

# Simulation of Large Sets of Metabolic Reactions on Graphics Processing Units (GPUs)

G. R. O. Kobori

**Abstract**—Simulating large sets of metabolic reactions is a very time-consuming process. Here CUDA enabled GPUs are utilized to simulate such metabolic reactions in an attempt to reduce simulation execution times. By using metaprogramming and libSBML biochemical models can be extracted from SBML files, and a CUDA program can be generated to simulate this model in a GPU. Two approaches were tested, one based on the Euler Method and another based on the Gillespie Algorithm, the latter being the one used in the final program. The simulation results and execution times were compared to another simulation software, COPASI, which does not have parallel characteristics. The CUDA application presented mixed results, for a well-behaved model it presented a better execution time, when extrapolating, for a number of simulations of 220 and higher; for a model with highly variable simulation time it did not present a speed up when comparing to a linear approach. This study showcases challenges when trying to implement the Euler Method approach in GPUs and the reasons why the Gillespie Algorithm was chosen over it. The final application represents a first step in simulating biochemical models in GPUs and shows main development points when creating similar applications.

**Index Terms**—biochemical models, CUDA, Euler method, Gillespie algorithm, GPU, libSBML, metabolic reactions, SBML

## I. INTRODUCTION

**S**IMULATING sets of metabolic reactions can give us a better understanding on how complex biochemical processes work. By being able to completely monitor the progression of all the chemical elements involved, it enables the process of making educated guesses on their effects on each other and the analysis on how the entire biological model functions on a deeply fundamental level, the chemical reactions. Furthermore, it enables simulating a large number of different situations, enabling observations on how the behavior of the biochemical model alters with changes in initial parameters, and beyond these initial parameters, it is also possible to simulate external interferences to the model by modifying parameters when certain conditions are met (e.g., simulating a drug administration on a certain disease development stage).

This work attempts to develop an application capable of executing such simulations using a GPU. It tries to achieve feasible results in an attempt to show that the generated program could be used in future investigations, or at least serve as a support for future GPU centered simulation software. It is also a study of the challenges, and possible solutions to these challenges, encountered when developing such programs, and how it can help future developers of similar applications.

The kinetic equations that model a chemical system can describe deterministic or stochastic systems or even both. The deterministic systems are described by differential equations which can be solved by a wide range of methods, for example the Euler Method. On stochastic systems one example is the Gillespie's Direct Method.

At first the Euler Method was considered but, after encountering some complications with properly fitting the method with the GPU, the stochastic method was elected as the next attempt. Given that the Gillespie Algorithm [3] was a better fit for the GPU implementation it was adopted in the final application developed.

### A. Deterministic Simulations

In deterministic simulations there are no random elements during the execution, a given set of inputs always produces the same set of outputs. One advantage over stochastic approaches is a faster execution time, since there is no need to repeat the simulations (assuming that the model is correct) a single run provides the expected results.

### B. Stochastic Simulations

Stochastic simulations are often used to represent physical phenomena, using the randomness present on its core to generate different results for a simulation, and using statistics to obtain the expected behavior using these distinct results. Multiple simulations need to be executed in order to guarantee statistical significance.

The Gillespie Algorithm approaches the simulation at the molecular level, as a container with a number of molecules, where the randomness comes from the propensity of the molecules to collide with each other which may then trigger a

July 2019

G. R. O. Kobori is with the Instituto Superior Técnico, Universidade de Lisboa, Portugal (e-mail: [guilherme.kobori@tecnico.ulisboa.pt](mailto:guilherme.kobori@tecnico.ulisboa.pt))

single chemical reaction. After a simulation step only one chemical reaction occurs, therefore only a few species values are updated.

## II. DETERMINISTIC METHOD APPLICATION

During the first steps of development the objective was to experiment with simulation possibilities and to obtain a better understanding of all the tools needed. As a starting point the intent was to create a serial and a parallel version of the same algorithm and compare how their performance values correlate with different magnitudes of biochemical models.

The Euler Method was chosen due to its widespread usage when dealing with deterministic approaches, as an added benefit, it was also a comfortable method to work with, due to its familiarity.

The steps in which the application progressed will focus on two main versions in order to better organize the process of development.

### A. Serial implementation with libSBML support

The very first version of the application was a straightforward serial implementation of the Euler Method, using libSBML [2]. It contributed in being a firsthand experience on handling SBML [1] files with the aid of the libSBML resources. A great resource provided by libSBML is an evaluation function that could calculate the mathematical expressions making the development uncomplicated.

### B. Parallel implementation

With a working serial version, the code was analyzed in order to identify possible parallelizable regions to be implemented using CUDA [4]. However, executing CUDA code and using libSBML in the same file is not possible, because CUDA programs require specific libraries created exclusively for GPUs, and libSBML does not support CUDA.

