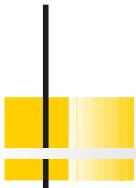


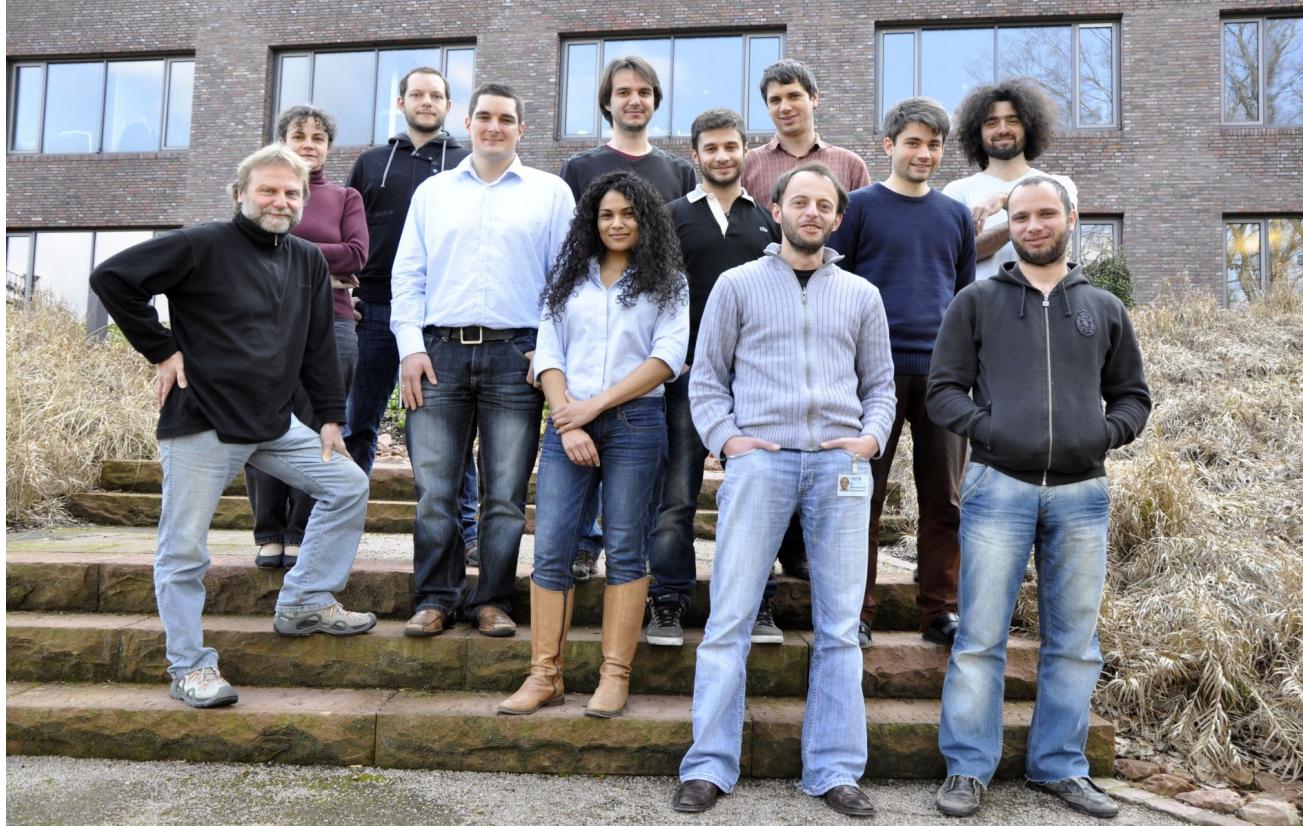
# High Performance Phylogenetics & Population Genetics



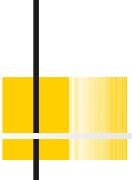
**Alexandros Stamatakis**

**Scientific Computing Group  
Heidelberg Institute for Theoretical Studies**

**[Alejandro.Stamatakis@h-its.org](mailto:Alejandro.Stamatakis@h-its.org)  
[www.exelixis-lab.org](http://www.exelixis-lab.org)**

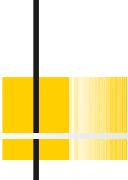


- IT
  - Cluster with 2500 cores
  - Archiving of Scientific Data
  - 90 TB parallel I/O system
  - All other IT services
- Research
  - Method development
  - SW development
  - Emerging Parallel Architectures



# Please ask Questions!

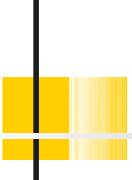
---



# Why is this important?

---

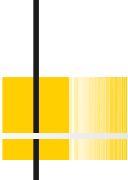
- Datasets are growing
- Need to use supercomputers
- Transformation into a computational science
- Software/methods are widely used when they are fast
- Software/algoritm engineering is required



# However

---

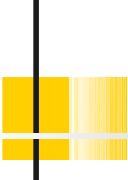
- Smarter algorithms are always better than brute-force HPC approaches



# Toward a Computational Science

---

- **Example** 1Kite Project [www.1kite.org](http://www.1kite.org) to sequence 1000 insect transcriptomes
- **Challenges**
  - Data transfer
  - Data storage
  - Code correctness
  - Supercomputer resources
  - CO<sub>2</sub> footprints
- **Software Properties**
  - Checkpointing
  - Scalability → Exascale/ExaFlop systems
  - Reduce memory footprints: 1500 taxa & 20,000,000 sites → 1TB
  - Low-level optimizations: 256-bit AVX and FMA intrinsics
  - Load Balance Issues



# Toward a Computational Science

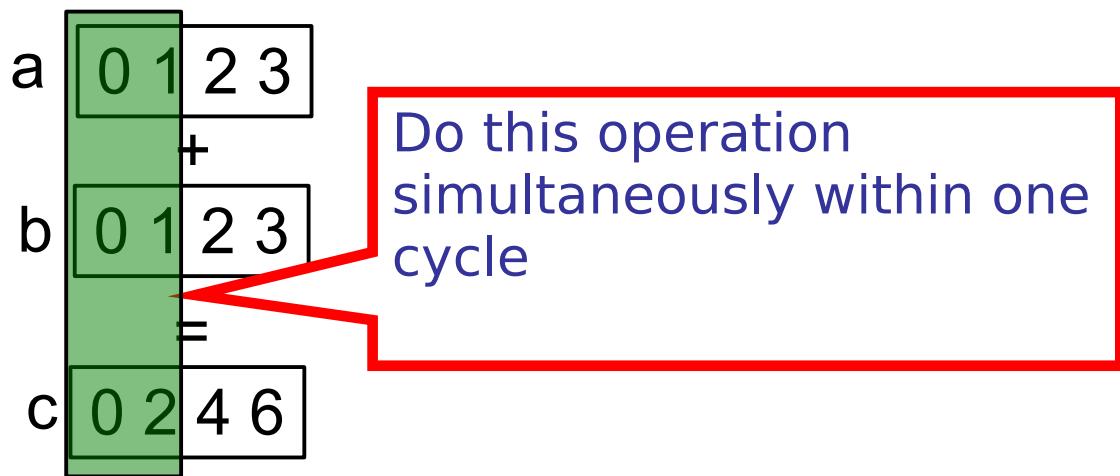
---

- **Example** 1Kite Project [www.1kite.org](http://www.1kite.org) to sequence 1000 insect transcriptomes
- **Challenges**
  - Data transfer
  - Data storage
  - Code correctness
  - Supercomputer resources
  - CO<sub>2</sub> footprints
- **Software Properties**
  - Checkpointing
  - Scalability → Exascale/Exaflop systems
  - Reduce memory footprints: 2500 taxa & 20,000,000 sites → 1TB
  - Low-level optimizations: 256-bit AVX and FMA intrinsics
  - Load Balance Issues

Who knows what this is?

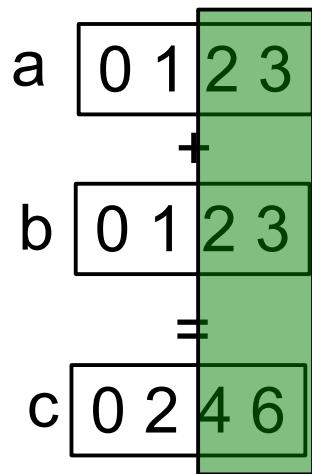
# Vector Instructions

- e.g. RAxML SSE3/AVX versions
- A clock tick: execute one instruction
- 2.2 GHz  $2.2 * 10^9$  instructions per second

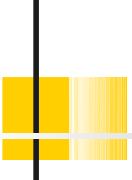


# Vector Instructions

- e.g. RAxML SSE3/AVX versions
- A clock tick: execute one instruction
- 2.2 GHz  $2.2 * 10^9$  instructions per second



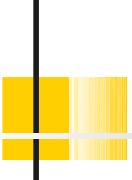
... and this operation  
simultaneously within one  
cycle: only two clock  
cycles (ticks) required



# The algorithmic problem

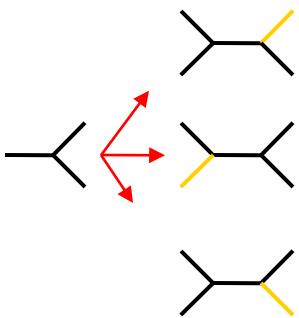
---

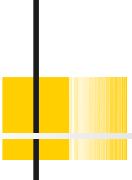




# The number of trees

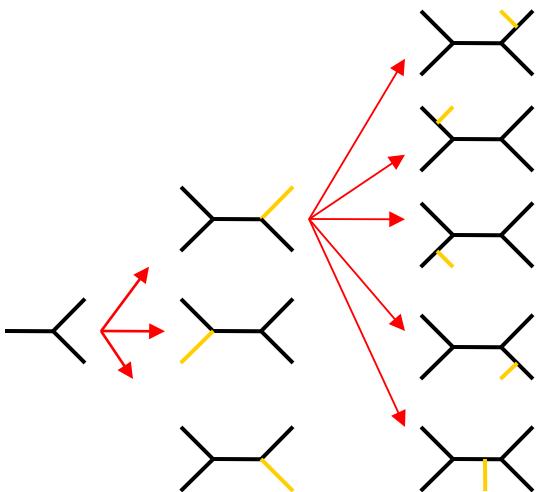
---

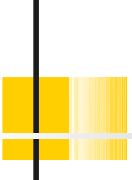




# The number of trees

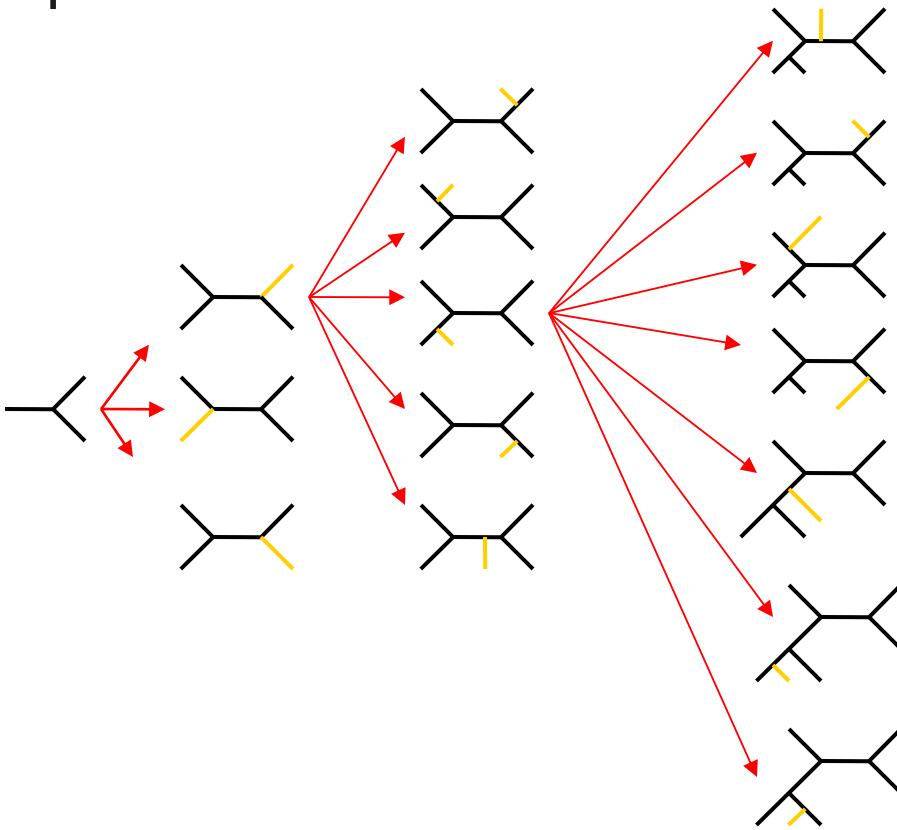
---



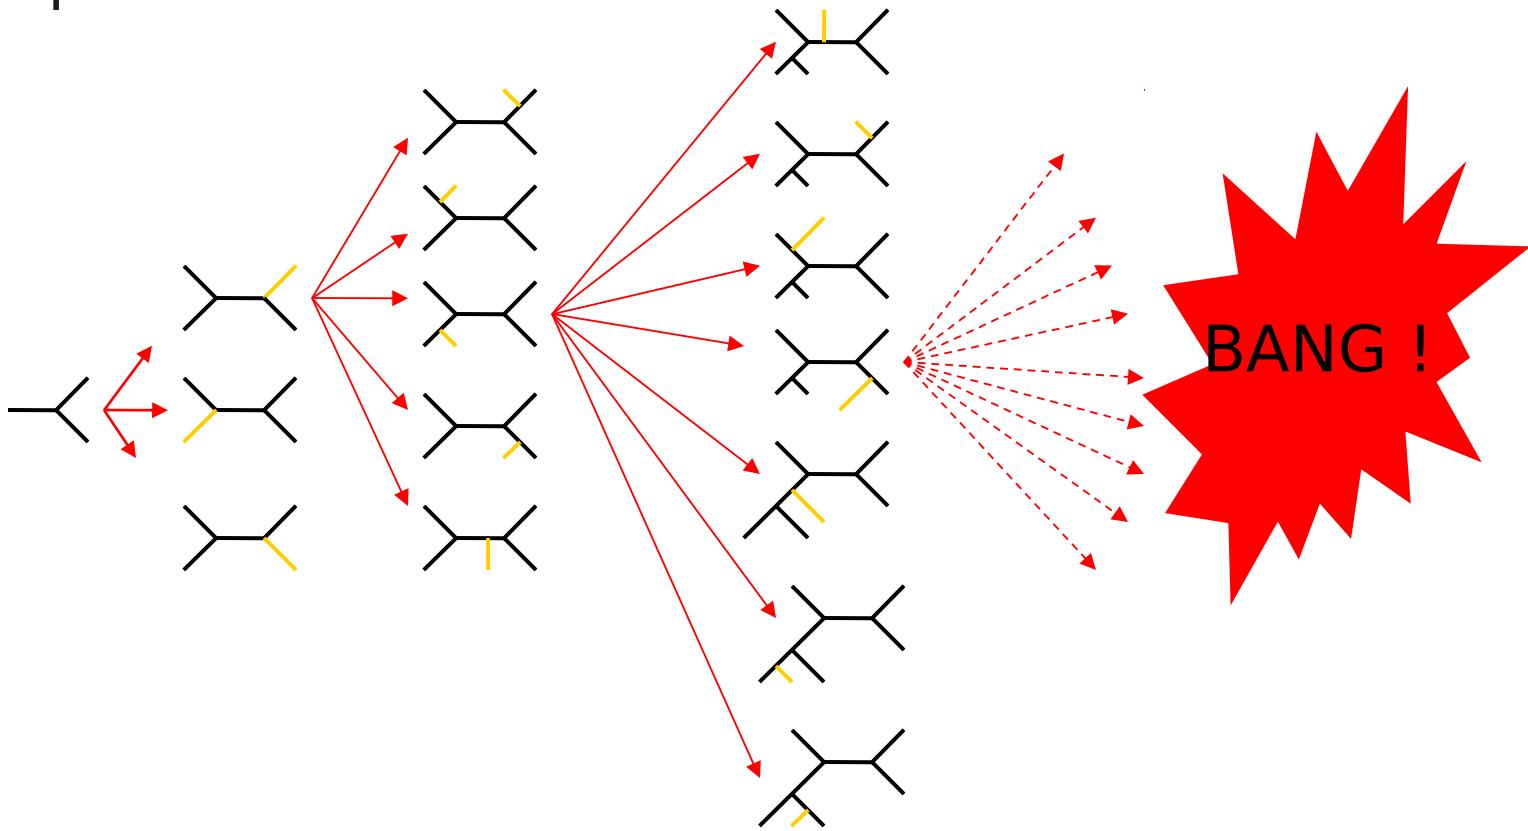


# The number of trees

---



# The number of trees explodes!



# # trees with 2000 tips

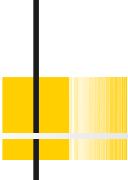


```
stamatatak@exelixis:~/Desktop/GIT/TreeCounter$ ./treeCounter -n 2000
```

GNU GPL tree number calculator released June 2011 by Alexandros Stamatakis

Number of unrooted binary trees for 2000 taxa: 30049638174211656151632910065681814981377232074237013089504954043012636525258308210827685996688247000464352735214265634288295  
8915023446000631493969130632970436056184861877465482277991223536809233455563199910834597693126756525012899867433187752811401960991631522367030609121735709762379847705467667  
7795324797182614385273338226727784250737252849916669687584403510579587020686505817687044666318123742901021438506432471360934491667021135969756940300666252646479269124551031  
01661366319554282281877625114848758254581227914289801132648902674033761294712745767036267579086843169660718609847941818865957214557044744572288661729053583520744253688123124  
2822975974655377102257672676842858108772249501062181173405232082653973429622735253659051586563138327203111984198746759973864631829032038325230859799792216101227215780805  
2481458312068440167606239306009716167297155047284877996343375313489942303724373478791319890859537640701348494461138775275695240870246172010874297380462275052545706689372  
31941820644070689188400387059028977197516454495975821662130620506461776169948566373416818358498932907699338207680105243728461492403422961155182609778228619192672071295189  
8936009951309742303721316382518428110305017144115688403051318658775443763085003114511107238370397046518223204040615470827307862995754933103127520861670066791298014262230  
0565123522718063819509335872651728623589020520016144361756075654286471422126613004438070840675015892476731663415395405750704474994909831496473031080411401891849735912811228  
3787740498843406561202405664244638600938996508574296519472699543015281237526510965815284699770367921112903568098108791695879516141592810495281798558472925344478644244359  
9880531532704796819694659917686145337018195829587715748245594337724369582576242663016946320482495182255939287403177623433881048604630975191556923871767513095213415098816  
7154643078623526062377864068386804246902527491139319276802611515990582603886733172930713673903403618637463980605764836474670274446727880885337074254421922726677747003329403  
32010328805311268902625518309679194835867892937016376817530482063389438714979311512353698229625111630714829459921162080302684762013356904410896814543615095155877581167  
9770012563912151116237444170497371704604029481104114228646613191882197957138336835207252605520276982397461321849524926489705079039836025625560628985288839561357874156576  
64889992608732866126306425432602489792291135600716405739845163752452433769437558573847255455643975996043755914640112221144755235573176239973057747183956531217416532295986675  
90129411612392407220932503696731248844915537592160560565015416720774159236240868667675348286512964888739059707578802473393463740848115901163977279747480417316268700916728735  
61216422684681606831989598120367346851531278161689587215123123308760034733810972531184233394039093737839506683557873530788635864640563299494906311874240290972779272693  
0303224453775957972248734568159114585570783850514681667674258113019580636321907500790250310882099727174813643698947397107932777706067301736017566538739726033777173008441  
343940512366905554949324861650825399577950363267049478442934988531727973481777971465671751511788763964306933245807634611073421432819504990968087402739768891470451747205543  
8969396667426017742189469321290453317331882677319465354113302100866575081713240324675804891862366364160795593740205163951540949806486242309796497211696659615808041  
78839922322642849894235527692391124360767508997717896834593378907243465354571933065615359902779162656386361861974048590938234620325345979733137213659343717585590664439  
64613283001136726019340687064423394891992153043852815416596301185494236348635245857466428346090629162795649265847230036085559898976191629324814009459248989468468862322586  
017055146890564983972760039274870695463788171469942904910728531804107765160072633897322700138435995973026572179804324664312359758553721796950143921010226596324  
7887839774053331315176278152880718652603176327264828944993805305849799570082893798080149029010052938984722799471678042168942415911828425769646478650731  
5325178302336307298251692210346584265894447464916123854689718507968172990932182834111848213847677283165486532123173820041319905105189670220188704958678180509590730360  
6930402937216038968176555872655382318093705826257083898387409098468656634271397500013291835105943321729879825243707508272087959859437157667660155782699660343197752623308  
89896258780062800956094416932173945544103369586261556256010669390302038789709836737860756641433585106111658314520424513208508589994932364816869711949516716195672  
70709697388958855795624666415365617235493018073940476052980172177139168678800027785196617307006128451730758250373564312065112443730825229625040453160590741343881872563  
4779138306650993188022531008340176840261401539616989192075147108033757708849740141834599753972059878682064879116064969858177601153972058492269897018134943269180182173  
31880636539108936898117148913574566805428074851701758526663963537018934494832669782385092657922201746372190723119641751489944010079638760178267410701994547321888732742  
8088967424371574132406009370425130989363053475974827998403132398947216649229042573095836853441621564055729082066224003863237526380910233298978388604237596250165795276  
6950798639881042948323316026721655517812089926467780493574132638713740842388554653833615864345130543962481397279559725995110706314305992615495622958320232708057681156690  
489568610522030573252984721874782136713666586927109487553697485849795108197203387828448647435642009581619303147234956724337957243664289363321244850  
5971992536852924930534625276413785312308943128905152378092556049587099051349458007404773143388800274532321747030695741967114441531273090205  
101004767101435063885795347844725538980154192331702751989618063515268254317319383292589193153016413054897231128666465492971930479296432829556719092881692091042334122007454  
242049900872585046208051104875883059495990311188736668509414882172573457635523396403848131821316740835900691640005326258184783765067804511771732865818989215358309447765  
350341796875

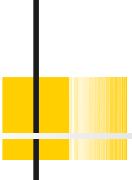
Approximately 3.00 times  $10^{6328}$



# Outline

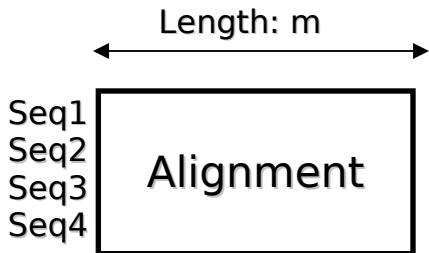
---

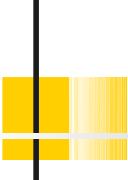
- Computing the Likelihood
- Optimizing & Parallelizing Likelihood Computations
- Saving Memory in Likelihood Computations
- HPC population genetics



# Maximum Likelihood

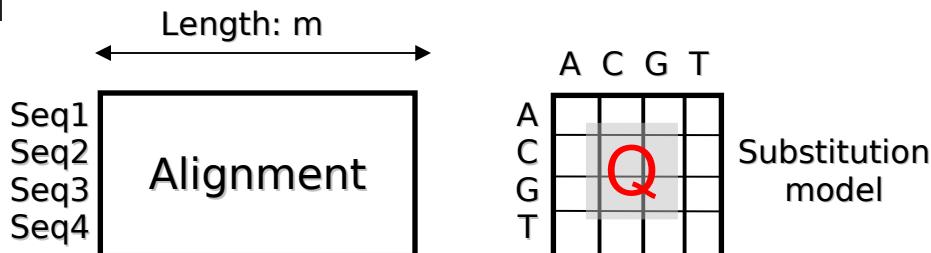
---





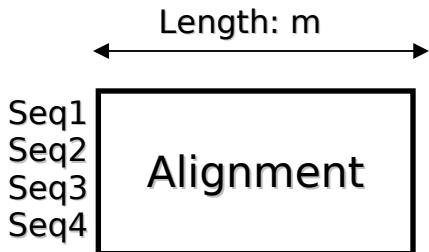
# Maximum Likelihood

---

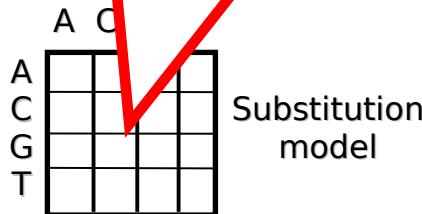




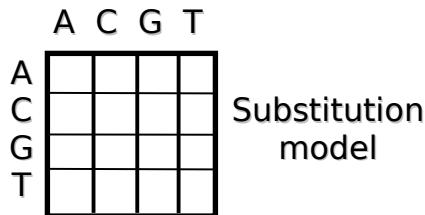
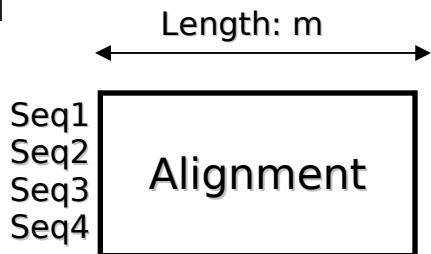
# Maxim



$P(t) = e^{Qt}$   
Important for load  
balance!



