

Structure and Long Time-Scale Kinetics of Nanoparticles

Matthew Welborn

July 31, 2010

1 Background

With the onset of the energy crisis, it has become increasingly important to pursue alternative sources of energy. One such potential source is the fuel cell. Unfortunately, current proton exchange membrane fuel cells are hamstrung by lack of good catalysts for oxygen reduction[1]. Nanoparticles (figure 1) are an attractive option for solving this problem[2].

Nanoparticles are small clusters of metal atoms, ranging in diameter from about one to one hundred nanometers. They have recently gained attention as potential catalysts for numerous reactions, including oxygen reduction, carbon monoxide oxidation, alkene hydrogenation, and p-nitrophenol reduction[3,4]. There are two main benefits to using these particles as catalysts. First, they have an increased surface-area-to-volume ratio relative to the bulk. Because catalysis occurs at surfaces, this produces a catalyst with greater activity per unit mass. Further, nanoparticles have been shown to exhibit qualitatively different reactivities than their bulk counterparts.

Unfortunately, little is known with certainty about the structure of these particles. Such information is vital to modeling the reactivity of these particles. More fundamentally, knowledge of the structure is key to developing intuition about their properties and synthesis.

Three groups of platinum nanoparticles of diameter 1.7 nm (140-147 atoms) were synthesized. The first group of particles, DENs, had no ligands bound to their surface. The second and third group had amine and thiol ligands bound to their surface. Unfortunately, the structure of these particles cannot be readily probed. Direct methods like TEM either do not have high enough resolution to resolve the atomic ordering of the particles, or will cause reduction of the particles. However, x-ray powder diffraction can be used to

probe the structure of the particles, yielding a pair distribution function (PDF) shown in figure 1. The data is difficult to deconvolute, but a qualitative interpretation suggested that ligand-bound particles had high disorder at their surface.

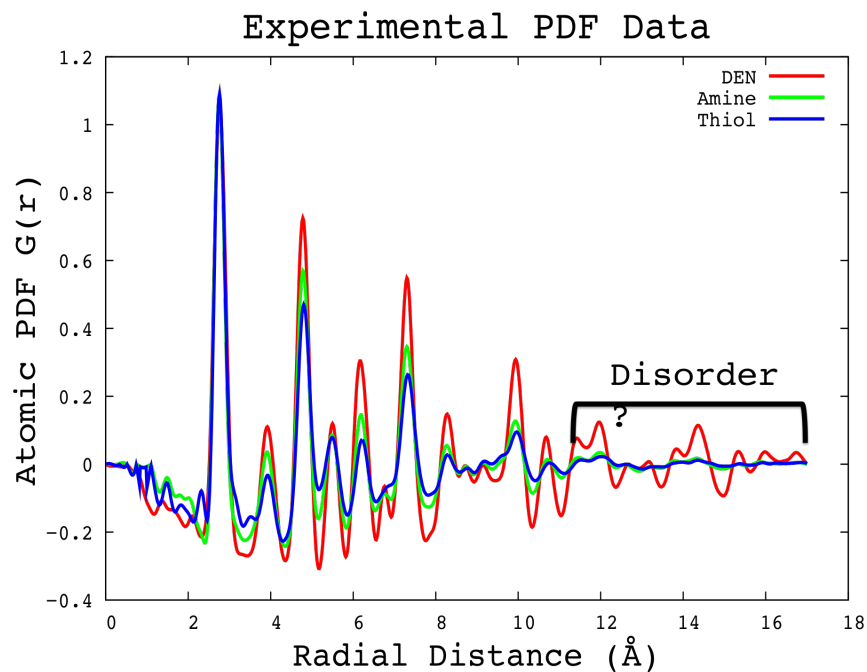


Figure 1: Experimental PDF data. The ligand-bound particles appear to exhibit surface disorder.

PDF data can be matched to structures fairly easily for periodic systems. For nanoparticles, however, fitting a structure is not straightforward. Previous attempts to fit data for gold nanoparticles have resulted in highly disordered structures. In these methods, a Monte Carlo optimization approach was applied to fit the data. These predictions conflict with theory, which predicts highly ordered structures. However, when theory alone is used to predict the structure of the particles, the resulting structures poorly fit the PDF data. As such, we took an approach that combined theory and experiment.

