

Implementation of NAMD molecular dynamics non-bonded force-field on the Cell Broadband Engine processor

Guochun Shi (gshi@ncsa.uiuc.edu), Volodymyr Kindratenko (kindr@ncsa.uiuc.edu)

National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign

Cell Processor and IBM Cell Blade

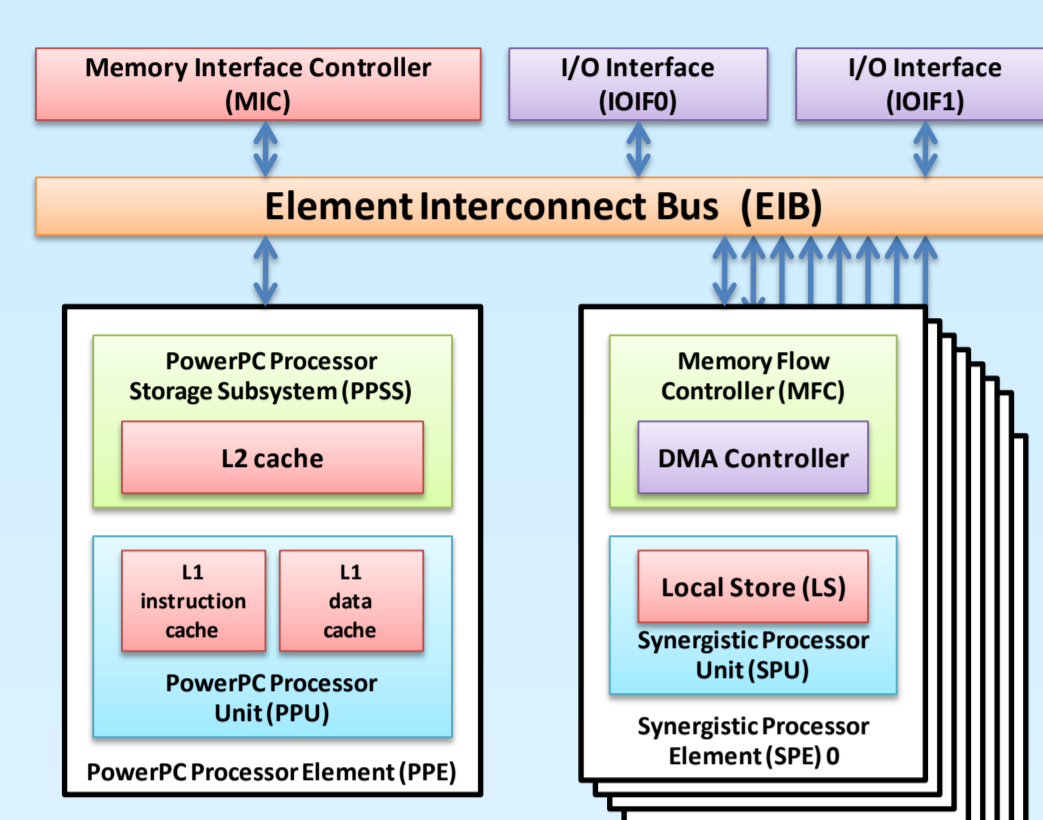


Figure 1: Cell/B.E. processor architecture

- One Power Processor Element (PPE) and eight Synergistic Processing Elements (SPEs)
- SPE's local storage (LS): 256 KB
- Processor clock speed: 3.2 GHz
- 25.6 GB/s processor-to-memory bandwidth
- 205 GB/s EIB sustained aggregate bandwidth
- Theoretical peak performance: 204.8 GFLOPS (SP) and 14.63 GFLOPS (DP)

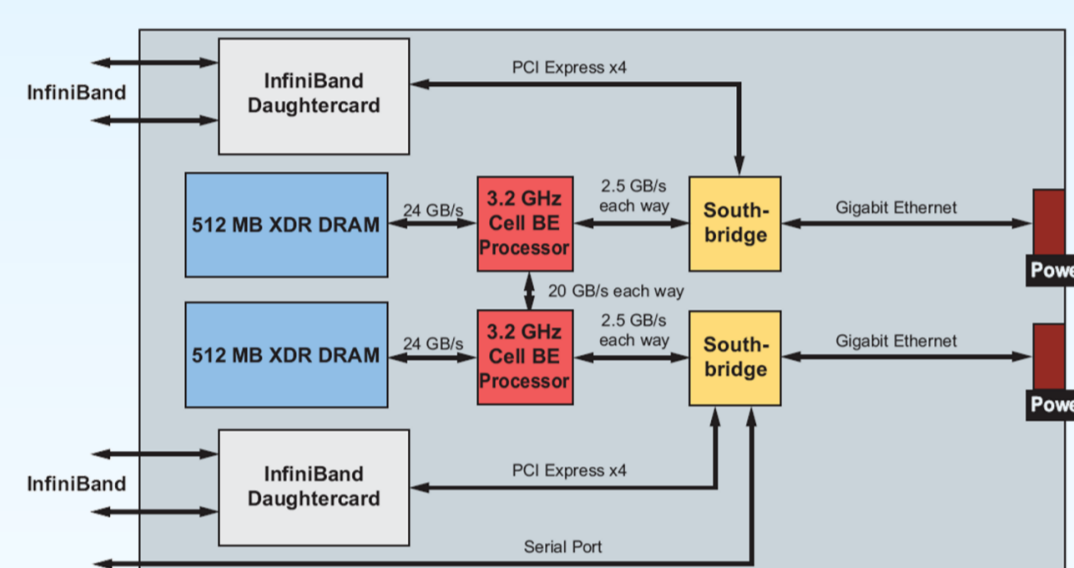


Figure 2: IBM dual-Cell/B.E. blade

- 410/28 GFLOPS (SP/DP) peak
- 512 MB of XDR DRAM per Cell/B.E. processor
- Gigabit Ethernet