# Maximum Likelihood



Prior probabilities,  
Empirical base frequencies

$\pi_A \ \pi_C \ \pi_G \ \pi_T$

# Maximum Likelihood

Length: m

Seq1  
Seq2  
Seq3  
Seq4

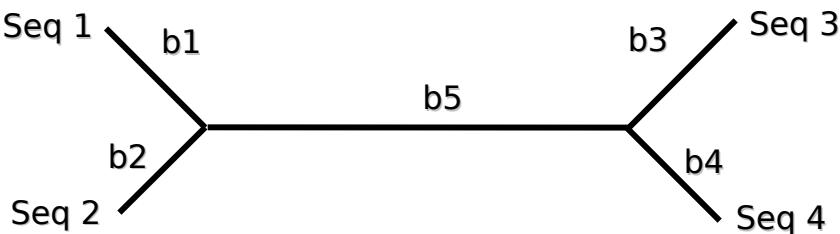
Alignment

	A	C	G	T
A				
C				
G				
T				

Substitution  
model

Prior probabilities,  
Empirical base frequencies

$\pi_A \pi_C \pi_G \pi_T$



# Maximum Likelihood

Length: m

Seq1  
Seq2  
Seq3  
Seq4

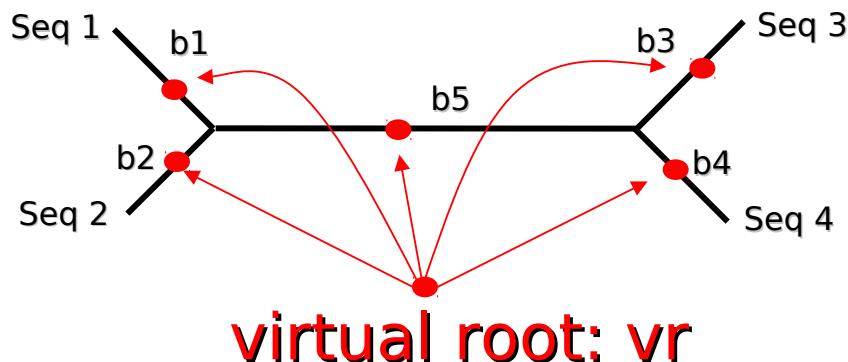
Alignment

	A	C	G	T
A				

Substitution  
model

Prior probabilities,  
Empirical base frequencies

$\pi_A \ \pi_C \ \pi_G \ \pi_T$



# Maximum Likelihood

Length: m

Seq1  
Seq2  
Seq3  
Seq4

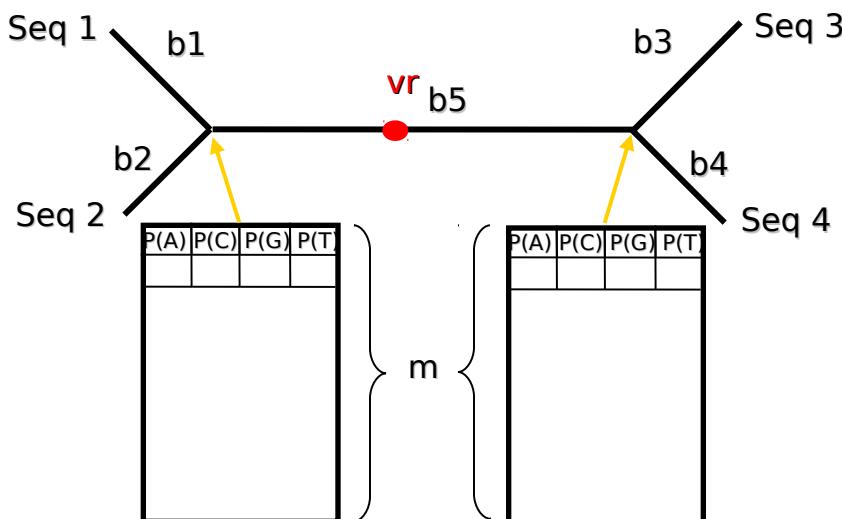
Alignment

	A	C	G	T
A				
C				
G				
T				

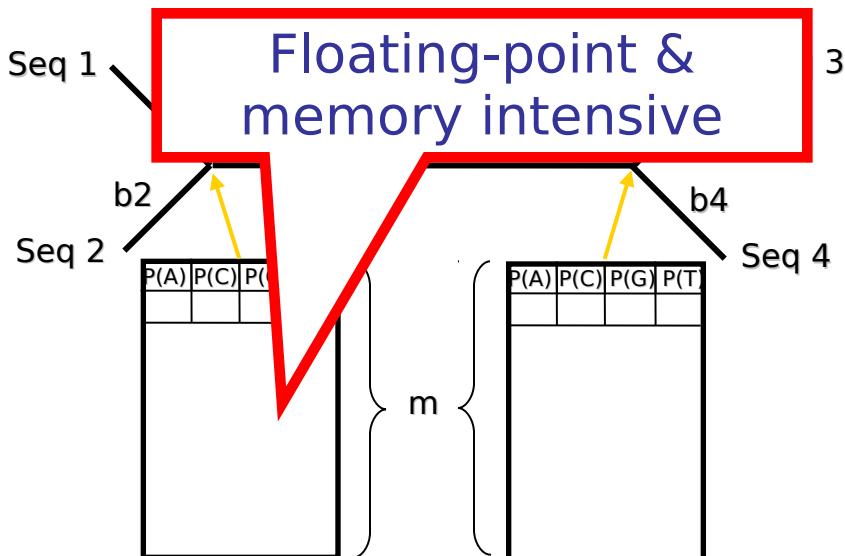
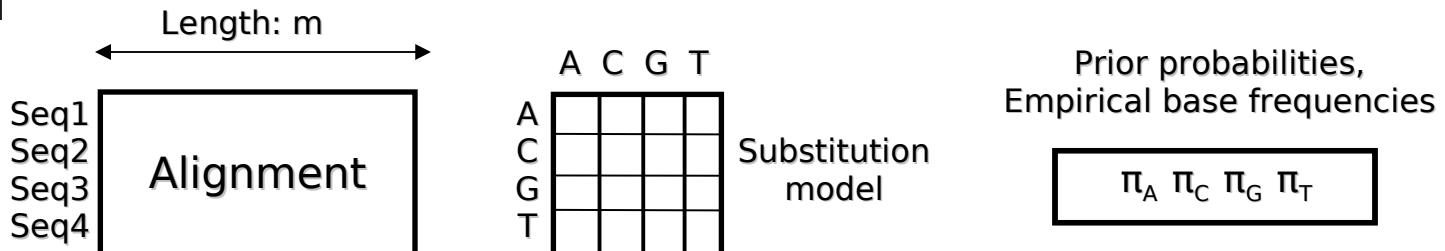
Substitution  
model

Prior probabilities,  
Empirical base frequencies

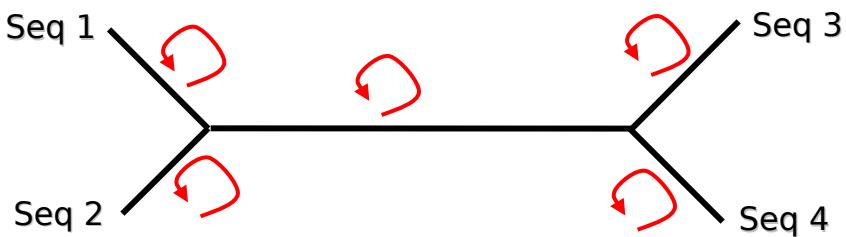
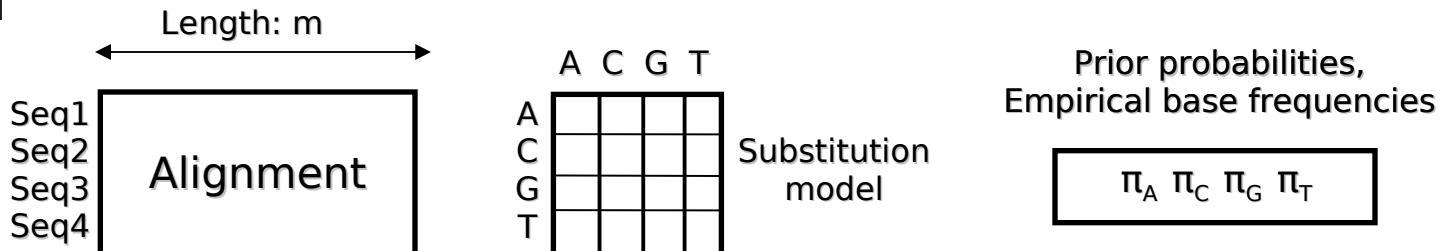
$\pi_A \ \pi_C \ \pi_G \ \pi_T$



# Maximum Likelihood

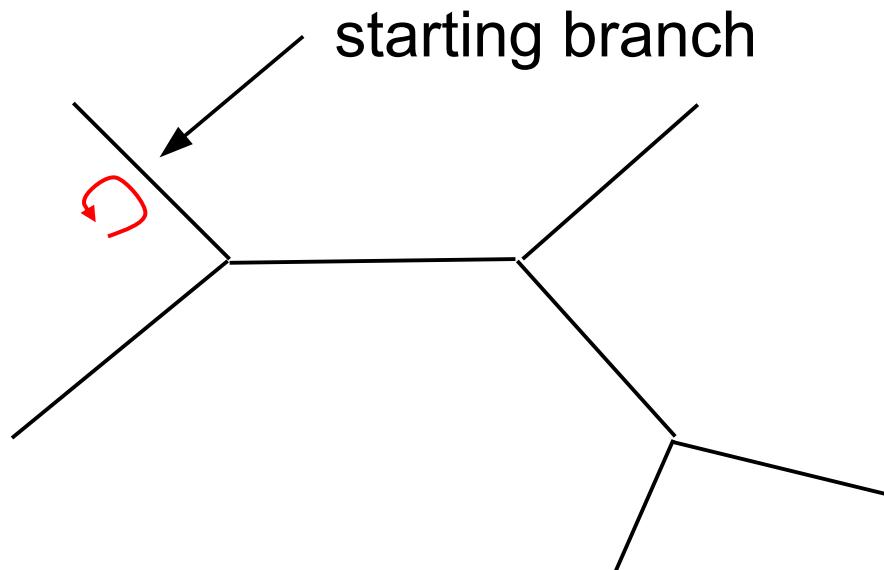
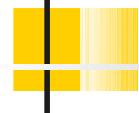


# Maximum Likelihood

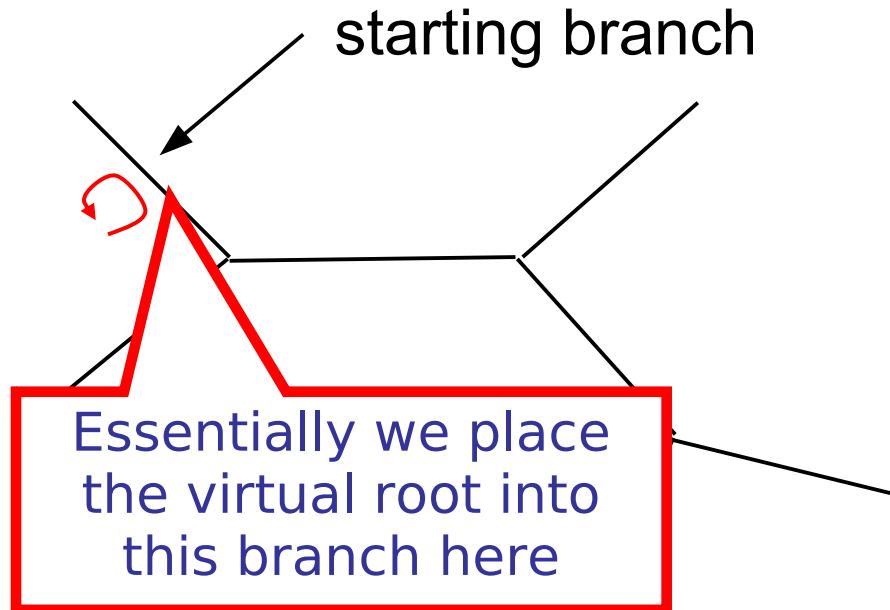


**optimize branch lengths**

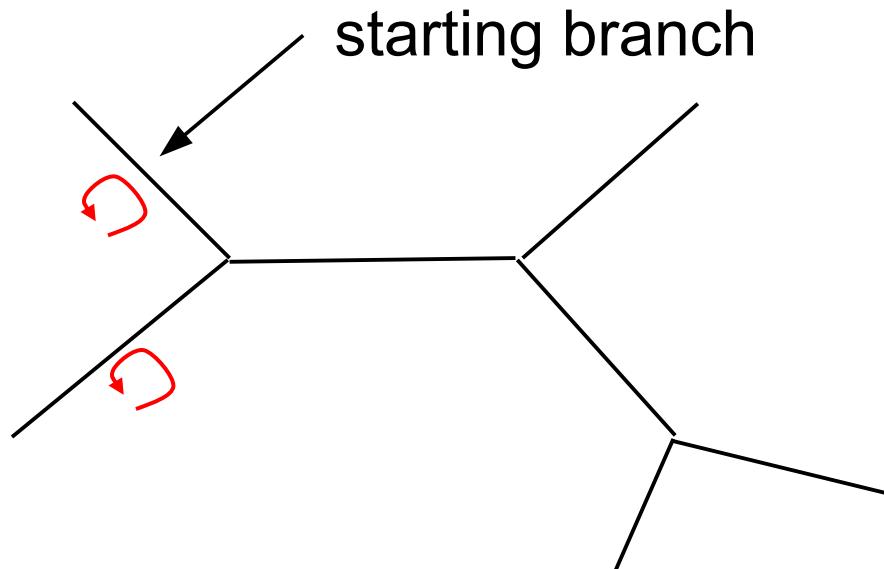
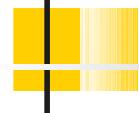
# Branch Length Optimization



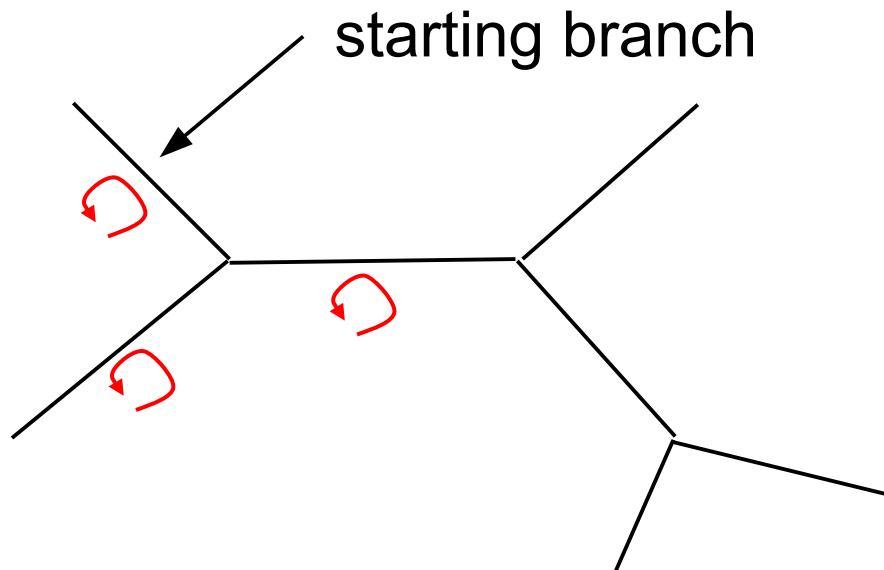
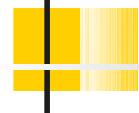
# Branch Length Optimization



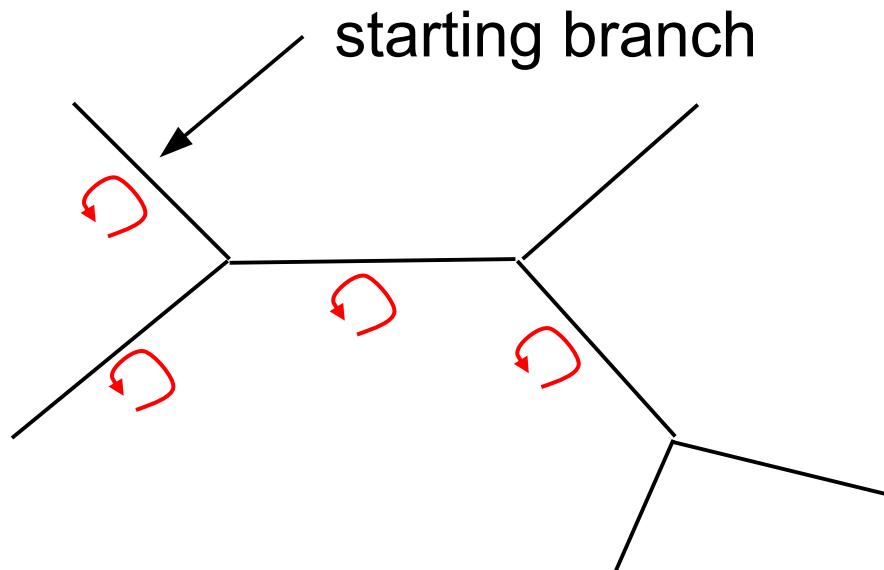
# Branch Length Optimization



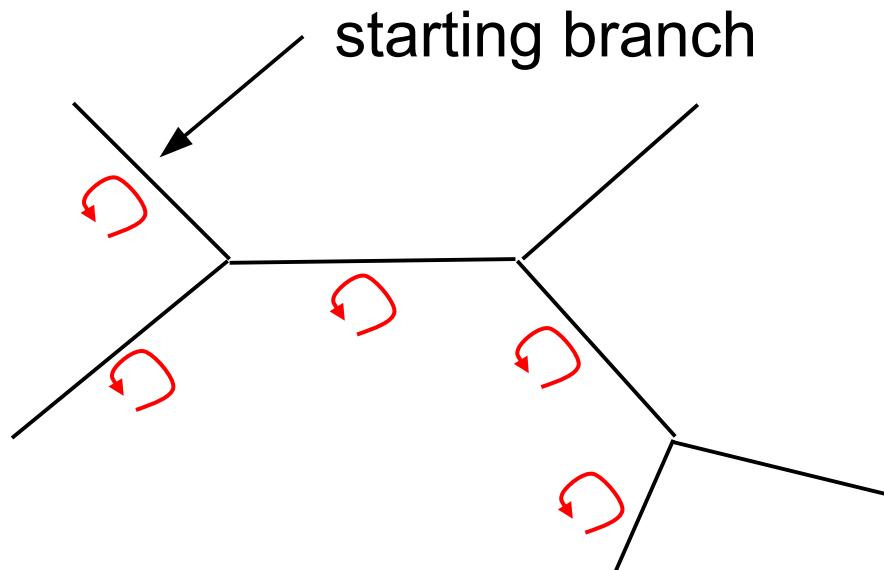
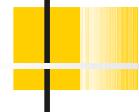
# Branch Length Optimization



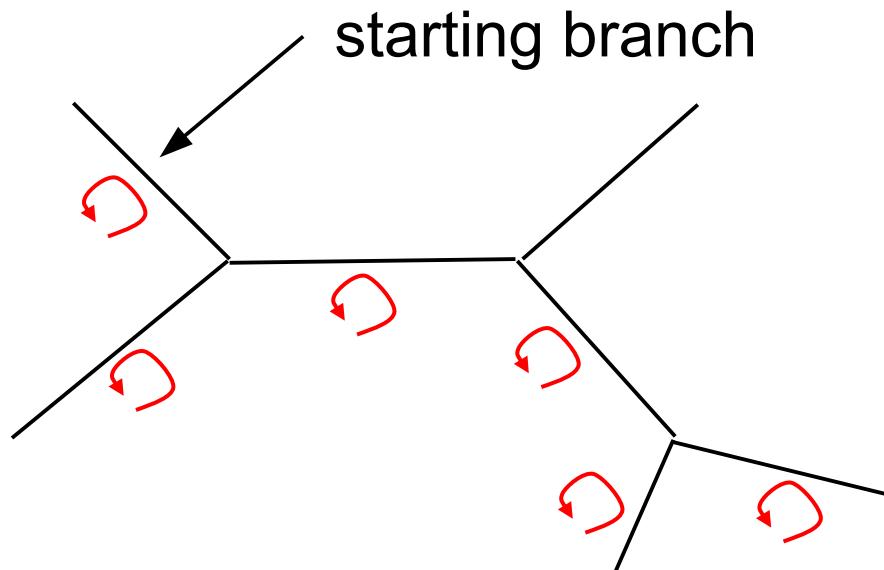
# Branch Length Optimization



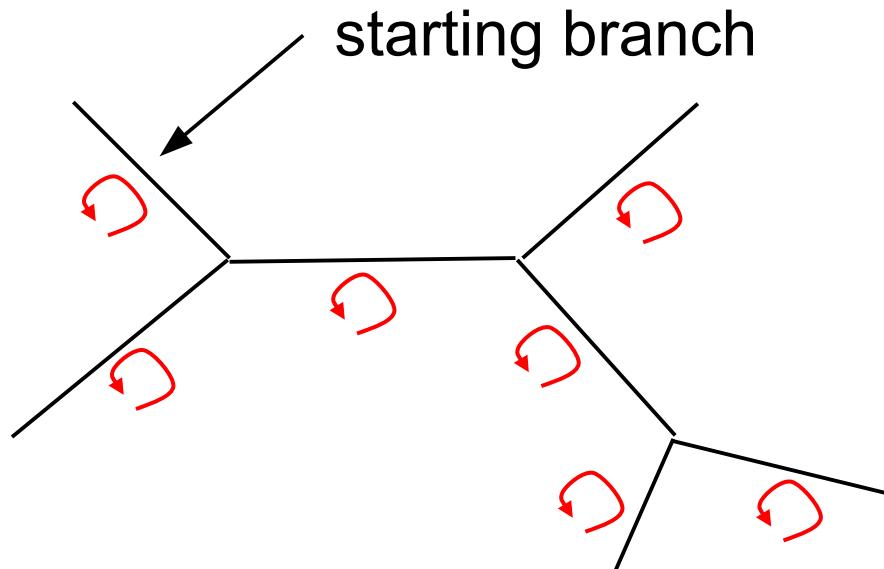
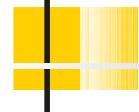
# Branch Length Optimization



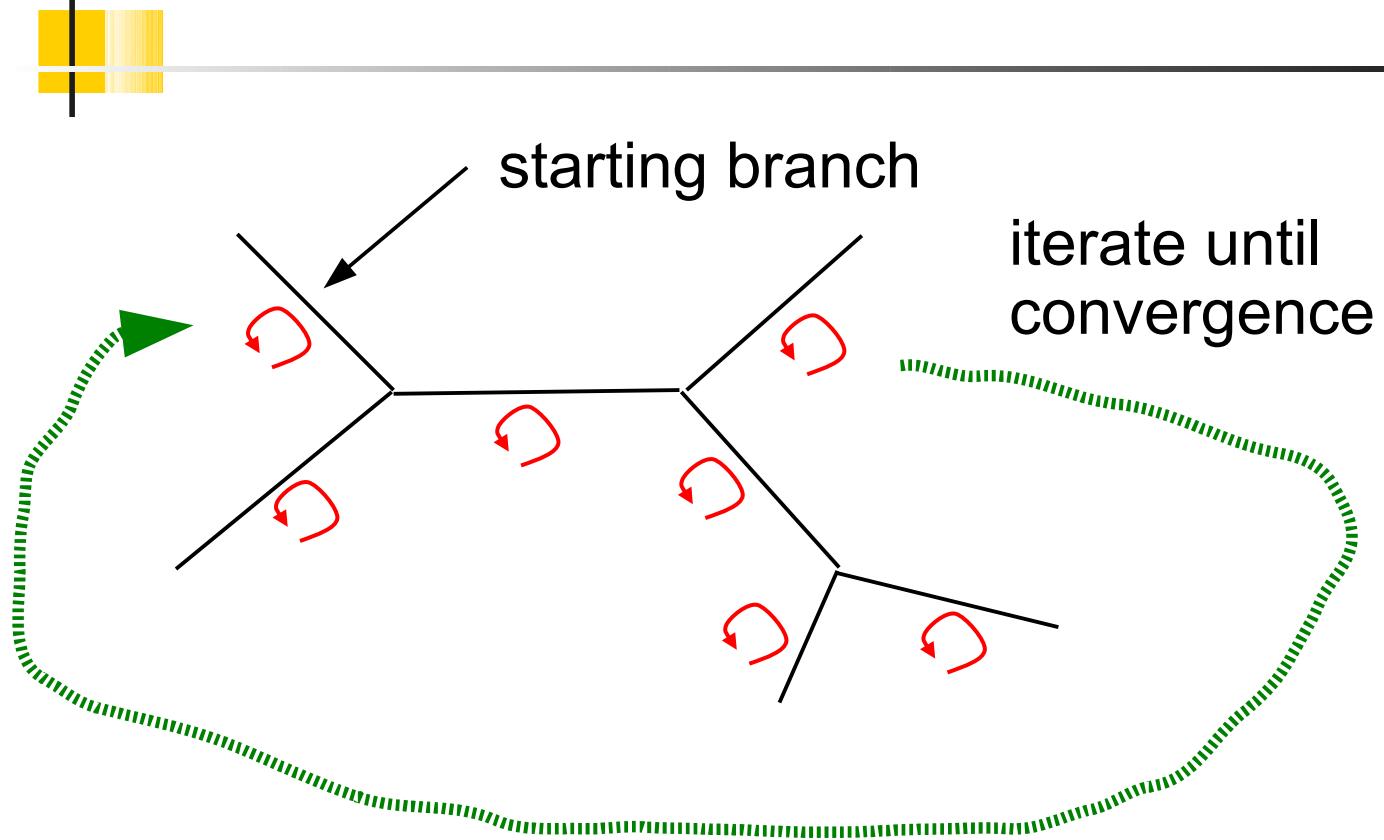
# Branch Length Optimization



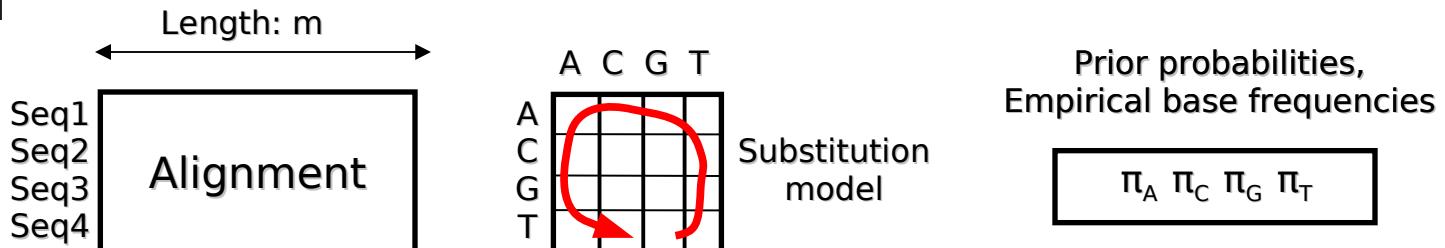
# Branch Length Optimization



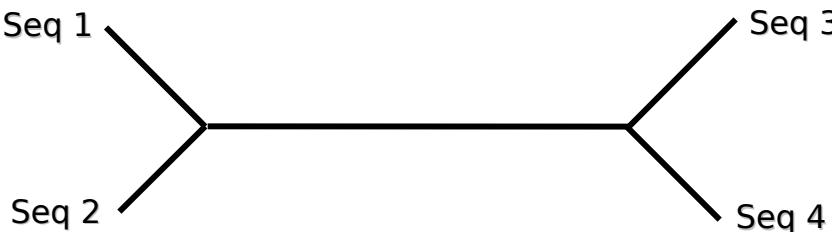
# Branch Length Optimization



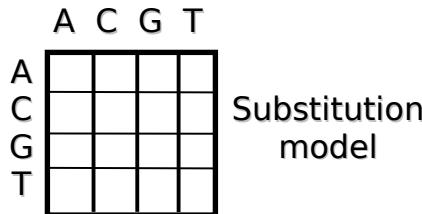
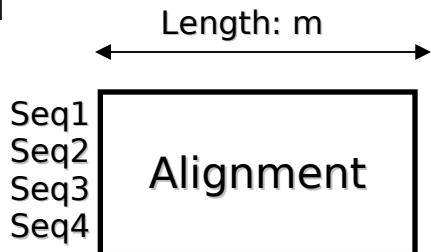
# Maximum Likelihood



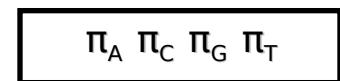
**optimize model parameters**



# Memory Requirements



Prior probabilities,  
Empirical base frequencies



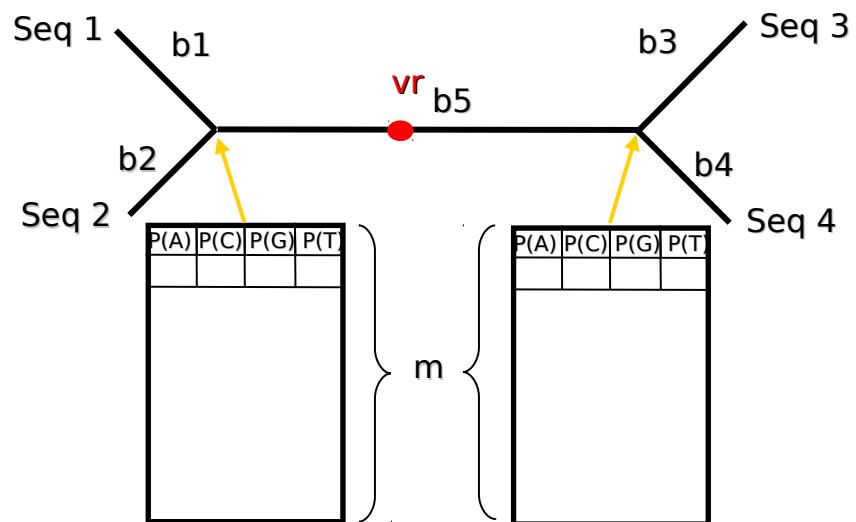
Memory Consumption:  
 $(n-2) * m * 4 * 1 * 8$  bytes

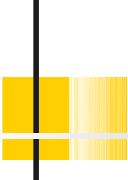
DNA

rate heterogeneity

double precision arithmetics

Three arrows point from the text "Memory Consumption:" to the factors  $n-2$ ,  $m$ ,  $4$ ,  $1$ , and  $8$  bytes. Arrows also point from the words "DNA", "rate heterogeneity", and "double precision arithmetics" to the corresponding factors in the formula.





