The Advanced Tokamak Modeling Environment (AToM-2019) for Fusion Plasmas

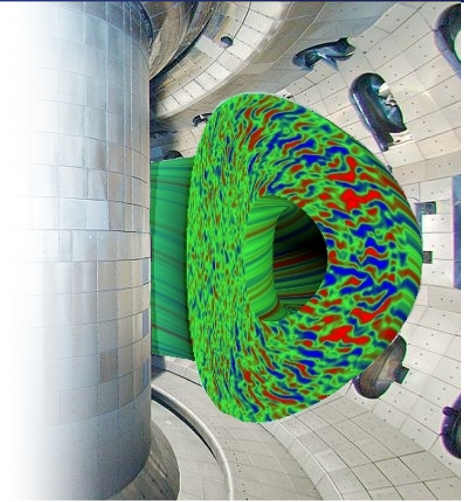

by
J. Candy[1] on behalf of the **AToM team**[2]

[1]General Atomics, San Diego, CA
[2]See presentation

Presented at the
2019 SciDAC-4 PI Meeting
Rockville, MD
16-18 July 2019

AToM (2017-2022) Research Thrusts

- AToM^0 was a **3-year SciDAC-3 project** (2014-2017)
- AToM is a **5-year SciDAC-4 project** (2017-2022)
- 6 thrusts address **code integration** and **workflow management**:

AToM (2017-2022) Research Thrusts

- AToM^0 was a **3-year SciDAC-3 project** (2014-2017)
- AToM is a **5-year SciDAC-4 project** (2017-2022)
- 6 thrusts address **code integration** and **workflow management**:

1. AToM environment, performance and packaging
2. Physics component integration
3. Validation and uncertainty quantification
4. Physics and scenario exploration
5. Data and metadata management

AToM (2017-2022) Research Thrusts

- AToM^0 was a **3-year SciDAC-3 project** (2014-2017)
- AToM is a **5-year SciDAC-4 project** (2017-2022)
- 6 thrusts address **code integration** and **workflow management**:

1. AToM environment, performance and packaging
2. Physics component integration
3. Validation and uncertainty quantification
4. Physics and scenario exploration
5. Data and metadata management
6. Liaisons to SciDAC partnerships

AToM Team

Institutional Principal Investigators (FES)

Jeff Candy	General Atomics
Mikhail Dorf	Lawrence Livermore National Laboratory
J-M. Park	Oak Ridge National Laboratory
Chris Holland	University of California, San Diego
Jai Sachdev	Princeton Plasma Physics Laboratory

Institutional Principal Investigators (ASCR)

David Bernholdt	Oak Ridge National Laboratory
Milo Dorr	Lawrence Livermore National Laboratory
David Schissel	General Atomics

AToM Team

Funded collaborators (subcontractors in ***green****)*

O. Meneghini, S. Smith, P. Snyder, D. Eldon, E. Belli, M. Kostuk	GA
W. Elwasif, M. Cianciosa, D.L. Green, G. Fann, K. Law, D. Batchelor	ORNL
N. Howard	MIT
D. Orlov	UCSD
J. Sachdev	PPPL
M. Umansky	LLNL
P. Bonoli	MIT
Y. Chen	UC Boulder
R. Kalling	Kalling Software
A. Pankin	Tech-X

Outline of this talk

Please note

- Skip many introductory slides presented in previous years

Outline of this talk

Please note

- Skip many introductory slides presented in previous years
- Skip energetic particle, edge kinetic (COGENT), RF, IPS activities

Outline of this talk

Please note

- Skip many introductory slides presented in previous years
- Skip energetic particle, edge kinetic (COGENT), RF, IPS activities
- Focus on present **integration activities**

Outline of this talk

Please note

- Skip many introductory slides presented in previous years
- Skip energetic particle, edge kinetic (COGENT), RF, IPS activities
- Focus on present **integration activities**
- For past presentations see:

`scidac.github.io/atom/literature.html`

Outline of this talk

1. Project scope and vision
2. Data management: OMAS and IMAS
3. Examples of fast-prediction workflows
4. Merging and regression
5. Fidelity hierarchy
6. Fusion simulation use cases

Outline of this talk

1. Project scope and vision
2. Data management: OMAS and IMAS
3. Examples of fast-prediction workflows
4. Merging and regression
5. Fidelity hierarchy
6. Fusion simulation use cases

Outline of this talk

1. Project scope and vision
2. Data management: OMAS and IMAS
3. Examples of fast-prediction workflows
4. Merging and regression
5. Fidelity hierarchy
6. Fusion simulation use cases

Outline of this talk

1. Project scope and vision
2. Data management: OMAS and IMAS
3. Examples of fast-prediction workflows
4. **Merging and regression**
5. Fidelity hierarchy
6. Fusion simulation use cases

Outline of this talk

1. Project scope and vision
2. Data management: OMAS and IMAS
3. Examples of fast-prediction workflows
4. Merging and regression
5. **Fidelity hierarchy**
6. Fusion simulation use cases

Outline of this talk

Project scope and vision

AToM Modeling Scope and Vision

Present-day tokamaks	Upcoming burning plasma	Future reactor design
DIII-D	**ITER**	**DEMO**

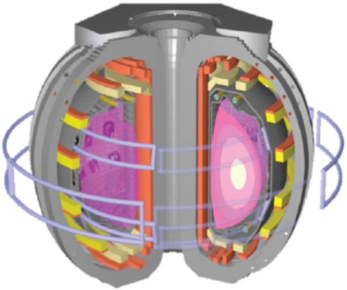

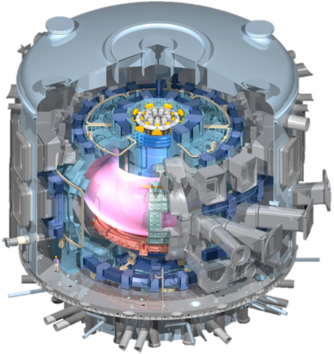

AToM Conceptual Structure

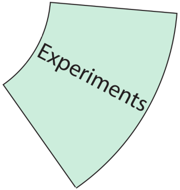

1. Access to experimental data
2. Outreach (liaisons) to other SciDACs
3. Verification and validation, UQ, machine learning
4. Support HPC components
5. Framework provides glue

AToM Conceptual Structure

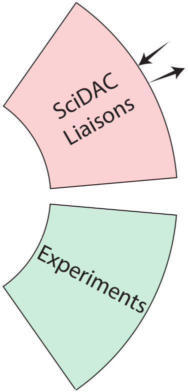

1. Access to experimental data
2. Outreach (liaisons) to other SciDACs
3. Verification and validation, UQ, machine learning
4. Support HPC components
5. Framework provides glue

