Theor Chim Acta (1993) 85: 345–361

Theoretica Chimica Acta

Massively parallel direct SCF calculations on large metal clusters: Ni_5–Ni_{481}

Lars G. M. Pettersson[1] **and Torgny Faxen**[2]

[1] Institute of Theoretical Physics, University of Stockholm, Vanadisvägen 9, S-113 46 Stockholm, Sweden

[2] Alliant Computer Systems Corp, One Monarch Drive, Littleton, MA 01467, USA

Received June 9, 1992/Accepted September 22, 1992

Summary. The convergence of the cluster model with respect to excitation energies, ionization potentials and hydrogen chemisorption energy in the four-fold hollow site of the Ni(100) surface is studied for a sequence of cluster models from Ni_5 up to Ni_{181}. For the largest, Ni_{481}, cluster studied, only the structure of the occupied levels for one state is obtained. The concept of bond-preparation is found to be essential for the evaluation of chemisorption energies also for clusters with more than 100 atoms. The cluster excitation energies show a slow decrease such that even for Ni_{181} the step between the lower excited states is still 0.1–0.2 eV. The effect of *p*-functions on surrounding cluster atoms is found to be 3–4 kcal/mol independent of cluster-size. The direct SCF program DISCO was parallelized using the TCGMSG toolkit in order to perform the calculations. The easy strategy utilized is analysed and exhaustive timings on the Alliant Campus/800 MPP system with 200 CPU's are presented.

Key words: Cluster models – Ni_5–Ni_{481} – Excitation energies – MPP

1 Introduction

The cluster model of metal surfaces has become a very useful tool for quantum chemical studies of local interactions at the surface, such as atomic or molecular chemisorption, and many successful applications of this model have been made over the last few years. Essential to this success has been the understanding of the limitations of the model, such as edge effects and, in particular, the discreteness of the energy levels of the cluster as contrasted with the band structure of the solid and the effects this will have on results obtained from cluster model studies. The oscillations in computed chemisorption energies as function of cluster-size has been of particular concern and for a long time lead to the conclusion that chemisorption energies should not be computed using this model or that particular clusters should be selected for the calculations. This situation has been rectified through the work of Siegbahn, Wahlgren and coworkers [1–7] with the concept of bond-preparation [4], i.e. the cluster should, when it interacts with an adsorbate, be regarded as a molecule which must have singly occupied orbitals of the appropriate symmetry available to form the required number of bonds to the adsorbate. An alternative viewpoint which applies to a more general class of adsorbates is that the adsorbate should "fit into" the electronic structure of the cluster, i.e. no new cluster orbitals should be generated as a result of the chemisorption. For certain clusters these requirements are fulfilled already for the ground-state electronic structure, others may require an excitation to an excited state to form the bonding. This excitation, which is only required for the interacting system, enters the evaluation of the chemisorption energy for these clusters. Regarding the cluster as a model of the extended surface, where the excitation energy may be expected to be negligible, the effects of the discrete level structure may thus be reduced by evaluating the chemisorption energy relative to the already excited cluster. This approach has been very successful in explaining the origin of, and correcting for, cluster size oscillations for adsorption of atomic hydrogen [4], oxygen [5] and fluorine [6] and also molecules such as CH_3 and CH_2 [7].

One assumption in this approach to clusters as models of surfaces is that the required excitation energy really decreases as the clusters become substantially larger. As a result one would expect the effects of bond-preparation to become smaller and eventually disappear for sufficiently large clusters. An interesting question then becomes for what size cluster this can be expected to occur. The largest clusters previously studied have consisted of less than 100 atoms, but it would be desirable to do calculations for substantially larger clusters. In particular, for the largest cluster, Ni_{70}, in the study of fluorine chemisorption [6] some questions were raised as to the applicability of the concept of bond-preparation to very large clusters; bond-preparation seemed to lead to too large a chemisorption energy and perhaps for very large clusters sufficient flexibility in bondformation is already present making explicit bond-preparation unnecessary.

To enable *ab initio* calculations on clusters containing a few hundred atoms, direct techniques must be employed and furthermore some consideration of computation time has to be made, even though only low-angular momentum basis functions will be employed. For large metal clusters this entails a careful consideration of sparsity in the evaluation of integrals in order to reduce the scaling of the compute time with the size of the cluster. Furthermore, parallelization of the calculations must be considered.

Direct SCF techniques are particularly well-suited for parallelization since each Fock-matrix element may be written as a sum of contributions from several processes, each of which has access to a copy of the required density matrices and evaluates its part of the total set of two-electron integrals. The contribution from each process is then a partial (two-index) Fock-matrix. This approach has been used by several workers previously. Most notably Lüthi and Almlöf [8] who performed parallel calculations using the direct SCF program DISCO [9] on several CRAY computers situated in Switzerland and the USA, but also by Feyereisen and Kendall [10] and by Lüthi et al. [11], all using DISCO.

In the present work we will present a somewhat different approach to parallel calculations, which we feel substantially simplifies the parallelization of existing programs. The required number of processes are started up on the system, all processes go through all lines of code, but in the parallel parts the work is divided up between processes. The communication then is reduced to the summation of contributions to the desired matrices, while all processes generate their own starting data and perform the diagonalizations. The construction of a parallel version of DISCO using the TCGMSG [12] message passing toolkit in this approach took negligible time and mostly concerned the I/O handling. This approach is suitable for cases where the parallel sections strongly dominate the computation, since the serial parts will still be executed by all processes.

In the present work we will present our results for excitation energies, ionization potentials, electron affinities and the hydrogen chemisorption energy for clusters of size Ni_5–Ni_{181}. For the largest cluster studied, Ni_{481}, only the structure of the occupied levels of the assumed lowest state was obtained. We will also discuss the strategy employed for parallelization and present exhaustive timings for the Ni_{181} cluster using the Alliant Campus/800 Massively Parallel Processing (MPP) computer with 200 processors.

