D10.1

First Report on CoE governing bodies, KPI, and infrastructure setting up

Luisa Neri, Maria Bartolacelli, Carlo Cavazzoni, Andrea Ferretti, Maria Celeste Maschio, Elisa Molinari, and Daniele Varsano

Due date of deliverable: 31/05/2019 (**month 6**)
Actual submission date: 31/05/2019
Final version: 31/05/2019

Lead beneficiary: CNR (participant number 1)
Dissemination level: PU - Public

Document information

Project acronym:	MAX
Project full title:	Materials Design at the Exascale
Research Action Project type:	European Centre of Excellence in materials modelling, simulations and design
EC Grant agreement no.:	824143
Project starting / end date:	01/12/2018 (month 1) / 30/11/2021 (month 36)
Website:	www.max-centre.eu
Deliverable No.:	D10.1

Authors: Luisa Neri, Maria Bartolacelli, Carlo Cavazzoni, Andrea Ferretti, Maria Celeste Maschio, Elisa Molinari, and Daniele Varsano

To be cited as: Neri et al., (2019): First Report on CoE governing bodies, KPI, and infrastructure setting up. Deliverable D10.1 of the H2020 project MAX (final version as of 31/05/2019). EC grant agreement no: 824143, CNR, Modena, Italy.

D10.1 First Report on CoE governing bodies, KPI, and infrastructure setting up.

Content

1. Executive Summary

The deliverable D10.1 summarises the setting up of governing bodies, the definition of KPIs and the building up of infrastructures of MAX. The Centre of Excellence is taking advantage of the experience acquired during the first phase of MAX with the consolidation and strengthening of the same management structure while implementing *ad hoc* corrections to take into account the involvement of a larger partnership.

2. First Report on CoE Governing Bodies

As defined in the DoA (section B.3.2) and agreed upon by the partners in the Consortium Agreement (art. 6.1) the organisational structure of this Consortium is composed by:

1. the General Assembly (GA)
2. the Executive Committee (ExC)
3. the Coordinator (also named Director)
4. the Management team
5. the International Advisory Board (IAB)

1. The General Assembly (GA) is the ultimate decision-making body of the consortium, it takes strategic and formal decisions concerning the Centre, it is composed of one representative for each Party and coordinated by the Coordinator. The first formal GA was held on May 27, 2019 by teleconference. All the PIs, WP leaders or delegates attended the meeting. A first informal GA, in preparation of the official one, was held during the MAX kick-off meeting in Modena (December 14, 2018). Two additional informal GA meetings were held on April 29 and May 6, 2019 by teleconference.

2. The Executive Committee (ExC) is the supervisory body for the execution of the Project which supports the Coordinator; it is composed by Andrea Ferretti (CNR Nano), Daniele Varsano (CNR), Daniel Wortmann (Juelich), Giovanni Pizzi (EPFL), Zeila Zanolli (ICN2), Carlo Cavazzoni (CINECA), Stefano de Gironcoli (SISSA), selected by the General Assembly upon proposal of the Coordinator. The ExC is committed to: (i) implement the decisions taken by the GA, taking the decisions needed to support their implementation; (ii) coordinate the actions of the WPs towards common objectives;

(iii) monitor the development of the activities and interactions with stakeholders; (iv) analysing any difficulties / new opportunities in the CoE development and determining the necessary actions.

3. The Coordinator (art. 6.4) (also named Director) is the legal entity acting as an intermediary between the Parties and the Funding Authority and coordinating the execution of the decisions of the GA. The Coordinator is Prof. Elisa Molinari (CNR).

4. The Management Team supports the work of the Coordinator and of the Consortium for the day-to-day management of the project (DoA section B.3.2) and is set in the Coordinator's node. The management team is in charge of supporting activities of the governing bodies, ensuring an efficient internal communication; overseeing reporting deadlines and quality control; collecting and (e-)archiving the project documents. The management procedures are aimed to the effective implementation of the work plan and the successful completion of MAX deliverables and milestones. The team is led by the WP10 leader, Luisa Neri.

5. The International Advisory Board (IAB), composed by leading experts from academia and industry, will be set up to advise the Coordinator and the GA (DoA section B.3.2). The IAB is supposed to steer major scientific directions of the CoE and to advise the Coordinator and the GA on the main strategic choices needed during the life of MAX, evaluate the overall progress with respect to the objectives, and help shaping the evolution of MAX after the end of the project. A first discussion on the IAB composition was held at the GA meeting of May 27, consultations are ongoing to finalize the IAB composition shortly.

Note that the present Governance and Management structure is designed to improve over the first MAX phase especially for what concerns the close interaction and coordination among WPs, and the emphasis on actions of MAX as a CoE, with tight coordination of all partners on given strategic targets. In particular, the new executive body, the ExC has this as one of its specific missions.

3. First Report on MAX Management team

MAX as an infrastructure was basically set up during the first phase of the project (2015-2018): the management activities will build and improve on the experience gained in those three years. The main features are described in this chapter. The effort in M1-6 has been put on the

reorganization and strengthening of the infrastructure in consideration of the new partnership.

