vored physicists.

EDWARD GERJUOY
University of Pittsburgh
Chairman, APS Committee on the International Freedom of Scientists, and the following past and present members of the Committee:
PIERRE HOHENBERG
Bell Telephone Laboratories
HERMAN FESHBACH
Massachusetts Institute of Technology
JULIAN HEICKLEN
Pennsylvania State University
J. DAVID JACKSON
University of California
JOEL L. LEBOWITZ
Rutgers University
PETER S. PERSHAN
Harvard University

5/84

In taking the position that in the light of present US–USSR relations it is more prudent to restore lines of communication with Soviet scientists than to maintain boycotts as a response to Soviet human-rights violations, we did not intend to convey in any way the impression that this position is official policy of the American Physical Society. (Opinions presented in PHYSICS TODAY *editorials generally do not express the official policy of AIP or its member societies.) We are in enthusiastic agreement with the Committee on the International Freedom of Scientists that those physicists who choose to reopen contact with Soviet physicists can and should, whenever the opportunity arises, express their outrage about the human-rights violations to their Soviet counterparts and to the Soviet government.*

—The Editor

Physics on microcomputers

In his interesting article, "Doing Physics with Microcomputers," Per Bak (December, page 29) arrives at a number of thought-provoking conclusions. As a person whose background is in physical science (astronomy), who is the owner of a microcomputer and whose position is director of a computing center, I would like to comment on some of these conclusions.

That the microcomputer is more than a toy and can be used for serious scientific research is undeniable. My latest paper was based largely on calculations made on my micro. Only the most demanding calculations were performed on the computing center's VAX-11/780. Interestingly, in the advantages attributed to the micro Bak fails to mention that computations performed on a micro can be *more* accurate than those made on a large computer. If the latter truncates the results of floating-point arithmetic operations (as does an IBM 4331/4341, for instance) and the former rounds, there will be less statistical accumulation of round-off error in the micro. With rounding, the total error in a numerical integration, for example, is proportional to the square root of the number of arithmetic operations executed, rather than linearly proportional to the number, as with truncating.

Having said this, I nevertheless feel that Bak stacks the cards too much in favor of the micro. His economic considerations do not include the time a physicist needs to develop a correctly working program. Large computers come equipped with more powerful text editors, interactive debuggers and more flexible and reliable software than is encountered on micros, and the computing center offers access to standard subroutine packages and consultations with computer professionals, all of which significantly decrease program-development time. This is an important consideration if we assume that a physicist should primarily do physics rather than programming. Furthermore, Bak's estimate of a ratio of 40 to 1 between a micro's and a large computer's CPU processing time appears too optimistic. In an actual comparison between the VAX-11/780 and my Radio Shack Model-II (the benchmark being an n-body solar system integration), I found the ratio to be nearly 1600 to 1 for computations performed in double-precision floating-point with machine code produced by FORTRAN compilation. A typical VAX four-hour run would need 267 days on the micro.

It is no longer true that large computers' speeds are adversely affected because they use "inefficient FORTRAN or similar programs." Modern compilers generate extremely efficient machine code. One may be able, as Bak states, to minimize program-execution time by coding in machine language (which is *not* simple if the computer has a large instruction set) or, better, assembler program sections where instructions are executed many times, inside nested DO loops for example. But one should not get carried away and code the entire program in assembler out of misplaced concerns for "efficiency." The few seconds of CPU time typically saved will not compensate for the enormously increased program-development time compared to that needed when one codes in a high-level language.

RICHARD BRANHAM
Centro Regional de Investigaciones Scientíficasy Tecnológicas
Mendoza, Argentina

1/84

•

As manager of a physics-department computing center, I would like to comment on several points made in Per Bak's article. First, I'd like to state that I know absolutely nothing about physics but have been intimately involved in scientific computing for more than 7 years.

letters

The first point I'd like to comment on is Bak's implication that microcomputers are "accessible to physicists 24 hours a day." If this only were true. However, microcomputers have the same problems as any other piece of electronic equipment—they break. While it may be true that microcomputers break down less often than larger computers due to their relative simplicity, they still break.

