

HW 4 - SSA, Optimization & UPC++

Issued: April 06, 2020

Due Date: April 27, 2020, 08:00am (new)

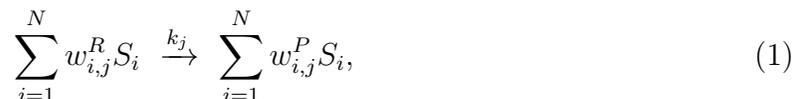
1-Week Milestone: Solve tasks 1 and 2

The skeleton codes for this homework are located in the gitlab repository: <https://gitlab.ethz.ch/hpcse20/exercise>.

Task 1: Stochastic Simulation Algorithm

The Stochastic Simulation Algorithm (SSA, also known as Gillespie algorithm) is a method that uses Monte-Carlo sampling to generate trajectories according to the chemical master equation.

Here, we consider the evolution of the quantity of N species $\{S_i\}_{i=1}^N$ within a well mixed volume which evolve according to M possible reactions:



where $j = 1 \dots M$, k_j are the reaction rates and $w_{i,j}^R, w_{i,j}^P$ are the stoichiometric coefficients for reactants and products.

For large S_i , this system would evolve according to the following ordinary differential equations:

$$\frac{ds_i}{dt} = \sum_{j=1}^M \left(w_{i,j}^P - w_{i,j}^R \right) r_j, \text{ with } r_j = k_j \prod_{i=1}^N s_i^{w_{i,j}^R}, \quad (2)$$

where $i = 1 \dots N$. The concentrations are defined as $s_i = S_i/\Omega$, where Ω is a scaling parameter dictating the number of particles in the system.

SSA solves the system by explicitly modelling the probability of collisions between reactants. It can therefore only be applied for reactions involving at most two molecules. Reactions involving more molecules can be modelled as sequences of binary reactions.

In the following description, we use $i \in [1, N]$ for the i -th species and $j \in [1, M]$ for the j -th reaction. Before computing the trajectory we need to set initial conditions S_i for our N species and set up reaction channels $RC_{i,j} = w_{i,j}^P - w_{i,j}^R$. We then evolve the species in time as follows:

Step 0: Evaluate propensities a_j and total propensity $a_0 = \sum_j a_j$.

Step 1: Draw $\tau \sim \text{Exp}(a_0)$.

Step 2: Draw reaction $r \in [1, M]$ from a discrete random variable X with $\text{Pr}(X = r) = a_r/a_0$.

Step 3: Advance the process by evaluating $t \rightarrow t + \tau$ and $S_i \rightarrow S_i + RC_{i,r}$.

Trajectory binning. A trajectory from SSA shows one possible realization of the time evolution of the system. Trajectory binning is a simple way to generate average trajectories over a number of sample trajectories. Since we do not have the same time steps for all samples, we need to average over time and over the samples. We assume that the simulation is run for times $t \in [0, t_{end}]$. We can then compute averaged quantities $\overline{S_{i_k}}$ for realizations of S_i in the time interval $[k\delta t, (k+1)\delta t]$ with $k = 0 \dots \lceil t_{end}/\delta t \rceil - 1$. We can therefore set $k = \lfloor t/\delta t \rfloor$ and update $\overline{S_{i_k}}$ for each time step of each sample trajectory.

For **Task 1** you have to implement the direct SSA algorithm for the following reaction system:



which for large S_i would evolve according to the following ordinary differential equations:

$$\begin{aligned}
 \frac{ds_1}{dt} &= -k_1s_1 + k_3s_1s_2, \\
 \frac{ds_2}{dt} &= -k_2s_2 - k_3s_1s_2 + k_4.
 \end{aligned} \tag{4}$$

We choose the rate constants $(k_1, k_2, k_3, k_4) = (1, 1, 1/50, 2)$ which leads to the propensities $(a_1, a_2, a_3, a_4) = (k_1S_1, k_2S_2, \frac{k_3}{\Omega}S_1S_2, k_4\Omega)$. The initial concentrations are defined to be $(S_1, S_2) = (4\Omega, 0)$.

For this settings the ODEs will reach a steady state solution at $(s_1, s_2) = (15, 5)$. A second steady solution exists at $(s_1, s_2) = (0, 20)$ and some trajectories in SSA might reach those values.

