Chapter 6

Molecular Dynamics Simulation of Disordered Phases

> Ludwig Boltzmann, who spent much of his life studying statistical mechanics, died in 1906, by his own hand. Paul Ehrenfest, carrying on the work, died similarly in 1933. Now it is our turn to study statistical mechanics.
>
> – David L. Goodstein [132]

It is with well-deserved reservations that we now embark on a study of disordered phases of silicon with the latest version of EDIP. Although our experience with cohesive energy curves may foreshadow problems with the liquid, a metal of greater density than the solid, the theory behind EDIP certainly addresses metallic bonds and overcoordination, so we may hope for a reasonable liquid. The success of the current EDIP for defect structures and the bulk crystal indicates an exceptionally good description of sp^3 hybrid bonds, so it would seem likely that the same version would perform well for the amorphous phase, which is made up of a random network of distorted tetrahedra. These expectations can only be checked by molecular dynamics simulations, using a sufficiently large number of atoms to minimize the influence of periodic boundary conditions and long enough times to obtain accurate thermal averages. In the first section we outline some of the techniques required to perform such simulations on high performance computers.

It is customary in the literature to describe these kinds of studies as computer simulations of a real material (*e. g.* silicon), which is somewhat misleading. It is more accurate to say that one is exploring the properties of a fictitious material characterized by a particular empirical potential, which may only bear some resemblance to the real material. In this chapter, we examine disordered phases of the current EDIP material through molecular dynamics simulation, including various liquid, glassy and amorphous specimens. The purpose of these of kinds of studies is to explore the transferability of the potential and to develop large-scale simulation techniques. With such knowledge, we may hope to improve the potential and as well as our ability to perform computer experiments to the point where reliable theoretical predictions for real silicon might be possible. The chapter closes with an outlook on the future of EDIP as a transferable model for silicon condensed phases and bulk defects.

6.1 Computational Methods

The molecular dynamics simulations discussed in this chapter were performed with program written in C originally designed for simulations of inhomogeneous systems of atoms interacting via short-range forces on the Connection Machine 5, a distributed-memory, massively-parallel architecture[1]. For the purposes of developing a transferable

[1]The initial work was done at Thinking Machines Corporation (a practicum experience supported by a Computational Science Graduate Fellowship from the Office of Scientific Computing of the Department of Energy) in collaboration with B. Larson, who later ported the SW force routine to the Silicon Graphics Power Challenge. The author would also like to acknowledge discussions with N. Bernstein, E. Kaxiras, P. Tamayo and J. F. Justo that were invaluable in bringing the code to its present state.

potential (a prerequisite for meaningful large-scale simulations), smaller simulations on serial machines suffice, but the foundations of large-scale parallelism have been laid. The code has data structures and some subroutines supporting the Single-Program-Multiple-Data (SPMD) parallel programming model where each processor "owns" a subset of the particles corresponding to a particular region of physical space, and continuously keeps track of "shared" particles owned by other processors that exert forces on its particles [139, 140, 141, 142, 143]. The original goal was to efficiently load balance and optimize communication patterns on for inhomogeneous atomic systems (with large density fluctuations) where the usual method of assigning each processor an identical rectangular chunk of physical space is inefficient, but such sophistication is not necessary for simulations of bulk phases of silicon, because density fluctuations are small.

In this section, we discuss the capabilities of the code for medium-scale simulations ($N \sim 10^3 - 10^4$), paying special attention to the efficiency of force computation with EDIP. The current version of the code is optimized for various serial workstations (Sun Sparc 5 and IBM RS-6000), and also can run in serial or parallel mode on a Silicon Graphics Power Challenge with four RS-8000 processors. On RS-10000 processors, the increased sophistication for out-of-order instructions increases speed by almost a factor of two, since molecular dynamics involves significant memory indirection. We also discuss the necessary modifications for larger ($N \sim 10^5$) parallel simulations.

6.1.1 Scaling with System Size

Within the program, a "particle" is a C structure living in a global array, containing flexible attributes like its mass and position, velocity and acceleration vectors. Linear $O(N)$ scaling of the force computation with the number of particles, is achieved using Verlet neighbor lists that store the interaction topology[144]. This involves keeping an evolving list of the "neighbors" of each atom, including not only atoms which are currently within the cutoff (and exerting forces) but also atoms beyond the cutoff by less than some safety distance δr, chosen large enough that more distant atoms cannot cross the safety distance and enter the interaction range before the list is reconstructed. The overhead in memory and computational time for creating and using the Verlet list increases as the safety distance grows, but a larger safety distance allows for less frequent list updates. Thus, there is an optimal choice of δr and the update interval of m time steps (of length δt) for every simulation. For example, $\delta r = 0.03$ Å, $m = 40$ and $\delta t = 0.005$ reduced units are good choices for simulation of an EDIP liquid at $T = 2500$ K, $P = 0$.

In the current version of the code, Verlet lists are constructed using the naive $O(N^2)$ algorithm (a double loop checking every pair of atoms), which dominates the linear scaling of the (much more intensive) force computation once the system size exceeds 10^5 particles. In order to simulate larger systems, it will be necessary to introduce an $O(N)$ algorithm to build the neighbor lists. For systems with homogeneous density, a successful algorithm involves sorting particles into cells based on their physical location[140, 142, 143]. The fairly small simulations presented here, whose primary purpose is validation of the potential, do not require such sophistication, but once we are ready for large-scale production runs with a reliable model, it will be beneficial to achieve linear scaling. Considering the speed of our force routine described in the Section 6.1.3, simulations of more than 10^5 atoms should be possible on our four-processor machine.

6.1.2 Dynamics and Measurements

Time integration of Hamilton's equations of motion is accomplished with a fifth order Gear predictor-corrector algorithm[144][2]. Several statistical mechanical ensembles may be probed: Of course, the microcanonical ensemble (constant NVE) is available, in

[2] Faster, but less accurate, lower order methods with smaller memory requirements are also available: second order velocity Verlet and fourth order Gear predictor-corrector.

which the classical trajectory is followed, exactly conserving energy to within numerical error. Contact with a heat bath is simulated by rescaling of velocities to set the kinetic temperature $\mathcal{T}$ using the equipartition theorem. To minimize the artificial influence on the dynamics, rescaling is only done intermittently, typically once every 50 time steps or equivalently several times during a phonon vibrational period. In this way thermodynamic averages can be measured in the constant $NV\mathcal{T}$ ensemble can be measured, which differ from the canonical ensemble (constant NVT) averages by $O(N^{-1/2})$. Finally, the isobaric (constant NPE) and isothermal-isobaric (constant NPT) ensembles can be simulated using the Andersen piston, an extended system method in which the volume is considered an additional degree of freedom that controls isotropic expansion of the system[145][3]. The kinetic pressure is measured using the virial theorem, with the internal virial computed efficiently within the force subroutine. The program also supports simultaneous integration of tangent space trajectories for measurement of the Lyapunov spectrum using the method of Bennetin[147, 148], which can be useful in studying connections between chaotic, classical dynamics and statistical mechanics.

In order to save disk space, measurements of statistical quantities are accumulated as the program runs, thus eliminating the need to periodically store the state of the system (except once at the end for purposes of restarting the simulation). In addition to the usual thermodynamic averages (T, P, E, V) and their root-mean-square fluctuations, several properties of the EDIP potential energy function are monitored, namely the coordination number and the pair and the three-body energies, which provide information about the type of chemical bonding present in the system. Average local structure in time and space is measured with the pair correlation function and bond angle distribution. Thermal averages of these structural quantities are usually taken at

[3]This method is useful for situations in which no spontaneous spatial symmetry-breaking is expected, as in the liquid and amorphous phases. For crystal phase transitions, methods like that of Parinello and Rahman are required to allow for changes in the shape of the simulation box as well as external shear stresses[146].

500–1000 time step intervals over 10,000–50,000 steps of dynamics. For larger systems, less time is needed for good averages due to the equivalence of temporal and spatial averaging of local physics in the thermodynamic limit.

For nonequilibrium simulations, the temperature and pressure may be controlled dynamically in several ways. The external pressure is imposed with the Andersen piston, and the temperature is controlled with velocity rescaling. Each can be held constant or ramped up or down smoothly using a cosine or linear profile in time. Instead of controlling the temperature explicitly, which involves sometimes adding and sometimes removing heat (kinetic energy), it is also possible to add or remove heat at a constant rate by an appropriate choices of velocity rescalings, which permits fluctuations in kinetic temperature. Following a crystal lattice start, the persistence of crystal structure is measured with the translational order parameter,

$$\rho(\vec{k}) = \frac{1}{N} \sum_{i=1}^{N} \cos(\vec{k} \cdot \vec{r}_i), \tag{6.1}$$

where $\vec{k}$ is a reciprocal lattice vector, like $(2\pi/a)(-1,1,-1)$ for FCC or diamond, where a is the lattice constant. For solids the order parameter is of order unity, and for liquids it oscillates about zero by the usual scaling with system size, $O(N^{-1/2})$. Equilibration following a lattice start or other nonequilibrium period is achieved once the sample distribution (mean and variance) of each thermodynamic variable converges to a fixed distribution.

6.1.3 Efficient Force Computation

Force computation is the primary bottleneck in molecular dynamics, even with the simplest interaction model (a pair potential). With an efficient $O(N)$ code, force computation takes around 70% of the total time for a large scale simulation with a pair potential and can take over 90% of total time with a many-body potential. These percentages reflect a fully optimized code, which is often faster than one's naive first attempt by several orders of magnitude.

m-Loop Potentials

A many-body potential like EDIP offers a challenging computational problem. The term "many-body" (or "N-body") typically refers to a potential containing nonlinear combinations of sums over neighbors, but from a computational point of view there are varying degrees of "many-body-ness". The embedded atom potentials[101] used for metals often have the form of the EDIP pair interaction, where coordination Z_i (a pairwise sum of radial quantities) is interpreted as a background embedding electron density, which, of course, is roughly proportional to the number of neighbors. Although all neighbors contribute nonlinearly to the energy, such a potential is effectively a three-body potential as far as computation is concerned: there can be forces on a third atom k from bond (ij) due to its changing the electron density (or coordination) of the bond. Therefore, let us introduce "three-loop" as a more descriptive term for the complexity of such a potential, rather than "many-body". The Tersoff potentials are also three-loop potentials, but the three nested loops arise from explicit angular terms in the bond order: a third atom k interacts with bond (ij) based on its angle θ_{ijk} as well as its distance from the bond r_{ik}. EDIP is more sophisticated, because it is a four-loop potential. The pair interaction requires three-loops, but due to its coordination dependence, the three-body interaction requires four nested loops: a fourth atom l generates forces on triplet (ijk) by changing the coordination of the central atom i.

