Finding the atomic configuration with a required physical property in multi-atom structures

2007 J. Phys.: Condens. Matter 19 402201

(http://iopscience.iop.org/0953-8984/19/40/402201)

IOP FTC

IOP PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **19** (2007) 402201 (7pp) doi:10.1088/0953-8984/19/40/402201

FAST TRACK COMMUNICATION

Finding the atomic configuration with a required physical property in multi-atom structures

Mayeul d'Avezac and Alex Zunger

National Renewable Energy Laboratory, Golden, CO 80401, USA

Received 22 June 2007
Published 21 September 2007
Online at stacks.iop.org/JPhysCM/19/402201

Abstract

In many problems in molecular and solid state structures one seeks to determine the energy-minimizing decoration of sites with different atom types. In other problems, one is interested in finding a decoration with a target physical property (e.g. alloy band gap) within a certain range. In both cases, the sheer size of the configurational space can be horrendous. We present two approaches which identify either the minimum-energy configuration or configurations with a target property for a fixed underlying Bravais lattice. We compare their efficiency at locating the deepest minimum energy configuration of face centered cubic Au–Pd alloy. We show that a global-search genetic-algorithm approach with diversity-enhancing constraints and reciprocal-space mating can efficiently find the global optimum, whereas the local-search virtual-atom approach presented here is more efficient at finding structures with a target property.

(Some figures in this article are in colour only in the electronic version)

Along with the traditional study of the properties of a *given* molecular or crystalline system, recent years have witnessed the development of the corresponding *inverse problem*, where one inspects large databases of molecules or periodic solids in search of a structure with a given target property [1–5]. Such quests often focus on materials derived from a given skeletal structure, the sites of which can be *decorated* with different substituents—chemical groups in chemistry [4, 5] and different atoms in solid solutions [1–3]. There are M^N possible configurations for a cell of N lattice sites occupied by M substituents. These configuration spaces feature broad ranges of physical properties, as illustrated by the differing optical properties of random versus ordered zincblende semiconductor alloys at fixed chemical composition [6, 7], by the range of Curie temperatures in Mn substitution patterns on the cation sublattice of binary semiconductors [2], by the thermoelectric response of structures with the same chemical composition but with different effective-mass anisotropy [8]. Perhaps the best-known manifestation of the dependence of properties on configuration is the difference in total energies between various decorations of structures with the same chemical formula, as evidenced by the phenomenon of isomerism in molecular chemistry and in solids, e.g. the

0953-8984/07/402201+07$30.00

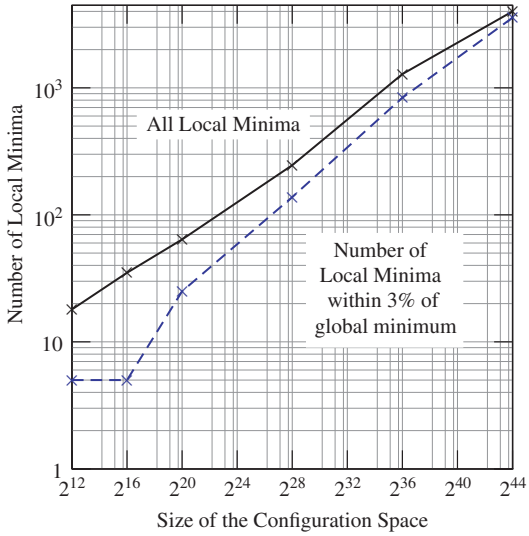

Figure 1. Total number of local minima of an AuPd fcc alloy in $2 \times 2 \times n$ supercells. The energy is obtained from a first-principles cluster expansion [18] $\mathcal{O}_{\mathrm{CE}}(\sigma)$. The solid line represents all local minima. The dashed line accounts only for the minima located within 3% of the deepest minimum energy configuration $L1_0$. The total number of local minima is seen to grow exponentially; most of these are located close in energy to $L1_0$.

existence of the many long-period CuPd superstructure variants [9] and numerous polytypes of SiC [10]. Indeed, the search for atomic decorations of a given lattice type (e.g. fcc, bcc) which gives the minimum total energy is one of the classic problems of alloy lattice theory [11–13].