- a) [5pts] In preparation for the following subquestion, derive the inverse transformation for the sampling
- from an exponential distribution with rate τ , i.e. $p(x) = \tau e^{-\tau x}$,
 - and the sampling from a multinoulli distribution with 4 different outcomes $\Pr(X = r) = a_r/a_0$ for $r = \{1, 2, 3, 4\}$ and $\sum_{r=1}^4 a_r/a_0 = 1.0$.

Document your solution.

- b) [20pts] Implement **Step 0**, **Step 1**, **Step 2** and **Step 3** from the SSA algorithm in the file `SSA_CPU.cpp`. In Step 1 and Step 2, use the inverse transform to draw samples from the corresponding probability density function. Please reuse the allocated random numbers from the `r48` array. Other files don't need to be modified.
In order to validate your implementation visualize your results output `.txt` with the python script `plot_output.py` and compare with Figure 1.
- c) [5pts] (Optional) Estimate the number of FLOPs, the memory accesses (read and writes) and operational intensity of the algorithm. For this, Implement your estimates in the methods `getFlops` and `getTransfers`. Only consider floating point operations and memory accesses that happen in the while loop given by `while (bNotDone) { ... }`. Note that

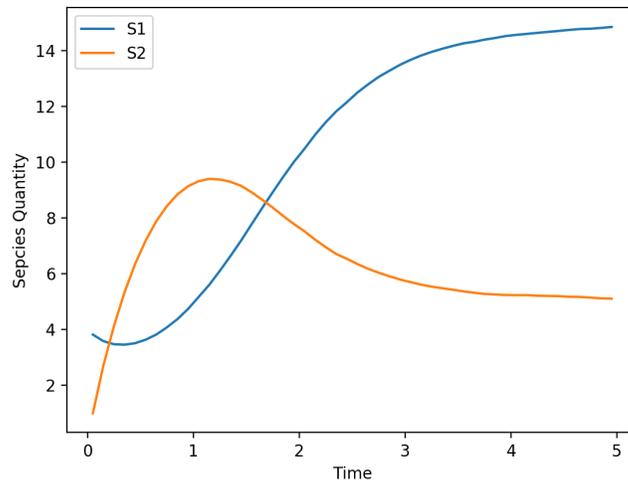


Figure 1: Evolution of Species S_1 and S_2 .

your result depends on the number of simulations stored in variable `numSamples`, number of bins stored in variable `nBins`, and the number of time steps per simulation, which can be read from the array `ntraj`. What performance in *GFLOP/S* do you achieve? Try different values for the environment variable `#OMP_NUM_THREADS`. Document your answer.

Note: This code will be reused in Homework 6 as a baseline for your CUDA implementation, please make sure that you don't fall behind. The maximal number of points in Task 1 is 25.

Task 2: Optimization with CMA-ES

In this exercise we will reuse the code from **Task 1**. We will use the CMA-ES solver from the *korali* framework to search the parameter set $(k_1, k_2, k_3, k_4) = (1, 1, 1/5, 20)$ that yields $(\overline{S}_{1k_{MAX}}, \overline{S}_{2k_{MAX}}) = (15, 5)$.

a) [0pts] Install the *korali* framework. Installation steps are documented here

- <https://www.cse-lab.ethz.ch/korali/docs/using/install.html>.

A documentation overview is available here

- <https://www.cse-lab.ethz.ch/korali/docs/>.

b) [20pts] Implement and maximize the objective function given in `objective.cpp` which returns the negative sum of squared errors (SSE), i.e.

$$f(k_1, k_2, k_3, k_4) = -(\overline{S}_{1k_{MAX}} - 15.0)^2 - (\overline{S}_{2k_{MAX}} - 5)^2. \quad (5)$$

Use the getter-methods `getS1` and `getS2` to read the terminal values $\overline{S}_{ik_{MAX}}$ from the `SSA_CPU` class. You can follow the `SSABenchmark.hpp` file on how to initialize and run the SSA simulations. Other files don't need to be modified. An exemplary *korali* optimization script and objective function can be found here:

- <https://github.com/cselab/korali/blob/master/tutorials/advanced/4.running.cxx/run-cmaes-direct.cpp>,
- <https://github.com/cselab/korali/blob/master/tutorials/advanced/4.running.cxx/model/direct.hpp>.

c) [5pts] Does CMA-ES converge to $(k_1, k_2, k_3, k_4) = (1, 1, 1/5, 20)$? Why (not)? Do we obtain a good fit? Also play with the parameter Population Size in the file `main_cmaes.cpp`. Document your findings. If you like to visualize optimization details, you can run the command `python -m korali.plotter`. This will read and visualize values from the `_korali_results` folder (must be present in the current directory).