Nested loops greatly increase computational time; the naive algorithm for an m-loop potential is $O(N^m)$. With neighbor lists the scaling is $O(N)$, but there is a possibly large prefactor n^{m-1} multiplying the computational time, where n is the average number of neighbors per atom. Thus, the scaling of three-loop potential is $O(Nn^2)$ with neighbor lists, and naïvely, a four-loop potential like EDIP should scale like (Nn^3). However, we shall see that the increased sophistication of EDIP comes at very little computational cost. The reason is that the coordination changes only when atoms lie in the range $c < r < b$ where $df/dr \neq 0$. The number of such atoms is usually very small, often zero, in any simulation. Thus, the computational efficiency of EDIP, $O(Nn^2n_c)$, is essentially the same as a simple three-body potential like SW, where n_c is the number of neighbors with $0 < f < 1$ in the range $c < r < b$.

Neighbor Lists

The choices of book-keeping methods, loop structures and factorizations can significantly affect the overall speed. First consider the book-keeping of neighbors, the key to linear scaling. For a pair potential, each "bond" or pair of atoms is stored once in the overall neighbor list[4], and in the force subroutine each atom gets and equal and opposite force (computed once per bond) by Newton's third law. This case has been discussed by many authors, but little has been published about the computation of many-body forces. For a three-body potential like SW, it might seem reasonable to store a list of every unique triplet of atoms, but this turns out to be rather inefficient and cumbersome, for many reasons. With such an approach, neighbor lists must store atoms at twice the cutoff distance, because an interacting triplet can be as elongated as a linear chain. This means that many bond lengths stored in the neighbor list will not make any force contribution. The size of the list is also increased compared to a simple Verlet pair list. This approach also poses major problems for SPMD parallel programming. As stated above, the most successful approach for short-range atomistic dynamics is to assign atoms to processors, with communication only required for interactions between atoms owned by different processors. The problem with a triplet-neighbor list is that it is more difficult to assign forces to the correct atoms and structure communication because there are multiple possibilities for triplet ownership, like own-own-own, own-own-shared, own-shared-own, shared-own-own, own-shared-shared, ... These problems grow quickly with the order of the interaction. The best solution for many-body forces

[4]These comments also apply to the cell-based algorithms, in which case the role of the Verlet list is played by the cell interaction graph[140, 142, 143].

is to store a list of *all* neighbors within the cutoff range of *each* atom. This causes the force routine to visit each interacting pair twice, each triplet three times, ..., but it turns out to be far more efficient. This is particularly so in the case of EDIP, because forces are not symmetric and depend on the environment (coordination) of only one atom at a time in each pair or triplet.

Force Algorithm

The design of an optimal loop structure and factorization of the derivatives is much more subtle than the choice of book-keeping method. Since the pair and three-body interactions depend on the coordination, a prepass over pairs of atoms is required. Because separation distances are checked in computing the coordination, it makes sense to precompute all radial quantities in the initial pair loop, which are stored in small temporary arrays (reused for each atom). The subsequent loops for many-body forces are made over these temporary arrays, thus eliminating the need to unnecessarily check if neighbors lie within the cutoff. In factoring the derivatives, the guiding principle is to push every computation into the outermost possible loop. With the optimal decomposition, the innermost (fourth) loop for three-body coordination forces only involves a single addition operation, performed only for previously identified neighbors in the range $c < r < b$. All forces from changing coordination are applied in an extra pair loop after the pair and three-body force loops. Pseudo-code for the algorithm is given in Fig. 6.1. In summary, the optimal algorithm consists of a sequence of four passes over the neighbors of each atom: (1) a pair loop to compute the coordination Z and store radial quantities sorted by interaction distance; (2) a pair loop to compute $V_2(r, Z)$ with a nested three-body loop to accumulate terms for coordination forces; (3) two nested three-body loops for $V_3(r_1, r_2, l_{12}, Z)$ with a nested four-body loop for coordination terms; (4) a pair loop to apply coordination forces. We shall see that this algorithm handles the sophisticated 4-loop EDIP forces with remarkable efficiency.

```
for each (own) atom $i$
   $Z = 0$
   for each neighbor (own or shared) $j$ of $i$
      if $r_{ij} < a$
         store radial factors in $V_2$ and $dV_2/dr$ (index $k$, $nk[k] = j$ )
         if $r_{ij} < b$
            store $g$, $dg/dr$ and $\vec{r}_{ij}$ (index $m$, $nm[m] = j$)
            if $r_{ij} < c$
               $Z = Z + 1$
            else
               $Z = Z + f$
               store $df/dr$ and $\vec{r}_{ij}$ (index $n$, $nn[n] = j$)
   compute $p(Z)$ and $dp/dZ$
   for each $n$, zero a scratch array $sz[n] = 0$
   for each $k$
      accumulate energy $V_2$
      if $i < nk[k]$, apply forces $dV_2/dr$ to $i$ and $nk[k]$
      compute $dV_2/dZ$
      for each $n$
         $sz[n] = sz[n] + dV_2/dZ$
   compute $w(Z)$, $dw/dZ$, $\tau(Z)$ and $d\tau/dZ$
   for each $m_1$
      for each $m_2$
         compute cosine $l_{12} = \vec{r}_{i1} \cdot \vec{r}_{i2}$
         accumulate energy $V_3$
         apply forces $(dg_1/dr)g_2h_{12}$ to $i$ and $nm[m_1]$
         apply forces $g_1(dg_2/dr)h_{12}$ to $i$ and $nm[m_2]$
         apply forces $g_1g_2(dh_{12}/dl)$ to $i$, $nm[m_1]$ and $nm[m_2]$
         compute $dV_3/dZ = g_1g_2dh/dZ$
         for each $n$
            $sz[n] = sz[n] + dV_3/dZ$
   for each $n$
      apply forces $sz[n] \cdot (df/dr)$ to $i$ and $nn[n]$
```

Figure 6.1: Outline of an efficient algorithm to compute EDIP many-body environment-dependent forces. Indentation specifies the scope of a loop or conditional statement. Interpret "apply forces" throughout as "apply equal and opposite forces using Newton's third law". "Own" and "shared" refer to the SPMD parallel programming model.

6.1.4 Benchmarks

In order to test the performance of our code, consider a variety of representative simulations using the SW potential on a single 90 MHz Silicon Graphics RS-8000 processor with a 4 MB cache. The results are displayed in Table 6.1. For small simulations involving a few thousand atoms or less, the code scales like $O(N)$ on a serial machine, indicating dominance of the total time by the $O(N)$ force computation. The benchmarks around 50 μs per atom per time step show that the dynamics of a few thousand atoms can be followed for about a million time steps in just over a day of CPU time on a typical workstation. In physical terms, this means that we can simulate a $35 \times 35 \times 35$ Å^3 chunk of silicon for over 50 ns without needing a supercomputer. With this capability, we can easily investigate a number of interesting systems of experimental relevance, such as solid phase epitaxial growth, radiation damage and plastic deformation. The factor limiting the feasibility of such simulations, however, is not computational expense, but rather transferability of the empirical potential (which is insufficient with the SW potential for the aforementioned applications).

Much larger simulations are also possible with our force routine, but some work is needed to prepare the rest of the code for large-scale parallelism. The force subroutine possesses reasonable parallel efficiency once the number of atoms exceeds a few thousand, meaning that the overall scaling is $O(N/P)$, where P is the number of processors, for small-scale, shared-memory parallelism. Unfortunately, once the system gets that large, the poor $O(N^2)$ scaling of the serial neighbor list algorithm begins to dominate, diminishing any advantage gained by the efficient, parallel force routine. Once N exceeds 10^4, the simulation speed is already reduced by a factor of two, over $N = 10^3$, and speed decreases as $1/N^2$ for larger systems. These problems can be eliminated by implementing an $O(N)$ cell-based neighbor list algorithm, which would make possible simulations of up to 10^5 atoms on the Power Challenge.

Since EDIP is considerably more complex than the SW potential, and certainly more

Phase	N	P	Total Time	Force Time	Forces	Lists	Integration
Solid	216	1	56	53	94%	2%	3%
Solid	1,728	1	62	52	84%	13%	3%
Solid	1,728	4	55	31	56%	33%	11%
Solid	12,288	4	138	28	20%	75%	5%
Liquid	1,728	1	125	166	75%	21%	4%
Liquid	1,728	4	83	59	70%	22%	8%

Table 6.1: Timing analysis of our molecular dynamics program. The effects of thermodynamic phase (equilibrium liquid at $T = 2500$ K and solid at $T = 300$ K), system size (N particles) and small-scale parallelism ($P = 1$ or $P = 4$ processors) are demonstrated for systems in the microcanonical ensemble interacting via the SW potential on a Silicon Graphics RS-8000 Power Challenge. The total simulation time and force time are in μs per particle per time step, and percentages of the total time are given for force computation, neighbor list construction ($m = 100$ for solid, $m = 50$ for liquid, $\delta r = 0.03$ Å) and time integration (velocity Verlet scheme).

Phase	LJ	SW	IK	EDIP
Solid	49	55	46	54
Liquid	43	135	130	98

Table 6.2: Comparison of the speed of force computation with the Lennard-Jones (LJ), Stillinger-Weber (SW), Ismail-Kaxiras (IK) and EDIP potentials for typical solids and liquids (in μs per atom per time step on a Silicon Graphics RS-8000 processor).

than a pair potential like Lennard-Jones (LJ), we may worry that some of benefits of large-scale simulation just described might be lost with EDIP, but that is not the case at all. As shown in Table 6.2, the optimized EDIP force routine described earlier outperforms both the SW and IK potentials in a fair comparison of equivalent simulations. The solids mentioned in the table (for SW, IK and EDIP) are diamond structures at $T = 300$ K (velocity rescaling) and $P = 0$ (Andersen piston), and the liquids are at $T = 2500$ K and $P = 0$. There are two reasons for the impressive speed. The first is that the many-body coordination forces do not arise often and are handled very efficiently when they do. The second is that the fitted EDIP has a smaller cutoff (3.45 Å) than SW (3.77 Å) and IK (3.73 Å). The result is that many fewer interactions need be considered, especially in the liquid, since the SW and IK cutoffs lie just short of the second neighbor distance (3.84 Å) in the solid. With a many-loop potential, speed is very sensitive to the number of neighbors.

The effect of the short cutoff of EDIP is seen more dramatically in comparison with the LJ potential. Of course a 2-loop pair potential is much faster to compute that a 4-loop many-body potential, but we are lucky that, although silicon bonding is more complex, it involves many fewer neighbors than the Van-der-Waals bonding described by the LJ potential. For a fair comparison with the silicon potentials, we must consider representative simulations of noble elements with the LJ potential. The LJ solid mentioned in Table 6.2 is argon in the FCC structure, using the commonly used cutoff of 2.5σ for the potential at room temperature and zero pressure. The liquid is in equilibrium just above the melting point ($T^* = 1, P^* = 0$). Remarkably, in such a fair comparison of typical simulations, EDIP competes with LJ in computational speed. For the solids, EDIP is only slightly slower than LJ (as are SW and IK), and for liquids it is slower by a factor of two (while SW and IK are three-times slower). In conclusion, the EDIP functional form offers greatly increased sophistication without paying any price in computational speed. In applications for which the model is reliably validated (*e.g.*

dislocations), very large and realistic simulations of silicon are possible. For example, if our EDIP force subroutine were incorporated into an existing massively parallel code like SPaSM [142], then simulations involving 10^8 atoms would be possible.