The second point is the cost of computing. The $500-per-hour figure mentioned is disturbing. It must be for time on Cyber or large IBM computers, because time can be bought on VAX 11/780 computers, another popular computer in physics, for less than $100 an hour from several commercial profit-making companies here in Santa Barbara. The VAX 11/780 under my charge costs only $6.70 an hour, because only maintenance and operations costs are being recovered. I've noticed that high computing costs are usually due to central computing facilities, which have higher overhead than departmental computing centers. It is also important to point out that when a departmental or central computer is used, such expensive and high-speed devices as disk and tape drives, printers and graphics devices are also being paid for by the seemingly high hourly cost. Adding such devices to a microcomputer often costs more than the microcomputer itself. When such factors are considered, the expense of computer centers becomes more reasonable.

Third, Bak claims that machine language, his main tool in getting adequate performance from his microcomputer, is "really not difficult to use." Many experienced computer programmers, from many different backgrounds, would disagree. In fact, one of the reasons high-level languages, such as BASIC, exist at all is exactly because of the inherent difficulty in programming in machine language. Although Bak admits that the Ising model is particularly well-suited to the microcomputer, I hope that readers realize that many of the programming-language features they have come to expect in FORTRAN, BASIC or PASCAL—such as arrays, floating point, strings, and "structured" control structures—have no direct counterpoints in machine language. How many physics problems are there that do not use at least one of these features? Bak mistakenly dismisses the use of high-level languages because they are "inefficient." While this statement can itself be disputed, I believe that the time needed to write, debug, maintain and document a program written in machine language is so much greater than the time needed to write the equivalent program in a high-level language that it is the machine-language approach itself that must be labeled inefficient.

My goal in mentioning these points has not been to claim that physics cannot be done on microcomputers. It obviously can. In fact, with the advent of special floating-point processors such as the Intel 8087, it will be reasonable to expect Vax-like performance from desk-top computers. What many people are failing to realize is that the user will be forced to take on many of the tasks now performed by their computing center's support staff. Doing so will lower the cost per hour of computing but will mean that time that could be spent doing science will instead have to be put into supporting a computer. In other words, physicists should keep in mind that the cost of computing is more than simply the cost of the computing machinery being used.

Jon Forrest
University of California at Santa Barbara
1/84 *Santa Barbara, California*

•

The article by Per Bak about using microcomputers for large-scale calculations in physics is full of misinformation and misconceptions. As users of micro-, super mini- and nearly supercomputers, we feel obligated to indicate the relative power of the various machines. Before detailing Bak's mistakes, let us show some comparisons of an LNW 80 (4MHz Z80), IBM PC (4.77 MHz 8088), IBM PC + 8087 floating-point chip, VAX 11/780 with floating-point accelerator, and an FPS 164 attached processor (5.5 MHz). First, we simply performed many double precision floating-point multiplications, additions and divisions in a FORTRAN program with the results:

	Time (μsec)	
Operation	**LNW 80**	**VAX 11/780**
addition	500	5.5
multiplication	1300	6.3
division	2000	11.1

A more realistic test case was run using the following FORTRAN program:

```
Implicit Real *8 (A-H, O-Z)
DO 10 I = 1,100000
A = I
B = A*A
C = B + 2.0 DO
D = C/2.0 DO
E = D − 1.0 DO
F = DSQRT (E)
10 Continue
Stop
End
```

continued on page 100

letters

continued from page 15

This yielded the results

Computer	Time (sec)
LNW 80	17 760 (1730[a])
IBM PC[b]	1222
IBM PC + 8087[b]	194
VAX 11/780	7.25
EPS 164	0.90

[a] single precision time
[b] courtesy of M. Gordon (Chemistry, North Dakota State Univ.)

The above values imply a few points. First, clockspeed is *not the only or even dominant variable* determining calculational speed, as can be seen by comparing the LNW at 4 MHz, IBM PC at 4.77 MHz and the FPS 164 at 5.5 MHz. Indeed, wordlength (that is LNW single versus double precision times) and the number of operations done in a clockcycle are critical, as are instruction pre-fetching and decoding and other forms of parallelism typically present in larger machines but not micros (yet).