2 Methods

The clusters were constructed by starting with a Ni_5 C_{4v} cluster representing the four-fold hollow site of the Ni(100) surface. This cluster was then systematically extended by including all atoms generated by translations of one lattice step in the x- and y-directions, with the z-axis normal to the surface (Figs. 1 and 2). C_{4v} symmetry was maintained and used throughout except for states of E-symmetry where C_{2v} symmetry was used instead. This gives the sequence of clusters Ni_5, Ni_{25}, Ni_{61}, Ni_{113} and Ni_{181}. In addition, the substantially larger Ni_{481} cluster was generated according to the same scheme. Two even-numbered clusters, Ni_{66} and Ni_{118}, were constructed by adding an Ni_5 unit beneath the central four-fold

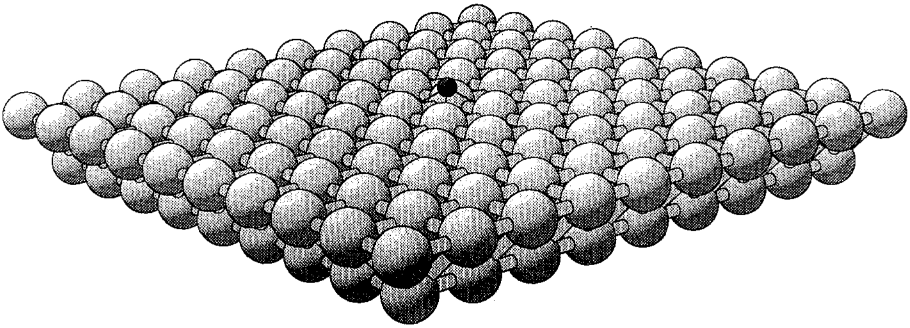

Fig. 1. The Ni_{181} cluster model of the Ni(100) surface with chemisorbed hydrogen in the central four-fold hollow site

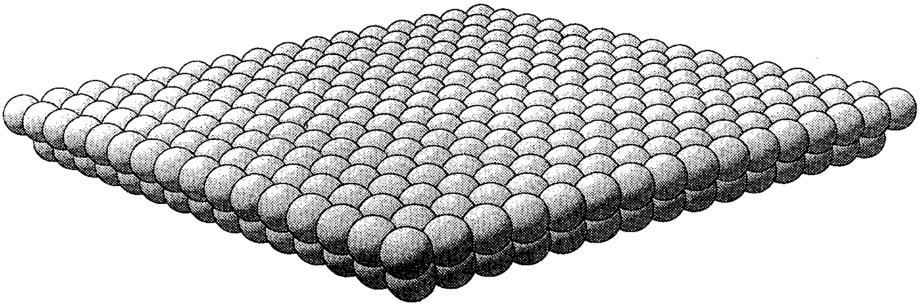

Fig. 2. The Ni_{481} cluster

hollow site, i.e. a four-layer cluster. For the smaller clusters, i.e. Ni_5–Ni_{181} the lowest excitations of each symmetry type were computed in addition to ionization potential and electron affinities. Finally, the chemisorption energy of atomic hydrogen at a fixed position 2.25 a_0 above the central four-fold hollow site was computed for these clusters. For Ni_{481} the orbital energies of the occupied levels were obtained to give a qualitative picture of the development of the density of states distribution.

The Ni atoms in the cluster were described using the relaxed core one-electron ECP description of Panas et al. [1] where the Ni 4*s* orbital is described by a (6*s*)/[2*s*] basis set. This basis was used for all atoms except for the five atoms in the central four-fold hollow site for which the basis was extended with a diffuse *p*-function ($\alpha = 0.112$). The hydrogen basis was the same as in [4], i.e. the (5*s*) basis of Huzinaga [13] extended with one diffuse *s*-function and two *p*-functions and contracted to [4*s* 2*p*].

The calculations have been performed at the SCF level of approximation using a modified version of the direct SCF program DISCO [9]. The modifications include an efficient implementation of effective core potential (ECP) integrals [14], modified prescreening of integrals, improved use of sparsity and diagonalization of a super-CI Hamiltonian [15] with Quasi-Newton extrapolations and line-search to obtain the improved orbitals. Furthermore, the evaluation of both two- and one-electron integrals was parallelized using the TCGMSG package [12] and a somewhat different strategy than previously applied. These modifications were essential to enable the performance of these calculations.

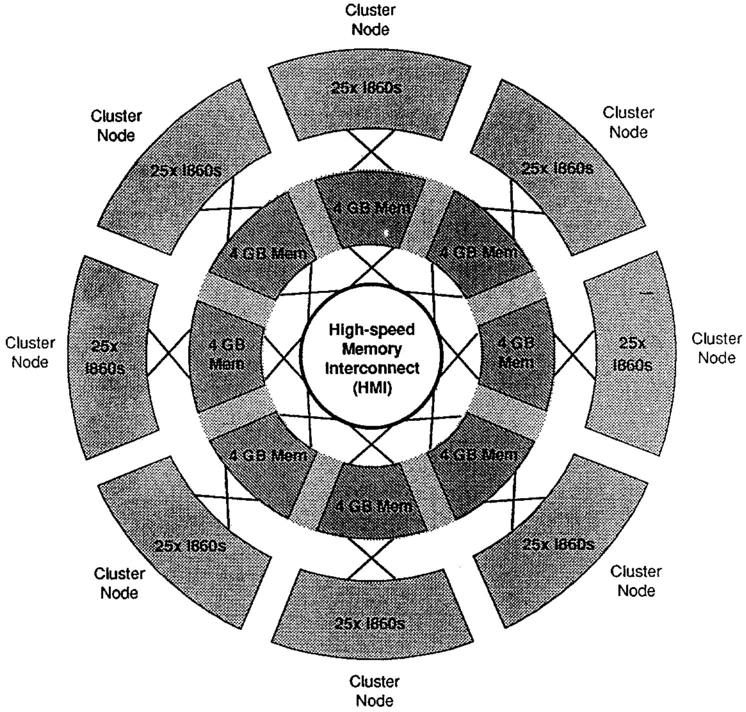

Fig. 3. The Alliant Campus/800 MPP architecture used in the calculations. The system used consisted of eight cluster-nodes each with 25 i860 processors and 0.5 Gbyte of memory

The calculations were performed on an alpha version of the Alliant CAMPUS/800 massively parallel computer (Fig. 3). The basic elements of the computer are the cluster-node composed of 25 i860 microprocessors sharing up to 4 GB of memory and a two-level switch for intra and inter cluster node communication. The first level switch supports symmetrical memory-access to all processors within the cluster-node, the second level switch connects cluster-node memories and is based on the High Performance Parallel Interface (HIPPI) standard.