# Memory Requirements

---

- For lazy people: mem. calculator at [www.exelixis-lab.org/software.html](http://www.exelixis-lab.org/software.html)

## RAxML Memory Requirements calculator by Simon Berger

taxa (n):  pattern (m):  DNA+GAMMA

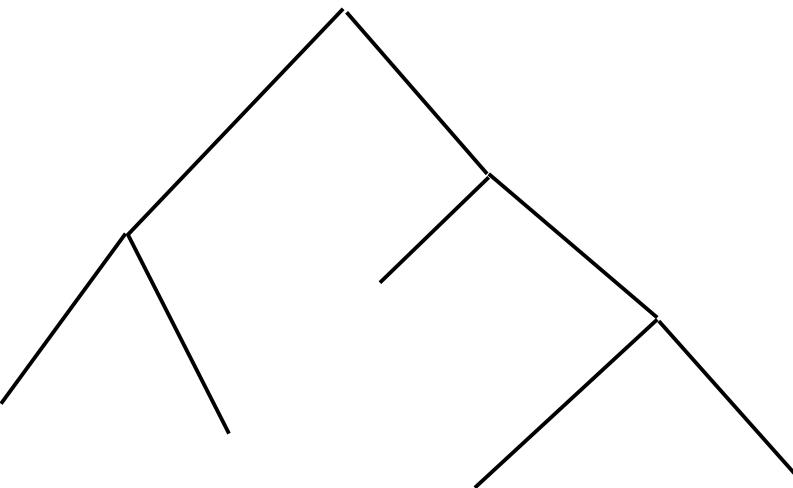
Required size:

$$(10-2) * 1000 * (16 * 8) \text{ bytes} = 1\text{MB (1024000 bytes)}$$

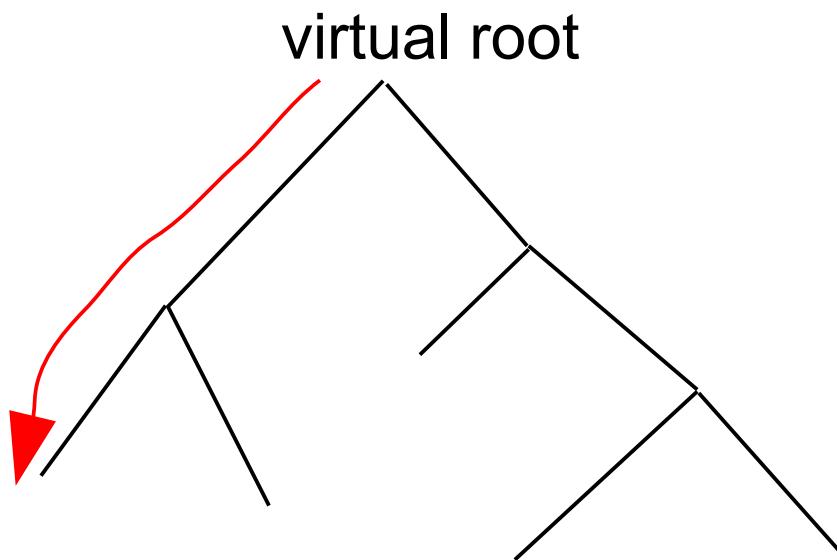


# Post-order Traversal

virtual root

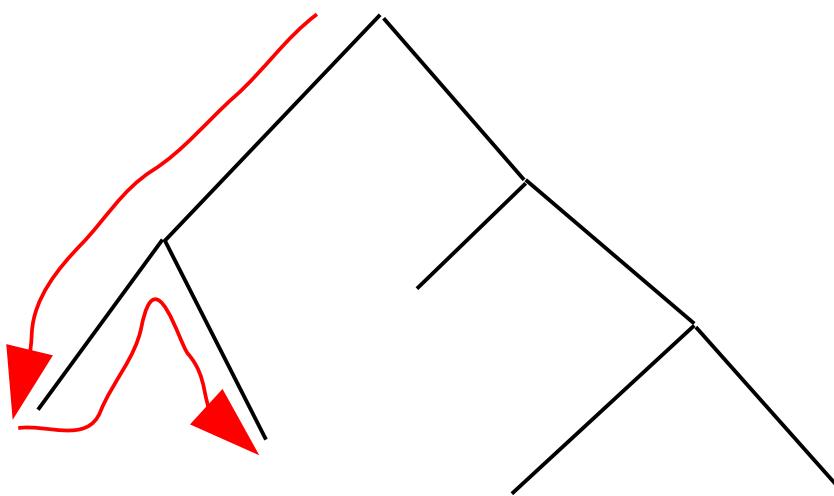


# Post-order Traversal

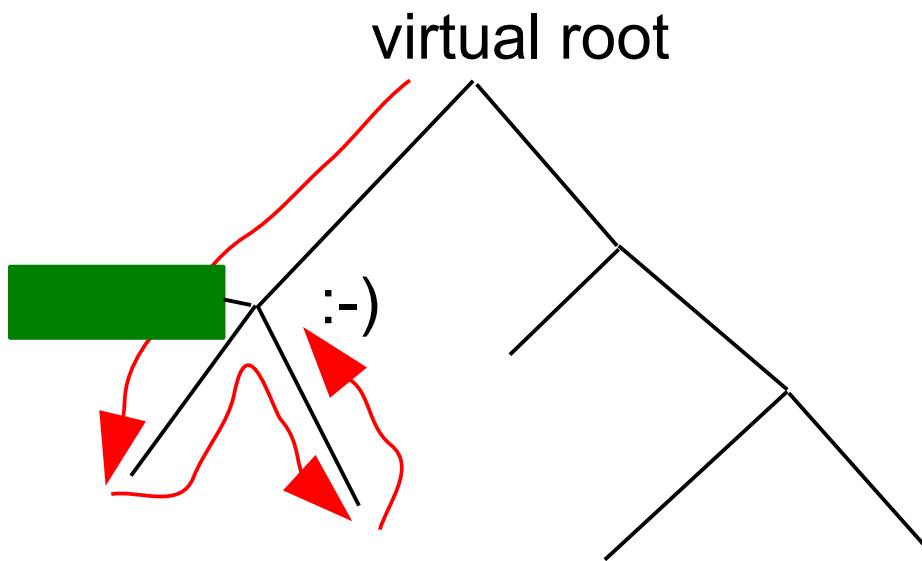


# Post-order Traversal

virtual root

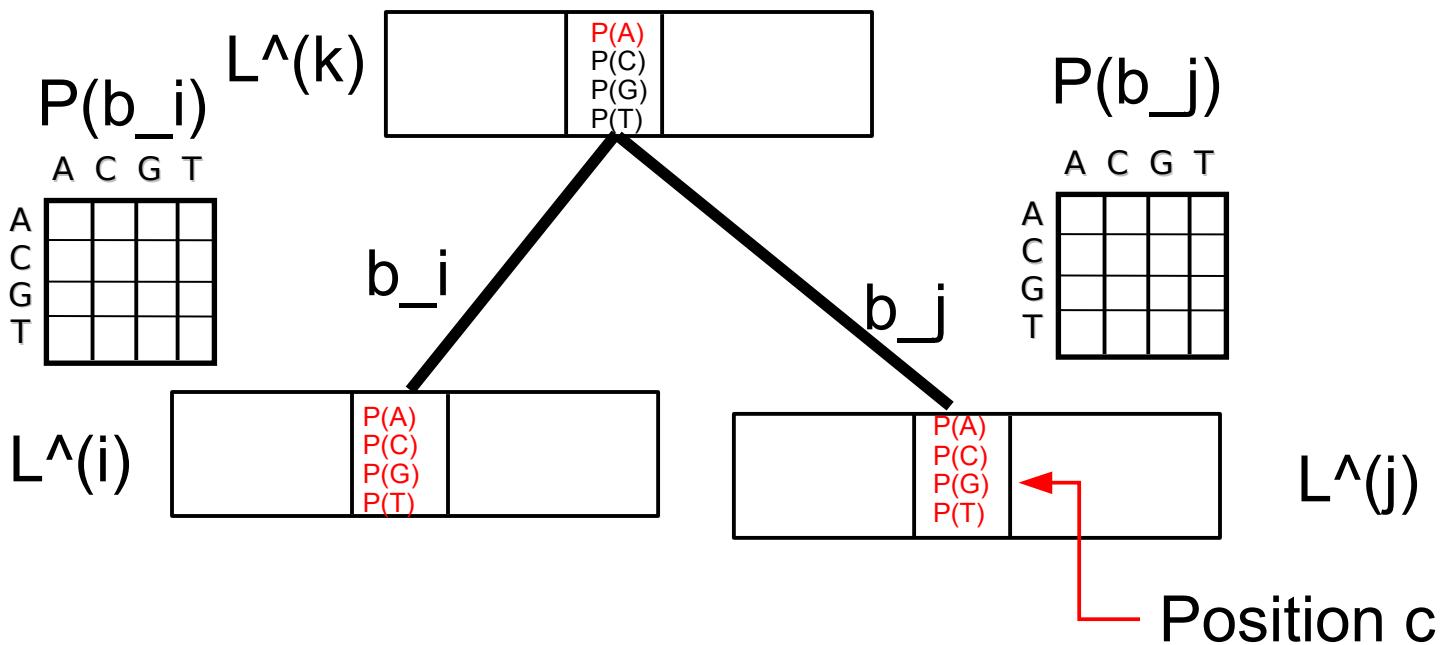


# Post-order Traversal

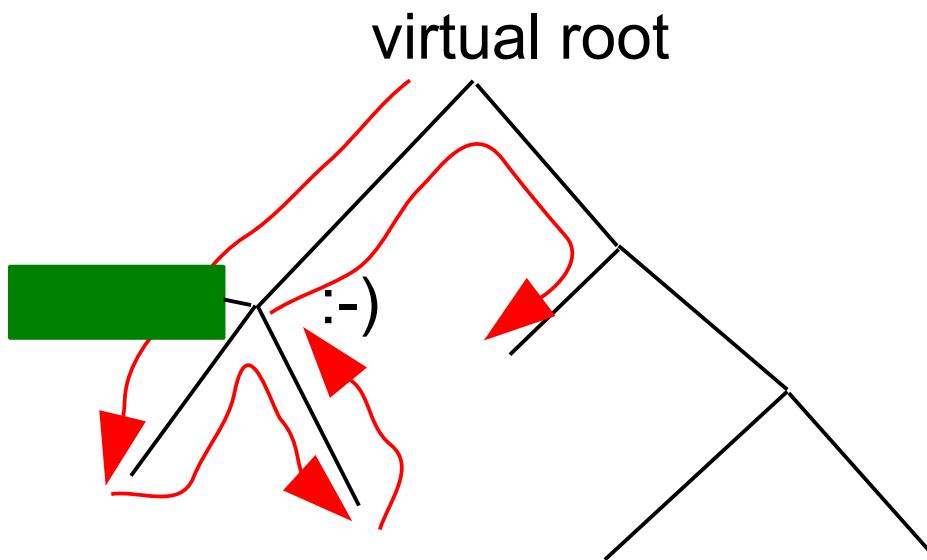


# What happens when we compute this inner vector?

$$\vec{L}_A^{(k)}(c) = \left( \sum_{S=A}^T P_{AS}(b_i) \vec{L}_S^{(i)}(c) \right) \left( \sum_{S=A}^T P_{AS}(b_j) \vec{L}_S^{(j)}(c) \right)$$