Second, assuming that Bak's VIC 20 is comparable to an LNW, and that the FPS 164 is comparable to a CDC 7600, then a reasonable ratio of computer speed (CDC 7600: VIC 20) is between 2000 and 20 000 to 1. Thus, Bak's 160-CPU-hour calculation on the VIC 20 would take between 4.8 and 0.48 min on the CDC 7600, which at \$500 per CPU-hour would cost between \$40 and \$4. A more realistic assumption of VIC 20 $\simeq \frac{1}{4}$ LNW 80 and FPS 164 $\simeq \frac{1}{2}$ CDC 7600 yields a cost estimate of \$5 to 50¢. Even though the costs are comparable (and not skewed toward the VIC 20 as Bak states), the CDC 7600 provides more convenience *and* availability: The VIC 20 is tied up on one job for 160 hours, or about one week, while the 7600 can run many of the jobs in one day. Because the results of one calculation often indicate what should be done next, the one-week "turnaround" on the VIC 20 is impossibly long. Misleading comments about machine coding are irrelevant since this may be done on mainframes as well as micros. In addition, most people realize machine coding is generally not cost-effective, because the speed-up over a good FORTRAN compiler is usually less than a factor of 3, while requiring large investments of programming time.

We would like to emphasize our fondness for micro-computers, of which we own three. They are excellent for many purposes such as word processing, graphics and so on. However, they are not adequate for many areas of scientific computing at present, and even the certainty that they will be much more powerful in the future does not really help the situation. Remember, supercomputers have also increased greatly in speed, and now are being designed to reach a speed of 10 gigaflops (10^{10} floating-point operations per second) within a decade. Since scientists generally try to treat systems that are as complex as possible, they will require the fastest, most powerful machines. It is unlikely that micros can be of much use, irrespective of Bak's self-professed expertise and zealousness.

Andrew E. DePristo
Stephen T. Elbert
Iowa State University
12/83 *Ames, Iowa*

•

I would like to comment on Per Bak's article in December. I got the general impression from some of his remarks that he is not fully aware of the important role that supercomputers have played in the process of scientific discovery and applications within the broad discipline of physics. It is unfortunate that his article has furthered two popular misconceptions about the use of microcomputers in doing physics calculations. In particular I wish to dispute the statements that "Most calculations in physics will be done on these 'home' computers" and that "The cost will be but a fraction of that for central-processor time at a typical computer center."

At the National Magnetic Fusion Energy Computer Center (where I work as a computational physicist), 80% of our computer use is by physicists doing large-scale calculations of plasma-physics phenomena. Our present mainframe computers, Cray-1s, are barely capable of simulating these physical phenomena because of limitations on both memory and speed.[1] This is in spite of the fact that these computers are among the fastest available and possess memories an order of magnitude larger than available with microcomputers. Memory limitations alone prevent one from doing large-scale calculations on the microcomputers. Therefore, I would replace the first statement with the statement that microcomputers are not able to perform most large-scale calculations in physics because they have insufficient memory.

I would agree that there is a class of problems—we might label them as "small-scale physics calculations"—for which the microcomputer is capable, albeit inefficient. To estimate these efficiencies, one must consider the computer's speed and its cost, as well as duty-cycle utilization. Central computer facilities typically have a very small fraction of idle time—corresponding to a high duty-cycle utilization. On the contrary, the utilization of microcomputers is rarely larger than half and more likely to be in the range from 10 to 20%. These high utilization rates are most easily realized in a central computer facility where there are large numbers of users.

Real-time considerations—turnaround—are important and also favor the mainframes. Otherwise the physicist may need to wait too long for an answer, thereby not using his own time effectively.