The system we used for the calculations had 8 cluster-nodes, each with 25 processors, 500 MB of memory and 4 MB of global cache. Production runs were scheduled on 2 or more nodes, depending on the size of the problem. The Ni_{481} calculations were typically performed on 6 or 8 nodes which gave an iteration time of 3–2.5 hours.

A point of clarification. We use the term "cluster" for three purposes here which may cause some confusion. There are metal clusters of atoms, cluster-nodes in the CAMPUS system and clusters of processors in the individual cluster-nodes. We have tried to use "metal cluster", "cluster-nodes" and "processor cluster" where appropriate.

3 Results and discussion

The present work concerns both the cluster model for chemisorption as extended to substantially larger clusters than previously considered and an in-depth study of parallelization of the direct SCF method for massively parallel computers built on a number of nodes each with a large number of processors sharing a large memory. We will divide the discussion of our results into two sections beginning with our study of the cluster model for chemisorption in Sect. 3.1 and then discussing our approach to parallel calculations and our specific applications to the Alliant Campus/800 MPP system in Sect. 3.2.

3.1 Cluster results

The clusters studied have been selected in order to provide a sequence of clusters which converges smoothly in terms of structure towards the infinite system. Clearly the present clusters will converge towards a two-layer infinite slab and as such should not be compared to real gas-phase clusters, for which the structures will be quite different. The computed excitation energies and hydrogen chemisorption energies should rather be compared with results from surface studies and the present study is thus a study of the performance of the cluster model as such and in particular the applicability of the concept of bond-preparation as the cluster models become very large. It should also be noted that dynamical correlation, which has not been included in the present study, makes important contributions to the actual chemisorption energies and also tends to diminish cluster oscillations compared with results at the SCF-level [2].

The electron occupations of the clusters were generated according to the aufbau principle based on the orbital energies of the partially occupied clusters. This was in general done using an automatic aufbau, but each time performing a few iterations to stabilize the solution before introducing additional electrons. The open-shell electrons were then distributed to give the different, high-spin, excited states shown in Table 1. In particular for the larger clusters the

Table 1. Excitation energies (ΔE in eV) for Ni_x and $Ni_x + H$ and hydrogen chemisorption energy (D_e) in kcal/mol for different size clusters. Chemisorption energies computed relative to ground state of the Ni_x cluster are given in parentheses when different from result from bondprepared state

	Cluster				$Ni_x + H$[a]	
Cluster	Closed shell occupation[b]	Open shells	State	ΔE	ΔE	D_e
Ni_5	(1, 0, 0, 0, 1, 0)[c]	a_1	$^2A_1$	4.03		
	(1, 0, 0, 0, 0, 0)	b_1e^2	$^4B_2$	2.56		
	(1, 0, 0, 0, 0, 0)	b_2e^2	$^4B_1$	1.51		
	(1, 0, 0, 0, 1, 0)[c]	e	^2E	0.39		
	(1, 0, 0, 0, 0, 0)	a_1e^2	$^4A_2$	0.00	0.00	34.0
Ni_{25}	(4, 0, 1, 1, 3, 3)	a_1	$^2A_1$	1.85		
	(4, 0, 1, 1, 3, 3)	a_2	$^2A_2$	1.59		
	(4, 0, 1, 1, 3, 3)	b_1	$^2B_1$	1.54	1.11	
	(3, 0, 1, 2, 3, 3)	a_2	$^2A_2$	0.93	1.28	
	(4, 0, 1, 1, 3, 3)[c]	e	^2E	0.53	1.40	
	(4, 0, 1, 1, 3, 3)	b_2	$^2B_2$	0.26	2.07	
	(3, 0, 1, 1, 3, 3)	a_1e^2	$^4A_2$	0.13	0.24	
	(3, 0, 1, 2, 3, 3)	a_1	$^2A_1$	0.01	0.00	29.1 (28.9)
	(3, 0, 1, 2, 3, 3)[c]	e	^2E	0.00	1.29	
Ni_{61}	(7, 1, 4, 4, 7, 7)	b_1	$^2B_1$	2.18	2.80	
	(7, 1, 4, 4, 7, 7)	a_1e^2	$^4A_2$	1.72	2.29	
	(7, 1, 4, 4, 7, 7)	a_1	$^2A_1$	1.54	1.80	
	(7, 1, 3, 4, 7, 7)	$a_1a_2b_1$	$^4B_2$	1.30	0.40	
	(7, 1, 4, 4, 7, 7)	a_2	$^2A_2$	0.39	1.06	
	(7, 1, 3, 4, 7, 7)[c]	a_1a_2e	^4E	0.37	0.00	54.9 (46.4)
	(7, 1, 4, 4, 7, 7)	b_2	$^2B_2$	0.29	0.94	
	(7, 1, 3, 4, 7, 7)	b_1e^2	$^4B_2$	0.06	0.09	
	(7, 1, 4, 4, 7, 7)	e	^2E	0.00	0.40	
Ni_{66}	(8, 2, 3, 4, 8, 8)	—	$^1A_1$	2.02	2.05	
	(8, 1, 3, 4, 8, 8)	a_1b_1	$^3B_1$	1.49	0.95	
	(8, 1, 3, 4, 8, 8)	a_2b_1	$^3B_2$	0.78	0.77	
	(7, 1, 4, 4, 8, 8)	a_1a_2	$^3A_2$	0.69	0.00	46.2 (30.3)
	(8, 1, 3, 4, 8, 8)	b_1b_2	$^3A_2$	0.61	0.56	
	(8, 1, 4, 4, 8, 8)	—	$^1A_1$	0.00	0.01	
Ni_{113}	(10, 3, 6, 6, 15, 15)	$a_1a_2b_2$	$^4B_1$	1.65		
	(10, 3, 6, 7, 14, 14)	$a_1a_2b_2e^2$	$^6B_2$	0.86		
	(10, 4, 6, 7, 14, 14)	a_1b_2e	^4E	0.58	0.55	
	(11, 4, 6, 7, 14, 14)	a_2	$^2A_2$	0.57	0.78	
	(11, 4, 6, 7, 14, 14)	a_1	$^2A_1$	0.56	0.00	52.1 (39.2)
	(11, 4, 6, 7, 14, 14)	b_1	$^2B_1$	0.55	0.65	
	(11, 4, 6, 7, 14, 14)	b_2	$^2B_2$	0.37	0.57	
	(11, 4, 6, 7, 14, 14)[c]	e	^2E	0.00	0.15	
Ni_{118}	(13, 4, 6, 7, 14, 14)	e^2	$^3A_2$	0.44	0.67	
	(12, 4, 6, 7, 14, 14)	$a_1a_2e^2$	$^5A_1$	0.31	0.30	
	(12, 4, 6, 7, 14, 14)	$a_1b_1e^2$	$^5A_2$	0.22	0.28	
	(12, 4, 6, 7, 14, 14)	$a_1b_2e^2$	$^5B_1$	0.00	0.00	48.2
Ni_{181}	(15, 7, 10, 12, 23, 23)	b_2	$^2B_2$	0.46	0.86	
	(15, 7, 10, 12, 23, 23)	a_2	$^2A_2$	0.32	0.79	
	(15, 7, 10, 12, 23, 23)	b_1	$^2B_1$	0.17	0.62	
	(15, 7, 10, 12, 23, 23)	a_1	$^2A_1$	0.00	0.00	48.4
Ni_{481}	(38, 23, 29, 30, 60, 60)	a_1	$^2A_1$			

