ELSEVIER

Journal of Alloys and Compounds 444–445 (2007) 314–319

Journal of ALLOYS AND COMPOUNDS

www.elsevier.com/locate/jallcom

Accelerated dynamics study of vacancy mobility in δ-plutonium

Blas Pedro Uberuaga*, Steven M. Valone, M.I. Baskes

Materials Science and Technology Division, Los Alamos National Laboratory, Mail Stop G755, Los Alamos, NM 87545, USA

Received 29 June 2006; received in revised form 12 October 2006; accepted 17 October 2006
Available online 28 November 2006

Abstract

Using parallel-replica dynamics and temperature accelerated dynamics, we extract the rates for mono- and di-vacancy diffusion in δ-plutonium (Pu) using two parameterizations of the modified embedded atom method (MEAM). We find that mono-vacancy diffusion is faster in "pure" Pu than in δ-stabilized Pu. Also, at higher temperatures, the rate of double jumps is nearly the same as single jumps in pure Pu. Since these double jumps contribute four times as much as single jumps to the diffusion constant, models incorporating mono-vacancy diffusion must account for this mechanism to predict mass transport in Pu. While di-vacancies are energetically only slightly preferred compared to mono-vacancies, they are significantly more mobile. Surprisingly, this enhanced mobility is due to the prefactor; the migration barrier is essentially identical. The di-vacancy dissociates at a rate similar to the mono-vacancy hop rate.

Keywords: Diffusion; Point defects; Kinetics; Nuclear reactor materials; Computer simulations

1. Introduction

The importance of plutonium (Pu) as a material central to our modern society cannot be overstated. From its role in the nuclear weapons program to its use in modern fission reactors, it holds a prominent place in many important technologies. Thus, a complete understanding of the properties of Pu is crucial to being able to fully utilize its potential. Such an understanding begins at the atomic level, where point defects – vacancies and interstitials – dominate its behavior. Not only do these defects govern the self-diffusivity of Pu, but also larger scale phenomena such as bubble growth. The goal of this work is to more fully understand the properties of one of these point defects, vacancies. These properties can then be input into continuum level models that track radiation damage evolution [1,2].

Past experimental work has yielded a consistent picture of self-diffusion in δ-Pu via vacancies. Experimental methodology has included tracer diffusion and creep experiments (see Ref. [3] and references therein). As summarized by Fluss et al. [4], the measured activation energy for the vacancy component of self-diffusion is 1.3 ± 0.3 eV. This activation energy is the sum of the vacancy formation energy and the energy barrier for vacancy migration. Kinetic Monte Carlo simulations, in an effort to reproduce experimental results, find the migration energy of vacancies to be about 0.55 eV [4]. However, more recent experiments have found a slightly higher value for the migration energy of 1.1 ± 0.3 eV [5].

2. Methodology

Molecular dynamics (MD) simulation is an important tool for understanding materials at the atomic scale. MD allows us to follow the motion of individual atoms on a time scale of about 1 ns. In doing so, we can measure the properties of vacancies and interstitials, as well as other defects, directly, providing input to higher levels of modeling. One limitation of MD is that the accessible time scales are relatively short: ps to ns. To overcome this time scale limitation, we employ accelerated dynamics techniques [6], in particular parallel-replica dynamics [7] and temperature accelerated dynamics (TAD) [8] to reach time scales relevant for mono-vacancy diffusion. These simulations were done on periodic cells containing 256 atoms.

Parallel-replica dynamics involves a straight-forward parallelization of time simulated on multiple processors. Each processor evolves an independent replica of the entire system until a transition to a new state is detected on any of the processors. If the rare events that govern motion from state-to-state in the material are first order processes, then parallel-replica

* Corresponding author. Tel.: +1 505 667 9105.
E-mail address: blas@lanl.gov (B.P. Uberuaga).