Task 3: Parallel Tasking in UPC++

[20pts]

You are given a sequential program that approximates the value of π up to any desired accuracy. The approximation is based on a finite series, which is summing up factors according to the formula:

$$\pi \approx 4 \sum_{n=1}^N F(n). \quad (6)$$

The more factors N are included in the series, the lower the approximation error, i.e.

$$\lim_{N \rightarrow \infty} 4 \sum_{n=1}^N F(n) = \pi. \quad (7)$$

The formula $F(n)$ that specifies how the factors are computed, are given to you as a black-box function, to which you have no access. **You are not allowed to alter the formula, or the way the factors are computed in any way.**

Before continuing make sure that you use the following configuration for UPC:

```
$ export UPCXX_GASNET_CONDUIT=smp
$ export UPCXX_THREADMODE=seq
$ export UPCXX_CODEMODE=O3
```

and you have loaded the necessary compiler:

```
module load new gcc/6.3.0
```

Throughout this exercise, the UPCXX tutorial provided in the exercise class as well as the User guide can be very helpful.

- a) [1pts] The provided sequential program computes an approximation $\tilde{\pi}$ based on the first $N = 240 \times 10^3$ factors of the series, and then writes the result to two text files in the folder `./Results`. **You have to create the folder before running the scripts, else the results will not be saved.** The program also reports the absolute approximation error $\epsilon = |\pi - \tilde{\pi}|$. Run the program and report the absolute error and the speed. Make sure that the textfiles are saved in the results folder after the execution.
- b) [8pts] In order to parallelize the execution of the sequential program, we follow a divide-and-conquer strategy using UPC++, distributing the factors among ranks and having each rank evaluate it's set of factors in parallel. Build your parallel version of the sequential program by filling the TODO: comments on the file `divideAndConquer.cpp` and run it using 24 ranks on an Euler compute node. Request an allocation (in an interactive mode) using:

```
bsub -W 24:00 -n 24 -R fullnode -Is bash
```

For simplicity, assume that the number of ranks divides exactly the number of factors the program computes. Take into account the following steps:

- Initialize UPCXX.
- Assuming that the number of ranks is N_{ranks} , each rank needs to compute N/N_{ranks} factors.

- After computing each factor, each rank needs to save the computed factor on an array that is accessible by all ranks in the global address space. This can be achieved by defining a global pointer. Each rank has its own private pointer to the shared memory. This pointer is named `upcxx::global_ptr<double> factorArray` in the skeleton code provided. The master rank then allocates a new array in the shared memory (global address space). This array is yet invisible to the rest of the ranks. The master rank has to **broadcast** the address of this array in all ranks (see next item), so that all private global pointers point to the same address in the global address space. The main rank creates the global array and broadcasts its address to all other ranks.
- The global factor array needs to be initialized by rank zero (`upcxx::new_array<double>`), and broadcasted to all other ranks.
- After broadcasting the array, each rank needs to compute the portion of the factors it is assigned, and then **place** the result back to the `factorArray`. Do not use RPCs in this question, use the `upcxx::rput` command.
- The ranks need to evaluate their factors and update the global result array in an asynchronous fashion, overlapping communication and computation. If we do not overlap communication and computation (e.g. use the `.wait()` method) then we need to wait on all completions, which can be costly and tricky in case of many asynchronous ranks running in parallel. This can be achieved using a so-called **conjoined** future (composing many futures) in `upcxx`, where we are gathering all futures together and wait for their completion only at the end.
- After the aforementioned step, the time of each rank computation is tracked.
- Finally, rank zero needs to compute the approximate value $\tilde{\pi}$ and save it to the results file, along with the total time. **Downcast** the global pointer to a local one and use it to compute the final approximation. Compute also the absolute error.
- Did you wait for the futures to be completed before rank zero is summing up the series? Do you need wait and synchronize the ranks?
- Finalize UPCXX.

