Institute of Physics Publishing Journal of Physics: Conference Series **16** (2005) 257–268
doi:10.1088/1742-6596/16/1/035 SciDAC 2005

New insights into high temperature superconductivity from a computational solution of the two-dimensional Hubbard model

T.A. Maier[1], J.B. White[1], M. Jarrell[2], P. Kent[3], T.C. Schulthess[1]

[1]Oak Ridge National Laboratory, P.O. Box 2008, Mail Stop 6164, Oak Ridge, TN 37831
[2]Department of Physics, University of Cincinnati
[3]Joint Institute of Computational Sciences, University of Tennessee

E-mail: `schulthesstc@ornl.gov`

Abstract.
The Cray X1 in the Center for Computational Sciences at Oak Ridge National Laboratory as well as algorithmic improvements over the past decade enable significant new science in the simulation of high-temperature "cuprate" superconductors. We describe the method of dynamic cluster approximation with quantum Monte Carlo, along with its computational requirements. We then show the unique capabilities of the X1 for supporting this method and delivering near optimal performance. This allows us to study systematically the cluster size dependence of the superconductivity in the conventional two-dimensional Hubbard model, which is commonly believed to describe high-temperature superconductors. Due to the non-locality of the d-wave superconducting order parameter, the results on small clusters show large size and geometry effects. In large enough clusters, converged results are found that display a finite temperature instability to d-wave superconductivity. The results we report here demonstrate for the first time that superconductivity is possible in a system of strongly correlated electrons without the need of a phonon mediated attractive interaction.

1. Introduction

Superconductivity is, from an energy point of view, usually mentioned in the context of power transmission. The fact that metallic wires have non-zero electrical resistance currently imposes constrains on how electric power is used and distributed. For example, in order to minimize energy losses alternating currents (ac) at high voltage are used in transmission, while the majority power sources and consumption operate with direct currents (dc) and at low voltage. Availability of wires with zero resistance that can sustain large enough electrical currents, would allow to directly connect power sources to end-users and open new avenues to power generations. For example, photovoltaic energy sources could be concentrated in remote locations while making the generated power available in major urban agglomerations. These and many other very appealing applications in electronics and medicine seem to explain the continued interest in superconductivity, a state of mater with zero electrical resistance. But the implications on science and technology are potentially much broader. The superconducting state is one of few known macroscopic quantum state, where of the order of 10^{23} electrons are coherent (or entangled in the terminology of quantum computing). In high temperature superconductors

© 2005 IOP Publishing Ltd

(HTSC), this state persists up to 150 degrees Kelvin. Understanding the mechanism how this state forms could, for example, fundamentally impact the way quantum computing and quantum teleportation is approached.

In 1913, Kamerling Onnes managed to liquefy helium and thereby reached low enough temperatures to observe a sudden drop in the resistance of mercury. He was trying to confirm experimentally a transition to an insulating state at zero temperature that was predicted by Kelvin, but instead discovered a transition to a perfect conductor. The work[1], which was awarded the 1913 Nobel Prize, evolved into what is now known as conventional superconductivity. It took half a century for a microscopic theory of superconductivity to emerge, when in the 1950s, Bardeen, Cooper, and Schriffer (BCS)[2] showed that in metallic conductors: (1) electrons form pairs due to an attractive interaction mediated by lattice vibrations (phonons); and (2) these so-called Cooper pairs condense into a new thermodynamic state with zero resistance. BCS theory successfully explained all known phenomena of superconductivity and predicted the existence of a gap in the electronic excitation spectrum that was subsequently measured experimentally - BCS were awarded the 1972 Nobel Prize. All known superconductors until the early 1980s had transition temperature below 25 K, which appeared consistent with BCS theory. Stronger random lattice vibrations at higher temperatures would destroy the attractive interaction between electrons and hence suppress the formation of Cooper pairs. It was believed that 25 K represents a fundamental limit, above which superconductivity could not be observed. The discovery[3] in 1986 by Bednorz and Mller of superconductivity with transition temperature above 30 K in Ba doped LaCuO2, a poor conductor (LaCuO2 is an insulator), marked the beginning of a new era in condensed matter research. Within a few years new transition metal oxides were discovered with superconducting temperatures as high as 150 K - Bednorz and Mller were awarded the Nobel Prize in 1987)[4]. It was soon recognized that these new materials were fundamentally different from conventional metallic superconductors. While the superconducting state is still due to Cooper pairs and hence the second part of BCS theory holds, the prevalent view today is that the paring mechanism is not phonon mediate. However, despite two decades of intense research, the pairing mechanism in high temperature superconductors remains a mystery and represents one of the most important outstanding problems in condensed matter science today. This is largely because HTSC cuprates are strongly correlated electronic materials, a subject that in itself represents one of the most challenging problems in physics today.