Configuration search faces two problems. First, evaluating the property $\mathcal{O}(\sigma)$ of the relaxed configurations σ can be onerous in terms of computer resources. In that case, it is impractical to evaluate every single candidate [13] of a space comprising M^N structures. Second, many physical properties exhibit multiple *local* minima when the atoms are swapped on lattice sites. Finding the *global* minimum-energy configuration is always a challenge for search methods. This can be illustrated for an alloy problem using the total energy of fcc AuPd configurations calculated from first principles. Figure 1 shows the total number of *local minima* (configurations for which any Au $\Leftrightarrow$ Pd atom swap raises the formation enthalpy) in $2 \times 2 \times n$ supercells of an fcc AuPd alloy with respect to the size of the configuration space. We see that the number of configurational local minima grows exponentially with the number of sites. The complexity of the configuration space, as evidenced by figure 1, raises the question of the reliability of a search procedure. We will address this issue by quantifying the average number of $\mathcal{O}(\sigma)$ evaluations which are necessary for obtaining the deepest minimum energy configuration within each $2 \times 2 \times n$ supercell. As will be seen, this approach yields two results: (i) it allows us to identify two very different types of search application; (ii) it allows us to rigorously compare different search methods. Indeed, we find that search problems can be separated into a 'ground-state configuration search' where the precise global minimum is the required outcome, and a 'target property configuration search' where one is not seeking the exact minimum-energy configuration but rather a set of pertinent configurations which fulfill a certain fuzzy requirement. An example of the first case is the identification of thermodynamically stable atomic configurations at $T = 0$ K. An example of the second case would be the search for a (Ga, Mn)As alloy configuration [2] with a 'high' Curie temperature. The configuration with maximum Curie temperature may not be an

adequate solution, since synthesizing it precisely may be impractical. The issue of a rigorous comparison between different search methods is very much lacking in the current material design literature. Indeed, the related field of 'general space group optimization' (in which configurational degrees of freedom as well as positional degrees of freedom are part of the search) has seen the introduction of genetic algorithms [14–16] as well as relatively simple random search approaches [17]. It is not certain at which point (with respect to the size of the system) the complexity of genetic algorithms is warranted.

Generally, optimization methods are separated into *local*-search methods (such as the conjugate gradient approach)—which travel from known configurations to better neighboring configurations until a local minimum is found—and *global*-search methods—which have the ability of finding better configurations via global transformations (such as swapping blocks of atoms between structures) of known configurations. In this paper we demonstrate that: (i) the 'deepest minimum energy configuration search' can be efficiently solved for using a *global*-search method, here a genetic algorithm (GA) where motifs passed on from parents to offspring are identified in reciprocal space and where the size of the region of explored configuration space is enhanced via a diversity constraint, and (ii) the 'target property configuration search' problem can be solved efficiently with a continuous *local*-search method. In the past, solving *discrete* configuration search problems—for which gradients do not exist *a priori*—were limited to *discrete* optimization methods such as simulated annealing. We discuss a procedure, the *virtual-atom* approach (VA), based on the mapping of a discrete configuration space onto a continuous configuration space, as introduced in Wang *et al* [4].

Both GA and VA search methods are applied to finding the structure with minimum formation enthalpy $\mathcal{O}_{\mathrm{CE}}(\sigma = \{S_0, \ldots, S_N\})$ of an $\mathrm{Au}_{1x}\mathrm{Pd}_x$ alloy. Each search is performed within a single cell-shape at a time, e.g. cell vectors are not optimized. All results in this paper are presented for single cell-shape searches. In practice, one is interested in the configuration with minimum-energy, or optimal target-property, *across all cell-shapes with N or fewer atoms*. As shown in [28], such a search can be performed by exploring a finite set of cell-shapes with N or fewer atoms. The occupation by an Au or Pd atom of each lattice site i in the supercell is represented with a discrete spin $S_i = \pm 1$. The enthalpy $\mathcal{O}_{\mathrm{CE}}(\sigma)$ was previously parameterized [18] using cluster expansion. This functional describes configurational energetics by fitting an Ising-like Hamiltonian of pair and many-body interactions to a set of ∼40 first-principles total-energy calculations. The formation enthalpies of *relaxed* fcc AuPd structures outside the fitting set are predicted to within 3 meV per atom of their LDA value. As a result, we expect that the behavior of optimizers which rely on $\mathcal{O}_{\mathrm{CE}}(\sigma)$ only will behave quantitatively in the same way when applied directly to an *ab initio* functional. We next discuss the main features of our GA and VA methods.

