E&TR

UCRL--52000-84-3

DE84 008877

Energy and Technology Review

Lawrence Livermore National Laboratory

March 1984

Brief	iii
The Centralia Partial-Seam CRIP Test	
Visualizing Chromosomes in Human Sperm	1
New techniques for fusing human sperm with hamster eggs enable us to directly observe structural and numerical abnormalities in the chromosomes of human germ cells.	
High-Speed Cell Sorting: A Tenfold Increase in Sorting Rate	13
Our improved high-speed cell sorter can identify and select samples for subsequent biochemical analysis at rates above 20 000 particles per second.	
Preparing for the Super Supercomputers	22
Our modeling studies indicate that to achieve computational speeds a thousand times faster than those of present supercomputers, only a tenfold increase in speed of the gain will come from improved chips and circuits; the rest will come from greater parallelism.	
Abstracts	31

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

UCRL-52000-84-3
Distribution Category UC-2
March 1984

About the Cover

To achieve computational speeds a thousand times faster than those of present supercomputers, drastic changes in computer architecture will be necessary. The diagram illustrates schematically one possible approach, with a large number of processors (outer layer), each with its own communication interface (second layer), all with equal access to a gigantic common memory (central sphere). The article on p. 22 shows that significant advances are required in computational techniques as well as in computer technology.

Scientific Editor

Frank A. Morrison, Jr.

General Editors

Robert A. Condouris
Richard B. Crawford
Lauren de Vore
Mary Ann Esser
Robert D. Kirvel
Palmer T. Van Dyke

Graphics Designer

John Schuster

Compositor

Louisa Cardoza

Production Staff

Linda Davis
Kimberly Smith
Linda Geniesse

Art Credits

Art Francheschi (pages iii,
24, and cover)
Steve Wright (pages 15, 17,
and 28)

Brief

The Centralia Partial-Seam CRIP Test

The partial-seam test of the Controlled Retraction Injection Point (CRIP) system is the latest in a series of underground coal gasification (UCG) tests we have been conducting at the WIDCO coal mine near Centralia, Washington. (See the November 1982 issue of *Energy and Technology Review*; UCRL-52000-82-11, p. 1, for a discussion of previous UGC tests at the WIDCO mine.) Underground coal gasification offers a way to obtain energy from the vast reserves of coal that are inaccessible to conventional mining methods. The basic idea of UGC is to burn the coal in place underground, supplying oxygen from the surface through an injection well and drawing off through a production well the large volumes of combustible gases evolved in the gasification process.

Our CRIP system is a refinement that extends the useful life of the injection well. Without the CRIP system, roof collapse soon cools the reaction, reducing the heating value of product gas. With the CRIP system, we can ignite a series of burn cavities without boring a new injection well for each one.

We began the partial-seam CRIP test on October 16, 1983. The coal seam was ignited with our silane igniter, and oxygen-enriched air was supplied throughout the day; the product gas thus generated had a heating value of 7.5 to 8.3 MJ/m³ (200 to 220 BTU/scf). We began steam/oxygen injection on October 17 and experienced smooth operation, except for a two-day shutdown to service equipment partially plugged by coal dust, obtaining very good product gas with a heating value of about 9.3 MJ/m³ (250 BTU/scf). We injected air for brief periods to keep the fire burning during power outages or steam-boiler repairs.

We used a vertical production well until day 12 of the experiment (October 28), when we switched to a slant

production well (Fig. 1). The entire changeover went smoothly, and we observed an improvement in product-gas quality. On October 30, we executed the controlled retraction maneuver and established a new cavity, with the expected improvement in product-gas quality. We continued steam/oxygen gasification using the new cavity and the slant production well until we shut down the experiment on November 14.

Figure 2 shows the product-gas heating value as a function of time for the burn period, and Fig. 3 shows some of the product-gas chemistry. Interaction of the burn with the roof can be seen during days 6-12 (October 22-28) for the first cavity and during days 22-29

Fig. 1

Layout of instrumentation wells (I-1 to I-22), production wells (P-1, PRD-2, and slant), and the injection well for our partial-seam CRIP test at the WIDCO mine near Centralia, Washington. The burn was started at point A; the gases were drawn off through vertical production well PRD-2. After a few days, we switched over to the slant production well. When the quality of the production gases began to decline because of roof collapse, we pulled back the injection to point B and started a new burn, restoring gas quality.

Fig. 2

Product gas quality (heat of combustion) as a function of time for the entire test, including power outages and shutdowns for mechanical maintenance (plugged production line).

Fig. 3

Variation in the proportions of the three principal combustible components in the product gas. Gas composition is strongly dependent on reaction temperature. The heating value of the product gas is strongly influenced by the amount of hydrocarbons (e.g., methane) present.

(November 7-14) for the second. Other significant changes in operation, such as the use of the slant production well and the CRIP maneuver, are indicated in the figure.

We gasified about 500 tonnes of coal with the first cavity and 1000 tonnes with the second, for a total of 1500 tonnes. Good thermal responses were obtained in several instrumentation wells. Our major conclusions from this highly successful underground coal gasification test are:

- The Centralia site is favorable for UCG.

- The large partial-seam test produced gas similar in quality to that obtained in the much smaller-scale tests we conducted two years ago at the same site, proving that the small-scale experiments have predictive capabilities for UCG.

- Use of a slant production well is advantageous for UCG.

- The CRIP method was successfully tested, improving UCG performance as expected.

For further information contact Douglas R. Stephens (415) 422-7330. □

Visualizing Chromosomes in Human Sperm

New and refined techniques for fusing human sperm with hamster eggs are being used to create hybrid cells in which the sperm chromosomes are fully visible and can be directly examined under a light microscope for precise determination of structural and numerical anomalies.

Sperm and egg cells are the link between one generation and the next. When one sperm and one egg unite during conception, the new embryo receives half of its chromosomes from each parent so that a full complement of chromosomes is reconstituted in the offspring. Even before birth, chromosomal constitution can profoundly influence reproductive outcome. Whether pregnancy is spontaneously lost or is achieved and maintained and whether a healthy or defective offspring is born are determined, in large part, by the embryo's complement of chromosomes.

Chromosomes contain the genetic material, deoxyribonucleic acid (DNA), that holds the blueprint for the individual. Some physical traits, such as eye color, are a direct expression of the individual's genetic constitution. Other traits, such as height, are influenced by environmental factors as well. Even behavioral predispositions, such as intelligence and musical talent, are determined by the interaction of both genetic and environmental factors. Thus,

the chromosomal constitution of both sperm and egg is of the greatest possible significance, not only during pregnancy but also in the expression of adult characteristics.

Geneticists have long suspected that human sperm and egg cells may be characterized by a relatively high frequency of chromosomal abnormalities. Such deviations include abnormal numbers of chromosomes and aberrations in chromosome structure. Abnormalities in chromosome number are apparent when one or more whole chromosomes are absent or extra chromosomes are present. Aberrations in chromosome structure are evident when pieces of a chromosome break off and are either lost completely or translocated to a new position on another chromosome.

Until recently, research on the frequency of chromosomal errors in human sperm and eggs has been confined largely to two indirect sources of evidence. One source is the chromosomal constitution of both normal and abnormal offspring. The other is the

For further information contact
Brigitte Brandriff (415) 423-0758.

chromosomal constitution of tissues obtained from the spontaneously lost products of conception, especially during the early stages of gestation. From such studies, researchers have estimated that from 7% (Ref. 1) to 50% (Ref. 2) of all human conceptions involve a chromosomal abnormality. In a high proportion of cases involving an abnormality, conceptions do not reach full term, partly because of maternal selection against the affected fetus.³ Among those pregnancies with some chance of surviving are those involving the condition known as Down's syndrome. In Down's syndrome, the offspring has an extra number-21 chromosome and exhibits mental retardation. Modern staining techniques reveal that about 20% of the offspring with Down's syndrome inherit the extra chromosome from the father.⁴

Studies of live offspring and of spontaneously aborted fetuses have been among the very few methods by which researchers could infer chromosomal constitution of human sperm and eggs. Studies of human eggs will always remain limited because of ethical and practical considerations. Human sperm, on the other hand, are readily available, but the chromosomes are so tightly packaged within the mature head that they cannot be distinguished even by electron microscopy. Thus, a new method is urgently required, one by which the material forming the chromosomes within human sperm can be made visible.

Six years ago, researchers at the University of Hawaii⁵ pioneered a new system for the direct study of chromosomes in human sperm. These workers discovered that human sperm can be induced to fuse with hamster eggs in a tissue culture medium. After several hours of culture, the eggs have the capacity to unfold the sperm DNA and then to repack it into visible chromosomes. Replication of these techniques in other laboratories proved to be difficult. However, researchers at the University of Calgary⁶ and at LLNL⁷ have recently overcome the technical difficulties. Thus, we are now able to observe directly the chromosomes within

human sperm and to establish the frequencies of structural and numerical errors.

At LLNL, our research program on the identification of chromosomal abnormalities in human sperm was supported initially by the Environmental Protection Agency and is now funded by the Department of Energy. Investigations of sperm cells are an integral part of the LLNL Biomedical Sciences Division's study of human genetics and genetic toxicology just as chromosomal errors in sperm may lead to untoward effects in offspring, similar errors in other cells of the body, called somatic cells, may produce untoward effects such as cancer. Thus, our studies of sperm cells complement those of somatic cells, which are being carried out by other investigators in the Biomedical Sciences Division.

Using these new techniques, we began by gathering baseline data on a small group of normal men (i.e., men with no known reproductive problems). To date, we have analyzed more than 1000 human sperm complements from five healthy volunteers to determine both the kinds and frequencies of abnormalities and to establish a basis for comparison in future work. In the future, we will extend our investigations to men who have been subjected to chemotherapy, radiotherapy, or environmental hazards. We also plan to analyze sperm from men with clinical reproductive problems to further our understanding of the basic biology of human reproduction.