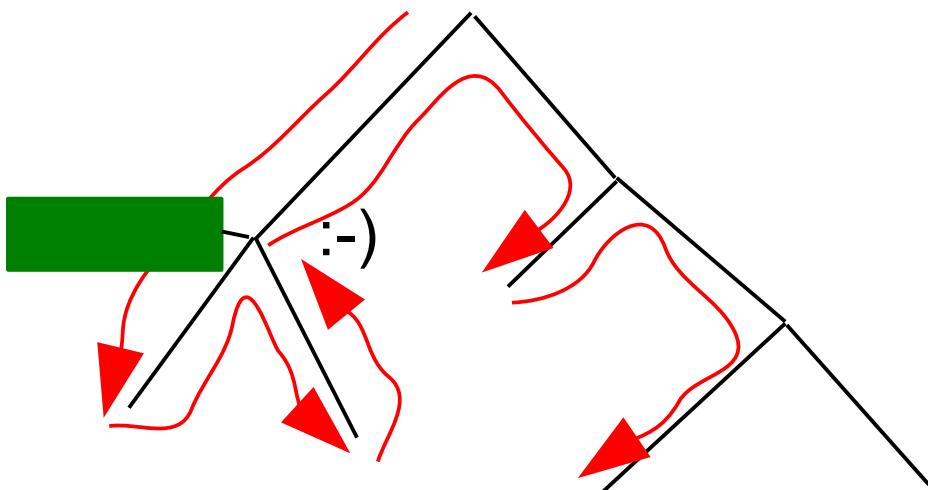


# Post-order Traversal

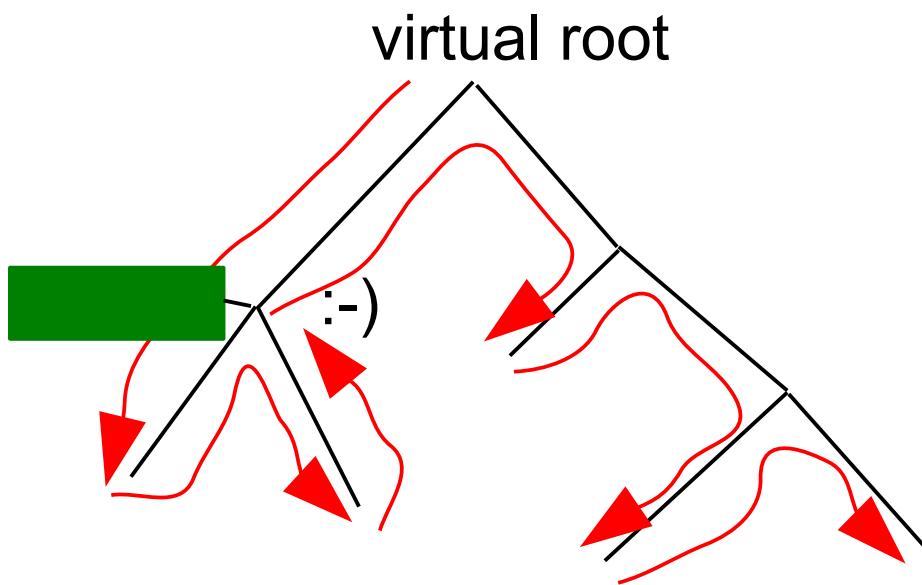


# Post-order Traversal

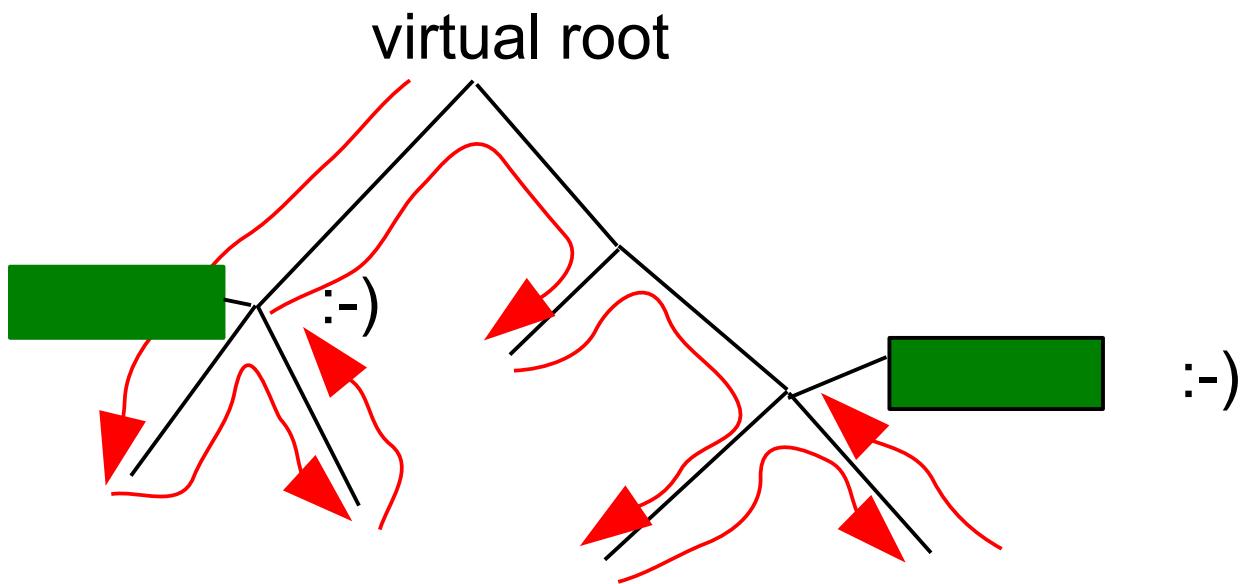
virtual root



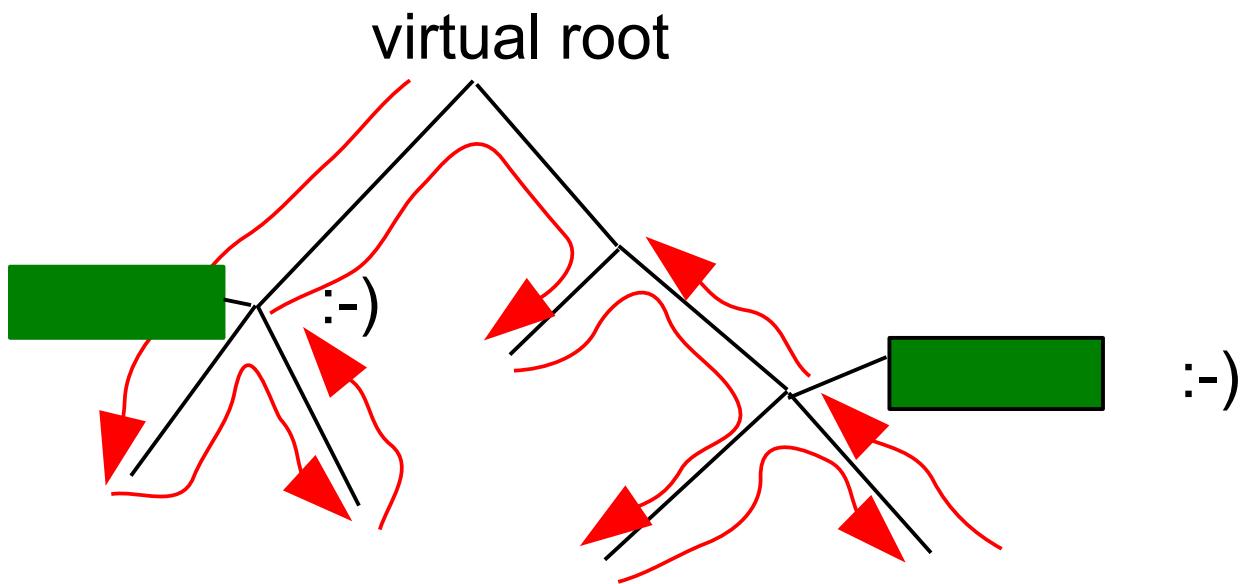
# Post-order Traversal



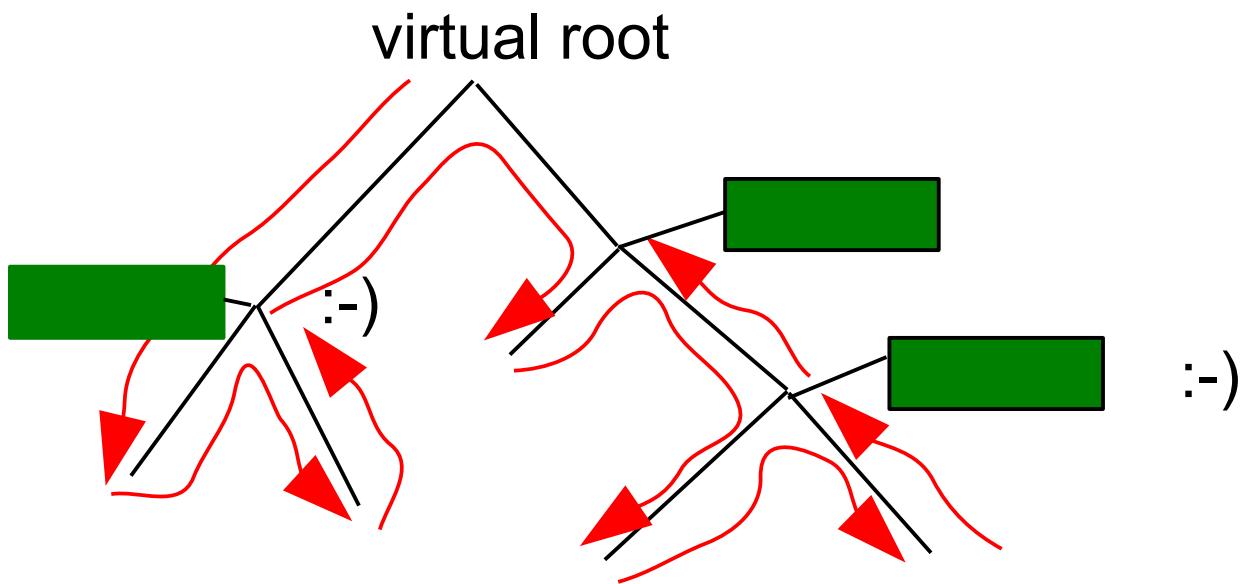
# Post-order Traversal



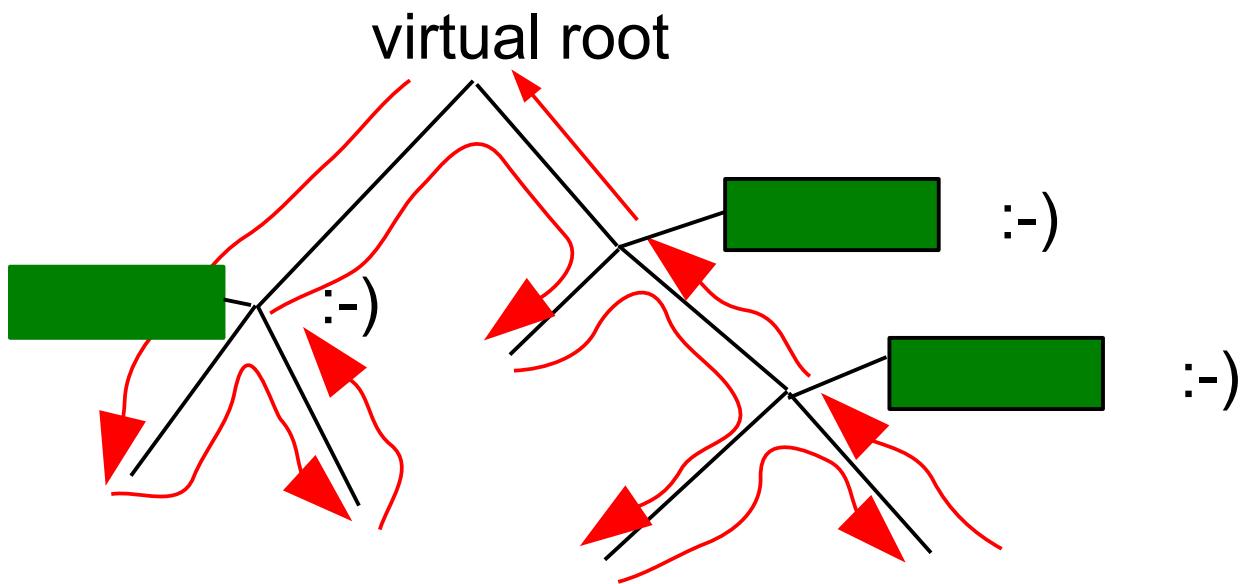
# Post-order Traversal



# Post-order Traversal

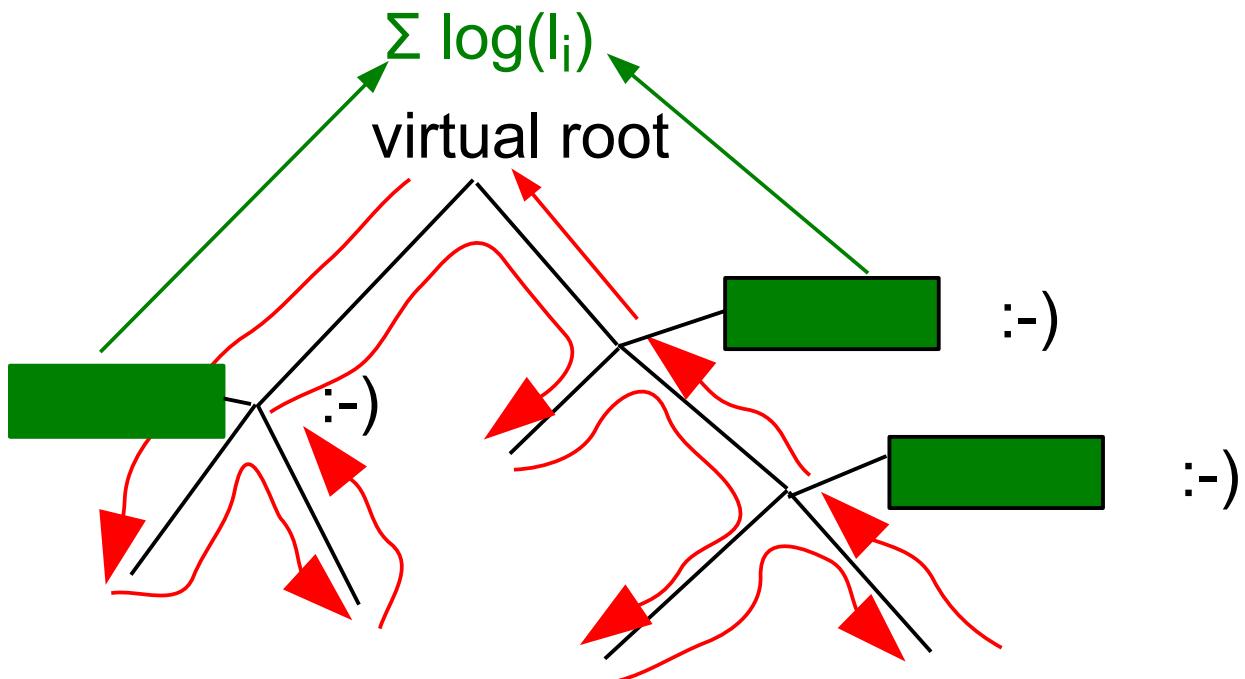


# Post-order Traversal

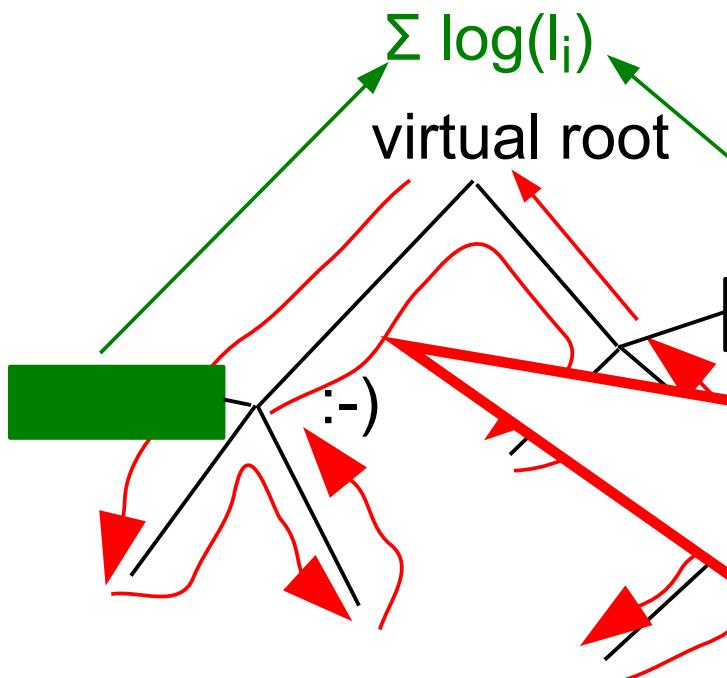


# Post-order Traversal

tree score: sum over per-site log likelihoods



# Post-order Traversal



We need to communicate the traversal order to all parallel processes:

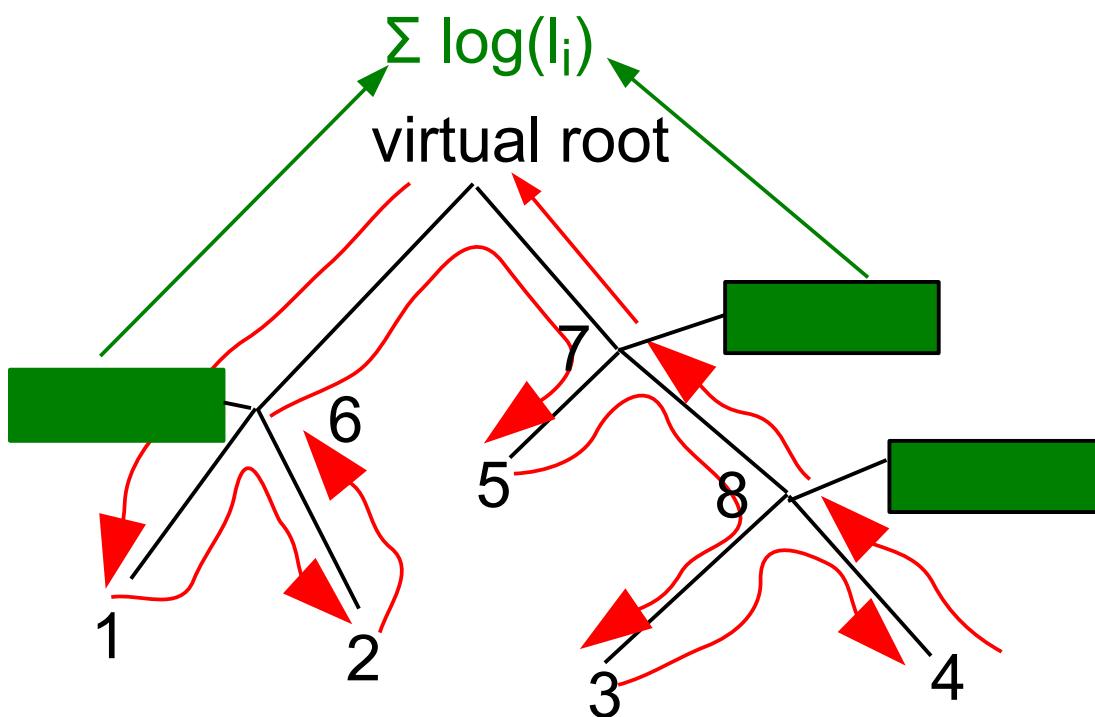
**Shared-Memory:**  
master thread computes and stores it: workers read it

**Distributed-Memory:**  
master process computes and broadcasts it as message to workers!

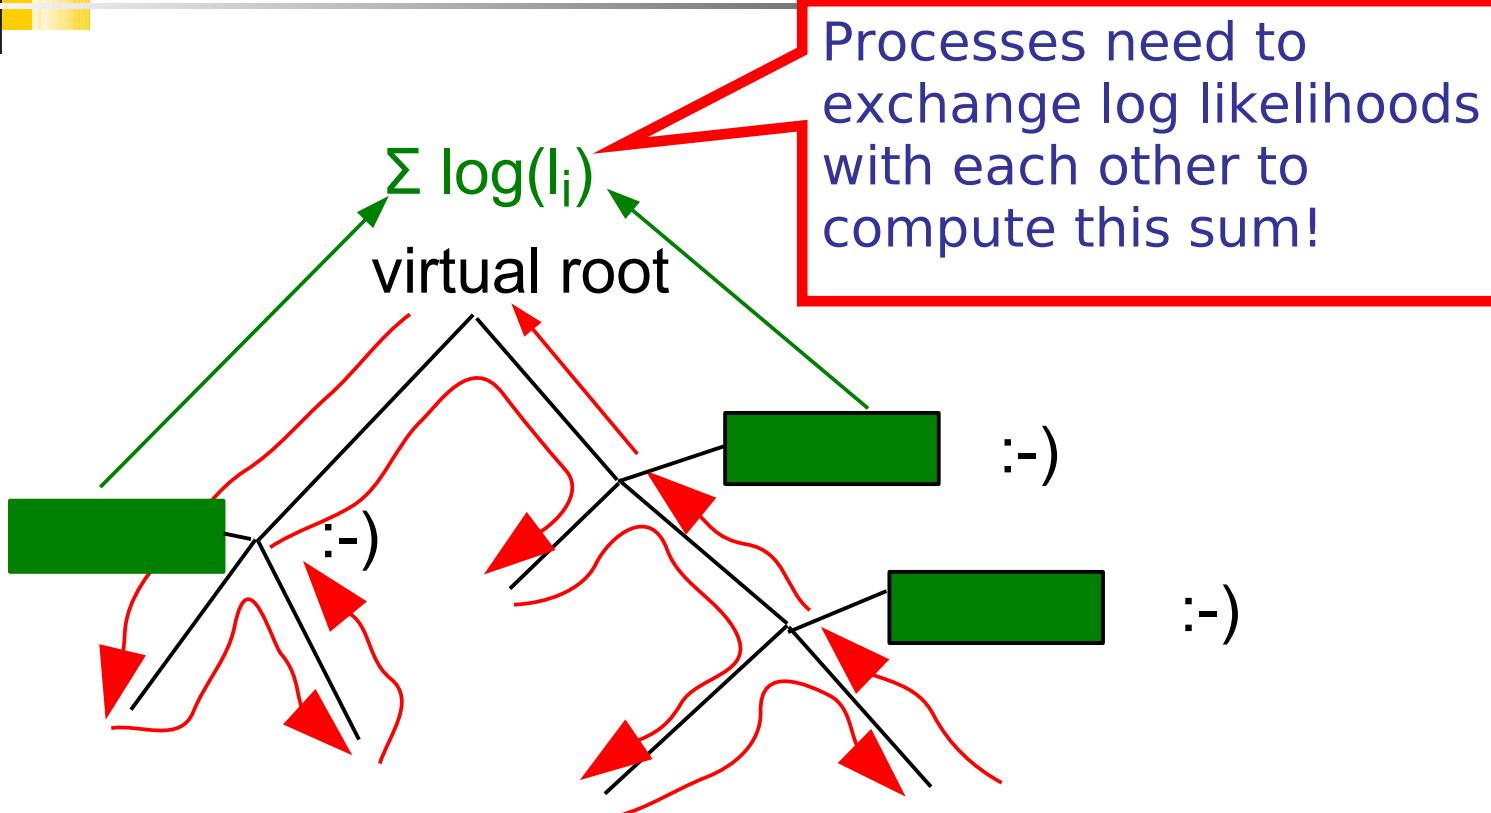
# Post-order Traversal

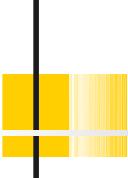
Traversal Descriptor:

$$\{ (6 \leftarrow 1,2), \\ (8 \leftarrow 3,4), \\ (7 \leftarrow 5,8), \\ (L \leftarrow 6,7) \}$$



# Post-order Traversal



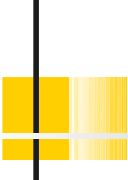


# Basic Operations

## Maximum Likelihood

---

- Compute Conditional Likelihood Array at an inner node
- Compute Likelihood at Virtual Root
- Optimize a Branch Length for a given Branch
- Optimize all Branch Lengths
- Optimize other Model Parameters



# Basic Operations

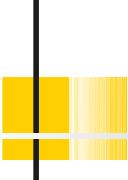
## Maximum Likelihood

---

- Compute Conditional Likelihood Array at an inner node
- Compute Likelihood at Virtual Root
- Optimize a Branch Length for a given Branch
- Optimize all Branch lengths
- Optimize other Parameters



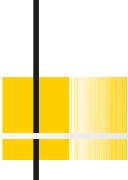
Bayesian programs  
only require two  
operations



# Outline

---

- Computing the Likelihood
- Optimizing & Parallelizing Likelihood Computations
- Saving Memory in Likelihood Computations
- HPC population genetics



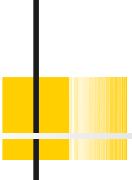
# Optimization

---

- Use vector intrinsics SSE3/AVX
- Special implementations for:



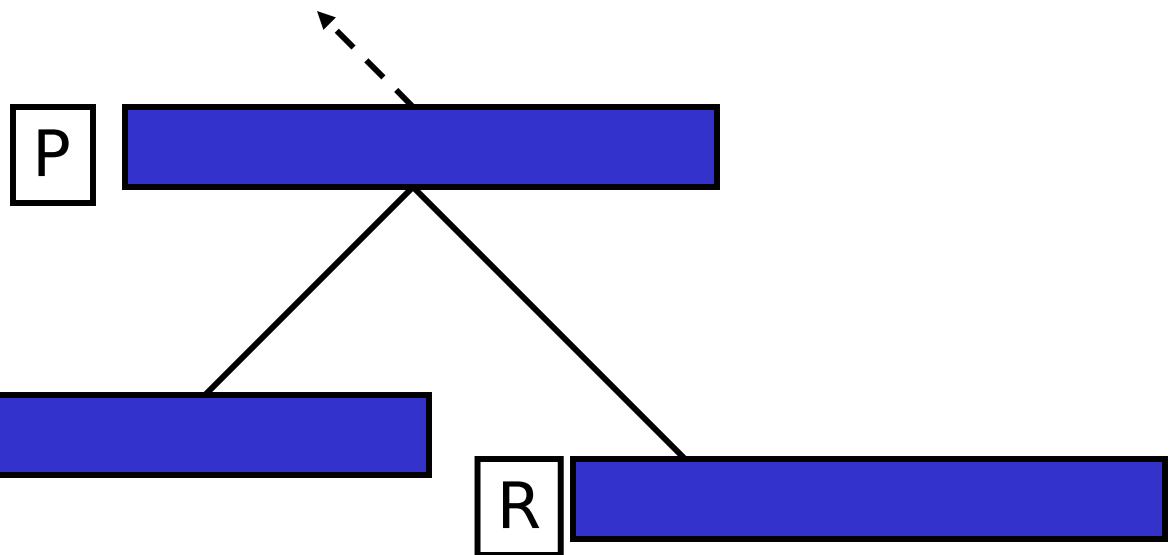
- I will spare you the details
- But, avoid redundant computations



