Accelerating Scientific Model Optimization with a Pipelined FPGA-Based Differential Evolution Engine

Manuel de Castro[1], Roberto R. Osorio[2],
Yuri Torres[1], Diego R. Llanos[1]

[1]Universidad de Valladolid, Spain, [2]Universidade da Coruña, Spain

{manuel|yuri.torres}@infor.uva.es

roberto.osorio@udc.es diego.llanos@uva.es

FCCM 2025, May 2025

Introduction

- Differential Evolution (DE) is a genetic algorithm that, when combined with Numerical Integration of scientific model on application-specific hardware:
  - Produces deep floating-point pipelines
  - Requires minimal external memory bandwidth
  - Benefits from large internal memory bandwidth

  } **FPGA's strong points!** 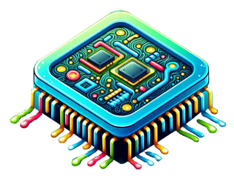
- Proposal: **Generic architecture for Diferential Evolution with Numerical Integration methods**
  - Evaluated with 2 scientific models: Hodgking-Huxley (with non-adaptive 4th order Runge-Kutta) and Circadian (with adaptive Runge-Kutta-Fehlberg)

Differential Evolution

1. **Mutation:**
   Alter each solution's parameters, using the others (with randomness)
   $$m_{i,G+1} = s_{r_1,G} + F \cdot (s_{r_2,G} - s_{r_3,G})$$
2. **Crossover:**
   New candidate solution, hybrid from original solution and mutant
3. **Limiting:**
   Ensure all parameter values are in valid range
4. **Selection:**
   Compute cost function of new candidate; keep solution with lower cost

Scientific models used for assessment

- **Hodgking-Huxley**: action potentials in neuron axons
  - Multiple non-linear differential equations
  - Always converges with non-adaptive RK4
  - 11 unknown parameters per solution
- **Circadian**: biological clock of plants
  - Multiple coupled stiff differential equations
  - Requires adaptive numerical integration: RKF
  - 27 unknown parameters per solution

× Complex, real-world models, not tackled previously in the literature in a fully-pipelined way

Resource usage

- Synthesized for Virtex Ultrascale+ XCVU13P

	LUT	FF	BRAM	DSP
Hodgkin-Huxley	220 123 (13%)	230 438 (7%)	114 (4%)	2 272 (18%)
Circadian	648 976 (38%)	659 083 (12%)	293 (11%)	4250 (37%)

- More than $\times 10$ LUT, BRAM; $\times 100$ FF, DSP than previous works

Architecture overview

- Generic: Adjustable via meta-parameters (number of variables to optimize, number of solutions to consider...)
- Modular design, for ease of adaptation
- Fully pipelined, high pipeline utilizaion through Islands of Solutions
- *Numerial integration* coded with HLS ⇒ ease of development, hassle-free FP computations
- Other modules coded in VHDL ⇒ high control and performance