The solution found to circumvent this problem was the usage of metaprogramming. This approach uses libSBML first, extracting all the data necessary from the SBML file, and then uses this information to generate another program, a CUDA program, the resulting code does not have any dependencies with libSBML functionalities. Metaprogramming is characterized by having a code generating another code.

The deterministic method was parallelized in this version by distributing each reaction increment values calculation to one thread, attempting to calculate all update values in parallel and therefore obtaining all updated species values in one instruction cycle.

### C. Limitations of the model

When further analyzing the performance of the aforementioned implementation an issue was found due to the ramifications in processing done by that approach. In order to assign each reaction to each thread to execute their designated calculations a switch case statement was used; however, this technique creates the worst possible ramification tree for the GPU. Since each specific reaction calculation is done once and only once, during this step in the execution only one of these calculations are done at any given moment, implying that only

one thread is executing while all the others are idle, waiting for their opportunities.

This happens due to the choice made during development to use only one block of the GPU, which grouped threads into warps and made these warps have a longer execution time due to the limitation of executing the same instruction per warp group per cycle. Having this same instruction per warp limitation with the ramifications in the code instructions caused the execution in the GPU to have only one thread executing its desired instruction while all the others in the warp group were idle waiting for their turns, where they would be running exclusively as well.

This represented a roadblock in development. While researching other approaches for simulating the model a suggestion was made to have a look on stochastic methods and figure out if they were more fit for the results this endeavor aims to obtain.

### D. Lessons learned

Although the program version developed during this pursue of the deterministic model approach is not used in the final application the experience provided by this ordeal was fundamental to reach the final version, which was inspired by lessons learned in this implementation.

The usage of metaprogramming is a characteristic that was kept in the stochastic version, fundamental for making the libSBML and CUDA work together. And the learning experience with the problems of ramification gave a hands-on lesson on how architectures that are suitable for CPU execution sometimes must be avoided when developing for GPU applications.

## III. STOCHASTIC METHOD APPLICATION

Given the independent nature between simulations in a stochastic process the usage of a GPU and its capabilities is quite fitting, the only bottleneck that happens between simulations is the access to the global memory in order to register the results of the simulation, but that represents an unavoidable cost that is an integral part of the process. However, in any other instance of a simulation, only local memory is involved, which is the fastest memory available for the graphics chip.

With that in mind the development of the program was done, the resulting application being not very user friendly, but having in its core the functionalities envisioned, and bearing in mind the possibilities on future stochastic simulation software.

In order to join functionalities of the libSBML [2] and CUDA [4], which do not support each other, a metaprogramming approach was used. A C program (`generateCUDAGillespie.c`) utilizes the libSBML functionalities to extract the information present in a given SBML [1] file in order to generate a CUDA program (`kernelGill.cu`) as the output. The resulting CUDA application then can be run on a CUDA enabled GPU, returning the results of the simulation of the biological model in a CSV file (`results.csv`).

In a regular use, the file `generateCUDAGillespie.c` is compiled, and the resulting executable is run, passing as parameters the name of an SBML file, the sample interval size

in seconds, the final simulation time in seconds, and the number of different simulations. After the program finishes it will have produced several text files, as such, the folder where the executable was run from will now contain 16 code segments.

Afterwards, in this same folder, the script `concatGill.bat` must be run. This script executes a simple operation joining all the segments that were created and generating a new file `kernelGill.cu`, which is the CUDA program that can subsequently be compiled and run on the GPU.

In this application two program files are involved, first, `generateCUDAGillespie.c` is a C program that uses `libSBML` and `metaprogramming` in order to read a given SBML file and generate the second program, `kernelGill.cu`, which is a CUDA program that simulates the SBML model read by the first program in a GPU.

#### A. CUDA program description - `kernelGill.cu`

The execution starts at the CPU, where the arrays `dev_output` and `species_global` are allocated in the global memory, which is shared between the CPU and the GPU. The initial species values (the quantity of the molecules) are copied to the `species_global` array, while the `dev_output` array is initialized with all zeros (this is the array where the simulation results will be registered).

Another array `devStates` is allocated in the global memory, this array is responsible for saving the state of the pseudorandom generator for each thread. The kernel function `initCurand` is called, which simply initializes the random number generator with a different seed for each thread (default seed is 23).

The main simulation is now ready to start, the simulation is partitioned in fixed intervals (arbitrary default interval of 10000s), this has the purpose of giving a feedback of the execution times to the user. In addition, it reduces the kernel run time avoiding problems with hardware drivers. A kernel run is performed for each of these intervals.

