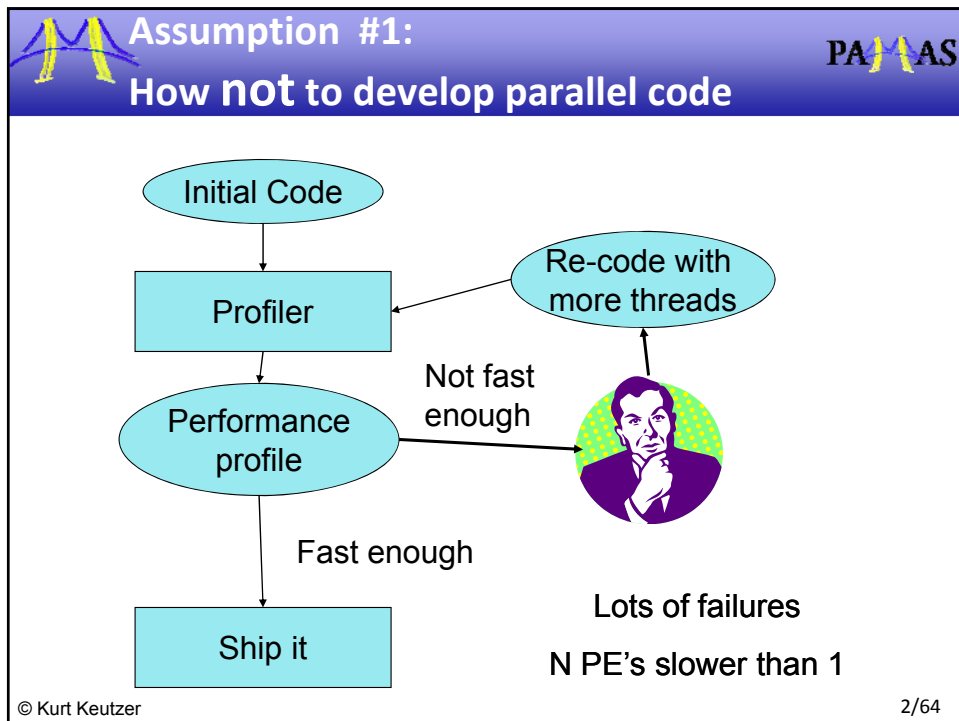



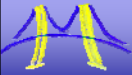
Architecting Parallel Software with Patterns

Kurt Keutzer


the PALLAS group,

Michael Anderson, Bryan Catanzaro, *Jike Chong*,
Katya Gonina, Dorothea Kolossa, Chao-Yue Lai,
 Mark Murphy, David Sheffield, Bor-Yiing Su,
 Narayanan Sundaram
 with thanks to Tim Mattson



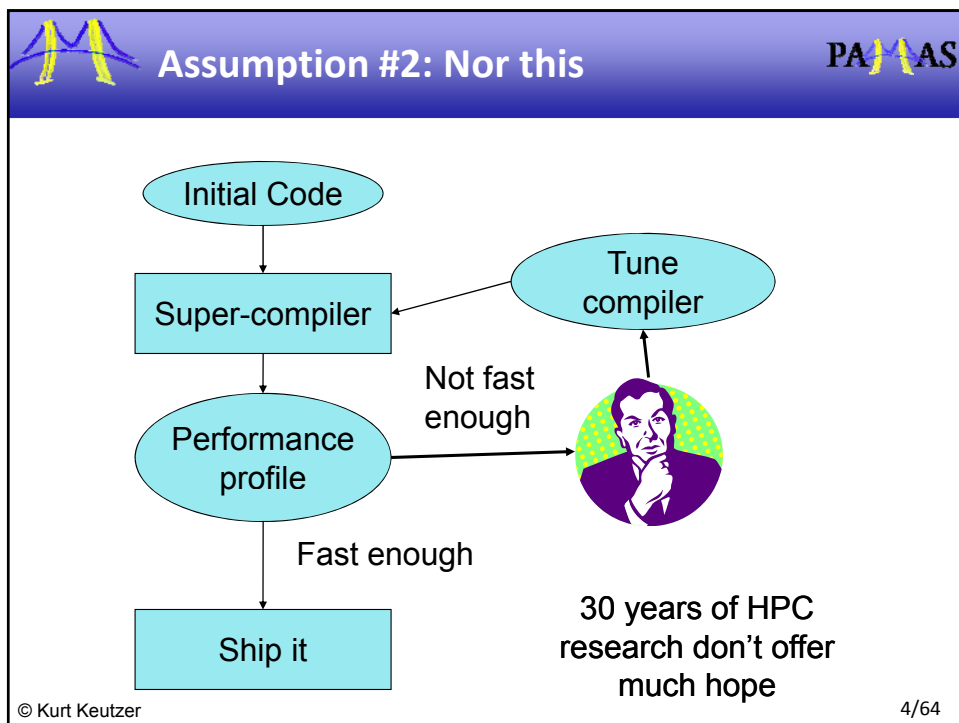


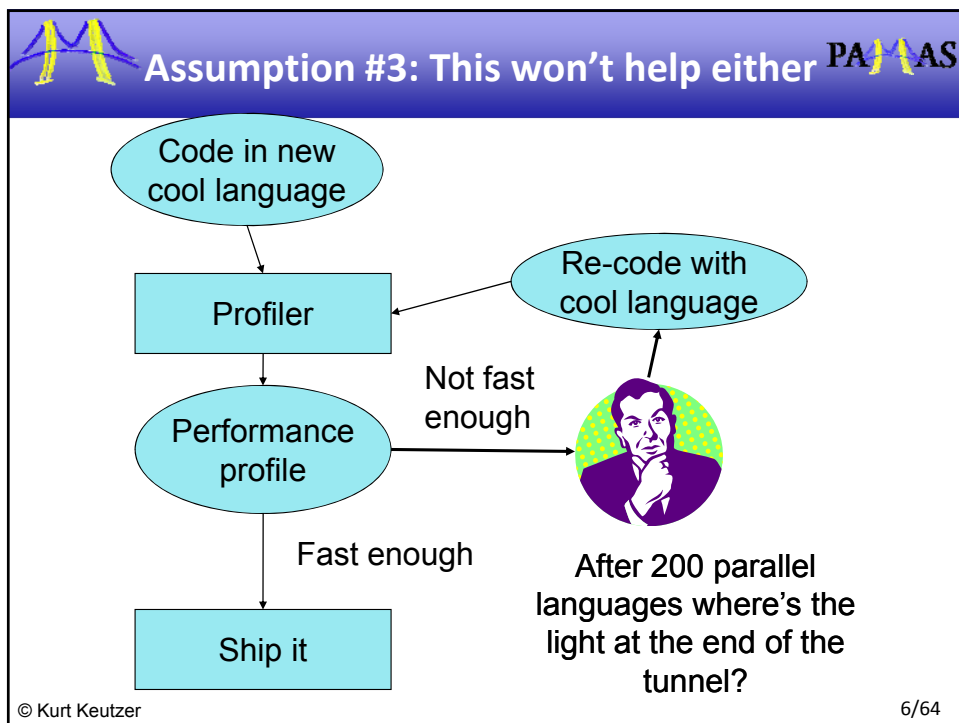
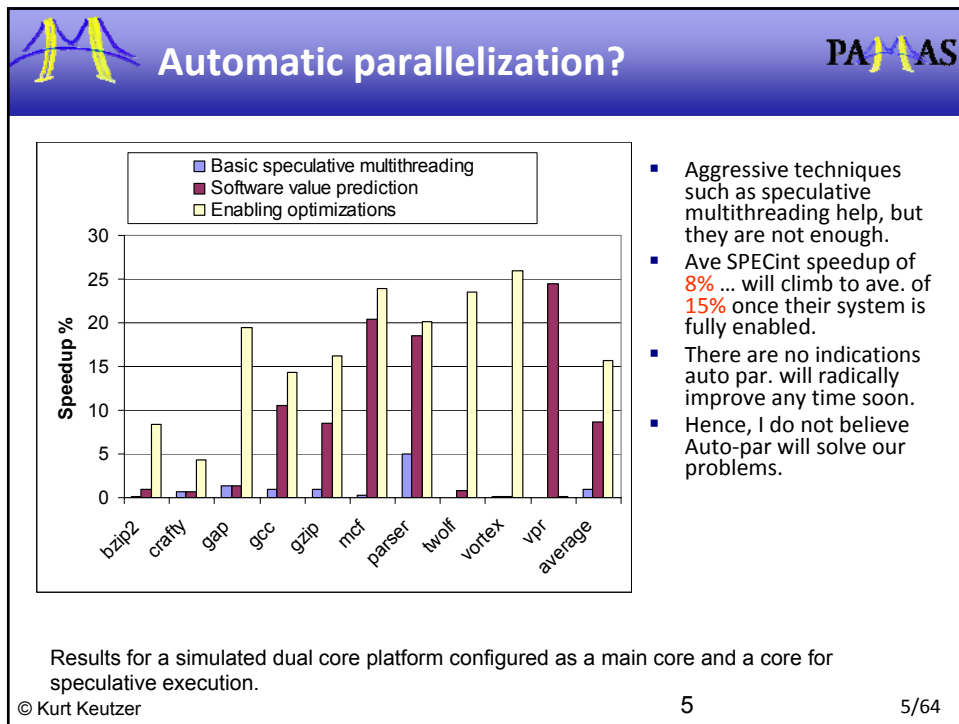
Steiner Tree Construction Time By Routing Each Net in Parallel