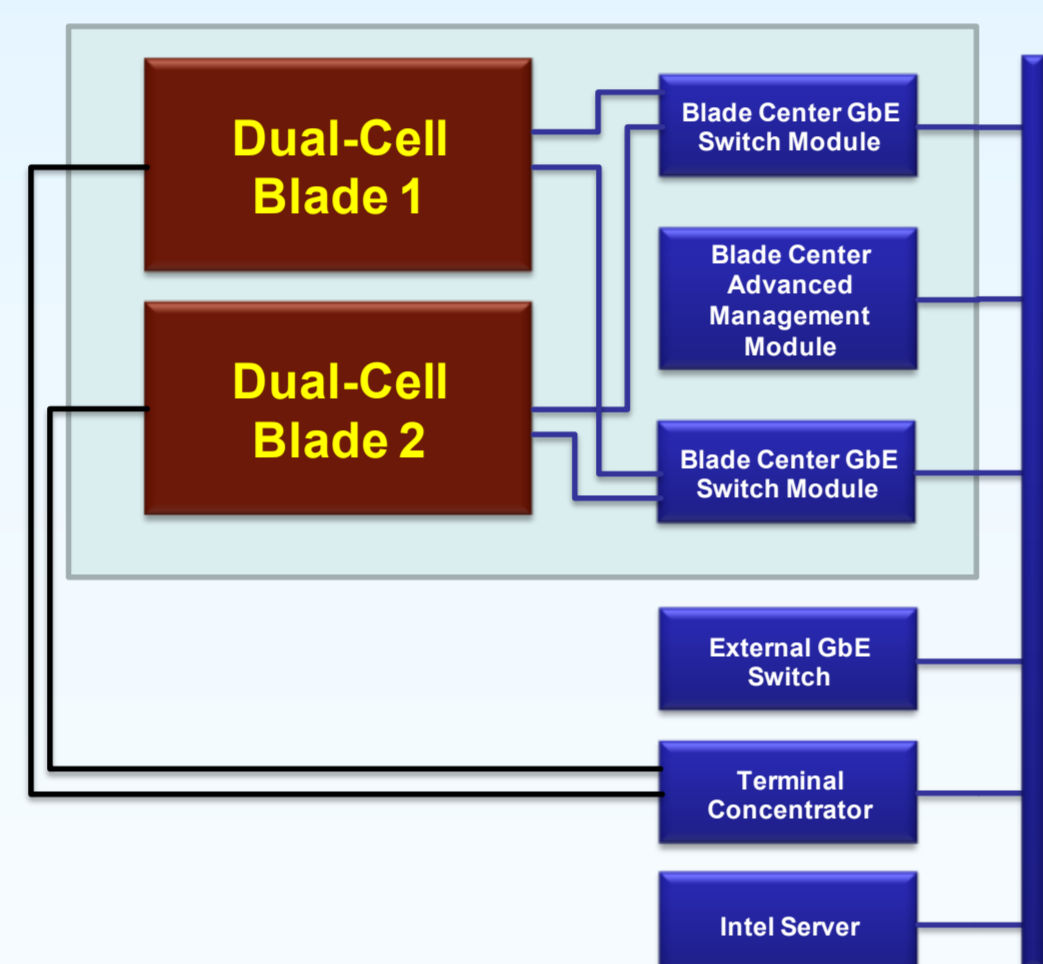


Figure 3: IBM Cell/B.E. blade center

NAMD kernel

NAMD is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems.

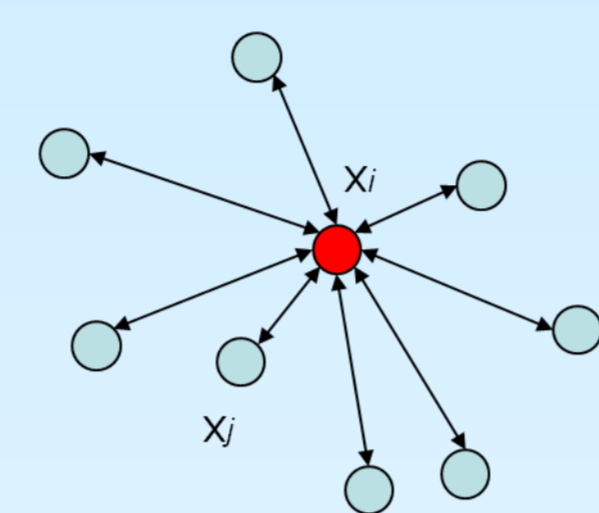
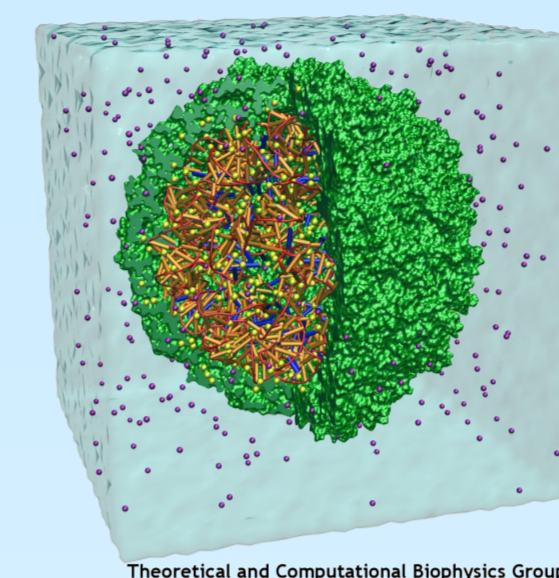
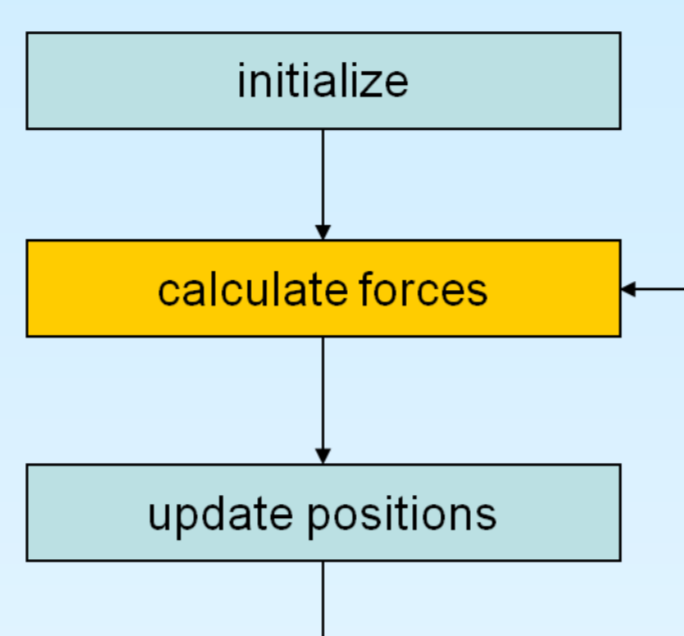