3.1 MAX managing team

MAX management is ensured by WP10, led by CNR with the collaboration of all nodes. Key is the CNR management team, that has gained experience in the first phase (WP9).

The CNR staff includes Luisa Neri, Project Manager and WP10 leader, Maria Bartolacelli, MAX Technical Secretariat; Maddalena Scandola, press office; Maria Grazia Angelini and Paola Corezzola, administration; Maria Celeste Maschio, training support. Further support and contributions are coming from the staff of CNR or other units.

The WP10 leader acts as Project Manager, to coordinate the Management staff, and to flank the GA, the ExC and the Coordinator from the managing point of view, in close connection with the WP leaders, the nodes financial offices and the EC.

3.2 Management objectives and principles

The main objectives of the management action are:

- to steer, coordinate and manage the project implementation in close coordination with the EC;
- to foster the effective cooperation and integration of the teams and activities;
- to manage and control the schedule, resources and quality;
- to ensure the financial, legal and administrative project management;
- to ensure the knowledge and data management of the project;
- to define the standards, tools and procedures for the internal communication;
- to manage the ethical and gender issues related to the project;
- to coordinate with the ecosystem of e-infrastructures, especially the other CoEs.

The principles inspiring MAX management are:

- focussing on user/business/development needs;
- delivering on time;
- collaborating with partners, industry and users;
- attention on quality;
- building up incrementally;
- developing iteratively, and
- communicating continuously and clearly.

The management will focus on people, interaction, community, skills, and talents.

Risk management will be an integral part of the project lifecycle. A detailed list of critical risks for implementation was included in the proposal with the definition of the WP involved and some measures of risk-mitigation. The management will follow closely the evolution of all WP activities and the implementation of contingency actions when needed.

Table 1. Critical Implementation risks and mitigation actions

Risk n.	Description of risk	WP n.	Proposed risk-mitigation measures
1	The collection of code components identified and slated for encapsulation, optimisation, and reuse does not cover the full functionality of MAX codes. (*medium risk*)	WP1	Target also code-specific modules for optimisation and interfacing work.
2	Specific architecture our codes are tuned for, will come to an end of lifetime due to vendor decisions. (*medium risk*)	WP2	Our work will not only focus on performance but also on portability to minimise the dependence on a single specific architecture. Together with the activities in WP4 we follow the developments pursued by hardware vendors to avoid this to happen.
3	The release of specifications of the planned pre-exascale machines and their deployment is delayed, we will not be able to adjust our efforts accordingly. (*medium risk*)	WP2	Through our connections to PRACE-partners we will aim at using the most advanced European supercomputers available. Furthermore our focus on performance portability should also allow us to produce codes easily adjustable on the final pre-exascale machines.
4	Code components identified in WP1 take longer to be extracted. (*medium risk*)	WP3	Focus the initial developments to the implementations that require less restructuring of the host application. Provide immediate feedback to WP1 work-groups should a design problem in the application components arise.
5	In T4.1 the available implementations of programming models to be tested have a TRL too low for use with main trunk codes and full applications. (*medium risk*)	WP4	Use mini-apps or relevant kernels to validate the paradigm, before considering the full application. Consider different implementations (e.g. at least two compiler technologies), preserve fallback solutions if appropriate.
6	In T4.2 there is risk of not being able to assess memory usage and then not being able to find solutions to exploit new memory for exascale systems. (*low risk*)	WP4	In all the assessment and co-design tasks we will consider at least two MAX flagship codes. Well established profiling tools will be used and we will ask for consultancy to experts from FET projects dealing with memory technology and other CoEs (e.g. POP or SAGE2), as already occurred in the first implementation of MAX.