0925-8388/$ – see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jallcom.2006.10.083

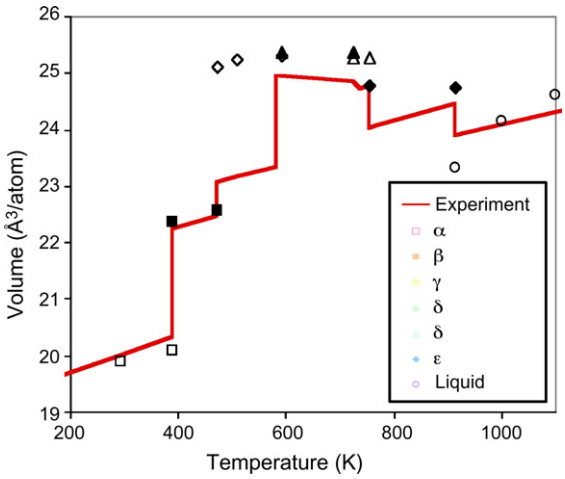

Fig. 1. Comparison of the predicted and experimental volume versus temperature of equilibrium Pu phases at 0 pressure. For these calculations, the monoclinic angles for the α and β phases were held at the experimental values. The MEAM prediction for the relative volume between the α and δ phases well approximates the experimental results.

dynamics is exact, even describing correlated events correctly. By parallelizing on M processors, one can reach a maximum speedup compared to standard MD of M times. Reaching this theoretical maximum efficiency depends on how infrequent the events are, as there is some computational overhead involved each time an event is observed.

TAD is potentially much more powerful, but also involves more assumptions about the nature of the system. In particular, if the system obeys harmonic transition state theory, TAD gives exact state-to-state dynamics. The algorithm involves running MD, constrained to the current basin or state of the system, at a temperature T_{high} higher than that of interest, T_{low}. The behavior at T_{high} is then extrapolated to T_{low} and, once a stopping criterion is met, an event is accepted based on this extrapolation. Using this algorithm, experimental time scales are achievable for certain systems [9,10].

Pu is an extremely complicated material, exhibiting more allotropes at ambient pressure (6) than any other element. Correctly describing this behavior theoretically has proven challenging, and even the best density functional theory methods have problems accurately describing the δ-phase. In the following work, we use the modified embedded atom method (MEAM) [11,12] to describe the Pu–Pu interaction. This potential is remarkable in that it correctly predicts the relative volume of each of the six allotropes of Pu, as shown in Fig. 1. To simplify the analysis somewhat, and to simulate the stabilizing effects of alloying Pu with elements such as Ga or Al, we have used both the parameters published in Ref. [13], as well as a modified parameter set in which $t_3 = 0$. This change in t_3 corresponds to removing a sensitivity to inversion symmetry in the crystal structure which stabilizes the δ-phase of Pu, much as Ga and Al do in real Pu. In what follows, these two parameterizations will be referred to as Pu^{I} and Pu^{II}, respectively.

The 0 K lattice constant predicted by the Pu^{II} parameterization if the system is held in a perfect FCC structure is 0.464 nm. This actually corresponds to the lattice constant of real Pu at a temperature between 650 and 700 K. However, we note that around 2 at.% Ga, the thermal expansion coefficient of δ −Pu is nearly zero over its temperature range of stability [14]. If Pu^{I} is allowed to relax isotropically (i.e., all 3 lattice constants are constrained to be equal), the volume is reduced by about 10–15%. If all three orthogonal lattice constants are allowed to be independent, the volume collapse is greater, reaching about 25%. That is, the ground state of α −Pu has a volume about 25% smaller than δ −Pu. For simplicity, we used a lattice constant of 0.464 nm for all simulations reported here as we are interested in mobility in δ −Pu.