Finally, on the GPU a local array `species` is created and initialized with the values saved in the `species_global` array previously mentioned. These values are stored in local memory because they are unique for each thread. Also, they are extensively used which requires a fast access time. Next two local matrices are created `reactionSpecies` and `reactionValues`, they work as a key value pairs, enabling a quick access to the species involved in a reaction and their stoichiometric values (the number of molecules involved in a chemical reaction) and they are essential in order to avoid ramifications in the execution, guaranteeing a linear approach.

After initializing the local variables, the simulation loop begins, it goes on until the segmentation time or the final execution time is reached. In this loop, it is checked if the sampling interval was reached, if so, the species array values are added atomically to the output array in the corresponding position and the next sampling interval is updated, this is done in order to calculate the average values of the species at each sampling time, and avoids data hazards of multiple threads updating the same memory position at the same time. Next, the propensity of each reaction is calculated cumulatively (which are represented by the reaction rates in the SBML file), in other words, the next reaction propensity is their own propensity

formula plus the sum of all the previous calculated propensities values. Then the time step is randomized, it uses an exponential distribution with  $\lambda$  equal to the total sum of propensities. The reaction is chosen by using a uniform distribution with the interval between zero, included, and the total sum of propensities, excluded, and by finding the index where this random number is lower than the cumulative propensity of that reaction index. In order to find such index a binary search is implemented, exploiting the fact that the cumulative propensities are in order. Unfortunately this method adds a corner case where if a reaction propensity is zero it has the same cumulative value as the previous one and therefore it can be selected by the binary search; this issue is resolved with an extra verification where it is required that the random number generated is also higher or equal to the accumulated propensity of the previous index. However, this generates another corner case where the reaction with index zero can never be validated, this is resolved by adding a verification step that checks if the selected reaction is the first one and if not then the binary search is made. This generates some level of ramification in the execution, but represents an unavoidable part of the algorithm. A search must be done in order to obtain the selected reaction, and the binary search is a great fit in an attempt to reduce ramification issues, in comparison to simpler, more naive search methods.

With the time step and the reaction index values at hand, by using the reaction index the matrices `reactionsSpecies` and `reactionsValues` are used to update the quantities of the species and then the simulation time is updated with the time step.

In the case where a reaction is selected and the number of reactants consumed is higher than one, it could be possible to cause a negative specie quantity. To avoid that the application verifies the resulting species quantities after the update and reverts their values if a negative quantity is found. The simulation then goes on as if, on that interaction, no reaction occurred.

At this stage it is also verified if any event was activated. Events represent external interferences to the system, and are declared in the SBML file with a condition and an update formula, when the condition is met the update formula is executed. For example, in one of the models presented on the simulation results section, 50 molecules of the species `antiAb` are added when the simulation reaches the fourth day, simulating a drug administration.

At the end of one simulation loop (kernel run) the species array values are saved in the global array `species_global`, in order to be able to restart the simulation in the same condition that it was in the case of this simulation being one of the intermediate partitions. Each thread has enough space in the global memory in order to register all its species values, since all these values are unique for each thread.

After executing all the segments, the program continues on the CPU. The data registered in the global array `dev_output` is saved in the hosts local memory, these values are then divided by the number of threads that were run, in order to obtain the average values of each species in each time sample. Finally, they are exported to an external file using the CSV format.

The application was developed using the SBML file provided together with the paper "Investigating Interventions in Alzheimer's Disease with Computer Simulation Models" [6].

Some of the constants used are optimal for this model and need to be changed to better fit other biological models. Namely, these constants are the sampling time interval, the final simulation time and the segmentation time size. Later other models [7] were used to validate the accuracy and precision of the application, and these constants had to be adjusted for each.

#### B. C program description - generateCUDA Gillespie.c

This is the program responsible for the metaprogramming segment of the application. It generates 16 text files (listed on table 1) with code segments that are fused together by a batch script, resulting in the CUDA file kernelGill.cu.

The application receives as execution parameters the SBML filename, the sampling time interval in seconds, the final simulation time in seconds and the number of simulations to be run in parallel. These parameters must be adjusted accordingly to each model characteristics. Another parameter that may be adjusted is the constant `SEGMENT_SIZE` that defines the segment size for the simulation partitioning, but this characteristic is only needed for the user feedback behavior on the application and is not crucial for the simulation to run properly.

At the start of execution, the SBML file is read and a model is created by libSBML that enables manipulating all the information present in the SBML file in an object-oriented fashion (in C this is done in the coding level by the use of Structs).

With the model in a format where the data can be easily extracted the code generating part begins. While iterating on some information in the model it is common for the application to edit multiple of the output text files at once in order to utilize the iterations as much as possible. This does not mean, however, that the code is completely optimized, since the main interest is the CUDA program generated by this one it was not a high priority to fully optimize this code, and since it is a quite simple data extraction application, its execution times are quick and not an issue.

