

Published: 12/16/66

Identification

The System Checkpoint Dump Decision Modules
S. H. Webber

Purpose

The system checkpoint dumping process produces a distinct body of detachable storage which contains one complete hierarchy skeleton (all file system directory segments) and the minimal set of accounting and system segments required for normal Multics operation. This dump constitutes a checkpoint measure whose results facilitate a swift return to Multics operation following a secondary storage catastrophe. The latest system checkpoint dump is the last data processed by the hierarchy reconstruction process (section BH.3.01) during such a reload. Data is placed on detachable storage such that after processing by the reconstructor, directory hierarchy restoration is complete although many non-directory segments may not yet be available to users. This section describes the system checkpoint dump decision modules and when read in context with section BH.2.00 gives a complete description of the system checkpoint dumping process.

Introduction

The system checkpoint dumper executes in two stages.

First - The entire directory hierarchy is scanned and every directory segment encountered is dumped. This data constitutes the hierarchy skeleton.

Second - The required set of system and accounting subtrees is dumped. The dump control module is supplied with a list of the names defining this set and for each, the scan module is called. All non-directory segments present in each sub-tree are dumped. The composite dump operates in two distinct stages and hence a different decision module must be specified for each.

Skeleton Dump Decision Module

Each time any directory branch is presented to the skeleton dump decision module by the scan module, the directory dump switch is set ON for that entry. This ensures that a copy of the directory defined by the entry is dumped

when examined by the dump module. Then, in order to insure that all inferior directories are searched for more directory branches, return is made to the scan module instructing it to call itself recursively. All other switches, availing other dump options, are turned OFF.

The decision logic is reflected on Figure 1.

This module is called from the scan module by:

```
call skeleton_dump (dsw, psp, deeper);
```

In this call deeper is the return to be made if the current entry is a directory branch, psp points to the base of the position segment and dsw is the call-dump-module switch. The PL/I declaration of the arguments in the following.

```
decl deeper label,
      dsw bit(1),
      psp ptr;
```

System Dump Decision Module

Following the skeleton dump, all system and accounting sub-trees are dumped. Since the hierarchy scan only enters those areas of the file system which are interesting, the system dump decision module is a procedure which recognizes all non-directory segments. For each non-directory branch encountered, the segment dump-switch (SDW) corresponding to that entry in the position segment is turned ON. This insures that the segment is dumped subsequently by the dump module. Each time a directory branch is found no action is taken but to return to the scan module instructing it to scan the inferior directory. In this way, all segments in the sub-tree are dumped.

The segment dump-switch is the only decision switch used by the module.

The system dump decision logic is illustrated in Figure 2.

The system dump decision module is called by

```
call system_dump (dsw, psp, deeper);
```

where the arguments are as described above.