In this paper, we report recent results which for the first time show that superconductivity is possible in a system of strongly correlated electrons without the need of a phonon mediated attractive potential. This work was enabled by algorithmic improvements over the past decade and resulted in the DCA/QMC code[5], as well at the computing capabilities available at the National Leadership Computing Facility (NLCF) at ORNL, which allows us to systematically solve the two-dimensional (2D) Hubbard model for the HTSC cuprates. The 2D Hubbard model is believed to contain all the necessary ingredients to describe HTSC in the cuprate[6], but remains unsolved until now. With the calculation we report here, we establish for the first time, that superconductivity is possible in a system of strongly correlated electrons without the need of a phonon mediated attractive potential. The paper is organized as follows. We introduce the model and the DCA algorithm in section 2 and 3, respectively. In section 4 we discuss results that were previously obtained on the IBM p690 and in section 5 we analyze the performance of the DCA/QMC code on the Cray X1 at ORNL. In section 6 we present the results obtained on the Cray X1 and discuss the new insight we obtained. Section 7 concludes the paper with a summary and a discussion of outstanding mathematical and physics issues.

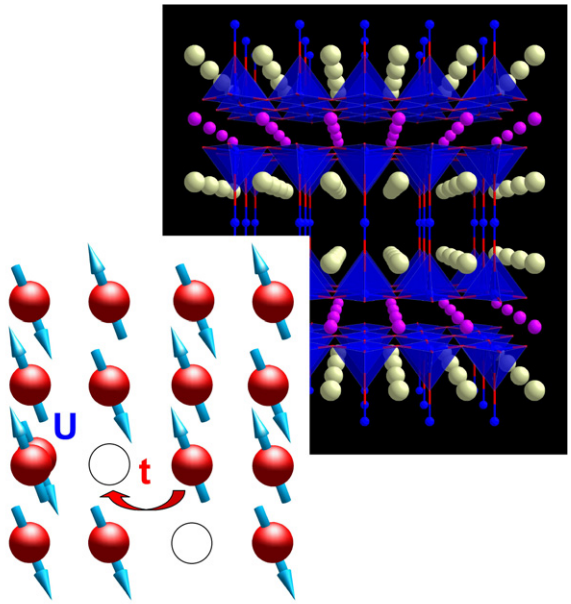

Figure 1. Crystal structure of $YBa_2Cu_3O_5$ and two-dimensional Hubbard model of the hole doped CuO_2 planes, with nearest neighbor hopping integral t and on-site Coulomb interaction U.

2. Hubbard model

The characteristic feature of all HTSC is a strongly anisotropic layered perovskite-like crystal structure with conducting CuO_2-planes separated by insulating layers of other elements (see right part of Fig. 1). Superconductivity takes place within the two-dimensional CuO_2 layers with the insulating barriers only providing charge carriers, usually holes to the layers and thus controlling the doping of CuO_2 planes.

First-principles calculations for HTSC compounds provide evidence that the band which crosses the Fermi surface has mainly CuO_2 character (see e.g. [7] and references therein). To reduce the complexity of the problem it thus seems reasonable to restrict calculations to a two-dimensional model with electrons moving in a single CuO_2 layer. Justified by the strong in-plane CuO bonds, the complexity may be further reduced by constructing a model that treats a whole CuO_2 plaquette as a single site. The resulting two-dimensional Hubbard model [8] is believed to capture the essential physics of HTSC [9, 10, 11]. A schematic of its Hamiltonian,

$$H = -t \sum_{\langle ij \rangle} c^{\dagger}_{i\sigma} c_{j\sigma} + U \sum_{i} n_{i\uparrow} n_{i\downarrow} \tag{1}$$

is illustrated in Fig. 1. The fermionic operator $c^{\dagger}_{i\sigma}$ ($c_{i\sigma}$) creates (destroys) an electron on site i with spin σ, and $n_{i\sigma} = c^{\dagger}_{i\sigma} c_{i\sigma}$ is the corresponding number operator. The first term describes the hybridization between sites with amplitude t, and the second term the Coulomb repulsion between two electrons residing on the same site. Because of screening, the magnitude of longer-ranged interactions is believed to be small compared to the on-site interaction.

Despite decades of intensive studies, this model remains unsolved except in one or infinite dimensions. Analytical methods based on a perturbative approaches suffer from the large magnitude of U, which renders these calculations at least questionable. Many theorists have turned to numerical approaches to close the gap between the model defined by its Hamiltonian and its properties. A large body of work has been devoted to a direct (numerically) exact solution

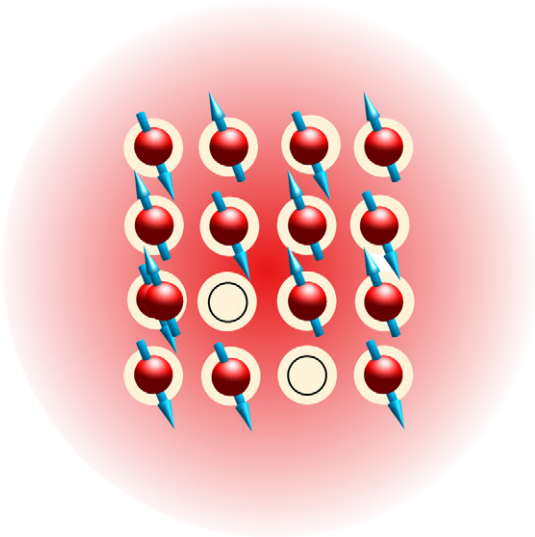

Figure 2. Schematic illustration of the DCA formalism. The model is mapped onto a finite-size cluster self-consistently coupled to a mean-field host. Correlations within the cluster are treated accurately while the physics on length scales beyond the cluster size is described on the mean-field level.

of finite-size systems using exact diagonalization or Quantum Monte Carlo (QMC) methods (for a review see [12]). Exact diagonalization, however, is severely limited by the exponential growth of computational effort with system size, while QMC methods suffer from what is known as "the sign problem" at low temperatures. Another difficulty of these methods arises from their strong finite-size effects, often ruling out the reliable extraction of low-energy scales, which are important to capture the competition between different ground states often present in correlated electron systems.