6.2 The Liquid Phase

Using the computational machinery of the previous section, it is straightforward to study disordered phases of EDIP. In the following sections we discuss computer experiments on specimens of EDIP material to test whether current fitted potential behaves anything like real silicon when melted and quenched. These phases are quite far from the fitting database, but lie within the theoretically predicted range of validity of the functional form. We begin with the liquid phase, obtained by melting the solid.

6.2.1 Crystal Melting

Experimentally, silicon melts at $T_m = 1685$ K, undergoing a phase transition with a latent heat of 50.7 kJ/mol from a tetrahedral, semiconducting solid to a metallic liquid with a 10% smaller volume and just over six-fold coordination [149, 150]. The detailed structure of the liquid, revealed by experiments and *ab initio* calculations [151], is described in the next section, and here we investigate thermodynamic properties associated with the phase transition. Although the structure of liquid silicon is rather difficult to predict with empirical potentials, most existing models have failed to predict even the thermodynamic properties. The Tersoff potentials (T2 and T3) predict melting temperatures almost twice the experimental value[5]. The BH potential also overpredicts the melting temperature at 3000K. The only models that describe the melting transition are SW and MFF, but it is these are also the only ones that were explicitly fit to the melting temperature. The density increase from solid to liquid is underestimated by

[5] This can be improved by changing the cutoff distance [40], but it is not clear that other important properties can be preserved.

SW by a factor of two [135]. In their original paper, SW report a melting temperature of 2013 K [13], but subsequent studies find values in the range 1665–1750 K (although if the SW potential is scaled to set the correct bulk cohesive energy, then the melting temperature should scale into the range 1775–1865 K [19]). The discrepancies among the calculated melting temperatures depend partly on varied simulation techniques, but mainly on available computational resources, which restricted early studies to rather small systems (64–216 atoms) and short times (0.1–10 ps). The MFF potential, which has not been thoroughly tested overall, is reported to predict a melting temperature of 2050 K (although their simulations of a 64 atom periodic cell are probably not very accurate) [55].

In our study of the melting transition with EDIP, we are able to achieve greater accuracy using recent advances in high performance computation, as described in the previous section. Following the method of Luedtke and Landman [134], we melt the crystal at zero pressure by slowly adding heat at a constant rate. In this way we can determine the heat capacity of the liquid and solid states and the latent heat of the phase transition, as well as the melting temperature. Starting with perfect, 1728-atom diamond crystal at $T = 300$ K, we add 1.1 eV/atom over 16.5 ns (215,000 time steps) of dynamics. The heat flux of 38,600 eV/atom·sec is, of course, much faster than experiment, but it is three orders of magnitude slower than in early studies (and the system is ten times larger). We could achieve even slower rates and larger systems on parallel supercomputers, but the discontinuities associated with the phase transition are seen quite clearly with these modest simulations, which can run overnight on a serial Silicon Graphics RS-8000 processor. The first order transition takes around 6.6 ns with these parameters, including superheating of the solid, melting, and recovery from supercooling of the new liquid.

The results of the simulation are shown in Fig 6.2. Because heat is added at a constant rate, the total energy is proportional to time (from the first law of thermodynamics,

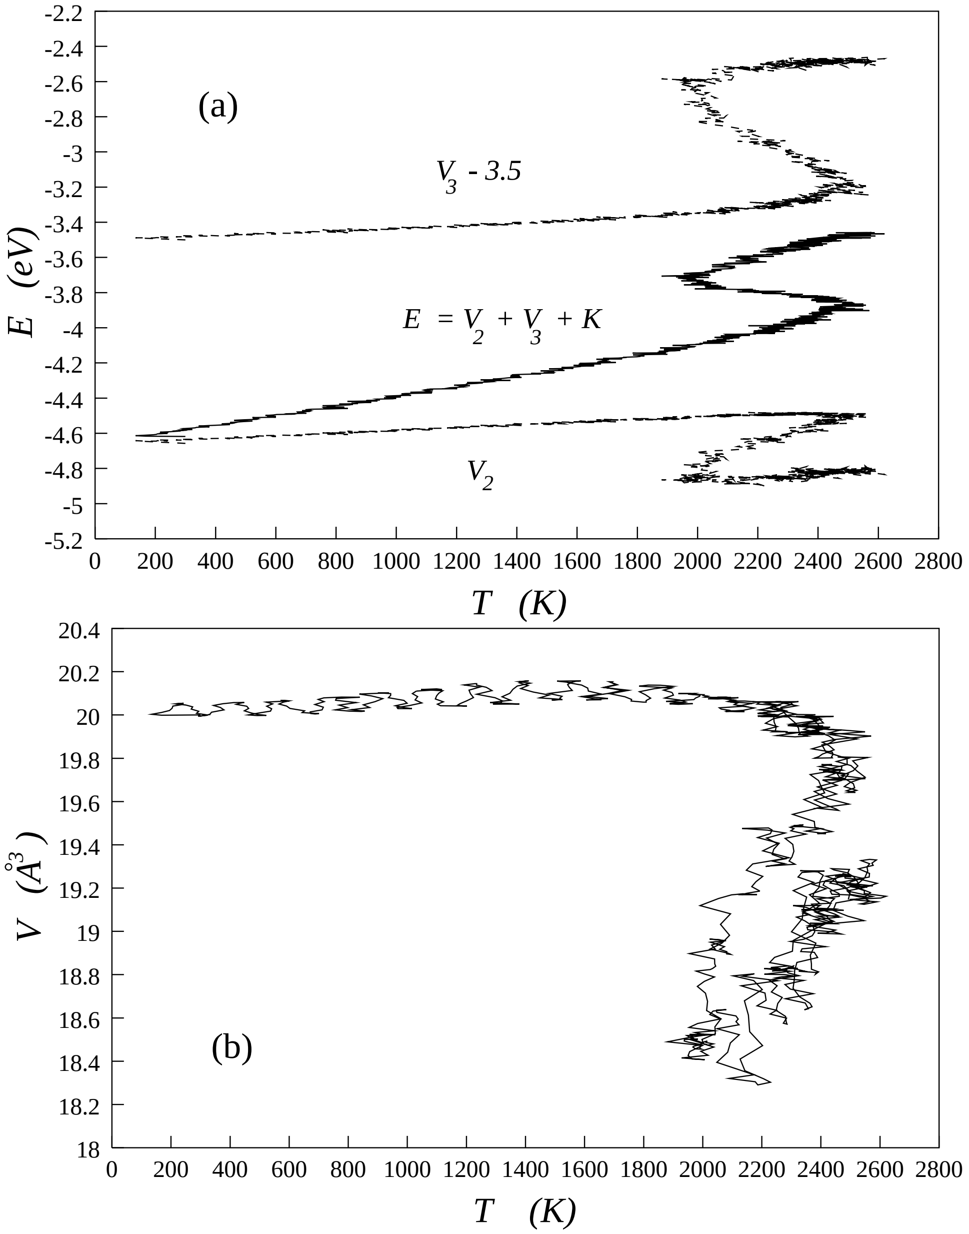

Figure 6.2: Melting of a 1728-atom EDIP solid with a constant heat flux of 38.6 eV/atom·ns. The total (E), pair (V_2) and three-body (V_3) energies as a function of temperature are shown in (a), and the volume per atom versus temperature is shown in (b).

$dE = dQ$ since $P = 0$), but the temperature fluctuates. In a unique phase (solid or liquid), the temperature rises at a roughly constant rate, but during the phase transition the temperature temporarily lowers as the potential energy rises. In the EDIP model, the pair energy actually reduces from liquid to solid (due to the increasing coordination), but the three-body energy increases to overcompensate. The heat capacity (slope of the E versus T curve) has a small temperature dependence away from the transition, with $C_p = 25$ kJ/mol·K for a solid at $T = 1000$ K, in good agreement with experiment (like the other potentials [126]). The heat capacity of the liquid $C_p = 36$kJ/mol·K at $T = 2500$ K is quite close to the experimental value 31 kJ/mol·K. (MFF predicts 28 kJ/mol·K [55].)

The unphysical temperature variation of the total energy during the phase transition reflects the finite system size and fast dynamics. The interfacial liquid-solid tension is not neglible for small nuclei, and, because there less likelihood of a nucleation event in a small, perfect crystal during a short time, the solid is easily superheated. The volume variation with temperature during the transition in Fig 6.2 (b) shows collapse into a supercooled liquid state of higher density and (large coefficient of volume expansion) before recovering to the equilibrium liquid under subsequent heating. Clearly the system is not large enough (or the melting slow enough) to allow coexistence of the two phases in equilibrium. The density increase from solid to liquid is underpredicted by EDIP at 4.4%, compared with 5.5% for SW. Coincidentally the supercooled liquid that appears during our phase transition has nearly the correct density.

From our bulk simulation, the melting temperature is between 2000 and 2500 K. It is tempting to make an equal area construction and conclude that the melting temperature is 2230 K, but such analysis is not valid in this data. A more accurate determination of the melting temperature can be made by introducing two (100) surfaces into the system by switching off periodic boundary conditions in one direction, and repeating the same melting procedure. These large defects prevent the need for small liquid nuclei

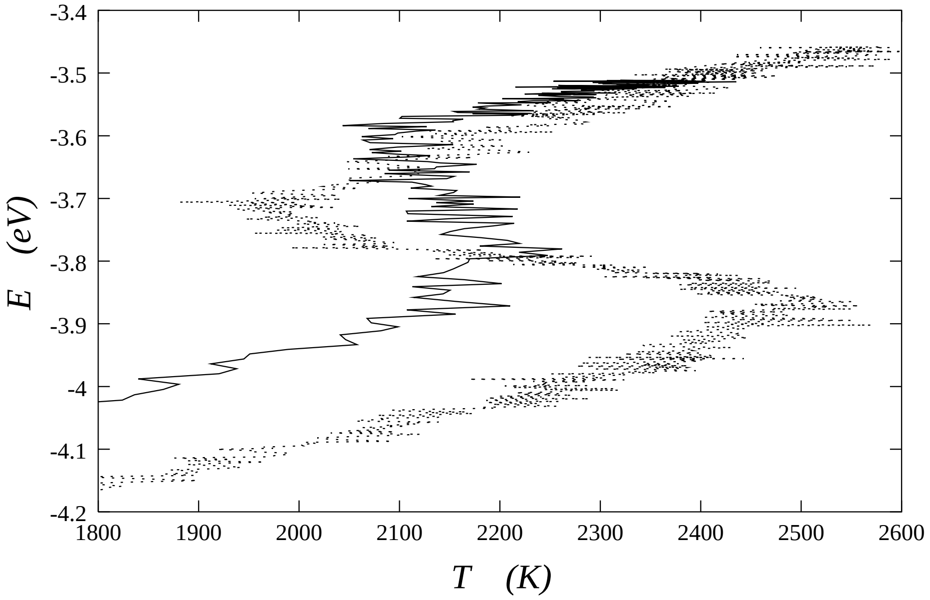

Figure 6.3: The melting transition (total energy versus temperature at constant pressure) under a constant heat flux for 1728-atom EDIP solids with (solid line) and without (dotted line) free surfaces.

growing in the bulk crystal to overcome interfacial tension, and provide numerous seeds of disorder as the surface premelts. In Fig. 6.3, we see that a fairly sharp phase transition is achieved, within the range suggested by the periodic bulk simulation, at a temperature of $T_m = 2150 \pm 50$ K. This result is in fairly good agreement with experiment, rivaled only by SW and MFF.

