Journal of Statistical Physics, Vol. 43, Nos. 5/6, 1986

Monte Carlo Studies with the ST-100 Array Processor

R. L. Sugar[1]

Monte Carlo simulations with the 100 Mflop ST-100 array processor are described. The architecture of the array processor and its applicability to large-scale numerical simulations is discussed. Results are presented for the Abelian Higgs model, a charge density wave transition in a quasi-one-dimensional system and the finite temperature phase transition in $SU(3)$ lattice gauge theory.

KEY WORDS: Quantum Monte Carlo; $SU(3)$ lattice gauge theory; Abelian Higgs model; charge density wave transition.

1. INTRODUCTION

As we heard at the conference, Monte Carlo simulations are presently being used to study a wide variety of physical phenomena. This field is clearly going through a very productive period. One of the major reasons for recent progress has been the marked increase in computing power available to its practitioners. The increased access to supercomputers has been most important, and the development of special purpose computers explicitly designed for numerical simulations holds great promise. Another development, which has perhaps been less widely noted, is the advent of a new generation of computers and array processors that have a significant fraction of the speed of a supercomputer but are economical enough so that they can be owned and operated by a small group of scientists. I believe that these machines can be very useful for large-scale simulations, and I would like to describe the experience our group has had with one of them, the Star Technologies, Inc. ST-100 array processor.

The ST-100 that I will describe is operated by a group of University of California faculty members: Jorge Hirsch, Julius Kuti, and Doug Toussaint of U.C. San Diego, and Doug Scalapino and myself of U.C. Santa Barbara.

[1] Department of Physics, University of California, Santa Barbara California 93106.

0022-4715/86/0600-1147$05.00/0

Our interests are primarily focused on the numerical simulation of field-theory problems in condensed matter and high-energy physics. I will discuss three of the problems we have studied with the ST-100 during the last year: the Abelian Higgs model, the development of correlations in a quasi-one-dimensional electron system, and the finite temperature phase transition in $SU(3)$ lattice gauge theory. First, however, I will briefly describe the ST-100. Numerical simulation is in a real sense an experimental field, and it is traditional for an experimentalist to describe new equipment the first time he or she uses it.

2. THE ST-100 ARRAY PROCESSOR

The ST-100 is an attached processor. We use a VAX11/750 as the host. All program development including editing, compiling, and linking the array processor programs is done on the VAX. In order to run a job on the ST-100 one runs a Fortran program on the VAX which prepares data for the array processor, sends the data and array processor program to the ST-100, and instructs the ST-100 to begin operation. The VAX then continues with other tasks until the array processor signals that it is finished and ready to return its results. Typically we use three to four minutes of VAX CPU time for each 24 hours of ST-100 operation. The ST-100 and its host VAX are located in Santa Barbara. The VAX is connected by a dedicated 9600 baud line to a similar computer in San Diego which also runs the ST-100 program development software. Thus the San Diego members of the group can do much of their software development on their own VAX and send programs to Santa Barbara for final testing and production runs on the array processor.

A schematic diagram of the ST-100 is shown in Fig. 1. It has a large main memory, which on our machine consists of 4 million, 32-bit words, but which could be expanded to 8 million words. Data can be stored in main memory in full words, half words, or bytes. There are four primary arithmetic units: two adders and two multipliers. The adders can also perform subtraction, logical operations, and conversion between fixed and floating point formats. The adders and multipliers have three-stage pipelines, and are each capable of producing one result per machine cycle. Since the cycle time is 40 ns, this comes to 25 million floating-point operations per second (25 Mflops) for each unit or a total 100 Mflops if the programmer is able to keep all four units in operation. There is also a divide/square root unit which is not pipelined and which requires 12 machine cycles per result. It is, of course, to be avoided whenever possible.

In order for the arithmetic units to operate on data in main memory, the data must first be moved into the cache memory. Cache consists of

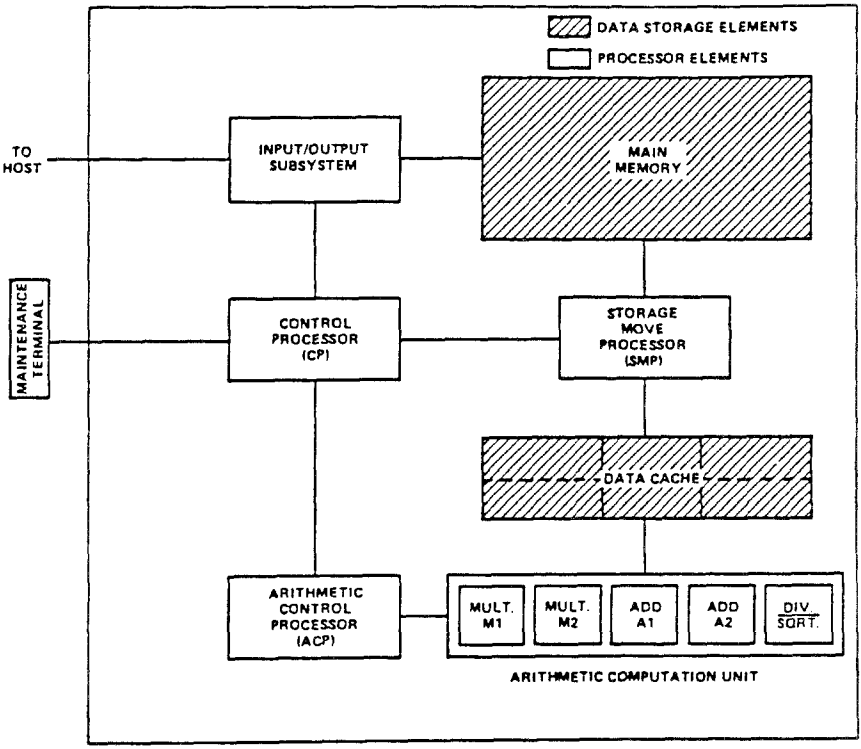

Fig. 1. A schematic diagram of the ST-100 array processor.

48,000 32-two bit words. It is divided into three sections of equal size, and each of these sections can in turn be divided into two subsections. Data movement between main memory and cache is controlled by the Storage Move Processor (SMP). It is possible to move one word per machine cycle to or from main memory.

The Arithmetic Control Processor (ACP) controls data flow between cache and the arithmetic units as well as the operation of these units. It is possible to move three words per cycle between cache and the arithmetic units, but each word must go to or from a different cache section. It is also possible to move output from the arithmetic units directly into their input registers without returning it to cache, and to use input data in more than one arithmetic operation without rereading it from cache. Given the limitation on data flow from cache, this is essential if one is to keep both adders and both multipliers in simultaneous operation.

The overall operation of the ST-100 is managed by the Control Processor. It is programmed in Array Processor Control Language

(APCL), which is at once a subset and an extension of Fortran. Among the extensions is the ability to call subroutines or macros that run on the SMP and ACP. The manufacturer supplies a large library of such macros for vector data movement and vector arithmetic operations. It is also possible to write one's own macros for the SMP and ACP, but this must be done in assembly language. For problems that are highly vectorizable it is possible to obtain excellent performance using APCL and the manufacturer provided macros, but it is often necessary to do some assembly language programming in order to obtain peak performance.

The SMP and the ACP operate independently so it is possible to overlap the movement of data between main memory and some subsections of cache with arithmetic operations on data in other subsections. APCL has the tools for doing so, and for synchronizing the two processors when that is necessary.

It should be clear from this brief description that developing software for the ST-100 provides a somewhat greater challange than programming a conventional computer or even a supercomputer. However, the reward for this extra effort is a 100 Mflop machine which is economical enough so that it can dedicated to a small number of computationaly intensive problems.

3. THE ABELIAN HIGGS MODEL

The first large-scale numerical simulation carried out with our ST-100 was a study of the Abelian Higgs Model by Doug Toussaint and myself.[1] This model describes the interaction of charged scalar particles with the quantized electromagnetic field.

