

# Notes on Thread Models in Mach 3.0

Bryan Ford  
Mike Hibler  
Jay Lepreau

UUCS-93-012

Department of Computer Science  
University of Utah  
Salt Lake City, UT 84112, USA

April, 1993

## **Abstract**

During the Mach In-Kernel Servers work, we explored two alternate thread models that could be used to support traps to in-kernel servers. In the “migrating threads” model we used, the client’s thread temporarily moves into the server’s task for the duration of the call. In the “thread switching” model, an actual server thread is dispatched to handle client traps. Based on our experience, we find that the migrating threads model is quite complex and difficult to implement in the context of the current design of Mach and the Unix single server. The thread switching model would fit more naturally and would probably be much simpler and more robust than migrating threads, making it a valuable approach to explore in the near future. However, we believe migrating threads inherently to be faster than thread switching, and ultimately to be the best long term direction.<sup>1</sup>

---

<sup>1</sup>This research was sponsored by the Hewlett-Packard Research Grants Program.