I/O Implementation and Evaluation of Parallel Pipelined STAP on High Performance Computers

Alok Choudhary

Wei-keng Liao,
Pramod Varshney, and
Donald Weiner

ECE Department
Northwestern University

EECS Department
Syracuse University

Outline

- Space-Time Adaptive Processing
- Parallel Pipelined Computational Model
  - Throughput and Latency
- Design and Implementation of a PRI-Staggered Post-Doppler STAP Algorithm
- Performance Evaluation
- Conclusions

Space-Time Adaptive Processing

- STAP refers to a class of signal processing methods which operate on data collected from a set of sensors over a given time interval
- STAP is used to detect weak target returns embedded in strong clutter, interference, and receiver noise in airborne radar
- STAP filters out the interference signals in two domains: frequency domain and direction of arrival domain
- Computational characteristics:
  - STAP contains a sequence of processing steps
  - A sequence of signal data arrives in a continuous fashion
  - Throughput and latency requirements

Parallel Pipelined Computational Model

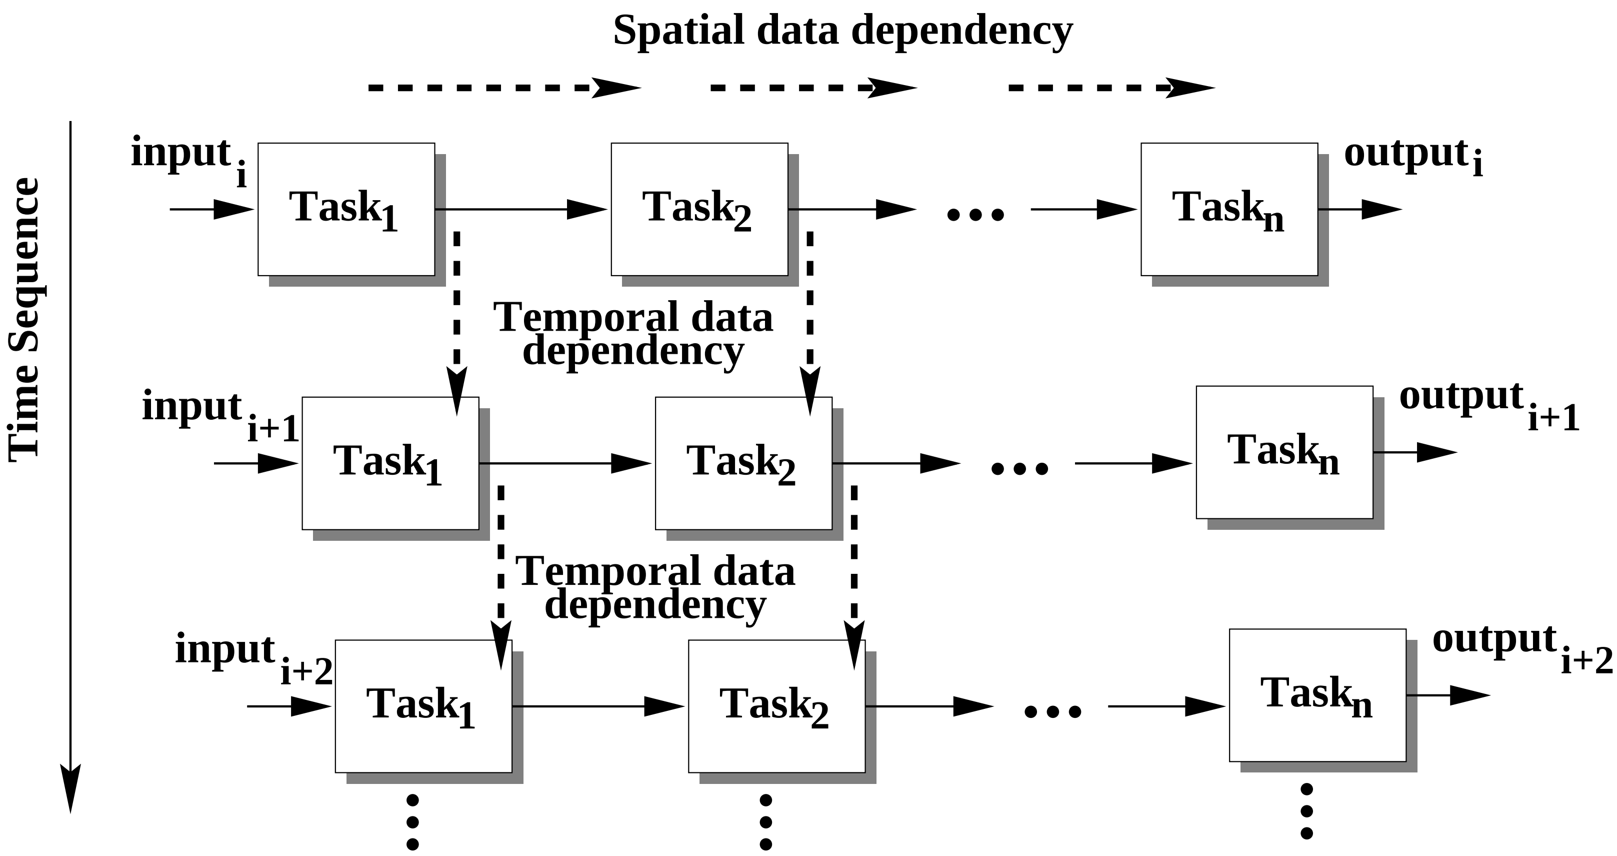

Throughput and Latency

- Throughput is determined by the task with the maximum execution time

$$throughput = \frac{1}{\max\limits_{0 \le i < n} T_i}$$

where T_i is the execution time of Task i

- Latency does not include the temporal data dependency (TD)