In a well-known paper[2] Coleman and Weinberg showed that in the weak coupling limit, that is the one-loop approximation, this model has a first-order phase transition from a Coulomb phase in which the charged particles interact via the standard long-range electromagnetic interaction. and a Higgs phase in which the photon acquires a mass and the interactions among the charged particles are short-range. Such a first-order transition plays a fundamental role in the inflationary scenario of the early universe. Of course in the early universe one is interested in the non-Abelian version of the model, but it seems reasonable that phenomena found in the Abelian model will be among those found in the more complex non-Abelian one. We, therefore, focused our attention on the Coulomb-Higgs phase transition in our study.

The continuum Lagrangian for the Abelian Higgs model is

$$L = -\frac{1}{4g^2} F^{\mu\nu} F_{\mu\nu} + (D^{\mu}\varphi)^* (D_{\mu}\varphi) - m^2 |\varphi|^2 - \lambda |\varphi|^4 \qquad (1)$$

where

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$$

and

$$D_\mu \varphi = (\partial_\mu - iA_\mu)\varphi$$

In order to carry out a numerical simulation it is necessary to put the model on a space-time lattice. There are, of course, many ways of doing so. Our choice for the lattice action is

$$S = \frac{\beta}{2} \sum_{plaquettes} \left[\sum_{\substack{links\ in \\ plaquette}} \pm \vartheta_{link} \right]^2 - \sum_{\substack{links \\ i\ and\ j}} \{ |\varphi_i|\,|\varphi_j| \cos[arg(\varphi_i) - arg(\varphi_j) + \vartheta_{ij}] \} + \sum_{sites} (4 + m^2)|\varphi_i|^2 + \sum_{sites} \lambda |\varphi_i|^4 \tag{2}$$

where $\beta = 1/g^2$. The variables are a real number ϑ, the gauge field associated with each link on the lattice, and a complex number φ, the Higgs field associated with each site.

Note that we have used a noncompact formulation of the gauge field. The compact version has a strong coupling phase in which electric charge is confined, and we did not want this phase to interfere with our study of the Coulomb–Higgs phase transition in the strong coupling regime. In some of the earlier studies of this model the magnitude of the Higgs field was kept fixed.[3] It is believed that this does not change the universality class of the model, but we expect the theory to approach the continuum limit more rapidly when the magnitude of the Higgs field is retained as a dynamical variable as we have done. Simulations in which the magnitude of the Higgs field was allowed to fluctuate have been carried out by Munehisa[4] and by Koutsoumbas.[5] Although considerable light was shed on the model in the earlier simulations, the order of the phase transition remained unclear.

We used the Metropolis algorithm to generate field configurations for the simulation. Separate algorithms were needed to update the gauge variables and the magnitudes and phases of the Higgs field. By making use of skew periodic boundary conditions it was possible to write highly vectorized code in which many sites or links of the lattice were updated simultaneously. The code was written in the array processor's high-level language, APCL, making extensive use of the vector subroutines for the

SMP and ACP provided by the manufacturer. The only assembly language routines which we wrote were a vector random number generator and routines for moving vectors with periodic boundary conditions between main memory and cache. The resulting code ran 150 to 200 times faster than a Fortran program performing the identical calculation on a VAX 11/750. The larger ratio was obtained for large lattices where the vector length reached its maximum value of 4000.

We studied the Coulomb–Higgs phase transition for a range of values of the gauge coupling β and the Higgs coupling λ. For all values for which the order is clear, this transition is first-order in agreement with expec-

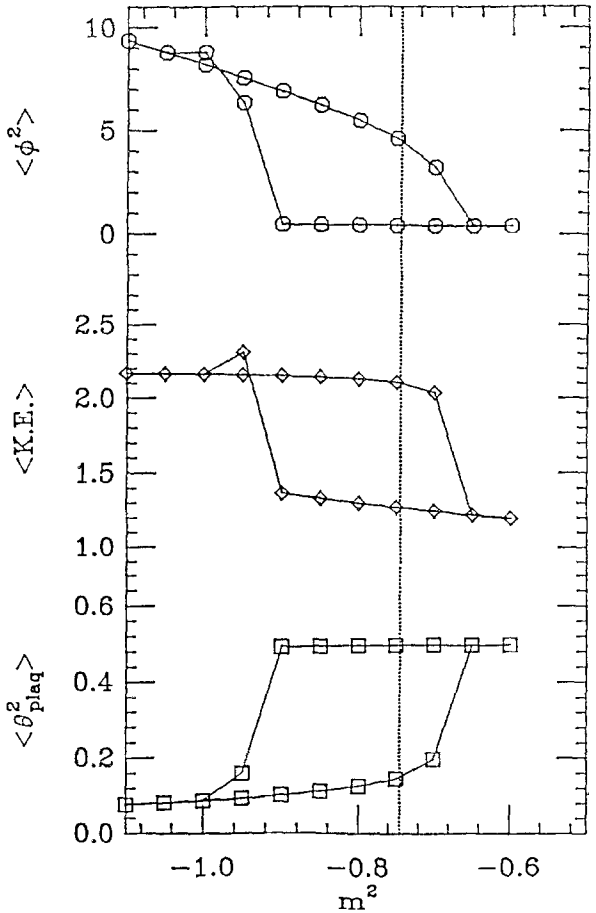

Fig. 2. Hysteresis loops for the plaquette energy (per plaquette), the kinetic energy (per site), and φ^2. Here $\beta = 0.2$ and $\lambda = 0.05$. The dotted line is the critical value of m^2.

tations from the work of Coleman and Weinberg.[2] Our results are most striking for small values of β and λ. In Fig. 2, I show the hysteresis loops for the plaquette and link energies and for $\langle\varphi^2\rangle$ obtained by varying m^2 for $\beta=.2$ and $\lambda=.05$ on a $7^3\times 8$ lattice. The existence of a large hysteresis loop is not, by itself, a proof of a first-order phase transition. It is necessary to show that the two branches of the hysteresis loop are well-defined metastable phases, in particular that the quantities measured on the two branches are independent of lattice size. This is demonstrated in Fig. 3 were I show the hysteresis loop for $\langle\varphi^2\rangle$ at $\beta=.2$ and $\lambda=.05$ on lattices of dimension $5^3\times 6$, $7^3\times 8$, and $9^3\times 10$. It is clear that our results are independent of lattice size. The points with large error bars are values of m^2 for which the lattice tunneled from a metastable state to a stable state. The value shown is simply the average values in the two phases weighted by the time spent in each phase.

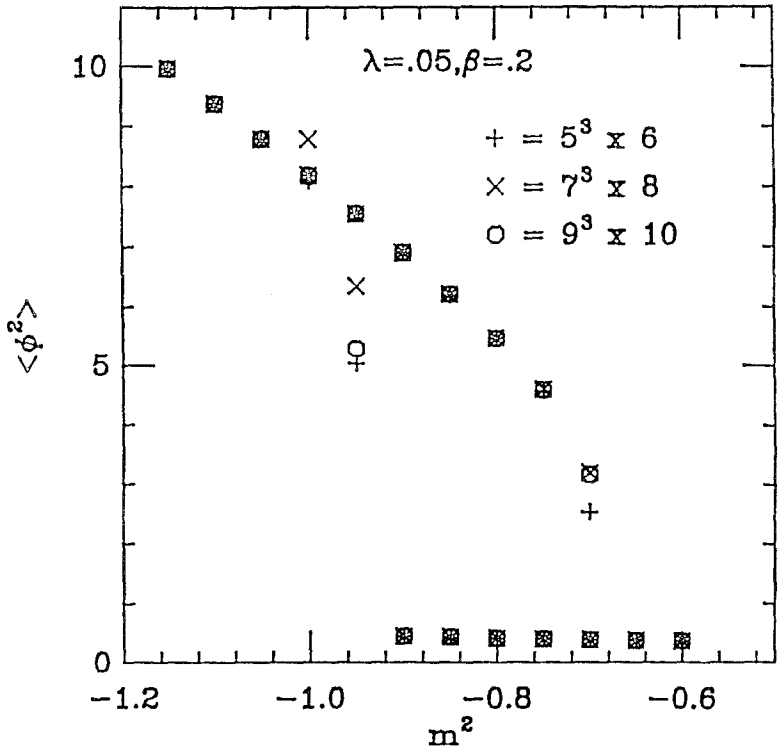

Fig. 3. Effect of lattice size on the value of $\langle\varphi^2\rangle$. The values of λ and β are the same as in Fig. 2.