Putting the issues of utilization and real-time performance aside, let us assume that one is a very clever and dedicated physicist who is able to keep his microcomputer busy all of the time. Most physics calculations I am aware of spend most of their time doing floating-point arithmetic with a minor fraction being devoted to input and output. I have statistics, based on typical costs of purchasing computers, that give the cost per installed megaflops capability that one might regard as a figure of merit to use to compare the cost efficiency of these machines. This value is just the ratio of the computer's cost to its speed in megaflops. With the exception of the VIC-20, the speed data were obtained by Jack Dongarra[2] in a rather exhaustive study of over 50 computer–software combinations. The VIC-20 result, not provided by Dongarra, is an extrapolation of the IBM-PC results based on intrinsic CPU cycle times (1 microsec versus 0.25 microsec, respectively). In each case the best available FORTRAN compiler was used and a

Ratio of computer cost to speed compared

Computer	Cost of Computer (\$)	Memory (bytes)	Mflops on LINPAC with FORTRAN	Relative cost (10^6 \$/Mflops)
VIC-20	1 k	20 k	0.00009	11.1
IBM-PC No 8087 Chip	4 k	256 k	0.00037	10.8
IBM-PC with 8087 Chip	5 k	256 k	0.0073	0.68
VAX 11/750	80 k	8 m	0.057	1.4
DEC KL-10	480 k	1.1 m	0.18	2.7
CDC-7600	4.5 m	3.7 m	4.6	0.98
Cray-1S	11 m	16 m	18	0.61
Cray XMP (2 CPU's)	15 m	32 m	53	0.28

Cost per unit installed computing power is estimated for various personal and mainframe computers. Only the computer's initial cost is used, including its disk drives. Inclusion of other costs of running a computer facility would not significantly change the results.

precision equivalent to 64-bit words was obtained. The accompanying table displays some of these comparisons.

Comparing the Cray-1 to the VIC-20 we find (in a FORTRAN environment) the mainframe is more than an order of magnitude less expensive per installed megaflops. I would expect a similar cost ratio if both machines were programmed directly in machine language. Thus the second statement should be revised to say that the mainframes are more economical even for small-scale physics calculations.

Other comparisons may be gleaned from the table. Some of the newer home computers, such as the IBM-PC with the 8087 chip, perform at a cost somewhat more than a Cray-1. However, the Cray-1 is eight years old and is therefore not a contemporary of the IBM-PC. The Cray-XMP is contemporary with the IBM-PC and is significantly less expensive per installed megaflops, as seen in the table.

Looking ahead, one sees microcomputers and the mainframe supercomputers both becoming much more capable, with bigger memories and faster CPUs. The supercomputers are rapidly becoming much smaller and much less expensive. In the next decade or so, mainframe computers will be literally table-top size (as dictated by signal-speed constraints).

Economics of calculation, rather than empire building, has most often dictated the choices to be made in the selection of many current mainframe computers. In the future, I suppose economics, rather than speed per se, will continue to dictate computer acquisitions. It is my prejudice that the faster machines will remain the most economical; the old adage, "time is money," will continue to be pertinent.

References

1. Bruce I Cohen and John Killeen, PHYSICS TODAY **36** 54 (May 1983).
2. Jack J. Dongarra, Argonne National Laboratory Report ANL-MCS-TM-23, (December, 1983).

DAVID V. ANDERSON
National Magnetic Fusion Energy Computer Center
1/84 *Livermore, California*

•

The recent article by Per Bak on the use of personal computers for physics in science, while informative, left out some important particulars. I have used computers since the old IBM 7040 days and still found the purchase of a suitable microcomputer somewhat bewildering. For a start, the number of bits per register in the microprocessor is an important factor. The microcomputers mentioned in the article have 8-bit microprocessors. There are now a number of 16-bit PCs and even a few full 32-bit "supermicros." All perform numerical computations with the same accuracy, representing a single-precision real number by 32 bits. However, an add between two reals, for example, takes about four times as long with 8 bits than with 16 bits. Why? Because the eight-bit internal architecture requires four eight-bit loads for the first real and four eight-bit loads for the second, and then does $4 \times 4 = 16$ add operations. The 16-bit micro requires two loads per real and $2 \times 2 = 4$ add operations. Does this mean the IBM PC with 16 bits is four times faster than an Apple II? No. The IBM has an 8-bit data path—as opposed to a 16-bit path on other micros—which affects the load operations. A further consideration regarding speed is whether a given micro supports a numerical coprocessor. The latter can give (realistically) a speed enhancement of two to five—significant, but far less than what is sometimes suggested by the manufacturers.