Benchmark	Serial	2 Threads	3 Threads	4 Threads	5 Threads	6 Threads
adaptec1	1.68	1.68	1.70	1.69	1.69	1.69
newblue1	1.80	1.80	1.81	1.81	1.81	1.82
newblue2	2.60	2.60	2.62	2.62	2.62	2.61
adaptec2	1.87	1.86	1.87	1.88	1.88	1.88
adaptec3	3.32	3.33	3.34	3.34	3.34	3.34
adaptec4	3.20	3.20	3.21	3.21	3.21	3.21
adaptec5	4.91	4.90	4.92	4.92	4.92	4.92
newblue3	2.54	2.55	2.55	2.55	2.55	2.55
average	1.00	1.0011	1.0044	1.0049	1.0046	1.0046


© Kurt Keutzer 3/64







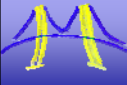
Parallel Programming environments in the 90's



ABCPL	CORRELATE	GLU	Mentat	Parafraze2	pC++
ACE	CPS	GUARD	Legion	Parallelation	SCHEDULE
ACT++	CRL	HAiL	Meta Chaos	Parallel-C++	SeTL
Active messages	CSP	Maskell	Midway	Parallaxis	POET
Adl	Cthreads	HPC++	Millipede	ParC	SDDA
Adsmith	CUMULVS	JAVAR	CparPar	ParLib++	SHMEM
ADDAP	DAGGER	HORUS	Mirage	ParLin	SIMPLE
AFAP	DAPPLE	HPC	MoC	Permanc	+++
ALWAN	DataParallel C	IMPACT	MOSIX	Parti	SISAL
AM	DC++	ISIS	Modula-P	pC	distributed smalltalk
AMDC	DCE++	JAVAR	Modula-2*	pC++	SML
AppLeS	DDD	JADE	Multipol	PCN	SONIC
Amosba	DICE	Java RMI	MPI	PCP	Split-C
ARTS	DIPC	javaPG	MPC++	PH	SR
Atlagas-on-9b	DOLIB	JavaSpace	Mumin	PEACE	Sthreads
Aurora	DOME	JIDL	Nano-Threads	PCU	Swand.
Automap	DORMOS	Joyce	NESL	PET	SUIF
bb_threads	DRL	Khoros	NetClasses++	PETSc	Synergy
Blaze	DSM-Threads	Karma	Nexus	PENNY	Telegraphos
BSP	Ease	KOAN-Fortran-9	Nimrod	Phosphorus	SuperPascal
BlockComm	ECO	LAM	NOW	POET	TCGMSG
C*	Effel	Lilac	Objective Linda	Polaris	Threads.h++
"C" in C	Eilean	Linda	Occam	POOMA	TreadMarks
C++	Emerald	JADA	Omega	POOL-T	TRAPPER
C++/OS	FPI	WWW++-A	OpenMP	PRATO	uc++
Cashmere	Encalibur	ISETL-Linda	Orca	P-RIO	UNITY
C4	Express	ParLin	OOF90	Prospero	UC
CC++	Falcon	Eilean	P++	Proteus	V
Chu	Filaments	P4-Linda	P3L	QPC++	V/C++
Charlotte	FM	Glenda	p4-Linda	PVM	Vistfold V-NUS
Charm	FLASH	POSYBL	Pablo	PSI	VPE
Charm++	The FORCE	Objective-Linda	PADE	PSDM	Win32 threads
Cid	Fork	LPS	PADRE	Quake	WinPar
Cik	Fortran-M	Locust	Panda	Quark	WWWinda
CM-Fortran	FX	Lparx	Papers	Quick Threads	XENOOFS
Converse	GA	Lucid	AFAP	Sage++	XPC
Code	GAMMA	Maisie	Para++	SCANDAL	Zounds
COOL	Glenda	Manifold	Paradigm	SAM	ZPL

© Kurt Keutzer

7/64

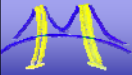


So What's the Alternative?




© Kurt Keutzer

8/64

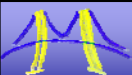


Principles of SW Design




- After 15 years in industry, at one time overseeing the technology of 25 software products, my two best principles to facilitate good software design are:
 - Use of modularity
 - Definition of invariants
- Modularity helps:
 - Architect: Makes overall design sound and comprehensible
 - Project manager:
 - As a manager I am able to comfortably assign different modules to different developers
 - I am also able to use module definitions to track development
 - Module implementors: As a module implementor I am able to focus on the implementation, optimization, and verification of my module with a minimum of concern about the rest of the design
 - Identify invariants and key computations

© Kurt Keutzer 9/64




Non-Principles of SW Design



- What's life like without modularity?
 - Spaghetti code
 - Wars over the interpretation of the specification
 - Waiting on other coders
 - Wondering why you didn't touch anything and now your code broke
 - Hard to verify your code in isolation, and therefore hard to optimize
 - Hard to parallelize without identifying key computations
- Modularity will help us obviate all these
- Parnas, "On the criteria to be used on composing systems into modules," CACM, December 1972.

© Kurt Keutzer 10/64



Modularity is important But ...

PAAS

Pop quiz: Is software more like?


a) A building

b) A factory



© Kurt Keutzer

11/64



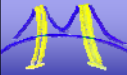
What computations we do is as important than how we do them

PAAS


Apps Dwarves	Embed	SPEC	DB	Games	ML	HPC	CAD	Health	Image	Speech	Music	Browser
Graph Algorithms												
Graphical Models												
Backtrack / B&B												
Finite State Mach.												
Circuits												
Dynamic Prog.												
Unstructured Grid												
Structured Grid												
Dense Matrix												
Sparse Matrix												
Spectral (FFT)												
Monte Carlo												
N-Body												

© Kurt Keutzer

12/64

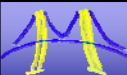


Architecting Parallel Software




- Putting computation and structure together:
 - We believe the key to productively building efficient and correct parallel software is *software architecture*
- A *software architecture* is a hierarchical composition of:
 - Computational patterns – the atoms
 - Structural patterns – the molecular bonds
- This software architecture naturally gives:
 - Modularity
 - Efficient management
 - Efficient implementation
 - Efficient verification
 - Identifies key computations, invariants, and interfaces

© Kurt Keutzer 13/64

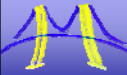


Outline




- ➔ ■ Architecting Parallel Software
- Structural Patterns
- Computational Patterns
- Examples
- Summary

© Kurt Keutzer 14/64

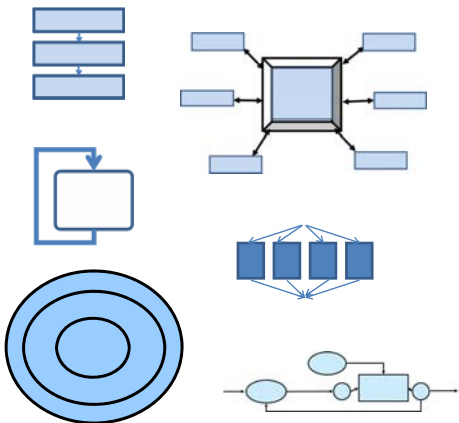


Identify the SW Structure



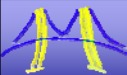
Structural Patterns

- Pipe-and-Filter
- Agent-and-Repository
- Event-based
- Layered Systems
- Model-view-controller
- Arbitrary Task Graphs
- Puppeteer
- Iterator/BSP
- MapReduce




