269
2/15/68
leg

ORNL-4153
UC-48 - Biology and Medicine

THE MEASUREMENT OF CLINICAL RADIOACTIVITY

D. A. Ross and C. C. Harris

OAK RIDGE NATIONAL LABORATORY

operated by

UNION CARBIDE CORPORATION

for the

U.S. ATOMIC ENERGY COMMISSION

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

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.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Printed in the United States of America. Available from Clearinghouse for Federal
Scientific and Technical Information, National Bureau of Standards,
U.S. Department of Commerce, Springfield, Virginia 22151
Price: Printed Copy \$3.00; Microfiche \$0.65

LEGAL NOTICE

This report was prepared as an account of Government sponsored work. Neither the United States, nor the Commission, nor any person acting on behalf of the Commission:

- A. Makes any warranty or representation, expressed or implied, with respect to the accuracy, completeness, or usefulness of the information contained in this report, or that the use of any information, apparatus, method, or process disclosed in this report may not infringe privately owned rights; or
- B. Assumes any liabilities with respect to the use of, or for damages resulting from the use of any information, apparatus, method, or process disclosed in this report.

As used in the above, "person acting on behalf of the Commission" includes any employee or contractor of the Commission, or employee of such contractor, to the extent that such employee or contractor of the Commission, or employee of such contractor prepares, disseminates, or provides access to, any information pursuant to his employment or contract with the Commission, or his employment with such contractor.

ORNL-4153

MASTER

Contract No. W-7405-eng-26

THERMONUCLEAR DIVISION

THE MEASUREMENT OF CLINICAL RADIOACTIVITY

D. A. Ross and C. C. Harris*

FEBRUARY 1968

LEGAL NOTICE

This report was prepared as an account of Government sponsored work. Neither the United States, nor the Commission, nor any person acting on behalf of the Commission:

A. Makes any warranty or representation, expressed or implied, with respect to the accuracy, completeness, or usefulness of the information contained in this report, or that the use of any information, apparatus, method, or process disclosed in this report may not infringe privately owned rights; or

B. Assumes any liabilities with respect to the use of, or for damages resulting from the use of any information, apparatus, method, or process disclosed in this report.

As used in the above, "person acting on behalf of the Commission" includes any employee or contractor of the Commission, or employee of such contractor, to the extent that such employee or contractor of the Commission, or employee of such contractor prepares, disseminates, or provides access to, any information pursuant to his employment or contract with the Commission, or his employment with such contractor.

*Present address: Division of Nuclear Medicine, Duke University Medical Center, Durham, N.C. 27706

OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee
operated by
UNION CARBIDE CORPORATION
for the
U. S. ATOMIC ENERGY COMMISSION

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

PREFACE

This discussion of clinical gamma-ray spectrometry is addressed to physicians, x-ray technicians, and paramedical laboratory personnel who may need to operate pulse-height spectrometers with some understanding of what they're doing. Excellent accounts of pulse-height spectrometry are, of course, available, but they are nearly all written for physicists and engineers and are not likely to prove either readable or comprehensible to the medical people. The accompanying account makes little attempt to be scholarly in language or presentation. It is hoped that its informal style will make it more readable to the people who need it.

CONTENTS

INTRODUCTION: MEASURING RADIATION FROM A PATIENT.....	1
Emitted Rays Measure Grams	1
Counting by Comparison.....	2
Interaction, Absorption, and Deflection.....	2
Detector Geometry.....	4
Inverse-Square Law	5
Cosine Law	5
Effect of a Surrounding Medium	6
Patient Phantoms	6
Energy-Selective Counting	7
PULSE-HEIGHT SPECTROMETER	7
Basic Principle: Pulse Height Measures Gamma Energy.....	7
Overall Counting System	7
Detector	7
Processing Unit	8
Scintillation Crystal.....	8
Light Production.....	8
Intrinsic Crystal Efficiency	9
Photomultiplier Tube	9
Light Conversion	9
Electron Multiplication	9
High-Voltage Supply.....	10
Crystal-Phototube Unit.....	10
Reflector and Optical Coupling	10
Canning and Magnetic Shielding.....	11
Phototube Output Circuit.....	11
Integration	11
Pulse Shaping	11
Double Differentiation	12
Preamplifier	13
Voltage Amplifier	13
Pulse-Height Selector or Analyzer	14
Circuit Operation	14
Window Controls	15
Pulse Counting	15

THE PULSE-HEIGHT SPECTRUM AND ITS LIMITATIONS.....	16
Basic Nature of a Pulse-Height Spectrum.....	18
Resolution and Its Impairment.....	20
Measurement of Resolution	21
Further Distortion: Energy Escape from the Crystal	21
Further Distortion: Predetector Scattering.....	23
Minimizing Scattering Effects	24
Crystal Efficiency.....	25
Special Kinds of Escape	26
Organic Scintillators.....	27
MULTICHANNEL ANALYZERS	30
The Multiscaler	31
OPERATING POLICIES	32
E-Dial Calibration (Gain Adjustment)	32
Amplifier Nonlinearity	33
E-Zero Offset	33
Window Operation	33
Signal and Noise: the S/N Ratio.....	34
Count Rate and "Enclosed Area"	36
Window Opening and Count Rate	37
Window Opening and S/N Ratio	37
Window Opening and Statistics: S^2/B	38
Window Opening and Count-Rate Stability	40
Prevention of Drift	43
Instability of Window Opening	43
Asymmetrical Window Settings	44
Plotting a Pulse-Height Spectrum	44
Multiple-Nuclide Counting.....	45
Large Windows and High Energies	46
Threshold or "Integral" Counting	46
Operating Policies	47
Virtues and Vices	48
CONCLUSION.....	48
REFERENCES	49

THE MEASUREMENT OF CLINICAL RADIOACTIVITY

D. A. Ross and C. C. Harris

INTRODUCTION: MEASURING RADIATION FROM A PATIENT

Since radionuclides can be administered harmlessly to patients if the dose is kept small, they are coming into use more and more widely in medical research and diagnosis. Their major advantage is that the emitted radiation can be detectable even when the amount of the radionuclide present, by actual weight, is almost nothing. For example, iodine-131 weighs about 8 *millionths* of a gram per curie, and we need only a dozen or so millionths of a curie (μc) for a thyroid test. We could administer as much as 100 μc of carrier-free iodine-131 – an amount that emits 3 million useful gamma rays every second – with the assurance that we would be adding less than a billionth of a gram (10^{-9} g) to the body iodine “pool” and that the dose will therefore melt unobtrusively into the prevailing iodine cycle without disturbing it in the least. The feasibility of working with these infinitesimal quantities has provided new methods for the diagnostic and research laboratories.

Although these new toys are pretty, there is a catch: ordinarily it is no trick to *detect* the emitted radiation, but it may be difficult to *measure* it quantitatively, for there are booby traps lying in wait for the unwary, the naïve, and the wishful. The measurement of the radiation coming from a small sample is fairly easy, but if the source is buried in a patient, the problem of getting a good measurement may tax the experimenter's resources. The following discussion is offered to newcomers in the field in the hope that it may help them to avoid errors.

Emitted Rays Measure Grams

When we measure a nuclide through its radiation, we are making the tacit assumption that the num-

ber of rays emitted per second is proportional to the amount of the radionuclide present. This is a good assumption. The law of radioactive decay says that in each second the same fraction of the material then present will disintegrate. The fraction is characteristic for the particular nuclide: in the case of iodine-131, for example, it is one atom in a million. Expressed in mathematical terms, the law is:

$$\text{disintegration rate, } \frac{dN}{dt} = -N\lambda,$$

where N is the number of radioactive atoms present and λ is the “decay constant,” which has a characteristic value for each radionuclide. Thus if we can measure the disintegration rate and can look up λ in the tables,¹ we will have a measure of quantity. Furthermore, each radionuclide has a characteristic number of rays emitted per disintegration (it is 1:1 only exceptionally); so if we measure rays emitted per second, it will give us disintegrations per second, which in turn will give us the amount.

The difficulty is that when we set up a detector and scaler we don't count the rays emitted per second, for they leave the exploding atoms in all directions and many of them will miss the detector entirely. Moreover, if the “ray” is a photon it won't necessarily be counted even if it does reach the detector, for reasons that we will come to later. Consequently the *count rate* represents only a fraction of the *disintegration rate*, and unless we can relate the two the count rate won't tell us much about quantity.

¹If your table gives half-life instead of λ there is an easy conversion:

$$\lambda = \frac{0.693}{\text{half-life in seconds}}$$

Counting by Comparison

There is a dodge that circumvents this problem, and in medical work it is used all the time. We measure by comparison, counting first the unknown and then a standard containing a known amount of the same nuclide. Then, if the instrument is counting the same fraction of the generated radiation each time, we can set up a proportion:

$$\frac{\text{unknown } \mu\text{c}}{\text{standard } \mu\text{c}} = \frac{\text{count rate from unknown}}{\text{count rate from standard}},$$

and therefore

$$\begin{aligned} &\text{unknown } \mu\text{c (see footnote}^2\text{)} \\ &= \text{standard } \mu\text{c} \times \frac{\text{count rate from unknown}}{\text{count rate from standard}} \end{aligned}$$

With this comparative approach we don't have to know what percentage of the emitted radiation the system is counting, but **it must be the same percentage for the standard as for the unknown!** This may be difficult to achieve when the unknown is located inside a patient, for the fraction of the radiation that reaches the detector depends not only on the position of the source but on what surrounds the source: namely the patient's tissues, through which the outgoing radiation must pass. It is important for an operator to understand why this is so, since the validity of his work hinges on his success in coping with the tricks that radiation may play on him.

Interaction, Absorption, and Deflection

The only radiation that concerns an external detector is gamma radiation, for practically all

alpha or beta particles will be stopped before reaching the detector. For our purposes, gamma rays are best considered not as waves but as particles called *photons*: electrically neutral, submicroscopic bullets shot out of the metastable nuclei, traveling with the speed of light (= 1000 ft/ μ sec in a vacuum), and each carrying a tiny load of energy.³ These photons may interact with matter as they pass through it: they collide with atomic electrons and often knock them loose — hence the name *ionizing radiation*. The denser the material, the more of these collisions there will be; thus there are many in lead, fewer in bone, still fewer in the soft tissues, very few in air, and none in a vacuum. Whenever a collision occurs, the photon is either *deflected* or *destroyed*.

The various ways in which photons can interact with matter may be familiar to the reader, but it might do no harm to review them briefly here. When a photon collides with an electron, the latter acquires some of the photon's energy and comes flying out of its atom with a velocity that depends on the energy received. If the photon's energy is not much greater than the binding energy of the electron, all of the energy is likely to transfer to the electron. We call such a process a *photoelectric interaction*; the photon vanishes and a *photoelectron* emerges. If the photon's energy is somewhat higher, on the other hand, the electron is more likely to receive only part of it; in such a case there will be two products of the interaction: (1) a high-speed electron and (2) a secondary photon carrying less energy than its parent primary. This process is called *Compton scattering*, after the physicist who discovered its laws; the dislodged electron is a *Compton electron*, and the secondary photon is a scattered or *degraded photon*. The high-speed electrons produced in both types of interaction collide repeatedly with other atoms and produce ionizations of their own, thus dissipating the kinetic energy. Similarly the scattered

²In situations where the actual number of microcuries is not critical, the standard can merely be a known fraction of the dose given the patient, in which case the unknown is rated as a percentage of the administered dose:

% of dose in unknown =

$$\% \text{ of dose in standard} \times \frac{\text{count rate from unknown}}{\text{count rate from standard}}$$

³The amounts of energy are so small that the usual cgs unit, the dyne-centimeter or *erg*, is much too large to be practical. In atomic work, therefore, the basic unit is the electron volt (ev): 1 million electron volts = 1.6 millionths of an erg (1 Mev = 1.6×10^{-6} erg). For the energy range of medical interest, the most convenient unit is the kev (=1000 ev). Ordinary x rays and many of the gamma rays fall in the range 1–300 kev; other gamma rays and the more penetrating or "super-voltage" x rays have higher energies — e.g., the cobalt-60 gamma rays at about $1\frac{1}{4}$ Mev. In sharp contrast are the weak little photons that transmit visible light — e.g., orange light (600 m μ) at 2 ev.

photons may undergo further interactions and produce more ions. If the absorbing material is thick enough, it will absorb practically all of the liberated electrons and scattered photons, in which case all the energy contained in the primary photon is expended in ionizing the absorber. The greater the primary energy, the more ionized atoms there will be.

We should note that the photoelectric process involves *bound electrons*, especially those in the inner electron shells (*K*, *L*, etc.), and the interaction is most probable when the energy of the approaching photon is just above the binding energy of the electron in its path. The binding energy depends on the nuclear charge (= proton number, *Z*), and accordingly the probability of interaction, called the photoelectric cross section, depends very much on the element involved. This is much less true for a Compton interaction, which involves "*free*" electrons, namely those actually unattached or so weakly bound that the binding energy is negligible. The Compton cross section, therefore, depends mainly on the density of free electrons that the photon sees in its path, and apart from this it depends very little on the nature of the absorbing material.

There is another important difference between the two processes. Immediately after a photoelectric interaction the violated atom finds itself with one of its inner electrons missing, and an x ray is generated as the vacancy is filled in from the outer shells. The x ray will have approximately the energy of the electron shell where the vacancy occurred and thus will be characteristic of the atom.⁴ After a Compton interaction, on the other hand, there will be no x ray, for the dislodged electron had no binding energy worth mentioning.

We must also note one very important side effect: in a Compton collision the photon must bounce off at an angle if it is to transmit any of its energy to an electron.⁵ Thus in a scattering medium some of the photons will appear to go around corners, and only in a vacuum will they all travel in straight lines. These changes of direction bother the locating instruments, of

course, but they also bother the quantitative counting equipment. We will return to this shortly.

A third way in which a photon can interact with an absorbing medium is by the process of *pair production*. Part of the photon's energy is converted into matter, creating an electron-positron pair.⁶ This requires, for each member of the pair, 511 kev, which is Einstein's energy equivalent for an electron's resting mass. Pair production can never occur, therefore, unless the incident energy exceeds $2 \times 511 = 1022$ kev, and the process remains quantitatively unimportant until the energy reaches 2 Mev or more, which is infrequent in medical work. At any rate, a positron and an electron are created, using up 1022 kev, and the energy left over is carried away by the two particles in the form of energy of motion.

The electron goes on its way producing ions in the usual manner. The positron, however, is short-lived; its ionizing collisions quickly slow it down to a relative walk, whereupon it attracts an unattached electron in the neighborhood and the phenomenon of *annihilation* occurs. The two particles neutralize each other electrically and vanish, and in their place appear two new photons, which shoot off in opposite directions. Each has an energy of 511 kev, which just matches the resting-mass energies of the annihilated positron and electron. Accordingly, when a high-energy photon is absorbed by pair production, the end products are a number of ionized atoms, and two 511-kev *annihilation photons*, which leave the scene of their generation in opposite directions. In some of the scanning instruments ingenious use is made of this curious 180° relationship.

Note that here also the original photon and the resulting annihilation radiation will almost always have different directions.

Interactions Within the Patient. — It will be clear from the foregoing that when some organ in

⁵ Compton worked out the relationships thoroughly (2). He found that if *E* represents the primary energy (in kev), *E'* the energy of the scattered photon, and ϕ the angle through which it is deflected off course, then

$$E' = \frac{E}{1 + (E/511)(1 - \cos \phi)}$$

Thus if there is no deflection, $\phi = 0$, $\cos \phi = 1$, and there will be no change in energy.

⁶ A positron is a particle having the same mass as an electron, but the charge is positive instead of negative.

⁴ Actually there will be a shower of x rays, for each shell is filled in mainly from the shell just outside it. Following a *K* vacancy, therefore, the x-ray energies will mostly be *K - L*, *L - M*, and so on. The most important x ray is the one with the *K - L* energy; it is loosely called the "*K* x ray."

a patient contains radioactive material, the radiation coming out of the body will differ both in quality and quantity from that emitted by the exploding atoms. Many of the emerging photons will be degraded in energy and will not be traveling in their original directions. Some photons will not come out at all, having undergone a series of Compton collisions followed by a final photoelectric interaction in which the photon disappeared. Finally, of course, some photons will come out undegraded, having failed to collide with anything. It is important to recognize that only these last can provide a reliable guide to (1) the position of the source and (2) the energies characteristic of its decay process. So if you need to find out *where* a source is or *what* it is, you must try to look only at the undegraded photons.

If you also need to know *how much* is there, you may have a difficult problem on your hands, but it will be made easier if you can establish the energies of the photons that arrive at the detector. These are some of the reasons why gamma spectrometry is important in nuclear-medical work.

Interactions within the patient will have a bearing on the fraction of the generated radiation that reaches the detector, and the probability of interaction will vary with the nature and thickness of the tissues through which the radiation must pass. It will also depend on the photon energy. Thus the fraction received by the detector will differ from one radionuclide to another, from one patient to another, from one part of a patient to another part of the same patient, and from one direction of exit to some other direction. Since, in a quantitative measurement, the fraction received must be the same for the unknown as for the standard of comparison, one must be very careful how the standard is counted. It is important, therefore, for an operator fully to understand what may be going on between the source and the detector. Accordingly let us examine this question in some detail, first establishing the fundamental principles involved and then discussing the problems brought on by their operation in a practical counting situation.

Detector Geometry

For the moment let us abolish the patient and consider a small radioactive source surrounded only by air. Figure 1 shows such a source, pictured at the center of a spherical shell, with the detecting crystal on the surface. The photons

Fig. 1. Illustrating detector geometry (=geometrical efficiency = extrinsic detector efficiency). The source is imagined at the center of a spherical shell (radius R) with the crystal (radius r) resting on the surface of the sphere. The photons are emitted in all directions, and the fraction reaching the crystal will be:

$$\text{Geometry} = \frac{\text{area covered by crystal}}{\text{total area of shell}} \approx \frac{\pi r^2}{4\pi R^2}.$$

Often 99% of the radiation is lost. From (13), redrawn.

come out in all directions, and clearly many of them will miss the detector completely. For descriptive convenience let us classify them into "sheep" and "goats," the sheep being accepted into the detector's fold while the goats are cast into outer darkness, with weeping and gnashing of teeth. The teeth are the experimenter's, for if there are few sheep and many goats he will have to give more radionuclide to the patient in order to obtain a workable count rate. This involves more ionizations within the patient, and may precipitate a dust-up with the Radiation Safety Officer. At any rate, in Fig. 1 there will be a *cone of entry*, with its apex at the source and its base at the crystal, and this cone will contain

the sheep. In the illustration the sheep are heavily outnumbered by the goats, and we say that "the detector geometry is poor." Geometry is simply the fraction of the radiation that reaches the detector.

It is not hard to estimate just how poor it will be, for clearly the ratio of the sheep to the total emission will be (area of front face of crystal) ÷ (total area of spherical shell); namely $\pi r^2 / 4\pi R^2$, where r is the radius of the crystal and R that of the sphere. It is worth a moment to throw typical figures into this formula, for many people have little realization of how inefficient an external detector may be. Suppose the source is suspended in the air, 5 in. in front of a 2- by 2-in. cylindrical crystal (Fig. 2A). The calculation shows that only 1% of the radiation will find the detector — an ironical ninety-and-nine goats for every sheep.