By measuring energy differences at the melting temperature in Figs. 6.2 and 6.3, we can obtain the latent heat of melting. EDIP predicts a latent heat of $\Delta E = 42.5$ kJ/mol (0.44 eV/atom), which is considerably closer to experiment (50.7 kJ/mol) than the SW (31.4 kJ/mol) [134] and MFF (30 kJ/mol) [55] values. The EDIP latent heat can be broken down into separate contributions of -0.38 eV/atom from the pair energy and 0.83 eV/atom for the three-body energy.

In summary, EDIP does rather well in predicting the thermodynamic properties of the melting transition. Its performance for the melting temperature, latent heat, heat capacities and volume change is comparable to SW and MFF, which were both explicitly fit to the transition, and EDIP outperforms other potentials that were not fit to the liquid. However, thermodynamic properties are only rough indicators of the realism of the model, and to better understand the physical relevance of the EDIP liquid we must look closely at its structure.

6.2.2 Liquid Structure

In order to study structural properties, the liquid is created as follows. The structure generated by this method is quite similar to that obtained by melting, but is ensured to be better thermalized, with no remnants of crystalline order. First, a perfect diamond lattice of 1728 atoms is given random initial velocities (uniform distribution for each component), rescaled to set a kinetic temperature of 5000 K. This is a very unphysical situation, an extremely super-heated solid, but it serves well to generate a random mixture of atoms with reasonable separations at nearly the right density. The system violently raises its entropy (and thus lowers its free energy) by melting, and the order parameter drops to zero in less than 50 ps. The temperature is maintained at 5000 K for 0.3 ns to thermalize the velocity distribution and to equilibrate the volume at constant pressure. The temperature is then gently ramped down to 2500 K over a period of 2.5 ns. Finally, the system is allowed to equilibrate at $T = 2500$ K and $P = 0$ for another 1.5 ns.

The average local structure of the EDIP liquid is shown Fig. 6.4. The pair correlation function $g(r)$ has a sharp peak that nearly overlaps with the first neighbor peak of the diamond solid. The first maximum is at $r = 2.38$ Å. Integrating up to the first minimum at $r = 2.84$ Å, we find that the first peak contains 4.24 neighbors. Thus, EDIP melts into a primarily covalent liquid, with four distorted sp^3 hybrid bonds per atom. In contrast, the experimentally observed liquid is metallic [149] with a coordination of 6.4 [150]. The *ab initio* liquid has coordination 6.5 up to a distance 3.10 Å [151], and a pair correlation function $g(r)$ quite similar to the SW result, shown in Fig. 6.4, although SW technically overestimates the coordination, as described below.

Although the *ab initio* liquid is metallic and overcoordinated, *ab initio* dynamical studies reveal that tetrahedral fluctuations do exist in the liquid, which helps it find the amorphous network upon cooling [151]. An analysis of *ab initio* charge densities shows that covalent bonds (charge transfer from atomic orbitals to bonding orbitals) always form between atoms closer than $r_c \approx 2.5$ Å, and larger separations break these bonds, leaving a weak metallic interaction. The EDIP liquid contains a similar mixture of covalent and metallic bonding, but the percentage of covalent bonds is overestimated. The *ab initio* liquid has roughly two first neighbors in the covalent range, $r < r_c$, while the EDIP liquid has four. The other four first neighbors in the *ab initio* liquid are in the range of metallic bonds.

EDIP also has metallic near-neighbors, but they are concentrated in a narrow, anomalous peak of $g(r)$ with a maximum at $r = 3.11$ Å in between the first and second

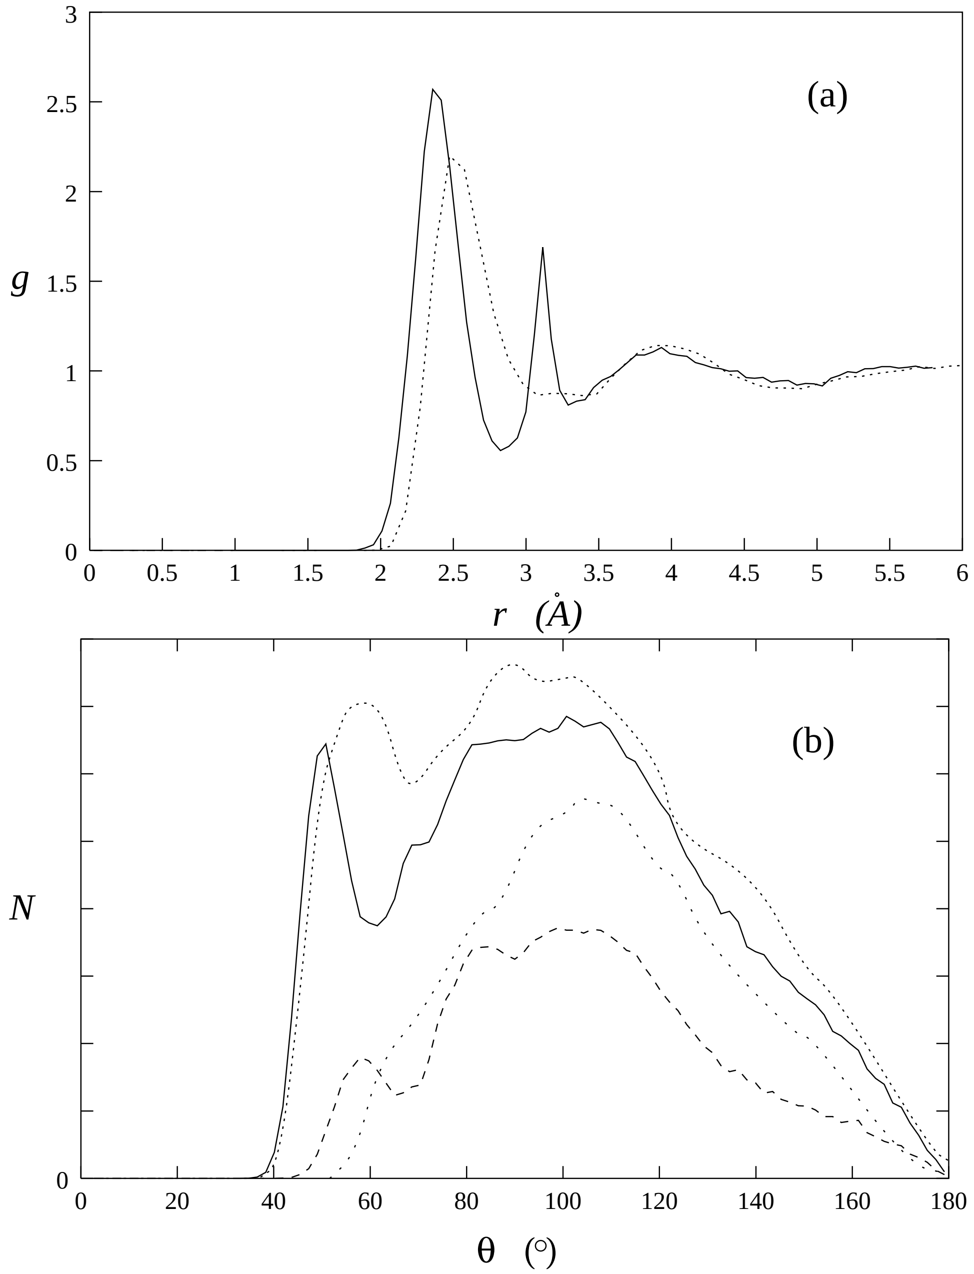

Figure 6.4: Structure of the EDIP liquid. The pair correlation (solid line) is shown in (a) and compared with the SW liquid (dotted line), which is close to *ab initio* [151]. The bond angle distribution is shown in (b) for coordination neighbors with $r < 3.31$ Å (solid line) and also $r < 2.84$ Å (dashed line), and compared with the *ab initio* distribution for $r < 3.10$ Å (dotted line) and $r < 2.50$ Å (widely spaced dotted line) [151].

neighbors in the diamond lattice. The second peak contains roughly one more neighbor. To be precise, integration up to the second minimum at $r = 3.31$ Å yields a total coordination of 5.02. Recall that we have encountered the distance $r = 3.11$ Å before; the unphysical dips in all the EDIP cohesive energy curves occur when neighbors pass through this distance and gain pair energy without paying any penalty in three-body energy due to the disparate cutoffs $b < a$. Clearly the same effect is seen in the liquid. Nevertheless, in spite of the unphysical splitting of the first neighbor peak, EDIP may be the first potential to predict a liquid with a clear mixture of covalent and metallic bonds, albeit in the wrong proportions, as evidenced by three-body correlations.

Our assessment of the EDIP liquid is refined by consideration of the bond angle distribution, shown in Fig. 6.4 (b). The atoms from first neighbor peak of $g(r)$ (dashed line) have a broad maximum around the tetrahedral angle, consistent with the sp^3 character of these bonds suggested by $g(r)$. The *ab initio* bond angle distribution (widely spaced dotted line) for atoms in the covalent range $r < r_c$ has a similar distribution. The overall *ab initio* bond angle distribution (dotted line) is broadly peaked at 90°, where EDIP first neighbors also have a small peak. If we include atoms from the anomalous second peak of $g(r)$ in the bond angle distribution (solid line), then a sharp peak develops at 50° along with a bulge at large angles, much like the *ab initio* distribution.