[a] Electronic structure of $Ni_x + H$ obtained by adding one a_1 orbital to cluster structure. The spin is coupled low-spin to open a_1 cluster-orbital if existing, otherwise high-spin coupled to cluster open orbitals

[b] C_{4v} closed-shell occupations given as $(a_1, a_2, b_1, b_2, e, e)$

[c] Calculation done in C_{2v} symmetry

possibilities to consider many different distributions of the closed-shell electrons in addition to the open-shell ones are greatly diminished and it is difficult to ascertain that the true ground state of the cluster has been found. However, the computed excitation energies are still expected to reflect the level-structure of the respective clusters.

The general picture obtained from a study of the cluster excitation energies in the Table is the expected decrease with size of the cluster both in terms of maximum difference and in terms of the steps between levels. For the smallest cluster, Ni_5, the ground state at the SCF-level is the $^4A_2$ state with the lowest excited ^2E state 0.39 eV above. The remaining states considered are all more than 1.5 eV higher in energy. For the largest cluster that was possible to study in detail, Ni_{181}, the highest excited state of the four states considered is the $^2B_2$ state which is found 0.46 eV above the lowest $^2A_1$ state. Here the difference between states has thus decreased to of the order of 0.15 eV.

For the Ni_{25} cluster two different occupations of the closed shells were considered as shown in the Table. For the lowest closed-shell occupation there are two different distributions of open-shell electrons that give states that lie within 0.13 eV of the lowest state found. In addition, the next closed-shell occupation that is obtained by moving two electrons from the b_2 symmetry to the a_1 symmetry gives rise to states lying between 0.26 and 1.8 eV from the lowest state. There are thus available several states that are low in energy compared with the lowest $^2A_1$ and ^2E states.

The Ni_{61} cluster shows a spread in energy levels that is very similar to that found for Ni_{25}. Here the lowest state found is a ^2E state with a very close-lying $^4B_2$ state at 0.06 eV and with additional excited states some 0.2–0.3 eV above. As is the case also for the similar Ni_{66} cluster, the highest state obtained from redistributing the open shells with fixed closed-shell occupations is about 2 eV higher in energy.

Increasing the size of the clusters by adding another shell of nearest neighbors to generate the Ni_{113} and Ni_{118} clusters gives a substantial reduction in excitation energies. We now find maximum excitation energies for redistributions of open-shell electrons of the order of 0.4–0.6 eV with the lowest excited state 0.2 eV above the lowest state found and with the remaining excited states in the Table very similar in energy. When the cluster is then increased to give Ni_{181}, the results are surprisingly equal to the situation for Ni_{113} and Ni_{118} with a spread in excitation energies of 0.5 eV, which is surprisingly large for such a big cluster. From the point of view of bond-preparation one could thus still expect sizable effects on the computed chemisorption energies. In cases where the bonding state (with an open cluster a_1 orbital) is not the lowest, the effect of bond-preparation would be to increase the computed chemisorption energy by the excitation energy which could then amount to as much as 5–10 kcal/mol even for clusters as large as 200 atoms.

The interaction between the hydrogen atom and the cluster splits the cluster-levels further (Table 1). For the Ni_{61}, Ni_{66} and Ni_{113} clusters the lowest cluster states found do not include an open a_1 orbital and in computing the chemisorption energies we thus consider both the lowest cluster state to lowest chemisorption state (values in parentheses) and the bond-prepared cluster state. The present values have been obtained with hydrogen at a fixed distance above the surface and do not include correlation. In earlier work [2] the height above the surface of the hydrogen atom has been found to vary between 1.4 and 2.5 a_0 with cluster-size when the valence electrons on cluster and hydrogen are explicitly treated and correlated. Dynamical correlation and polarization of the 3*d* "core" orbitals through the use of a core-polarization potential [16, 17] or explicit correlation [18] gives an additional decrease of the distance above the surface by some 0.5 a_0. The distance used in the present work then represents a value in the range obtained for smaller clusters (without correlation) and was selected with the expectation that for very large clusters the height above the surface would become insensitive to the number of atoms in the cluster.

The average effect of correlation on the hydrogen chemisorption energy in the size-range Ni_5 to Ni_{50} has been found [2] to be 15.8 kcal/mol with a standard deviation of 1.3 kcal/mol, where ten different clusters have been included. Since the experimental value for the hydrogen chemisorption energy on the Ni(100) surface is 63 kcal/mol [19] the SCF limit with respect to cluster-size would then be expected to be 48 kcal/mol under the assumption that the correlation effects can be regarded as constant. There are some systematic deviations included in the error bars for the correlation energy contribution that should be pointed out here. In [2] it was shown that the correlation energy contribution to the chemisorption energy in the case of hydrogen chemisorption depends on the type of bonding, i.e. the admixture of ionic and covalent components in each particular case. Clusters with a small SCF-level chemisorption energy due to longer bond-distance and weaker covalency were shown to have larger correlation energy contributions to the chemisorption energy, while for clusters where hydrogen comes closer to the surface and is more covalently bound the effects of correlation are smaller. In view of these results it is thus likely that small deviations (a few kcal/mol) in the SCF-level chemisorption energy from the expected 48 kcal/mol, due to geometry optimization or basis set extensions, may be coupled with a corresponding increase or decrease in correlation energy. In the present work we do not explicitly consider correlation and will thus limit our discussion to the SCF limit with respect to cluster-size for the chemisorption energy which we will take as 48 kcal/mol.