3. Dynamical Cluster Approximation

Mean-field theories are defined in the thermodynamic limit and therefore do not face the finite-size problems. Generally, mean-field theories divide the infinite number of degrees of freedom into two sets. A small set of degrees of freedom is treated explicitly, while the effects of the remaining degrees of freedom are summarized as a mean field acting on the first set. The Dynamical Mean-Field Theory (DMFT) [13, 14] (for a review see [15]) for itinerant correlated systems (such as the HTSC or systems described by the model Eq. (1)) is analogous to the coherent potential approximation for disordered systems [16, 17, 18]. It retains the dynamics of local degrees of freedom by mapping the lattice onto an impurity self-consistently embedded in a dynamical mean-field host.

Despite its success in the description of many correlated phenomena such as the Mott-Hubbard transition, the DMFT and CPA share the critical flaw of neglecting the effects of non-local fluctuations. Thus the DMFT is unable to capture the effects of *e.g.* spin-waves in magnetic systems, localization in disordered systems, or spin-liquid physics in correlated electron systems. Furthermore it cannot capture phase transitions to states with non-local order parameters, such as the d-wave superconducting phase in the HTSC. Non-local corrections are required to treat even the initial effects of these phenomena.

Here we use the Dynamical Cluster Approximation (DCA) [19, 20, 21, 5] (for a review see [22]) to study the properties of the Hubbard model, Eq. (1). The DCA extends the DMFT

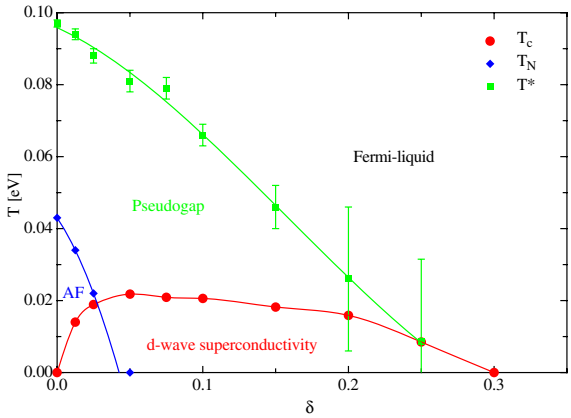

Figure 3. DCA/QMC Temperature-doping phase diagram of the two-dimensional Hubbard model when $t = 0.25$ eV, $U = 2$ eV for a 4-site cluster. Consistent with experiments on HTSC, regions of antiferromagnetism, d-wave superconductivity and pseudogap behavior are found.

by non-local correlations. Instead of mapping the lattice onto a single impurity, the system is mapped onto a periodic cluster of size N_c coupled to a mean-field host representing the remaining degrees of freedom (see Fig. 2). As a result, dynamical correlations up to a range limited by the cluster size are treated accurately, while the physics on longer length scales is described on the mean-field level. Translational invariance of the original system assures that the quantity describing the mean-field host can be self-consistently determined from the solution of the cluster problem. The complexity of the original problem with an infinite number of degrees of freedom is thus reduced to a self-consistent finite-size cluster problem with N_c degrees of freedom. The remaining cluster problem may then be solved numerically by a number of techniques including the QMC method [5] used here.

4. Small Clusters

Computations with a cluster of only four sites, the smallest cluster that can capture superconductivity with a d-wave order parameter, on the IBM p690 at the Center for Computational Sciences (CCS) show very good general agreement with HTSC. These results are summarized in the temperature-doping phase-diagram shown in Fig. 3 (see also [23, 24]). At low doping, δ, the system is an antiferromagnetic insulator below the Neél temperature $T_{\rm N}$. At finite doping, $\delta \leq 0.3$, an instability is found at the critical temperature T_c to a superconducting state described by a $\mathrm{d}_{x^2-y^2}$-wave order parameter. In the normal state, low-energy spin excitations become suppressed below the crossover temperature T^*. Simultaneously the electronic excitation spectrum displays a pseudogap, *i.e.* a partial suppression of low-energy spectral weight. Consistent with optical experiments, computations for a four-site cluster show that the superconducting transition is accompanied by a lowering of the electronic kinetic energy [25]. This result further shows the unconventional character of superconductivity in these systems. It is fundamentally different from the BCS theory for conventional superconductors [26], where pairing occurs through a reduction of the electronic potential energy accompanied by a slight increase in kinetic energy.