The structure of the EDIP liquid can be understood in terms of a simple model of the atomic arrangements. The dominant structure is a distorted tetrahedron with bond lengths close to the covalent distance 2.35 Å (first split peak of $g(r)$) and one extra atom at a metallic distance of 3.1 Å (second split peak) centered along a tetrahedral edge. The extra atom then lies at the center of a neighboring tetrahedron, with the atom at the center of the first tetrahedron playing the role of the extra atom for the second. So, the model is that of two distorted tetrahedra joined at a common edge. The four-atom ring created between the two tetrahedra is shown in Fig. 6.5. The angle formed between first neighbors at the edge is 75°, which is consistent with the skewing of the tetrahedral

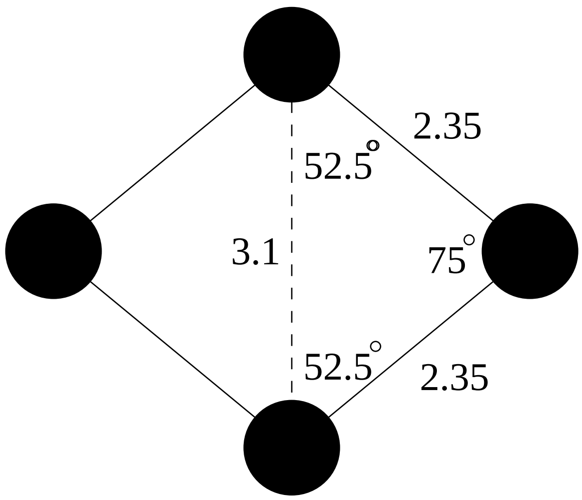

Figure 6.5: A typical structure in the EDIP liquid, which may also appear in the *ab initio* liquid, possessing a mixture of covalent and metallic bonds.

peak in the bond angle distribution toward the range 80–90° in Fig. 6.4 (b). The angle between the extra atom and the edge atom is 52.5° which explains the sharp peak just above 50° in the bond angles between the two split peaks in the $g(r)$. Similarities with the *ab initio* bond angle distribution suggest that this local structure may appear in the real silicon liquid as well.

Another sensitive measure of liquid structure is the distribution of local coordinations, defined as the number of neighbors closer than the first minimum (beyond the first peak) of $g(r)$. The numerical results are given in Table 6.3 and plotted in Fig. 6.6. The *ab initio* distribution resembles a bell curve of half-width 1.5 centered at 6.5. For a fair comparison with EDIP, we must include in the coordination neighbors from the first two peaks with $r < 3.31$ Å, which clearly represent a splitting of the first peak (into covalent and metallic subshells as discussed above). The EDIP distribution is remarkably close to *ab initio* across the entire range of local coordinations, a feat which has not been reported for any other potential. The table also gives coordinations for neighbors in the inner subshell $r < 2.83$ Å, which is peaked at 4, in agreement with the

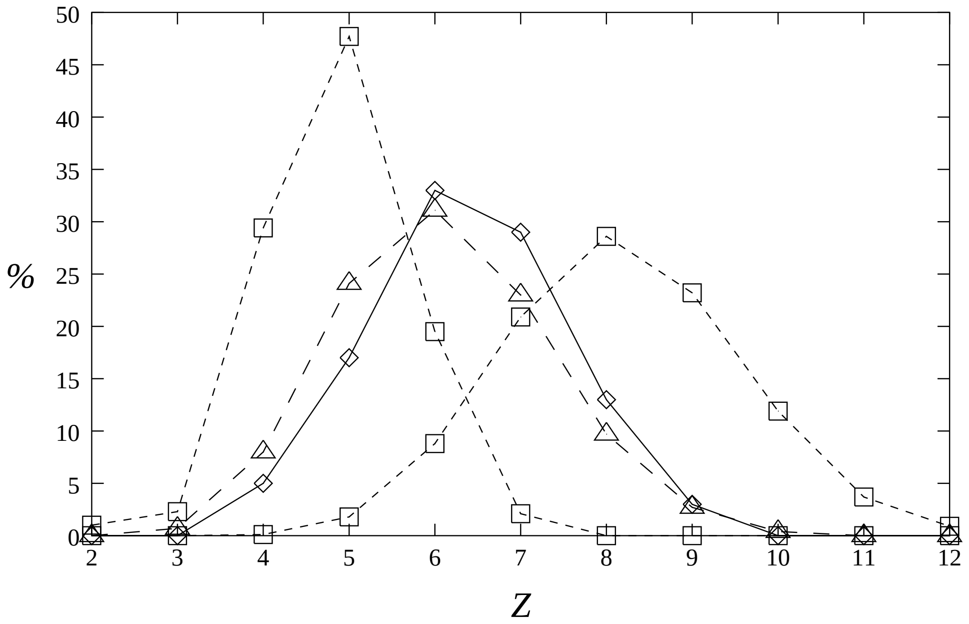

Figure 6.6: The distribution of atomic coordinations in the *ab initio* (diamonds), EDIP (triangles) and SW (squares) liquids. The higher Z SW curve is for the true coordination cutoff 3.43 Å, and the other is for a shorter cutoff of 3.0 Å [134].

Z	LDA	EDIP	SW	SW 1	EDIP 1	EDIP Z
2	0	0	0	1	0	0
3	0	1	0	2	16	5
4	5	8	0	29	48	36
5	17	24	2	48	30	45
6	33	31	9	20	5	13
7	29	23	21	2	0	1
8	13	10	29	0	0	0
9	3	3	23	0	0	0
10	0	0	12	0	0	0
11	0	0	4	0	0	0
12	0	0	1	0	0	0

Table 6.3: Distribution of local coordinations (in %) for the *ab initio* (LDA) [151], SW ($r < 3.43$ Å) and EDIP ($r < 3.31$ Å) liquids at 2000 K. For EDIP, separate statistics are given for neighbors under the inner split first-neighbor peak of $g(r)$ (EDIP 1), and for SW we also show published data [134] (SW 1) presumably using a cutoff around 3.0 Å. The distribution of effective coordination numbers (EDIP Z) is also given, rounded to the nearest integer.

ab initio analysis of atoms inside the covalent cutoff r_c [151]. Once again, we see the coexistence of covalent and metallic bonds in both the EDIP and *ab initio* liquids.

The best known description before EDIP is by SW, but as seen in Fig. 6.6 it is not very accurate. However, the comparison depends on the definition of coordination. There is some discrepancy in the literature regarding the coordination of the SW liquid, which we may attribute to the fairly flat shoulder in $g(r)$ between the first and second neighbor peaks. The inner edge of the shoulder is at 3.0 Å, and our calculations suggest that Broughton and Li must have used this distance in getting their coordination statistics (SW 1 column in Table 6.3), yielding an average coordination of 4.97 [133]. Their (inner) coordination distribution for SW is sharply peaked at 5, a fact attributed by Kaxiras and Boyer to the low energy of the BCT5 phase predicted by SW [106]. The first minimum of $g(r)$, however, is at the other end of the shoulder at 3.43 Å, as seen clearly in the data of Broughton and Li [133]. This is the more appropriate cutoff for the coordination, yielding an average of 8.16 with our data, which is consistent with other studies reporting values of 7.7 [135] and 8.0 [13]. The full distribution for the SW liquid greatly overestimates coordinations, as seen in the figure.

Table 6.3 also shows the distribution of effective coordination numbers in the EDIP liquid, defined in Chapter 5.1. The true coordinations are severely underestimated by the EDIP coordination number ($\overline{Z} = 4.23$). We have already pointed out the underestimation of coordination when studying cohesive energy curves, but here the situation is worse. Incorrect effective coordination numbers undermine the theoretical basis of the potential, and might lead to problems. This has occured because the fitting database of bulk defect structures has very few neighbors in the range $3.0 < r < 3.2$ Å, which are plentiful in the liquid.

6.2.3 Discussion

In summary, the EDIP liquid has a number of reasonable features, even though the pair correlation function is clearly artificial. In particular, the bond angle distributions for the covalent and metallic bonds of the EDIP liquid are quite similar to their *ab initio* counterparts, a rather subtle effect. In contrast, the SW liquid bond angle distribution does not contain angles with metallic character at any distance, having simply a broad peak at the tetrahedral angle for all first neighbors [133, 134]. EDIP also captures the statistics of local coordinations much better than SW. It is often said that SW offers an excellent a description of the liquid state, indeed the best of any of the popular potentials. While SW does predict an accurate melting temperature and pair correlation function, it is worth emphasizing that upon closer inspection, the local structure of the SW liquid is wrong, as indicated by the coordination statistics and bond angle distribution. The behavior of EDIP is encouraging because it was not fit to the liquid (as was SW) or any other overcoordinated structure. The performance of EDIP for the liquid is purely an extrapolation from crystalline bonding states, made with the theoretically motivated coordination dependence of the functional form.

So, perhaps the EDIP liquid has some new and relevant physics not contained in existing models, which blindly try to enforce tetrahedral order in every situation. Although it has numerous flaws, EDIP is at least capable of modeling multiple bonding states, which opens the possibility of a new degree of transferability. It is interesting to note that if the first two peaks of $g(r)$ were merged at the weighted average distance of $r = 2.5$ Å (coincidentally, the location of the *ab initio* peak), then the structure of the EDIP liquid would be fairly realistic. Perhaps, if the abrupt rise in the three-body radial function we noted earlier were smoothed by increasing the fitting parameter γ, the liquid structure might be improved.

6.3 Amorphous Phases

Most potentials predict a quench from the liquid into a glassy phase characterized by frozen-in liquid structure. Real (and *ab initio*) silicon, on the other hand, quenches into an amorphous phase consisting of a random tetrahedral network of distorted sp^3 covalent bonds. It has been proposed that the seeds of crystalline order are created from the fluctuating covalent bonds within the first neighbor shell of the liquid described above ($r < r_c$) [151]. Using empirical potentials like SW [133, 134, 135] and BH [152], it has always been necessary to make artificial changes, usually strengthening the three-body interaction, in order to guide the system into an amorphous phase when quenching from the liquid. In these simple models, there is only one physical principle at work: the balance between the three-body forces that favor tetrahedral angles and thermal fluctuations that tend to break them. The angular preference must be weak enough to be overcome by thermal vibrations at the melting temperature, but apparently this level of angular force is not enough to overcome the free energy barrier to recover the tetrahedral structure upon cooling. The amorphous is regained by simply adjusting this balance, controlled by a single parameter, the strength of the three-body interaction.

The EDIP interaction model is more complicated, and hence has a more complex phase diagram. This make it possible to describe more subtle physical properties like the bond angle and coordination statistics in the liquid and the structure of crystalline defects, but it opens more ways for the potential to bypass tetrahedral bonding at low temperatures. Unlike other potentials, EDIP quenches into an amorphous phase that is qualitatively very different from the liquid, containing quasi-crystalline short-range order. Unfortunately, the local order is not tetrahedral.

6.3.1 The Quenched Liquid

The quench is performed by gently ramping the temperature of a well-equilibrated 1728-atom liquid down from 2500 K to 300 K in 50 ps at zero pressure, resulting in an amorphous structure we may call a-EDIP-I. The liquid-amorphous transition is second order (or perhaps even a sequence of second order transitions), with a discontinuity in heat capacity $(dE/dT)_p$ and coefficient of volume expansion $(dV/dT)_p$ around $T_a = 670 \pm 30$ K. The energy at $T = 300$ K and $P = 0$ of the EDIP amorphous phase is -4.497 eV/atom, only 0.16 eV/atom higher than the ground state crystal and lower than any other crystal phase from Chapter 5.2.3. The EDIP amorphous volume is quite low, only 16.46 Å^3/atom. The corresponding density, 0.0608 Å^{-3}, is 22% larger than the equilibrium diamond solid, while experimentally amorphous silicon has a 1% *smaller* density than the crystal. This surprising result foreshadows highly unphysical properties of the EDIP amorphous phase, which we may expect to be related to overcoordination and metallicity.

