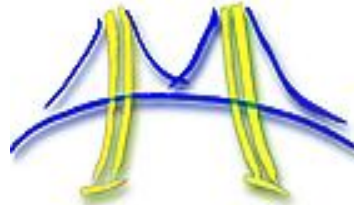


PARLab Parallel Boot Camp

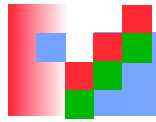


Architecting Parallel Software with Patterns

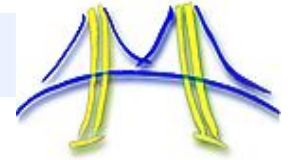
Kurt Keutzer

Electrical Engineering and Computer Sciences
University of California, Berkeley

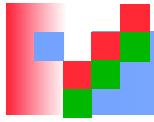




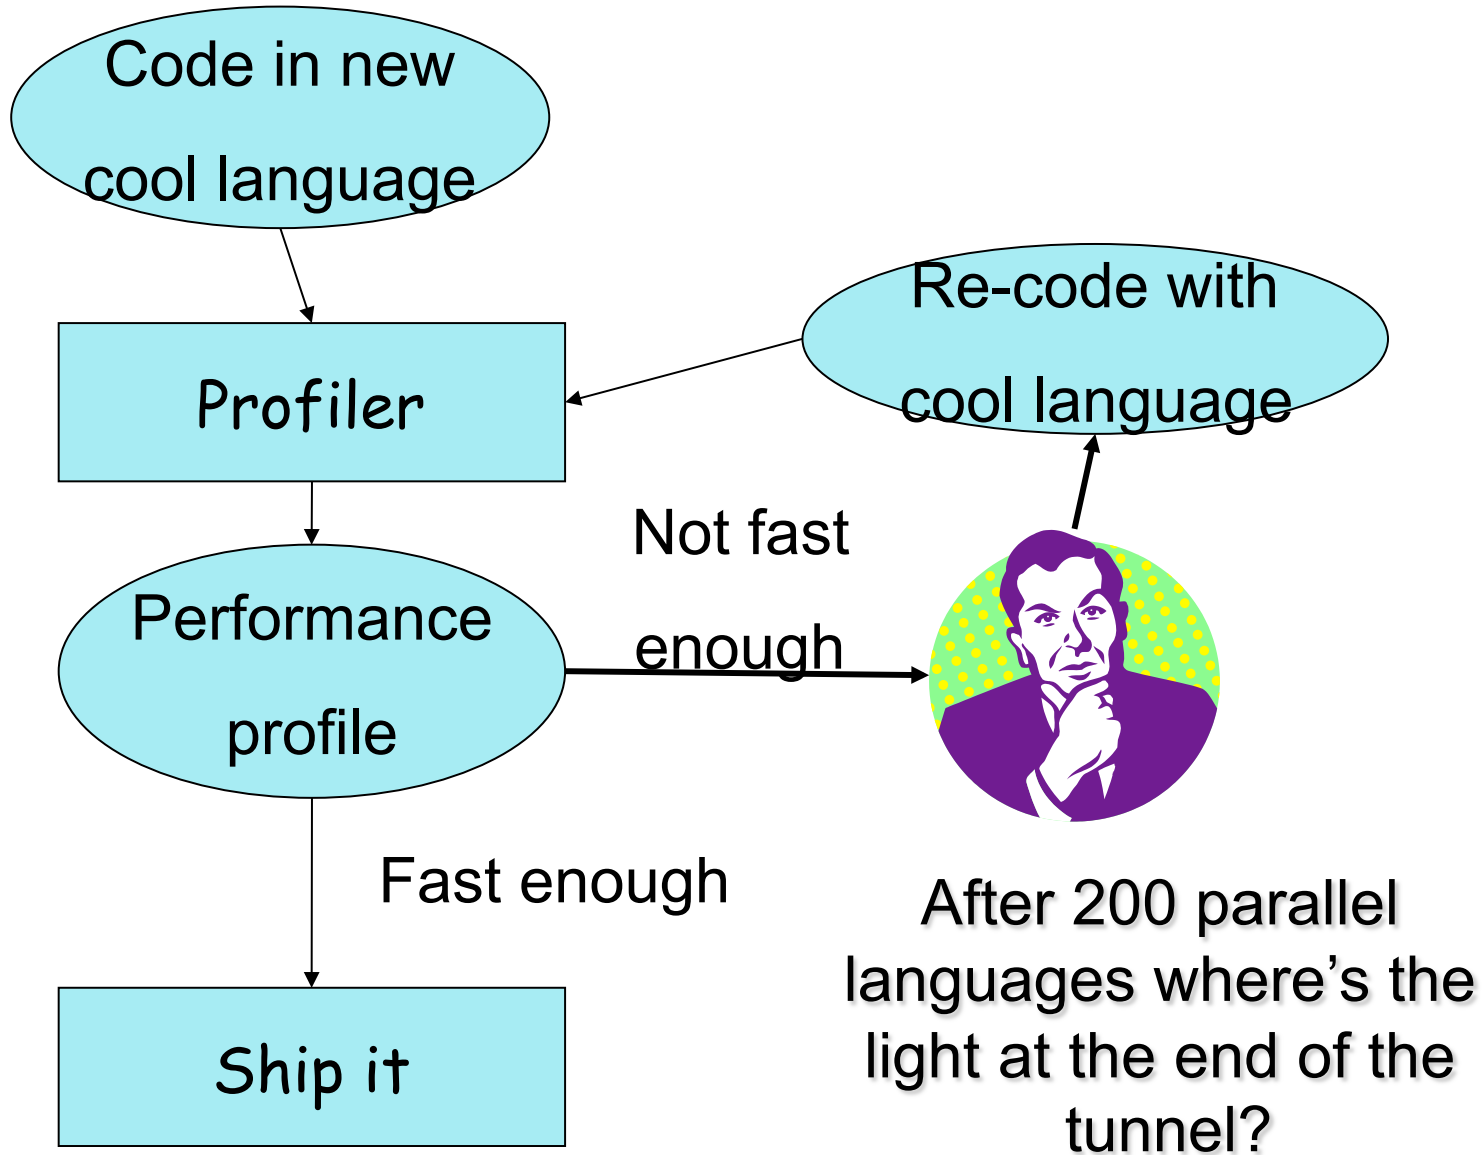
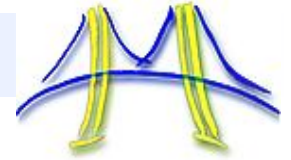
Steiner Tree Construction Time By Routing Each Net in Parallel

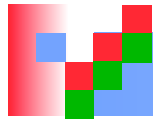


Benchmark	Serial	2 Threads	3 Threads	4 Threads	5 Threads	6 Threads
adaptec1	1.68	1.68	1.70	1.69	1.69	1.69
newblue1	1.80	1.80	1.81	1.81	1.81	1.82
newblue2	2.60	2.60	2.62	2.62	2.62	2.61
adaptec2	1.87	1.86	1.87	1.88	1.88	1.88
adaptec3	3.32	3.33	3.34	3.34	3.34	3.34
adaptec4	3.20	3.20	3.21	3.21	3.21	3.21
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

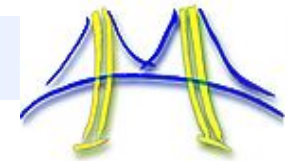


Assumption #2: This won't help either

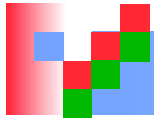




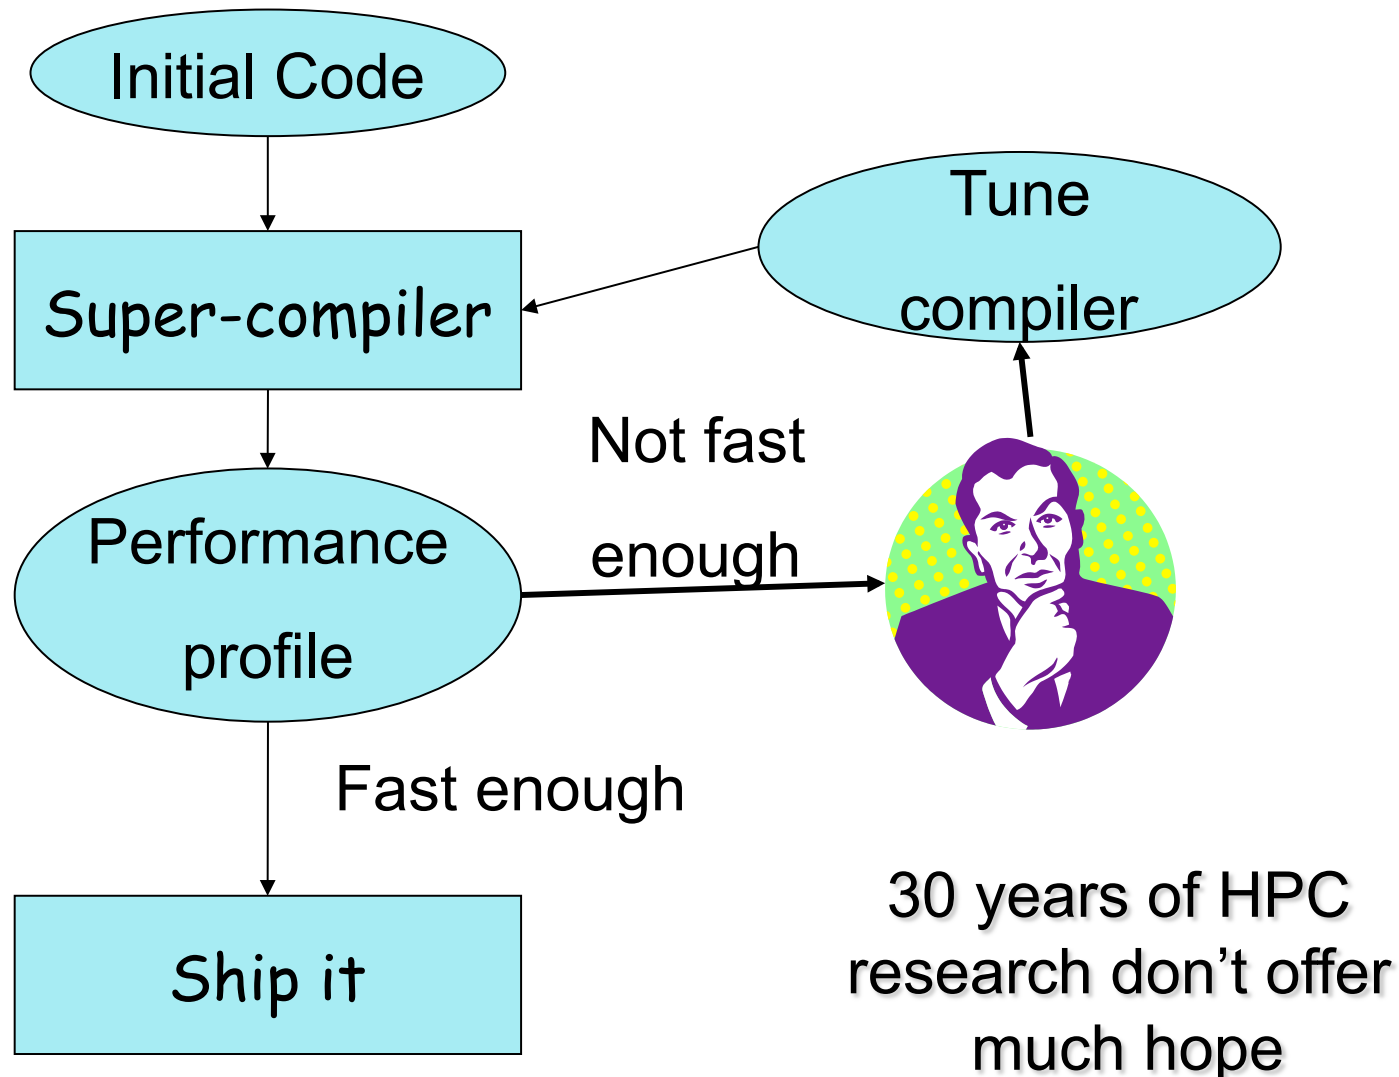
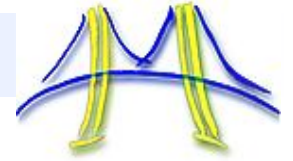
Parallel Programming environments in the 90's

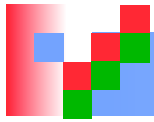


ABCPL	CORRELATE	GLU	Monter	Parafix	pC++
ACE	CPS	GUARD	Legion	Parallel	SCHEDULE
ACT++	CEL	HAIL	Meta-Chance	Parallel-C++	SetTL
Active messages	CSP	Masked	Midway	ParLan	POET
Adi	Others	HPC++	Midpipe	ParC	SODA
Adisynth	CUMULUS	JAVAR	OpasPar	ParLib++	SHMEM
ADDAP	DAGGER	HORUS	Mirage	ParLin	SDMPLE
AFAP	DAPPLE	HPC	MpC	Parman	Sua
ALTRAN	DataParallel C	IMPACT	MOSIX	Paru	SISAL
AM	DC++	ISIS	Module-P	pC	distributed smalltalk
AMDC	DCE++	JAVAR	Module-2*	pC++	SOIL
AppLeS	DDD	JADE	MultiPol	PCP	SOONC
Arise	DICE	JenaRMII	MPI	PCP	Split-C
ARTS	DIPC	javaPO	MPC++	PE	SR
Asynchronous-Or	DOLIB	JenSpace	Movins	PEACE	Threads
Aurora	DOME	JDC	Nano-Threads	PCU	Strand
Automap	DOSMOS	Joyce	NEEL	PET	STUP
bb_threads	DRL	Khoros	NetClass++	PETSc	Synapse
Baze	DSS-Threads	Karma	Nexus	PENNY	Telegraph
BSP	Ease	KOAN Fortran-8	Nimrod	Phosphorus	SuperParallel
BlockComm	ECO	LAM	NOV	POET	TOGMSG
C*	EdHel	Lilac	Objective Linda	Polaris	Threads++
C* in C	Eileen	Linda	Ocean	POOMA	ThreadMask
C++	Emerald	IADA	Omega	POOL-T	TRAPPER
CarLOS	EPL	WWinda	OyeshOP	PRESTO	uC++
Castmore	Excelsior	IBETL-Linda	Orca	P-SDO	UNITY
C4	Express	ParLin	OOF90	Prospero	UC
CC++	Falcon	Eileen	P++	Proteus	V
Cho	Filaments	P4-Linda	PIL	QPC++	VIC*
Chadone	FM	Glenda	p4-Linda	PVM	Winfield V-MUS
Charm	FLASH	POSTBL	Pablo	PSI	VPE
Charm++	The FORCE	Objective-Linda	PADE	PSDM	Win32 threads
Cid	Fork	LPS	PADRE	Quake	WinPar
CG	Fortran-M	Locust	Panda	Quick	WWinda
CM-Fortran	FX	Lparc	Papers	Quick Threads	XENGOOPS
Cosmos	GA	Lucid	AFAPL	Sage++	XPC
Code	GAMMA	Maisie	Par++	SCANDAL	Zounds
COOL	Glenda	Manifold	Paradigm	SAM	ZPL

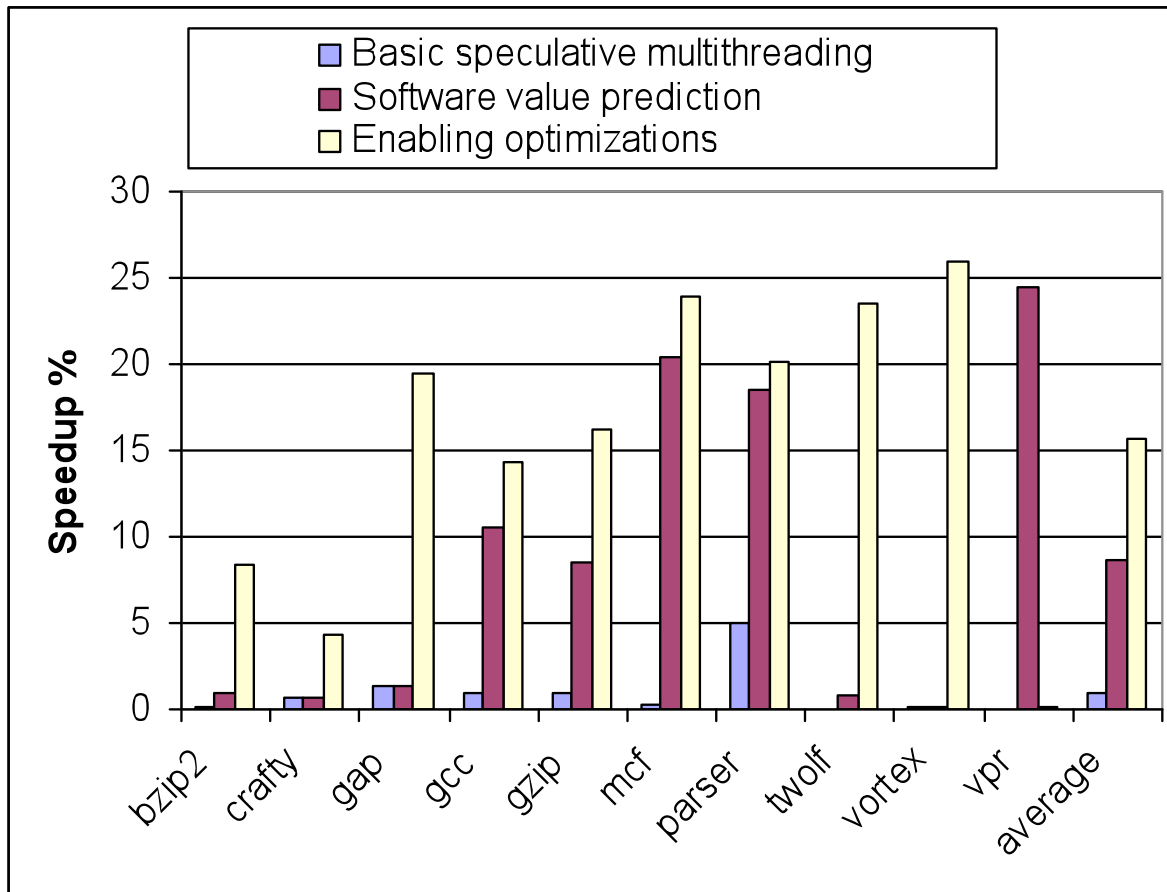
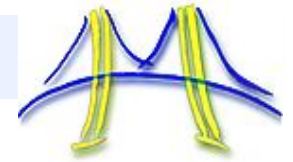


Assumption #3: Nor this



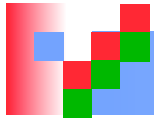


Automatic parallelization?

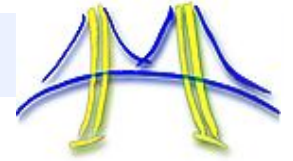


- Aggressive techniques such as speculative multithreading help, but they are not enough.
- Ave SPECint speedup of **8%** ... will climb to ave. of **15%** once their system is fully enabled.
- There are no indications auto par. will radically improve any time soon.
- Hence, I do not believe Auto-par will solve our problems.

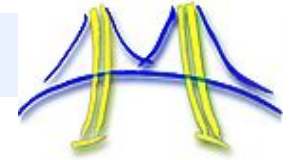
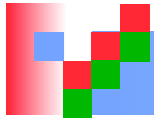
Results for a simulated dual core platform configured as a main core and a core for speculative execution.



Outline

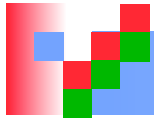


- ⇒ ■ Intro to Kurt
 - General approach to applying the pattern language
 - Detail on Structural Patterns
 - High-level examples of composing patterns

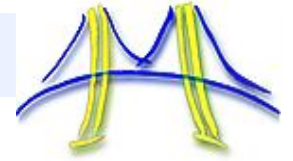


Key Elements of Kurt's SW Education

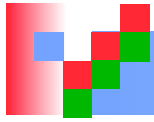
- AT&T Bell Laboratories: CAD researcher and programmer
 - Algorithms: D. Johnson, R. Tarjan
 - Programming Pearls: S C Johnson, K. Thompson, (Jon Bentley)
 - Developed useful software tools:
 - » Plaid: programmable logic aid: used for developing 100's of FPGA-based HW systems
 - » CONES/DAGON: used for designing >30 application-specific integrated circuits
- Synopsys: researcher → CTO (25 products, ~15 million lines of code, \$750M annual revenue, top 20 SW companies)
 - Super programming: J-C Madre, Richard Rudell, Steve Tjiang
 - Software architecture: Randy Allen, Albert Wang
 - High-level Invariants: Randy Allen, Albert Wang
- Berkeley: teaching software engineering and Par Lab
 - Took the time to reflect on what I had learned:
 - Architectural styles: Garlan and Shaw
 - » Design patterns: Gamma et al (aka Gang of Four), Mattson's PLPP
 - » A Pattern Language: Alexander, Mattson
 - » Dwarfs: Par Lab Team



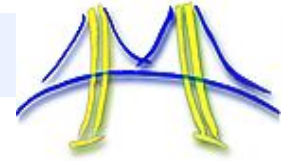
What I learned (the hard way)



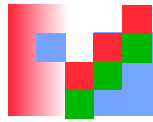
- Software must be **architected** to achieve productivity, efficiency, and correctness
- SW architecture >> programming environments
 - >> programming languages
 - >> compilers and debuggers
 - (>>hardware architecture)
- Discussions with superprogrammers taught me:
 - Give me the right program structure/architecture I can use any programming language
 - Give me the wrong architecture and I'll never get there
- What I've learned when I had to teach this stuff at Berkeley:
 - Key to **architecture** (software or otherwise) is **design patterns** and a **pattern language**
- Resulting software design then uses a hierarchy of software frameworks for implementation
 - Application frameworks for application (e.g. CAD) developers
 - Programming frameworks for those who build the application frameworks



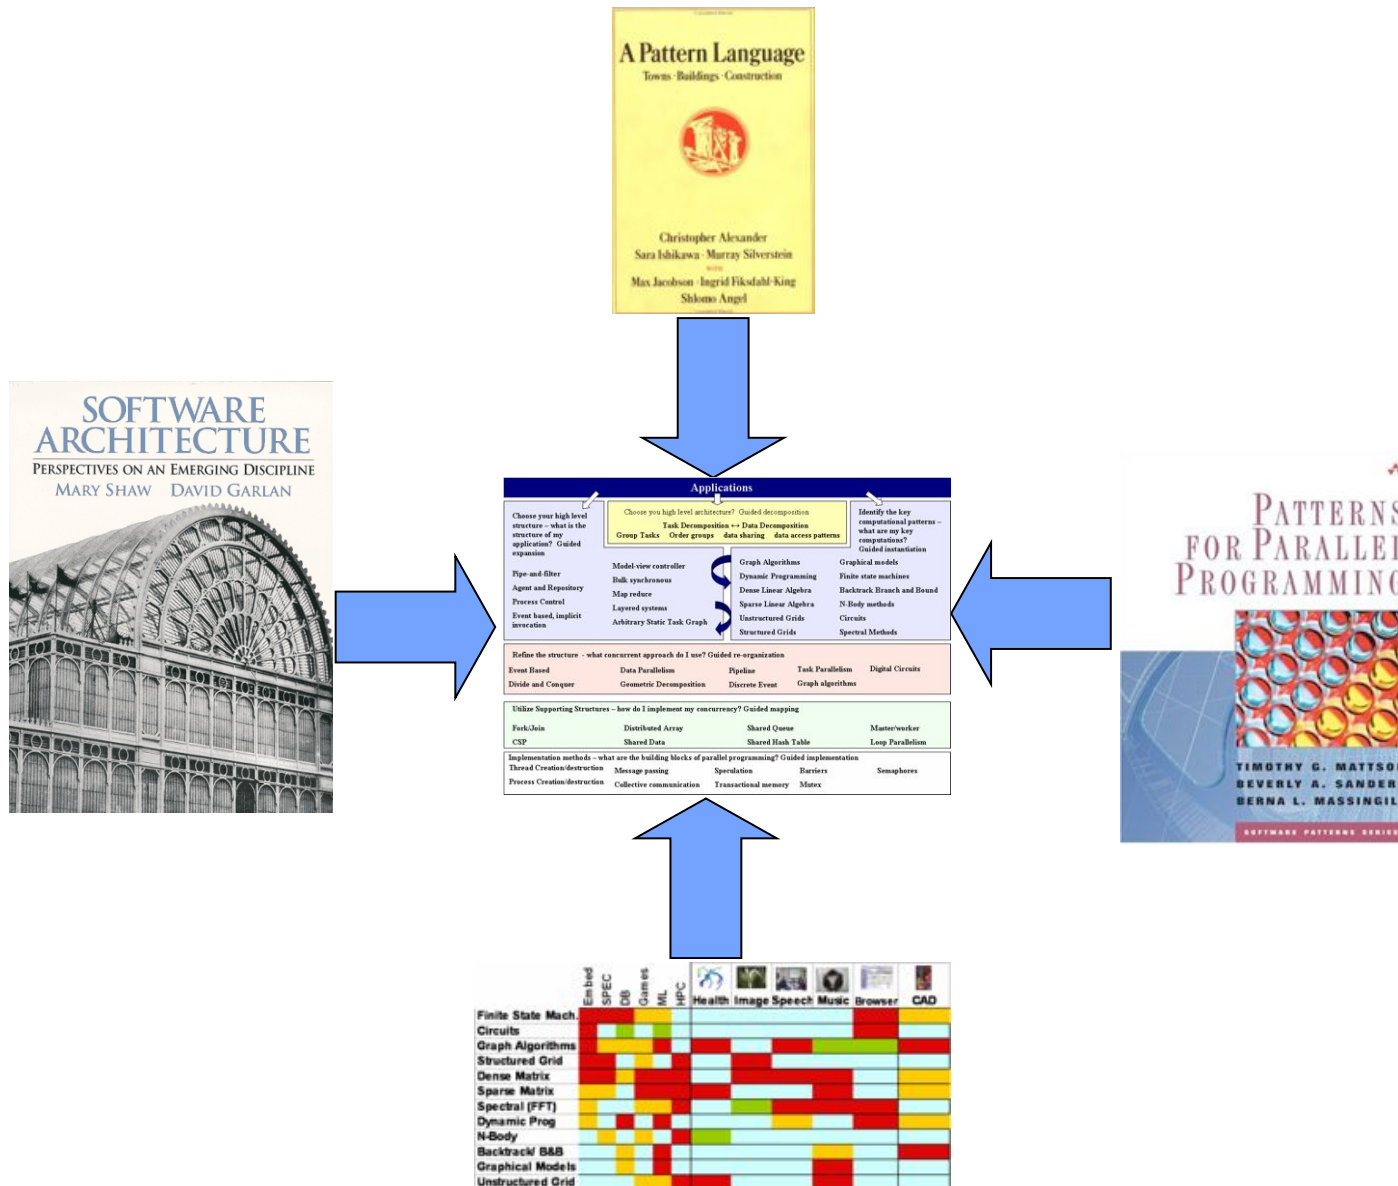
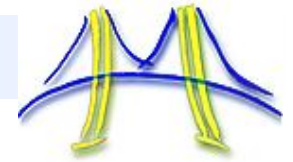
Outline

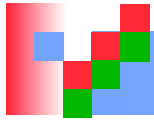


- Intro to Kurt
- ⇒ ■ General approach to applying the pattern language
- Detail on Structural Patterns
- High-level examples of composing patterns

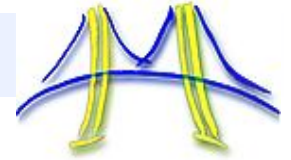


Elements of a pattern language

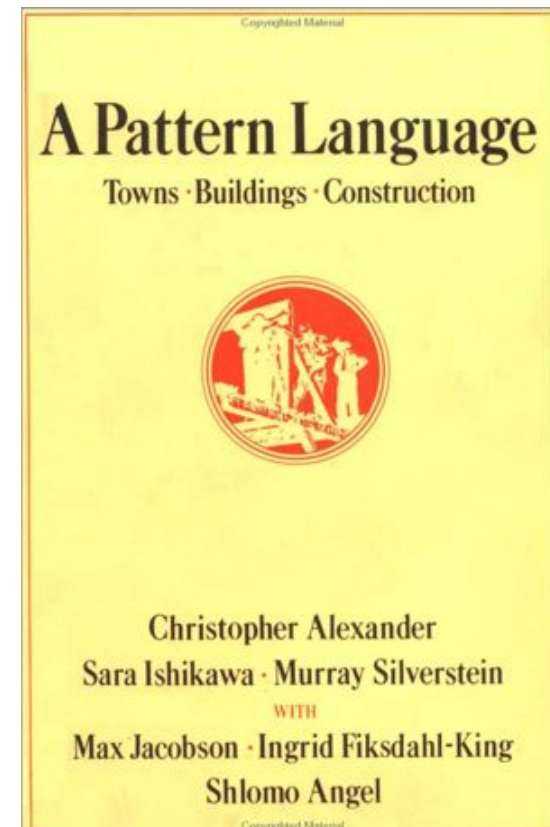


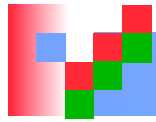


Alexander's Pattern Language

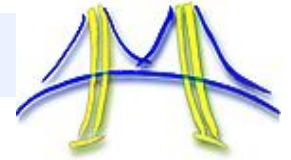


- Christopher Alexander's approach to (civil) architecture:
 - "Each **pattern** describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice." *Page x, A Pattern Language, Christopher Alexander*
- Alexander's 253 (civil) architectural **patterns** range from the creation of cities (2. distribution of towns) to particular building problems (232. roof cap)
- A **pattern language** is an organized way of tackling an architectural problem using patterns
- Main limitation:
 - It's about civil not software architecture!!!