In order to make an accurate measurement of the critical mass at which the phase transition takes place, we initialized the lattice with half of it in the Coulomb phase and half in the Higgs phase. The system was then allowed to evolve until it was entirely in one phase. Clearly the critical value of m^2 lies between the smallest value for which the system becomes disordered and the largest value for which it becomes ordered. In making these measurements we used elongated lattices of dimension $7^3 \times 16$ or $7^3 \times 32$, dividing the lattice along the longest dimension. In Fig. 4 I show the time history of $\langle \varphi^2 \rangle$ for a variety at values of m^2 at $\beta = .2$ and $\lambda = .05$. Using this approach we were able to make accurate measurements of the critical mass with a rather modest expenditure of computer time. A plot of the critical mass versus the gauge coupling constant for $\lambda = .05$ is given in Fig. 5.

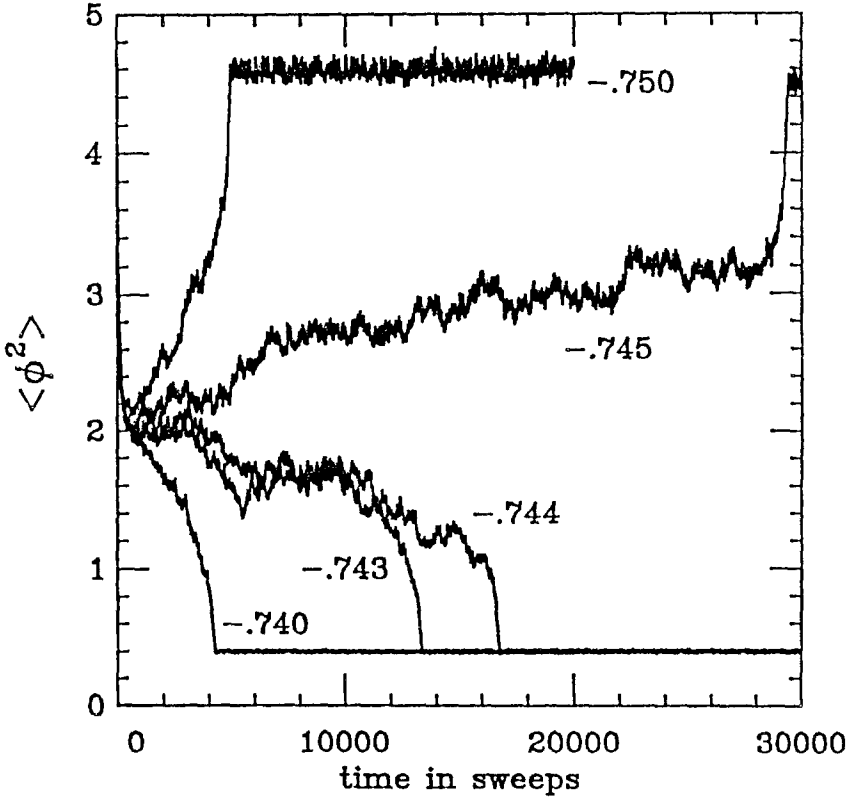

Fig. 4. Evolution of the average of φ^2 over the lattice from a mixed start. Results are shown for several values of m^2 bracketing the critical point. Here the lattice size was $7^3 \times 32$, $\beta = 0.2$, and $\lambda = 0.05$.

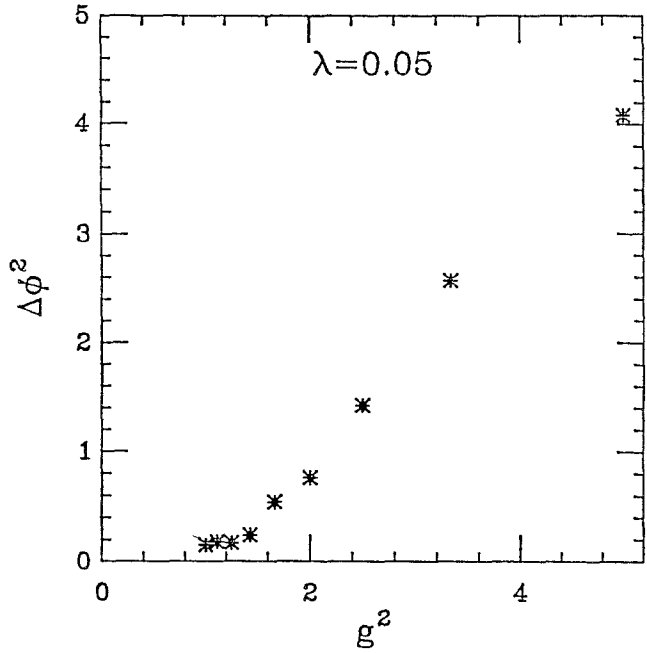

Fig. 5. The critical *mass*^{2} as a function of β for $\lambda = 0.05$.

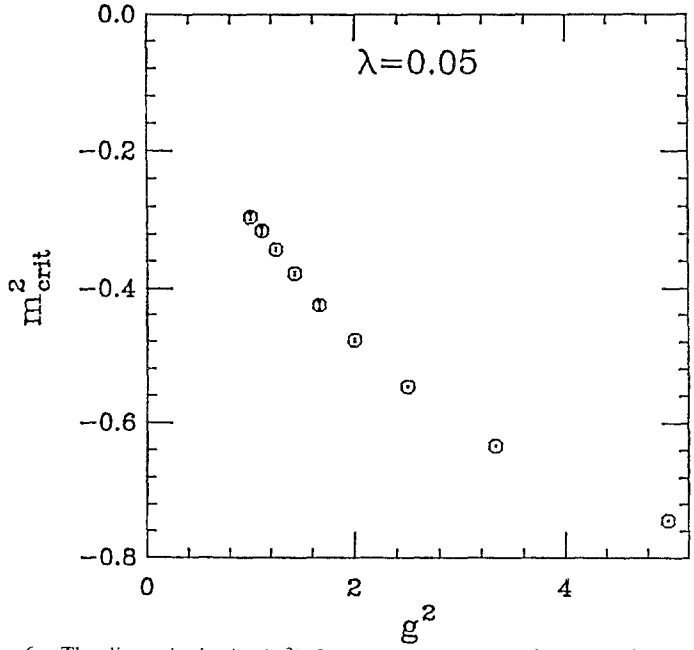

Fig. 6. The discontinuity in $\langle \varphi^2 \rangle$ for the same values of β ınd λ as in Fig. 5.

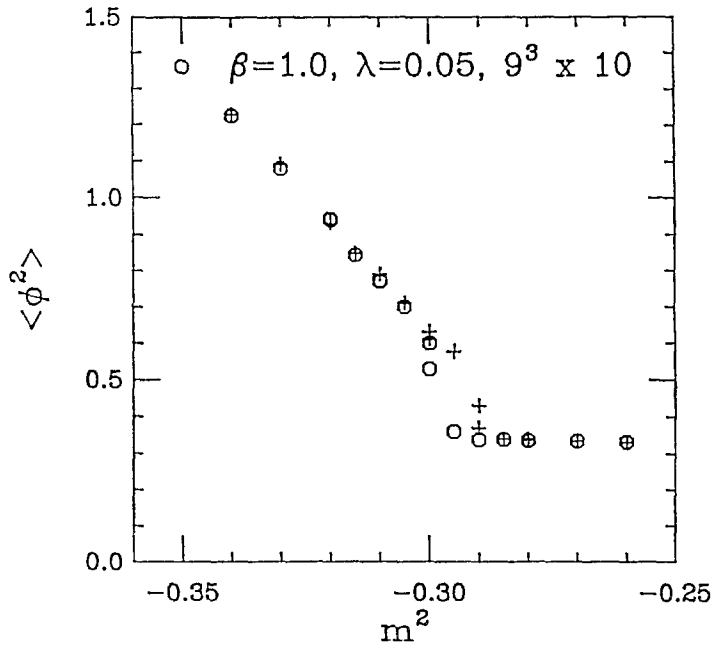

Fig. 7. Hysteresis loop for $\langle \varphi^2 \rangle$ for $\beta = 1.0$ and $\lambda = 0.05$, showing how the discontinuity and size of the metastable region decrease as β increases.

