DRAFT:

Identification

Multilevel Storage Support

L. J. Lambert, K. J. Martin

Discussion

Multilevel Storage Management (BH.4) determines where in secondary storage a file in the Multics file hierarchy shall reside. The primary goal of Multilevel is to arrange file residence for the most efficient overall operation of Multics. This is accomplished by placing frequently accessed files on fast storage devices and less frequently accessed files on slower devices. In addition, very large files are most efficiently placed on slower devices, leaving the space on faster devices for several small files.

The owner of a file (that is, the user who pays for its residence on secondary storage) knows the most about his needs regarding his file. In nearly all cases the owner of a file will need no special handling for his files. The files which he accesses frequently serve him best when residing on fast access devices. His infrequently accessed files need not be on fast access devices and are more economical when residing on slow-access devices. Occasionally, although rarely, a user will find that Multilevel's treatment of his files is not appropriate. For instance, he may know that although a file is large and used infrequently, when it is needed, the access must be prompt. The criteria normally used by Multilevel causes such a file to be placed on a device where access time would be prohibitive.

As illustrated above, it is desirable that the owner of a file be able to influence the actions of Multilevel Storage Management when necessary. The two factors which a user would want to influence are:

1) the length of time required to access a file residing in secondary storage, and

2) the amount he is willing to pay for secondary storage residence of a file.

Two ways of charging users for their use of secondary storage are considered here. One method is to charge more for time that a file resides on a fast, expensive device than for time it resides on a slower, cheaper device. This method has the advantage that it charges the user for the actual cost of storing his file. The second method is to charge uniformly for storage of files regardless of the devices on which they reside. The facts that Multilevel constantly changes the place of residence of files, and that users have no absolute final control over residence of files, makes this second method attractive. Each Multics installation may

choose whichever method it likes. It is assumed that most owners of files want the best possible service from Multics for the least possible money, and that a user does not object to getting better service than he needs if the better service is not reflected in increased cost. Using either method of charging, factors mentioned above reduce to the requests:

1) to keep a file on a device at least as fast as device X, and

2) to keep a file on a device as slow as or slower than device Y in return for lower residency costs.

The interpretation of the requests on a system which uses the actual cost method of charging is as follows. Request 1 means that the user must have fast service and is willing to pay the increased cost at all times rather than paying for whatever Multilevel gives him. Request 2 means that the user is willing to pay no more than specified at all times and is also willing to accept poorer service. Under the constant cost method, request 1 means that the user must have fast service and is willing to pay a premium above the standard cost. Request 2 means that the user is willing to receive poorer service for a cost below the standard.

The differences in interpretation are slight. Under the actual-cost method the user will still pay for the actual residence of the file within the limits which he designated. Under the constant cost method the user pays either a constant but higher-than-normal or a constant but lower-than-normal cost.

These requests are useful to system administrators in another way. They serve as a guide in supplying equipment to the users of Multics. If the number of requests to keep files on the fastest available device exceeds the capacity of that device, then more units of that device are needed. If most users are willing to sacrifice speed of access for a low charge, then the need for large amounts of fast access storage is questionable.

Usage

To service the first demand (to keep a file on a device at least as fast as device X) the Transactor needs to know two things:

1) the file to be affected;
2) the storage device which is to be considered the lower limit for residency.

To service the second demand (a reduced charge for storage of a file) the Transactor needs to know:

1) the file to be affected;
2) the storage device which is the maximum for which the user is willing to pay residency charges.

Both requests may be made at once in which case three items are needed by

the Transactor. (The file name need not be stated twice.) When both requests are made, the maximum-cost device must be either the same as or more expensive than the minimum-speed device.

The implementation of these requests to the Transactor will be as a command. When the command is issued, the Transactor goes to the directory which contains the branch for the specified file and attempts to reset the Multilevel parameter(s) which are located in the file branch. These parameters indicate the lower limit and the upper limit for file residency as requested by the owner. The parameters are numeric values which are assigned to devices.

In order for the Transactor command to reset these parameters, the user issuing the command must be able to write in the directory. Access control (BG.9) automatically checks whether the user has write access for the directory. This serves as a check on the user's right to set the Multilevel parameters for the file, since write access is required to enter the branch in the directory originally.

In case only one of the parameters is being set, the Transactor checks that the new setting is not inconsistent with the other parameter. That is, the upper-limit device for cost must be at least as high as the lower-limit device for residence.

Interpretation by Multilevel Storage Management

As stated previously, the primary responsibility of Multilevel is to make over-all operation of Multics as efficient as possible. Since this is the case, a user's requests for placement of his files will be regarded as priorities for consideration rather than as guarantees. The requests are considered as follows:

Multilevel first decides where all files should reside as a function of their size and use. If some device has room on it for more files, directories on slower-access devices are checked for files which requested to be on a faster-access device. These files are the first to be moved to the faster-access device. Files on slower-access devices for which the owner is not willing to pay residence on the faster-access device are the very last to be moved up to the extra space.

Conversely, when a device begins to get too crowded, those files which have an upper-limit on charges which is below that device are the first to be moved to a slower-access (and cheaper) device. Files which have a request to be kept on a device at least that fast are the last to be moved down.

Multilevel will investigate recent usage of the file, then decide on the basis of that usage and the user-set limits (as interpreted above) where to place the file.