

**13th International Workshop on Programming Models and Applications
for Multicores and Manycores (PMAM'22)**

A performance-oriented comparative study of the Chapel
high-productivity language to conventional programming environments

Guillaume Helbecque¹ (speaker), Jan Gmys¹, Tiago Carneiro², Nouredine Melab¹, Pascal Bouvry³

¹Université de Lille
CNRS/CRISTAL
Inria Lille Nord-Europe
France

²Université du Luxembourg
FSTM
Luxembourg

³Université du Luxembourg
DCS-FSTM/SnT
Luxembourg

Background

- Heterogeneity (CPU-GPU) and size of **modern supercomputers** (millions of cores) [1]:

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	Sierra - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438



Top 3 of the Top500 bi-annual ranking (Nov. 2021)

- Emergence of **high-productivity programming languages** [2]:
 - Chapel, X10, Fortress, *etc.*

Motivations and objectives

- **Compare Chapel to conventional parallel programming libraries**, in terms of performance...
 - on both shared- and distributed-memory systems;
- Illustrate the programming effort in each parallel environment... and **provide a sense of "productivity"**;
- **Provide a useful data point** using shared- and distributed-memory multi-core systems for supercomputer programmers...
 - through a well-known and complete parallel application.

Outline


- Formulation of the test-case
- Parallel programming environments and implementations
 - OpenMP
 - Chapel
 - MPI
 - Hybrid MPI+OpenMP
- Experimental evaluation
 - Testbed
 - Shared-memory experiments
 - Distributed-memory experiments
 - Parallel overheads
- Conclusions and future works

Test-case: the Mandelbrot set

It is defined as the set of complex numbers $c=a+ib \in \mathbb{C}$ such that the sequence $(z_n)_{n \in \mathbb{N}} \subset \mathbb{C}$ defined by

$$z_0 = 0, \quad z_{n+1} = z_n^2 + c, \quad (1)$$

remains bounded in \mathbb{C} .



Monochrome Mandelbrot set

```
1 function Compute_pixel(a, b) :  
2   x = y = 0;  
3   n = 0;  
4   while x2 + y2 < 4 and n < N do  
5     t = x;  
6     x = x2 - y2 + a;  
7     y = 2ty + b;  
8     n = n + 1;  
9   end  
10  I(a, b) = n/N;  
11 end
```

Pseudo implementation of the Mandelbrot
set computation

- Embarrassingly parallel application, due to the **independency between pixels**;
- **Domain decomposition** method, **along the lines**;
- **Static decomposition**;
- Lines are mapped in **round-robin** fashion.

Open Multi-Processing (OpenMP)

- OpenMP is an **application programming interface**;
- It is designed for:
 - **Ease of programming**;
 - **High performance**, and;
 - **Portability**;... only on shared-memory systems;
- OpenMP supports a **multithreaded execution**, through a fork-join model;
- It provides **simple high-level constructs**, for work-sharing among threads for example;
- For more details [3]: <https://www.openmp.org/>.

```
1 function Compute_image_omp() :  
2   #pragma omp parallel for schedule(static, 1);  
3   for a = 0 to nb_lines do  
4     for b = 0 to nb_columns do  
5       | Compute_pixel(a, b);  
6     end  
7   end  
8 end
```

Pseudo OpenMP implementation of the
Mandelbrot set computation

Cascade High Productivity Language (Chapel)

- Chapel is a **PGAS-based language**;
- It is designed for:
 - **Ease of programming**;
 - **High performance**, and;
 - **Portability**;
- Chapel supports a **multithreaded execution model** and allows:
 - **Data/task parallelism**;
 - **Locality control**, *etc*;
- Here Chapel follows the **SPMD execution model**;
- For more details [4]: <https://chapel-lang.org/>.

```
1 function Compute_image_chpl() :  
2   coforall loc = 0 to nb_locales do  
3     on loc do  
4       coforall rank = 0 to nb_tasks do  
5         for a = loc.id + rank * nb_locales to  
6           nb_columns by  
7             nb_tasks * nb_locales do  
8               for b = 0 to nb_columns do  
9                 Compute_pixel(a, b);  
10              end  
11            end  
12          end  
13 end
```

Pseudo Chapel implementation of the
Mandelbrot set computation

Message Passing Interface (MPI)