With the exception of a tall and narrow second peak[6], the pair correlation function for the EDIP amorphous is fairly close to the *ab initio* amorphous [153] as shown in Fig. 6.7. The first peaks are almost identical, with a maximum just above the diamond crystal bond length of 2.35 Å, but the EDIP peak is not quite as sharp and contains a few too many neighbors, 4.32 instead of 4.00. The non-glassy nature of the EDIP amorphous phase is seen in the vanishing pair correlation in the regions between the first three peaks. As in the liquid, there is an anomalous second peak at $r = 3.12$ Å, with the next minimum at $r = 3.26$ Å defining the true coordination distance. The second peak is much larger than in the liquid, containing around four more neighbors for a total coordination of 8.23. The eight-fold coordination may suggest a BCC-like random network. The positions of neighbors in a perfect BCC crystal at the same density are shown in Fig. 6.7 (a). The second neighbors perfectly overlap with the split second peak of the EDIP $g(r)$, but the populations are wrong. The first two BCC peaks contains 8 and 6 neighbors, respectively, while the EDIP peaks contain roughly four each. So, if

[6] Once again, due to the small radius of the second peak, it is really a splitting of the first neighbors into two peaks, one covalent and one metallic.

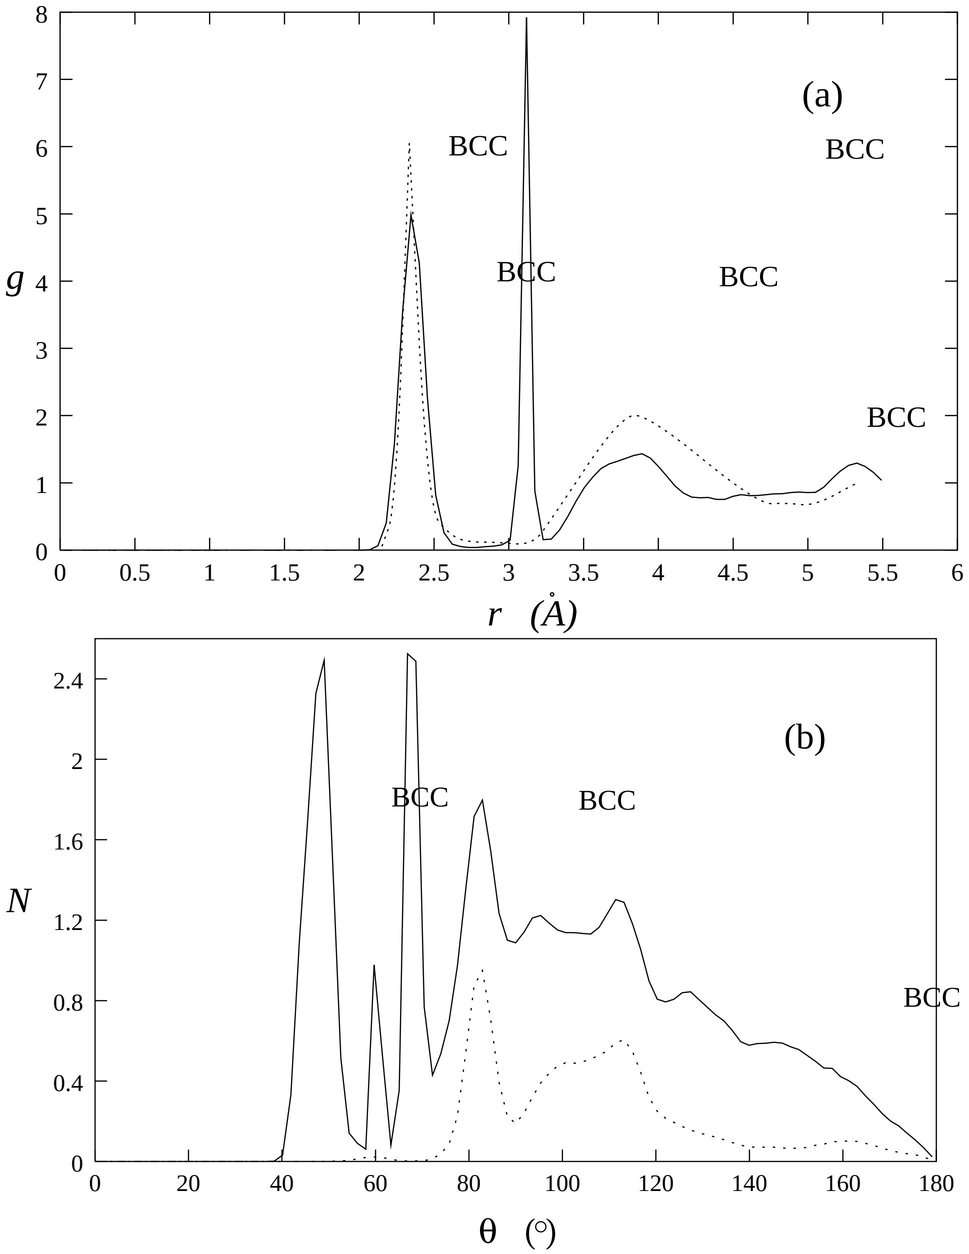

Figure 6.7: Structure of the EDIP amorphous phase a-EDIP-I obtained from quenching the liquid. In (a), the EDIP (solid line) and *ab initio* (dotted line) [153] pair correlation functions are compared, along with peaks in $g(r)$ for a BCC crystal at the same density. In (b), the EDIP bond angle distributions for $r < 2.83$ Å (dotted line) and $r < 3.26$ Å (solid line) are shown, and compared with BCC angles.

the EDIP amorphous has any BCC character, it must come from a splitting of the first peak into two four neighbor subshells.

In spite of the bizarre coordination shells, the pair correlation of the EDIP amorphous bears some resemblance to *ab initio* at larger distances. The second neighbor peak corresponding to the diamond second neighbor distance $r = 3.84$ Å is clearly seen, and the diamond third neighbor peak at $r = 4.5$ Å is absent. These physically correct features are also predicted by SW, but only EDIP predicts the diamond fourth neighbor peak at $r = 5.43$ Å in agreement with *ab initio* [134, 153]. However, note that the shapes and sizes of the peaks are incorrect, and EDIP predicts a flat plateau in $g(r)$ for $4.2 < r < 5.1$ Å.

In the liquid, the bond angles were quite well described in spite of the anomalous first peak splitting of $g(r)$, but bond angles reveal the EDIP quenched amorphous to have a fascinating, but completely unphysical structure. The *ab initio* bond angle distribution for a-Si is simply a bell curve of half-width 20° centered at the tetrahedral angle (with a tiny extra peak at 60°) [153]. The EDIP amorphous, on the other hand, has a remarkable degree of local orientational order, characteristic of a metallic crystal, as seen in Fig. 6.7 (b). As in the liquid case, the splitting of the first peak of $g(r)$ contains qualitatively different physics in the bond angle distributions as well. The bond angles for neighbors in the first peak of $g(r)$ with $r < 2.76$ Å, have an asymmetric peak at the tetrahedral angle, with a sharp decay on the large angle side, and a shoulder at 100° on the small angle side, indicating some sp^3 covalent bonding in the inner subshell, as in the liquid. There are also, however, a much larger and narrower peak at 83° and a small peak at at 160°, showing that the tetrahedra are systematically distorted to fit into a larger structural unit. The local coordinations in the first peak are 59% four-fold and 36% five-fold. The local structure of the five-fold coordinated inner subshell atoms may be similar to BCT5, which involves angles 86°, 106° and 148° (the first two are consistent with the data).

The full bond angle distribution, including the second subshell, is quite complicated. The distribution involves very sharp and tall peaks at small angles (48°, 60°, 68°, 83°) characteristic of metallic quasi-crystalline order. The sharp peak at 68° lies close to the BCC angle of 70.53° we mentioned earlier as an atom sitting on a tetrahedral face. The other strong small angle peak is close to the angle of 54.74° for an atom on a tetrahedral edge. The tetrahedral angle itself, also in the BCC crystal, is present in the outer subshell, and large angles also appear in the range $120 - 160°$, which are more or less absent in the first subshell. The structure does not have all the features of BCC, however, because angles near 180° are conspicuously missing. The many subsidiary peaks are probably related to averaging over a number of characteristic local structures. This is indicated by the broad distribution of local coordinations, which is peaked at 8 and 9 at 25% and 22%, respectively, with moderate numbers of atoms, around 15% each, at 7 and 10. The sharpness of the small angle peaks, however, suggests that these many different local structures contain similar building blocks. A typical structure, consistent with the data, might involve a BCC eight-atom coordination shell with four tetrahedral atoms pinched inward (first split peak of $g(r)$) and the other four moved outward (second split peak). In summary, although EDIP has not found the familiar tetrahedral structure of a-Si, it has, nevertheless, quenched into a curious random network with a different and more complicated local structure.

6.3.2 Another Amorphous Phase

The inability of the EDIP liquid to quench into the correct amorphous structure is not a major setback, since no other potential can do it either. Based on the extensive theoretical and fitting input involving sp^3-bonded structures, however, we would expect EDIP to be capable of describing a-Si if it could be coaxed into the correct structure. It seems reasonable that the complexity of the potential, which allows for strange phases like the diamond-BCC amorphous, would cause the free energy barrier between the quenched liquid and the highly-ordered (low entropy) tetrahedral network to be too large to overcome in a fast quench. A much less stringent test than quenching is to see what happens when a realistic amorphous sample is annealed with EDIP forces. Unfortunately, and perhaps surprisingly, the present version of EDIP fails this test: Not only is the correct structure dynamically unstable, but a spontaneous transition to another overcoordinated amorphous network occurs.