Reciprocal-space genetic algorithm: any genetic algorithm mimics natural evolution and survival of the fittest in order to optimize a *fitness* or physical prospect $\mathcal{O}(\sigma)$ of interest. The efficiency of genetic algorithms depends primarily on mating, i.e. the ability of parents to pass on favorable traits to their offspring. For instance, the approaches of [14, 19] swap blocks of atoms between configurations, thus combining the motifs in each block into the offspring. Unfortunately, it also creates artificial boundaries across which 'motifs' are not inherited from parents, but are rather an artifact of the mating procedure. We introduce reciprocal-space mating (see figure 2(a)), where an intermediate offspring configuration is created by swapping the amplitudes of the finite set of structure factors $\mathcal{S}(\mathbf{k})$ of the periodic supercell ($\mathbf{k}$ are the allowed reciprocal-space vectors). These amplitudes constitute a seamless set of motifs spanning the whole supercell. Indeed, the $\mathcal{S}(\mathbf{k})$ represent concentration waves of A/B atomic material. In contrast to prior real-space mating procedures [14, 20, 19], with our approach, the occupation of each site is correlated with the occupation of all other sites within the supercell. A

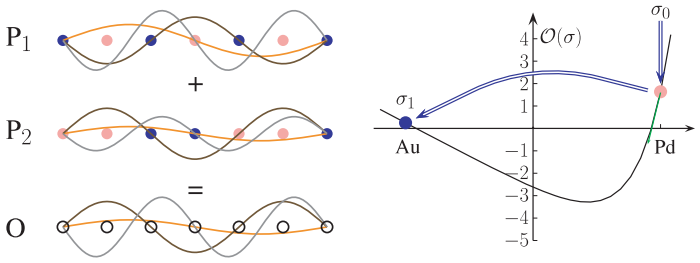

(a) Reciprocal-Space GA (b) Virtual Atom

Figure 2. Schematics of (a) the reciprocal-space mating, and (b) the virtual-atom algorithm. In the former, pale (red) and dark (blue) circles denote the occupation of a lattice site by an Au or Pd atoms in the parent configurations P_1 and P_2. The intermediate offspring configuration O_i is constructed from the concentration waves (solid lines) of AuPd material present P_1 and P_2, i.e. from their structure factors. A physical configuration is obtained *via a normalization step*. The virtual atom (b) implements a local-search method where the decision to explore a neighboring discrete configuration is deduced from the gradient at the original discrete configuration.

final offspring configuration is obtained by normalizing the spins of the intermediate offspring to $S_i = \pm 1$ depending on the sign of their real value. This approach can be readily expanded to problems with more than two substituents by adequately partitioning the complex plane. Genetic algorithms generally have the unfortunate propensity of 'locking' onto a set of more or less favorable traits, until eventually the population is composed entirely of clones of a single individual. To avoid this problem, we enforce diversity constraints ensuring that each individual configuration is unique within the existing population.

The virtual-atom approach is an alternative *local-search* approach for obtaining the global optimum which uses the information contained within the gradient of the function $\delta\mathcal{O}(\sigma)$. The challenge is to transform the discrete problem of decorating lattice sites with either A-type or B-type atoms into a continuous problem where discrete atoms are replaced by 'virtual atoms' which change continuously from pure A to pure B. This type of approach first surfaced in structural design engineering as early as 1988 [21], but has been used only recently for material design by Wang *et al* [4]. The virtual-atom approach replaces the discrete spins $S_i = \pm 1$ by continuous spins $\tilde{S}_i \in [-1, 1]$, which represent mixtures of x_i A-type atoms with $(1 - x_i)$ B-type atoms (with $\tilde{S}_i = 2x_i - 1$) at *each* site i (in the following, tilde indicates continuous quantities). This approach differs from the virtual-crystal approximation [22] in that each site i has a different concentration x_i. The derivative $\delta\tilde{\mathcal{O}}(\sigma)/\delta\tilde{S}_i$ takes on the meaning of how desirable it is for site i to be occupied by one type of atom rather than the other. With this derivative in hand, *gradient-driven local*-search methods can be used to optimize *discrete*-configuration problems[1].