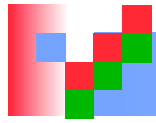




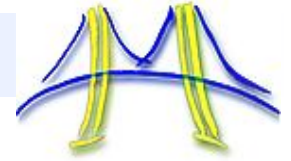
Alexander's Pattern Language (95-103)



- Layout the overall arrangement of a group of buildings: the height and number of these buildings, the entrances to the site, main parking areas, and lines of movement through the complex.
- 95. Building Complex
- 96. Number of Stories
- 97. Shielded Parking
- 98. Circulation Realms
- 99. Main Building
- 100. Pedestrian Street
- 101. Building Thoroughfare
- 102. Family of Entrances
- 103. Small Parking Lots



Family of Entrances (102)



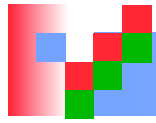
- May be part of **Circulation Realms (98)**.
- Conflict:
- When a person arrives in a complex of offices or services or workshops, or in a group of related houses, there is a good chance he will experience confusion unless the whole collection is laid out before him, so that he can see the entrance place where he is going.

Resolution:

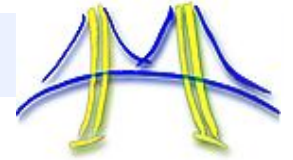
Lay out the entrances to form a family. This means:

- 1) They form a group, are visible together, and each is visible from all the others.
- 2) They are all broadly similar, for instance all porches, or all gates in a wall, or all marked by a similar kind of doorway.
- May contain **Main Entrance (110)**, **Entrance Transition (112)**, **Entrance Room (130)**, **Reception Welcomes You (149)**.

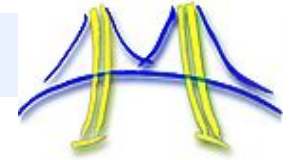
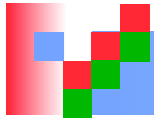




Family of Entrances

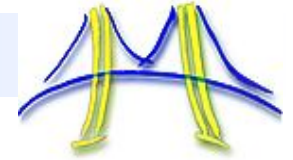
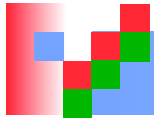


<http://www.intbau.org/Images/Steele/Badran5a.jpg>



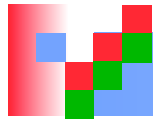
Elements of a Pattern - 1

- Name
 - It must have a meaningful name useful to remember the pattern and when it is used.
- Problem
 - A statement of the problem ... a one-line preamble and the problem stated as a question.
- Context
 - The conditions under which the problem occurs. Defines when the pattern is applicable and the configuration of the system before the pattern is applied.
- Forces
 - A description of the relevant *forces* and constraints and how they interact/conflict with one another and with goals we wish to achieve. Defines the tension that characterizes a problem.



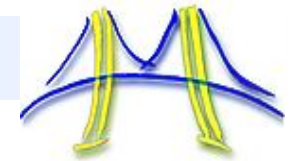
Elements of a Pattern - 2

- **Solution**
 - Instructions used to solve the problem. When done right, it resolves the tension defined in the forces section; flowing from the context and forces. We also define the new context for the system following application of the pattern.
- **Invariant**
 - What must be invariantly true for this pattern to work. May be stated in the form of Precondition, Invariant, Post-condition
- **Examples**
 - Examples to help the reader understand the pattern.
- **Known Uses and frameworks**
 - Cases where the pattern was used; preferably with literature references.
- **Related Patterns**
 - How does this pattern fit-in or work-with the other patterns in the pattern language.

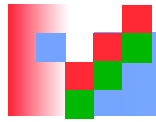


Computational Patterns

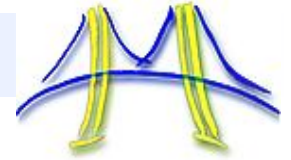
(Red Hot → Blue Cool)



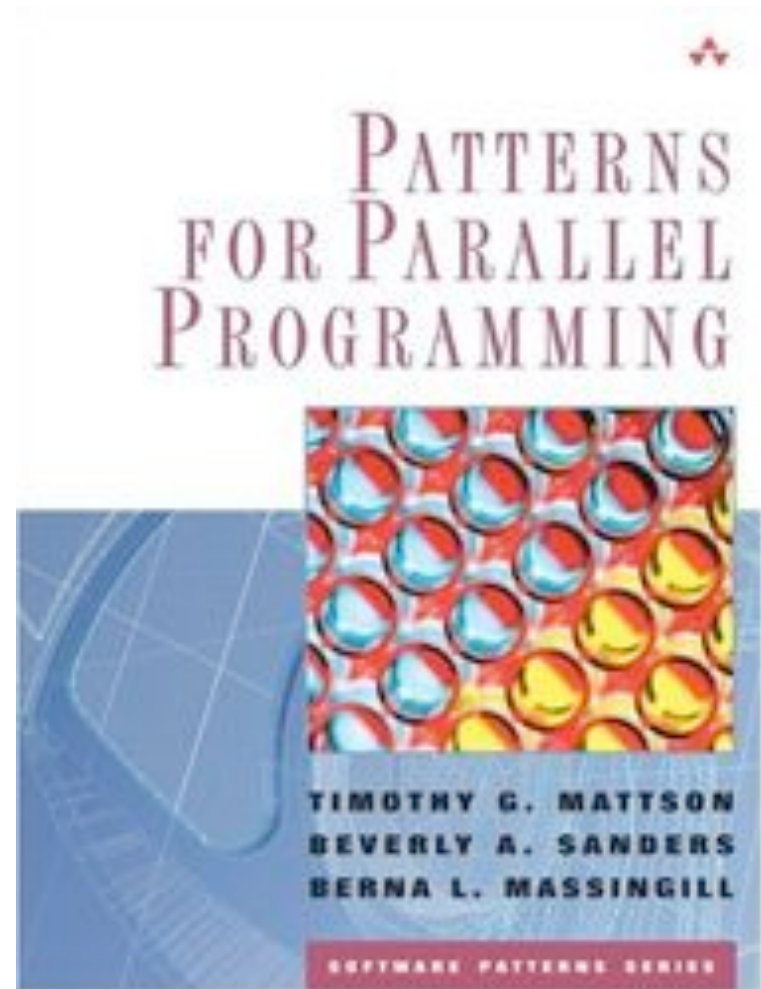
	Embed	SPEC	DB	Games	ML	CAD	HPC	Health	Image	Speech	Music	Browser
1 Finite State Mach.	Red	Red	Red	Yellow	Yellow	Yellow	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Red
2 Circuits	Red	Light Blue	Green	Light Blue	Green	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Red
3 Graph Algorithms	Red	Yellow	Yellow	Yellow	Red	Red	Light Blue	Red	Light Blue	Red	Green	Green
4 Structured Grid	Red	Red	Light Blue	Yellow	Light Blue	Light Blue	Red	Light Blue	Red	Light Blue	Light Blue	Light Blue
5 Dense Matrix	Red	Red	Yellow	Red	Red	Red	Red	Light Blue	Red	Red	Red	Light Blue
6 Sparse Matrix	Yellow	Yellow	Light Blue	Red	Red	Red	Red	Red	Light Blue	Light Blue	Red	Light Blue
7 Spectral (FFT)	Yellow	Light Blue	Light Blue	Yellow	Yellow	Yellow	Red	Light Blue	Green	Red	Red	Red
8 Dynamic Prog	Yellow	Light Blue	Red	Light Blue	Red	Red	Light Blue	Light Blue	Light Blue	Yellow	Light Blue	Red
9 Particle Methods	Light Blue	Yellow	Light Blue	Yellow	Light Blue	Light Blue	Red	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue
10 Backtrack/ B&B	Light Blue	Light Blue	Yellow	Light Blue	Red	Red	Light Blue	Light Blue	Light Blue	Light Blue	Yellow	Light Blue
11 Graphical Models	Light Blue	Light Blue	Yellow	Light Blue	Red	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Red	Light Blue
12 Unstructured Grid	Light Blue	Light Blue	Light Blue	Yellow	Yellow	Yellow	Red	Red	Light Blue	Light Blue	Red	Light Blue

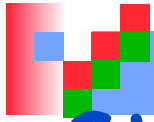


Patterns for Parallel Programming

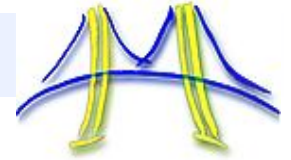


- PLPP is the first attempt to develop a complete *pattern language* for parallel software development.
 - PLPP is a great model for a pattern language for parallel software
 - PLPP mined scientific applications that utilize a monolithic application style
 - PLPP doesn't help us much with horizontal composition
-
- Much more useful to us than: *Design Patterns: Elements of Reusable Object-Oriented Software*, Gamma, Helm, Johnson & Vlissides, Addison-Wesley, 1995.

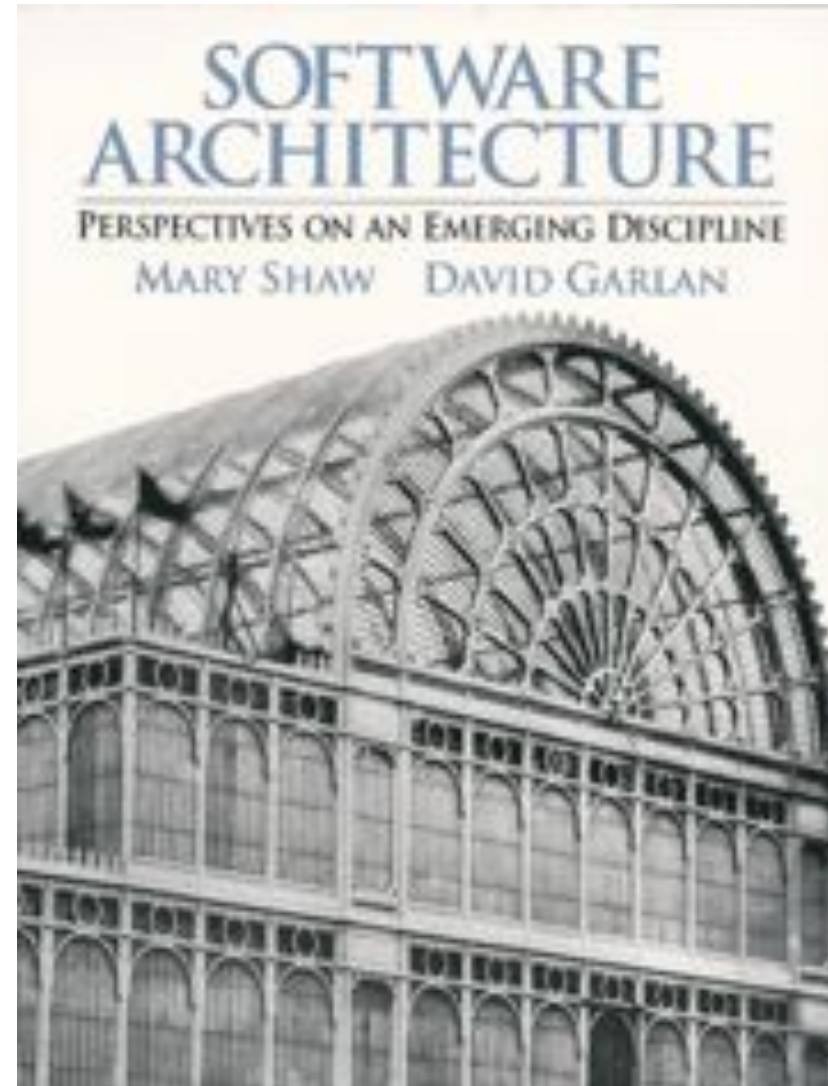


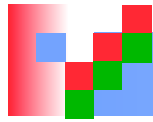


Structural programming patterns

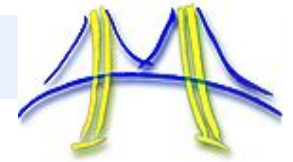


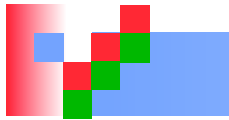
- In order to create more complex software it is necessary to compose programming patterns
- For this purpose, it has been useful to induct a set of patterns known as "architectural styles"
- Examples:
 - pipe and filter
 - event based/event driven
 - layered
 - Agent and repository/blackboard
 - process control
 - Model-view-controller



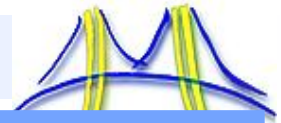


Put it all together

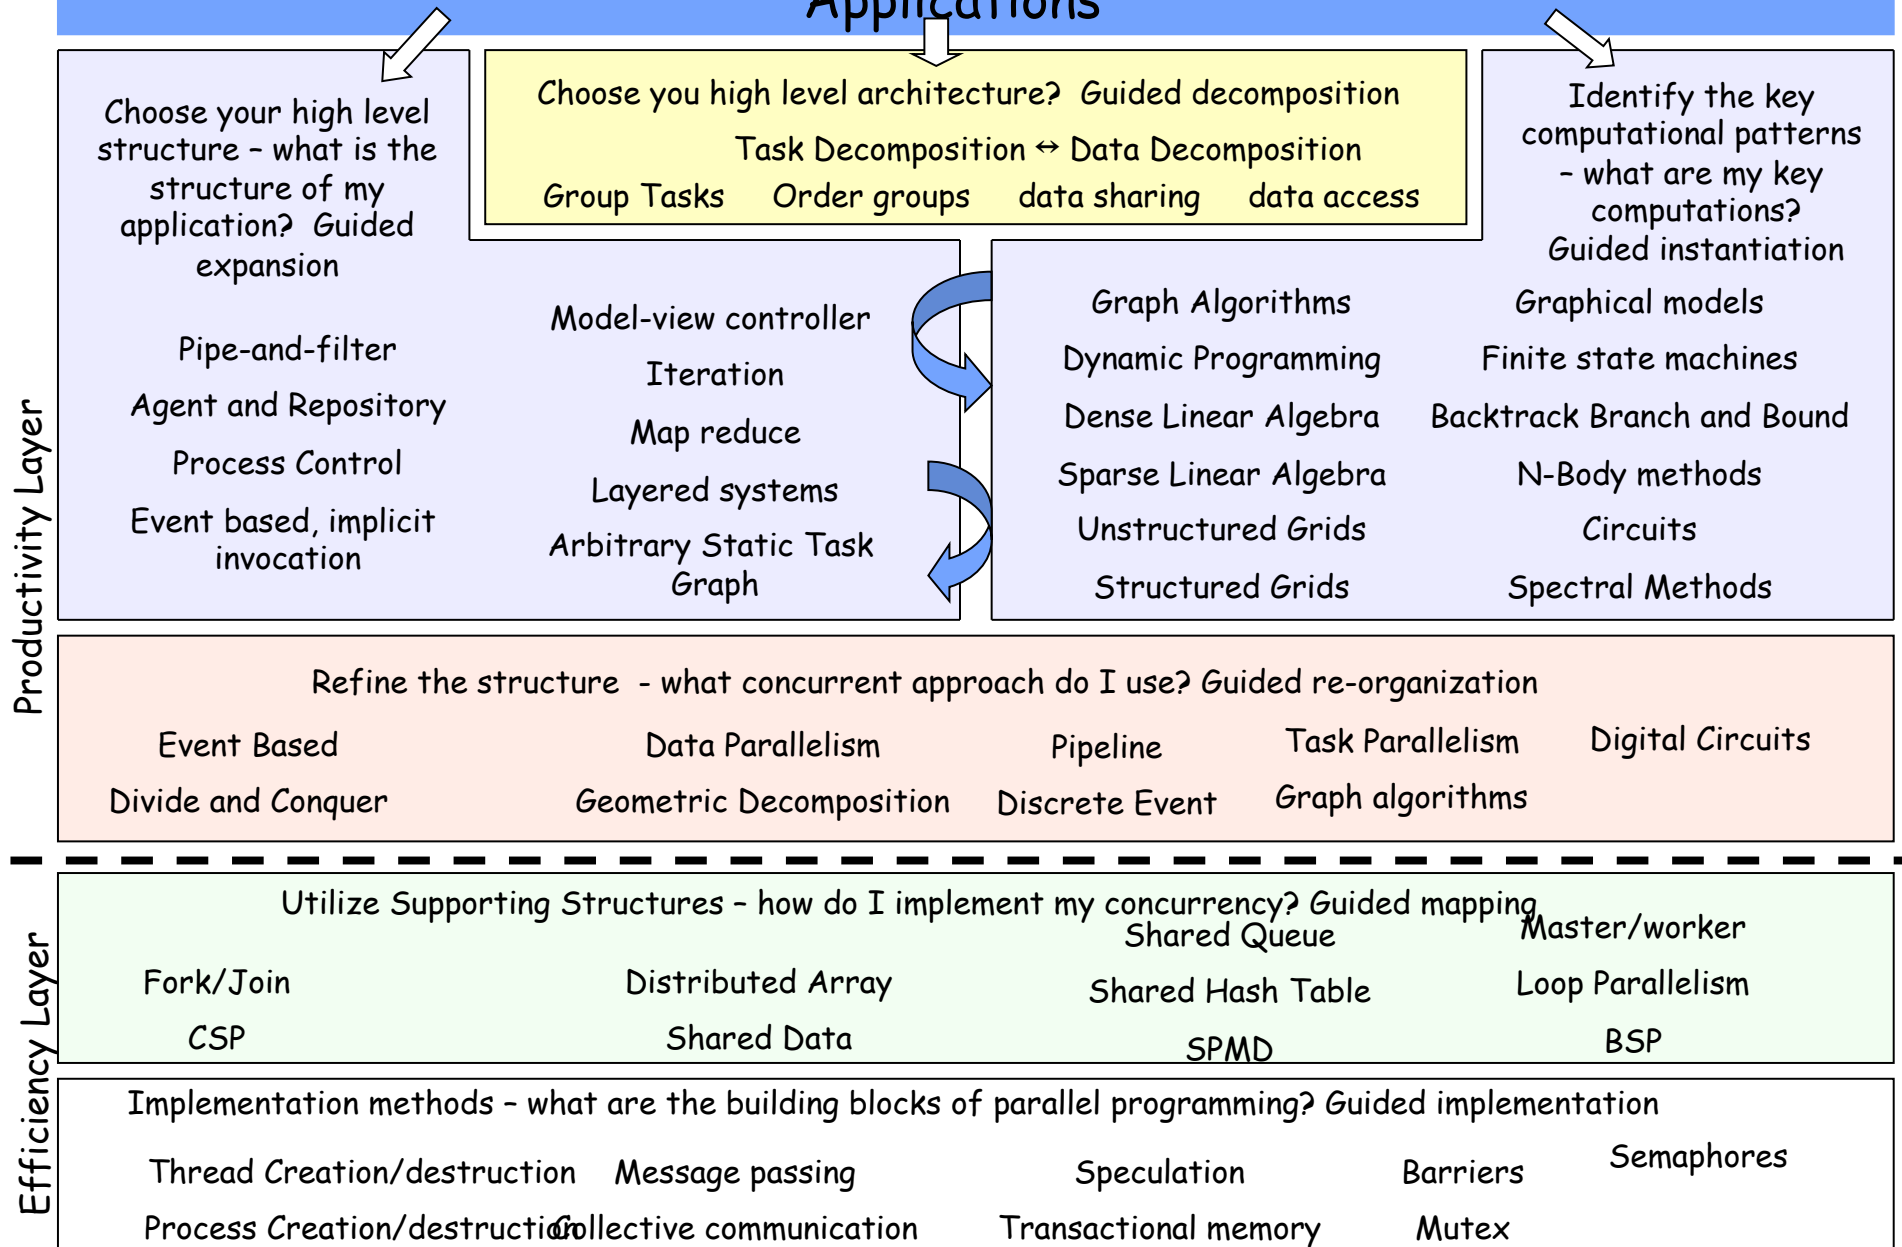


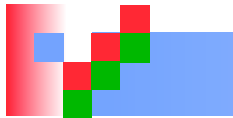


Our Pattern Language 2.0: Keutzer and Mattson



Applications





Our Pattern Language 2.0



Applications

Choose your high level structure - what is the structure of my application? Guided

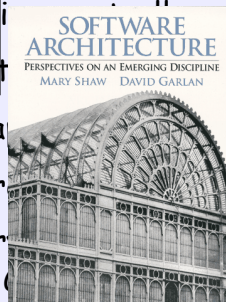
Choose you high level architecture? Guided
decomposition
Task Decomposition → Data Decomposition
Group Tasks Order groups data sharing data access

Identify the key computational patterns - what are my key computations?

Garlan and Shaw

Architectural Styles

Agent and Repository
Process Control
Event based, implicit invocation



Graph Algorithms

Dynamic Programming



Berkeley View

13 dwarfs

Track Branch and Bound

N-Body methods

Circuits

Spectral Methods

Productivity Layer

Refine the structure - what concurrent approach do I use? Guided re-organization

Event Based
Divide and Conquer

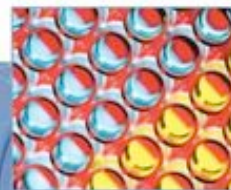
Data P
Geometric

he
Event

Task Parallelism
Graph algorithms

Digital Circuits

PATTERNS
FOR PARALLEL
PROGRAMMING



Utilize Supporting Structures

Fork/Join
CSP

Distrib
Shar

Concurrency? Guided mapping
Shared Queue

Master/worker

Shared Hash Table

Loop Parallelism

SPMD

BSP

Implementation methods - what are the b

Thread Creation/destruction Message

Process Creation/destruction Collective communication

programming? Guided implementation

ulation

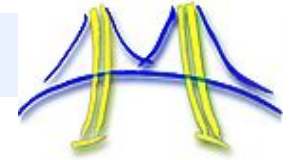
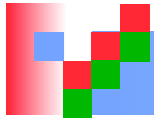
Barriers

Semaphores

Transactional memory

Mutex

Efficiency Layer



Architecting Parallel Software

Decompose Tasks/Data

Order tasks

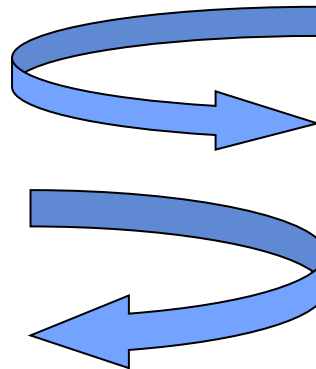
Identify Data Sharing and Access

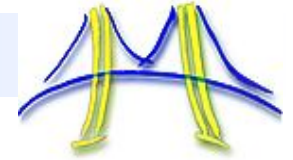
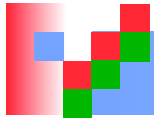
Identify the Software Structure

- Pipe-and-Filter
- Agent-and-Repository
 - Event-based
- Bulk Synchronous
 - MapReduce
- Layered Systems
- Arbitrary Task Graphs

Identify the Key Computations

- Graph Algorithms
- Dynamic programming
- Dense/Sparse Linear Algebra
 - (Un)Structured Grids
 - Graphical Models
- Finite State Machines
- Backtrack Branch-and-Bound
 - N-Body Methods
 - Circuits
- Spectral Methods



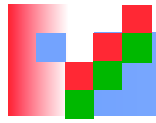


Pop Quiz: Software is More Like ...

