Practical Deployment: Hardware, Operating Systems, and Other Considerations

Sam Clark
CPFD, LLC

2019 Barracuda Virtual Reactor Users Conference
June 27-28, 2019
Chicago, IL

Topics to Cover

Current hardware developments and recommendations

- CPUs and GPUs

Operating system compatibility issues

- Windows + Cygwin / Virtual Machine for post-processing
- CentOS 7 + GMV and xanim

Cloud computing

Sources of information and support

Open floor for Q&A

Hardware Strategy for Barracuda: CPU vs GPU

Barracuda versions pre-17.4.0

- Most functions parallelized onto GPU
- *Except* volume-average chemistry, which could use parallel CPU (OpenMP)

Barracuda 17.4.0

- All functions parallelized onto GPU
- Parallel CPU chemistry no longer advantageous

Best strategy for Barracuda performance

- Fastest available CPU, don't need high core count
- Fastest available NVIDIA GPU, newest generation

See Barracuda Hardware Timing Problem support site post

- Collected data from multiple CPU and GPU combinations

General CPU Options

CPFD uses Intel® Core™ i7 processors in our calculation machines

Intel® Core™ i9 processors

- More expensive than i7 processors
- Higher core count, but this is not an advantage for Barracuda

Intel® Xeon® processors

- More expensive that both i7 and i9 processors
- Higher memory capacity, but this is usually not a limitation for Barracuda
- Higher core count, but this is not an advantage for Barracuda

64-bit AMD processors are also compatible with Barracuda

Specific CPU Recommendations

Processor	Release Date	Current Price	Cores	Cache (MB)	Base & Max Turbo Frequency (GHz)	Max Memory (GB)	PCIe Lanes
i9-9920X	Q4'18	$1,200	12	19.25	3.50 / 4.40	128	44
i9-9900X	Q4'18	$1,000	10	19.25	3.50 / 4.50	128	44
i9-7920X	Q3'17	$1,000	12	16.5	2.90 / 4.30	128	44
i9-7900X	Q2'17	$950	10	13.75	3.30 / 4.30	128	44
i7-9800X	Q4'18	$600	8	16.5	3.80 / 4.40	128	44
i7-7800X	Q2'17	$400	6	8.25	3.50 / 4.00	128	28

Specific GPU Recommendations

GPU Card	Release Date	Price	Architecture	RAM (GB)	CUDA Cores
Quadro RTX 8000	Q4’18	$5,500	Turing	48	4608
Quadro GV100	Q1’18	$9,000	Volta	32	5120
TITAN RTX	Q4’18	$2,500	Turing	24	4608
Quadro RTX 6000	Q4’18	$4,000	Turing	24	4608
Quadro P6000	Q4’16	$3,700	Pascal	24	3840
Quadro RTX 5000	Q4’18	$2,150	Turing	16	3072
Quadro P5000	Q4’16	$1,700	Pascal	16	2560
TITAN Xp	Q2’17	$1,650	Pascal	12	3840
RTX 2080 Ti	Q3’18	$1,200	Turing	11	4352

General System Recommendations

Install as much RAM as your motherboard supports

- Watch out for motherboard + CPU compatibility in this respect

Use a fast SSD for the Operating System drive

- 250 GB to 500 GB options are available in the $50 to $100 price range

Install multiple large data drives for running Barracuda

- Format hard drives as individual drives, not RAID
- 4 TB to 8 TB drives are currently available at a good price vs capacity

Most common hardware failure we see is power supply (PSU)

- Replace with exactly the same model for easiest swap-out

Current Operating System 1st-Choice Recommendation

CentOS 6 / RedHat Enterprise Linux 6

Pros:

- Best compatibility with GMV, xanim, and Xmgr
- Stable servers for VNC, NoMachine, and AutoFS

Cons:

- Older software versions compared with newer distros
- Maintenance updates planned to end November 30, 2020
- Cannot run Docker containers

Challenges with CentOS 7 and Other Newer Linux Distros

GMV performance is poor for large models

- No “Reduced Resources” setting in Gnome 3
- GMV continuously redraws window during resize

xanim is half-transparent when playing animations

- Cannot disable compositor in Gnome 3
- Workaround: use a different video player

MATE desktop can help, but has its own issues

Our experience has been that some system services are unstable:

- VNC and NoMachine
- AutoFS

Challenges with Windows

GMV cannot open large model results

- 32-bit Cygwin imposes a 4 GB memory limit

Scripts run much slower in Cygwin terminal

- Reading/writing to the file system is slower in Cygwin than in native terminal

Support for some scripting languages is limited

- Scripts that usually work: Python, Perl, simple BASH scripts
- Scripts that often have problems: Tcl, complex BASH scripts

Workaround: use Virtual Machine for post-processing

Cloud Computing with Barracuda Virtual Reactor

CPFD is testing running Barracuda in the cloud

- Amazon Web Services (AWS)
- Google Cloud

Advantages

- On-demand compute and storage capacity
- Accessible from anywhere in the world
- No calculation machines to purchase, configure, and maintain

