



PROGRAMME
DE RECHERCHE
NUMÉRIQUE
POUR L'EXASCALE

Brainstorming on Exa-DoST work-plan

Report on the solutions selected in Exa-DoST to answer the Data challenges at Exascale & Design document for illustrators

Work Packages in Exa-DoST

WP1: Exascale
I/O and
storage

WP2: Exascale
In-situ data
processing

WP3: Exascale
ML-based data
analytics

WP4: Shared building blocks
& integrated illustrators

WP5: Management, dissemination and training

Previously in Exa-DoST

What is an application motif?

- This is the description of a use-case from applications
 - As specific as need be to clearly identify the problem (for real, no vague, general description)
- This is a description of the problem tools & libraries we develop will solve
 - Irrelevant details can be removed (is the name of the application relevant to the use of our library?)
=> If you are used to agile methods, this is a “user story”: a problem, not a solution
- The application motif is the description of the problem our tools will solve for applications
 - It should be specific enough that if another application can relate to the motif, the tools and libraries we develop are useful for them to take on the shelf
- An application motif should make it possible to answer the question by another application:
 - “Your work on this application, is it relevant to me?”

Application motif example #1

Large checkpoint motif:

“The writing of one large file (about 1/4 of the RAM) by one process per node, all nodes at the same time, every few hours of simulation, including as the very last step of the job. The data is a multi-D array, not necessarily contiguous in RAM (ghosts cells typically), but can be at the cost of writing a larger block of data. Only the latest write will be read and all previous writes can be disposed of as soon as a following write is securely written. The read happens with the same distribution as the writes and each process reads exactly what was written by a single specific process.”

- This example is inspired by Gysela and quite specific technically
- This can drive our work on IO optimization
- The Gysela name, the type of simulation that is not relevant to the checkpointing optimization has been removed

Application motif #2

Steering motif:

“A simulation is built by coupling two equally costly parts modeling two physical processes. One of the two parts is irrelevant at the beginning of the simulation (part A) and only becomes important when some specific conditions are met. A python script can be used to analyse the state of the simulation and give a probability of whether the part of the simulation should be started. The script requires access to two variables distributed over all processes and taking about 1% of the nodes RAM. The goal is to execute both the simulation and the script concurrently and to start part A in the simulation only when the probability it is required goes above 75%.”

- This motif is (freely) inspired by Coddex
- More information should be added to allow the design an in situ workflow
 - But description of the solution, the workflow is not part of the motif.

And now, your new episode of Exa-DoST

What it a work-plan?

- This is our plan to address the motifs
 - The bottleneck we already identified, and
 - Not the solution, but the way we plan to work on the solution
- What contribution by each WP will be used in each illustrator?
 - What libraries, how they will be combined?
 - What libraries from the project, but also what external libraries?
 - What is already available and needs to be integrated in applications?
 - What new features should be designed to solve the problems?
- Who will need to work with whom, on what subject
- Everything we wanted to say last time, but I asked you to delay

Work-plan example #1

Large checkpoint work-plan:

“Writing large checkpoints works well with HDF5 and one file per process when supercomputers are empty. Performance degrade catastrophically when external load compete for IO resources. We will work and apply the tools for performance analysis developed in WP1 to better understand what really impacts performance. We will keep using HDF5 through PDI but will explore various IO strategies (excluding ghost cells or not, using separate or parallel HDF5). WP1 will work on the development of a monitoring or “IO weather” tool that can be installed on supercomputers to monitor the state of the elements previously identified. We will develop a plugin for PDI that can delay IO when the machine is too loaded. We will integrate and evaluate this strategy in Gysela.”

- We identify the illustrator, libraries & tools we plan to put together
- We identify some goals and know who should collaborate to reach them
- This is an initial plan and will have to be updated regularly

Work-plan example #2

Steering work-plan:

“We will couple the simulation and analysis part in Codex. We will use Damaris through PDI to do so, as well as an synchronous in-process coupling. We will implement a system enabling asynchronous feedback from the analysis in Damaris. This mechanism will be used to provide the result of the analysis back to Codex as a boolean that the simulation will use to start the second module when required. We will compare the performance achieved using three configurations: synchronous analysis, on dedicated cores or on dedicated nodes. We will design an algorithm based on these experiments to determine the best placement in production. If time allows we will compare the result to a solution based on Deisa.”

