



# Initial Kernel Timing Using a Simple PIM Performance Model

Daniel S. Katz<sup>1\*</sup>, Gary L. Block<sup>1</sup>, Jay B. Brockman<sup>2</sup>,  
David Callahan<sup>3</sup>, Paul L. Springer<sup>1</sup>, Thomas Sterling<sup>1,4</sup>

<sup>1</sup>*Jet Propulsion Laboratory, California Institute of Technology, USA*

<sup>2</sup>*University of Notre Dame, USA*

<sup>3</sup>*Cray Inc., USA*

<sup>4</sup>*California Institute of Technology, USA*

\*Technical Group Supervisor  
Parallel Applications Technologies Group  
<http://pat.jpl.nasa.gov/>  
[Daniel.S.Katz@jpl.nasa.gov](mailto:Daniel.S.Katz@jpl.nasa.gov)



# Purpose of this Poster

- Discuss initial results of paper-and-pencil studies of 4 application kernels applied to a processor-in-memory (PIM) system roughly similar to the Cascade Lightweight Processor (LWP)
- Application kernels:
  - Linked list traversal
  - Vector sum
  - Bitonic sort
- Intent of work is to guide and validate work on Cascade in the areas of compilers, simulators, and languages



# Poster Topics

- Generic PIM structure
- Concepts needed to program a parallel PIM system
  - Locality
  - Threads
  - Parcels
- Simple PIM performance model
- For each kernel:
  - Code(s) for a single PIM node
  - Code(s) for multiple PIM nodes that move data to threads
  - Code(s) for multiple PIM nodes that move threads to data
  - Hand-drafted timing forecasts, based on the simple PIM performance model
- Lessons learned
  - What programming styles seem to work best
    - Looking at both expressiveness and performance