Inverse-Square Law. — Five inches is a rather close working distance; what will happen if you back the detector off to 10 in. (Fig. 2A)? The area of the detector, πr^2 , will remain unchanged, but the area of the sphere, $4\pi R^2$, will be multiplied by 4. The fraction of photons in the cone will therefore fall to one-quarter of its former value. This is the basis of the *inverse-square law*, which states that the count rate varies inversely as the square of the source-to-detector distance. Triple the distance, it says, and you will divide the count rate by 9, and so on.

Cosine Law. — An alert reader will realize that this inverse-square argument is valid only if the angle of the cone is small;⁷ so for short distances and/or large crystals the inverse-square law gets into trouble. For example, if the distance is reduced to zero (Fig. 2D) the count rate will certainly not become infinite, as the law says it should; in fact the crystal will still let half of the radiation get away. A more reliable law, which holds good even when the angle of the cone is large, says that the geometrical efficiency will be $\frac{1}{2}(1 - \cos \theta)$, where θ (see Fig. 2A) is the "angle of generation" of the cone: namely, half the angle of the apex as seen from one side. Note that at zero distance, θ becomes 90° and its cosine is therefore zero; this law says, accordingly, that at zero distance the geometry will be 50%, which is correct. The cosine law is less convenient than the

Fig. 2. Three stages of detector geometry. (A) With the source $2\frac{1}{2}$ crystal diameters away, the geometry is down to 1%. Double the distance, and it becomes 0.25% — *inverse-square law*. (B) With the source against the crystal the geometry will be 50% — see *cosine law*. (C) With a well crystal the geometry approaches 100% (solid angle $\approx 4\pi$ steradians), but the sample must be small.

other⁸ but is more widely applicable. It holds wherever the detector presents a circular outline to the source, and that is about the only restriction.

⁸In ordinary practice θ will be obtainable not through its cosine but through its tangent, the ratio of detector radius to source distance. For convenience let us revise the formula to make it avoid θ altogether. And since it is awkward to measure a radius, let us use d = crystal diameter. R = source distance as before. Then

$$\text{geometry (\%)} = 50 - \frac{100 R}{\sqrt{4R^2 + d^2}}.$$

The sphere formula comes within 2% of the cosine law for source distances greater than three crystal diameters and within 5% for two diameters.

⁷Otherwise πr^2 , the area of the flat face of the crystal, will not provide a good approximation for the corresponding curved area on the surface of the sphere.

One can further improve the geometry of Fig. 2B by using the "well" principle, inserting the sample into a hole drilled part way into the crystal (Fig. 2C). Here there might be 95% sheep and 5% goats, so that a switch from Fig. 2A to Fig. 2C would improve the geometry nearly 100-fold. This excellent geometry is what gives a well counter its high sensitivity, in terms of counts per minute per microcurie. The drawback is that an ordinary well crystal can admit only small samples. Note in passing that the number of escaping photons will depend on how far down in the well the source is placed. Thus the volume of the sample will have a bearing on the geometry, and for comparisons to be valid, the volumes must be the same — unless a predetermined correction factor is applied. Because of leakage through the crystal, such a factor will depend on the energy.

Effect of a Surrounding Medium. — On the basis of this cone principle one can take almost any counting situation involving a small source and a cylindrical crystal and calculate the fraction of the total emitted radiation that will reach the crystal — as long as the source is in the air. When it is surrounded by a scattering medium, however, the pretty, straight-line pictures of Fig. 2 lose some of their value. Any photon that incurs a Compton collision will be deflected out of its original course, and thus what should have been a sheep may become a goat, as shown in Fig. 3. Alternatively a goat may become a sheep, and there are so many goats that this possibility is a very real one. Furthermore, a photon that starts off as a sheep may get killed in a photoelectric crash. In short, *absorption and scattering pervert the detector's geometry* and modify the percentage of the generated photons that will reach the crystal. Thus while the tidy diagrams of Fig. 2 are valuable in orienting our thinking about detectors in general, they can't be trusted too far in the patient-counting situation. Counts per second per microcurie can be one thing in one patient and something else in another. Or — which is not only worse but more probable — counts per second per microcurie may be one thing when you count a patient and something else when you count the standard representing a known fraction of the administered dose. Yet similar efficiency is essential if a quantitative comparison is to be valid, for the counter must register the same fraction of all disintegrations both times.

Fig. 3. If the source is contained in a patient, the detector geometry becomes disturbed. The "sheep" should stay within the cone subtended by the detector, while the "goats" fly off into outer space; but scattering can deflect a sheep out of the cone (top) or a goat into it (bottom).

Patient Phantoms. — The usual way out of this difficulty is to construct a phantom in which the standard of comparison will be counted, the phantom's job being, essentially, to duplicate the patient's disturbed geometry. When making thyroid measurements, for example, we use a neck phantom so contrived that the absorption and scattering in the phantom will closely imitate what happens in a human neck containing a radionuclide-loaded thyroid gland in the usual position. But the phantom's assignment is a difficult one. It might be able to imitate one particular neck with reasonable success, but in the clinic the test will involve many necks, and they will come in all sizes and shapes. Moreover, the gland itself is apt to vary in size, shape, and position, especially if it's sick. A neck phantom, therefore,

has too many problems, and would welcome anything that would ease the strain. If one could persuade the counting system to disregard the scattered radiation (goats in sheep's clothing, Fig. 3), the phantom could concentrate on imitating the missing sheep and thus do a better job. The scattered photons, we recall, are typically degraded in energy and so, in a sense, are painted red. Could one use this telltale earmark to sort them out and reject them?

Energy-Selective Counting. — A gamma-ray spectrometer is an instrument that will count photons with due regard to their energies and can be made to accept only those energies that the operator wants it to. Though it doesn't work by resonance, it behaves, in effect, like the tuning circuit in a radio, which can let in a desired signal and suppress all others. This feature permits a spectrometer operator, in many cases, to reject the degraded photons, whose geometry is under suspicion, and count only the undeflected primary photons. Note that we say "in many cases" rather than "always," for when the primary energy is low (20–130 keV) the energy change accompanying a Compton collision may be so small that the spectrometer can't distinguish between the primary and the degraded radiation.⁹ Given a reasonable primary energy, however, a spectrometer can reject much of the scatter, thus greatly enhancing the reliability of a comparison between a patient count and the standard count, the latter being made with a sample of the same radionuclide counted outside the patient in a realistic phantom. If the phantom could be an accurate facsimile, there would be no geometry problem; but humanity is variable and a phantom can't match everybody.

There are other advantages offered by energy-selective counting (see *Conclusion*, p. 48), and they make a clinically oriented spectrometer a valuable instrument generally as well as a necessity in many counting and scanning situations. Any measuring process, however, depends for its success on the collaboration of three ingredients: (1) the method, (2) the instrument, and (3) the operator. In this case the interaction between operator and instrument is unusually important,

for either can make a monkey out of the other. An operator will do a much better job if he understands what is going on in his instrument; accordingly let us inquire into its mechanism.

PULSE-HEIGHT SPECTROMETER

Basic Principle: Pulse Height Measures Gamma Energy

The goal that we are striving for is an answer to the question, "How many photons per second?" This will be delivered by some device that counts electrical pulses; namely, a scaler or (less desirably) a rate meter. Accordingly we must convert an incoming photon into an electrical pulse, but if we are to count the photons with due consideration for their energies ("How many photons per second in the energy band between p and q keV?"), the pulse must carry some kind of energy tag all the way from the detector to the point where the energy selection is made. The energy tag is the pulse's height (voltage). All through the following discussion of the spectrometer's mechanism, therefore, we will find the requirement that *pulse height should be proportional to gamma energy*.

Overall Counting System

A pulse-height spectrometer consists of several components, each of which contributes its bit to the business of counting a selected fraction of the photons that interact in the detector. The system has two main subdivisions: the *detector* and the *processing unit*. The detector is a "transducer": it accepts the incoming photons and produces their electrical analog, a stream of output pulses, much as a microphone accepts incoming sound waves and changes them into electrical waves. The processing unit is the analyzer: it measures and counts the pulses and rearranges their jumbled information into an intelligible, overall picture.

The components and their functions are as follows:

Detector (Fig. 4). — *Scintillation Crystal.* — Generates a flash of light, the *scintillation*, for each absorbed photon, the intensity of the flash being proportional to the amount of gamma energy absorbed. The crystal has a *shield* to keep out background radiation, and a *collimator* that permits it to "see" only in the desired direction.

⁹For example, it is clear from Compton's equation (footnote 5) that the most energy a 30-keV photon can lose in a Compton collision is 3 keV, and this presents the spectrometer with an impossible discrimination problem. Consult ref. 12 for further discussion (in English).

Photomultiplier (or Multiplying Phototube or simply Phototube). — Produces, in response to each scintillation, a burst of electrons whose number is proportional to the intensity of the flash. An important accessory for the phototube is the *high-voltage supply*, which provides suitable potentials for the tube's electrodes (Fig. 5).

Processing Unit. — Integrator. — Adds up the number of electrons delivered by the phototube, producing a pulse whose voltage is proportional to this charge. Integration is accomplished most simply by a capacitor (Fig. 5), but often a feedback amplifier is used.

Pulse Shaper (or Differentiator, or Clipper). — A device that cuts off the long tails on the pulses derived from the integrator, so that a new pulse can't stand on its predecessor's tail and thus look taller than it really is.

Voltage Amplifier. — Amplifies the shaped pulses from millivolts to volts, so that they can be sorted effectively. A *gain control* provides manual adjustment of the amplification, or voltage gain, so that pulse height may be related to gamma-ray energy by a convenient ratio. A *preamplifier* may contribute to the gain, but its main function is to provide enough power so that the pulses can be transmitted from the detector unit to the main chassis without incurring heavy losses in the connecting cable. The preamplifier may also perform the integrating function. Throughout the amplifying system, *linearity* between input and output is important; so also is *rapid recovery* from the effects of a large input pulse.

Pulse-Height Selector (or Differential Discriminator). — Provides an electronic window analogous to the slit in an optical spectroscope. This component generates an output signal for each input pulse whose tip falls within the window, and ignores all others. Sill height and window opening are adjustable: the *E dial* (or base) controls the sill height, while the ΔE dial (window opening) controls the spacing between sill and top (Fig. 8). For special purposes the sill and top may be independently adjustable.

Scaler. — Counts the pulses that come through the analyzer. Together with a *timing clock* this gives the average number per second or minute. Less satisfactorily, the counting rate can be indicated directly by a rate meter. The newer "digital rate meters," used for dynamic studies, give the effect of a succession of scalers. A

multichannel analyzer (p. 30) can do this job very neatly.

Recorder. — Although not an essential part of a spectrometer, recorders are used in many applications to record the average count rate continuously as a function of time.

Of the foregoing components, it is those in the detector that will need the most careful consideration, for the transformation of gamma photons into electrical pulses is not a perfect process, and unless the instrument's operator understands what is going on in the detector, he can be misled by the data that his machine turns out (Fig. 9; pp. 16 ff.). Accordingly we must consider the detecting components in some detail and then proceed to the analyzing unit.

Scintillation Crystal

As we have seen, ionizations occurring in the patient are a nuisance: they deposit potentially injurious electron energy, they steal this energy from the photon, and often they steal the whole photon. Ionizations produced in a detector, on the other hand, are a necessity, for there is no way of measuring or even detecting a gamma photon except through the ions that it produces. We do our best, accordingly, to encourage ion production in the detector. In particular we would like each photon to expend *all* its energy in the detector; any that escapes in a scattered photon (Fig. 11) will cause the spectrometer to underestimate the incoming energy.

Light Production. — Certain substances will generate a flash of light when they stop a high-speed electron. Through the three processes discussed earlier, an incoming gamma photon produces high-speed electrons, and these rip through the nearby material, knocking other electrons out of their orbits. The orbital vacancies attract unbound or loosely bound electrons, and as each of these falls into a vacant orbit, thus losing potential energy, a photon of equivalent energy is released. The resulting multitude of practically simultaneous, tiny flashes of light is the *scintillation*.

In clinical service most of the detecting crystals are made of sodium iodide, activated with a trace of thallium. The abbreviation is "NaI(Tl)," but ordinarily the thallium is not mentioned, since everyone understands that it is there. Sodium iodide is dense (3.67 g/cc), and the iodine's proton number of 53 gives it a good photoelectric

cross section; so incoming photons have a good chance of interacting one way or another. Besides having high stopping power, this material produces generous amounts of light to which the crystal itself is transparent (as it must be), and the blue-violet light is well suited to many good phototubes. Sodium iodide is a medium-speed scintillator,¹⁰ which means that whereas it is fast enough to permit high counting rates, it does not unduly complicate the design of the processing circuitry.

One of the most valuable properties of sodium iodide is that the total light in the scintillation is fairly closely proportional to the energy deposited in the crystal by the photon. This is because the higher the incoming energy, the more ions it will produce, the more numerous the little subflashes will be as the atoms deionize, and therefore the more total light there will be in the scintillation. This is important, because if the entering photon can be persuaded to deposit all of its energy in the crystal, the amount of light in the scintillation will then provide a measure of the photon's energy.

Intrinsic Crystal Efficiency. — It has been pointed out that in most detecting situations many of the photons emitted from the source will miss the detector entirely, and the fraction that reaches the detector is known as the *geometry*. We must also note that not all of the rays actually incident upon the crystal will produce ions; some will go straight through, leaving the detector unaware of their existence. If we know the thickness of the crystal and the half-layer of sodium iodide for the energy concerned, we can calculate, approximately, the percentage of the incident rays that will be lost in this way by "transparency." A measured figure is better, however, because a photon arriving near the outer edge of the crystal may not see the whole crystal thickness in its path, and so tends to cut through the corner and be lost. An actual measurement, therefore, delivers a better efficiency figure, defined as follows:

$$\text{intrinsic total efficiency} = \frac{\text{photons that cause scintillations}}{\text{all incident photons}}$$

¹⁰In sodium iodide, light production starts as the gamma ray interacts. A quarter-microsecond later, 63% of the total light, on the average, has emerged. At the end of 1 μ sec, practically all the light has been emitted (see Fig. 6B).

This figure will depend, of course, on the energy of the incoming radiation. We will explain later (pp. 21–25) why we say *total* efficiency; there's a good reason.

Photomultiplier Tube

Light Conversion. — The next step is to convert the scintillation into an electrical pulse. If we allow light (especially violet light) to strike a substance whose outer electrons are only weakly bound, some of the electrons will be knocked loose. This is the photoelectric process in the ordinary sense, referring to visible or near-visible light. The detached electrons can be pulled away by the electrostatic attraction of a positively charged conductor located nearby, and thus a current is set up. Figure 5 shows how this principle is put to work in the spectrometer. A photomultiplier tube is used: this is a vacuum tube with a glass envelope that permits the light to enter. Deposited on the inside of the entering window is the photocathode, a very thin layer of a compound (e.g., cesium and antimony) whose electrons are easily dislodged. Near it is a structure called a dynode, charged about 300 v positive to the photocathode. Electrons liberated from the photocathode are therefore pulled toward the dynode, and the result is an electron current. If the illumination is very brief, as it is with a sodium iodide scintillation, a burst of electrons will leave the photocathode and the tube will then become quiescent again.¹¹ Understandably, the more intense the scintillation, the more electrons the burst will contain.

Electron Multiplication. — If the gamma photon is low in energy, so that the scintillation is weak, there might be no more than a dozen or so electrons knocked loose from the photocathode, so clearly the system will need all kinds of amplification before a scintillation becomes useful. The first part of the amplification is carried out inside the multiplier tube, and the method is to use a series of ten or more dynodes, each about 100 v more positive than the one before it (Fig. 5). The first dynode is so placed, and the voltage it carries is so chosen, that the liberated electrons are accelerated toward it at high speed, and

¹¹Actually there is a low-grade *dark current*, which need not concern us here.

they splash into its specially prepared surface so forcibly that other electrons are knocked loose. The augmented cloud of electrons now comes under the attractive influence of the second dynode; they speed toward it, splash, and multiply their number again. At each dynode the number of departing electrons is three to five times the number that arrived, so the amplification achieved by the string of dynodes can be enormous. With ten dynodes, for example, $4^{10} \approx 1$ million.

Beyond the last dynode there is a still more positive collecting plate, the *anode*. The arrival of the cloud of electrons causes the anode's potential to jump a few millivolts in the negative direction, thus providing a pulse that can be amplified and counted.

These are the processes of light sensing and electron multiplication that give these tubes the name of "photomultiplier." We loosely call them "phototubes" or "multiplier tubes" for short, though clearly neither term is accurate.

High-Voltage Supply. — The violence of the electron splash at the surface of a dynode depends upon how fast the electrons are traveling when they strike, and this depends in turn on the strength of the electrostatic field responsible for accelerating them. The electron multiplication or gain is therefore dependent on the voltage difference between dynodes, and in fact the gain is roughly proportional to the voltage difference. The dynodes derive their potentials from a string of voltage-dividing resistors whose high end is connected to a dc voltage supply (Fig. 5). If the high voltage should increase by 1%, the voltage on each dynode will increase by 1%, and therefore the electron gain at each stage will also increase by about 1%. This means that the overall gain will increase by something like 10% (since $1.01^{10} \approx 1.10$), and this is why the high-voltage supply for a photomultiplier tube must be very carefully stabilized if constant gain is to be maintained. Constant gain is important because the pulses are to be accepted or rejected, farther down the line, on the basis of pulse height. In these spectrometers, then, the demands on the high-voltage supply are particularly exacting, and special care must be taken in designing the regulating circuits. Some of them may need a warmup time of an hour or more to allow them to reach temperature equilibrium.

Crystal-Phototube Unit

The crystal and phototube must work very closely together, and each has to suit the other. For example, the photocathode must be sensitive to the kind of light produced by the scintillator; again, if the scintillator is of the type that produces less light than usual (e.g., some of the organic scintillators), the phototube must have an especially low noise level, and so forth. A recent trend has been to procure the crystal and phototube as an integral unit, assembled by the manufacturer and sealed together permanently in his shop — a process that requires a good deal of skill. These units have generally worked out very well.

Reflector and Optical Coupling. — Figure 4 shows how the components of a typical detector are arranged. The crystal must be optically coupled to the sensitive end of the photomultiplier tube, using a cement that transmits ultraviolet light well. There must also be an approximate match of refractive index between the cement and

ORNL DWG. 65-12644

Fig. 4. Typical arrangement of the components of a sodium iodide detector.

the surfaces in contact with it, to minimize the reflection of light back into the crystal. Care must be taken to exclude all air bubbles, for they would provide reflecting surfaces. This *exit surface* of the crystal, however, is the only one where the light should get out; at all other surfaces it should reflect, for if some of it fails to reach the phototube the latter will report, falsely, a substandard scintillation. Current practice is to have only the exit surface highly polished; the other surfaces are rough-ground, and between them and the protective can there is a brilliantly white reflecting powder such as alpha-alumina (aluminum oxide). This causes most of the light to be reflected back and encourages it to leave properly through the exit surface. The internal path length will vary, however, so the crystal must be perfectly clear or the phototube will not always receive the full amount of light.

Canning and Magnetic Shielding. — The crystal and its reflecting powder are enclosed in a thin aluminum can that lets gamma photons in easily but excludes room light (Fig. 4). Note that in a low-energy detector the can must be unusually thin.¹² It must keep water vapor away from the sodium iodide, which is deplorably hygroscopic. Often a magnetic shield partially surrounds the can (Fig. 4), since stray magnetic fields influence the electron paths between the various elements inside the photomultiplier tube and thus can change the electron gain. All it may take is a screw-driver (many tools are magnetized these days) carelessly left near the phototube. A weak magnetic field could be tolerable if it were *constant*, but this may not be so in a scanner, where the detector moves around, or in a test where the detector is switched from a vertical to a horizontal orientation between the counts for the unknown and the standard. Again, a power transformer (particularly one of the voltage-regulating type) can generate a strong 60-cycle magnetic flux. The best defense against these sources of error is good magnetic shielding.