Hamster Eggs and Human Sperm Chromatin

Our technique takes advantage of the fact that eggs (ova) are "programmed" to undergo a series of physiological changes after fusion with sperm (see the box on p. 3 for a description of the egg's developmental program and the various structures associated with the egg). In our work with hamster eggs, we mimic the signal that initiates development by introducing a foreign sperm. For our research, a human sperm provides that signal. However, we must first prepare

Sperm and Eggs: Interactions during Fertilization

Fertilization involves the union of specialized cells, the gametes. The common names for gametes are ova or eggs, produced by the female, and sperm or spermatozoa, produced by the male. Both ova and sperm are also called germ cells to differentiate them from all other cells of the body, called the somatic cells.

Germ cells may be distinguished from somatic cells on the basis of chromosome number. Each germ cell contains one-half the number of chromosomes found in a somatic cell. The number of chromosomes within a germ cell is the haploid number, which in humans is 23. Upon fertilization, one egg and one sperm unite. In this way, the embryo receives half of its chromosomes from the mother and half from the father. The full complement of 46 chromosomes (the diploid number) is thereby restored in the offspring.

Each human sperm contains one of the two types of sex chromosome, either the X or the Y chromosome; the human egg contains only one type of sex chromosome, the X. If the egg is fertilized by an X-bearing sperm, the offspring will be female; if it is fertilized by a Y-bearing sperm, the offspring will be male.

Fertilization is a multistep process and may be divided into three sequential stages through which the gametes pass. The three stages are approach, fusion, and response. During the approach stage, sperm progress from the point of deposition in the female reproductive tract to the vicinity of the egg. During this progress, sperm undergo capacitation, a physiological change that enables them to fuse with ova. Fusion begins when the first sperm contacts an egg. The sperm binds to the extracellular coat, the zona pellucida, that surrounds the egg membrane. After a sperm passes through the zona

pellucida (a, figure below), hairlike projections of the egg cell membrane, the microvilli, engulf the sperm head. The membranes of both gametes then fuse, and the sperm is taken into the interior of the egg.

The egg's responses to fusion with a sperm are manifold and may be collectively referred to as the egg's early developmental program. The egg gives off a small amount of cytoplasm, containing a set of chromosomes, called the second polar body. The extrusion of this second polar body, together with the first polar body that was extruded during ovulation, ensures that the haploid number of chromosomes remains within the egg. In response to fusion, the egg also unfolds the sperm head by partly removing the sperm-specific proteins that are associated with its DNA. The sperm chromosomes assume an extended form and cannot be individually resolved; they are, in fact, barely visible under a microscope.

The egg cytoplasm then synthesizes two nuclear envelopes: one for the sperm chromosomes and a second for the egg chromosomes, which are also in an extended form. The two nuclei, called pronuclei (b, figure below) remain separate until the first division of the newly formed embryo. In preparation for the first division, the chromosomes duplicate their DNA and contract so that they become visible via light microscopy. Male and female chromosomes mingle and collect on the metaphase plate (c, figure below), a structure that specifically functions to divide the chromosomes so that equal numbers are distributed to the two daughter cells (d, figure below). With the mingling of the male and female chromosomes, fertilization is completed, and the newly formed embryo begins its developmental journey.

(a) Sperm penetrates zona pellucida

(b) Two pronuclei are formed

(c) Chromosomes collect on metaphase plate

(d) Chromosomes are distributed to daughter cells

both the hamster eggs and the human sperm for fusion.

From golden hamsters, *Mesocricetus auratus*, we obtain 40 to 50 ova per female by stimulating the ovaries with hormone injections. We place the eggs in a special medium and apply the enzyme trypsin to remove the egg's specialized extracellular layer. This layer, called the zona pellucida, normally acts as a barrier against foreign sperm. Eggs of the golden hamster are unusual, however, in that they can fuse with many different species of sperm, including those from humans. Other rodent eggs, such as those from mice and rats, will not fuse with human sperm even after the zona pellucida is removed. To instigate

capacitation (see the box on p. 3) and the ability to fuse with ova, we incubate the human sperm for several hours at 37 °C in a medium containing salts and metabolites. Once the sperm and ova are suitably prepared, we pipette the ova into suspensions of sperm and incubate the gametes together for 1 to 3 hours at 37 °C.

The fusion of hamster egg and human sperm sets in motion the early developmental program of the ovum. The development continues just long enough to render the chromosomes visible. The fertilized ovum may even divide once; thereafter, development is arrested and the hybrid dies. Indeed, rodent eggs that are fertilized with sperm

Fig. 1

Hamster egg after fusion with four human sperm. The DNA of the fused sperm is unfolded, and the heads have expanded to many times their original volume. Tails remain attached to the expanded heads within the egg. Compare the much smaller size of sperm attached to the outside of the egg cell membrane. This is a phase-contrast micrograph image of living cells.

from the same species and studied *in vitro* frequently die after the first division, a phenomenon that is referred to as the two-cell block.

Upon fusion, the egg extrudes the second polar body, and decondensing sperm heads begin to appear within the egg cytoplasm after 1 to 2 hours of co-incubation. The hamster egg cytoplasm is able to remove the human sperm proteins (protamines) even though the chemical composition of these protamines differs from that of hamster sperm. (Protamines occur in the sperm of many animal groups and in some plant groups as well. Various theories on the function of protamines have been proposed, including the packaging of DNA inside the sperm head and the protection of sperm DNA during transmission to the female.)

Figure 1 shows four human sperm that have entered one hamster egg. This condition is known as polyspermy and is always lethal in mammalian eggs that are fertilized *in vivo*. In our case, the hamster egg's normal block to polyspermy is eliminated by dissolving the zona pellucida. During actual experiments, we remove the eggs from the sperm suspension to prevent fusion with more than 1 or 2 sperm because the presence of additional sperm reduces the likelihood of further development.

Figure 2 shows the fully developed pronuclei that appear at about 3 hours after fusion. One pronucleus contains the hamster's complement of chromosomes; the other contains the human complement. Both complements are in the extended conformation, called interphase, and individual chromosomes remain indistinguishable.

Figure 3 shows the egg at about 10 hours after fusion. At this time, the DNA with its associated proteins, collectively referred to as chromatin, begins to condense into individual fibers. At 12 to 14 hours, the chromosomes are fully condensed in preparation for the first cell division. At this point, we chemically arrest the eggs so that the male (human) and female (hamster) chromosomes cannot mingle on the metaphase plate of the first division. We then spread the eggs on microscope slides and stain them so that we can identify each chromosome (see Fig. 4).

Identification of Individual Chromosomes

In most chromosome studies, material is derived from the metaphase stage of cell division because it is then that the chromosomes are ideally condensed and clearly visible under a light microscope. During metaphase, the sister chromatids are not yet separated but are joined at a

Fig. 2

Photomicrograph of fixed and stained hamster and human pronuclei shortly after fertilization. The sperm head has now fully expanded and is indistinguishable from the hamster pronucleus.

(a) Hamster chromosomes

(b) Human chromosomes

Fig. 3

Photomicrographs of fixed and stained hamster (a) and human (b) chromosomes in the early stages of metaphase. The hamster complement of 22 chromosomes is in a more advanced stage of condensation than the human complement of 23 chromosomes.

(a) Hamster chromosomes

(b) Human chromosomes

Fig. 4

Photomicrographs of fixed and stained hamster (a) and human (b) chromosomes that are fully condensed. The degree of condensation is equal in the two complements.

(a) Metacentric

(b) Submetacentric

(c) Acrocentric

(d) Telocentric

Fig. 5

Metaphase chromosomes are classified according to their size and the position of the centromere. A chromosome's shape is (a) metacentric when the centromere is in the middle, (b) submetacentric when the centromere is nearer one end, (c) acrocentric when the centromere is very near one end, and (d) telocentric when the centromere is at the end of the chromosome.

structure or constriction called the centromere. Sister chromatids are identical chromosomes that are the product of DNA replication. The size of the sister chromatids and the position of the centromere provide the century-old basis upon which chromosomes are identified. Figure 5 illustrates the general shapes of metaphase chromosomes. A chromosome's shape is described as metacentric when the centromere is in the middle, submetacentric when the centromere is nearer one end, and acrocentric when it is very near one end. Telocentric chromosomes, in which the centromere is at the very end of the chromosome, do not occur in humans but are found in other mammals.

Since 1969, new techniques using fluorescent dyes have made it possible for researchers to identify individual chromosomes. The dyes are derivatives of quinacrine and produce characteristic bands of fluorescence along the chromosome. Each of the 22 pairs of chromosomes and the two sex chromosomes of the human, when stained with quinacrine, shows a specific pattern of alternating light and dark bands when viewed through a light microscope under ultraviolet illumination. Although the quinacrine-banding technique was developed for and extensively applied to human blood cells, it also can be used with other kinds of cells. Fortunately, the technique is suitable for analyzing the chromosomes derived from human sperm after they have been processed by hamster eggs.

Human chromosomes may be classified into 22 pairs. The two sex chromosomes that determine gender bring the total to 46 chromosomes per cell. The cells of golden hamsters have 44 chromosomes, but the size distributions are different from the human complement. Figure 6a shows

Fig. 6

(a) Normal human sperm chromosome spread stained with quinacrine to yield characteristic dark and light bands. (b) Karyotype in which the same chromosomes are arranged into their respective groups and numbered systematically.

...we focused our research on obtaining larger numbers of karyotypes (from 160 to 315) from each donor so that detailed comparisons among donors could be made.

a human array, called a metaphase chromosome spread, that we derived from a hamster egg. Figure 6b shows the same chromosomes arranged into their respective groups, called a karyotype. The karyotype represents a chromosomal complement that is characteristic of a species. Although the patterns of alternating dark and light bands can be clearly seen in the photograph, sister chromatids cannot be seen in Fig. 6 because the space separating them is obscured by our fixation and staining procedures.

Frequency of Chromosomal Abnormalities in Human Sperm

To establish a general picture of the kinds and numbers of abnormalities that might be found in human sperm, we examined sperm chromosomes from five normal, healthy men. The men's ages ranged from 34 to 39, and three were of proven fertility, each having fathered at least one child. In contrast to workers at the University of Calgary, we focused our research on obtaining larger numbers of karyotypes (from 160 to 315) from each donor so that detailed comparisons among individuals could be made. To date, we have completed work on 1133 karyotypes.