7	In T4.3 there is risk of not being able to properly feed the co-design cycle. (*low risk*)	WP4	Consult experts in co-design coming from architecture specialists (Arm and external contacts) and integration companies (E4) as well as involve code developers. We will consider many co-design vehicles and codes, so we have a high probability of having enough code base and test cases to keep the co-design cycle alive.
8	In T4.4 there is risk of not having enough computational resources to properly run all benchmarks and profiling campaigns. (*low risk*)	WP4	Define at least one benchmark with large code coverage for each application that will be prioritised. We will calibrate the input dataset and limit benchmark run to few iterations, in order to fit in the in-kind resources available through HPC centres (e.g. CINECA commit to deliver 50K node/hours each year to run benchmarks at scale). Moreover, the benchmark team will apply for CPU grants at national and EU level (e.g. through PRACE preparatory program).
9	In task 5.1, supercomputing centres could adopt new fast schedulers optimised for HTC on exascale machines that have a very different implementation logic from the existing ones and require significant changes to the AiiDA API. (*low risk*)	WP5	Periodic interactions of AiiDA developers with supercomputer centres to know well in advance novel scheduler models and have time to adapt the API.
10	In task 5.2, some codes could be difficult to fully automate or might have high non-convergence or failure rates. (*medium risk*)	WP5	Minimise impact: provide turn-key solutions that give the user more input options when their choice is hard to automate, document their meaning; implement workflows to recover from failed runs. Continuous iterations between high-throughput researchers and code developers to highlight critical cases, giving developers detailed problematic cases and enough time to make the codes more robust.
11	In task 5.4, the data needs of the community might grow much faster than expected and the running costs of long-term storage exceed the current targets. (*medium risk*)	WP5	Implement per-user/per-project quotas to ensure sustainable limits. Monitor size growth of submissions to act on time in order to minimise risk impact.
12	In task 5.6 pilot 1, the amount of data might not be enough to train a reliable neural network for performance prediction. (*low risk*)	WP5	Complement the neural-network predictions with performance models based on code flow analysis, using hardware performance information as input. Act early to implement automated processes to collect data statistics of runs.
13	The viability of the realisation of these MAX Demonstrators relies on the availability of supercomputing resources. (*low risk*)	WP6	We have provisioned access to adequate resources *via*: –The integration of supercomputing centres in the MAX consortium. The team of each MAX flagship code is closely linked to a supercomputing centre participating in

			MAX (QUANTUM ESPRESSO and YAMBO with CINECA; SIESTA with BSC; FLEUR with Juelich; BIGDFT with CEA; CP2K with CSCS). These centres routinely provide computational resources for development, benchmarking and debugging, and also large amounts of CPU time during the testing periods of newly installed supercomputers. – Access to PRACE resources, through their continuous calls for Preparatory Access projects. – As a backup plan, CINECA has committed to provide the computational resources needed to perform the activities in MAX Demonstrators.
14	Difficulty to find dedicated personnel for service provisioning in case of large custom development. (*medium risk*)	WP7	A number of actions aimed at the training of SW developers has been identified and put in place in MAX. This should provide a larger base of developers thereby minimising the risk.
15	The planned training offer is insufficient with respect to the requests. (*medium risk*)	WP8	Apply for additional support from other sources in order to increase the domain-specific offer, further coordination with PRACE and Focus CoE to address requests. Training materials available online.
16	There is limited interest in the education and training offered by the CoE. (*low risk*)	WP8	Focused actions in event advertising in collaboration with WP9.
17	Users and stakeholders not responsive to the selected communication & dissemination channels. (*medium risk*)	WP9	Multiple different media and communication & dissemination strategies will be used, making the action flexible.
18	Users and stakeholders not responsive to the selected communication & dissemination channels. (*medium risk*)	WP10	Multiple different media and communication & dissemination strategies will be used, making the action flexible.

3.3 Management Activities

Several activities have been performed since the beginning of the project in 3 specific domains: 1. Governance, monitoring and quality; 2. Administrative, legal and financial management; 3. Knowledge and communication management.

Activities performed:

- consortium Agreement discussion and definition;
- day-to-day coordination of the project activities and of the secretariat activities;
- monitoring of WPs thematic work and deadlines respect;
- monitoring of the internal periodic report in order to have a complete overview of the project status of advance;

- definition of the guidelines, internal procedures and templates for deliverables and other project documents;
- support for Consortium meetings, Governing Bodies meetings and WPs meetings;
- overall legal, administrative and financial management (budgeting and financial reporting, distribution of the EU financial contribution among partners, support to partners regarding financial and administrative issues and duties);
- reorganization of the MAX internal repository and an internal IT infrastructure using Google Apps.

At the Kick-off meeting, the management presented a detailed plan, setting dates for the collection of inputs, reports and meetings. All WP teams, nodes and committees are invited to respect such internal deadlines and the process of management is going on pretty smoothly.

3.4 Management tools and repository structure

The management tools that were chosen are the Google Apps.

A complete renewal of the MAX internal IT infrastructure based on Google Apps was done in order to have a wider and more participated use of this important tool. It included:

- a re-organization of folders in the Google Drive repository, that has been defined in order to collect relevant documents in real time and have updated lists related to meetings, events and activities. This enhances the quality of the management as it lowers the request for papers and information, and helps the cross-checking of records;
- Google Calendar (with different sub-calendars) is now used for deadlines and meetings overview;
- Google accounts have been created for all the project participants;
- Google Groups have been created to facilitate internal communication;
- Google Meet is the channel used for periodic formal and informal meetings.

Furthermore we setup other MAX software repositories on GitHub, GitLab, DockerHub, and defined other communication channels (MAX-centre Slack).