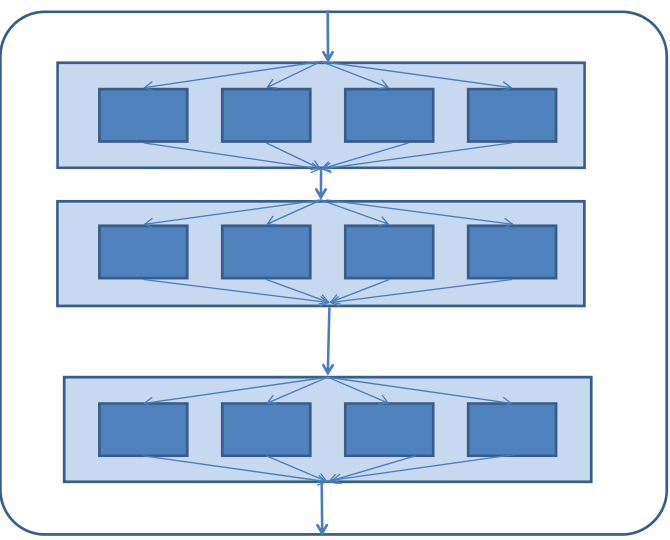
These define the structure of our software but they *do not* describe what is computed

© Kurt Keutzer
15
15/64

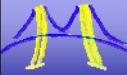



Analogy: Layout of Factory Plant





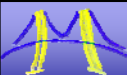

© Kurt Keutzer
16
16/64

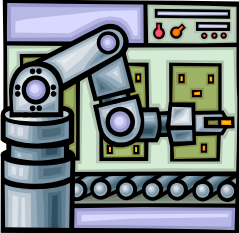

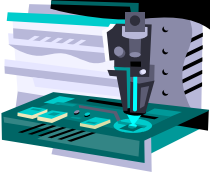
 **Identify Key Computations** 

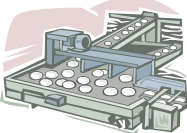

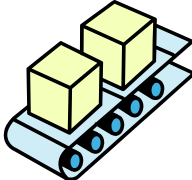
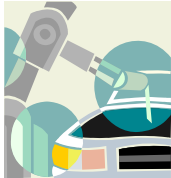
Apps Dwarves	Embed	SPEC	DB	Games	ML	HPC	CAD	Health	Image	Speech	Music	Browser
Graph Algorithms												
Graphical Models												
Backtrack / B&B												
Finite State Mach.												
Circuits												
Dynamic Prog.												
Unstructured Grid												
Structured Grid												
Dense Matrix												
Sparse Matrix												
Spectral (FFT)												
Monte Carlo												
N-Body												

■ Computational patterns describe the key computations but not how they are implemented

© Kurt Keutzer 17/64

 **Analogy: Machinery of the Factory** 

© Kurt Keutzer 18 18/64

Architecting the Whole Application PAMAS

SW Architecture of Large-Vocabulary Continuous Speech Recognition

- Raises appropriate issues like scheduling, latency, throughput, workflow, resource management, capacity etc.

Analogous to the design of an entire manufacturing plant

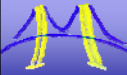

© Kurt Keutzer 19 19/64

Outline PAMAS

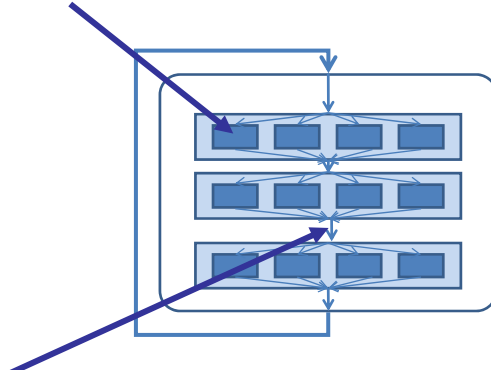
- Architecting Parallel Software
- ➔ ■ Structural Patterns
 - Pipe and filter
 - Iterator
 - Map Reduce
- Computational Patterns
- Examples
- Summary

© Kurt Keutzer 20/64

Elements of a structural pattern

- Components are where the computation happens

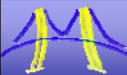



- A configuration is a graph of components (vertices) and connectors (edges)
- A structural patterns may be described as a family of graphs.

Connectors are where the communication happens

© Kurt Keutzer
21/64

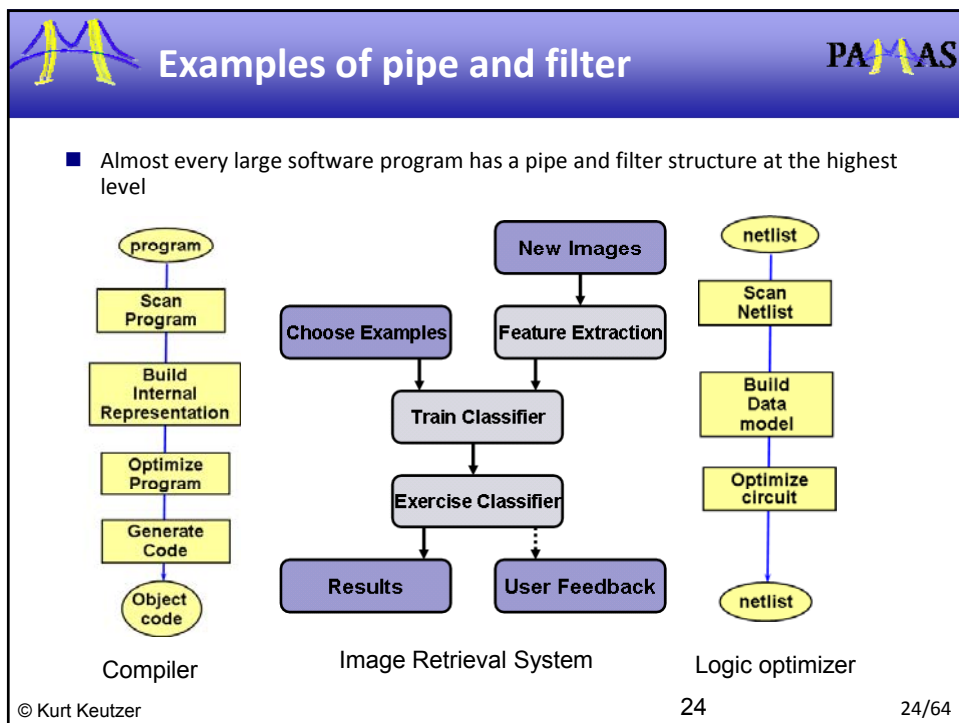
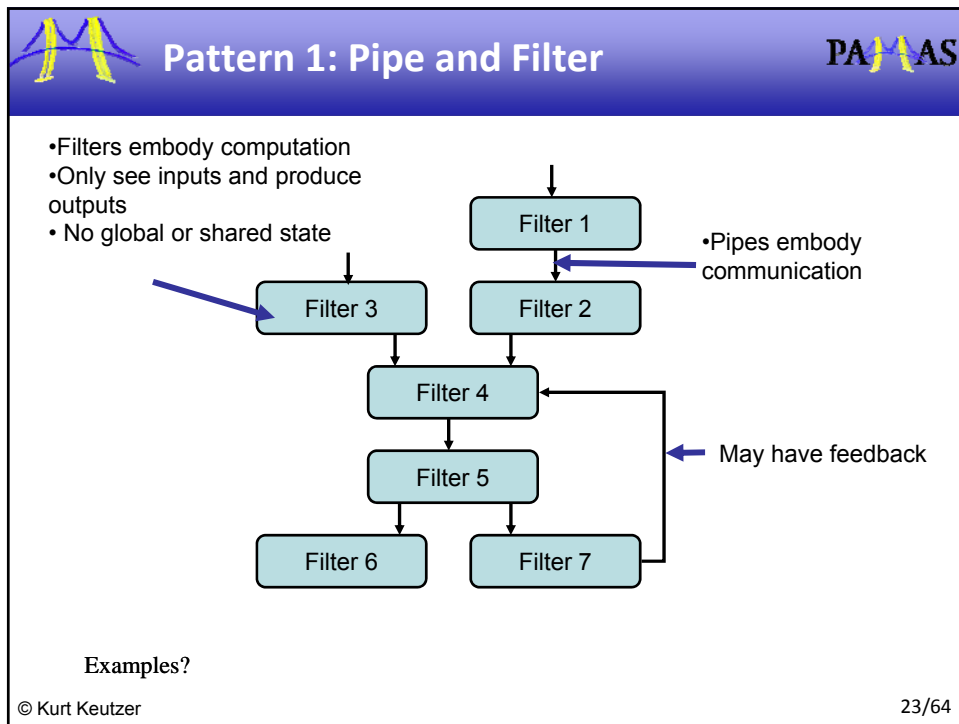
Inventory of Structural Patterns