We classified structural aberrations according to the international system for human cytogenetic nomenclature (ISCN).⁸ The kinds of structural aberration we observed were chromosome breaks, chromatid breaks and exchanges, acentric fragments, and deletions. In a chromosome break, both sister chromatids are severed at the same site; in a chromatid break, only one of the sister chromatids is broken. We recorded chromosome breaks when two pieces, one with and the other without a centromere, could be ascribed to a specific chromosome. In addition, we recorded an acentric fragment when a piece could not be assigned to a specific chromosome of origin.

Figure 7 shows an abnormal chromosome spread from one of our donors. This spread contains a chromatid exchange between chromosomes 1 and 3. On the basis of work with other kinds of

cells, we know that such a cruciform configuration develops during the preceding DNA replication. Therefore, our illustrated chromatid exchange must have developed while the sperm chromosomes were in the hamster egg.

In addition to the chromatid exchange, the spread in Fig. 7 also reveals a break in chromosome 4. The centric and distal pieces are designated 4A and 4B, respectively. In our study, we observed that specific break points sometimes were repeated two or three times in a given donor. Of the repeat breaks in four of our five donors, most were recovered from separate ejaculates collected from 1 to 9 months apart. Repeat breaks occurred more frequently than would be expected by chance alone. We do not yet understand the significance of this intriguing phenomenon and expect to gather more information through further research.

Figure 8a shows an abnormal chromosome spread that contains an extra chromosome. The error in chromosome number becomes obvious when the chromosomes are arranged into the karyotype in Fig. 8b. The presence of the extra chromosome, number 21, is characteristic of Down's syndrome.

In Table 1, we summarize the incidence of errors in both chromosome structure and number among our five donors. The percentage of sperm cells with structural aberrations varied widely. In fact, structural aberrations varied as much as eightfold from one subject to another, with an average rate of 6.5% (the range in Table 1 is from 1.3% to 30.4%). The overall frequency of aberrations in chromosome structure was much higher than the incidence of aneuploidy, the deviation from normal chromosome number. Aneuploidy denotes one of two conditions—either the presence of one or more extra chromosomes (hyperhaploidy) or the absence of one or more chromosomes (hypohaploidy). The overall rate of both hyperhaploidy and hypohaploidy in our donors was low, about 1%. The range in Table 1 is from 0% to 2.6%, and no significant differences were observed among our five donors.

Fig. 7

Abnormal human spread containing a chromatid exchange between chromosomes 1 and 3 and a break in chromosome 4. (a) The spread is stained with quinacrine to identify individual chromosomes. (b) The spread is stained with Giemsa (a nonfluorescent stain) to confirm abnormalities. Note the distal piece (4B), the fragment of chromosome 4 containing the centromere (4A), and the cruciform configuration of the chromatid exchange between chromosomes 1 and 3.

In Table 2, we compare the frequencies of structural aberrations in sperm and lymphocyte chromosomes from the same individuals. The frequencies of such aberrations in sperm, surprisingly, are higher than those recorded for blood cells, for which a substantial body of data is available. Indeed, a normally low incidence of aberrations was characteristic of lymphocyte chromosomes from these donors, the only individuals studied so far in this manner. In donors A, C, and D, the frequencies of structural aberrations were significantly higher in sperm than in lymphocytes.

Of our 1133 karyotypes, 52.1% contained an X chromosome and 47.9% contained a Y chromosome. The ratio

of our observed frequencies does not significantly differ from the expected 50:50 ratio.

Discussion

Our first finding of interest is the elevated average frequency of structural aberrations (6.5%) relative to aneuploidy (2.0%) in sperm chromosomes. Such data may be compared with those from another study of human sperm karyotypes. In a survey of 1000 karyotypes from 33 men,⁶ the reverse relationship was seen, with a 5.2% rate of aneuploidy and a 3.3% rate of structural aberrations. Because the high incidence of structural aberrations in four of our five donors was unexpected, more data from the human sperm and hamster

Fig. 8

(a) Abnormal human spread containing an extra chromosome. The arrows point to the two number-21 chromosomes. (b) The karyotype of the same chromosomes shows that the sperm carried an extra chromosome 21, characteristic of Down's syndrome.

egg system are needed before we can provide a definitive explanation for the disparate results.

Among the chromosomal errors identified in tissue examined after the spontaneous loss of the products of conception, 96.2% were numerical and only 3.8% were structural.² It is reasonable to ask why a much lower proportion of structural abnormalities was identified in aborted material compared to the relatively high proportion in our sperm samples. To answer this question, we must remember that the reported percentages are derived from recognized, spontaneously aborted pregnancies that occur at an overall rate of about 15%. Estimates of the number of early, spontaneously aborted pregnancies that include those remaining unrecognized are, however, far higher and range from 40 to 50% in one study⁹ to nearly 60% in another.¹⁰ Early embryos containing the kinds of breaks and exchanges observed in our donors probably would not survive beyond the very earliest stages of cleavage. Because such embryos would not be included in abortion samples, we would expect to see far lower rates of structural defects in those samples than in the sperm we studied. The higher rates of sperm with structural aberrations we have verified may lend support to the hypothesis that many spontaneous abortions remain unrecognized.

In three of our donors, the proportions of cells with structural aberrations were significantly higher in sperm than in lymphocytes. Some researchers have extrapolated data on human fibroblasts and lymphocytes to germ cells.¹¹ However, conclusions from studies on animals concerning the success of predictions based on somatic cells have been contradictory.^{11,12} Our comparisons of actual chromosomal errors in lymphocytes and sperm from the same individuals are the first data of their kind to be reported in men and suggest that extrapolations from somatic cells to germ cells should be used with reservations and caution.

Another issue raised by our research is whether some selection process determines which sperm will interact

with the hamster eggs and ultimately appear in our chromosome spreads. Experiments on mouse gametes show that sperm with either numerical¹³ or structural¹⁴ errors function normally during fertilization. However, a variety of factors may operate as selection mechanisms in our interspecies system. Only further research will determine whether or not our human sperm karyotypes do, in fact, represent an unbiased sample.

Future Research

We have shown that the frequencies of chromosomal aberrations in sperm from normal men can vary by as much as a factor of eight. We have also shown that the frequency of structural aberrations in sperm chromosomes can

Table 1 Abnormalities in sperm chromosomes.⁷ Values in parentheses are percentages of the respective totals.

Donor	Total karyotypes studied ^a	Cells with structural aberrations ^b	Cells with more than 23 chromosomes	Cells with fewer than 23 chromosomes
A	315	15 (4.8%)	4 (1.3%)	2 (0.6%)
B	160	2 (1.3%)	0 (0.0%)	2 (1.3%)
C	222	23 (10.4%)	0 (0.0%)	3 (1.4%)
D	212	19 (9.0%)	1 (0.5%)	2 (0.9%)
E	224	15 (6.7%)	6 (2.6%)	2 (0.9%)
Total	1133	74 (6.5%)	11 (1.0%)	11 (1.0%)

^aIncludes all spreads in which chromosomes were identified individually or by group (no more than one chromosome was identified tentatively).

^bIncludes chromosome breaks, chromatid exchanges and breaks, acentric fragments, and deletions.

Table 2 Comparison of structural aberrations in sperm and lymphocyte chromosomes.⁷

Donor	Number of cells analyzed		Structural aberrations, % ^a	
	Sperm	Lymphocytes	Sperm	Lymphocytes
A	315	310	4.8	0.65
B	160	200	1.3	2.0
C	222	200	10.4	1.0
D	212	300	9.0	1.0

^aThe *P* values for frequencies of aberrations in sperm vs lymphocyte chromosomes were 0.0008, 0.29, 0.00002, and 0.000006 for donors, A, B, C, and D, respectively.

be much higher than that in blood-cell chromosomes from the same individual. To further our understanding of the baseline frequencies of structural and numerical chromosomal errors, we are planning to extend our analysis to sperm from several additional men in good health. We also are adapting methods that will allow us to transport sperm samples from other facilities so that we can more easily gain access to donors of interest from the greater San Francisco Bay area. Eventually, we will concentrate our efforts on men in whom spermatogenesis is perturbed. These donors will include men with structural aberrations that have been ascertained from lymphocyte studies as part of a genetic counseling program, older men, heavy smokers, radiotherapy and chemotherapy patients, and men who have been exposed by either occupation or accident to various environmental hazards.

Key Words: aneuploidy; chromatin; chromosome—abnormality; break; cytogenetics—human; fertilization—in vitro; hyperhaploidy; hypohaploidy; karyotype; ova—hamster; spermatozoa—human.

