Abstract

As society becomes more and more reliant on ML and AI technologies, we are increasingly turning to large amounts of data and more efficient data storage methods. One such method that has been proposed is quantum computing, where typical bits are replaced with qubits, or quantum bits. The big advantage of qubits, and of quantum computing in general, is that while bits can only be solely in a 0 or 1, qubits can be in some combination of both at the same time, which is a physical concept known as superposition. As such, a classical computer needing an exponential number of bits to represent an exponential number of possibilities can be represented with the base number of qubits, since they can represent all possibilities at the same time. However, a major challenge in creating this environment is ensuring that it is isolated from outside influence or factors, since qubits are especially sensitive to noise. One common measure of the quality of a system is the coherence, or the time during which a quantum system can maintain a quantum state. To succeed in quantum computing, we need to develop materials with high quantum coherence, and one such class of materials that has been suggested for this purpose is the use of rare earth metals.

The more a qubit interacts with the outside world, the more likely it is to decohere. From both my chemical knowledge and the research I've done, it looks like there are 5 main chemical properties that can influence coherence: electron shielding, electron-phonon interaction, hyperfine interactions, spin orbit coupling, and optical coherence.

My goal for this project was to calculate the values for these characteristics, but as we all know, that didn't go exactly as planned. I'll touch on all of these later, but re more practical learning, I learned that simulations really aren't great at reflecting real world behavior. Conceptually, I learned a lot about basis sets and which are used in which scenarios. Practically, for the smaller version of this project with the limited features I was able to complete, it turns out that the coordination number of the center element seems to affect both the size of the system and the resulting amount of electron shielding.

Simulation Methodology/Methodology Discussion

Due to the nature of my project almost being iterative, I've decided to combine my simulation methodology and methodology discussion sections. It will allow me to more easily discuss what I've tried, why it didn't work, and what I could have done differently on the way to the final calculations that were done.

For many of my simulations, I chose to use DFT with B3LYP and the 6-31G(d), if (d) applicable, basis set. I chose DFT since it gave a balance of computational efficiency and accuracy, and since I felt a method based on electron density might give more accurate results when it comes to metals. While other more complex wave function based methods might perform better, I also knew that my project was already going to be computationally intense from the start due to the focus on transition metals. I chose B3LYP because it was a hybrid functional and would thus likely perform well on both the water molecules and the transition metals themselves. However, knowing what I do now about the problems that arose when I tried to use

these settings, I'm wondering if a more accurate method (MP2?) or one which focused specifically on weaker interactions might have given better results.

First Idea: individually optimize the water molecule and then add it to the rare earth. For my water molecules, I ran a DFT-B3LYP 6-31G optimization, which worked. Unsurprisingly, this was quick and easy, but I soon realized that due to the constraints of building a system in Gaussian (as well as the time + effort it would have taken to painstakingly create this system), that it would be better to optimize a rare earth element first and then add the water molecules. I think this idea could be improved upon if I was a bit more experienced in the methods in Gaussian- copying, rotating, and placing molecules at exact measurements is still pretty tricky for me, and while I don't know if this would have solved the problem, I'm sure it would have helped.

Second idea: optimizing the RE within a water molecule cage. For this step, I did some research on hydration numbers and coordination numbers to determine what configuration would give me the greatest chance of success. In most cases, it turns out to be a coordination number of 4,6, or 8, and it looks like most atoms within these systems have a charge of +3. Initially, I tried cubic, radial, and octahedral structures, where each water molecule was about 2.5 angstroms (that seems to be an average of a lot of the numbers I saw) away from the rare earth metal. But, even after learning from the literature that an octahedral structure seemed to be best, I ran into quite a few problems with my basis set.

Initially, I was planning to use the LAN2LDZ basis set, since I thought (wrongly) that this was the basis set most typically used for transition metals. I was running into a lot of convergence errors, so I tried upping the SCF cycles. However, upon closer inspection, it only seems to work through La. I first tried to use the SDD basis set, seeing as it dealt with core potentials and would hopefully simplify the computational process, but that didn't work either. This is something I should have realized when starting this project, but Gaussian natively supports very few basis sets that work on transition metals. As such (as everyone who is writing/reading this paper knows), I researched, learned, and tried running the calculation with a custom basis set specifically designed for transition metals. While that did end up running, I continued getting SCF errors even after increasing the number of cycles and the memory. If I were to do this differently next time, I'd make this calculation alone my whole simulation report. I think it could have benefited from more time, trial/error, and research on my part, since I had no idea when I started just how complicated this step would be.