The apparent violation of the Mermin-Wagner theorem [27], according to which no phases with conventional long-range order can occur at finite temperatures in the two-dimensional Hubbard model, is a consequence of the small cluster size, and hence large mean-field character, in these simulations. In the case of antiferromagnetism, the Mermin-Wagner theorem thus necessarily translates to $T_{\rm N} = 0$ for the two-dimensional system. Superconductivity however can

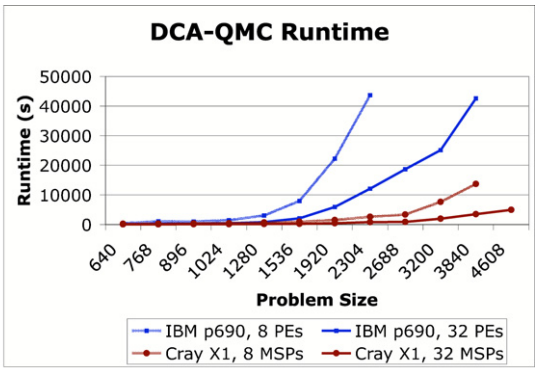

Figure 4. Runtimes for a series of DCA/QMC production runs. Each run is indicated by its value of $N = N_c N_l$. The lines connecting the data points are only guides to the eye.

exist even at finite temperatures as topological order below the Kosterlitz-Thouless transition temperature [28]. Therefore, larger-cluster-size studies are needed to see if the simulations recover the Mermin-Wagner theorem and if superconductivity survives as topological order in the infinite cluster size limit where the DCA becomes exact.

In the HTSC, on the other hand, a small but finite coupling between the two-dimensional CuO_2 layers induces long-range order at finite temperatures.

5. A performance boost with the Cray X1

The central quantity of the DCA code is the single-particle cluster Green function $\mathbf{G}_c$, which is a matrix of size $N \times N$ [5]. Here $N = N_c \times N_l$ where N_l is the number of "time-slices" in the time direction. The majority of the CPU time is spent in the inner loop of the QMC simulation, which updates the Green function matrix according to the vector outer product

$$\mathbf{G}' = \mathbf{G} + \mathbf{a} * \mathbf{b}^T \,, \tag{2}$$

where $\mathbf{a}$ and $\mathbf{b}$ are two vectors of dimension N. This computation is handled by the BLAS [29] call DGER, which performs a double-precision rank-one matrix update representing $O(N^2)$ operations. Each iteration requires N such calls, however, resulting in $O(N^3)$ operations.

Another CPU-intensive task is the evaluation of two-particle correlation functions. In the QMC technique this reduces to evaluating products of Green functions and thus to computing matrix products. This is done by using the BLAS call CGEMM, which performs single-precision complex matrix-matrix multiplication, and one call again is $O(N^3)$.

Porting and tuning the DCA/QMC implementation on the Cray X1 was straightforward. The port required no changes beyond the "`Makefile`", and tuning involved performance profiling and adding "`concurrent`" directives to one file. This file contains a number of nested loops using indirect addressing, or index arrays. The bulk of the tuning effort was in determining which loops did and did not iterate over repeated indices.

The Cray X1 has a number of advantages over general-purpose systems in performing DCA/QMC computations, particularly with increasing cluster size. This advantage is demonstrated in Fig. 4, which compares runtimes of some early DCA runs on the X1 and the IBM p690 in the CCS, using 8 and 32 processors (MSPs) on each.The figure shows runtimes for production runs with a fixed value of $N_c = 64$ and increasing values of N_l and thus N, where the value of N is shown. Eight X1 MSPs easily outperform thirty-two 1.3-GHz Power4 processors for the larger problem sizes.

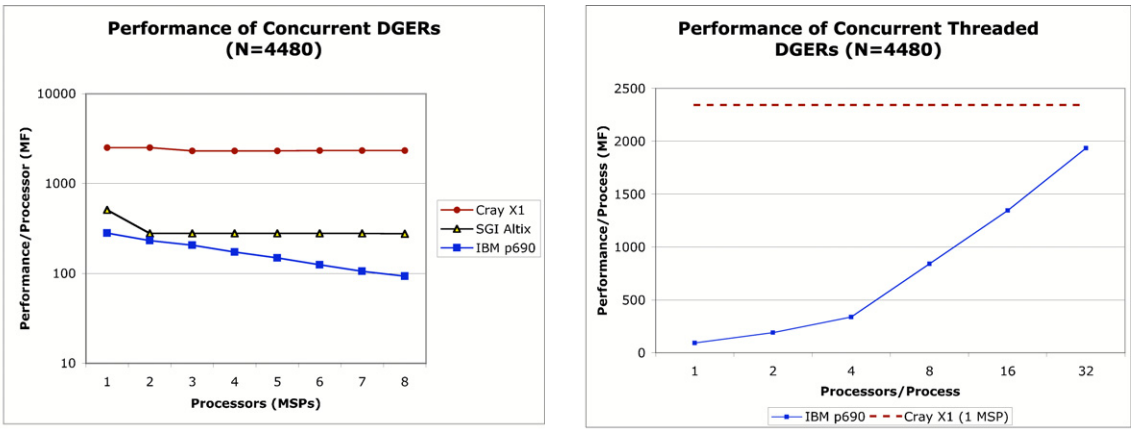

Figure 5. Left panel: per-processor performance of concurrent DGER calls using $N = 4480$ matrices. Right panel: DGER performance of a fully loaded IBM p690 using different numbers of processors per threaded process. The dotted line is the per-MSP performance of loaded Cray X1 nodes.

