Section 10. Resilient Iterative Methods

Whatever Has Been Done Can Be Outdone

HPC Top500 Computers:

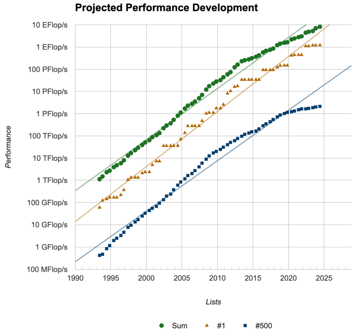

Obstacles to scalable simulation

- Model complexity
- Extreme concurrency
- Multiple levels of parallelism
- Complex memory hierarchies
- More costly data movement
- Hardware failures and soft error

Challenges to Keep Up With the Moore's Law: [Waldrop, Nature, 2016]

- Instruction-level parallelism (ILP) wall: availability of parallel instructions for a multi-core chip
- Power wall: the chip's overall temperature and power consumption

$$\text{Dynamic Power} = K \cdot (\text{Capacitive Load}) \cdot (\text{Voltage})^2 \cdot (\text{Frequency})$$

- Memory wall: bandwidth/latency of the channel b/w CPU and RAM

Avoiding Communications

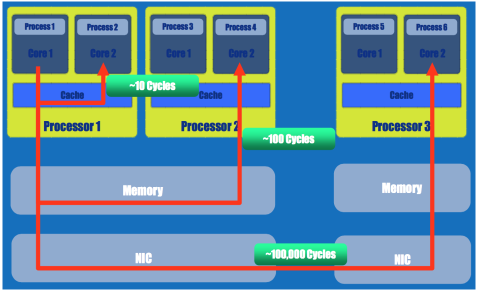

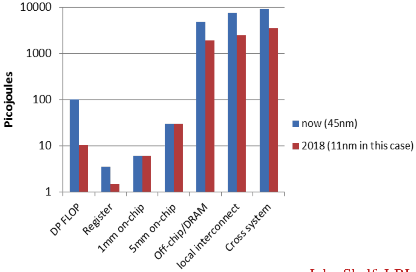

John Shalf, LBL

Floating-point time	≪	1 / Memory bandwidth	≪	Memory latency
59%		23%		5%
Floating-point time	≪	1 / Network bandwidth	≪	Network latency
59%		26%		15%

- Avoid communications (data movements) to save time $\Longrightarrow$ Redesign algorithms
- Linear algebra, LAPACK/ScaLAPACK, ... J. Demmel and collaborators

Improve Resilience Using Subspace Corrections

Improving Mean Time To Failure (MTTF)

- ASCI Q computer (12,288 EV-68 processors) in the LANL experienced 26.1 radiation-induced CPU failures per week [Michalak et al. 2005]
- BlueGene/L (128K processors) experiences one soft error in its L1 cache every 4–6 hours due to radioactive decay [Bronevetsky-Supinski 2008]

Dependability of Computer and Software

Properties:	availability, reliability, safety, integrity, ...
Threats:	fault, error, failure, ...
Resilience:	forecasting, preventing, removal, ...

Improving Resilience: Redundant Subspace Correction Method

- Maintain convergence when error occurs assuming it is detectable
- Introduce low computational overhead when no error occurs
- Require only small amount of point-to-point communication (locality) and maintain good load balance (granularity)

Redundant Subspace Corrections

Redundant subspace correction (RSC) methods

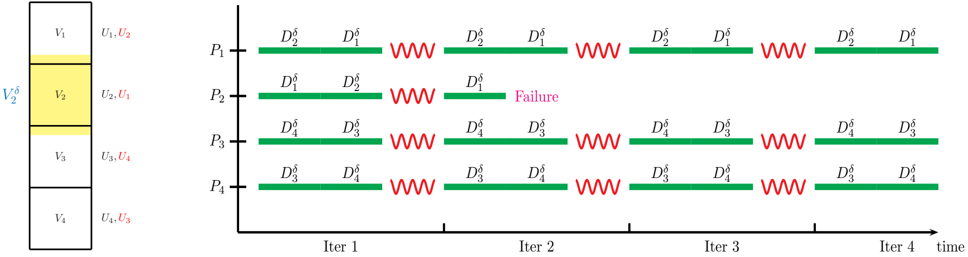

Our design goals:

- Light overhead: non-global communication needed
- Resilient to temporary or permanent hardware failures
- Redundancy to maintain convergence when components fail
- Globally parallel and locally successive subspace correction

Convergence Rate Analysis of RSC

Theorem (Xu, Zikatanov 2002)

SSC is convergent if each subspace solvers are convergent. In particular, if the subspace solvers are exact, then

$$\|I - BA\|_A^2 = \Big\| \prod_{i=1}^{J} (I - P_i) \Big\|_A^2 = 1 - \Big(\sup_{\|v\|_A=1} \inf_{\sum_i v_i = v} \sum_{i=1}^{J} \|P_i \sum_{j \geq i} v_i\|_A^2 \Big)^{-1}.$$

Theorem (Cui, Xu, Z. 2016)

If an error occurs during computation, the convergence rate of the successive RSC satisfies

$$\|I - B_{\mathrm{SRSC}} A\|_A \leq \|I - B_{\mathrm{SSC}} A\|_A.$$

If there is no error during computation, the convergence rate satisfies that

$$\|I - B_{\mathrm{SRSC}} A\|_A \leq \|I - B_{\mathrm{SSC}} A\|_A \; \|I - \tilde{B}_{\mathrm{SSC}} A\|_A.$$

Parallel Scalability of RSC for Poisson

DOFs	#Cores	Error-Free			With Error		
		#Iter	Time	Efficiency	#Iter	Time	Efficiency
1,335,489	16	12	8.09	—	13	8.13	—
2,146,689	32	13	8.64	75.25%	15	8.99	72.68%
4,243,841	64	14	8.91	72.13%	16	9.37	68.93%
10,584,449	128	19	12.87	62.27%	20	13.95	57.73%
16,974,593	256	23	18.01	35.68%	25	19.13	33.76%
33,751,809	512	25	20.90	30.57%	27	26.11	24.59%

#Cores	Standard B_{PSC}			B_{PRSC} Error-Free			B_{PRSC} With Error		
	#Iter	Time	Speedup	#Iter	Time	Speedup	#Iter	Time	Speedup
16	42	162.9	—	21	163.4	—	24	165.3	—
32	48	82.71	1.97	25	82.79	1.97	27	83.35	1.98
64	48	41.14	3.96	26	42.13	3.88	28	43.62	3.79
128	50	20.95	7.78	27	22.66	7.21	29	23.95	6.91
256	51	11.84	13.8	27	13.46	12.1	29	14.03	11.8
512	52	6.91	23.6	28	7.43	21.9	29	7.79	21.2

Overhead of Redundant Subspace Correction

Error-Free	Poisson (2,146,689 DOFs)		Maxwell (1,642,688 DOFs)		Elasticity (823,875 DOFs)	
	#Iter	Time	#Iter	Time	#Iter	Time
Yes	44	70.73	63	68.76	73	223.14
No	48	81.01	67	74.28	74	229.21

Table: Convergence of the colorized SRSC (16 cores)

Example	DOFs	Standard B_{PSC}		B_{PRSC} Error-Free		B_{PRSC} With Error	
		#Iter	Time	#Iter	Time	#Iter	Time
Poisson	1,335,489	23	7.92	12	8.09	13	8.13
Maxwell	468,064	42	4.09	21	4.23	24	4.48
Elasticity	436,515	16	10.18	9	11.01	10	11.35

Table: Convergence of the PRSC preconditioner (16 cores)

Tests with Multiple Failures

#Cores	1 failure		2 failures		4 failures		8 failures		16 failures	
	#Iter	Time	#Iter	Time	#Iter	Time	#Iter	Time	#Iter	Time
16	24	165.3	32	192.5	38	225.8	46	261.2	—	—
32	27	83.35	37	107.1	42	119.5	45	122.1	50	129.8
64	28	43.62	35	50.4	36	52.1	42	57.9	49	65.5
128	29	23.95	31	24.2	36	27.1	41	30.3	47	34.6

Table: Convergence of the PRSC with failures for Poisson (16,974,593 DOFs)

A few comments:

- Light algorithmic overhead: non-global communication needed
- Resilient to temporary or permanent hardware failures
- Redundant subspaces to maintain convergence when components fail
- Globally parallel and locally successive subspace correction