Once the critical mass is known it is straightforward to measure he discontinuity in any of the quantities of interest by making measurements at the critical point starting the system alternately in the ordered and disordered states. As an example, the discontinuity in $\langle \varphi^2 \rangle$ is plotted in Fig. 6 for $.2 \leqslant \beta \leqslant 1.0$ and $\lambda = .05$.

As β is increased the hysteresis loops become narrower and the discontinuities become smaller. An example is shown in Fig. 7. This result is to be expected since, as $\beta \to \infty$, the gauge degrees of freedom are frozen out and we obtain a lattice φ^4 theory with a second-order phase transition. Similarly the hyseresis loops shrink as λ increases for fixed β. Again this behavior is expected since the effective potential will only have broadly spaced minima for small values of λ. Thus, it is difficult to study the Coulob–Higgs phase transition for large values of β and λ. However, within the domain that we have studied, our results point to a first-order phase transition. We cannot, of course, rule out the possibility of the phase transition becoming continuous for large, but finite, values of β and λ.

4. CHARGE DENSITY WAVE TRANSITION IN A QUASI-ONE-DIMENSIONAL SYSTEM

A number of quasi-one-dimensional systems exhibit an anisotropic growth of correlations as their temperature is lowered. Examples are KCP and TTF-TCNQ, which undergo metal-insulator Peirls transitions. These materials can be thought of as consisting of sets of one-dimensional chains which are weakly coupled to each other. Such systems can be studied experimentaly through scattering experiments and through the measurement of thermodynamic properties such as the specific heat. Scalapino, Toussaint, and I have recently used the ST-100 to simulate a model of such systems.[6]

The model we have studied consists of an array of chains along which spinless fermions can hop. The fermions interact with each other via near-neighbor, intra- and inter-chain Coulomb interactions. The Hamiltonian of the system is given by

$$H = \sum_{i} \left[-t(c^{+}_{i+\hat{x}} c_i + c^{+}_{i} c_{i+\hat{x}}) + V_x (n_{i+\hat{x}} - \tfrac{1}{2})(n_i - \tfrac{1}{2}) \right.$$
$$\left. + V_y (n_{i+\hat{y}} - \tfrac{1}{2})(n_i - \tfrac{1}{2}) + V_z (n_{i+\hat{z}} - \tfrac{1}{2})(n_i - \tfrac{1}{2}) \right] \tag{3}$$

The chains run along the x axis. c^{+}_{i} and c_i are the creation and annhilation operators for fermions on the ith lattice site and $n_i = c^{+}_{i} c_i$ is the occupation number at this site. To date we have studied the case of a half-filled band and repulsive interactions, although it is straightforward to simulate other fillings and attractive interactions. I report on data taken for a 10×10 array of chains with 30 sites per chain. The couplings are taken to be $V_x = 2.0$, $V_y = V_z = 0.2$, and $t = 1.0$. Results for other couplings and other size systems are reported elsewhere.[6]

The simulation is carried out using the canonical ensemble, so the average value of an operator A is given by

$$\langle A \rangle = \frac{\mathrm{tr}[A e^{-\beta H}]}{\mathrm{tr}[e^{-\beta H}]} \tag{4}$$

where β is the inverse temperature. The imaginary time interval $0 \leqslant \tau \leqslant \beta$ is divided into L subintervals of width $\Delta\tau = \beta/L$, with L chosen large enough so that $\Delta\tau \cdot V_x$, $\Delta\tau \cdot V_y$, $\Delta\tau \cdot V_z$, and $\Delta\tau \cdot t$ are all small parameters. At each imaginary time interval a complete set of intermediate states, in which the fermions are localized on individual lattice sites, is introduced. The sum over intermediate states is equivalent to sum over all possible configurations of the fermion world lines, and it is this sum that is carried out by Monte Carlo methods. Details of this approach are given in Ref. 7. The

well-known difficulties which ordinarily arise in the numerical simulation of three-dimensional fermion systems are absent in this problem because the fermions are restricted to remain on individual chains.

As in the case of the Abelian–Higgs model, the ST-100 code was written in APCL primarily using the manufacturer's macro library. A special problem arose because of the need to use periodic boundary conditions in the space directions and antiperiodic boundary conditions in the imaginary-time direction for the fermions. These boundary conditions led to considerable indirect addressing in the data movement for which we wrote our own macros. The resulting code ran approximately 150 times faster than equivalent code on our VAX 11/750.

We focused our attention on the specific heat, the structure function

$$S(\vec{q}) = \sum_{\vec{l}} e^{i\vec{q}\cdot\vec{l}} \langle (n_{\vec{i}+\vec{l}} - \tfrac{1}{2})(n_{\vec{i}} - \tfrac{1}{2}) \rangle \tag{5}$$

and the operators

$$O_{\parallel} = \langle (n_{\vec{i}+\hat{x}} - n_{\vec{i}})^2 \rangle \tag{6}$$

and

$$O_{\perp} = \langle \tfrac{1}{2}(n_{\vec{i}+\hat{y}} - n_{\vec{i}})^2 + \tfrac{1}{2}(n_{\vec{i}+\hat{z}} - n_{\vec{i}})^2 \rangle \tag{7}$$

$O_{\parallel}$ and $O_{\perp}$ measure the order along and between chains, respectively. Each takes on the value of 0.5 in a completely disordered state and 1.0 in an ordered one.

The following picture emerges from our simulation. At high temperatures the system acts as a collection of independent one-dimensional chains. As the temperature is lowered the correlations along the chains increase, but not the correlations between the chains. At sufficiently low temperatures the chains begin to lock into each other, and a three-dimensional phase transition to a charge density state occurs.

In Fig. 8 the specific heat is plotted as a function of temperature. The crosses and their associated error bars are our Monte Carlo results, and the solid curve is the $N \to \infty$ extrapolation of the single-chain specific heat obtained by Bonner and Fisher[8] for $V_x/t = 2.0$. Notice how closely the specific heat follows the one-dimensional result at high temperatures. The sharp peak at $T \approx 0.44$ is associated with the three-dimensional ordering of the chains. Further evidence for this picture is given in Fig. 9 where we plot $O_{\parallel}$ and $O_{\perp}$ for the same lattice and couplings. Notice that $O_{\parallel}$, which measures correlations along the chains, increases steadily with a marked

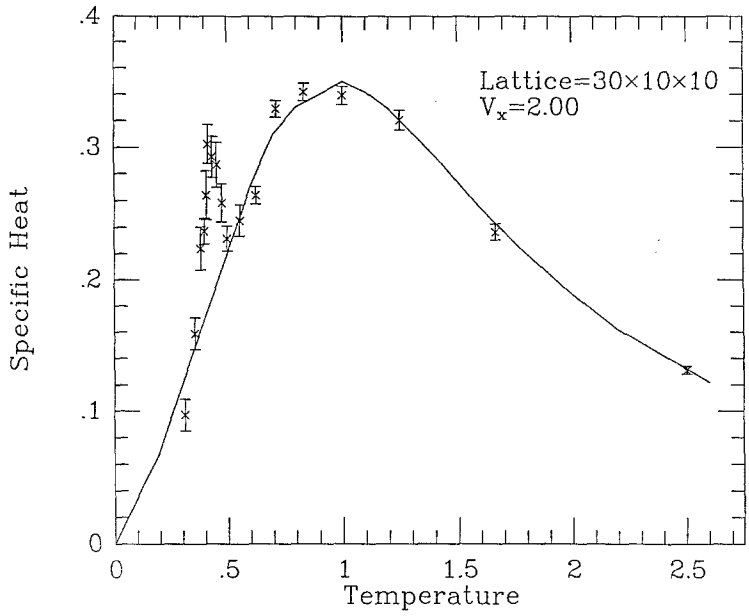

Fig. 8. The specific heat per site versus temperature for $V_x = 2.0$, $V_y = V_z = 0.2$, on a $30 \times 10 \times 10$ lattice. The solid curve is the Bonner–Fisher[8] result for a single chain extrapolated from finite-chain calculations.

break at the phase-transition point. On the other hand, $O_\perp$, which measures interchain correlations, remains near its uncorrelated value of 0.5 until the temperature is lowered to the phase-transition point.