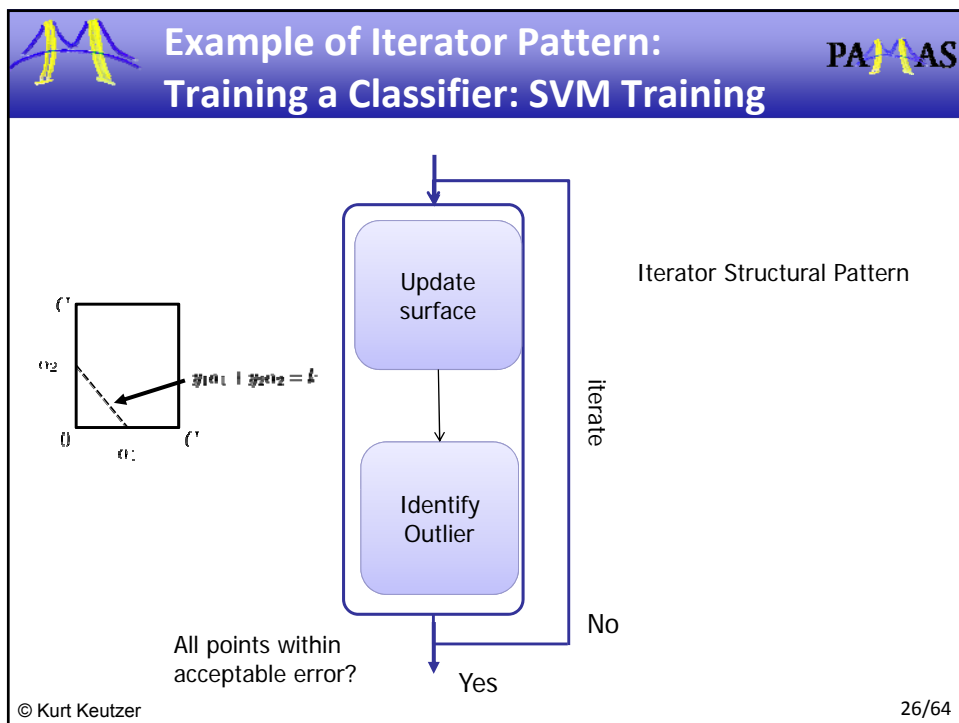
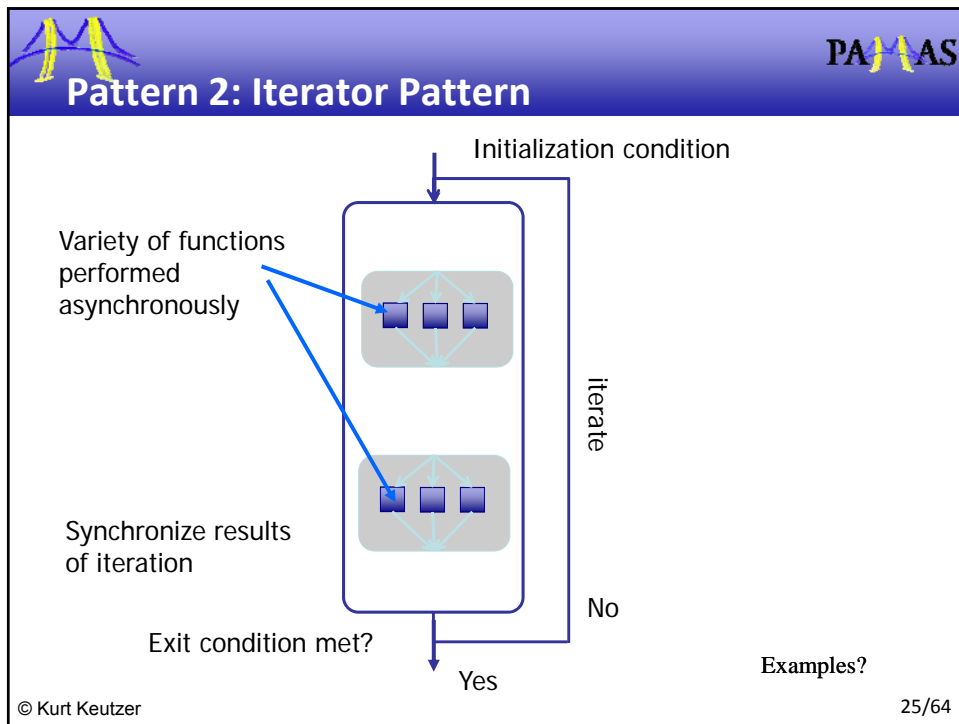



- Pipe-and-Filter
- Agent-and-Repository
- Event-based
- Layered Systems
- Model-view-controller
- Arbitrary Task Graphs
- Puppeteer
- Iterator/BSP
- MapReduce

- We build arbitrarily complex software structures out of these nine patterns

© Kurt Keutzer
22/64

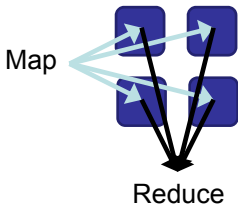




Pattern 3: MapReduce

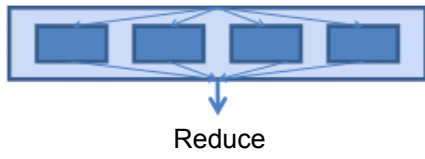
❖ To us, it means

- A map stage, where data is mapped onto independent computations
- A reduce stage, where the results of the map stage are summarized (i.e. reduced)



Map

Reduce



Map

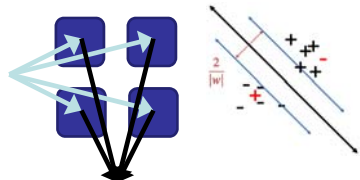
Reduce

Examples?

© Kurt Keutzer 27/64

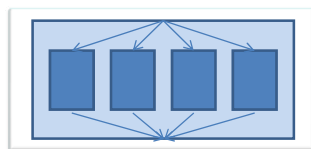
Examples of Map Reduce

- General structure:
- Map a computation across distributed data sets
- Reduce the results to find the best/(worst), maxima/(minima)



Support-vector machines (ML)

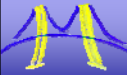
- Map to evaluate distance from the frontier
- Reduce to find the greatest outlier from the frontier




Speech recognition

- Map HMM computation to evaluate word match
- Reduce to find the most-likely word sequences

© Kurt Keutzer 28/64

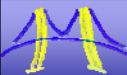


Outline




- Architecting Parallel Software
- Structural Patterns
- ➔ ■ Computational Patterns
 - Linear Algebra
 - Spectral Methods
 - Dynamic programming
- Examples
- Summary

© Kurt Keutzer
29/64



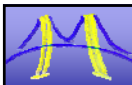

Inventory of Key Computations



Apps Dwarves	Embed	SPEC	DB	Games	ML	HPC	CAD	Health	Image	Speech	Music	Browser
Graph Algorithms												
Graphical Models												
Backtrack / B&B												
Finite State Mach.												
Circuits												
Dynamic Prog.												
Unstructured Grid												
Structured Grid												
Dense Matrix												
Sparse Matrix												
Spectral (FFT)												
Monte Carlo												
N-Body												

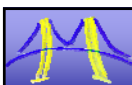

■ We build arbitrarily complex computations out of these thirteen computational patterns

© Kurt Keutzer
30/64


CP1: Linear Algebra


- **Vector Space:** A set closed under + has identity and inverse elements, scalar multiplication
- **Linear Map:** Operator T on vectors u, v , scalar α s.t.
 $T(u + v) = Tu + Tv$, and $T(\alpha v) = \alpha T(v)$
- **Matrix:** An $m \times n$ array of numbers representing a Linear map from R^n to R^m
- **Linear Equations:** $Ax = b$
- **Eigenvalues/vectors:** $Ax = \lambda x$

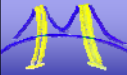

© Kurt Keutzer
31/64

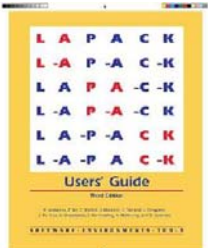

Basic Linear Algebra Subroutines (BLAS)


- Three "Levels", known as BLAS, characterized by intrinsic ratio of computation to memory movement

