ADVANCED GRAPHIC INTELLIGENCE LOGICAL COMPUTING ENVIRONMENT (AGILE)

Frequently Asked Questions

Question #	Question	Answer
	1. General Questions about the Program	
1.1	Are foreign nationals eligible to participate?	Refer to the BAA Section C.1 "Elligible Applicants."
1.2	Are quantum or analog devices in scope?	Quantum devices are not in scope. Analog devices are in scope. Models of the behavior of analog devices must be delivered to the T&E teams for evaluation in either Verlog-A or SPICE format.
	2. Questions about workflows, kernels and algorithms	
2.1	Are we allowed to modify the application source code including language, programming model etc?	Answer: Yes. You should express the workflows, derived kernels, and industry standard kernels as appropriate for your architecture; however, you must get the same answers as the reference implementations do for each data set. Note, the answers for some workflows and derived kernels will not be exact. For those workflows and derived kernels, we will work with you to insure that your answers are of the same quality.
2.2	Are we allowed to use different data structures than the ones used by the reference implementations?	Yes. You can use, organize and store the input data as you wish.
2.3	Can the derived kernels use different data structures than the workflows?	No. The source code of a derived kernel must replicate faithfully the data organization and machine state at the call site in the workflow. Any reorganization of the data that occurs in the workflow before the kernel is called must be included in the derived kernel code and must be timed.
2.4	Can an implementation of a standard industry kernel use a different data organization than the reference implementation?	Yes, but any reorganization of the data from its input order you do must be timed. Note the execution time of the reference implementation will include any reorganization from input order that it does.
2.5	Can we get different input sizes to benchmark that you believe are representative?	Yes. We will release different size data sets for each workflow and derived kernels.
2.6	For some of the AGILE analytics problems, are non-exact answers allowed? If so, how much deviation from the solutions provided are we allowed?	We will identify workflows and derived kernels that return exact answers and those that do not. For the former, non-exact answers are not allowed. For the latter, we will work with you to insure that your answers are of the same quality.

2.7	What programming language and inter-node communication library should performers use?	The T&E team will provide the performers with GFI examples of the applications written in C++17 and some supporting libraries. However, these are only to provide a concrete example of the target workflow and the performers are by no means confined to using the languages or communication libraries used for the workflow examples so long as the performer supply all compilers and supporting libraries to the T&E teams so that the resulting workflows can be evaluated.

3. Questions about emulation environment and baseline system

3.1	Am I allowed to use other hardware for FPGA runs (on our own hardware), e.g. can I use Altera FPGAs or other non-FireSim based frameworks?	You can use any FPGA hardware you want for your own testing and development, but artifacts that are delivered to T&E must be platform neutral (e.g. IEEE standard Verlog/SystemVerilog that can be compiled on any FPGA platform or logic synthesis tool.)
3.2	Can performers have access to the FPGA emulation platform for testing and scaling their models.	FireSim (an open source FPGA emulation platform based on Xilinx U2xx series FPGA) is available and openly accessible on Amazon F1 instances so that the performers have equal access to FPGA emulation resources. The program is also considering opening up access to a private emulation cluster comprised of Xilinx U250 and U280 instances.
3.3	What is the baseline system that our performance improvements will be measured against.	The baseline will be a contemporary CPU and/or GPU accelerated cluster system. The NERSC Perlmutter system (https://www.nersc.gov/systems/perlmutter/) and the FireSim RISC-V cluster will serve as a concrete example of specific systems, but any contemporary platform could serve as a baseline.

4. Questions about A-SST simulation environment

4.1	Are we allowed to modify the A-SST element source code when constructing models?	Yes, provided any modifications are supplied as a diff/patch file against a specific SHA to the AGILE T&E team so the changes can be reviewed and used during evaluation.
4.2	Are performers allowed to develop sub-components for A-SST elements to demonstrate their models?	Yes, performers are free to create any components, sub-components, models or modules for A-SST provided all source code to these models is provided to the T&E team for evaluation.
4.3	Are we allowed to link other simulators to A-SST?	Yes, provided the full source code is made available to the T&E team. Performers may not supply pre-compiled binaries for any elements of their simulation tools or models.
4.4	Do performers have to use a specific version of A-SST during AGILE?	Updated releases and bug fixes for A-SST will be supplied throughout the AGILE program. Performers are encouraged to utilize the latest releases wherever possible. If version updates conflict with implementations, specific details should be discussed with the T&E team.