Systematics of Nuclear Surface Vibrations in Deformed Nuclei

Jonathan Engel (PI)
Jun Terasaki (co-PI)

Abstract

Modern nuclear ground-state density-functional theory (DFT) works best in heavy strongly deformed nuclei — better than in spherical nuclei, which are apparently simpler. But the question of how well the theory works for *excitations* (e.g. surface vibrations) in deformed nuclei is still completely open. Because of computational demands, attempts to address the question have all proceeded by eliminating particle degrees of freedom at some point and treating the nucleus as a kind of liquid drop. We therefore propose the systematic application of the Quasiparticle Random Phase Approximation to surface vibrations in the rare-earth region of the isotopic chart, where the nuclei are strongly deformed and a lot of data exist. The results will not only tell us how well DFT works with excitations, but will also test the quality of existing density functionals and suggest ways to improve them. We request 10,000,000 Kraken SU's, and a small amount of advanced user support.

1 Proposed Research

Because this is our first proposal to Teragrid (and at the risk of overemphasizing already funded science) we give a fairly detailed description of the importance of our proposed computations, including some background. We then present information on the qualifications and funding of the investigators, a description of our methodology, an analysis of code performance and scaling, and a statement/justification of our computing request. We close with a brief request for advanced support.

Nuclear-structure theory is in a renaissance. Figure 1 below displays an isotopic chart, with the x and y directions representing the numbers of neutrons and protons, and the yellow section containing the set of known nuclei. The job of nuclear-structure theory is to reproduce the data from that section of the chart, and to predict unmeasured properties of nuclei all over the chart, particularly in the neutron-rich region near the bottom that is populated in r-process nucleosynthesis. Though the "strongly interacting many-body problem" that governs nuclear structure is difficult, computers are now powerful enough so that nuclear energy levels and wave functions can be calculated for any of these nuclei, with an accuracy that ranges from very good in light nuclei (total nucleon number $A < 20$) to reasonable and quickly improving in heavy nuclei (up to $A > 200$).

The increasing accuracy of calculations promises to improve our understanding of phenomena in both astrophysics and particle physics, as well as nuclear structure itself. The outcome of supernova explosions, for example, depends critically on the nuclear equation of state, as does the structure of neutron stars that are often the result of those explosions. The concomitant heavy-element synthesis (in the r process, for example)

Figure 1: The isotopic chart.
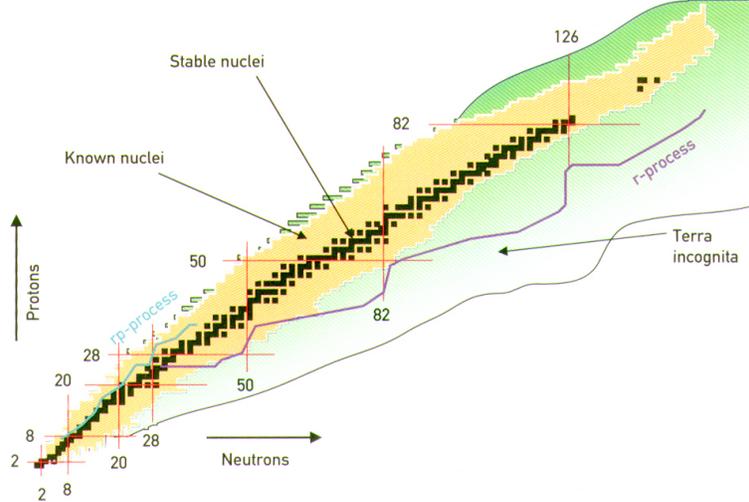
depends on the structure of heavy neutron-rich nuclei not accessible in the laboratory, and on the ability of neutrinos to excite those and other nuclei. The rate of nuclear double beta decay, from which we hope to learn about neutrinos and their masses, depends delicately on the structure of the initial nucleus and the nucleus to which it decays. Nuclear-structure theory will be better able to address all this physics in the next few years.