Phototube Output Circuit

So far, then, the arriving photon has caused a swarm of electrons to land on the phototube's

anode. We need to know how many there were in the swarm, for with a constant photomultiplier gain of (say) half a million, the number of electrons in the swarm tells us how many left the photocathode, and that tells us how bright the scintillation was, and that tells us the energy of the photon that started the chain of events going.

Each electron carries a minute electrical charge; it takes $6\frac{1}{4}$ billion billion of them to make up one coulomb.¹³ When we measure an electrical charge, therefore, we are really counting electrons. Let us see if we can measure the charge arriving at the phototube's anode.

Integration. — If the anode is tied directly to the high-voltage supply, the arriving electrons will quickly disappear into the supply, leaving no useful signal. One could make them generate a voltage by placing a resistance in their path, but this would tell us only the *rate* of flow (current), not the total charge in the swarm. A better plan is to use a capacitor (condenser), which is a kind of storage tank for electrons. Just as the pressure in an oxygen tank can tell you how much oxygen is there, so the voltage (= electrical pressure) between the plates of a capacitor measures the impressed charge — in other words, the number of electrons. Taking advantage of this, we place a condenser in the phototube's anode line (Fig. 5), and lead off the voltage surge to the amplifying system. Each time a swarm of electrons arrives, the voltage jumps a few millivolts in the negative direction, and the size of this "step" is a measure of the number of electrons. Thus through a series of proportionalities, the step voltage measures gamma energy.

For reasons of electronic design, it is not always expedient to use the simple capacitive method to add up the electrons in the swarm; often this function is combined with that of the preamplifier by using a feedback trick. The end result, however, is the same.

Pulse Shaping

A stream of incoming photons could cause a series of negative steps at the phototube's anode, the voltage getting lower each time. If this were allowed to happen, the tube would soon stop functioning; so we let the anode's electrons drain

¹²The can may be made of beryllium, which, because of its very weak absorbing power, stops few photons even at low energies.

¹³More accurately, 6.242×10^{18} electrons per coulomb.

Fig. 5. The crystal-photomultiplier unit develops an electrical pulse for each interacting photon. The blue-violet light of the scintillation liberates electrons from the photocathode. These are multiplied enormously while cascading through a series of dynodes to the final collecting plate, or anode. The electrodes receive their potentials from a resistor string powered from the system's high-voltage supply. Voltages shown are approximate only.

away into the high-voltage supply through a resistance placed in parallel with the integrating capacitor (Fig. 5). This runoff must not be too fast or it will hamper the development of the step voltage, which takes one or two microseconds to reach completion (Fig. 6B). One therefore proportions the resistance and capacity in such a way that the step will be followed by a long "tail" that returns the system to equilibrium, as shown in Fig. 6A.

There is one difficulty, however. All the information we want is contained in the step; the tail is a handicap, for if another pulse comes along soon enough, it can stand on its predecessor's tail and thus look taller than it deserves to. We would do better to amputate the tail just after the step, leaving a zero base line ready for the next pulse.

This is done by a circuit called a clipper or differentiator. It may be built from resistors and capacitors, but if large pulses are anticipated (e.g., from cosmic rays hitting a large crystal),

it is better to use a device called a *delay line* (Fig. 7). The step-and-tail pulse is applied to the open end of the line, whose far end is carefully short-circuited. The pulse travels down the line, reflects from the shorted end with reverse polarity, and returns to the input end about a microsecond after it entered. During the delay the original step is free to drive the amplifier, but after one microsecond the reflected pulse cancels out the tail. The result is a brief pulse whose leading edge is the original step and whose trailing edge is a mirror-image step that restores a clean base line (Fig. 7, insert, and Fig. 6B).

Figuratively speaking, then, we have caught the anode's swarm of electrons in a bucket, counted them, and emptied the bucket to get ready for the next swarm — all in two or three millionths of a second.

Double Differentiation. — It is almost always desirable to differentiate the integrator's output twice instead of once. This produces a double-sided pulse with approximately equal areas above and below the base line. Such pulse shaping helps the following amplifier cope with overloads and sudden changes in pulse rate. In addition, it improves the performance of pulse-sorting circuits by eliminating shifts due to high counting rates. Single differentiation is used primarily for simplicity and economy; the problems mentioned above must then be dealt with in other ways.

Fig. 6. (A) The electrical pulse delivered by the phototube has a steeply rising step followed by a more slowly declining tail. We amputate the tail, since it carries no information and can deceive the analyzing circuits (see section on *Delay line*). Normally this pulse is negative in sign. (B) The step, viewed with expanded time scale. The numerous subflashes that make up the scintillation are over in a microsecond or so. The dashed curve shows what happens when the delay line cancels out the tail (see insert, Fig. 7). Redrawn from ORINS-30 (6).

ORNL DWG. 65-12639

Fig. 7. Showing the detecting, pulse-shaping, and amplifying components of a typical pulse-height analyzer. Redrawn from ORINS-30 (6).

Preamplifier

Since a movable detector is desirable, there will usually be several feet of electrical cable between the detector and the main electronic unit. The capacities in such a cable can attenuate a phototube's output pulses considerably, so an operator who tries to drive a long coaxial cable directly from the phototube is, figuratively speaking, trying to drive a nail with a sponge.¹⁴ The preamplifier solves the problem. It is placed on the detector side of the connecting line (Fig. 7), and its chief function is to prevent losses in the line by putting adequate current into a pulse at the detector end. The pulse then reaches the main chassis in good health.

Note in passing that the delay line's driving circuit (which is critical) includes the preamplifier and the interconnecting cable, and this makes their characteristics important for effective delay-line performance. For this reason the detector

¹⁴ Ancient proverb: You can't drive a nail with a sponge, no matter how hard you soak it.

and the main chassis must be properly matched in order to work in harmony, and it is advisable to procure them from the same manufacturer. He who interchanges components between manufacturers does so at his own peril.

The preamplifier, in spite of its name, may or may not increase the actual *voltage* of a pulse, depending on the overall design. Sometimes the preamplifier is asked to perform the integrating function, in which case it is called a *charge-sensitive preamplifier*.

A typical arrangement is shown in Fig. 7. In this configuration, integration is furnished by a capacitor at the input of the preamplifier, and differentiation, by delay-line clipping, at the input of the main amplifier.¹⁵

The high-voltage supply is not a functional part of the preamplifier, but often it is most practical to feed the high voltage in through the preamplifier to the string of voltage-dividing resistors (Fig. 5) mounted on or near the phototube's socket.

Voltage Amplifier

The pulses always need considerable amplification, for after integration and clipping, and even some preamplification, they are typically only a few millivolts in height. Since pulse height must measure energy, the amplification must be linear over a wide range of input heights: that is, large and small pulses must have the same amplification factor, up to some reasonable limit. In vacuum-tube systems the sorting circuits usually operate on pulses up to 100 v in height; in transistorized systems a typical figure would be 10 v. The former therefore have voltage gains in the range 3000 to 10,000 and the latter about a tenth of this.

Another requirement is that the amplifier must recover quickly after overload: that is, after receiving an input pulse too big to be amplified linearly. Moreover, it must not put out false pulses during recovery. The overload require-

¹⁵ In the language of nuclear instrumentation, the term *linear amplifier* includes the integrator, the clipper, and the voltage amplifier. This system has the definite requirement that it must produce a voltage pulse whose height is proportional to the total charge in an input pulse. Any preamplifier used must fit properly into this arrangement of functions. This system is not the same as a "linear amplifier" used in other applications.

ment is necessary because cosmic rays carry very high energies and therefore make bright scintillations, especially in a large crystal. It is important, therefore, that the amplifier recover quickly and cheerfully after a severe input jolt. Ordinary amplifiers often sulk for a short time after being abused.

The voltage amplifier is usually built in two or more sections, each provided with generous negative feedback. The feedback helps to keep the characteristic linear, and it also stabilizes the gain from day to day.

The system's voltage gain (= amplification) can be controlled by adjusting the phototube's supply voltage (p. 10), but it is often more convenient to control the gain at the amplifier. This can be done by placing a voltage divider between two amplifier sections, as shown in Fig. 7. The control usually works in steps, each of which cuts the gain in half; thus a four-step control would provide full gain and one half, one fourth, and one eighth of full gain. The 1:2 ratios are not intended to be precise, so the divider is usually made of low-cost 5% resistors. This is tolerable because the dial readings are not to be trusted anyway, and there must always be an independent check on the overall gain. The method of gain adjustment, using a gamma source of known energy, is described in the section on operation, p. 32.

Pulse-Height Selector or Analyzer

The pulses from the linear amplifier are large enough to be presented to the device that decides which are to be counted and which are not. Because of the linear nature of the system, the peak height of each pulse is a measure of energy deposited in the crystal by the incoming photon. So when these pulses are sorted out by height, it is the equivalent of sorting the photons according to their energies. One would like to be able to pick out any band of gamma energies at will and count the photons incident per second within that band, for in that way one could count the honest photons and shut out those that carry misleading information. If it is desired to plot a gamma-ray spectrum, one may select a narrow pass band and sweep it slowly over the whole range of energies, recording the counting rate continuously. This is similar to the plotting of a visible spectrum, in which a narrow slit is moved through the various

colors of the rainbow, recording intensity as the sweep proceeds.

To understand how the pulse-height selector performs its job, it might be helpful to consider an analogy. Suppose we have a large herd of giraffes and are interested in learning how tall they are: how many are between 9 and 10 ft tall, how many between 10 and 11 ft, and so on. This could be done by providing a high window with an observer looking out through it from a position well back inside the room. The giraffes could then be driven past the window in single file, the observer being instructed to count all those whose ears he could see, but no others. This would tell us how many giraffes fell in the particular height category defined by the sill and top of the window. For other height categories one could move the window and observer up or down and count again. The categories could be made coarse or fine by increasing the window's sill-to-top opening or closing it down.

A pulse-height analyzer operates on the same principle. If we want to count all pulses whose maximum heights are between 60 and 70 v, we build a kind of electronic window with controls by which we can set the window sill at 60 v and provide a sill-to-top opening of 10 v. This window, following our "instructions," will ignore all pulses whose peak voltage falls below 60 v or above 70. The 10-v band of pulse heights will correspond, of course, to a certain energy band among the incoming photons (Fig. 21A); so by selecting a windowsill level and a window opening we can, within reason, select any group of gamma energies in which we may be interested. The whole spectrum may be obtained by setting a narrow window to a succession of adjacent locations and taking a count for each; alternatively the continuous-recording method mentioned above may be used. Examples of spectra taken point by point with a single window may be seen in Figs. 12, 14, and 16, while those in Figs. 21-23, 25, 29, and others were continuously recorded. The spectra in Figs. 17, 18, and 20 were made with multichannel analyzers (see p. 30), in which a couple of hundred windows can accumulate counts concurrently.

Circuit Operation. — Let us examine the basic principles of operation without getting bogged down in electronic details. The fundamental element is a "trigger" circuit or discriminator, which will fire only if the voltage arriving at its

input reaches a preset threshold; short of that level the circuit remains quiescent.¹⁶ Like the switch that turns on the instrument's power, it has no in-between states: either it fires completely or not at all. The discriminator therefore passes all pulses with heights at or above the threshold level and ignores all the smaller ones.

A single discriminator provides only a sill for the window; to get the top takes a bit more doing. Going back to the giraffes, we imagine that the observer we have posted inside the room is rather stupid: he can be trusted to shout "count!" if he sees a giraffe's ears, but we can't trust him to keep quiet if he sees the neck but not the ears. So we must employ two stupid observers instead of one. Each operates on a simple window-sill basis; each is instructed to shout if he sees *any* part of a giraffe. But we give them separate window sills to sight over, one higher than the other. We then post a "checker" to tally the results, and instruct him to reject all giraffes seen by the upper observer, who is thus given a kind of veto power over the lower observer. Under such a policy the only giraffes counted will be those whose heights were between the two observers' window-sill levels, and thus we achieve the effect of a window with both sill and top by using two simple-minded threshold devices with an appropriate difference in level between them. In the pulse-height selector (Fig. 8) we do a similar thing: we use two discriminators with one threshold set higher than the other and with their out-

puts subject to review by a veto or "anticoincidence" circuit. The veto circuit passes the output of the lower discriminator only if there is no simultaneous output from the upper discriminator. In this way any scaler or rate meter connected to the output of the veto circuit counts only those pulses whose peaks fall between the thresholds of the upper and lower discriminators — namely, "fall within the window."

Window Controls. — The panel dial that controls the lower discriminator, or window sill, is called the "*E* dial" (*E* for energy), "base," "lower level," etc. The window top is usually tied to the sill so that the two move up or down together when the *E* dial is turned; the spacing between them then has a panel control called the " ΔE dial," "window opening," "slit width," etc. This arrangement is useful for finding the regions of maximum count rate that indicate the peaks in a spectrum, and also for plotting the spectrum. The constant window width ensures that we see a similar energy band no matter where the window sill is set. For special purposes the upper discriminator may have an independent control.

Pulse Counting

The information is now in the form of a stream of pulses coming through the window, with random time intervals between them. The veto circuit makes these pulses all substantially the same size, for we have already made the energy selection, and the proportional relationship between pulse height and energy is no longer needed. The remaining problem is to count the pulses over a known time in order to obtain the average rate per second or per minute. This is normally done by an ordinary scaler and timer (Fig. 8), and readers are undoubtedly acquainted with these devices already. Here we will merely say a few words about the two competing approaches: *preset time* and *preset count*.

The most direct way to measure counts per minute is to set up a timer for one minute and let the count run until the clock shuts it off — hence the name "preset time." The only concern is to accumulate enough counts to provide the desired statistical validity, and the operator must

Fig. 8. The "electronic window" and the pulse-counting components of a pulse-height analyzer. Redrawn from ORINS-30 (6).

¹⁶ Compare the *all-or-nothing law* of nerve and muscle.

lengthen the preselected time suitably if the count rate is low. Since this involves guesswork, some of his results will have reasonable reliability, others will have more than necessary, and others will have too little. Result: inefficient utilization of the available counting time, and since most laboratories are under pressure to get out their data quickly, a remedy would be welcome.

"Preset count" was designed as the answer. Here it is a count that shuts off the clock, instead of the other way around. Statistical theory says that

$$\text{fractional standard deviation (\%)} \approx \frac{100}{\sqrt{N}},$$

where N is number of counts accumulated. If you want about 3% accuracy in a count, therefore, you must aim for 1000 counts or more, letting the clock measure whatever time is required to bring the count up to the selected figure. You then divide the count by the indicated time to get count rate. Some of the samples will need a long counting time, while others will need less; but the accuracy will be about the same for all. If you feel you need 1% accuracy, you preset for 10,000 counts, and so on.

There are two drawbacks to preset counts. One is the practical nuisance of dividing by a time that won't come out even. This nuisance is considerable if the clock is calibrated in minutes and seconds only. The other failing is that you can preselect the statistical accuracy only for a gross count, which is what shuts off the clock, and not for the net count, which involves a background component as well. It is the accuracy of the net count, of course, that you are concerned about. We will deal with this question more fully in the section on operating policies (p. 38). Because of these drawbacks most operators are satisfied with a preset-time system, although sometimes satisfaction is possible only through ignorance of the errors being made. This should not be; it is easy for an operator to estimate his probable standard deviation (S.D.) (p. 39), and a careful worker will not only do this but will actually measure the S.D. from time to time, by making 20 or 30 repeat counts on the same sample, to see whether his system is behaving according to the book. If it is not, he must smoke out the gremlin that is spoiling his results — instrument malfunction, background variability, etc.

A measured standard deviation is most readily found by the formula

$$\text{S.D.} = \sqrt{\frac{n \cdot \Sigma(x^2) - (\Sigma x)^2}{n(n-1)}},$$

where n is the number of observations (say 30), and x is the value of an individual observation. You will need the sum of the individual readings, Σx , and also the sum of their squares, $\Sigma(x^2)$. A multiplying calculator is a great help.

A third way to obtain a count rate is by means of a rate meter, although here it is difficult for the operator to read the indicating pointer with acceptable accuracy. The pointer wavers about unsteadily, and one has no choice but to watch it for a time and then make a guess as to its average position. If the rate meter's time constant ($= RC$) is made long, the reading will be less erratic, but it will also take a longer time to reach equilibrium, and — human impatience being what it is — the technician is tempted to read it before it has properly settled down. In general, rate meters are a pestilence, and for most purposes we prefer the scaler-timer combination.

The above refers to the use of a rate meter to obtain a static count rate. If a continuous record of a changing count rate is needed, the relatively low-cost conventional rate meter can do a fair job if the rate is not changing too rapidly, but some distortion is always present. It can be minimized by intelligent handling of the time constant. Reference (7) compares several methods of recording a changing count rate and discusses the pros and cons for each.

THE PULSE-HEIGHT SPECTRUM AND ITS LIMITATIONS

The rainbow is the visible spectrum, an array of colors derived from sunlight, ranging from red at the low-energy end ($\sim 1\frac{1}{2}$ electron volts) to violet at the high-energy end (~ 3 ev). The eye gives only a rough idea of the physical intensities of the various colors, and a much better guide is a graph that plots intensity (vertical axis) against energy (horizontal axis). Since "intensity" can mean different things to different people, we had better be more specific: we want a graph that answers the question, "Of all the visible photons reaching the eye, what are the percentages for

red, yellow, green, and so on?" — bearing in mind that we will plot them against energy, which is physical and objective, rather than against color, which is in the observer's mind and is much harder to deal with. This kind of graph is an *optical spectrum*, and it will have a characteristic contour for each kind of colored light.

A gamma-ray spectrum is also a photon spectrum, similar to its optical counterpart except that the energies involved are thousands of times higher. Each gamma emitter has a characteristic spectrum for its radiation: a plot that answers the question, "Of all the gamma photons emitted by these dis-

integrating atoms, what percentage have such-and-such an energy, what percentage have this other, ...?" and so on until all the emitted photons are accounted for. A few radionuclides emit only one kind of gamma photon; many more emit several kinds. Two familiar examples are shown at the top of Fig. 9. Practically 100% of the chromium-51 gamma emission shows a single energy; 322 kev. Iodine-131 emits gamma photons with five principal energies and other less important ones. In Fig. 9 the location of a vertical line gives the energy of the photon it represents, while the height of the line gives the abundance as a per-

ORNL DWG. 65-12633

Fig. 9. Top: *Emission spectra* for chromium-51 (left) and iodine-131 (right). Each line's position gives the energy for the photon in question, and the height of the line (with figure above) gives the abundance as a percentage of all gammas emitted. Note that this is not per 100 *disintegrations*; chromium-51, for example, emits this photon only one-tenth of the time, but practically 100% of the gammas have this energy. Bottom: *Pulse-height spectra* for the above nuclides, NaI crystal, 2 x 2 in.; scattering materials near source as noted. Redrawn from ORINS-30 (6).

centage of all photons emitted. In contrast with beta emission, it is characteristic of gamma emission that the energies are discrete and discontinuous; thus the true emission spectrum is a series of lines with nothing in between, not a continuous graph as in a beta spectrum.