Is this the whole story on speed? No, there's much more. The compiler can be very important, up to a factor of two in speed (not to mention the headaches if the compiler is filled with bugs, as is one well-known FORTRAN compiler for the IBM PC). Some specifics: A recent Douglas Aircraft study found the IBM PC with the 8087 arithmetic co-processor using the Microsoft 3.1 compiler to be about 1/3 the speed of a PDP 11/44 and 1/9 the speed the speed of a VAX 11/780, both of the latter with floating-point hardware. This study used the standard Whetstone benchmark. This is not to endorse either IBM or Microsoft; better PCs and compilers are available. My own "benchmark"—a DO loop with a million floating-point multipliers—gave a ratio of 35-to-1 for the relative speeds of a VAX 11/780 with floating-point hardware and an IBM PC using the 8087 chip. So even a given benchmark is not the whole story.

Are these all the "gotchas"? Not by a long shot. Unfortunately, neither the major retail outlets nor many of the major PC manufacturers have given much, if any, emphasis to supporting scientific users with appropriate software; there's too much money in first-time business and educational users. After investing a considerable sum on my IBM PC, I was pained to discover that, for example, all the marvelous graphics shown off by the salesperson were not available in either of my native languages—FORTRAN or PASCAL. I have since learned 8086-8088 assembly language, and have developed a FORTRAN–PASCAL graphics interface that provides elementary screen operations as well as a capability for 2-dimensional and 3-dimensional plotting and modeling of solids. One may obtain these routines by writing me for further information.

Another issue is addressable memory. Some of the newer micros even come with a virtual memory. I could go on, but many of these "gotchas" come and go with a time constant of about six months. Unfortunately, the stability and accoutrements expected by the old-time mainframe user must be reassessed for the new micro generation.

DONALD ESTERLING
School of Engineering and Applied Science
The George Washington University
2/84 *Washington, D.C. 20052*

•

We read with great interest the article by Per Bak and must congratulate the author for presenting the case of microcomputers for full-fledged research. Working in a dual institute (one of us at the University of Pisa and the second at the Council of National Research, both housed under one roof) and given our computational environment (six 3270-type terminals connected to a computer center equipped with a 370/168 and a 3033N from IBM, a SEL 32/87 at our total disposal, a Texas Instruments DS990/1, a Digital Professional 350, 3 Apples, a VDS EC01 (a CP/M machine), 2 Hewlett-Packard 16-bit 'mini' computers, plus a number of home computers for personal use), we could not resist setting down some of our own findings on the above matter. We think there may be quite a few science institutes where the researcher has to cope with the question of *what* machine to use.

Bak suggests that computational physics can—actually should—be performed on microcomputers, because of their easy availability, low price and high performance.

However, based on our own experience, we tend to disagree on some of the points he raises. Of course, microcomputers today have an almost irresistible appeal, since at nominal cost they offer high-level language, adequate memory and instant availability at any moment, every day. This should not mislead us to believe that every computation fitting the available memory (surprisingly few fail to fit) is worth doing on the microcomputer. The fact is, not every computational task lends itself to the elegant treatment Bak adopted for the Ising model: Most computations are computation intensive, requiring large amounts of floating-point operations: No tricks will substitute.

Thus, for instance, we have integrated the nonlinear differential equation[1,2]

$$M\ddot{x} + M\dot{x}\tau^{-1} + \alpha x - dx^3 = F\cos(\omega t)$$

using a simple but accurate program, following a cue from Davis's book on nonlinear differential equations.[13] We ran this both with a microcomputer as well as a mainframe. A computation that required less than 2 seconds on the IBM 3033N mainframe took more than six hours on the microcomputer. This equation is a classic model used in physics and chemistry to understand the dynamics of molecules interacting with electromagnetic field. Our experience shows how the case for microcomputers, for a simple problem having hardly any storage requirement, need not be so favorable as one would be tempted to conclude.