Notes and References

1. P. A. Jacobs, "Chromosome Abnormalities and Fertility in Man," in *Proc. Intl. Symp. on the Genetics of the Spermatozoan*, R. A. Beatty and S. Gluecksohn-Waelsch, Eds. (University of Edinburgh Press, Edinburgh, 1971), pp. 346-358.
2. J. Boue, A. Boue, and P. Lazar, "Retrospective and Prospective Epidemiological Studies of 1500 Karyotyped Spontaneous Human Abortions," *Teratology* **12**, 11-26 (1975).
3. Z. Stein et al., "Spontaneous Abortion as a Screening Device," *Amer. J. Epidemiol.* **102**, 275 (1975); M. S. Golbus, "Chromosome Aberrations and Mammalian Reproduction," in *Fertilization and Embryonic Development*, L. Mastroianni, Jr., and J. D. Biggers, Eds. (Plenum Press, 1981), p. 257.
4. J. del Mazo, A. P. Castillo, and J. A. Abrisqueta, "Trisomy 21: Origin of Non-Disjunction," *Hum. Genet.* **62**, 316-320 (1982); A. Hansson and M. Nikkelsen, "The Origin of the Extra Chromosome 21 in Down's Syndrome," *Cytogenet. Cell Genet.* **21**, 194-203 (1978).
5. E. Rudak, P. A. Jacobs, and R. Yanagimachi, "Direct Analysis of the Chromosome Constitution of Human Spermatozoa," *Nature* **274**, 911 (1978).
6. R. H. Martin et al., "Direct Chromosomal Analysis of Human Spermatozoa. Preliminary Results from 18 Normal Men," *Amer. J. Hum. Genet.* **34**, 459 (1982); R. H. Martin et al., "Sperm Chromosome Abnormalities in Men of Proven and Unproven Fertility," *Fertil. Steril.* **39**, 441 (1983); R. H. Martin et al., "The Chromosome Constitution of 1000 Human Spermatozoa," *Hum. Genet.* **63**, 305 (1983).
7. B. Brandriff et al., *Chromosomal Abnormalities in Human Sperm: Comparisons Among Four Healthy Men*, Lawrence Livermore National Laboratory, Rept. UCRL-89826 (1984); *Hum. Genet.* (1984, in press).
8. ISCN, "An International System for Human Cytogenetic Nomenclature," *Cytogenet. Cell Genet.* **21**, 309 (1978).
9. R. H. Glass and R. J. Ericsson, *Getting Pregnant in the 1980s. New Advances in Infertility Treatment and Sex Preselection* (University of California Press, Berkeley, 1982), p. 58.
10. D. K. Edmonds et al., "Early Embryonic Mortality in Women," *Fertil. Steril.* **38**, 447-453 (1982).
11. B. Dutrillaux et al., "Risk of Chromosomal Disease due to Radiation. Tentative Estimate from the Study of Radiation-Induced Translocations in Human Fibroblasts," *Mutat. Res.* **119**, 343 (1983).
12. R. J. DuFrain et al., "Evaluation of Chemically Induced Cytogenetic Lesions in Rabbit Oocytes. II. A Comparison of Streptonigrin Effects on Somatic and Germ Cells," *Mutat. Res.* **94**, 103 (1982); P. P. W. van Buul, "Absence of Correlation Between the Chromosomal Radiosensitivity of Peripheral Blood Lymphocytes and Stem-Cell Spermatogenesis in Mammals," *Mutat. Res.* **95**, 69 (1982).
13. C. J. Epstein and B. Travis, "Preimplantation Lethality of Monosomy for Mouse Chromosome 16," *Nature* **290**, 144 (1979); T. Magouso, S. Smith, and C. J. Epstein, "The Development of Monosomy 19 Mouse Embryos," *J. Embryol. Exp. Morph.* **69**, 223 (1982).
14. C. E. Ford, "Gross Genome Unbalance in Mouse Spermatozoa: Does it Influence the Capacity to Fertilize?" in *Proc. Intl. Symp. on the Genetics of the Spermatozoan*, R. A. Beatty and S. Gluecksohn-Waelsch, Eds. (University of Edinburgh Press, Edinburgh, 1972), pp. 359-369.

High-Speed Cell Sorting: A Tenfold Increase in Sorting Rate

Three significant engineering improvements in our high-speed cell sorter allow us to quickly isolate cells or subcellular particles that occur at low frequencies within a population. System resolution and efficiency remain high even at processing rates above 20 000 particles per second, while sorted cells and chromosomes are recovered with little or no damage.

Flow cytometry has become an important technique with wide applications in the biomedical sciences. Using cytometric methods, researchers can automatically identify and sort cells or subcellular material for subsequent assessment. Samples to be sorted are suspended in an aqueous buffer solution, stained with a fluorescent dye, and then processed by a fluorescence-activated cell sorter. The machine scans the samples with a laser beam to accurately select specific subsamples of interest.

During the last 15 years, cell sorters have been used to purify a variety of cell and organelle types. Leukocytes (white blood cells) and bone marrow cells can be distinguished by their differential affinity for antibodies.¹ Human chromosomes can be sorted by their differential uptake of one or more dyes that are specific for deoxyribonucleic acid (DNA).² Rare, mutant erythrocytes (red

blood cells) and other peripheral blood cells can be identified by their altered affinity for antibodies against a variant protein.³ Malignant cells can be differentiated by their altered ability to scatter light and bind DNA-specific stains or antibodies.⁴ Once the objects of interest are sorted, they become available for visual study, biochemical analysis, or functional assessment.

Most commercial cell sorters can scan several thousand cells per second and then select specific cells to be purified. Unfortunately, a processing rate in this range is actually too slow for many applications. For example, bulk quantities of a single type of chromosome are required both for gene mapping and for the production of recombinant DNA libraries. When objects are sorted at a rate of 2000 s^{-1} , more than 120 hours are required to purify $1\text{ }\mu\text{g}$ of DNA from the human Y chromosome.⁵ Another

For further information contact
Donald C. Peters (415) 422-5628 or
Joe W. Gray (415) 422-5610.

example of the limitation imposed by slow processing rates is the measurement of the frequency of rare, mutant erythrocytes as an indication of mutagen-induced genetic damage. Such mutant erythrocytes occur at frequencies as low as one in every ten million; thus, extended analysis times are required to process over a billion cells to locate 100 mutant cells for visual identification and study.

Researchers in the Biomedical Sciences Division at LLNL have recently developed a prototype high-speed sorter that promises to overcome many of the limitations imposed by the relatively slow processing speed of conventional instruments. This research has resulted in a system that operates more than ten times faster than conventional systems and is capable of processing and sorting cells at rates in excess of $20\,000\text{ s}^{-1}$. When completed, the high-speed sorter will substantially improve the capabilities of the Biomedical Sciences Division to purify chromosomes and mutant cells. The instrument will permit us to study the biochemical composition of individual chromosomes, to determine the chromosome specificity of chemical carcinogens or mutagens, to facilitate the production of chromosome-specific recombinant DNA libraries, and to simplify quantification of frequencies of mutant cells in humans as a measure of mutagen-induced genetic damage.

In this article, we describe the engineering efforts that were necessary to develop our high-speed sorter. We then demonstrate the utility of our system for sorting cells and chromosomes. Specifically, we present the outcomes of several evaluations, including resolution of total fluorescence, sorting efficiency, and maximum sorting speed. Finally, we discuss the effects of the sorting process on cellular and chromosomal damage.

Principles of Operation

Most fluorescence-activated cell sorters, including our high-speed system shown in Fig. 1, operate according to the same general principles.⁶ Samples of interest are first stained with a fluorescent dye that is specific for a

cellular component such as DNA, RNA, or a particular protein. Individual particles are suspended and carried single file within a fast-moving aqueous stream. The stream intersects a focused laser beam, the wavelength of which is adjusted to excite the fluorescent dye. The resulting fluorescent light from each cell or particle is collected by an optical system and projected onto a photodetector. The detector and associated electronics convert the fluorescence intensity to an electrical pulse, the amplitude of which is proportional to the total fluorescence emitted as the particle crosses the laser beam.

Each particle to be selected is recognized by its distinctive fluorescence intensity. After fluorescence measurement, each particle flows within the aqueous jet-in-air until the jet breaks into discrete droplets. The distance from the point at which the particle is illuminated by the laser beam to the region where the droplet is formed (called the droplet breakoff point) is adjusted and made stable by vibrating the entire flow chamber from which the jet is formed. Vibrations are induced by a piezoelectric crystal that is mounted axially on the top of the chamber.

To separate a cell or particle for purification, a voltage pulse is applied to the liquid jet at a time equal to the transit time from the point at which fluorescence was measured to the droplet breakoff point. The width of the voltage pulse is selected according to the number of droplets to be sorted for a given object. To allow for timing errors, three droplets usually are charged after a desired particle is detected (ideally, the droplet containing the desired particle and one on either side).

The droplets containing the cells or particles to be sorted separate from the end of the jet during the application of voltage pulses and, therefore, retain an electrostatic charge that is proportional to the voltage at the end of the jet and the capacitance between the droplets and ground. Droplets forming at other times remain uncharged. In this way, the cells or particles to be studied are encapsulated in an electrically charged

droplet. Uncharged droplets are either empty or contain cells or particles of no interest. Both charged and uncharged droplets then fall through a high-voltage electric field (see the deflection plates in Fig. 1) and the charged droplets are separated electrostatically. The charged droplets are recovered in a collection vessel, and the uncharged droplets are discarded.

Design Considerations

In typical applications, it is essential that the sorted subsamples of cells or particles be pure. Purity is ensured by withholding a charging pulse on certain occasions. If, for example, unwanted particles will reach the droplet breakup point during the charging period for a desired object, the charging pulse is inhibited. Because three droplets are usually charged to allow for timing inaccuracies, and because particles appear at random, the optimal processing rate in cell sorters is roughly ten times lower than the rate of droplet production. Thus, final processing is often limited by the rate at which droplets can be formed.

We reduced this limitation in our high-speed sorter by increasing the frequency of droplet production to $200\,000\text{ s}^{-1}$ (200 kHz). This required several substantial engineering developments, including a high-pressure fluid system capable of operating at pressures well in excess of 1.4 MPa, a low-turbulence flow nozzle to produce a stable jet and droplet, and an electronics package able to analyze and sort more than 20 000 particles per second.

The frequency of droplet production in any sorter jet is determined largely by two parameters—the velocity and the diameter of the jet. Rayleigh⁷ has shown that droplets are produced most efficiently when the frequency f between droplets equals $v/4.5 d$, where v is the jet velocity and d is the jet diameter. We selected a nozzle diameter of $75\text{ }\mu\text{m}$ from our experience with nozzles of this diameter in conventional sorters. Studies suggested that a $75\text{-}\mu\text{m}$ -diameter nozzle produces little damage to flowing cells or chromosomes and that, with some care, the nozzle rarely becomes plugged.

The diameter of the jet as it emerges from our $75\text{-}\mu\text{m}$ -diameter nozzle orifice

Fig. 1

Schematic diagram of the high-speed cell sorter. Samples, delivered by a motor-driven syringe, are surrounded by a conductive sheath of liquid. Individual particles within the liquid stream intersect a tightly focused argon-ion laser beam inside the quartz flow chamber. Fluorescence and scattered light are collected by a television camera lens, focused through an adjustable pinhole, and projected through spectral filters onto a pair of photomultipliers. Particles of interest are identified as they pass through the laser beam according to the signal derived from the photodetectors and associated circuitry. After identification, particles exit the flow chamber through a $75\text{-}\mu\text{m}$ -diameter synthetic ruby orifice. Vibration of the flow chamber (at about 200 kHz) by a piezoelectric crystal causes the liquid jet to break up into droplets of equal size and spacing. The helium-neon laser, camera, and monitor are used for viewing the jet and droplets. Droplets containing objects of interest are charged, deflected by an electric field, and collected for subsequent analysis.

is approximately 50 μm , so that a jet velocity of about 50 m/s is required for the most efficient droplet production at 200 kHz. To produce a jet with a higher velocity, we increased the operating pressure on the sample and its liquid sheath. According to Bernoulli's equation, velocity is related to pressure according to the equation $v = \sqrt{2P/\rho}$, where P is water pressure and ρ is the density of the sheath liquid. Therefore, to produce the desired 50-m/s aqueous jet-in-air, we increased the system operating pressure to 1.4 MPa.

