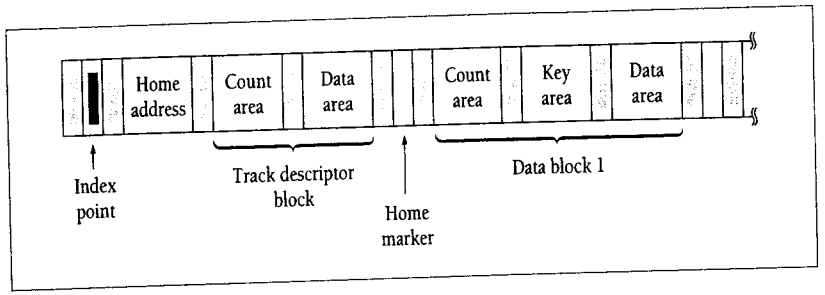

Figure 27 Count-Key-Data Track Organization

key length field shows the length of the key in the count-key-data track format. The data length field contains the length of the data area for the block. The cyclic check field is used to detect errors, as in the home address.

The first block in a track is special and is called the **track descriptor block**, designated block zero or R0. The count area of the track descriptor block contains the same information found in a typical count area. Here, the block number and the key length are zero. Also, the number of unused bytes after the last block on the track may be stored in the count area of R0. The data area of the track descriptor block describes the content of the track in terms of the number and length of blocks on the track.

All other blocks contain user data. The relative block number is indicated in the count area. The key area is optional. If specified, the organizational key for the file is duplicated in this area. This duplication costs storage, but somewhat speeds up processing. This is because the system does not have to go through additional processing to isolate the key in the main memory.

The data area in a block stores logical records. The number of logical records and the logical record length have to be specified by the programmer. Logical records in data blocks may be unblocked or blocked, and may have fixed or variable lengths. In variable-length logical records, for a fixed blocking factor the length of the data area in each block will be different. In variable-length record formatting, in addition to the logical records, the data area of each block contains a block length field, and for each logical record in the block, a record length field.

The programmer has the flexibility to specify a variety of record formats for logical records. The records are structured in each block according to the requested record format. In general, the following record formats may be supported:

1. Unblocked fixed-length (designated F)
2. Blocked fixed-length (FB)
3. Unblocked variable-length (V)
4. Blocked variable-length (VB)
5. Unblocked variable-length spanned (VS)
6. Blocked variable-length spanned (VBS)
7. Undefined (U)

In the **unblocked fixed-length (F)** format, each data block contains a single fixed-length data record (Figure 28). **Blocked fixed-length (FB)** record arrangement is shown in Figure 29. Here, two or more logical data records of the same length are blocked together and stored in the data area. The key of each record is kept as part of the record in the data area, and the logical records are arranged in increasing order by key. If the key area is requested for track formatting, the largest value is duplicated in the key area.

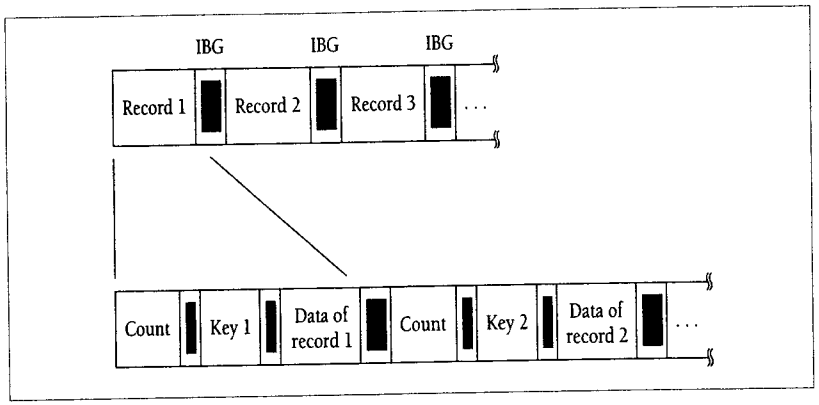

Figure 28 Unblocked Fixed-Length Record Format (F)

Figure 30 shows the **unblocked variable-length** (V) format. There is only one variable-length record in each data block. In the data area of each block, in addition to the actual logical record, block length (BL) and record length (RL) are recorded.

In the **blocked variable-length (VB)** format (Figure 31), the data area of each block contains two or more logical records of varying length. It also contains a block length field and a record length field for each logical record. The computer system uses these fields to separate logical records in a block in the main memory. Again, as in the FB format, the records are arranged in each block in increasing order of their keys, and the largest key is duplicated in the key area.

The **unblocked variable-length spanned** (VS) format is used if a logical record is very long. In such a case, it may be preferable to break it into two or more segments and store each segment in a separate block. When a logical record continues across block boundaries, we call it **spanned**. Usually, spanned formatting requires variable-length records. In the data area of each block are a

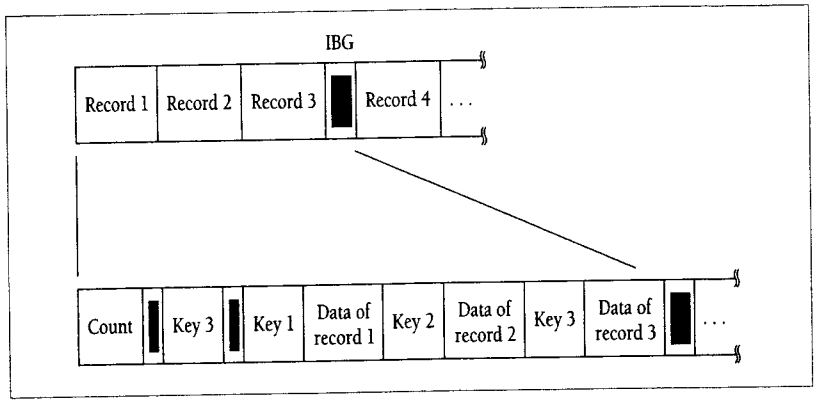

Figure 29 Blocked Fixed-Length Record Format (FB)