As discussed above, the DCA implementation includes two $O(N^3)$ computations built on the BLAS calls CGEMM and DGER. CGEMM is a BLAS3 call, which implies that it can be blocked effectively for cache memory, and many modern general-purpose processors can perform the operations near their peak. The X1 processors can also, but they have the added benefit of a very high peak rate augmented by the ability to perform single-precision operations at twice the rate of double-precision.

The Cray X1 has a more significant advantage over the prevailing cache-dependent architectures in the DGER operations. Each call depends on the results of the previous call, so the operations cannot be interleaved. DGER is a BLAS2 call, which implies that it does much fewer computations per memory access than CGEMM, and thus is typically limited by memory bandwidth.

We conducted separate DGER benchmarks to measure the advantage of the Cray X1 in this operation, and results for the CCS Cray X1, SGI Altix (1.5 GHz Itanium2), and IBM p690 (1.3 GHz Power4) are in Fig. 5 (left panel). The vendor-optimized DGER was used for each system. The figure shows the performance of DGER for a matrix of size $N = 64 \times 70 = 4480$, which is representative of large DCA runs. Separate DGER instances were run concurrently across increasing numbers of processors (MSPs), mimicking the processes of a Monte-Carlo simulation. The X1 memory system is able to maintain performance and efficiency with added processors, while the p690 steadily degrades. The Altix degrades going from one to two processors because memory bandwidth is shared between processor pairs. The X1 maintains 8–25 times the performance and 4–10 times the efficiency of the other systems.

Despite the Monte-Carlo nature of the DCA/QMC algorithm, the Cray X1 also has an important scalability advantage over systems with weaker processors. Each DCA/QMC process has a significant fixed start-up cost, which favors splitting the Monte-Carlo iterations across fewer, faster processors.

Another option would be to multithread each Monte-Carlo process, effectively using an SMP as a large single "processor". We explore this possibility in Fig. 5 (right panel), which shows the performance of IBM's multithreaded DGER on a p690, again using a matrix size of $N = 4480$. The dashed line shows the per-MSP performance of an X1 performing concurrent DGER operations on 32 MSPs, thus simulating a loaded system. The solid line shows the

performance of a 32-processor p690 loaded with concurrent DGER computations, but using different numbers of processors per DGER process.

The left-most point thus shows the performance of a single processor when all 32 processors of the p690 are performing independent DGER operations, while the right-most point shows the aggregate performance of dedicating all 32 processors to a single DGER. The figure indicates that dedicating a full IBM p690 to each DGER does not match the performance of a single X1 MSP. No threaded version of vendor-optimized DGER was available for the Altix or the Cray X1 at the time of this test. Tests of untuned DGER implemented with Fortran loops and OpenMP showed little improvement on the X1 for matrices of size 4480, and the Fortran/OpenMP implementation on the Altix was not competitive with the single-threaded vendor-optimized DGER.

The significant performance advantage of the Cray X1 for DCA/QMC computations, as illustrated by its dominance in DGER performance, has allowed us to perform simulations that are out of the reach of other systems, all without having to resort to hybrid parallelization. In particular, the X1 has provided the capability needed to perform DCA/QMC computation with much-larger cluster sizes.

6. New insights with large cluster results

The numerical complexity of larger cluster simulations restricts the ability to explore both a large parameter space and different cluster sizes. The parameters are therefore chosen to favor superconducting and antiferromagnetic order and cluster geometries are carefully selected. Much can be learned from simulations of finite size systems, where periodic boundary conditions are typically used. Betts and Flynn [30] systematically studied the 2D and 3D Heisenberg models on finite size clusters and developed a grading scheme to determine which clusters should be used in finite size simulations. The main qualification is the perfection of the near-neighbor shells: a measure of the completeness of each neighbor shell compared to the infinite lattice. In finite size calculations of the order parameter and the ground-state energy, they found that the results for the most perfect clusters fall on a scaling curve, while the imperfect clusters generally produce results off the curve. We employed some of the cluster geometries proposed by Betts (see Fig. 6) to study the antiferromagnetic transition at half filling and generalized Betts' arguments to generate a set of clusters appropriate to study d-wave superconductivity.