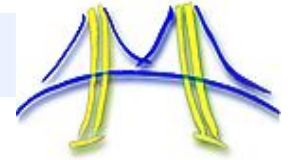
a) A building

b) A factory



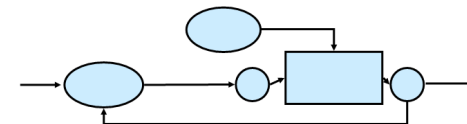
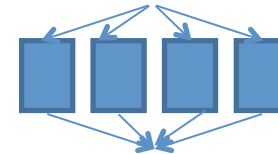
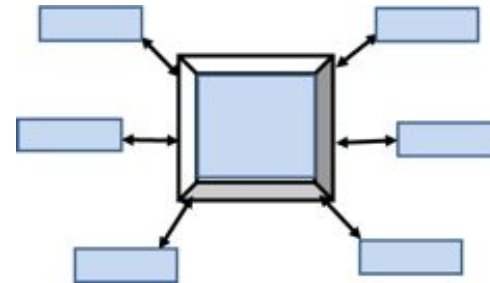
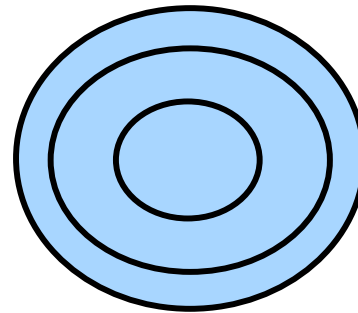
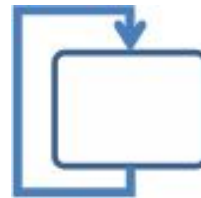
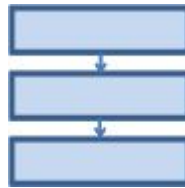


Identify the SW Structure

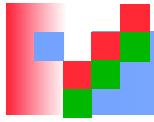


Structural Patterns

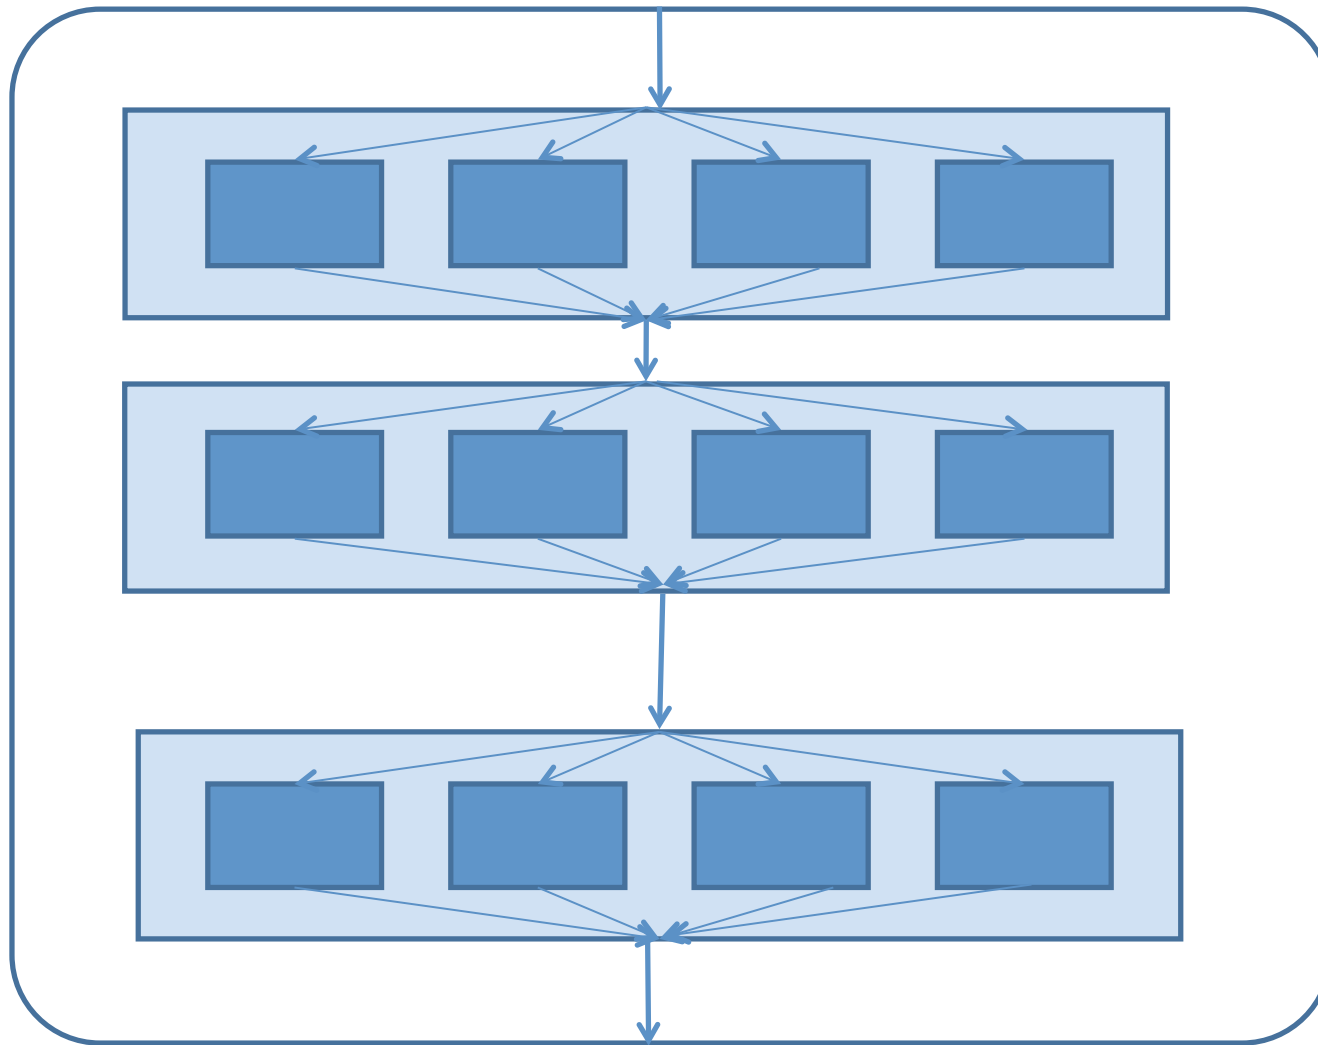
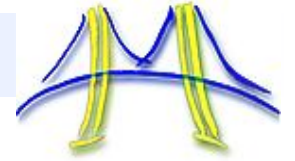
- Pipe-and-Filter
- Agent-and-Repository
- Event-based coordination
 - Iterator
 - MapReduce
- Process Control
- Layered Systems

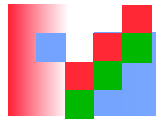


These define the structure of our software but they *do not describe* what is computed

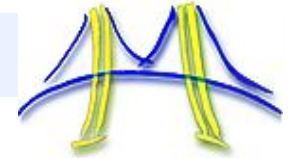


Analogy: Layout of Factory Plant











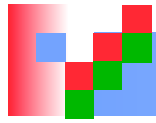
Identify Key Computations



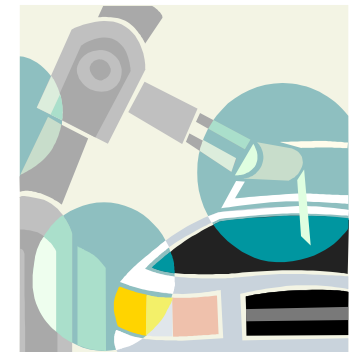
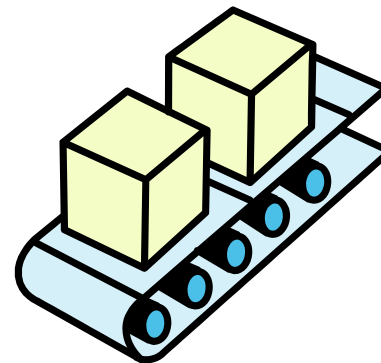
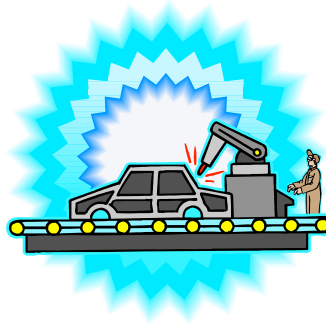
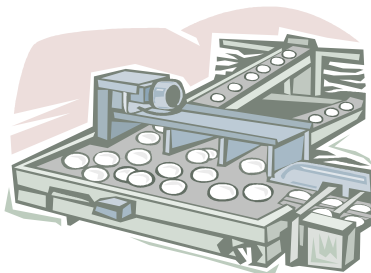
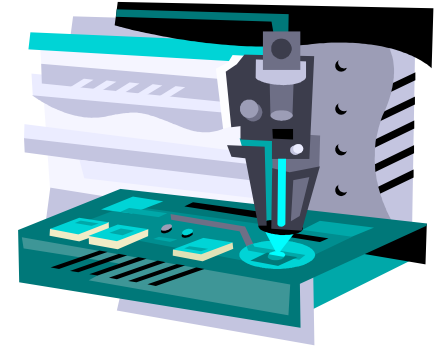
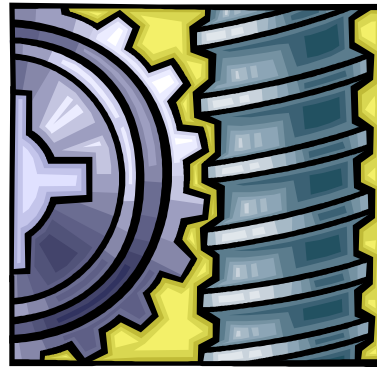
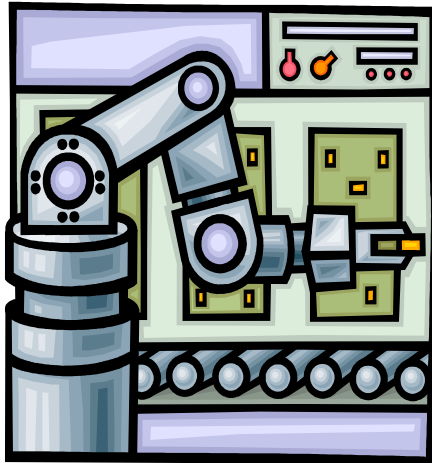
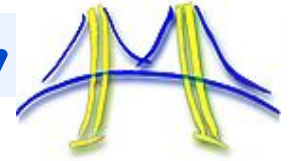
Computational Patterns

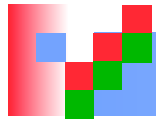
	Embed	SPEC	DB	Games	ML	HPC	 Health	 Image	 Speech	 Music	 Browser	 CAD
Finite State Mach.	Red	Red	Red	Yellow	Yellow	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Red	Yellow
Circuits	Red	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Light Blue	Red	Light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

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

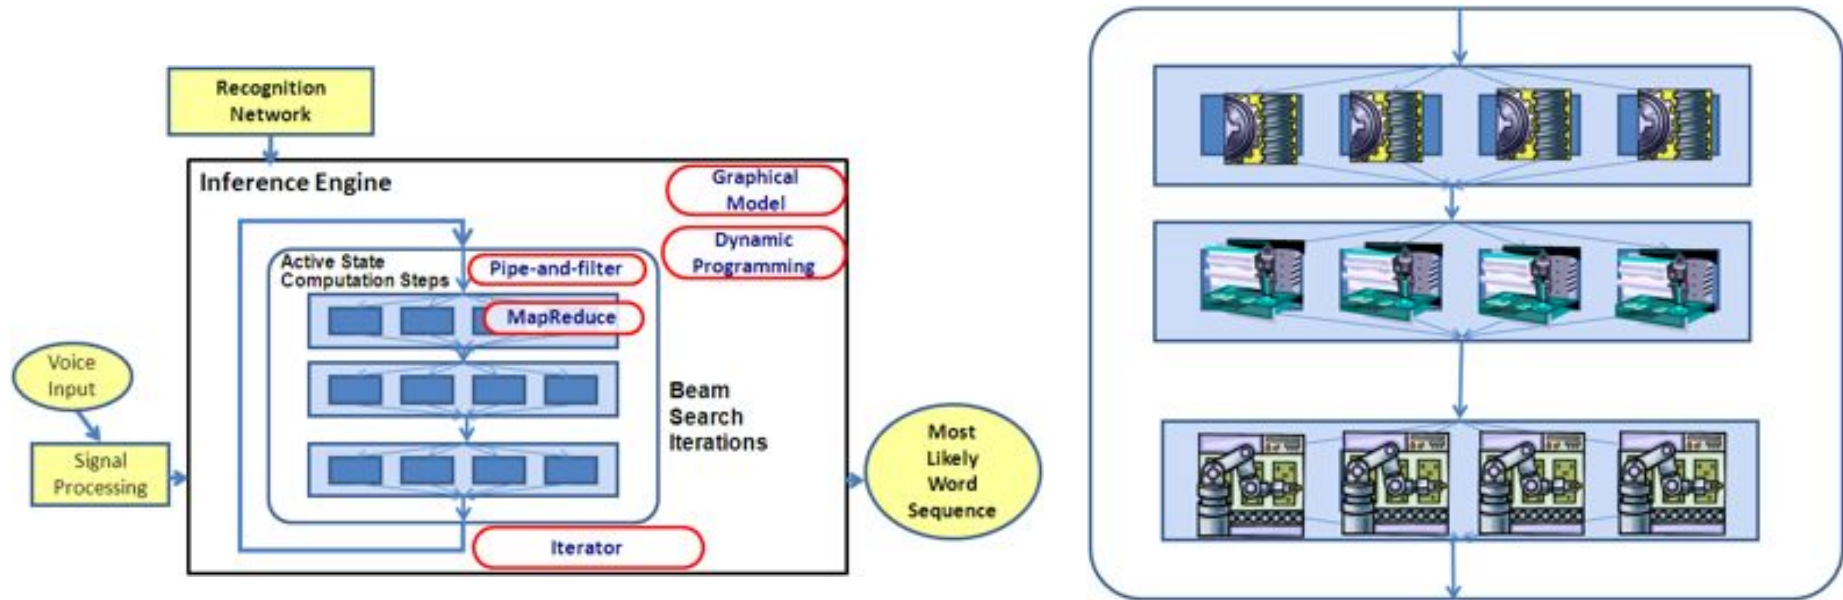
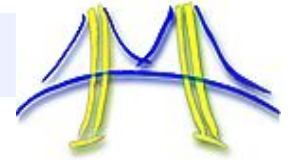


Analogy: Machinery of the Factory





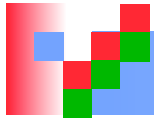
Architecting the Whole Application



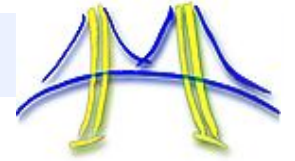
- **SW Architecture of Large-Vocabulary Continuous Speech Recognition**

Analogous to the design of an entire manufacturing plant

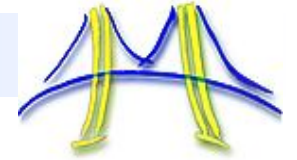
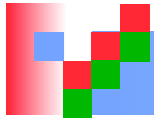
- Raises appropriate issues like scheduling, latency, throughput, workflow, resource management, capacity etc.



Outline

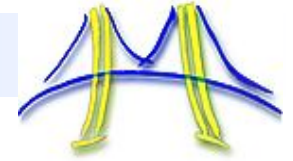
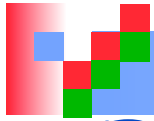


- Intro to Kurt
- General approach to applying the pattern language
- ⇒ ■ Detail on Structural Patterns
- High-level examples of composing patterns



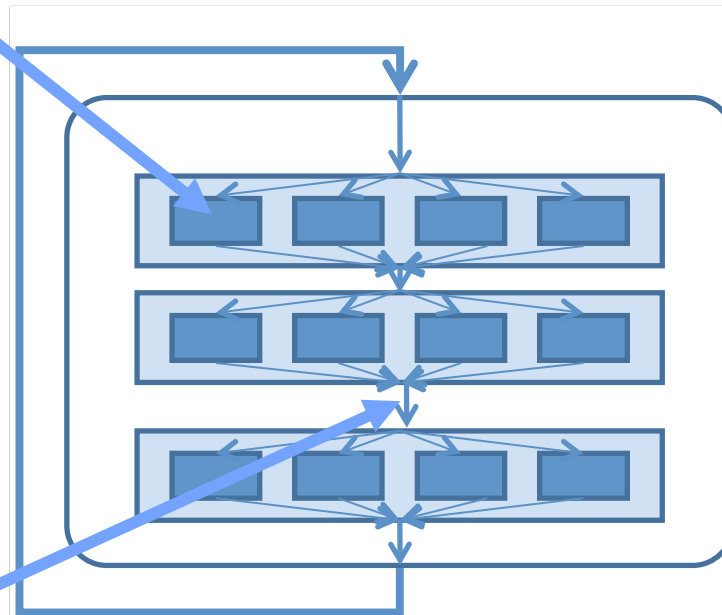
Inventory of Structural Patterns

- pipe and filter
- iterator
- MapReduce
- blackboard/agent and repository
- process control
- layered
- event-based coordination
- puppeteer
- (call-and-return/arbitrary task graph)



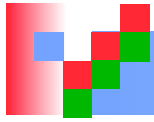
Elements of a structural pattern

- Components are where the computation happens

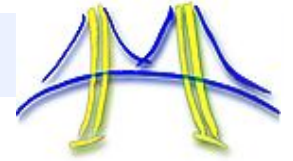


- A configuration is a graph of components (vertices) and connectors (edges)
- A structural patterns may be described as a family of graphs.

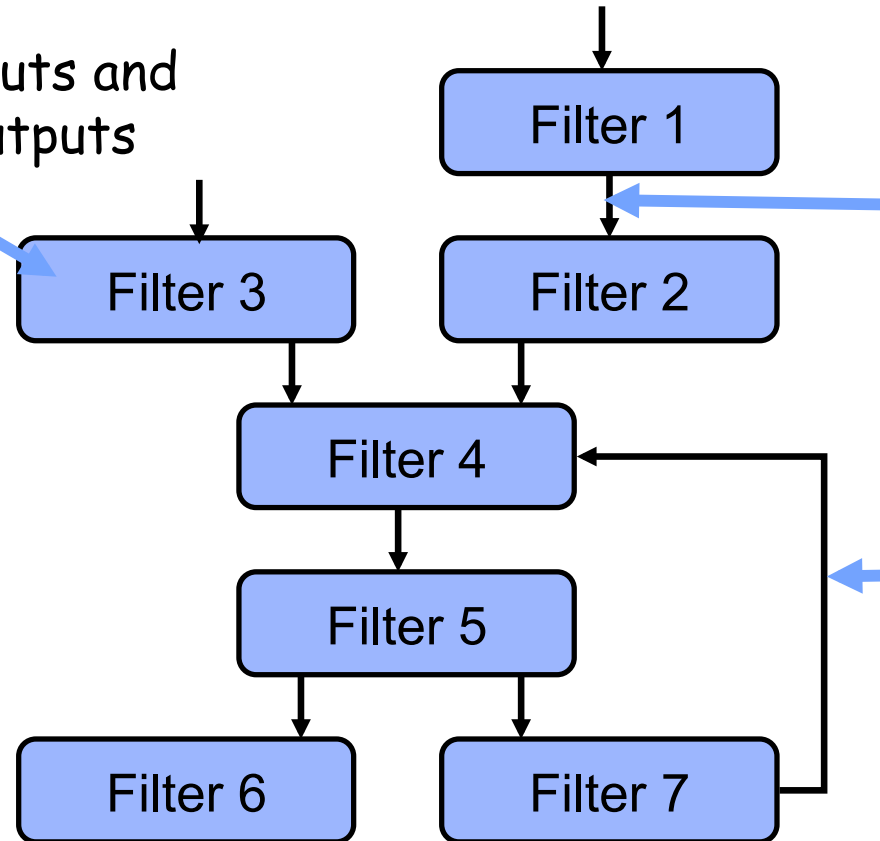
Connectors are where the communication happens



Pattern 1: Pipe and Filter



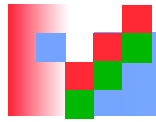
- Filters embody computation
- Only see inputs and produce outputs



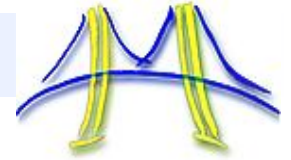
• Pipes embody communication

May have feedback

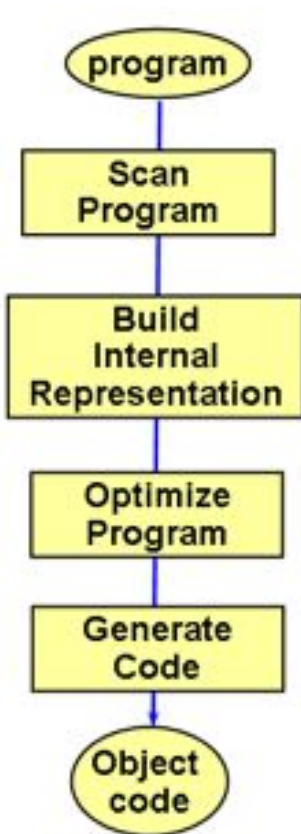
Examples?



Examples of pipe and filter



- Almost every large software program has a pipe and filter structure at the highest level



Compiler

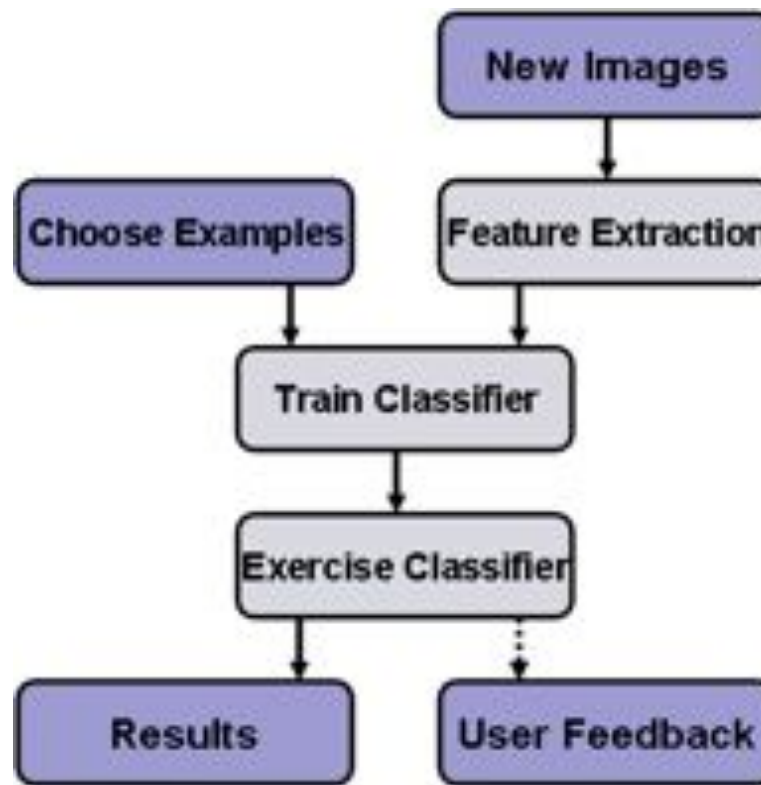
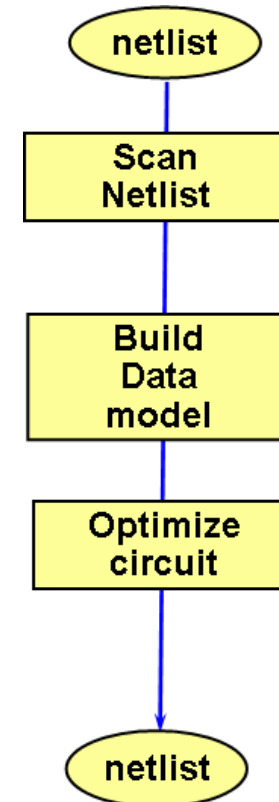
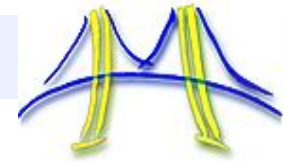
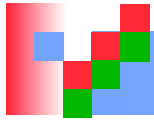


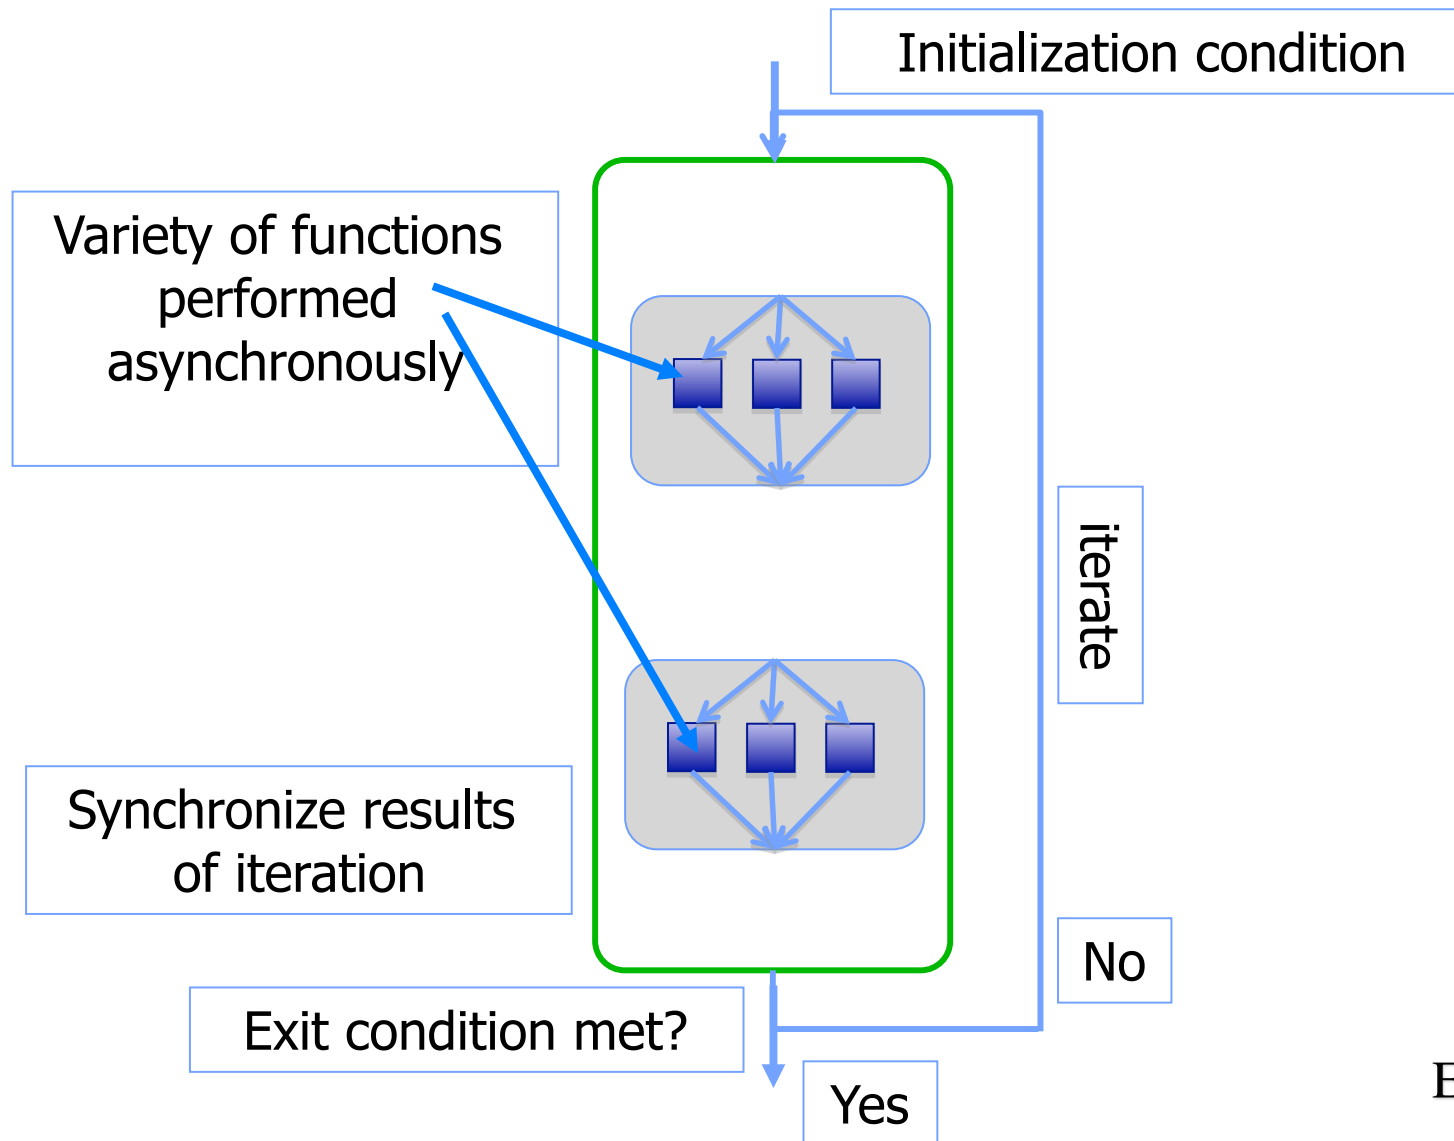
Image Retrieval System



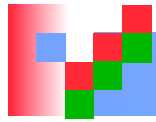
Logic optimizer



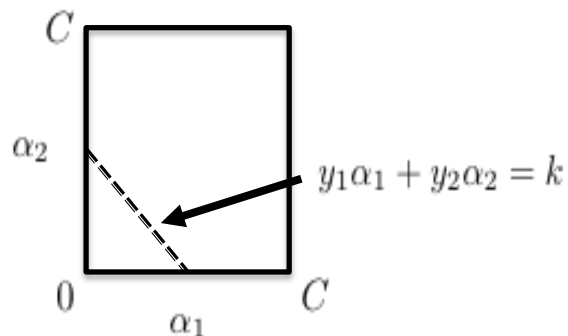
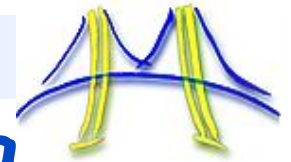
Pattern 2: Iterator Pattern