In Fig. 10 the evolution of the structure function $S(q_x, q_y, \pi)$ is shown as a function of temperature. The data is again from a $30 \times 10 \times 10$ lattice with $V_x = 2.0$ and $V_y = V_z = 0.2$. As the temperature is lowered charge density wave correlations start to develop along the chains as can be seen from the increase in $S(q_x, q_y, \pi)$ with q_x. In Fig. 10(a), $S(q_x, q_y, \pi)$ is virtually flat in q_y, indicating an absence of correlations among the chains. This corresponds to the lines in diffuse x ray scattering. As T approaches T_c, interchain correlations develop rapidly, as is seen in Fig. 10(c), (d). Below T_c a Bragg spot develops in the x ray scattering associated with the spike at $\vec{q} = (\pi, \pi, \pi)$. In Fig. 11, $S(\pi, \pi, \pi)$ is plotted as a function of temperature.

Symmetry arguments indicate that the phase transition should be in the same universality class as the three-dimension Ising model, and our results are consistent with that prediction.

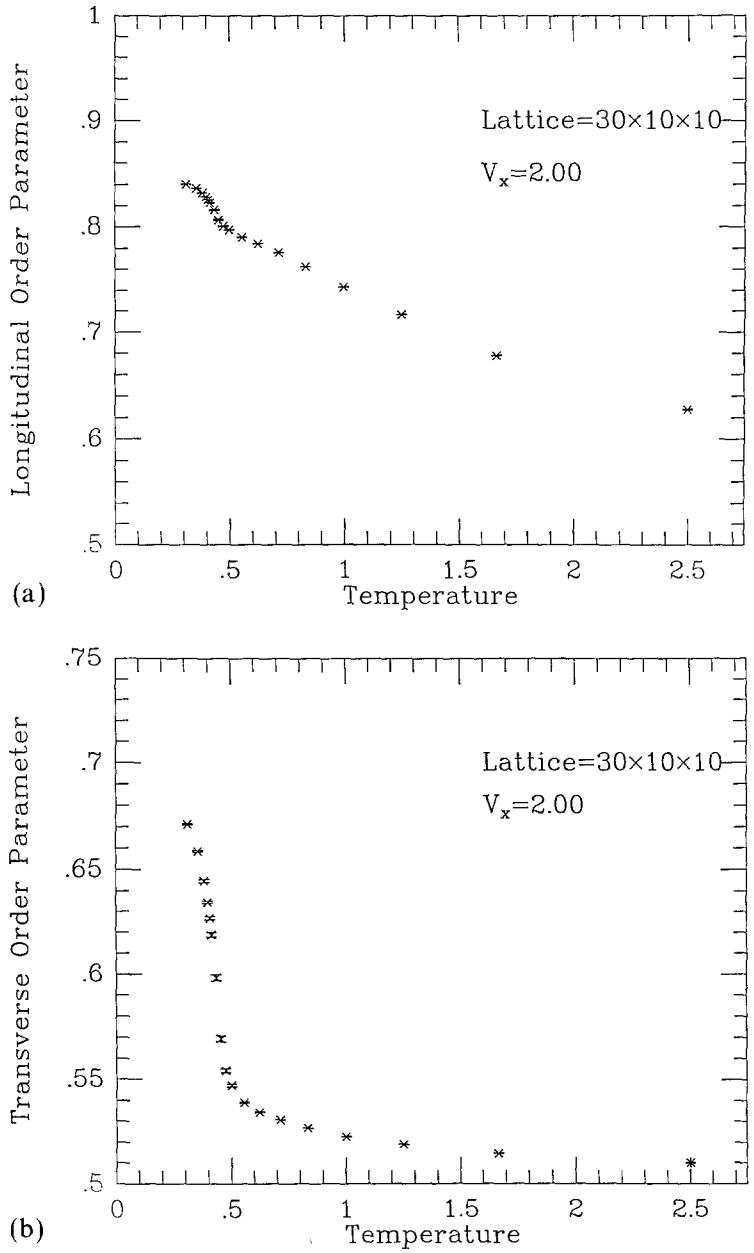

Fig. 9. The longitudinal correlation function $O_{\parallel}$ versus temperature. (b) The transverse correlation function $O_{\perp}$ versus temperature. In both cases the lattice size and couplings are the same as in Fig. 8.

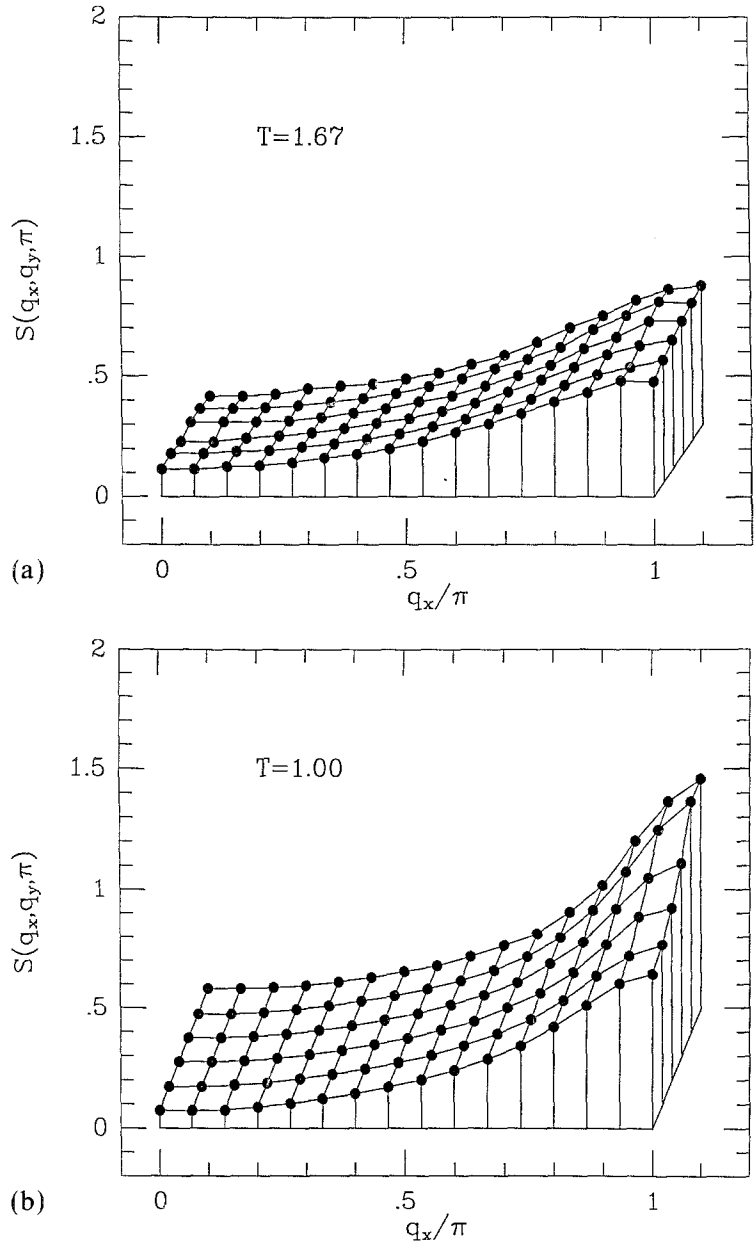

Fig. 10. The structure function $S(q_x, q_y, \pi)$ versus q_x and q_y for differnt temperatures. (a) $T=1.67$, (b) $T=1.00$, (c) $T=0.556$, (d) $T=0.455$. The growth of correlations along the chains is evident in Figs. 10(a) and (b), while the peak at $q_x=q_y=q_z=\pi$ in (c) and (d) indicate the development of three-dimensional order.

