PRACE Peer Review Process: How to prepare a good Project Access proposal

Workshop, Cenaero
04 October 2018

John Clifford, Peer Review Officer, PRACE

Outline

- PRACE Peer Review process combines **scientific** and **technical** assessments
- Two groups of experts: **scientific** review **builds** on the **technical** one
- Key elements of a successful proposal, considering both scientific and the technical aspects
- New proposal **template** for Call 18

Project Access

Scientific Assessment

- The scientific review is performed by internationally recognised experts.
- Allocations can be increased/decreased by the prioritization panel (PRACE Access Committee)
- The proposals must address the following scientific criteria:

  - **Scientific excellence**
  - **Novelty and transformative qualities**
  - **Relevance to the call**
  - **Methodology**
  - **Dissemination**
  - **Management**

Scientific Assessment

- The scientific review is performed by internationally recognised experts. Allocations can be decreased or increased by the prioritisation panel (PRACE Access Committee).
- The proposals must address the following scientific criteria:
  - **Scientific excellence:** the proposed research must demonstrate scientific excellence

Scientific Assessment

- The scientific review is performed by internationally recognised experts. Allocations can be decreased or increased by the prioritisation panel (PRACE Access Committee).
- The proposals must address the following scientific criteria:

  - **Novelty and transformative qualities:** proposal should develop transformative topics of major relevance to **European** research

Scientific Assessment

- The scientific review is performed by internationally recognised experts. Allocations can be decreased or increased by the prioritisation panel (PRACE Access Committee).
- The proposals must address the following scientific criteria:

  - **Relevance to the call:** describe how the research is addressing the **scope** of the **call** if a specific scope is stated in the call

Scientific Assessment

- The scientific review is performed by internationally recognised experts. Allocations can be decreased or increased by the prioritisation panel (PRACE Access Committee).
- The proposals must address the following scientific criteria:

  - **Methodology:** the **mathematical numerical** methodology should be described and be appropriate to achieve the goals of the project

Scientific Assessment

- The scientific review is performed by internationally recognised experts. Allocations can be decreased or increased by the prioritisation panel (PRACE Access Committee).
- The proposals must address the following scientific criteria:

  - **Dissemination: channels** and **resources** for dissemination should be described.

Scientific Assessment

- The scientific review is performed by internationally recognised experts. Allocations can be decreased or increased by the prioritisation panel (PRACE Access Committee).
- The proposals must address the following scientific criteria:

  - **Management:** there must be a **solid management** structure in place

Role of Reviewers

- PRACE staff responsible for ensuring high quality **technical** and **scientific** reviews for each proposal
- **Technical** reviewers are experienced technical advisors
- **Technical reviewers** report on the “technical quality”, the suitability of the requested architecture, and the size of the allocation request
- **Scientific reviewers** from academia or industry with a proven track record
- **One reviewer** may be chosen by the applicant

Role of Reviewers

- **Scientific reviewers** report on the novelty and quality of the research proposed, potential impact and on the applicant's ability to carry it out
- **Scientific reviewers** also assess the proposal based on its fulfilments of any **pre-defined** criteria for a particular call
- **Scientific** reviewers will be known only to the PRACE peer-review staff
- Reviewer's **must** indicate they agree with the PRACE Conflict of Interest and Confidentiality agreement

Proposal template

Detailed Proposal information

- You should provide information on all of the subsections listed
  - If you wish to leave a section empty - provide a reason!
  - The information should be suitable for expert review in your field
  - …but also appropriate for a broader audience

Key contribution of the proposal

Scientific / Societal / Technological contribution of the proposal

- Outline the scientific / societal importance of your project
- How will HPC help you achieve your goals?
- What are the major expected outcomes?

Importance of the scientific problem and project overview

- **Justify the scientific importance of the problem and the request**
  - Describe main scientific / technical advances
  - Industrial partners should also summarize the potential economic or strategic business impact
  - Justification of the requested resources must be clearly linked to the software performance evaluation

Importance of the scientific problem and project overview

- **Overview of the Project**
  - Motivation, objectives and scientific challenges of the problem
  - Computational methods
  - Advances enabled through the requested Tier-0 PRACE award (e.g.: impact on community paradigms, new insights, etc.)
  - Expected outcomes of your proposal
  - Interdisciplinary value of your proposal

Validation, Verification, State of the art

- Describe the **validity** of simulations and predictions resulting from resources
- Address issues of reproducibility and highlight the predictive capabilities of your simulations.
- Provide **references** to relevant publications
  - **Validation**
  - **Verification**
  - **Sensitivity analysis and uncertainty quantification**
  - **Comparison with state of the art**