In heavy nuclei, the framework for improving calculations is a generalization of mean-field theory referred to as nuclear density-functional theory (DFT). Like atomic DFT, it involves the use of generalized mean-field-like equations, the solution of which can yield near exact energies, densities, and other observables if the right density functional is used in their derivation. Where such a density functional is too complicated, nuclear DFT supplements it with extensions to mean-field equations.

This approach is at the center of the SciDAC UNEDF (Universal Nuclear Energy Density Functional) project [1, 2], a 15-institution collaboration of nuclear theorists and computer scientists funded by the Department of Energy. The main goal of the collaboration is to benchmark existing density functionals, which are the result of a useful but far from perfect fitting procedure, and to develop and benchmark new ones. Our understanding of nuclear forces should soon be good enough to build new functionals from first principles, starting from the behavior of two- and three-nucleon systems and proceeding with controllable approximations. As of right now, the state of the art is a more systematic kind of phenomenology than is traditional.

Thus far, the framework has been applied mainly to ground states; the Kohn-Sham equations at the center of DFT are structured to reproduce ground-state energies and densities, and any fitting of functionals is to these properties. One result of the ground-state calculations is that the theory works better for deformed (ellipsoidally shaped) nuclei than for apparently simpler spherical nuclei. The reason appears to be that shape changes can be largely captured by deformed mean fields.

But ground-states aren't the whole story; DFT can be extended to excitations. Figure 2 below displays the web of topics under investigation in UNEDF; the oval on the lower left represents work on dynamics and excitaions rather than on ground states. The oval lists a number of methods by acroynm. Of these, the most straightforward extension of ground-state DFT is the Quasiparticle Random Phase Approximation (QRPA). The QRPA is the adiabatic (slow-motion) limit of *time-dependent* density functional theory. A density functional developed for static ground states can be applied without modification to slow oscillations if the density doesn't change too much from its ground state form — that is if oscillations are small as well as slow. In many nuclei, quadrupole surface vibrations, which change the shape of the roughly ellipsoidal nuclear boundary (i.e. the amount and nature of the deformation) seem to be good examples of slow small-amplitude oscillation.

Will the high-quality description of ground states in deformed nuclei translate to excitations? We don't really know because QRPA calculations in deformed nuclei have, until now, required too large a computational effort. In nuclei with spherical ground states that is not the case, and surface vibrations have been explored systematically [3]. The QRPA does pretty well in reproducing energies and time-dependent vibrational densities in these nuclei, even with existing non-ideal density functionals, but is far from perfect. Figure 3, from Ref. [3], shows the quality of the calculations with the functional SLy4 (the acronym indicates that the functional is of "Skyrme" type, i.e., semilocal, and that it was the fourth in a series of functionals developed in Lyon, France). The figure displays calculated electromagnetic decay rates (scaled by a phase-space factor to give a nuclear quantity called the $B(E2)$ value) versus the experimental rates (similarly scaled) for a large number of spherical nuclei. In a perfect calculation the points would lie on the diagonal. The deviation from the diagonal is a measure of the quality of the density functional and