MaX 2018-2021 > MaX Repository

	Name	Contents
	1. MaX Official Documents	Contains documents for all members Grant Agreement Consortium Agreement Periodic reports AOD
	2. MaX Continuous Reporting	MaX Communication – Press repository (to collect docs, images, file related to communication & press) Open Access (contains guidelines about OA) MaX Publication file (to collect publications) MaX Events file (to collect dissemination events) MaX Communication - Press file (to collect the list of press releases, articles.... Soft copies can be put in the folder) ...
	3. MaX Events	related to MaX-organized events
	4. MaX Deliverables	Templates for deliverables Collection of final deliverables
	5. MaX General Assembly	Minutes List of meetings Any Official Documents issued or needed
	6. MaX WPs (one folder per WP)	Minutes List of meetings ...
	7. MaX Partners (one folder per Partner)	Minutes List of meetings ...
	8. MaX Executive Committee	Minutes List of meetings ...

Fig 1. Scheme of the MAX repository organization in Google Drive

3.5 Coordination with Pan-european and national HPC ecosystems

MAX aims at being one of the driving forces within the European HPC ecosystem. In order to do so, we planned to actively pursue effective interfaces with its main components:

1. Interface with the EC and with European initiatives on policy and technical issues (PRACE, EuroHPC JU, FENIX, ETP4HPC, the other CoEs and the FocusCoE CSA, etc.).

Since the beginning of MAX, two of our members have become part of the Euro HPC Research & Innovation Advisory Board (RIAG): Elisa Molinari - CNR, nominated by the EuroHPC Governing Board, and Carlo Cavazzoni - CINECA, nominated by the ETP4HPC. This will greatly facilitate the interactions of MAX with the whole ecosystem and offer opportunities for MAX contributions.

In these early months, MAX members have taken part in a number of events (partially as a consequence of the above commitments):

<u>EuroHPC JU:</u>

- three RIAG meetings (Brussels on March 1, 2019 and April 5, 2019; Poznan May 13,

2019) (E. Molinari, C. Cavazzoni)

- EuroHPC Summit Week Poznan, May 13-17, 2019 (E. Molinari, L. Neri - CNR, C. Cavazzoni - CINECA, S. Huber - EPFL, F. Magugliani - E4), with presentations in the Codesign Workshop (Cavazzoni) and the Data Workshop (Huber)
- contributions to the draft RIAG Strategic Plan 2019 and the Multiannual Strategic Agenda (ongoing) (E. Molinari, C. Cavazzoni): MAX PIs are offering input for these strategic documents

All MAX partners have been especially active within national communities and institutions to promote their commitments towards EuroHPC.

<u>CoEs coordination activities:</u>

- CoEs Workshop held in Brussels (BE) on December 3, 2018 (E. Molinari - CNR)
- CoE meetings organized with FocusCoE: FocusCoE kick-off meeting, Frankfurt (DE) on February 21, 2019 (E. Molinari, L. Neri - CNR; S. Muscella, F. Osimanti - TRUST-IT); informal General Assembly and General Assembly, Poznan (PL) (E. Molinari, L. Neri - CNR) May 13-17, 2019. E. Molinari has been elected vice-chair of the "HPC CoE Council" (as the GA has been renamed). She will be in office for three years.

Fig. 2 Group Picture of FocusCoE attendees, Frankfurt (DE) on February 21, 2019

2. Interface with national initiatives within Europe, including countries which are not yet fully integrated/supported in the European HPC system; interface with international HPC materials research ecosystem.

Pan-European interactions:

- EOSC and EOSC-Hub:
  - telco to design collaborations and the pilot role of MAX in EOSC-Hub a pilot EOSC-Hub (19/02/2019 T. Ferrari, S. Garavelli EOSC, D. Testi EOSC-Hub, G. Pizzi, L. Talirz, E. Passaro, E. Molinari, L. Neri, M. Bartolacelli MAX); followed by a technical telco of WP5 and EOSC
  - EOSC Week, Prague, CZ - April 9-12, 2019 (G. Pizzi - EPFL, invited talk on "Pathways for EOSC-Hub and MAX collaboration" April 11, 2019)
- Activities to broaden the MAX impact in other countries, especially Eastern Europe:
  - several informal contributions in the HPC Competence Centers debate and construction; participation in the HPC Competence Centers Workshop, Brussels, BE - November 28, 2018 (E. Molinari - CNR)
  - an informal brainstorming session of the MAX General Assembly to collect ideas for a MAX "Pan-European action plan". The most immediate result was planning a number of events taking place in Eastern Europe Countries and involving organizers from those countries:
    (i) the Summer School on Quantum ESPRESSO that will be organized by MAX in Slovenia in September 2019 (13-20/09/2019), jointly with the Jožef Stefan Institute (Ljubljana), the Ruder Boskovich Institute (Zagabria) and the Institute of Physics (Belgrade) ;
    (ii) the School on First Principles Simulation of Materials with SIESTA, organized by MAX in Belgrade in November 2019 (4-9/11/2019), in collaboration with the University of Belgrade and the Barcelona Supercomputing Centre, with the support of CECAM;
    (iii) the Aiida tutorial, that will be organized by MAX in Vilnius, Lithuania, in Spring-Autumns 2020, jointly with the Vilnius University in collaboration with CECAM and MARVEL. Specific target attendance from Lithuania and its neighbouring countries (Latvia, Estonia, Belarus, Poland).