Validation, Verification, State of the art

- Describe the **validity** of simulations and predictions resulting from resources
- Address issues of reproducibility and highlight the predictive capabilities of your simulations.
- Provide **references** to relevant publications

  - **Validation:** validate your model against experiments or other established reference data (if available)

Validation, Verification, State of the art

- Describe the **validity** of simulations and predictions resulting from resources
- Address issues of reproducibility and highlight the predictive capabilities of your simulations.
- Provide **references** to relevant publications

  - **Verification:** verify the numerical consistency of your method or provide evidence of existing verifications

Validation, Verification, State of the art

- Describe the **validity** of simulations and predictions resulting from resources
- Address issues of reproducibility and highlight the predictive capabilities of your simulations.
- Provide **references** to relevant publications

  - **Sensitivity analysis and uncertainty quantification**:
    - Provide **sensitivity analysis** of your methods
    - Provide estimates of the **uncertainty** of your predictions
    - Multiphysics / multiscale problems - **uncertainty** of **methods** and **software** is desirable

Validation, Verification, State of the art

- Describe the **validity** of simulations and predictions resulting from resources
- Address issues of reproducibility and highlight the predictive capabilities of your simulations.
- Provide **references** to relevant publications

  - **Comparison with state of the art:** Place project in the context of competing work.

Software

- Describe the software that will be used including a discussion of the state of the art in the field

Software

- Describe the software that will be used including a discussion of the state of the art in the field
  - **Software:** all codes you are using in the proposal. Justify your choices and describe alternatives

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field
  - **Particular libraries:** required by production analysis software, algorithms and numerical techniques, programming languages

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field
  - **Parallel programming models:** for example MPI, OpenMP/Pthreads, CUDA, OpenACC, etc.

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field

  - **I/O requirements:**
    - I/O requirements (e.g. amount, size, bandwidth) for execution, input files, restart and other output
    - Describe I/O strategy (number of files, frequency, read/write size) and I/O behaviour of code during execution
    - Specify the restart overhead, not only for I/O; (e.g. a code may have to perform a costly domain decomposition first)

Data Storage, Analysis and Visualization

- **Project workflow:** include **role and timeline** of data analysis and visualization
- **Software workflow solution:** pre- and post-processing scripts that automate run management
- **I/O requirements:** amount, size, bandwidth, etc. for data analysis and visualisation
  - Highlight any exceptional I/O needs
  - provide data for (one or several) precise systems that will be simulated

Software Performance

- **Information on software performance is mandatory**
  - code should be tested on the requested machine.
  - If the preparatory host machine is different from the target machine, specify why the data is relevant
  - Report briefly the conversion factor (e.g. ratio of time to solution, flops or requested core hours) from the preparatory-test machine

Software Performance

- **Quantify the HPC performance of your project**
  - Data must be representative of the entire workflow of the project proposed and refer to the **main application code**
  - Scalability must be used to set most efficient job size for planned simulations
  - Performance must be linked to the request of the computing resources
  - Estimates based on related codes and/or data related to parts of production codes will not be accepted
  - All data must refer to the targeted systems in production runs or a system with comparable size, software stack and with the same architecture and network
  - Contact PRACE Tier0 centres if in doubt about the portability of your code

Performance results

Performance results

- **Strong and weak scalability:**
  - starting with the minimum size of the computer necessary to run the problem (usually 1 core or 1 node)
  - **justify** the minimum size for your scaling if it is larger than 1 core or 1 node (e.g. due to memory limitations)
  - **justify** if **weak** or **strong** scaling is not relevant for the project.

Performance results

- **Precision reported**:
  - Either single precision, double precision, or mixed precision. Only the precision you use in the simulation is relevant

Performance results

- **Time-to-solution**
  - normalized / averaged per iteration, number of cores and size of the problem: T_i^* = (Time-per-iteration) x (No. of cores) / (No. of computational elements)

    AND

  - normalized as total time to solution, number of cores and size of the problem: T_f^* = (Total-time-to-solution) x (No. of cores) / (No. of computational elements)
  - Justify the choice of your code (e.g. comparison with existing codes, methods or any other scientifically rigorous argumentation)

Performance results

- **System scale**
  - Either measured on full-scale system or projected from results of smaller system

Performance results

- **Measurement mechanism**:
  - Either timers, FLOP count, static analysis tool, performance modelling; specify if other

Performance results

- **Memory usage**:
  - specify requirements per node or core depending on the size of the problem

Performance results

- **OPTIONAL:**
  - Percentage of available peak performance
  - Alternatively provide code specific metrics for the requested machine (FLOPS, etc.)