Therefore, we compared benchmarks for a number of available micro- and personal computers. We adopted an extremely simple BASIC benchmark, coded in one or two lines, requiring many (typically 10 000) executions of a multiplication such as $B = A^*B$. B can be preset to anything, but A must be close to unity to prevent over- or underflow. This can be done in a minute—at a computer show or shop, for instance—and gives a simple, effective way to obtain an index-free, although loop-dependent, multiplication loop. Most current micros employing 8-bit data, 16-bit address CPUs (such as all Commodore, Apple, Sinclair and Radio Shack machines, plus HP-85 and all CP/M-equipped micros) yielded between 160 and 230 loops/sec. Few true 16-bit machines were available and tested, among them the Olivetti M20 and the Digital Professional 350; they gave a two- or three-fold speed increase. On some machines a BASIC compiler was available; it improved performance by 30% as a maximum, confirming that the simple multiplication loop takes longer to be executed than to be interpreted. This is typical of any computation-intensive task, so that the "20 to 100 times speed increase" claimed by compiler vendors only applies to sorting, word processing and so on.

Floating-point hardware is optional on some costly (above $5000) personal computers, such as the Digital Professional 350 and the IBM PC. We tested the Digital machine with this hardware: BASIC, either interpreted or compiled, so far does not support the floating-point hardware, requiring a switch to FORTRAN. Under control of the much less convenient FORTRAN translator, a double-precision multiplication took about 200 microsec (on all machines using 32-bit single precision numbers, double precision is a must for meaningful computation-intensive tasks; this includes all CP/M machines).

Coming now to large machines usually installed at major computing centers, we can quote results obtained on the IBM 3033N computer available locally at the CNR computing center of Pisa. The BASIC benchmark given above yielded about 300 000 loops/sec. Of course, even on a large computer, a BASIC compiler is not expected to be highly optimized; in fact, the same benchmark, recoded in FORTRAN, yielded about 1 500 000 loops/sec using an optimizing compiler and double precision. Thus, a large mainframe, which is nevertheless somewhat slower than the CDC 7600 quoted by Bak, outperforms typical personal computers by a factor of 2000 to 10 000 on computation-intensive tasks. This is sufficiently larger than the factor of 200 arrived at by Bak to offset almost all the cost advantage he claims for the microcomputer. The performance ratio for the BASIC integration program mentioned before is particularly favorable to the mainframe, reflecting the fact that more complex programs lend themselves to more effective optimization on the part of the compiler; furthermore, because the micro operated with CP/M, double precision had to be used, further degrading the micro's perfomance much more than that of the mainframe.

Computer type	Model	Time (secs)	Comments
16-bit 'mini'	HP 21MX-E	50	Floating-point firmware
16-bit 'mini'	DEC PDP11/60	8	RSX11-M, FIV+, floating-point hardware
32-bit 'midi'	Perkin-Elmer 3250	2	Without global optimization
32-bit 'midi'	Data General MV10000	1.5	
32-bit 'midi'	SEL 32/87	1	
32-bit 'maxi'	IBM 3033N	0.3	VM/FORTRAN HX, opt = 2

For a floating-point-intensive benchmark (diagonalization of a 26 × 26 matrix by Jacobi algorithm in double precision format. About 500 000 operations required on double-indexed data arrays)

The table gives results of a more complex benchmark more typical of computational physics. We see that almost every true minicomputer on the market using FORTRAN has a decided performance advantage over any personal computer (note that HP-E series CPUs are now superseded by the considerably move efficient HP1000A series; PDP 11/60 has been effectively replaced by the newly introduced 11/73 processor). Also, truly outstanding performance has been arrived at by the most sophisticated 32-bit "midi" computers.

Will this level of performance become available in the personal computer of tomorrow? We obviously join Bak in hoping so. However, we are not convinced it will arrive soon, not because it is technically impossible, but precisely because it could have been achieved even by yesterday's technology. The fact is, technical choices lag beyond commercial ones; and computational speed, while important for the scientific community, is not required by the large market of home and office users—for them, 200 operations per second is very high speed.