There are two problems with such an approach. First, the virtual configuration $\tilde{\sigma}_0$ obtained from a gradient-driven method is not necessarily a physically realizable configuration (where each site is occupied by either an A-type or a B-type atom). Second, the physical minimum-energy configuration is not necessarily the configuration closest to the continuous minimum-energy configuration, nor indeed does it have to be in its neighborhood. Since a discrete set of points is targeted, imposing physical constraints via Lagrange multipliers will only lead the optimization procedure to the nearest-lying physical configuration, irrespective of the value of $\mathcal{O}(\sigma)$.[2] For this reason, we introduce a local-search method which only uses physical atoms,

[1] *Computation alchemy* [27] constitutes implementation of such fictitious gradients *ab initio*; this, to the best of our knowledge, has never been used for material design.

[2] We thank Peter Graff for discussing this point.

yet still makes use of the virtual-atom *gradient* to go from one physical point to the next. This method is described in figure 2(b) and proceeds as follows. Starting from a physical configuration σ_0, a spin direction S_i is picked at random and its gradient $\delta\mathcal{O}(\sigma_0)/\delta\tilde{S}_i$ computed. If the gradient is positive, we iterate with another spin direction. Otherwise, the prediction offered by the gradient is checked by evaluating the physical prospect $\mathcal{O}(\sigma_1)$ of the neighboring *physical* structure σ_1 in direction $\tilde{S}_i$. We then iterate from the better structure σ_0 or σ_1. Convergence is deemed achieved when all spin directions have been explored without finding a better physical structure. In this way, the algorithm hops from one physically realizable structure to the next, avoiding the pitfalls associated with unphysical structures. A similar algorithm has been proposed recently [5], with the difference that all gradient directions are evaluated at each step, and that the algorithm moves in the direction indicated by the lowest gradient.

Importance of a statistical comparison: since $\mathcal{O}(\sigma)$ contains many local minima (see figure 1), any local-search method, such as VA, must be restarted a number of times from different positions in the configuration space to ascertain that a global optimum has been found. As for genetic algorithms, they are stochastic in essence. Hence both GA and VA efficiencies are meaningful only on a statistical basis. To this end, we document $\langle N\rangle$, the average number of evaluations $\mathcal{O}(\sigma)$ before which 380 out of 400 independent runs have found target structures (i.e. 95% confidence). Since the VA approach contains two types of evaluation, $\mathcal{O}(\sigma)$ and $\delta\mathcal{O}(\sigma)$, we will present VA results 'including gradient cost' and 'excluding gradient cost'. For practical purposes, we set the cost of computing $\delta\mathcal{O}(\sigma)$ as equivalent to the evaluation of $\mathcal{O}(\sigma)$. Actual cost will differ depending on the functional explored and on implementation.

We study the efficiency of both VA and GA at finding the target structures in $2 \times 2 \times n$ supercells of the fcc $Au_{1-x}Pd_x$ system, with $n = 3$–12. The global optimum of each of these supercells is a multiple of the $L1_0$ decoration, with a formation enthalpy of -84.28 meV/atom. Note that $L1_0$ is not a true ground state of AuPd, but only of the specific $2 \times 2 \times n$ supercells that we investigate [18]. This result was obtained as part of the deepest minimum-energy search reported below, by running the genetic algorithm for a very long time.

To establish the expectation for a good search method, a random search approach was first applied where candidate structures were picked randomly until one satisfied the target property requirement. We find that solving for a space composed of 4096 configurations requires as many as 6000 evaluations of $\mathcal{O}(\sigma)$ (neither in this approach nor in the VA or GA approach did we keep track of the previously computed $\mathcal{O}(\sigma)$).

Deepest minimum energy configuration search: figure 3(b) gives the average number of evaluations necessary for obtaining the exact global minimum with respect to the size of the configuration space. We observe a crossover between the efficiencies of the two approaches as the configuration space increases in size, and hence in complexity (as evidenced by figure 1). Indeed, for a system consisting of 48 atoms, reciprocal-space GA will investigate as few as 6000 configurations before finding the minimum, whereas the virtual-atom approach needs as many as 10 times more evaluations. It should be noted that the computational effort necessary to find the deepest minimum energy configuration grows exponentially with the number of atoms N for both search procedures. This is not surprising since solving the Ising problem in three dimensions is 'NP-hard' [23, 24]. More to the point, it is not expected that a search procedure exists (for a deterministic computer) for which the required computer effort would grow at most polynomially with N.

