



P-Aevol: an OpenMP Parallelization of a Biological Evolution Simulator, Through Decomposition in Multiple Loops

T. Gautier, C. Perez, J. Rouzaud-Cornabas, L. Turpin

16th International Workshop on OpenMP, IWOMP 2020
Austin, TX, USA, September 22–24, 2020




A Field in Biology: Experimental Evolution



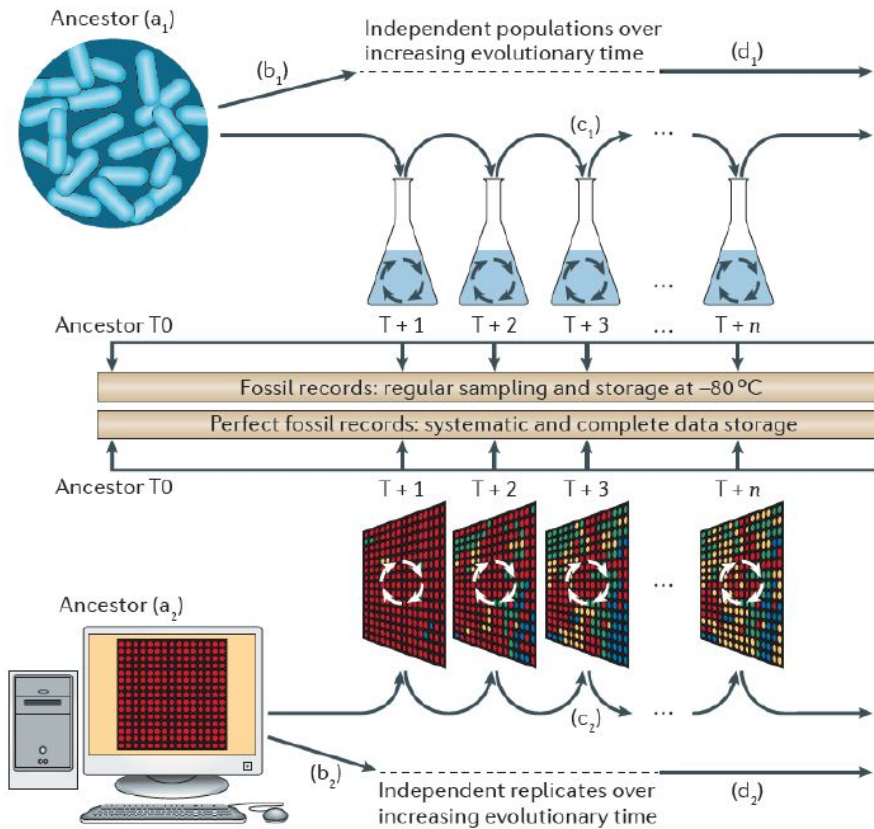
Long Term Evolution Experiment

- Richard Lenski (MSU/Beacon Center, USA)
- running since 1988
- 73,500 generations

