

P150 Chen - Increase productivity for Matlab computations

Ideas: Expand computational capabilities (power and/ or number of computers), increase efficiency of code, streamline workflows, etc.

See also

- [Chen - Consult on scaling up image processing work](#)

Purpose of page

Write down concerns, ideas, and efforts as understood by Chemistry IT so research group members can review and correct.

Chen Computational Server

1) Confirm support (draft)

Chemistry IT to do or figure out

Where in our 248 Baker Lab server farm? CPU and UPS. Power, switch/ networking.

Still need to confirm we can "turn off" logging into Windows using NetIDs.

- System is AD-based Windows using group's instrument account. And VNC remote access tool uses AD group to enable access by NetID, so confirm that still also works with any changes made to restricting Windows log-on.

Chen Group to do or figure out

Confirm no storage backup since can rerun computations potentially lost. If hard drive breaks, Chemistry IT can re-install and configure OS and applications. But all group data and results will be lost.

Does the analogy of using the server as a "toaster" resonate and work? The only time a failure of the server impacts work is when a job is being calculated. Once the data has been analyzed, the server should no longer have a copy of the original data analysed, nor the results. Those are like the bread, in a toaster.

Chemistry IT's select roles and responsibilities

- Be responsive and respectful of requests from the Chen group, including requests for software installations. Includes:
 - Installing or removing applications, configuring CPU threading.
- Conduct regular maintenance by Chemistry IT
 - At least every 3 months, work requires the system be taken off-line. Usually able to accomplish updates within a morning. Scheduled for every 3 months for a specific place-holder date and reminders that all jobs will be stopped are sent 2 weeks, then 1 week, then 1 day before the maintenance. On day-of work, any remaining jobs are terminated by Chemistry IT so they can accomplish the required work.
- Chemistry IT installs and maintains the OS and application updates and upgrades. Configures and maintains remote access technology.

Chen group's select roles and responsibilities

- Be responsive and respectful of requests from Chemistry IT, including planning around maintenance or emergency technical work.
- Understand that access will require using a VPN if accessing the server from personal computers, or from eduroam/ RedRover, or from home.
 - VPN is not required if server is accessed from current group's computational computers.
- Understand server is for processing data, not for on-going storage of data or results. Maintain data hygiene on server's storage drive to ensure that:
 - Server never has collected data that is not primarily available elsewhere. Once data has been analyzed, that data should be deleted from the server.
 - Group members do not use the server for long-term storage of results. Results should be copied elsewhere (example: Group's file share) and then deleted from the server.

2) Deciding on the final specs

This step enables Chemistry IT to get a "final" quote for Peng's review. If approved, Chemistry IT can place the order.

Below on this page is an outline on the costs and trade-offs of likely computation bottlenecks, for Peng's consideration.

Server: Dell R440

\$4,016.14 itself, plus:

- Any upgrades of the server's processors or memory (below).
- Server's storage drive (below): \$100-900 (not necessarily Dell).

Other anticipated, related costs:

- Uninterruptible power supply (UPS):
- Licensing: MatLab
 - Move from existing workstation (no change in annual costs), or add a license?
- Licensing: VNC
 - \$xx (per year?)
- Unlikely: Switch or other networking. Use existing infrastructure, paid by others.
- Unlikely: Rack or other physical provisioning costs. Use small portion of existing infrastructure, paid by others.

1) Processors

The server will have two Xeon Intel processors.

Confirm specifically which ones to invest in.

Intel model	Cores (actual)	Clock speed (GHz)	Cache (MB)	Cost change from tested	Notes
Silver 4110	8	2.1	11	(As quoted)	As tested
Silver 4114	10	2.2	14	\$350.90	
Silver 4116	12	2.1	16	\$919.60	

Many, many more options, for more money with gradations similar to the above, up to a 22-core (which is about \$6,000 more).

- See price sheet if open to spending more for processors.

2) Memory (RAM)

Total memory	Price	Notes
64 GB	subtract: (\$1,462.90)	As tested
128 GB	(As quoted)	

3) Storage (hard drive)

For performance, support, and maintenance, purchase a server-class, single, larger SSD drive.