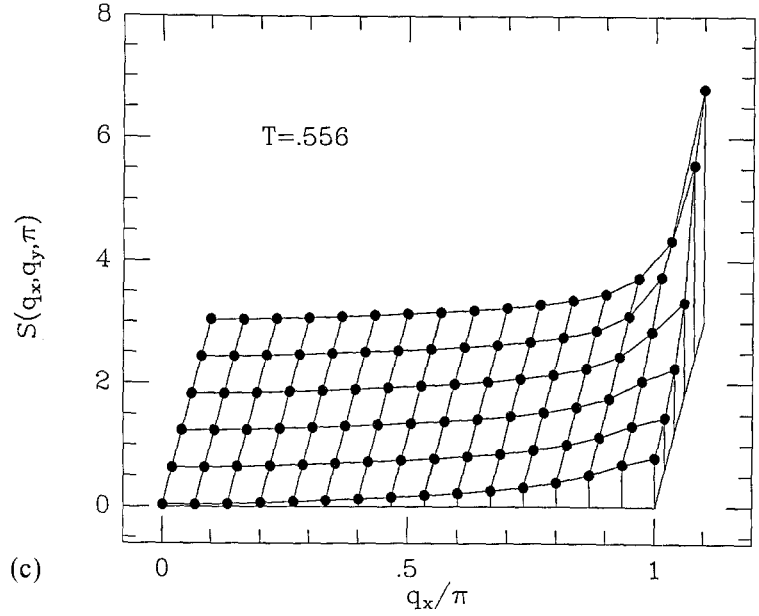

(c)

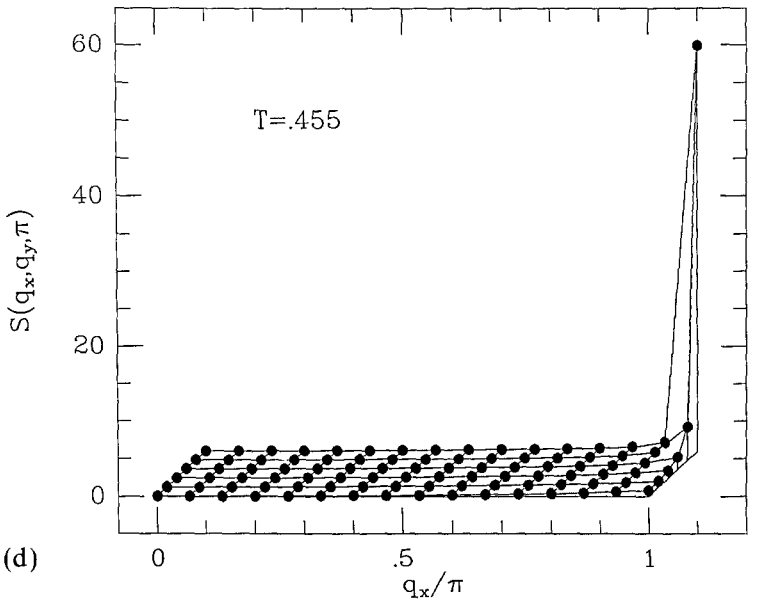

(d)

Fig. 10 (*continued*)

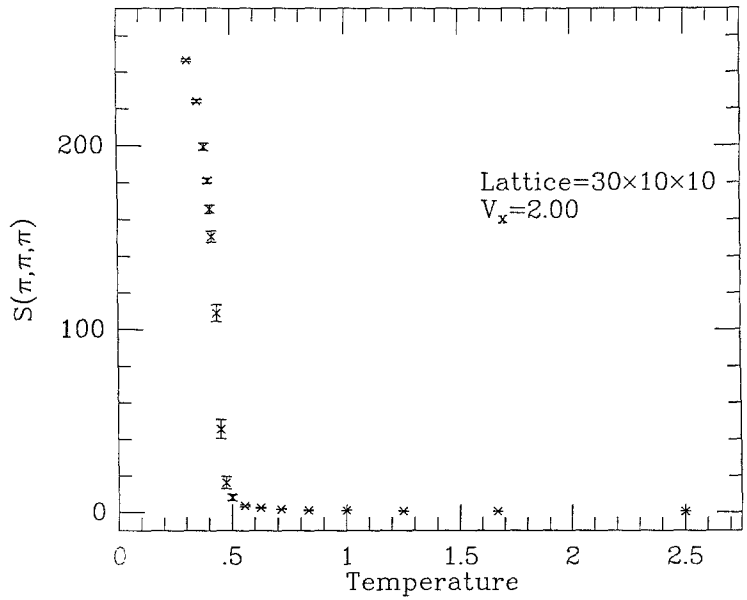

Fig. 11. $S(\pi, \pi, \pi)$ versus temperature for the canonical lattice and coupling constants.

5. THE DECONFINING PHASE TRANSITION IN *SU*(3) LATTICE GAUGE THEORY

The largest simulation that we have undertaken to date with the ST-100 is a study of the deconfining phase transition in $SU(3)$ lattice gauge theory.[9] One of the central problems in lattice gauge theory is to extract the physical quantites of the continuum theory from calculations performed on lattices with finite spacings. This can only be done with confidence in the scaling regime of the theory where the renormalization group β function is universal and calculable from perturbation theory.

The deconfining phase transition is a particularly good tool for studying the continuum limit. Locating the critical temperature T_c is relatively easy because the system undergoes a sharp first-order phase transition. Further, we expect that thermodynamic quantites such as T_c can be studied on smaller lattices than correlation functions at the same value of the lattice coupling g.

The transition temperature is a renormalization group invariant which depends on the lattice coupling constant through the relation

$$T_c = \text{const} \frac{1}{a} \exp\left(- \int_0^g \frac{dg'}{\beta(g')} \right) \tag{8}$$

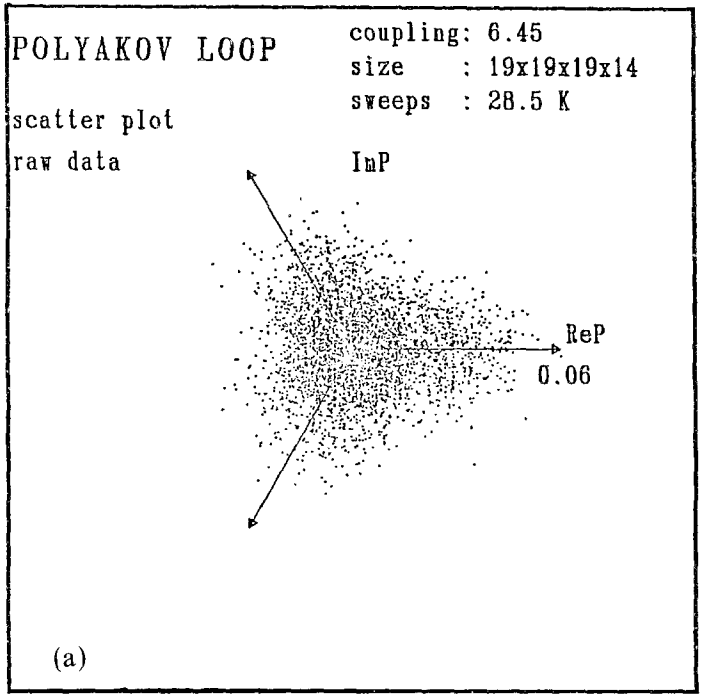

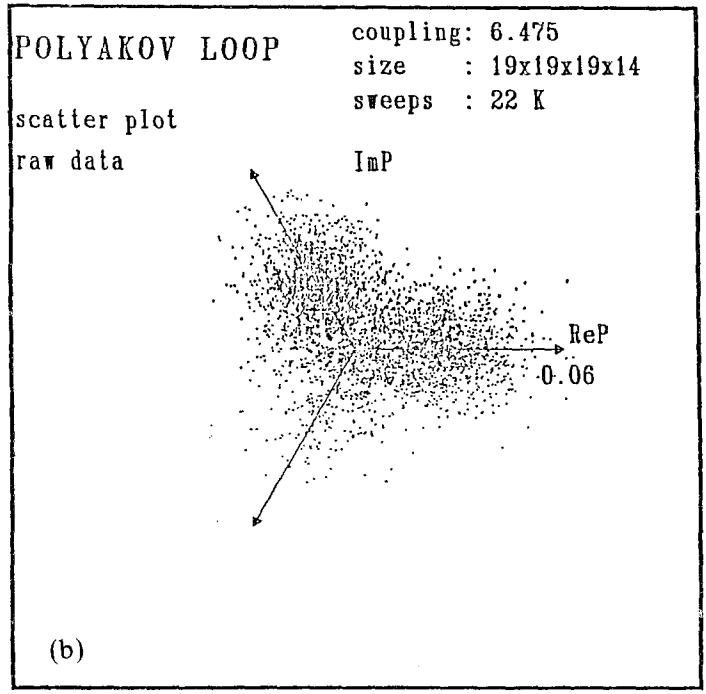

Fig. 12. Scatter plots for the Polyakov loop on a $19^3 \times 14$ lattice. The plots show the system (a) in the confined phase, (b) near the transition point with the confined and deconfined phases coexisting, and (c) in the deconfined phase.