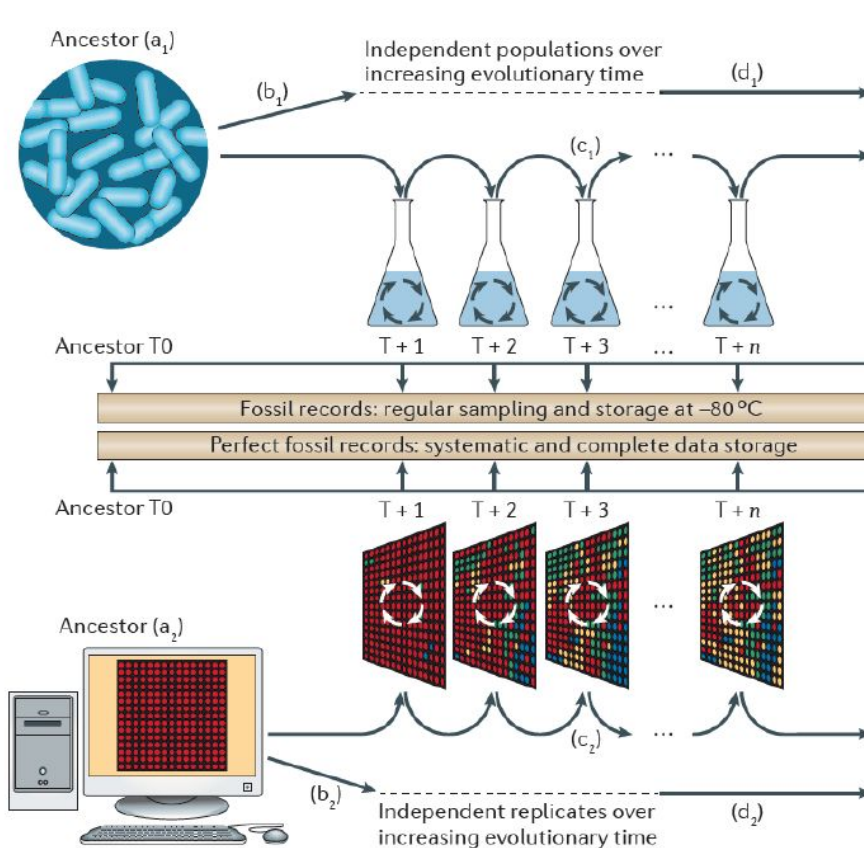
From *in-vitro* to *in-silico*



From *in-vitro* to *in-silico*



From *in-vitro* to *in-silico*

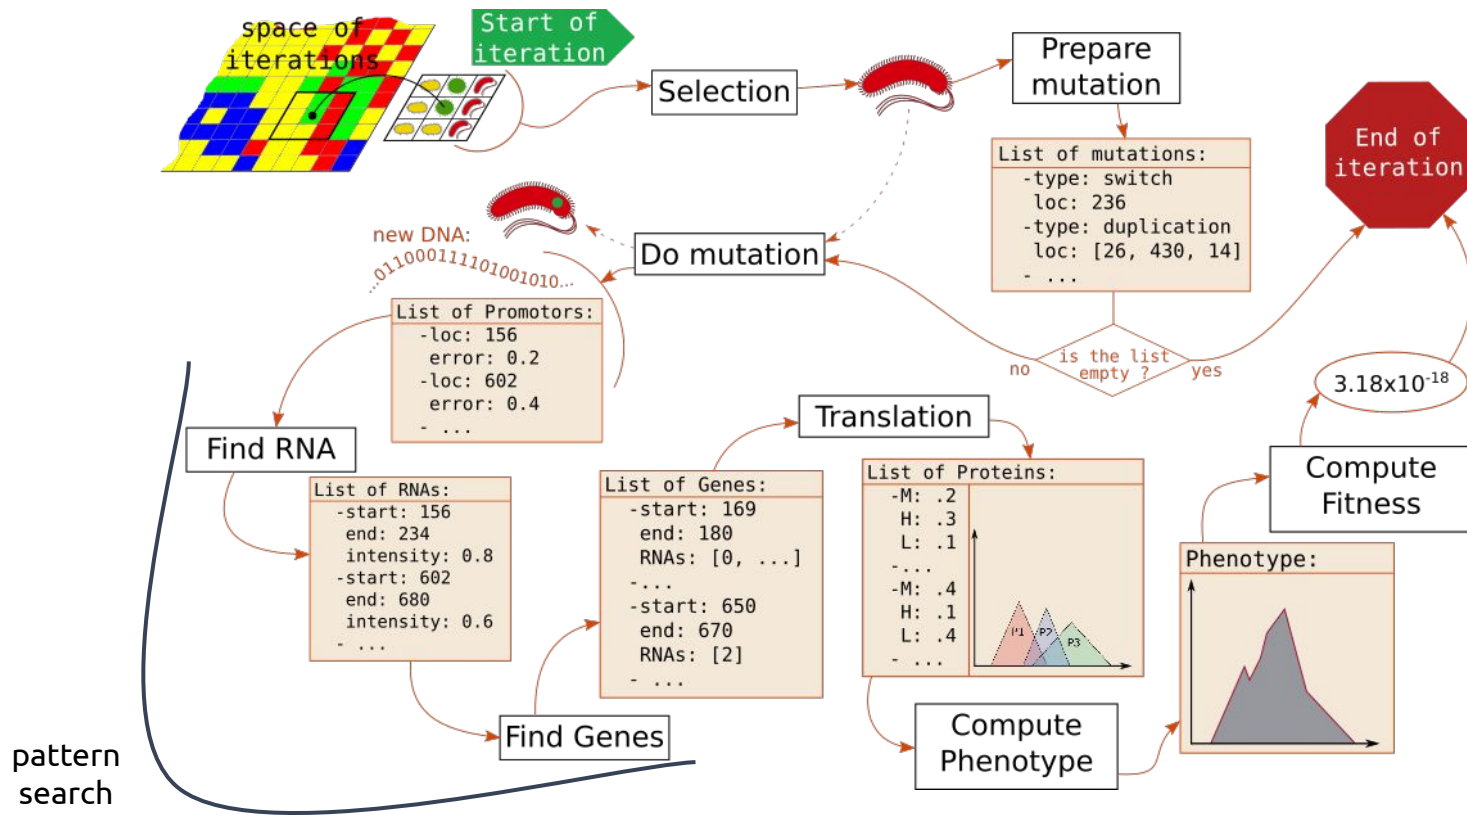


available at
aevol.fr

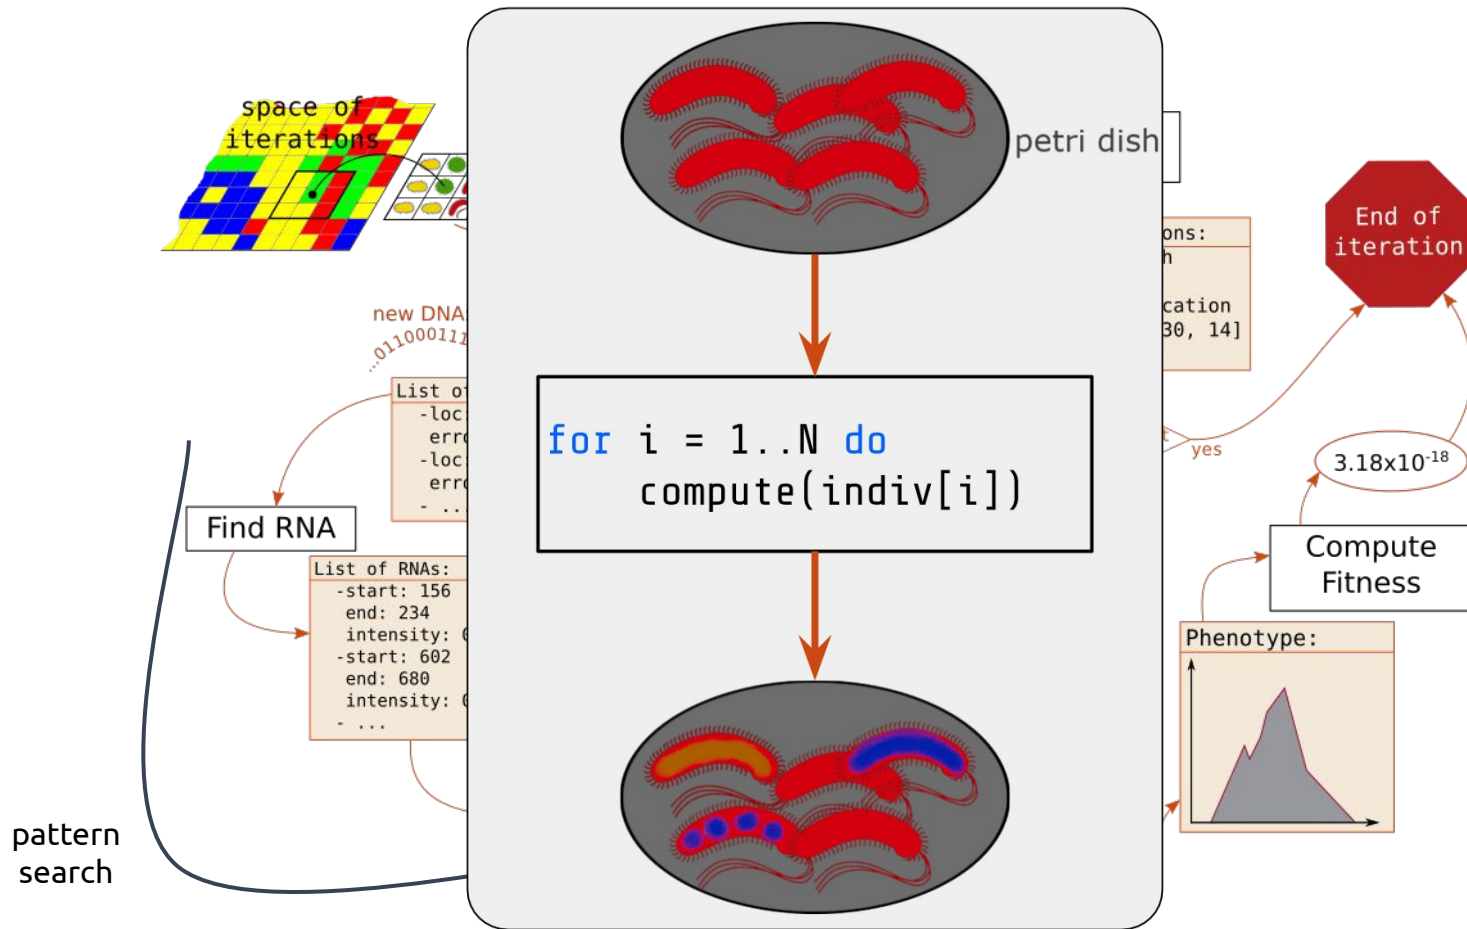
A biological simulator

- Around 50,000 C++ LOC on mono-node arch
- Simulate the evolution of micro organisms
- Compute sequentially one generation after the other
- For one experiment:
 - thousands of hours of computation
 - Terabytes of data (not I/O intensive though)
- Simulate millions of generations
 - around 30ms per generation (1024 individuals)
- **Goal** : Accelerate the computation of a generation