The design of the flow chamber and nozzle proved to be critical for regular droplet production, an essential requirement for accurate sorting. The samples are introduced into the central axis of a square-cross-section quartz flow chamber from a motor-driven syringe. Within the flow chamber, particles are surrounded by an electrically conductive liquid (aqueous) sheath. The sheath, which confines the particles to a small (5- to 10- μm -diameter) central core, flows from a stainless steel tank that is pressurized to 1.4 MPa with dry nitrogen. The particles travel single file through the flow chamber. There, they are illuminated by a precisely focused argon-ion laser that operates at wavelengths of 351, 364, 458, or 488 nm, depending on the fluorescent stain that has been used. The particles and sheath exit the flow chamber through a 75- μm -diameter orifice.

Figure 2 is an illustration of the final design for our flow chamber. A key feature is the length-to-diameter ratio of the orifice through which the fluid exits the chamber. Our objective in the design of this orifice was to form a jet-in-air that breaks into droplets in a highly regular fashion. We found that instabilities in droplet formation decreased significantly as the length-to-diameter ratio of the orifice was reduced, thereby minimizing the turbulence induced by the interaction between the flowing liquid and the orifice walls.

Cytometric analyses of the particles are made while they are within the quartz flow chamber. In this way, we take advantage of the high optical quality of the quartz and can make more

effective use of high-quality optical components for laser-beam shaping and fluorescence collection. In addition, the flow velocity in the quartz region is only about one fifth the velocity of the jet-in-air. The relatively low velocity allows the cells to remain in the measurement region longer so that more accurate measurements of total cellular or particle fluorescence are possible.

Optical and Electronic Components

The shaping optics for the laser beam consist of the three components shown in Fig. 1: a beam expander that increases the beam to a 2-cm diameter, a 400-cm focal length cylindrical lens, and a 6-cm spherical lens. The lenses focus the laser beam so that it has an elliptical cross section with a 3- μm minor axis and a 300- μm major axis upon reaching the stream of particles or cells. The stream intersects the tightly focused 3- μm axis of the ellipse to minimize the probability that more than one object at a time will be illuminated by the beam.

Fluorescence and scattered light arising from the interaction of a particle and the laser beam are collected by an $f/1$ television camera lens. The light is focused through an adjustable pinhole and projected, through one or more spectral filters, onto a pair of photomultipliers. The filters select the specific wavelength of light that will pass to a photomultiplier. For example, one type of filter may pass the fluorescence emission from a particle and block the scattered laser light while the other filter may pass the scattered laser light and block the fluorescence.

Particles are sorted according to the nature of the signal derived from the photodetectors as each particle passes through the laser beam. The output of each photomultiplier (Fig. 3) is integrated and shaped by a bandpass amplifier to optimize the signal-to-noise ratio and to minimize pulse pileup that might otherwise occur as particles pass in rapid succession through the laser beam. The amplitude of pulses from the bandpass amplifier is proportional to the total light from the particle. Pulses are processed by a pulse-height analyzer and a window

discriminator. The pulse-height analyzer generates a histogram of the distribution of pulse amplitudes among the particles of the measured population. Particles are selected for sorting by the window discriminator, which detects pulses with amplitudes that are within a preselected range.

Droplet Formation, Charging, and Deflection

After analysis by the electronics package, all particles entrained in the aqueous liquid stream exit the quartz flow chamber through the 75- μm -diameter ruby orifice. Vibration of the entire flow chamber at approximately 200 kHz is induced by a piezoelectric

Fig. 2

Design of the flow chamber. Sheath liquid is injected axially through a tube and through a piezoelectric crystal mounted on top of the flow chamber. The liquid flows downward into a tapered quartz section where the sample is injected. Particles are entrained in the liquid and flow in single file through a 250- μm square quartz channel where they intersect an argon-ion laser beam. The laser light excites one or more fluorescent dyes in each particle and is scattered. The fluid then exits through a 75- μm -diameter orifice. The enlargement shows the orientation and shape of the orifice that is formed by a polished, synthetic ruby jewel.

crystal mounted axially on the top of the chamber. The regular vibrations create pressure waves that are transmitted to the liquid jet. In turn, the pressure waves cause regularly spaced cylindrical depressions on the jet surface, thereby forming droplets of equal size and spacing. In this manner, we ensure that the time for a cell or particle to travel from the measurement region to the droplet breakoff point is constant.

Droplets are charged by applying a voltage to the end of the liquid jet for the duration of the number of drops to be charged. The charging voltage is actually applied to the liquid within the flow chamber. The voltage at the end of the jet, which is capacitively coupled through the air to ground, is limited by the resistance of the jet and the short time that a charge is applied.

Droplet charging and deflection are complicated by the shape of the jet at the breakoff point and by the velocity of the droplet during passage through the deflecting electric field. Figure 4 is a

photograph of the liquid jet in the vicinity of the droplet breakoff region. The region is approximately 2 cm downstream from the point at which particles are measured, and the transit time from the measurement region to the breakoff point is about 500 μ s.

There are two important differences between our high-speed sorter and conventional sorters in this droplet breakoff region. First, the liquid neck between droplets is extended in the high-speed sorter (see Fig. 4), thereby increasing the electrical resistance of the jet. Second, the jet velocity is higher. The diameter of the liquid neck is sufficiently small (approximately 2 μ m) to cause interference with the induction of charge on the separating droplets, thereby interfering with proper droplet deflection during sorting. At higher jet velocities and droplet production rates, the neck diameter becomes even smaller and eventually prevents proper droplet charging. To compensate for our higher velocity, we use deflection plates that are

Fig. 3

Electronics block diagram. A photomultiplier tube converts the fluorescence intensity of an illuminated particle to a current that is integrated as the particle passes through a laser beam. The output of the integrator is shaped with a bandpass amplifier to optimize its signal-to-noise ratio and to reduce pulse-pileup problems from closely spaced particles. The output of a master discriminator, which detects all particles passing through the sorter, is used to recognize closely spaced objects and to inhibit sorting when an undesired one might be sorted simultaneously. The computer monitors the threshold levels of the sorting window discriminators and displays them as intensified markers on the pulse-height analyzer. The computer also is used for setting the delay between particle detection and droplet charging and for selecting the number of droplets to be charged. The piezoelectric transducer used for vibrating the flow chamber is driven with a square wave from the frequency generator. An acousto-optic modulator, mounted on the output of a 3-mW helium-neon laser diffracts out a 1- μ s pulse of light at the droplet frequency for stroboscopic illumination of the droplet breakoff region. Jet stability is observed via a television camera and monitor.

about 18 cm long, longer than those in conventional sorters. This increases the time the droplets spend in the deflecting field. We also raised the potential difference between the plates (from the approximately 5 kV of previous designs) to 10 kV to increase the deflecting force on each charged droplet.

To induce a charge on the droplets containing objects to be sorted, we apply a 120-V pulse to the liquid jet for a time equal to that required for one or a few drops to break off. Each drop has a charge proportional to its capacitance to ground and to the voltage at the end of the jet.

As the charged droplets are deflected into a collection device, we take additional precautions to avoid loss of material due to aerosol formation or splattering that can occur when droplets strike an aqueous surface at high velocity. We use different types of collection device depending on the intended use of the sorted material. For example, we have successfully used vacuum-backed, small-mesh filters, frozen or heated microscope slides, and large test tubes with small entrance holes.

System Performance

We conducted several studies to determine the operating characteristics of our high-speed sorter. Specifically, we evaluated the resolution of total fluorescence, the sorting efficiency, and the damage (if any) resulting to cells or chromosomes sorted at high speed. As test objects, we used microspheres, living cells obtained from hamster ovaries, ovary cells fixed in ethanol, and hamster and human chromosomes.

The system resolution was determined by measuring fluorescence from uniform, 1.8- μm -diameter fluorescent polystyrene microspheres and fluorescence from propidium-iodide-stained human chromosomes. We selected these test particles because of their relatively constant dye content and size (coefficient of variation less than 1%). The fluorescence distributions measured for microspheres at processing rates of 1000, 5000, 10 000, and 25 000 s^{-1} increased only from 1.0 to 1.4% with increasing

processing rates, an acceptable variation. The fluorescence distribution measured for human chromosomes processed at 20 000 s^{-1} (Fig. 5) was similar to the distribution measured for chromosomes processed at 2000 s^{-1} , confirming the feasibility of high-speed sorting.

We determined the sorting efficiency of the system by sorting known numbers of microspheres (typically from 20 to 50) onto filters with a small pore size (about 0.2 μm). The measured sorting efficiency was routinely greater than 90% for a three-droplet sort (the droplet containing the microsphere and one droplet on

Fig. 4

Photograph of an aqueous jet at the droplet breakoff region. In this experiment, the frequency of droplet production was 215 kHz. The jet was illuminated by a helium-neon laser modulated at the frequency of droplet production during the several-second exposure required to produce the picture. Notice the long, thin neck that connects droplets prior to breakoff.

either side were charged and deflected). When the particles were sorted at a very low rate (about 1 s^{-1}), the collection efficiency was 100%. This slow sorting rate allows enough time between events for the droplets to be absorbed into the filter, thereby eliminating droplet-on-droplet splattering. Sorting efficiency was 70% for a one-droplet sort, at a sorting rate of about 1 s^{-1} . This lower sorting efficiency arises because of the small variations in time required for particles to travel from the measurement region to

the droplet breakoff point. In addition, the occurrence of a particle in the neck between droplets results in some uncertainty as to which droplet will eventually contain the desired particle.

