DR. NORMAN CHRIST
Columbia University

Many physicists view matter and in particular the strongly interacting particles as made up of quarks. The behavior of quarks is actually described by a theory that is specified in some detail called QCD (Quantum Chromo Dynamics). Figure 1 illustrates some of the strongly interacting particles such as the neutron, proton, and pi meson and contrasts them to the structureless point-like electron, muon, and photon. These strongly-interacting particles have a size and are believed to be made up of elementary, presumably point-like constituents, called quarks. The quarks interact with each other through a field, the gluon field, very much in analogy to the electromagnetic field that describes the interaction between electrons. Figure 2 illustrates the interaction between two electrons and the interaction between two quarks showing that there are many similarities. The gluon field between the two quarks is quite different from the electromagnetic field in that it obeys a non-linear equation and presumably is squeezed into a tube of flux between the two quarks so that the energy increases linearly with the separation of the two quarks. This suggests that a free quark is something that one will never see. This is quite a complicated problem. Classically one has a non-linear version of Maxwell's equations. Quantum mechanically one has a problem involving strong coupling.

Q C D

(QUANTUM CHROMO DYNAMICS)

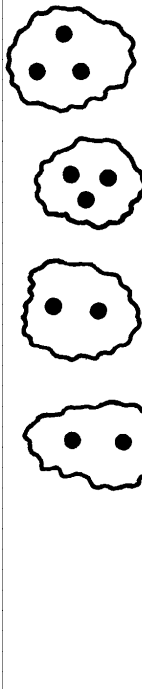

STRONGLY INTERACTING PARTICLES

NEUTRON

PROTON

π - MESON

ρ - MESON

•
•
•

NOT STRONGLY INTERACTING PARTICLES

ELECTRON •

MUON •

PHOTON •

•
•
•

Fig. 1. Strongly interacting particle contrasted to point-like particles.