- Similar to the previous example: illustrator, libraries & tools, goals, people, to be updated regularly
- Important: every-time some work is identified, the person in charge or contact should be clearly listed
- Warning: both examples are totally artificial and do not reflect the work that will actually be done

How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups

Gysela

SKA

Coddex, Dyablo
& others

WP1:
Exascale
I/O and
storage

WP2:
Exascale
in-situ data
processing

WP3:
Exascale
ML-based
data
analytics

How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups
- Remember about distributed matrix multiplication?

Gysela

SKA

Coddex, Dyablo
& others

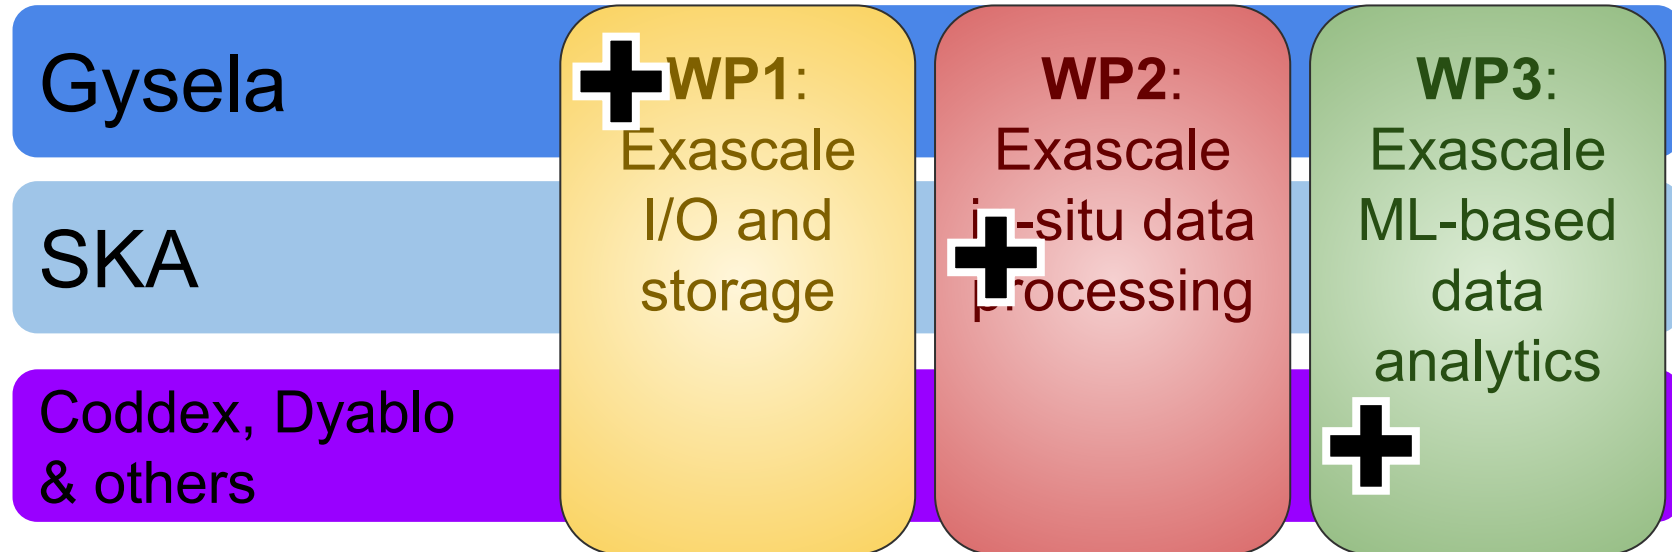
WP1:
Exascale
I/O and
storage

WP2:
Exascale
in-situ data
processing

WP3:
Exascale
ML-based
data
analytics

How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups
- Remember about distributed matrix multiplication?



How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups
- Remember about distributed matrix multiplication?

Gysela

SKA

Coddex, Dyablo
& others

WP1:
Exascale
I/O and
storage

+ WP2:
Exascale
in-situ data
processing

WP3:
Exascale
+ ML-based
data
analytics

How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups
- Remember about distributed matrix multiplication?