AToM Conceptual Structure

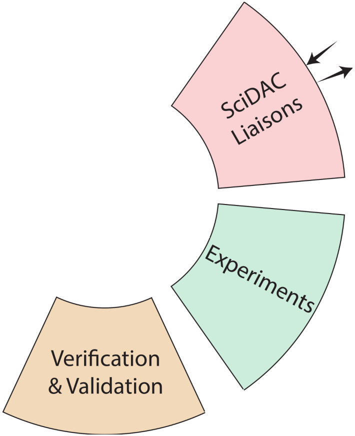

1. Access to experimental data
2. Outreach (liaisons) to other SciDACs
3. Verification and validation, UQ, machine learning
4. Support HPC components
5. Framework provides glue

AToM Conceptual Structure

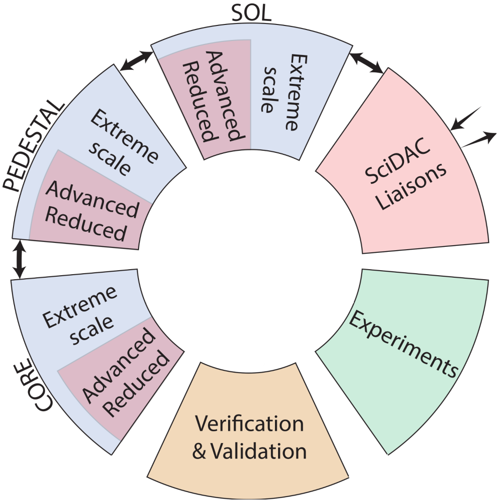

1. Access to experimental data
2. Outreach (liaisons) to other SciDACs
3. Verification and validation, UQ, machine learning
4. Support HPC components
5. Framework provides glue

AToM Conceptual Structure

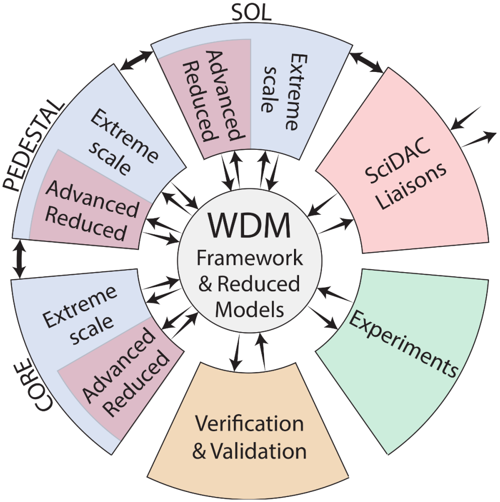

1. Access to experimental data
2. Outreach (liaisons) to other SciDACs
3. Verification and validation, UQ, machine learning
4. Support HPC components
5. Framework provides glue

Adapted from Fig. 24 of
Report of the Workshop on Integrated Simulations for Magnetic Fusion Energy Sciences (June 2-4, 2015)

Data management: OMAS and IMAS

ITER Integrated Modeling and Analysis Suite (IMAS)

Data schema and storage infrastructure to support ITER operations

imas

- ambitious European effort to build a standard fusion format
- **IMAS data schema: Interface Data Structure (IDS)**
  - Data organized into 48 IDSs (tree) for different physics
  - Store both experimental and simulated data
- **IMAS storage infrastructure: Access Layer (AL)**
  - Layer that passes data between components and to/from storage
  - C/C++, Fortran (F95), Java, Matlab, Python
- **Significant effort** underway to make IMAS a standard
  - ITER data will be available **only through IMAS**
  - European tokamaks making notable progress adopting IMAS

IMAS is challenging for developers

IMAS is challenging for developers

1. access layer (AL) **tightly linked** to data-schema
2. requires **recompile of IMAS** and components for each data-schema release
3. proposed new HDC API to be independent of data-schema
4. IMAS infrastructure is **heavy**, and **hard to install** and manage
5. API does not provide any useful functionality besides data storage

Solution: do not rely exclusively on IMAS

- Shortcomings and evolving IMAS infrastructure demand a solution
  - want to **decouple AToM** environment from IMAS
  - want to **ensure IMAS** compatibility

Solution: do not rely exclusively on IMAS

- Shortcomings and evolving IMAS infrastructure demand a solution
  - want to **decouple AToM** environment from IMAS
  - want to **ensure IMAS** compatibility
- Solution:

OMAS

  - python package to organize data in **compliance with IMAS schema**
  - **fast, stable, portable**

Simplified use of IMAS through OMAS (O. Meneghini, S. Smith)

via access to ITER IMAS database (requires ITER account)

```python
import omas
ods = omas.load_omas_iter_scenario(shot=130010, run=1)
plot( ods['core_profiles']['profiles_1d'][256]['electrons']['pressure'] )
```

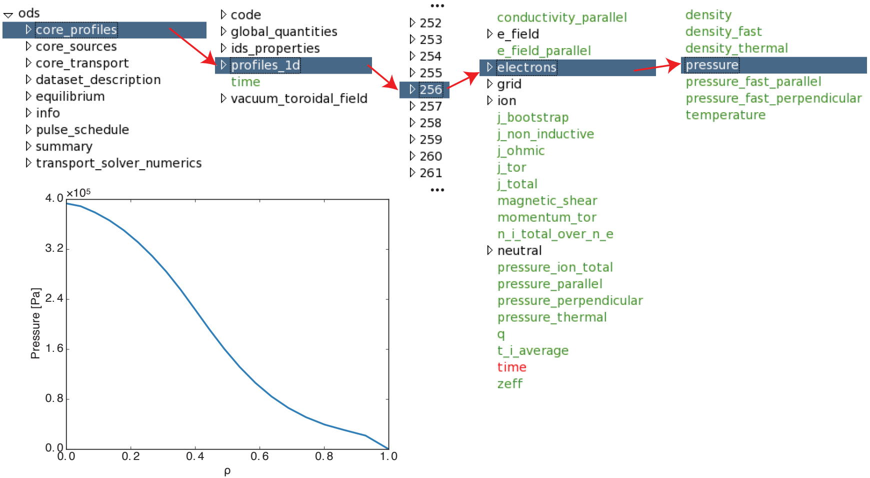

OMFIT classes `.to_omas()` and `.from_omas()`

provide an effective way to simplify code integration

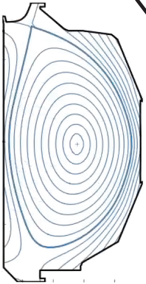

OMFIT + OMAS
Mapper

```
from omfit.classes.omfit_eqdsk import OMFITgeqdsk, OMFITsrc
from omas import *

# read gEQDSK file in OMFIT
eq = OMFITgeqdsk(OMFITsrc + '/../samples/g133221.01000')

# convert gEQDSK to OMAS data structure
ods = eq.to_omas()

# save OMAS data structure to IMAS
paths = save_omas_imas(ods, tokamak='DIII-D', new=True)

# load OMAS data structure from IMAS
ods1 = load_omas_imas(user, tokamak='DIII-D', paths=paths)

# generate gEQDSK file from OMAS data structure
eq1 = OMFITgeqdsk('g133221.02000').from_omas(ods1)

# save gEQDSK file
eq1.deploy()
```

