Application of Recent CFD Advancements to the Modeling of Chemical Looping Systems

James PARKER*# and Andrew LARSON

CPFD, LLC, 10899 Montgomery Blvd NE, Albuquerque, NM

**Corresponding Author, james.parker@cpfd-software.com, #Presenting Author*

Abstract – Computational fluid dynamics (CFD) has become a common engineering tool for the research, design, and optimization of chemical looping processes. Like many gas-solid and liquid-solid chemical processes, chemical looping presents a significant modeling challenge for traditional CFD due to the industrial scale of the process units and the complex coupled particle-fluid dynamics, reaction chemistry, and heat transfer. With recent advancements in CFD modeling tools, chemical looping simulations can be more accurate and efficient than what was possible just a few years ago, in large part due to increases in model calculation rates afforded by improved computing hardware and improved parallelization of CFD software. In this work, the impact of these improvements is assessed using a chemical looping reactor model as a benchmark.

1 Introduction

Computational fluid dynamics (CFD) has become a common engineering tool for the research, design, and optimization of chemical looping processes. Like many gas-solid and liquid-solid chemical processes, chemical looping presents a significant modeling challenge for traditional CFD due to the industrial scale of the process units and the complex coupled particle-fluid dynamics, reaction chemistry, and heat transfer calculations within the unit. The size and complexity of a model have a corresponding computational time and cost which often determines the feasibility of using CFD as an engineering tool for an application or forces the size and accuracy of a model to be reduced to fit into an engineering timeframe.

Recent advancements in CFD are demonstrated in this work using a model of the Chemical Looping Reactor (CLR) located at the National Energy Technology Laboratory (NETL) in Morgantown, West Virginia. This model was originally developed by CPFD Software and NETL in 2012 [1] and has subsequently been used to demonstrate later CFD improvements [2]. Because of the complexity of the CLR model – it is a three-dimensional model of a full loop where particle-fluid hydrodynamics, heat transfer, and reaction chemistry are simulated – and the previous history of contemporaneous reporting of model setup and calculation rates, the CLR model is a good benchmark model for demonstrating current CFD capabilities and studying the recent history of model calculation rates.

2 Chemical Looping Reactor Model

The CLR model simulates the circulation of ilmenite particles with a particle size distribution between 13 and 322 microns which transport oxygen through the reaction of Iron (II) Oxide (FeO) with oxygen to form Iron (III) Oxide (Fe_2O_3). The particles are assumed to have an oxygen transport capacity of 3.3 wt% of the oxidized particle, based on the measurements of Abad et al. [3]. To achieve this capacity, the fully-oxidized particles are assumed to consist of 33 wt% Fe_2O_3 and 67 wt% inert whereas the fully-depleted particles have a composition of 30.7 wt% FeO and 69.3 wt% inert. Oxidation and reduction by CH_4, H_2, and CO reactions are implemented using the kinetics of Abad et al.

Coal particles with a diameter between 50 and 150 microns are used as fuel in the CLR model. The coal is initially composed of 51 wt% carbon char, 34 wt% volatile organics, 10 wt% ash, and 5% moisture with an initial particle density of 1333 kg/m^3 and a higher heating value of 32 MJ/kg on a dry ash-free basis. The volatile organics are assumed to devolatilize into 66.8% methane, 24.4% carbon monoxide, 5.9% water, and 2.8% carbon dioxide on a weight basis and the moisture and is allowed to evaporate. The water-gas shift, CO_2 gasification and steam gasification reactions are modeled, as detailed by Parker [3].

The CLR model simulates particle transport through the full loop of the approximately 12-foot-tall chemical looping reactor at a target circulation rate of 363 kg/hr. This system consists of an air reactor and riser, cyclone, loop seal, fuel reactor, and L-valve, as shown in Figure 1. In the air reactor, the ilmenite particles are oxidized as they are transported to the top of the riser by air. The oxidized particles then move to the cyclone, where the particles are separated from the air and pass to the loop seal. The loop seal is fluidized with nitrogen which flushes any remaining oxygen gas from around the particles. After leaving the loop seal, the oxidized particles move to the fuel reactor, where the fuel gases reduce the oxidized ilmenite back to the original state. After leaving the fuel reactor, the particles pass through an L-valve, which controls the circulation rate, and return back to the air reactor. The flow rates, compositions, and temperatures for all gas feeds are listed in Table 1. All gas outlets are assumed to be at atmospheric pressure.