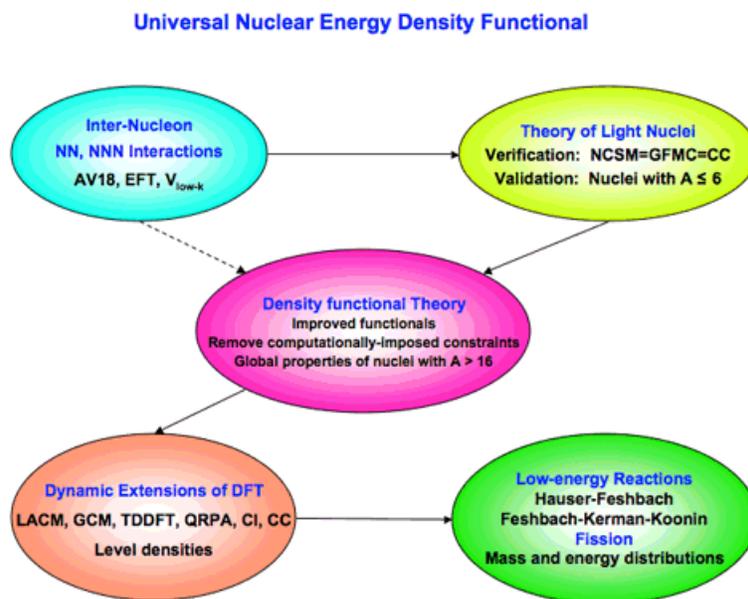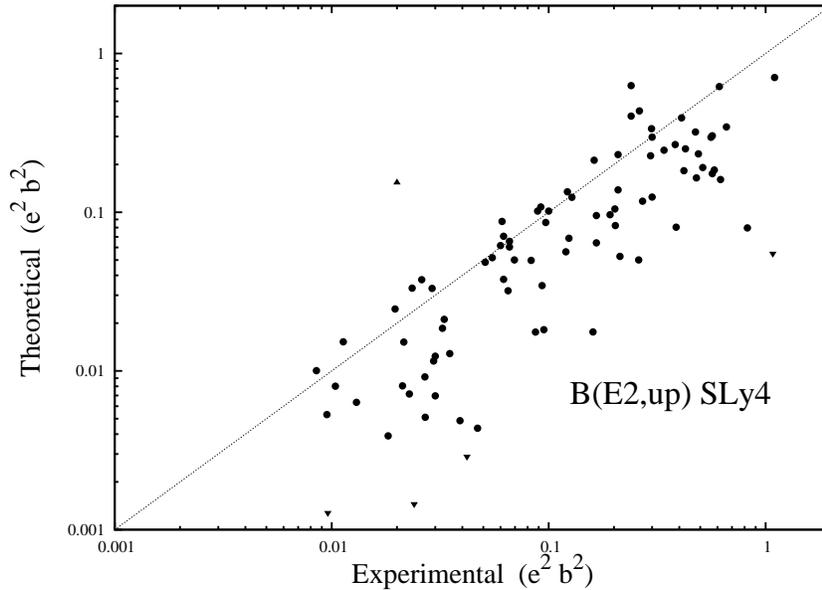Figure 2: Scope of the SciDAC UNEDF project.

Figure 3: Calculated photon emission rates (scaled) of shave-vibrational states in many nuclei vs. measured rates (also scaled), from Ref. [3].

the QRPA. We are able to get a pretty good idea of which is most at fault for deviations in particular nuclei by comparing our results to calculations in a different scheme — a method based on a marriage of the liquid drop model and "generator coordinates" [4, 5] — that is less computationally intensive but avoids the assumptions of slow oscillation and small amplitude inherent in the QRPA.

With a code we have written to solve the QRPA equations in deformed nuclei, we are now finally ready to address excitations in deformed nuclei. In view of the questions raised above, our goals for this project are

1. To test how well the QRPA reproduces energies and strengths of surface vibrations in heavy deformed nuclei, the systems and oscillations for which it appears to be best suited.

2. To uncover weaknesses in existing functionals and suggest improvements.

We propose to use our code to calculate the energies and decay rates of quadrupole surface vibrations in a large number of deformed nuclei in and around the rare-earth region of the isotopic chart — the nuclei with between 58 and 71 protons in Fig. 1 — with one or more density functionals.

We choose the rare-earth nuclei for several reasons. First, they are heavy (atomic mass around 160). There have been a few DFT-based QRPA calculations in deformed nuclei here and there (see, e.g., Ref. [6]), but none in any nearly this heavy. Second, they are strongly deformed. Third, there are data on many isotopes for each atomic charge Z. We will thus be able to explore the variation in calculated vibrational properties with neutron number, which will tell us in turn about a piece of the density functional that

is not well constrained by ground-state data. Comparing our results with those in the complementary generator-coordinate-based framework [4, 5] will allow us to assess the quality the QRPA, and the strengths and weaknesses of existing density functionals. Our results will be invaluable for the UNEDF collaboration and its attempt to develop better functionals.