The values obtained in the present work show a fair convergence towards the expected result. Bond-preparation is required in three cases and in each case gives a substantial improvement of the predicted chemisorption energy, even though the resulting values for Ni_{61} and Ni_{113} become somewhat too high compared with the expected SCF limit. The values for both Ni_{118} and Ni_{181} are the expected 48 kcal/mol, but in view of the value for the Ni_{113} cluster it is difficult to claim that this represents a converged result. However, it is very clear that the fluctuations around the final results are much smaller for these larger clusters, and that bond-preparation is still essential even for clusters in the size-range 61–113 atoms.

The computed (ΔSCF) ionization potentials and electron affinities for the odd-numbered clusters are given in Table 2. These may be compared with the experimental photoelectric workfunction of 5.15 ± 0.1 eV [20]. The contribution to the cluster IP from valence correlation and core-valence correlation has been investigated for smaller nickel clusters in the size-range four to nine atoms [21]. Excluding the Ni_4 cluster, for which the results seem more uncertain, the average value for these clusters for the valence correlation contribution to the IP was 1.08 eV. The core-valence correlation was estimated using a core-polarization potential [16, 17] and was found to give an average contribution of 0.61 eV to the IP. All resulting values were lower than experiment by between 0.1 and 0.5 eV. Thus, if these contributions are assumed transferrable to the present clusters the resulting IP's would be around 6 eV which is somewhat high compared with the

Table 2. Computed (ΔSCF) ionization potential (IP) and electron affinities (EA) for some Ni clusters. All energies in eV

Cluster	IP	EA
Ni_{25}	3.71	1.07
Ni_{61}	4.18	2.31
Ni_{113}	4.05	2.14
Ni_{181}	4.24	3.01

experimental value. The clusters for which the correlation effects have been computed in [21] are rather small compared with the present systems. It does not seem unlikely that the effects on the IP of, in particular, the valence-correlation should decrease with size of the cluster and that the assumed corrections for correlation effects thus are too high. Using the present approach these effects could in principle be investigated at the MP2 level also for the larger clusters included here. This is, however, beyond the scope of the present work and will be left for future studies.

The present clusters basically only contain s-functions in their description. For the case of describing chemisorption into the four-fold hollow site it has been demonstrated [2] that the effect of including p-functions on more than the central, directly interacting atoms is small or of the order of less than 1 kcal/mol on the computed chemisorption energy. However, the inclusion of p-functions on the remaining cluster atoms is essential if the bonding in the cluster is to be described. In fact, for the present clusters the only cluster with a positive atomization energy is the Ni_5 cluster where all the atoms have a polarizing p-function in addition to the normal s-basis. The inclusion of p-functions in the basis can also be expected to affect the level-spacing and thus the computed excitation energies. This effect was investigated for the Ni_{25} cluster (Table 3) by computing the excitation energies using three different cluster descriptions; no p-function on any atoms, a diffuse p-function on the five central atoms and on all atoms, respectively. Substantial effects are found for the cluster atomization energies, which were computed to be -5.3 eV, -1.5 eV and $+2.6$ eV given in the same order as above.

Table 3. Effect of p-functions on Ni_{25} cluster excitation energies (eV)

		Number of centers with p-functions		
Occupation[a]	State	0	5	25
(4, 0, 1, 1, 3, 3)	$^2B_1$	1.05	1.53	1.86
(3, 0, 1, 2, 3, 3)	$^2A_2$	0.98	0.92	1.17
(4, 0, 1, 1, 3, 3)	^2E	-0.56	0.52	0.69
(4, 0, 1, 1, 3, 3)	$^2B_2$	0.10	0.25	0.08
(3, 0, 1, 1, 3, 3)	$^4A_2$	-1.50	0.12	0.57
(3, 0, 1, 2, 3, 3)	^2E	-0.69	-0.01	0.46
(3, 0, 1, 2, 3, 3)	$^2A_1$	0.00	0.00	0.00

[a] C_{4v} closed-shell occupations given as $(a_1, a_2, b_1, b_2, e, e)$

Addition of p-functions on the central five atoms has a very large effect on the spectrum compared with the case with only s-functions present. Without p-functions the ground state is found to be a $^4A_2$ state which is 1.5 eV lower in energy than the ^2E state that becomes the ground state upon addition of p-functions. Other shifts of up to 1 eV in the energy levels are evident in the Table. The main effects are obtained already with polarization functions only on the central atoms and only smaller shifts are found from adding a diffuse p-function on the remaining 20 atoms in the cluster. However, the changes are still of the order of 0.2–0.5 eV such that this could indeed have an effect on the computed chemisorption energies. Computing the chemisorption energy for the Ni_{25} cluster with diffuse p-functions on all centers and hydrogen at the fixed 2.25 a_0 above the surface we obtain a chemisorption energy of 34.7 kcal/mol, i.e. an increase of 5.6 kcal/mol over the results with p-functions only on the central atoms. A small part (1 kcal/mol) of this effect may be ascribed to increased linear dependencies in the basis, but the remaining 4.5 kcal/mol effect is substantially larger than the 0.2 kcal/mol reported previously [2] for a Ni_{13} cluster in the case of oxygen chemisorption.

In the case of the Ni_{25} cluster the optimized height above the surface was obtained in [2] as 1.37 a_0 and with an SCF-level chemisorption energy of 35.3 kcal/mol. These results were obtained using a somewhat different description of the central five atoms for which a frozen 3s orbital was included in the ECP description, while the remaining cluster was described using the same ECP model as in the present work. Using this description to compute the interaction energy at 2.25 a_0 height above the surface, a binding energy of 27.8 kcal/mol is obtained compared with the 29.1 obtained in the present work without frozen 3s orbitals. At 1.37 a_0 an interaction energy of 32.6 kcal/mol is obtained compared with 35.3 from [2]. Thus, the present description and that of [2] give very similar results for the chemisorption energy, but with a somewhat different distance dependence. Adding p-functions to all centers gives a chemisorption energy of 40.4 kcal/mol at 1.37 a_0. The results are thus dependent on the size of the cluster basis set to a larger degree than previously noted. In order to investigate this effect further, we performed additional calculations on the Ni_{66} cluster with p-functions on all centers. The computed hydrogen chemisorption energy in this case was 49.5 kcal/mol at 2.25 a_0 above the surface compared with 46.2 kcal/mol in the standard cluster model. The sensitivity to the description of the surrounding atoms is thus larger than previously assumed, but there does not seem to be any dependence on cluster-size. However, an effect of the order of 3–4 kcal/mol on chemisorption energies must be assumed from expansions of the cluster basis set. These investigations were done for the bond-prepared states, but the results are expected to have a more general applicability.