The model is initialized with approximately 10 kg of reduced ilmenite in the air reactor, 29 kg of oxidized ilmenite in the fuel reactor, 12 kg of oxidized ilmenite in the loop seal, and 6.5 kg of reduced ilmenite in the L-valve. The domain is created from 168,000 cells and contains 713,000 computational particles, representing approximately 1e10 real particles with a mass of 57.5 kg.

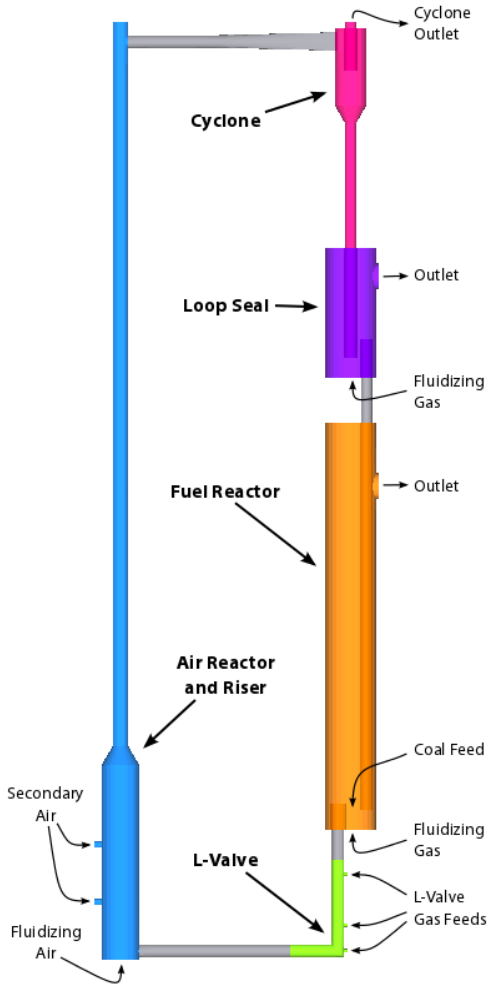

Figure 1: Diagram of CLR model.

Table 1: CLR model boundary conditions

Unit	Boundary	Gas	Temp. (K)	Rate (lb/hr)
Air Reactor	Fluidizing air	Air	1273	40.0
	Secondary air (upper)	Air	773	39.1
	Secondary air (lower)	Air	773	39.1
Fuel Reactor	Fluidizing gas	N2	1273	16.4
	Coal feed gas	N2	298	0.029
Loop Seal	Fluidizing gas	N2	1273	11.9
L-valve	Stripper (upper)	N2	298	0.50
	Aeration (middle)	N2	298	0.14
	Eductor (lower)	N2	298	0.10

3 Modeling Approach

Barracuda Virtual Reactor® software was used to model the original 2012 CLR model and in this work as well. Barracuda Virtual Reactor was developed for the modeling of industrial-scale particle-fluid systems using the Multiphase Particle-In-Cell (MP-PIC) approach [5-8], an Eulerian-Lagrangian method, to solve the fluid and particle momentum equations in three dimensions with bidirectional coupling between the fluid phase (Eulerian) and solid particles (Lagrangian). In the MP-PIC method, a computational particle is defined as a Lagrangian entity in which particles with the same properties such as composition, size, density, and temperature are grouped, allowing industrial-scale systems containing massive numbers of particles to be analyzed using millions of computational particles without losing the advantages of discretizing the solid phase in a Lagrangian frame of reference. The particle-fluid heat transfer and reaction chemistry are coupled with the particle-fluid hydrodynamics to model the interdependencies between the hydrodynamics, temperature, and composition within the CLR unit.

4 Parallelization Improvements

Starting with the release of Barracuda Virtual Reactor in 2013, the software has been parallelized onto the GPU in stages, starting with the most computationally-intensive functions of the code. In version 16.0, the first parallelized release of Barracuda Virtual Reactor, all major particle functions and the linear pressure solver were GPU-accelerated which resulted in speed increases of 2.5 to over 6 times, depending on the size and complexity of the model [9]. For the CLR model, the use of version 16.0 on GPU increased the calculation rate of the CLR model from 0.6 sec/day to 3.1 sec/day, a 5.2X increase [2]. In the next release, version 17.0, major improvements were made to the GPU memory management which resulted in another 50% increase in speed over version 16.0 [10], depending on model size and complexity. In the third GPU parallelization release, version 17.1, the non-linear pressure, velocity, and energy solvers were parallelized which resulted in up to another 60% increase in calculation speed.

