they are called today. What counts in high energy collisions is the *energy of collision* observed in that frame of reference in which the total centre of mass of the colliding particles is at rest. When a beam of particles strikes a stationary target, then at least half of the energy of the incident particle is used to move the centre of mass of the two colliding particles ahead, while only the remaining fraction of the energy is available for the collision itself. According to the laws of relativistic collisions, this useful fraction decreases continuously as the energy of the incident beam becomes larger. For example, the energy of collision of a 500 GeV proton striking a stationary target is only 31 GeV. If, on the other hand, two 500 GeV particles undergo a head-on collision, then the total collision energy available is 1000 GeV.

This enormous advantage in collision energy is somewhat offset by the fact that the density of practical particle beams is very much lower than that of ordinary matter. Thus while colliding beams produce dramatic gains in collision *energy*, they lead to large losses in collision *rate*. However, one cannot have everything, and the loss in collision rate can to some extent be compensated by surrounding the points of collision with detectors sufficiently large to catch almost all the fragments from the events which do occur.

I have digressed to this brief outline of the evolution of accelerating devices during the past fifty years in order to emphasize that, while our studies of the fundamental nature of matter have focused on structures of smaller and smaller size, the actual rate of progress in this field has been determined by the much more mundane matter of successive accelerator technologies. However, as with all exponential growth patterns, the dramatic evolution shown in fig. 15 must sooner or later slow down.

Among the colliding beam devices, the most productive has been the electron–positron collider. An example of such a device was shown in fig. 14. Here electrons and positrons counter-rotate in a ring, and the orbits of these particles are confined by a group of electromagnets. The more powerful recent electron–positron storage rings all use strong focusing magnets; in addition, large radiofrequency power systems are required to compensate for the loss of energy caused by the electromagnetic radiation emitted by the electrons as they traverse their circular orbits.

The advantage of carrying out high energy physics experiments with electron–positron collisions is the inherent simplicity of the annihilation process. As was indicated in fig. 13, the electrons and positrons indeed annihilate, resulting in what physicists call a virtual photon, which describes a state of pure electromagnetic energy. This electromagnetic energy in turn can rematerialize into any combination of particles which conserves both the energy involved and certain other of the symmetry characteristics of the initial collision. Since, unlike collisions in which protons strike material targets, the initial particles have completely disappeared, the final state can be particularly simple and therefore relatively easy to analyse. In particular, the final state can be pairs of particles and their antiparticles, whatever these particles may be. This, in retrospect simple, situation led to the spectacular discoveries of November 1974, aptly called the November Revolution.

7. The November Revolution

The November Revolution began with the simultaneous publication of two experimental results. The work of Samuel Ting and collaborators at the Brookhaven National Laboratory demonstrated that pairs of electrons produced from a beryllium target bombarded by 25 GeV protons in the Alternating Gradient Synchrotron exhibited a peculiar distribution. The correlation in angle and energy of the electron