Coordination with other projects/organizations include:

- HPC Europa3: discussions and collaboration on several proposed visits; a Memorandum of Understanding in preparation;
- EU H2020-NMBP-TO-IND-2018 INTERSECT project (Grant Agreement n. 814487); coordination and participation to the INTERSECT Kick-off meeting (February 5, 2019, Modena IT) and definition of a Memorandum of Understanding;
- Psi-k organization: agreement about coordinated dissemination activities (workshops and 2020 Psi-k conference). N. Marzari - EPFL serves as Psi-k Chair, and E. Molinari - CNR and S. Blügel - JUELICH are in the Board of Trustees;

- CECAM: agreement about coordinated training activities;
- Graphene Flagship: coordination and co-organization of a HPC section within the Graphene Week (September 23-27, 2019, Helsinki Finland) on September 24, 2019;
- Battery2030+: participation in preparatory meetings for the Ramp-up projects; brainstorming on challenges and opportunities within the EuroHPC ecosystem.
- Materials Modeling Market Place for Increased Industrial Innovation - H2020-NMBP-25-2017 (Grant Agreement n. 760173): brainstorming for collaboration, MoU in preparation.

3.6 Kick-off meeting

The two-day kick-off meeting was held in Modena on December 13-14, 2018 and it was organised by CNR in the beautiful setting of the San Carlo Theatre. All the partners attended the meeting with their PIs, researchers and students (60 people in total).

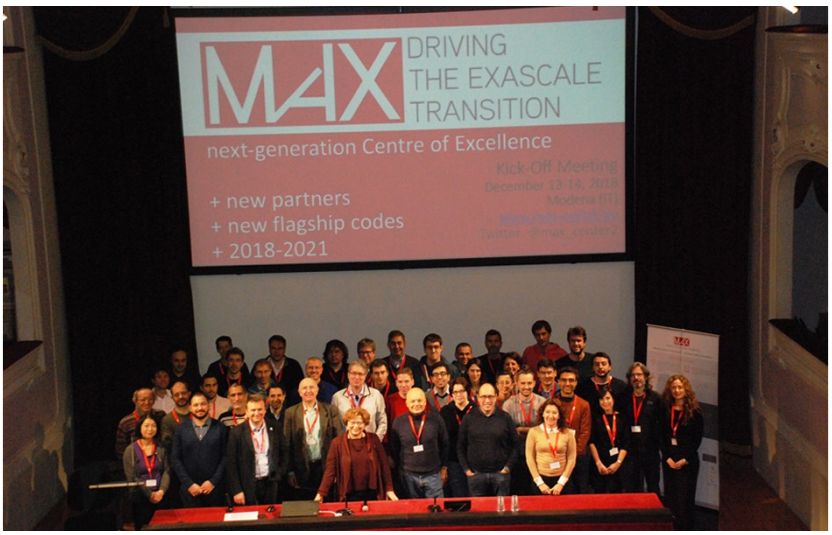

Fig 3. The MAX kick-off meeting official picture

The agenda included public presentations and internal discussions organized per WP, working groups and thematic focuses. An overview of expected work was given by PIs, and thematic groups gathered to discuss activities across WPs. An informal GA meeting was organized at the end of the Kick-off.

MAX KICK-OFF MEETING 13-14 DECEMBER 2018

DAY 0			
from 18.00	*GET TOGETHER*		

DAY 1	**December 13, 2018**		
09.00-10:30	Public presentation: European context, MaX		
10:30-11:00	*coffee break*		
11:00-12:30	Public presentation: European context, MaX		
12:30-13:30	*lunch*		
13:30-14:30	Introduction to the work plan, and to some new issues and challenges		
14:45-16.00 PARALLEL SESSIONS	WP6 demonstrators (chiar: P. Ordejon)	WP4 co-design (chair: C. Cavazzoni)	WP8 training (D. Varsano)
16:00-16:30	*coffee break*		
16.30-18.00 PARALLEL SESSIONS	WP9 communication (chair: S. Muscella)	WP4 co-design (chair: C. Cavazzoni)	
18.00-18:30	Wrap up/summary from each working group		
20:00	*social dinner*		

DAY 2	**December 14, 2018**		
09:00- 10:30 PARALLEL SESSIONS	WP5 Data (chair: N. Marzari)	WP1+WP2+WP3: SW structure, portability, algorithms (chairs: S. Baroni, D. Wortmann, L. Genovese)	WP9 and WP10 Engagement, Uptake, Interaction with broader EU ecosystem
10:30-11:00	*coffee break*		
11.00-12.30 PARALLEL SESSIONS	WP5 Data (chair: N. Marzari)	WP1+WP2+WP3: SW structure, portability, algorithms (chairs: S. Baroni, D. Wortmann, L. Genovese)	WP7: Services and users (chair: M. Ippolito)
12:30-1300	Wrap up/summary from each working group		
12:30-14:00	*lunch*		
14:30-16:00	more group work (except for GA members)		General Assembly - GA (team leaders)
16.00-16.30	Wrap up and farewell		

Fig 4. The MAX kick-off meeting agenda

For more info please refer to the dedicated web page http://www.max-centre.eu/max-kick-off-meeting-december-13-14-2018-in-modena-it/

4. Definition of Key Performance Indicators

To measure the impact of MAX actions a set of quantitative indicators has been identified. MAX Key Performance Indicators (KPIs) are designed in order to be readily measurable and adequate in evaluating MAX performance in achieving its strategic and operational goals.