The payoff of this project for physics outside nuclear structure is also potentially large. One project long anticipated by nuclear astrophysics is an accurate calculation of the β decay rates needed to understand r-process nucleosynthesis. Most of the important nuclei involved are deformed and cannot be studied experimentally (because they are so short lived). If the QRPA proves an accurate tool for surface modes in deformed systems, the prospects for its accuracy in β decay are high. We have begun the development of a β-decay version of our code, and a calculation of r-process rates is next on our list of projects.

2   Project Team Qualifications/Funding

Dr. Terasaki is a postdoc whose position is funded entirely by a five-year ($85,000/year) grant to Dr. Engel from the Department of Energy, as part of the $3,000,000/year UNEDF project. The UNEDF collaboration is beginning year four of its five-year life (though it could live longer). Dr. Terasaki has done most of the development of the present code, primarily on the NERSC Cray XT4, Franklin (as part of a UNEDF allocation), on the smaller NERSC machines Bassi and Jacqard, and on the NCCS XT5 Jaguar. He has also used a startup account on the NICS XT5 Kraken. Prior to his work on this project, he made extensive use of the now decommissioned ORNL machines Eagle and Cheetah. He has been working on the current project for more than two years, using about 1,000,000 core hours, under the supervision of Dr. Engel.

Over the last several years here at UNC, Drs. Engel and Terasaki, with collaborators, have applied DFT-based QRPA to a wide range of phenomena in spherical nuclei. In Ref. [7], we introduced our method and applied it to low-lying "pygmy" dipole resonances — so called because they are excited much less strongly than the well know "giant" resonances at higher energy — that appear only in nuclei with many more neutrons than protons. We showed that as the number of neutrons gets really large, close to the neutron drip line, the term pygmy no longer always applies because the states can get extremely excitable. In Ref. [8] we investigated the dependence on neutron number of surface vibrations, dipole resonances of neutrons against protons, and other more complicated modes in several isotope chains, and with several state-of-the-art (at the time) density functionals. In Ref. [9] we looked at the density distributions of excited states, showing that certain kinds of collective resonances are localized inside the nucleus and emit nucleons very slowly, even though they are high in the continuum. In Ref. [3] we made a systematic study of surface vibrations (again, in spherical nuclei), showing where the QRPA worked well and where it didn't, and comparing the quality of two popluar density functionals. In Ref. [10] we used charge-changing resonances, which are related to β decay, to constrain new density functionals. Dr. Terasaki was the lead author on most of these publications.

Dr. Engel and collaborators have also applied the "spherical" QRPA to astrophysics and particle physics. In Ref. [11] we calculated β-decay rates in spherical nuclei through which r-process nucleosynthesis proceeds. These rates have a large affect on the final nucleosynthetic abundances, and our results were quite different from those of schematic calculations that came before. In Refs. [12] and [13] we used the QRPA to calculate neutrino cross sections in supernova detectors; those results will help us learn about neutrino oscillation parameters in the event of a supernova in our galaxy. In Ref. [14] used a variant of the method to understand the connection between CP violation among fundamental particles and atomic electric dipole moments (the experimental limits on which now provide some of the best constraints on CP violation from outside the standard model).

We are not the only group with ambitions to apply the QRPA to heavy deformed nuclei; several groups in Europe and Japan work on the same subject. Some of them employ alternatives to the matrix form of the QRPA equations that we use here (and describe below). One alternative, involving coordinate-space Green's functions [17, 18], is well developed, but equally intensive numerically. Other new approaches [15, 16] promise to represent the matrix equation more efficiently, but are not yet fully developed and have limited accuracy. Furthermore, we are the only group to do truly systematic studies over wide sets of nuclei [3, 8], and now have considerable experience with the fastest supercomputers. We therefore believe that we are more qualified than any other group to attempt this first systematic study of vibrations in deformed heavy nuclei, and in the near term our matrix-based code is the best hope doing so.