In version 17.1, the Eulerian chemistry calculation was the only major un-parallelized functionality. Eulerian chemistry is solved in Barracuda Virtual Reactor with a general-purpose ODE solver which integrates the changes in chemical composition from reactions over a single timestep. Parallelization of the solvers is challenging due in part to the large differences in solution sub-timesteps from cell to cell which disrupts the thread synchronization necessary for efficient parallelism. Two options for chemistry acceleration have since been developed: a multicore CPU parallelization (OpenMP) that was released in version 17.2; and a full GPU acceleration of chemistry that is still under development by CPFD Software. The most recent release of Barracuda Virtual Reactor, 17.3.1 (released July 2018) contains all cumulative GPU-parallelization improvements as well as the OpenMP option for accelerating Eulerian chemistry. Both version 17.3.1 and an unreleased development version of the solver containing GPU-acceleration of the chemistry, (version “Future”), were used to simulate the CLR model in this work.

4.1 Computing Hardware

The computing hardware used to simulate the model, which has seen significant increases in speed, capacity, and flexibility in recent years, is a determining factor of the calculation rate of the simulation. The original CLR simulation was executed on a single core of a desktop computer containing a 3.33 GHz Intel i7 [4]. The version 16.0 simulation, was executed on a desktop computer with a slightly faster CPU (i7-3770K @ 3.5 GHz) and a Nvidia GeForce GTX Titan, a Kepler-architecture GPU card. For this work, simulations were also executed on the Google Cloud Platform (http://cloud.google.com) on three separate virtual machine instances: an instance with a single Nvidia K80; an instance with a single Nvidia P100; and an instance with a single Nvidia V100. All instances had 8 Intel Xeon CPU cores available, operating at 2.2 GHz. The computing hardware is summarized in Table 2. Additionally, the cloud computing costs are listed for uninterruptible GPU instances on the Google Cloud Platform without any continuous use discounts, as of July 2018 [11].

Table 2 - Computing Hardware used for CLR simulations

GPU card	GPU architecture	GPU cores	GPU memory (MB)	CPU type	CPU clock rate (GHz)	RAM (GB)	Form	Compute Cost ($ / day)
None	-	-	-	i7	3.33	NA	Desktop	N/A
GTX Titan	Kepler	2,688	6,143	i7-3770K	3.50	16	Desktop	N/A
K80	Kepler	2,496	11,441	Xeon	2.20	31	Cloud	$19.92
P100	Pascal	10,752	16,281	Xeon	2.20	31	Cloud	$44.16
V100	Volta	15,360	16,160	Xeon	2.20	31	Cloud	$68.64

4.2 Calculation Rates of CLR model

The calculation rate was measured in this work by simulating the execution of the model for 3 seconds at an average time step of 0.1 milliseconds which resulted in approximately 28,000 timesteps. Sampling started at 10 seconds of operating time, at which point the particle and gas distributions have been established throughout the loop and calculation rates were stable. The measured calculation rate for each hardware and chemistry acceleration configuration are presented in Table 3. The calculation cost, also listed in Table 3, is the cost per simulation second of executing the simulation in the cloud. This value is obtained by dividing the compute cost for each instance by the calculation rate.

Table 3: Calculation rates of CLR model

Solver	Eulerian Chemistry Acceleration	GPU card	Calculation Rate (sim. sec/day)	Calculation cost ($/sim. sec)	Source
15.0	None	None	0.6	NA	[4]
16.0.5	None	GTX Titan	3.1	NA	[2]
17.3.1	None	GTX Titan	6.6	NA	This work
		K80	5.6	$3.55	
		P100	8.8	$5.02	
		V100	9.0	$7.63	
17.3.1	OpenMP on 4 cores	GTX Titan	8.6	NA	This work
		K80	6.8	$2.93	
		P100	12.2	$3.62	
		V100	13.2	$5.20	
Future	GPU acceleration	GTX Titan	9.4	NA	This work
		K80	7.84	$2.54	
		P100	18.0	$2.45	
		V100	20.4	$3.36	

Analysis of the calculation rate data in Table 3 yields two important conclusions. First, the maximum calculation rate has increased dramatically from 0.6 seconds per day on a single CPU to 13.2 seconds per day on a V100 GPU using version 17.3.1 and OpenMP acceleration of the chemistry, as shown in Figure 2. Looking forward, simulations running on a V100 GPU card with GPU acceleration of chemistry could see calculation rates as high as 20.4 seconds / day.