KPIs are defined for each Work Package and are also divided into functional areas. In fact, WP1-WP3 are focused on Software Development, while WP4-WP6 are devoted to Exascaling, Co-design, and Data handling. The improvement of Infrastructure and Services is related to the work of WP7 and WP8 and the outcomes of Dissemination, Exploitation and Management derive from WP9 and WP10.

In the following we present the description and the rationale behind the choice of each KPI.

Software Development

K1.1: Refactored code components

WP1 is focused on the definition of an exascale-enabling software architecture to empower MAX flagship codes. This will lead to a major refactoring of the code structure, having a number of components (libraries and modules) disentangled from the codes and re-assembled, with the final goal of implementing the principle of "separation of concerns" in the SW architecture. As the specific details of exascale machine architectures are still unknown, MAX codes will then target high flexibility. The key elements of MAX codes will be identified and will be encapsulated following specific criteria into system-, mathematical- and domain-specific libraries. In view of the above discussion, K1.1 is defined as the "*number of software components (libraries, modules, building blocks) identified and refactored as self - standing objects*". We expect a tentative target of at least 15 components fully refactored and released as self-standing object by M36.

K1.2: Impact of MAX libraries

The refactoring process of MAX codes will lead to a number of exascale ready and self-standing components (libraries, modules, etc) to be defined and publicly released. These components can be adopted either by the MAX flagship codes themselves or by other quantum-simulation community codes from outside the MAX consortium. In order to measure the impact of the MAX libraries, K1.2 is defined as the "*average (over libraries) number of codes (inside/outside MAX) using each MAX library*". We aim at a target value equal-larger than 2 on average.

K2.1: Portability on different architectures

Performance portability is one of the key aspects of MAX work plan. Within WP2, we aim at porting the essential components of the MAX flagship codes to the new and diverse hardware (HW) architectures envisioned/foreseen in the pre- and exascale era. While emerging, relevant HW architectures will be identified and targeted for porting. Therefore, K2.1 is defined as the "*number of new ports of MAX codes' components to different hardware architectures*". As an example (considering some currently available cases) Intel KNL, Intel + NVIDIA GPU, IBM Power + NVIDIA GPU, etc would all qualify as different target architectures. Overall, each component will be ported to at least 2 to 3 different HW architectures while a subset of 8 relevant components, chosen by their computational weight, will be ported to at least 3 to 4 different HW architectures.

K2.2: Auto-tuning of code components

Besides the portability of components to diverse HW architecture, WP2 also deals with making the performance available to users. Within MAX, one strategy identified to reach this goal is to equipe libraries/components with auto-tuning capabilities. In this perspective, K2.2 is defined as the "*number of intra-code components that have been tuned*", with a target value of at least 4-5 auto-tuned components by M36.

K2.3: Code exploitation of HPC machines

MAX flagship codes span a variety of different approaches to compute the electronic structure properties of materials, while all address computational performance in the HPC environment. Moreover, because MAX targets pre- and exascale readiness for its flagship codes, HPC performance is a crucial quantity to measure. In this view, K2.3 has been defined as the "*number of MAX codes able to exploit in a single run a partition of a Tier0 machine which provides at least 10% of the maximum available performance (depending on whether the code is compute- or memory- bound)"*. We expect by M36 all MAX flagship codes to be able to achieve this target.

K3.1: New code functionalities

WP3 deals with exascale-oriented algorithms and functionalities in electronic structure and materials science. Here we address both the calculation of properties that were not possible to compute before the exascale era (out-of-reach, not fully developed) as well as the use of exascale-enabled algorithms to make calculations faster or more reliable. K3.1 provides a measure for the first part of this activity and is defined as the "*number of new physical functionalities enabled to run in massively parallel environment*", with a target at M36 of at least 10 new functionalities over all codes.

K3.2: Algorithmic improvements

As stated above, an important part of the WP3 activity deals with the implementation and exploitation of novel and/or exascale-oriented algorithms in electronic structure and materials science calculations. K3.2 is defined as the "*number of code functionalities enabled by algorithmic developments from WP3*", with a target of at least 7 major algorithmic implementations by M36 over all codes.

Exascaling, Co-design & Data handling

K4.1: Profiling and benchmarking

The activity of continuously monitoring and evaluating the performance of its flagship codes is central to MAX, and is the subject of Task 4.4. This is crucial to drive the development actions of WP1-2-3 and to generally validate/demonstrate the improvements encoded in the flagship codes. K4.1 is defined as the "*number of profiling and benchmarking campaigns*" and we expect to have at least 2 campaigns per year per code (e.g. one on mini-apps/kernels and one on full applications using updated code versions).