Main metaprogramming advantages explored include the usage of the species IDs as variables in the CUDA code, making a quite legible product that if needed can be modified directly since navigating is done easily. Another advantage is the direct usage of the mathematical formulas obtained by using libSBML in the CUDA program, this avoids the need of an infix notation evaluator in order to calculate these formulas.

Some functions that exist in SBML models that were stumbled upon and were not implemented are:

1) *non-constant parameters, where the reaction parameters constants that are used can be changed;*

2) *function definitions, where mathematical functions that can be called with some parameters to define a mathematical formula can be defined;*

3) *initial assignments, where initial amount values can be overridden, and species values can be changed by assigning them a formula to be calculated at instant zero;*

4) *verification of measurement units, the application assumes that all values have coherent units, so that no unit conversions is required.*

## IV. IMPLEMENTATION DECISIONS

In this section decision choices are discussed. Some aspects of the stochastic simulation needed clever solutions in order to try to optimize the application.

### A. Metaprogramming usage

Using the resources provided by the libSBML library were of great help manipulating SBML files but integrating this library directly with the CUDA programming is a difficult task, CUDA requires specific libraries implemented specifically for GPUs, which is not the case for libSBML. By using metaprogramming these issues are avoided, and the resulting program is independent from any libSBML resource, being a pure CUDA program that can be modified on itself. Metaprogramming also benefited the processing of the formulas present in the files, removing the necessity of implementing an evaluator that interprets formulas given in strings.

### B. Segmentation of the simulation

The CUDA program is executed in partitions (default 10 000s for the base model). At first this was done as an attempt to avoid a problem that caused a timeout in the communications between the device and the host. As this problem was resolved the functionality was kept as a convenience at first, but which later has been proven to be useful on debugging steps, and providing a more user friendly interface, sacrificing some small execution time in order to give some feedback to the user about the progress of the simulation.

If, for some reason, the user wishes to run the entire simulation in only one GPU call this can be easily done by setting the constant `SEGMENT_SIZE` time to a higher value than the final simulation time, effectively creating only one partition to be run on the GPU.

### C. Usage of the reaction update matrices

In order to improve the linearization of the code in the GPU an implementation was developed using an approach with matrices, available in local memory for quick access, these matrices `reactionsSpecies` and `reactionsValues` act as a key value pair in such a way that given the index of a reaction, accessing the data containing which species are involved and how their quantities change can be done directly, avoiding the need of any kind of search algorithm to find the information.

This implementation has the benefit of being a more linearized approach and therefore a quicker approach for the GPU. However, it consumes more memory in order to store the update matrices, and most of this memory is redundant, since

each thread has their own local copy of the matrices in order to improve memory access times (since these matrices are accessed every iteration of the simulation). One approach that was not tested was using the constant memory, by declaring a constant variable in the GPU memory, this could save local memory and lead to improvements.

If, by some reason, the local memory cannot be used the reactionsSpecies and the reactionsValues matrices could be stored in shared or global memory as well, but the program will have a higher runtime due to the slower memory access.

Other option, that does not require extra memory, is using different codes for each chosen reaction. However, this would result in code ramification as in the Euler Method implementation and in a bad performance in the GPU.

#### D. Binary search

In the step of finding which reaction needs to be chosen for one iteration of the model simulation it is required to make a search in the values of the propensity of each reaction, finding in which index the value created by the random number generator falls. This propensity data is organized in a cumulative way, meaning that the value of the next propensity is its own value plus the sum of all the previous propensities that were calculated, creating a naturally organized data array, which is made in such a fashion with the purpose of facilitating finding the reaction.

As a naive implementation one could simply sweep the array from index zero until it finds the threshold value where the value at a given index is higher than the value searched, obtaining such value as the desired result. However, in a worst-case scenario, for an array of size N at max N calculations will be made, and for large arrays this approach is quite time consuming and it does not fully utilize the organization structure of the data array.

### V. VALIDATING THE APPLICATION

In order to validate the data outputted by the program, SBML [1] model files were used to compare the results between the COPASI [5] output and the developed application output.

The models used are:

1) *Proctor2013 - Effect of A $\beta$  immunization in Alzheimer's disease [6];*

2) *Proctor2005 - Actions of chaperones and their role in ageing [7].*

The COPASI software allows to do another simulation of the SBML model and compare and validate the results.

#### A. Comparing the results

Here results from different available SBML models are compared between the output produced by the COPASI software and the output produced by the developed application. The analysis is done by visual inspection to verify the main behavior of the system and also by checking the maximum error percentages for each species during the simulation.