- Dell pricing: ~\$0.8 to \$1/GB
- Non-Dell (for same drives?): Q: \$0.5 to \$0.6/GB?

Strategy: To get a significant price reduction, consider buying the drive separately from Dell server. Even though more cost-effective, still will be serviced under (a non-Dell) warranty.

- Quote server with cheapest drive available (and don't use), if can't buy without a hard drive.

Solid State (SSD) storage	Dell pricing (full price)	Price we'd likely pay elsewhere
960 GB	846.40	~\$500?

Many, many more options, whether for more or less storage.

- Use \$/GB ratios above to estimate price for other desired sizes.
- See price sheet for specific options with Dell's pricing.

4) Warranty

Most research groups expect to use server investments for 5 years. And hope for more years, beyond the 5.

Normal warranty is for 3 years, but can instead get for 4 or 5 years.

- Question: Worth getting quotes for adding 4th and 5th year warranty, vs. base of 3 years?

Most groups take a risk for years 4 and 5, taking into account (1) rapid technological obsolescence and (2) historically low failure rates.

- Thus, groups are prepared to pay for fixes, or buy new, rather than put initial money into a more extended warranty.
- Can manage risk by buying more years up-front, up to 5.

Request from group: Buy a server to process multiple Matlab calculation simultaneously

Considerations, expectations, and outstanding questions:

- Windows Server or Linux server OS?
 - Past testing by researchers have yielded faster processing on servers running Linux than Windows. Some researchers only comfortable working within Windows.
 - If Windows Server used, buy ~\$100 software on server (free clients) to enable moving large amounts of data to server more speedily.
 - If Windows Server used, Chemistry IT must work out the "how", with group's input. Examples of unknowns:
 - How will user accounts and their access work? How many simultaneously? Per user accounts or shared accounts? How manage contention, if desired by group?
- There will be monthly (or maybe every 3 months) shutdown periods to ensure baseline patching and OS, file-share and hardware checking.
 - Chemistry IT will coordinate with group, as we do for many other servers we manage for other research groups.
- Chemistry IT expects group to purchase and maintain an adequately sized, rack-mount uninterruptible power supply (UPS) for servers we manage.
 - Approximate cost is **\$200-450**; need to size correctly. Expect batteries to last 3-4 years easily.
- Budget about **\$200** (reality-check amount with Michael Hint) for hardware costs related to the rack and networking required.
 - Chemistry IT can often contribute some used items "free", too, if available and of adequate quality.

Server spec suggestions and costs

Dell's help for choosing processors, memory, and hard drives/ SSD's (respectively) for their R440 server, valid as of 8/9/2018:

- <https://www.dell.com/learn/us/en/04/help-me-choose/hmc-processor-intel-xeon-14g>
 - Pricing via KFS: [Dell-R440-Proc.JPG](#)
- <https://www.dell.com/learn/us/en/04/help-me-choose/hmc-memory-14g>
 - Pricing via KFS: [Dell-R440-RAM-128GB.JPG](#)
- <https://www.dell.com/learn/us/en/04/help-me-choose/hmc-server-hard-drive>
 - Pricing via KFS: [Dell-R440-SSD.JPG](#)

		Minimum theoretical server				
Criteria	Current desktop's specs	Represents ~4 TIMES the desktop specs	Cost increase if go up a level or so	FIRST borrowed Dell Server specs, PowerEdge R430 (1 U) "CHEN-21"	July/ August 2018 SECOND borrowed Dell Server specs, PowerEdge R440 (1 U) "CHEN-23" Cost: \$3,135.69	Notes

