Power-Aware Computing

MEGASCALE

http://www.para.tutics.tut.ac.jp/megascale/

HPC Cluster System Design by Adaptive Power Control

Conventional Design

- Number of nodes is decided to satisfy
  Sum of Peak Power ≤ Cooling Limit

Proposed Design

- Key observation : Effective power seldom reaches peak power
- Number of nodes is decided to satisfy
  Sum of Effective Power ≤ Cooling Limit
- More nodes integrated ⇒ Higher performance achieved
- To satisfy the constraint, runtime controller adjusts effective power by using DVFS

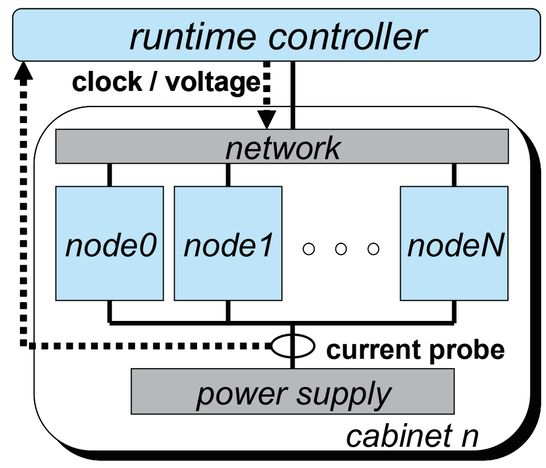

Runtime Control Algorithm

```
while(application_run){
  Pnow = Get_Power();

  if(Pnow > Plimit){
      Down_DVS();
      Tsafe = 0;
  }else if(Tsafe == Titvl){
      Up_DVS();
      Tsafe = 0;
  }
  Sleep(Tsampling);
  Tsafe += Tsampling;
}
```

Monitor power consumption every Tsampling

Select lower clock / voltage

Select higher clock / voltage every Titvl to keep high performance

Power Profile

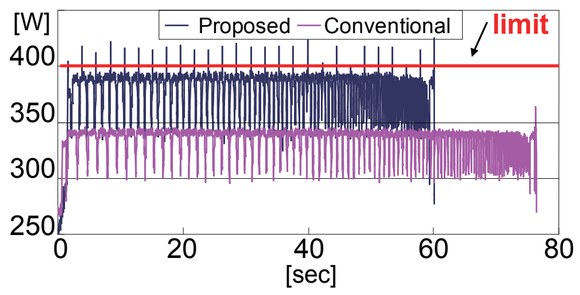

- 56.7% performance improvement (on average)
- Proposed system can use up to 10 nodes (conventional: 8 nodes) and higher clock frequency

Experimental Result

- Effective power of conventional design is far below the power limit
- Proposed system can fully utilize the available power budget

Performance Improvement

Energy Reduction Algorithm Using Slack Reclamation

- An energy reduction algorithm for parallel program represented in directed acyclic task graph (DAG).
- Our alogorithm shifts the gear for a non-crithical path as low and uniformly as possible reclaiming its slack.