Many legacy codes share the same file formats

imas

Open source: `pip install omas`

documentation at `http://gafusion.github.io/omas`

Installation — OMAS

https://gafusion.github.io/omas/install.html

OMAS | Concept | Examples | Installation | ITER | OMFIT | Data schema | API | In this page

Google: omas imas

Installation

OMAS runs both with *Python2* and *Python3*.

Pypi: To install OMAS with pip (for users):

```
pip install --upgrade omas
```

where *upgrade* is used to update the omas installation to the latest version.

The development version of omas can also be installed with pip:

```
pip install --upgrade -e git+git@github.com:gafusion/omas#egg=omas
```

Conda: To install OMAS with conda (for users):

```
conda install -c conda-forge omas
conda update  -c conda-forge omas
```

GitHub To clone OMAS from GitHub (for developers):

```
git clone git@github.com:gafusion/omas.git
cd omas
sudo pip install --upgrade -e .[build_structures, build_documentation]     # Add this `omas` directory to your $PYTHONPATH
                                                                           # The [build_structures,build_documentation] options
                                                                           # install packages required for extra development purposes
```

OMFIT STEP module (O. Meneghini, others)

IMAS-compliant modeling workflows

- OMFIT **STEP module**

OMFIT STEP module (O. Meneghini, others)

IMAS-compliant modeling workflows

- OMFIT **STEP module**
- couples components (*steps*) to support workflows
  - open-loop prediction
  - control
  - optimization
- **Data exchanged** between steps via OMAS
- Can write data to IMAS at any stage

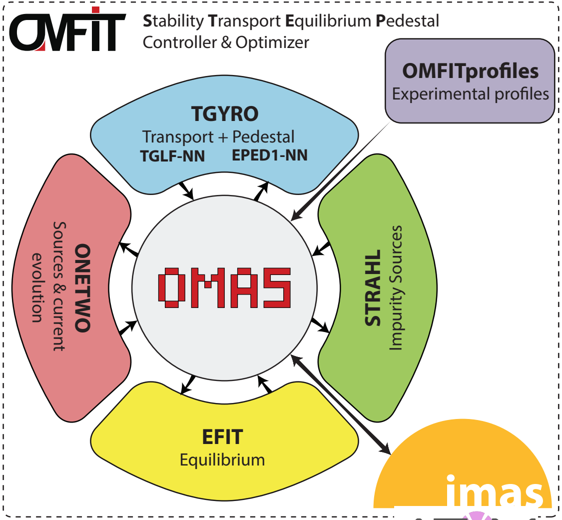

TRANSP Collaboration (J. Sachdev, B. Grierson)

PPPL support of AToM/GACODE

- AToM seeks **synergy with TRANSP** usage
- PPPL staff assist with
  - maintenance of **TRANSP OMFIT module**
  - development of the **Plasma State** code including OMAS/IMAS translators
- AToM **reduced model** development feeds into TRANSP
- TRANSP modules to be deployed via **git**, accessible by community
  - PSPLINE - nearly complete
  - Plasma State - ongoing investigation

TRANSP Collaboration

PPPL support of AToM/GACODE

TRANSP usage: over 62k simulations performed since 2010

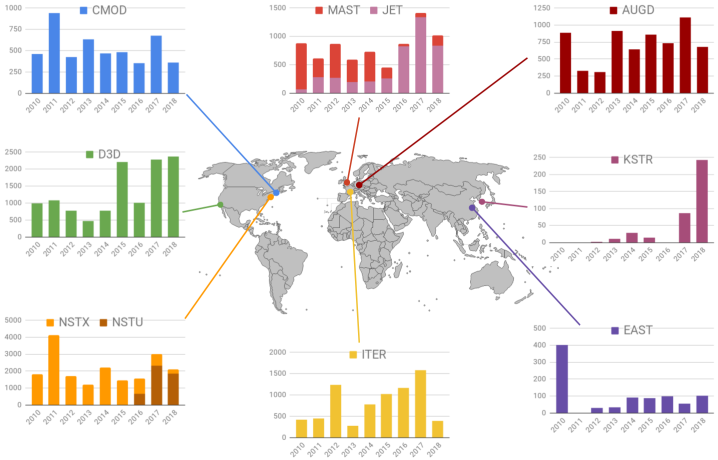

TRANSP used with several other devices including: ARIES, DEMO, FNSF, HI2A/HL2M, IGTR, JT60, KDMO, LTX, MST, RXFM, STEP, TCV, TFTR, WRK