Cores, hyper-threading (HT)	4 cores (one processor) i5-6500, 3.20 GHz (Not HT capable)	Total 16 cores (Two, 8-cores each) Each: Xeon Silver 4110, 2.10 GHz 11MB Cache (85W) HT-capable	16 total cores => 20 total cores <ul style="list-style-type: none"> • +\$460: 2.20 GHz • +\$1,628: 2.40GHz 16 total cores => 24 total cores <ul style="list-style-type: none"> • +\$1,200: 2.10GHz • +\$1,960: 2.30GHz* • +\$3,060: 2.60GHz* • +16,060; 3.00GHz* *Availability delay, as of 12/22 /17.	Dual Intel Xeon E5-2620v4 2.1ghz processors with 8 real cores / 16 virtual cores per processor More details: 20MB Cache (85W)	Dual Intel Xeon Silver 4110 2.1ghz processors with 8 real cores / 16 virtual cores per processor	Hyper-threading (HT) is useful to the group. Group tested performance between HT being turned on and off on "CHEN-21". With HT, test system showed 32 cores. 25 of these could be used without saturation, and thus without delaying processing time. Without HT, only 16 cores available so maximum much less than 25 effectively used under HT. Notes: <ul style="list-style-type: none"> • See chart below for options with more cores and faster speeds. • See write-up below chart for why we are pricing two-processor server options instead of single- or four-processor servers.
Storage All SSD.	500 GB SSD	2TB is 4 times the space: 2.0TB Samsung 960 PRO M.2 PCIe 3.0 x4 NVMe Solid State Drive (cost is \$1,499) Above configuration is the fastest option since combines OS and data on fast bus.	If need less space, easy. Save about \$375 per 500GB. Replace 2.0TB at left with: <ul style="list-style-type: none"> • 512GB, at \$399. • 1.0TB, at \$749. If more space is needed, more complicated and slower, but of course doable. Please ask.	400 GB SSD (380GB usable)	RAID-0 (striped) dual 240GB SSD drives acting as a single 480GB SSD for performance	Large amounts of data not needed to be stored on server, nor moved from server. Instead, simply deleted from server after processing.
RAM	32 GB	128 GB	+\$1,440: 128 GB => 256 GB	32 GB	64GB (128GB is \$880.55 more)	
Warranty	3 year, NBD Dell on-campus "parts locker"	3 year, Advanced Parts Replacement Warranty	Warranty upgrades (all non-NBD): <ul style="list-style-type: none"> • +\$150: 3 => 4 years • +\$375: 3 => 5 years 	5 year, NBD, 5x10 on-site.		
Other	n/a	n/a	Option: Redundant power supply unit (+\$224). Option: Any value in Dual and/or 10-Gigabit Ethernet? (+\$150-500)	Network card: Quad-port, 1Gb.		
Total cost, approx.	~\$1,000 each (*4 => \$4,000)	~\$5,500 for just the server	Easy to add \$500 - \$18,000 to SAME server!	~\$3,700. Compared to minimum recommendation: <ul style="list-style-type: none"> • Older, cheaper processors • Much less RAM • Much smaller storage • Warranty is 2 year longer and on-site • Dual power supplies • Better network card 		

Core prices snap-shot

- Prices taken 12/22/2017
- Prices go up dramatically with both core-count and processor speeds.

- Cache sizes and power consumption also increase with capabilities.
- This chart of estimated marginal costs is a sub-set of options pulled from <<https://www.thinkmate.com/system/rax-xs4-21s1-sh>> and <<https://www.thinkmate.com/system/rax-xs4-21s1-10g>>, focusing on SLOWEST procs to reduce costs. Other, more expensive-per-proc (faster) options available.