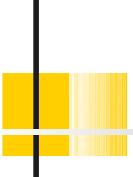
# Loop Level Parallelism

---

virtual root



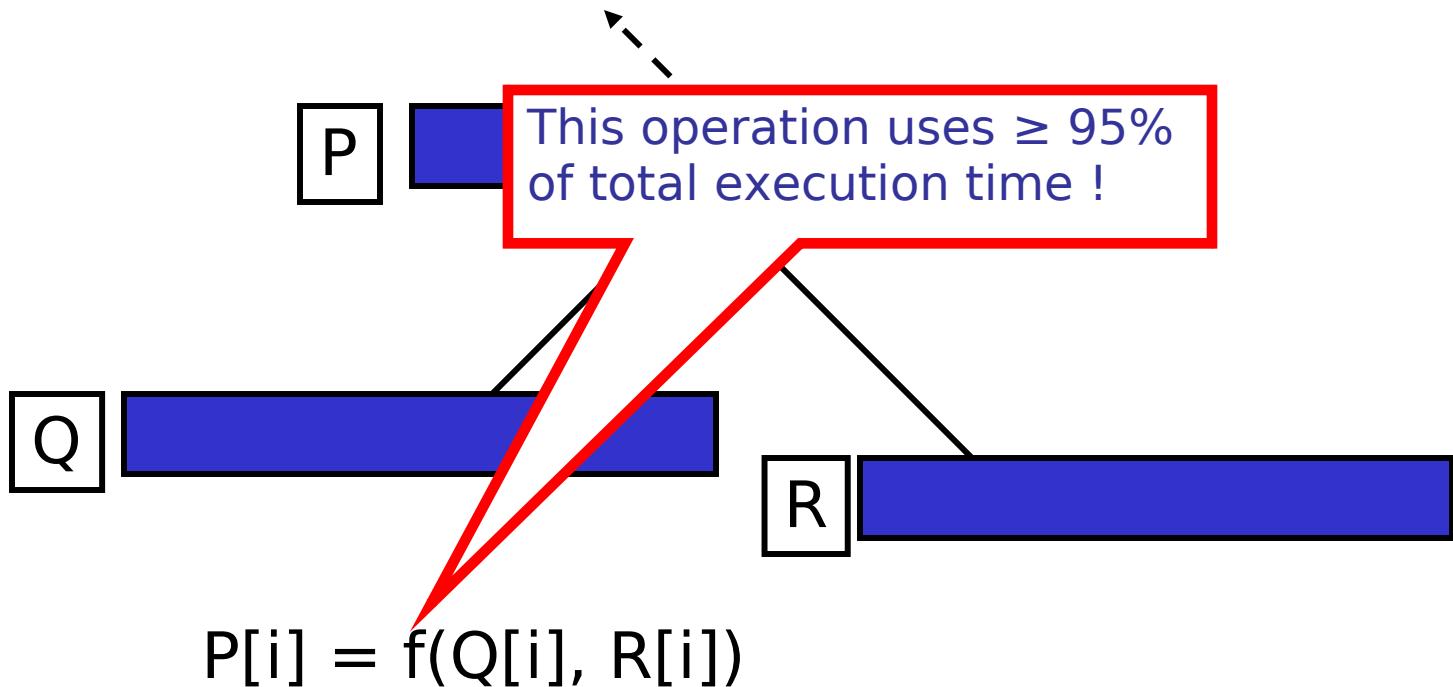
$$P[i] = f(Q[i], R[i])$$

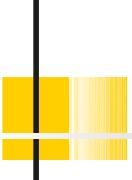


# Loop Level Parallelism

---

virtual root

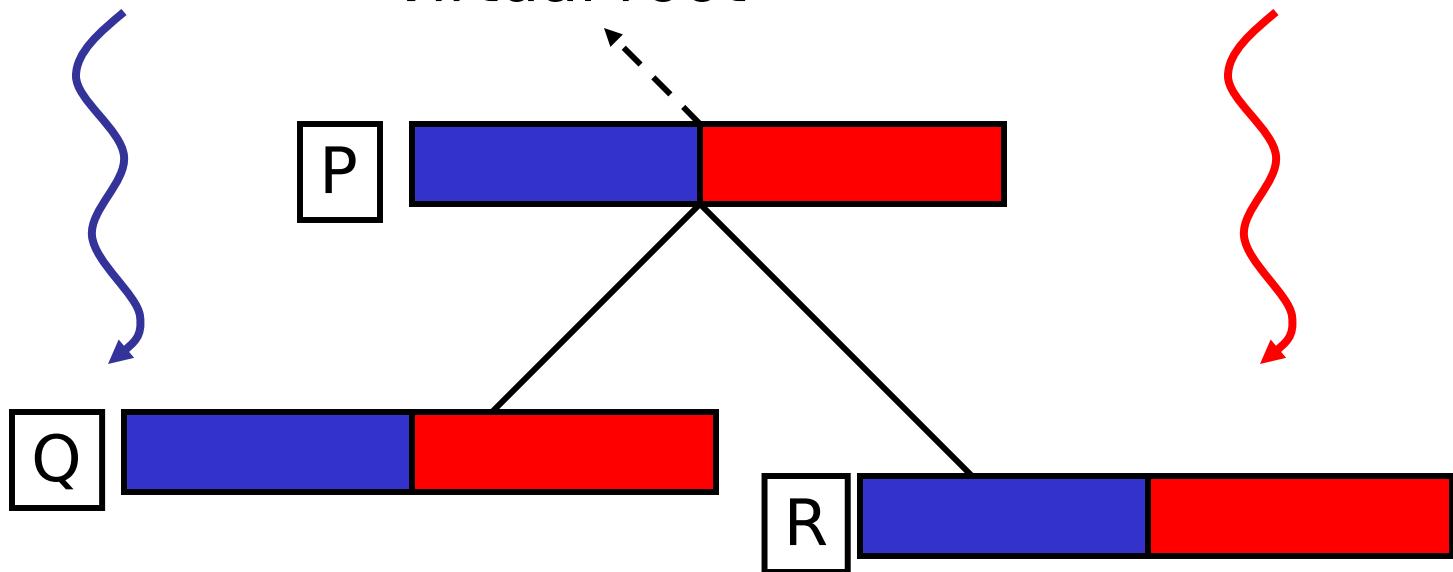




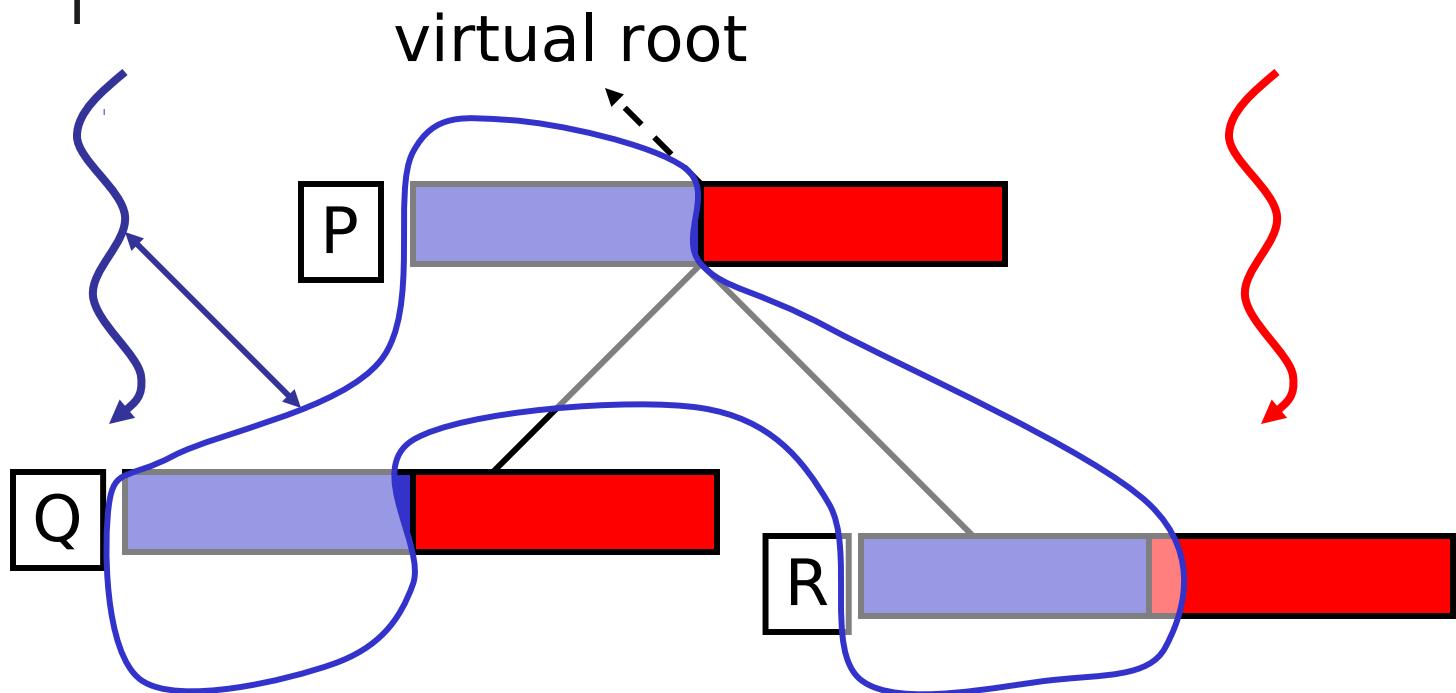
# Loop Level Parallelism

---

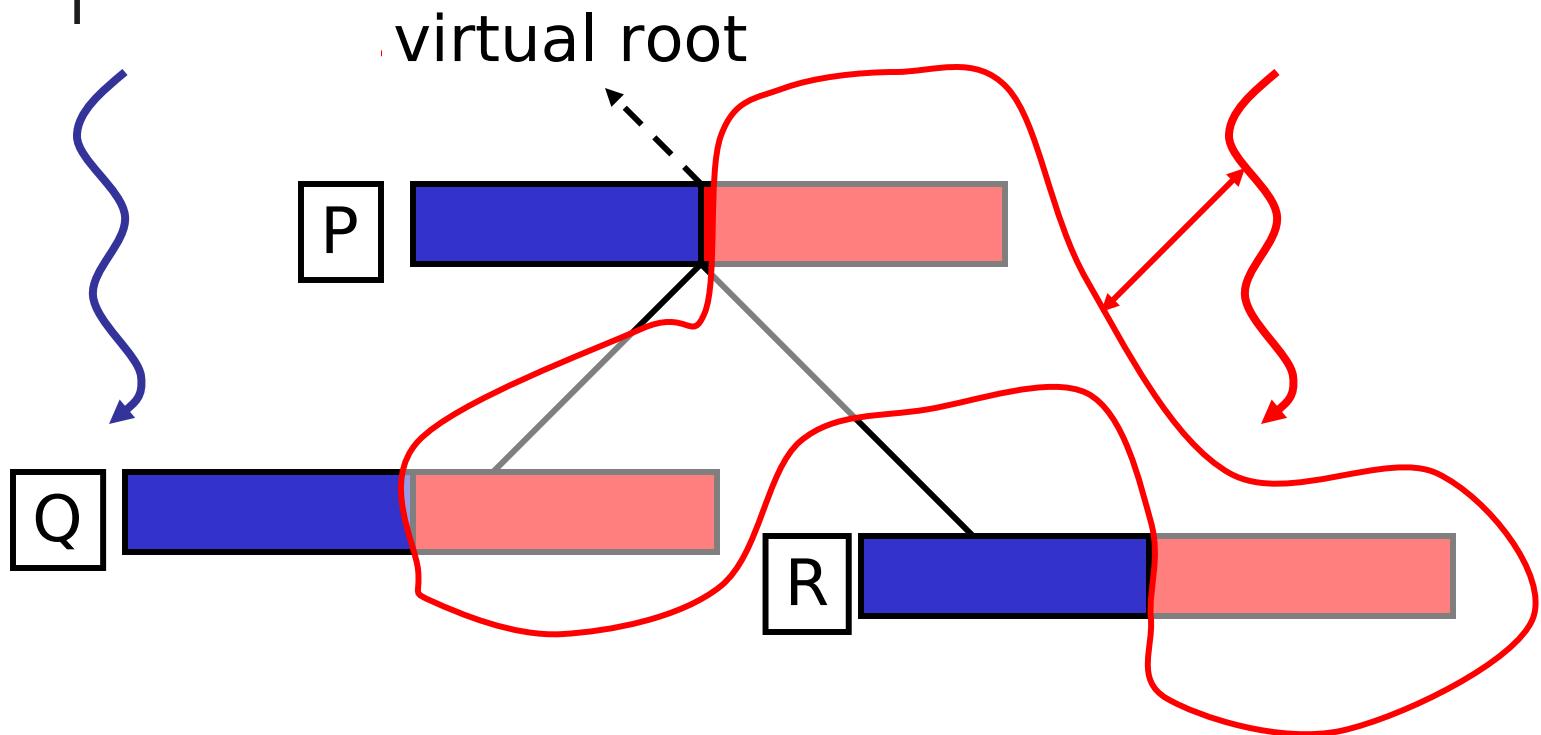
virtual root



# Loop Level Parallelism

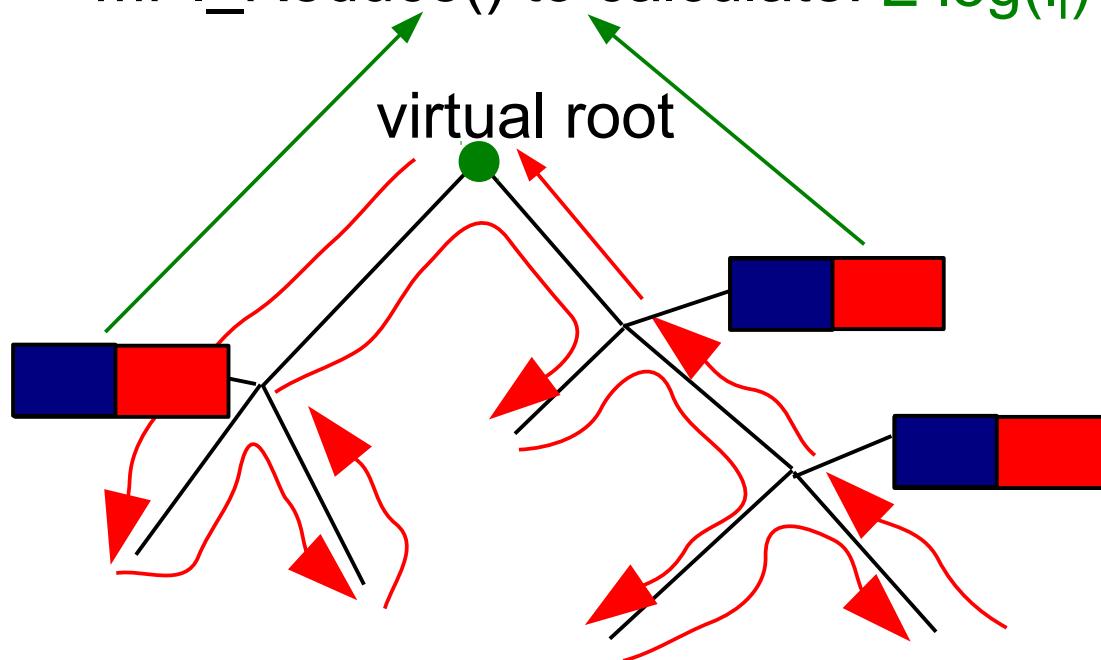


# Loop Level Parallelism

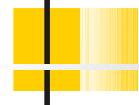


# Parallel Post-order Traversal

Only need to synchronize at the root  
→ MPI\_Reduce() to calculate:  $\sum \log(l_i)$



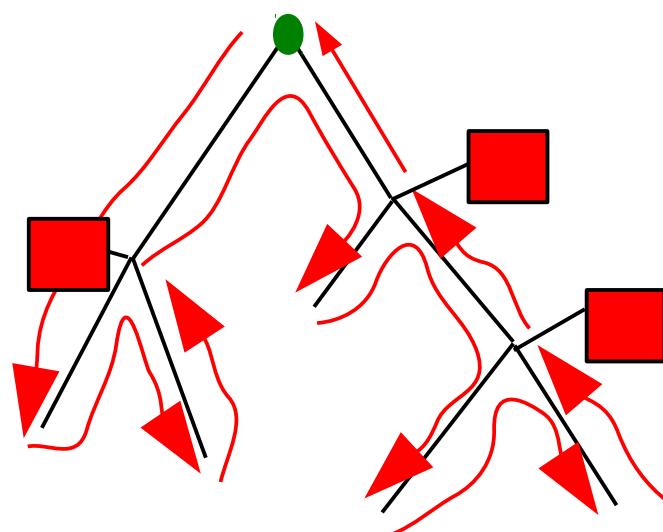
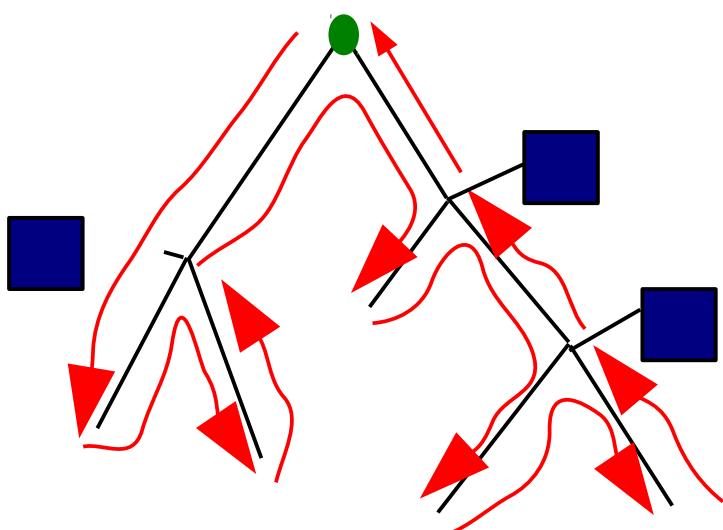
# Parallel Post-order Traversal



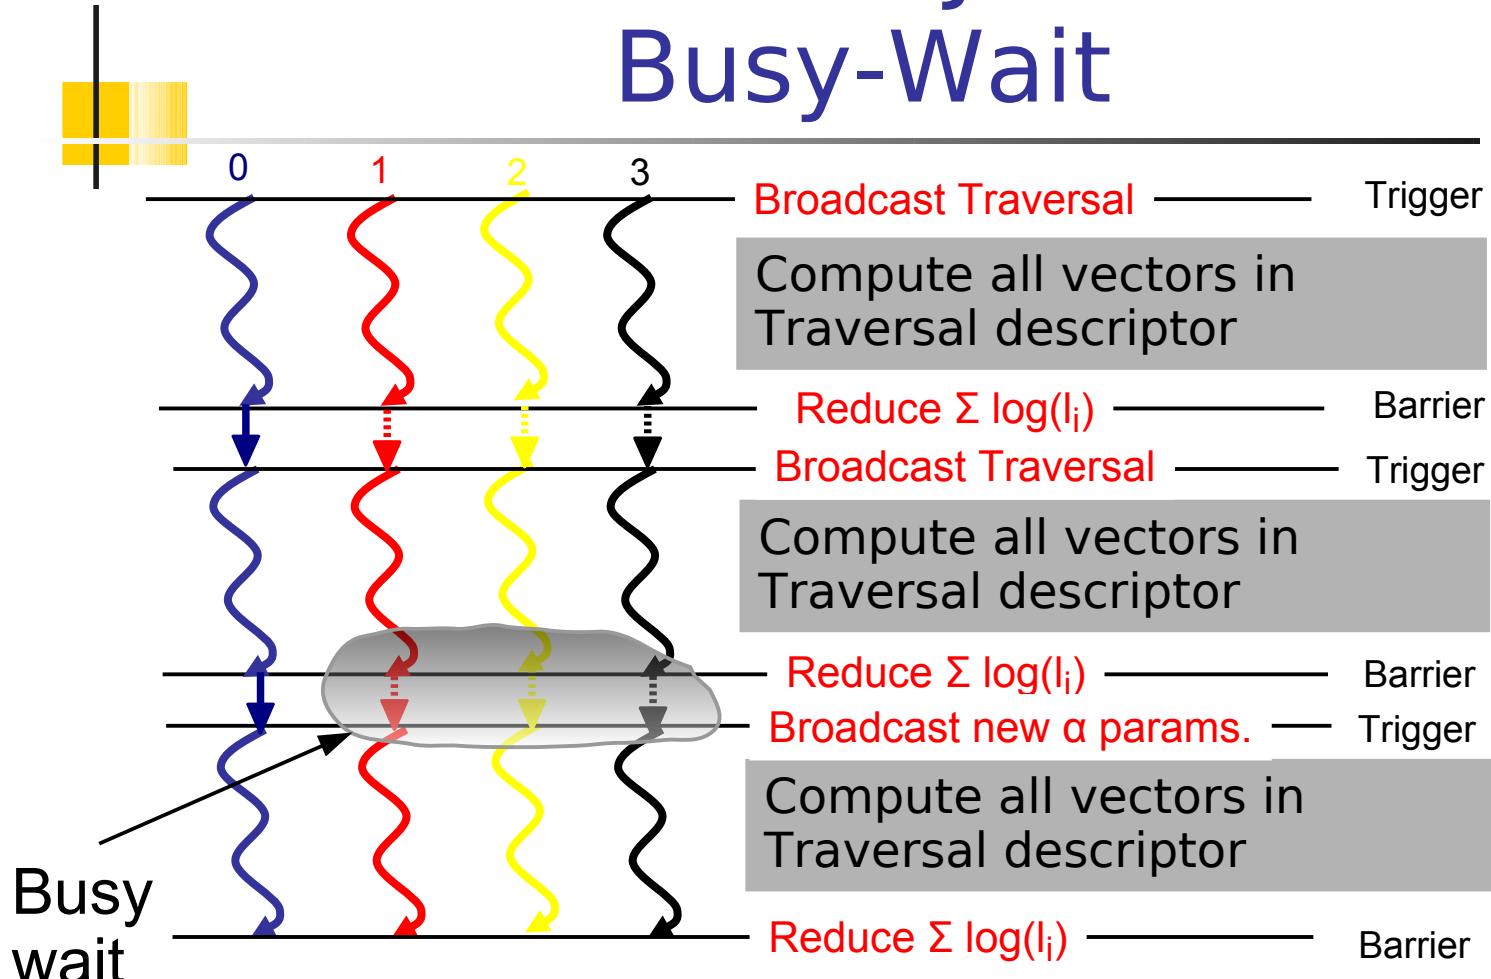
Overall Score

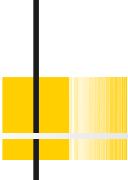
$$\Sigma \log(l_i)$$

$$\Sigma \log(l_i)$$



# Classic Fork-Join with Busy-Wait



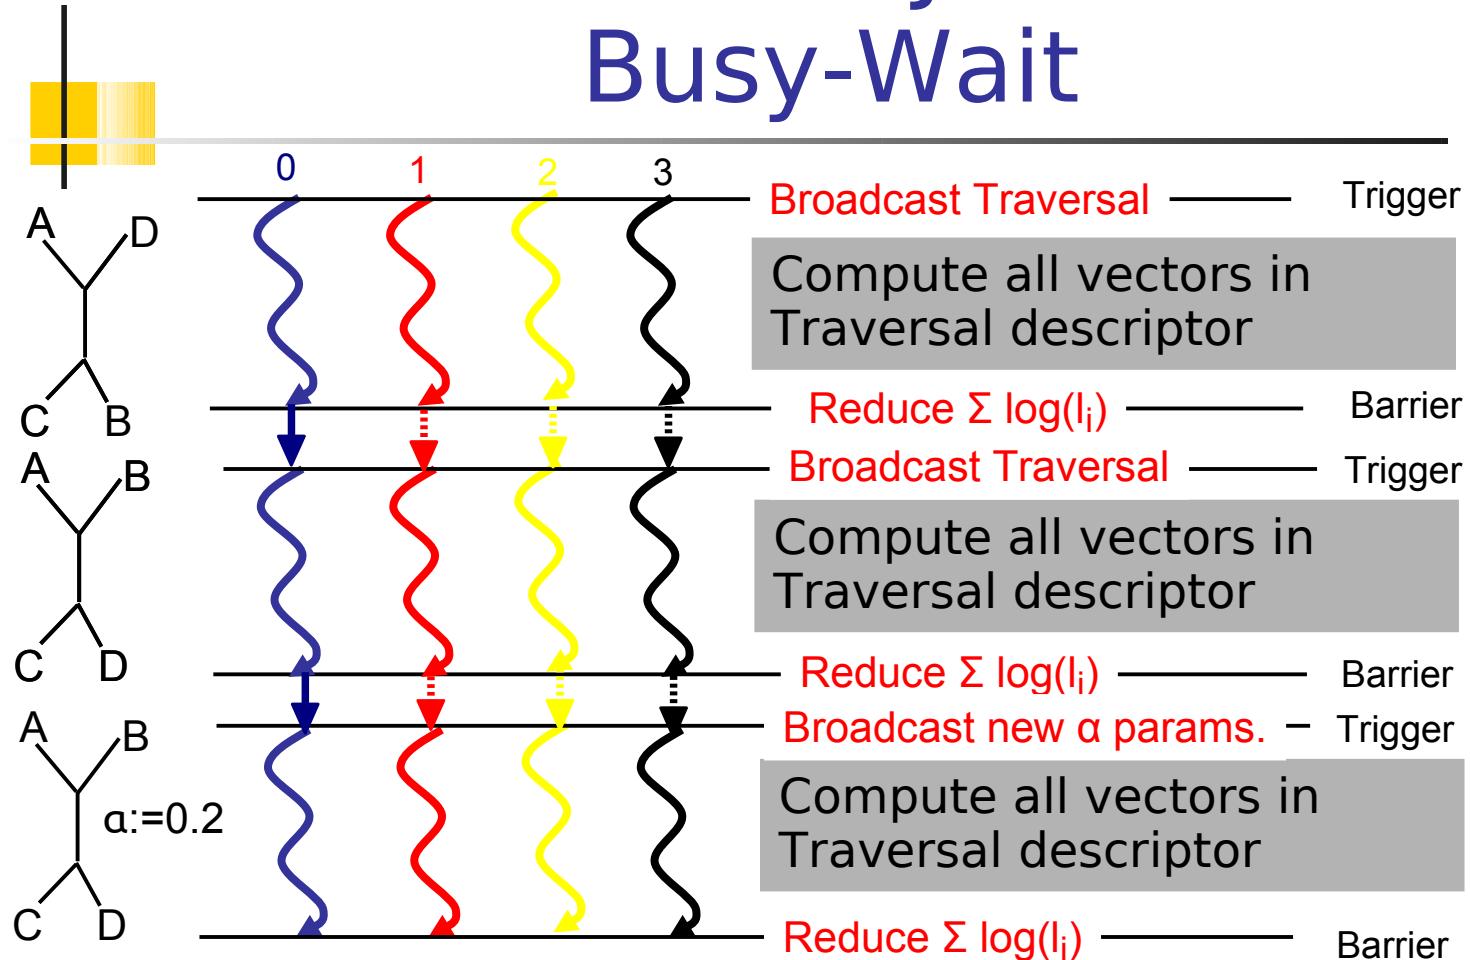


# Synchronizations in RAxML with Pthreads

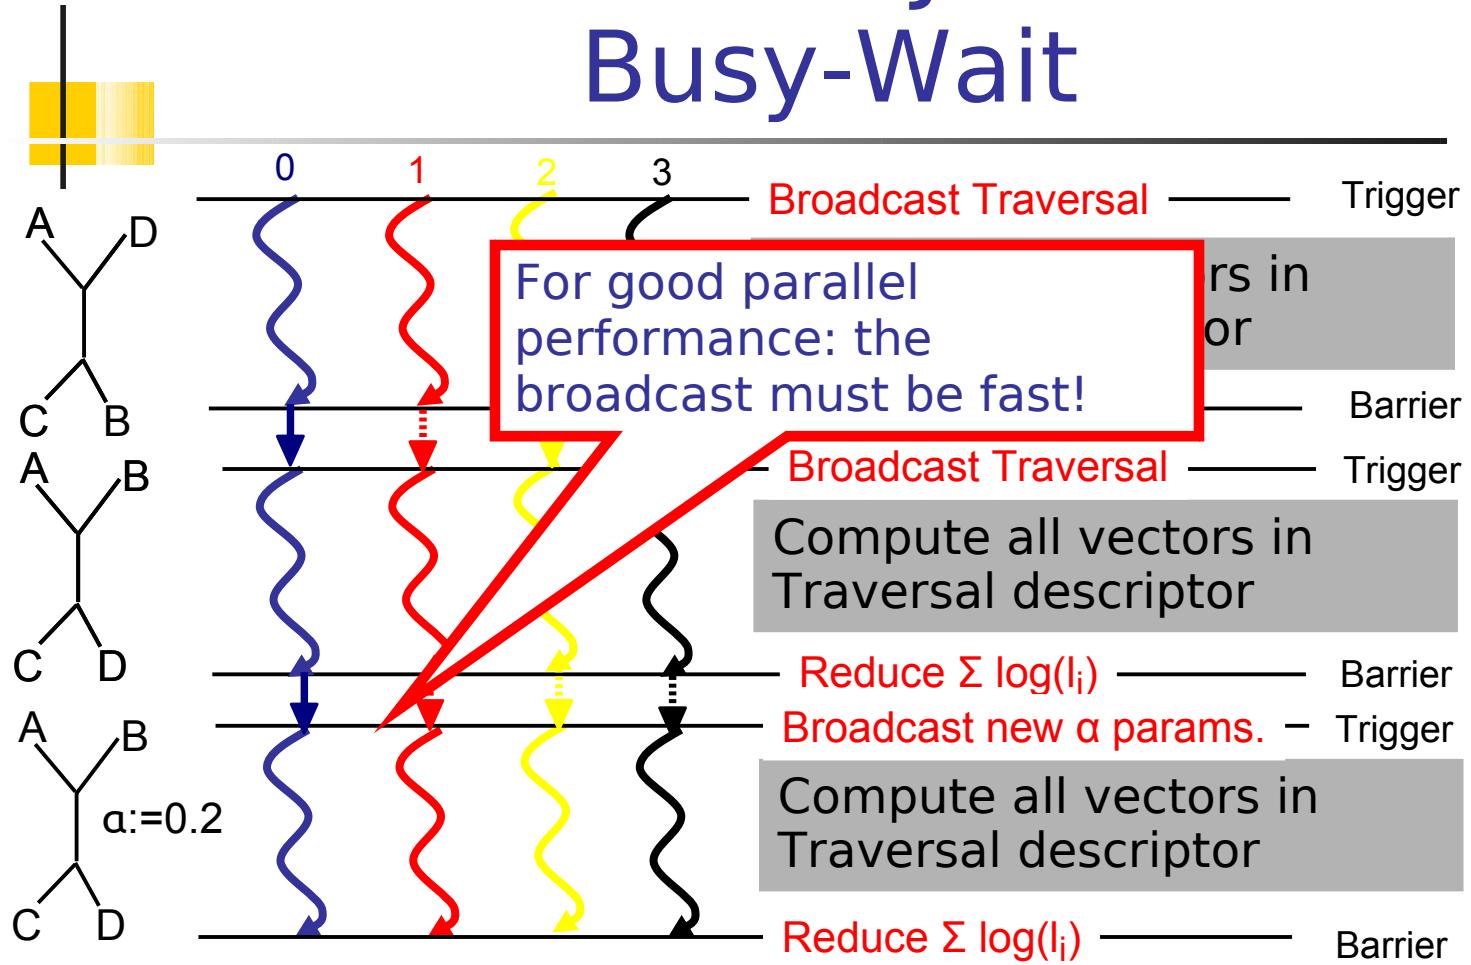
---

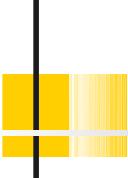
- RAxML Pthreads for a run time of about **10 seconds** on 16 Cores/16 Threads
- 404 taxa 7429 sites: **194,000** Barriers
- 1481 taxa 1241 sites: **739,000** Barriers

# Classic Fork-Join with Busy-Wait



# Classic Fork-Join with Busy-Wait



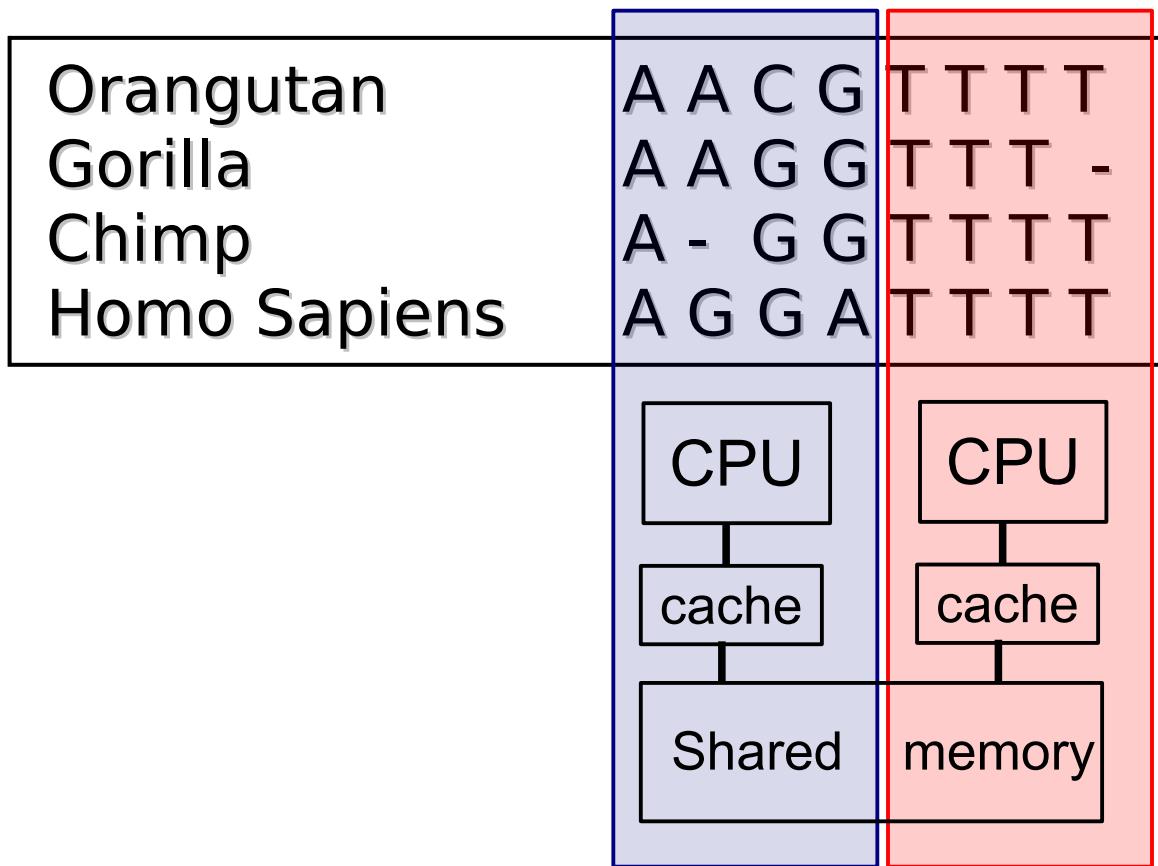


# Parallel Performance Problems

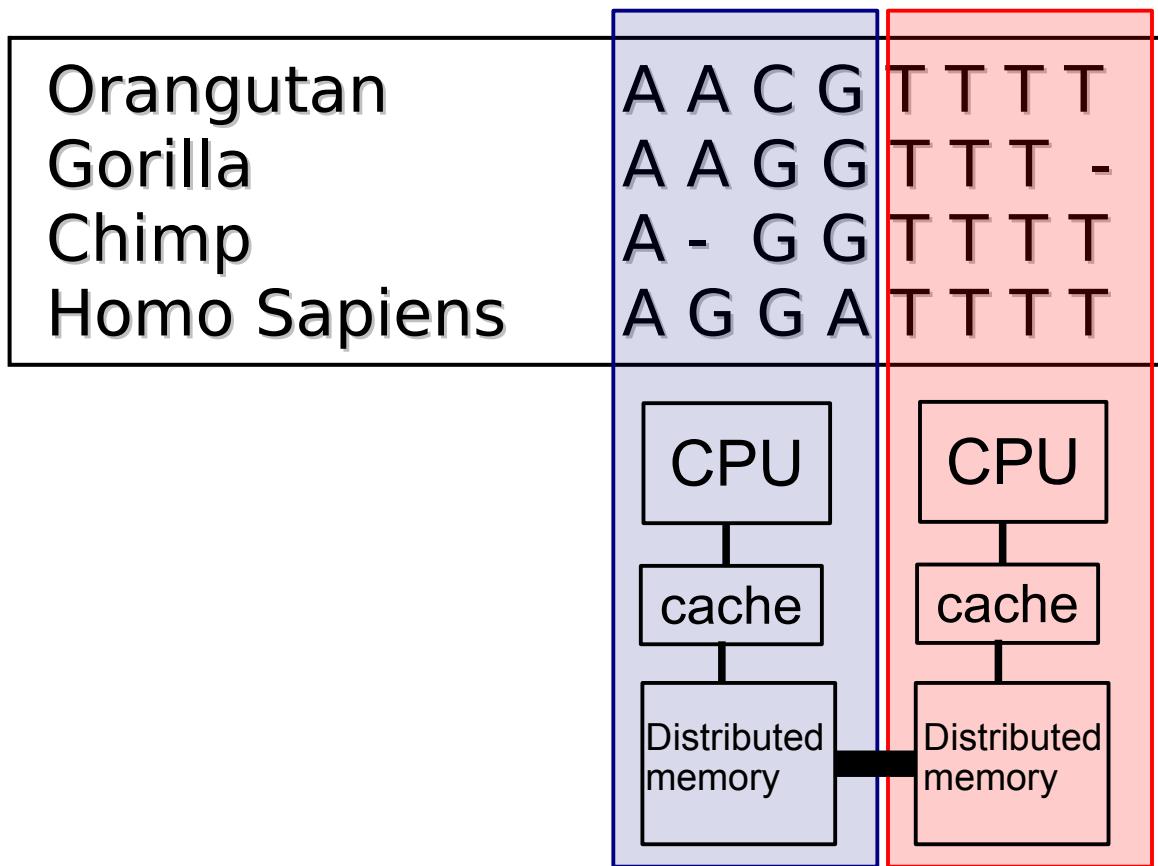
---

- They all start with partitioned datasets!
- How do we distribute partitions to processors?
- How do we calculate parameter changes?
- How much time does our Broadcast take?
- Goal: Keep all processors busy all the time → reduce amount of communication and synchronization!