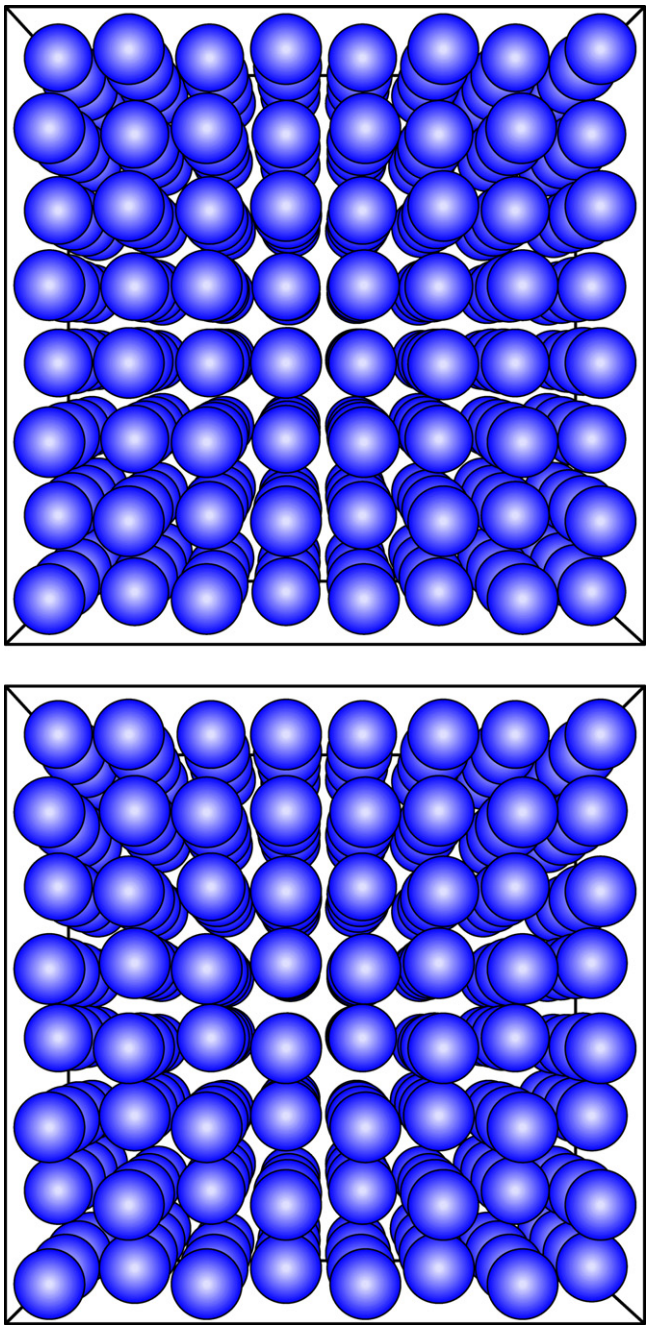

Fig. 2. The perfect FCC crystal and the distorted structure, predicted as the ground state structure by the Pu^{I} parameterization.

In the accelerated dynamics simulations, we are using temperatures above the melting point of the potential (T_m = 913 K). We do not observe melting in these simulations because the periodic boundary conditions force the system to be volume stabilized, preventing melting. However, the high temperatures may introduce anharmonic errors in the rates. As shown below, where we can compare migration energies from dynamics to static calculations, they agree very well. We thus do not believe we are introducing strong anharmonic effects by using such high temperatures.

As described in Ref. [17], the volume-stabilized 0 K ground state structure predicted by the Pu^I parameterization is not a true FCC structure, but rather is distorted. In fact, there are many distorted structures that are nearly degenerate in energy. The lowest energy structure we have found in the simulations described here is slightly different than that reported in Ref. [17] and is shown in Fig. 2. In the perfect FCC structure, at the 0 K lattice constant, each atom has 12 neighbors at 0.33 nm. In the distorted structure – which is lower in energy by 0.013 eV/atom – atoms pair up so that each atom has one neighbor at 0.29 nm, one at 0.37 nm, and the rest at 0.33 nm. At high temperature, the crystal is able to move between the 12 symmetry equivalent distortions available to the crystal. The existence of these degenerate states means that the potential energy surface for Pu^I is decorated with local minima. This will have consequences for the behavior of mono-vacancies in Pu.