Target property configuration search: Figure 3(b) gives the number of evaluations necessary for obtaining at least one structure with a formation enthalpy lower than -81.75 meV, i.e. within three per cent of the minimum-energy configuration ($L1_0$) obtained above. The local-search virtual-atom approach is quite equal to resolving this particular search problem. Indeed,

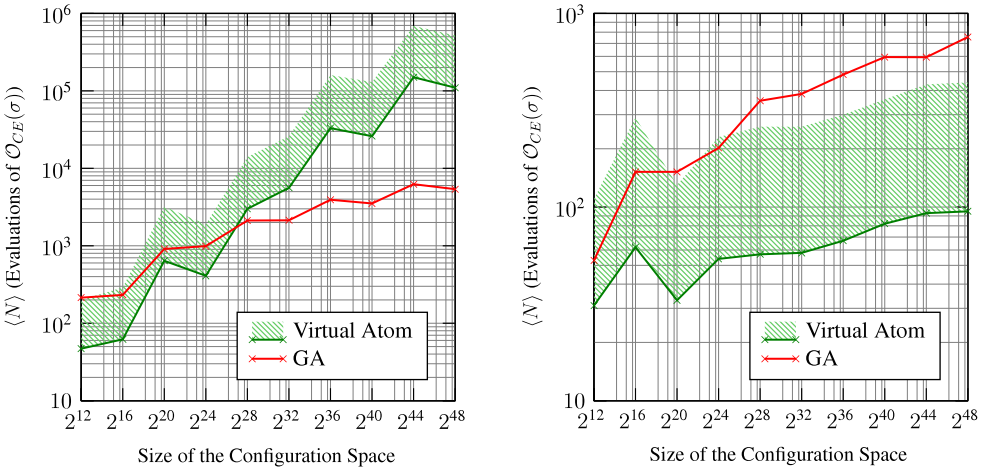

(a) Deepest Minimum Search (b) Target-Property Search

Figure 3. Deepest minimum (a) and target property (b) configuration search for an AuPd fcc alloy in $2 \times 2 \times n$ supercells. The energy is obtained from a first-principles cluster expansion [18] $\mathcal{O}_{\mathrm{CE}}(\sigma)$. Shown are the average number of evaluations $\langle N \rangle$ of $\mathcal{O}_{\mathrm{CE}}(\sigma)$ necessary to obtain (a) the deepest minimum energy configuration, and (b) at least one target configuration in 95% of the runs. A target configuration is defined as any configuration within 3% of the deepest minimum energy configuration $L1_0$. The shaded area represents the virtual-atom results 'including' (upper boundary) and 'excluding' (lower boundary) the cost of evaluating the derivatives $\delta\mathcal{O}(\sigma)$, where the cost of one derivative in one direction is set equal to the cost of one $\mathcal{O}(\sigma)$ evaluation. GA results have been optimized with respect to population size.

even with a 48-atom supercell, a pertinent structure is found on average after evaluating less than a hundred configurations out of the 2^{48} possible candidates[3]. This result is of import since it confirms that VA and related approaches [4, 5] put configuration searches well within the reach of first-principles functionals, thus allowing the quantitative prediction of new structures with desired properties.

In summary, we have studied two different types of search problem: (i) the 'deepest minimum energy configuration search' where finding the exact optimum decoration suffers no compromise, and (ii) the 'target property configuration search' where one is interested in target structures within a certain range of the physical prospect $\mathcal{O}(\sigma)$. In the case of complex configuration spaces with many local minima, the more stringent expectations of the former search require a global-optimization method such as the reciprocal-space genetic algorithm presented here. In the latter case, a local-search method such as the virtual-atom approach, with its ability to directly focus on local minima, proves more than adequate, since the problem can be solved within the first hundred calls to the functional for systems up to 48 atoms. This puts such studies well within the range of more computationally intensive functionals like density functional theory.

This work was supported by DARPA, Defense Sciences Office, under NREL Contract No. DE-AC36-99GO10337. We acknowledge the use of the Cluster Expansion ATAT package [25] and the genetic-algorithm library EO [26].

[3] This result may be specific to the metallic alloy under consideration since, as shown in figure 1, a majority of the local minima fit the pertinent-structure target property requirement.