The error average was calculated using the absolute difference between the COPASI and the CUDA [4] application

results for each time in relation to the maximum species value reached by the COPASI simulation.

For each species and each time sample:

$$\text{AverageError} = \frac{\frac{1}{N} \sum_{i=1}^N |\text{COPASI}(i) - \text{APPLICATION}(i)|}{\max_i \text{COPASI}(i)} \quad (1)$$

With  $i$  representing each simulation executed and N representing the number of simulations executed.

#### B. Proctor2013 - Effect of A $\beta$ immunization in Alzheimer's disease [6]

The article "Investigating Interventions in Alzheimer's Disease with Computer Simulation Models" [6] simulates some metabolic pathways related to Alzheimer's Disease using the Gillespie Algorithm, and has an SBML file attached that enables the reproduction of the results obtained in this study, making it a great fit for testing and validating the application. Its BioModels ID is MODEL1704060000.

Some highly sensitive species show a high percentage error average, this occurs due to the fact that these species are present in very low quantities, with some being unitary at sometimes. This model also has a high variance between simulations, for example, the species damDNA represents damaged DNA, and changes in its quantity imply drastic changes in the simulation behavior.

Due to this high variance in the species values a better comparison can be done visually and, when inspecting the graphical representation, it can be clearly seen that the main behavior of the model is present. The COPASI graph has less noise than the CUDA program graph, both were generated using 128 simulations.

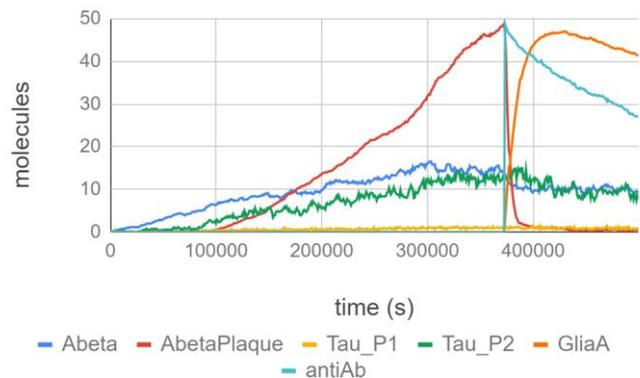


Fig. 1. Plot of quantity of species Abeta, AbetaPlaques, Tau\_P1, Tau\_P2, GliaA, and AntiAb versus time obtained using the developed application, executing 128 simulations in the GPU.

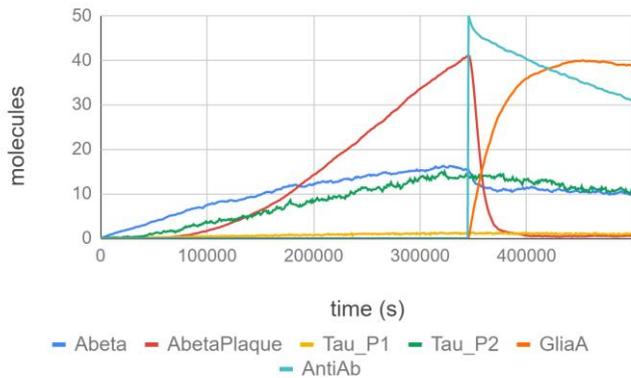


Fig. 2. Plot of quantity of species Abeta, AbetaPlaques, Tau\_P1, Tau\_P2, GliA, and AntiAb versus time obtained via simulation in the COPASI software, using the averages of 128 simulations.

TABLE I

AVERAGE PERCENTAGE DIFFERENCE FOR SPECIES WHEN COMPARING COPASI AVERAGE VALUES WITH THE DEVELOPED APPLICATION AVERAGE VALUES IN THE "EFFECT OF  $A\beta$  IMMUNIZATION IN ALZHEIMER'S DISEASE" MODEL

Species	Average Difference
Mdm2	8.1%
p53	7.4%
Mdm2_p53	1.3%
Mdm2_mRNA	7.2%
p53_mRNA	2.4%
ATMA	7.5%
ATMI	0.48%
p53_P	7.6%
Mdm2_P	8.6%
ROS	9.6%
E1	0.98%
E2	1.1%
E1_Ub	0.98%
E2_Ub	1.2%
GliaI	0.055%
GliaM1	5.9%
GliaM2	31%
Mdm2_p53_Ub	22%
Mdm2_p53_Ub2	21%
Mdm2_p53_Ub3	19%
Mdm2_p53_Ub4	27%
Mdm2_P1_p53_Ub4	3.2%
Mdm2_Ub	23%
Mdm2_Ub2	29%
Mdm2_Ub3	26%
Mdm2_Ub4	7.2%
Mdm2_P_Ub	11%
Mdm2_P_Ub2	10%
Mdm2_P_Ub3	15%
Mdm2_P_Ub4	8.6%
Mdm2_Ub4_Proteasome	11%
Mdm2_P_Ub4_Proteasome	9.3%
GSK3b	0.25%
GSK3b_p53	7.5%
GSK3b_p53_P	9.1%
AggAbeta_Proteasome	9.4%
Abeta	6.3%