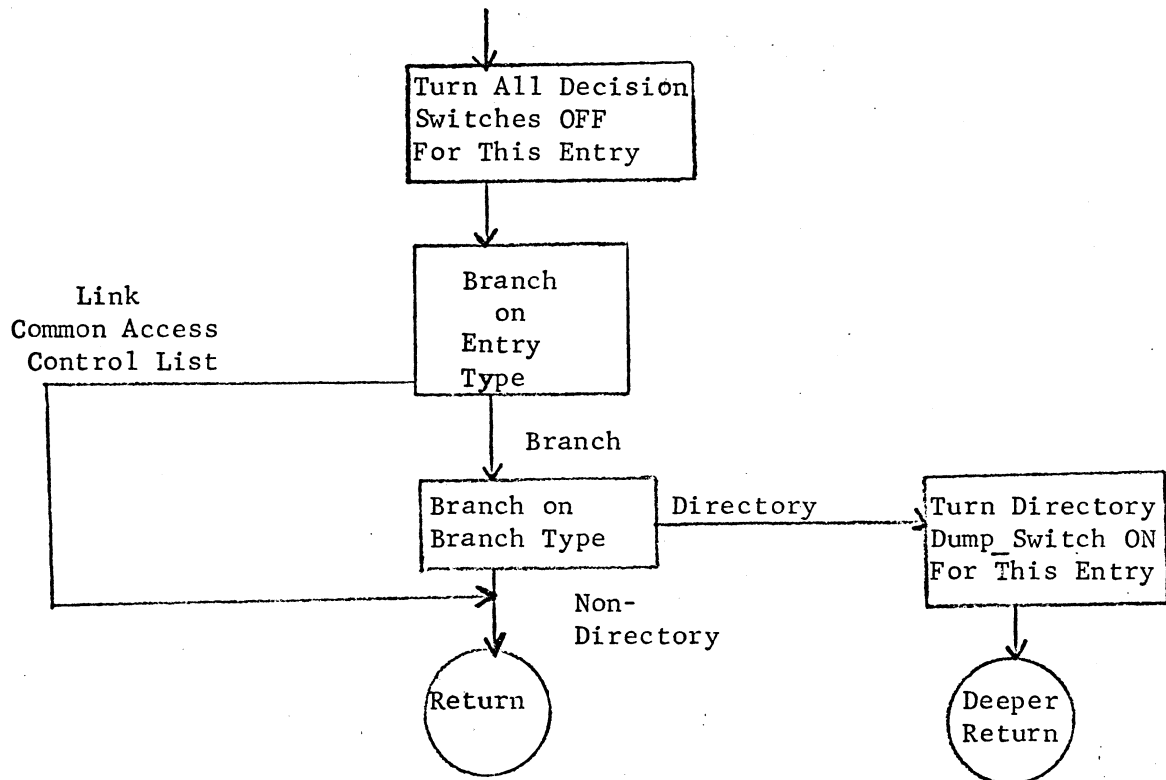


Figure 1 Skeleton Dump Decision Module

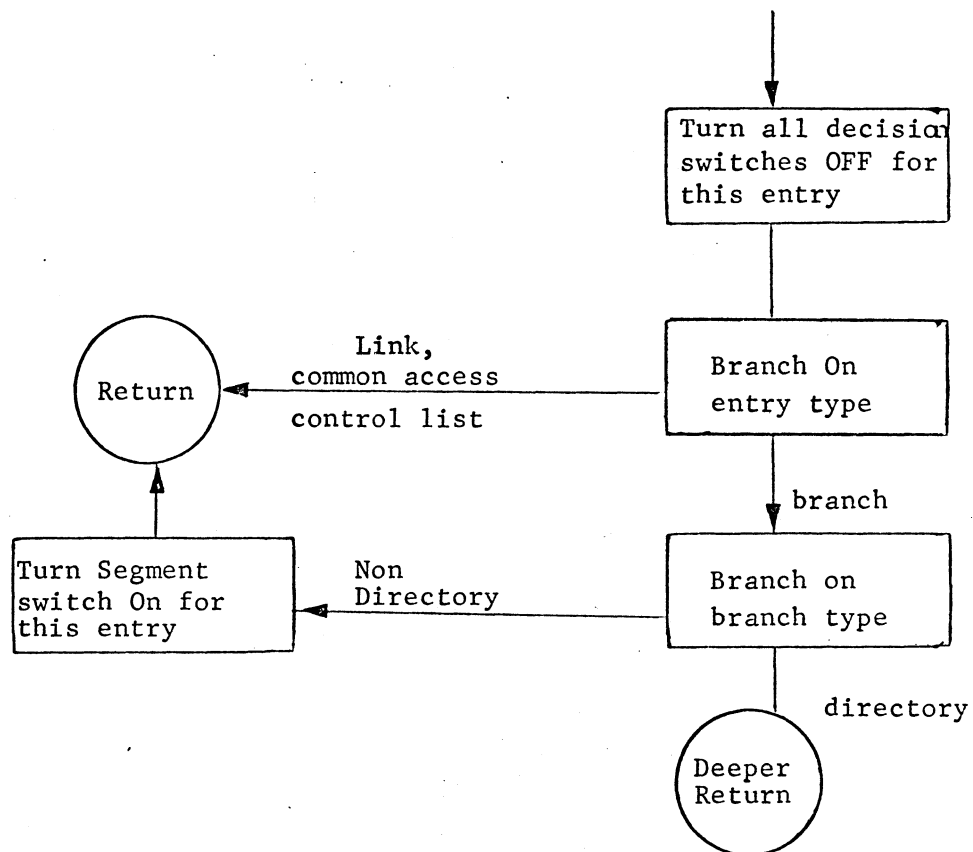


Figure 2 System Dump Decision Module