To determine the extent to which the system might damage whole cells, we examined Chinese hamster cells following sorting. In one series of experiments, cells were fixed in ethanol and stained with chromomycin A_3 , a dye that binds principally to the DNA in the cell nucleus. The fixed and dyed samples were injected into the high-speed sorter, and all cells were sorted (charged and deflected). The sorted cells were collected and processed through a conventional flow cytometer to determine the effect of high-speed sorting on the total fluorescence distribution. Figure 6 shows the total fluorescence distribution for samples both before and after high-speed sorting. It is clear that the effect of the sorting process is negligible, suggesting that our high-speed system does not produce mechanical forces sufficient to disrupt the cells. Through microscopic examination, we confirmed that the sorted cells were not physically damaged.

Living Chinese hamster cells were also sorted and analyzed. We anticipated that there might be problems in processing living cells because they are more fragile than cells that are fixed and stained with fluorescent dyes. The living cells were sorted according to their large-angle light scatter, the intensity of which is proportional to a cell's size. The cells were sorted directly into plastic culture dishes and cultured to determine the plating efficiency (the frequency with which cells can grow to form colonies of 50 cells or more). The plating efficiency for our sorted cells relative to that for unprocessed cells was 95%; therefore, most of the cells were unaffected by passage through the sorter. The missing 5% might have been killed by the sorting process or might have splashed out of the aqueous culture medium in the petri dishes.

We also examined the effects of our high-speed sorter on chromosomes derived from Chinese hamster cells and from human lymphocytes. We assessed

Fig. 5

Fluorescence distribution for human chromosomes stained with propidium iodide and measured while processing at a rate of $20\,000 \text{ s}^{-1}$. This distribution is almost identical to fluorescence distributions measured at 2000 s^{-1} .

Fig. 6

Fluorescence distributions for Chinese hamster cells. Cells were fixed in ethanol and stained with chromomycin A_3 , a fluorescent dye that binds principally to DNA in the cell nucleus. Distributions before (a) and after (b) sorting suggest that our high-speed system does not produce mechanical forces sufficient to disrupt the cells.

three possible types of chromosome damage following the sorting process: changes in chromosome morphology, total fluorescence distribution following staining with a DNA-specific fluorescent dye, and alteration in molecular weight of extracted DNA. Upon microscopic examination, the morphology (i.e., the gross shape) of sorted chromosomes appeared to be unaltered. The total fluorescence distribution of the sorted chromosomes is identical to the fluorescence distributions of unsorted chromosomes, indicating that stained chromosomes are unaffected by passage through the high-speed sorter.

We also used gel electrophoresis to investigate the integrity of the DNA molecule in sorted chromosomes. Gel electrophoresis is a technique for separating molecules of differing molecular weight. Material is introduced into a column of gel, along which is applied a static electric field. The velocity of migration along the gel is inversely related to the size of the DNA molecule. Preliminary analyses suggest that the molecular weight of our extracted DNA remained acceptably high (i.e., the DNA molecules were unbroken) following the high-speed sorting process.

Summary and Future Applications

We have constructed a high-speed cell sorter that is capable of operating at a pressure of 1.4 MPa and at a droplet production frequency greater than 200 kHz. Our high-speed sorter represents a substantial improvement over conventional sorters in several respects. Particles can be processed and sorted at rates more than ten times faster than was previously possible. Both sorting efficiency and measurement accuracy are comparable to conventional sorters. The system does not induce significant damage to either fixed or living cells or to chromosomes.

For future applications, our sorter is well suited to the purification of cells and organelles in biochemical studies that require large amounts of material for analysis and, thus, high sorting rates. In addition, the system is ideal for purifying objects that occur at low frequencies

within a population. Indeed, we envision a variety of applications that were, heretofore, too time consuming and costly. For example, the system can be used to purify bulk quantities of chromosomes of a single type for gene mapping,⁸ to produce recombinant DNA libraries,⁹ and to concentrate variant peripheral blood cells that occur at low frequencies. With respect to the latter application, we can use our system to more precisely estimate the frequency of variant cells as a measure of mutagen-induced genetic damage. **□**

Key Words: cell sorter—fluorescence-activated, high-speed; chromosome—Chinese hamster, human; flow cytometry; laser—argon-ion, helium-neon.

Notes and References

1. R. A. Hoffman and W. P. Hansen, "Immunofluorescent Analysis of Blood Cells by Flow Cytometry," *Intl. J. Immunopharmacol.* **3**, 249-254 (1981).
2. A. V. Carrano *et al.*, "Measurement and Purification of Human Chromosomes by Flow Cytometry and Sorting," in *Proc. Natl. Acad. Sci. USA* **76**, 1382-1384 (1979).
3. W. L. Bigbee *et al.*, "Cell Sorter Immunofluorescence Detection of Human Erythrocytes Labeled in Suspension with Antisera—specific for Hemoglobin S and C," *J. Immunol.* **128**, 117-127 (1981).
4. N. S. Kruth *et al.*, "Simultaneous Analysis of DNA and Cell Surface Immunoglobulin in Human B-Cell Lymphomas by Flow Cytometry," *Cancer Res.* **41**, 4895-4899 (1981).
5. C. Cremer, J. W. Gray, and H. H. Ropers, "Flow Cytometric Characterization of a Chinese Hamster X Man Hybrid Cell Line Retaining the Human Y Chromosome," *Hum. Genet.* **60**, 262-266 (1982).
6. M. Melamed, P. Mullaney, and M. L. Mendelsohn, *Flow Cytometry and Sorting* (J. Wiley and Sons, New York, 1979).
7. J. W. S. Rayleigh, *The Theory of Sound* (Dover Publications, Inc., New York, 1945), 2nd ed., Vol. II, Ch. XX.
8. R. V. Lebo *et al.*, "Assignment of Human α , β , and γ -Globin Genes to the Short Arm of Chromosome 11 by Chromosome Sorting and DNA Restriction Enzyme Analysis," in *Proc. Natl. Acad. Sci. USA* **76**, 5804-5808 (1979).
9. R. Krumlauf, M. Jeanpierre, and B. D. Young, "Construction and Characterization of Genomic Libraries from Specific Human Chromosomes," in *Proc. Natl. Acad. Sci. USA* **79**, 2971-2975 (1982).

Preparing for the Super Supercomputers

Our modeling studies indicate that to achieve computational speeds a thousand times faster than those of present supercomputers, an increase of a hundred will come from improved parallelism in hardware and in computational methods.

For further information contact
Timothy S. Axelrod (415) 422-4314.

Several groups, both at LLNL and elsewhere, are working to create, within a decade, computing systems a thousand times more powerful than the current "supercomputers." Because this ambitious $1000\times$ goal can be achieved only through revolutionary changes in the nature of computers and of computation itself, there is only limited agreement on the technology to be employed in these new computers, how they will be organized, or what kind of software they will need. As is often the case on the verge of a revolution, the number of possible paths seems large and their relative merits unclear.

Although we are aiming for $1000\times$ faster systems, the generally predicted large-system improvement in speed by 1990 is less than $10\times$. While we recognize that technological breakthroughs may yield additional improvements, increased parallelism in computer architecture must play a dominant role in reaching the $1000\times$ goal.

Here, we sketch some of these possible paths and explore in some detail one of the most promising. This latter path involves making major advances in the design of computational physics algorithms for parallel computers. It is essential to begin work on these

advanced concepts now, even while knowledge of the future computers and their characteristics is incomplete, because it takes years to develop a new class of computational methods.

Architectural Considerations and Difficulties

It is useful to recall the role that parallelism has already played in improving the performance of large scientific computers. For example, the Cray 1, first delivered in 1976, is about a million times faster on specialized tasks such as matrix multiplication than the EDSAC¹ of the 1950s.¹ Of this factor of a million, only about 200 is attributable to increased clock speed, which went from 2000 ns to 12.5 ns. The rest is attributable to increased parallelism.

This parallelism was not added in a single step, of course, but in a long series of steps. The first of these, the change from bit-serial to bit-parallel arithmetic, is now taken so completely for granted that it is seldom recognized as the major increase in parallelism that it was. The subsequent evolution of parallelism has been capably discussed elsewhere.² It is important to recognize, however, that the progression from bit-parallel arithmetic to ever more comprehensive pipelining and the use of multiple functional units

has made parallelism more and more visible to the programmer. Any serious user of a "vector" machine (e.g., a Cray 1 or Cyber-205) must take careful account of the parallelism of the machine when designing his algorithms. In contrast to this, the great increase in parallelism between the EDSAC1 and the IBM 7094, for example, would generally appear to the user only as increased speed, without requiring any algorithm redesign.

In our search for further substantial increases in parallelism, the areas of interest appear to be:

- Further increasing the parallelism in executing a single instruction stream (SI) acting on a single data stream (SD), as in a "scalar" SISD machine.
- Increasing the number of data streams (multiple data or MD) operated on by a single instruction stream, as in "vector," or "array" SIMD machines.
- Replicating both instruction and data streams, as in "multiprocessor" MIMD machines.

Considered for the moment to be mutually exclusive alternatives, only the third one offers any realistic hope of attaining the required factor of 100. The first alternative already has been exploited nearly to its limit and, thus, can yield probably no more than a twofold improvement. The second alternative can easily attain the factor of 100 but only for specialized applications (e.g., image processing) that match the architecture closely. For the more general scientific applications in which we are interested, we must develop the MIMD architectures.

Considered simply, the MIMD approach is to partition a single computational problem into n smaller pieces and assign each of these to an autonomous computing engine, producing an aggregate processor with n times greater speed. To attain our $1000\times$ goal, we could use 100 of the fastest available processors, 1000 current-speed processors, or even 1 000 000 very simple processors. There are advocates for each of these positions, but we have good reasons to believe, first, that for any fixed goal, n should be as small as possible, consistent with achieving the goal, and, second, that the task of designing an

effective computing system of this nature is far more difficult than it initially appears.

One of the first issues to be faced is how to organize and interconnect the many processors. Where, for example, should we keep the data that each processor will use? Should we use a regular network of processors, each with its own local memory and a few channels for communicating with its neighbors (Fig. 1), or should we use a system in which a large number of processors without individual memories all have equal access to the data in a large central memory (Fig. 2)? A processor in the first or "loosely coupled" arrangement can use data from anywhere in the network, but it may take many communication steps to get it. On the other hand, a processor in the second or "closely coupled" system may find the section of memory it needs busy when it requires data.