References

[1] Franceschetti A and Zunger A 1999 The inverse band-structure problem of finding an atomic configuration with given electronic properties *Nature* **402** 60
[2] Franceschetti A, Dudiy S V, Barabash S V, Zunger A, Xu J and van Schilfgaarde M 2006 First-principles combinatorial design of transition temperatures in multicomponent systems: the case of Mn in GaAs *Phys. Rev. Lett.* **97** 047202
[3] Dudiy S V and Zunger A 2006 Searching for alloy configurations with target physical properties: impurity design via a genetic algorithm inverse band structure approach *Phys. Rev. Lett.* **97** 46401
[4] Wang M, Hu X, Beratan D N and Yang W 2006 Designing molecules by optimizing potentials *J. Am. Chem. Soc.* **128** 3228
[5] Keinan S, Hu X, Beratan D N and Yang W 2006 Designing molecules with optimal properties using the linear combination of atomic potentials approach in an am1 semiempirical framework *J. Phys. Chem.* A **111** 146
[6] Wei S-H and Zunger A 1990 Band-gap narrowing in ordered and disordered semiconductor alloys *Appl. Phys. Lett.* **56** 662–4
[7] Zunger A and Mahajan S 1994 *Handbook on Semiconductors* vol 3 (Amsterdam: Elsevier) p 1399
[8] Dresselhaus M S, Chen G, Tang M Y, Yang R G, Lee H, Wang D Z, Ren Z F, Fleurial J-P and Gogna P 2007 New directions for low-dimensional thermoelectric materials *Adv. Mater.* **19** 1043
[9] Huang P, Menon S and de Fontaine D 1991 *J. Phase Equilib.* **12** 3
[10] Cheng C, Needs R J, Heine V and Jones L 1989 I. A theory of the origin of SiC polytypes *Phase Transit.* **16/17** 263
[11] Ducastelle F 1991 *Order and Phase Stability in Alloys (Cohesion and Structure)* (Amsterdam: North-Holland) chapter 4
[12] Kanamori J and Kakehashi Y 1977 *J. Physique Coll.* **38** C7 274
[13] Ferreira L G, Wei S-H and Zunger A 1991 Stability, electronic structure, and phase diagrams of novel inter-semiconductor compounds *Int. J. Supercomput. Appl.* **5** 34–56
[14] Deaven D M and Ho K M 1995 Molecular geometry optimization with a genetic algorithm *Phys. Rev. Lett.* **75** 288
[15] Stucke D P and Crespi V H 2003 Predictions of new crystalline states for assemblies of nanoparticles: perovskite analogues and 3-d arrays of self-assembled nanowires *Nano Lett.* **3** 1183
[16] Oganov A R and Glass C W 2006 Crystal structure prediction using *ab initio* evolutionary techniques: principles and applications *J. Chem. Phys.* **124** 244704
[17] Pickard C J and Needs R J 2006 High-pressure phases of silane *Phys. Rev. Lett.* **97** 45504
[18] Barabash S V, Blum V, Müller S and Zunger A 2006 Prediction of unusual stable ordered structures of Au–Pd alloys via a first-principles cluster expansion *Phys. Rev.* B **74** 35108
[19] Abraham N L and Probert M I J 2006 A periodic genetic algorithm with real-space representation for crystal structure and polymorph prediction *Phys. Rev.* B **73** 224104
[20] Johnston R L 2003 Evolving better nanoparticles: genetic algorithms for optimising cluster geometries *Dalton Trans.* 4193–207 doi:10.1039/b305686d
[21] Bendsøe M P and Kikuchi N 1988 Generating optimal topologies in structural design using a homogenization method *Comput. Methods Appl. Mech. Eng.* **71** 197–224
[22] Nordheim L 1931 The electron theory of metals *Ann. Phys.* **9** 607
[23] Barahona F 1982 On the computational complexity of Ising spin glass models *J. Phys. A: Math. Gen.* **15** 3241–53
[24] Istrail S 2000 Three-dimensionality and NP-completeness: I. universality of intractability of the partition functions of the Ising model across non-planar lattices *STOCOO: Proc. 32nd ACM Symp. on the Theory of Computing (Portland, OR, May 2000)* (New York: ACM) pp 87–96
[25] van de Walle A and Asta M 2002 Self-driven lattice-model Monte Carlo simulations of alloy thermodynamic properties and phase diagrams *Modelling Simul. Mater. Sci. Eng.* **10** 521–8
[26] Keijzer M, Merelo J J, Romero G and Schoenauer M 2001 Evolving objects: a general purpose evolutionary computation library *Artif. Evolution* 231–44
[27] de Gironcoli S, Giannozzi P and Baroni S 1990 *Phys. Rev. Lett.* **66** 2116
[28] Trimarchi G, Graf P and Zunger A 2006 Exploring the configurational space of binary alloys: different sampling for different cell-shapes *Phys. Rev.* B **74** 014204