- Message-passing **application programming interface**;
- MPI is widely used in the academic and industrial areas:
 - **Portability**;
 - **High performance**, and;
 - **Standardization**;
- MPI defines a **communication protocol among processes** running on distributed-systems:
 - Point-to-point or **two-sided**;
 - **One-sided**;
- Here MPI follows the **SPMD execution model**;
- For more details [5]: <https://www.open-mpi.org/>.

```
1 for l = rank to nb_lines by commsize do
2   Compute_line(l);
3   if rank ≠ 0 then
4     MPI_Send(/*args*/);
5   else
6     MPI_Recv(/*args*/);
7   end
8 end
```


Pseudo MPI two-sided implementation
of the Mandelbrot set computation

```
1 MPI_Win_create(win);
2 MPI_Fence();
3 for l = rank to nb_lines by commsize do
4   Compute_line(l);
5   if rank ≠ 0 then
6     MPI_Put(/*args*/);
7   end
8   MPI_Fence();
9 end
```

Pseudo MPI one-sided implementation
of the Mandelbrot set computation

Hybrid MPI+OpenMP

- OpenMP and MPI are complementary...
 - OpenMP is used for intranode parallelism;
 - MPI is used for distribution across nodes (1 MPI process/node).



Hybrid MPI+OpenMP approach

Experimental environment


- Grid'5000 French national testbed [6]:
 - **6 computer nodes** allocated;
 - **32 AMD EPYC 7301 CPUs @2.20GHz** / nodes;
 - **25 Gbps Intel Ethernet Controller XXV710 network**;
- gcc 10.2.1, Open MPI 4.1.0, OpenMP 4.5, Chapel 1.25.0;
- Compile options:
 - gcc `-O2` optimization flag;
 - Chapel `--fast` optimization flag;
- Chapel multi-locale configuration:



Variable	Value
CHPL_RT_NUM_THREADS_PER_LOCALE	64
CHPL_TARGET_CPU	<i>native</i>
CHPL_HOST_PLATFORM	<i>linux64</i>
CHPL_LLVM	<i>none</i>
CHPL_COMM	<i>gasnet</i>
CHPL_COMM_SUBSTRATE	<i>udp</i>
GASNET_PSM_SPAWNER	<i>ssh</i>

Shared-memory experiments


- Fixed image size : 768x1024 pixels;
- **1 to 32 CPU cores (hyperthreading enabled, 2 threads/core);**
- N controls the granularity;
- 5 different implementations.



Speed-up achieved by all five shared-memory implementations

Distributed-memory experiments


- Fixed image size : 3840x5120 pixels;
- **1 to 192 CPU cores** (hyperthreading disabled);
- N controls the granularity;
- 5 different implementations.



Speed-up achieved by all five distributed-memory implementations

Parallel overheads

- **N=1** to highlight the communication overheads;
- Fixed image size: 768x1024 (shared) and 3840x5120 (distributed);
- Chapel relies by default on the *qthreads* tasking layer [7].



Computational overhead measured by all implementations in shared- (left) and distributed-memory (right) experiments

Conclusions and future works

- Chapel...
 - **outperforms its counterparts in shared-memory**, and;
 - **competes with hybrid MPI+OpenMP** in distributed-memory.
- The *qthreads* default tasking layer of Chapel could explain these performances...
 - although it seems to **suffer from the lack of high-performance network** between nodes.
- We plan to investigate...
 - **more complex benchmarks**, involving message aggregation and data replication;
 - other Chapel features, such as **distributed iterators**, **specific data structures**, *etc.*

Some references

- [1] TOP500 ranking. <https://www.top500.org/>
- [2] E. Lusk and K. Yelick. 2007. Languages for High-Productivity Computing: the DARPA HPCS Language Project. *Parallel Processing Letters* 17 (2007), 89–102.
- [3] OpenMP API. *The OpenMP API specification for parallel programming*. <https://www.openmp.org>
- [4] Open MPI. *Open MPI: Open Source High Performance Computing*. <https://www.open-mpi.org>
- [5] Chapel. *The Chapel Parallel Programming Language*. <https://chapel-lang.org>
- [6] Grid'5000. <https://www.grid5000.fr>
- [7] K. Wheeler, R. Murphy, D. Stark, and B. Chamberlain. 2011. The Chapel Tasking Layer Over Qthreads.

Thank you for your attention !

Any question ? Any remark ?

`guillaume.helbecque@univ-lille.fr`