

# Technical NEWS

from  
Dynamic Concepts Inc.

July 1997

## c-tree

### Error 240

In both UniBasic and dL4, the index portion of an Indexed Contiguous file is a c-tree file. Each index in a c-tree file is a balanced tree (B-tree) with one or more levels. Each c-tree level corresponds to a disk access, which is composed of a disk seek (a slow and painful process), and a disk read or write (a relatively fast process). Therefore, optimal system performance can be achieved by having minimal c-tree levels. There are two ingredients that affect the number of c-tree levels. They are 1) maintaining a balanced tree and 2) inserting as many keys as possible in each B-tree block. Since the c-tree algorithm ensures a balanced tree, inserting maximum keys in each B-tree node is the only user configurable option. Using the following formula can approximate the number of keys that may be inserted in each node:

$$\# \text{ of keys/node} = (\text{ISAMSECT} * 128) / (\text{key length} + 4)$$

As you can see, each B-tree node can hold fewer keys as the key size increases, thus requiring more c-tree levels. The number of c-tree levels can be reduced as the key length increases only by increasing the value of ISAMSECT. The table below illustrates the relationship between node capacity, total

number of keys, and the maximum B-tree levels.

Increasing ISAMSECT value may require multiple sector read or write operations. However, multi-sector read/writes are relatively fast on most systems and hence a visible performance penalty should not be of much concern.

Where UniBasic uses four (4) as its default ISAMSECT value, dL4 uses eight (8) as its value. Using 8 for ISAMSECT increases the node capacity and thereby reduces the number of required levels. Because c-tree is limited to nine (9) levels, this also increases the maximum number of keys that can be inserted into an index. If the ISAMSECT value is too small, a c-tree error 240 may be reported.

Those files created with a larger ISAMSECT cannot be read with a smaller ISAMSECT. For example, files created with ISAMSECT set to 8 cannot later be accessed with ISAMSECT set to 4. However, the inverse is possible, i.e., files created with ISAMSECT set to 4 can later be accessed with ISAMSECT set to 8. Also, typically, a file's size is initially larger when it is created with larger ISAMSECT, but is actually smaller as it grows.

Keys/Node	Total Number of keys/Number of Levels			
	1,000	10,000	100,000	1,000,000
10	4	6	7	9
20	3	4	5	6
40	3	3	4	5
80	2	3	3	4
160	2	2	3	3



**Did you  
know that  
c-tree is  
used with  
dL4, IQ  
and  
UniBasic?**

Although the ISAMSECT value may be configured as an ISAMSECT environment variable for UniBasic, the same is not true with dL4. This is because an ISAMSECT value is properly set to eight to accommodate large files.

*\* Excerpts for this article taken directly from FairCom's c-tree*

Telephone:  
800.822.2742  
714.448.8600



Email:  
support@dynamic.com