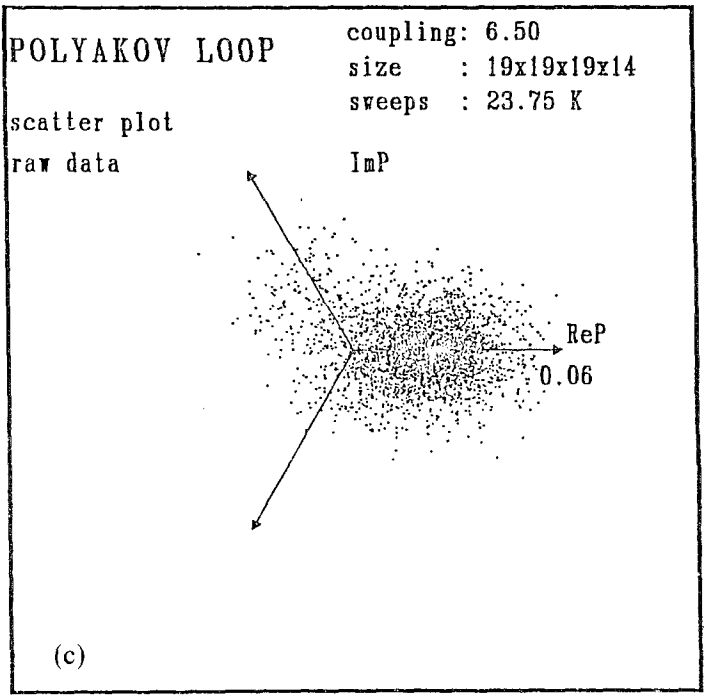

Fig. 12 (*continued*)

where a is the lattice spacing and β has the weak coupling expansion

$$-\beta(g) = b_0 g^3 + b_1 g^5 + \cdots \tag{9}$$

with $b_0 = 11/16\pi^2$, $b_1 = 102/(16\pi^2)^2$. Thus, the behavior of T_c is known for small values of g. Monte Carlo measurements of this quantity can be used to study the approach to the continuum limit and to determine the β function.

The order parameter of the deconfining phase transition P is the spatial average of the Polyakov loop

$$P(x) = \mathrm{tr} \prod_{t=1}^{n_t} U_t(\vec{x}, t) \tag{10}$$

where $U_t(\vec{x}, t)$ is an $SU(3)$ matrix along a timelike link at spatial location $\vec{x}$ and time t. n_t is the number of time slices on the lattice. In addition to the local gauge symmetry, the theory has a global $Z(3)$ invariance. Under the action of this symmetry the order parameter transforms as $P \to zP$, where z is an element of $Z(3)$. The expectation value of the Polyakov loop vanishes

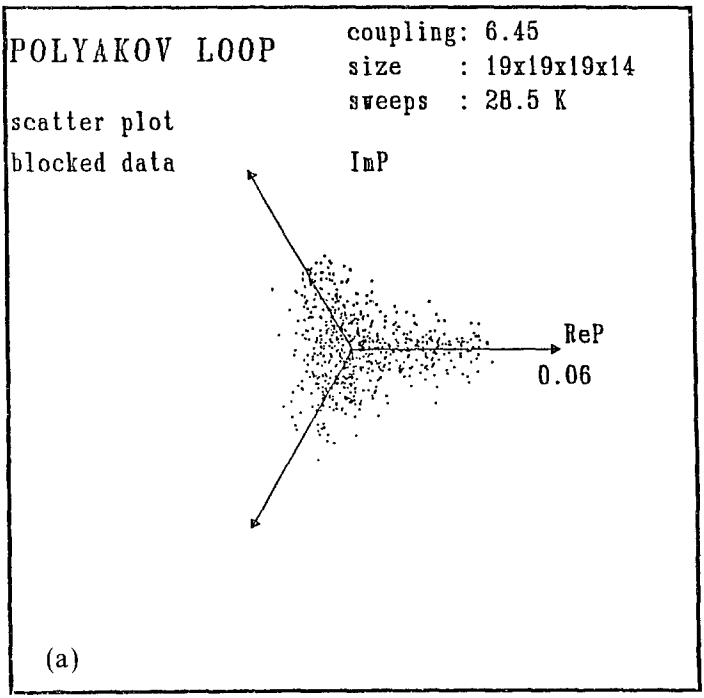

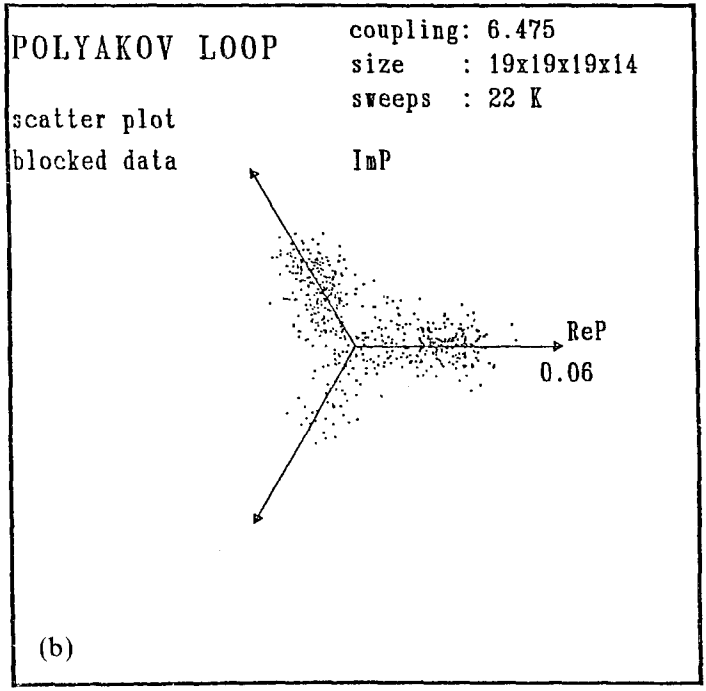

Fig. 13. (a–c) The same data as in Fig. 12. Five successive values of P have been averaged to reduce high-frequency scatter.

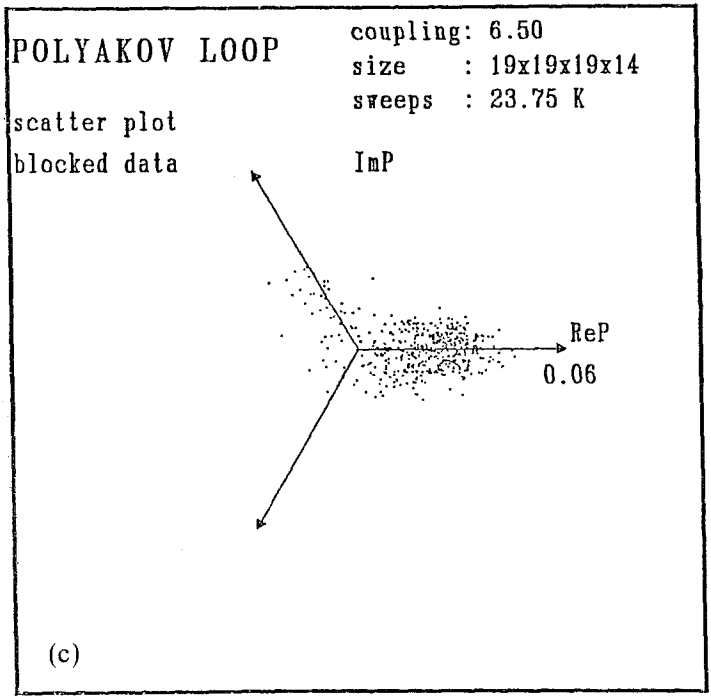

Fig. 13 (*continued*)

in the confined phase. In the deconfined phase, the $Z(3)$ symmetry is broken and the ground state is three-fold degenerate.