Finally, we present in Fig. 4 a plot of the occupied orbital energies as function of cluster-size beginning with Ni_{25} and going up to the largest cluster, Ni_{481}, studied in the present work. It should be noted that since only the central five atoms here have p-functions this really only gives a very qualitative picture of how the density of levels increases with cluster-size. However, it may be noted that the width of the "s-p-band" comes out remarkably constant with only small increases in going from Ni_{113} to Ni_{481}. For smaller clusters similar results have been found in e.g. [22] in calculations on Cu_n clusters up to Cu_{34}. For the largest cluster in the present work we find the highest occupied level (the cluster "Fermi-level") around −3.9 eV (Ni_{481}) and the lowest level at −15.2 eV giving a "width" of 11.3 eV. This can be compared with the computed s-p-band width

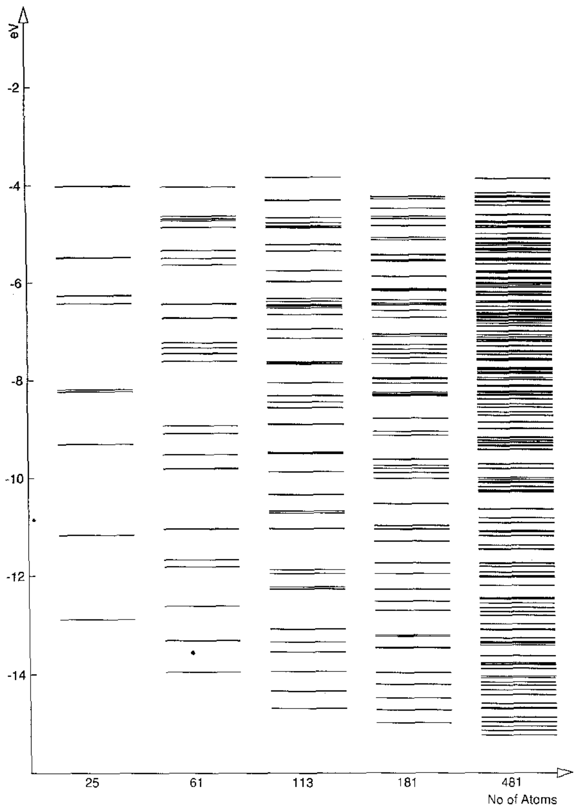

Fig. 4. Orbital energy level diagrams (in eV) for the occupied orbitals of clusters Ni_{25} to Ni_{481}

of 9.2 eV from band structure calculations [23]. The present SCF-level s-band thus overestimates the width of the band by about 2 eV, but the approach to a band structure in the form of an increased density of occupied levels is nicely illustrated by the present calculations.

3.2 Parallelization

The structure of the direct SCF method is extremely well suited for parallel processing and with the present approach of allowing all processes to each perform all of the serial parts of the code as well as its part of the parallel work the actual code-changes required turned out to be quite few and straightforward. This approach may be seen as a waste of compute cycles but as long as the parallel section of the code accounts for a large percentage (90–95%) of the compute time this approach significantly simplifies the whole porting procedure without losing any performance. In fact it may actually improve performance since communication between the processes is kept at a minimum. We believe that simplicity in porting to an MPP system is vital, it should be a matter of days rather than weeks to get a code up and running. The final tuning can and will probably take longer time but that is another issue.

We used the TCGMSC library from Argonne [12] to control, synchronize and send messages between the parallel tasks. The use of the library was straightforward, easy and the performance was very robust. Loadbalancing was achieved through use of the "nxtval" function in TCGMSG. This function simulates a shared counter by communicating with a dedicated server process. Using this simple approach we achieved an easy and efficient way of distributing the tasks with fully adequate loadbalancing. To illustrate the simplicity in parallelizing this code using the present approach we show the modified loop structure for the two-electron integral part in Table 4. Since each process generates all of its local data based on the input there is no message passing before starting the parallel part. Each process obtains its unique value of the linearized loop counter from the shared counter through the function nxtval, enters the loop structure and performs its assigned part of the work. The size of each block may vary strongly. In particular, the test on the density matrix elements may eliminate the block of integrals entirely whereupon a new value

Table 4. Parallel construction of the Fock matrix using TCGMSG. Code modifications for parallelism in italics. All steps executed by all processes

```
1. Construct Density Matrices
2. Zero Out Fock Matrices
C*** Get number of processes
3. n_nodes = nnodes()
C*** Set linear loop-count
4. icount = 0
C*** Get next available value from shared counter
5. next = nxtval(n_nodes)
C*** Start Loops
  Loop over Atoms I
    Loop over Atoms J
      Loop over Atoms K
        Loop over Atoms L
          Loop over Shells i on Atom I
            Loop over Shells $j_s$ on Atom J
              Loop over Shells $k_s$ on Atom K
                Loop over Shells $l_s$ on Atom L
                  if(next. eq. icount) then
                    Loop over Symmetry Operator $\hat{R}$
                      Loop over Symmetry Operator $\hat{S}$
                        Loop over Symmetry Operator $\hat{T}$
                          CALL DTST(DENTST,...)
                          IF (DENTST. GE. TOLS) THEN
                            CALL INT2(FOCK,DENS,...)
                          ENDIF
                    End Loops over Symmetry Operators
                    next = nxtval(n_nodes)
                  endif
                  icount = icount + 1
  End All Loops
C*** Terminate global counter block
6. junk = nxtval(-n_nodes)
C*** Sum Fock-matrix contributions over processes
7. call dgop(1 + msgdbl, fock,n,'+',work)
8. Add one-electron contribution
9. Diagonalize
```

will be requested immediately. At the end of the parallel section the global shared counter is turned off and the summation of the partial matrices is performed through the dgop routine.