3   Methodology

Solving the QRPA equations to high accuracy in deformed nuclei (as one must do here) in a wide range of nuclei is straightforward in principle, but difficult in practice. One must first solve Hartree-Fock-Bogoliubov (HFB) equations — the mean-field-like equations that determine the ground state around which the nucleus oscillates — with high-accuracy. (The Bogoliubov at the end of the method's name means that BSC-style pairing is included.) Next one must diagonalize a self-consistent residual effective two-body interaction in an infinite-dimensional space of "two-quasiparticle" states that make up the collective vibrations. Although the space must be truncated at some level, enough of it must be included to adequately represent the continuum of excited states in which a proton or neutron escapes the nuclear potential. Any number of approximations that make the equations easier to solve (and have therefore been used at one time or another) fail to do all this.

To carry out a QRPA calculation in a given deformed nucleus, we take three distinct steps. The second uses by far the most computing time, but we briefly describe the other as well.

1. The first step is an HFB self-consistent mean-field calculation, which involves determining an optimum set of single-quasiparticle wave functions (the analogs of single-particle orbitals in a calculation that treats pairing in mean-field theory) with

high accuracy. Since the nucleus is not spherical, but *is* axially symmetric, the wave functions depend on two continuous variables, the usual cylindrical coordinates ρ and z, with boundary conditions corresponding to rigid walls at $\rho = 20$ femtometers (fm) and $z = \pm 20$ fm. The code iterates the construction and diagonalization of a matrix representing the Hamiltonian in the basis of single-quasiparticle states until it achieves self consistency, i.e. until the mean-fields generated by the nuclear state reproduce that same state on diagonalization. The matrix elements of the Hamiltonian are double integrals of products of quasiparticle wave functions; the integrands are represented by B splines on a 40×40 mesh and the integrals done by Gaussian quadrature. The matrix is broken up into about 20 blocks with particular values of parity and angular-momentum z-projection. Each block, of typical size 2,000 by 2,000, is diagonalized by the LAPACK95 generic routine `geev` on a separate processor, and the results gathered for the construction of the next iterate of the Hamiltonian. Though the total number of core hours for this process is small, the code we are using, which we did not write, is limited in the number of processors it can employ so that the wall-clock time can be fairly long. Our advanced-support request is for help to remedy this problem.

2. The most time-consuming step, and the one that intensively uses processors, is calculation of matrix elements of the QRPA Hamiltonian from quasiparticle wave functions. Since the QRPA basis states involve *two* quasiparticles, there are many more states than in the HFB step, and vastly more matrix elements. In the 2^+ channel we treat here (the channel associated with surface quadrupole vibrations) there are no states corresponding to motion of the nuclear center of mass, and as a result we can get away with matrices of about 50,000 by 50,000 (and because of symmetry we need compute only about half of these). Each matrix element involves integrals over z and ϕ of products of four quasiparticle wave functions or wave-function derivatives. We also have to do a nominally quadruple integral (which we reduce to a series of double integrals via a multipole expansion) for each matrix element to include the effects of the Coulomb interaction among protons. As in the HFB code, we represent the wave functions with B splines, and use Gaussian quadrature for integration.

   The procedures in step 2 have been speeded up considerably in the past year. Initial versions of the code represented wave functions on a 100×100 equidistant mesh and used a form of Simpson's rule for quadrature. They also expressed wave functions in the "quasiparticle basis". By moving to the "canonical basis", obtained by diagonalizing the one-body HFB density matrix (see Ref. [11] for a brief discussion of these bases), we were able to reduce the number of integrals by about a factor of four. When assigning matrix elements to processors, we now use block cyclic distribution [19] for load balancing. And we have spent a lot of time finding optimum values for parameters such as the number of mesh points, the number of quasiparticle wave functions included, the number of multipoles included in the Coulomb expansion, etc.

3. The third step is the diagonalization of the QRPA Hamiltonian constructed in step

two. The matrix is transformed to an equivalent real symmetric matrix and diago-nalized by the parallel ScaLAPACK [19] routines `pdsyevx` and/or `pdsygvx`.

4 Code Performance and Scaling

The exploration of performance and scaling presented here was carried out on the Cray XT5 "Jaguar", to which we had brief access through an INCITE nuclear-structure-theory collaboration. We were not on the original grant proposal, however, and no longer have that access. Our tests for steps 1 and 3 are on the nucleus ^{172}Yb, and the scaling analysis of the really time-consuming part of the calculation in step 2 is on subsets of the QRPA matrix elements discussed above. All the codes referred to are written in Fortran 90.