The calculations were carried out on Cyber 205 supercomputers and on the ST-100. The quasi-heatbath algorithm[10] was used on the former and the Metropolis algorithm on the latter. The fact that we obtained identical results from two different machines running different algorithms added to our confidence in the results.

Unlike the two previous problems that I have described, extensive use of assembly language was made in preparing the ST-100 code. Because of the size of the calculation this extra effort seemed warranted to obtain maximum performance from the machine. The ST-100 is very well-suited to the problem of lattice gauge theory. The even balance between adders and multipliers makes it especially efficient for complex arithmetic, and the three-stage pipelines make the multiplication of $SU(3)$ matrices particularly simple. The resulting code updated a link in 105 μs. This was an average time which included 15 Metropolis hits per link, measurements at every fourth sweep of the lattice, and reunitarization after every second sweep. For the larger lattices, the $SU(3)$ matrices were stored in main memory in

16-bit precision. Results with this precision agreed with those using full precision to within statistical errors. Although the ST-100 was approximately four times slower than the Cyber 205, it was able to contribute to this calculation on an equal footing because substantially more running time was available on it. This of course is the great advantage of a dedicated computer.

The number of time slices n_t on the lattice is related to the temperature through $T = a \cdot n_t$. To search for the critical point, we study a lattice with fixed n_t and vary the coupling constant g. The phase transition can be identified in a variety of ways. If one makes a scatter plot of the values of the Polyakov loop over a number of independent field configurations, then in the confined phase the points will be symmetrically distributed about $\Re P = \Im P = 0.0$. In the confined phase if one runs for sufficiently long times on a finite lattice, then three lobes will develop on the scatter plot corresponding to the three degenerate ground states in which $|P| \neq 0.0$. On the large lattices that we have studied tunneling is difficult, and all three lobes are not always filled in. However, it is still straightforward to identify the deconfined phase. As an example, scatter plots for the largest lattice we have studied to date, $19^3 \times 14$, are shown in Fig. 12. Figure 12(a) shows the system in the confined phase; (b) shows it near the transition point with the confined and deconfined phases coexisting, and (c) shows it in the deconfined phase. The plots can be clarified by averaging the Polyakov loop over several successive measurements to average out high-frequency scatter in the data. Figure 13 shows the same data as Fig. 12 with five successive points having been averaged.

Estimates of T_c can be made from visual inspection of the scatter plots and plots of the magnitude of the order parameter versus sweep number. A quantitative measurement of the degree of confinement is given by the "triality"

$$\Pi = \langle \cos(3 \arg P) \rangle \tag{11}$$

Here arg P is the phase of the order parameter, and the brackets indicate an average over Monte Carlo samples. Clearly Π vanishes in the confined phase and is one in the deconfined phase for an infinite spatial lattice. We found that all three measures of T_c gave consistent values.

Our main results are shown in Fig. 14 where the ratio of the measured values of T_c to the prediction of the two-loop β function is plotted. Our new results are for $n_t = 8$, 10, 12, and 14. For completeness, points at $n_t = 2$, 4, and 6 from an earlier work[11] are included. Notice that after the apparent early scaling between $5.1 < 6/g^2 < 5.7$ there is strong violation of scaling for $6.15 < 6/g^2 < 6.1$. Asymptotic scaling is observed in this study for

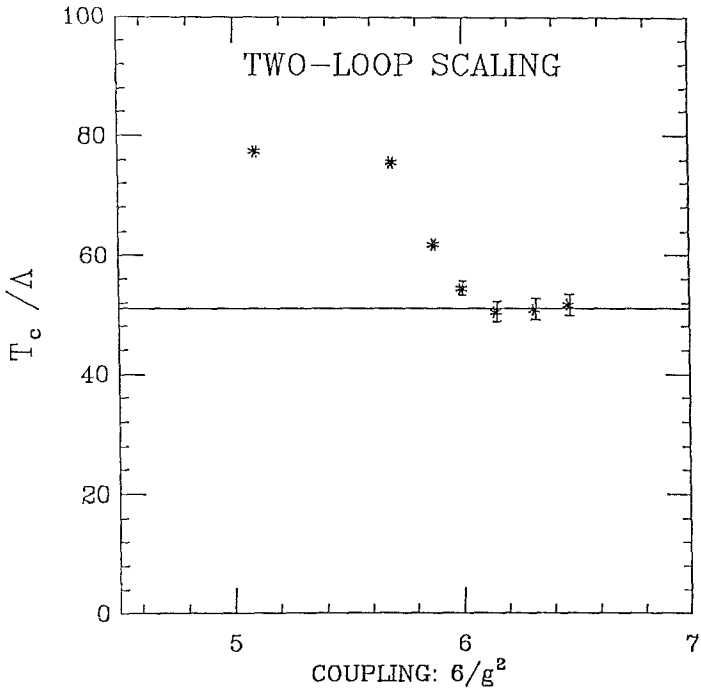

Fig. 14. The ratio of the measured values of T_c to the prediction of the two-loop β function. Our new results are for $n_t = 8$, 10, 12, and 14. Points at $n_t = 2$, 4, and 6 are taken from Ref. 11. Asymptotic scaling sets in for $6/g^2 > 6.15$.

$6.15 < 6/g^2 < 6.50$. Thus, the onset of scaling is at a much weaker coupling than early optimistic expectations. This means that either a significant increase in computer power or a substantial improvement on the Wilson action is needed for practical calculations of hadronic properties. However, our work does indicate that continuum quantites in pure $SU(3)$ lattice gauge theory can be obtained from Monte Carlo calculations with $6/g^2 > 6.15$ on sufficiently large lattices.

6. CONCLUSION

I hope that it is clear from this brief survey that the ST-100 array processor is a powerful tool for carrying out large-scale numerical simulations. In addition to the work reported on here, the ST-100 has been used to study $U(1)$ lattice gauge theory and the quantum Heisenberg model. Code is presently being developed to study the three-dimensional Hubbard model; to include phonons in the quasi-one-dimensional fermion

model described above; to study alternatives to the Wilson action in $SU(3)$ lattice gauge theory; and to include fermion in the $U(1)$ and $SU(3)$ lattice gauge theory codes.

ACKNOWLEGMENTS

Work reported on in this paper was supported by National Science Foundation grants DMR83-20423, PHY83-13324, and PHY82-17853 as well as Department of Energy grants DE-AT03-81-ER40029 and DE-FG0385ER45197 and contract DE-AT03-83ER45008. Support from Control Data Corporation, E.I. Dupont de Nemours and Company, and Xerox Corporation is gratefully acknowledged. Thanks are also due to the Rechenzentrun der Universitat Karlsruhe and the Florida State Supercomputer Research Institute for use of their Cyber 205s and to the Physics Division of Argonne National Laboratory for use of its ST-100 array processor.

REFERENCES

1. W. D. Toussaint and R. L. Sugar, *Phys. Rev. D* **32**:2061 (1985).
2. S. Coleman and E. Weinberg, *Phys. Rev. D* 7:1888 ((1973).
3. D. J. E. Callaway and L. J. Carson, *Phys. Rev. D* **25**:531 (1982); K. C. Bowler, G. S. Pawley, B. J. Pendelton, D. J. Wallace, and G. W. Thomas, *Phys. Lett. B* **104**:481 (1981).
4. Y. Munehisa, *Phys. Rev. D* **30**:1310 (1984); Y. Munehisa, *Phys. Rev. D* **31**:1522 (1985).
5. G. Koutsoumbas, *Phys. Lett. B* **140**:379 (1984).
6. D. J. Scalapino, R. L. Sugar, and W. D. Toussaint (to be published).
7. J. E. Hirsch, D. J. Scalapino, R. L. Sugar, and R. Blankenbecler, *Phys. Rev. B* **26**:5033 (1982).
8. J. C. Bonner and M. E. Fisher, *Phys. Rev. A* **135**:610 (1964).
9. S. A. Gottlieb, A. D. Kennedy, J. Kuti, S. Meyer, B. J. Pendelton, R. L. Sugar, and W. D. Toussaint, *Phys. Rev. Lett.* (to be published).
10. N. Cabibbo and E. Marinari, *Phys. Lett. B* **110**:387 (1982).
11. A. D. Kennedy, J. Kuti, S. Meyer, and B. J. Pendelton, *Phys. Rev. Lett.* **54**:87 (1985).