Finally, we come to the essential question: Which computations to perform on the micro, which at the computing center? Of course, the elapsed time between the moment at which we feel the need for some number and the moment at which the number is under our eyes is for us "computing time." Due attention should therefore be given to factors (such as connection and turnaround time, as well as closedown days or hours) on which the personal machine has decided advantages. Although individual preferences will certainly play a considerable role, we think that, at least on working days, 100 000 to 500 000 floating-point operations (meaning typically 10 minutes to 1 hour run time) is the range above which we would seriously consider having recourse to the large computing center. We should also consider the extreme ease with which we can, in that case, have the computation repeated, without additional delay, using different starting parameters, smaller integration steps and the like. Let us hope for the performance jump that will add an order of magnitude to this limit.

References

1. N. Bloembergen, Opt. Commun. **15**, 416 (1975).
2. B. A. Huberman and J. P. Crutchfield, Phys. Rev. Lett. **43**, 1743 (1979).
3. H. T. Davis, "Introduction to Nonlinear Differential and Integral Equations," Dover, 1962.

ROBERTO AMBROSETTI
Council of National Research
Pisa, Italy
NASEEM K. RAHMAN
University of Pisa
3/84 *Pisa, Italy*

•

Per Bak claims to have shown in his article that "state-of-the-art" simulations on the three-dimensional Ising model can be performed efficiently on a home computer. But it is not that easy to beat the state of the art.

A large computer such as the CDC 7600 (mentioned by Bak for comparison) performs more than a million Monte Carlo steps per second on a three-dimensional Ising model, if programmed in efficient FORTRAN,[1] and may be even faster with assembler programming.[2] Thus the speed gain over his home computer's 5000 steps per second is much more than a factor 100—not much less, as the author thinks. (The Cyber 205 "supercomputer" of CDC even flips about ten million Ising spins each second,[3] which may be a more valid example for a comparison with the latest home computer model.) Also, the quality of the calculation may be compromised, contrary to Bak's hopes, if $32\times32\times32$ is the maximum system size for simulations of critical phenomena. If all reviewers for NSF were to share the feeling that big computers are too costly for this type of simple model, and that most of these computations should be done on a Commodore, physics in the US would not be helped. (A minor point: Bak's estimate 4.50 ± 0.02 for the critical temperature should be compared with[4] the 200-fold more accurate 4.5115 ± 0.0001, not with old results.)

Fortunately, I can live a reasonably happy life as a "parasite," to mention Bak's catchy picture, feeding off the resources of a CDC Cyber "dinosaur." A small machine makes me independent of bureaucracies but more dependent on my brain; some people around here doubt that this would be an advantage. However, since there might be more to physics than just the simulation of huge three-dimensional Ising models, a well-balanced diet of large, medium and small computers is what NSF and others should feed us with.

References

1. C. Kalle, V. Winkelmann, J. Statist. Phys. **28**, 639 (1982).
2. B. A. Freedman, G. A. Baker, J. Phys. **A 15**, L 715 (1982).
3. S. Wansleben, J. G. Zabolitzky, preprint for J. Statist. Phys.; C. Kalle, Master's thesis, Cologne University, to be written.
4. J. Adler, J. Phys. **A 16**, 3585 (1983).

DIETRICH STAUFFER
Cologne University
1/84 *Cologne, West Germany*

THE AUTHOR COMMENTS: I am delighted about the extensive debate on computations in physics provoked by my article. Most of the letters express interesting and valid points of view, and I shall restrict myself to correcting a few misunderstandings.

The reason for hand-picking a calculation involving machine language and integer operations only was to permit a comparison of the potential capabilities of the computers on an equal footing, not that I particularly enjoy writing machine programs. It makes no sense to compare floating-point performance of a microcomputer using a BASIC interpreter and an integer-number processor with that of a mainframe computer using a compiled language and floating-point processor, as was done by David Anderson, Richard Branham, and Andrew DePristo and Stephen Elbert, when estimating a performance ratio of the order of 2000 to 10 000. Naseem Rahman and Roberto Ambrosetti arrive at roughly the same estimate when comparing the Digital Professional 350 with the IBM 3033N.