standard (the fastest) gear

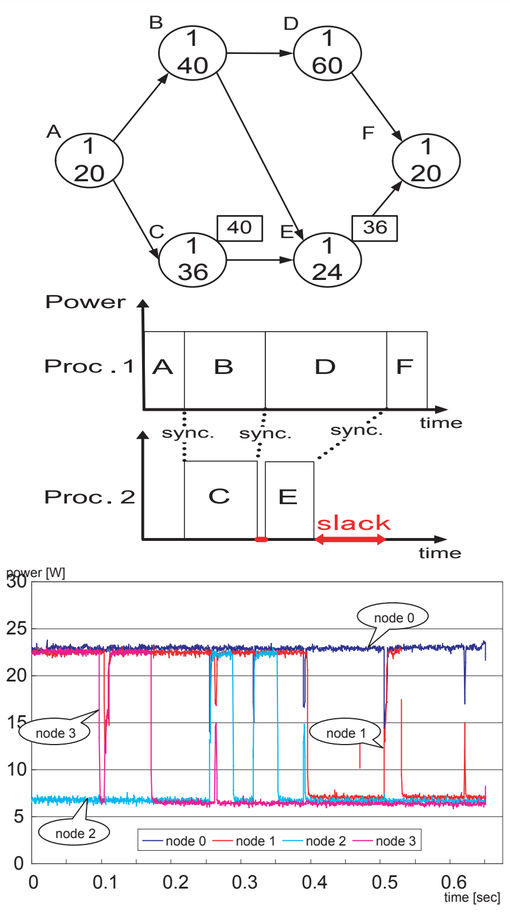

optimized

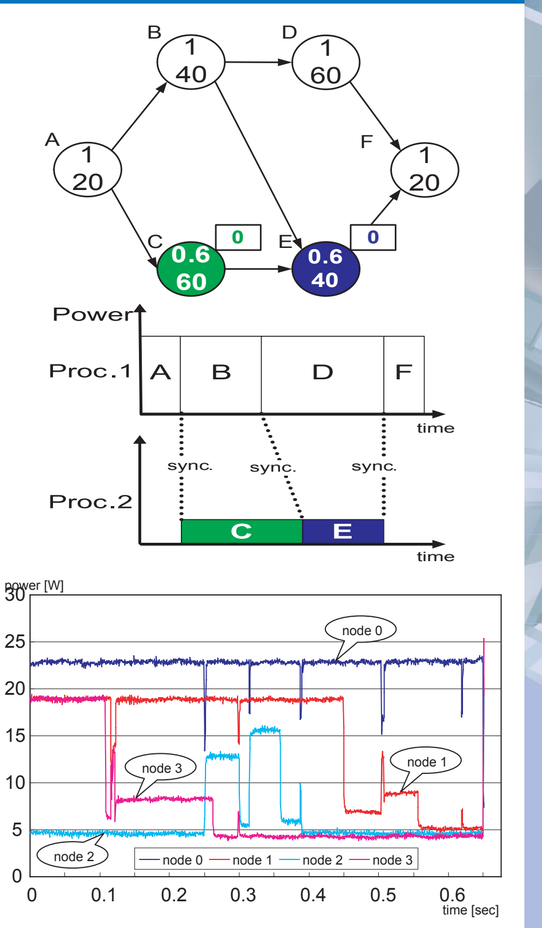

- Power profile of tree-based parallel program.
- Energy reduction by 11.8% without performance impact.

Profile-based Power-Performance Optimization by Using DVFS

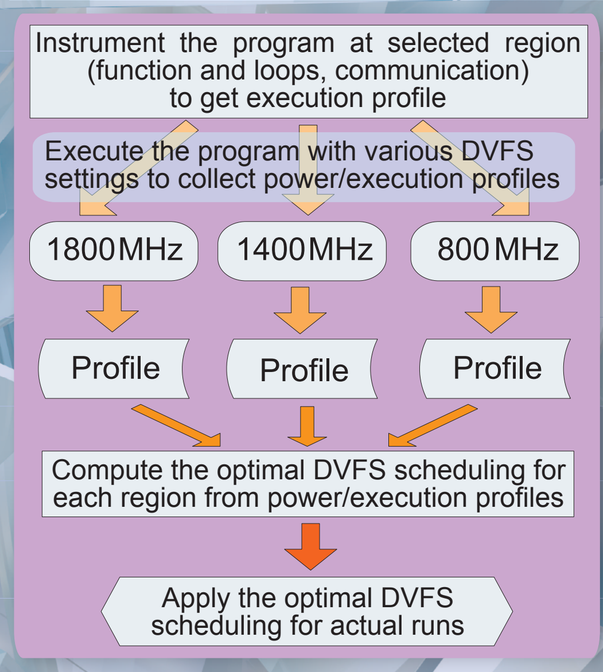

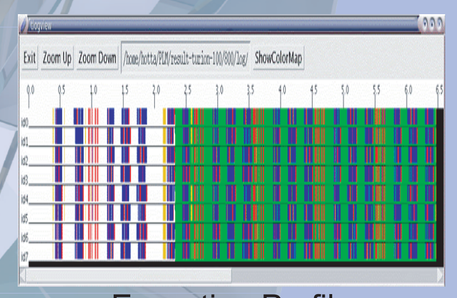

Execution Profile

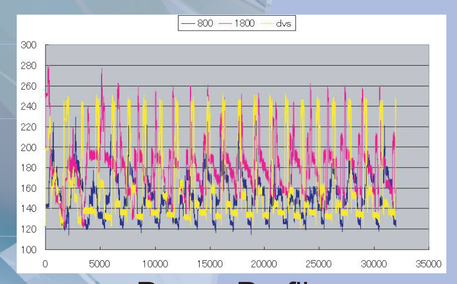

Power Profile

Algorithm & Experimental Result

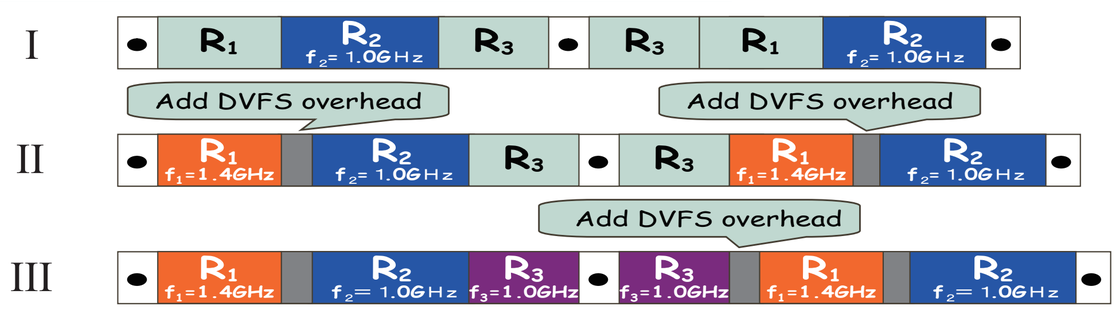

- Our algorithm takes DVFS overhead for transition into acount in a HPC cluster.
- Both PDP and EDP can be reduced compared with those at standard (the fastest) gear.