Gysela

SKA

Coddex, Dyablo
& others

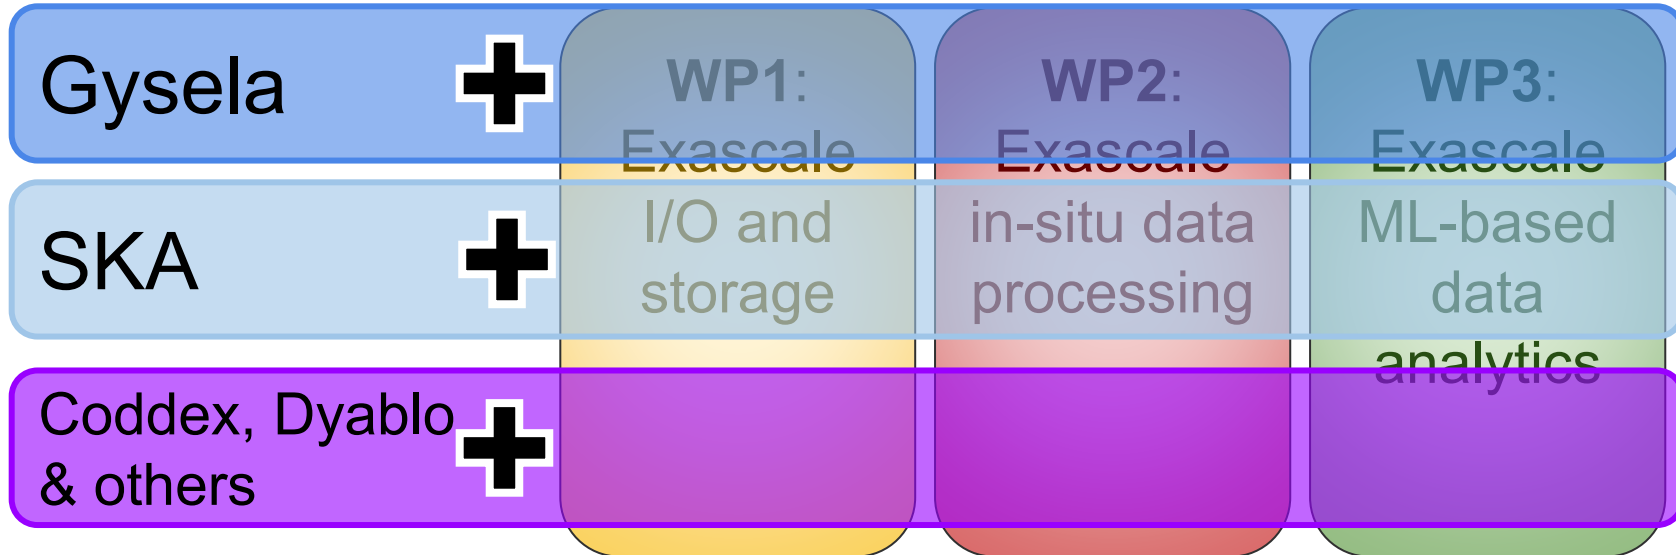
WP1:
Exascale
+ I/O and
storage

WP2:
Exascale
in-situ data
processing
+

+ WP3:
Exascale
ML-based
data
analytics

How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups
- Remember about distributed matrix multiplication?



How we'll work

- We have 3 “libraries & tools” Work-packages
- We have 3 “illustrator” groups
- Remember about distributed matrix multiplication?

- Session leaders (WP1-2-3 leader + illustrators representative)
 - Should come prepared, with some slides to launch the work
 - Will animate the sessions
- Please feel free to use a collaborative note taking system
 - But identify one responsible for the notes of each session
- We'll have feedback after each breakout
 - To keep everyone synchronized