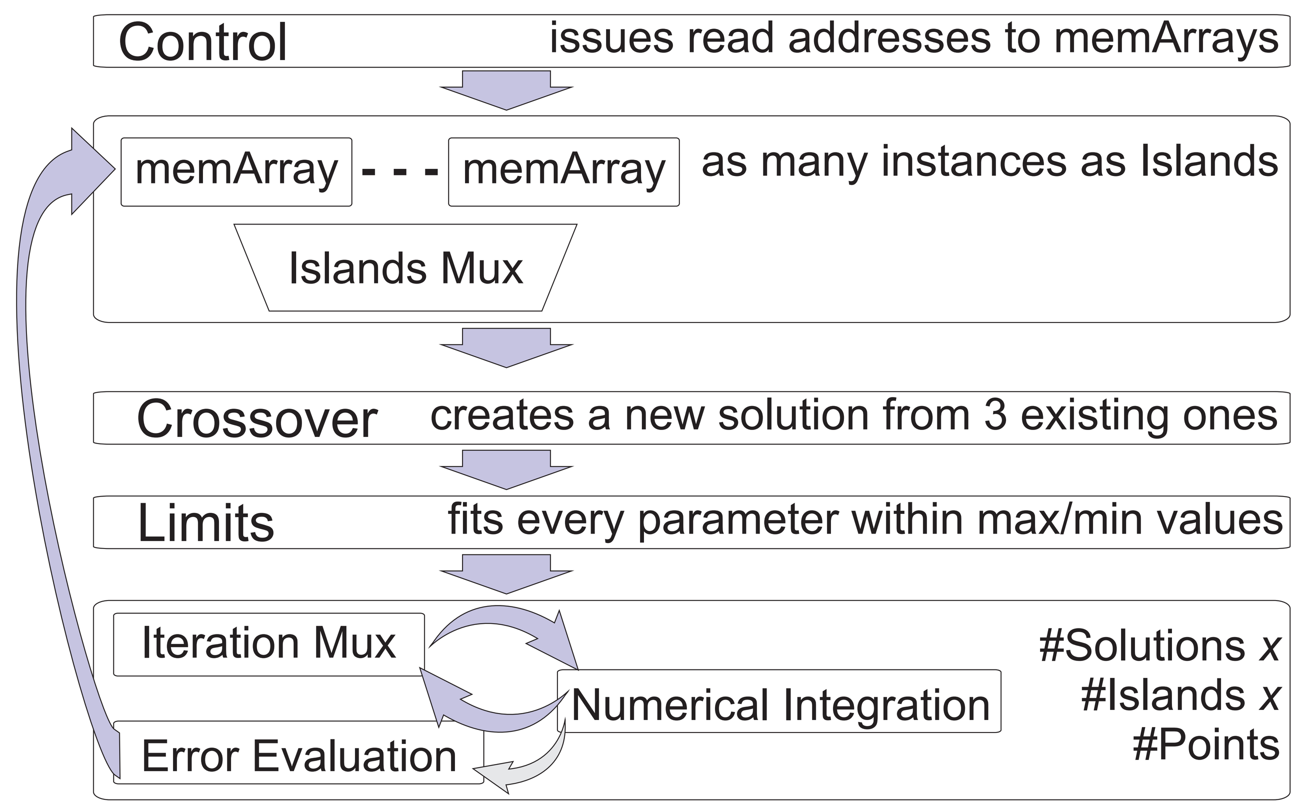

× *Control*: Initializes state, manages *memArrays*, reads results

× *memArray*: Stores parameters values and solution cost; one per solution

× *Crossover*: Hybridizes solutions (**Mutation** and **Crossover** steps)

× *Limits*: Manages out-of-range values from *Crossover* (**Limiting** step)

× *Iteration Mux*: Manages iterations of *Numerical Integration*

× *Numerical Integration*: Obtains models' results from candidate solutions

× *Error Evaluation*: compares results against observed data (**Selection** step)

Performance

- Tested: Multiple number of candidate solutions, multiple amount of numerical integration iterations
- Speedup with respect to single-threaded CPU implementation, executed in an AMD EPYC 7713 @ 3.67 GHz, compiled with *gcc* and *-O3* optimizations
- Comparison against **NVIDIA A100 GPU**, using a GPU-optimized high-performance implementation: One thread per candidate solution

