Longest-prefix matching (LPM) is a key processing function of Internet routers. This is an important step in determining which outbound port to use for a given destination address. The time required to look up the outbound port must be less than the minimum inter-arrival time between packets on a given input port. Lookup times can be reduced by caching address prefixes from previous lookups. However all misses in the prefix cache (PC) will initiate a traversal of the routing table to find the longest matching prefix for the destination address. This table is stored in memory so a traversal requires multiple (perhaps many) memory references. These memory references become an increasingly serious bottleneck as line rates increase. In this paper we present a novel second level of caching that can be used to expedite lookups that miss in the PC. We call this second level a dynamic substride cache (DSC). Extensive experiments using real traffic traces and real routing tables show that the DSC is extremely effective in reducing the number of memory references required by a stream of lookups. We also present analytical models to find the optimal partition of a fixed amount of memory between the PC and DSC.