Figure 5: MD simulation principles



$$F(x_i^k) := \sum_{j=1}^N f(x_i^k, x_j^k)$$

$$x_i^{k+1} = x_i^k + f(F(x_i^k))$$

NAMD SPEC 2006 CPU benchmark kernel:

- for each atom i in patch p_i
- for each atom j in patch p_j
- if atoms i and j are bonded, compute bonded forces
- otherwise, if atoms i and j are within the cutoff distance, add atom j to the i 's atom pair list
- end
- for each atom k in the i 's atom pair list
- compute non-bonded forces (L-J potential and PME direct sum, both via lookup tables)
- end

We implemented a simplified version of the kernel that excludes pairlists and bonded forces:

- for each atom i in patch p_i
- for each atom j in patch p_j
- if atoms i and j are within the cutoff distance
- compute non-bonded forces (L-J potential and PME direct sum, both via lookup tables)
- end
- end

NAMD on Cell: program flow

Different vectorization schemes are applied in order to get best performance:

- Self-compute:** fill zeros in unused slots
- Pair-compute:** save enough pairs of atoms, then do calculations

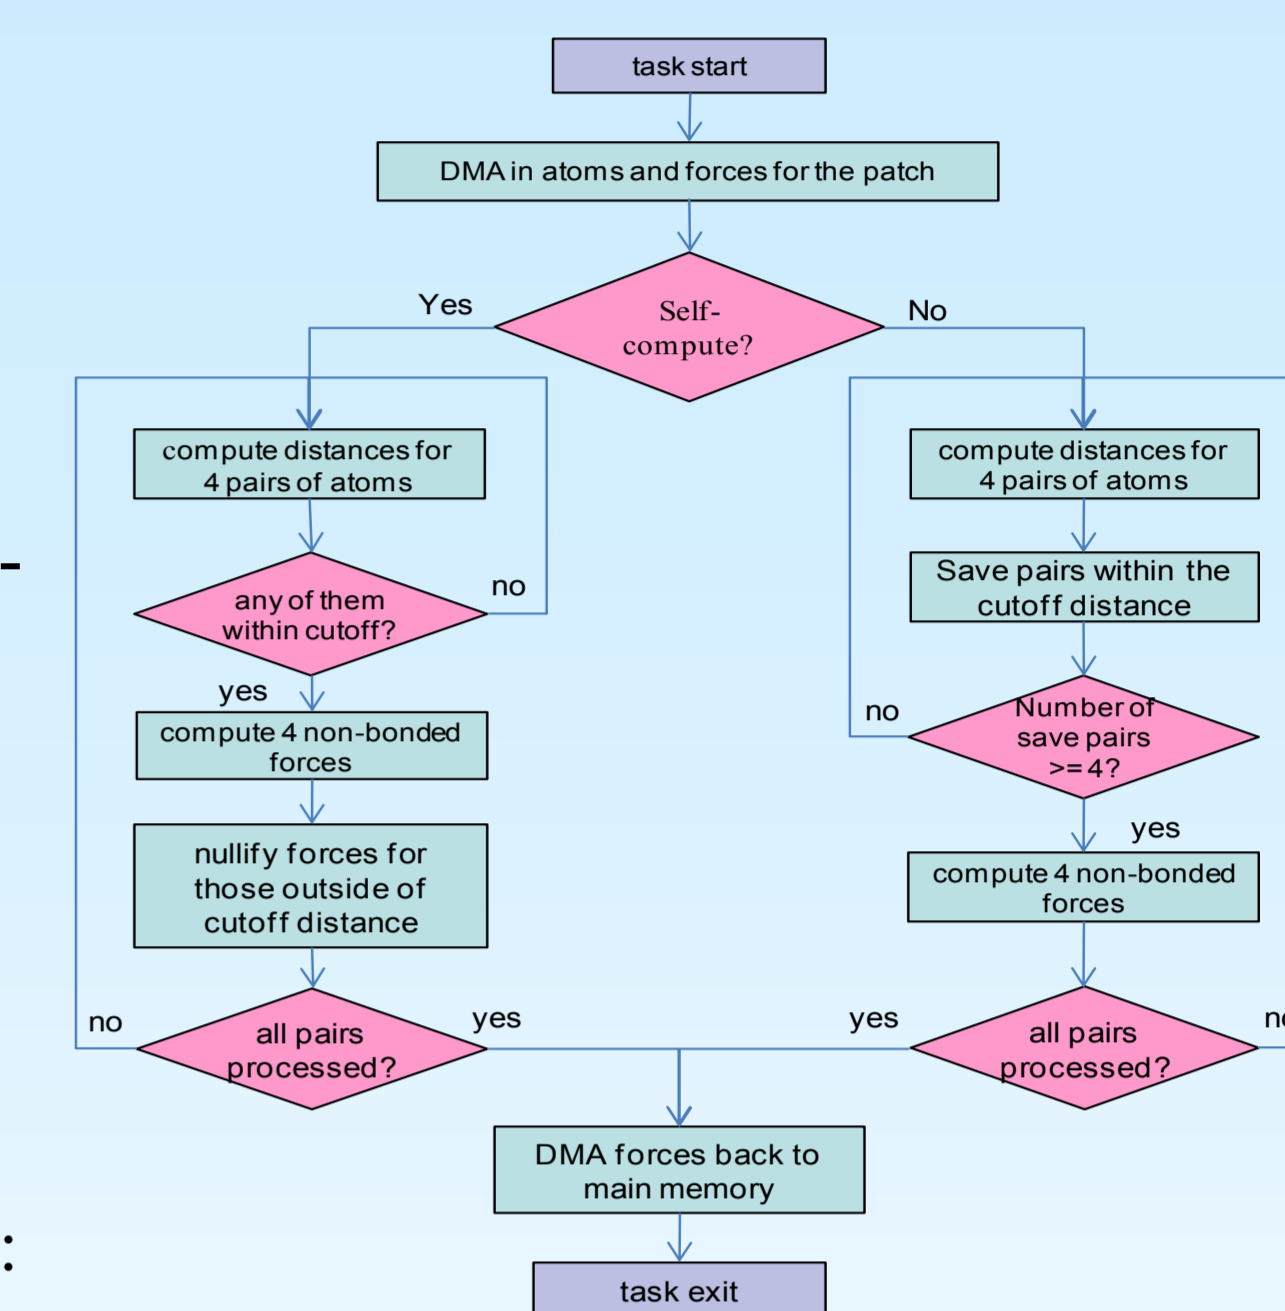


Figure 7: SPE Program flow

PPE runs a task dispatch system:

- Find dependency-free task
- Schedule the task on the idle SPE

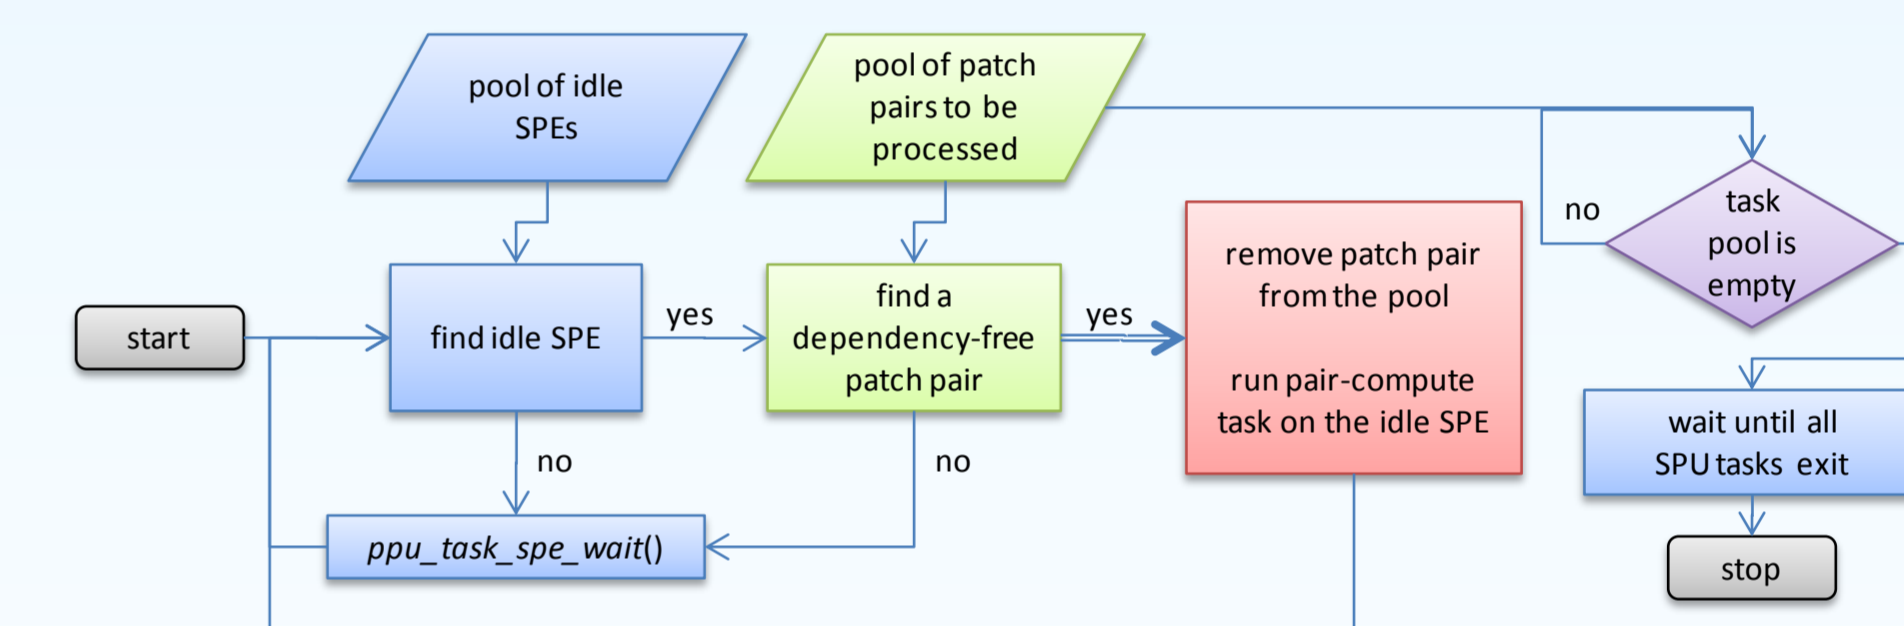


Figure 8: Task dispatch system running on the PPE

Function offload programming model

- User compute task structure inherits the common `task_t` structure:

```
typedef struct task_s {
    int cmd; // operand
    int size; // size of task structure
} task_t;
```

```
typedef struct compute_task_s {
    task_t common;
    <user_type1> <user_var_name1>
    <user_type2> <user_var_name2>
    ...
} compute_task_t;
```

