

larger graphs if enough main memory is available.) Both general and specific (bipartite, grid, planar, disk) types of graphs can be generated. A collection of useful operations for testing graph properties, manipulation of the graph and automatic embedding of the graph is provided as well.

The program also supports automatic generation of up to 100 services with random integer sizes in the range $[0, 100]$ and generation of different types of random requirements. (Again, these limits are only a pragmatic restriction of input parameters; internally, the system can handle service lists of arbitrary length and with arbitrary sizes.)

Currently, the following algorithmic tools are provided:

- Computation of four kinds of approximate LLB's ("Largest Lower Bounds"): The *sum* and *FF LLB's* provide a lower and upper bound for the clique packing number, respectively, in which the size $p_M(W)$ of a clique W is approximated using $\lceil \mu_{R_W}/M \rceil$ and $FF(R_W, M)$, respectively. Analogous bounds (*sum* and *FF strict LLB's*) can be computed for the *strict* clique packing number, which is defined as the maximum of $\max_{w \in W} (p_M(R_w) + p_M(R_W \setminus R_w))$ over all cliques W of G . We remark that this number is at most twice the ordinary clique packing number, and that the lower (sum) bound on the strict clique packing number is at most one more than the ordinary sum LLB. The strict clique packing number gives a lower bound for the number of colors required for a strict ensemble assignment, just like the ordinary clique packing number yields a lower bound on arbitrary ensemble assignments. These tools are implemented using the standard greedy (FF = First-Fit) bin packing algorithm (see, e.g., Garey/Johnson 1979) and a generalized version of the Carraghan/Pardalos clique algorithm. Note that the latter algorithm has an exponential running time in the worst case (but usually it works reasonably fast in practice, at least for graphs which are not too dense).
- Computation of SFF (Simultaneous-First-Fit) and GFF (Global-First-Fit) ensemble assignments satisfying the given requirements; implemented using the FF algorithm.
- Computation of arbitrary and "regular" colorings for the computed ensemble assignments. *Regular colorings* assign the same color to each instance of an ensemble in different areas; they are obtained by coloring the *regular ensemble graph* which has \mathcal{B}_V (the set of all ensembles) as its vertex set and edges between all distinct ensembles B and C occurring in the same or interfering areas. These tools are implemented using the sequential coloring algorithm using different kinds of orders. An iterative version (Culberson algorithm), in which the ensemble graphs are repeatedly recolored using an order determined from the previous coloring, is also provided.
- Computation of the maximum clique size of the ordinary and the regular ensemble graph for the current ensemble assignment. These numbers provide lower bounds for the corresponding colorings and can thus be used to assess the quality of the computed colorings. Implemented using the Carraghan/Pardalos algorithm.

The LLB and SFF/GFF tools can both be applied to the whole graph or a subset of its nodes. This allows the user to construct mixed types of ensemble assignments in an incremental fashion. Service sizes, requirements and ensemble assignments can also be edited directly by the user.

The graph, together with the associated services, requirements, ensemble assignments and colorings, can be saved to and reread from files in a special format. Reading and writing files