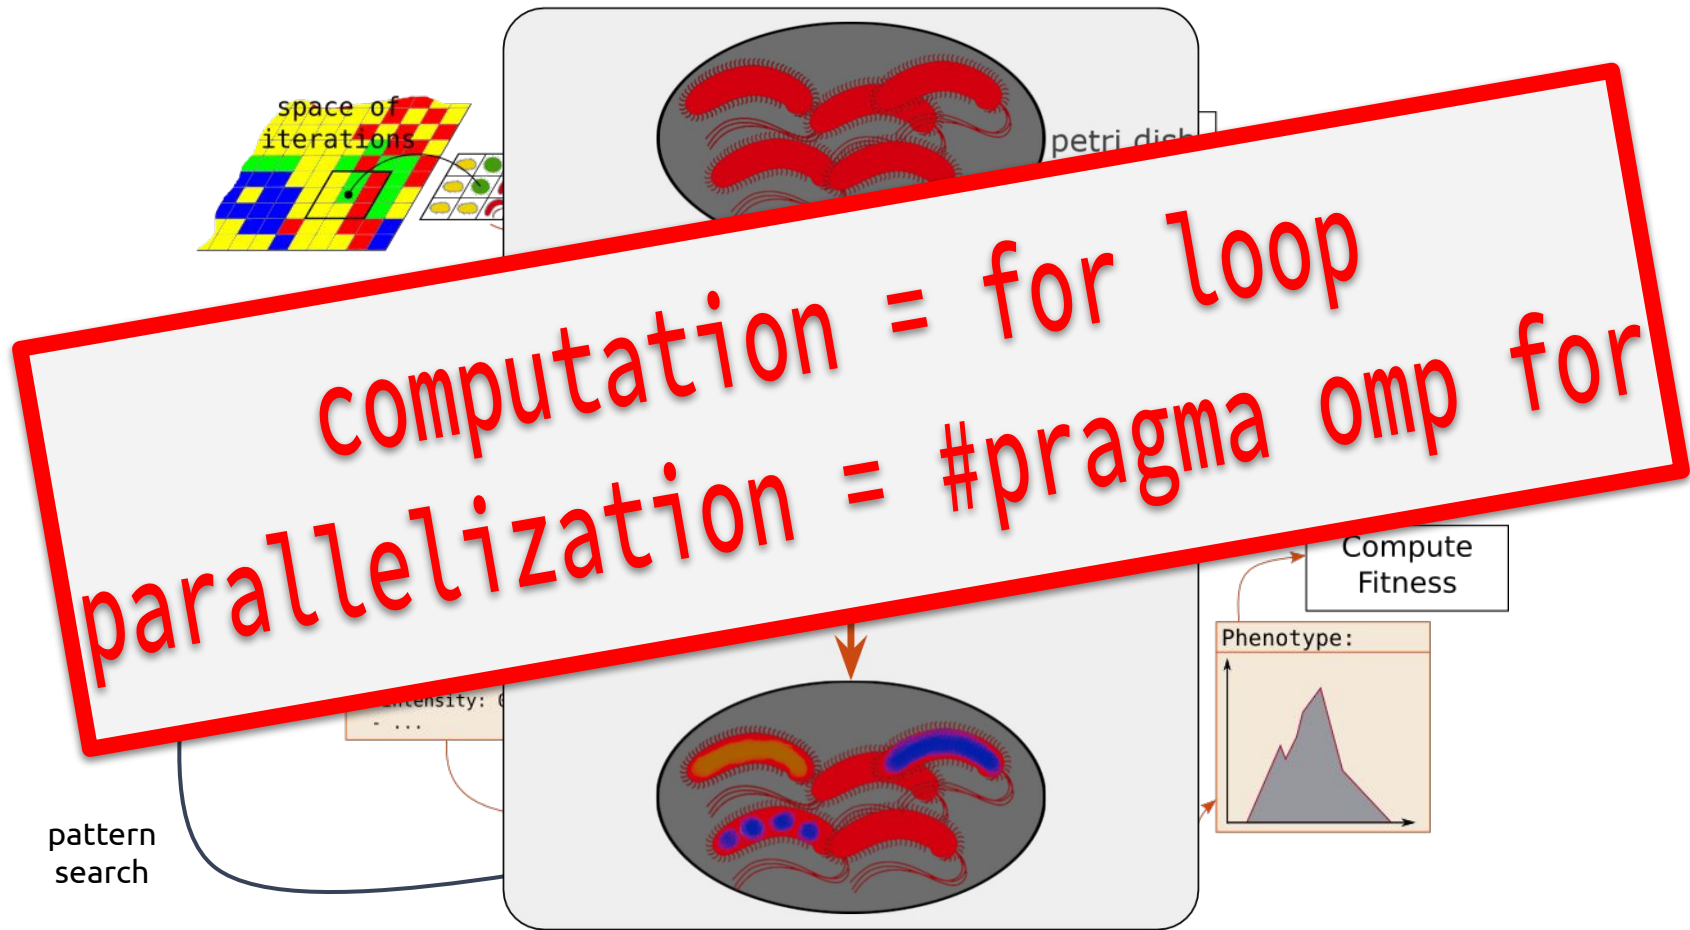
AEVOL: Workflow of a Generation



AEVOL: Workflow of a Generation



AEVOL: Workflow of a Generation



Disappointing Performance

How?


- OpenMP // loop + dynamic schedule
- GCC/libGOMP 8.3
- 64 core single node machine

Result

- More than 20% of idle time