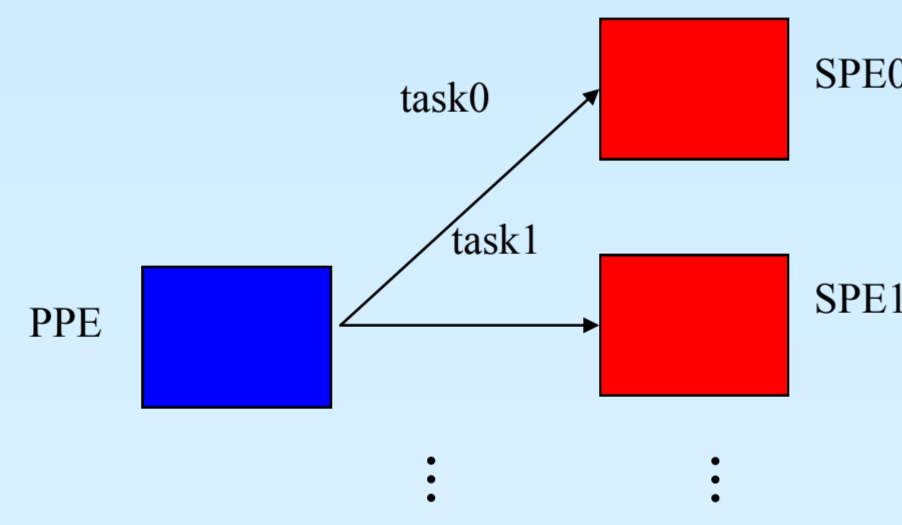


Figure 4: Task based function offload model

- API for the PPE

```
int ppu_task_init(int argc, char **argv, spe_program_handle_t); // initialization
int ppu_task_run(volatile task_t * task); // start a task in all SPEs
int ppu_task_spu_run(volatile task_t * task, int spe); // start a task in one SPE
int ppu_task_spu_wait(void); // wait for any SPE to finish, blocking call
void ppu_task_spu_waitall(void); // wait for all SPEs to finish, blocking call
```

- API for the SPEs

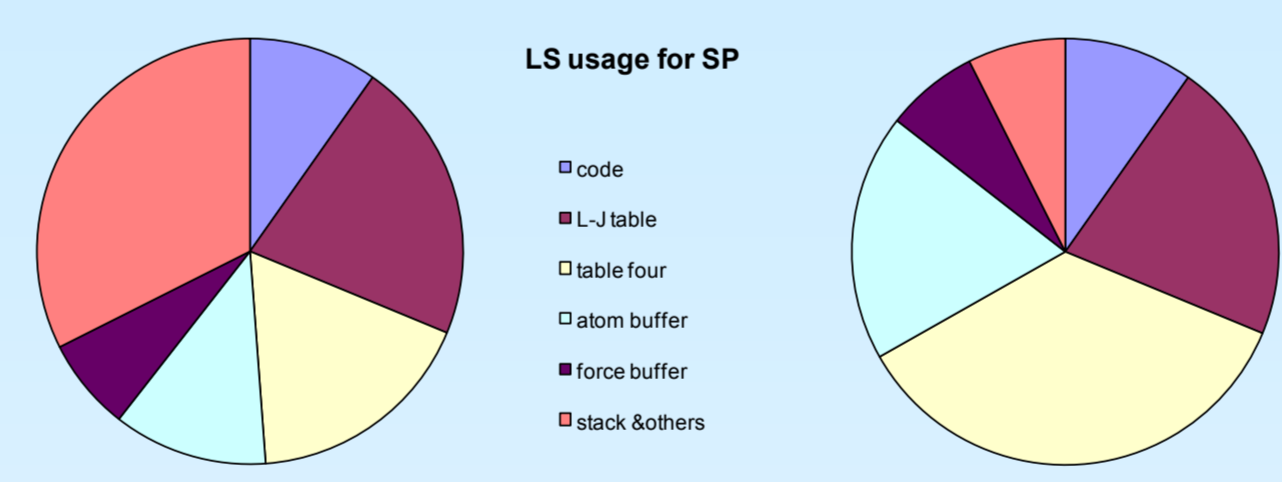
```
int spu_task_init(unsigned long long);
int spu_task_register(dotask_t, int); // register a task
int spu_task_run(void); // start the infinite loop, wait for tasks
```

- Two programming models are supported: *single SPE* and *multiple SPEs*

SPU NAMD Kernel implementation

Kernel type	Code	L-J table	table_four	atom buffer	force buffer	stack & other
single-precision	25KB	55KB	45KB	30KB	18KB	83KB
double-precision	25KB	55KB	91KB	48KB	18KB	19KB

Table 1: SPE LS memory usage for different kernel types



- Entire patch is loaded into SPE's local storage
- Double-precision case just fits, single-precision still have some space left
- Substantial number of data movement operations are needed