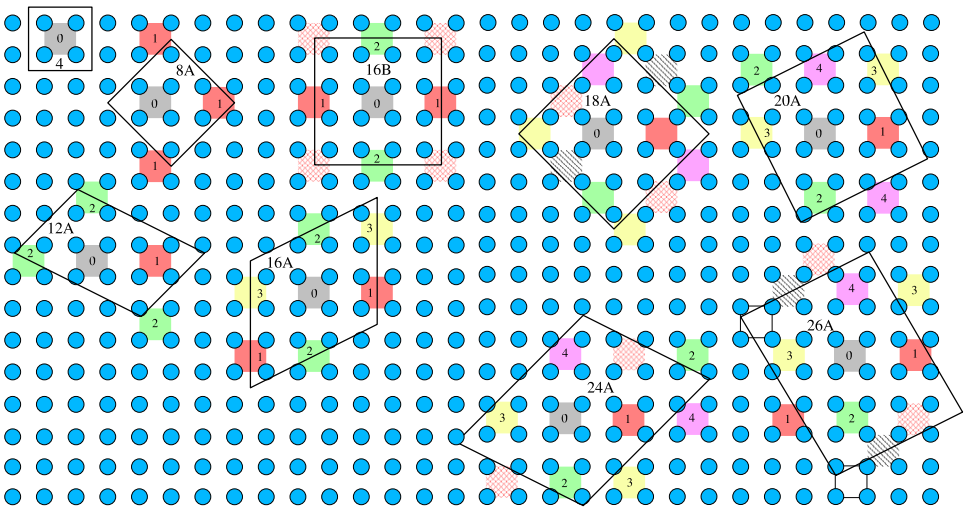

Figure 6. Cluster sizes and geometries used in our study. The shaded squares represent independent d-wave plaquettes within the clusters. In small clusters, the number of neighboring d-wave plaquettes z_d listed in table 1 is smaller than four, i.e. than that of the infinite lattice.

To illustrate that the DCA recovers the correct result as the cluster size increases, we plot in Fig. 7 the DCA results for the Neél temperature $T_{\rm N}$ at half-filling as a function of the cluster size N_c. $T_{\rm N}$ decreases slowly with increasing cluster size N_c. As spin-correlations develop exponentially with decreasing temperature in 2D, the $N_c > 4$ data falls logarithmically with N_c, consistent with $T_{\rm N} = 0$ in the infinite size cluster limit. Thus, the Mermin-Wagner theorem is recovered for $N_c \to \infty$. The clusters with $N_c = 2$ and $N_c = 4$ are special because their coordination number is reduced from four. For $N_c = 2$ the coordination number is one and for $N_c = 4$ the coordination is two. Hence, fluctuations of the order parameter are overestimated, antiferromagnetism is suppressed and their corresponding $T_{\rm N}$ does not fall on the curve.

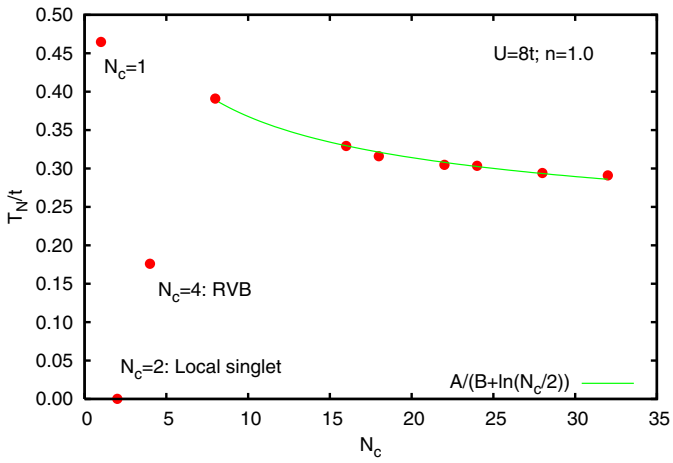

Figure 7. Neél temperature at half-filling when U=8t versus the cluster size. $T_{\rm N}$ scales to zero in the infinite cluster size limit. The solid line represents a fit to the function $A/(B + ln(N_c/2))$ obtained from the scaling ansatz $\xi(T_N) = L_c$. For $N_c = 2$ a local singlet and for $N_c = 4$ the RVB state suppress antiferromagnetism.

The transition to a superconducting state with d-wave symmetry is indicated by the divergence of the pair-field susceptibility P_d, or equivalently by the node of P_d^{-1}. In contrast to the order parameter indicating magnetic order, the d-wave order parameter is non-local and involves four bonds or sites. Thus, large size and geometry effects have to be expected in small clusters. Similar to the cluster grading scheme Betts developed for magnetic order, we can classify the different clusters according to their quality for d-wave order. At low temperatures, local d-wave pairs will form, but phase fluctuations of the pair wave-function prevent the system from becoming superconducting. Since the DCA cluster has periodic boundary conditions, each four-site d-wave plaquette has four neighboring d-wave plaquettes. However, as illustrated in Fig. 6, in small clusters, these are not necessarily independent and hence the effective dimensionality is reduced. The number z_d of neighboring independent d-wave plaquettes in a given cluster is then a measure for the strength of phase fluctuations which act to suppress the pairing correlations and hence T_c.

Table 1. Number of independent neighboring d-wave plaquettes z_d.

Cluster	4	8A	12A	16A	16B	18A	20A	24A	26A
z_d	0 (MF)	1	2	3	2	1	4	4	4

Fig. 6 shows the arrangement of independent d-wave plaquettes in the clusters used in this study and the corresponding values of z_d are listed in table 1. The $N_c = 4$ cluster encloses exactly one d-wave plaquette ($z_d = 0$). When a local d-wave pair forms on the cluster, the system becomes superconducting, since no superconducting phase fluctuations are included. Thus, the $N_c = 4$ result corresponds to the mean-field solution. In the 8A cluster, there is room for one more d-wave pair, thus the number of independent neighboring d-wave plaquettes $z_d = 1$. Since this same neighboring plaquette is adjacent to its partner on four sides, phase fluctuations will be overestimated as compared to the infinite system. The situation is similar in the 16B cluster, where only two independent (and one next-nearest neighbor) d-wave plaquettes are found ($z_d = 2$). In contrast, $z_d = 3$ in the oblique 16A cluster. We thus expect d-wave pairing correlations to be suppressed in the 16B cluster as compared to those in the 16A cluster. With the exception of the 18A cluster, where neighboring d-wave plaquettes share one site and thus are not independent, the larger clusters 20A, 24A, and 26A all have $z_d = 4$ and are thus expected to show the most accurate results.