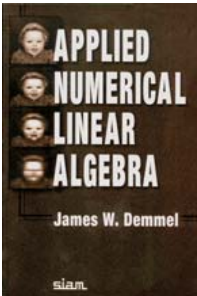
Level	Example	# mem refs	# flops	q
1	xAXPY: $y = y + \alpha x$	$3n$	$2n^1$	$2/3$
2	xGEMV: $y = y + Ax$	n^2	$2n^2$	2
3	xGEMM: $C = C + AB$	$4n^2$	$2n^3$	$n/2$

© Kurt Keutzer
32/64

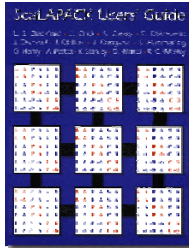
 **Linear Algebra Resources** 




www.netlib.org/lapack




gams.nist.gov



www.netlib.org/scalapack

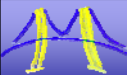



www.netlib.org/templates


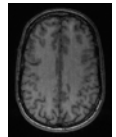


www.cs.utk.edu/~dongarra/etemplates

© Kurt Keutzer 33/64

 **CP2: Spectral Methods Pattern: MRI Reconstruction** 

- "Spectral Methods" are a broad class of numerical algorithms for solving PDEs, but notions of Spectral Analysis (i.e. convenient changes of basis) are important in every application area
- In Magnetic Resonance Imaging (MRI), images are collected in "k-space" -- i.e. an MRI scan produces a Fourier Domain image
- Fourier and Wavelet representations are different Spectral analyses that expose different properties of images convenient for solving our problems

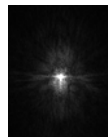
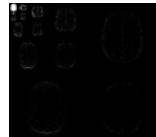



MRI Scan
== DFT


→

Wavelet
Xform

→

© Kurt Keutzer 34/64



Spectral Methods Pattern: Fast Transforms


PAAS

- Spectral Methods rely on representations of data in "convenient" bases that produce working, computationally *feasible* algorithms
- Changing a basis is, in general, an $O(N^2)$ matrix-vector multiplication. The matrices representing "convenient" bases factor into $O(N \log N)$ fast transforms!

$$U_k = \sum_{j=0}^{N-1} x_j \omega^{jk}$$


$$F = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & \omega & \omega^2 & \omega^3 \\ 1 & \omega^2 & \omega^4 & \omega^6 \\ 1 & \omega^3 & \omega^6 & \omega^9 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -i \\ 1 & i & -1 & -i \end{bmatrix}$$

$$F_N x = \begin{bmatrix} I_{N/2} & D_{N/2} \\ I_{N/2} & -D_{N/2} \end{bmatrix} \begin{bmatrix} F_{N/2} x_{\text{even}} \\ F_{N/2} x_{\text{odd}} \end{bmatrix}$$



© Kurt Keutzer


35/64



Spectral Methods Pattern: Libraries

PAAS

- Fast transform algorithms like the FFT are notoriously difficult to optimize:
- Luckily, implementations of the FFT exist for every platform. E.G:
 - FFTW and SPIRAL: Highly successful auto-tuners for FFT (and others) on PCs and workstations
 - CUFFT for Cuda on Nvidia GPUs



© Kurt Keutzer

36/64

CP3: Dynamic Programming PAAS

- Class of problems for which the optimal solution can be built up from optimal solutions to sub-problems
- Principle of optimality: Optimal cover for a tree consists of a best match at the root of the tree plus the optimal cover for the sub-trees starting at each input of the match

Best cover for this match uses best covers for x, y, z

Choose least cost tree-cover at root

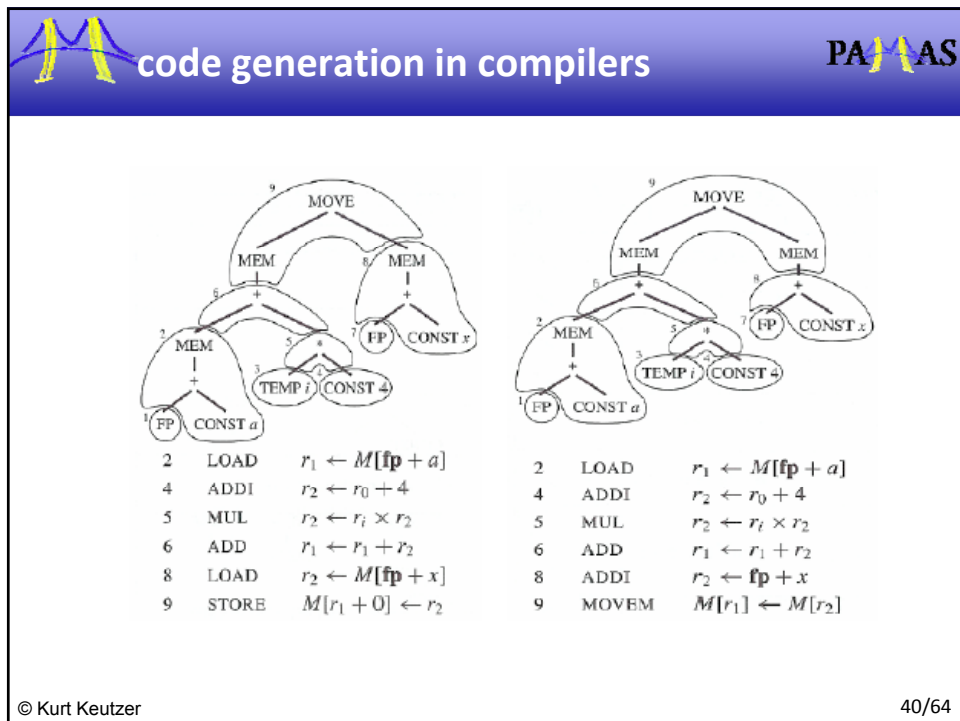
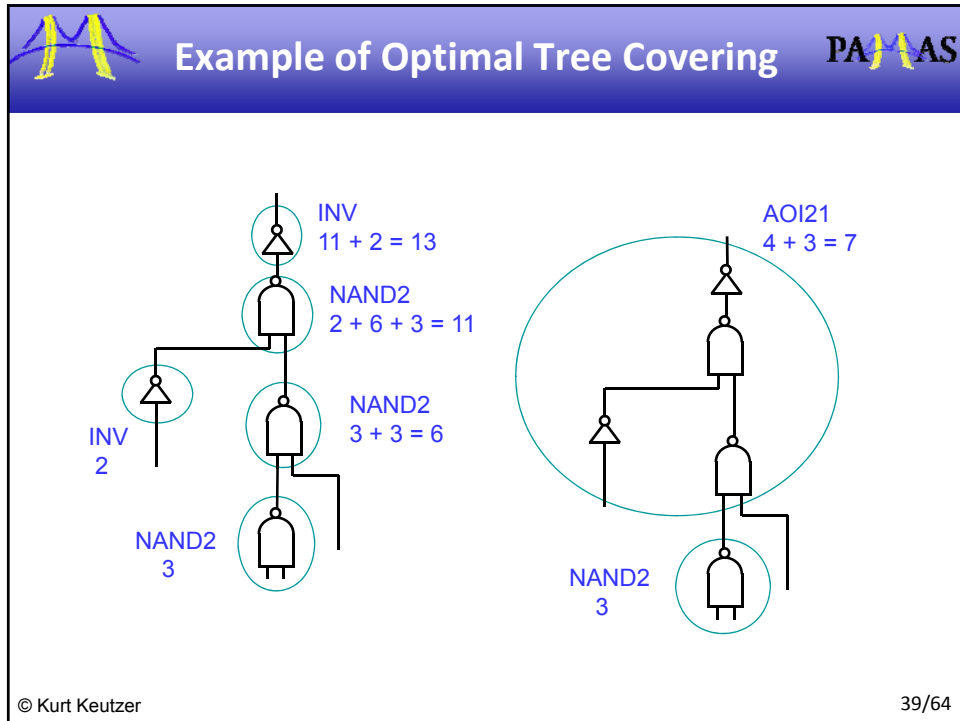
Best cover for this match uses best covers for p, z