K4.2: Programming models

With the emergence of new and disrupting HW architectures, novel programming models will be developed in order to ease (or simply permit) the exploitation of the available computing power (e.g. OpenMP 5.0 as a programming model to address accelerated HW). The testing of advanced programming models on selected kernels of interest to the MAX community is at the core of WP4, and the indicator K4.2, defined as the "*number of kernels tested with innovative programming models*", provides a measure of it. We expect to test at least 10 different kernels over the MAX time span.

K5.1: Public workflows and turn-key solutions

WP5 is focused on the delivery of powerful workflows for the calculation of materials properties using the MAX flagship codes. High-level workflows can be equipped with all the logic and heuristics to make them robust and to require a minimum user intervention in order to compute a complex property end-to-end. This is what we call a "turn-key" solution. Ultimately, these turn-key solutions offered to the users will also permit to choose from a set of codes which workflows run, depending on code accuracy, efficiency and availability. The creation of such ecosystem of workflows and turn-key solutions is at the core of MAX WP5.

K5.1 is defined as the "*total number of workflows and turn-key solutions publicly released*". By M36 we expect at least a dozen turn-key solutions publicly released and production ready.

K5.2: Quantum Mobile releases

Quantum Mobile is a downloadable virtual machine deploying all MAX flagship codes, together with a collection of tools, all properly set up and ready to be used through a local AiiDA installation and the respective plugins and workflows. Each periodic version of the Quantum Mobile (QM) will be updated with the newest releases of MAX software. Quantum Mobile is a crucial tool for the dissemination of MAX software by lowering the access barriers to it. In turn this enhances the engagement of users, with a special focus on the industry environment. K5.2 is defined as the "*number of periodic releases of updated versions of the Quantum Mobile*". We expect at least 2 releases of QM per year.

K5.3: Repository of materials data and workflows

MaterialsCloud (MC, http://www.materialscloud.org) is the hub where AiiDA plugins, workflows, curated materials science data (and more) can be hosted and shared among users (if AiiDA were git, MC would be GitHub or GitLab). This community infrastructure, developed at EPFL, is at the core of MAX data management plan and, with its 5 sessions (LEARN, WORK, DISCOVER, EXPLORE, ARCHIVE) is offered to the scientific (both academic and industrial) community active in the field of materials science as a central hub for hosting and sharing knowledge. K5.3 is defined as the "*total number of validation runs (number of structures times number of codes), and of highly curated data present on the MaterialsCloud*". We plan to host at least 50.000 highly curated datasets by M36.

K6.1: Tested demonstrators

The advances obtained by WP1-WP5 will be used in WP6 to demonstrate the readiness of MAX flagship codes for an efficient use of the pre-exascale machines (and to prepare for the exascale ones). For this purpose, a number of demonstrators will be defined, addressing specific scientific challenges that would not be possible without the technology improvement and development of the codes provided by MAX. The gaps towards the exploitation of pre- and exascale machines will be identified and addressed, giving fundamental feedback to the work made by WP1-WP5. K6.1 is defined as the "*number of Extreme Scale Pilot projects run and analysed*", and we expect to set up at least 6 complete scientific demonstrators by M36 to investigate and test the advances of MAX codes.

K7.1: High level domain-specific support

The activity of WP7 is mainly focused on offering materials science domain-specific support to EU scientists from both academia and industry. High level support actions range from consultancy to custom developments and include: (i) evaluation of the performance of MAX codes on different architectures, (ii) priority bug fixing on selected features, (iii) determination of best performance parameters for a given system/architecture, and (iv) execution of scientific studies in materials science. The user requests are managed through the MAX users portal and a dedicated ticketing system. K7.1 is defined as the "*total number of high-level support requests addressed by the CoE (trouble tickets, consultancy)*", and we statistically expect to have more than 100 requests per year over the 3 year time span of MAX.

K8.1: User uptake in MAX training events

The training of a new generation of developers and users is the core of WP8. Within the wide plan of domain-specific schools, courses and hackathons, the effectiveness of these actions is quantified through the number of all the MAX trainees. K8.1 is defined as the "*number of people trained in MAX training events (expressed in person days / year)*". Overall, we aim at having at least 500 person-days of face-to-face training per year (this corresponds to e.g. 5-6 schools lasting 3 days and having 30 participants on average, per year).

K8.2: Training quality

Being a user driven centre of excellence, users feedback is of uttermost importance to MAX. Calculating the average on the evaluation assigned by participants to MAX training events by anonymous feedback surveys, this KPI measures the perceived and reported quality of the training and face-to-face support activities put in place by MAX WP8. K8.1 is defined as the "*average evaluation assigned by participants in training events in anonymous questionnaires (values normalized between 1, negative and 5, optimum)*" and, given the results of the previous MAX phase, we expect an average value at least larger than 4.0.