To be able to perform calculations on clusters all the way up to the Ni_{481} case three requirements had to be met:

(a) CPU-power. From smaller runs we estimated that each iteration would take about 2–3 hours for the largest metal cluster and this was considered to be adequate.

(b) Communication. After each iteration a global summation of the FOCK-matrix is done. All the nodes communicate through the "dgop" call in TCGMSG that will add all the contributions from each process and return the total value. To make this more efficient "dgop" was rewritten to make a local sum over each cluster-node and then a global sum over all cluster-nodes, which proved to be very efficient and in fact can be viewed as a parallelization also of the summation since this then was distributed over all eight nodes simultaneously. About 16 MB of data was shared between all the nodes in the largest calculation. Using the rewritten "dgop" together with the high speed memory interconnect over HIPPI, and the fact that we used a two-level parallel structure (local looplevel/global tasklevel, see below) communication turned out to be almost negligible. It should be noted that even using an ethernet connection the communication did not present any problems.

(c) Memory requirements. In the largest calculations each process would need between 70 and 110 MB of real memory. To avoid excessive paging and swapping we chose to use a two-level parallel structure for DISCO:

- The lower parallel structure was implemented through compiler-generated concurrency using loop-level parallelism in the code. This was only done in the two-electron part of the code which traditionally is responsible for almost all of the compute time. Each process would run on a processor cluster with 4 processors. This type of parallelism usually works excellently up to 4 processors and has the advantage that the memory requirements usually stay close to the 1 processor runs. The speedup observed was typically in the range of 1.7–2.0. We expect that this could be tuned up to a factor of 3.0 with some more effort.
- The second higher level of parallelism was implemented through the use of the TCGMSG library by dividing large parallel tasks in DISCO into smaller tasks. Each smaller task would then run on a processor cluster of 4 processors and communicate with other tasks through message-passing via shared-memory or sockets. Totally this meant that we could have a maximum of 6 processes on each cluster-node, giving a total of 48 processes on the whole CAMPUS system, all running almost entirely in real memory doing minimal paging/swapping and at the same time minimizing communication overhead.

This approach is actually very attractive since it reduces significantly memory requirements and communication, both of which will be major bottlenecks for any MPP system.

Table 5 shows the times from running DISCO on the Alliant MPP system as a function of the number of processor-clusters we were using. The times are wall clock times for the whole run as well as for the one- and two-electron parts of the code. Since we deal with wall clock, time for communication, I/O and serial parts are included. We should also remember that each cluster-node has 6 processor-clusters with 4 processors in each cluster, so 8 cluster-nodes means a

Table 5. Wall clock times for DISCO in seconds for three iterations on the $\mathrm{Ni}_{181}\ {}^2A_1$ cluster with 377 basis functions. Different configurations and number of nodes (*n*) of the MPP system used with either detached CPU's (*p*) or clusters (*c*) of four CPU's assigned to each process. The one-electron part (1-EL) is executed once and the two-electron integrals (2-EL) computed three times. Each process requires about 15 MB of memory and writes about 4 MB of disk files. C_{2v} symmetry used

Config:	1*p*	1*c*	3*c*	12*p* (det)[a]	6*c*	24*p* (det)[a]	12*c*	24*c*	48*c*	CRAY X-MP
Equals:		(4*p*)		(3*c*)	(1*n*)	(6*c*)	(2*n*)	(4*n*)	(8*n*)	
1-EL:	8562	8448	2980	741	1586	511	755	412	293	965
2-EL:[b]										
//-part:	160511	96750	33867	19231	23383	14430	11716	5728	3313	
Comm:	—	—	24	27	27	27	33	33	36	
Serial:	165	168	156	180	236	236	225	216	237	
Total:	160676	96918	34047	19438	23646	14693	11974	5977	3586	14239
Total runtime										
Wall:	170571	105875	38207	21961	25426	15818	12955	6593	4153	15242
Speedup	1.0	1.6	4.5	7.8	6.7	10.8	13.2	25.9	41.1	

[a] Detached processor (det)

[b] Separate timings for the parallel region, communication (summation of Fock-matrices over processes), and serial parts (diagonalisation)

total of 48 processor-clusters and a total of 192 processors. The metal cluster used for the timings was the Ni_{181} cluster with diffuse p-functions on the central four-fold hollow atoms using three iterations only. The number of contracted functions was 377.

As we can see we have a very good parallel speedup both for the two-electron and one-electron parts as well as for the program as a whole. In particular we notice that

(a) For a single node we have an almost linear speedup going from 1 to 3 clusters but this tails off when we go from 3 to 6 clusters. The reason for this is mainly that we start seeing the effect of a shared memory and a shared cache that the different processes compete for. The advantage of the 2 level memory hierarchy in the CAMPUS system is clearly seen when we go from 6 to 12 clusters (1 to 2 nodes) where we again get an almost linear speedup. This speedup continues to be almost linear until we reach 8 nodes.

For a single node we also did the run using all 24 detached processors to get an idea of the maximum speedup we could get. As we can see we lose about a factor of 1.6 by using the processor clusters instead of detached processors. This is the price we have to pay to be able to run large calculations with large memory requirements.

(b) From 4 to 8 nodes the parallel speedup starts to decline. What we now see is Amdahls law in two different ways. The first is a loadbalancing effect where a few processes continue to run when most of the processes have finished. This is due to the granularity of the tasks. They can be made smaller but then communication overhead and waiting time will increase instead. The second way is a pure Amdahls law effect where the serial portions (communication, I/O, serial parts of the code) finally start to come into play. For a larger metal cluster this effect would have occurred later, i.e. for a larger number of processor-clusters.

In the Table we also have timings from making the same run on a single Cray X-MP/416 processor (9 ns clock). The performance obtained immediately on the CAMPUS/800 was 3.7 times that of the Cray X-MP, but at this point no particular tuning of the program had been performed.

Further tuning on a single node gave the numbers shown in Table 6. The performance improved substantially due to compiler-improvements, further

Table 6. Wall clock times for three iterations on the Ni_{181} cluster in C_{2v} symmetry. Effect of cache and further optimization. Two different processor types are used, p = model 400 processor (4 on a board) or t = model 200 processor (2 on a board + onboard cache)

Config:	1*p*	12*p* (det)	24*p* (det)	12*t* (det)	CRAY X-MP
One-electron part					
// Part	7710	698	464	556	965
Two-electron part					
// Part	144085	13515	12484	10994	
Comm		27	30	30	
Serial	165	180	243	200	
Total	144250	13722	12757	11224	14239
Total run	151960	14430	13231	11790	15242
Speedup	1.0	10.5	11.5	12.9	

tuning and by allotting larger blocks of work for each call to the global counter. The latter represents simply moving the test on *next* and *icount* in Fig. 3 up a few levels to the loop over basis functions *i*. This time all the processors ran in detached mode. By this we mean that all the processors run a separate thread, not using any loop-level concurrency. Further tuning for concurrency would be expected to improve the cluster times as well.