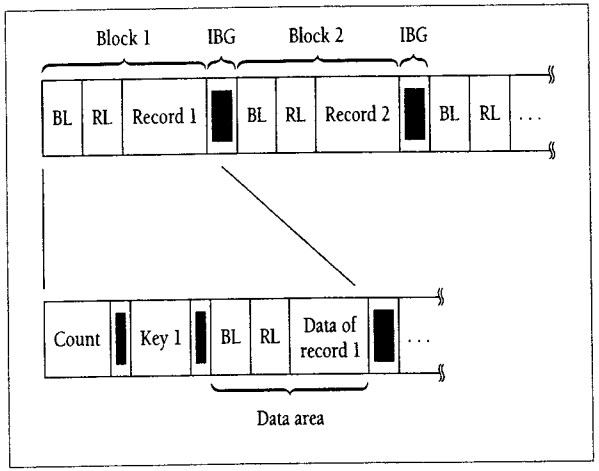

Figure 30 Unblocked Variable-Length Record Format (V)

block length field, a record length field, and a logical record, either in its entirety, or a segment of it (Figure 32).

Blocked variable-length spanned (VBS) format is basically a combination of the arrangements for VB and VS formats, as shown in Figure 33.

In the **undefined (U)** format, the blocks are unstructured. Each block is identified by the preceding and following interblock gaps. The content of a block varies from one to another. A block is treated as a character string, brought into the memory, and its content is analyzed and identified. This record format is rarely used today, but may come in handy if we have an old unstructured file and would like to interpret its contents.

Before we conclude our discussion of rotating magnetic storage devices, let us consider a hypothetical device in terms of its standard specifications, and compute some timing and storage parameters.

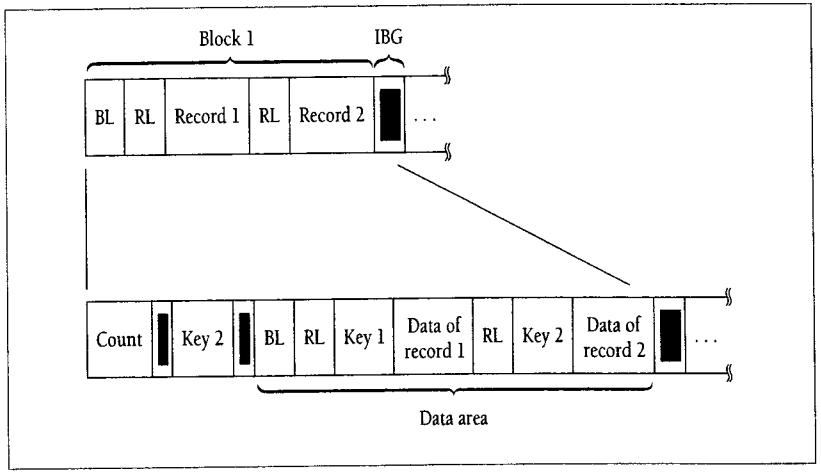

Figure 31 Blocked Variable-Length Record Format (VB)

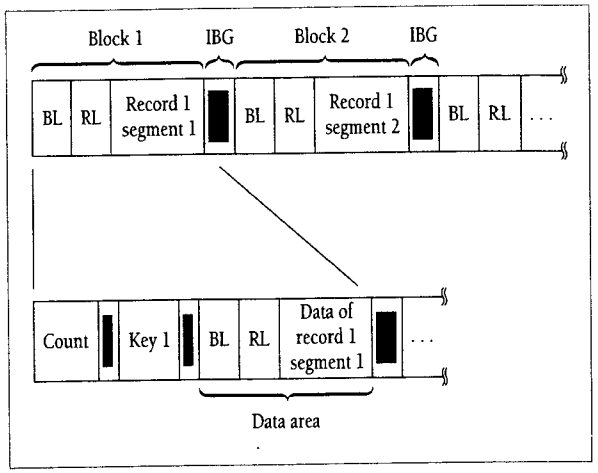

Figure 32 Unblocked Variable-Length Spanned Record Format (VS)

Example The device is an exchangeable disk pack drive with the following characteristics:

- Number of usable surfaces per pack = 19
- Number of tracks per surface = 500
- Net track capacity = 15,000 bytes
- Rotation speed = 3600 rpm

Using these specifications, we can compute the following:

- Rotation time = (1000 * 60/3600) = 16.7 msec
- Average latency = 1/2 * rotation time = 8.3 msec
- Number of cylinders = number of tracks per surface = 500
- Full cylinder capacity = (19 * 15,000) = 285,000 bytes
- Pack capacity = 285,000 * 500 = 142,500,000 bytes = 142.5 megabytes

File Directories and Labels

As many files are created by and maintained for different users in a computing environment, each one must be properly identified. Information that identifies files is kept in file directories. There is a file directory entry for every file on magnetic mass storage. In addition, as magnetic tape drives or disk drives can store many files, they have their own directories, known as **volume directories.**

Information included in volume and file directories and its organization vary from one operating system to another, and are somewhat different for tape and disk storage devices. Both can be tailored to meet the particular needs of a computer installation.

For tape storage, the operating system creates and maintains a volume directory in a special record called **volume label.** A single volume label is generated for a tape reel, written as the first record on the tape immediately after the tape index marker. It is approximately 80 bytes long, and contains information about the contents of a tape reel. This includes volume serial number, name of the tape owner, whether the volume is write-protected or not, and so on.