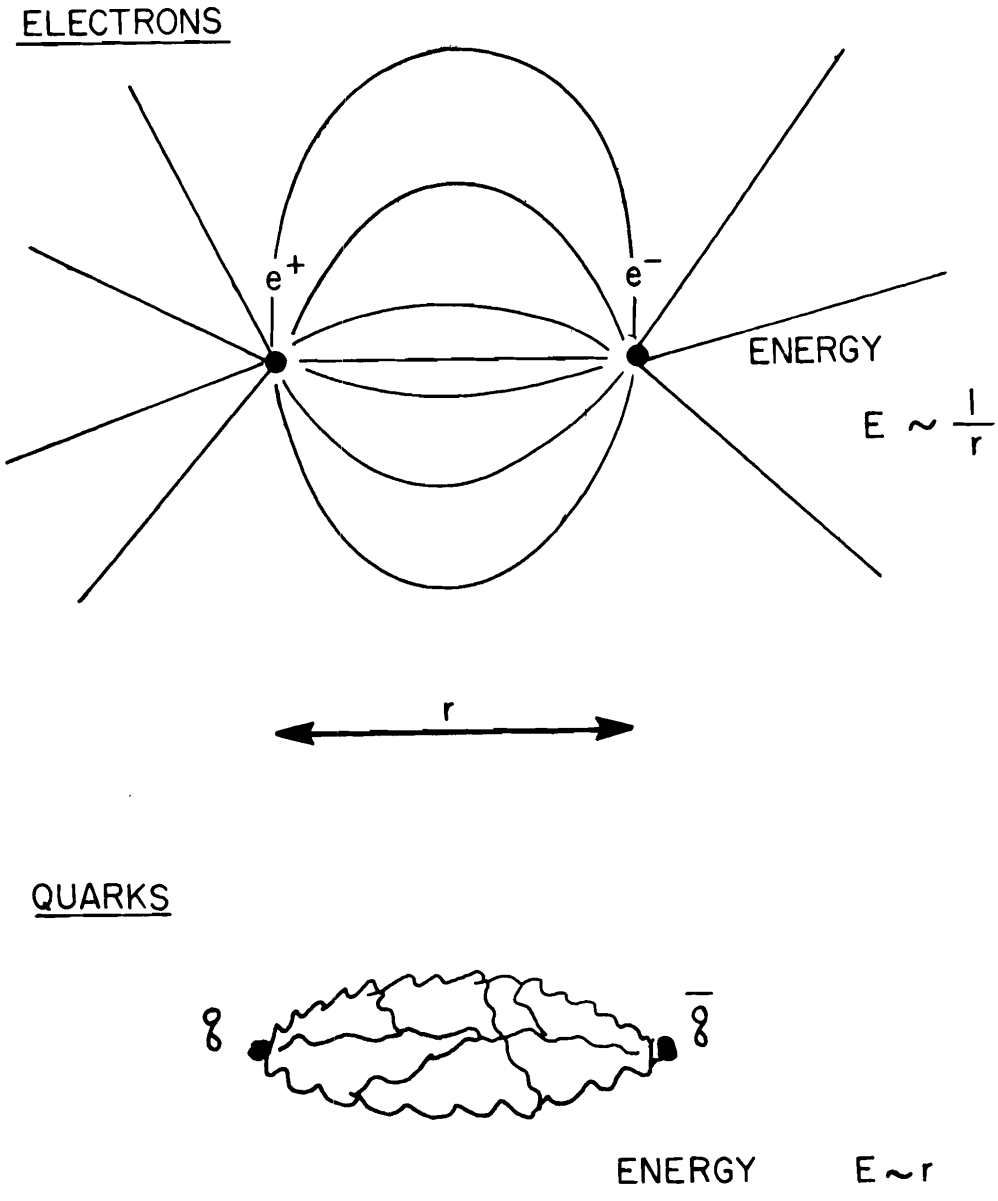

Fig. 2. Interactions between electrons compared to interactions between quarks.

Now consider a quantum mechanical state that starts off with some number of quarks and then add to it an array of anti-quarks to give a sea of various kinds of particles as shown in Fig. 3 (a). Then if the gluons are included, the system is very complicated indeed. These particles interact strongly. For example, the usual weak coupling approximation that's made to analyze the interaction of electrons and photons doesn't work. This is a problem where conventional theoretical techniques have made very little progress but in the last four years there has been significant progress using numerical methods. These methods begin by replacing the space-time continuum by a rectangular grid so one imposes on the problem a lattice structure and requires that all of these particles lie on the lattice. Finally the picture looks something like Fig. 3 (b) with the quarks on the vertices of the lattice and the gluons going on the links between them.

The quantum mechanical problem involving all of these degrees of freedom is best approached by the Feynman path integral, that is the Feynman sum over histories. The actual quantity that must be computed is the rather simply specified, but in fact quite complicated, integral shown in Fig. 4. The idea is that to each link in the lattice one associates a three by three matrix. Imagine that you want to measure a physical observable (O).

GLUONS, QUARKS AND ANTIQUARKS

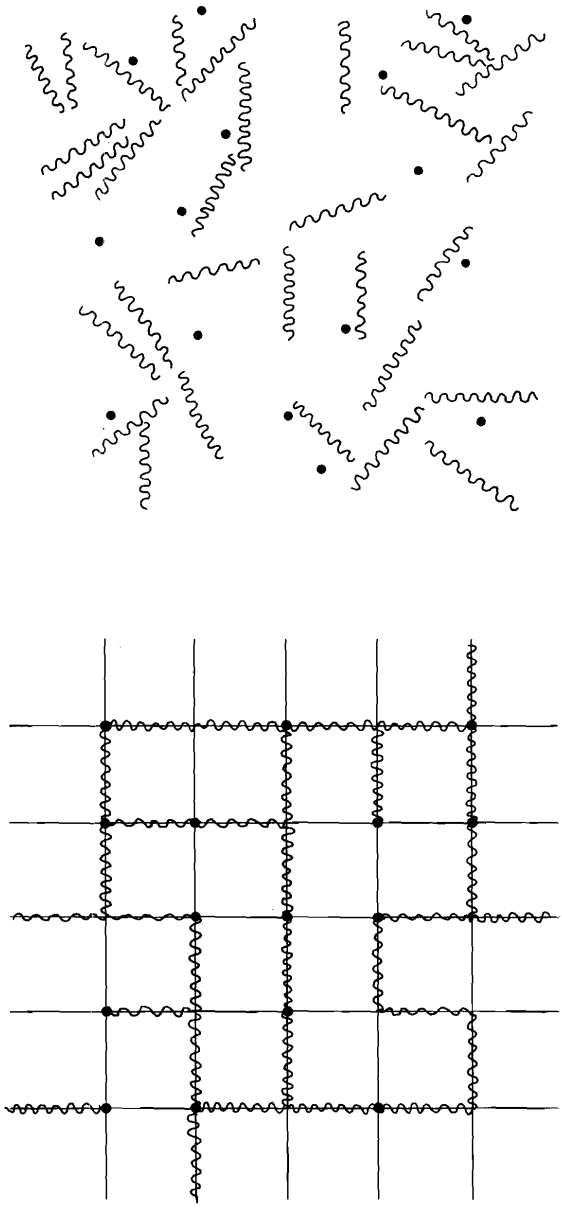

ig. 3. Solving QCD. (a) a sea of quarks and antiquartks nteracting through gluons; (b) the system is modeled by a attice structure.