© Kurt Keutzer 37/64

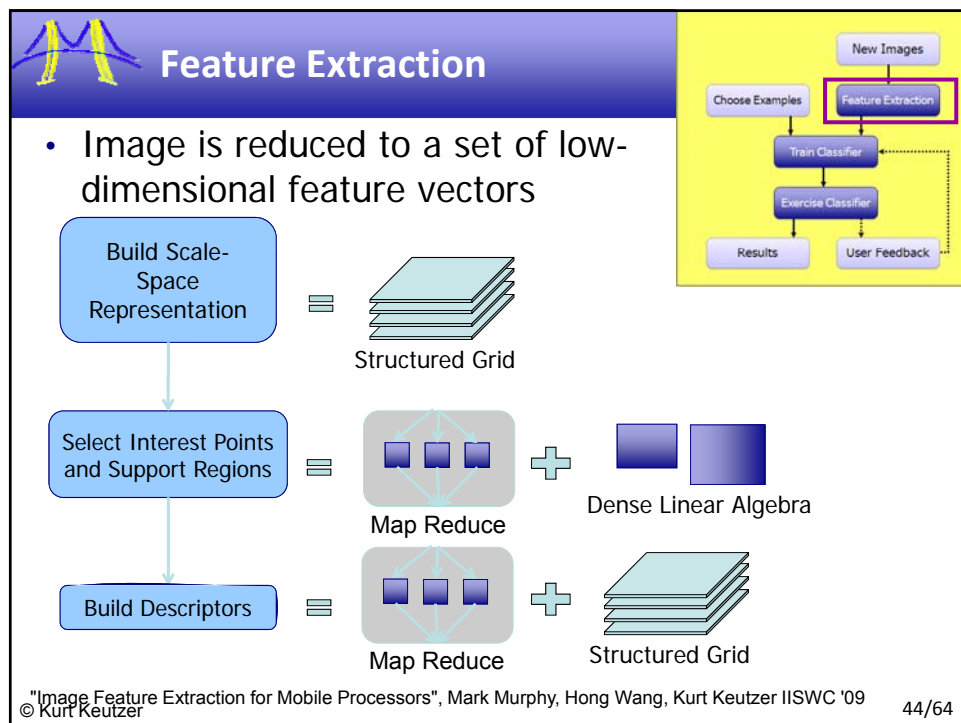
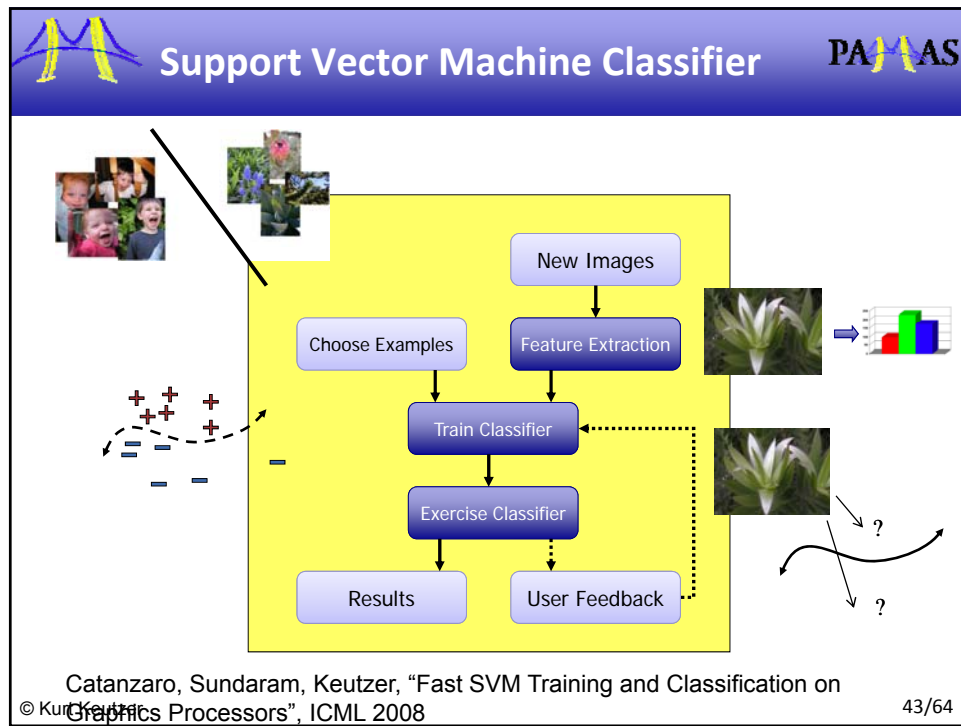
Mapping a circuits into a Cell Library PAAS

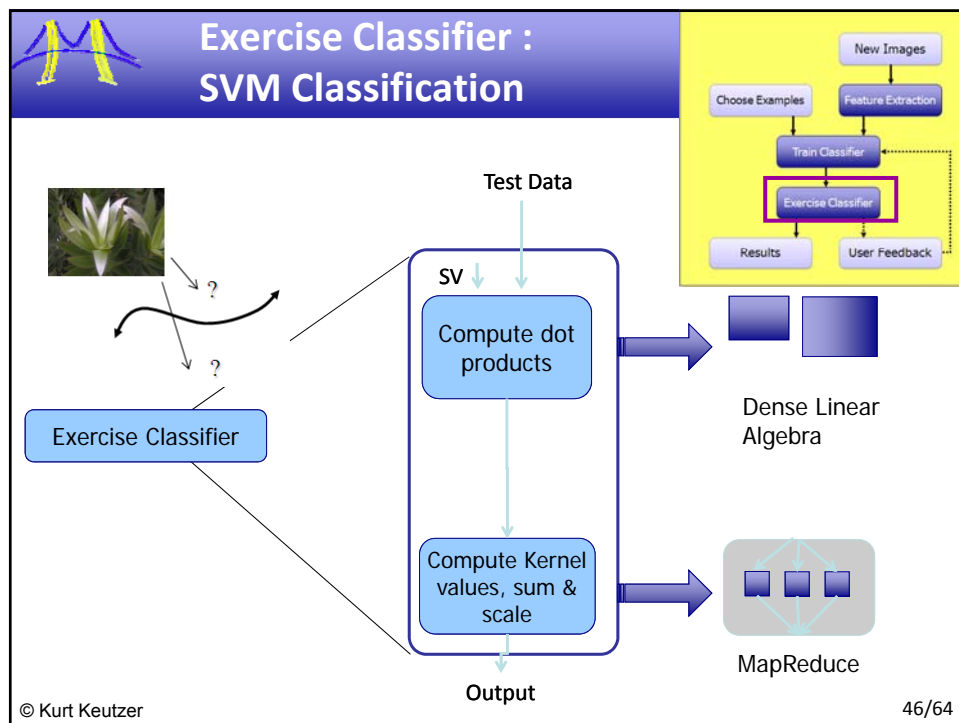
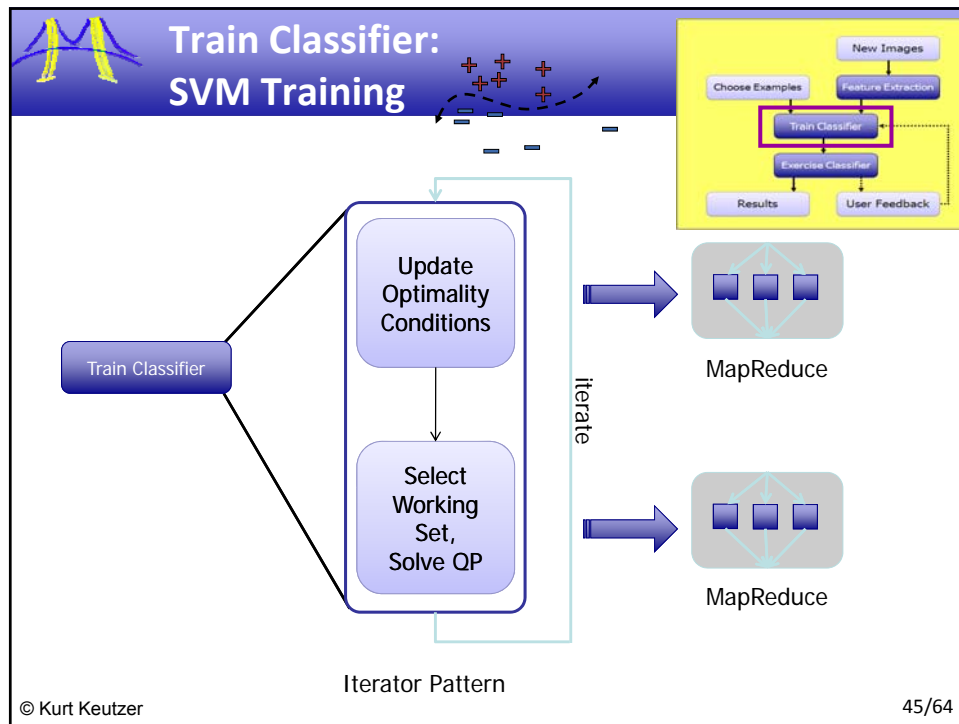
subject tree

