F. J. MUNIZ, "Parallel Load-Balancing on Message Passing Architectures"

This thesis investigates and develops dynamic load mechanisms on distribute-memory MIMD machines . Load-balancing is essential for efficient computation on parallel processors. The primary objective of load-balance mechanisms is to achieve a balanced and therefore efficient utilisation of processing resources. ln particular, the investigations described here facilitate portability of the user application, since the investigation, since the thesis concentrates on system- levelload-balancing mechanisms. The load-balancing mechanisms studied are also suitable for systems that can vary in size, concentrating on methods with potential for T800 transputers programmed in the OCCAM and 'C' languages and an automatic routing package communication software mechanism (the Virtual Channel Router). Careful consideration was given to develop generic mechanisms that would be applicable to a wide range of processing architectures, not just those built from transputers. The research was initially focused on user transparent dynamic migration load-balancing mechanisms with a subsequent investigation of dynamic placement load- balancing approaches. Synthetic user application models together with a processor 'busyness'- leveI measurements were implement were implemented to provide a tested in order to evaluate these proposed mechanisms. The application model utilising the dynamic placement based on remote procedural calls. This spawning mechanism enables dynamic verification of processing availability, as well as the characterisation of global placement space and competition for resources. A real application chosen from the field of computer graphics (ray tracing) was also implemented, within the framework presented above,. to provide a further example of the applicability of the of the dynamic placement load-balancing mechanism. This thesis therefore contains an investigation into the fundamentaIs of dynamic load-balancing mechanisms on distributed-memory MIMD machines as well a demonstration that system-Ievel user-independent dynamic load-balancing is feasible, practical and can significantly improve performance.