What every simulator knows / checks / tests when debugging

1. What is the expected number of particles, dimensions, and densities/concentrations based on the interaction cutoffs and system sizes? What is the minimum size?
2. What is the performance (e.g., ns per day), and how does it scale with size? Would parallelization help or hurt?
3. Is the choice of ensemble the most appropriate to the problem (or alternatively the experimental data to which I'm comparing), or merely the one I am most familiar with? What do I expect to happen in various ensembles?
4. What are the expected orders of magnitude, signs, and trends (with conditions) of various properties?
5. How do properties scale with system size? Am I at a "large scale" limit where finite size effects are minimal? Do relative fluctuations decrease as expected with N?
6. Are the results for properties in agreement with the simulation literature for the same model? For related models? With the experimental literature?
7. In the literature, what are people in my field known for (and good at)? Am I identifying established experts by verifying results, for new and emerging papers, on a regular basis?
8. What are reasonable uncertainties to expect in order to deem that things are "reproducible"?
9. What are the statistical errors in calculated properties?
10. What are the relevant correlation times? Am I averaging and saving frames at the right frequency?
11. How do I assess whether or not the simulation has reached equilibrium? What is a typical relaxation time? Is this confirmed by mean squared displacements or correlation times in various properties?
12. Are calculated equilibrium properties insensitive to the initial condition?
13. Are tail and long-range corrections correctly set up? Are the electrostatics set up appropriately?
14. Are the boundary conditions and neighbor list set up appropriately?
15. Are the results thermodynamically consistent? Does an NPT simulation give the same density as an NVT simulation with the same pressure, for example? Do fluctuations in energy, particles, volume, other things scale as expected, e.g., with temperature? Are Maxwell relations satisfied?
16. If I turn off interactions, or change conditions appropriately, do I retrieve ideal behavior (ideal gases, hard spheres, ideal solutions, etc.)?
17. As I turn off different parts of the force field, do the resulting simulated properties make sense?
18. If there is a bug, what happens for a minimalist system? What is the smallest version of the system suitable for debugging? Two particles?
19. If I run exactly the same starting configuration, random number seed, velocity distribution, etc., do I get exactly the same set of results?
20. Did I check initial and final configurations visually for overlaps and other issues/behaviors?
21. Is energy conserved if a molecular dynamics simulation without a thermostat is used?
22. What is the effect of the timestep? Do problems go away as it shrinks to zero?

23. Where exactly does the coding error occur? Are there print statements or frequent checkpoints in the code? Is the code modified so that data / checkpoints are saved at every single time step until the bug is pinpointed?
24. Do errors or warnings arise if I compile and run with full debug, verbose flags, checkbounds, etc.?
25. If an error occurs at very long simulation times, is it a precision issue or a mathematical rarity? For example, could it be acos(1.00000001)?
26. Has the trajectory been visualized? Can the exact timestep that a bug occurs be pinpointed by saving every single time step?
27. Do the properties of test configurations make sense? Do different codes give the same energy and can the energy components be decomposed?
28. Are my simulations plentiful and long enough that I should consider optimizing bottleneck routines? Would I benefit from coding certain routines in another, faster language (e.g., Fortran)?
29. Am I spending a lot of effort making my code very general, when it would be more efficient just to make a copy of it and insert the modifications I need for each new project?
30. Am I too attached to and limited by a particular code/software package/algorithm/model when something else is more appropriate? Am I wasting time with a non-optimal setup when I should just learn a new code or technique?
31. Do I lack a useful functionality in the package or code that I'm using? Should I edit its source code and recompile? Or should I simply write my own code?
32. Am I comfortable only working with certain kinds of systems – isotropic versus directional interactions, single versus multiple component, electrostatics-free versus Ewald based models, liquids versus crystals, homogeneous bulk versus interfacial, equilibrium versus nonequilibrium, etc.? How can I make time to explore outside of my comfort zone?
33. Am I being fearless, juggling several simultaneous but complementary approaches to the problem, and jumping into new techniques/approaches/codes/etc. as I need them? Or am I becoming a simulation homebody that is too married to a particular approach and reluctant to try alternatives?