K9.1: Dissemination

A set of actions made by WP9 must ensure the coverage of stakeholders and adequate visibility across Europe. For example, the dissemination of MAX CoE results will be done through publications in peer-reviewed scientific international journals, presentations at workshops, conferences meetings, as well as via tutorials on the MAX flagship codes. Importantly, the MaterialsCloud repository will also be exploited to host and disseminate MAX results. In terms of a measure, K9.1 is defined as the "*total number of invited talks to conferences and schools (including event organisation) plus the total number of scientific*

publications on international journals". Considering the wide community working on the topics and the large amount of materials science dedicated events, we believe that, respectively, the scientific production will be greater than 25 papers per year and the invited talks will be more than 75 per year.

K9.2: Communication and Outreach

Engagement, Communication, Dissemination, & Uptake are at the core of the actions of WP9. A number of activities and channels (web sites, newsletters, social media, organization of dedicated events, webinars, videos, brochures, press releases, etc) to address the relevant stakeholders will be considered and exploited. In particular, regular social media activities will be carried out and continuously managed and results tracked on the progress and outreach. The just mentioned actions are used to define K9.2 as the "*total number of people in the MAX network (social media followers, newsletter subscribers)*". As a target, K9.2 is estimated to overcome 1000 followers by M12.

K10.1: Coordination with pan-European and national HPC ecosystems

As part of the European HPC ecosystem, MAX mission emcompasses the coordination and collaboration with all relevant stakeholders. In particular the following actions are at the core of WP10: (i) interfacing with EC and other European initiatives on policy and technical issues (PRACE, EuroHPC JU, FENIX PPP, ETP4HPC, FocusCoE CSA and other CoEs); (ii) interfacing with national initiatives within Europe, including countries which are not yet fully integrated/supported in the EU HPC ecosystem. K10.1 is defined as the "*number of collaborative actions (e.g. participation to working/advisory groups, EuroHPC meetings, etc) to support the HPC ecosystem*". We expect to participate to at least 8 collaborative actions per year.

	#	Key Performance Indicator	KPI definition	Start	Target at M36
Software Development	K1.1	Refactored code components	Number of software components (libraries, modules, building blocks) identified and refactored as self-standing objects	M1	≥ 15
	K1.2	Impact of MaX libraries	Average (over libraries) number of codes (inside/outside MaX) using each MaX library	M1	≥ 2
	K2.1	Portability on different architectures	Number of new ports of Max codes' components (eg KNL, intel+GPU, power+GPU, ARM) to different hardware architectures	M1	≥ 2-3/component
	K2.2	Auto-tuning of code components	Number of intra-code components that have been tuned	M1	≥ 4-5
	K2.3	Code exploitation of HPC machines	Number of MaX codes able to exploit in a single run a partition of a Tier0 machine which provides at least 10% of the maximum available performance (depending on whether the code is compute- or memory- bound)	M1	6
	K3.1	New code functionalities	Number of new physical functionalities enabled to run in a massively parallel environment	M1	≥ 10
	K3.2	Algorithmic improvements	Number of code functionalities enabled by algorithmic developments from WP3	M1	≥ 7
Exascaling, Co-design & Data handling	K4.1	Profiling and benchmarking	Number of profiling and benchmarking campaigns	M1	≥ 2/year
	K4.2	Programming models	Number of kernels tested with innovative programming models	M1	10
	K5.1	Public workflows and turn-key solutions	Total number of workflows and turn-key solutions publicly released	M1	≥ 12
	K5.2	Quantum Mobile releases	Number of periodic releases of updated versions of the Quantum Mobile	M1	≥ 2/year
	K5.3	Repository of materials data and workflows	Total number of validation runs (number of structures times number of codes), and of highly curated data present on the MaterialsCloud	M1	50,000
	K6.1	Tested Demonstrators	Number of Extreme Scale Pilot projects run and analysed	M1	≥ 6
Infrastructure & Services	K7.1	High level domain-specific support	Total number of high-level support requests addressed by the CoE (trouble tickets, consultancy)	M1	≥ 100/year
	K8.1	User uptake in MaX training events	Number of people trained in MaX training events (expressed in person days)	M1	≥ 500/year
	K8.2	Training quality	Average evaluation assigned by participants in training events in anonymous questionnaires (values normalized between 1, negative and 5, optimum)	M1	> 4 .0
Dissemination, Exploitation & Management	K9.1	Dissemination	Total number of invited talks to conferences and schools (including event organization) + total number of scientific publications on international journals	M1	≥ 25 papers/year + 75 invited/year
	K9.2	Communication and Outreach	Total number of people in the MaX network (social media followers, newsletter subscribers)	M1	≥ 1000 by M12
	K10.1	Coordination with pan-European and national HPC ecosystems	Number of collaborative actions (eg participation to working/advisory groups, EuroHPC meetings, etc) to support the HPC ecosystem	M1	≥ 8/year

Fig 5. The MAX KPI table