- Speed up less than 36 on 64 cores
- Worse results with other OMP schedulers
 - Static or Guided

Why?

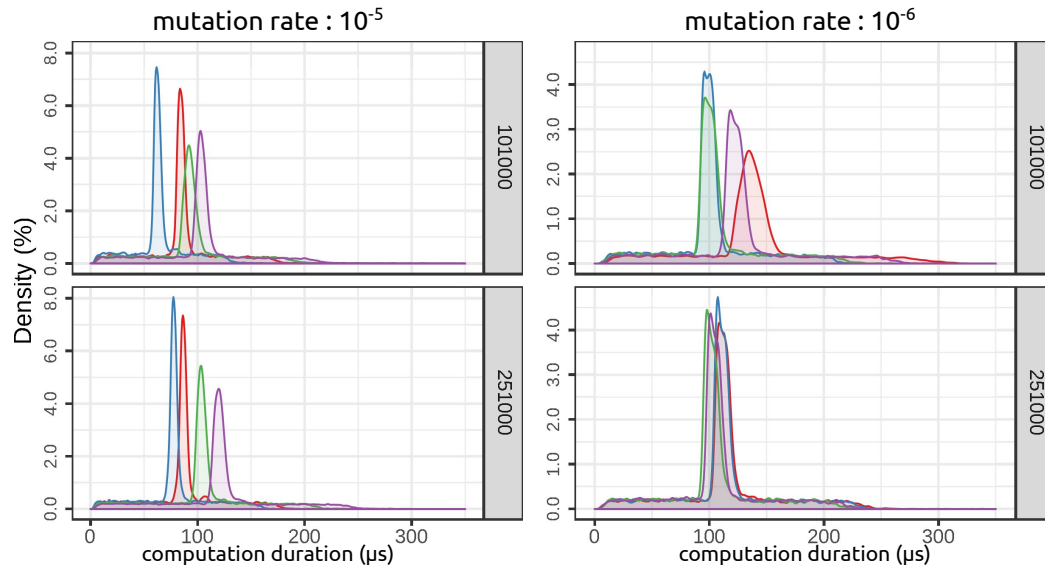
- Individual computation times
 - **irregular** (from 1 μ s to 1,000 μ s)
 - **can vary** from one generation to the other
 - **unpredictable** (stochastic simulation)