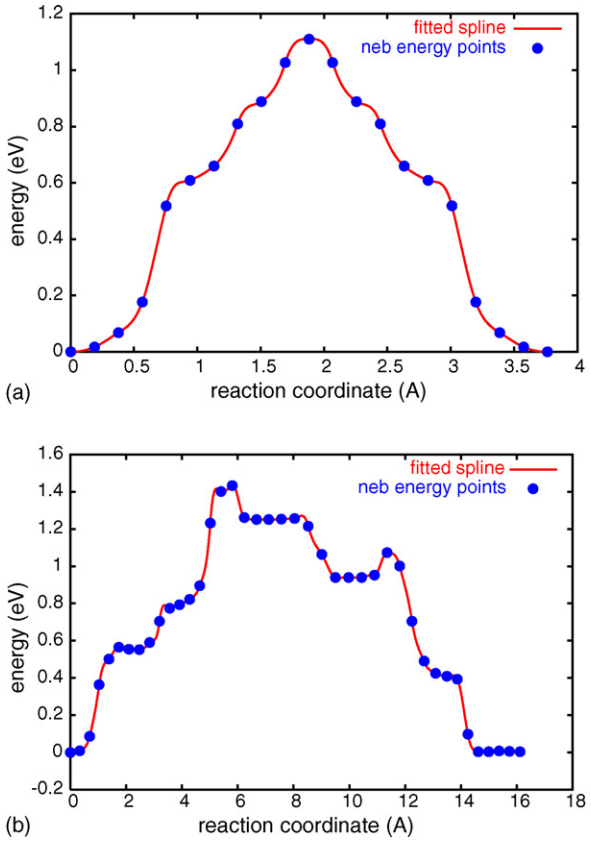

Fig. 3. Energy profile of minimum energy path for (a) a single mono-vacancy hop and (b) a double mono-vacancy hop in Pu^I.

3. Results

Because of the complicated landscape of the Pu^I parameterization, we have had more success applying parallel-replica dynamics than TAD. Parallel-replica dynamics was used to study both the mono- and di-vacancy in Pu^I at T = 900–1200 K. For a typical simulation running on 38 processors, we achieved a speedup of 36 times.

At these temperatures, mono-vacancy hops are relatively rare: in a simulation of Pu^I at 900 K for 3.0 μs, we saw only 12 events. Eight of these were simple single vacancy moves, in which one Pu atom moves directly into the vacancy site, resulting in the mono-vacancy moving one lattice spacing. The energy profile of the minimum energy path (MEP) for such an event has been calculated with the nudged elastic band (NEB) method [15,16] and is illustrated in Fig. 3a for Pu^I. The barrier for a single hop is about 1.11 eV, which agrees with the estimates of previous work using this potential [17] and with the value from the dynamics of 1.06 ± 0.01 eV.

Surprisingly, we also saw four events at 900 K in which the mono-vacancy executes a double hop, effectively moving a distance of 0.656 nm. Two Pu atoms move in a concerted motion and the mono-vacancy moves two lattice spacings during one event. The best MEP we have found for this process in Pu^I is illustrated in Fig. 3b. It has a barrier of about 1.43 eV which agrees well with the value extracted from the dynamics of 1.38 ± 0.13 eV. The MEP for this process is much more complicated than for the single hop.

By performing similar simulations at other temperatures, we have been able to extract the rates for these two processes at temperatures between 900 and 1200 K. Assuming that the prefactors are temperature independent, we find the rates listed

Table 1
Rates for the various mechanisms for the mono-vacancy (V) and di-vacancy (V_2)

Parameter set	Defect	Event	Rate prefactor (s^{-1})	Migration energy (eV)
Pu^I	V	Single hop	5.0×10^{12}	1.06 ± 0.01
Pu^I	V	Double hop	6.5×10^{13}	1.38 ± 0.13
Pu^I	V_2	Single hop	6.4×10^{13}	1.00 ± 0.09
Pu^I	V_2	Dissociation	5.9×10^{13}	1.26 ± 0.16
Pu^{II}	V	Single hop	4.1×10^{13}	1.55
Pu^{II}	V	Double hop	3.7×10^{14}	2.18
Pu^{II}	V	Triple hop	n/a	2.7

The rates for events in Pu^I are extracted from waiting times found during parallel-replica simulations. Those for events in Pu^{II} were found via TAD using a combination of NEB to find the migration energy and average waiting times to find the rate prefactor. We estimate that the prefactors are good to within a factor of two, while the uncertainty in the migration energies are given in the table. There are no errors for the migration energies in Pu^{II} as those are from static calculations.

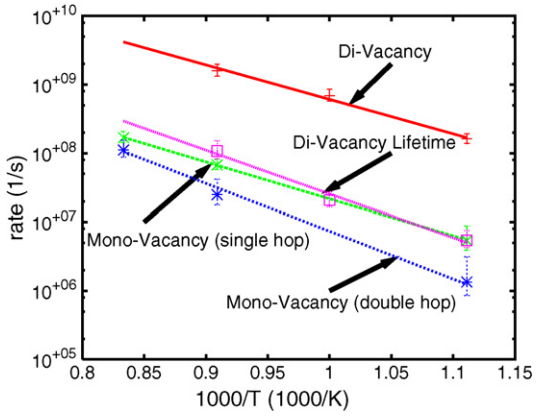

Fig. 4. Rates for the mono-vacancy single and double hop mechanisms, the di-vacancy diffusion mechanism, and the dissociation of the di-vacancy as a function of temperature for Pu^{I}. Error bars represent standard errors of the parallel-replica simulations.

in Table 1. The extracted activation energy for the single hop of the mono-vacancy, 1.06 eV, is nearly identical to the barrier calculated with NEB and shown in Fig. 3a. Thus, the dynamical path at finite temperature is likely the same as that found on the 0 K surface using the NEB and the mono-vacancy diffuses by exchanging with the atom that it is "paired" with in the distorted ground-state structure of δ-Pu predicted by the Pu^{I} parameterization. While in principle, there are twelve neighbors for the mono-vacancy to jump to, at any given time, only one is available for a hop.