Challenges

- Running simulations in the cloud can be much more expensive than buying equivalent hardware
- Downloading large amounts of data from the cloud can be expensive (egress fees)
- Requires knowledge of cloud system administration and security
- To check out an RLM license, the cloud system must be able to communicate with your RLM server

Cloud Computing Cost Analysis

Older/slower GPU cards are available for a lower cost per hour

Newer/faster GPU cards cost more per hour

Barracuda parallelization improvements reduce the total time to complete simulations

The “optimum” card to choose for cloud computing may not always be the newest/fastest one

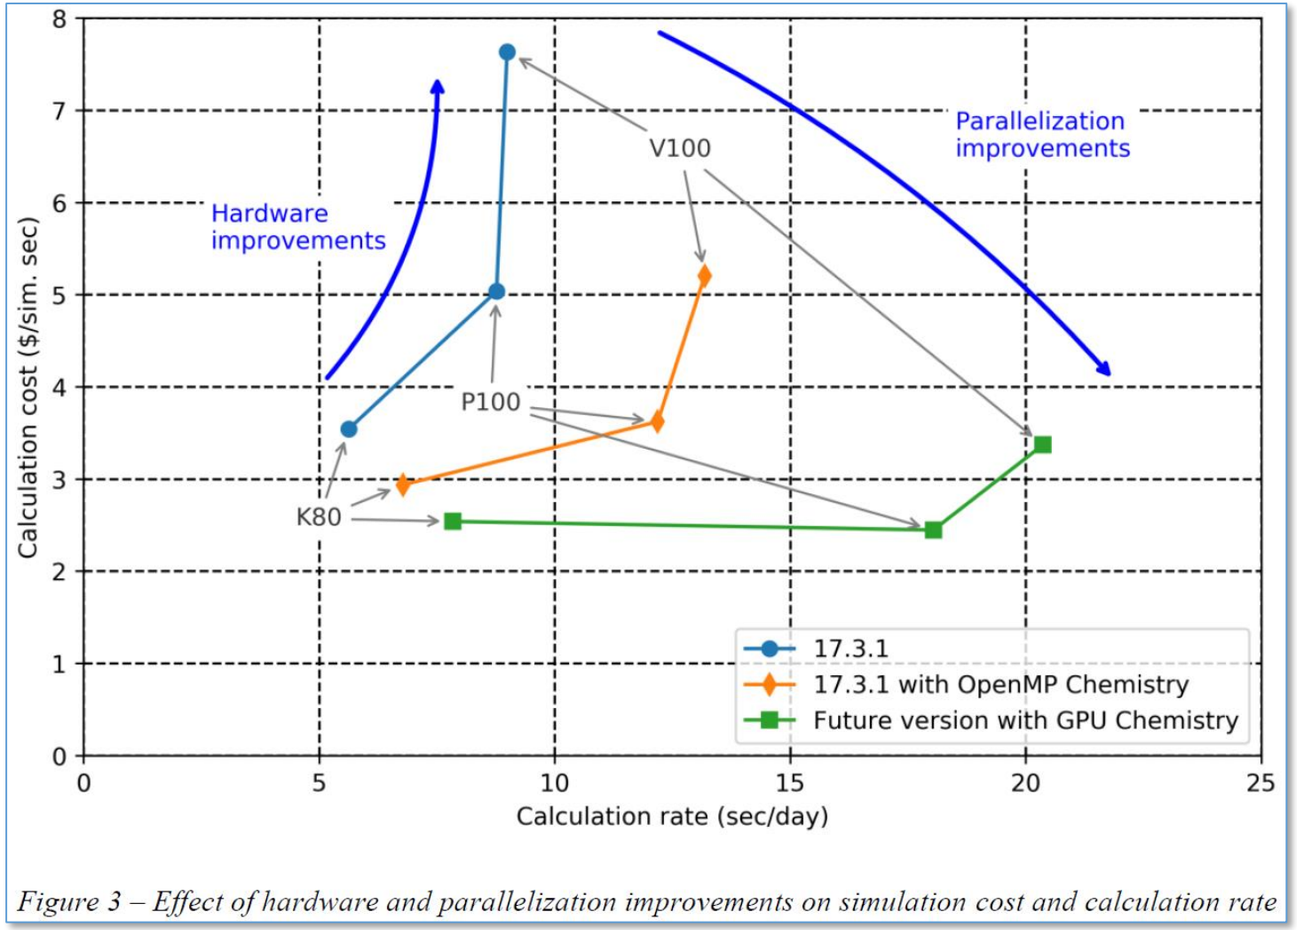

Figure 3 – Effect of hardware and parallelization improvements on simulation cost and calculation rate

Parker & Larson, “Application of Recent CFD Advancements to the Modeling of Chemical Looping Systems”, 2018.

Sources of Information and Support

User Manual

- Updated with each release of Barracuda Virtual Reactor
- Most up-to-date version is always on our support site

Support Site

- Register for an account if you don't have one

Email support@cpfd-software.com

Schedule a web meeting – sharing screens can be very helpful!

Call us at +1 (505) 275-3849 during USA business hours (8 am – 5 pm MDT)

Questions & Open Floor for Discussion