Characterization of the irregularity

Distribution of the computation time of mutants

Multiple colors stand for multiple experiments with different random number generator seeds



Generation
#101000

Generation
#251000

Mutations

- Stochastic events on the DNA (char array)
- Probability of occurrence depends on the **mutation rate** input argument and the current size of the DNA

2 Types of Population

- Non-Mutants
 - Take only 1% of the computation time
- Mutants
 - Take 99% of the computation time

Our Problematic

- Goal
 - Improve performance of irregular, varying, and unpredictable iterations
- Scheduling challenges
 - Reduce idle time and work inflation
- How to tackle it?

Scheduling independent iterations

- List scheduling [Graham 1966]: Dynamic
- Base on bin-packing [Hochbaum & Shmoys 1987]
 - Complex to implement outside the runtime
- LPT [Graham 1969]
 - Simplicity and Robustness [Coffman & Seti 1976]

OpenMP loop scheduler

- Internal modification of an OpenMP runtime (libGOMP) [Durand 2013, ...]
- Passing information from application to OpenMP loop scheduler [Penna 2019]
 - Authors promote application with almost constant workload


Loop regularisation through decomposition

```
1 #pragma omp for schedule(dynamic, 1)
2 for i = 1..N do
3   fitness[i] = f_n ∘ ... ∘ f_2 ∘ f_1(indiv[i])
```

Speed up ~ 33

Suppositions

- Output from f_k may predict time of f_{k+1}
- Output from f_k may predict time from f_{k+1} to f_{k+j}





Loop regularisation through decomposition

```
1 #pragma omp for schedule(dynamic, 1)
2 for i = 1..N do
3   fitness[i] = fn ∘ ... f2 ∘ f1(indiv[i])
```

Speed up ~ 33

Suppositions

- Output from f_k may predict time of f_{k+1}
- Output from f_k may predict time from f_{k+1} to f_{k+j}




```
1 for i = 1..N do [in parallel]
2   r[i] = fk ∘ ... f2 ∘ f1(indiv[i])
3   schedule = compute_schedule(r)
4   for i = 1..N do [in parallel with schedule]
5     fitness[i] = fn ∘ ... fk+2 ∘ fk+1(r[i])
```

Decomposition in 2 loops

- 1st loop to predict the load for the iterations of the 2nd loop
- Compute a schedule based on the prediction with a more clairvoyant strategy
 - LPT (Graham, 1966)
- 2nd loop is executed with the previously computed schedule

One data to explain them all




Case of Aevol

- Linear relation between duration and size of DNA after mutation
- Linear relation depends on the generation
- Sufficient for an LPT schedule generation after generation
- Let's call it LDNA

Confirmation by Postmortem Simulation

- Comparing three scheduling strategies
 - (Non clairvoyant) Dynamic and LDNA
 - Clairvoyant LPT
- LDNA almost as good as LPT



Sketch of the final solution

```
1  vector<Mutant> mutant_list; // In global scope
2  #pragma omp parallel
3  {
4  #pragma omp for schedule(static)
5  for (auto i = 0; i<N; ++i) {
6      indiv[i] = prepare_mutation ◦ selection(cell[i])
7      if has_mutate(indiv[i])
8          mutant_list.push_back(i) // Concurrent access to the list
9  }
10 << synchronize_sort(mutant_list) >> // Sorted by new DNA size
11 #pragma omp for schedule(monotonic: dynamic(1))
12 for (auto i: mutant_list)
13     fitness[i] = do_fitness ◦ ... ◦ do_mutation(indiv[i])
14 }
```