Third idea: just run the calculations on a la system that (yay!) had actually converged. In troubleshooting the LAN2LDZ basis set, I tried running it on La, the highest element it purported to be capable of. But, when I tried doing population analysis on the La system, I got errors that my basis set couldn't find the "f type valence orbital". Thanks to a nudge from Frank, I started trying different basis sets, but once again was met with a fair number of errors (with LAN2LDZ, SDD, and the custom ones as well). Once again, I think something that could have helped here was truly understanding how these basis sets worked and what the problem at hand was. From a

higher level, I understand what population analysis is, but I couldn't tell you how it interacts with basis sets and what steps I would take to create a basis set that wouldn't cause this problem. I think doing just this one simulation could again make up a whole simulation report.

Fourth idea: run the systems with metals that were higher up on the periodic table. I knew at this point that the transition metals were the atoms causing the problems in my simulations, so I decided to replace them and try systems with Va, Cr, Ti, Mn, and Fe. While the first three continually had SCF convergence failures (even after several restarted iterations), the second two worked! I then proceeded to start with the simplest simulation I had initially proposed- mulliken/npa population analysis. I ended up choosing the NPA analysis since my research told me Muliken wasn't as accurate on metals.

At this point, I wasn't able to get anything else to run, but I also tried doing NMR spectroscopy, to see which system had a larger shift. Once again, I was met with the error message that NMR doesn't work well on metals, which is a bit frustrating, and something I wish I had known when I proposed many of these calculations. The rest of the simulations were things I didn't even attempt- they were more complicated versions of custom basis sets, or analyses that were even more computationally intensive than the NMR I had just tried to troubleshoot, and as such, likely weren't going to run either.

Overall, parts of the project that went well: my research + project idea. Parts of the project that didn't go so well: project scope, actual calculations, and Gaussian debugging. For this project, I was able to come up with an idea that was both interesting to me and relevant to the field I'll likely go into, so that was a big win. I also did (or at least thought I did) a lot of research to inform this project, and even if it still didn't turn out to be enough, I really enjoyed learning about this topic a bit more in depth. Conversely, though, I don't think I did anywhere near enough research to even understand the surface of what I was doing. My project was too large in scope, I wasn't experienced enough to debug the simulations I was proposing, and had no idea just how annoying the calculations I proposed would turn out to be.

Analysis of the Simulation Results

Molecule	Graphs
La	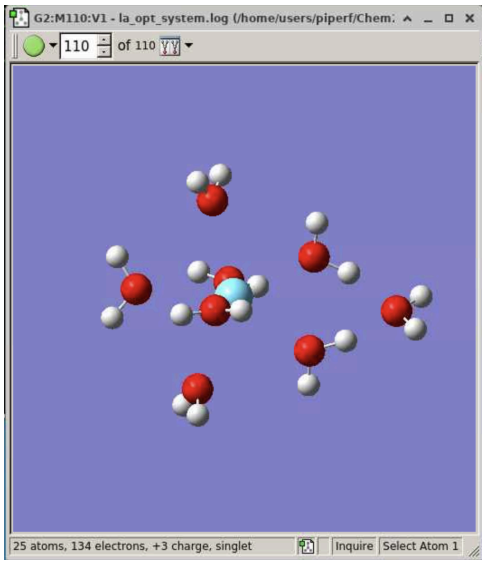
Fe	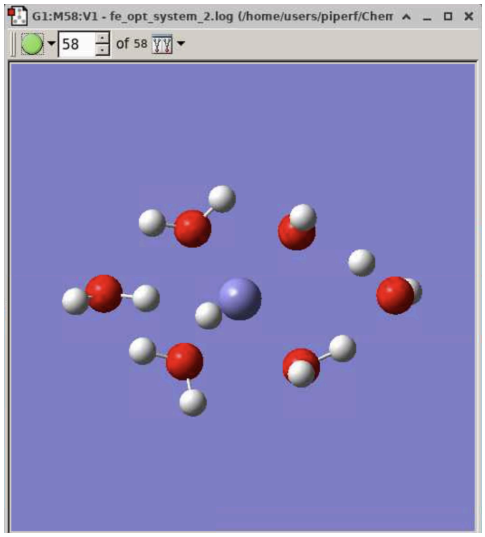 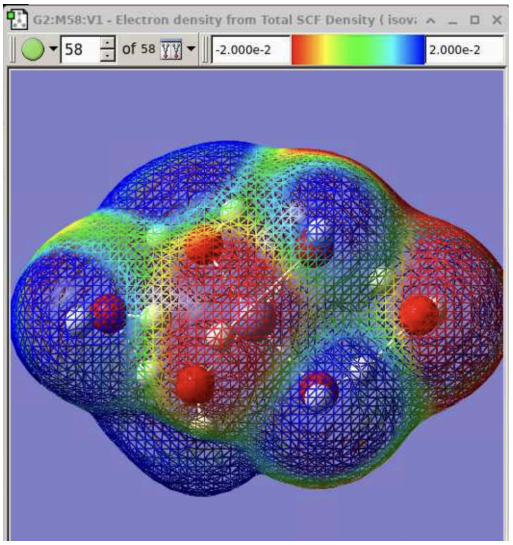