© Kurt Keutzer 38/64




21





Support-Vector Machine Mini-Framework



```

graph TD
    NewImages[New Images] --> FeatureExtraction[Feature Extraction]
    FeatureExtraction --> TrainClassifier[Train Classifier]
    TrainClassifier --> ExerciseClassifier[Exercise Classifier]
    ExerciseClassifier --> Results[Results]
    ExerciseClassifier --> UserFeedback[User Feedback]
    UserFeedback --> TrainClassifier
    FormQuery[Form Query] --> ExerciseClassifier
  
```

- Support-Vector Machine Framework used to achieve:
 - 9-35x speedup for training
 - 81-138x for classification
 - 1100 downloads since release

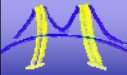
Fast support vector machine training and classification , Catanzaro, Sundaram, Keutzer, International Conference on Machine Learning 2008

© Kurt Keutzer 47/64


Outline

- Architecting Parallel Software
- Structural Patterns
- Computational Patterns
- Examples
 - CBIR
 - ➔ ■ MRI
 - Speech Recognition
- Summary

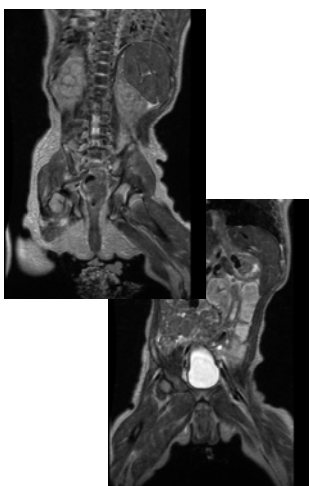
© Kurt Keutzer 48/64



Compelling Application: Fast, Robust Pediatric MRI

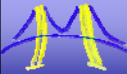


- Pediatric MRI is difficult:
 - Children cannot sit still, breathhold
 - Low tolerance for long exams
 - Anesthesia is costly and risky
- Like to accelerate MRI acquisition
 - Advanced MRI techniques exist, but require data- and compute- intense algorithms for image reconstruction
- Reconstruction must be fast, or time saved in accelerated acquisition is lost in computing reconstruction
 - Slow reconstruction times are a non-starter for clinical use




© Kurt Keutzer


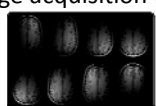
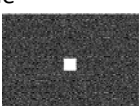
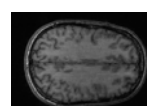
49/64



Domain Experts and State-of-the-Art Algorithms



- Collaboration with MRI Researchers:
 - Miki Lustig, Ph.D., Berkeley EECS
 - Marc Alley, Ph.D., Stanford EE
 - Shreyas Vasanawala, M.D./Ph.D., Stanford Radiology
- Advanced MRI: Parallel Imaging and Compressed Sensing to dramatically reduce MRI image acquisition time

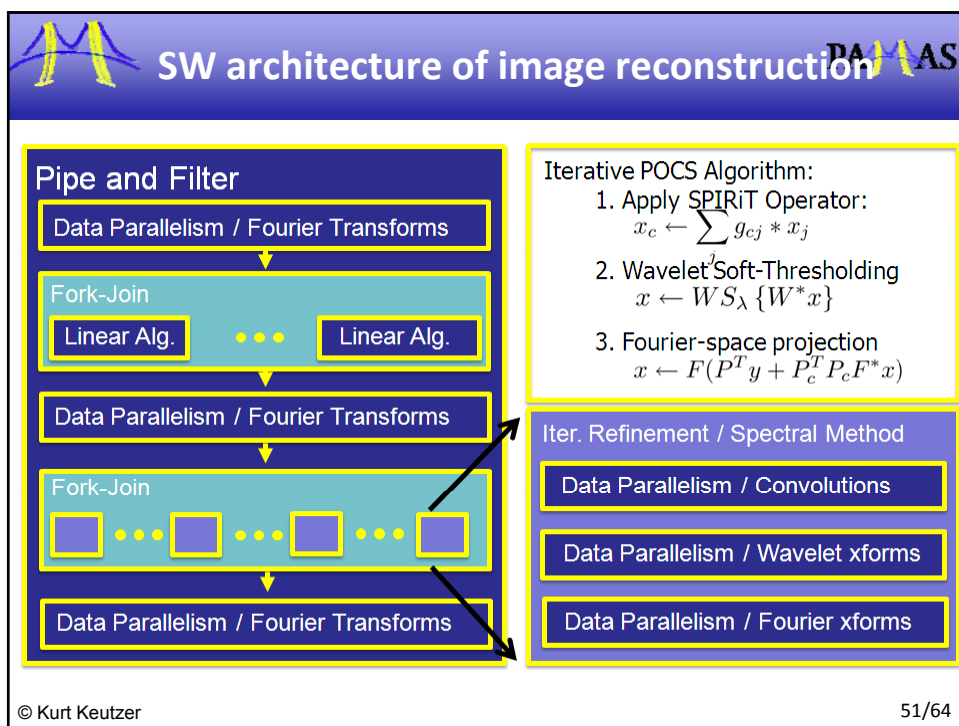
- Computational IOU: Must solve constrained L1 minimization

$$\begin{aligned}
 &\text{minimize } \|Wx\|_1 \\
 &\text{s.t. } \mathbf{F}_\Omega x = y, \\
 &\quad \|Gx - x\|_2 < \varepsilon
 \end{aligned}$$

© Kurt Keutzer

50

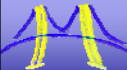
50/64




Game-Changing Speedup

- 100X faster reconstruction
- Higher-quality, faster MRI
- This image: 8 month-old patient with cancerous mass in liver
 - 256 x 84 x 154 x 8 data size
 - Serial Recon: 1 hour
 - Parallel Recon: 1 minute
- Fast enough for clinical use
 - Software currently deployed at Lucile Packard Children's Hospital for clinical study of the reconstruction technique

© Kurt Keutzer
52
52/64



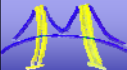
Outline




- Architecting Parallel Software
- Structural Patterns
- Computational Patterns
- Examples
 - CBIR
 - MRI
- ➔ ■ **Speech Recognition**
- Summary