Purely based on OpenMP Standard

- LPT thanks to dynamic scheduler
- monotonic modifier since OpenMP 4.5

Remaining issue

- Handle the list of mutants
- Efficient sort

Concurrent List of Mutants

Omp_Reduc

```
1 #pragma omp declare\  
2     reduction(merge: vector<Mutant>:\  
3         sort_merge_lists(omp_out, omp_in))  
4 #pragma omp parallel  
5 {  
6     #pragma omp for reduction(merge: mutant_list)  
7     for (auto i = 0; i < N; i++)  
8     {... mutant_list.push_back(i) ...}  
9     #pragma omp for schedule(monotonic: dynamic, 1)  
10    for (auto i: mutant_list) {...}  
11 }
```

vs

DIY

```
1 #pragma omp parallel  
2 {  
3     #pragma omp for nowait  
4     for (auto i = 0; i < N; i++)  
5     {... local_mutant_list[p_id].push_back(i) ...}  
6     sort(local_mutant_list[p_id],  
7         [](m1, m2){ return size(m1) > size(m2); });  
8     #pragma omp barrier  
9     #pragma omp single  
10    mutant_list = merge_lists(local_mutant_list)  
11  
12    #pragma omp for schedule(monotonic: dynamic, 1)  
13    for (auto i: mutant_list) {...}  
14 }
```

Concurrent List of Mutants

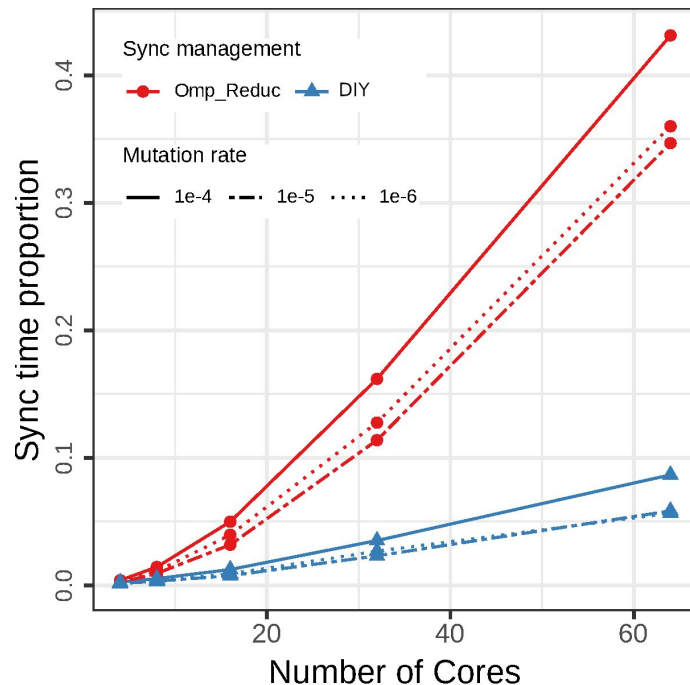
Omp_Reduc

```
1 #pragma omp declare\  
2   reduction(merge: vector<Mutant>:\  
3       sort_merge_lists(omp_out, omp_in))  
4 #pragma omp parallel  
5 {  
6 #pragma omp for reduction(merge: mutant_list)  
7 for (auto i = 0; i < N; i++)  
8 {... mutant_list.push_back(i) ...}  
9 #pragma omp for schedule(monotonic: dynamic, 1)  
10 for (auto i: mutant_list) {...}  
11 }
```

vs


DIY

```
1 #pragma omp parallel  
2 {  
3 #pragma omp for nowait  
4 for (auto i = 0; i < N; i++)  
5 {... local_mutant_list[p_id].push_back(i) ...}  
6 sort(local_mutant_list[p_id],  
7     [](m1, m2){ return size(m1) > size(m2); });  
8 #pragma omp barrier  
9 #pragma omp single  
10 mutant_list = merge_lists(local_mutant_list)  
11  
12 #pragma omp for schedule(monotonic: dynamic, 1)  
13 for (auto i: mutant_list) {...}  
14 }
```




Proportion of time taken for the **synchronization** in one generation

Final Results



LDNA

- Solution purely based on OpenMP Standard
- ~ 20% of gain over Dynamic scheduler
- Ad Hoc solution for Aevol
 - Mix between biological model and parallelization model : no separation of concerns



AEVOL: Original computation pattern

- Highly irregular and varying application with unpredictable behavior over generation
- Fine grain computation
- Manual decomposition in two loops with specific scheduler
 - ~20% of improvement
- Need to analyse the biological model and its implementation to find an efficient solution to schedule the application

More realistic biological simulations need much more computation
Revisit Aevol parallelization to improve its performance

- Target multi-CPU/GPU node

OpenMP conclusion & perspective

AEVOL: Original/Unique computation pattern?