Mn	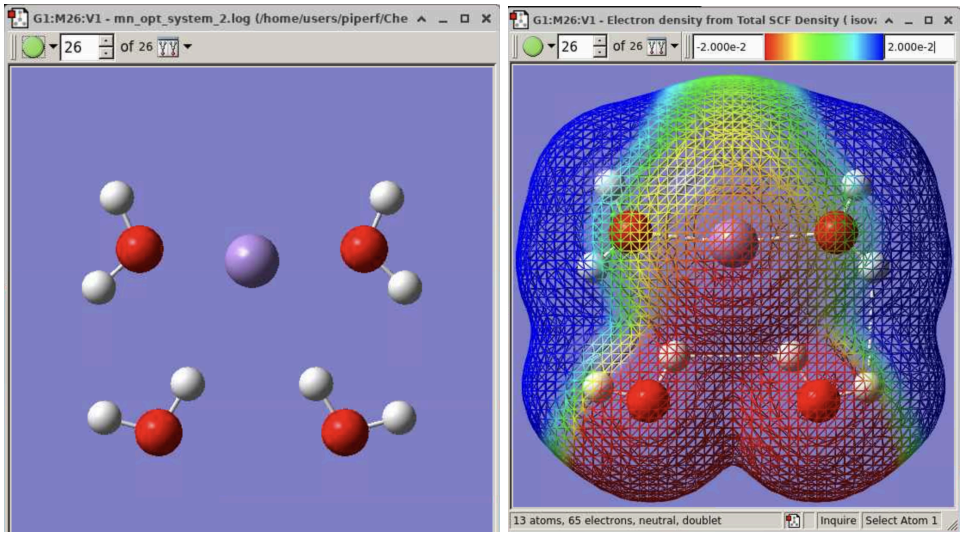

While I couldn't get the La NPA analysis to actually run, judging by the other two simulations, my guess is that it would actually be more similar to the Fe system, considering that it has a higher coordination number. Unsurprisingly, in both cases, we see the electrons cluster more towards the hydrogens and around the edges of the systems. While this is far too little data to actually extrapolate from, if higher coordination numbers do typically correlate to greater shielding, we'd want to choose metals with larger coordination numbers (to a point, of course). Having more shielding means that electrons in quantum computing are less likely to decohere and lose their states, meaning they'd hold onto their information for longer. We see this in the greater shielding in the Fe system, where it's not only bigger, but takes up space in several dimensions.

One other thing I'll comment on, even from this small sample of data, is that I was quite surprised to see the La system get optimized the way it did. My intuition was that it would likely stay mostly octahedral, since the solvent would be the most spread out that way. In addition, an octahedral system would allow the system to more easily crystalize and pack with other systems, which is important for creating solid crystal ligands in quantum computing. While I knew from my previous research that La wasn't one of the elements in consideration for these systems, this result makes me wonder if the reason why is the lack of density in the solid form.

