need both quick and dirty, and award winning viz

Lessons learned from initial work that are relevant for FSP (continued)

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT PSFC
BERKELEY LAB
IU
NYU
Lodestar

- It is difficult to estimate time for code componentization. Red flags:
  - How extensively portable the code is initially
    - 32bit/64bit
    - Dynamic/shared
    - Compilers (esp. fortran)
    - Operating systems
  - Does it use explicity typing or rely on flags to promote to correct type
  - To what extent is code generation used
  - Flexibility with regards to I/O file names
  - Global variables can cause problems (common block as interface)
- FACETS has benefited from the earlier NTCC effort
  - Disentanglement of NUBEAM
  - Availability of flux models for FMCFM with documented inputs/outputs

FACETS has made significant progress in its period of performance

TECH
Argonne
Colorado State
GENERAL ATOMICS
ornl
ParaTools
PPPL
UCSD
MIT PSFC
BERKELEY LAB
IU
NYU
Lodestar

- First parallel core transport solver written and validated
- UEDGE made a parallel component with new solvers, new capabilities, and better portability
- NUBEAM componentized and working within FACETS framework
- GYRO componentized and made available through same interface as the currently working GLF23, MMM95 (NCLASS will be available as well)
- Wall module developed - componentization underway
- Preliminary physics calculations using new core-edge coupled simulation capability