A Sustainable Pathway Towards Successful Integrated Modeling

*S. P. Smith (GA), O. Meneghini (GA)
* Corresponding author: smithsp@fusion.gat.com

Request Oral Presentation: Yes
Primary Panel: G **Secondary Panel:** C

In this paper we provide a review of the authors' experience with the development of the OMFIT [gafusion.github.io/OMFIT-source] integrated modeling (IM) framework and the AToM project. In particular, we lay out what we believe are the key elements that are driving the success of these projects. The most important of these elements is what we call a `bottom-up` approach to integrated modeling, which has proven to be very effective at tackling practical and pressing physics problems, while securing a strong user base. Finally, we provide recommendations on how such an approach can be supported and encouraged throughout the Office of Fusion Energy Sciences (FES).

To start, we define **user adoption** and **scientific impact** (i.e. publications) as the prime metrics of OMFIT or AToM project success. Other metrics, such as the number of coupled codes and High Performance Computing (HPC) usage and scalability will naturally follow as user adoption and scientific impact increase. Rapid progress in developing OMFIT has been guided by this definition of success. Today OMFIT is an integrated modeling (IM) framework that is regularly used by over one hundred (and growing) fusion scientists worldwide in support of their research. Several major milestones of the AToM project were delivered after only six months. Through OMFIT, the capabilities and workflows enabled by the AToM project are made available to a large pool of scientists.

User adoption and scientific impact go hand in hand. Scientists (users) are the drivers for discovery, but without a record of proven scientific impact, they have no reason to adopt an IM suite. Since one cannot force scientists to adopt a specific framework, focus must be given right away on delivering a simple but working framework capable of tackling at least one physics application for which scientists have a tangible and pressing need. For OMFIT this was the generation of kinetic equilibrium reconstructions and edge-stability analyses of DIII-D plasmas (it's a good idea to engage with experimentalists, as they are the largest demographic and they provide a reality check against real world data). This automatically implies that a successful fusion science framework must be more than a mere computer-science project. Early users can then be involved from the early stages of the project, providing valuable feedback that can shape and improve the framework development. In striving to adapt to user requirements, the mission-critical applications are the first ones to be completed and refined. As new and improved features attract more scientists, a virtuous cycle can then be established that bootstraps the framework's adoption. The framework design and its capabilities can then organically evolve to support the progressive integration of the components that are required to solve problems of increasing complexity. Such an incremental approach is the same principle that is at the basis of an **agile development methodology [http://agilemethodology.org/]**.

Leveraging and integrating existing research within the US magnetic fusion program has been critical for creating solutions that have user traction and are at the forefront of fusion scientific research. The most important reasons to embrace such a philosophy are:

- To save time and effort by not reinventing the wheel, which ensures the IM development effort remains productive and sustainable with finite resources

- To take advantage of collective wisdom that is embodied in existing tools, which have been developed and validated over years
- To continue to develop and improve existing tools, without diverging development paths (forks)
- To minimize changes to users' existing research workflows (thus reducing "friction" and allowing the framework to easily and incrementally become part of the research)

These defining characteristics are what we call a **`bottom-up` approach to integrated modeling**.

Other important aspects that favor strong user adoption are the availability of **powerful GUIs**, and a **responsive user support** (an automatic email error report allows issues to be fixed usually within minutes). A public version of OMFIT is maintained and made available to scientists at GA, PPPL, ASIPP, and NERSC. In addition to the usual channels, for **outreach and documentation** we are maintaining an online website, offering hands-on workshops, and providing Google docs and YouTube videos for training. Collective and distributed software development, revision control, issue tracking, and the overall progress of the project are coordinated (even for multiple geographically separated institutions) with high efficiency by **taking advantage of modern online development tools** (namely GitHub). Both OMFIT and AToM feature a **nimble organizational structure,** which can respond, adapt, and refocus quickly as needed by a dynamic research environment.

Recommendations:

- User adoption and scientific impact as metrics for success
- Support rather than subvert existing integrated modeling research
  - Use existing codes, frameworks, and data structures and formats
  - Support integration among different frameworks: 1) they have different capabilities; 2) more efficient and sustainable than integrating tens of individual codes (e.g. OMFIT and IPS within AToM); 3) division of labor and expertise
  - Leverage integration in frameworks as an opportunity for performance engineering of legacy codes (e.g. IPS-EPED workflow within AToM).
- Promote an ecosystem fostering bottom-up development strategy
  - Encourage use of a unified online distributed revision control system: a FES hosted GitHub account (or equivalent) available to all institutions. Think of it as a modern NTCC repository, but with the advantage of always being up to date.
  - Facilitate sharing by simple FES licensing for use of codes (developed using FES funds) within the US and outside the US, but within the ITER collaboration.
  - Promote availability of publicly available codes at different institutions.
  - Unified access to data and physics codes with one account across facilities
  - Provide resources for user support and code maintenance