Today—as anticipated in the article—floating-point processors and compilers (see Jassby's letter) are available for several microcomputers, and we can compare the floating-point performance directly. As correctly pointed out by Donald Esterling, Anderson, and DePristo and Elbert, an IBM personal computer with an 8087 floating-point processor has about $\frac{1}{10}$–$\frac{1}{30}$ of the speed of a VAX, and about $\frac{1}{100}$–$\frac{1}{300}$ of the speed of a CDC 7600. Certainly, less expensive and more powerful microcomputers than the IBM PC (such as the Apple Mackintosh with the 68000 processor) are available today, making the gap even narrower. I agree with Jon Forrest that we are near the point where desktop computers will perform like a VAX.

Forrest also made the important observation that the main reason for high computing cost at computer centers is the large overhead; Anderson forgets this when comparing only prices of installed computer power.

DePristo and Elbert estimate (by comparing different calculations on different computers) that my Monte Carlo calculation would take 0.1–1 minute on a CDC 7600 at a cost of \$5 to 50¢. Actually, it would take two hours,[1] with 0.5×10^6 flips per second. Dietrich Stauffer argues that supercomputers may perform 10^7 steps per second. Again, if this performance is compared with that of the most sophisticated home computer, not more than a modest factor of 100 in speed is gained. Moreover, we are not restricted to a $32\times32\times32$ lattice in a modern microcomputer with about 100 K memory with room for more than a million spins.

In summary, I certainly disagree with Elbert and DePristo's statement that "it is unlikely that micros can be of much use." Already now, microcomputers are mushrooming at physics departments everywhere.

Reference

1. L. Jacobs and C. Rebbi, J. Comp. Phys. **41**, 203 (1981).

PER BAK
Brookhaven National Laboratory
4/84 *Upton, New York*

Physics News in 1983

I, for one, applaud the publication of "Physics News in 1983" as a special section in January. As adviser of our SPS chapter, I have received a single copy of Physics News for the last several years and have been at a loss about a method to make it uniformly available to our department. When I have placed it in our student conference room, it disappears quickly, probably to a student or faculty member who "borrows" it "overnight." This way, all the faculty and SPS members have their own copies. With the number of potential borrowers significantly reduced, it should now be safer to leave our copy available.

F. W. PROSSER
The University of Kansas
1/84 *Lawrence, Kansas*

•

I refer to John Croat's article on the new class of high-performance iron–rare-earth–boron permanent magnets (January, page S-20, "Physics News in 1983").

I agree with Croat that this is a major discovery in the world of permanent magnets and it will probably revolutionize the electric motor business. These magnets do not contain any cobalt, a metal considered to be critical and stragetic; instead they consist mostly of iron with some additions of light rare-earths and boron, which are abundant and less expensive materials.

However, for the completeness of the story I would like to mention the appropriate references that were left out in Croat's article. The very first data on the magnetic hardening of iron–rare-earth–metalloid alloys were announced at the Intermag meeting in Philadelphia on 21 April 1983, where coercivities higher than 23 kOe and energy products of around 10 MGOe were reported in Fe(Co)PrB(Si) samples.[1] The first official paper on these materials was published[2] on 15 October 1983, and it showed an energy product of 13 MGOe in a melt-spun sample of $Fe_{76}Pr_{16}(BSi)_8$. The same paper included transmission electron microscope data and x-ray diffraction studies that showed the presence of a new anisotropic phase with the composition $Fe_{21}R_3B$, which has a tetragonal structure[3] similar to that reported by Croat.

Our data, as well as those of Croat, were obtained on melt-spun samples. The only difference was that our samples had to be annealed at around 700 °C to obtain magnetic hardening, while those of Croat showed magnetic hardening in the as-quenched state after being slowly quenched.

Very recently, the Sumitomo Co. of Japan announced the production of the first commercial magnet made out of these materials with a record of "energy product" of 36 MGOe. Sagawa and collaborators[4] have attributed the hard magnetic properties of these magnets to the presence of a tetragonal phase that appears to be the same phase