A realistic sample of a-Si can be generated using the bond-switching algorithm of Wooten, Winer and Weaire [155] or a molecular dynamics technique involving modified SW potentials (which generates the so-called "indirect amorphous") [133, 134]. In our tests we use a 216-atom sample created with the latter method (by N. Bernstein). The purity of the sample is improved by running at 300 K with double-strength three-body energy, and then annealing down to 100K with the unmodified SW potential. This structure is equilibrated with EDIP at 600 K and quenched down to 100 K in 18 ps, resulting in a phase we may call a-EDIP-II. The energy of a-EDIP-II, -4.45 eV/atom, is higher than a-EDIP-I by 0.05 eV/atom and higher than the diamond ground state by 0.21 eV/atom. The unphysical nature a-EDIP-II is suggested by the density *increase* of 5% versus the equilibrium crystal. As seen in Fig. 6.8, the pair correlation function retains some of the original structure, with a sharp isolated first peak at 2.38 Å, containing 4.17 neighbors, but as in the all disordered structures we have encountered, a sharp second peak (really a splitting of the first peak) appears at 3.12 Å, containing 1.71 more neighbors for a total coordination of 5.88. The distribution of local coordinations, given in Table 6.4, shows the predominance of five-fold and six-fold coordination with many larger coordinations also present, which accounts for the increased density. Of course, the coordination distributions from more realistic models are sharply peaked at 4. The a-EDIP-II bond angle distribution in Fig. 6.8 (b) is quite similar to that of a-EDIP-I, with all the same narrow, small-angle peaks, and the inner subshell structure resembling the liquid. The main difference between I and II is that the former has 4

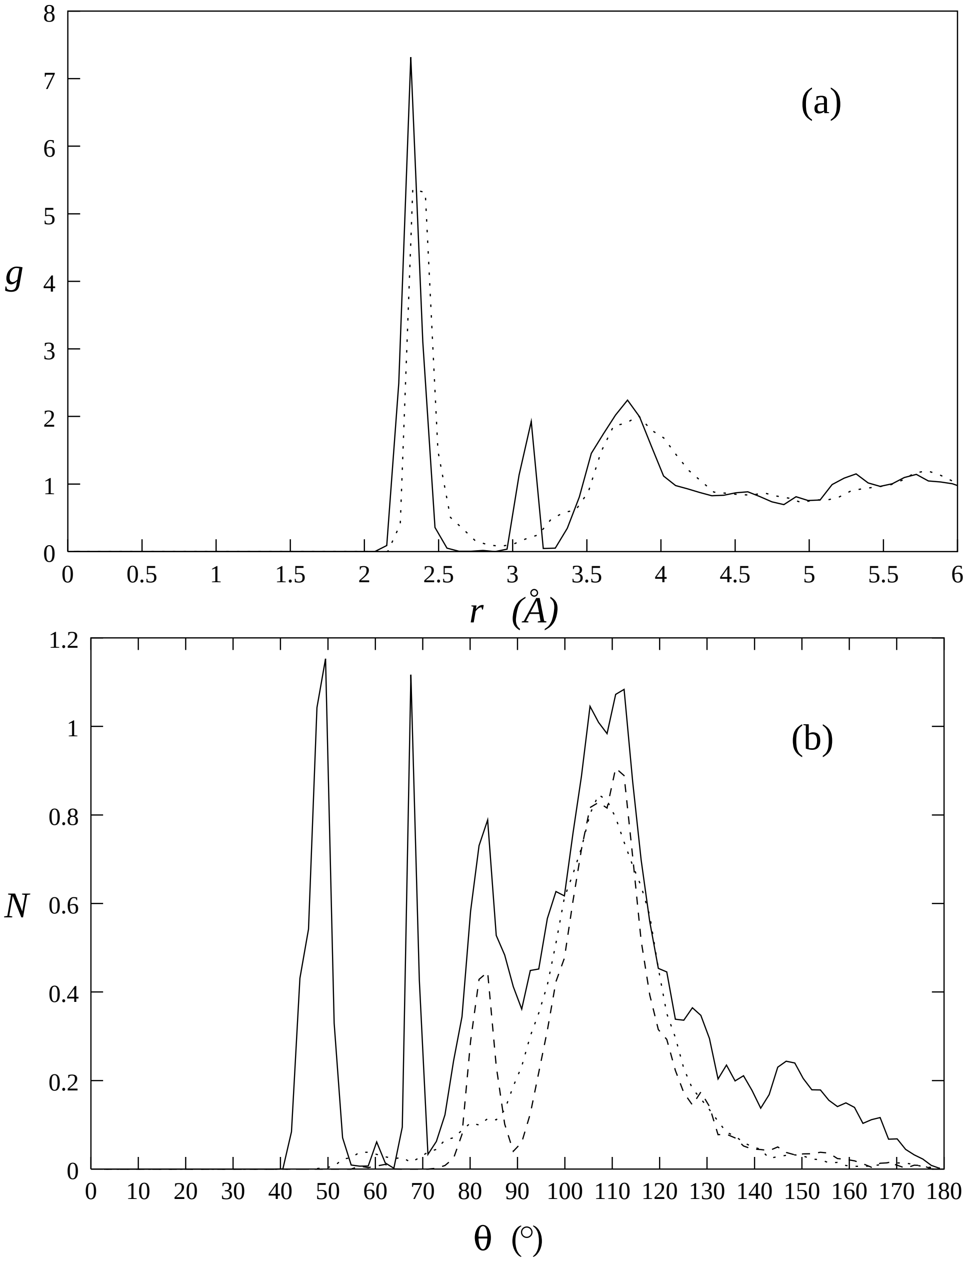

Figure 6.8: Structure of the second EDIP amorphous phase a-EDIP-II obtained from annealing a SW amorphous sample (dotted lines). The pair correlation function is shown in (a), and the bond angle distributions for the first (solid line) and first two (dashed line) peaks of $g(r)$.

Z	LDA	WWW	SW	EDIP
3	0.2	1.2	0.0	0.0
4	96.6	86.6	81.0	12.1
5	3.2	11.8	18.1	29.7
6	0.0	0.2	0.9	31.1
7	0.0	0.0	0.0	14.2
≥ 8	0.0	0.0	0.0	12.9

Table 6.4: Coordination statistics for amorphous structures generated by *ab initio* (LDA) dynamics with quenching [153], the Wooten-Winer-Weaire algorithm (WWW) [155], the indirect SW method (SW) [133], and annealing of the indirect SW structure with EDIP, a-EDIP-II.

instead of 2 neighbors on average in the outer, metallic subshell, and thus also has a much smaller volume ($V_I = 16.45$, $V_{II} = 19.00$ Å^3/atom). Otherwise various bonding arrangements are quite similar in the two amorphous phases.

As a last hope for EDIP to describe a-Si, we can try to build the amorphous phase by annealing a structure with fewer overcoordinated defects than the indirect SW sample. However, it turns out that a rather wide range of structures resembling a-Si transform into a-EDIP-II upon annealing, indicating that it is a free energy minimum with a broad basin of attraction. For example, a-EDIP-II emerges from annealing the low-density, four-fold coordinated structure generated by quenching the SW liquid with the SW potential modified for double-strength three-body energy (without the final step of annealing with the true SW potential). This starting point has greater volume and fewer overcoordinated atoms than the indirect SW amorphous, but it still relaxes into a-EDIP-II with EDIP forces, even at room temperature.

6.4 Thermal Stability of Bulk Defects

The results of the previous section lead us to question the range of transferability of the current version of EDIP because it sometimes transforms into unphysical structures in the presence of disorder, even at low temperatures. For pervasive long-range disorder, we have documented the collapse into a-EDIP-I from quenching the liquid and into a-EDIP-II from annealing a variety of low-temperature, amorphous states, characterized by tetrahedral short-range order. Consistent with the latter result, a-EDIP-II is formed when an amorphous-crystalline interface (created by N. Bernstein) is relaxed with EDIP at $T = 1000$ K, which should instead lead to crystallization in real silicon. These results are troublesome because they suggest that bulk crystalline defects, whose energies are exceptionally well described at $T = 0$, may be thermally unstable. Fortunately, that is not the case, as evidenced by a number of tests.

First of all, note that the splitting of the first peak of $g(r)$ present in all the unphysical disordered phases does not appear when heating a perfect diamond crystal all the way up to the melting point. Hence, the diamond structure is thermally stable against the kinds of unphysical, structural relaxations we have been examining. The same is true of some isolated point defects. For example, the unphysical rearrangements do not occur when a sample with a vacancy concentration of 0.05% is heated to the melting point. This concentration is many orders of magnitude larger than any observed in experiment. If the vacancy concentration is increased even further, then eventually there is sufficient disorder that regions of a-EDIP-II are nucleated. At 5% vacancy concentration, a-EDIP-II domains appear even at $T = 100$ K, but this is not problematic because such a high concentration implies non-isolated defects and voids, characteristic of a (very unphysical) foam. On the other hand, extensive tests of partial dislocation structures have not uncovered these kinds of problems, and in general, there appears to be no problem with finite temperature simulations of crystal defects right up to the melting point. Thus, we may claim EDIP to be a reliable potential for the simulation of bulk defects.

6.5 Prospects for Increased Transferability

Of course, we would like to be able to claim more. The theory behind the functional form should be capable of describing all the bulk phases, both covalent and metallic, ordered and disordered. Let us begin by addressing the amorphous, the most serious problem with the current version.

The only way to create and stabilize a more realistic amorphous (and avoid a-EDIP-I and a-EDIP-II) is apparently to modify the potential. We need to somehow eliminate the outer part of the split first peak of $g(r)$. Since this peak arises where the three-body force nearly vanishes but the pair attraction does not, it seems reasonable to try setting $b = a$ in the potential. This extends the range and increases the strength of three-body forces. Sure enough, when the indirect SW amorphous state is annealed with such a modified EDIP (without any refitting of the other parameters), a more reasonable structure results, as shown in Fig. 6.9 (dashed lines). It is a great improvement over the I and II phases described above, but is less realistic than the indirect SW amorphous. The isolated first peak contains 4.50 neighbors with covalent bond lengths, and the metallic second split peak is absent. An improvement over the indirect SW amorphous include the vanishing pair correlation in between the first two peaks, indicating less disorder in the random network, and the density, which is almost the same as the equilibrium diamond crystal. The bond angles are nicely peaked around tetrahedral, and the small angles peaks have mostly vanished. However, some new unphysical features appear, namely a broad peak at $60°$ and a small shoulder at $145°$. Nevertheless, with a simple change (suggested by the cohesive energy curves), namely increasing the three-body cutoff to the pair cutoff distance, we have salvaged reasonable behavior without any refitting of the potential. Of course, all other properties are altered, including elastic constants and defect formation energies, and there is no guarantee that the desirable