Fig. 1

Diagram of a network multiprocessor in which each unit processor can communicate with every other unit (in one or more steps) to share data stored in the individual memories.

As presented so far, of course, these two examples are mere caricatures of real computer systems. There is no way to assess their performance without specifying a great many details such as the internal organization of the shared central memory and how many processors can tap it concurrently. Nonetheless, we can compare the two systems in terms of what happens to them as we increase the number of processors.

In the case of the loosely coupled processor network, the answer depends greatly on the communication pattern of the application problem. Unless this pattern is highly localized, system performance will increase only slowly as processors are added (and it may even decrease) because, if each processor is

equally likely to call up any other in the network, the bigger the network, the greater the fraction of "long distance" calls that will involve a large number of communication steps. Thus, the average time required to obtain a piece of data will increase for larger and larger networks. Unless some other factor offsets this inexorable trend, sooner or later, most of the processors will be idle, waiting for the "overseas operator."³⁻⁵

Furthermore, if we partition a problem of fixed size into an ever-greater number of pieces, these pieces will need to communicate with one another ever more frequently. In the common case, where the "pieces" are contiguous groups of zones in a two-dimensional mesh, for example, the intercommunication rate grows as the ratio of the perimeter to the area.

At first glance, the shared memory of the closely coupled multiprocessor seemingly avoids these communication-time problems because all of the processors are equally near to all the data. Unfortunately, this is true only so long as we neglect the shared memory's internal structure. With realistic choices here, similar though less severe problems arise.

The main point is that, in all existing and planned multiprocessors, the time to access the memory (in CPU clock cycles) increases rapidly, largely because memory size must grow roughly in proportion to computer speed. Large memories are slower than small ones, both because they contain more gate delays and because it is impractical, economically and technically, to use superfast logic chips in them. This trend makes it ever more difficult to design a true closely coupled multiprocessor in which it is quicker to communicate a result between processors than to recompute it wherever it is needed.

Both classes of systems, then, exhibit the same basic problem: the component pieces of a computational task must communicate with one another, but this communication is expensive and tends to limit the computing system's speed. As the number of computing elements increases, this problem intensifies rapidly.

Fig. 2

Diagram of a shared-memory parallel processor, in which each processor, with no memory of its own, has equal access to a large central memory.

Analysis of a Possible Approach

The way around this difficulty apparently has two main components. The first, well known and much used in current machines, is to exploit whatever "locality" there is in the computation by giving each processor a small but very fast memory that can satisfy local references without involving the slow shared memory system. The vector registers of the Cray 1 and the cache memory of the S-1 computer⁹ are two quite different ways to organize the local memory, and there are a number of others. Fast local memory alone is not enough, however. Communication must still take place and it is inevitably slow.

The second component of the solution relies on a further large increase in parallelism. The basic idea is simple: give each processor a long list of parallel processes, any of which it can work on at any particular time, and give each processor the ability to switch very rapidly from one process to another. Then, whenever the active process needs data from the slow shared memory, the processor immediately puts it aside and goes ahead on some other process. If there are enough parallel processes available, we can "hide" the arbitrarily long communication delays. This strategy is being used in the current MIMD computer called the HEP-1.^{7,8}

The success of a computing system built along these lines depends not only on correct choices among a large number of technological alternatives but also, most critically, on the amount of parallelism in the computations that will use it. A real computation always involves dependencies between parallel tasks, which reduces parallelism. The question then becomes one of how much intertask dependence can be tolerated without seriously compromising system performance. Clearly, the technological decisions must be made in light of what is possible in parallel computation, and developments in parallel computation must be guided by what is possible and desirable technologically.

As an aid to sorting through the many possibilities for the design of such

parallel systems, and particularly as a guide for the development of future parallel algorithms, we have created a simplified model of an idealized MIMD computer (Fig. 3). The model parametrizes both the hardware and the application program being run, so that we may study the interaction of the two. In this model, the shared central memory consists of a large number of independent memory banks connected, through a pipelined multistage switch, to a large number of communication ports, a degree of complexity required to supply sufficient parallelism in its operation. Much work has recently been done on the design and analysis of such memory systems.^{9,10}

Unlike a crossbar switch, which is maximally conflict free (used currently by the S-1 computer, among others), the pipelined multistage switch used in our model permits some conflicts (in which case one of the requests must be delayed). This design is appropriate for

Fig. 3

Layout of a possible architecture for a supercomputer a thousand times more powerful than those now existing.

very large numbers of memory banks and communication ports because its complexity grows relatively slowly as the number of banks and ports is increased. Both the HEP-1 and the proposed NYU Ultracomputer¹¹ are broadly similar in design to this system.

Our group's main tool for analyzing the performance of such a system is the theory of queuing systems,^{12,13} which governs any network of queues and servers inhabited by some population of customers. We test the system as a stochastic one and look for steady-state averages of the variables that describe the system's status. The approach is quite analogous to that of statistical mechanics. Its principal advantage is that it can accurately predict some aspects of the behavior of systems that are too complex to simulate or otherwise analyze in detail.

To model our proposed computer system, we treat the processors, switch elements, and memory banks as servers and consider all of the parallel processes that make up the application program as their customers (Fig. 4). Despite the general structural similarity of the model to the physical organization of the system, the model behaves in ways that are not evident from the system's physical structure. There are two classes of customer processes, for example, those that produce information and those that consume information. The model explicitly shows the differing behavior of

producers and consumers and the ways in which a particular process can switch from one class to the other.

Let us follow the progress of an arbitrary process as it circulates through the computing system, together with hordes of other processes. Suppose we begin our observation when it is in one of the processors (i.e., its instruction stream is being executed). A unit time later, the processor evicts our marked process and begins executing instructions from the next process waiting in its queue (if any). At this point, the process we are observing can either rejoin the queue waiting for processor service (if the data it wants reside in its processor's local memory) or, if it needs to read or write data to the shared memory system, it can join one of the communication channel queues. The model makes this decision at random, with a probability β that it rejoins the processor queue and a probability $(1 - \beta)$ that it seeks communication service.

If our process requires communication, it joins the queue of one of the many communication channels associated with its processor, selected at random. Thereafter, it progresses from stage to stage within the switch as it travels toward a (randomly selected) memory bank where its communication request will finally be completed. At each stage within the switch, the process may have to wait until other processes in front of it have released a needed switch path.

Fig. 4

Representation of a queuing model used for analyzing the performance of super multiprocessors and deriving the necessary characteristics of programs designed to run on them. This diagram shows only one processor and associated communication channels. The model assumes a large number of processors working in parallel on a large number of parallel tasks.

P = processor
C = communication channel
M = memory
 α = consumer fraction
 β = processor evict fraction
 f = consumer memory probability

When the the process finally reaches the queue for the memory bank, its subsequent behavior depends on whether it is a producer or a consumer. If the process is a producer, its communication request can be fulfilled independent of the actions of other processes. If the process is a consumer, however, it may have to wait for some other process to perform an action before it can proceed. This action may be to supply data (the most common case in parallel algorithms), but the same basic situation can also arise in other ways, for instance with "critical sections" or "locks." In any case, the model decides on the basis of probability whether a process is a producer or a consumer and whether a consumer must wait to have its request filled. If our process must wait, it is waiting for some *particular* other process to arrive as a producer. The parameter input to the model is then the average number of communications that other process must make before it gets around to serving our waiting marked process.

To complete the life history of our marked process, its communication request is eventually satisfied at the memory bank. It then migrates back through the switch stages until it arrives at its processor queue. When it finally reaches the head of this queue and its instruction stream once again begins executing, the whole probabilistic cycle of events begins anew.

Even though our queuing model is extremely simple, compared with the envisioned computing system, it is still a formidable task to derive an exact solution when hundreds of processors and thousands of processes are involved. Hence, we used decomposition techniques¹³ to obtain an approximate solution. For example, we replaced the multistage pipelined switch with a single equivalent server whose load-dependent service rate includes the effects of conflicts between the requests of different processors.

Given values for the model input quantities, we obtain probability distributions for the queue lengths at the various servers of the model. From these probability distributions, we can easily

determine how much of the time the various hardware components are in use. We are, of course, most immediately interested in the utilization of the processors, since these perform all the arithmetic operations required to advance the calculation. Only if we can keep the processors busy most of the time can we obtain the maximum performance of which the hardware is capable.

To illustrate the results obtained from our model, let us assume an advanced computing system whose processors have a 1-ns cycle time, making each roughly 10 times as powerful as the Cray 1. We further assume that the cycle time of the memories and switch elements is 100 ns and that 95% of the memory references are satisfied locally. We neglect, for the moment, all interprocess dependencies, so that there are no consumer processes.

Figure 5a shows the number of *effective* processors the system delivers as a function of the number of processors and the number of communication channels per processor if the application program supplies 1000 parallel processes. This may seem like a large number, but even with 15 channels per processor, the maximum number of effective processors obtainable is only about 13. Thus, even if we have hundreds of processors available, most of them are idle most of the time. Figure 5b illustrates the effect of increasing the number of parallel processes to 10 000. This permits us to realize about 70 effective processors with 80 physical processors and 10 communication channels per processor, bringing us close to our 1000 $\times$ goal.

Admittedly, however, we reached (modeled) this state by *ignoring* interprocess dependency, which may severely reduce the effective power of these computing systems. We are currently investigating the technological constraints on the number of communication channels per processor and the number of processes per processor, as well as ways to model the dependency characteristics of various application programs. One conclusion seems secure: the application program must supply many thousands of parallel

processes if this approach is to deliver computing systems with a 1000X improvement in performance.

Implications for Algorithm Development

One serious limitation of this and other queuing models is that they only predict steady-state performance. Parallel

algorithms tend to have many transients, however, and it is quite possible for a system's performance to be determined by its reaction to these transients.