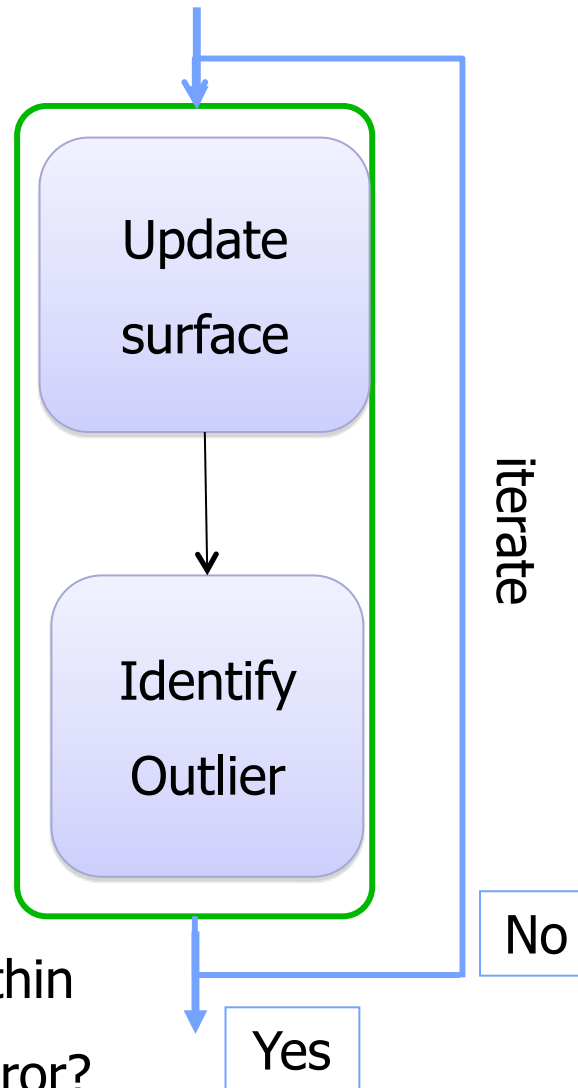
Examples?



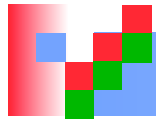
Example of Iterator Pattern: Training a Classifier: SVM Training



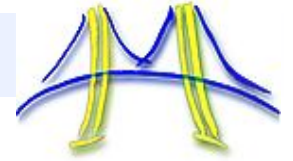
All points within
acceptable error?



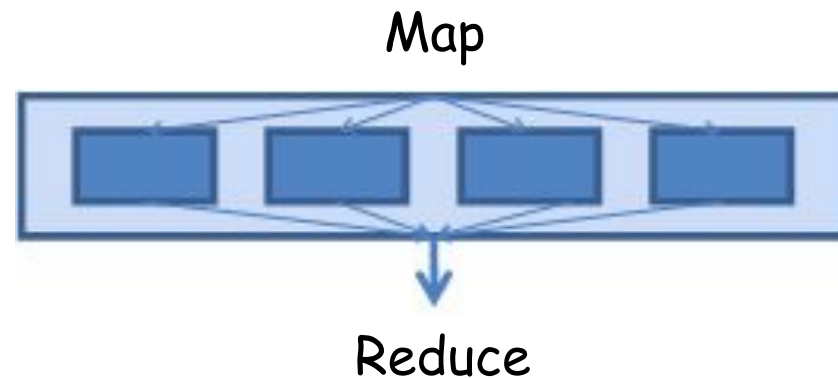
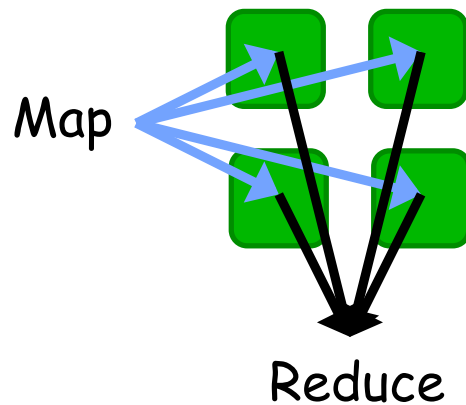
Iterator Structural Pattern



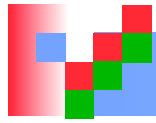
Pattern 3: MapReduce



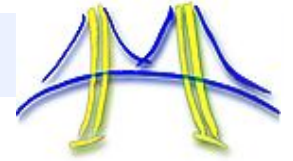
- To us, it means
 - A map stage, where data is mapped onto independent computations
 - A reduce stage, where the results of the map stage are summarized (i.e. reduced)



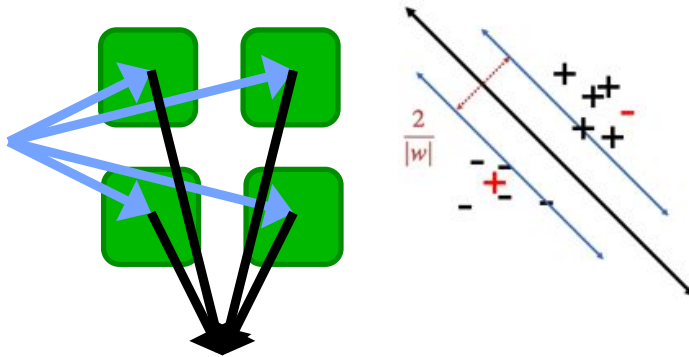
Examples?



Examples of Map Reduce

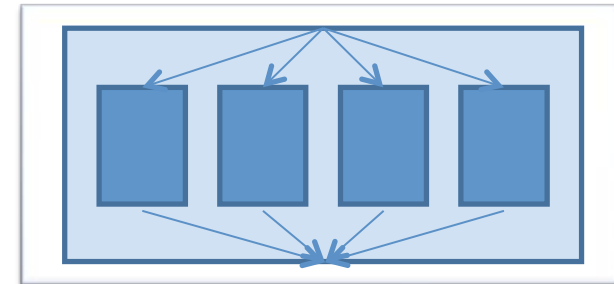


- General structure:
 - Map a computation across distributed data sets
 - Reduce the results to find the best/(worst), maxima/(minima)



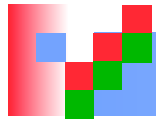
Support-vector machines (ML)

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

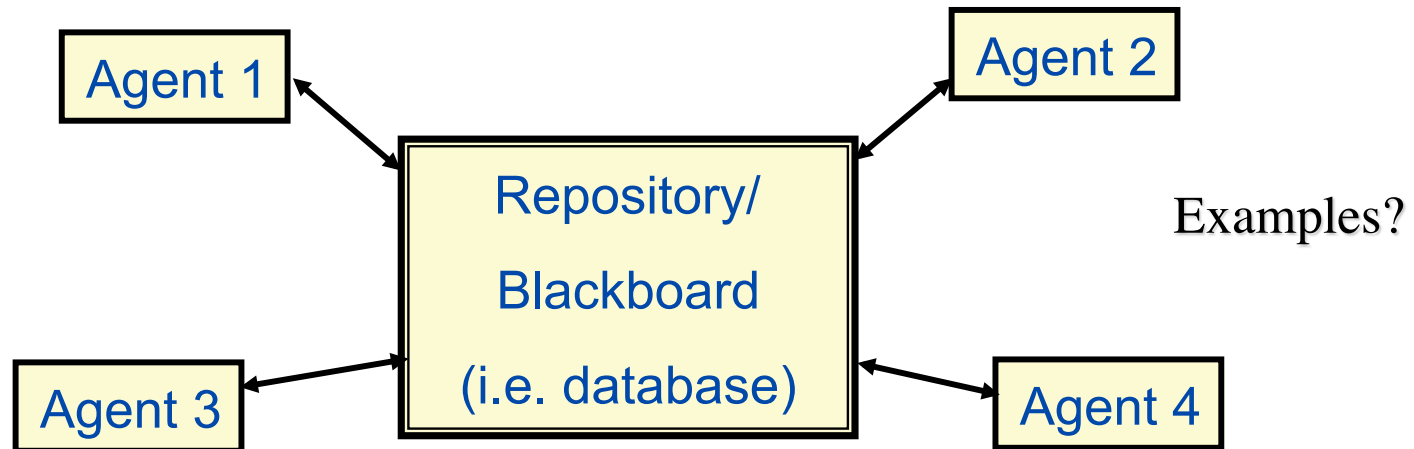
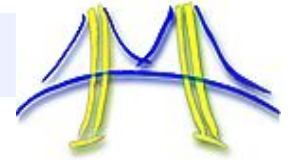


Speech recognition

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



Pattern 4: Agent and Repository

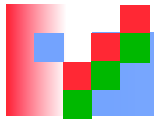


Agent and repository : Blackboard structural pattern

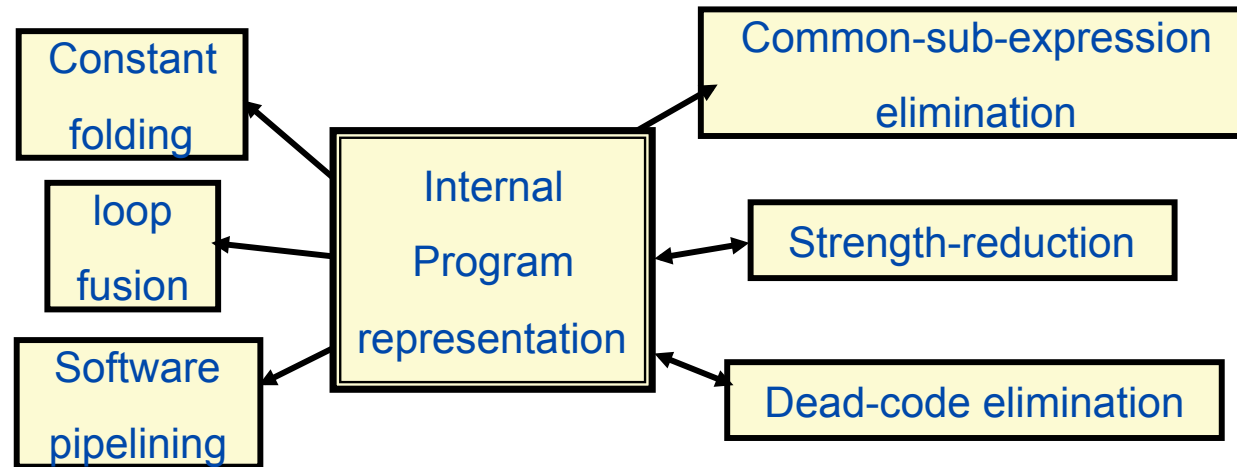
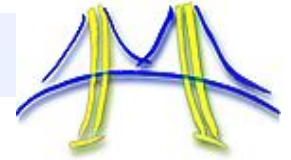
Agents cooperate on a shared medium to produce a result

Key elements:

- **Blackboard**: repository of the resulting creation that is shared by all agents (circuit database)
 - **Agents**: intelligent agents that will act on blackboard (optimizations)
- **Manager**: orchestrates agents access to the blackboard and creation of the aggregate results (scheduler)

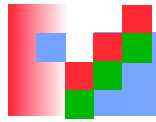


Example: Compiler Optimization

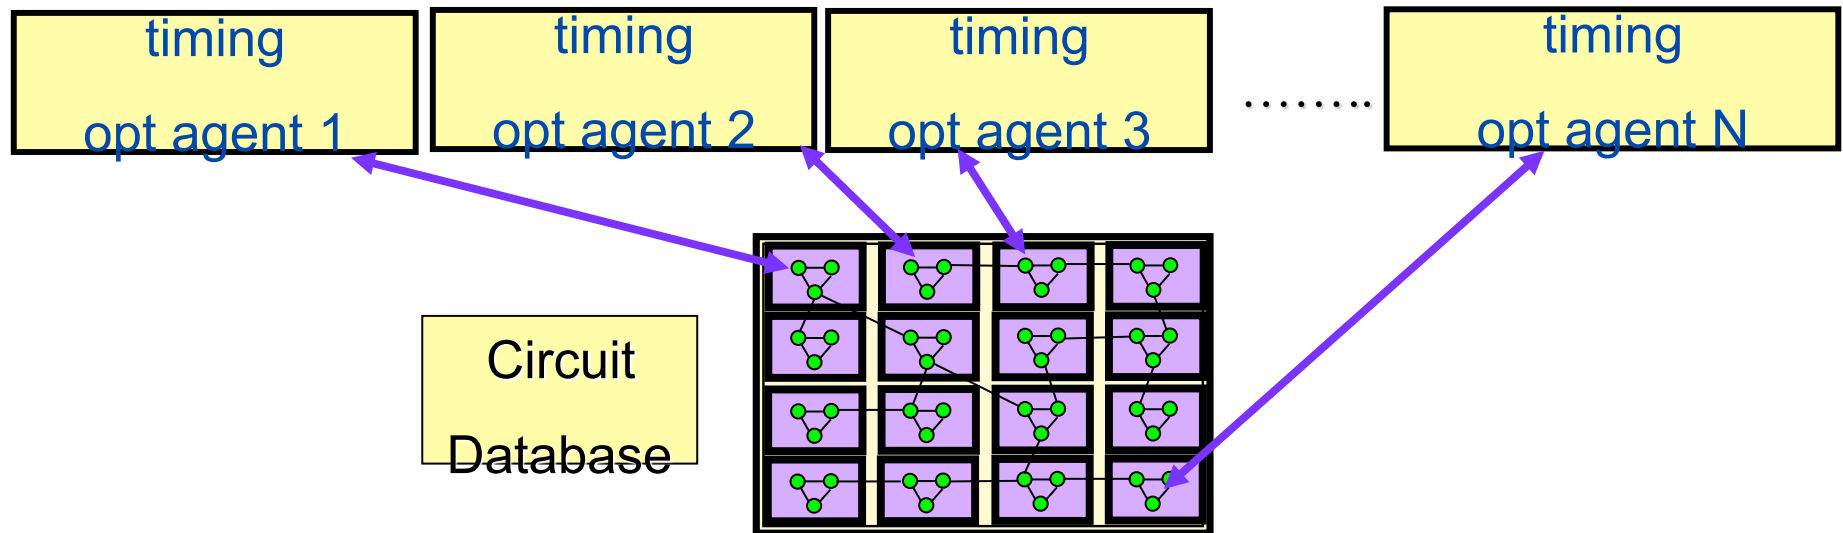
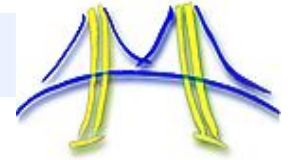


Optimization of a software program

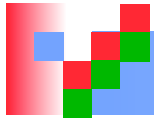
- Intermediate representation of program is stored in the repository
 - Individual agents have heuristics to optimize the program
- Manager orchestrates the access of the optimization agents to the program in the repository
 - Resulting program is left in the repository



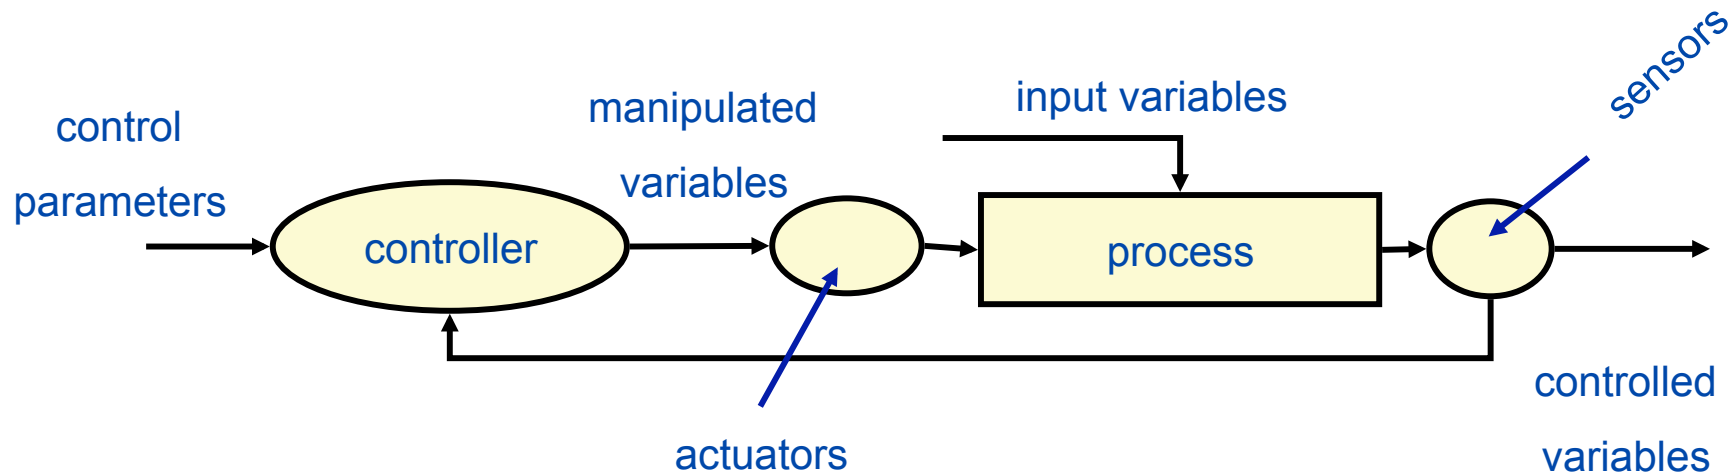
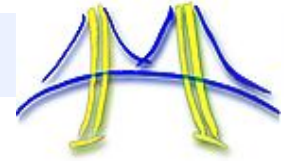
Example: Logic Optimization



- Optimization of integrated circuits
 - Integrated circuit is stored in the repository
 - Individual agents have heuristics to optimize the circuitry of an integrated circuit
 - Manager orchestrates the access of the optimization agents to the circuit repository
 - Resulting optimized circuit is left in the repository



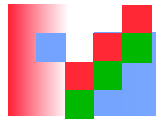
Pattern 5: Process Control



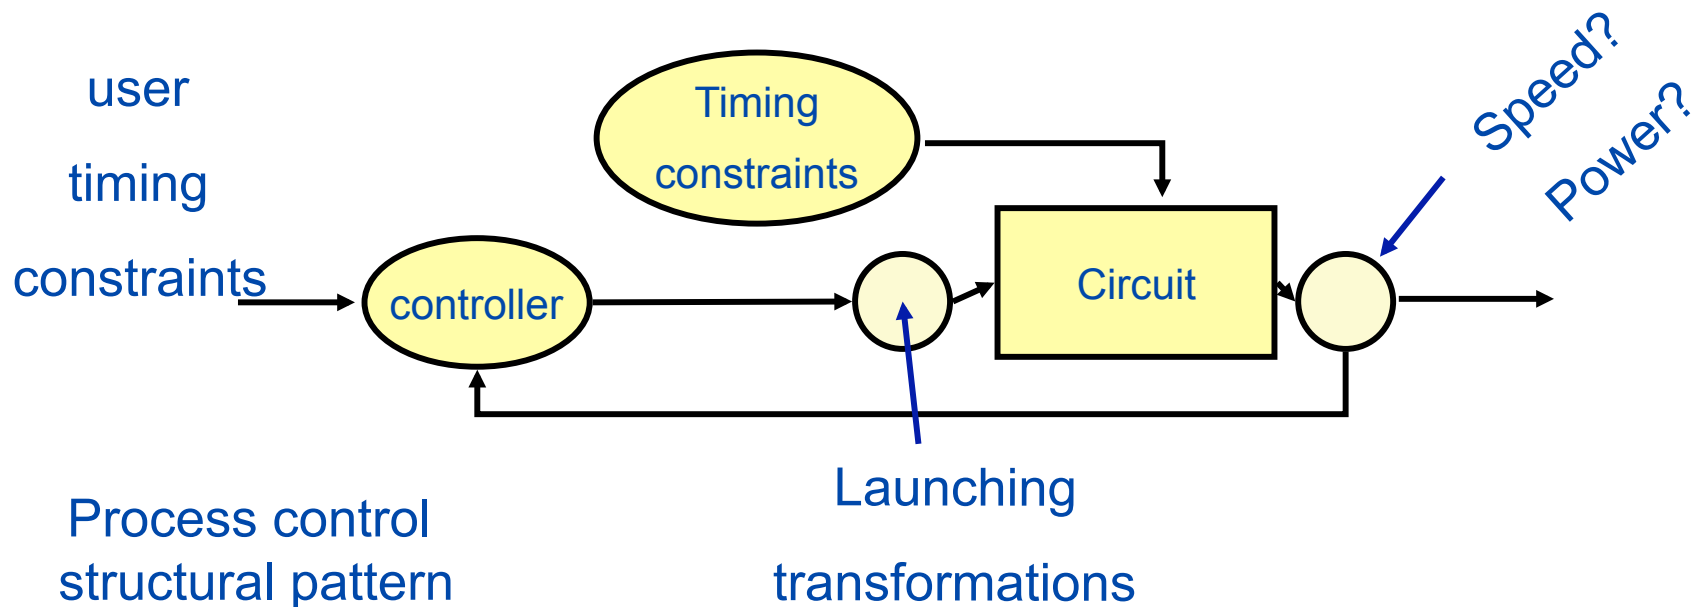
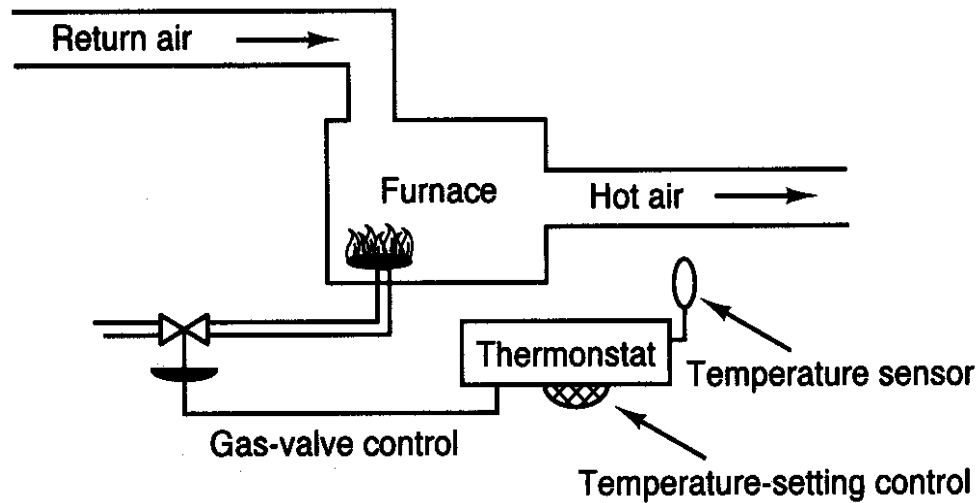
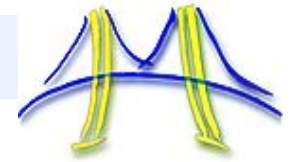
Source: Adapted from Shaw & Garlan 1996, p27-31.

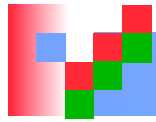
- **Process control:**
 - **Process**: underlying phenomena to be controlled/computed
 - **Actuator**: task(s) affecting the process
 - **Sensor**: task(s) which analyze the state of the process
 - **Controller**: task which determines what actuators should be effected

Examples?