$$\langle O \rangle = \frac{\prod_{\ell} \int dU_{\ell} e^{-\beta \sum \mathrm{tr} U_{\square}} \det [\not{\partial}] O(U)}{\prod_{\ell} \int dU_{\ell} e^{-\beta \sum \mathrm{tr} U_{\square}} \det [\not{\partial}]}$$

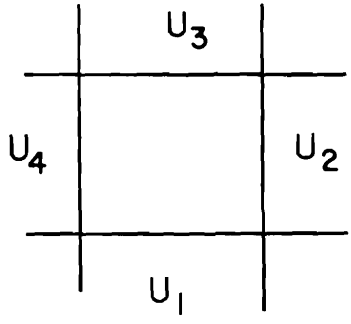

$$U_{\square} = U_4 \; U_3 \; U_2 \; U_1$$

Fig. 4. The equation for a QCD observable.

This may be an energy or mass or correlation function. It is necessary to take that observable (O) as it depends on the degrees of freedom and integrate over all the degrees of freedom, that is all of the 3 by 3 matrices corresponding to all of the links in the lattice. This integral is weighted with the exponential of a sum of traces. Each term is a trace of a 3 by 3 matrix, one matrix for each elementary square in the lattice, where that matrix is constructed by multiplying together the four matrices corresponding to the four links that bound the square. For a big lattice there are a lot of squares and a lot of traces. The worst thing is the determinant det. Here is an operator defined on the lattice in the discrete approximation. It is also a matrix but a matrix whose number of rows and numbers of columns equals the number of vertices in the lattice. Interesting preliminary results that are not at all satisfactory have been obtained by using lattices as large as 10 by 10 by 10 by 10. This is in four dimensional space time, so that such a lattice has ten thousand sites and 40 thousand links. There are eight variables in each of these matrices, 320 thousand degrees of freedom in this integral, and finally the determinant of a matrix which is 120 thousand by 120 thousand.

The problem has now gone into a regime where the number of degrees of freedom are so large that very good use can be made of statistical techniques. The integrals here are really quite successfully treated, it appears, by using a Metropolis technique, that is a Monte Carlo algorithm of the Metropolis type. One generates samples of configurations, assignments of

matrices to links, distributed according to the product of exponential and determinant in the integrand of Fig. 4. The expectation value, that is the value of a measurable quantity O is gotten by averaging O over the ensemble. The problem is such that recent calculations have used 10 hours of Cray time and in one case 100 hours. This is just touching the surface of the problem. Because the technique is statistical, it is necessary to run the program 100 times longer to get 10 times the accuracy. In addition, it is desirable to deal with much bigger lattices.

The problem, then, requires two or three orders of magnitude increase over the amount of power that is being devoted to it today. Also this problem, the physics of strongly interacting particles, may not be the most interesting one. This is a class of phenomena that experimental physicists have studied for the past 20 or 30 years. Both the theory and the experimental results are known and here one is just making the connection. However this type of theory, these strongly coupled gauge theories, are believed to explain perhaps all of reality and there are very large areas where the theory is not yet known, where the experimental results aren't known and the calculations are much harder.

Two of us at Columbia, Tony Toronto and myself, have designed and are building a special purpose computer intended to give this needed increase in computer power. The computer takes advantage of special properties of this particular problem. The interactions, the physics of the particles on the lattice, are local so that we could easily do with the the kind of

architecture that David Wallace just described. This is a grid of processors, arranged in two dimensions with only nearest neighbors in communication. The array is homogeneous. In fact the same physics is going on at every node. At least in one mode the processors could conceivably operate in lock step with the same calculation being done at every site. The matrix multiplication, which is the big difficulty, is heavily arithmetic but it is very organized so that it can be easily pipe-lined. Finally, because the whole problem is statistical and the answers are not very precise, the method is one which doesn't require high numerical precision. So what we propose is an array of processors, perhaps in the end a 16 by 16 array, capable of doing this kind of arithmetic very fast. The structure is shown in Fig. 5. The square boxes are memories, each containing the data for those sites and links with a group of x and y coordinates but all values of z and t. The circles are processors. Each neighboring pair of memories is connected by a single processor. The design of the processors is quite straightforward. One begins with a microprocessor, the Intel 80286, that is really a quite fast and sophisticated, and also general purpose. A specially designed arithmetic unit is added to that. The memories are divided into two independent halves from which two arguments can be simultaneously fetched to perform the multiplication. The result of the previous multiplication can be accumulated with that of the previous additions and finally the result written back into one of the memories. This is all done in a pipeline fashion at 8 megahertz so 16 million

floating point operations can be performed per second. The whole process is controlled by a microprogram which can contain the instruction for doing one of these matrix multiplications.