$$latency = \sum_{i=0}^{n-1} T_i - \sum_{i \in TD} T_i$$

PRI-Staggered Post-Doppler STAP

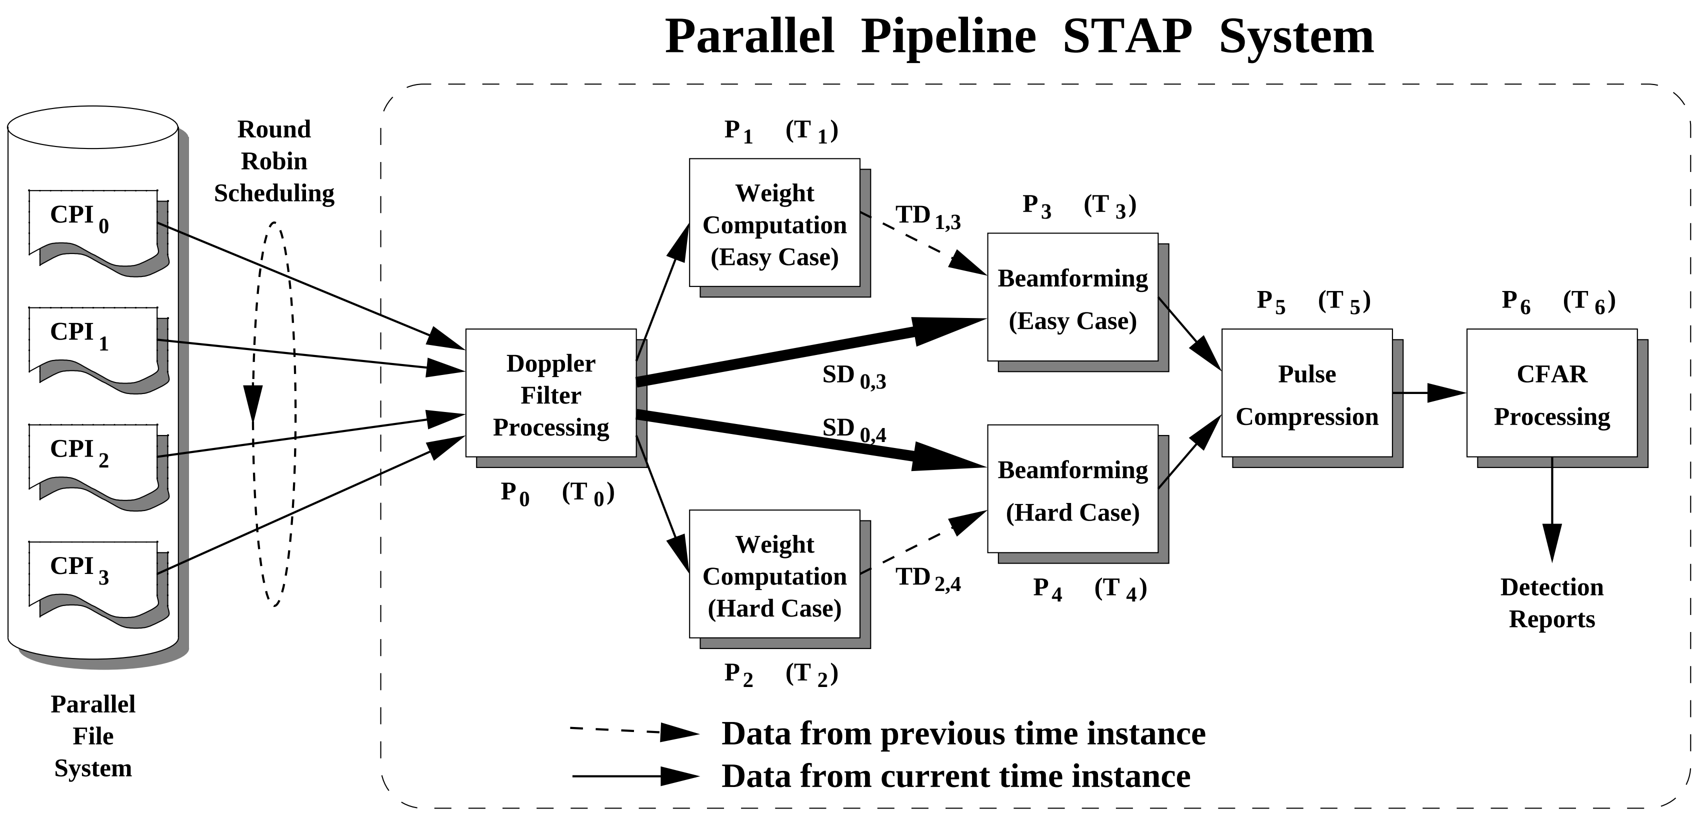

HPC System Platforms

	AFRL Paragon	ANL IBM SP
Configuration	MPP†	Cluster
CPU Type	i860 RISC	P2SC‡
RAM (MByte)	64	256
MFLOPS/proc	100	480
MHz /proc	40	120
No. nodes	232	80
Execution mode	dedicated	dedicated
File System	PFS	PIOFS
Stripe Factor	16, 64	80
Default Stripe Unit Size	64K bytes	64K bytes
I/O operation	Async / Block	Block Only