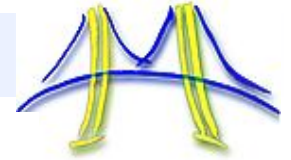


Examples of Process Control

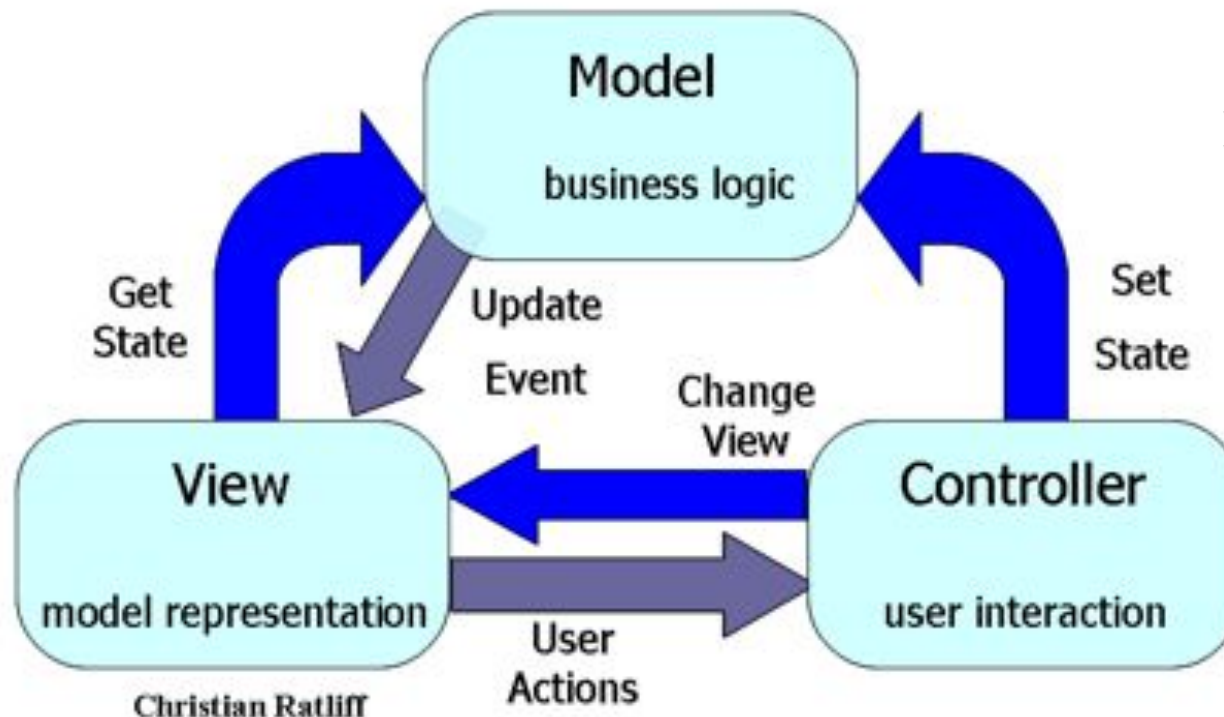




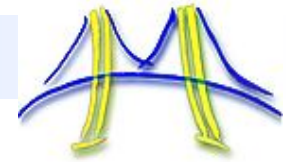
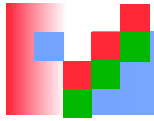
Pattern 6: Model-View-Controller



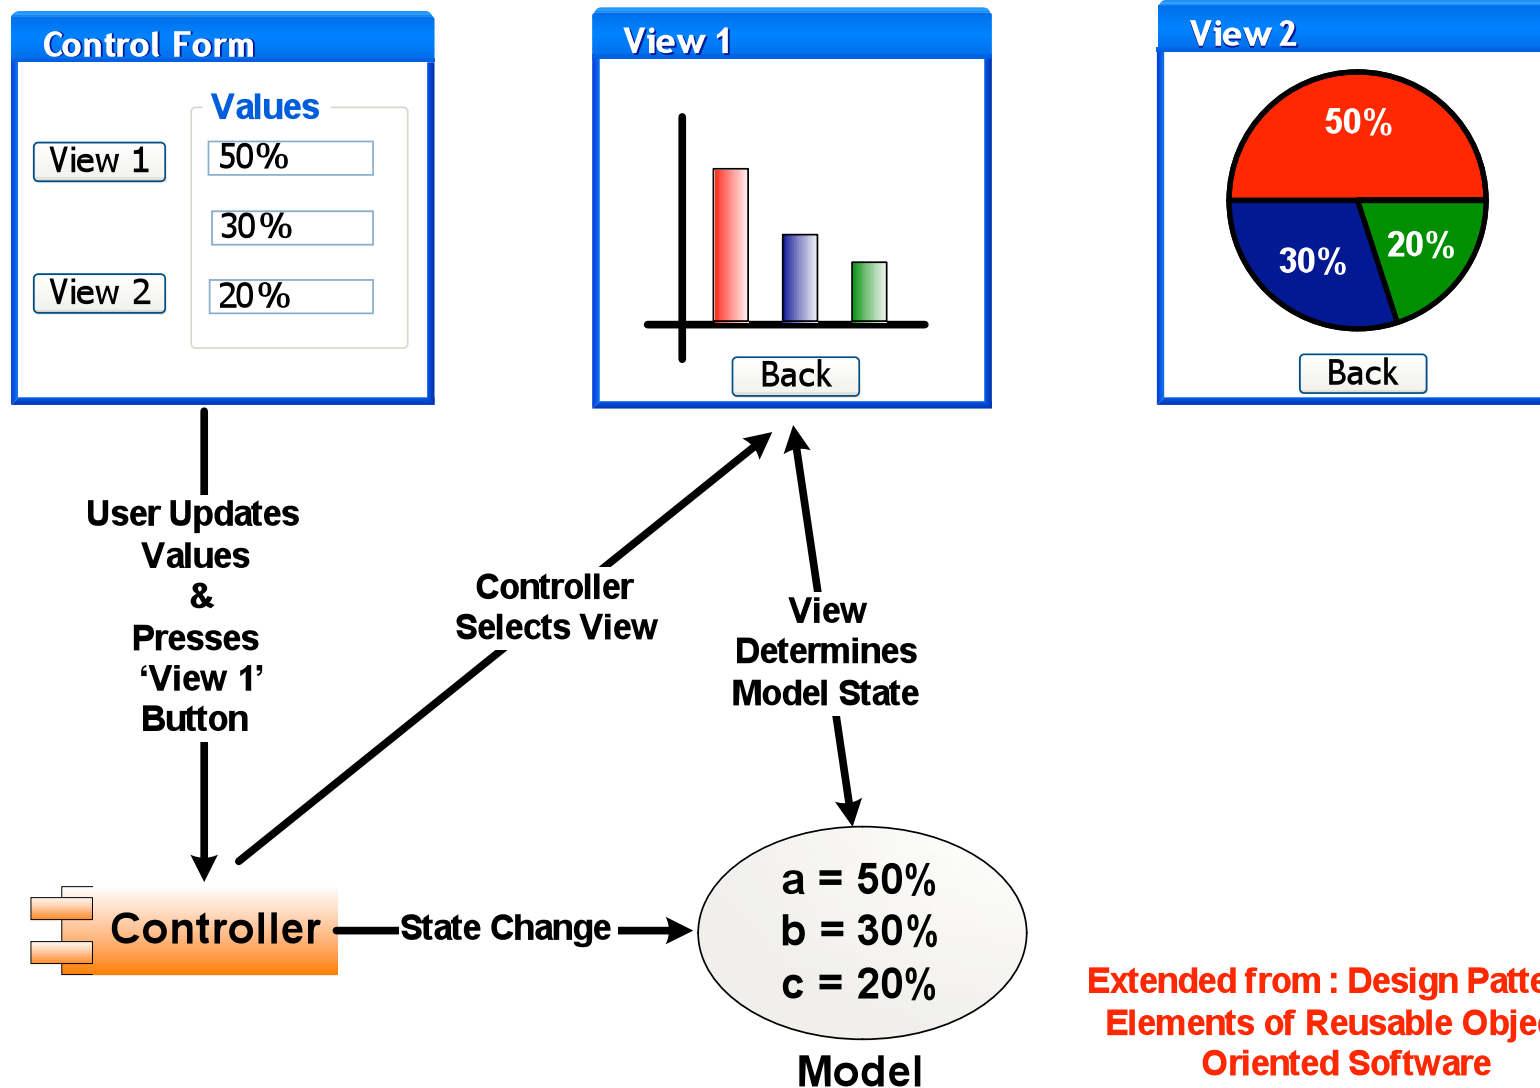
Examples?



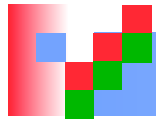
- **Model:** embodies the data and "intelligence" (aka business logic) of the system
- **Controller:** captures all user input and translates it into actions on the model
- **View:** renders the current state of the model for user



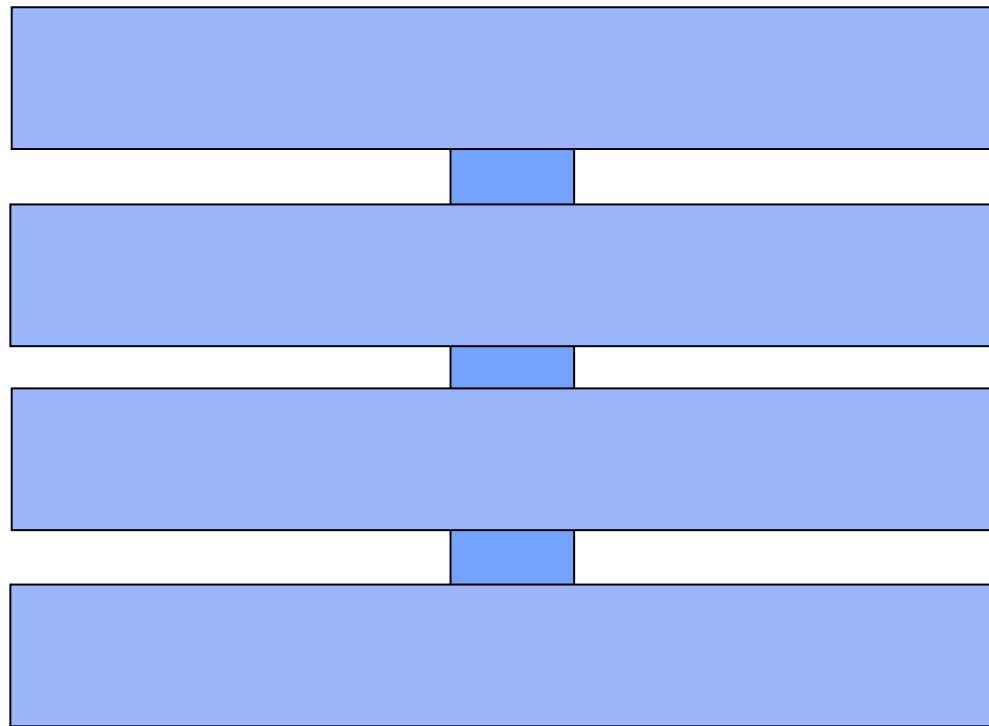
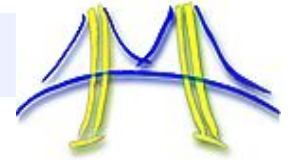
Example of Model-View Controller



Extended from : Design Patterns
Elements of Reusable Object -
Oriented Software

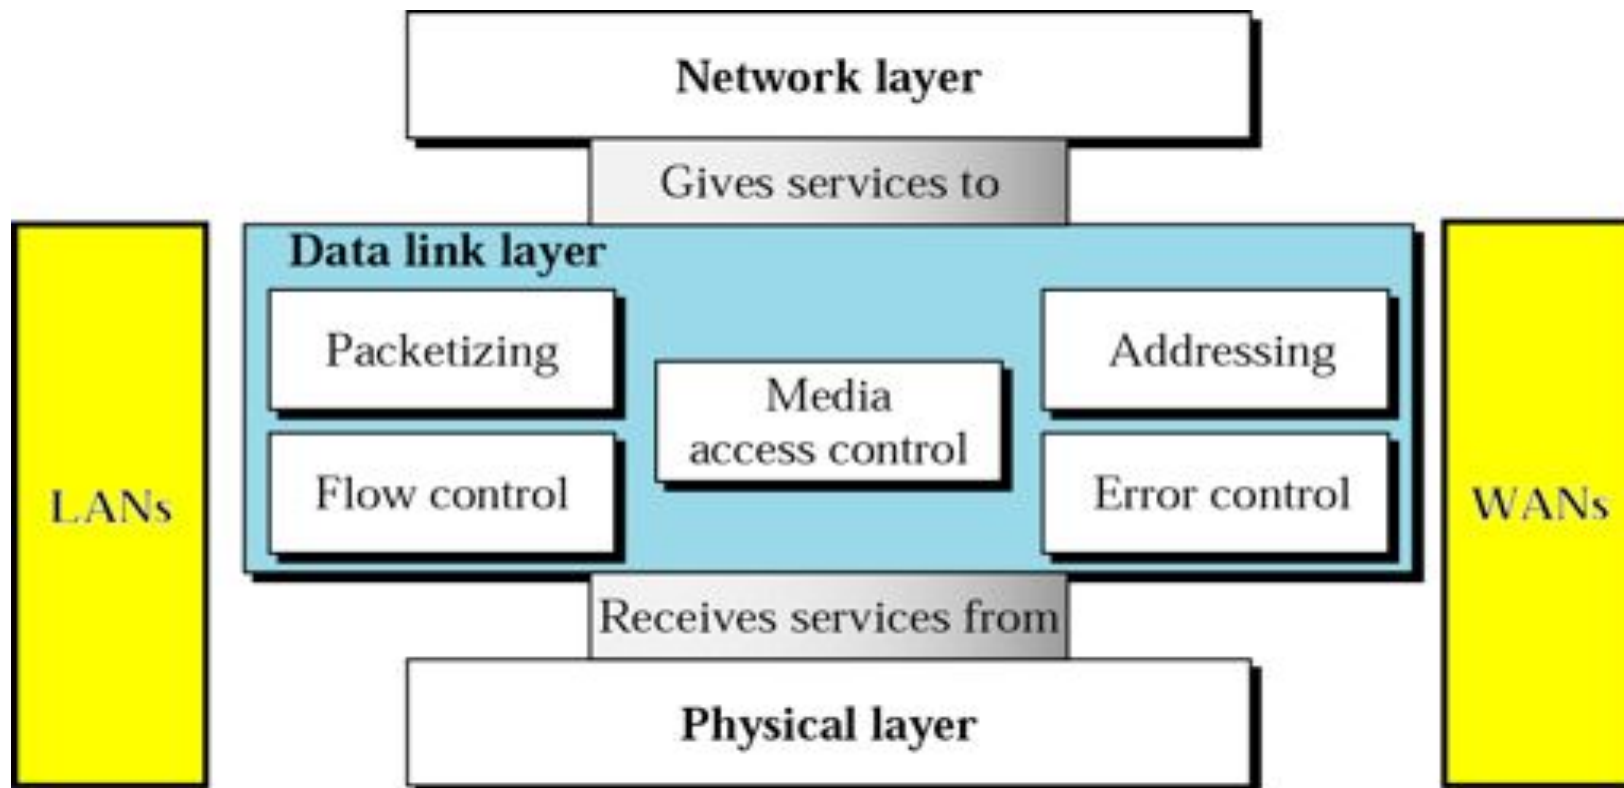
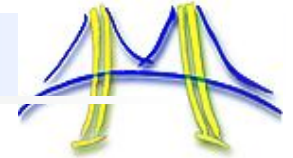


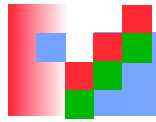
Pattern 7: Layered Systems



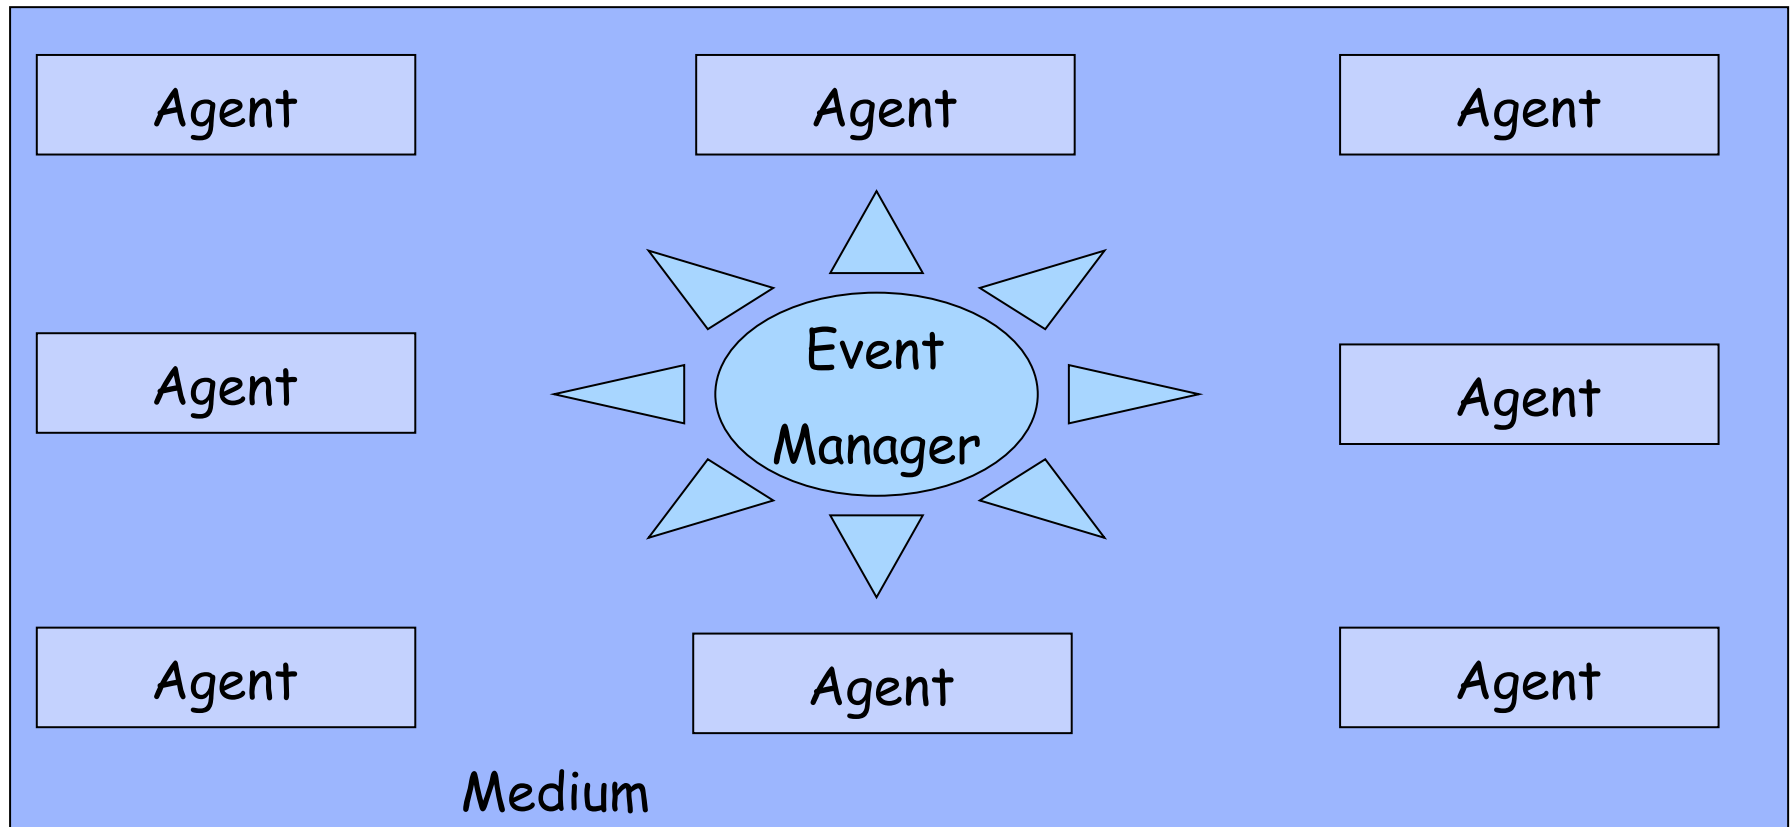
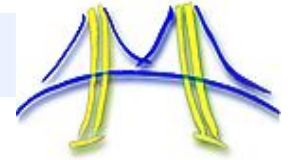
- Individual layers are big but the interface between two adjacent layers is narrow
- Non-adjacent layers cannot communicate directly. ^{Examples?}

Example: ISO Network Protocol



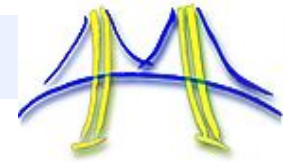
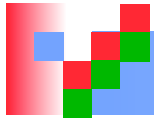


Pattern 8: Event-based Systems



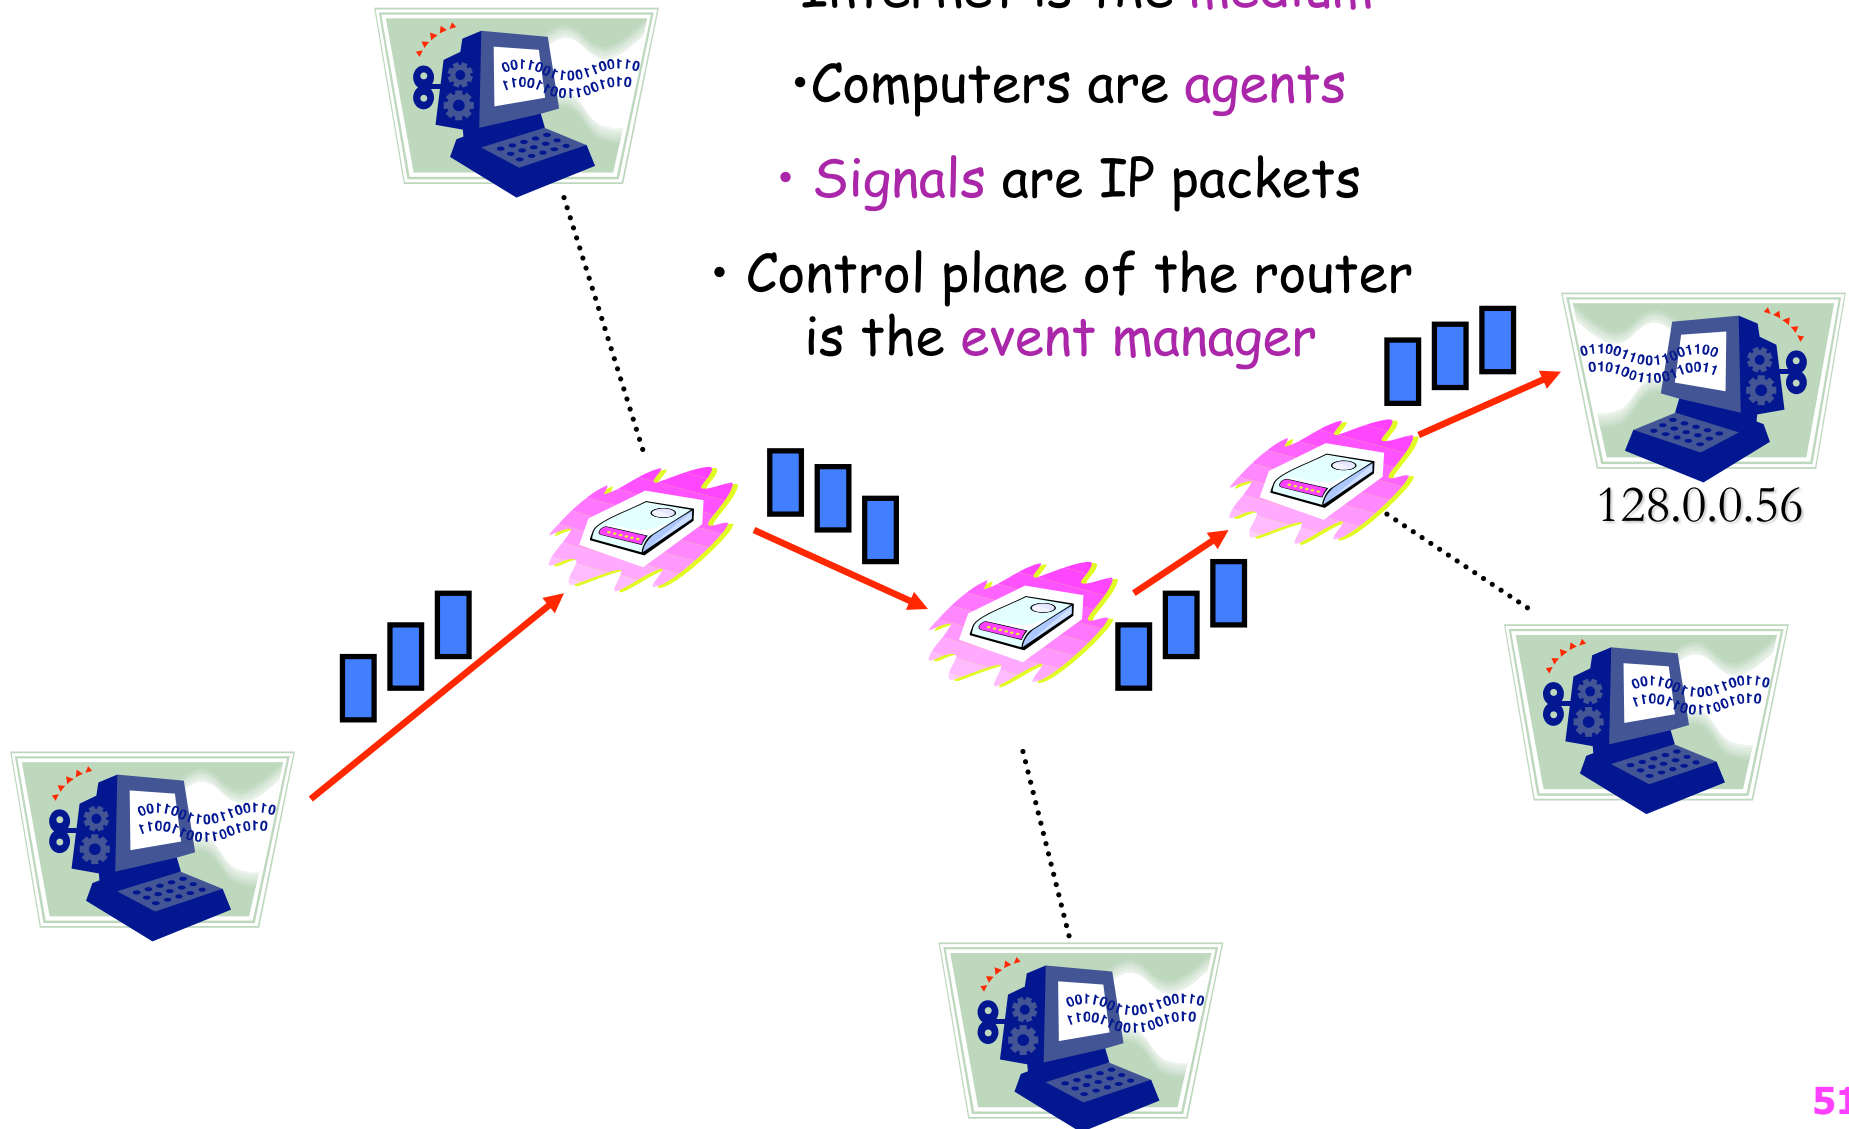
- Agents interact via **events/signals** in a **medium**
- **Event manager** manages **events**
- Interaction among **agents** is dynamic - no fixed connection

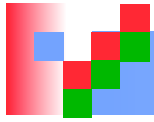
Examples?