Below these line spectra in Fig. 9 are pulse-height spectra for the same two nuclides, and it is clear at once that the picture has become distorted. Each gamma ray is now represented by a hump instead of a line, and there may be a broad hump, or even a plateau, that has no counterpart at all in the line spectrum. The distortion doesn't make pulse-height spectra worthless, but it does require us to interpret them with insight and care. This section of the account is concerned with why a pulse-height spectrum looks the way it does.

Basic Nature of a Pulse-Height Spectrum

We have tried to design the detector in such a way that the height (voltage) of the phototube's output pulse will be proportional to the energy of the gamma photon that caused the scintillation. If this effort succeeds, we can sort the pulses in terms of height, and this will be the equivalent of sorting the photons in terms of energy. The original question, "What percent of the incoming photons have an energy of (say) 500 kev?" will then become the analogous question, "What percent of the pulses have a height of (for example) 50 v?"

Just a minute: do we mean *precisely* 50 v – say plus or minus a thousandth of a percent? Of course not, for you might wait a long time before observing even one pulse that could meet such an exacting specification. One must be reasonable about this sort of thing; 50 v plus or minus half a volt would be more realistic. So we set about sorting the pulses into a series of categories, each being (in this example) one volt wide, and when the results are plotted out, the ordinates will tell us what percent of the pulses will be found in each one-volt category.

Such a plot is called a *distribution*. We can discuss its properties more easily if we take a more down-to-earth example; so let us consider the distribution, in terms of height, of the giraffes in a certain herd. Suppose we select one-foot categories and ask how many giraffes are between 10 and 11 ft tall, how many between 11 and 12 ft,

and so on. The plotted results could be called a giraffe-height spectrum, and it might look something like Fig. 10, top, where each bar shows how many giraffes will be found in the height range covered by the bar. These same data could have been used to plot the distribution by coarser intervals, say 10 to 12 ft, 12 to 14, etc., but since this would approximately double the height of each bar (and halve the number of bars), we must state the interval used or the vertical readings won't mean much. If we should be interested in all the giraffes whose heights lie between two wide limits, say 12 to 17 ft, we would simply add up the numbers indicated by all the bars included in this range. In doing so we would be adding sub-populations to obtain a population. To get the total number of giraffes investigated, we would add up the numbers indicated by all the bars. It

Fig. 10. Top: A giraffe-height spectrum: a bar graph (*histogram*) showing the distribution of giraffes into a series of 1-ft height categories. Bottom: The same data as above, plotted as a continuous curve. With linear coordinates, the area under the curve gives the total number of giraffes in the population studied.

is easy to see, therefore, that the shaded area in the bar graph represents the whole population, and that if the data had been plotted as in Fig. 10, bottom, using points instead of bars, the total population would be represented by the area lying under the curve.

There are a couple of important comments to make about these distribution curves. The first is that if you make the bars too wide you will obscure the contained information, for its ups and downs, slurs, notches, and flat places may be averaged out of sight. So the curve becomes more informative as the intervals are narrowed. However, a point will be reached beyond which further narrowing won't improve matters, for the details of the curve's shape will already be established. Further narrowing, moreover, not only does no good, but may actually do harm, for with very narrow bars the number of individuals in a bar gets so small that the figures become statistically ragged. Consider, for example, a giraffe-height spectrum for which we have used height intervals only $\frac{1}{8}$ in. wide — supposing that you can figure out how to measure a giraffe's height to within an eighth of an inch. There might well be 3 giraffes in one interval, 17 in the one above it, 9 in the one below, and so forth, which would make the curve very jagged. In collecting data for a distribution curve, therefore, one must steer a middle course between intervals so narrow that the curve is impossibly rough and intervals so wide that important details are concealed.

We note that statistical roughness can be reduced by studying a larger total population, although the additional time and effort can become expensive. A proposal to study all the giraffes in Africa is a pure pipe dream.

It is clear that if you use different height intervals you will get different giraffe-height spectra from a study of the same herd. The plots of Fig. 10 would be high above the base line if the intervals used were wide (because each bar would then contain more individuals), and, conversely, low if the intervals were narrow. The area under the curve (= herd population) would therefore look different in each case, and it is a pity that the plots can't be alike if they are really delivering the same information. A simple trick, however, can make them comparable: instead of plotting the actual number of individuals in a given bar, we can plot the number of individuals *per unit interval width*, for example, per inch of height in

the case of the giraffes. We might run the actual counts in terms of 6-in. intervals, and let us suppose we found 1326 giraffes in one of the height categories. We divide this number by 6 before plotting it, since a 1-in. category would contain, on the average, one-sixth as many giraffes as a 6-in. group centered at the same height. So, whether the measurements are done with 6-in. intervals, dividing by 6, or with 3-in. ones, dividing by 3, etc., one would in any case be plotting a point saying that at the indicated height level there are 221 giraffes per inch of height. If the intervals are narrow, there will be more points, but they will all lie very nearly on the same curve.

There is one further fly in the ointment. We have emancipated the curve from the height interval used (as long as it is sensible), but the height of the curve above the base line is still dependent on the total number of giraffes studied, and again this is a pity. If I say that there are 2700 giraffes between 12 and 13 ft tall, this will be an astonishing finding if the herd contains 3000 individuals, whereas it would be nothing remarkable in a herd of 20,000. The important question is *what fraction of the total population* will be found in the height category mentioned. Accordingly, the best way to plot a giraffe-height spectrum is to scale the vertical axis in terms of *percent of the herd per inch of height*. If we agree on this policy, you and I will produce the same distribution curve whether you measure 100,000 giraffes and I 30,000, or whether you use a 3-in. interval while I use 2 in. The curve will be generalized and will show what it is supposed to show, namely the distribution of the heights, emancipated as far as possible from dimensions. Admittedly we won't get the same curve (although we will get the same information) if you use inches while I use centimeters. This is why you should have used centimeters.

Somewhere on the graph, of course, there should be an indication of what the total population was and the width of the height categories used. If these choices have been made intelligently, there need be no fear of telling the reader what they were; whereas if they have been unwise he has a right to be warned.

We can now return to the pulse-height spectra (Fig. 9) and inquire into the origins of the little white lies that they tell, for we need to know how little and how white. These spectra are intended

to indicate the distribution of energies in the photons arriving at the detector; actually they show the distribution of heights in the amplifier's output pulses. In times past we have sometimes scaled off the horizontal axis in terms of "pulse-height units," these being dial readings related, in an unspecified manner, to pulse voltage (see Fig. 16). Their only virtue was that they were easily procurable. Since this kind of scale might just as well not be there, our present practice is to convert pulse height to the equivalent energy and plot against this — say kev or Mev. An ideal spectrum, then, will be plotted in terms of percent of the incoming photons per kev energy interval, and somewhere the actual interval used (e.g., 5 kev) will be stated.

Most of the medical spectra that we see around these days are plotted against energy all right, but it is unusual to see the ordinate in terms of percent. The reason is purely a matter of convenience: the relative abundances of the various photons, which is what we need to know, can be measured directly by the rates at which they arrive. Usually, therefore, we simply plot "count rate per kev," or some suitable equivalent, which gives us all the information we need unless the circumstances are special. If you want to plot percentages you will need an additional measurement, namely the total count rate for all the photons. Most of the time we don't bother with it.

Resolution and Its Impairment

Now for the things that distort the spectral picture. First of all, why does a line (Fig. 9, top) become a hump (Fig. 9, bottom)?

Let us imagine a stream of photons, all of the same energy, entering the detecting crystal. Some go right through it, leaving no ions; some are completely stopped in the sodium iodide, depositing all of their energy; and some give part of their energy to an electron in a Compton collision, but the deflected photon escapes from the crystal, taking the rest of the energy with it (Fig. 11). Let us talk first about the second group, the photons that undergo *total absorption*. Since these all bring the same energy into the crystal, and all of it gets absorbed, the electrical pulses should all have the same height.

In real life it doesn't work out that way. Not all of the photon's energy appears in the flash

Fig. 11. If a pulse-height spectrum is to be reliable, all the energy of the incoming photon must be absorbed in the crystal, the final event being a photoelectric one (left). If the crystal is small, or the energy high, some photon energy may escape (right), causing a sub-standard scintillation. Redrawn from ORNL-2808 (6).

of the scintillation; a small and somewhat variable fraction is turned into heat.¹⁷ So the scintillations produced by the incident monoenergetic photons have nearly the same intensities, but not quite. Moreover, not all of the light of a scintillation succeeds in entering the phototube. The white powder between the crystal and the can is not a perfect reflector, and if the light has to bounce several times before arriving at the crystal's exit face, it will be somewhat weakened. There may also be some absorption by the crystal itself; this can be serious if the crystal is old, for there is likely to be a gradual release of free iodine, which makes yellow patches here and there. The optical coupling between the exit face of the crystal and the phototube's entrance window must be good: the optical cement must provide a good match of refractive indices or the light will reflect, and great care must be taken to exclude air bubbles, which, of course, cause strong reflection. For a number of reasons, therefore, the amount of light reaching the phototube is not always precisely the same, even with monoenergetic incident photons.

¹⁷For an NaI(Tl) crystal a plot of scintillation intensity vs energy becomes somewhat nonlinear as the energy comes down below 100 kev.

The phototube adds to the dispersion. The photocathode can't be made completely uniform in sensitivity throughout its area, so the number of photoelectrons released will depend, somewhat, on the part of the cathode that is closest to the scintillation. The nonuniformity is apt to be more prominent in tubes of large diameter. If the scintillation is weak, as it will be when the incident energy is low, the photocathode releases perhaps only a dozen or so electrons, and such a small number will be subject to statistical variations. We must expect, therefore, that the spectrometer's resolving power will be poor at the lower energies. Again, there will be statistical variations in the paths that the electrons take in bouncing from dynode to dynode. These paths can be severely disturbed if there is a 60-cycle or other alternating component in the magnetic field that surrounds the phototube — one good reason for proper magnetic shielding. The high-voltage supply for the phototube may not be adequately filtered; this will produce a 60- or 120-cycle fluctuation in the electron gain. Finally, if either the phototube or the amplifier is noisy, one pulse may stand on the crest of a noise wave while another stands in a trough, and this causes random variations in the heights of pulses that should all be alike.

For reasons such as these, monoenergetic gamma radiation does not produce pulses that are all the same height, and the ideal line is therefore spread out into a hump. With a well-made detector most of the pulses will have nearly the same size, but on either side of this *mode* there will be other pulses, some taller and some shorter than they should be. The more nearly alike they are, for a given incoming energy, the better is the *resolution* of the detector. Since resolution is important, we had better see how it is measured.

Measurement of Resolution

The various sources of dispersion just mentioned spread a spectral line out into a hump or peak. This is called the *total-absorption peak*, since the pulses constituting this group all come from scintillations caused by total absorption of the incoming gamma energy. (We will come to partial absorption in a moment.) The total-absorption peak approximates the familiar bell-shaped "normal curve of error." If it is narrow and sharp, we know

that the total-absorption pulses all have very nearly the same height, which means that the resolution is good. The sharpness of the peak can be used, therefore, as an index of resolution. Draw a horizontal line through the peak at a level half the peak height; then measure the width of the peak along this line (= "full width at half maximum," or FWHM — see Fig. 23). The resolution is this width (kev) expressed as a percentage of the peak energy (also kev). Since the resolution is energy-sensitive, deteriorating as the energy goes down, it is understood to be measured at 662 kev (^{137}Cs) unless otherwise specified. Note that a *low* figure indicates *good* resolution; 8% is considered pretty fair.

One should mention in passing that good instrumental resolution can be spoiled by a mistake in operating policy; namely by the use of too wide a window opening (= slit width, pass band, channel width, etc.) in the analyzer when the data for the spectrum are being taken. Here, as in the case of the giraffes, one must steer a middle course between a window so wide that it "smears" the details of the spectrum, and one so narrow that it generates statistical troubles — and some others as well (pp. 38 ff.).

Further Distortion: Energy Escape from the Crystal

Let us move on now to some of the other sources of distortion in the pulse-height spectra of Fig. 9. The chromium-51 curve is not just a single bell-shaped peak, centered on 322 kev and showing some 10% or so of spread. It also shows a broad plateau extending downward from about 200 kev, with a prominent trough between there and the main peak. Since the emission spectrum (top) shows nothing at all over the plateau energies, we conclude that in this region the spectrometer is lying shamelessly. Why the collapse of technological morality?

This curious feature is due largely to Compton scattering, whence the name *Compton plateau*. We recall that in sodium iodide the photoelectric process is highly probable only at rather low energies; therefore when a medium-energy photon enters the crystal, the most likely event is a Compton collision. Usually this is followed by another, and perhaps others, until the energy is brought down to a level where a photoelectric

snap swallows the rest of it, and absorption is total (Fig. 11, left). But anywhere in the pre-photoelectric stage – particularly after the first collision, when the energy is still fairly high – the scattered photon may make its way to the surface of the crystal and escape, thus robbing the detector of some of the energy that ought to have been deposited (Fig. 11, right). The result is a substandard scintillation, a stunted electrical pulse, and a false count chalked up in the spectrum somewhere to the left of the total-absorption peak. If there are many such events, a Compton plateau will result. This kind of energy escape will be more frequent, of course, if the crystal is small.

The trough just left of the main peak says that nearly total absorption is unlikely; either it will be total or it will fall short of total by a considerable margin. Compton's equation (footnote 5) predicts this phenomenon. Remember that the scintillation reports only the energy of disturbed electrons; energy remaining in a photon stays hidden. A colliding photon can give all, or only some, of its energy to the electron in its path, but it can't give nearly all. Going back to the equation, E'/E is the fraction of the original energy that the scattered photon carries away; therefore $1 - (E'/E)$ is the fraction that the electron receives. This quantity increases as ϕ increases (i.e., as $\cos \phi$ decreases), but the fraction reaches a maximum as ϕ approaches 180° , for there $\cos \phi$ reaches -1 and won't go any more negative. The maximum fraction is therefore $E/(E + 256)$ (Fig. 21B), and unless the scattered photon collides again, the fraction of the energy delivered to the crystal cannot lie in the gap between $E/(E + 256)$ and 1; it will either be less than the former or 100%. The result is a trough, whose theoretically sharp edges (see Fig. 12) have been rounded off by (1) multiple Compton collisions and (2) imperfect resolution.

Figure 12 illustrates a number of other points. It is a pulse-height spectrum for cesium-137, plotted in terms of counts per second per 20 kev (the bandwidth used) vs energy. Here the vertical axis has a logarithmic scale, which is a convenient way of making the small details look larger, so that they can be studied more effectively. (If the plot had been linear, the plateau would have been so low as to be almost featureless.) Cesium-137 decays to a barium-137m daughter that emits monoenergetically at 662 kev; 11% conversion of this photon produces a 32-kev barium

Fig. 12. Solid curve: Pulse-height spectrum from cesium-137, adjusted to take out the barium K x-ray (32 kev) and predetector scatter. NaI crystal, $1\frac{1}{2} \times 1$ in. Total count rate about 700 counts/sec. Other curves: Theoretical data, as indicated. Redrawn from ORNL-2808 (6).

x ray, but we have removed this component from the spectrum to avoid complicating the discussion. Here the detecting crystal was rather small ($1\frac{1}{2}$ in. diam by 1 in.), permitting considerable loss of crystal-generated scatter.

This spectrum consists of a total-absorption peak and a Compton distribution. If we add up the values for all the (imaginary) bars here,¹⁸ we get the total count rate, which is about 700 per second.

¹⁸ Be careful about areas when working with semilog plots. Zero is at infinite distance in the downward direction (because $\log 0 = -\infty$), so the "area under the peak" will be infinite. This is why we add up bars, in this case, instead of measuring areas – see discussion of Fig. 10, p. 18.

If there were no distortion in this monoenergetic spectrum, then, it would consist of a single line at 662 kev, reading 700 counts/sec on the vertical scale (Fig. 12).

Now let us retain total absorption but introduce the distortion due to imperfect resolution. The ideal line is spread out into a narrow peak, bell-shaped if plotted with linear coordinates, but still adding up to 700 counts/sec. Figure 12 shows this peak as the dot-dashed curve. Finally turn to the real experiment, where escape phenomena are prominent: the total-absorption peak now contains only 240 counts/sec, for partial absorption has thrown the other 460 counts/sec into the Compton continuum lower down.

It was pointed out earlier (p. 4) that the only "honest" photons coming from a patient are the primary or undegraded ones; only they give reliable information about both the position and the energy of the hidden radionuclide. This tempts us to use only the counts in the total-absorption peak of the spectrum and ignore the deceptive remainder. But if we did this while using the small crystal of the Fig. 12 experiment, we would be working with a deplorably poor payload, throwing away nearly two-thirds of the counts.

The moral is that one should encourage total absorption by all means possible, and one good way is to discourage escape by using a large crystal.

Further Distortion: Predetector Scattering

In the foregoing section the pulse-height spectrum was accused of flagrant mendacity in the region of the Compton plateau, and this was not wholly fair. Monoenergetic emitters exist, but by the time the radiation reaches the detector, some scattering has always occurred, so that many of the photons are already degraded in energy. They have bounced off other atoms in the source, in the material surrounding the source, in the air between there and the detector, in the walls of the detector's collimator or shield, or in the can or the alumina reflector surrounding the crystal (Fig. 4). The radiation may be lily-pure, therefore, as it leaves the atoms where it was born, but by the time it enters the crystal it is likely to be contaminated with scatter, and one can't fairly blame the detector for not being able to guess what the degraded photons were like before they lost their virtue. In the pulse-height spectrum

for chromium-51, then (Fig. 9), some of the substandard pulses are due to *pre-detector*, rather than *intra-detector*, scattering. The results of this process are even more prominent in the iodine-131 pulse-height spectrum at lower right in Fig. 9, for there the emitter was contained in a thyroid gland surrounded by its owner's neck. The scatter pile between 100 and 200 kev is one of the consequences.

It is important to recognize that diametrically opposing roles are played by predetector and intradetector scattering — for our purposes, scattering within the patient and within the crystal. The energy reported by the scintillation is electron energy: energy transferred to the crystal's electrons by any or all of the three modes of interaction between a photon and matter. The energy remaining in a crystal-scattered photon is still out of sight as far as the spectrometer is concerned. In a succession of scattering interactions in the crystal, each collision increases the probability that the absorption will become total, since the photon energy comes down each time. This is favorable; so we try to procure a large crystal in order to leave plenty of room for a series of Compton sideswipes, leading up, we hope, to the final photoelectric crash.

When we turn to the patient, on the other hand, these same processes become drawbacks, for the only kind of energy we can detect coming from the patient is *photon* energy; the electron energy doesn't get out. Each Compton collision in the patient, therefore, decreases the experimenter's chances of learning something useful: it changes the photon's direction (important in the patient though not in the crystal) and increases the probability that not even a substandard photon will be available for inspection. Moreover, it raises the spectre of possible radiation injury to the patient, which is important emotionally if not otherwise. Scattering within the patient, therefore, is a curse.

Since the information remaining after predetector scattering will reside in photons rather than in electrons, the trough-and-plateau relationships, described above for the crystal, will be exactly reversed in predetector scattering: the trough will be near zero and the plateau to its right. Each may, and probably will, be masked by other, concomitant processes. This is one more thing to bear in mind when a patient's spectrum is being studied.