When the rates are plotted as a function of temperature (see Fig. 4), it is clear that the rate for single hops is higher than that for double hops. However, each double hop results in the mono-vacancy moving twice as far. This increases the contribution to diffusion by the double-hop mechanism by a factor of four, resulting in a comparable contribution as the single-hop mechanism. Thus, any model which includes a description of mono-vacancy migration must take into account both single and double hops, as both are important.

Again, by performing parallel-replica runs at temperatures between 900 and 1200 K, we extract the rate of the di-vacancy diffusion mechanism given in Table 1. This rate, as illustrated in Fig. 4, is significantly higher than that for the mono-vacancy. The activation energy is nearly identical, but the rate prefactor is about two orders of magnitude larger. That the activation energy is very similar to the single-hop mechanism for the mono-vacancy suggests that the mechanism is also similar, which is surprising as the di-vacancy presumably allows for more space for the atoms to diffuse. The similarity in the two activation energies, however, suggests that the atom moving in the di-vacancy hop process does not take advantage of the extra space afforded by the extra missing atom, remaining, instead, strongly bound to neighboring atoms. That is, the diffusion path of the atom moving in the di-vacancy hop mechansim is very similar to that in the mono-vacancy hop mechanism where there is no extra vacancy. Indeed, NEB calculations confirm this view of the di-vacancy motion.

We can only speculate on the origin of the larger prefactor. The di-vacancy does have more nearest neighbors with which it can exchange, but the difference compared to the mono-vacancy is less than a factor of 2, much less than the difference we see in the rates. It seems reasonable to posit that, because the di-vacancy configuration is more "open" – and the potential may then be flatter – than the mono-vacancy, the system moves more quickly from one local distortion minimum to another and can find a "paired" site more quickly to which it can jump. Note that a homologous Stage V annealing temperature of 0.5, where vacancy clusters begin to dissolve and all defects are expected to be mobile, is approximately 460 K, based on Pu melting point of 920 K. This would require that a species mobile at lower temperatures, such as this di-vacancy, be the dominant diffusing species.

There is a slight binding energy of between 0.1 and 0.2 eV for two vacancies to form a di-vacancy cluster. Once the di-vacancy forms, it diffuses until it dissociates, either by a single hop of one vacancy to a neighboring site that is not a neighbor of the other vacancy, or one of the vacancies executes a double hop similar to the mechanism seen for the mono-vacancy. We estimated this di-vacancy lifetime as a rate which is reported in Table 1. This rate is also shown in Fig. 4. Once a di-vacancy forms, even though the binding energy is not very strong, there is a kinetic limitation to the rate at which it dissociates. However, at equilibrium, as the binding energy is so small, the relative number of di-vacancies compared to mono-vacancies will be small. This ratio goes as $\exp[-(E_f - E_b)/k_B T]$, where $E_f = 0.7$ eV is the formation energy of the mono-vacancy, $E_b = 0.1 - 0.2$ eV is the binding energy of the di-vacancy and k_B is the Boltzmann constant (assuming entropic factors are the same for both species). At $T = 300$ K, this gives a relative concentration of di-vacancies to mono-vacancies of between 10^{-12} and 10^{-9}. Thus, at equilibrium, di-vacancies will contribute little to self-diffusivity in Pu. However, they may be critical in non-equilibrium conditions, such as those present in radiation damage environments.

