

Περίληψη

As CPU manufacturers move to a multi-core model, OS kernels must be ready to take advantage of additional processing units. Most kernels have elected to use fine-grained locking in order to allow parallel execution in the kernel. Instead, the DragonFlyBSD kernel partitions the main network-related data structures onto each cpu and moves network protocol processing into independant per-cpu kernel threads that communicate with applications using message passing. This approach has the potential to scale linearly with the processor count while providing a simpler programming model. Our work eliminates the need for synchronization between the protocol processing threads on the performance sensitive control paths. We also design and implement a socket buffer architecture which is used by the kernel threads in a lock and wait free fashion. Our initial measurements give promising results and we investigate directions for further improvements.

Keywords: SMP, Multi-core, Internet protocol, TCP, UDP, Producer-consumer, Fine-grained locking, Message passing, DragonFlyBSD