2 Approach

We created a function describing the goal of satisfying both experiment and theory by optimizing the following objective function:

$$U = \alpha Energy + (1 - \alpha)\Delta PDF \quad (1)$$

The above function combines theory (energy) and experiment (ΔPDF) to describe the fitting error for a given structure. It contains a parameter, α , determines how much emphasis is placed on theory or experiment. Large α places emphasis on theory, small α on experiment. This function was first minimized globally using the fast and relatively inaccurate embedded atom model to evaluate energy. Candidates from the global optimization were then examined in more detail using the more accurate and slower density functional theory to evaluate the energy.

3 Methods

3.1 PDF Error

The pair distribution function of particles was calculated using a modification of the method described by Neder and Korsunitsky. The following modified equation was used to calculate the pair distribution function, G:

$$G_{calc}(r) = \frac{A}{r} \sum_{i,j} \frac{1}{2\pi\sigma^2} e^{-\frac{r-r_{ij}}{2\sigma^2}} - 4\pi B \tanh(S(R-r)) \quad (2)$$

where r_{ij} is the distance between platinum atoms i and j , R is the largest distance between two atoms in the particle, $\sigma = 0.01$ (the resolution of the experimental data), $S = 0.132$, and A and B are fitting parameters. S corresponds to the overall shape of the particle. A scales the PDF to fit the data (because intensity is proportional to a bunch of uninteresting constants). B defines the size of the baseline correction. Here, a gaussian is substituted for the delta function used by Neder and Korsunitsky in order to both smooth the PDF and to make the derivative of G continuous. The PDF error was calculated using the following integral:

$$\Delta PDF = \frac{1}{R} \int_0^R (G_{exp}(r) - G_{calc}(r))^2 dr \quad (3)$$

with A and B chosen such that they minimize ΔPDF .

For the purpose of minimization, the derivative of the PDF error with respect to the atomic coordinates was also calculated. In this derivative, A and B were assumed to be constant.

3.2 Energy

Energy can be determined rapidly and approximately with the Embedded Atom Model. The Embedded Atom Model (EAM) developed by Voter[6] is an empirical potential for the set of FCC metals (of which platinum is a member). It has two terms. The first term is a pair potential which describes the two-body interactions between pairs of atoms. While this is a good approximation for rare gas clusters, a correction is needed to model the properties of solids. The second term is roughly related to the coordination of the atoms surrounding it; each atom feels a “density” from its neighbors, and there is a function which determines the effect this local density on the energy of that atom.

$$E = \sum_{i,j} V(r_{ij}) + \sum_i f(\rho_i) \quad (4)$$

where i and j index the atoms, V is the morse potential, and ρ_i is the local “density” of atom i .

EAM is fit to the bulk properties of platinum, so it is very approximate for nanoparticles. However, 140 atom nanoparticles are sufficiently large to be modeled effectively by EAM. EAM is unable to model ligand bound to the nanoparticle surface. As such, when EAM is used for the energy, all contribution to the structure of the particle comes from the Δ PDF term. While it is approximate, this model is incredibly useful because it is fast. Global optimization of the function requires an incredible number of evaluations of the energy, so EAM is used for the global optimization.

Density Functional Theory (DFT) can be used to get a far more accurate energy and also incorporate the effects of ligands. It is an ab initio quantum mechanical theory, and gives highly accurate energies. Unfortunately, DFT calculations are incredibly expensive; local minimization of a structure can take 384 processor hours on a supercomputer. (For comparison, EAM is about one billion times faster.) As such, DFT cannot be used to perform a global search. Instead it is used to more finely resolve the energies of the best candidate structures found by EAM.

3.3 Global Optimization

The objective function was minimized with a Metropolis Monte Carlo scheme with basin hopping. The general flow of this algorithm is:

1. Pick a random starting coordinate on the function.
2. Find the value of U at the local minimum corresponding to this point. Call it U_1 .
3. Displace from this coordinate in a random direction
4. Find the value of U at the local minimum corresponding to this new point. Call it U_2 .
5. If $U_2 < U_1$, assign the value of U_2 to U_1 .
6. If not, calculate $p = e^{-\frac{U_2 - U_1}{kT}}$ where T is the temperature of the simulation. Pick a random number on $[0,1]$. If p is larger than this number, assign the value of U_2 to U_1 . Otherwise, do nothing.
7. Return to step 3.