Run the divide and conquer program with $N_{nodes} = 24$ nodes using the command:

```
upcxx-run -n 24 ./divideAndConquer
```

and report the total runtime as well as the absolute error $\epsilon = |\pi - \tilde{p}_i|$. Did you manage to accelerate the sequential program? If yes, what is the speed-up S you achieved? Report the optimal speed-up S^{opt} and the speed-up efficiency S/S^{opt} .

- c) [1pts] Use the provided script `plot.py` to plot the rank runtimes of the parallel program. You can either copy the results folder locally and plot the figures locally, or load the python module in Euler with:

```
module load python
```

then run the script, and copy the figures to your local machine. The python script, receives two inputs, the number of ranks used by the program and the strategy. It outputs a plot with the runtime of each rank. The program, processes the result files from the folder `./Results` and saves the plot figure in the folder `./Figures`. To run the python script use:

```
python3 plot.py --NUM_RANKS 24 --strategy
    divide_and_conquer
```

What do you observe? Is the load balanced? Calculate the load imbalance ratio. (you can use the python script to calculate the imbalance ratio)

d) [8pts] In the following you decide to alter the parallelization strategy and try a producer consumer tasking strategy. Implement your program by filling the `TODO`: comments in the file `producerConsumer.cpp`. The following steps may assist you in your implementation.

- Initialize UPCXX.
- Similarly to the divide and conquer strategy, the master needs to initialize the `factorArray` array with the `upcxx::new_array<double>` command. Then the array needs to be broadcasted so that all ranks have access and point to the shared address space. Note that different to the skeleton code provided in the divide and conquer strategy, the global pointer is already defined in line 24.
- Then, the master rank is calling `master(rankCount)`, while every other rank (consumer rank) `k` is calling `worker(k)`.
- The master is initializing a queue with all the available workers. This queue is filled all the time **only** with the workers that are available to compute (not the busy ones, that are already computing a factor).
- Until all factors are ready, the master is waiting for any worker to be available. Whenever any worker is available, the master has to get his `workerId`, by popping out the first `workerId` from the queue of the available workers.
- After identifying an available worker, the master needs to send a fire and forget RPC `upcxx::rpc_ff`, the task that the consumer has to complete. The task the master is assigning to the next available worker is of course the evaluation of the next factor.
- Check that the factors are evaluated in reverse order. As we observed from the previous task that the most computationally heavy tasks are the ones with high k assigned to the bigger ranks in the consumer-producer strategy, so we prefer to evaluate them first.
- When all factor evaluations are finished, the master needs to use a RPC to notify the consumers that all evaluations are finished by setting the boolean variable `continueFactorEvaluations` to `False`.
- Complete also the `workerComputeFactor` function. This is the function sent from the master with a RPC to the worker with `rankId`. The worker, has to compute the associated factor, then use a RPC to update the master's `factorArray` with the computed factor, and notify the master that the consumer is available again to compute by pushing your `rankId` to the queue.
- Finally, when all factors are evaluated, the master rank needs to compute the approximate value $\tilde{\pi}$ and save it to the results file, along with the total time. **Downcast** the global pointer to a local one and use it to compute the final approximation. (Similarly to the divide and conquer strategy)

Report the total runtime and the associated absolute approximation error. Compute the achieved speed-up and the parallel efficiency. In your computation for parallel efficiency take into account only the ranks devoted for computation and not for communicating the tasks.

e) [2pts] Now, plot the runtimes of the producer-consumer strategy with the python program. What do you observe? Is the load balanced? Calculate the load imbalance ratio in this case. What is your conclusion, is this producer/consumer strategy preferable over the divide and conquer strategy in this specific example ? When is in general the producer consumer preferable over the divide and conquer strategy? Document your answers.

Guidelines for reports submissions:

- Archive your written pdf and source code (e.g.: .tar, .zip) and submit it via Moodle until April 27, 2020, 08:00am