Step 1 above makes use of a preexisting code from another source that allocates a fixed number of processors to each job and does not scale well with large numbers of cores. We are currently modifying the code to circumvent this problem, (and requesting advanced support for our efforts) but in any event step 1 consumes a small part (less than 3%) of the total time.

Step 2, the evaluation of QRPA matrix elements, is carried out as follows for a given nucleus: A small fraction of the cores (1/6th for large numbers of cores) store the quasi-particle wave functions needed for the calculation; each such core stores 1/4th of the wave functions and communicates with about 20 other cores. Each core, including those storing wave functions, is then assigned an equal number of matrix elements to calculate and is given a list containing the location (in the set of storage cores with which it communicates) of all the wave functions. A core then asks for the four wave functions it needs each time it comes to a new matrix element on its list of such, then returns the matrix element when each evaluation is complete.

Figure 4 shows how the procedure scales. It plots the time needed to calculate 1,000 matrix elements per core versus the number of cores. The first point actually corresponds to 6 cores. If scaling from 6 cores were perfect, the line would be flat. The scaling here, while not quite perfect, is very good.

For step 3 (the ScaLAPACK matrix diagonalization), we find poor scaling after about 500 processors. We will not use more than that for this step, but note that diagonalization consumes a tiny fraction (about 0.1%) of the total QRPA computing time.

5 Computing Request

5.1 Current Resources

As part of the UNEDF project we have about 350,000 core-hours for 2010 on the NERSC machine Franklin. Here at UNC, we have regular-priority access to a 4,000-core cluster used by researchers across campus. We can at most, though, use 512 cores simultaneously.

Figure 4: Scaling in calculation of QRPA matrix elements for 2^+ sates (step 2) in the heavy nucleus ^{172}Yb. The plot displays the time needed to calculate 1000 matrix elements per core vs. the number of cores used.
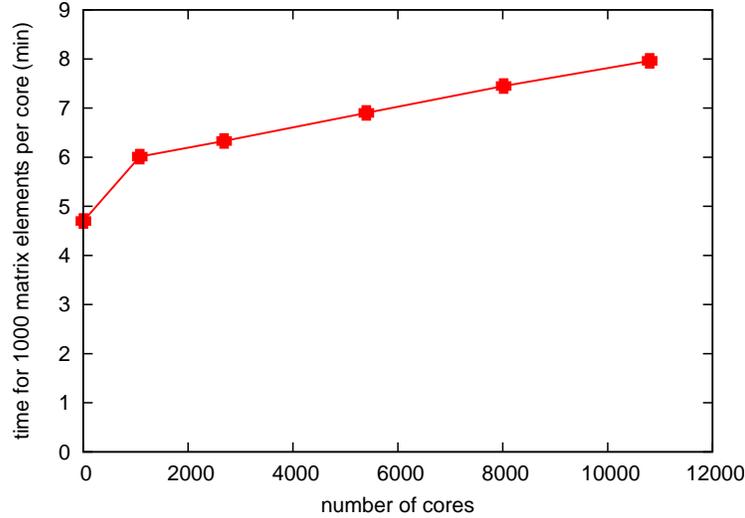
5.2   Justification

To estimate the amount of computation time we need, we ran a complete calculation in the 2^+ channel of ^{172}Yb on "Jaguar". The time consumed:

Step 1: 40 cores $\times$ 160 hours = 6,400 SU.

Step 2: 10,800 cores/job $\times$ 0.75 hours $\times$ 16 batch jobs = 129,600 SU. This is consistent with Fig. 4, which indicates about half a second per matrix element per core for large numbers of cores (with boundary conditions corresponding to a 20 fm cylindrical box). A nucleus like the one treated here requires a symmetric matrix of size about $50,000$ by $50,000$.

Step 3: 144 cores $\times$ 1.5 hours = 216 SU.

The total time was 6,400 + 129,600 + 216 = 136,216 SU. As already noted, step 2 takes almost all of that time.