For example, our parallel implementation of SIMPLE, a two-dimensional hydrodynamics and heat conduction code, uses 14 global synchronization barriers each cycle.¹⁴ At

Fig. 5

How the effective computational power of a multiprocessor (in multiples of a single effective processor, indicated by the numbers on each step) varies with the numbers of individual processors and of communication links for each processor, when the problem being computed provides a large number of completely parallel tasks. (a) With 1000 parallel tasks, the maximum computing power, even with hundreds of processors and dozens of links apiece, is only about 13 times that of a single processor. (b) With 10,000 completely parallel tasks, an effective computing power about 70 times that of each unit processor can be obtained with 80 processors and 10 links apiece. However, these calculations neglect intertask dependence, a large factor in most real algorithms.

each of these barriers, every process must wait until all other processes have arrived. The arrival of the last process triggers the start of a new race to the next synchronization barrier. Hence, the number of active processes climbs from a minimum of one as the last straggler arrives at a barrier up to a maximum, and then decreases as processes arrive and wait at the next barrier. Unless the barriers are far enough apart, the system will never achieve anything approaching steady-state performance.

Calculating this minimum distance between barriers is quite difficult, but we can estimate what the synchronization barrier will cost in terms of performance by considering the random spread of the process arrival times.¹⁵ For a fixed total number of operations n , this analysis shows that the barrier will slow performance by more than a factor of two unless

$$n > (\sigma P)^2,$$

where σ is the ratio of the width of the cycle time distribution to its average value, and P is the number of parallel processes. This indicates that the distance between synchronization barriers, as measured in arithmetic operations, must grow rapidly as P increases. Furthermore, this is only a lower bound; the delays inevitable in building up parallelism in a real computing system will further increase the required separation.

A similar kind of bottleneck forms in an algorithm that begins from a single root process, grows geometrically in parallelism to some maximum, and then reverses the process.¹ We can readily calculate the price we pay for such behavior on an idealized machine, one that always has more parallelism than required by the algorithm. Let us consider only the "buildup" phase of the algorithm: let the number of parallel processes increase by the factor k at each step and suppose there is a total of n steps. Suppose each active process performs the same amount of work each step. We take this as our "unit operation," and the time to perform it as our "unit time." Then, the total number

of operations O performed during this n -time-step buildup is

$$\begin{aligned} O &= 1 + k + k^2 + \dots + k^{(n-1)} \\ &= (k^n - 1)/(k - 1). \end{aligned}$$

For comparison, consider an idealized parallel algorithm that performs the same number of operations but all at a parallelism P equal to the maximum achieved by the other algorithm (i.e., $k^{(n-1)}$). The time it would take to complete the task, therefore, would be only O/P .

If we define the efficiency e of our algorithm to be the ratio of the ideal time to the actual time, we obtain

$$e \cong [k/(k-1)] \log k / (\log k + \log P),$$

where the approximation comes from assuming P very much larger than 1. If, for example, $k = 2$ and $P = 10\,000$, then $e = 0.14$, and we pay roughly a factor of seven for our algorithm's leisurely buildup of parallelism. On a realistic machine, however, the penalty could easily be still greater.

Work Under Way

The above considerations indicate that some revolutionary changes in parallel algorithms may be required to make effective use of 1990-class computing systems. Not only are very high peak levels of parallelism required, but this parallelism must be sustained through time with an absolute minimum of downward fluctuations. Very few such algorithms are known, but there are some clues to how more of them may be developed.

Parallelism is frequently limited by the need for exactly reproducible results, which are hard to achieve in a calculation performed on a computing system with partially random behavior. Often, however, the calculation consists of an iteration that keeps going until the estimated error is small enough. Therefore, if we can still guarantee a small enough error, we should be able to accept random behavior in our algorithms. In fact, if we are willing to accept all the implications of this, it may be possible to achieve the continuous

high level of parallelism we require. We are currently attempting to extend some early work on such so-called "chaotic algorithms"¹⁶ and apply them to a wide class of problems in linear algebra and computational physics.

We have shown that the development of computing systems 1000 times as powerful as our current supercomputers will require significant advances not only in computer technology but also in the actual techniques of computation. It may well be that the greatest surprises will come in the form of new algorithms and approaches to computational problems.

Key Words: algorithms—chaotic, parallel, improvements; computer—Cray 1, Cyber-205, EDSAC, HEP-1, IBM 3081, IBM 7094, MIMD, NYU Ultracomputer, S-1, SIMD, SISD; parallelism.

Notes and References

1. R. Hockney and C. Jesshope, *Parallel Computers: Architecture, Programming, and Algorithms*, (Adam Hilger, Ltd., Bristol, 1981).
2. H. Lorin, *Parallelism in Hardware and Software: Real and Apparent Concurrency* (Prentice-Hall, New York, 1972).
3. D. Reed and H. Schwetman, *IEEE Trans. Computers* C-32, 83 (1983).
4. D. Reed, in *Proc. Intl. Conf. Parallel Processing* (1983), p. 161.
5. D. Reed, in *Proc. IEEE Intl. Workshop Computer Systems Organization* (1983), p. 72.
6. For a discussion of the S-1 computer, see the September 1979 issue of *Energy and Technology Review* (UCRL-52000-79-9), p. 1.
7. B. Smith, in *Proc. Intl. Conf. Parallel Processing* (1978), p. 6.
8. H. F. Jordan, in *Proc. Intl. Symp. on Computer Architecture* (1983), p. 207.
9. J. H. Patel, *IEEE Trans. Computers* C-30, 771 (1981).
10. D. M. Dias and J. R. Jump, *Computer*, 14, 43 (December 1981).
11. A. Gottlieb et al., *IEEE Trans. Computers* C-32, 175 (1983).
12. L. Kleinrock, *Queuing Systems* (John Wiley and Sons, New York, 1976), Vol. 2.
13. K. M. Chandy, U. Herzog, and L. Wou, *IBM J. Res. Devel.* 19, 43 (January 1975).
14. T. S. Axelrod, P. F. Dubois, and P. G. Eltgroth, presented at *Intl. Conf. Parallel Processing* (1983), Bellaire, Michigan.
15. M. Dubois and F. A. Briggs, *IEEE Trans. Computers* C-31, 1083 (1982).
16. G. Baudet, *The Design and Analysis of Algorithms for Asynchronous Multiprocessors*, Ph.D. dissertation, Carnegie-Mellon University (1978).

The Centralia Partial-Seam CRIP Test (Brief)

The 1983 test of our Controlled Retraction Injection Point (CRIP) system for extending the useful life of the injection well in underground coal gasification was successful. The heating value of the product gas obtained with steam-oxygen injection was about 9 MJ/m³. Gas quality deteriorated when the burn became large enough to cause roof collapse, but execution of the retraction maneuver restored the product-gas quality as expected.

Contact: Douglas R. Stephens (415) 422-7330.

Visualizing Chromosomes in Human Sperm

New techniques for fusing human sperm with hamster ova enable us to create a hybrid cell in which the sperm chromosomes are fully visible. Normally, sperm chromosomes are too tightly packed to permit direct observation. In the hybrid cell, however, the clearly displayed chromosomes can be examined under a light microscope for determination of structural and numerical abnormalities. We are analyzing sperm chromosomes from normal, healthy men to establish a baseline frequency of chromosomal errors. We plan to extend this study to chromosome abnormalities in sperm of men who have been exposed to environmental hazards or who have clinical reproductive disorders.

Contact: Brigitte Brandriff (415) 423-0758.

High-Speed Cell Sorting: A Tenfold Increase in Sorting Rate

Three significant engineering improvements in our high-speed cell sorting equipment (a fluid system capable of operating at pressures in excess of 1.4 MPa, a low-turbulence nozzle to produce a stable jet and droplet, and an electronics package able to analyze and sort more than 20 000 particles per second) greatly reduce the time needed to accumulate usable samples of uncommon cells or subcellular particles. The sorted particles or cells are not damaged by the sorting process, even at the highest speed. Research applications include the purification of bulk quantities of chromosomes of a single type for gene mapping, the production of recombinant DNA libraries, and the concentration of variant peripheral blood cells for study and evaluation.

Contact: Donald C. Peters (415) 422-5628 or Joe W. Gray (415) 422-5610.

Preparing for the Super Supercomputers

Our modeling studies indicate that to achieve computational speeds a thousand times faster than those of present supercomputers (a goal being actively pursued in various computer design centers and whose attainment is expected within the next decade), only a factor of 10 of the speed increase will come from improved components. The remaining factor of 100 will come from improved parallelism, in both hardware and computational methods. We explore some of the implications of this finding, which appears to call for the development of an entirely new class of computational methods.

Contact: Timothy S. Axelrod (415) 422-4314.

Past Titles

Articles published in recent issues of the *Energy and Technology Review* are grouped below mainly according to their chief sponsors, the Assistant Secretaries of the U.S. Department of Energy. Research funded by other Federal agencies is listed under Work for Others.

ENERGY RESEARCH

Predictions of Atmospheric Carbon Dioxide: What Are the Uncertainties? (December 1983)
Detoxification of Environmental Copper in the Common Mussel (November 1983)
Evolution of the Magnetic Configuration for the MFTF-B (August 1983)
Tracking the Global Carbon Cycle (August 1983)

FOSSIL ENERGY

The Chemistry of Oil Shale Processing (January 1984)
Hydraulic Fracturing and Methane Recovery in Coalbeds (December 1983)
An Overview of Oil-Shale Retorting Processes (October 1983)

DEFENSE PROGRAMS

Inertial Fusion

Fusion Experiments Analysis Facility Completed (Brief—August 1983)

Military Application

Artificial Intelligence in Chemical Analysis (February 1984)
Providing Innovative System Monitoring and Reliability Assessment
Through Microengineering (February 1984)
Toward a Science of Welding (February 1984)
Modeling Dynamic Consolidation in Metal Powders (January 1984)
Abrasive Superpolishing of Metal Optics (December 1983)
Modeling EMP Damage to Semiconductors (November 1983)
Computer-Controlled Triple-Quadrupole Mass Spectrometry (October 1983)
Shocked Fluids at High Densities and Temperatures (October 1983)
Models of Nonequilibrium Nuclear Reactions (September 1983)
Reflection Seismology at the Nevada Test Site (September 1983)
Revealing Subsurface Structure with Borehole Gravimetry (September 1983)

WORK FOR OTHERS

Defense Advanced Research Projects Agency

Combustion Synthesis: A New Area of Research in Materials Science (January 1984)
Magnetic Compressors: High-Power Pulse Sources (August 1983)

Office of Naval Research

Compliant Surface Coatings to Reduce Turbulent Drag and Noise (November 1983)