Examples of fast-prediction workflows

Tokamak physics spans multiple space/timescales

Core-edge-SOL (CESOL) region coupling

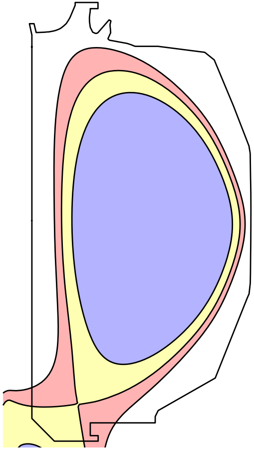

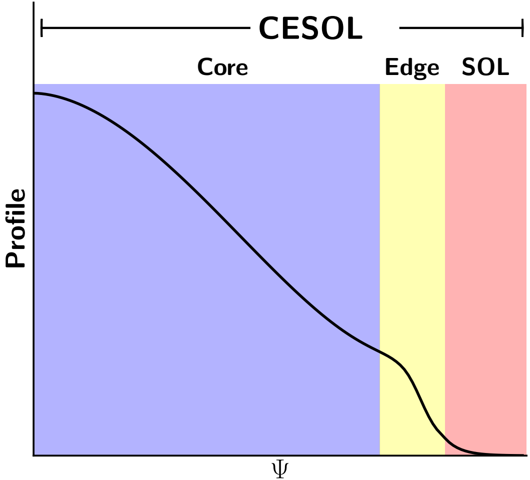

AToM OMFIT-TGYRO CE(sol) ITER prrdictions

impurities (Z_{eff}) improve performance despite core dilution

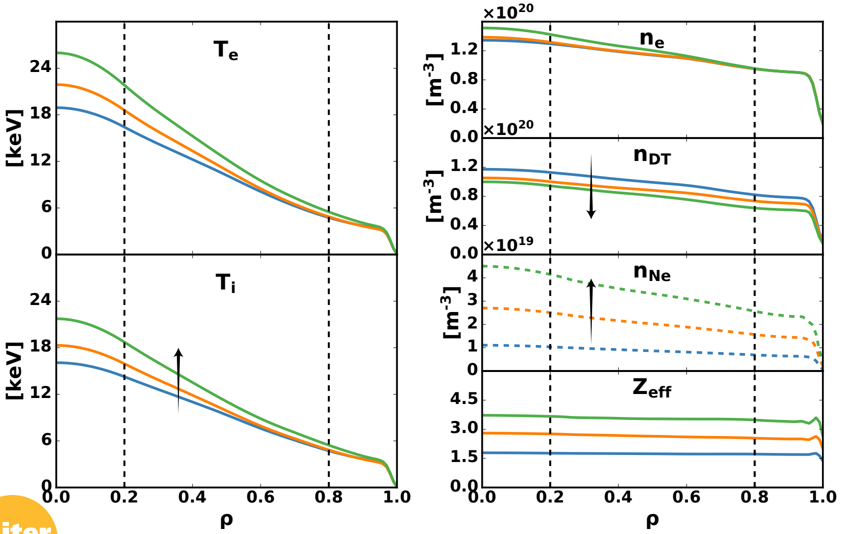

AToM IPS-FASTRAN CESOL is being Extended to Wall (JM. Park)

2D Impurity transport in entire tokamak volume

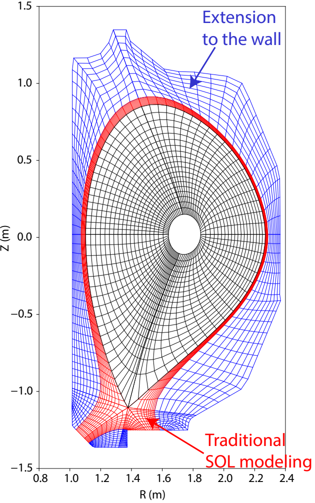

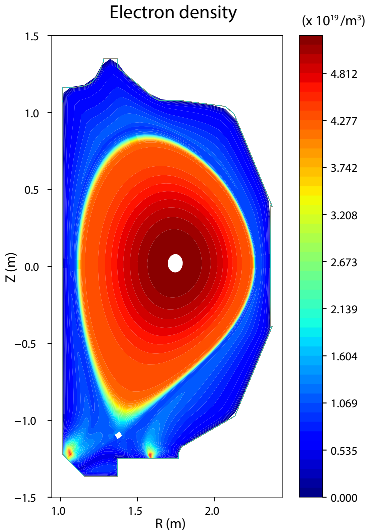

- FASTRAN (core)
- C2 (edge-SOL)
- Extending C2 mesh to wall

Fast predictive capability 1: TAUENN (J. McClenaghan)

(0+ϵ)-D capability

- Equilibrium and sources based on input global parameters:

$$R, a, B_T, I_p, n_{e,ped}, P_{aux}, \kappa, \delta, q_0, Z_{eff}$$

$$q = q_0 + (q_{95} - q_0)(r/a)^2$$
$$q_{95} = 5a^2BS/(RI_p)$$
$$\kappa = \kappa_o + (\kappa_s - \kappa_0)(r/a)^2$$
$$\delta = \delta_s(r/a)$$

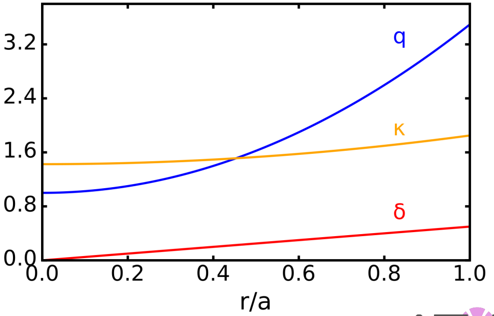

Fast predictive capability 1: TAUENN (J. McClenaghan)

(0+ε)-D capability

- Equilibrium and sources based on input global parameters:

$$R, a, B_T, I_p, n_{e,ped}, P_{aux}, \kappa, \delta, q_0, Z_{eff}$$

- Pedestal shape $(r = r_{ped})$ from EPED-NN: **Set α_{EPED}**

$$T(r) = f\left(r, \alpha_{TGLF}, \alpha_{EPED}\right)$$

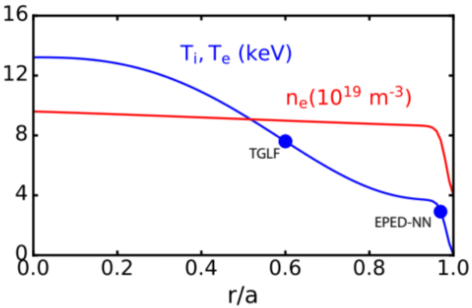

ATOM

Fast predictive capability 1: TAUENN (J. McClenaghan)

(0+ϵ)-D capability

- Equilibrium and sources based on input global parameters:

$$R, a, B_T, I_p, n_{e,ped}, P_{aux}, \kappa, \delta, q_0, Z_{eff}$$

- Pedestal shape $(r = r_{ped})$ from EPED-NN: **Set** α_{EPED}
- Core shape $(r = 0.6a)$ from TGLF: **Set** α_{TGLF}

$$T(r) = f\left(r, \alpha_{TGLF}, \alpha_{EPED}\right)$$

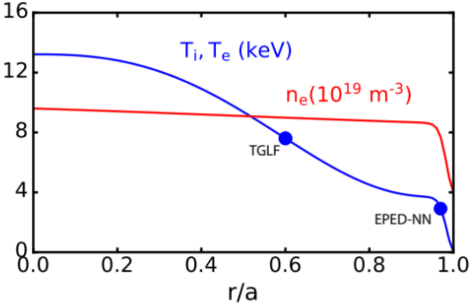

Fast predictive capability 1: TAUENN (J. McClenaghan)