AbetaPlaques	12%
Tau	17%
Tau_P1	10%
Tau_P2	9.9%
MT_Tau	0.26%
Proteasome_Tau	9.8%
Proteasome	0.18%
PPI	0%
NFT	4.0%
AbetaDimer	14%
disaggPlaques	13%
Ub	0.49%
damDNA	7.2%
p53_Ub4_Proteasome	3.6%

### C. Proctor2005 - Actions of chaperones and their role in ageing [7]

The article "Modelling the actions of chaperones and their role in ageing" [7] also has an attached SBML file. The main objective was to have a different model to validate the CUDA program, observing how it behaves with different models of the same type and showing that it is an adequate application for this type of model.

The average errors values are all good, with the exception of the MisP species, when analyzing this particular species, it is noticed that its values are quite close to zero, ranging from 0 to 0.5 at maximum. Given that the Gillespie Algorithm works with full molecule count it is expected to have a highly unstable average value when the averages are too close to zero and some misrepresentation of the model might occur.

The graphs inspection, on the other hand, shows that, on the interesting species where we can observe the main behaviors of the model, the data is very coherent. The other species not shown in the graphs have quite constant values, and those match between COPASI and the CUDA program as well.

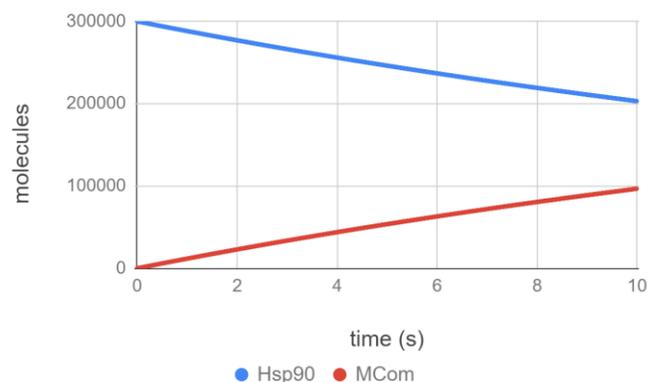


Fig. 3. Plot of quantity of species Hsp90 and MCom versus time obtained using the developed application, executing 128 simulations in the GPU.

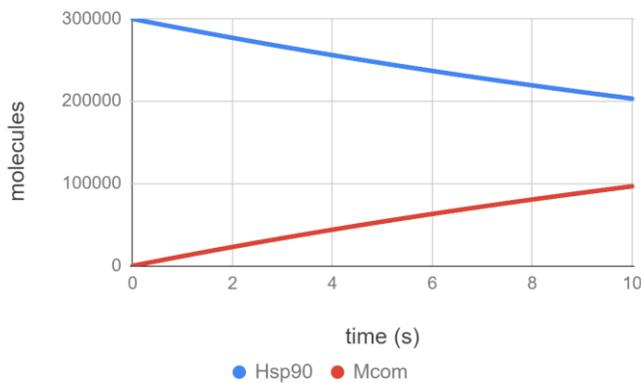


Fig. 4. Plot of quantity of species Hsp90 and MCom versus time obtained via simulation in the COPASI software, using the averages of 128 simulations.

TABLE II  
AVERAGE PERCENTAGE DIFFERENCE FOR SPECIES WHEN COMPARING COPASI AVERAGE VALUES WITH THE DEVELOPED APPLICATION AVERAGE VALUES IN THE “ACTIONS OF CHAPERONES AND THEIR ROLE IN AGEING” MODEL

Species	Average difference
Hsp90	0.040%
HCom	0.0060%
HSF1	0.36%
MisP	3300%
MCom	0.13%
NatP	0.0020%
X	2.2%
ROS	0.044%
ATP	0.022%
ADP	0.24%

## VI. PERFORMANCE ANALYSIS

In order to evaluate the performance of the developed application we are going to compare its execution times with an application that runs only on a CPU. This will be done using COPASI [5] and its feature that enables multiple sequential executions of the simulation.

All simulations were run on the same computer, a Windows 8.1 64-bit, with an Intel Core i3-2100 CPU, a clock rate of 3.10 GHz and 4.00 GB of RAM. The GPU is a GeForce GTX 970.