Approximate price increase buying TWO processors compared to 8-core Xeon Silver 4110 NOTE: Can also get a server with just ONE processor (at half the marginal cost), if core-count is sufficient.	Processor (all Intl Xeon)	Actual core count EACH processor Obtain total count by multiply by 2 since 2 procs. All HT-capable	Other specs
\$0 (base-line, for price comparison)	Silver 4110	8-core => 16	2.10GHz 11.00MB Cache (85W) (Same price as 4-core 2.60GHz 8.25 MB Cache (85W) version)
+460	Silver 4114	10-core => 20	2.20GHz 13.75MB Cache (85W)
+1,200	Silver 4116	12-core => 24	2.10GHz 16.50MB Cache (85W)
+1,628	Gold 5115	10-core => 20	2.40GHz 13.75MB Cache (85W)
+1,960 (availability delay)	Gold 5118	12-core => 24	2.30GHz 16.50MB Cache (105W)
+3,060 (availability delay)	Gold 6126	12-core => 24	2.60GHz 19.25MB Cache (125W)
+3,280	Gold 6130	16-core => 32	2.10GHz 22.00MB Cache (125W)
+5,060	Gold 6138	20-core => 40	2.00GHz 27.50MB Cache (125W)
+7,560	Gold 6152	22-core => 44	2.10GHz 30.25MB Cache (140W)
+6,560 (availability delay)	Platinum 8153	16-core => 32	2.00GHz 22.00MB Cache (125W)
+15,660 (availability delay), and +\$400 chipset	Platinum 8158	12-core => 24	3.00GHz 24.75MB Cache (150W)
+10,160, and +\$400 chipset	Platinum 8160	24-core => 48	2.10GHz 33.00MB Cache (150W)
+13,560, and +\$400 chipset	Platinum 8164	26-core => 52	2.00GHz 35.75MB Cache (150W)
+16,660, and +\$400 chipset	Platinum 8170	26-core => 52	2.10GHz 35.75MB Cache (165W)
+19,760, and +\$400 chipset	Platinum 8176	28-core => 56	2.10GHz 38.50MB Cache (165W)

Why server with 2 processors? Why not one or four, for example?

No savings buying just a one-processor server, especially if want more cores. JUST the price of the process jumps by the cost of an entire server!

- +\$5,080: 16 total cores => 24 total cores (one proc)

Consider four processors only if needed and can afford many more cores. Option is not cost-effective for the core-counts we are currently looking at. Other considerations:

- Currently, only older processors are currently as an option with four-processor servers.
- Server prices START at about \$12,000, for 32 total cores.
- Core counts and price upgrades are "times 4", not "times 2".

What if need more than 2TB total storage?

If need more than 2TB total, likely must use more complicated, multiple storage option.

- We can dig deeper to confirm this, if necessary to decision-making.
- More storage likely will be slower, even is sticking to SSDs.

An example of an option to get more than 2TB storage:

First: Choose boot drive, which can be small. For example:

- OS only: 512GB Samsung 960 PRO M.2 PCIe 3.0 x4 NVMe Solid State Drive (cost is \$399).

THEN: Choose a single main storage drive. Some examples and their costs:

- Data ONLY fast SSD:
 - 3.84TB => \$1,500 - 2,100

- 7.68TB => \$3,000
- Data ONLY slow spinning, compared to current desktop's much faster SSD:
 - 4TB => \$200
 - 8TB => \$350
 - 12TB => \$650

Status

10/5/17: Oliver met with Mahdi <mh2356> and Kushal <ks2285>. Action steps to have Peng review and refine:

(1) Group: Decide if worth having (select?) Matlab code reviewed by experts at CAC, focused primarily to increase efficiency. Secondary outcomes include:

- May result in the ability to run older code on current version of Matlab, expanding where code could run on non-group computers.
- May result in clarifying computational bottlenecks so the best fitted computational hardware is purchased. What does one prioritize when faced with choice to invest in: better processors, number of processors, number of cores per processors, bus speeds, SSD drives, and/ or RAM?
- May result in a confirmation whether or not problem lends itself to parallelization. If so, can increase efficiency with the right hardware and expands the locations to efficiently run the code (RedCloud, etc.).

(2) Oliver: Have group test their code on test server in 248, initially by-passing using the network to get the data to the server.

- Time comparisons of both single runs and simultaneous runs. Does the server reduce computational time for a single job, as compared to current workstations? To what degree does the server's performance drop as more jobs are added? Again, compare to current workstations.

(3) Oliver: Optimize getting data to test server in 248 via the network.

Other thoughts from Oliver:

- Confirm if any campus computing is a good fit for the group: CISER, RedCloud (likely only if code can be and is parallelized), David Botsh's cluster (?), others?
- Group may benefit from optimizing workflow at various workstations.

Prior conversations, for historical background

Chen - Consult on scaling up image processing work

Chen process many, many images using group-developed software (on Matlab?). It uses dedicated computers (all in the lab?) to meet current needs, but Peng is concerned that it will not scale well. In addition to processing power, effort demands and produces copious amounts of data, currently stored on one-off, external hard drives.