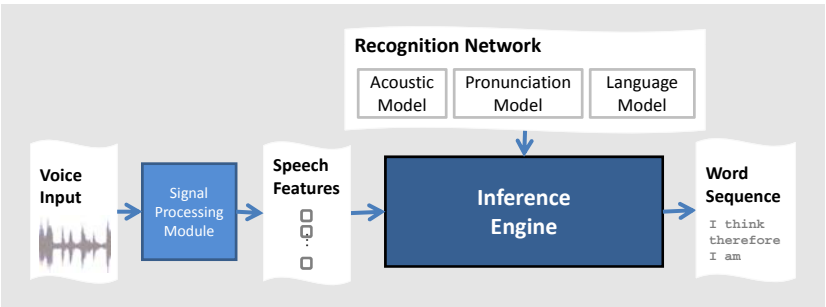
© Kurt Keutzer

53/64



Large Vocabulary Continuous Speech Recognition

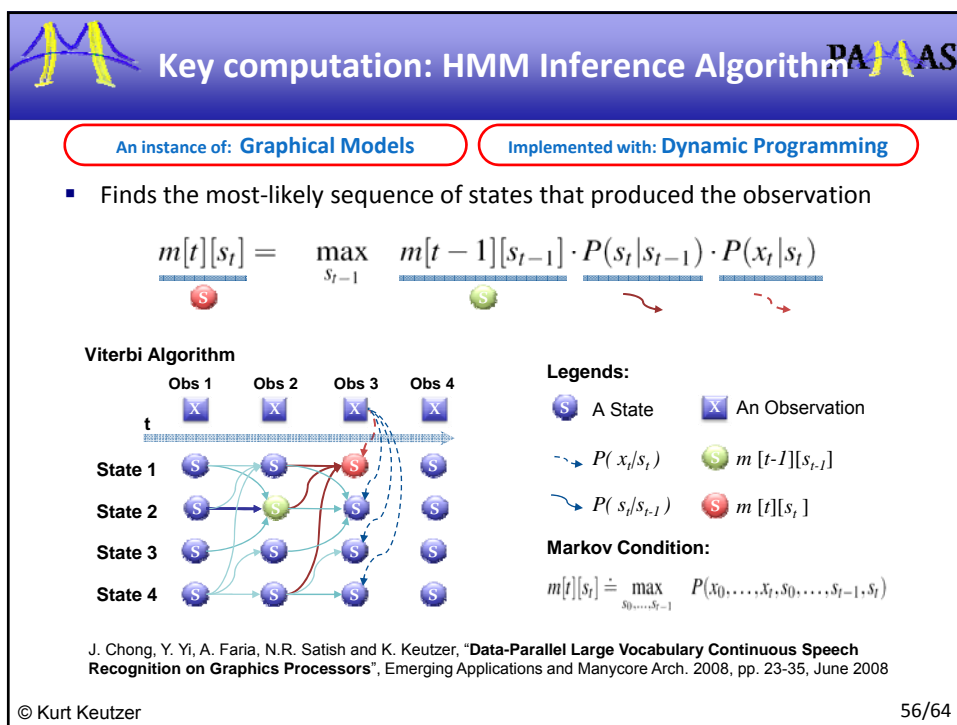
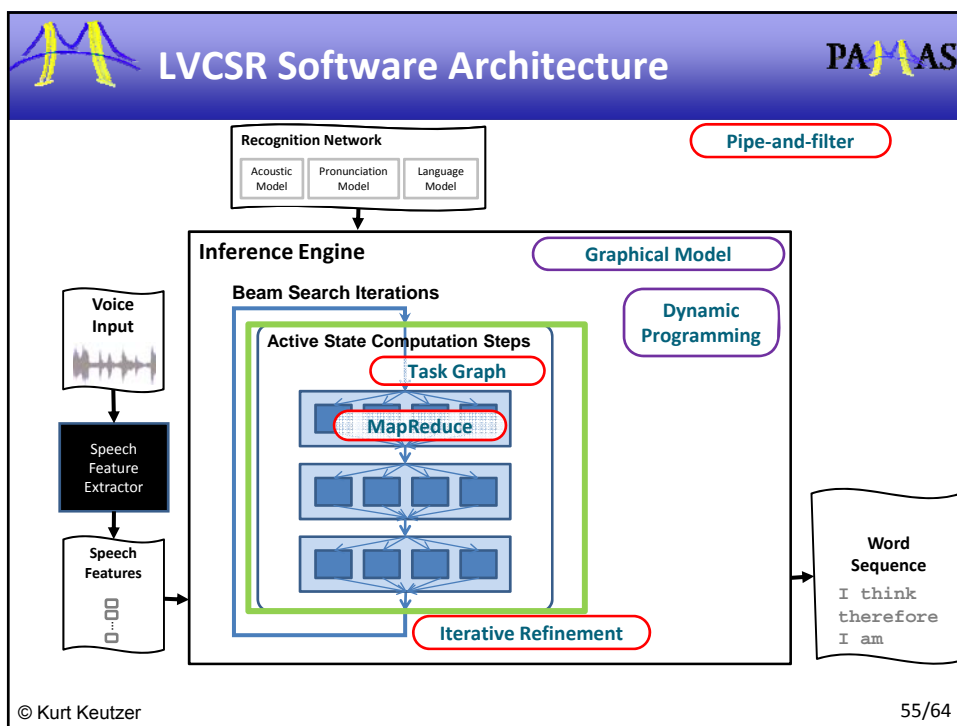



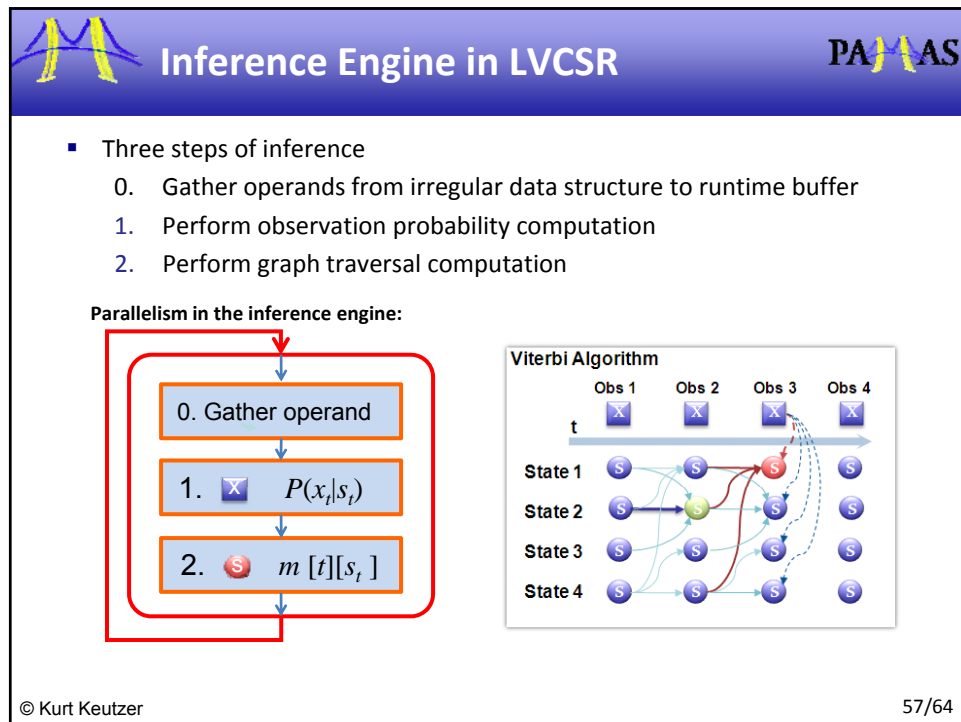


- Inference engine based system
 - Used in Sphinx (CMU, USA), HTK (Cambridge, UK), and Julius (CSRC, Japan) [10,15,9]
- Modular and flexible setup
 - Shown to be effective for Arabic, English, Japanese, and Mandarin

© Kurt Keutzer

54/64





Speech Recognition Results PAMAS

- Input: Speech audio waveform
- Output: Recognized word sequences

- Achieved 11x speedup over sequential version
- Allows 3.5x faster than real time recognition

- Our technique is being deployed in a hotline call-center data analytics company
- Used to search content, track service quality and provide early detection of service issues

Scalable HMM based Inference Engine in Large Vocabulary Continuous Speech Recognition, Kisun You, Jike Chong, Youngmin Yi, Ekaterina Gonina, Christopher Hughes, Wonyong Sung and Kurt Keutzer, IEEE Signal Processing Magazine, March 2010

© Kurt Keutzer
59/64

Multi-media Speech Recognition PAMAS

Prof. Dorothea Kolossa
Speech Application Domain Expert
 Technische Universität Berlin

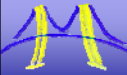
Extended *audio-only speech recognition* framework to enable *audio-visual speech recognition (lip reading)*

Achieved a **20x speedup** in application performance compared to a sequential version in C++


The application framework enabled a **Matlab/Java programmer** to **effectively utilize highly parallel platform**

Dorothea Kolossa, Jike Chong, Steffen Zeiler, Kurt Keutzer, "Efficient Manycore CHMM Speech Recognition for Audiovisual and Multistream Data", Accepted at Interspeech 2010.

© Kurt Keutzer
60/64

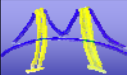


Outline




- Architecting Parallel Software
- Structural Patterns
- Computational Patterns
- Examples
- ➔ ■ Summary

© Kurt Keutzer 61/64

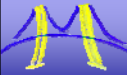


Other Interesting Results




- Patterns have helped the PALLAS research group publish papers in a diverse group of leading Computer Science conferences in the last few years:
 - Interspeech 2009, Interspeech 2010 (2)
 - IEEE Signal Processing Magazine 2009
 - European Conference on Computer Vision 2010
 - International Conference on Computer Vision 2009
 - Workshop on High Performance Computing in Finance at Super Computing 2009
 - Joint Annual Meeting of the International Society for Magnetic Resonance in Medicine, *ISMRM 2010*
 - International Conference on Machine Learning 2008
- What's the point?
 - Computational patterns give a new powerful viewpoint to efficiency programmers:
 - Enable us to disentangle the big fuzzy ball of yarn of computation
 - add 20 IQ points to our problem solving (as per Alan Kay)
 - Our Pattern language helps you to write good parallel code

© Kurt Keutzer 62/64



Summary



- The key to productive and efficient parallel programming is creating a good software architecture – a hierarchical composition of:
- Structural patterns: enforce modularity and expose invariants
 - I showed you three –seven more will be all you need
- Computational patterns: identify key computations to be parallelized
 - I showed you three –ten more will be all you need
- Orchestration of computational and structural patterns creates architectures which greatly facilitates the development of parallel programs:
 - I showed you three – there are *many* more

Patterns: <http://parlab.eecs.berkeley.edu/wiki/patterns/patterns>

PALLAS: <http://parlab.eecs.berkeley.edu/research/pallas>

CS194: Engineering Parallel Software: Fall 2010

© Kurt Keutzer 63/64