Fig. 8 shows the temperature dependence of the inverse d-wave pair-field susceptibility, $1/P_d$, in the 10% doped system. Since a proper error propagation is severely hampered by storage requirements, we obtain the error-bars shown on the 16A results from a number of independent runs initialized with different random number seeds. Error-bars on larger cluster results are expected to be of the same order or larger. As noted before, the $N_c = 4$ result is the mean-field result for d-wave order and hence yields the largest pairing correlations and the highest T_c. As expected, we find large finite size and geometry effects in small clusters. When $z_d < 4$, fluctuations are overestimated and the d-wave pairing correlations are suppressed. In the 8A cluster where $z_d = 1$ we do not find a phase transition at finite temperatures. Both the 12A and 16B cluster, for which $z_d = 2$, yield almost identical results. Pairing correlations are enhanced compared to the 8A cluster and the pair-field susceptibility P_d diverges at a finite temperature. As the cluster size is increased, z_d increases from 3 in the 16A cluster to 4 in the larger clusters, the phase fluctuations become two-dimensional and as a result, the pairing correlations increase further (with exception of the 18A cluster). Within the error-bars (shown for 16A only), the results of these clusters fall on the same curve, a clear indication that the correlations which mediate pairing are short-ranged and do not extend beyond the cluster size. The low-temperature region can be fitted by the KT form $P_d = A\exp(2B/(T - T_c)^{0.5})$. For all clusters with $z_d \geq 3$ we find a transition temperature $T_c \approx 0.025t$. We cannot preclude, however, the possibility of a very slow, logarithmic cluster size dependence of the form $T_c(N_c) = T_c(\infty) + B^2/(C + \ln(N_c)/2)^2$ where $T_c(\infty)$ is the exact transition temperature.

7. Summary and Conclusions

The Cray X1 in the Center for Computational Sciences at Oak Ridge National Laboratory has enabled significant new progress in the understanding of HTSC within a minimal microscopic model, the two-dimensional Hubbard model. DCA/QMC simulations at small cluster size $N_c = 4$ show very good general agreement with HTSC, including superconductivity at high temperatures. Due to the small cluster size however, the results violate the Mermin-Wagner theorem, according to which no long-range order is allowed at finite temperatures in the two-dimensional model. The significant performance advantage of the X1 for the DCA/QMC computations has provided the capability to study much larger cluster sizes. Consistent with the Mermin-Wagner theorem, the finite temperature antiferromagnetic transition found in the $N_c = 4$ simulation is systematically suppressed with increasing cluster size. In small clusters, the results for the d-wave pairing correlations show a large dependence on the size and geometry of the clusters. For large enough clusters however, we find converged results that display a finite temperature instability to a d-wave superconducting phase at $T_c \approx 0.025t$ at 10% doping when $U = 4t$. This established that superconductivity is possible in a system of strongly correlated

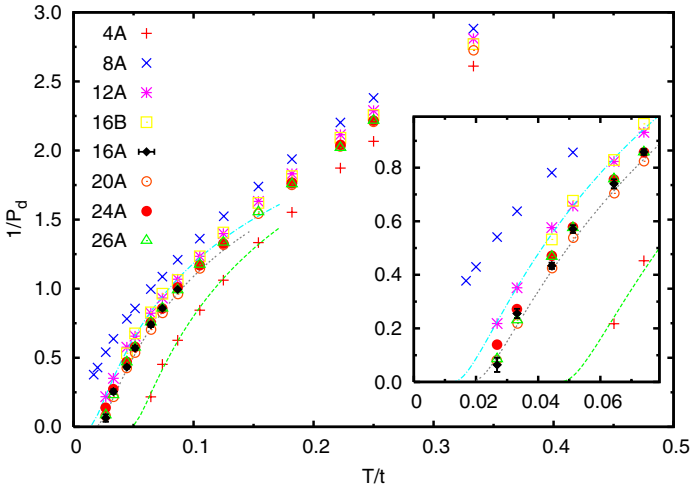

Figure 8. Inverse d-wave pair-field susceptibility as a function of temperature for different cluster sizes at 10% doping. The continuous lines represents fits to the function $P_d = A\exp(2B/(T-T_c)^{0.5})$ for data with different values of z_d. Inset: Magnified view of the low-temperature region.

electrons without the need of a phonon mediated attractive interaction.