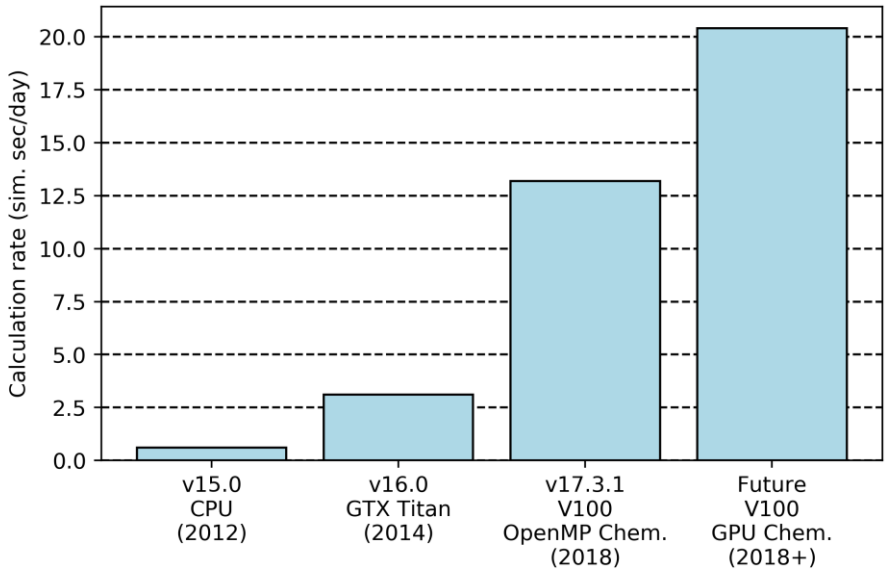

Figure 2 - Maximum calculation rates of CLR Model by release.

Executing simulations on a cloud platform allows the calculation rate and cost to be compared for different hardware configurations and parallelization advancements. Such a view is helpful in determining the ideal hardware for executing a simulation while also clearly showing the benefits of continued improvements to parallelization. As shown in Figure 3, while each GPU card upgrade comes with additional speed and cost, the calculation rate can become controlled by the un-parallelized parts of the code. For example, the calculation rates of 17.3.1 without any chemistry acceleration are seen to approach an asymptote at a rate of 9 seconds per day. Beyond this point, additional parallel computing power will not result in an appreciable increase in speed, as demonstrated by the upgrade from a P100 instance to a V100 instance, despite the over 50% increase in cost. Significant speed and cost savings can be obtained by eliminating the controlling un-parallelized functions, as seen in the cases with OpenMP acceleration and the future GPU acceleration which show both increased speed and lower costs compared with the 17.3.1 case alone. In the extreme, it could be cheaper and faster to run on an upgraded GPU card, as is the case with the simulation with GPU chemistry acceleration on a P100 instance which is 2.3 times faster than the K80 and $0.09 per simulation second cheaper. Furthermore, the continued parallelization improvements continue to shift the potential asymptote to higher calculation rates, which provides an expectation that increases in calculation rate can continue to increase with future hardware improvements, whereas older versions with less parallelization would stagnate.