Wednesday 5th November

Location	Start time	End time	Subject	People in charge	Additional info
Lecture hall of the LABRI	10:30	11:30	Inria-Bordeaux seminar open to NumPEX members who are already on-site: lab talk by Franck Cappello, member of the Scientific Board of Exa-DoST EAIRA: Establishing a Methodology for Evaluating AI Models as Scientific Research Assistants	Franck Cappello	Due to the shutdown situation in the US, Franck Cappello's travel has finally been cancelled. Unfortunately, this Inria Bordeaux seminar, which was facultatively open to the participants of our General Meeting, cannot take place.
Ada Lovelace	12:00	12:30	An introduction / a refresher to NumPEX and Exa-DoST (strongly recommended for new recruits)	Gabriel Antoniu Julien Bigot	Particularly useful for new recruits and people who want a refresher on the overall context
Next to Ada Lovelace	12:30	14:00	Standing buffet lunch with posters		
Ada Lovelace	14:00	14:15	A few introductory words for everyone	Gabriel Antoniu Julien Bigot	Explanations regarding the overall agenda, the aim of breakout sessions, etc.
	14:15	14:35	First results and 2 scientific focuses for WP1 - I/O and data storage	Francieli Boito François Tessier	5 min for WP leaders 5 min for each scientific focus (by Exa-DoST recruits) 5 min for Q&A
	14:35	14:55	First results and 2 scientific focuses for WP2 - In situ data processing	Yushan Wang Laurent Colombet	
	14:55	15:15	First results and 2 scientific focuses for WP3 - ML-based data analytics	Thomas Moreau Bruno Raffin	
	15:15	15:35	First results and 2 scientific focuses for WP4 - Shared building blocks and integrated illustrators	Virginie Grandgirard Damien Gratadour	
Next to Ada Lovelace	15:35	16:00	Coffee break with posters		
- Grace Hopper 2, B405 - George Boole 2, B205 - Alan Turing 2, B305	16:00	17:30	Breakout sessions to make applications talk with WPs (1/3) - Gysela x WP1 - I/O and data storage - SKA x WP2 - In situ data processing - Other applications: Coddex, Dyablo... x WP3 - ML-based data analytics	- Virginie Grandgirard x Francieli Boito & François Tessier - Damien Gratadour x Yushan Wang - Laurent Colombet x Thomas Moreau & Bruno Raffin	A few slides
TBD	17:30	18:30	Private exchange with the board People who are not concerned may organize work sessions on their chosen subjects		PC leaders + WP leaders + Board
Next to Ada Lovelace	19:30	21:30	Standing buffet dinner with posters		

Thursday 6th November

Location	Start time	End time	Subject	People in charge	Additional info
Next to Ada Lovelace	08:30	09:00	Welcome coffee		
	09:00	09:30	Talk by François Mazen (Kitware)	François Mazen	Overview of Kitware and potential links with Exa-DoST: in-situ, in-transit, and progressive analysis, Exascale perspectives...
	09:30	10:00	Talk by Xavier Delaruelle (TGCC)	Xavier Delaruelle	Overview of Alice Recoque
Ada Lovelace	10:00	10:10	Feedback on Gysela x WP1 breakout session	Virginie Grandgirard x Francieli Boito & François Tessier	
	10:10	10:20	Feedback on SKA x WP2 breakout session	Damien Gratadour (Shan Mignot) x Yushan Wang	
	10:20	10:30	Feedback on other apps (Coddex, Dyablo...) x WP3 breakout session	Laurent Colombet x Thomas Moreau & Bruno Raffin	
Next to Ada Lovelace	10:30	11:00	Coffee break with posters		
- Grace Hopper 2, B405 - George Boole 2, B205 - Alan Turing 2, B305	11:00	12:30	Breakout sessions to make applications talk with WPs (2/3) - Gysela x WP2 - In situ data processing - SKA x WP3 - ML-based data analytics - Other applications: Coddex, Dyablo... x WP1 - I/O and data storage	- Virginie Grandgirard x Yushan Wang - Damien Gratadour x Thomas Moreau & Bruno Raffin - Laurent Colombet x Francieli Boito & François Tessier	
Next to Ada Lovelace	12:30	14:00	Standing buffet lunch with posters		

Thursday 6th November

Next to Ada Lovelace	12:30	14:00	Standing buffet lunch with posters		
Ada Lovelace	14:10	14:20	Feedback on Gysela x WP2 breakout session	Virginie Grandgirard Yushan Wang Laurent Colombet	
	14:00	14:10	Feedback on SKA x WP3 breakout session	Damien Gratadour Thomas Moreau Bruno Raffin	
	14:20	14:30	Feedback on other apps (Coddex, Dyablo...) x WP1 breakout session	Laurent Colombet Francieli Boito François Tessier	
	14:30	16:00	Plenary session on modularization, centered around representatives from each software component. Application people may have free work sessions in parallel if they wish.	Julien Bigot Bruno Raffin	There are ongoing exchanges on Slack regarding modularization, do not hesitate to contribute ahead of the General Meeting! For this session to be fruitful, representatives of all software components involved in Exa-DoST should come prepared.
Next to Ada Lovelace	16:00	16:30	Coffee break with posters		
- Grace Hopper 2, B405 - George Boole 2, B205 - Alan Turing 2, B305	16:30	18:30	Breakout sessions to make applications talk with WPs (3/3) - Gysela x WP3 - ML-based data analytics - SKA x WP1 - I/O and data storage - Other applications: Coddex, Dyablo... x WP2 - In situ data processing	- Virginie Grandgirard x Thomas Moreau & Bruno Raffin - Damien Gratadour x Francieli Boito & François Tessier - Laurent Colombet x Yushan Wang	
Le Café du Port, 1 quai Deschamps, Bordeaux	19:30	22:30	Dinner in Bordeaux		