8. Acknowledgements

This research was enabled by computational resources of the Center for Computational Sciences and was sponsored by the offices of Basic Energy Sciences and Advanced Scientific Computing Research, U.S. Department of Energy as well as the laboratory research and development fund. Oak Ridge National Laboratory is managed by UT-Battelle, LLC under Contract No. DE-AC0500OR22725. The development of the DCA formalism and algorithm was supported by the NSF under Grant No. DMR-0312680 as well as through resources provided by San Diego Supercomputer Center under NSF cooperative agreement SCI-9619020.

9. References

[1] P.H.E. Meije. Kamerlingh onnes and the discovery of superconductivity. *Am. J. Phys.*, 62:1105 – 1108, 1994.

[2] J. Bardeen, L.N. Cooper, and J.R. Schrieffer. Theory of superconductivity. *Phys. Rev. B*, 108:1175–1204, 1957.

[3] J.G. Bednorz and K.A. Müller. Possible high t_c superconductivity in the ba-la-cu-o system. *Z. Phys. B*, 64:189–193, 1986.

[4] J.G. Bednorz and K.A. Müller. Perovskite-type oxides = the new approach to high-t_c superconductivity. *Rev. Mod. Phys.*, 60:585–600, 1988.

[5] M. Jarrell, Th. Maier, C. Huscroft, and S. Moukouri. A quantum monte carlo algorithm for non-local corrections to the dynamical mean-field approximation. *Phys. Rev. B*, 64:195130, 2001.

[6] F.C. Zhang and T.M. Rice. Effective hamiltonian for the superconducting cu oxides. *Phys. Rev. B*, 37:3759–3761, 1988.

[7] P. Fulde. *Electron Correlations in Molecules and Solids.* Springer Series in Solid-State Sciences. Springer Verlag, Berlin/Heidelberg/New York, 1991.

[8] J. Hubbard. *Proc. Royal. Soc. London*, 276:238, 1963.

[9] F.C. Zhang and T.M. Rice. Effective hamiltonian for the superconducting cu oxides. *Phys. Rev. B*, 37:R3759, 1988.

[10] F.C. Zhang and T.M. Rice. Validity of the t-j model. *Phys. Rev. B*, 41:7243, 1990.

[11] P. W. Anderson. *The Theory of Superconductivity in the High-T_c Cuprates.* Princeton University Press, Princeton, NJ, 1997.

[12] E. Dagotto. Correlated electrons in high-temperature superconductors. *Rev. Mod. Phys*, 66:763, 1994.

[13] W. Metzner and D. Vollhardt. Correlated lattice fermions in $d = \infty$ dimensions. *Phys. Rev. Lett.*, 62:324, 1989.

[14] E. Müller-Hartmann. Correlated fermions on a lattice in high dimensions. *Z. Phys.* **B**, 74:507, 1989.

[15] A. Georges, G. Kotliar, W. Krauth, and M.J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.*, 68:13, 1996.

[16] P. Soven. Coherent-potential model of substitutional disordered alloys. *Phys. Rev.*, 156:809–813, 1967.

[17] D.W. Taylor. Vibrational properties of imperfect crystals with large defect concentrations. *Phys. Rev.*, 156:1017, 1967.

[18] H. Shiba. A reformulation of the coherent potential approximation and its applications. *Prog. Theo. Phys.*, 46:77–94, 1971.

[19] M. H. Hettler, A. N. Tahvildar-Zadeh, M. Jarrell, T. Pruschke, and H. R. Krishnamurthy. Nonlocal dynamical correlations of strongly interacting electron systems. *Phys. Rev. B*, 58:R7475–R7479, 1998.

[20] M. H. Hettler, M. Mukherjee, M. Jarrell, and H. R. Krishnamurthy. Dynamical cluster approximation: Nonlocal dynamics of correlated electron systems. *Phys. Rev. B*, 61:12739–12756, 2000.

[21] Th. Maier, M. Jarrell, Th. Pruschke, and J. Keller. A non-crossing approximation for the study of intersite correlations. *Eur. Phys. J B*, 13:613, 2000.

[22] Th. Maier, M. Jarrell, Th. Pruschke, and M.H. Hettler. Quantum cluster theories. *preprint cond-mat/0404055*, 2004.

[23] Th. Maier, M. Jarrell, Th. Pruschke, and J. Keller. d-wave superconductivity in the hubbard model. *Phys. Rev. Lett.*, 85:1524, 2000.

[24] M. Jarrell, Th. Maier, M. H. Hettler, and A. N. Tahvildarzadeh. Phase diagram of the hubbard model: Beyond the dynamical mean field. *Europhys. Lett.*, 56:563, 2001.

[25] Th.A. Maier, M. Jarrell, A. Macridin, and C. Slezak. Kinetic energy driven pairing in the cuprates. *Phys. Rev. Lett.*, 92:027005, 2004.

[26] J.R. Schrieffer. *Theory of Superconductivity.* Addison Wesley, Reading, MA, 1993.

[27] N.D. Mermin and H. Wagner. *Phys. Rev. Lett.*, 22:1133, 1966.

[28] J.M. Kosterlitz and D.J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys. C*, 6:1181, 1973.

[29] J. J. Dongarra, J. Du Croz, I. S. Duff, and S. Hammarling. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Soft.*, 16:1–17, 1990.

[30] D.D. Betts, H.Q. Lin, and J.S. Flynn. *Can. J. Phys.*, 77:353–369, 1999.