Examples of Performance Reporting

Weak Scaling Example

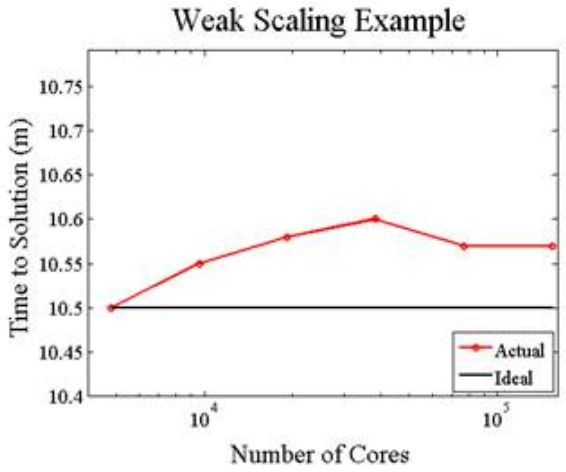

nProc	Time to Solution (m)	Ideal Time to Solution (m)
4800	10.50	10.50
9600	10.55	10.50
19200	10.58	10.50
38400	10.60	10.50
76800	10.57	10.50
153600	10.57	10.50

Strong Scaling Example

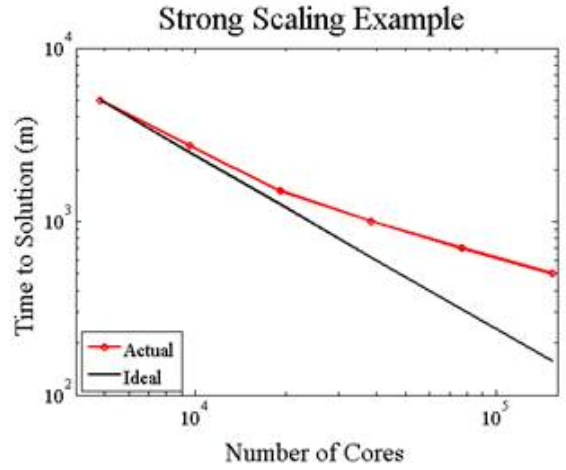

nProc	Time to Solution (m)	Ideal Time to Solution (m)
4800	5000.00	5000.00
9600	2725.00	2500.00
19200	1500.00	1250.00
38400	1000.00	625.00
76800	700.00	312.50
153600	500.00	156.25

Milestones

- Goals and milestones should articulate simulation and development objectives
- Provide clear connections between the project's milestones, planned simulations, and the computer time required
- Clarify any interdependencies
- Ensure that the core hour consumption is regular throughout the allocation or provide a requested schedule
- Provide a Gantt Chart of the simulation plan in production indicating job sizes and scheduling of computing tasks

Personnel and Management Plan

- **Present personnel overview**
  - personnel that will be hired for the project and their responsibilities
  - potential personnel turnover during the project and strategy for replacing them
  - Does proposal include a team of collaborators?
  - Outline the focus of each individual or subgroup and their interrelationships

- **Previous Allocations and Results**
  - Provide references to publications that **acknowledge PRACE resources**!

Further information

- PRACE Call Announcements (ToR, Technical Guidelines, etc.):
  - http://www.prace-ri.eu/call-announcements
- Application procedure:
  - http://www.prace-ri.eu/application-procedure
- Contact centres with technical inquiries:
  - http://www.prace-ri.eu/prace-resources
- Contact PRACE Peer-Review for assistance: peer-review@prace-ri.eu

Call 18 now Open (closing 30/10/2018)!

System	Architecture	Site (Country)	Core Hours (node hours)	Minimum request
Joliot Curie - SKL	BULL Sequana X1000	GENCI@CEA (FR)	142 million (3 million)	15 million core hours
Joliot Curie - KNL	BULL Sequana X1000	GENCI@CEA (FR)	101 million (1.5 million)	15 million core hours
JUWELS	BULL Sequana X1000	GCS@JSC (DE)	70 million (1.5 million)	35 million core hours
Marconi-Broadwell	Lenovo System	CINECA (IT)	36 million (1 million)	15 million core hours
Marconi-KNL	Lenovo System	CINECA (IT)	610 million (9 million)	30 million core hours
MareNostrum	Lenovo System	BSC (ES)	240 million (5 million)	30 million core hours
Piz Daint	Cray XC50 System	CSCS (CH)	510 million (7.5 million)	68 million core hours Use of GPUs
SuperMUC-NG	Lenovo ThinkSystem	GCS@LRZ (DE)	125 million (2.2 million)	35 million core hours

Thank you for your kind attention