# Data Distribution

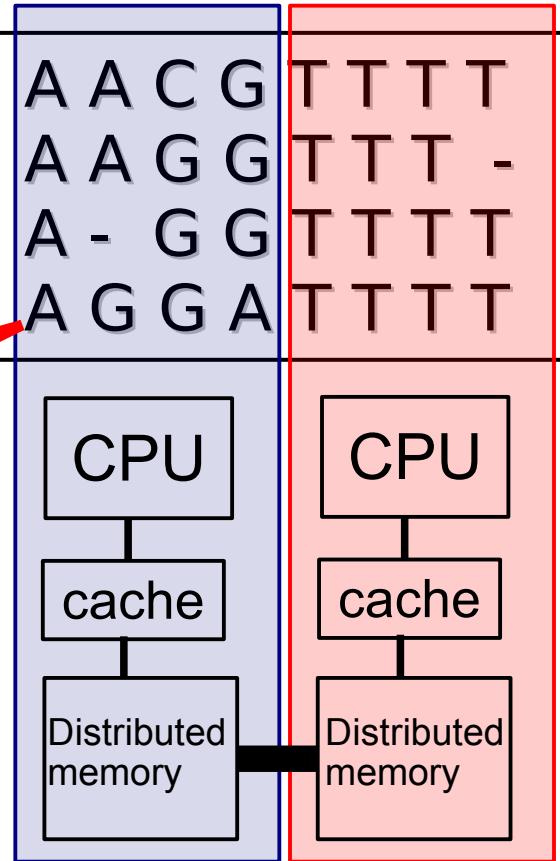


# Data Distribution



# Data Distribution

Orangutan  
Gorilla  
Chimp  
Homo Sapiens

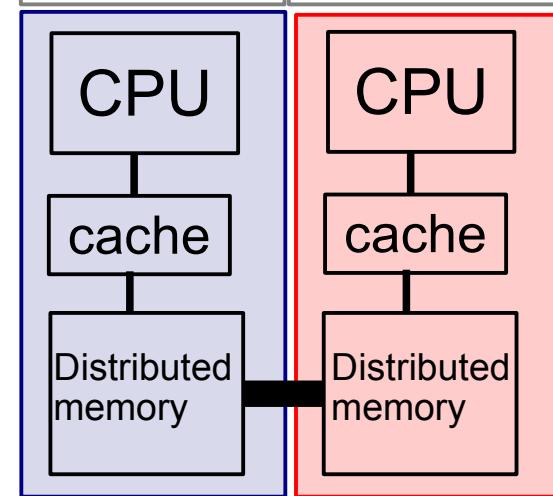


Partitioned data distribution is not that trivial!

# Data Distribution

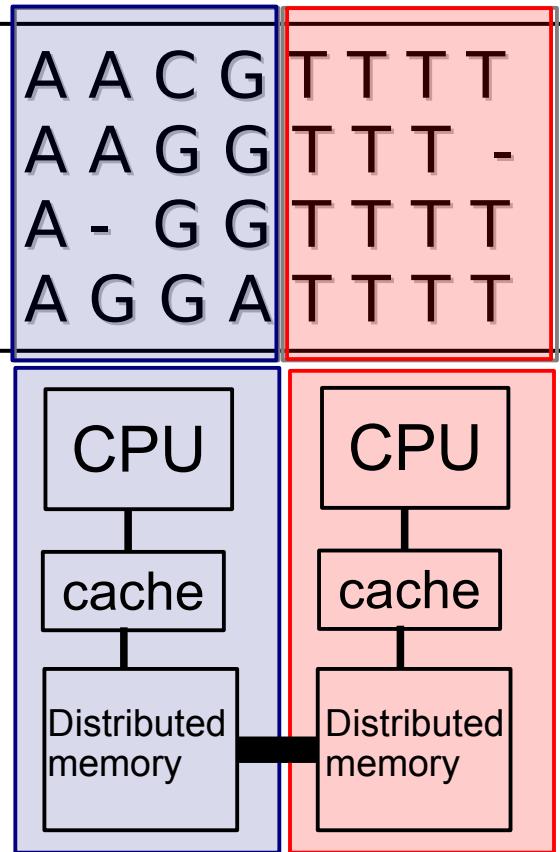
Orangutan  
Gorilla  
Chimp  
Homo Sapiens

A	A	C	G	T	T	T	T
A	A	G	G	T	T	T	-
A	-	G	G	T	T	T	T
A	G	G	A	T	T	T	T



# Data Distribution I

Orangutan  
Gorilla  
Chimp  
Homo Sapiens

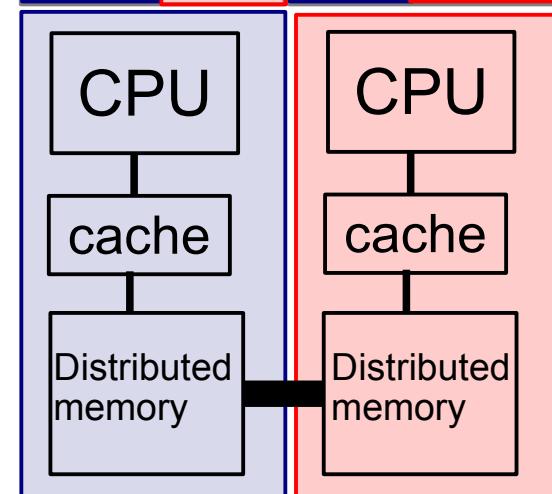


Works well when we have more partitions than processors:  
May lead to load imbalance not all processors obtain equal number of sites!

# Data Distribution II

Orangutan  
Gorilla  
Chimp  
Homo Sapiens

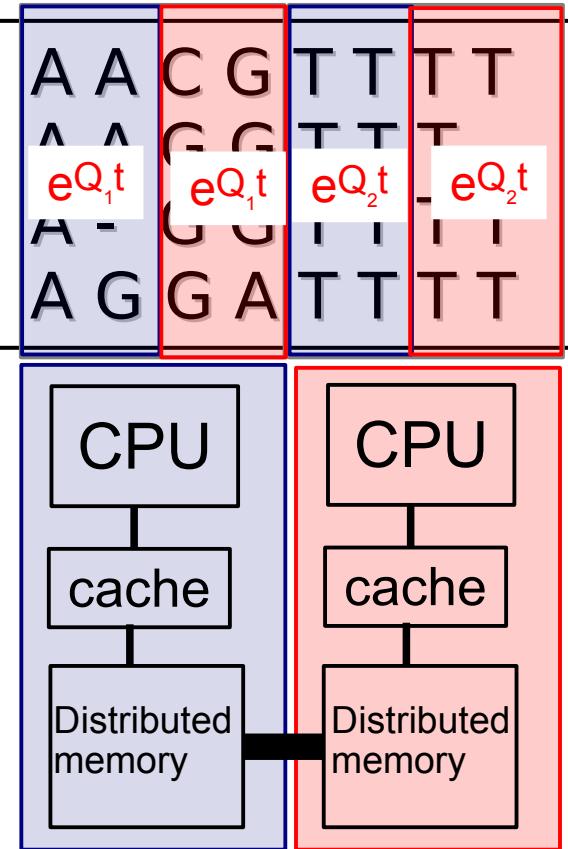
A A	C G	T T	T T
A A	G G	T T	T -
A -	G G	T T	T T
A G	G A	T T	T T



Works well when we have more processors than partitions:  
However we will need to compute:  $P(t) = e^{Qt}$  for each partition at each processor!

# Data Distribution II

Orangutan  
Gorilla  
Chimp  
Homo Sapiens



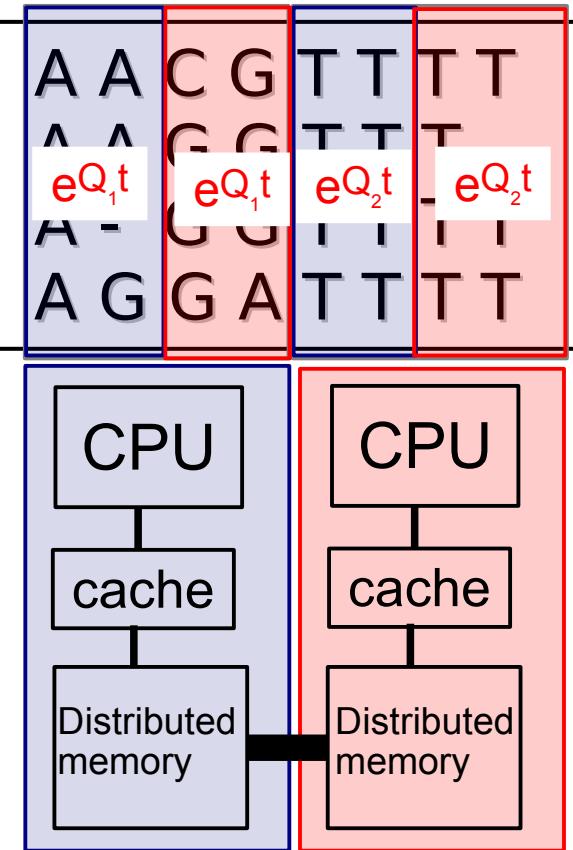
Works well when we have more processors than partitions:  
However we will need to compute:  $P(t)=e^{Qt}$  for each partition at each processor!

# Data Distribution II

Orangutan  
Gorilla

Performance impact  
depends on number of  
states in data/dimension  
of Q

Works well when  
we have more processors  
than partitions:  
However we will need to  
compute:  $P(t) = e^{Qt}$  for each  
partition at each processor!

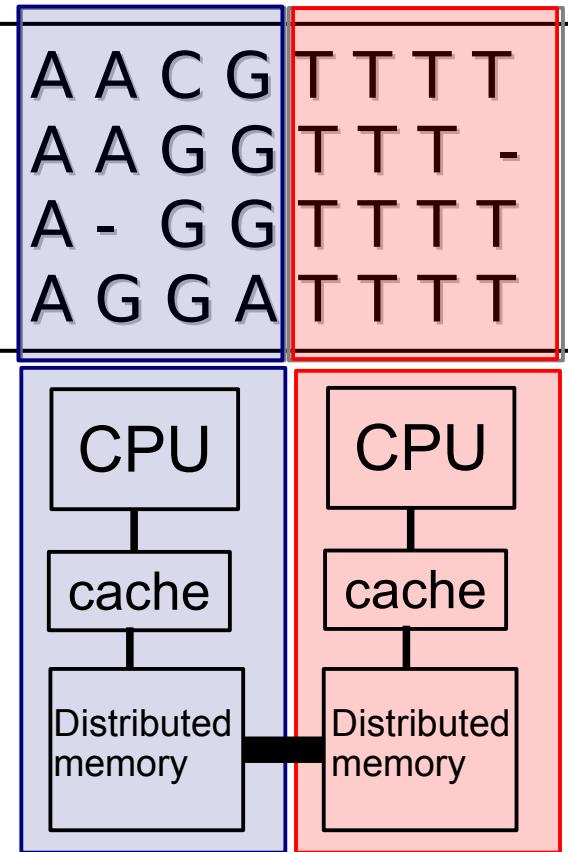


# Data Distribution I

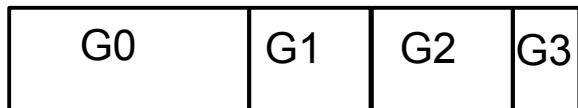
Orangutan  
Gorilla

How do we distribute partitions to processors?

Works well when we have more partitions than processors:  
May lead to load imbalance not all processors get equal number of sites!



# Load Balance I



P0



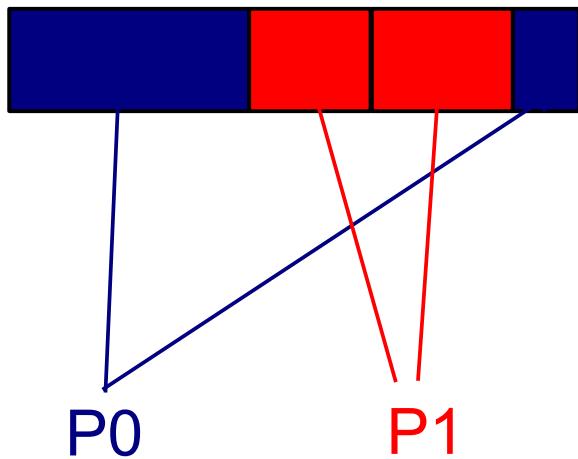
SOFTGATE.ch

P1

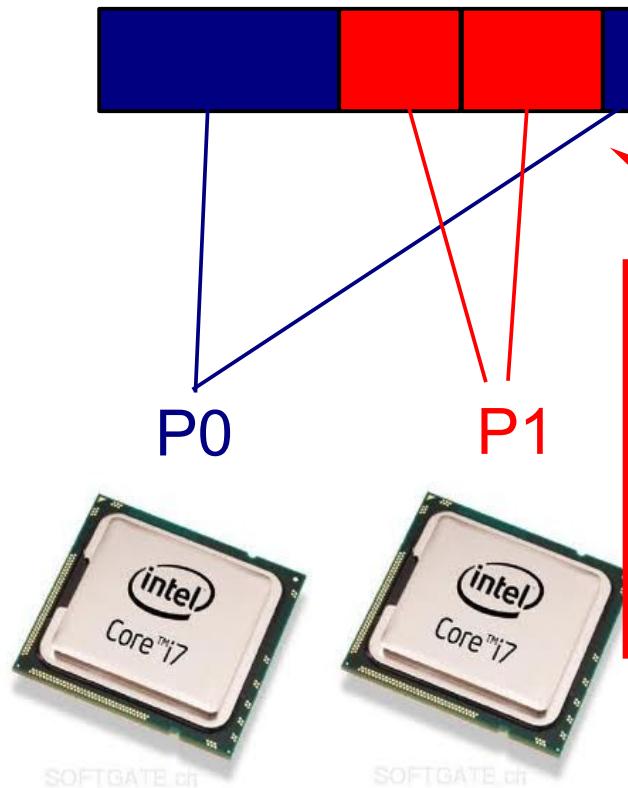


SOFTGATE.ch

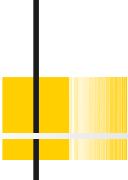
# Load Balance I



# Load Balance I



Find the partition-to-processor assignment such that the maximum number of sites per processor is minimized → this is NP-hard



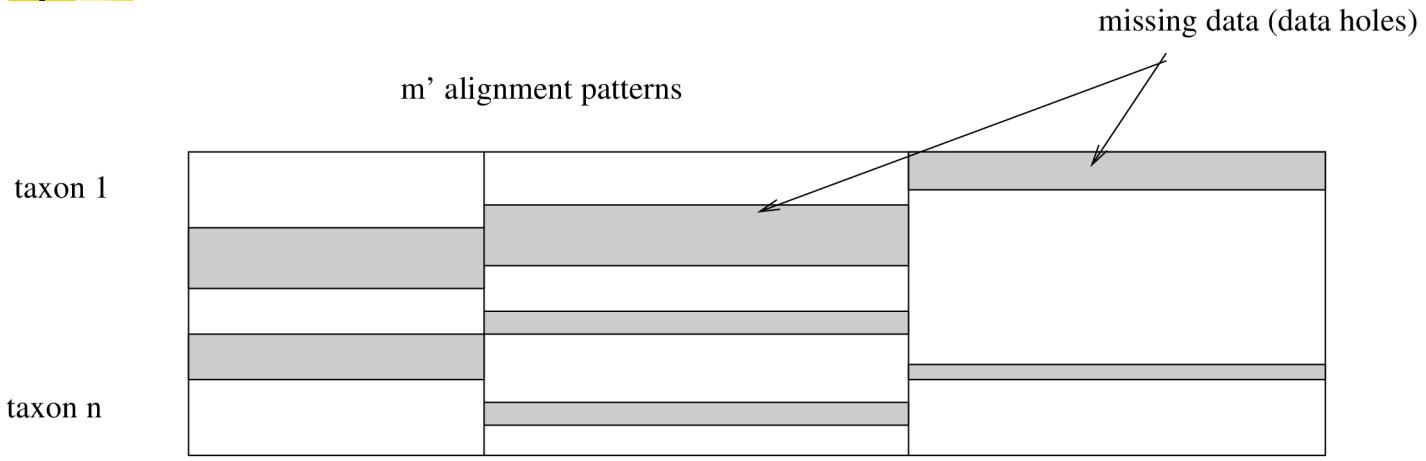
# Load Balance I

---

- The **multiprocessor job scheduling problem** in phylogenetics
  - Problem when #partitions  $>>$  #cores
  - Tested per-site data distribution versus Longest Processing Time (LPT) heuristics
  - 25 taxa, 220,000 sites, 100 genes
    - GAMMA model
      - naïve:      **613** secs
      - LPT:          **550** secs
    - CAT model
      - naïve:      **298** secs
      - LPT:          **127** secs
  - Larger protein dataset under GAMMA: 10-fold performance improvement!

J. Zhang, A. Stamatakis: "The Multi-Processor Scheduling Problem in Phylogenetics", 11th IEEE HICOMB workshop (in conjunction with IPDPS 2012).

# Partitioned Branch Lengths & other parameters



partition 0



separate estimate of  
Q-Matrix  
alpha-shape  
Branch Lengths

partition 1



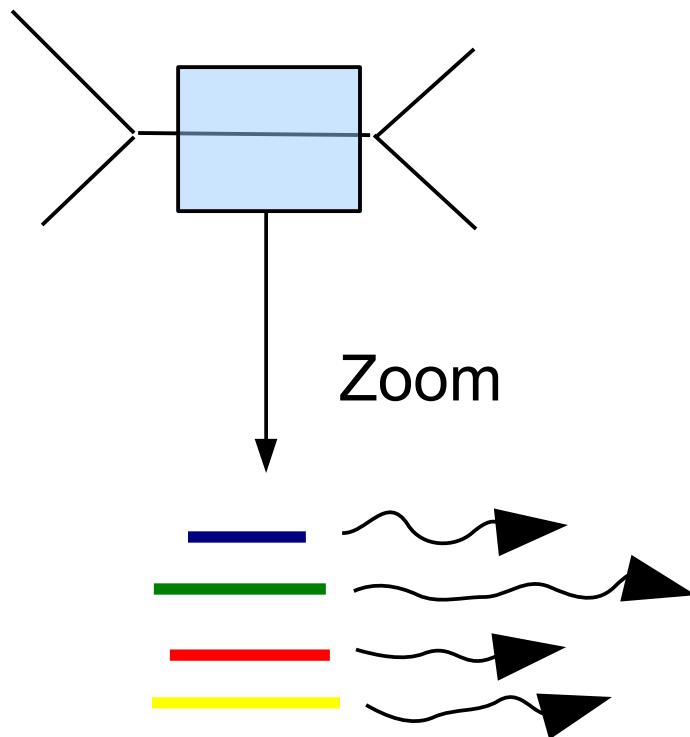
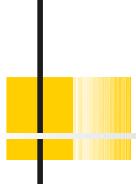
separate estimate of  
Q-Matrix  
alpha-shape  
Branch Lengths

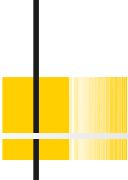
partition 2



separate estimate of  
Q-Matrix  
alpha-shape  
Branch Lengths

# Load-Balance II

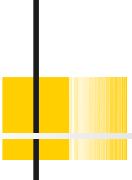




# Synchronization Points

---

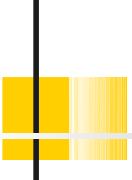
- Assume 10 branches
- Each branch requires 10 Newton-Raphson Iterations
- Each NR Iteration requires a synchronization via a reduction operation
- One branch/partition at a time: 100 sync. points, less work (only one partition) per sync. point
- All branches concurrently: 10 sync. points, more work per sync. point
- Branches will need distinct number of operations
- Add convergence state → bit vector



# Synchronization Points

---

Org1 AC GT  
Org2 AC TT



# Synchronization Points

---

Org1	AC	GT
Org2	AC	TT

---

---

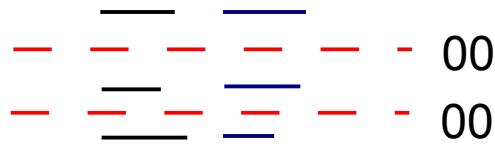
# Synchronization Points

Org1 AC GT  
Org2 AC TT

— — — — 00  
— — — —

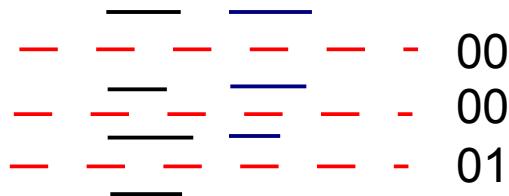
# Synchronization Points

Org1 AC GT  
Org2 AC TT



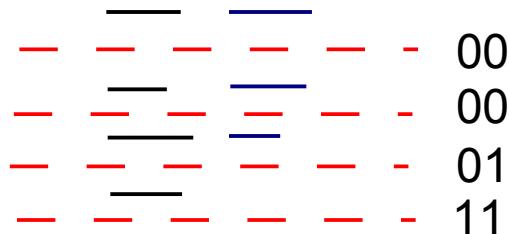
# Synchronization Points

Org1	AC	GT
Org2	AC	TT



# Synchronization Points

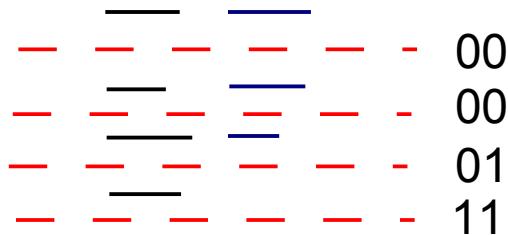
Org1	AC	GT
Org2	AC	TT



In this example: 4 instead of 7 sync points!

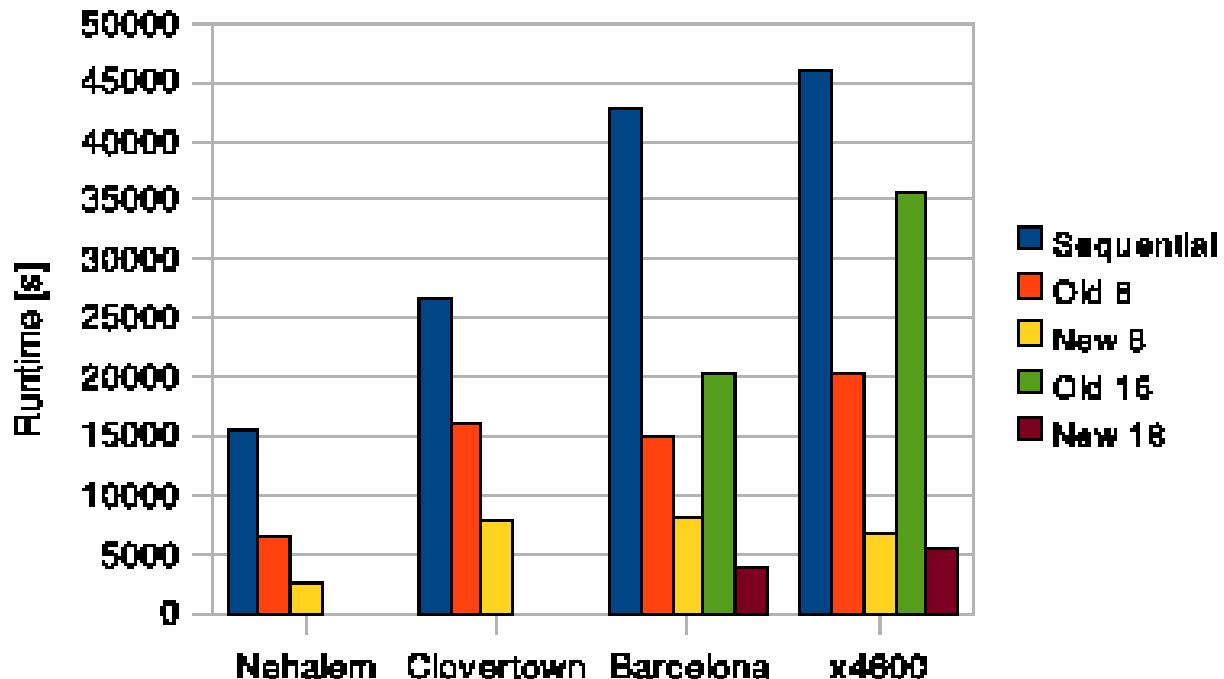
# Synchronization Points

Org1	AC	GT
Org2	AC	TT



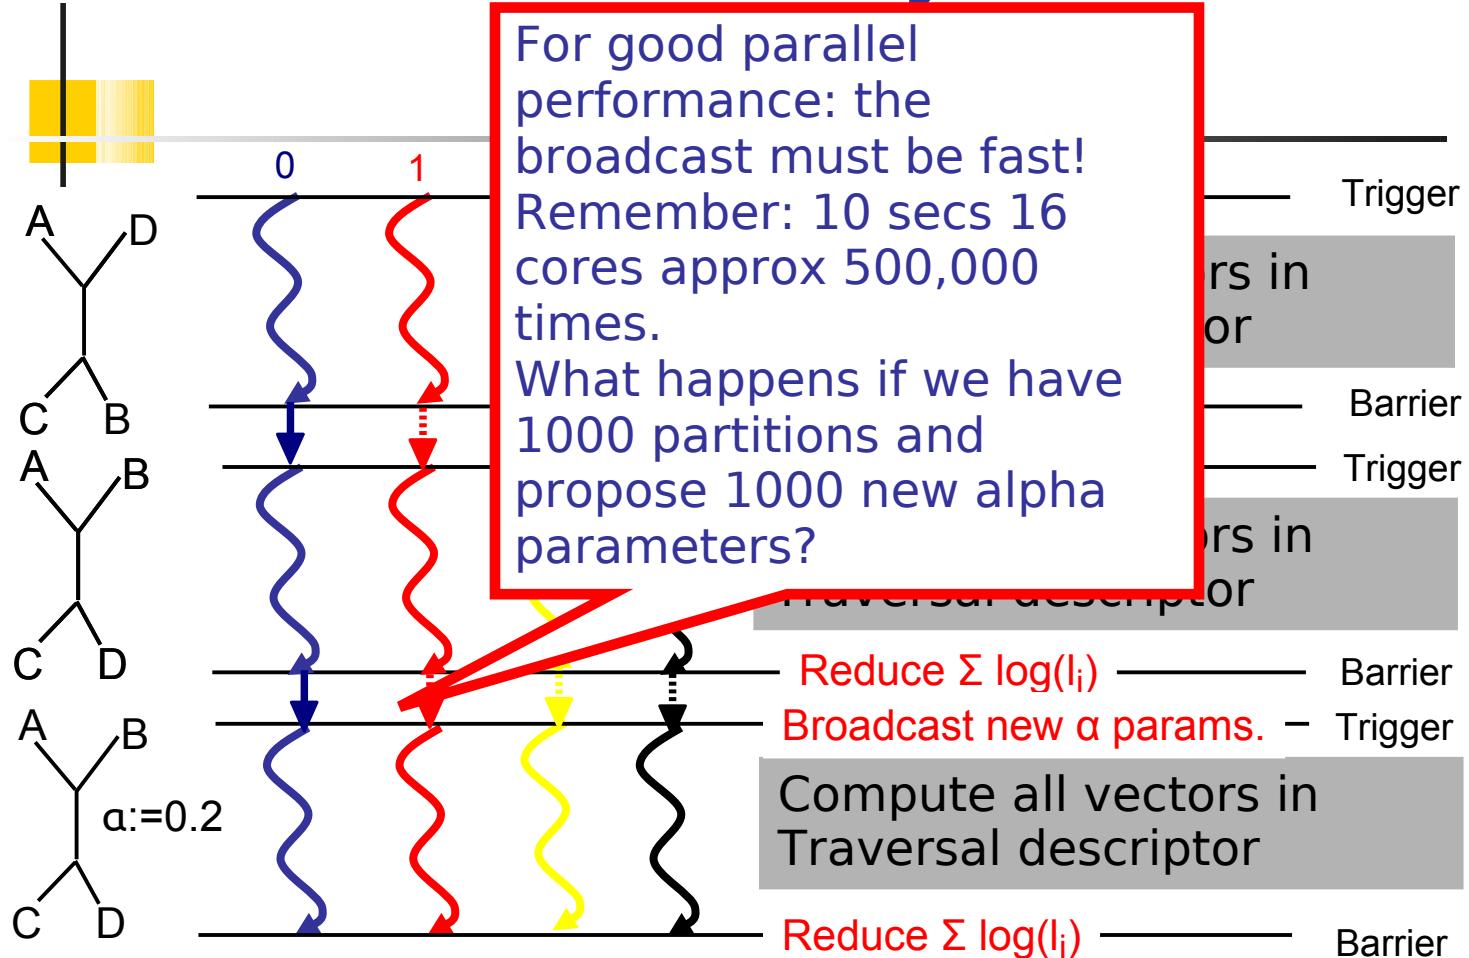
**Implications for Bayesian programs:** Propose e.g., new  $\alpha$  parameters or new branch lengths simultaneously for all partitions!