Example: The Internet

- Internet is the **medium**
- Computers are **agents**
- **Signals** are IP packets
- Control plane of the router is the **event manager**

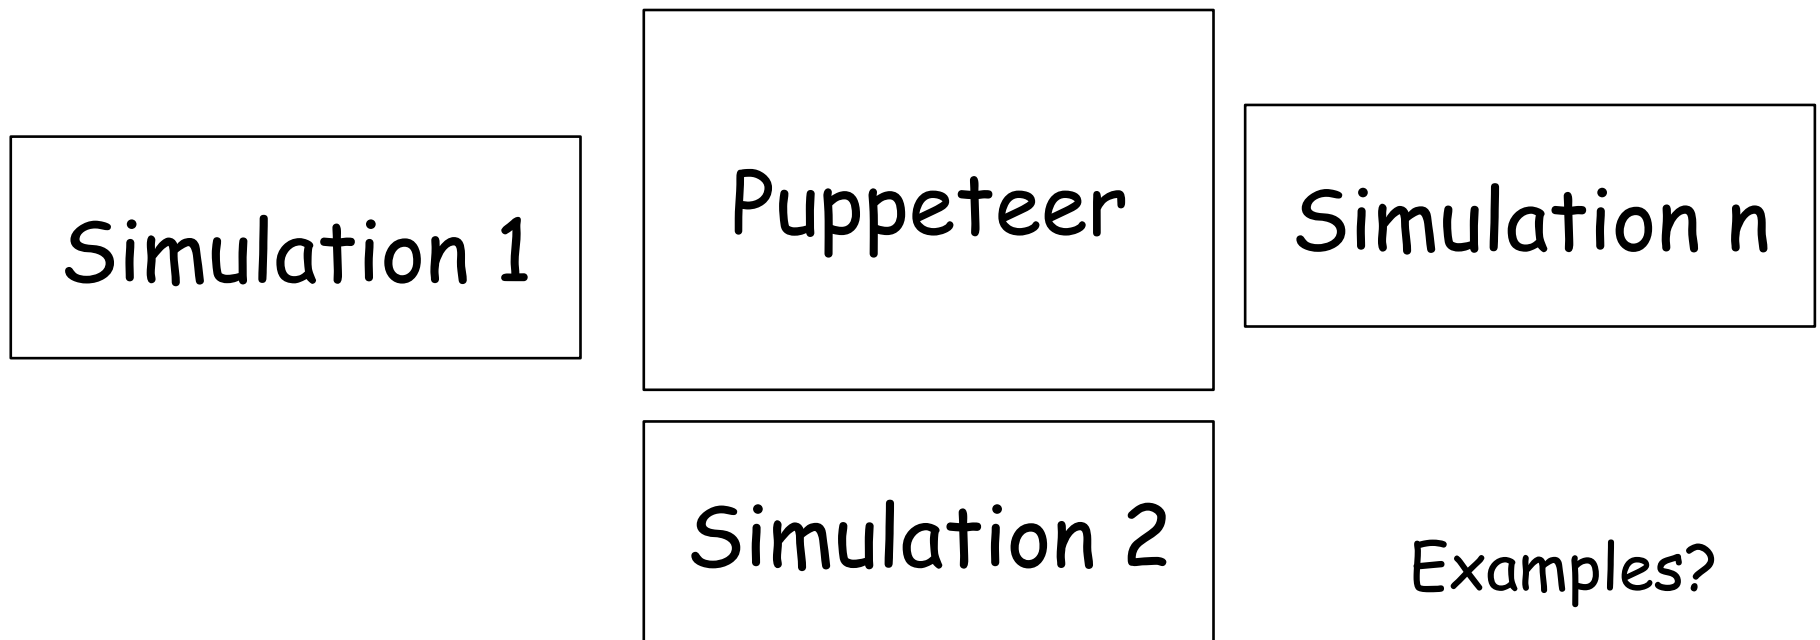


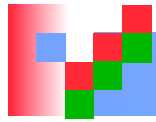


Pattern 9: Puppeteer

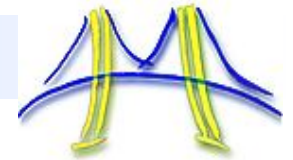


- Need an efficient way to manage and control the interaction of multiple simulators/computational agents
- **Puppeteer Pattern** - guides the interaction between the simulation codes to guarantee correctness of the overall simulation
- Difference with agent and repository?
 - No central repository
 - Data transfer between simulators

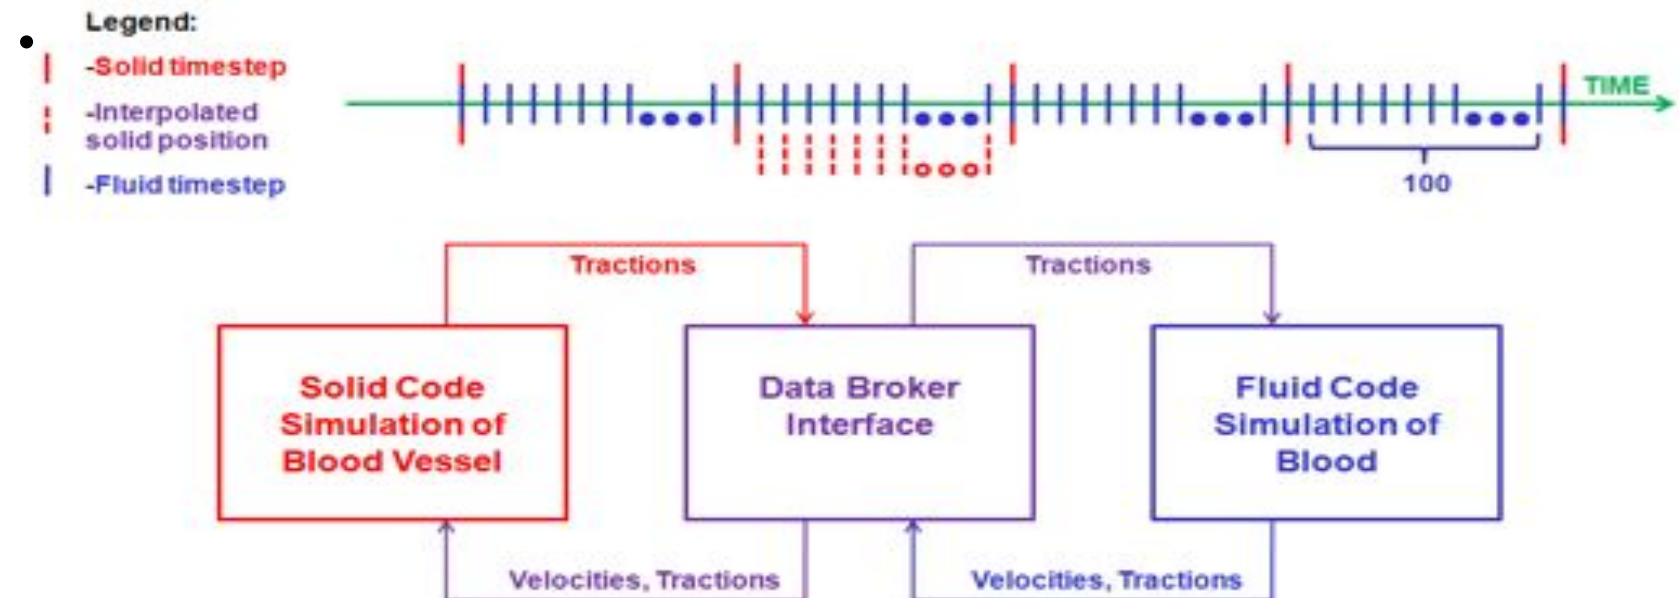


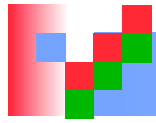


Overall Computation

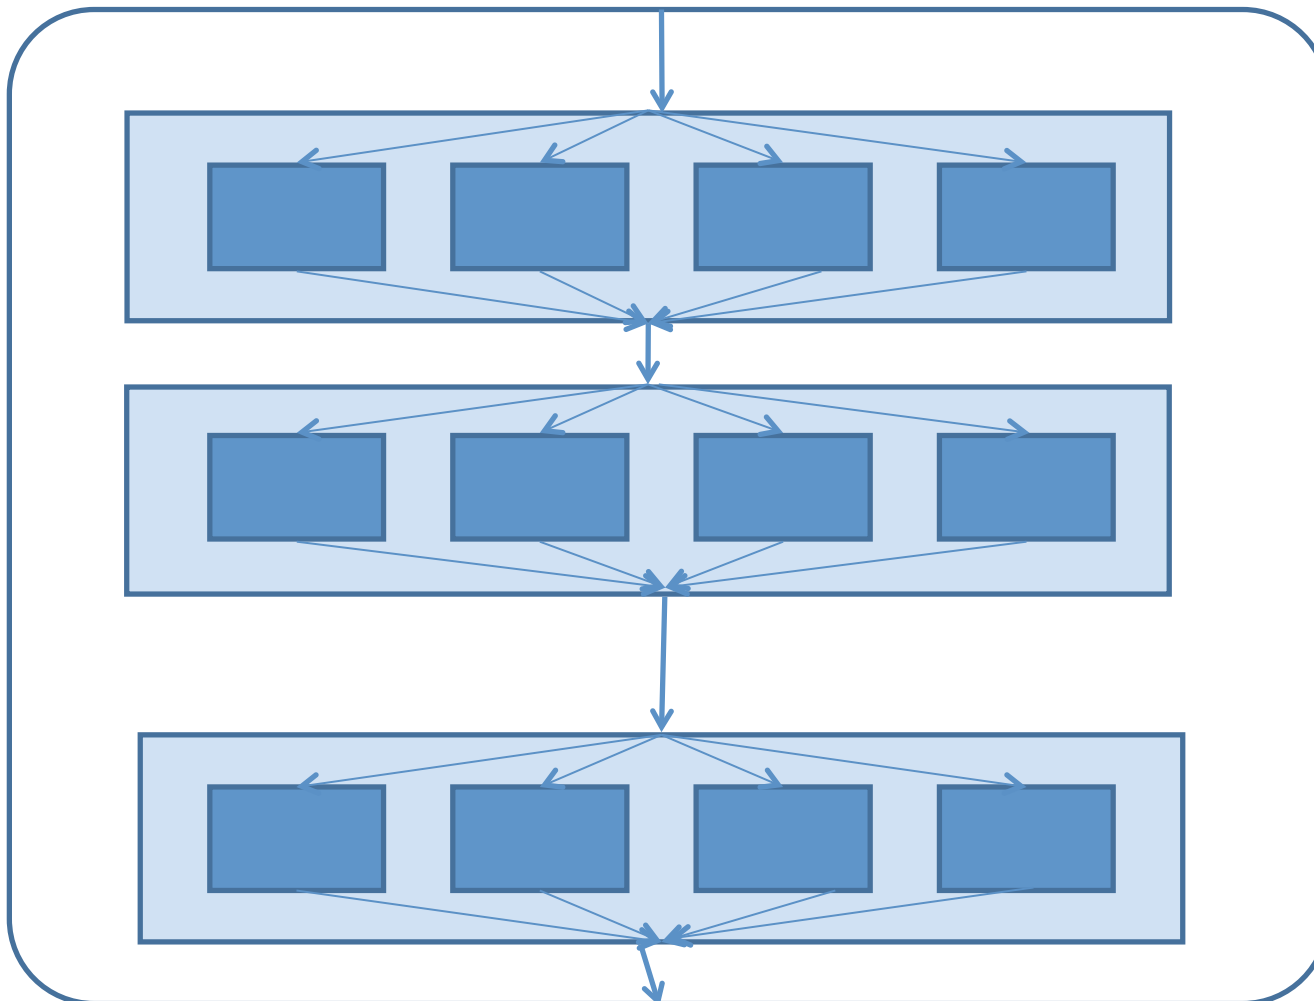
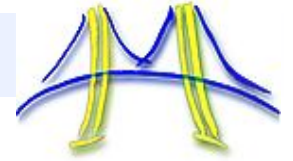


- Modeling of blood moving in blood vessels
- The computation is structured as a controlled interaction between solid (blood vessel) and fluid (blood) simulation codes
- The two simulations use different data structures and the number of iterations for each simulation code varies
- Need an efficient way to manage and control the interaction of the two codes

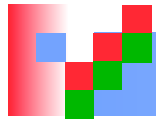




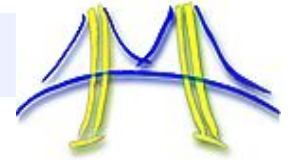
Remember the Analogy: Layout of Factory Plant



- We have only talked about structure. We haven't described computation.



Architecting Parallel Software



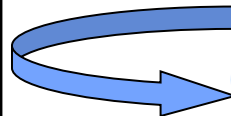
Decompose Tasks

- Group tasks
- Order Tasks

Decompose Data

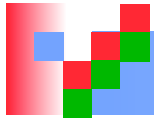
- Identify data sharing
- Identify data access

Identify the Software
Structure

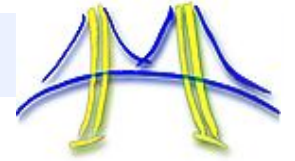


Identify the Key
Computations

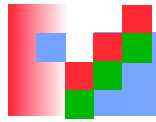
Friday: Computational Patterns of Parallel
Programming (James Demmel (UCB))
(8:45 - 10:45am)



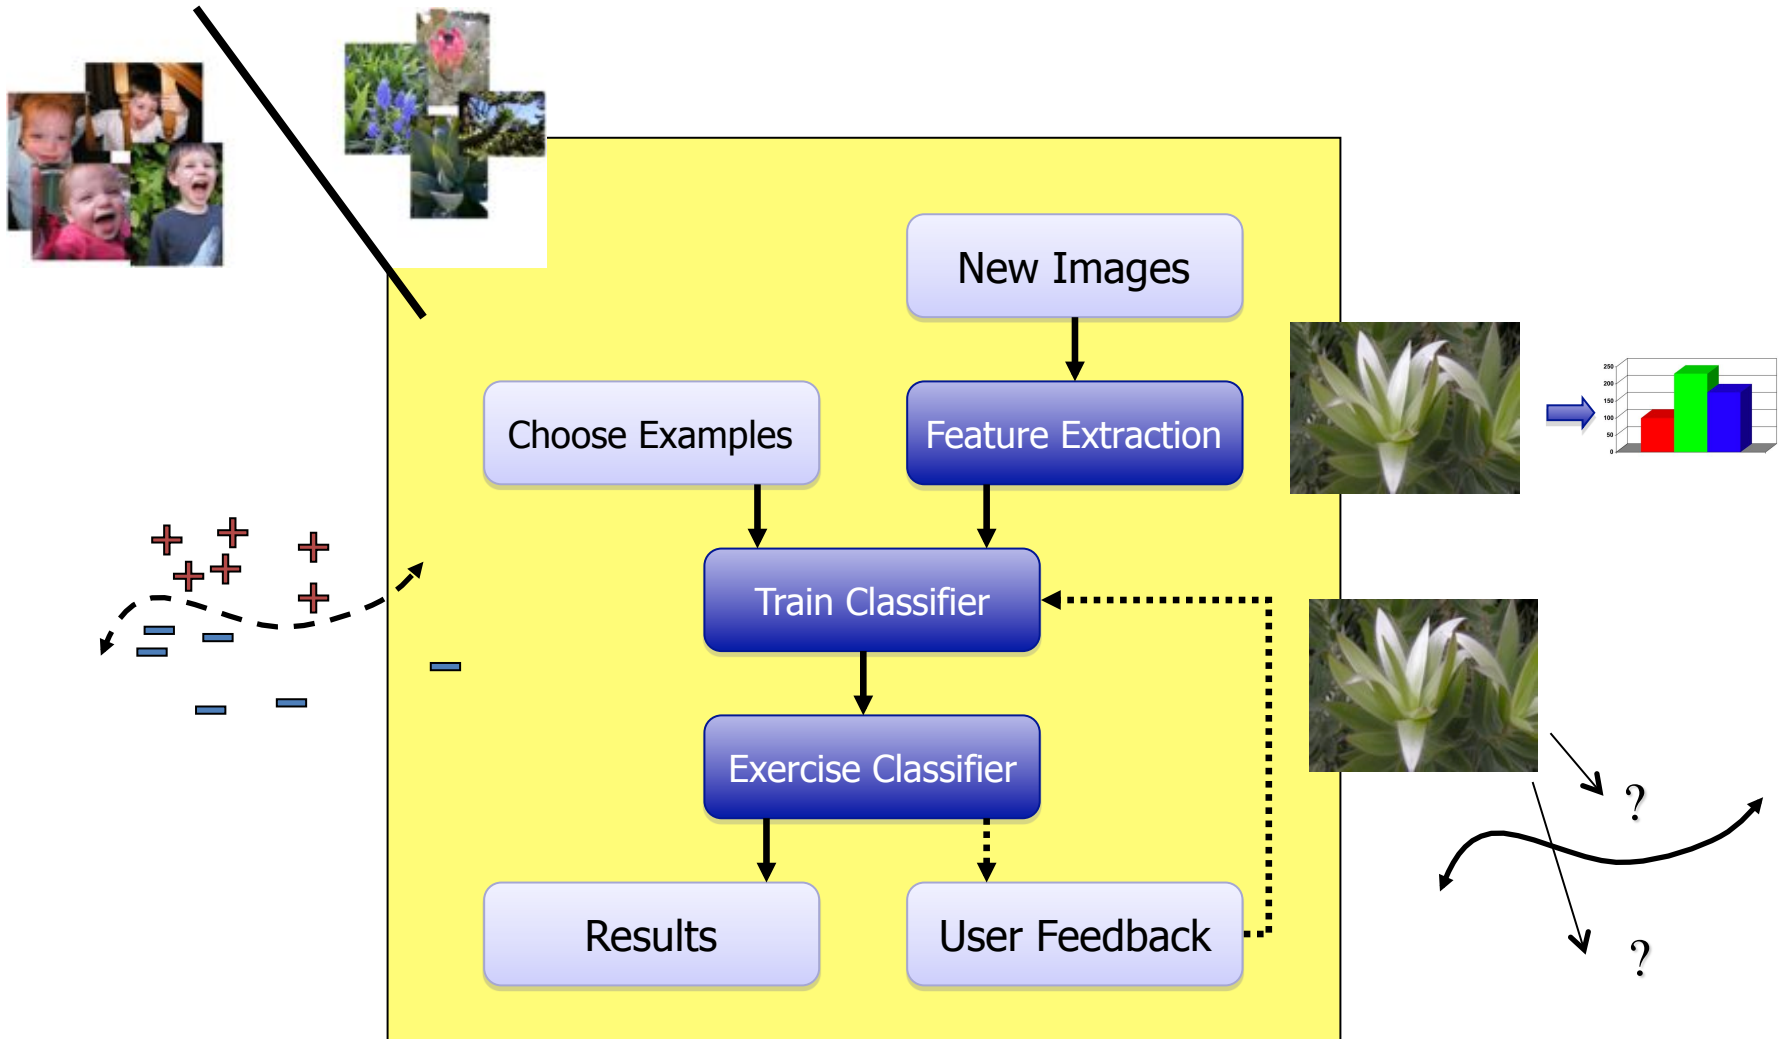
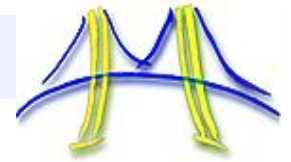
Outline



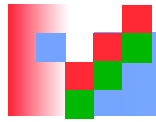
- Intro to Kurt
- General approach to applying the pattern language
- Detail on Structural Patterns
- ⇒ ■ High-level examples of composing patterns



CBIR Application Framework

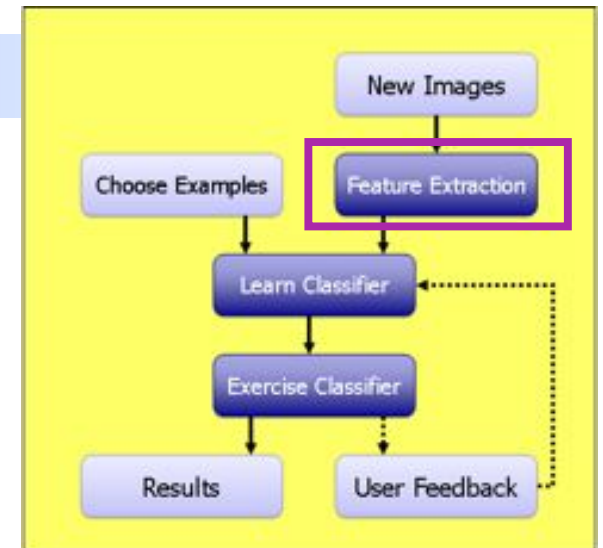
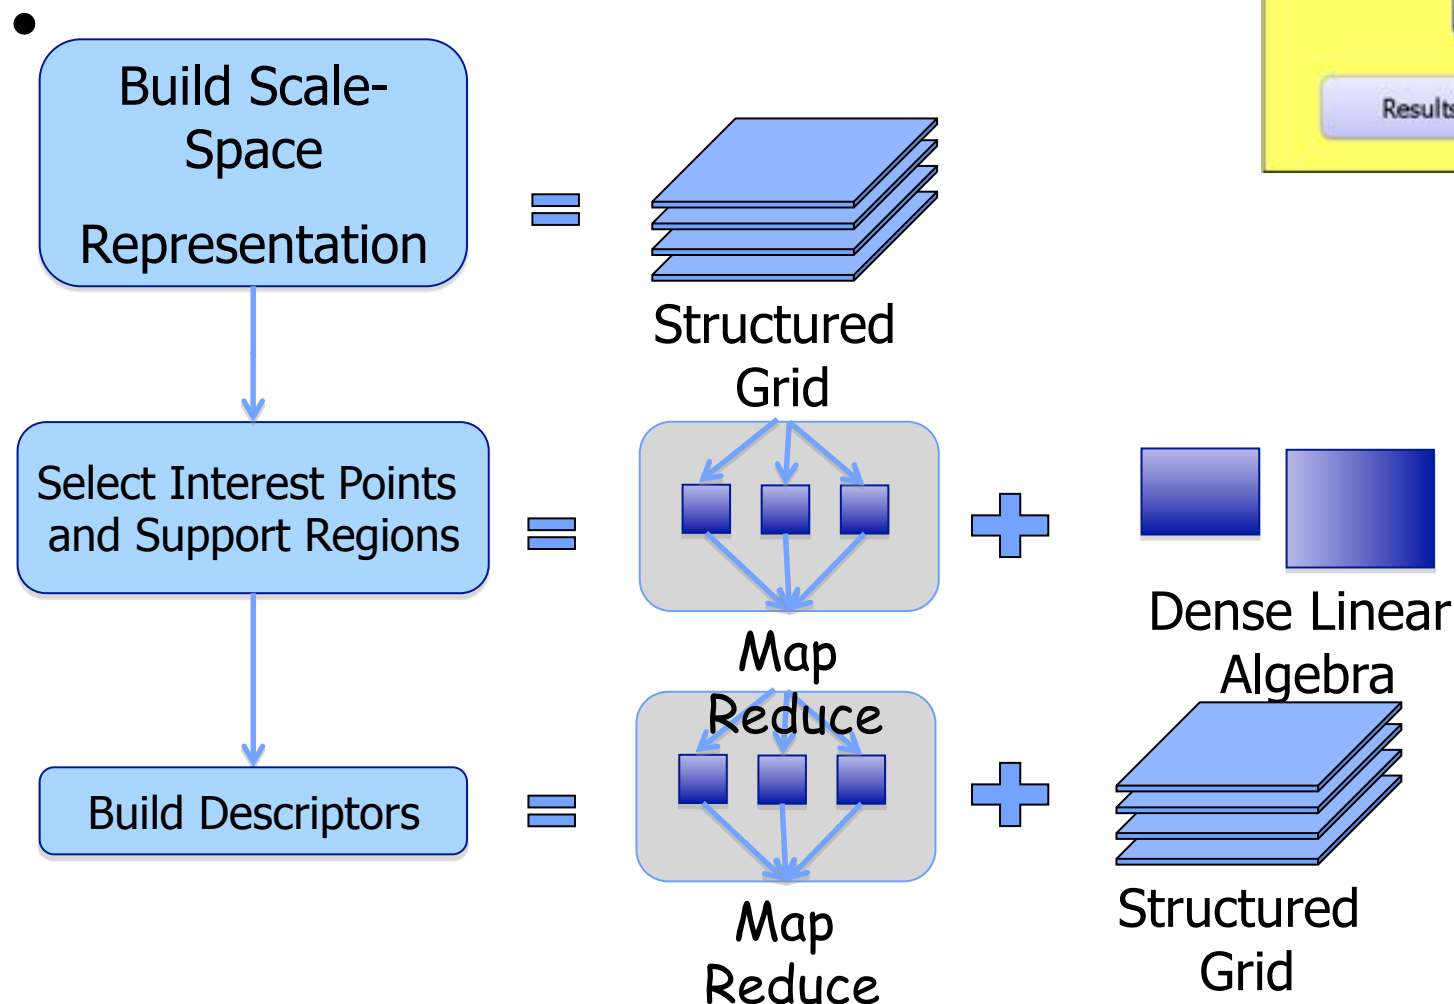


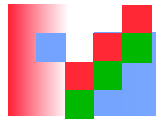
Catanzaro, Sundaram, Keutzer, "Fast SVM Training and Classification on Graphics Processors", ICML 2008