(0+ε)-D capability

- Equilibrium and sources based on input global parameters:

$$R, a, B_T, I_p, n_{e,\text{ped}}, P_\text{aux}, \kappa, \delta, q_0, Z_\text{eff}$$

- Pedestal shape $(r = r_\text{ped})$ from EPED-NN: **Set α_EPED**
- Core shape $(r = 0.6a)$ from TGLF: **Set α_TGLF**
- Match TGLF flux at $r/a = 0.6$

$$T(r) = f\,(r, \alpha_\text{TGLF}, \alpha_\text{EPED})$$

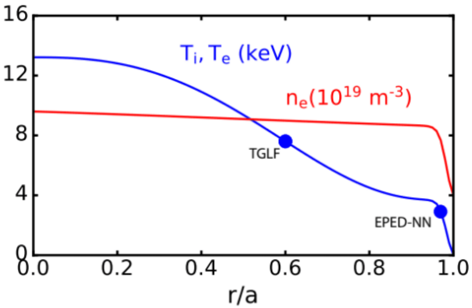

ATOM

Fast predictive capability 1: TAUENN (J. McClenagham)

Performance relative to power-law fit

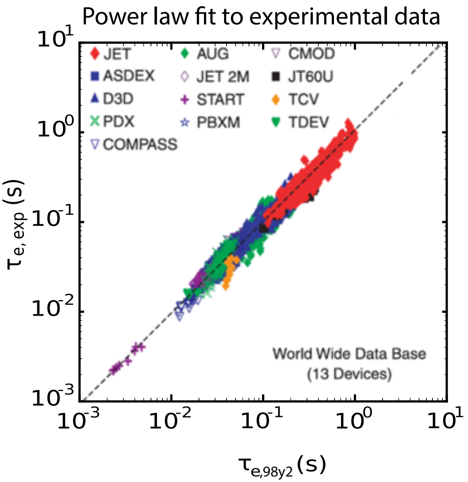

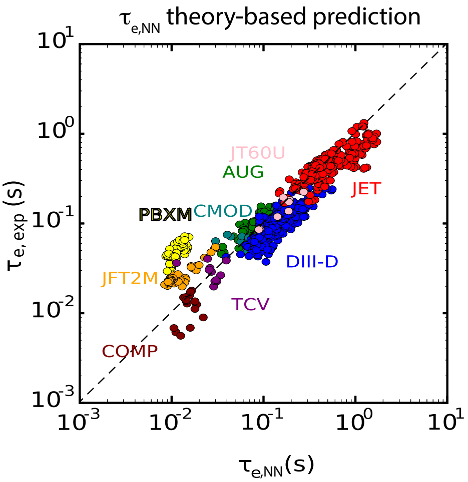

Fast predictive capability 1: TAUENN (J. McClenaghan)

Performance relative to power-law fit

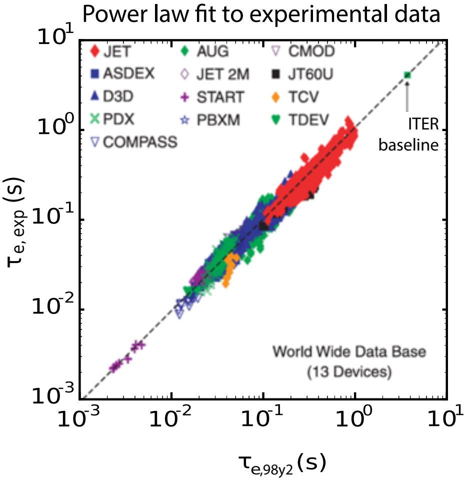

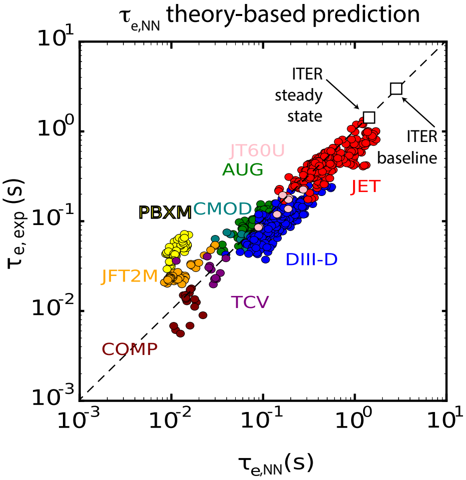

Fast predictive capability 2: MODEL-PROFILES (J. Kinsey)

1D model: scaling law plus fast equilibrium/heating

- Rapid equilibrium/profile estimation
- Data exchange:

  GACODE expro interface
- Sources:

  Ohmic, NBI, radiation
- Equilibrium:

  VMOMS
- Profiles:
  - Rotation (DeGrassie), scaling (Thomsen-Cordey)
- Compute time **less than 30 seconds**

Workflow diagram

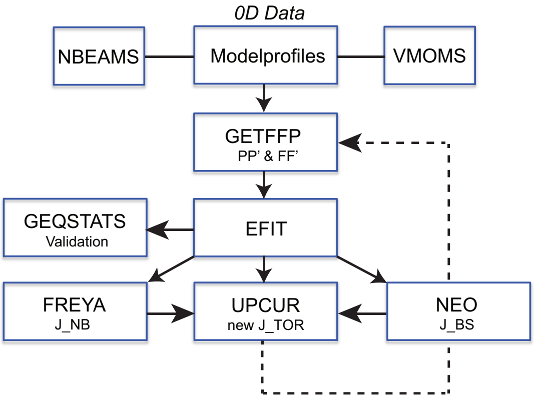

Fast predictive capability 2: MODEL-PROFILES (J. Kinsey)

1D model: scaling law plus fast equilibrium/heating

- Rapid equilibrium/profile estimation
- Data exchange:
  - GACODE expro interface
- Sources:
  - Ohmic, NBI, radiation
- Equilibrium:
  - VMOMS
- Profiles:
  - Rotation (DeGrassie), scaling (Thomsen-Cordey)
- Compute time **less than 30 seconds**