In summary, when a spectrum is obtained from a patient harboring a monoenergetic emitter, the Compton region will be registering both good and bad photons. The bad ones are patient-generated scatter: goats masquerading in sheep's clothing, with their change of direction threatening to compromise the detector's geometry. These goats can be recognized by an energy-sensitive counter and kept out of the count. But there are also sheep in the Compton plateau: photons that reached the detector undegraded, by going straight, but then suffered the misfortune of losing some of their energy out of the crystal and thus producing substandard pulses. To the spectrometer these sheep look like goats, and thus good information can be thrown away. A large crystal provides most of the true sheep with total absorption, which places them in the primary peak where they belong; there they can be counted, uncontaminated with goats, by a properly instructed pulse-height spectrometer.

Minimizing Scattering Effects

To reduce the distortions that we have been discussing we should try to avoid gamma interactions within the patient while encouraging them in the crystal. One thinks at once of the medium-energy photons, to which the patient will be nearly transparent while the crystal is nearly opaque. [But the low-energy emitters do have some redeeming virtues. Before espousing one, weigh the pros and cons carefully against each other (12, 13).] Apart from choice of energy, there isn't much one can do to get photons out of the patient, whose size, shape, and material are built-in hazards not lightly to be tampered with. When we turn to the detector, on the other hand, there is much that can be done, and the most important thing is to procure as large a crystal as is physically and fiscally practical. This will provide better geometry for any given distance, making more sheep and fewer goats (Fig. 1). Also, the crystal's thickness will reduce the loss of sheep via straight-through leakage. Furthermore, the loss of crystal-scattered photons will be minimized, making the spectrum report more truthfully on the energies of the incoming photons and permitting the sheep and the goats to be separated. In the medium-energy range, photoelectric absorption at the first collision is not a reasonable possibility; at the cesium-137 energy the Compton process is

Fig. 13. Showing how the stopping power of sodium iodide changes with energy. The top curve is the total linear attenuation coefficient (μ), excluding undegraded (coherent) scattering; the curves below show its three components. (Graphs by Harshaw Chemical Co., from data by White-Grodstein, NBS Circular 583.)

nine times more probable (Fig. 13). We must therefore give the incoming photon plenty of room in which to rattle around inside the crystal, for with each collision total absorption becomes more likely.

Figure 14 compares two crystals at 322 kev (^{51}Cr). Both have the same geometry, for the crystal with twice the diameter of the other is twice as far from the source. When the counts are added up, the larger crystal is found to give 43% more total counts, and, more important, it gives 69% more peak counts. With a higher energy the advantage of the 3- by 3-in. crystal would be still more striking. A 3- by 2-in. crystal is nearly as good, which accounts for its current popularity in general-purpose detectors.

Fig. 14. Pulse-height spectra from a chromium-51 source, using two NaI(Tl) crystals at the same geometry: 20 cm distance for the 3- x 3-in. crystal and 10 cm for the 1½ x 1. The larger crystal, in spite of the greater source-to-crystal distance, gives 43% more total counts than the smaller crystal, and 69% more peak counts.

Crystal Efficiency

For purposes of quick reference there are three indices describing a crystal's absorbing and analyzing performance, as follows:

Intrinsic peak efficiency

$$= \frac{\text{total-absorption scintillations}}{\text{total incident photons}}$$

Intrinsic total efficiency

$$= \frac{\text{scintillations of any kind}}{\text{total incident photons}}$$

Peak-to-total ratio (P/T)

$$= \frac{\text{total-absorption scintillations}}{\text{scintillations of any kind}}$$

The geometry, or extrinsic efficiency, will be:

Geometry

$$= \frac{\text{gamma photons reaching the crystal}}{\text{all gamma photons emitted from source}}$$

Note that the intrinsic peak efficiency is the product of the intrinsic total efficiency and P/T. As one would expect, the three intrinsic ratios depend on the entering energy and on the size and shape of the crystal. As one might not expect, they also depend on the position of the source, and thus are not purely intrinsic. The reason for this is "corner cutting." The escape of a scattered photon is more probable if the Compton collision occurs near one edge, but it will be still more probable if the photon enters the front face near the edge and also at an angle, tending to cut through the corner. Corner cutting will be minimal if the source is either close to the crystal or far away, and it will be prominent when the source is one crystal radius away from the front face. Large crystals will not only cut down leakage and have high absorption efficiencies, but they also will discourage corner cutting, for the large front face will present proportionately more interior and less edge to a source that is in the unfavorable, middle-distance position. Sometimes the middle-distance situation can't be avoided, for short distances may be impractical (e.g., with a honeycomb collimator), while the longer distances may involve intolerable sacrifice of geometry. A helpful trick, in such a case, is to bevel the front edge of the crystal; this takes away the thinnest part of the corner.

In summary, if the crystal is large: (1) there will be more incident rays, for the same source distance; (2) more of the incident rays will interact; and (3) a larger fraction will interact totally.

Figure 15 shows how crystal efficiency changes with energy, the upper set of curves describing intrinsic *total* efficiency while the lower set shows intrinsic *peak* efficiency. The crystal size (in inches) and the source-to-crystal distance (in centimeters) are indicated in each case. The total-efficiency curves illustrate the effects of corner cutting very nicely.

All the curves emphasize that the efficiency will improve, sometimes quite rapidly, as the photon energy comes down. This has important consequences, because it means that the crystal

Fig. 15. Intrinsic efficiency, as a function of energy, for a variety of sodium iodide crystals. Top: Intrinsic total efficiency for 3- x 3-in and 1 1/2 x 1-in crystals, with two source distances for each. Corner cutting is severe, reducing the efficiency, when the source is about half a crystal diameter away. Redrawn from ORNL-2808 (6). Bottom: Intrinsic peak efficiency for crystals with dimensions indicated (in inches), for source-to-crystal distances shown. See text for discussion. From (8).

will emphasize predetector scatter. Moreover, the smaller the crystal, the more prominent this selective effect will be, since the curves slope upward and to the left more steeply when the crystal is small. This makes the rejection of unwanted scatter more difficult. In fact, we can now see that it is not scatter per se that the spectrometer rejects; it is the Compton continuum, whether its pulses are unwanted or not. It would be a pity, therefore, to procure a good spectrometer, with excellent latent discriminating ability, and then poison its performance by marrying it to a small crystal. It is equally foolish to procure a large crystal and marry it to a counting system that swallows everything unselectively.

Special Kinds of Escape

We have been discussing ordinary, nonspecific escape of photon energy from the crystal. There are, in addition, two special kinds of escape that require mention. One occurs when the incoming energy is high enough to make pair production frequent. The annihilation radiation's two 511-keV photons are highly penetrating ($X_{1/2} \approx 2$ cm), and one or both of them may skip out of the crystal without contributing to the scintillation. In such a case the deposited energy is diminished by a definite amount, namely by 511 keV if one ray escapes or by twice that if both do. Thus in a spectrum of sodium-24 (Fig. 16) there is a prominent primary peak at 2760 keV, a small single-escape peak at 2250 keV ($= 2760 - 511$), and a still smaller double-escape peak at 1740 keV ($= 2760 - 2 \times 511$).

In the other special kind of escape, what would have been a total-absorption scintillation is substandard by about 28 keV, the energy of an iodine ($K - L$) x ray. This happens when the incoming photon's energy is fairly low, in which case (1) the gamma will interact close to the entrance surface, and (2) photoelectric interaction is probable. In this interaction, a photoelectron is ejected from the K shell of an iodine atom; this causes a scintillation. For complete absorption of the photon's energy, however, the x ray produced by the filling of the K -shell vacancy must also be absorbed within the crystal. This x ray may come off in any direction, of course, but if the interaction was only one or two millimeters deep and the x ray heads back toward the surface, it will have a fair chance of getting out of the crystal. Whenever

Fig. 16. Spectrum from a sodium-24 source placed 3 cm from a 3- $\times$ -3-in. NaI crystal. Each pulse-height unit corresponds to $32\frac{1}{2}$ kev. The primary spectral peaks at 2.76 and 1.38 Mev are easily recognized. Since pair production will occur at these energies, escape of some of the resulting annihilation radiation causes lesser peaks to appear: (a) for single escape, at 2.25 Mev ($=2.76 - 0.511$); and (b) for double escape, a still smaller peak at 1.74 Mev ($=2.76 - 2 \times 0.511$). Similar escape peaks are not found below the 1.38-Mev peak because this energy is not much more than the minimum required for pair production, namely $2 \times 511 = 1022$ kev. This photon therefore produces many Compton electrons but few positron-electron pairs - cf. Fig. 13. From ORNL-2808 (6).

that happens, the scintillation will represent the energy of the incoming ray minus the 28 kev lost in the escaping x ray. For this reason a pulse-height spectrum showing a total-absorption peak at 150 kev or less often shows, in addition, a smaller peak placed 28 kev to the left of the primary peak. Several examples of these escape

peaks may be seen in Fig. 17, where the primary energies are mostly below 150 kev. Other typical spectra are shown in Fig. 18; here the energies lie in the medium range (clinically speaking), and the primary peaks show little evidence of x-ray escape.

Organic Scintillators

Thus far our discussion of the detector's behavior has been based on the assumption that the scintillating material will be sodium iodide, as it usually is in medical work. There are many other kinds of scintillators, however, and one group in particular should be mentioned. These are the organic scintillators, some being solid (either crystalline or plastic) while others are liquid solutions. Both have a place in medical work; they have limitations, but they offer some compensating advantages.

The chief advantage is that they can be big. Sodium iodide crystals may be grown to an astonishing size, but a point is reached where the process becomes so troublesome that this material prices itself out of the market. Here the organic scintillators take over; not only do they cost much less per cubic centimeter, but they can be made much larger. The tank counters developed at Los Alamos provide a good example: these use well-type detectors big enough to accept a whole adult man (1). Smaller versions accept a leg, an arm, a dog, a rat, and so on. Since the scintillator largely surrounds the sample, the detector geometry is not too far below unity, the advantages of which the reader will appreciate. With a large detector, background is apt to be high, but the similarly high net count rates generally take much of the sting out of the high background, so that counting times need not be excessive and can often be surprisingly short.

The big drawback, as far as gamma detection is concerned, is poor stopping power. Sodium iodide has a density of 3.67 g/cm^3 , and the low density of the organic scintillators means that they will present considerably fewer electrons per square centimeter to the oncoming radiation. Thus the probability of Compton interaction is reduced to about one-third. To make matters worse, the low proton numbers of the elements involved ($C = 6$, $N = 7$, $O = 8$, etc., compared with $I = 53$) make photoelectric interactions unlikely until the energy becomes very low (<1 kev). With an organic

Fig. 17. Pulse-height spectra from four rather low-energy gamma emitters. Xenon-133 and mercury-197 are from Heath's data (9), with a 3- $\times$ 3-in. crystal at 10 cm. Cobalt-57 is ORNL's, run at 20 cm from a 3- $\times$ 3-in. crystal. Backscatter was minimal, since the crystal hung free in the laboratory air, some 8 ft above the floor and well away from the walls. The technetium-99m spectrum is also ours, run at 23 cm from a 3- $\times$ 3-in. crystal inside a 4-in.-thick lead box. All these spectra are net.

Exercise – For each of these nuclides, calculate the energy of the backscatter peak, $256E/(256 + E)$. Note how it crowds up toward the primary when the energy is low, so that all the first-collision scatter is contained in a narrow band. Look for the Compton edge at $E^2/(256 + E)$.

Fig. 18. Typical pulse-height spectra from four gamma emitters in the medium-energy range. Mercury-203, strontium-85, and gold-198 are plotted from data by Heath (9); all used a 3- $\times$ 3-in. NaI crystal at 10 cm; energy bandwidths are as shown on the ordinate. Selenium-75 was run in our laboratory using a 3- $\times$ 3-in. NaI crystal. Two source distances were used to show how summing effects can alter the configuration of the spectrum. The 401-kev state of the excited arsenic-75 daughter undergoes transition to ground state by (at least) two routes: (1) directly, with the emission of a 401-kev gamma ray, and (2) in a two-stage cascade, emitting photons at 137 and 265 kev. These two are very nearly coincident in time, since the half-life of the intermediate state is less than 15 μ sec. With a typical source-to-crystal distance — that is, with poor geometry — only one member of a synchronous pair is likely to enter the crystal, but if both of them should happen to do so, the scintillation will look like a 401-kev one. Accordingly, if the geometry is improved (upper curve), the 401-kev peak increases considerably in relative strength. To avoid overload problems, a weaker source was used for this curve.

Summing phenomena become particularly prominent in a well counter, owing to the good geometry. Whenever spectral data are used for a quantitative assay, the possible consequences of summing must be carefully considered.

Query: What has happened to the Compton shoulder in the mercury-203 spectrum?

scintillator, therefore, the absorption of gamma energy is rarely total, and the pulse-height spectrum thus lacks a total-absorption peak. It follows that these materials are useful as gamma detectors mainly in situations where good geometry is important but good spectrometry is not.

When it comes to alpha and beta counting, on the other hand, the organic scintillators take on a new look. The scintillator can be a liquid solution with the sample mixed right into it. This provides close to 100% geometry and 100% absorption for alpha and beta particles, since their ranges are typically short. The mixture of sample and scintillator fluid is placed in a cuvette that is inspected by two phototubes, and each scintillation produces a pulse whose height is a measure of the emitted energy. There are some difficulties, however. The organic scintillators do not produce much light, and if the emitter gives off only a weak beta particle (e.g., hydrogen-3: 18 kev maximum), the pulse spectra for signal and background may overlap disagreeably. "Dark noise" pulses generated within the phototube can also be troublesome, and to control them the phototubes may have to be specially selected, operated in a refrigerator, or provided with coincidence circuitry. Furthermore, comparison with a known standard is often upset because the sample may bother the scintillator chemically. Sometimes a sample will not dissolve easily in the scintillating fluid, and added solvents may interfere with light production. Or the sample may color the fluid, so that even if light is properly generated it may not reach the phototubes without considerable, and nonuniform, attenuation. These difficulties with light production and transmission are referred to collectively as "quenching"; they affect the pulse-height picture like a reduction in gain, crowding the spectrum to the left (Fig. 26) and – the matter of practical importance – making the sample and standard spectra dissimilar. Several special techniques, beyond the scope of this account, have been developed to cope with quenching problems: special solvents and scintillators, various tricks of window operation, standardization by "spiking" the unknown, and so on (11). However troublesome this technique may be, it nevertheless provides a powerful sample-counting tool, helping to bridge the embarrassing gap left by those biologically important elements (C, H, S, P, etc.) that provide no gamma emitter with a reasonable half-life.

MULTICHANNEL ANALYZERS

Spectrometers of the kind we have been discussing are called *single-channel* analyzers because they determine the count rate for a single band of energies. The counts lying outside this band are discarded, and potentially this could be wasteful, for it lowers the count rate and may impair statistical accuracy. One can throw away counts legitimately only when they represent uninformative or obscuring trash (background, for instance), and an effort should be made not to throw away "good" counts.

When a patient's liver is to be counted, therefore, the experimenter decides which energies are informative and sets up his window to catch this band, letting the rest of the radiation go. When a spectrum is needed, on the other hand, one must count all the energies without suppressing any, for the spectrum will be informative only so far as it reveals the relative populations of the various energy bands, as measured by their count rates. The procurement of a spectrum with a single-channel analyzer is a slow and tedious process: one must place a narrow window successively at the various pulse-height levels and take a count for each (e.g., Fig. 16). It can be made automatic by providing the *E* dial with a motor drive that sweeps the window slowly through the spectrum at constant speed while a recording rate meter makes a graph of the changing count rate (e.g., Fig. 29). But this is a wasteful approach, for at any one time the window is looking at only a small fraction of the radiation coming in, while most of it is being thrown away. The method is satisfactory, therefore, only if there are plenty of photons to play with, and one can afford to squander most of them.

Often this will not be the case. In particular it is not likely to be so in human spectrometry, where the "sample" is a whole patient. Here the available radiation is usually small in quantity, since nobody wants to give an experimental subject any more radionuclide than is strictly necessary. Sometimes, also, it is desirable to study normal people (we are all weakly radioactive), and in this case the radiation level will be very low indeed. It is true that in some situations one can compensate for a low count rate by increasing the counting time, but this won't work when you're counting live patients, who can't – or won't – lie still for hours on end. They aren't that patient.

There are times, therefore, when a wasteful spectrometer isn't good enough, and the single-channel instrument is out.

Faced with a low-activity source, how can we avoid throwing counts away? Instead of having a single scaler, counting pulses in a narrow energy band, we need a whole series of scalers representing a series of adjoining energy bands, and we should measure the voltage of each pulse as it comes along and then direct the count into the scaler whose band covers that voltage. Such a device would work somewhat like a post office, where each letter is scrutinized and then tossed into the proper pigeonhole. For good energy resolution we might need a hundred or more such receptacles. Translating into electrical jargon, we need a multichannel *memory*, preceded by an *addressing unit* that looks at the voltage of each pulse and then routes it into the correct memory channel. After an appropriate counting time we can ask each channel how many counts it has accumulated, and the answers will enable us to plot out the spectrum. With this approach none of the pulses are thrown away, and thus it becomes possible to obtain low-level gamma-ray spectra without an inordinate counting time.

The memory of a multichannel analyzer is usually made up of tiny magnetic rings, about 2 mm in diameter, each of which can be magnetized either "north" or "south," and thus can store either 0 or 1 in a binary code. Strung up on the proper wire grids, with suitable driving circuitry, a multitude of these little rings can provide 500 memory channels in a space no bigger than a lunch box, each channel being capable of storing up to a million counts. Accessory circuits can question the channels in succession, asking each how many counts it contains; they can do this nondestructively and repetitively, and — this is hard to believe — so fast that they run through the whole memory bank 30 times a second. Thus a standing spectrum can be made to appear on the face of a cathode-ray tube. Other circuits can make the multichannel memory print out its numbers through a printer or typewriter; still others can make it plot out the spectrum on a chart. Not long ago these miraculous machines were rare, but they are now fairly common.

The device that measures the height of each pulse is called an *analog-to-digital converter*, or ADC for short. The pulse height — the analog of energy — is converted into a number, and the

taller the pulse the larger the number will be. The row of memory channels is also numbered, and it is not too difficult, though hardly easy, to arrange things so that once a pulse gets its number it is routed into the memory channel carrying that number. This gets the sorting and addressing job done. With modern circuitry, the ADC works so fast that it can process a pulse in 100 μ sec or less, after which it is ready for a new pulse. With a dead time as short as this, few pulses will be thrown away unless the count rate is high, and high count rates won't occur when you're counting normal people or other weak samples.

The multichannel analyzer thus enormously reduces the time and labor required to obtain a gamma-ray spectrum, and since it throws very few of the counts away, its internal efficiency is high, making human spectrometry achievable within reasonable counting times. We should hasten to add, however, that the virtues of a multichannel analyzer don't make human spectrometry easy, for with a sample as large as an adult person, detector efficiency and background control provide plenty of problems for the designer. The analyzer simply leaves him with a bad headache instead of a lethal one.

The Multiscaler

When a laboratory buys a multichannel analyzer, it is often procuring an important accessory capability, for most of these instruments nowadays have a switch that provides for operation in what is called the "multiscaler mode." In this service the ADC is disconnected, and the business of the memory system is simply to count pulses without regard for their heights. (Normally, therefore, the multiscaler is preceded by a separate single-channel spectrometer.) A timing device opens up one channel after another to the oncoming pulses, leaving each channel accumulating counts for a constant, short time interval. With this arrangement the plotted-out contents of the memory will show the course of a changing count rate as a function of time. The advantage of this system over a conventional recording rate meter is that the memory has no inherent lag, as the storing circuit of a rate meter has (7). The multiscaler is therefore particularly valuable in studies of rapidly changing count rates, the only limitation being that enough counts must accumulate in each chan-

nel to make the statistics reasonable. The duration of the time intervals is chosen with this consideration in mind.