This algorithm has a number of tunable properties. The size of the displacement applied in state 3 has a strong effect on the efficiency of the optimization. If too small of displacements are chosen, the optimizer will not explore much space on the function. If too large of displacements are chosen, the energy correlation between states and the algorithm approaches uniform random sampling. The type of displacement has a large impact as well. Single-atom moves, multiple-atom move, short molecular dynamics simulations, etc. will sample the function in very different ways. (In fact, some [bad] methods aren't ergodic.)

The temperature also plays a role. Higher temperatures allow for greater coverage of the space, but have a danger of going to to high of energy. Lower temperatures keep good states with high probability, but sample the function very slowly. On top of this, the temperature can be changed as a function of simulation time in order to anneal the system down from melting to a frozen structure. Choice of the functional form and parameters of this cooling schedule is far from trivial.

4 Results

Figure 2 shows global optimization results for the DEN, amine, and thiol particles. Over a large range of alpha spanning 5kT in energy, the truncated octahedron structure best fits the objective function. Figure




Figure 2: Global optimization data. Left and down is better




Figure 3: Local DFT optimization data. Left and down is better

3 shows local optimization of candidates from the global optimization done with density functional theory providing the energy. The leftmost plot in the second row shows that the truncated octahedron also fits the DEN data when DFT is used to evaluate the energy. The global optimizations used the embedded atom model to evaluate energy, which did not take into account the effect of the ligands. For these, the effect of the ligands came entirely from the PDF data. Density functional theory can model these ligands. We expect that if we have the correct structure from global optimization, then adding ligands should make it match better when DFT is used. In the amine plot, this is the case; a particle with 18 ligands bound is a better fit than a particle with 6 ligands bound, which is a better fit than a particle with no ligands. This trend does not hold entirely for the thiol-bound particles. It is possible that the wrong structure is being used for the particle or that the ligands are bound at the wrong sites. Figure 4 shows that the PDF of these structures match the experimental data qualitatively as well as quantitatively. The green line is theory, and the blue line is experiment. The disordered and icosahedral structures do not match any of the data sets well. The cube octahedron and truncated octahedron can both be made to match any of the data sets with only minor surface changes (the structure remains in the same minimum on the potential energy surface). Both the cube

octahedron and the truncated octahedron are FCC structures, so they match the PDF in qualitatively the same way.

5 Transition State Simulation Environment

In addition to my above research, I also worked on the Transition State Simulation Environment (tsse). Tsse is an all-in-one structure manipulation, visualization, simulation, and methods development platform. I worked closely with one of the graduate students to rewrite large swaths of the code for clarity and correctness. I also added some methods (NVT molecular dynamics, normal modes calculation, adaptive kinetic monte carlo) and a large number of structure manipulation routines.

6 Future Work

(I've already done some work on this, but for the sake of formatting I've included the entire project here.)

Core-shell nanoparticles are an attractive option for catalysts . These nanoparticles consist of a roughly spherical core of one metal surrounded by a shell of another. These particles are already known to catalyze many other reactions such as the oxidation of carbon monoxide and the hydrogenation of alkenes . But there are potential problems with the long-term use of these particles as fuel cell catalysts. First, if the particle binds strongly to oxygen, it is possible that the catalyst will become poisoned over time, and thus rendered inactive. Second, the core-shell configuration is not necessarily the lowest free energy state of these particles; mixed configurations have higher entropy and, often, lower potential energy. If these particles are to be used as catalysts in fuel cells, it is important to know whether they will leave the core-shell configuration, and if so the timescale of the process.

I want to model the kinetic behavior of core-shell nanoparticles over the course of months or hopefully years, in order to see how long it takes for these particles to leave the core-shell configuration. This is important because the catalytic properties of the particles will change with configuration. I will run the simulation until the particle either changes conformation and leaves the core-shell state, or the simulation time reaches

a cutoff.