### A. Effect of $A\beta$ immunization in Alzheimer’s disease [6]

#### 1) CPU only run times

These times were obtained by doing a manual measurement. Given the human error involved and the imperfections of this approach these times should be taken with care and are only used as rough estimates to compare performances.

All simulations were done using the same base SBML [1] model and the Gillespie direct method for exact stochastic kinetics [3], the number of simulations used were 16, 32, 64 and 128 simulations, with an interval size of 900s and a duration of 100 000s. Each time was measured 10 times and the averages obtained are shown in table IV.

Quantity of simulations	Average execution time (s)
16	23
32	46
64	104
128	253

As one would expect the execution times obtained roughly follow the increase in the quantity of simulations. Doubling the number of executions doubles the execution time.

#### 2) GPU simulation run times

These times were obtained using the timing capabilities that are already part of the final application.

All simulations were executed in the same conditions as the ones used in the CPU run time measurement. The application was run for 16, 32, 64 and 128 simulations, with an interval size of 900s and a duration of 100 000s. Since the pseudo random number generator uses a seed value to initialize, this will be changed between measurements. Each time was measured 3 times and the averages are shown in table V.

Quantity of simulations	Average execution time (s)
16	1462
32	1589
64	2707
128	3368

These times are not what one would expect for a parallel program, similar values were expected for these execution times. This behavior happens due to the model characteristics, it is highly volatile, meaning its executions varies greatly. The model has a damDNA species that increases the number of molecules in the simulation proportionally to its value, making the simulation take longer to execute by decreasing the time step values. Due to this unstable behavior the time each simulation takes varies considerably; in a warp we can observe threads that have already finished simulating but that are waiting for another thread that is taking longer.

In the case of the GPU simulation, the running time is at best the worst-case time of all the simulations simulation times, while in the CPU is the sum of all simulation times. For this model the worst-case time gets worse as the number of simulations increases.

#### 3) Comparative analysis

By extrapolating the execution times obtained with a linear trend line we obtain Fig. 5, it shows that the CUDA [4] program always has a higher execution time than the COPASI. After a further analysis of the model it was found that this model in particular has an uneven execution, making a couple threads much slower than the others and compromising the entire parallel execution, since all other threads must wait for the slower ones.

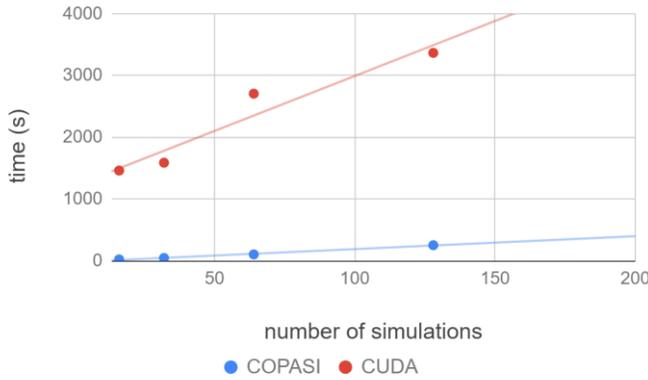


Fig. 5. Trendline comparing the COPASI and the CUDA program performances

### B. Actions of chaperones and their role in ageing [7]

This model was chosen because it is more stable than the previous, it does not have behavior changing random elements.

#### 1) CPU only run times

An interval size of 0.1s and a duration of 100s was used. The times were measured for 16, 32, 64 and 128 simulations.

Quantity of simulations	Average execution time (s)
16	43
32	85
64	175
128	344

#### 2) GPU simulation run times

The same parameters were used in the GPU simulation, using a single block.

Quantity of simulations	Average execution time (s)
16	501
32	512
64	541
128	547

#### 3) Comparative analysis

Once more the data was extrapolated, resulting in Fig. 6, but on this model the behavior of the simulations is more uniform, since it does not have an unexpected change between simulations behavior the extrapolation shows promising results, having a breakeven point around 220 simulations where after this point the CUDA application becomes faster than the COPASI.

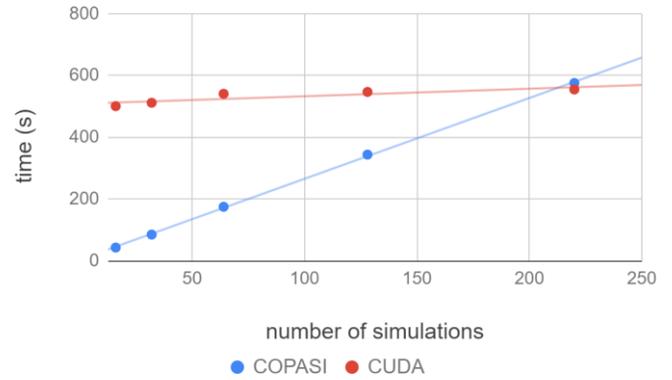


Fig. 6. Trendline comparing the COPASI and the CUDA program performances

## VII. CONCLUSION

The application developed with this document represents a first step in simulating large sets of metabolic reactions in a GPU. It successfully simulates SBML [1] models that are within its developed scope but has some shortcomings when dealing with all kinds of SBML files.