OPERATING POLICIES

As we have seen, a pulse-height spectrometer's window is basically a rejecting mechanism: it allows some of the pulses to be counted and throws the rest away. If this valuable feature is to be used intelligently, the operator must adjust the instrument's controls correctly, so that useful information is not lost.

E-Dial Calibration (Gain Adjustment)

If we are to place the window so that it sees some chosen energy band, we must establish the relationship between gamma energy and the divisions on the window's two controls (E and ΔE dials). The window makes its selection on a basis of pulse height, and the height of the pulses will depend on the overall electrical gain, including phototube multiplication. We need to adjust the gain so that each E -dial or ΔE -dial division represents a known number of kev. We normally make this ratio a small whole number; this is not only a convenience, it minimizes mental arithmetic, and therefore mistakes.

The E dial is usually a ten-turn potentiometer with 1000 scale divisions. Since 1000 kev spans most of the energies used in medical work, the usual gain adjustment is the one that gives 1 kev per E -dial division. Let us try for this relationship first.

The method is as follows: (1) pick a convenient monoenergetic source whose energy is known, (2) set up the E and ΔE dials so that the center of a narrow window falls at the number of divisions corresponding to the source energy, and (3) adjust the gain until the tallest pulses enter this window. The description of window position requires a little care. Ordinarily the E dial determines the height of the window *sill*, while ΔE controls the window *opening* (Fig. 8). The designers usually make E and ΔE divisions equivalent, in which case the window center will be at $E + \frac{1}{2}\Delta E$. Twenty divisions is a typical ΔE setting for gain adjustment; it gives adequate resolution without making the count rate too low, and the half-window has the convenient value of 10.

With ΔE at 20, set the E dial to read the source energy minus 10.

Cesium-137 is a favorite calibrating source. Its barium-137m daughter provides only one gamma photon, at 662 kev, and an x ray at 32 kev, the latter being useful for zero adjustment (see E -zero offset). The 30-year half-life makes decay unimportant. Another useful nuclide is tin-113, whose daughter emits a 393-kev photon; the half-life is 112 days. In the range below 100 kev, 1.3-year cadmium-109 gives a gamma photon at 88 kev and an x ray at 22. Calibrating sources should be unspillable, easily set up in air, and should have a minimum of scattering material around them. With cesium-137, 1 to 20 μc is about right for an external detector; a well crystal needs much less.

To calibrate with cesium-137, set the E dial to 652 divisions and ΔE to 20, making $E + \frac{1}{2}\Delta E = 662$; then present the source to the detector. The amplifier will be delivering a stream of pulses (Fig. 21A), and we need to adjust the gain so that the total-absorption pulses will find the chosen window. If we start with the gain too low and slowly increase it, the pulses will at first lie below the window sill, but as they grow tall enough to begin entering the window, the count rate will rise. If the gain becomes too high, the pulses will extend up above the window, and the count rate will fall again. When the gain has been adjusted to produce maximum count rate, 1 E division = 1 kev. Figure 19 represents the procedure

Fig. 19. Illustrating the calibration of a spectrometer's E dial with a cesium-137 source: adjustment of the high voltage to make one dial division equivalent (in this case) to 1 kev.

pictorially for a spectrometer that uses its high voltage for continuous gain control.

The means of controlling the gain will vary from instrument to instrument. It is best if the amplifier can have a stepwise coarse gain control and also a continuously variable fine control. Gain can also be adjusted through the high voltage supplied to the multiplier tube in the detector, and sometimes this HV knob is the only continuous gain control the system has. In any case, bear in mind that though the calibrating source may be monoenergetic, its spectrum is likely to contain more than one hump (e.g., a backscatter peak or an x ray), and it is important not to standardize on the wrong peak. This is the reason for starting with a gain that is too small and bringing the primary peak up into the window *from below*. If this is done, the first hump to come into the window will be the primary one. On the other hand, if you start with the gain high and bring it down, you may unwittingly tune in one of the peaks below the primary, and the calibration will then be badly in error. (This hazard is greatly increased if you use a polyenergetic source.) As a final check, therefore, leave the gain where you think it should be and run the *E* dial up toward the top of its range. There should be only a few background counts there, for the cesium-137 spectrum peters out quickly above 750 keV (Fig. 12).

Some indicator for the count rate is needed for this adjustment, and first choice would be a "howler." This is a device that produces an audible note whose pitch goes up or down with the count rate, and therefore when the pitch is highest the count rate is maximum. An audible signal has the advantage of leaving the operator's eyes free to watch the dials. Since the total-absorption peak is rounded at the top, the end point is not sharp; accordingly check the adjustment by offsetting the *E* dial perhaps 25 divisions, first below, then above, the presumed maximum. The two offset rates should be about equal. Precise equality is not needed because the sides of the peak slope steeply in the half-maximum region.

If a howler is not available, any other suitable count-rate indicator can be used. A rate meter will do the job, but (1) it requires visual attention, and (2) its action is jiggly if it reaches equilibrium quickly, or slow if one lengthens the time constant in order to smooth out the response. Another indication is the rate of flashing of the neon lights in a scaler or the rate of rotation of the pattern in

a dekatron tube. These are usually adequate to signal approximate equality between the two off-center count rates, and it should be possible to locate the primary peak where it is wanted within two or three divisions of the *E* dial.

It is well to mention two sources of error in the *E*-dial calibration. These are (1) amplifier nonlinearity and (2) *E*-zero offset.

Amplifier Nonlinearity. — This can distort the proportional relationship between gamma energy and pulse height. If the *E* dial and the amplifier are both working correctly, one should be able to peak cesium-137 at $E = 652$, $\Delta E = 20$, and then find that chromium-51 (for example) peaks at $E = 312$, $\Delta E = 20$. Remember, however, that the sodium iodide itself is somewhat nonlinear at energies below about 100 keV.

***E*-Zero Offset.** — The *E*-dial mechanism may get out of adjustment, so that zero on the dial does not represent zero volts. A plot of window-center reading vs energy should pass through zero when extended; if it does not, there is usually a screw-driver adjustment that will correct the situation. The cesium-137 peak and its barium x ray are useful in testing for zero offset.

Because these two troubles may occur without warning, regard the initial calibration as approximate only and make the final gain adjustment using the actual sample material. This policy, coupled with that of comparing with a known standard in a good phantom, will take care of many counting hazards.

Window Operation

Granted the privilege of energy selection, one must decide which energy band will be most effective: where the band should be centered (window position) and how wide it should be (window opening). The choices depend on what the window is expected to accomplish.

A spectrometer can be helpful in four main ways:

1. in achieving a good "signal-to-noise ratio" (S/N),
2. in minimizing statistical uncertainty,
3. in providing operation that is insensitive to drift,
4. in producing a gamma-ray spectrum.

Let us consider these in turn, for the policy varies with the need.

Signal and Noise: the S/N Ratio. — Remember that the window's basic function is to *reject*; it counts only a fraction of the incoming pulses and throws the rest away. Our job is to make the spectrometer accept the "good" counts that bring in reliable information, while it throws away "bad" counts that represent meaningless or confusing trash. In line with engineering practice we call the informative counts the *signal* and the rest *noise*. The signal will never be available in pure form, although the noise may not contaminate it seriously. The window's job, then, is to reject as much of the noise as possible, so that the signal stands out prominently.

It will be realized that what we are calling the signal *usually* corresponds to the net count rate: the gross rate minus room background. But there are other kinds of background as well, so each of them deserves a bit of discussion.

Room Background. — This kind of noise is always with us. It arises from cosmic-ray degradation products and from the minute amounts of radioactive material present in the surrounding building, in the detector itself, in its shield, and in some of the electronic components located within the shield. These nuclides include potassium-40, the various members of the uranium-238 and thorium-232 families, and perhaps some of the fallout elements. Room background has the virtue that it is usually constant; exceptions are caused by improperly shielded sources that people move around, by spilled radioactive materials, by "hot" patients who may walk about in nearby rooms (a wall may be poor protection), by x-ray and therapy machines, which get turned on and off, and so on. Even a watch with a radium dial can cause trouble.

Given proper shielding and good housekeeping, the spectrum of an ordinary background should be nearly *structureless* — that is, it should have no prominent peaks in it. It will rise toward the low-energy end because of the prevalence of scattered radiation, and conversely it will be weak at the higher energies. In the middle range it will present a reasonably smooth curve, representing degradation products more than primary photons. This is the usual picture. With exceptional shielding and a long count, on the other hand, some signs of structure may appear. This is illustrated in Fig. 20, where a count was run all night with a 3- by 2-in. crystal inside a 4-in.-thick lead box. The peak at 1460 kev is undoubtedly due to potassium-40, which occurs in natural potassium, one atom

Fig. 20. A laboratory background spectrum, run for 1000 min (overnight) to bring out its weak structure. A 3- x 2-in. NaI crystal inside a large, 4-in.-thick lead box; multichannel analyzer at 20 kev per channel; semi-log and linear plotting. Note very low count rates (scale at left). In the spectra we can recognize potassium-40 (probably in the detector), thallium-208 (a thorium daughter), and small peaks near 500 and 600 kev, perhaps due partly to pair production (511 kev) and partly to fallout contamination: for example, ruthenium-106 → rhodium-106 (513 and 625 kev).

In ordinary backgrounds, where the shielding is much thinner, these features are usually diluted out of sight.

in 8500. This element is troublesome because of its wide distribution and high solubility, and a manufacturer must exercise extreme care if he is to keep potassium-40 out of the crystal, the phototube, and the electrical components mounted on the phototube's socket, all of which will be inside the shield. Feeble indications of the radium or thorium daughters may also be seen in a background spectrum. In Fig. 20 there is little or no

evidence of radium, but the thorium family produces a weak thallium-208 peak at 2615 kev.

The peak near 500 kev might be due partly to annihilation radiation, resulting from pair production by any high-energy photon (p. 3), but it probably gets a lift from the fallout elements. One-year ruthenium-106, for example, produces a rhodium-106 daughter emitting at 513 and 624 kev. It must be emphasized, however, that these peaks, whatever their sources, are seen here only because of the good shielding and the long count; ordinarily they would be buried, unseen, in an unstructured, statistically ragged "background continuum" (Fig. 23, lower curve). As we shall see shortly (p. 37), when a spectrum is unstructured, the count rate is *approximately proportional to window opening*.

Body Background. — When some organ in a patient is being counted, the usual room background will be present and, in addition, a body background derived from those parts of the patient that one would prefer to exclude — that is, from the so-called *nontarget areas*. Body background will not be structureless, for it is derived largely from the administered radionuclide and will show peaks for most of the emitted energies. It is important to note, however, that the *body-background spectrum may not match that of the target organ*, because there will be different degrees of absorption and scattering in the various tissues that the detector can see. Thus the thyroid spectrum may differ from that of the nonthyroid neck background, the latter is likely to differ from a thigh spectrum, and so on. Many people seem unaware of this, and will gaily subtract a thigh count from a gross neck count without ever wondering about the spectra. Under certain rather extreme conditions this is legitimate, but more often it is not. One can't subtract 3 cm from 10 in. and expect the 7 to mean anything.

Cross-talk Background. — This occurs when two radionuclides are present together. The photons from the higher-energy nuclide produce cross-talk that complicates the measurement of the lower-energy nuclide. This problem is discussed on p. 45.

Scattered Radiation. — This is noise more often than not. It is undesirable when it is out of control, so that an unknown sample and the standard of comparison can't be counted in the same geometry (p. 6). In a well counter, on the other hand, the dimensions and position of the standard can accurately duplicate those of the unknown; here

the changes of direction caused by scattering do not upset the validity of the comparison, and the inclusion of the scatter counts will add to the count rate and may raise the statistical reliability.¹⁹ Sometimes it can be helpful to include scatter counts even when they are not under good control, simply because a scattered photon can compensate for a primary photon that has been lost through deflection or absorption. This is a rather tricky game, requiring specially selected thresholds or energy pass bands, determined experimentally for each radionuclide to be used, and for each configuration of the scattering medium (4).

Scattered radiation is bad when a position-sensitive detector is trying to locate a buried source. This is because a scattered photon, by definition, has been deflected out of its original course and therefore looks as if it is coming from a location other than the true one. The scanning people speak of the *target* region, at which the collimator is aimed, while everywhere else is *non-target*, and their need is to keep T/NT, the *target/nontarget ratio*, high. Thus T/NT in scanning is the counterpart of S/N elsewhere, for the non-target counts are detrimental and are therefore classed as noise.

Electrical Noise. — No amplifier is completely quiet, and on the resting base-line there will be small amounts of hissing, ac hum, power-line static, and so forth. A resting phototube similarly puts out small pulses called *dark noise*. A signal pulse that happens to stand on the crest of a noise wave will look taller than an identical pulse arising in a trough; thus electronic noise impairs resolution. When the signal pulses are small, the noise pulses will be proportionately larger, which is why quiet amplifiers and phototubes are particularly important when low-energy counting is to be done. Not all electrical noise is of low voltage, however; spike pulses can get into the amplifier through improperly shielded leads or through the power lines. Any device with sparking contacts can cause this kind of trouble: a commutator-type motor, the reversing dc motor on a scanner, the relays in an electric typewriter, and so on. It is sound policy to keep such devices off the power

¹⁹Statistical improvement is not guaranteed, since the inclusion of the scattered radiation will also bring in more room background. Apply the F.S.D. formula given on p. 40 to see whether the situation is really better.

line that supplies the amplifier and the high-voltage unit.

The foregoing kinds of noise can be classified into two basic types: the *constant* and the *variable*. The former is easier to deal with, for unless it is overpowering it can be measured carefully and subtracted out. In such a case, as we shall see shortly, the best policy is to make S^2/N [= (net sample count)²/background] as large as possible. With variable or intermittent noise, on the other hand, a different approach is required. This problem is likely to drive an operator out of his mind, for even if he can get a fair estimate of the average noise level, how is he to know what the level was while a particular count was being run? His only hope is to keep the S/N ratio high, by enhancing S or reducing N.²⁰ If this is done, it will increase the likelihood that the variations in the noise component will be insignificant, and the average can then be subtracted out.

S and N react differently to both the location and the opening of a window, so window operation is important. The choice of window *location* usually presents no problem; the spectrum for the signal (net) and that for the noise can be plotted one above the other, whereupon inspection will indicate where the best S and N are likely to be found. Ordinarily the window is placed over a prominent peak, for there the signal level rises whereas the noise usually does not. The influence of window *opening* on S and N, however, is less obvious, and before discussing it we must establish a very important basic principle.

Count Rate and "Enclosed Area."²¹ — The principle is that if one represents the window on a linear pulse-height spectrum as in Fig. 21B, with its sill to the left and top to the right, the count rate will be proportional to the *AREA* lying under the window: the area bounded on each side by the window's sill and top and extending from the base line up to the line of the spectrum. This law stems from the essential nature of the spectral

²⁰If a neighboring x-ray installation is the source of the trouble, you may have to move your laboratory to a distant location. X-ray machines not only produce radiation; they also throw brief, heavy loads on the power lines, making the line voltage dip from time to time, perhaps beyond the compensating capacity of the regulating circuits in your instruments. Therapy for line-voltage variations is discussed on p. 43.

²¹Whenever we speak of an area under a spectral curve, we are referring to a plot in which both axes have linear scales. In a semilog spectrum the "area under the curve" becomes infinite ($\log 0 = -\infty$).

Fig. 21. The electrical pulses (A) can be sorted out according to pulse voltage and thus according to gamma energy. The count rate per keV can then be plotted out into the spectrum (B). The shaded areas relate to the discussion of the effects of window opening on count rate (p. 37). For discussion of the plateau, trough, and Compton edge — see p. 21. Spectrum from ORINS-30 (6).

plot. The vertical axis is often labeled "count rate," whereas what is really meant is "count rate per keV of window opening," for if one widens the window it will accept more pulses, and therefore the width must be specified if the count rate is to mean anything.²² Let us suppose that the spectrum was constructed by taking counts with a 1-keV window, placed successively at the various energy levels covered by the horizontal-scale. If we need to know the count rate for a window spanning the range from 300 to 301 keV, we imagine a narrow, vertical strip standing with its lower end on the base line between 300 and 301 keV and reaching up until it hits the curve. The height of the strip is, of course, an index of the count rate, which can be read off against a vertical scale at the left.

Suppose that for the 300–301 strip the reading is 2000 counts/min, while that for 301–302 keV is 2100. If we plant a 2-keV window so that it spans from 300 to 302, clearly it will see all the pulses passed by the two 1-keV windows. Accordingly its rate will be $2000 + 2100 = 4100$ counts/min. Expanding the argument, for a 40-keV window we would be adding up 40 vertical strips instead of 2, and a consideration of the figure will show that the sum of all these strips will

²²See discussion of giraffe-height spectrum, pp. 18 ff.

be the area lying under the curve enclosed between the lower and upper limits of the window at 300 and 340 kev.

For any given spectrum plotted with linear coordinates, then, the conclusion is that no matter where a window is placed, or how large or small its opening is made, the count rate will be proportional to the area enclosed by the window — the area “at which the window is looking.”

Window Opening and Count Rate. — Let us apply the area concept to window opening, taking a simple case first. Referring to Fig. 21B, suppose we place a 40-kev window with its center at 140 kev; it will be looking at the heavily shaded area 40 kev wide. Now double the opening, keeping the window centered at 140 kev; the count rate will double, for you will be adding on at each side strips of area that match the two central ones pretty closely, and the area therefore doubles. This will happen wherever the spectrum is uncurved; it can be tilted and the area will still double, since a tall strip at the left, for example, will compensate for a short one on the right. In a region where the spectrum is straight, therefore, count rate will be proportional to window opening to a reasonably good approximation. This will be expected to hold for the Compton plateau or for a

background spectrum if the window is of no more than moderate opening, and it will be nearly true for narrow windows almost anywhere.

Now go back to Fig. 21B, and shift operations from the plateau to the primary peak. If you center a 40-kev window at 320 kev and then widen 20 kev on each side, you will not double the area, for the lightly shaded strips added on are not as tall as the dark, central ones. In this case, therefore, the doubled opening will see less than twice the area of the narrower one; count rate will *increase with* window opening, but it will *not be proportional* to it. Flatness or insignificant curvature in the spectrum makes for the proportional relationship, whereas “structure” destroys it. Over a hump, the count rate will increase more slowly, proportionately, than the window opening; over a trough the opposite will be true.