This year we have been talking about the physics of SU3. Next year it may be E6, SU5 or whatever looks interesting. The point is that the device must be some what general. Finally, these devices have to be coupled to their neighbors. We do this in the crudest possible way. All of the operations are supposed to be synchronous. When the communication between neighbors is occurring all of the processors have to be executing the same instructions in lock-step; there is no hand-shaking between units. The multiplier can get its arguments from its local memory or from its neighbor's memories at exactly the same rate, that is at 16 megabytes per second per node per operand. The final result is a fairly inexpensive node. We have one built and two-thirds working, at a cost of $2,500 for the single node. All the nodes there are identical so it's possible to make one board and then reproduce them. We hope to hook together 256 of them to achieve 4 billion floating point operations a second.

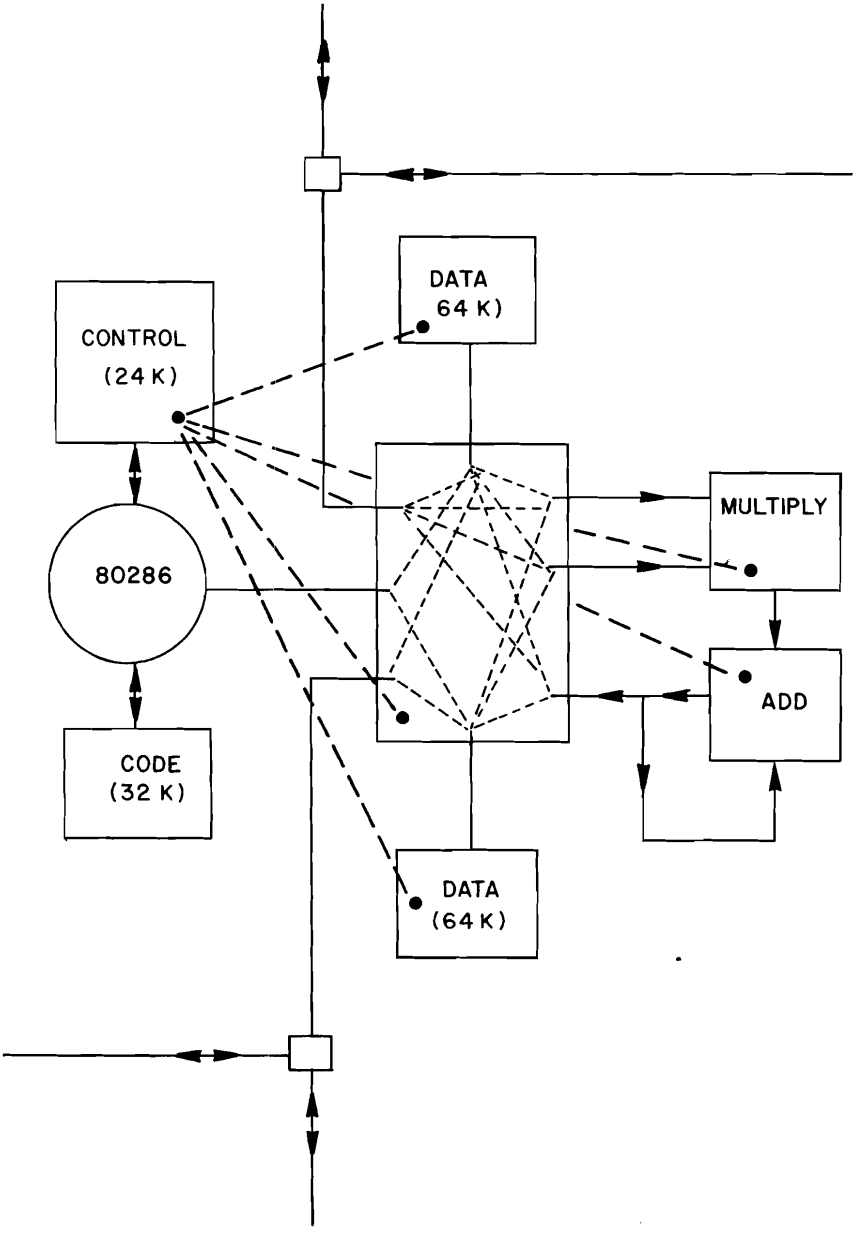

Fig. 5. The Columbia special purpose QCD computer.