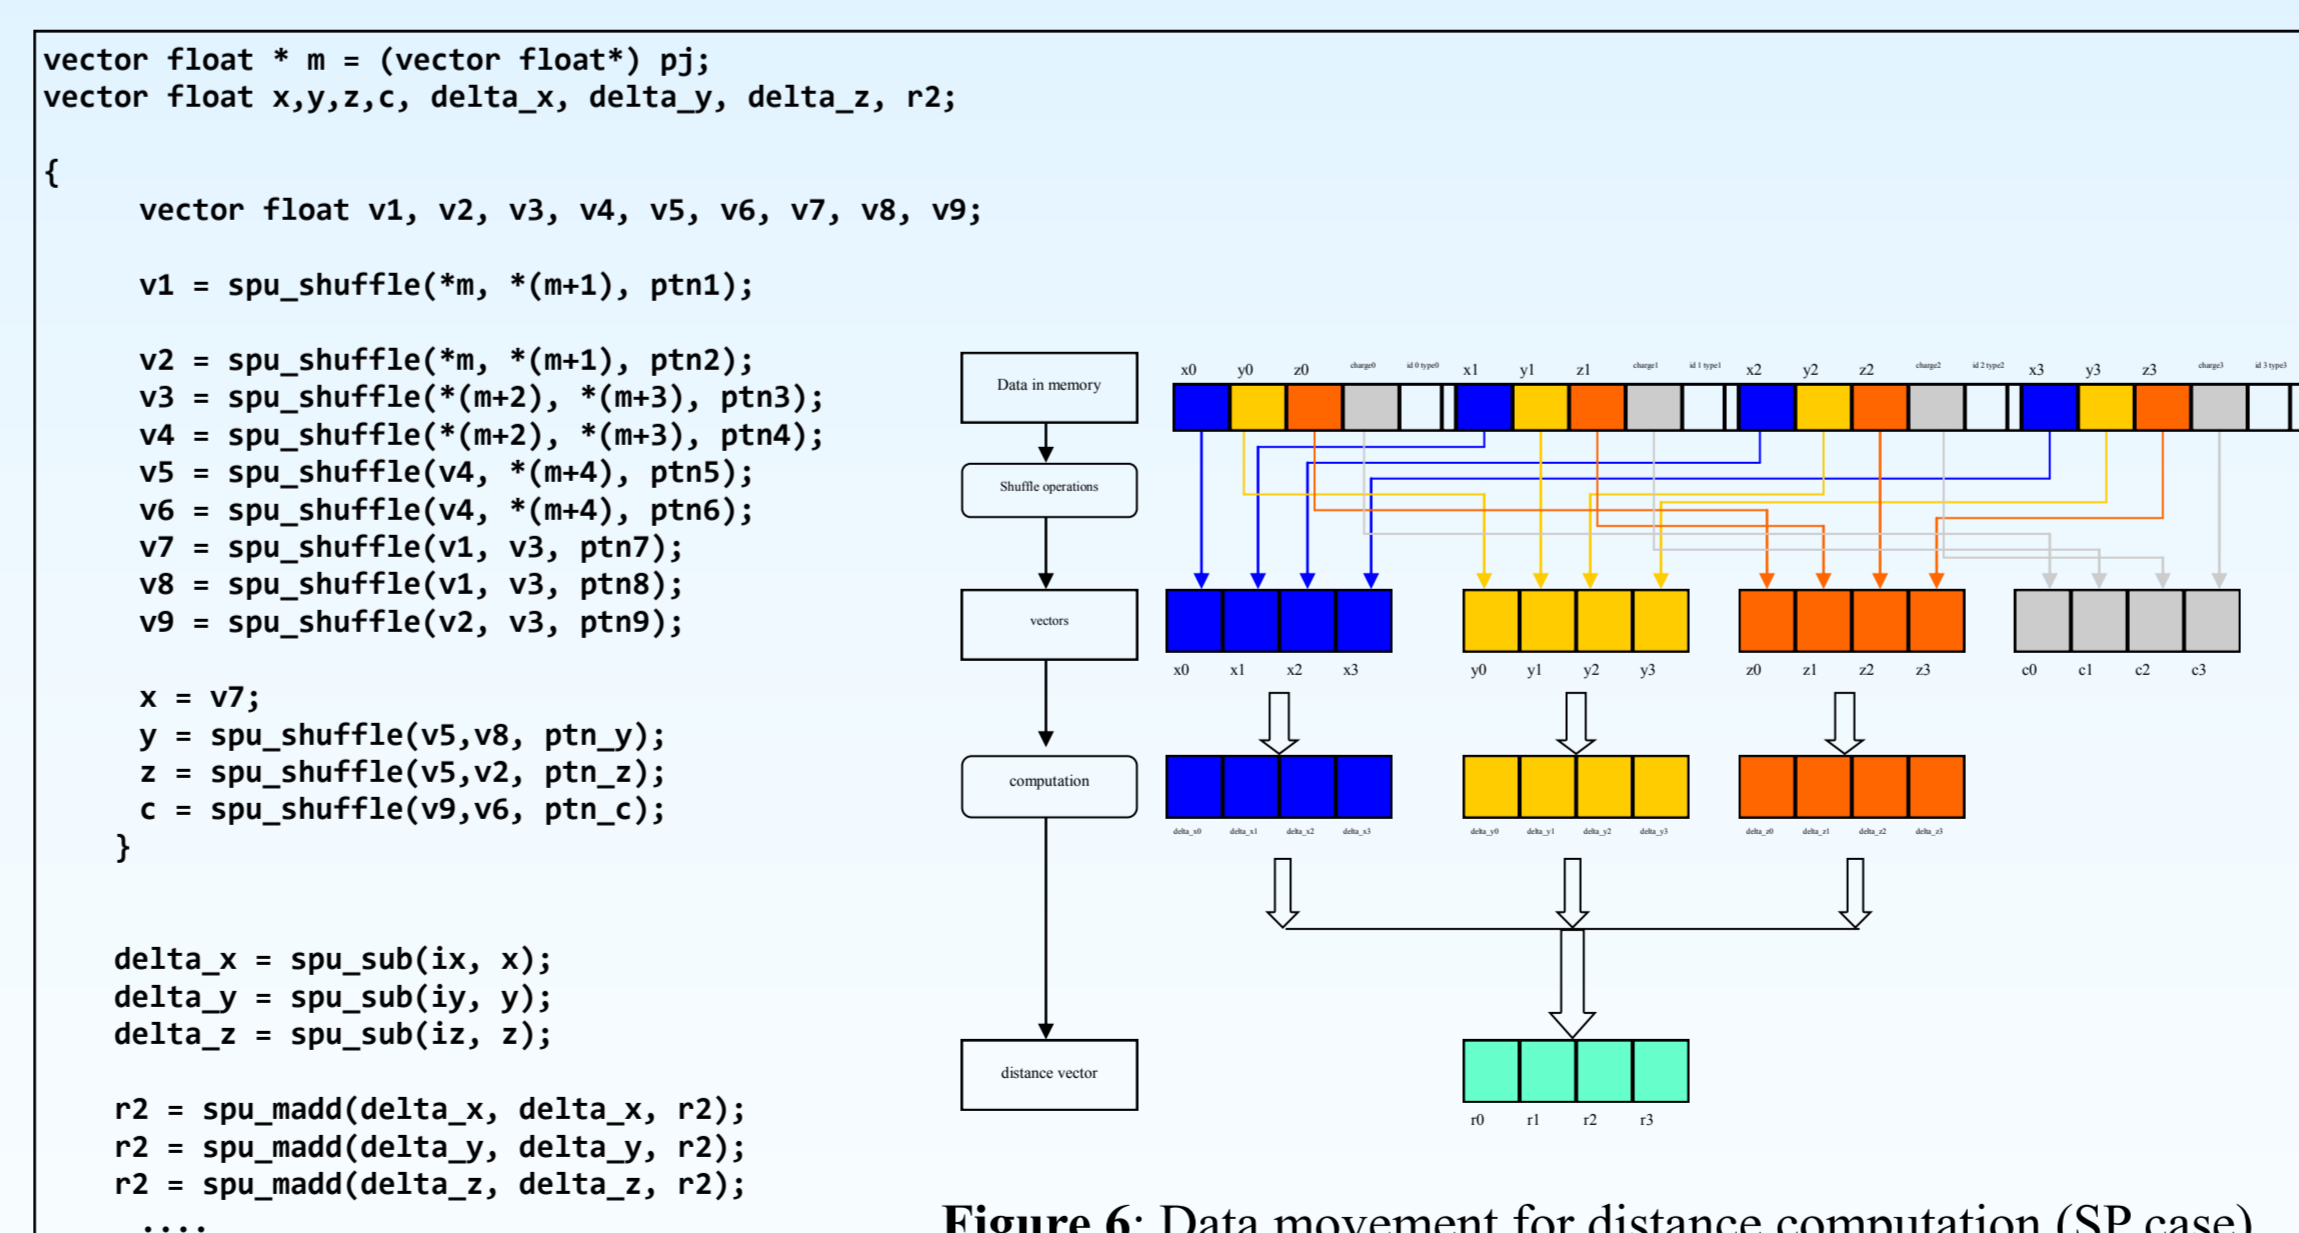


Figure 6: Data movement for distance computation (SP case)

Performance

- Distance computation code takes most of the execution time
- Data manipulation time is significant

