

TECHNICAL SPECIFICATIONS FOR THE PROCUREMENT OF THE SERVICE OF PORTING THE DYNINST LIBRARY TO THE RISC-V ARCHITECTURE FOR THE BARCELONA ZETTASCALE LABORATORY PROJECT.

File: CONSER02024036OP

SUBJECT OF THE CONTRACT, CONTENT AND FEATURES OF THE SERVICE

The content of these technical specifications derives from the **BSC-CNS Zettaescale Laboratory** project within the framework of the RECOVERY, TRANSFORMATION AND RESILIENCE PLAN forming part of the project for the development of technologies for the generation of prototypes based on RISC-V technology.

By submitting their offer, the bidding company accepts the technical specifications set out at this document.

1. Motivation

The purpose of the following specifications is to define the technical requirements for the procurement of the service for the development of the porting of the *DynInst* dynamic instrumentation library to the RISC-V architecture together with a minimum set of reference use cases.

The main objective of the **BSC-CNS Zettaescale Laboratory** project is the development of new RISC-V based architectures and the proposal of a set of software tools that allow the full use of the developed platforms. While optimisation of the code generated specifically for a target architecture generally takes place in the compiler (and in compile-time), there are opportunities to use tools in run-time that allow for the optimisation of the generated code.

Over the 42 months of the current project duration (between December 2022 and June 2026), the different work teams will be aligned to the milestones established by the *project's technical report* and the *performance reports on the activities carried out* so far. Its action consists of running an experimental development project to promote the generation of research and development relating to Spanish and European technology in the field of advanced processors. The result from this project will be two chip prototypes, technologically compatible with future European supercomputers. These also represent the main milestones for the project:

- **Milestone 1** (scheduled for the first quarter of 2025). This includes all the tasks that lead to the *tapeout* of first chip and its implementation on an evaluation board. In addition to these tasks, there will be a first version of the software adapted to the features of this first chip. The current plan is for this to include a design with a limited number of cores (between 2 and 4) that will allow for testing of the correct operation of all its hardware components: scalar processor, vector co-processor (VPU),

3-level cache hierarchy and other analogue and peripheral components.

- **Milestone 2** (scheduled for the first quarter of 2026). It includes all the tasks that lead to the *tapeout* of the second chip and its implementation on an evaluation board. In addition to these tasks, there will be an optimised version of the software adapted to the features of this second chip. In this case, the design will include more cores and will allow the scalability of the proposal to be demonstrated. The number final of cores will depend on the area available and the associated cost, but the current planning includes a minimum of 16 cores organised on a 4x4 grid. This prototype will allow HPC applications to be run using all the chip's cores, demonstrating the operation of the design on a large scale.

2. Description of the service

The dynamic instrumentation allows the modification of the binary code of a computer application, both while running and on the disk. The changes made can be used for different purposes, from mere instrumentation through other tools already included in the software stack (e.g., Extrae), to changes in the planning of instructions or parameters at run time.

In addition to the porting of the DynInst library to the RISC-V architecture, the project includes the implementation, by the bidder, of three pilot use cases to guide the development of the project software included in the software stack, and to evaluate and improve the performance of codes on the platform. The following are the possible applications, as well as the use cases of interest envisaged.

In relation to instrumentation through other tools, the Extrae instrumentation library has in the past used the existence of the DynInst library in other architectures for the generation of the Paraver trace. Porting of the DynInst component to the RISC-V architecture will enable this instrumentation mechanism. This instrumentation mechanism is not a use case to be developed since Extrae has previous experience in using the DynInst programming interface.

As for the instrumentation of code itself, DynInst would allow the generation of a performance profile of a given code based on basic blocks. This profile would include information such as: the number of times a basic block is run and its duration. It could also be elaborated with other information that the compiler team might consider of interest for use in a new, this time guided, compilation of the target code. This performance profile based on basic blocks constitutes the first use case of interest.

DynInst would also allow the re-planning of instructions in the binary code generated. An example could be to advance the memory instructions (load and store), increasing the performance of the code execution; especially on architectures with limited bandwidth, with limited out-of-order execution capabilities, and/or that only allow a few on-the-fly memory instructions. Re-planning of instructions is the second use case of interest.

Finally, DynInst may also allow the modification of code control values. An example could be the modification of the *vector length* parameter in the VLEN instruction set. In certain cases, the optimal value of this parameter does not have to match the maximum allowed by the architecture. The modification of code control values is the third, and final, use case of interest.

During the running of the project, further use cases may be agreed with the bidding company to allow the development of new forms of instrumentation and/or performance improvements. The tender process also envisages proposals for additional use cases.

It is estimated that the bidding company will need to allocate a minimum of two people to the proper execution of the project: a software analyst (supervisor) and a software developer. The formation of teams will be left to the discretion of the bidding company.

3. Procedure for the provision of the services, monitoring and control

The successful bidder will be responsible for the porting of the DynInst library as well as the development of the different use cases finally agreed in the contracting of this service.

The successful bidder will provide access to the necessary documentation for the DynInst library that facilitates the development and implementation of the tools that exploit the use cases finally agreed in the procurement of this service.

The successful bidder will provide access to the most updated stable version of the DynInst library from Milestone 2: "*Partial version of the porting of the library*" and to the end of the service.

The successful bidder will make all commercially feasible and reasonable efforts to deliver a fully functional version of the library at Milestone 3: "Final porting candidate version".

The monitoring and control of the development of the service will require coordination between the BSC-CNS technical team and the successful bidder. For this reason, the BSC-CNS will appoint a manager from within its organisation. In the same way, the successful bidder must have a service manager, who will be the contact point for the manager appointed by the BSC-CNS.

These people must meet at least on a MONTHLY basis to supervise, monitor and discuss any aspect related to the development of the contract, in order to ensure that it is being executed in accordance with these specifications.

4. Do No Significant Harm (DNSH) principle

The contractor is obliged to comply with the obligations imposed through the application of the “do no significant harm” (DNSH) principle. The contractor, in compliance with the provisions of the Recovery Plan, Regulation (EU) 2021/241 of 12 February 2021 establishing the Resilience and Recovery Mechanism, and its implementing regulations, in particular the Commission Notice (2021/C 58/01) Technical guidance on the application of “do no significant harm”, as well as with the requirements of the Council Implementing Decision on the approval of the assessment of the Spanish Recovery and Resilience Plan (CID), is obliged to ensure that all funded actions to be carried out under this contract must respect the principle of *Do No Significant Harm* to the environment (DNSH principle). This includes compliance with the specific conditions provided for in component 15, Investment C15.I5 in which it is framed and especially in the Annex to the CID and those set out in sections 8 and 8 of the Plan’s Component document.

5. Other special conditions

At the end of the contract, a report on all the actions carried out may be requested, which must be provided by the successful bidder within one month from the day following receipt of the request.

Xavier Teruel
*Established Researcher in
the Best Practices for Performance and Programmability group
from the Department of Computer Sciences*

Barcelona, 9 September 2024