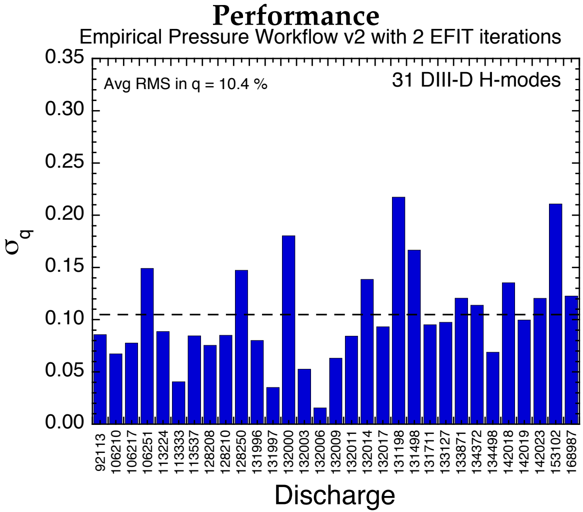

Merging and regression

OMFIT Update Workflow (R. Kalling)

OMFIT is mission-critical and complex

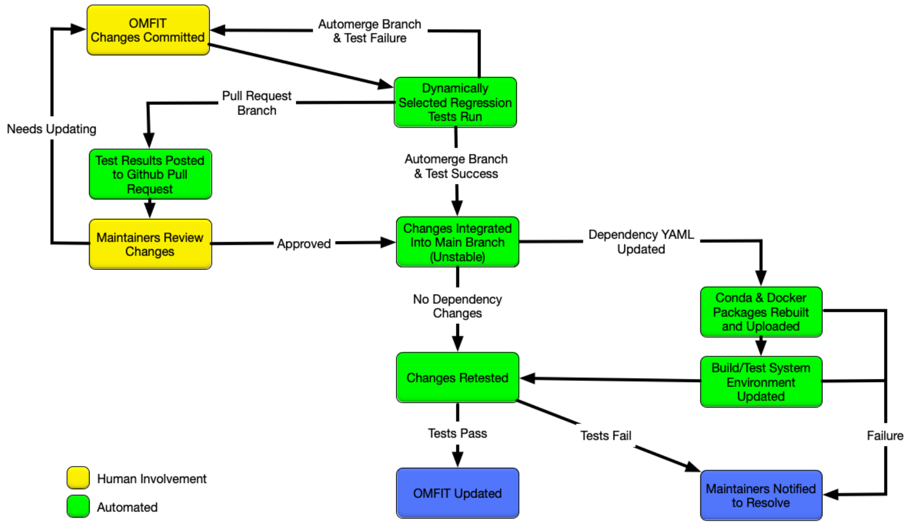

OMFIT Update Workflow (R. Kalling)

Automatic regression and update

- The automerge branch allows trusted developers to integrate features with a greatly **reduced risk of broken code** being distributed
- The regression test system uses a **labeling mechanism** to exclude tests that are not appropriate for a given test environment
  - for example, no gui, or a specific server not being available
- Regression test system automatically **selects relevant tests** given code changes in a commit to reduce testing time
- **Automatic package rebuild/upload** allows installations to stay up-to-date whenever a developer changes OMFIT dependency requirements

Fidelity hierarchy

Fidelity hierarchy

Key theme for the future of whole-device modeling

Fidelity Hierarchy is CRITICAL

Range of models from leadership codes to REDUCED MODELS

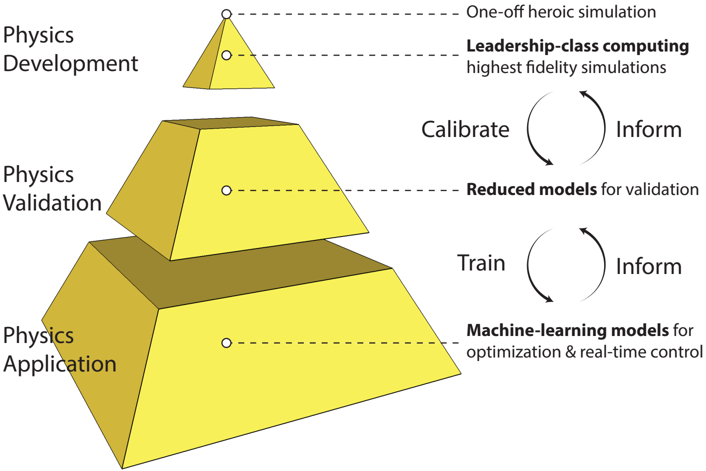

CGYRO ITER Baseline Simulation (N. Howard, C. Holland)

Electron-ion multiscale resolution

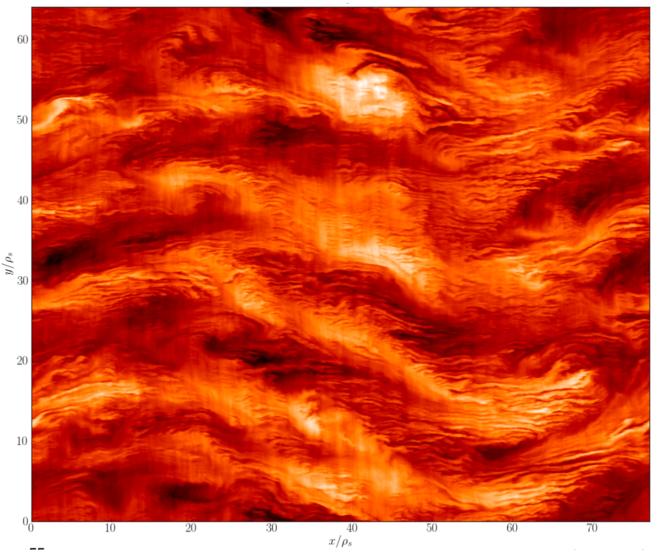

CGYRO ITER Baseline Simulation (N. Howard, C. Holland)

Electron-ion multiscale resolution

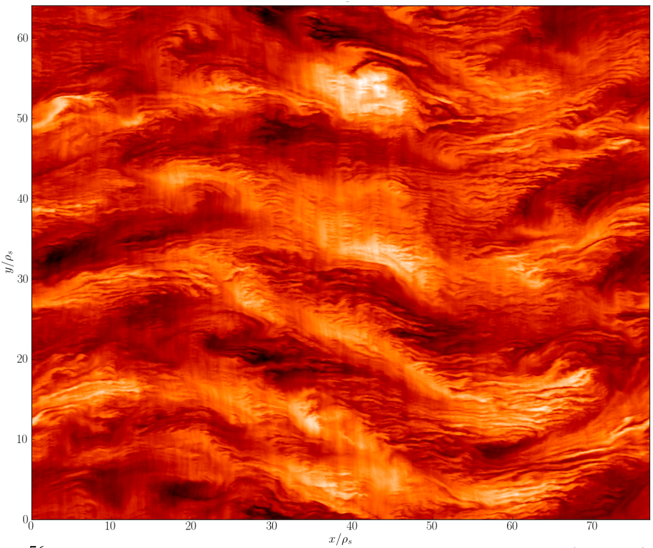

- $k_x \rho_i \leqslant 92$, $k_y \rho_i \leqslant 54$
- **Highest GK resolution ever**
- 280M core hrs on Titan
- Δt: 220K FFTs of length 5.6M
- 500K Δt