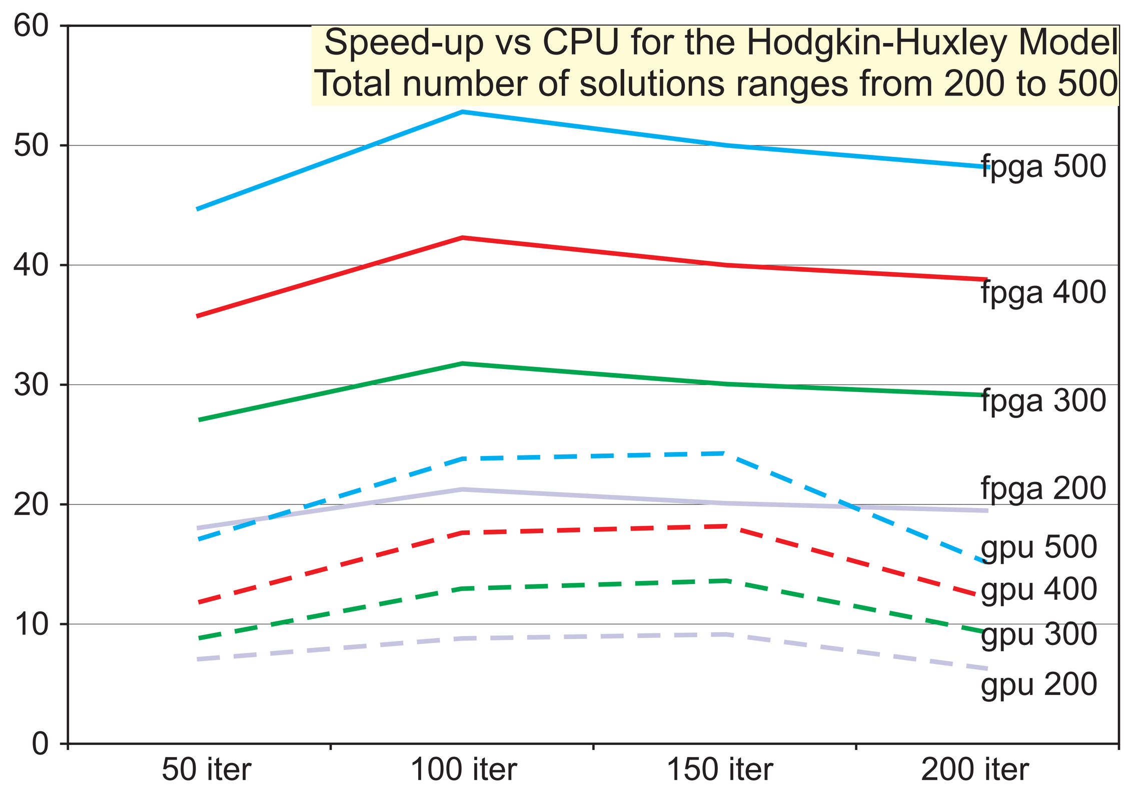

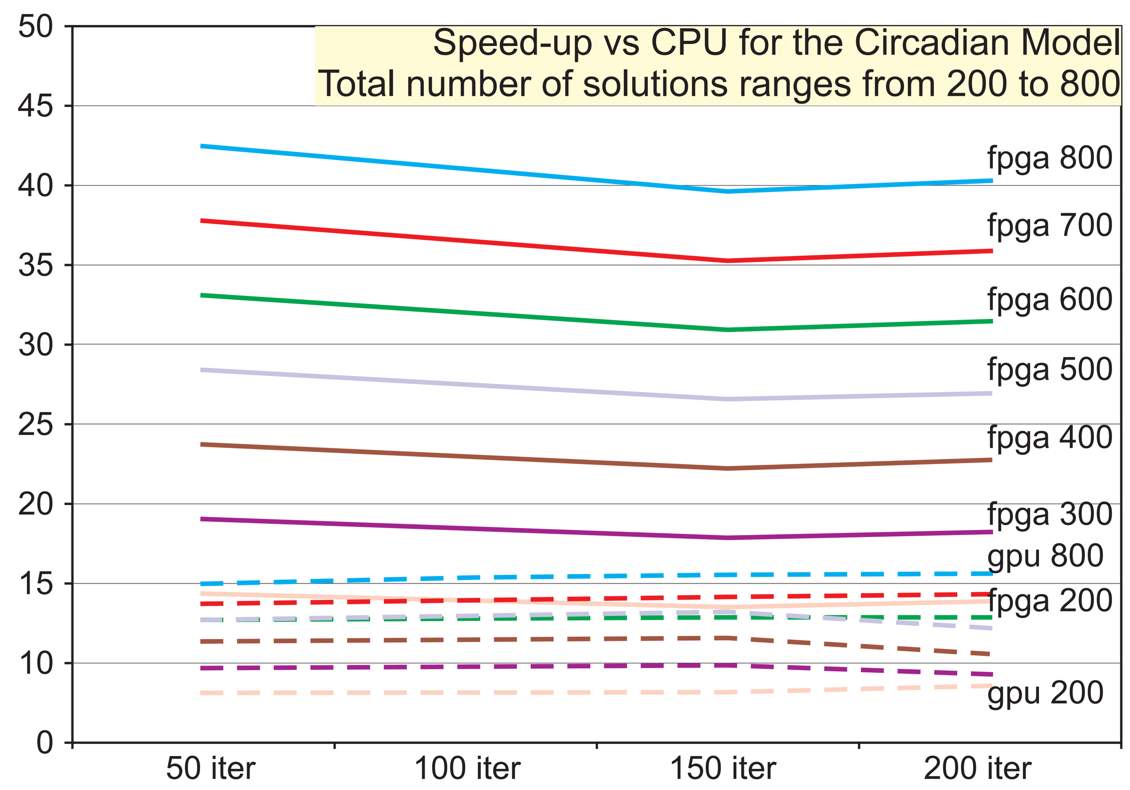

- FPGA and GPU performance scale with the number of candidate solutions
  - FPGA scaling is more linear
- Average energy savings vs GPU:
  - **Hodgkin-Huxley**: $\times 9.0$ more energy efficient (11.04% the energy consumption of the GPU)
  - **Circadian**: $\times 4.11$ more energy efficient (24.33% the energy consumption of the GPU)

Conclusion and future work

- Differential Evolution with Numerical Integration is well suited for FPGA-based CCMs: Parallelism, independent memory accesses, deep floating-point pipelines
- We present a generic, parameterized, modular architecture for Differential Evolution
- Our architecture presents superior performance to GPU-based high performance solutions
- **Future work:** optimizing engine performance and resource usage, investigating new fitness functions, and comparing against multicore CPU implementations

References

Acknowledgements

This work was supported by Spanish MICIU and the ERDF Program of the EU, (NATASHA Project, Grant PID2022-142292NB-I00); by Junta de Castilla y León-FEDER Grants, (PROPHET-2 Project), Junta de Castilla y León, Spain under Grant VA226P20; by MCIN/AEI/10.13039/ 501100011033 through "European Union NextGenerationEU/PRTR," (Grant TED2021–130367B–I00); and by MCIN/AEI/10.13039/501100011033 through the ERDF Program of the EU (Grant PID2022-136435NB-I00).

Manuel de Castro was supported by Spanish MICIU through Ayudas para la Formación de Profesorado Universitario FPU 2022.

This research was supported by grants from NVIDIA (A100). The authors would like to thank Sergio Alonso Pascual for his help developing the GPU software used.