# Load Balance II

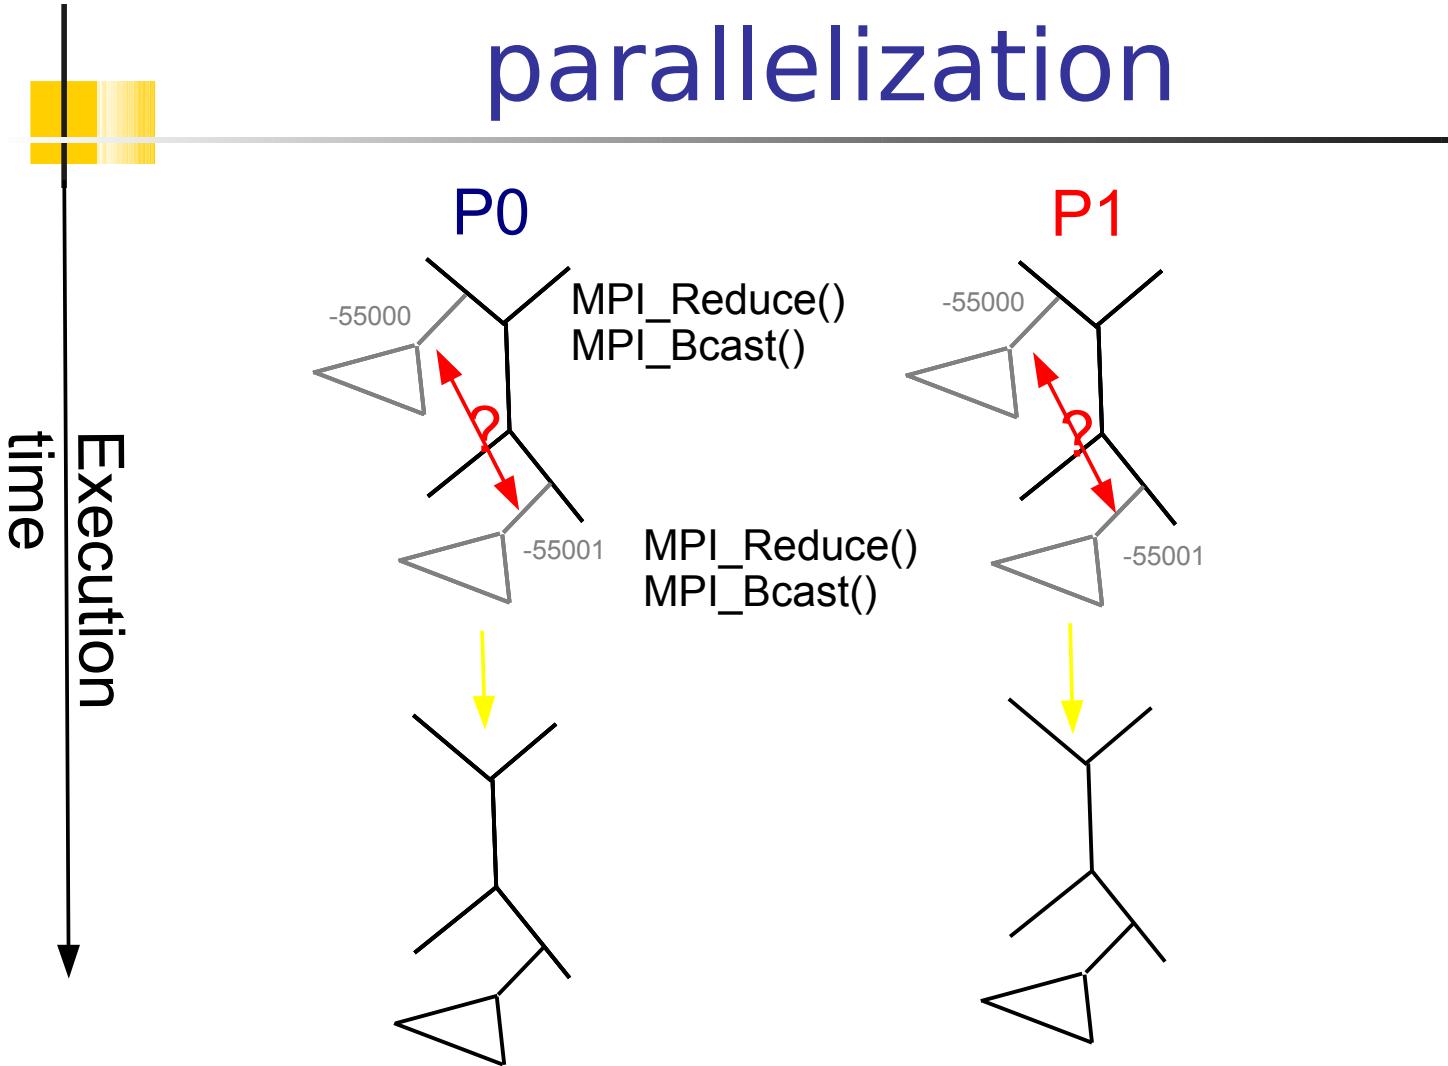


A. Stamatakis, M. Ott: "Load Balance in the Phylogenetic Likelihood Kernel".  
Proceedings of ICPP 2009, Vienna, Austria, September 2009.

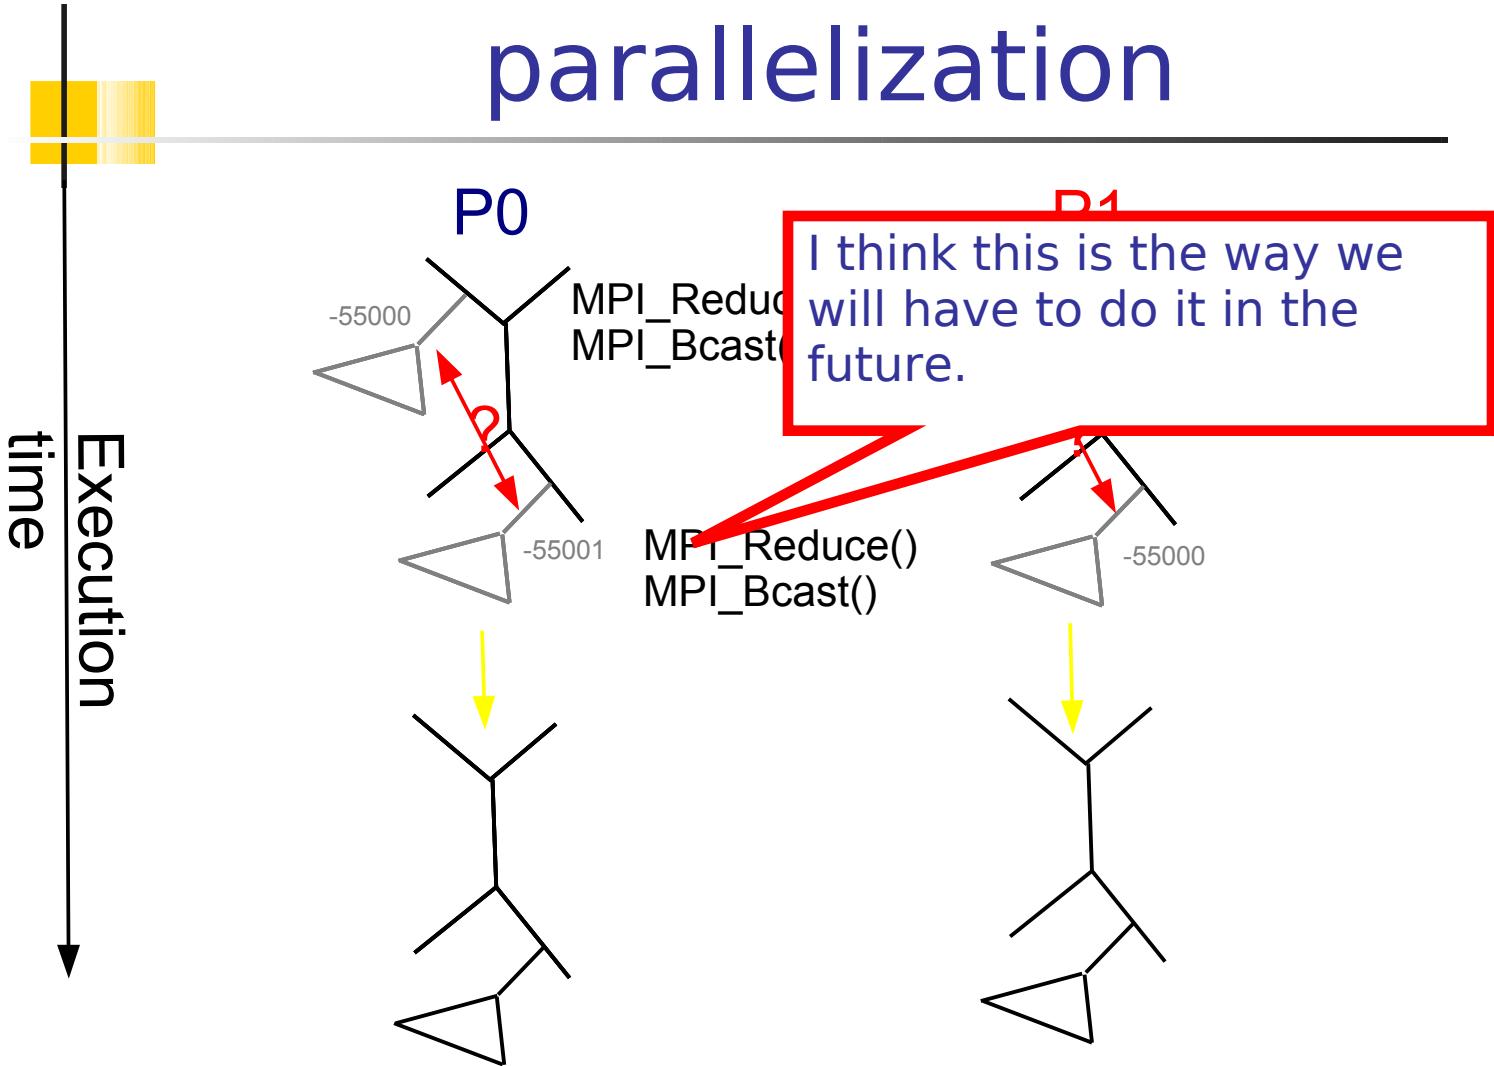
# Classic Fork-Join with

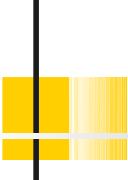


# Alternative MPI parallelization



# Alternative MPI parallelization

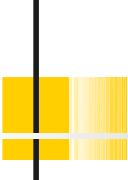




# Outline

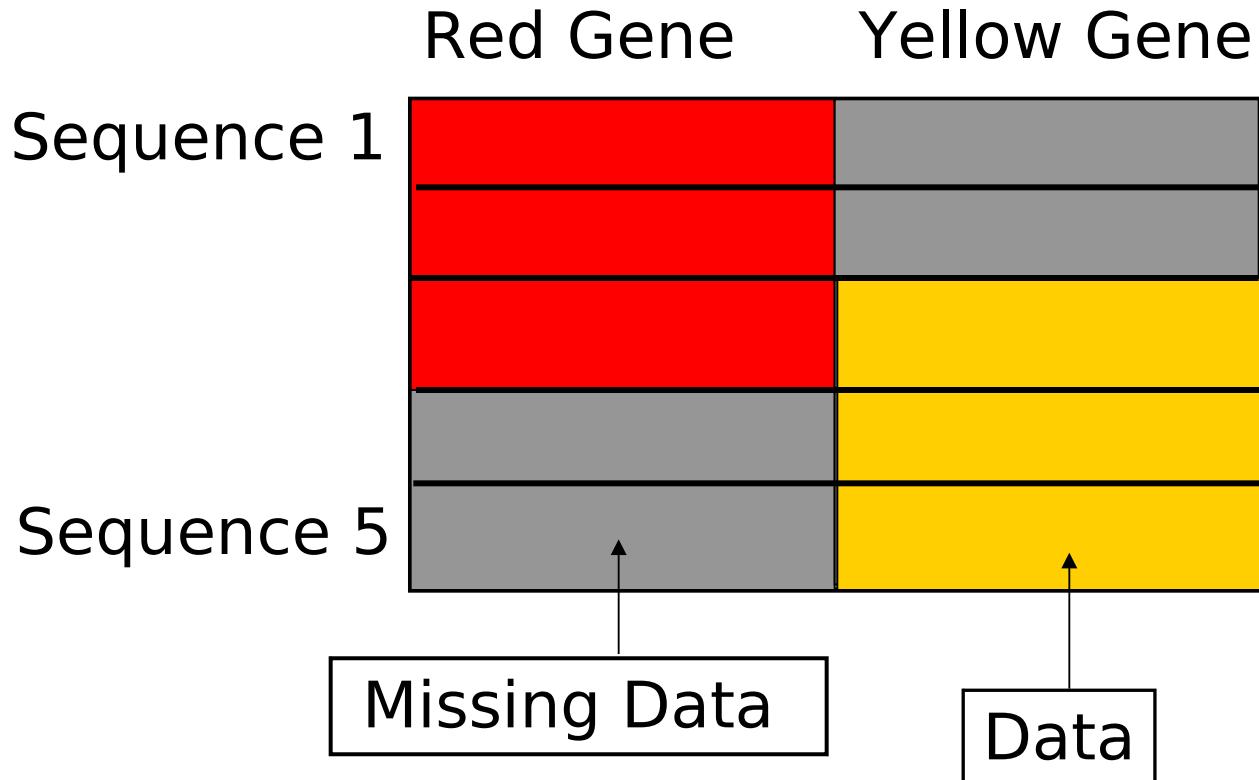
---

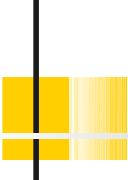
- Computing the Likelihood
- Optimizing & Parallelizing Likelihood Computations
- Saving Memory in Likelihood Computations
- HPC population genetics



# Gappy Phylogenomic Alignments

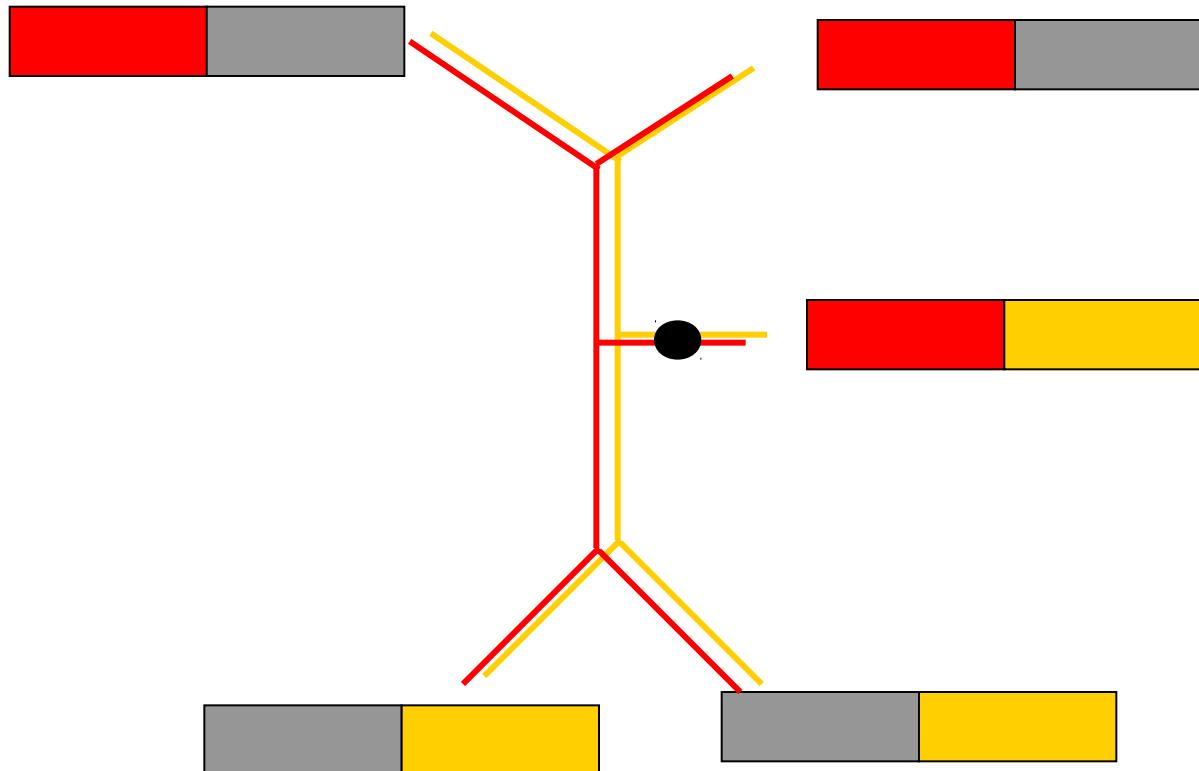
---

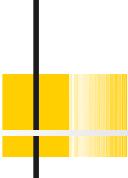




# A Multi-Gene Model

---





# Initial mesh-based Approach

---

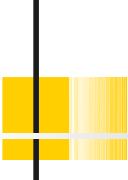
A. Stamatakis, N. Alachiotis: "Time and memory efficient likelihood-based tree searches on gappy phylogenomic alignments", Proceedings of ISMB 2010, Boston, Massachusetts, July 2010. In Bioinformatics, 26(12):i132-i139.

- Good performance
- Only works for per-partition branch length estimates
- Very high code complexity
  - only proof-of-concept implementation
  - production abandoned because too complex
- Does not work for dense datasets!
- How long will datasets stay gappy?

# Initial mesh-based Approach

A. Stamatakis, N. Alachiotis: "Time and memory efficient likelihood-based tree searches on gappy phylogenomic alignments", Proceedings of ISMB 2010, Boston, Massachusetts, July 2010. In Bioinformatics, 26(12):i132-i139.

- Good performance
- Only works on gappy datasets:  
Closely-related theoretical work on gappy datasets:  
M. Sanderson, M. McMahon, M. Steel:  
→ only  
→ produces  
“Terraces in Phylogenetic Tree Space”, *Science* 2011.
- Does not implement the RAxML prprof-of-concept
- How long implementation now prints out the number of terrace moves encountered during a search.



# Easier Memory-saving Techniques

---

- Resurrection of subtree equality vectors for gappy alignments

A. Stamatakis, T. Ludwig, H. Meier, M.J. Wolf: "AxML: A Fast Program for Sequential and Parallel Phylogenetic Tree Calculations Based on the Maximum Likelihood Method". In Proceedings of 1st IEEE Computer Society Bioinformatics Conference (CSB2002), 21–28, Palo Alto, California, August 2002.
- Out-of-core execution/external memory algos

F. Izquierdo-Carrasco, A. Stamatakis: "Computing the Phylogenetic Likelihood Function Out-of-Core", accepted for publication at IEEE HICOMB 2011 workshop (held in conjunction with IPDPS 2011), Anchorage, USA, May 2011.
- Trading memory for computations

F. Izquierdo-Carrasco, J. Gagneur, A. Stamatakis: "Trading Memory for Running Time in Phylogenetic Likelihood Computations", Bioinformatics 2012 conference, Vilamoura, Portugal, February 2012.

# Easier Memory-saving Techniques

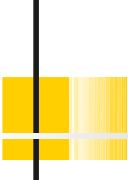
- Resurrection of subtree quality vectors for gappy alignments

A. Stamatakis, T. Ludwig, H. Meier, M.J. Voss, "Parallel Phylogenetic Tree Calculations Based on Subtree Quality Vectors", Proceedings of 1st IEEE Computer Society Bioinformatics Conference, Palo Alto, California, August 2002.

These techniques can be used in all likelihood-based programs!
- Out-of-core execution/external memory algos

F. Izquierdo-Carrasco, A. Stamatakis: "Computing the Phylogenetic Likelihood Function Out-of-Core", accepted for publication at IEEE HICOMB 2011 workshop (held in conjunction with IPDPS 2011), Anchorage, USA, May 2011.
- Trading memory for computations

F. Izquierdo-Carrasco, J. Gagneur, A. Stamatakis: "Trading Memory for Running Time in Phylogenetic Likelihood Computations", Bioinformatics 2012 conference, Vilamoura, Portugal, February 2012.



# Subtree Equality Vectors

## An old idea revisited

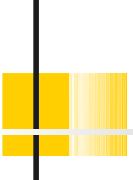
---

- Acceleration of Likelihood-function by detecting (and avoiding recomputation) of identical patterns in subtrees
- Detecting identical patterns may induce high bookkeeping overhead
- Conditional in innermost likelihood kernel loop may perturb prefetching and branch prediction
- Ancient papers

Stamatakis *et al.* “Accelerating Parallel Maximum Likelihood-based Phylogenetic Tree Calculations using Subtree Equality Vectors”, Supercomputing **2002**.

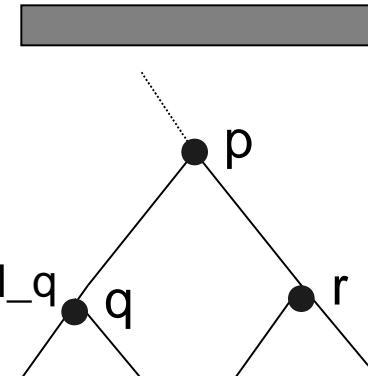
Stamatakis *et al.* “AxML: A Fast Program for Sequential and Parallel Phylogenetic Tree Calculations based on the Maximum Likelihood Method”, **CSB2002**.

Some old paper by Fredrik from the 90ies?