The results obtained show that the CUDA [4] application deals well with well-behaved models, where the randomness of the stochastic simulations does not change drastically the course of the model. On the same note, the results present an issue with the parallelization method implemented, it does not thrive with models with drastic, randomness related, changes, in which case the COPASI serial application showed better execution times.

The final version of this program has some points that could be improved in future versions, to create a more universal and robust version.

### A. Improvement points

Some shortcomings of the application as a whole, that if implemented would increase the quantity of models supported and improve the GPU usage.

#### 1) SBML related improvements

- a) support for non-constant parameters, the reaction parameters constants that are used can be changed;
- b) support for function definitions, mathematical functions that can be called with some parameters to define a mathematical formula can be defined;
- c) support for initial assignments, initial amount values can be changed by assigning a formula to be calculated at instant zero;
- d) support for verification of measurement units, the application assumes that all values have coherent units.

## 2) CUDA related improvements

- a) *support for usage of more than one block and for more than one grid in the GPU, increasing the number of simulations that can be run at once;*
- b) *support for scientific format values, allowing very small constants to be declared. At the moment only constants until 10<sup>-10</sup> are supported;*
- c) *support for a better distribution of threads between blocks, making use of the maximum capacity of the GPU.*

- [6] Proctor CJ, Boche D, Gray DA, Nicoll JA. Investigating interventions in Alzheimer's disease with computer simulation models. *PLoS One*. 2013;8(9):e73631, 16 Sep. 2013. doi:10.1371/journal.pone.0073631
- [7] Proctor CJ, Soti C, Boys RJ, Gillespie CS, Shanley DP, Wilkinson DJ, Kirkwood TB. Modelling the actions of chaperones and their role in ageing. *Mech Ageing Dev*. 2005 Jan;126(1):119-31. PubMed PMID: 15610770.
- [8] Hatakeyama M. Autocatalytic Reaction System by Gillespie's Algorithm (Direct Method) <http://masa.o.o7.jp/dm/>

## ACKNOWLEDGMENT

I would like to thank Professor Paulo Alexandre Crisóstomo Lopes for all the help in the process of developing this dissertation. For giving me the opportunity to work on this subject as well as for providing the computer equipped with the GPU needed for this work. Also, for the weekly guidance and help with better understanding concepts and guidance on the steps necessary for the project progression.

My thanks to Instituto Superior Técnico, for the opportunity given to me to obtain a master's degree on their exchange students programme, and for all the support provided in these two years of studying in this institution.

I would also like to express my gratitude towards Escola Politécnica da Universidade de São Paulo, also involved in this exchange opportunity, for providing me with the bases of knowledge that were of great help on this journey.

As important as the other institutions, I would like to thank INESC-ID and all its staff for providing a physical space where I could work on this dissertation on remarkably great conditions.

At last I would like to thank my pets, my friends, my family, for all the support throughout the years.

## REFERENCES

- [1] M. Hucka, et al. The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models, *Bioinformatics*, Volume 19, Issue 4, 1 March 2003, Pages 524–531, <https://doi.org/10.1093/bioinformatics/btg015>
- [2] Bornstein BJ, Keating SM, Jouraku A, Hucka M. LibSBML: an API library for SBML. *Bioinformatics*. 2008;24(6):880–881. doi:10.1093/bioinformatics/btn051
- [3] Gillespie, Daniel T. "Exact Stochastic Simulation of Coupled Chemical Reactions." *The Journal of Physical Chemistry* 1977 81 (25), 2340-2361 DOI: 10.1021/j100540a008
- [4] NVIDIA CUDA C Programming Guide [https://docs.nvidia.com/cuda/pdf/CUDA\\_C\\_Programming\\_Guide.pdf](https://docs.nvidia.com/cuda/pdf/CUDA_C_Programming_Guide.pdf)
- [5] Stefan Hoops, Sven Sahle, Ralph Gauges, Christine Lee, Jürgen Pahle, Natalia Simus, Mudita Singhal, Liang Xu, Pedro Mendes, Ursula Kummer, COPASI—a COMplex PATHway SIMulator, *Bioinformatics*, Volume 22, Issue 24, 15 December 2006, Pages 3067–3074, <https://doi.org/10.1093/bioinformatics/btl485>