A similar study of the mono-vacancy was done for Pu^{II}. Because of the stabilization of the FCC phase, Pu^{II} does not exhibit the global distortions in structure characteristic of Pu^{I}. However, local distortions around the mono-vacancy itself do exist and we have had to modify the TAD procedure to improve efficiency for such systems. First, instead of directly minimizing the system, we perform simulated annealing where the system is thermalized at 300 K and slowly cooled to 0 K. If, at that time, the geometry is greater than $r_{trans} = 0.03$ nm away, we declare a transition. Thus, typically, transitions are only seen if the system leaves the current "super"-basin and enters another, where the barrier separating these basins is high relative to the annealing temperature of 300 K. Second, we also ignore transitions to states that have small barriers to return. In essence, minima separated by very small barriers of 0.3 eV, such as the states labeled 'A' in Fig. 5, are treated as a single basin.

Using $T_{high} = 2000$ K and $T_{low} = 1000$ K, we have run long enough to simulate 0.43 μs. This took about 8 days of CPU time on a single processor and represents a speed up of about 6 times over MD. During this 0.43 μs, three single mono-vacancy hops were accepted. The MEP for such a hop is illustrated in Fig. 5. The barrier for the single hop is about 1.55 eV. While none

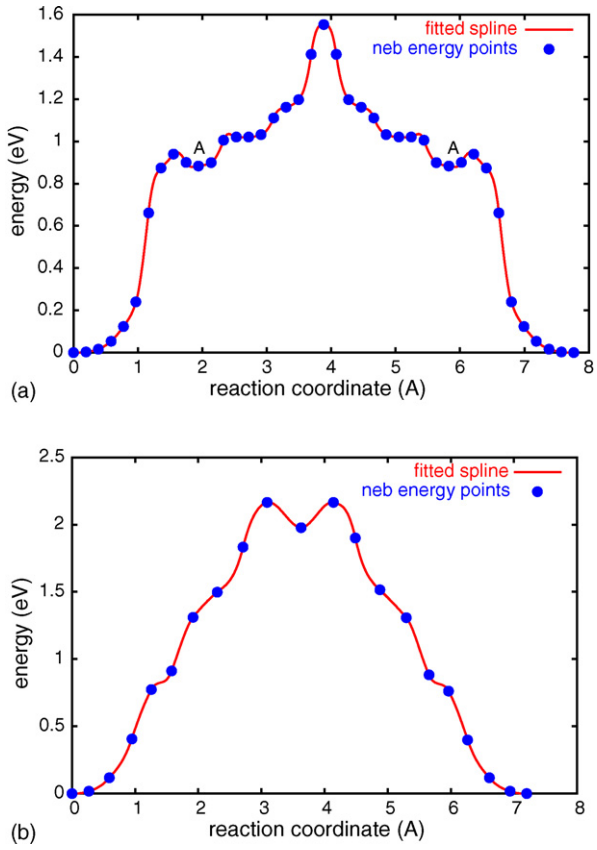

Fig. 5. Energy profile of minimum energy path for (a) a single mono-vacancy hop and (b) a double mono-vacancy hop in Pu^{II}. The label A in the path for the single-hop illustrates a local minimum that is ignored in the TAD procedure.

were accepted, the TAD simulation did see a number of double hops as well and the MEP for such events is also shown in Fig. 5. The barrier for this event is much higher – about 2.18 eV – than the single hop barrier. This is in contrast to the similarity seen between the two types of events in Pu^{II}. We have also seen the mono-vacancy execute a triple hop in Pu^{II} with a barrier of 2.7 eV. We have not seen enough of these events to extract a prefactor from the waiting times. No such event was seen in Pu^{I}.

Extracting rates from the waiting times at T_{high}, at which many events were observed, we find the rates given in Table 1. While the prefactors for both mechanisms are an order of magnitude higher for Pu^{II} than Pu^{I}, the barriers are much higher, resulting in mono-vacancy diffusion in Pu^{II} being much slower than in Pu^{I}. Also, mono-vacancy diffusion in Pu^{II} is completely dominated by single hops, while both single and double hops contribute in Pu^{I} (once the factor of four increase in the double-hop contribution to diffusion is accounted for).

4. Discussion

Isochronal annealing experiments have been performed previously on a Pu–Ga alloy [18]. In those experiments, a sample was held at ~ 20 K and allowed to self-irradiate. The sample was then returned to various target temperatures and held at those temperatures for a specified time. The resistivity of the sample was then measured after that time had elapsed. The estimated transition temperature for stage III annealing [19], where mono-vacancies are thought to become mobile, is 190 K. This can be compared to a homologous, stage III temperature for Pu of 183 K. However, because of the complicated phase diagram of Pu, referring to homologous temperatures should be viewed with caution. The implied diffusion barrier based on this transition temperature is estimated to be 0.7 eV [18]. As discussed by Fluss et al. [4], previous experimental data put the activation energy for self-diffusion via vacancies at 1.3 ± 0.3 eV [3]. If one assumes the typical FCC behavior that 60% of the activation energy is due to vacancy formation and the remaining 40% is due to vacancy migration, the vacancy migration energy in δ-Pu would be between 0.4 and 0.6 eV.