HFB calculations [20] indicate that there are 68 even-even nuclei having significant deformation in and around the rare-earth region of the isotopic chart. There are data on surface vibrations in at least 32 of these. We would like to calculate energies and transition rates for quadrupole vibrations in all the nuclei with measured transitions, with two separate energy-density functionals. To do so we will need about 9,000,000 SU. Adding some time for analysis of results and auxiliary calculations, we request a total of 10,000,00 Kraken SU. We would prefer to use Kraken as much as possible, because its many processors will allow us to use the computing time in a relatively short interval of real time.

9

If our request is too large, we can reduce the number of nuclei we calculate. A study of 20, while less constraining for functionals than one of 32, would still be useful and could be accomplished with about 7,000,000 SU.

6   Advanced Support Request

As mentioned above, the HFB procedure in step 1, while not using really large amounts of time, takes a long time to run (real time) because the repeated diagonalizaton of nonsymmetric matrices is not parallelized. We tried but failed, for unknown reasons, to make ScaLAPACK handle the problem through its supbroutines `pdgehrd`, `pdlahqr`, and `pzgesv`. We request the support needed to parallelize the diagonalization, either with ScaLAPACK or by other means.

Success here would have an impact on other research groups as well as our own. Groups at Vanderbilt and ORNL use similar HFB codes, also without parallel diagonalization, and could potentially save a lot of time when running with thousands of processors.

We would only require support until this relatively simple issue is resolved. We are not in a position to estimate the time necessary (it depends on the expertise of the personnel involved), but the task seems straightforward. We would contribute to the effort by familiarizing the support person with our code, running test calculations, and in general working with that person as much as is necessary.

References

[1] See `http://unedf.org` .

[2] `http://www.scidac.gov/physics/unedf.html` .

[3] J. Terasaki, J. Engel, and G.F. Bertsch, Phys. Rev. C **78** (2008) 044311 (8 pages).

[4] J.P. Delaroche, M. Girod, J. Libert, H. Goutte, S. Hilaire, S. Peru, N. Pillet, and G.F. Bertsch, arXiv:0910.2940 (2009).

[5] G.F. Bertsch, M. Girod, S. Hilaire, J.-P. Delaroche, H. Goutte, S. Pe'ru, Phys. Rev. Lett. **99** (2007) 032502

[6] Kenichi Yoshida, European Phys. J. **A** (2009) 583.

[7] J. Terasaki, J. Engel, M. Bender, J. Dobaczewski, W. Nazarewicz, and M. Stoitsov, Phys. Rev. C **71** (2005) 034310.

[8] J. Terasaki and J. Engel, Phys. Rev. C **74** (2006) 044301.

[9] J. Terasaki and J. Engel, Phys. Rev. C **76** (2007) 044320.

[10] M. Bender, J. Dobaczewski, J. Engel, and W. Nazarewicz, Phys. Rev. C 65 (2002) 054322.

[11] J. Engel, M. Bender, J. Dobaczewski, W. Nazarewicz, and R. Surman, Phys. Rev. C **60** (1999) 14302.

[12] J. Engel, G.C. McLaughlin, and C.Volpe, Phys. Rev. D **67** (2003) 013005.

[13] H. Ejiri, J. Engel, and N. Kudomi, Phys. Lett. B **530** (2002) 27.

[14] J.H. de Jesus and J. Engel, Phys. Rev. C **72** (2005) 045503.

[15] J. Toivanen, B.G. Carlsson, J. Dobaczewski, K. Mizuyama, R.R. Rodriguez-Guzman, P. Toivanen, and P. Vesely, arXiv:0912.3234 (2009).

[16] Takashi Nakatsukasa, Tsunenori Inakura, and Kazuhiro Yabana, arXiv:nucl-th/0703100 (2007).

[17] M. Matsuo, Nucl. Phys. **A**696 (2001) 371.

[18] E. Khan, et al. Phys. Rev. **C** 66 (2002 024309.

[19] L. S. Blackford et al., `http://www.netlib.org/scalapack/slug/index.html` .

[20] M.V. Stoitsov et al., Phys. Rev. C **68**, (2003) 054312 .