I intend to look at the formation of alloys instead of poisoning by oxidation. The reason for this is practical: in order to treat the oxygen in the poisoned catalyst, prohibitively expensive Density Functional Theory calculations would be needed. For systems only involving metals, far cheaper empirical potentials, such as the Embedded Atom Model, or Sutton-Chen potentials, can be used. There are many different core-shell configurations to choose from. I intend to look at a palladium shell, gold core nanoparticle of 147 atoms for a number of reasons. First, because both metals are noble the particle should bind oxygen fairly weakly. Because of this, poisoning will hopefully play a small roll in this system. Second, these particles have been synthesized in the Crooks Lab (figure 2), and have been used to catalyze certain reactions[4]. Third, there is evidence that they make good oxygen reducing catalysts , which could lead to their use in fuel cells, and thus make a simulation of their lifetime relevant. Finally, there is evidence[7] that mixing occurs in these particles. I would like to understand the mechanism for this rearrangement.

Achieving long time scales may be difficult. Were I to run simple molecular dynamics, the simulation would require on the order of 10^{22} time steps (for reference, an empirical potential might give 1000 time steps per second of computer time). Instead, I will use the adaptive kinetic Monte Carlo algorithm[5]. This algorithm treats vibrations statistically, and looks only at the transitions of the system from one potential energy minimum to another. This allows for much larger steps in time. The basic outline of the algorithm is:

1. Pick a starting state (in this case, a perfect core-shell PdAu particle).
2. Search for processes (paths from connecting potential energy minima) until some confidence probability has been reached that there are no more.
3. Pick a process to occur. The probability for selecting a transition is proportional to its rate.
4. Carry out the chosen process.
5. Increase the time by the reciprocal of the sum of the rates, where u is a random number between 0 and 1.
6. Return to step 2 and repeat.

There are two bottlenecks in this procedure: finding processes in step 2 and processes with a high rate in step 5. Finding processes is expensive, requiring about 400 energy evaluations (half a second for the potentials I plan to use). The number of processes is potentially very large, and characterizing them could take too much computational time. To make this more efficient, I intend to store processes as they are found so that known processes can be reused with very low computational cost.

When fast processes are present, the simulation time step is small. If the particle has fast, trivial transitions (such as switching of two identical atoms on the surface), the simulation will never reach the desired duration. Fortunately, a solution exists that involves combining states connected by shallow transitions into superbasins.

The combination of empirical potentials, the process database, and superbasins will hopefully speed up the simulation to the point where I can reach my large desired time. Even with these improvements, I still expect that this project will require large amounts of computational time. Completion of the project will likely require the use of the TACC supercomputers.

Although this simulation will require a large amount of computational time, the methods I will use vastly reduce the computational time that simulating kinetics on this time scale would otherwise require using standard molecular dynamics. These simulations will reveal whether a PdAu nanoparticle will become an alloy over long timescales. This could potentially help to determine whether the catalytic utility of core-shell nanoparticles is limited by their instability.

7 References

1. Multi-Year Research, Development and Demonstration Plan: Planned Program Activities for 2005-2015, <http://www1.eere.energy.gov/hydrogenandfuelcells/mypp/> (accessed April 2010).
2. Zhong, C.-J.; Maye, M. *Adv. Mater.* 2001, 13, 1507-1511.
3. Maye, M.; Luo, J.; Chan, W.; Han, L.; Kariuki, N.; Naslund, R.; Engelhard, M. H.; Lin, Y.; Sze, R.; Zhong, C. *Mat. Res. Soc. Symp. Proc.* 2003, 756, FF6.2.1-6.
4. Scott, R. W. J.; Wilson, O. M.; Oh, S.-K.; Kenik, E. A.; Crooks, R. M. *J. Am. Chem. Soc.* 2004, 126, 15583-15591.