Window Opening and S/N Ratio. — Figure 22 illustrates the operation of the foregoing theory in an actual case: the counting of a weak cesium-137 source in the presence of an unstructured background. The areas concerned were measured approximately with a planimeter, and the results are given, in arbitrary units, in the small table at the left of the figure. Rounding off the figures, which can't be too accurate anyway, it is seen

Fig. 22. Showing improvement in the signal-to-noise ratio (S/N) when a wide window is narrowed. The cost is high, however, in terms of lost counts — see table. The coordinates are linear. From ORINS-30 (6).

that when a narrow window was made six times as wide, the background became about six times as high; the signal count, on the other hand, went up to only three times its former level, so that the S/N ratio was cut nearly in half. The lesson is clear: *IF* the counting problem calls for improvement of the S/N ratio and the background is unstructured, center the window over a high primary peak and keep it reasonably narrow. This will throw away counts, of course, but the narrowing of a wide window throws away background faster than signal.²³

Window Opening and Statistics: S^2/B . — Now turn to the other situation in which noise is a menace: namely where it is well behaved, measurable, and constant in quantity, but simply too plentiful, so that it threatens to drown the signal. Here the gross count is only a little greater than the background, and the net count is small. The problem is how to measure, with acceptable accuracy, a small difference between two larger quantities, each of which is subject to random fluctuations that threaten the small difference with very large percentage errors. As an example, suppose we can measure two quantities with unusually good reliability (standard deviation 1% of average) and that one of them is 105 units and the other is 100. If the variations are random, statistical theory says that²⁴

$$(105 \pm 1\%) - (100 \pm 1\%) = 5 \pm 29\%.$$

Clearly the error can skyrocket when the difference is small.

²³ There must be plenty of incoming photons if a narrow window is to be practical. Try to get adequate activity into the source, to get good detector efficiency, to minimize background by effective shielding, and so forth. Moreover, don't narrow the window too far. When its edges cut the peak at 70% of the peak value, the S/N ratio has reached about 90% of its maximum. There is little to be gained, therefore, by narrowing further and much to be lost in terms of falling count rate.

²⁴ S.D. for a sum or difference:

$$(S.D._{a \pm b})^2 = (S.D._a)^2 + (S.D._b)^2.$$

Or, in pictorial form:

Let us review standard deviations for a moment. The S.D. is a measure of variability and is derived, properly speaking, by considering the fluctuations in a number of measurements of the same thing. One can't properly talk about the S.D. of a single measurement, but if the counts come in with random timing, as theory says they will, we can estimate the *expected* S.D. by taking the square root of the number of counts in the measurement concerned.

Variability doesn't mean much until it is related to the size of the thing that is varying. An S.D. of "50 counts" might mean 50 counts out of 100 (a very shaky measurement) or 50 counts out of 100,000 (a spectacularly good one). A more meaningful measure of reliability, therefore, is the *fractional standard deviation* (F.S.D.), which is simply the S.D. expressed as a fraction of the average. It can be a percentage if you like, as in the example just cited. The statisticians call it the *coefficient of variation*, but surely the self-explanatory term is better.

If we count a sample and get N counts, then, the square-root rule says that if the count were repeated 20 or 30 times, the S.D. would turn out to be $\sqrt{N}$, approximately. If we increase the counting time, so that N goes up, the S.D. will also go up, but the F.S.D. will come down, for

$$F.S.D. \approx \frac{\sqrt{N}}{N} \approx \frac{1}{\sqrt{N}},$$

which will decrease as N increases. Accordingly a long count, or anything else that increases the total number of counts accumulated, should improve the reliability, *other things being equal*.

What about the increased count rate delivered by a wide window? We remember that a wide window reduces the S/N ratio, but perhaps the increasing count rate could more than offset this drawback. There might be an optimum window opening for use when the gross count is only a little above background, a compromise between rising count rate and falling S/N. We should look into this.

First we will ask the statisticians what relationship we should achieve between gross count (G) and background (B) in order to minimize the uncertainty in the net sample count, $S = G - B$.²⁵

²⁵ Note that we are still dealing with accumulated counts, not count rates, and that $S = G - B$ only if the sample and background have the same counting time.

This done, we can worry about the kind of window opening that will best meet their requirements. The S.D. for S can be calculated if we know those for G and B ; we can't measure these with single counts but we can estimate what they would probably turn out to be by using the square-root rule. Accordingly,

(S.D. for net count)²

$$= (\text{S.D. for gross count})^2 + (\text{S.D. for background})^2$$

$$\approx \sqrt{G^2} + \sqrt{B^2} = G + B = (S + B) + B.$$

$$\approx S + 2B.$$

Therefore, expected S.D. $\approx \sqrt{S + 2B}$, and expected F.S.D. $\approx \sqrt{S + 2B}/S$, where the sample and background are counted for equal times, S being the net sample count and B the background.²⁶

Now the sample spectrum will contain prominent peaks, while the background should not; consequently S and B will behave differently toward window opening. Not long ago, when we were searching for a defense against a moderate but variable background, we looked for a window opening that would maximize S/N ; here, where the gremlin is a formidable but constant background, the need is to minimize $\sqrt{S + 2B}/S$.

One can tackle this problem with the help of two assumptions, both pretty reasonable. One is that the background count rate will be approximately proportional to the window opening, for reasons that have already been discussed (p. 37). The other assumption is that the shape of a total-absorption peak will approximate that of the bell-shaped "normal curve of error," whose mathematical properties are well known. This assumption will hold quite well for the right-hand slope of the peak, and it will also hold well enough for the left-hand slope unless escape phenomena and/or scattering markedly distort the low-energy side. With these premises, then, we can calculate S and B for a series of window openings, determine

the F.S.D. in each case, and find out whether it goes through a minimum.

It does indeed, and the statistically optimum window is shown in Fig. 23, looking at a chromium-51 peak. One might call it the "window of desperation," to be used when the signal is almost drowned in the background mud. The sill and top cut the primary peak at the two points where the count rate is about *one-third of maximum*. This makes the opening some 25% greater than FWHM, the full width at half maximum. Ironically, this window will catch only 85% of the counts in the total-absorption peak; if you chase after the other 15%, the small profit in signal is wiped out by the more rapidly mounting background, and the statistics get worse. When the signal is very weak, then, 85% window efficiency provides the best way out of a bad business.

By one of the few lucky breaks in this counting game, the optimum window setting is not at all critical: the opening can be made somewhat narrower or wider without much change in the reliability. This is shown in the curves of Fig. 24, where the estimated F.S.D. (%), representing the uncertainty of a count, is plotted as a function of window opening. The curves all start high where

"85% WINDOW" FOR BEST STATISTICS

Fig. 23. For a statistically optimum window, make the top and sill cut the peak at about one-third of the peak height. The statistics deteriorate — though not much — if you try to catch the remaining 15% of the total-absorption pulses (see Fig. 24). Spectra by courtesy of ORINS Medical Division; from (13).

²⁶ Arising from this relationship is the " S^2/B rule," which says that when the gross count is not much above background, the ratio of the *net* sample count, squared, to the background should be made as large as possible. Reason: when S is small, $S + 2B \approx 2B$, in which case F.S.D. approximates $\sqrt{2B}/S$. To minimize the F.S.D., therefore, maximize S^2/B .

Fig. 24. Statistical uncertainty, as gaged by the expected fractional standard deviation (%), is plotted against window opening, expressed in terms of the full width at half maximum (FWHM). Three pairs of curves are shown, each assuming an arbitrary signal-to-background ratio, as indicated. Solid lines assume equal sample and background counting times; dashed curves show the 20% or so improvement that is obtainable if background is counted five times as long as the sample. The primary peak is the normal curve of error, on linear coordinates; the F.S.D. curves are semilog.

the window is narrow and come down (i.e., get better) as the opening increases. They hit a broad minimum, however, where the opening is between 1 and $1\frac{1}{2}$ times FWHM; so anywhere in this range the window will be operating near its best. We shall see in the following section that wider openings may be preferred if stability is a problem.

The curves of Fig. 24 are plotted for three values of S/N as measured at maximum — that is, with a narrow window; the conclusions regarding the optimum opening are the same for all. The solid lines are based on equal sample and background counting times; the dashed curves show the improvement that can be obtained if the background is counted five times as long as the sample. A simple development of the F.S.D. formula permits

the use of count rates and counting times instead of total counts, thus providing for situations where sample and background times are different. The relationship is

$$\text{expected F.S.D.} \approx \frac{1}{R_s} \sqrt{\frac{R_s}{t_s} + \frac{R_b}{t_b} \left(1 + \frac{t_s}{t_b}\right)},$$

where R_s = net counts per minute for the sample, R_b = background counts/min, and t_s and t_b are the two counting times, also in minutes.

It is seen that the fivefold background time can reduce the uncertainty to around 80% of what it would be with equal times. But there is a law of diminishing return: a tenfold background time would not produce very much additional improvement. In practice, then, a "long background" should last five times as long as the sample counting time, unless additional background time costs you nothing. Don't forget that the longer you run the background, the greater is the chance that something may happen during the count to make the result unrepresentative of the background that was present while you were counting the sample. Long backgrounds are valid only when they are steady. If the background is variable, the best thing you can do is to track the gremlin down and slay him in his lair. Failing that, all you can do, when the sample is very weak, is to use equal sample and background counting times, taking half of the background count just before the sample, and the other half just afterward. If you can't make the two backgrounds agree, you quit and go home.

Window Opening and Count-Rate Stability. —

Count rate will necessarily depend on window position, but we would like it not to be *critically* dependent; so we might look for a window-width policy that makes the system unfussy about window position. Either the gain or the window position may drift, and some drifting is to be expected, since an electronic instrument can rarely hold its adjustment indefinitely. Line-voltage changes may exceed the instrument's compensating capacity; a photomultiplier tube may be temperature-sensitive; someone may move a magnetized screwdriver (too many tools are magnetized these days) close to an unprotected phototube, thus changing its gain; and so on. Consequently an unanticipated gain change is a very real possibility, and one should design operating policies to minimize the harm that might result.

First consider a drift in window position, bearing in mind the area principle (p. 36). If a window is kept centered over a primary spectral peak and the opening is widened, strips of area will be added on at each side, and the count rate will rise. If, on the other hand, the opening stays the same but the position shifts, area is gained at one side but lost at the other, and the effect on count rate will depend on whether or not the gain exceeds the loss. Count-rate stability will be best if we can make gain and loss about equal. Figure 25 shows an iodine-131 spectrum with a window centered, initially, on the main primary peak. But the window then drifts upward, gaining area on the right while losing it on the left. It is clear that the unhatched rectangle designated "net loss" represents the change in area, for the other area components approximately cancel each other. Our concern is to keep the net-loss rectangle small.

This rectangle has width and height. The width represents the amount of drift, and although we try to keep this small, the reason why it's there is because we couldn't prevent its occurrence; so let us turn our attention to the height, about which we might be able to do something.

For a given amount of drift the height of the net-loss rectangle depends on how steep the sides of

the peak are at the points where the curve is cut by the vertical lines representing the window's top and sill. Clearly we should keep away from the steep places in the spectrum and plant the window's top and sill where the curve slopes only gently — namely either well up near the top of the peak or well down near the bottom. Either of these should keep the rectangle small.

We can eliminate the high-up alternative right away, for the narrow window has two obvious faults. In the first place it will be poor statistically, since the F.S.D. will go up when the window is narrowed (Fig. 24). Second, to get good stability the net-loss rectangle should not only be small, it should be a small *percentage* of the total area seen by the window. If this area is small, as it will be when the window is narrow, even a small rectangle could represent a sizable fraction of the total. In contrast, the much larger area under a wide window makes a small rectangle still less important. There is little doubt, therefore, that the wide window will be better.

Fortunately there is no serious conflict between a window designed for stability and one giving best statistics; the broad minima in the F.S.D. curves (Fig. 24) provide a generous amount of leeway, and there will be little statistical pain if

ORNL DWG. 65-12645

Fig. 25. When a properly placed window drifts upward, area is gained at the right but lost at the left. Two triangular areas for gain and loss about balance each other; so do the gain and loss rectangles. Net loss is therefore the clear rectangle. Iodine-131 spectrum from a patient's neck; from ORINS-30 (6).

the window is made somewhat wider than the ideal "85%" one. Moreover, the statistics will be of major concern only if the signal is so weak that background threatens to drown it, so in most situations maximum window efficiency is unnecessary. If there is a bothersome stability problem, therefore, the window limits should be made to cut the curve low down on the peak, fairly near the base line. The system will then be unfussy about drift, the window efficiency will be high, and the statistics will still be close to the best obtainable under the circumstances.

Note that for a given window and peak, maximum count rate will occur when the window's top and sill cut the sides of the peak at equal heights above the base line, since with such a window a net-loss rectangle will appear no matter which way the window is shifted. Unless the peak is symmetrical, therefore, the center line of a wide window, adjusted for maximum count rate, will not pass through the tip of the peak. Stability is best, of course, when the window is adjusted for maximum rate, for that is the point where an imaginary spectrum, run with this wide window, will be almost flat.

Now let us turn from the problem of a drifting window to that of drifting amplification. In the former case, illustrated in Fig. 25, the window moves relative to the spectral peak, whereas when the gain changes the peak creeps out from under a stationary window. These two kinds of drift are not strictly comparable, although in most instances they are approximately so. If a drifting peak could move to one side without changing its size, one could fairly apply the same argument that is used for the drifting window, but the peak *does* change its size as it moves. If it shifts to the right, it becomes shorter and fatter; to the left it grows taller and thinner.

Figure 26 shows how the spectrum of chromium-51 changes when the gain is increased 33%. The 320-kev peak, originally at 32 v, moves out to 43 v, and its height is reduced at the same time. This happens because the total area under the curve must remain the same, since it represents the total number of scintillations occurring in the crystal per minute, irrespective of how tall or short the resulting electrical pulses may be. These scintillations depend only on the incoming radiation, and not on how the operator chooses to amplify the pulses. Although the gain may change, then, the area under the spectrum does

Fig. 26. Effect of increased amplification on a chromium-51 spectrum. The peak moves to the right in direct proportion to the gain, but, because the underlying area must remain the same, the peak height varies inversely as the gain. Measures designed to protect against drifting gain should take account of this. From ORINS-30 (6).

not. Consequently, if the peak moves to the right in the ratio of 4:3, its height must shrink in the ratio of 3:4.

This circumstance permits the gain-change problem to be explored mathematically. It turns out that for protection against drifting gain the window's sill and top should *not* cut the peak at equal heights above the base line; the sill should cross the curve at a higher level than the top, in such a way that the ratio between the count rates at the sill and top is the same as that between the voltages at the top and sill (Fig. 27, bottom). This ratio will be near unity if the window opening is not large compared with the sill height, and in such a case a setting designed to protect against gain change will also be pretty good for window drift. As the lower energies are approached, however, the spectrometer's resolution becomes poorer and the total-absorption peak therefore becomes proportionately broader. The ratio of V_2 to V_1 (Fig. 27) then becomes considerably greater than 1, and when this happens a single window placement will not be ideal for both kinds of drift. In such cases it is usually best to protect against gain change, since this is the more likely to occur. To this end, maximum count rate should be found by adjusting the gain rather than the window. It is important to bear this in

ORNL DWG. 65-12640

Fig. 27. Top: If window position is likely to drift, the top and bottom of the window would cut the peak at equal heights above the baseline ($R_1 = R_2$). Bottom: If the gain is likely to drift, the relationship shown in the lower figure should be used. This becomes important when the peak is broad — for example, at the lower energies, or in beta counting.

mind during beta scintillation counting, for there a broad low-voltage hump is the rule.

Prevention of Drift. — From what has been said it will be clear that it is far better to prevent drift than to try to live with it. Apart from competent instrument design, steady ac power is the most important preventive. A voltage-regulating transformer will be needed more often than not. Bear in mind, however, that to deliver a constant output voltage the transformer requires a constant load; so do not load it with devices that get turned on and off. If necessary, isolate the spectrometer and its high-voltage supply and run only them from the regulating transformer, using the raw ac line for uncritical, intermittently operating components such as a strip-chart recorder, printout typewriter, X-Y plotter, tape punch and reader, and so on. Also, take care to procure a transformer with the

proper volt-ampere capacity: for effective performance the transformer should be loaded to between 67 and 97% of its full rating.

If the line-voltage variations are too severe to be coped with in this way, try to power your counting instruments from a private line whose use is under your control, for then you can prevent its being subjected to heavy, intermittent loads — e.g., from electric ovens, refrigerators, air conditioners, x-ray equipment, elevator motors, air compressors or floor buffers.²⁷ Make the connection for this line as far back toward the power station as you can, for there the wiring will be heavy, and other people's intermittent loads, which you can't prevent, will be a smaller percentage of the total wattage available at the point where your line ties in. You may have to fight savagely to obtain such a line, for it will be a rare administrator who understands why you need it, and a rare electrician who won't hold out for a conveniently short connection.²⁸

Constant temperature in the counting room is also important, but *do not* operate either a heater or an air conditioner from the instrument line!

Instability of Window Opening. — We have been worrying about unsolicited changes in window position, relative to the primary peak of interest. What about similarly spontaneous changes in window opening? These are much harder to detect; can we devise any defense against them?

The defenses are two in number. The first is the *wide window*, which we have already decided to use anyway for most quantitative counting situations. When the window is wide, the enclosed area is large, and the top and sill cut the spectrum at a low level. Consequently an uncompensated area change produced by a change in window opening (1) will be small, because the rectangle involved (Fig. 25) is not tall, and (2) will be a

²⁷ Buffers are an intolerable abomination. They not only throw heavy loads on and off the power lines, making life miserable for any laboratory that needs a steady line voltage, but they scatter spore-bearing dust throughout the hospital, they greatly increase the slippery-when-wet hazard, and they merely make a floor look clean instead of being clean. They have no business in any hospital, and an outright ban is long overdue.

²⁸ With annoying stubbornness, one of us insisted on a restricted instrument line for the low-background, whole-body-counting installation at ORINS Medical Division, and this provision has paid off handsomely in terms of improved voltage stability. The 120-v, single-phase power is derived from a small step-down transformer fed directly from one of the high-voltage distribution lines outside the building.

small fraction of the total enclosed area. Thus the damage done will be minimal.

The second defense is the practice of *counting by comparison*, which, similarly, is a good idea for several reasons. The net count rate from an unknown sample is compared with that obtained from a known quantity of the same nuclide, counted under the same or equivalent conditions. If conditions are equivalent the spectra will be similar, in which case a minor shift in either the window or the peak will have nearly the same effect on sample and standard alike, and the ratio between the counts will remain unaltered. It is easy to understand, on the basis of the area concept, why the two spectra must be made similar. Sometimes, of course, it is difficult to equalize the two counting situations.

Whenever two counts are to be compared, the smaller the time lapse between them, the less important any drifting problem will be.

Sometimes an operator is interrupted between the unknown and standard count, so that the latter is delayed. Upon returning to the instrument, she may find that the peak has moved to one side. The proper procedure is to correct the gain so that the peak is brought back where it belongs, rather than to search out the displaced peak with the window. This is because the number of kev per *E* division changes when the gain changes; therefore when a window of fixed width is moved to the new peak position, it will not be covering the original energy span. The resulting error in count rate may be quite unimportant, of course, if the window opening is wide.

Asymmetrical Window Settings. — Sometimes special circumstances make it advisable to set the window somewhat higher than the normal, stability-oriented position in which the sill and top cut the primary peak at equal heights above the base line. The asymmetrical variant should be considered (1) if scattered radiation is highly objectionable, and (2) if the energy is sufficiently low to make it difficult for the spectrometer to reject the scatter (12). We have pointed out (p. 21) that a spectrometer's resolution deteriorates as the energy goes down; to make matters worse, the energy band into which a photon may be scattered in a single Compton collision becomes narrower as the primary energy gets lower, so that the scatter band snuggles up close to the primary line (ref. 12: Fig. 3, p. 112). It follows that differentiation between a primary and a scattered ray requires greater precision just at the time when

the spectrometer's precision is falling apart. The rejection of scattered energies lying close to the primary will be assisted if the window is shifted a little to the right, so that its sill rides well up on the left-hand limb of the peak. The price paid for the improved S/N ratio is a somewhat reduced count rate, plus increased vulnerability to anything that causes drifting; consequently, this expedient should not be used unless it is needed. By the same token, the primary energy should not be any farther below about 130 kev than is necessary, and the possible drawbacks should be borne in mind whenever a switch to a low-energy gamma emitter appears desirable for other reasons.