Three different runs were made using 12 and 24 detached processors of the same type as previously and with 12 detached processors using a different processor-board with two i860 + 256kb board-cache per processor. This board is called model 200. The earlier runs were made on a processor-board with 4 i860 processors and no board-cache, called model 400.

As we can see from the results we are getting a very good speedup up to 12 processors while we do not gain much from 12 to 24 processors. We believe the reason for this is that with more processes we introduce higher overhead from the operating system, they also all do some I/O and they all compete for the same cache and memory system. When all this is summed up the net gain is quite small.

The timings for the model 200 board, with the on-board cache, turns out to be the best of them all. Since this is a highly vectorizable code it makes good use of the on-board cache.

When we start to spread the code over multiple nodes we should get a speedup over nodes that is very similar to the one from Table 5. We can make this statement since the number of nodes is still rather small, the non-parallel parts are still very small and the fact that each node has its own memory system.

A simple extrapolation gives at hand that this should give a total runtime of about 1800 seconds over 8 nodes or roughly 8.5 times faster than a single Cray-XMP processor.

4 Conclusions

The computed excitation energies for the sequence of cluster models studied do decrease as the clusters get larger, but somewhat more slowly than expected; even for the largest cluster studied in detail, Ni_{181}, the lowest excitation energy is at 0.17 eV and with additional states at about 0.1–0.2 eV intervals. The hydrogen chemisorption energies do become more stable for the larger clusters and seem to converge to the expected SCF-limit of 48 kcal/mol. However, bond-preparation remains important even for the large clusters studied in the present work. The computed chemisorption energies for the Ni_{61}, Ni_{66} and Ni_{113} cluster models are all improved by bond-preparation. For the Ni_{113} cluster the effect is still as large as 13 kcal/mol so that this remains an important concept also for this size cluster models.

The computed chemisorption energies are shown to be more sensitive than previously assumed to the description of the surrounding atoms. For both the Ni_{25} and the Ni_{66} clusters the chemisorption energy is increased by 3–4 kcal/mol by including *p*-functions on all centers rather than just on the atoms in the central four-fold hollow site. This seems to be a constant effect that does not depend on the size of the cluster, however.

The large number of calculations on very large cluster models that have been performed in the present work would not have been possible without utilizing the possibility to parallelize the calculations on an MPP system and thus obtaining a substantial improvement in job turn-around. The port of DISCO to the Alliant CAMPUS MPP computer was done in a straightforward manner with very few actual changes in the code, using the TCGMSG library. Since the computer actually resembles many existing computer installations (clusters of shared-memory computers) we believe that this is a valid port for many other computer installations also.

Actually running the largest metal cluster of this type ever done also puts focus on things we believe will be important for all large applications to be run in production work at any MPP system. The need is for processor power, memory and fast communication. Most MPP systems will have abundant processor power but real applications will be limited by the amount of memory and the speed of communications available. Using an architecture with a two-level memory system, shared and distributed, the whole code could be divided into a set of parallel tasks where each task in its turn ran on 4 processors, using loop-level parallelism. This offers an interesting solution where both the memory and the communication requirements can be significantly reduced (in this case with a factor of 4) and thus allowing the port to be done in a straightforward manner without having to substantially rewrite the code.

Acknowledgment. We thank Prof. K. Hermann for a copy of the Balsac plot program and assistance in producing Figs. 1 and 2.

References

1. Panas I, Siegbahn P, Wahlgren U (1987) Chem Phys 112:325
2. Panas I, Siegbahn P, Wahlgren U (1988) Theor Chim Acta 74:167
3. Wahlgren U, Siegbahn P (1991) Quantum Chemical Models of Chemisorption on Metal Surfaces, in: Salahub D (ed) Metal-ligand interactions: From atoms, to clusters, to surfaces. Kluwer, Dordrecht
4. Panas I, Schüle J, Siegbahn P, Wahlgren U (1988) Chem Phys Lett 149:265
5. Siegbahn PEM, Wahlgren U (1992) Int J Quantum Chem 42:1149
6. Siegbahn PEM, Pettersson LGM, Wahlgren U (1991) J Chem Phys 94:4024
7. Schüle J, Siegbahn P, Wahlgren U (1988) J Chem Phys 89:6982
8. Lüthi HP, Almlöf J (1993) Theor Chim Acta 84:443
9. Almlöf J, Faegri Jr K, Feyereisen M, Korsell K, DISCO a Direct SCF and MP2 code; for reference see e.g. Almlöf J, Faegri K, Korsell K (1982) J Comp Chem 3:385
10. Feyereisen M, Kendall RA (1993) Theor Chim Acta 84:289
11. Lüthi HP, Mertz JE, Feyereisen MW, Almlöf JE (1992) J Comp Chem 13:160
12. Harrison RJ (1991) Int J Quantum Chem 40:847
13. Huzinaga S (1965) J Chem Phys 42:1293
14. Pettersson LGM, Wahlgren U, Gropen O (1987) J Chem Phys 86:2176
15. Roos BO (1980) Int J Quantum Chem S14:175
16. Müller W, Flesch J, Meyer W (1984) J Chem Phys 80:3297
17. Pettersson LGM, Åkeby H (1991) J Chem Phys 94:2968
18. Åkeby H, Pettersson LGM (1991) Chem Phys 155:197
19. Ertl G (1979) in: Rhodin TN, Ertl G (eds) The nature of the surface chemical bond. North-Holland, Amsterdam
20. Eastman DE (1970) Phys Rev B2:1
21. Nygren MA, Siegbahn PEM, Wahlgren U, Åkeby H (1992) J Phys Chem 96:3633
22. Hermann K, Bagus PS, Nelin CJ (1987) Phys Rev B35:9467
23. Moruzzi VL, Janak JF, Williams AR (1978) Calculated electronic properties of metals. Pergamon Press, New York