

But the f^* value of any node on an optimal path is equal to $f^*(s)$, the minimal cost, and therefore $f(n') \leq f^*(s)$. Thus, we have:

RESULT 2: At any time before A* terminates, there exists on *OPEN* a node n' that is on an optimal path from s to a goal node, with $f(n') \leq f^*(s)$.

Combining this result with our previous argument, that even the smallest f values of the nodes on *OPEN* of a nonterminating A* become unbounded, shows that A* must terminate even for infinite graphs. Thus,

RESULT 3: If there is a path from s to a goal node, A* terminates.

RESULT 3 has an interesting corollary, namely, that any node, n , on *OPEN* with $f(n) < f^*(s)$ will eventually be selected for expansion by A*. We leave the proof as an exercise for the reader.

Now it is a simple matter to show that A* is admissible. First, we note again that A* can either terminate by finding a goal node in step 5 or, after depleting *OPEN*, in step 3. But *OPEN* can never become empty before termination if there is a path from s to a goal node because, by RESULT 2, there will always be a node on *OPEN* (and on an optimal path). Therefore, A* must terminate by finding a goal node.

Next we would like to show that A* only terminates by finding an optimal path to a goal node. Suppose A* were to terminate at some goal node, t , without finding an optimal path, that is, $f(t) = g(t) > f^*(s)$. But, by RESULT 2, there existed just before termination a node, n' , on *OPEN* and on an optimal path with $f(n') \leq f^*(s) < f(t)$. Thus, at this stage, A* would have selected n' for expansion rather than t , contradicting our supposition that A* terminated. Therefore, we finally have

RESULT 4: Algorithm A* is admissible. (That is, if there is a path from s to a goal node, A* terminates by finding an optimal path.)

Each node selected for expansion by A* has an interesting property that follows directly from RESULT 2: Its f value is never greater than the cost, $f^*(s)$, of an optimal path. This result will be important to us later. To show that it is true, let n be any node selected for expansion by A*. If n

is a goal node, we have $f(n) = f^*(s)$ by RESULT 4; so suppose n is not a goal node. Now A* selected n before termination, so at this time (by RESULT 2) we know that there existed on *OPEN* some node n' on an optimal path from s to a goal with $f(n') \leq f^*(s)$. If $n = n'$, our result is established. Otherwise, we know that A* chose to expand n rather than n' ; therefore it must have been the case that

$$f(n) \leq f(n') \leq f^*(s).$$

Therefore, we have

RESULT 5: For any node n selected for expansion by A*, $f(n) \leq f^*(s)$.

2.4.4. COMPARISON OF A* ALGORITHMS

The precision of our heuristic function h depends on the amount of heuristic knowledge it possesses about the problem domain. Clearly, using $h(n) \equiv 0$ reflects complete absence of any heuristic information about the problem, even though such an estimate is a lower bound on $h^*(n)$ and therefore leads to an admissible algorithm.

Let us compare two versions of A*, namely, A_1 and A_2 using the following evaluation functions:

$$f_1(n) = g_1(n) + h_1(n)$$

and

$$f_2(n) = g_2(n) + h_2(n)$$

where h_1 and h_2 are both lower bounds on h^* . We say that algorithm A_2 is *more informed* than algorithm A_1 if for all nongoal nodes, n , $h_2(n) > h_1(n)$. This definition seems intuitively reasonable, since with h bounded from above by h^* for admissibility, one suspects that using larger values of h (and thus values closer to h^*) requires more accurate heuristic information.

As an example, consider the 8-puzzle solved in Figure 2.8. There we used the evaluation function $f(n) = d(n) + W(n)$. We can interpret the search process of that example as an application of A* with