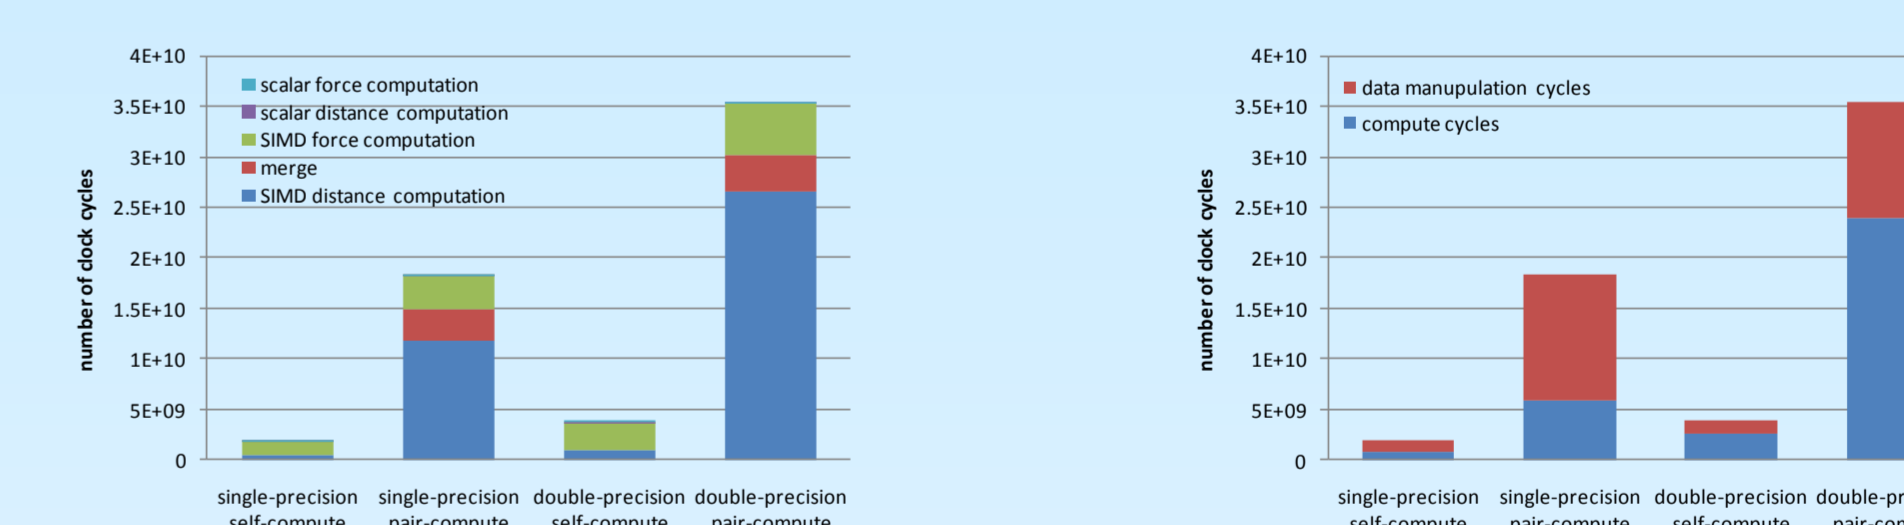


Figure 9: Static analysis of the computing kernels performed with the help of the timing tool, spu_timing

- Perfect scaling on multiple SPEs
- >10x speedup compared with Intel Xeon processor