Microtearing Turbulence (X. Jian, C. Holland)

Discovery of MTM-driven transport in high-β_p discharges

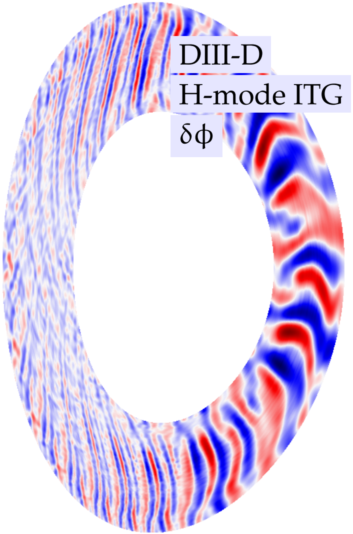

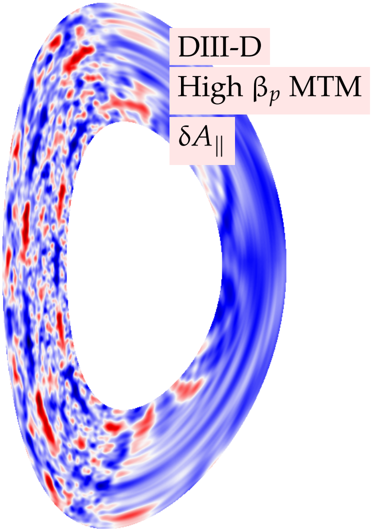

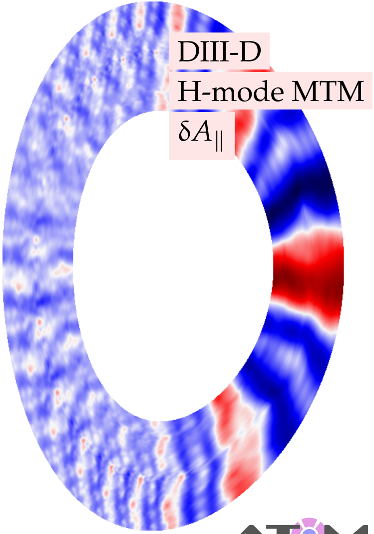

Performance on Leadership Systems (I. Sfilogoi, G. Fann)

GPU systems lack compute-communicate balance of CPU systems

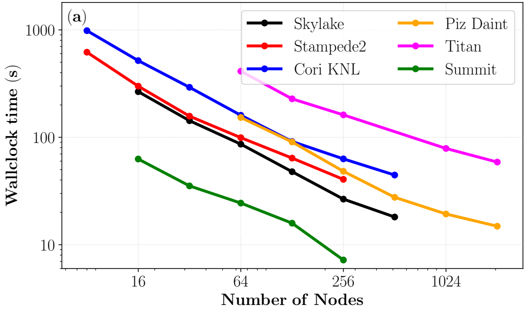

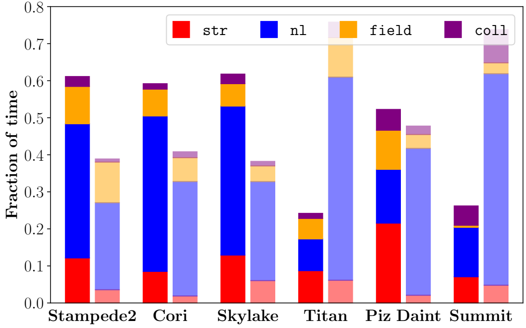

- LEFT: 6-platform (3 **CPU** + 3 **GPU**) strong-scaling comparison
- RIGHT: kernel-level analyses (compute time, communicate time)

OpenMP performance on KNL

High throughput/productivity on Cori

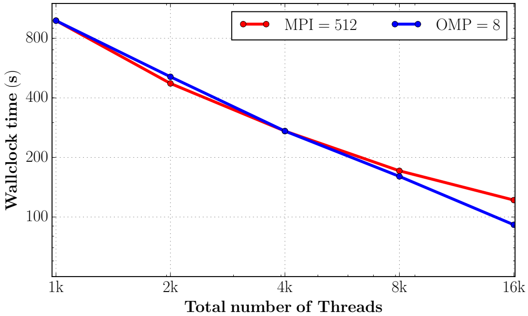

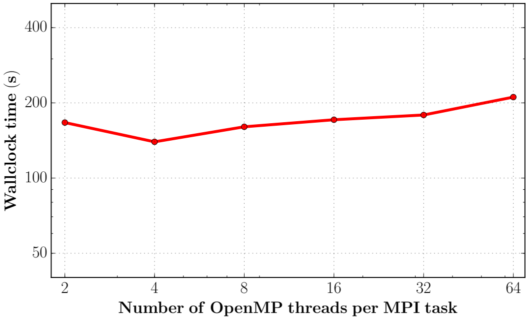

- **Significant loop-level work** (for OMP) left after MPI distribution
- Excellent scaling up to 64 OpenMP threads on Cori
- Hundreds of **CGYRO database** runs completed in 2019 ⟶ **reduced model**

GPU Performance (via cuFFT and GPUDirect MPI)

New Optimizations by Igor Sfiligoi (SDSC)

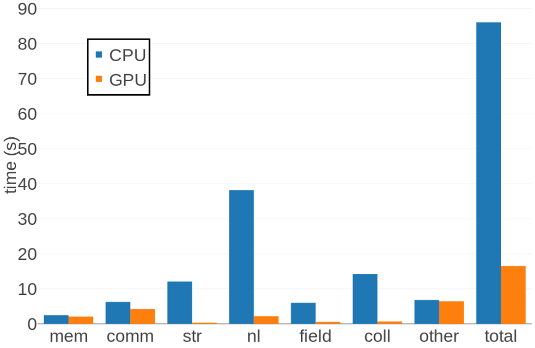

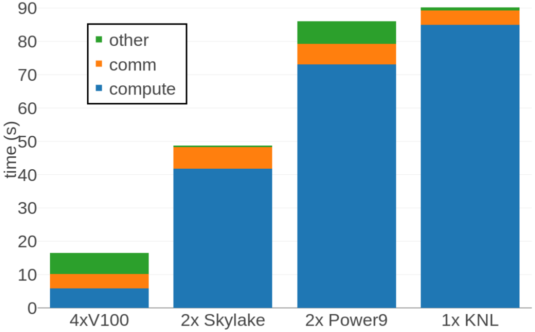

- Expensive kernels (`nl,coll`) **remarkably fast** on GPU
- Summit has **GPUDirect MPI bug** (IBM Spectrum MPI, libcoll complex)
- Underway: **embedded/adaptive timestepping** (G. Fann, ORNL)

