service will only be supplied in a single ensemble for all areas where it is requested). As already indicated, we would expect SFF to work best with sparse, and GFF with dense graphs. In fact, the performance bounds for the first-fit bin packing algorithm directly carry over to SFF and GFF solutions. That is, a (properly colored) SFF (resp. GFF) assignment on an independent (resp. complete) area graph using a service order by decreasing bandwidths will at most be about 22% off the optimum (asymptotically).

6 Lower Bounds

Whenever heuristics are employed to solve difficult problems in combinatorial optimization, quality control becomes an important issue. For a minimization problem, we need a good *lower* bound on the optimum value of the target function which can be computed in a reasonable amount of time. With such a bound at hand, we can estimate the maximum deviation of the heuristic solution from the optimum. For the ensemble planning problem, a useful lower bound is provided by an appropriate generalization of the clique number, which we discuss in this section.

As in the preceding section, we let $R_W = \bigcup_{w \in W} R_w \ \forall\, W \subseteq V$. Similarly, the set of all ensembles in a given area set $W \subseteq V$ is denoted $\mathcal{B}_W$, i.e., $\mathcal{B}_W = \bigcup_{w \in W} \mathcal{B}_w$. Our bound is based on the following simple observation:

Lemma 1 (Packing Lemma) *If $\mathcal{B}$ is an admissible ensemble assignment w.r.t. R and M, then $|\mathcal{B}_W| \geq p_M(R_W) \ \forall\, W \subseteq V$.*

Lemma 1 follows immediately from the fact that if $\mathcal{B}$ is admissible, then $R_W \subseteq \bigcup_{B \in \mathcal{B}_W} B \ \forall\, W \subseteq V$ and hence $\mathcal{B}_W$ is an "M-cover" of a superset of R_W (which is just like an M-packing, except that the members of $\mathcal{B}_W$ are not necessarily disjoint). Now consider the special case that W is a clique of G. In this case the subgraph of $G^{\mathcal{B}}$ induced by $\mathcal{B} \cap (W \times \mathcal{B}_W)$ always contains a clique of size $|\mathcal{B}_W|$. (For each $B \in \mathcal{B}_W$ choose some $w_B \in W$ s.t. $B \in \mathcal{B}_{w_B}$. Then $\{(w_B, B) : B \in \mathcal{B}_W\}$ is a clique of the requested size.) Hence

$$p_M(R_W) \leq |\mathcal{B}_W| \leq \omega(G^{\mathcal{B}}) \leq \chi(G^{\mathcal{B}}). \tag{7}$$

Now we define

$$\pi_M^R(G) = \max\{p_M(R_W) : W \text{ clique of } G\}. \tag{8}$$

By taking the maximum over all cliques W of G on the left-hand side of Equation (7), and the minimum over all admissible ensemble assignments $\mathcal{B}$ on the right-hand side, we obtain

$$\pi_M^R(G) \leq \chi_M^R(G). \tag{9}$$

The quantity $\pi_M^R(G)$, which we call the *clique packing number* of G w.r.t. R and M, is a weighted generalization of the clique number which, instead of merely counting the