Scans illustrating the cleanup value of the asymmetrical window are shown in ref. (12). Conventional and asymmetrical window settings have been tabulated for several nuclides of clinical interest (5).

Plotting a Pulse-Height Spectrum. — The intelligent way to tackle any new or difficult counting problem is to see first of all what the sample and background spectra are like in the actual counting situation that you propose to use — and, of course, to take remedial measures if they are unfavorable. The spectra are all the more important if you are trying to count more than one nuclide. Sometimes a spectrum is needed to check the radiochemical purity of supplied material or to establish the identity of an unknown nuclide when evidence of contamination suddenly appears, especially if the contamination is in ante-mortem personnel. In such situations a spectrum is almost always helpful and can be of vital importance.

The reader won't need to be told that for spectral plotting a narrow window is essential, for a large window sees a wide band of gamma energies and the details of the spectrum become "smeared," to borrow an unlovely term from the engineers. The narrow window brings with it the drawbacks of poor count rates and fussiness about drift, and the latter promptly contraindicates long-counting-time therapy for the low count rate. It is a good plan, therefore, to start with a sample containing a generous amount of activity²⁹ and to take all

²⁹Like many other things, generous source activity can be overdone. A high-activity source can be beating the brains out of the phototube and amplifier, while at the same time the pulse rate coming through the window may be low because of a narrow opening. Remember that the window can protect the rate meter from overload, but it doesn't protect the detector or amplifier (Fig. 8).

practicable measures to assure constant room temperature, adequate warmup time, well-stabilized ac power, and so forth. If appreciable background is present, it must be determined and subtracted out for each window position; this is such an endless chore that time devoted to reducing background to a negligible level is likely to be well spent.

If the spectrum is to be recorded automatically, using a rate meter and an electrically driven, moving window, be sure that the speeds of response in the rate meter and recorder are adequate to minimize distortion of the steep sides of a peak. Test this on the sharpest peak in the spectrum. A response that is too slow will (1) reduce the peak's true height and (2) shift its position toward the end of the recording. Both of these can be checked by means of one or two stationary counts, and if an error is found, the time constant of the recording system should be reduced. A whiskery spectrum, otherwise free from distortion, is preferable to a beautiful, smooth curve that lies. Too often there is deception behind beautiful, smooth curves.

If you need spectral information in the low-energy range, it is well not to trust the window at low *E*-dial settings. Use more amplifier gain, so that the pulses of interest can be seen at *E*-dial settings above 50 divisions (on a 1000-division scale). You can also use more high voltage on the phototube, but *first look up the maximum voltage* that your tube will tolerate. High gain, of course, can precipitate noise problems, and it may become necessary to find a quieter amplifier and/or phototube. Analyzers vary considerably in regard to noise and also to the *E*-dial reading below which the window will become unreliable.

Multiple-Nuclide Counting. — When two radio-nuclides are present in the same sample, there is all the more urgent need for careful use of the spectrometer's window. Corresponding to the two primary peaks selected for counting, there will be an *upper window* for the "upper nuclide" and a *lower window* for the "lower" one. For illustration we might take iron-59 and chromium-51, two nuclides that are frequently used in blood studies (Fig. 28). The high-energy photons from the iron component produce many substandard pulses, some due to escape from the crystal and others to predetector degradation. It is against this barrage of iron-derived pulses, thrown down

Fig. 28. Double-nuclide counting is illustrated by iron-59 and chromium-51. The lower window will see room background as usual. In addition, however, there will be a *cross-talk background* thrown into the lower window by the upper nuclide, and this cross-talk will be proportional to the amount of upper nuclide present. The room-background component should, of course, be constant. Spectra from ORINS-30 (6).

into the lower window from above, that the radiation from the chromium-51 must be counted. We may call this added noise the *cross-talk background* to distinguish it from the ordinary room background that will also contribute to the gross count in the lower window.

The cross-talk background will have the virtue of being steady, but quantitatively it will be unusual in being proportional to the amount of upper nuclide present. The way to deal with it, therefore, is to procure a pure sample of the upper nuclide and count it with both windows, subtracting out the upper and lower room backgrounds to get net counts. The measurements must be made using the same detector and geometry (e.g., a well counter) that will be used later in the actual experiment, for otherwise the spectra are likely to change. This determination tells the operator that every time he has 100 net counts in the upper window he will be throwing *x* counts of cross-talk background down into the lower window, and once this ratio has been established, the net upper-window count in an unknown mixture can be used to calculate the cross talk and subtract it out. The percentage of cross talk will depend, of course, on the spectrum of the upper nuclide and on the chosen positions and openings for the two windows. The latter must therefore be selected

with considerable care, perhaps after one or two pilot experiments. The windows will be different, naturally, for each pair of nuclides.

Much of the success of a double-nuclide determination will hinge on whether, by careful planning, one can arrange for a reasonable sample-to-background ratio in the lower window. The problem becomes additionally complicated if the selected peaks are so close together that the lower nuclide also cross-talks into the upper window, as ^{51}Cr might well do into a conventional ^{131}I window. This creates a situation so unfavorable that evasive tactics may well be the better part of valor — for example, the measurement of the ^{131}I via its topmost two gamma rays (Fig. 9), which are normally shunned because their abundance is low. Double counting is discussed in greater detail in the forerunner of this account, ORINS-30 (6).

Large Windows and High Energies. — Many spectrometers have ΔE dials that give a maximum window opening equal to 100 E -dial divisions, and sometimes a need arises for a pass band greater than 100 kev. This is easily obtained by reducing the amplifier gain so that each E division becomes equal to 2 kev instead of the usual 1 — or, for that matter, any other calibration ratio that the operator finds convenient. These reduced gain settings also permit a window sill to operate in the regions above 1000 kev, which is the top of its range in many spectrometers if one uses the ordinary 1:1 calibration. The only precaution to bear in mind is that usually the ratios on a step-wise gain control are not intended to be accurate, and one should always restandardize a new gain using a single gamma peak of known energy. Again, as recommended earlier, do the final trimming of the gain on the actual peak that you intend to count.

Threshold or "Integral" Counting

Some spectrometers are provided with a switch that paralyzes the window-top mechanism but not the window sill; this permits an operator to count all gamma rays having energies greater than a chosen threshold value. This feature will not be wanted often because the window type of operation is normally preferable. It may be helpful, however, to include here a discussion of the threshold mode of counting, since many commercial scintillation counters are not spectrometers

and will operate only on the threshold basis. In many situations these instruments are perfectly capable of doing a creditable job, provided only that they are operated with intelligence and insight. Ironically, they are less likely to be operated properly than a spectrometer is, for the simplicity of the controls on a threshold machine invites people to assume that operation is easy. Actually it is more difficult, since it is harder for the operator to find out what his instrument is doing — which energies it is accepting and which it is not. The ludicrous result is that here we have a widely popular instrument, being operated, more often than not, with high confidence and low reliability.

Figure 29 shows the differential (i.e., window) and integral (threshold) spectra for a small source suspended in air except for a plastic container 3 mm thick. The nuclide is barium-133, which emits several gamma rays plus a 31-kev x ray from the cesium daughter. For the differential spectrum the count rate was recorded continuously with a

Fig. 29. Differential and integral spectra for a small source of barium-133 in air. Note that the integral curve rises more steeply wherever the differential spectrum shows a maximum. The integral spectrum is not flat anywhere, which invites stability problems. It was recorded at about $\frac{1}{16}$ sensitivity to keep it on the chart. Spectra by courtesy of ORINS Medical Division; from (13).

rate meter while the window, 5 kev wide, was swept slowly to the left at a constant speed of 300 kev/hr. When the differential spectrum was finished, the window-top circuit was inactivated, and the window sill, now the threshold, was again swept downward at the same speed, thus providing the integral spectrum shown in the same figure. For the latter the sensitivity was cut to about $\frac{1}{16}$ of its former value in order to keep the curve on the chart, for removal of the window top always makes the count rate higher, especially when the threshold is low.

It is not difficult to figure out what is going on in a threshold spectrum. Let us return to the differential spectrum and center the window on the topmost peak, spanning 360 to 365 kev. It will be looking at a tall, narrow strip of area between this peak and the base line. Now leave the sill where it is and start moving the window top to the right; the area will increase and the count rate go up. This process will continue, but as the window-top line cuts the curve at lower and lower levels, there is less and less area to add on; so the count rate increases only up to a point and then levels off, when the window top is well over to the right. At this juncture the window top might just as well not be there, for all the difference it is making to the count rate. It rejects a few high-energy background pulses, and that is all.

For practical purposes, then, this is very nearly the situation that exists when a threshold is located at 360 kev and there is no window top; the system is seeing all of the area lying under the differential curve and to the right of the threshold line. If we now move the threshold toward the left, the area will increase, rapidly where the differential curve is high above the base line (because there the strips of area being added on are tall), and slowly where it is lower down. As the threshold moves to the left, accordingly, the area always grows, and therefore the integral curve is always climbing. This is why there are no peaks in the integral spectrum.

Note that the positions of the peaks in the differential spectrum are signaled in the integral spectrum by places where the curve rises steeply. This is because it is opposite the peaks that area is being added on most rapidly by the moving threshold line. Thus the integral curve is steep near 363 and 85 kev, where the differential spectrum is well above the base line; in fact, the integral curve is equally steep at these two

places because the two peaks in the differential spectrum are about the same height. The integral curve is still steeper in the 31-kev region, where the peak is higher. Conversely, it slopes only gently at 450, 240, and 10 kev, where the differential spectrum is low.

It is the absence of spectral peaks that makes the calibration of a threshold instrument's *E* dial (if one is provided) awkward. In a spectrometer all one has to do is search out a dial setting for a maximum count rate, whereas in a threshold instrument one has to find a setting where the rate is rising most steeply, and this is much more of a chore. It is such a nuisance, in fact, that many a technician simply won't bother with it; she runs her instrument on a basis of arbitrary dial settings, without having any good idea of what is going on. And the threshold, unchecked, will drift.

Operating Policies. — In some threshold instruments the only threshold control is a screwdriver potentiometer located inside the cabinet, but suppose you have an instrument with a proper control on the panel, and suppose that you know how to calibrate this dial in terms of energy equivalent. You have the privilege, then, of placing the threshold wherever you like. How do you decide where to put it?

As in the case of the spectrometer, the policy depends on what kind of problem is plaguing you. For *scatter rejection* a threshold system can be nearly as effective as a window: the threshold is placed in the trough just to the left of the primary peak, and nearly all the scatter will fall below this level. In a trough, of course, the slope of the integral curve will be gentle. To deal with a *background* problem, this same location will probably be about as good as any, although the only way to make sure is to pick several likely-looking threshold settings, determine S^2/B for each of them (p. 39) and then select the best one. Background is likely to be more of a hazard in a threshold instrument than in a spectrometer because the absence of a window top allows a wide band of high-energy counts to pour in. One can only hope that they won't be too numerous. *Stability* is likely to be the most urgent problem, for the continually climbing integral curve (Fig. 29) provides no flat place anywhere, and therefore stability will never be as good as in a properly operated spectrometer. The best you can do is to pick out the place where the *percent change in count rate*, per kev of threshold shift, is lowest.

Anything you can do to prevent threshold shift or gain change will be highly desirable.

In many commercial threshold counters there is no panel *E* dial, but there is an internal threshold, to be left at a constant setting, and a continuously variable high-voltage control. Since the high voltage controls the phototube gain, the higher the gain is, the more pulses will reach above the threshold and thus be counted. A plot of count rate vs high voltage will look somewhat like the mirror image of the integral spectrum shown in Fig. 29: the high count rates will be toward the right, since the higher the voltage is, the more counts there will be. But the relationship between count rate and high voltage is not linear (it is a power function — see p. 10); so the curve will be compressed laterally at the right-hand side. Nevertheless, control over effective threshold via the high voltage is better than nothing. Bear in mind that the gain changes very rapidly with high voltage.

Virtues and Vices. — With the foregoing considerations in mind we can summarize the advantages and drawbacks of the threshold counters. They are electrically simpler than the spectrometers, which makes for lower initial cost and less frequent component failure. The count rates with threshold machines are higher, other things being equal, for inevitably a wide band of energies is accepted. Against this, the background rates will also be higher; so the increased count rate will be helpful only if one can show that S^2/B has actually gone up. As in the spectrometer, a characteristic threshold position will be needed for each nuclide; but whereas dial calibration is quick and easy with a spectrometer, it is a time-consuming nuisance in the threshold instrument, and a run-of-the-mill operator is tempted simply to skip it, letting quality of performance look out for itself. Stability will be inferior in threshold operation, for the integral spectrum has no flat regions. Compared with spectrometry, therefore, correct operation in the threshold mode is simultaneously

simpler and more difficult. One is reminded of another simple instrument, the violin; it is easy to play badly, but difficult to play well.

CONCLUSION

The foregoing account is addressed to doctors and technicians who are interested in the medical uses of radionuclides. It is intended to help them understand what a pulse-height spectrometer is all about, what it can do for them that will be helpful, and what it should not be expected to do for them. We can summarize its virtues and vices as follows.

It offers the operator:

1. a means of finding out, easily, which gamma energies the instrument is counting, and the privilege of adjusting the energy pass band to suit the counting problem at hand;
2. relative freedom from errors due to uncontrolled scattered radiation;
3. an improved signal-to-background ratio, since much of the background is out of sight;
4. improved stability, when the instrument is operated correctly;
5. greater freedom in using two nuclides together, since the instrument can, within limits, pick out the characteristic radiations individually;
6. a powerful analytical tool that can check the radiochemical purity of supplied material, or establish the identity of an unknown emitter, by revealing the gamma-ray spectrum.

On the other hand, these advantages become available only at a price. A spectrometer is somewhat more expensive than an ordinary "threshold" scaler, which, in effect, has a window sill but no window top. There are more components in a spectrometer to get out of order. It requires careful and informed operation, for it is a booby trap to the slapdash technician. And its owner is vulnerable to prowling knob twisters.

REFERENCES

- (1) Anderson, E. C., R. L. Schuch, J. D. Perrings, and W. H. Langham, "The Los Alamos Human Counter," *Nucleonics* 14(1), 26 (1956).
- (2) Compton, A. H., "A Quantum Theory of the Scattering of X-Rays by Light Elements," *Phys. Rev.* 21, 483 (1923). Our equation derives from No. 3, p. 486, with wavelength converted to energy.
- (3) Crouthamel, C. E., "Applied Gamma-Ray Spectrometry," *International Series of Monographs on Analytical Chemistry*, Pergamon Press, New York, 1960.
- (4) Gibbs, W. D., "A New Large-Volume Radioactive Sample Counter," ORINS-43, p. 88 (1963). (Clearing House for Federal Scientific and Technical Information, National Bureau of Standards, Springfield, Va.). See also: "A Method for Precise Geometry-Independent Radioassay of Large Biological Samples," currently in preparation for the *Journal of Nuclear Medicine*.
- (5) Harris, C. C., "Certain Fundamental Physical Considerations in Scanning," p. 13 in *Scintillation Scanning in Clinical Medicine* (ed. by J. L. Quinn), W. B. Saunders, Philadelphia and London, 1964.
- (6) Harris, C. C., D. P. Hamblen, and J. E. Francis, Jr., *Basic Principles of Scintillation Counting for Medical Investigators*, ORNL-2808, bound together with: Ross, D. A., *Medical Gamma-Ray Spectrometry*, ORINS-30 (1959). (Clearing House for Federal Scientific and Technical Information, National Bureau of Standards, Springfield, Va.)
- (7) Harris, C. C., M. M. Satterfield, P. R. Bell, and D. A. Ross, "A Comparison of Various Methods of Rate Recording," in *Recent Advances in Nuclear Medicine* (ed. by M. N. Croll and L. W. Brady), Appleton-Century-Croft, New York, 1966.
- (8) Harris, C. C., M. M. Satterfield, D. A. Ross, and P. R. Bell, "Rectilinear vs. Stationary Scanners," presented at *A symposium on fundamental problems in scanning, Chicago, May 8-9, 1965*; Charles C Thomas, Springfield, Ill. (in press).
- (9) Heath, R. L., *Scintillation Spectrometry: Gamma-Ray Spectrum Catalogue*, TID-4500, 2d ed., vols. 1 and 2, (1964). (Clearing House for Federal Scientific and Technical Information, National Bureau of Standards, Springfield, Va.)
- (10) Neiler, J. H., and P. R. Bell, "The Scintillation Method," chap. 5, p. 245, in *Alpha, Beta, and Gamma-Ray Spectroscopy* vol. 1, (ed. by Kai Siegbahn), North-Holland Publishing Co., Amsterdam, 1965.
- (11) Rapkin, E., "Liquid Scintillation Counting 1957-1963: A Review," *Intern. J. Appl. Radiation Isotopes* 15, 69 (1964).
- (12) Ross, D. A., C. C. Harris, M. M. Satterfield, P. R. Bell, and J. C. Jordan, "Low-Energy Gamma Emitters in Scanning and Other Clinical Applications," *Radioaktive Isotope in Klinik und Forschung* 6, 108 (1965) (= *Strahlentherapie*, supplement vol. 60; Urban and Schwarzenberg, München-Berlin, 1965).
- (13) Ross, D. A., C. C. Harris, M. M. Satterfield, and P. R. Bell, "New Challenges in Clinical Counting," *Japan. J. Nucl. Med.* 2(2), 81 (August 1965).

**THIS PAGE
WAS INTENTIONALLY
LEFT BLANK**

INTERNAL DISTRIBUTION

1. Biology Library	74. H. C. Hoy
2-4. Central Research Library	75. G. S. Hurst
5-6. ORNL – Y-12 Technical Library	76. R. P. Jernigan
Document Reference Section	77. W. D. Jones
7. Thermonuclear Library	78. C. E. Larson
8-57. Laboratory Records Department	79. J. N. Luton, Jr.
58. Laboratory Records, ORNL R.C.	80. H. G. MacPherson
59. R. G. Affel	81. W. J. McClain
60. R. Baldock	82. R. E. Meadows
61. M. C. Becker	83. G. D. O'Kelley
62. H. Beutler	84. E. L. Olson
63. D. S. Billington	85. R. B. Parker
64. R. D. Birkhoff	86. J. A. Ray
65. C. W. Blue	87-91. R. H. Rohrer
66. T. A. Butler	92-141. D. A. Ross
67. D. L. Coffey	142. M. M. Satterfield
68. W. Davis, Jr.	143. M. J. Skinner
69. G. R. Dyer	144. A. H. Snell
70. J. S. Eldridge	145. D. A. Sundberg
71. B. R. Fish	146. W. E. Unger
72. J. E. Francis	147. A. M. Weinberg
73. W. A. Gibson	148. H. A. Wright

EXTERNAL DISTRIBUTION

- 149-272. Given external distribution
- 273-322. C. C. Harris, Division of Nuclear Medicine, Department of Radiology,
Duke University Medical Center, P.O. Box 3223, Durham, N.C. 27706
323. J. A. Swartout, Union Carbide Corporation, New York, N.Y.
324. Laboratory and University Division, AEC, ORO
- 325-589. Given distribution as shown in TID-4500 under Biology and Medicine
category (25 copies – CFSTI)