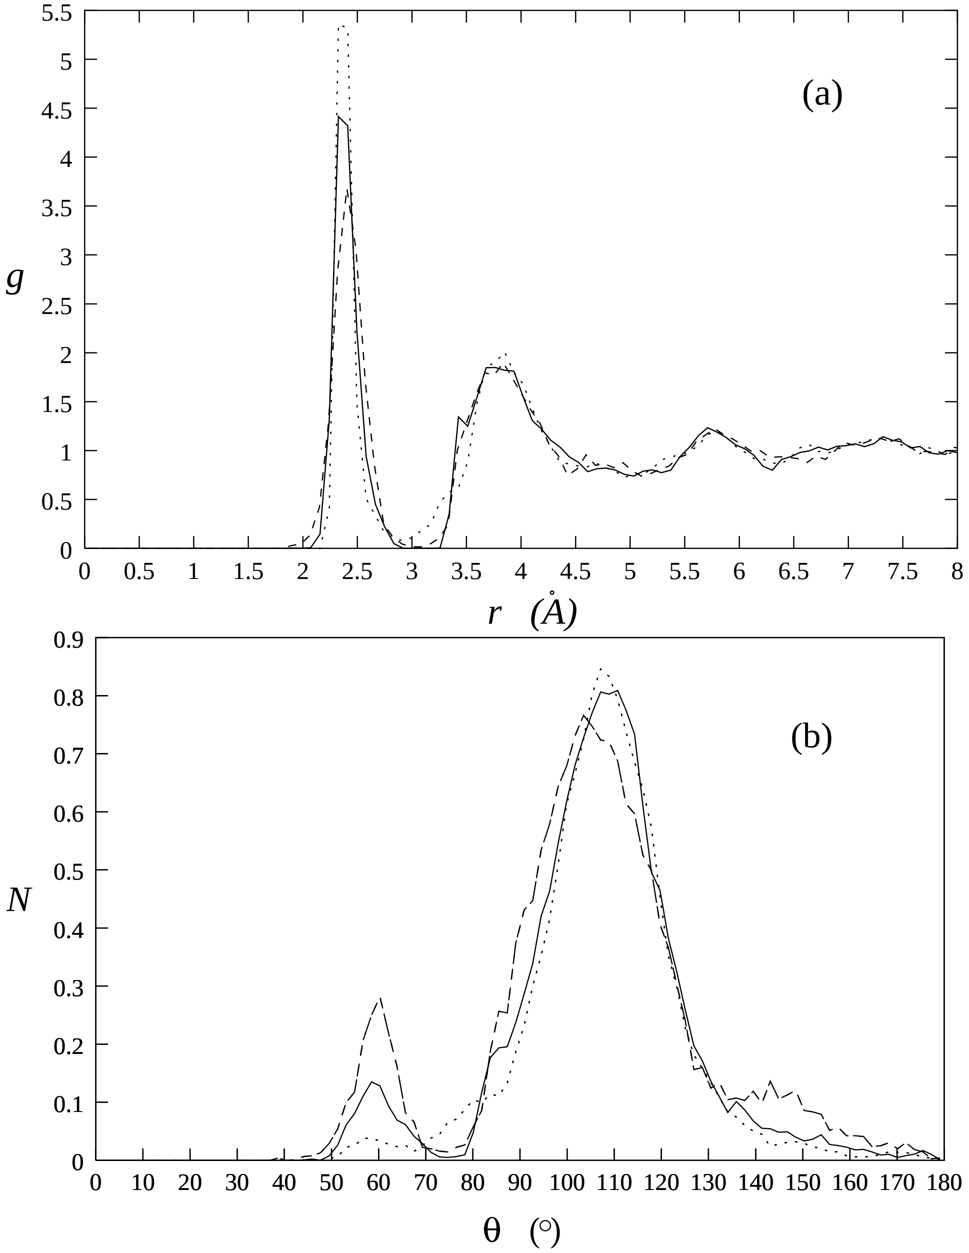

Figure 6.9: The pair correlation (a) and bond angle distribution (b) for amorphous phases of EDIP modified by using $H(x) = \lambda x^2$ with an extended three-body radial function (solid lines) and by simply setting $b = a$ (dashed lines), compared with the indirect SW structure (dotted lines).

features of the current version of EDIP can be preserved while fixing the problems with the amorphous phase. Indeed, recent experience in refitting the potential with the constraint of $b = a$ suggests that this goal is impossible with the current functional form.

To better understand the prospects for a simultaneous description of the amorphous phase and crystalline defects with EDIP, we explore modifications to the functional form that have minimal effect on bulk properties. The crucial gap in the theoretical foundation of EDIP explained in Section 5.1 is the shape of the angular function $H(x)$. The coordination dependence $x = (\cos\theta + \tau(Z))/w(Z)$ is supported by various arguments, but the particular choice $H_1(x) = \lambda(1 - \exp(-x^2))$, in spite of some appealing properties, has no real justification (aside from $H(0) = H'(0) = 0$ and $H''(0) > 0$).

Both a-EDIP-I and a-EDIP-II are characterized by the presence of small angles, $\theta < 90°$, which are mostly absent in the fitting database. If we could increase the penalty for small angles, then most of the database energies would not be affected. It turns out that (keeping all parameters the same) adding an additional penalty for small angles to the current version, $H(x) = H_1(x) + \eta x^m$ (for $m = 4, 6, 8$ and small η) does not solve the amorphous problems.

Drawing on the relative success of SW, it seems natural to try $H(x) = \lambda x^2$, which makes the EDIP form reduce to SW for $Z = 4$. This choice (with the original parameters of EDIP), keeps the same elastic constants and greatly increases the penalty for small angles. With this change, the amorphous improves dramatically (but not completely), indicating the great importance of the angular function. The small angle peaks of the bond angles reduce along with the split peak of the pair correlation, but none vanish.

A satisfactory amorphous can be produced with the EDIP functional form by combining the two helpful changes just described. First the angular function is changed to $H(x) = \lambda x^2$, and second the range of the three-body radial function $g(r)$ is extended. Unlike the exercise above where we simply set $b = a$, we now do the same with complementary changes to $g(r)$ to preserve values near $r = 2.35$ Å. Specifically, we choose a larger value $\gamma = 0.5$ (for reasons described in Section 5.2.3), and adjust the magnitude of $g(r)$ by multiplying the original λ by 1.225^2. This extends $g(r)$ to larger radii, $2.8 < r < 3.2$ Å, without affecting much the range, $2.2 < r < 2.6$ Å, present in most bulk defect structures. As shown Fig. 6.9 (solid lines), the pair correlation and bond angles are quite realistic with these modifications. The coordination under the first peak is 4.37, and 72% of atoms have four and 25% have five neighbors, which agrees with the *ab initio* data in Table 6.4 almost as well SW.

Although it has not been fitted to anything explicitly (other than the elastic constants and diamond structure, which we have preserved), it is interesting to test our amorphous-modified EDIP for the liquid. It does not perform well, but behaves differently than the original version. The structure at $T = 2500$ K and $P = 0$ shown in Fig. 6.10 is reminiscent of a hot amorphous phase. The pair correlation has a fairly well-isolated first peak containing 5.67 neighbors, and the bond angle distribution is bimodal with peaks at 60° and 90°. With this example we see the difficulty of simultaneously describing all the bulk phases, consistent with the observations of Ding and Andersen from their work in applying the SW potential to germanium [136].

The difficulty of the problem is also suggested by our analysis of the original EDIP liquid and amorphous phases. In the former, we have determined that the split first peak of $g(r)$ should be merged, preserving the mixture of covalent and metallic bonds, particularly the local structure of Fig. 6.5. In the latter, we must somehow achieve the opposite. In the amorphous the outer shell of first peak must be removed, leaving only covalent, tetrahedral atoms. Thus, in order for EDIP to describe both the liquid and amorphous phases, a delicate balance must be struck between the tendencies for metallic and covalent bonding.

Nevertheless, EDIP shows exceptional promise as general and transferable model for bulk phases and defects. We have seen that its functional form is much more flexible

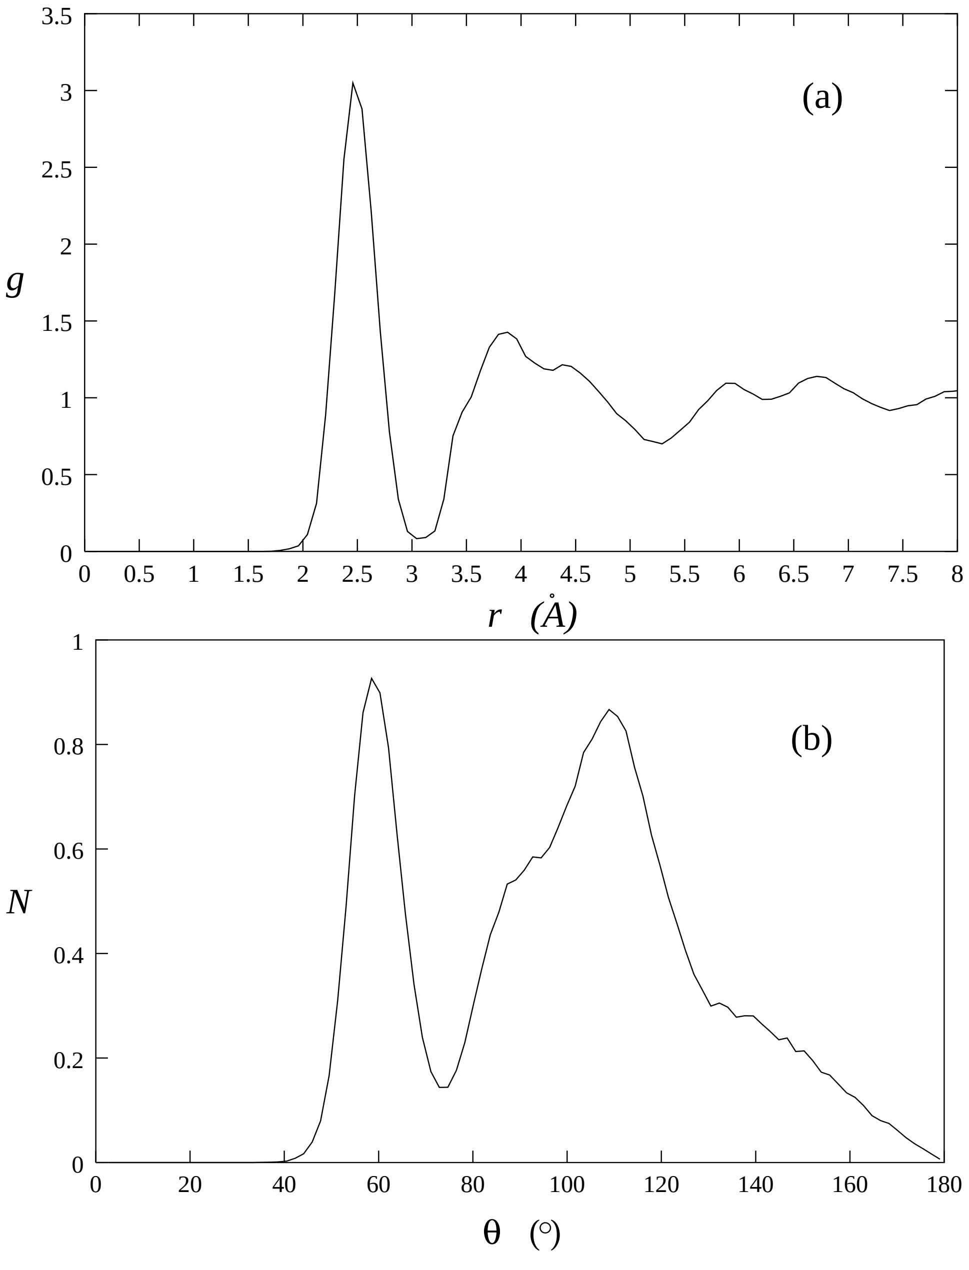

Figure 6.10: The liquid pair correlation (a) and bond angle distribution (b) for EDIP modified by using $H(x) = \lambda x^2$ with an extended three-body radial function.

than others, in spite of its comparable computational efficiency and small number of parameters. The examples in this section provide some guidance for refitting the potential for disordered phases, and it is likely that the amorphous and liquid can both be accommodated with minor changes to the functional form and some refitting.