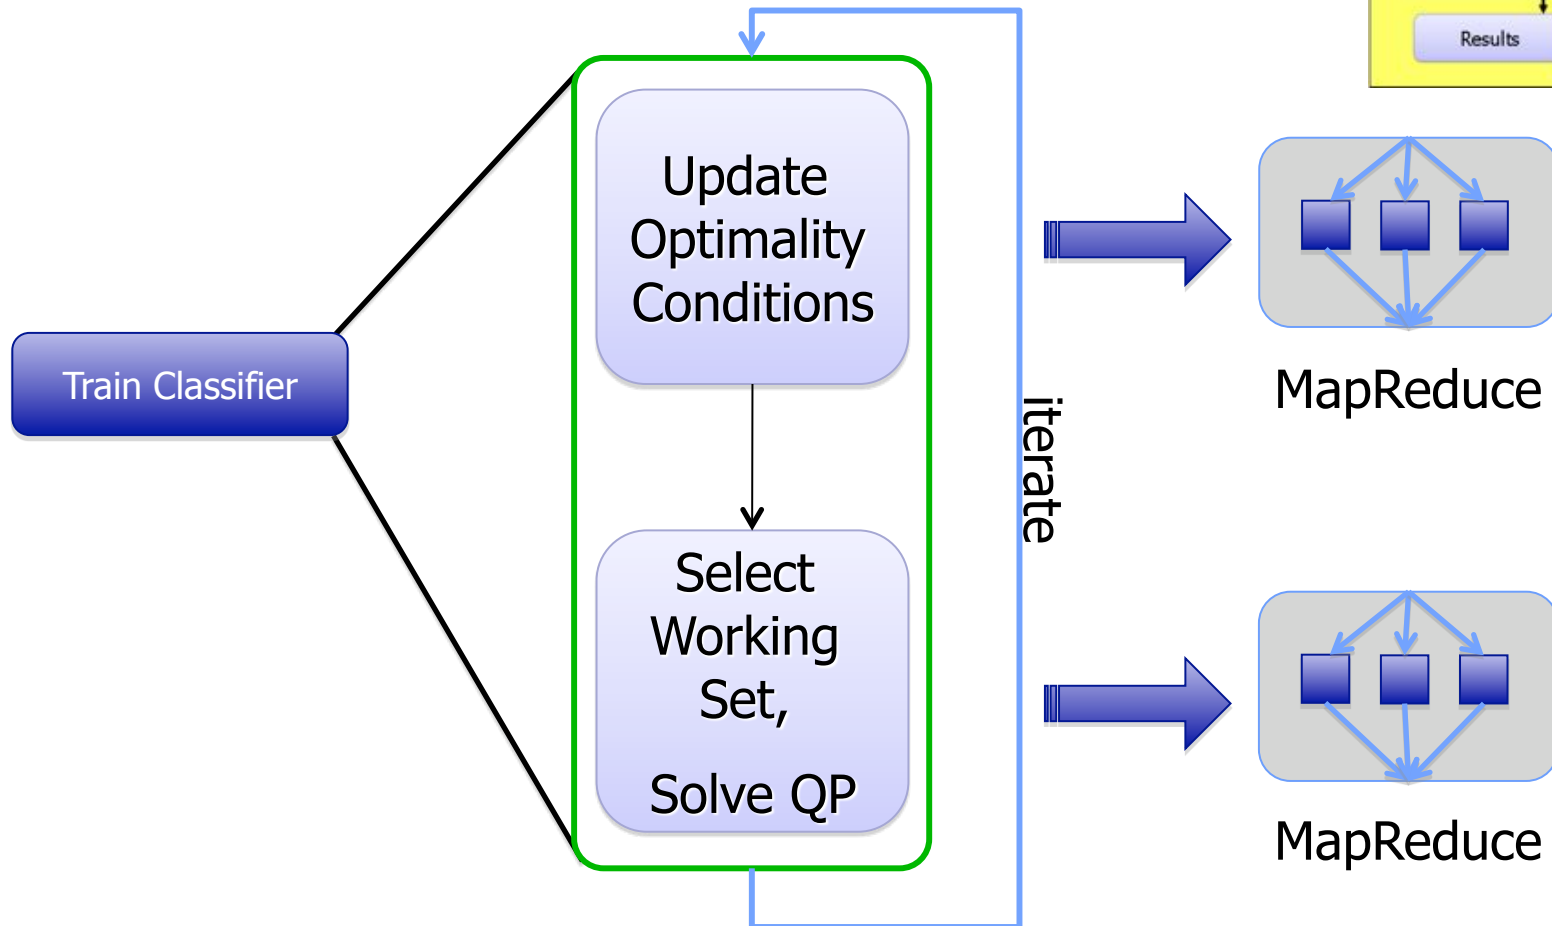
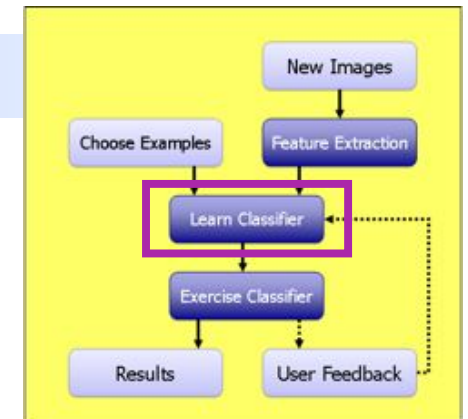
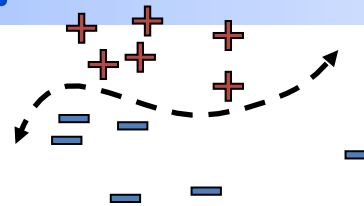
Feature Extraction

- Image is reduced to a set of low-dimensional feature vectors



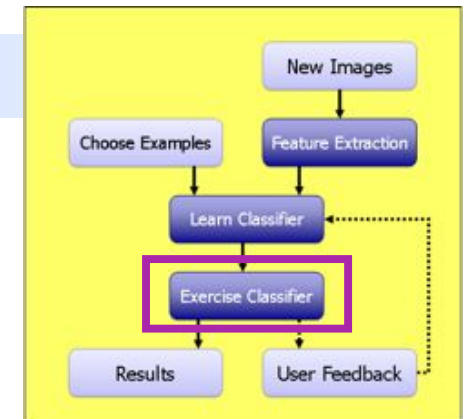
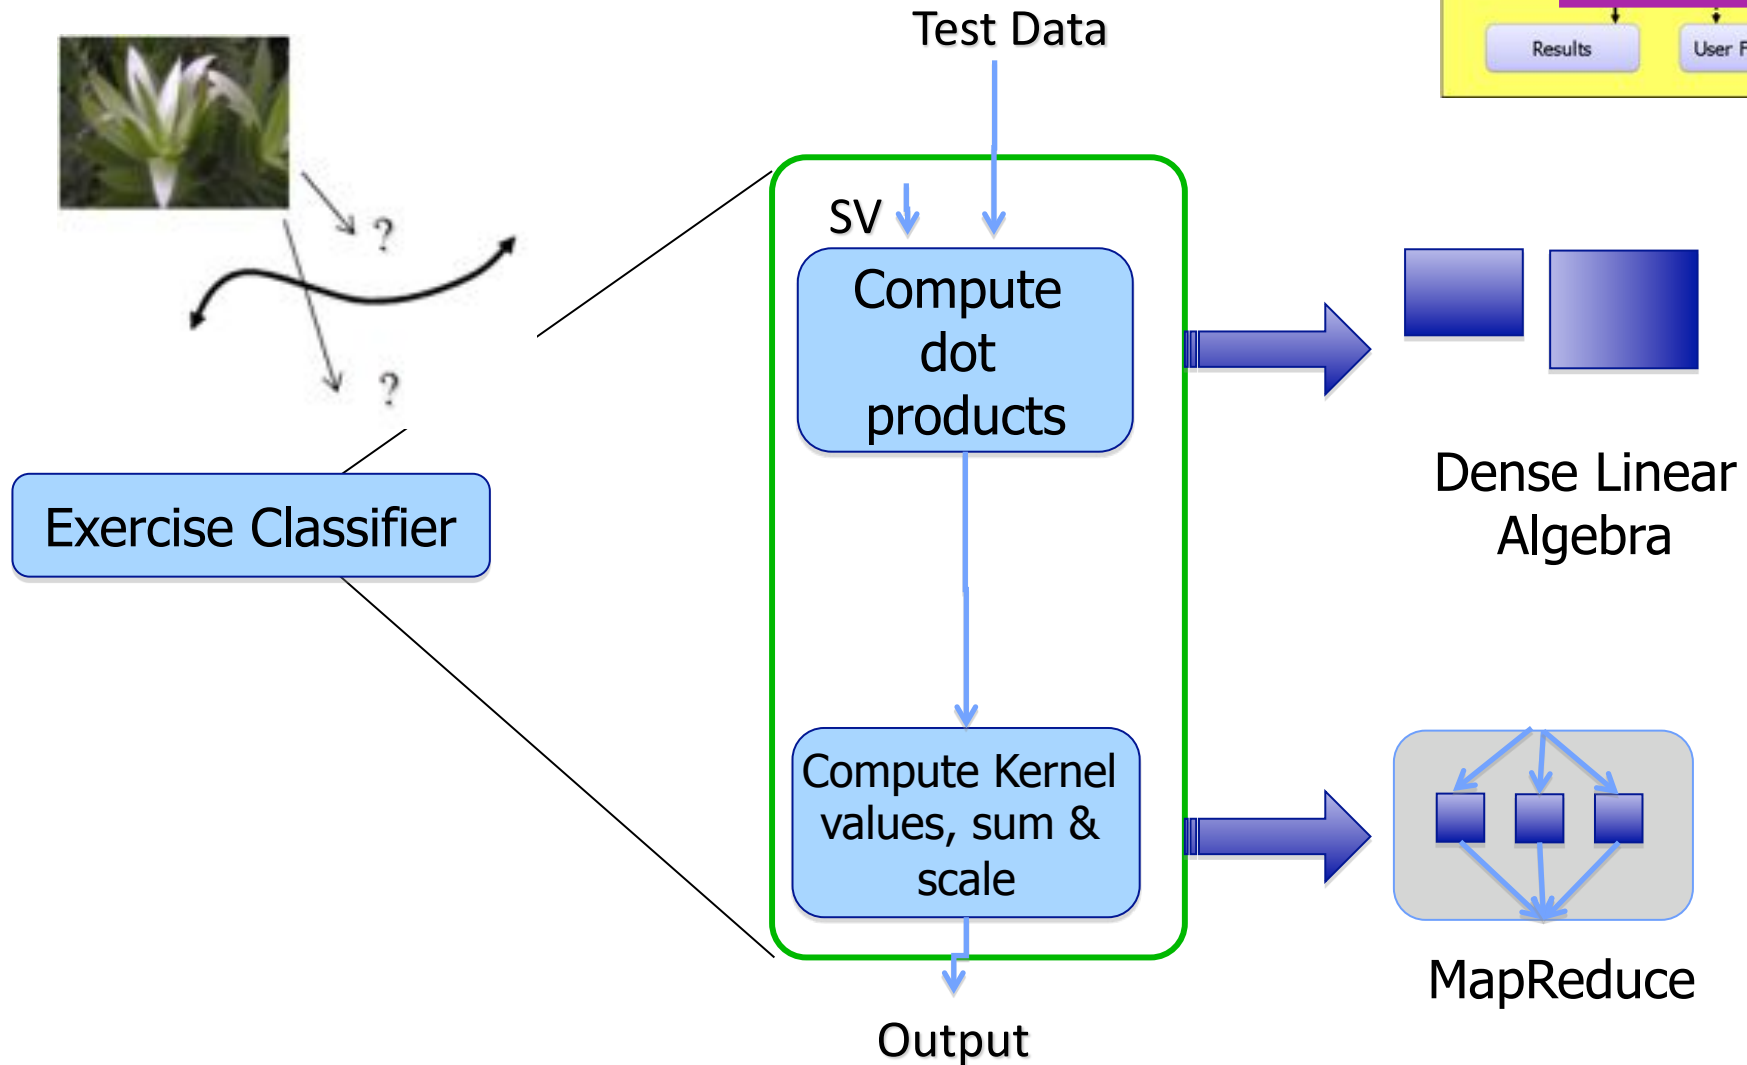


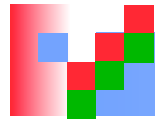
Train Classifier: SVM Training



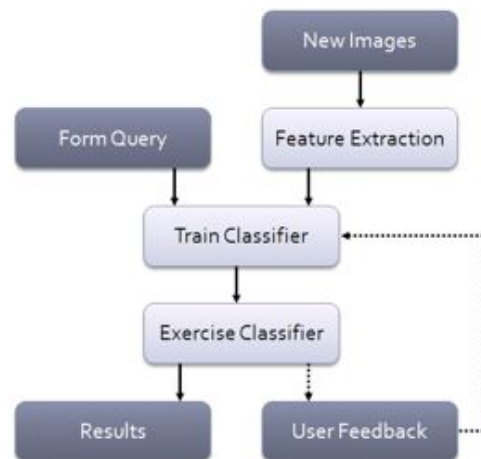
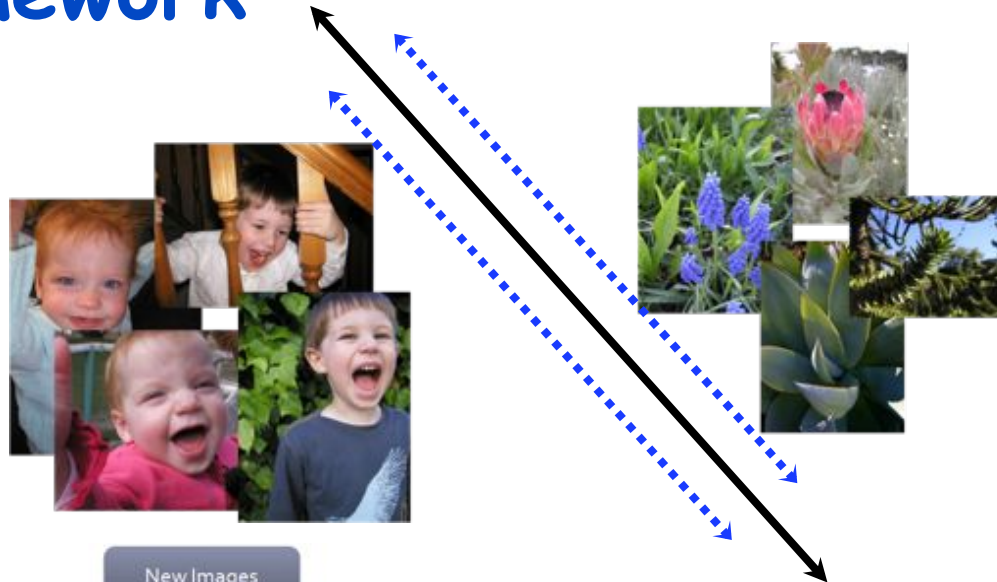
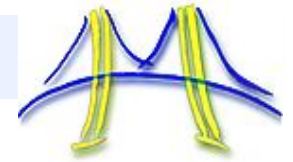
Iterator Pattern

Exercise Classifier : SVM Classification



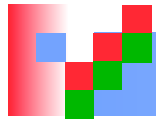


Support-Vector Machine Mini-Framework

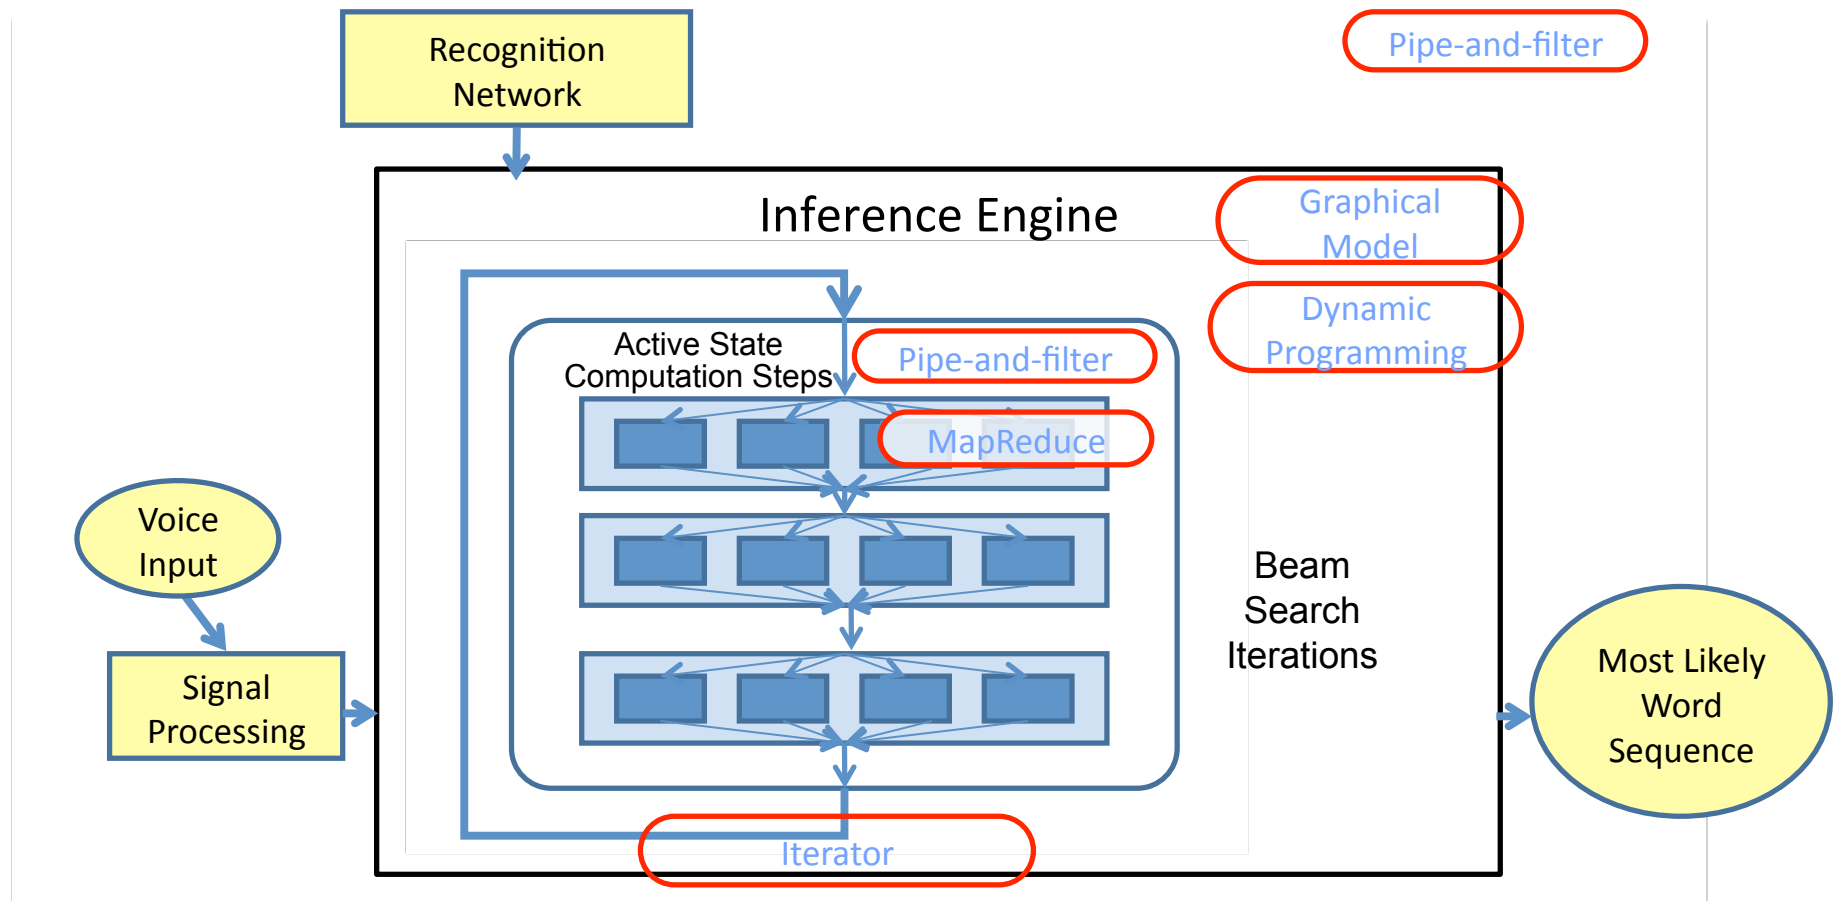
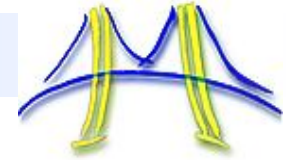


- Support-Vector Machine Framework used to achieve:
 - 9-35x speedup for training
 - 81-138x for classification
- 340 downloads since release

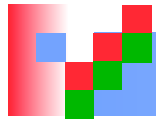
Fast support vector machine training and classification ,
Catanzaro, Sundaram, Keutzer, International Conference on
61 Machine Learning 2008



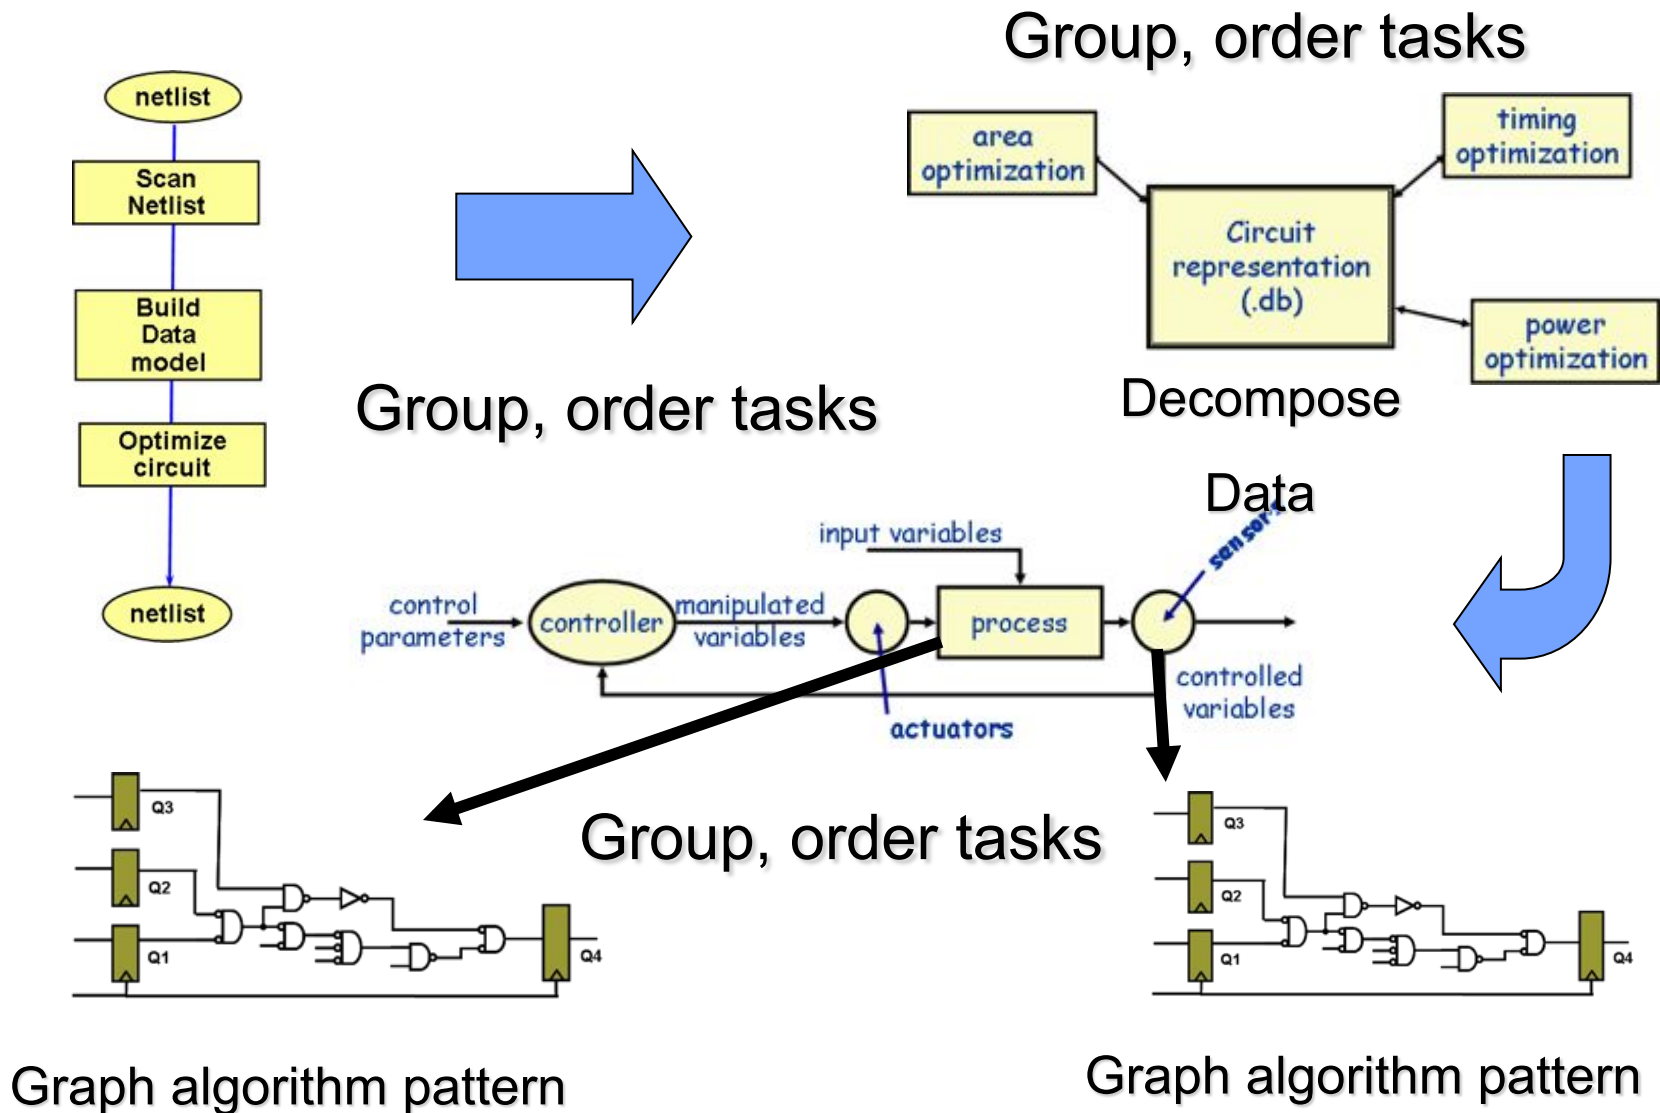
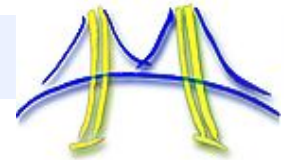
Architecting Speech Recognition

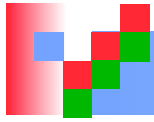


Large Vocabulary Continuous Speech Recognition Poster: Chong, Yi
Work also to appear at Emerging Applications for Manycore Architecture

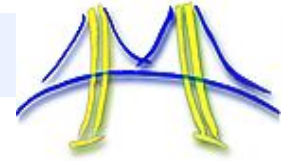


Architecture of Logic Optimization

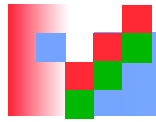




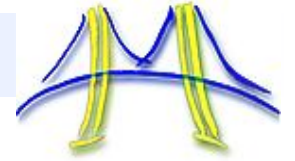
The take away



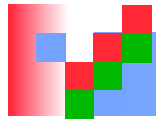
- My own experience has shown that a sound software architecture is the greatest single indicator of a software project's success.
- Software must be **architected** to achieve productivity, efficiency, and correctness
- SW architecture >> programming environments
 - >> programming languages
 - >> compilers and debuggers
 - (>>hardware architecture)
- Key to **architecture** (software or otherwise) is **design patterns** and a **pattern language**
- At the highest level our pattern language has:
 - Eight structural patterns
 - Thirteen computational patterns
- Yes, we really believe arbitrarily complex parallel software can built just from these!



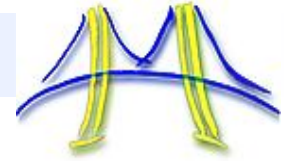
What's next ...