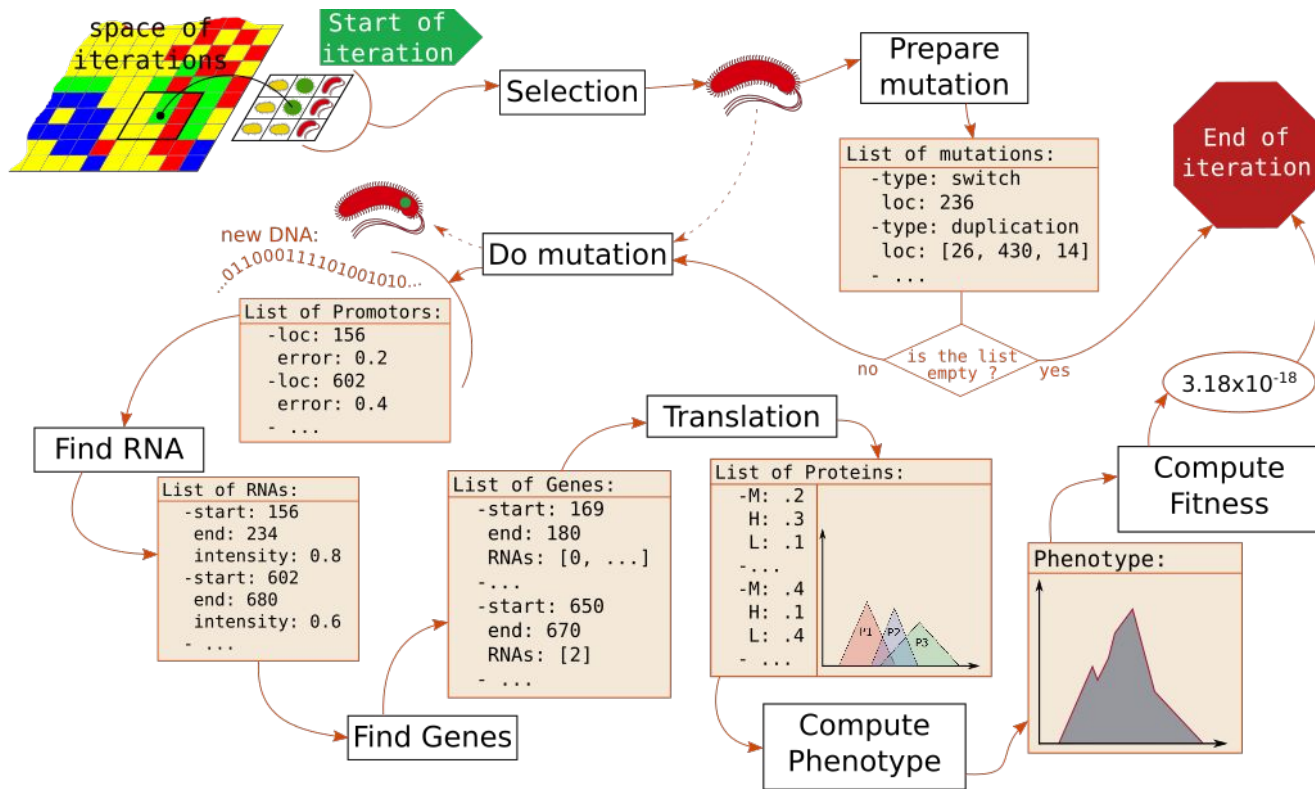
- Solution for CPU based purely on OpenMP standard
 - Through decomposition in 2 OpenMP for loops + specific LPT schedule
- **Code transformation depends on the scheduling solution!**
 - How to implement application dependent loop scheduler with code annotation only?
 - Have more clairvoyant schedulers in the OpenMP runtime
- Is this computational pattern frequent? Can a modification of the OpenMP standard help handle this kind of pattern with less effort?


We have a dream...

```
#pragma omp for schedule(static)
for (auto i = 0; i < N; i++) {
    R =  $f_k \circ \dots \circ f_2 \circ f_1(\text{indiv}[i])$ 
    #pragma omp barrier schedule_modifier(\
        schedule(monotonic: dynamic, 1)\
        sort: R, [](r1, r2){return size(r1.dna) > size(r2.dna);})
    fitness[i] =  $f_n \circ \dots \circ f_{k+2} \circ f_{k+1}(R)$ 
}
```

Thank you !

AEVOL: WORKFLOW OF A GENERATION





dynamic

OpenMP
reduction

Our final
Solution