Friday 7th November

Location	Start time	End time	Subject	People in charge	Additional info
Next to Ada Lovelace	08:30	09:00	Welcome coffee		
- Grace Hopper 2, B405 - George Boole 2, B205 - Alan Turing 2, B305	09:00	10:45	Breakout sessions on applications: integration of tools in applications and update of the application motifs - Gysela - SKA - Other applications: Coddex, Dyablo...	- Virginie Grandgirard - Damien Gratadour - Laurent Colombet	WP members can choose whichever application they want based on their previous discussions in the first rounds of apps/software breakout sessions.
Next to Ada Lovelace	10:45	11:05	Coffee break with posters	TBD	
Ada Lovelace	11:05	11:15	Feedback on modularization #2 breakout session	TBD	
	11:15	11:25	Last feedback on Gysela (x WP3)	Virginie Grandgirard x Thomas Moreau & Bruno Raffin	
	11:25	11:35	Last feedback on SKA (x WP1)	Damien Gratadour x Francieli Boito & François Tessier	
	11:35	11:45	Last feedback on other applications (Coddex, Dyablo...) (x WP2)	Laurent Colombet x Yushan Wang	
	11:45	12:00	Conclusion	Gabriel Antoniu Julien Bigot	
Next to Ada Lovelace	12:00	14:00	Packed lunch		

Wednesday 5th November

Location	Start time	End time	Subject	People in charge	Additional info
Lecture hall of the LABRI	10:30	11:30	Inria-Bordeaux seminar open to NumPEX members who are already on-site: lab talk by Franck Cappello, member of the Scientific Board of Exa-DoST EAIRA: Establishing a Methodology for Evaluating AI Models as Scientific Research Assistants	Franck Cappello	Due to the shutdown situation in the US, Franck Cappello's travel has finally been cancelled. Unfortunately, this Inria Bordeaux seminar, which was facultatively open to the participants of our General Meeting, cannot take place.
Ada Lovelace	12:00	12:30	An introduction / a refresher to NumPEX and Exa-DoST (strongly recommended for new recruits)	Gabriel Antoniu Julien Bigot	Particularly useful for new recruits and people who want a refresher on the overall context
Next to Ada Lovelace	12:30	14:00	Standing buffet lunch with posters		
	14:00	14:15	A few introductory words for everyone	Gabriel Antoniu Julien Bigot	Explanations regarding the overall agenda, the aim of breakout sessions, etc.
Ada Lovelace	14:15	14:35	First results and 2 scientific focuses for WP1 - I/O and data storage	Francieli Boito François Tessier	5 min for WP leaders 5 min for each scientific focus (by Exa-DoST recruits) 5 min for Q&A
	14:35	14:55	First results and 2 scientific focuses for WP2 - In situ data processing	Yushan Wang Laurent Colombet	
	14:55	15:15	First results and 2 scientific focuses for WP3 - ML-based data analytics	Thomas Moreau Bruno Raffin	
	15:15	15:35	First results and 2 scientific focuses for WP4 - Shared building blocks and integrated illustrators	Virginie Grandgirard Damien Gratadour	
Next to Ada Lovelace	15:35	16:00	Coffee break with posters		
- Grace Hopper 2, B405 - George Boole 2, B205 - Alan Turing 2, B305	16:00	17:30	Breakout sessions to make applications talk with WPs (1/3) - Gysela x WP1 - I/O and data storage - SKA x WP2 - In situ data processing - Other applications: Coddex, Dyablo... x WP3 - ML-based data analytics	- Virginie Grandgirard x Francieli Boito & François Tessier - Damien Gratadour x Yushan Wang - Laurent Colombet x Thomas Moreau & Bruno Raffin	A few slides
TBD	17:30	18:30	Private exchange with the board People who are not concerned may organize work sessions on their chosen subjects		PC leaders + WP leaders + Board
Next to Ada Lovelace	19:30	21:30	Standing buffet dinner with posters		