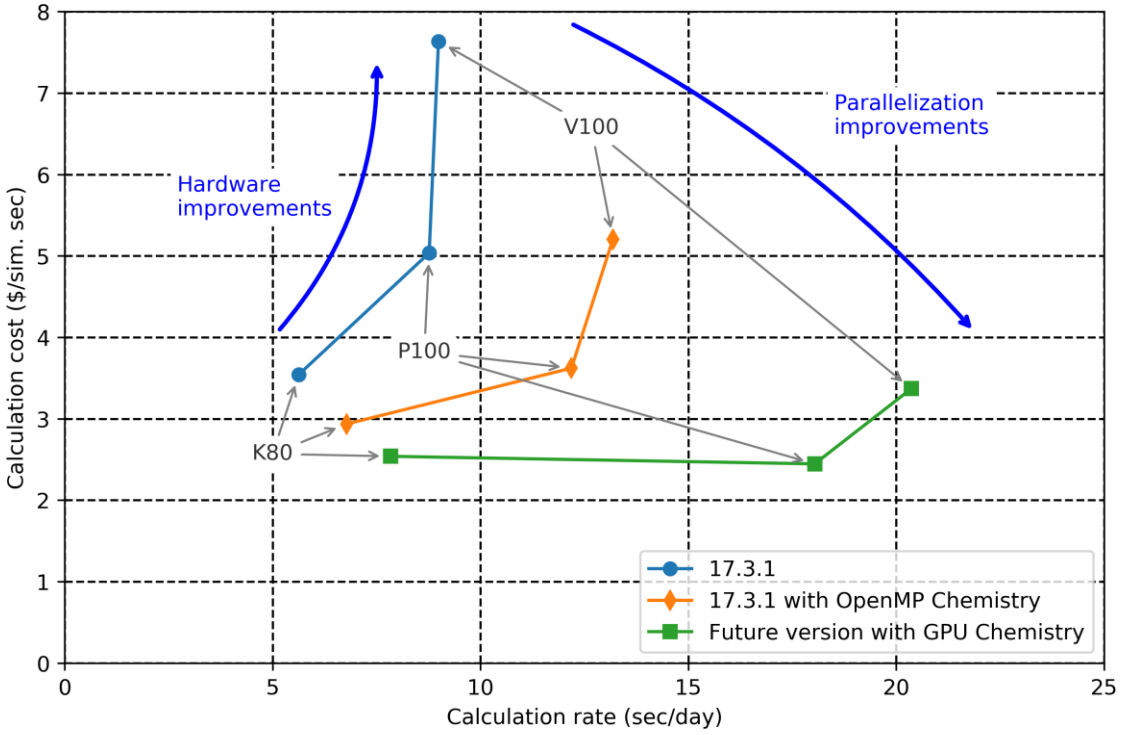

Figure 3 – Effect of hardware and parallelization improvements on simulation cost and calculation rate

5 Conclusions

Calculation rates of a CFD model are an important factor that drives the size and accuracy of a model and the feasibility of using the model as an engineering tool. Fortunately, recent improvements to computing hardware and software parallelization have increased the calculation rates dramatically for industrial-scale particle simulations. In this work, the increase in calculation rates is demonstrated using the NETL Chemical Looping Reactor model as a benchmark. Since the original development of this model in Barracuda Virtual Reactor, the calculation rate has increased from 0.6 seconds per day in 2012 to 3.1 seconds per day in 2014, to 13.2 seconds per day running on the latest generation GPU card and the most recent software release. A future release of the Barracuda Virtual Reactor, containing GPU acceleration of Eulerian chemistry, could see calculation rates over 20 seconds per day for this model. These increases in calculation rate are driven in large part by efforts to parallelize all compute-intensive functions in the solver, shifting the maximum calculation rate higher and allowing better utilization of the computing hardware.

References

[1] Parker, J. (2012) Simulation of Coal Particles in a Full Chemical Looping Combustion System. *Final report submitted to US Department of Energy, National Energy Technology Laboratory*. March 2012.

[2] Parker, J., Thibault, S., Williams, K. (2014) Advancements in the CFD Modeling of Clean Coal Technologies. *Clearwater Clean Coal Conference 2014*, Clearwater, FL.

[3] Abad, A., Adanez, J., Cuadrat, A, Garcia-Labiano, F., Gayan, P., and de Diego, L. (2011) Kinetics of redox reaction of ilmenite for chemical-looping combustion, *Chemical Engineering Science* **66**, 689-702

[4] Parker, J. (2014) CFD model for the simulation of chemical looping combustion, *Powder Technology* **265**, 47-53

[5] Andrews, M. and O'Rourke, P. (1996) The multiphase particle-in-cell (MP-PIC) method for dense phase particulate flows, *International Journal of Multiphase Flow*. **22**, 379-402.

[6] Snider, D., O'Rourke, P., and Andrews, J. (1998), Sediment flow in inclined vessels calculated using multiphase particle-in-cell model for dense particle flows, *International Journal of Multiphase Flow* **24**, 1359-1382.

[7] Snider, D. (2001). An incompressible three-dimensional multiphase particle-in-cell model for dense particle flows, *Journal of Computational Physics* **170**, 523-549.

[8] O'Rourke, P. J., Zhao, P., and Snider, D. (2009) A model for collisional exchange in gas/liquid/solids fluidized beds, *Chemical Engineering Science* **64**, 1784-1799.

[9] CPFD Software, LLC (2013). *Barracuda Virtual Reactor 16-GPU Release Guide*. Albuquerque, NM.

[10] CPFD Software, LLC (2015). *Barracuda Virtual Reactor Series 17 Release Guide*. Albuquerque, NM.

[11] Google, Inc. (2018). *Google Compute Engine Pricing.* Retrieved July 29, 2018, from https://cloud.google.com/compute/pricing.