- Friday: Computational Patterns of Parallel Programming (James Demmel (UCB))
(8:45 - 10:45am)

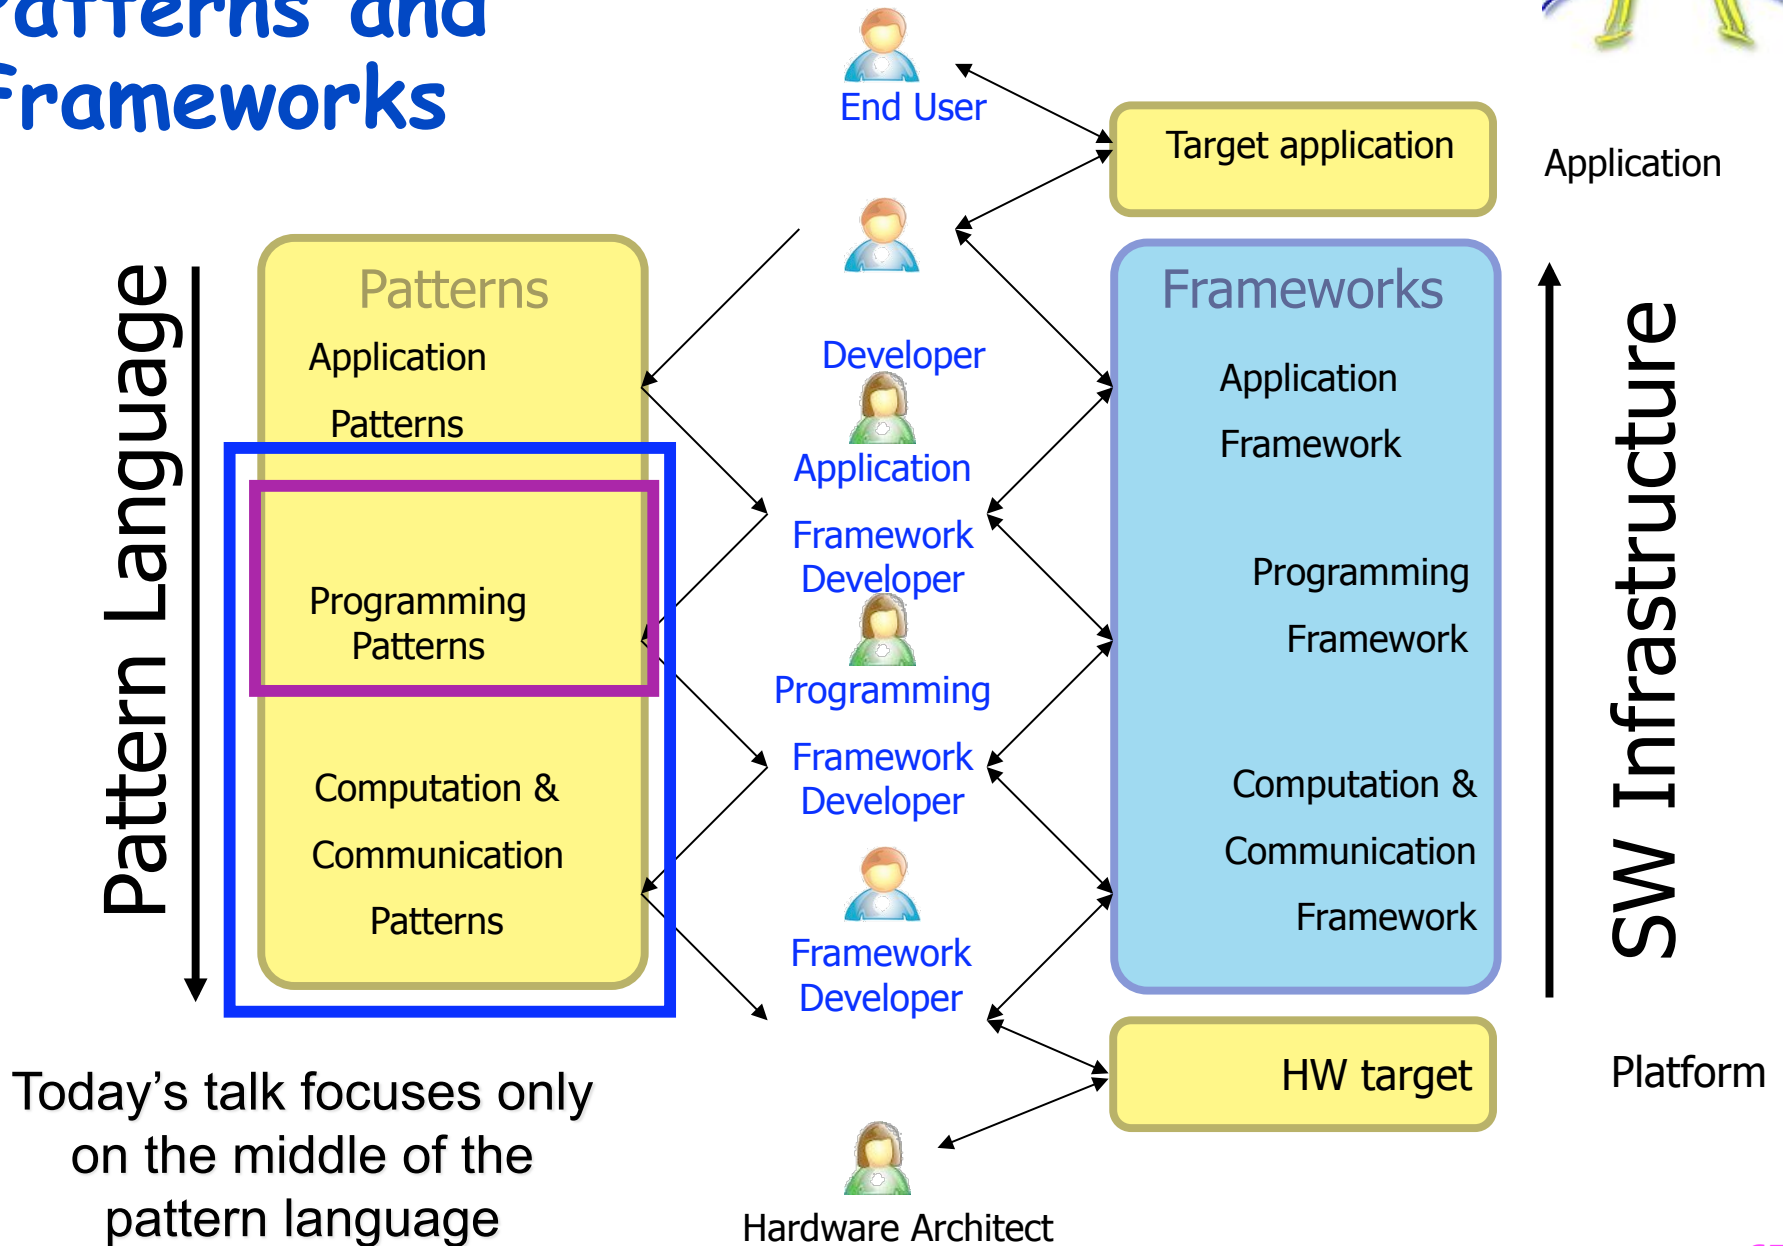
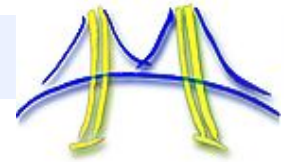


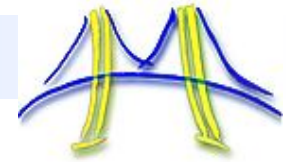
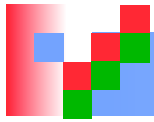
Extras



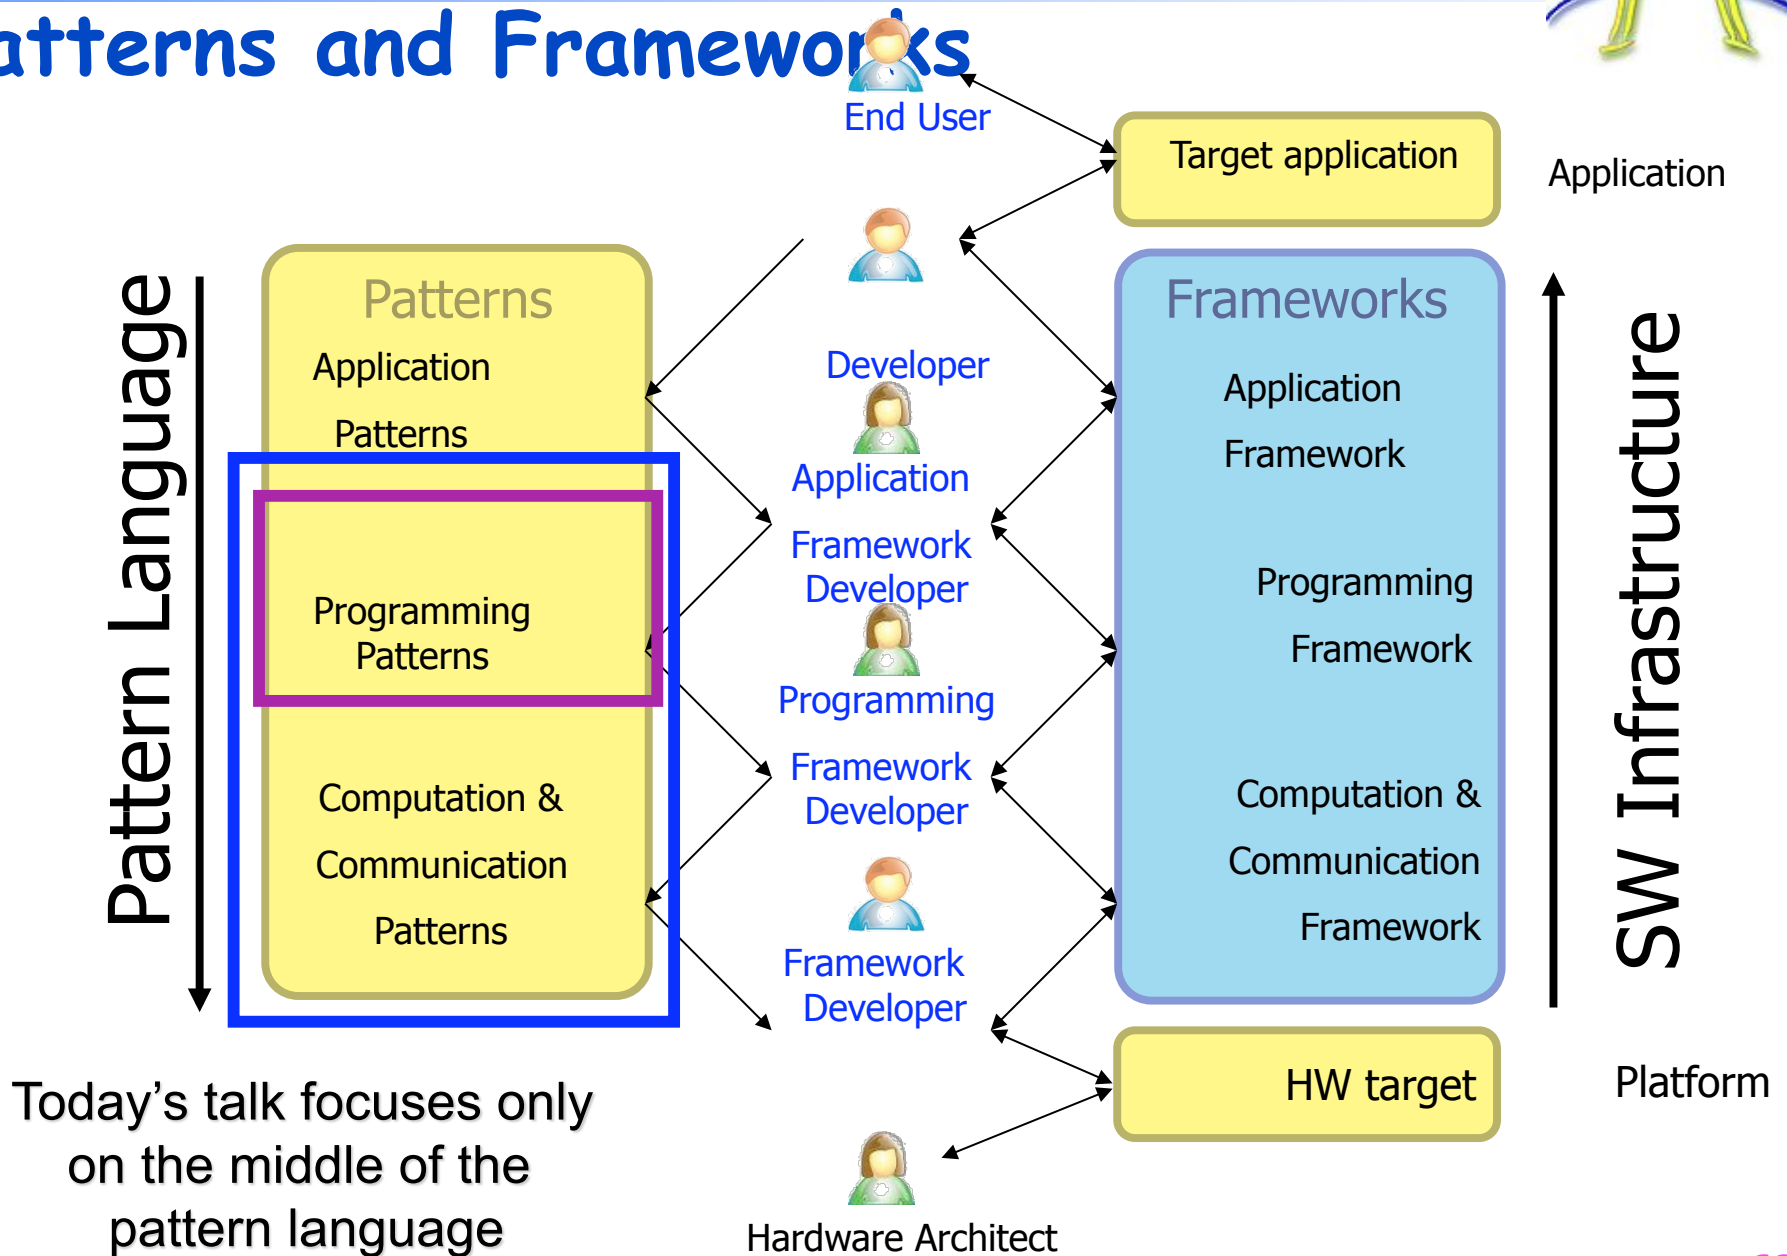


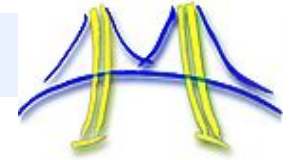
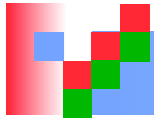
Patterns and Frameworks





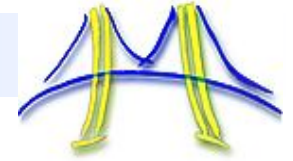
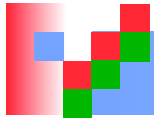
Patterns and Frameworks





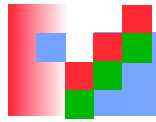
Elements of a Pattern - 1

- Name
 - It must have a meaningful name useful to remember the pattern and when it is used.
- Problem
 - A statement of the problem ... a one-line preamble and the problem stated as a question.
- Context
 - The conditions under which the problem occurs. Defines when the pattern is applicable and the configuration of the system before the pattern is applied.
- Forces
 - A description of the relevant *forces* and constraints and how they interact/conflict with one another and with goals we wish to achieve. Defines the tension that characterizes a problem.

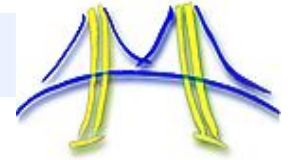


Elements of a Pattern - 2

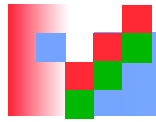
- Solution
 - Instructions used to solve the problem. When done right, it resolves the tension defined in the forces section; flowing from the context and forces. We also define the new context for the system following application of the pattern.
- Invariant
 - What must be invariantly true for this pattern to work. May be stated in the form of Precondition, Invariant, Post-condition
- Examples
 - Examples to help the reader understand the pattern.
- Known Uses and frameworks
 - Cases where the pattern was used; preferably with literature references.
- Related Patterns
 - How does this pattern fit-in or work-with the other patterns in the pattern language.



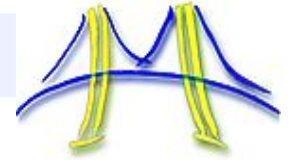
Definitions - 1



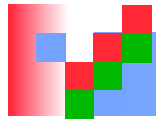
- **Design Patterns:** "Each design pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice." *Page x, A Pattern Language, Christopher Alexander*
- **Structural patterns:** design patterns that provide solutions to problems associated with the development of program structure
- **Computational patterns:** design patterns that provide solutions to recurrent computational problems



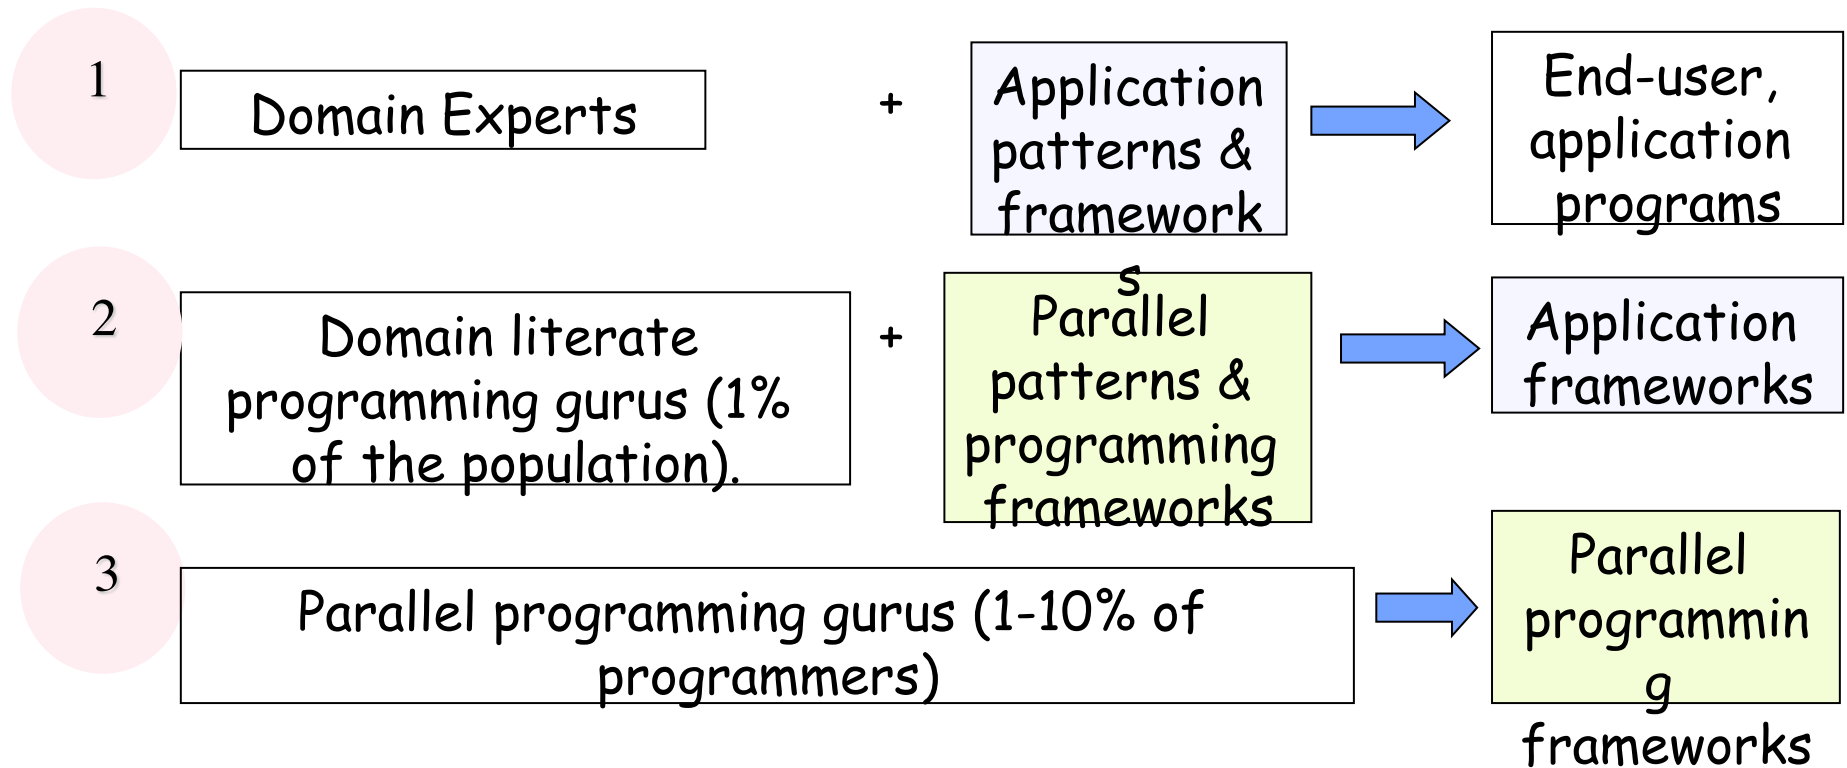
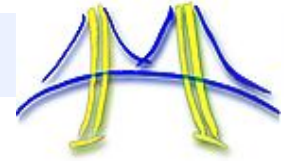
Definitions - 2



- **Library**: The software implementation of a computational pattern (e.g. BLAS) or a particular sub-problem (e.g. matrix multiply)
- **Framework**: An extensible software environment (e.g. Ruby on Rails) organized around a structural pattern (e.g. model-view-controller) that allows for programmer customization *only* in harmony with the structural pattern
- **Domain specific language**: A programming language (e.g. Matlab) that provides language constructs that particularly support a particular application domain. The language may also supply library support for common computations in that domain (e.g. BLAS). If the language is restricted to maintain fidelity to a structure and provides library support for common computations then it encompasses a framework (e.g. NPClick).

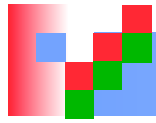


Eventually

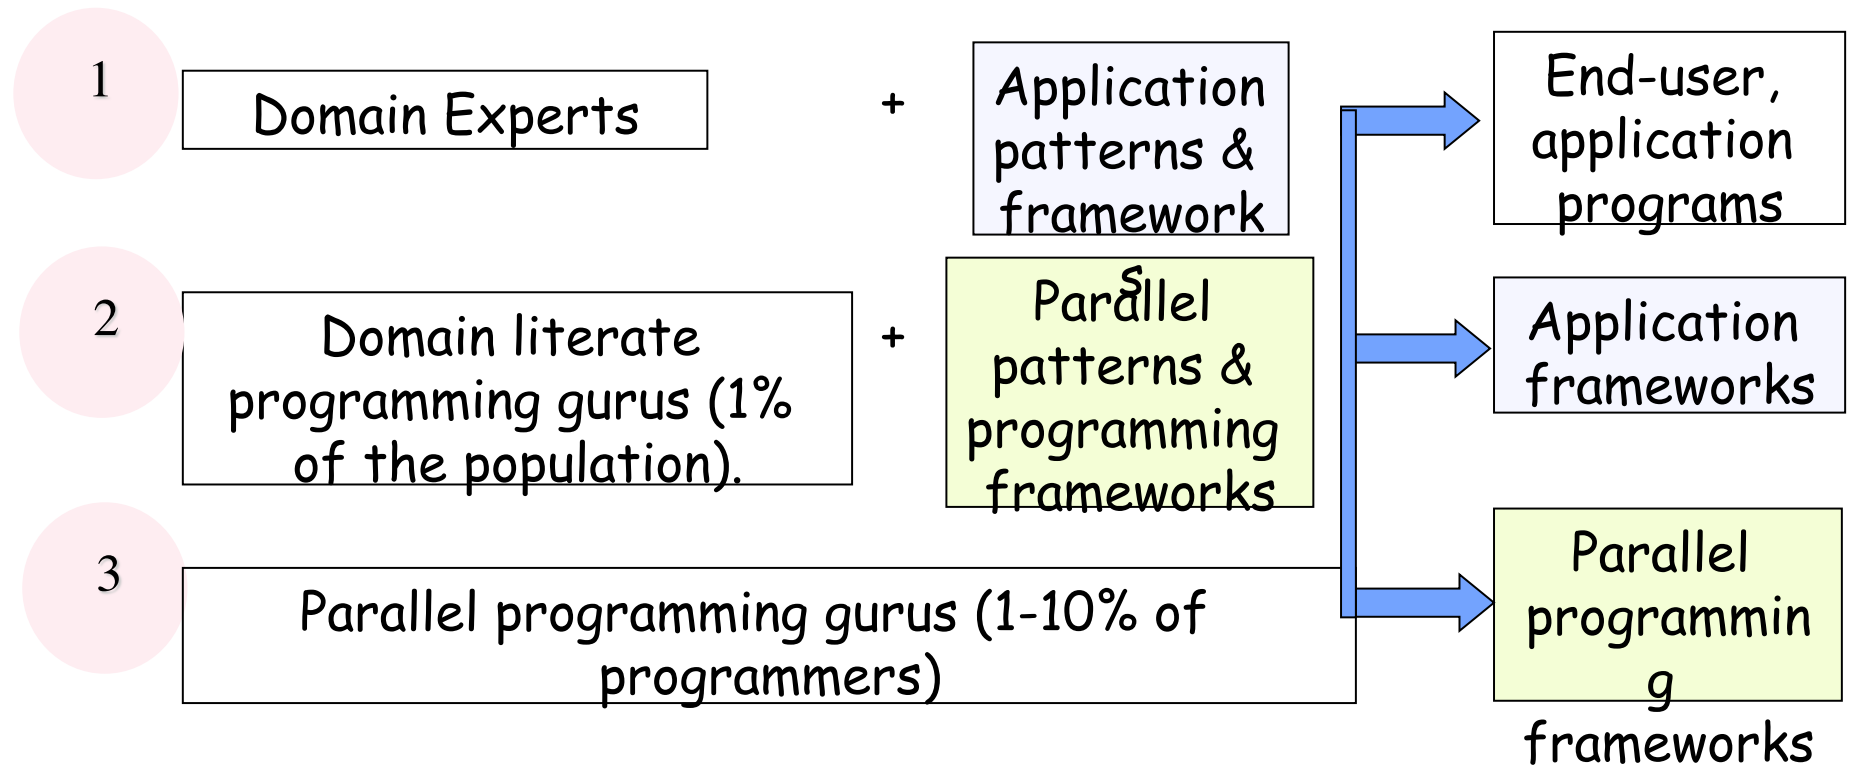
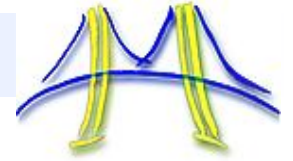


The hope is for Domain Experts to create parallel code with little or no understanding of parallel programming.

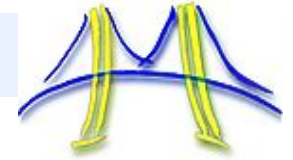
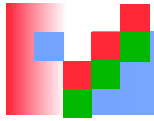
Leave hardcore “bare metal” efficiency layer programming to the parallel programming experts









Today



- For the foreseeable future, domain experts, application framework builders, and parallel programming gurus will all need to learn the entire stack.
- That's why you all need to be here today!



People, Patterns, and Frameworks

	  Design Patterns	  Frameworks
Application Developer 	Uses application design patterns (e.g. feature extraction) to architect the application	Uses application frameworks (e.g. CBIR) to develop application
Application-Framework Developer 	Uses programming design patterns (e.g. Map/Reduce) to architect the application framework	Uses programming design patterns (e.g. MapReduce) to develop the programming framework