

NAME

`ccomps` – connected components filter for graphs

SYNOPSIS

`ccomps` [`-sxvenzC?`] [`-X[#]s[-f]`] [`-ooutfile`] [*files*]

DESCRIPTION

`ccomps` decomposes graphs into their connected components, printing the components to standard output.

OPTIONS

The following options are supported:

- `-e` Do not induce edges in the connected components.
- `-s` No output graph is printed. The return value can be used to check if the graph is connected or not.
- `-x` Only the connected components are printed, as separate graphs.
- `-v` Counts of nodes, edges and connected components are printed.
- `-z` Sort components by size, with the largest first. This is only effective if either `-x` or `-X#` is present. Thus, `-zX#0` will cause the largest component to be printed.
- `-C` Use clusters in computing components in addition to normal edge connectivity. In essence, this gives the connected components of the derived graph in which nodes top-level clusters and nodes in the original graph. This maintains all subgraph structure within a component, even if a subgraph does not contain any nodes.
- `-n` Do not project subgraph structure. Normally, if `ccomps` produces components as graphs distinct from the input graph, it will define subgraphs which are projections of subgraphs of the input graph onto the component. (If the projection is empty, no subgraph is produced.) If this flag is set, the component contains only the relevant nodes and edges.
- `-X node_name`
Prints only the component containing the node *node_name*, if any.
- `-X# start`
- `-X# start-`
- `-X# start-last`
Prints only components in the given range. In the first form, only the component whose index is *start*, if any, is printed. In the second form, each component whose index is at least *start* is printed. In the last form, only those components whose indices are in the range [*Istart*,*Ilast*] are printed. Thus, the flag `-x` is equivalent to `-X#0-`.
- `-o outfile`
If specified, each graph will be written to a different file with the names derived from *outfile*. In particular, if both `-o` and `-x` flags are used, then each connected component is written to a different file. If *outfile* does not have a suffix, the first file will have the name *outfile*; then next *outfile_1*, then next *outfile_2*, and so on. If *outfile* has a suffix, i.e., has the form *base.sfx*, then the files will be named *base.sfx*, *base_1.sfx*, *base_2.sfx*, etc.

By default, each input graph is printed, with each connected component given as a subgraph whose name is a concatenation of the name of the input graph, the string `"_cc_"` and the number of the component.

OPERANDS

The following operand is supported:

- files* Names of files containing 1 or more graphs in dot format. If no *files* operand is specified, the standard input will be used.

RETURN CODES

Unless used to extract a single connected component, `ccomps` returns `0` if all the input graphs are connected; and non-zero if any graph has multiple components, or any error occurred. If just extracting a single component, `ccomps` returns `0` on success and non-zero if an error occurred.

BUGS

It is possible, though unlikely, that the names used for connected components and their subgraphs may conflict with existing subgraph names.

AUTHORS

Stephen C. North <north@research.att.com>

Emden R. Gansner <erg@research.att.com>

SEE ALSO

gc(1), dot(1), gvpr(1), gvcolor(1), acyclic(1), sccmap(1), tred(1), libgraph(3)