Our lowest barrier is 1.00 eV, significantly higher than the experimental estimates, with the exception of Timofeeva [5], which found a value of 1.1 ± 0.3 eV. Comparing to the bulk of experimental data, consideration of the mobility of the di-vacancy does not resolve the discrepancy between our result and experiment. This discrepancy has several possible sources that we can not conclusively discriminate at this time. One possibility is that the potential is not describing the transition region of the interactions correctly. Another is that the annealing stages in the experiments have been assigned incorrectly. A third possibility is that the diffusion mechanism in the Ga-stabilized alloy is different than in the metal. Additional investigation is needed to resolve this issue.

5. Conclusions

To summarize, we are applying accelerated dynamics methods to Pu in an effort to understand the behavior of mono- and di-vacancies. Using parallel-replica dynamics, we have identified two different mono-vacancy diffusing events – a single hop and a double hop mechanism – in the original parameterization for Pu for the MEAM potential. The di-vacancy diffuses via a mechanism similar to the mono-vacancy single-hop mechanism and dissociates by a combination of both mono-vacancy mechanisms. We find that the contribution to diffusion for both mono-vacancy diffusion mechanisms are similar. TAD was applied to a modified potential which stabilizes the δ phase. Again, we see both types of hops, but single hops clearly dominate the behavior for this material. As might be expected, mono-vacancy diffusion occurs much more quickly in "pure" Pu (PuI) than in the "stabilized" Pu (PuII).

Acknowledgements

We would like to acknowledge helpful discussions with A.F. Voter. This work was funded by the DOE Enhanced Surveillance Campaign.

References

[1] C.M. Schaldach, W.G. Wolfer, in: M.L. Grossbeck, T.R. Allen, R.G. Lott, A.S. Kumar (Eds.), The Effects of Radiation on Materials, ASTM STP 1447, ASTM International, West Conshohocken, PA, 2003.
[2] H. Trinkaus, B.N. Singh, S.I. Golubov, J. Nucl. Mater. 283–287 (2000) 89.

[3] W.Z. Wade, J. Nucl. Mater. 38 (1971) 292.
[4] M.J. Fluss, et al., J. Alloy Compd. 368 (2004) 62.
[5] L.F. Timofeeva, in: K.K.S. Pillay, K.C. Kim (Eds.), International Conference on Plutonium-Future—The Science: Topical Conference on Plutonium and Actinides, AIP Conference Transactions, 2000, p. 11.
[6] A.F. Voter, F. Montalenti, T.C. Germann, Ann. Rev. Mater. Res. 32 (2002) 321.
[7] A.F. Voter, Phys. Rev. B 57 (1998) R13985.
[8] M.R. Sorensen, A.F. Voter, J. Chem. Phys. 112 (2000) 9599.
[9] J.A. Sprague, et al., Phys. Rev. B 66 (2002) 205415.
[10] B.P. Uberuaga, et al., Phys. Rev. Lett. 92 (2004) 115505.
[11] M.I. Baskes, J.S. Nelson, A.F. Wright, Phys. Rev. B 40 (1989) 6085.
[12] M.I. Baskes, Phys. Rev. B 46 (1992) 2727.
[13] M.I. Baskes, Phys. Rev. B 62 (2000) 15532.
[14] A.C. Lawson, et al., Phil. Mag. B 82 (2002) 1837.
[15] G. Henkelman, B.P. Uberuaga, H. Jónsson, J. Chem. Phys. 113 (2000) 9901.
[16] G. Henkelman, H. Jónsson, J. Chem. Phys. 113 (2000) 9978.
[17] S.M. Valone, et al., J. Nucl. Mater. 324 (2004) 41.
[18] B.D. Wirth, et al., MRS Bull. 26 (2001) 679.
[19] P. Ehrart, in: H. Ullmaier (Ed.), Atomic Defects in Metals, Landolt-Bornstein New Series Group III/25, Springer-Verlag, Berlin, 1991.