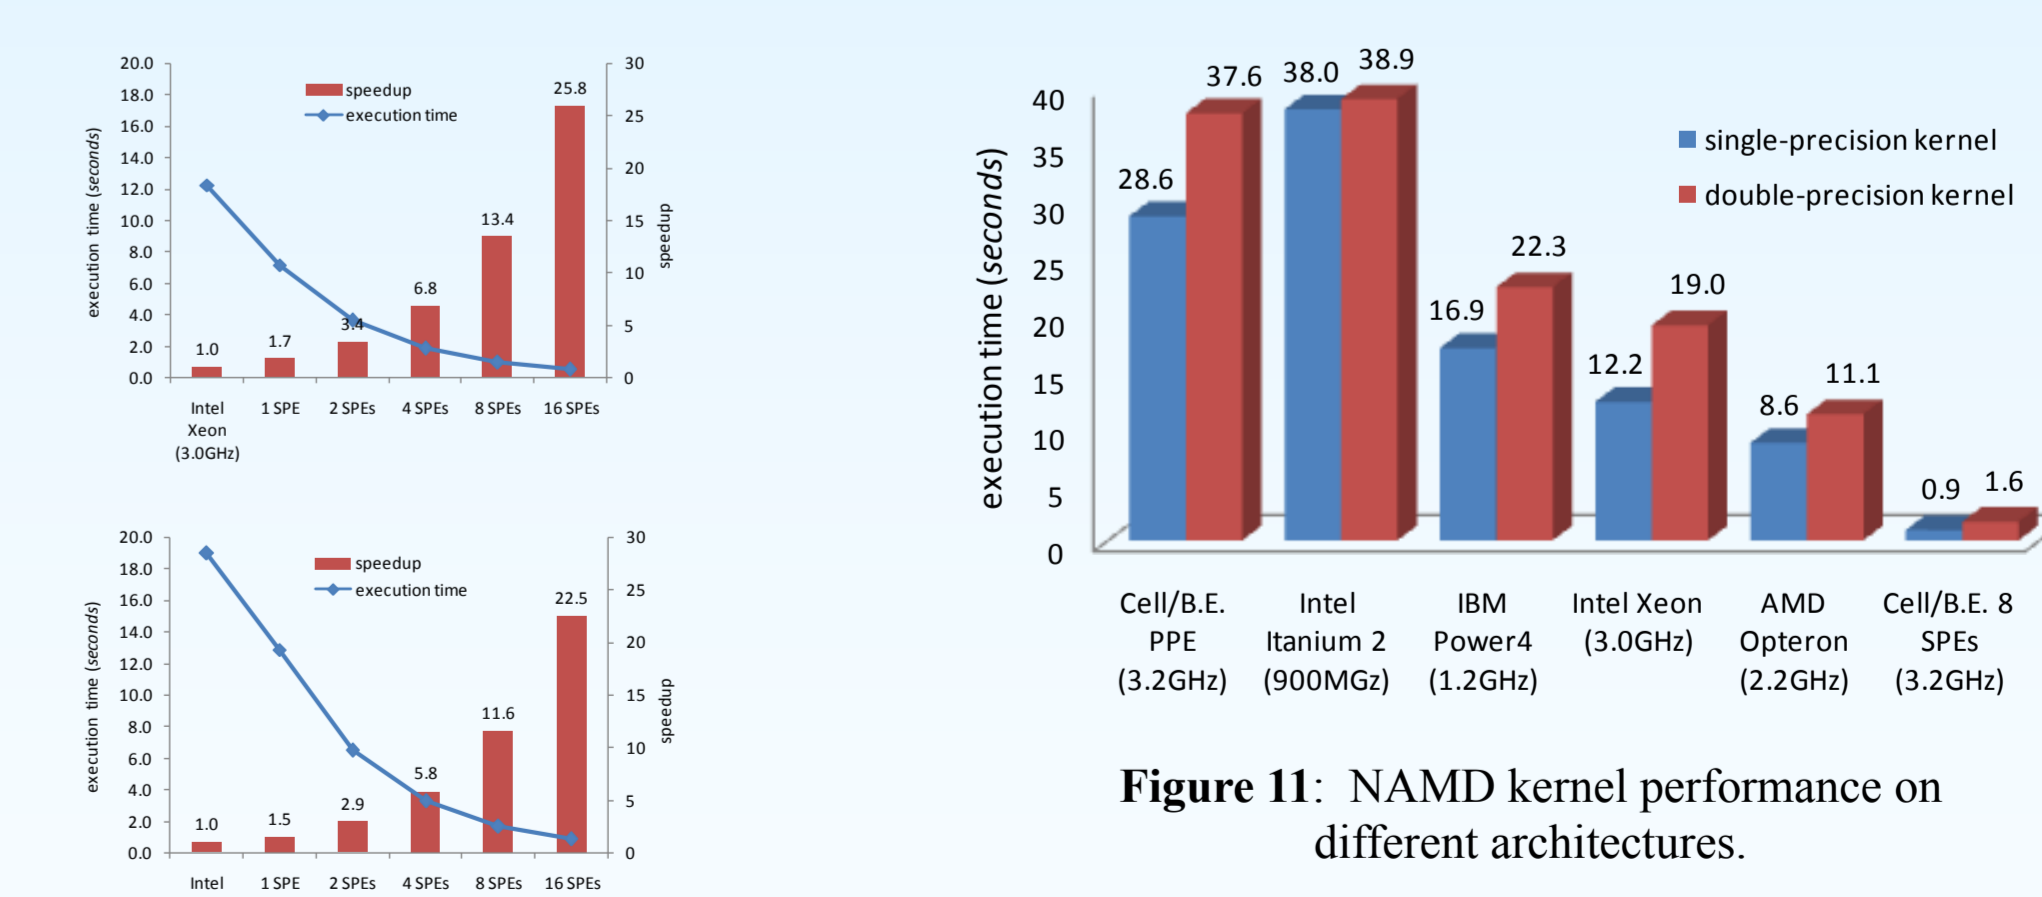


Figure 10: Scaling and speedup of the force-field kernel implementation as compared to a 3.0 GHz Intel Xeon processor.



This work was funded by National Science Foundation grant SCI 05-25308 and by the IBM Linux Technology Center. NAMD was developed by the Theoretical and Computational Biophysics Group in the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign. We are grateful to Dr. James Phillips from the Theoretical and Computational Biophysics Group for his help with NAMD. We acknowledge Georgia Institute of Technology, its Sony-Toshiba-IBM Center of Competence, and the National Science Foundation for the use of Cell Broadband Engine resources that have contributed to this research. We are also thankful to Dr. Paul Woodward from the University of Minnesota for providing access to the Cell blade system. Special thanks to Jeremy Enos from NCSA's Innovative System Lab (ISL) for installing and administering the Cell/B.E. blade system at NCSA, Dr. John Larson from ISL for useful comments on performance evaluation, and Trish Barker from NCSA's Office of Public Affairs for help in preparing this publication.

Acknowledgements