†MPP: Massively Parallel Processor

‡P2SC: Power 2 SuperScalar chip

Performance Results on Intel Paragon

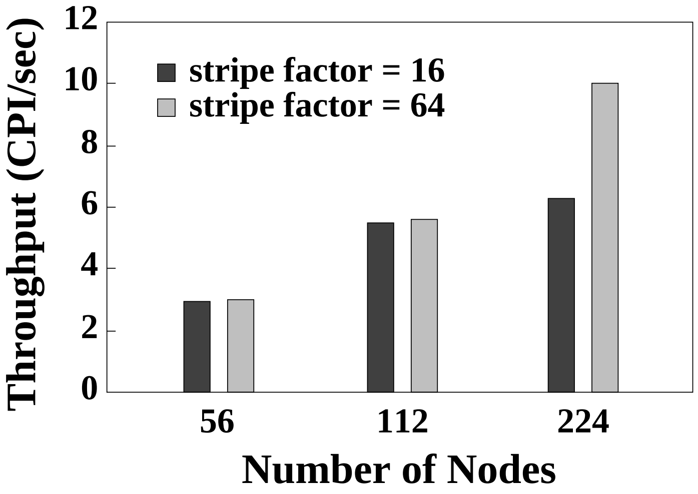

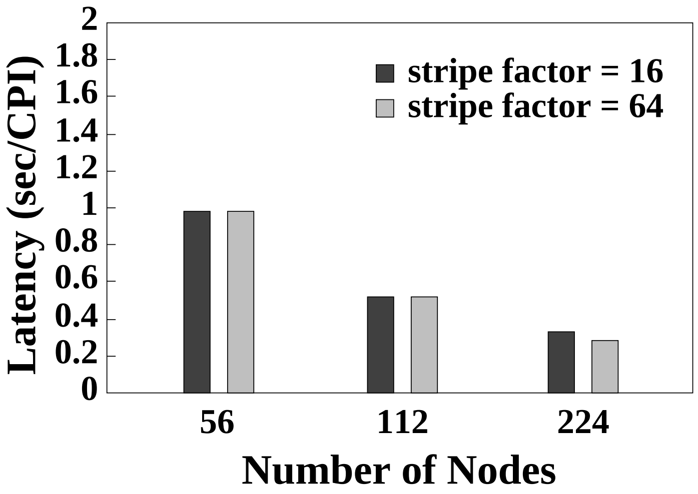

- Use two Caltech Paragon PFS systems: 16 and 64 striped directories
- CPI data files are stored across all striped directories with striped unit size of 64K bytes
- I/O bottleneck occurs on 224 nodes when using the PFS with 16 striped directories
- Throughput improved when the PFS with 64 striped directories is used

Performance Results on IBM SP-2

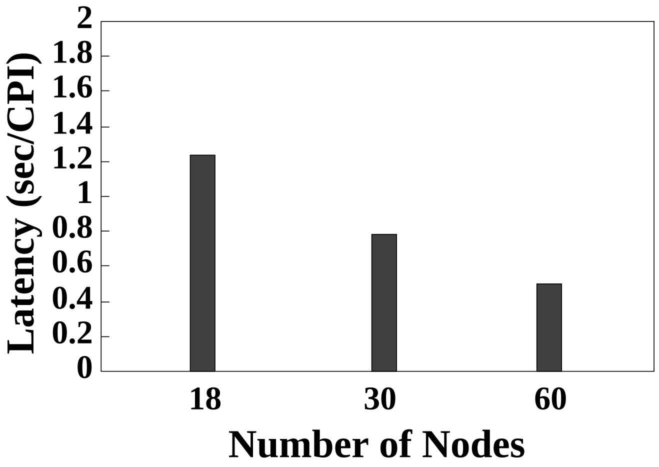

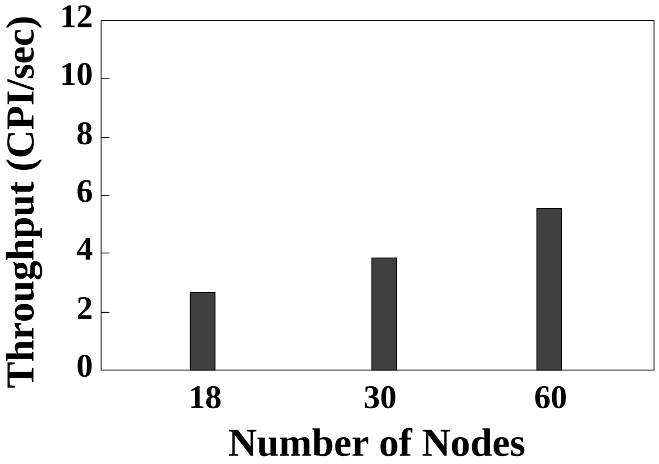

- Use the PIOFS at ANL with 80 striped directories
- CPI data files are stored across all striped directories with *default* striped unit size of 64K bytes
- No asynchronous I/O subroutine is available on SP, standard UNIX C read is used instead
- Performance limited by the inability to overlap I/O with computation and communication

Conclusions

- Parallel pipelined computational model designed for STAP applications on commercial HPC systems
- This model allows us to improve performance in terms of throughput and latency
- Full STAP application parallelized and implemented
- Disk I/O implemented using parallel file systems
- Scalable performance for both throughput and latency