Curiously enough, when building these systems, I spent a lot of time orienting the oxygen atoms inwards, since I knew that they'd likely be oriented towards the positively charged metal ions. However, when we look at the Mn result and graph, we see that the bottom two molecules aren't facing in that same orientation, thereby creating the relatively charged positive region that we see at the bottom. This is a bit counterintuitive, considering that it puts the oxygen farther away from the regions of electron density.

Future Work

After all of that simulation and trial and error, here's what I would do differently next time:

1. Do far more research on the scientific background of what I was doing. One of the areas in which this project really suffered is the lack of similar published literature about this project. This was mostly something that I thought up and thought would be cool to explore, but it turns out that the reason there's very little about this topic is because it's frustrating to research, and to the best of my knowledge, most of the tools we're currently using aren't powerful enough to support this kind of research and simulation. If I had chosen a topic more grounded in current research, I would have likely had similar papers, other simulations, and other resources to turn to when I got stuck.
2. Extensively minimize the scope of my project. One of the things I learned during this project (and during this class!) is that the starting configuration, basis set, functional, and so many other factors can have a massive impact on whether your system converges or not. If I were to spend more time on this project, I'd probably start with a simple system/basis set convergence test to figure out which basis sets worked, why, and which gave the best results. One of the things I learned from our paper discussions is that a lot of tradeoffs happen when you're considering time and computational resources, and I think I had to make some of those same decisions here. I say this because I found that there were a lot of gaps in my knowledge- while we had learned about the basis sets in class, I wasn't familiar with the applications of them to transition metals, or how to debug or conceptualize the errors I was getting.
3. Learn more appropriate tools for this job. When researching the properties that lead to quantum coherence, I was making a lot of assumptions based on what Gaussian calculations might be able to approximate these properties, and have now realized that most of those calculations were both pretty hard to do in Gaussian, and wouldn't have approximated the property well anyways. I'd be curious to research and learn more about what tools are typically used to calculate these properties, and hopefully employ them to more success than I did in Gaussian.

References:

Quantum Computing, Progress and Prospects, Mark Horowitz, Quantum Coherence and Nonlocality of Two Qubits in the Presence of a Common Dephasing Environment, https://www.sciencedirect.com/science/article/pii/S221137972300459X#:~:text=We%20display%20the%20effect%20of%20the%20qubit%E2%80%93qubit,and%20quantum%20correlations%20and%20show%20that%20quantum

Quantum networks using rare-earth ions, https://arxiv.org/html/2501.06110v1

Rare-Earth Doped Materials for Applications in Quantum Information Storage and Signal Processing, https://www.sciencedirect.com/science/article/pii/S002223131000534X, Roadmap for Rare-Earth Quantum Computing, https://arxiv.org/pdf/2103.15743

Coherent Control of Rare Earth 4f Shell Wavefunctions in the Quantum Spin Liquid Tb2Ti2O7, https://www.nature.com/articles/s41467-024-51339-0, Enhancing Spin Coherence in Optically Addressable Molecular Qubits through Host-Matrix Control (ZFS), https://journals.aps.org/prx/abstract/10.1103/PhysRevX.12.031028

Electron Spin Coherence in Optically Excited States of Rare-Earth Ions for Microwave to Optical Quantum Transducers, https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.122.247401

Selective Optical Addressing of Nuclear Spins Through Superhyperfine Interaction in Rare-Earth Doped Solids, https://arxiv.org/abs/1709.01959

Universal Quantum Computing Using Electronuclear Wavefunctions of Rare-Earth Ions, https://journals.aps.org/prxquantum/abstract/10.1103/PRXQuantum.2.010312

Enhancing Spin Coherence in Optically Addressable Molecular Qubits through Host-Matrix Control (ZFS), https://journals.aps.org/prx/abstract/10.1103/PhysRevX.12.031028

Long Optical Coherence Times and Coherent Rare-Earth Magnon Coupling in Rare-Earth Doped Anti-Ferromagnet, https://arxiv.org/abs/2411.05182