5. Xu, L.; Henkelman, G.; J. Chem. Phys. 2008, 129, 114104
6. Voter, A. F.; Los Alamos Unclassified Technical Report #LA-UR 93-3901 (1993).
7. Knecht, M. R.; Weir, M. G.; Frenkel, A. I.; Crooks, R. M.; Chem. Mater. 2008, 20, 1019-1028

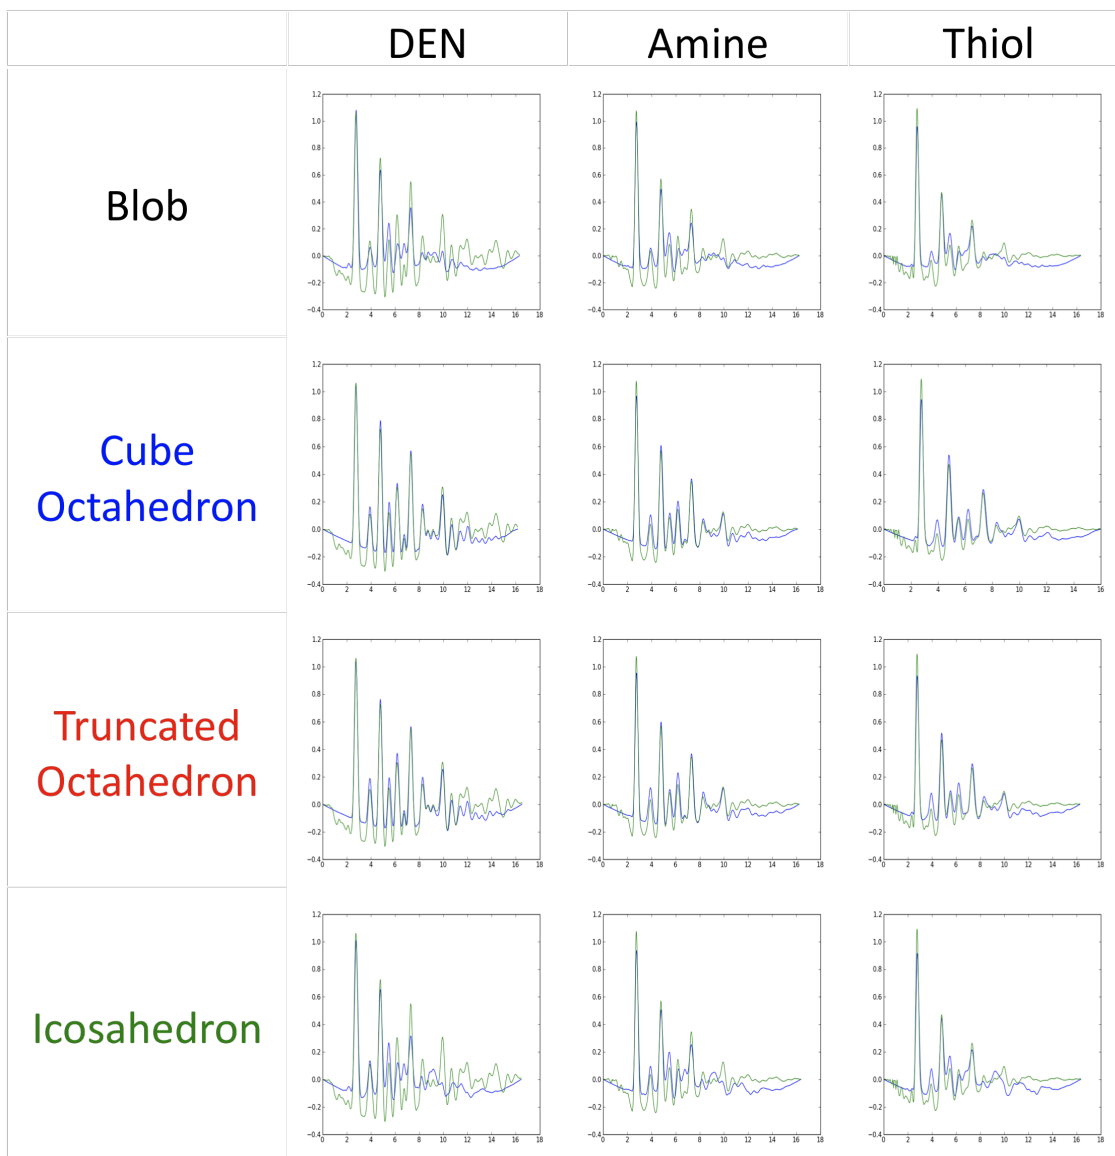


Figure 4: Qualitative PDF plots showing comparison of experiment to theory. Green is experiment. Blue is theory.