Next level of fidelity hierarchy: TGLF

Centerpiece of all AToM predictive modeling workflows

- **Reduced model** of nonlinear gyrokinetic flux (1 second at 1 radial point)

Next level of fidelity hierarchy: TGLF

Centerpiece of all AToM predictive modeling workflows

- **Reduced model** of nonlinear gyrokinetic flux (1 second at 1 radial point)
- Determines **quality of profile prediction**

Next level of fidelity hierarchy: TGLF

Centerpiece of all AToM predictive modeling workflows

- **Reduced model** of nonlinear gyrokinetic flux (1 second at 1 radial point)
- Determines **quality of profile prediction**
- TGLF is the heart of AToM profile-prediction capability
  - linear gyro-Landau-fluid **eigenvalue solver**
  - coupled with sophisticated **saturation rule**
  - evaluate quasilinear fluxes over range $0.1 < k_\theta \rho_i < 24$

Next level of fidelity hierarchy: TGLF

Centerpiece of all AToM predictive modeling workflows

- **Reduced model** of nonlinear gyrokinetic flux (1 second at 1 radial point)
- Determines **quality of profile prediction**
- TGLF is the heart of AToM profile-prediction capability
  - linear gyro-Landau-fluid **eigenvalue solver**
  - coupled with sophisticated **saturation rule**
  - evaluate quasilinear fluxes over range $0.1 < k_\theta \rho_i < 24$
- Saturated potential intensity
  - derived from a **database** of nonlinear GYRO simulations
  - database resolves only long-wavelength turbulence: $k_\theta \rho_i < 1$

Next level of fidelity hierarchy: TGLF

Centerpiece of all AToM predictive modeling workflows

- **Reduced model** of nonlinear gyrokinetic flux (1 second at 1 radial point)
- Determines **quality of profile prediction**
- TGLF is the heart of AToM profile-prediction capability
  - linear gyro-Landau-fluid **eigenvalue solver**
  - coupled with sophisticated **saturation rule**
  - evaluate quasilinear fluxes over range $0.1 < k_\theta \rho_i < 24$
- Saturated potential intensity
  - derived from a **database** of nonlinear GYRO simulations
  - database resolves only long-wavelength turbulence: $k_\theta \rho_i < 1$
- 10^7 **times faster** than nonlinear gyrokinetics

TGLF

Ongoing calibration with CGYRO leadership simulations

- **Theory-based approach** – must be calibrated with nonlinear simulations
- Predictions validated with ITPA database
- **Discrepancies:** L-mode edge, EM saturation
- **CGYRO multiscale simulations needed**

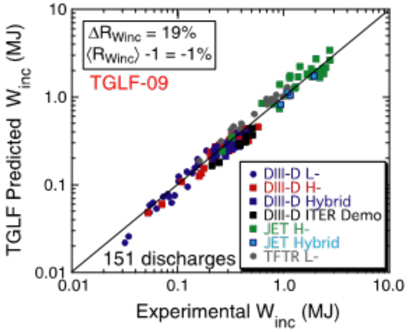

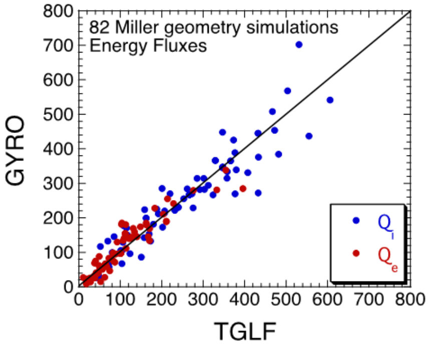

Fusion simulation use cases

AToM Use Cases (C. Holland, P. Bonoli, others)

Coordination of validation/physics studies

- **Observation**
  - Most every modeling effort eventually **settles on certain set of inputs**

AToM Use Cases (C. Holland, P. Bonoli, others)

Coordination of validation/physics studies

- **Observation**
  - Most every modeling effort eventually **settles on certain set of inputs**
  - these inputs provide **benchmark points** for regression testing and physics studies
  - can be, but not necessarily, drawn from actual experiments
- **Plan**
  - organize AToM validation and scenario modeling work about **uses cases**
  - provide comprehensive, organized, documented datasets
  -
- **Long-term vision**
  - development use cases through **iterative process**
  - start simple and **grow as needed** by maturity of physics and validation workflows
  - will grow to provide a **community knowledge-base**

AToM Use Cases

Tentative examples

Use case description	Machine	Shot	Time (ms)	B_T (T)	I_p (MA)	P_{RF} (MW)	P_{NBI} (MW)	T_{inj} (N-m)
L-mode shortfall	DIII-D	128913	1500 ± 100	2.1	1.0	0	2.6	2.14
ITER I_p ramp	DIII-D	161129	400 ± 30	2	0.5	0	1.5	1.1
ITER I_p ramp	DIII-D	161129	700 ± 30	2	0.8	0	1.6	1.3
ITER I_p ramp	DIII-D	161129	1500 ± 30	2	1.5	0	1.6	1.2
H-mode stiffness	DIII-D	145456	1775 ± 100	2.1	1.2	0	3.2	1.5
H-mode stiffness	DIII-D	145452	1665 ± 100	2.1	1.2	0	7.2	1.4
H-mode stiffness	DIII-D	145937	1825 ± 100	2.1	1.2	0	6.9	5.9
ITER baseline	DIII-D	153523	3380 ± 400	1.7	1.3	3.4	2.8	0.6
ITER baseline	DIII-D	155196	3000 ± 200	1.7	1.3	0	2.8	1.5
ITER baseline	DIII-D	155196	2200 ± 200	1.7	1.3	3.3	2.7	2.3
ITER baseline	DIII-D	171534	4200 ± 500	1.7	1.3	3.5	2.8	1.5
ELMy H-mode	C-Mod	1120815026	1025 ± 75	5.6	0.9	1.1	0	0
I-mode	C-Mod	1120907028	1005 ± 45	5.8	1.1	2.1	0	0
Inductive Q=10	ITER	---	---	5.2	15	17	33	34


FIN