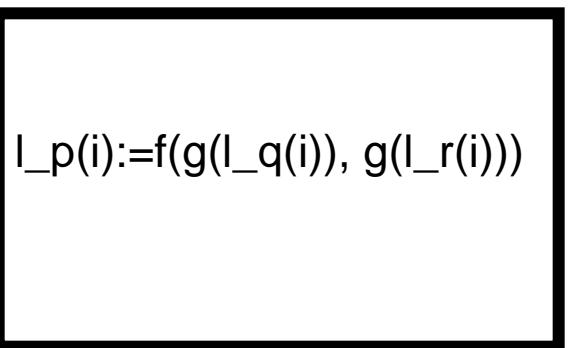
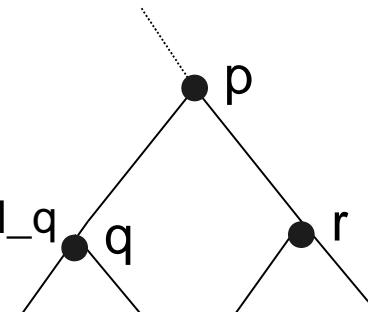


# Subtree Equality Vectors

---



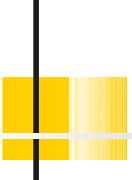
$l_p$


$$l_p(i) := f(g(l_q(i)), g(l_r(i)))$$


$l_q$

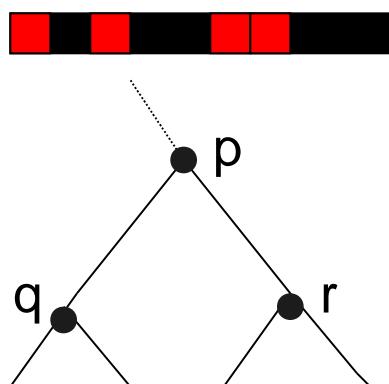


$l_r$

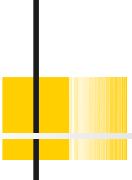


# Subtree Equality Vectors

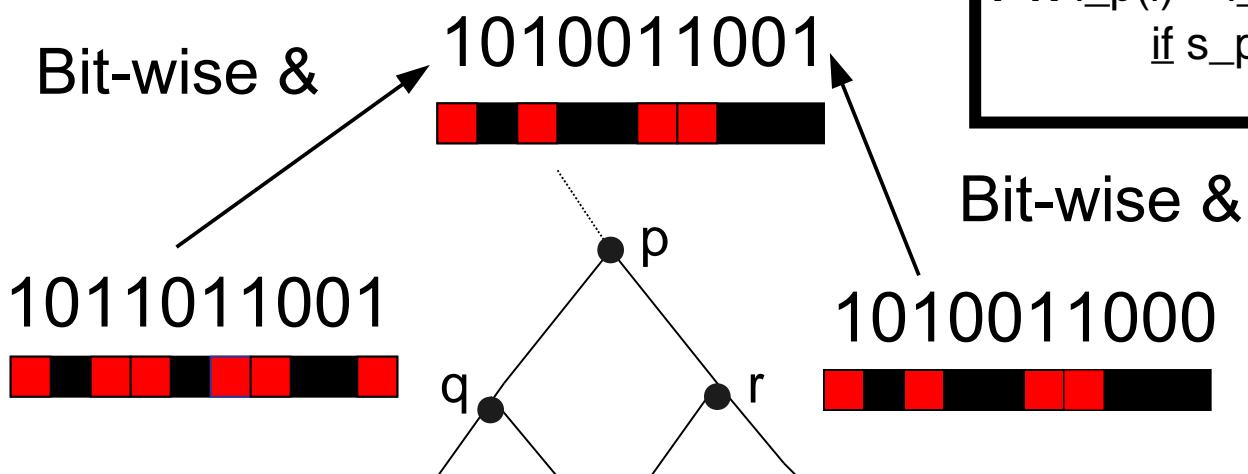
---


$$l_p(i) := f(g(l_q(i)), g(l_r(i)))$$

**P1:**  $l_p(i) = l_p(j)$   
if  $s_p(i) = s_p(j)$



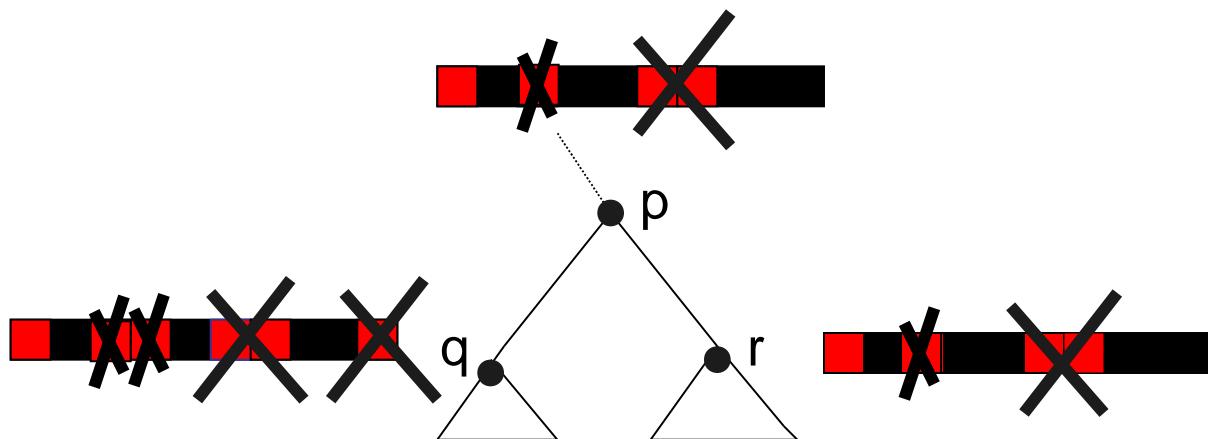
# Subtree Equality Vectors



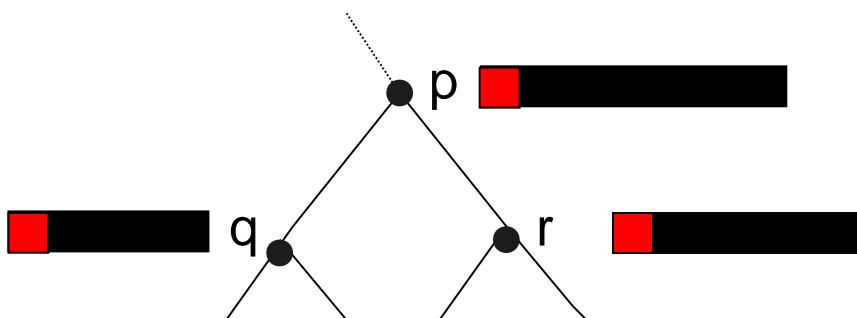
$$l_p(i) := f(g(l_q(i)), g(l_r(i)))$$

**P1:**  $l_p(i) = l_p(j)$   
if  $s_p(i) = s_p(j)$

# Subtree Equality Vectors Implementation Option 1



# Subtree Equality Vectors Implementation Option 2



# Initial Results

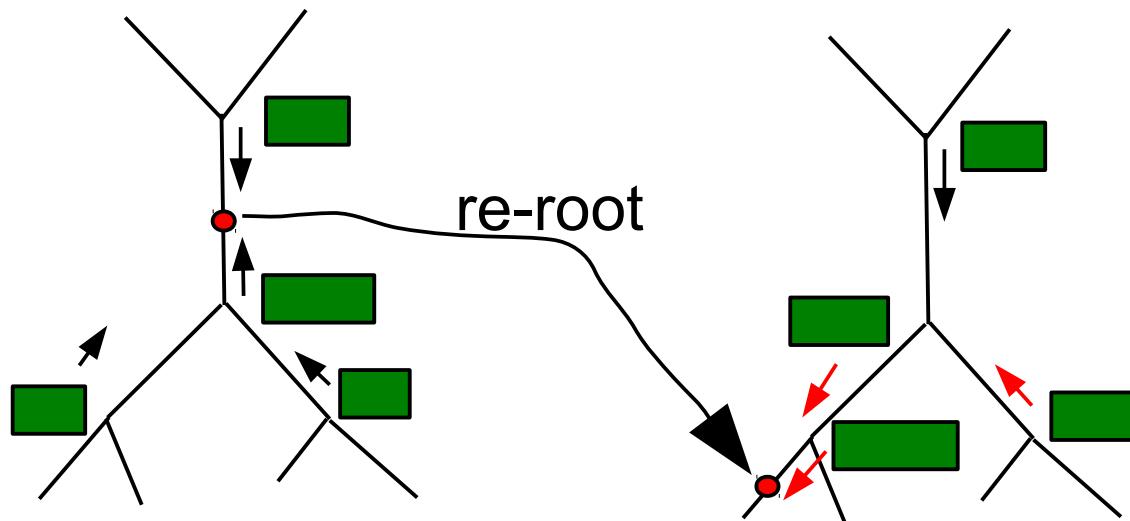
## Dataset Gappyness $\approx 80\%$

---

37831 taxa			
	SEVs	SEVs with memory saving	standard
Runtime (s)	4125.1	4116.8	6541.1
Memory (GB)	42	15	41
LogLikelihood	-5528590	-5528590	-5528590

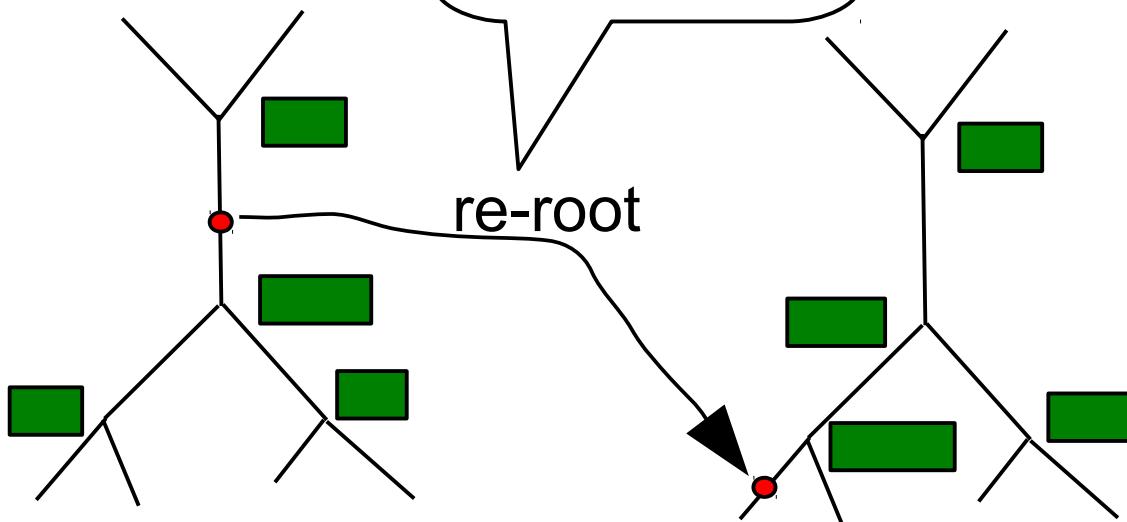
55593 taxa			
	SEVs	SEVs with memory saving	standard
Runtime (s)	7145.2	8095.1	11181.4
Memory (GB)	67	29	67
log likelihood	-7059556	-7059556	-7059556

# Memory Allocation SEVs



# Memory Allocation SEVs

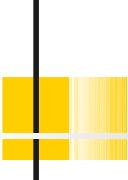
Need to free() and  
malloc() the vectors  
again



# Frequent free() and malloc() in Pthreads version

---

- Can lead to performance degradation → global lock
- Solution: use multi-core allocators
  - AMD Magny 48-core server 256GB
  - 33,000 species, 8000 sites
  - Standard malloc(): 33,400 secs
  - Facebook jemalloc(): 18,200 secs

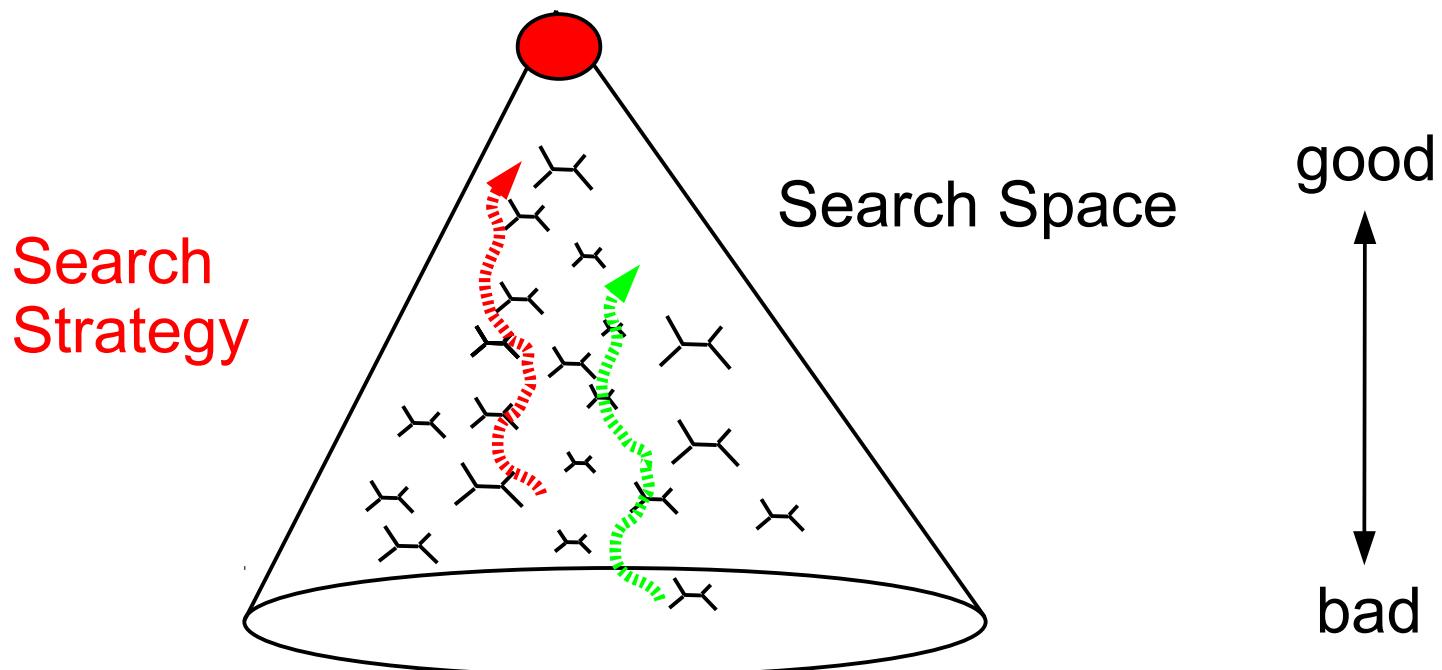


# Subtree Equality Vectors Summary

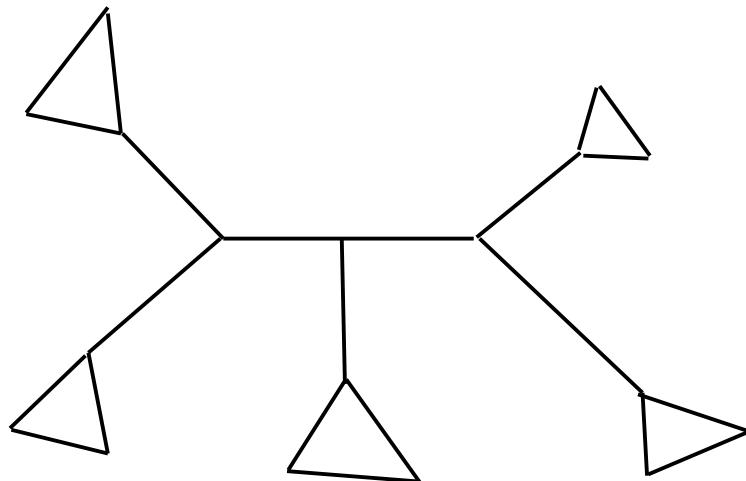
---

- Memory savings depend on
  - Proportion of missing data in input
- Speedups depend on
  - Flops saved versus bookkeeping overhead
  - Implementation of malloc() and free() on multicores when memory saving is used in the Pthreads version  
≈ 50% faster with dedicated allocators

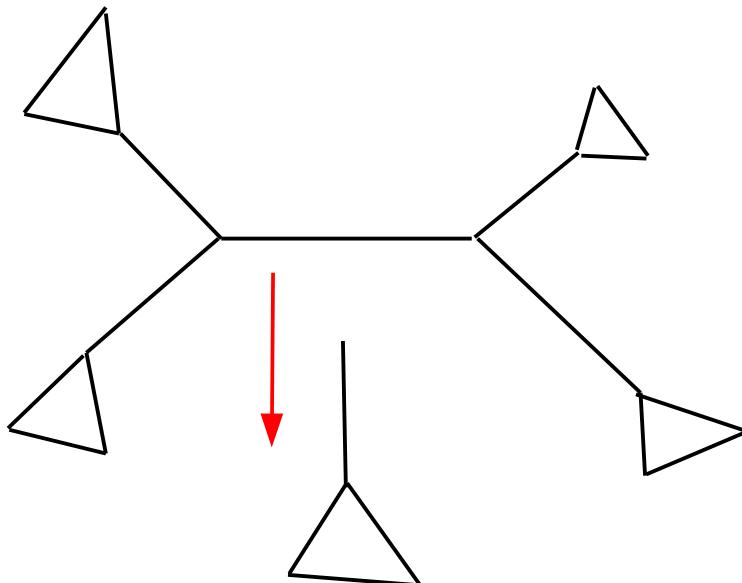
# Search Strategies ML Analyses



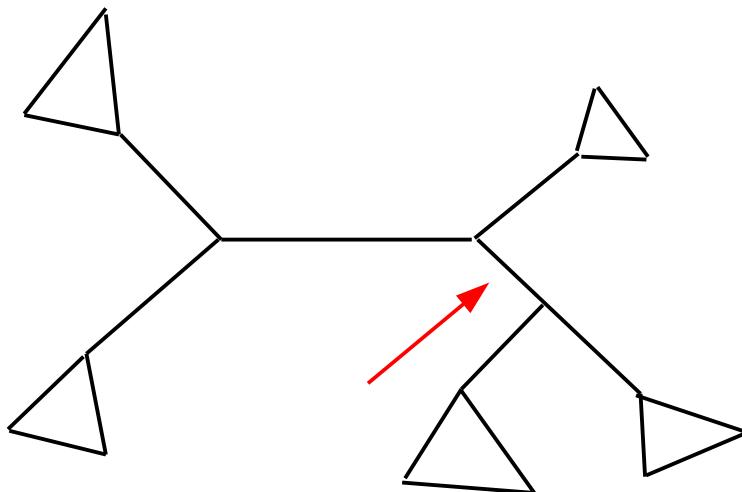
# Tree Search



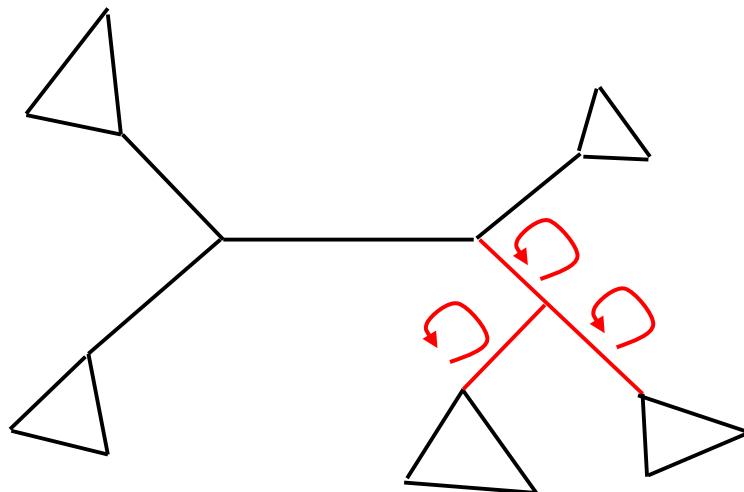
# Tree Search



# Tree Search

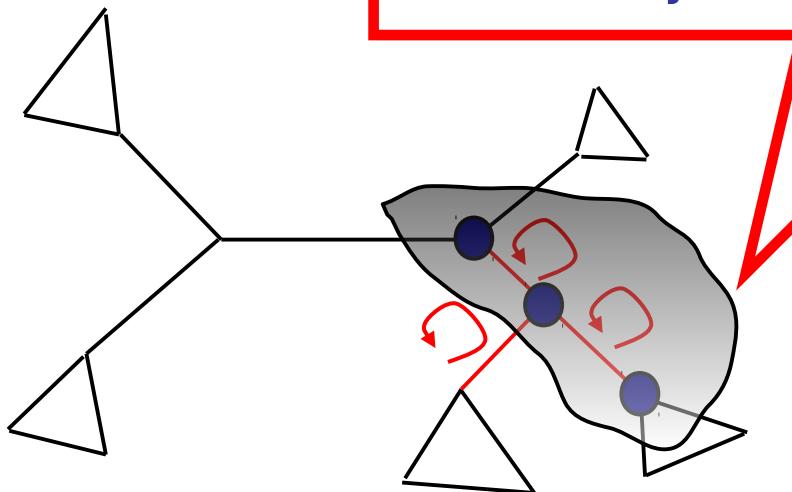


# Tree Search



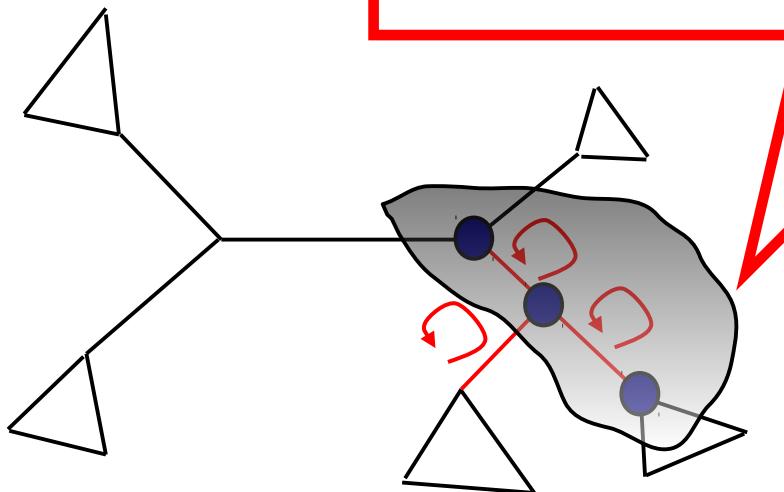
# Tree Search

Only three vectors are affected by the change!

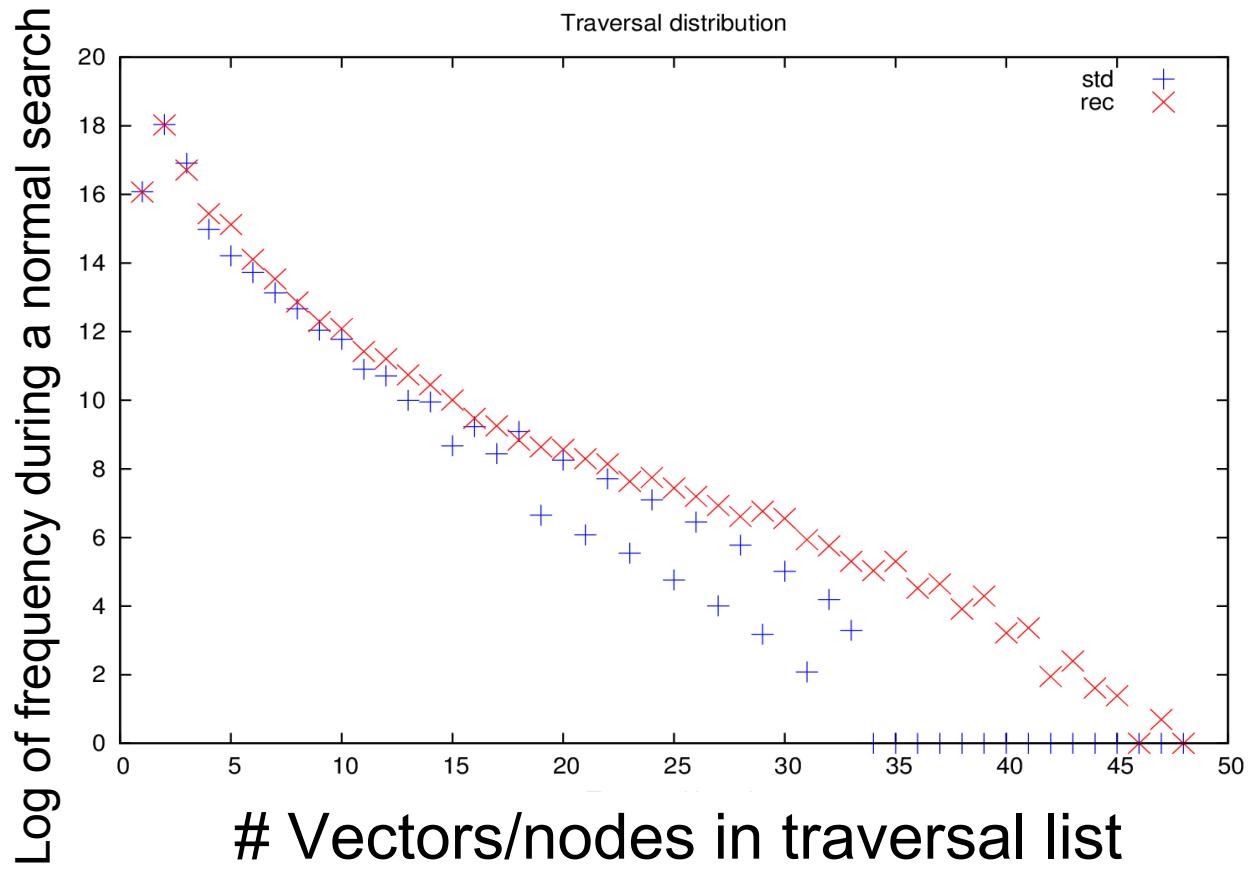


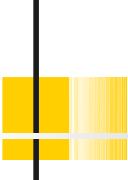
# Tree Search

Most operations are local  
to the tree!



# Traversal Lengths



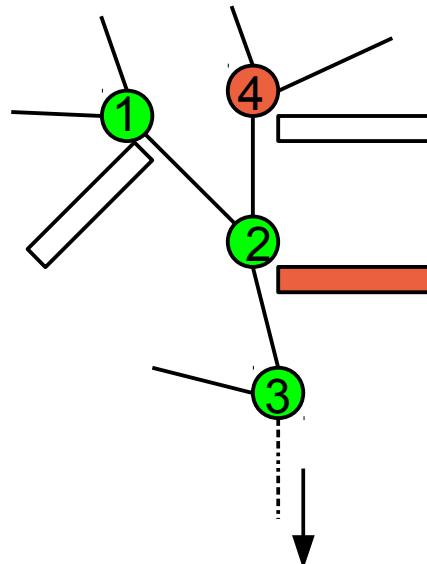


# Out-of-core

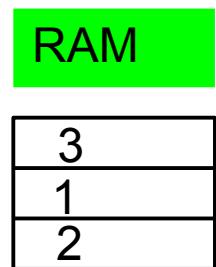
---

- Likelihood computations exhibit high data locality
  - mostly short local traversals
  - expect low miss rate

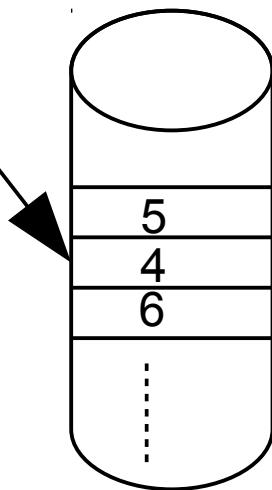
# Out of core



Toward  
the root



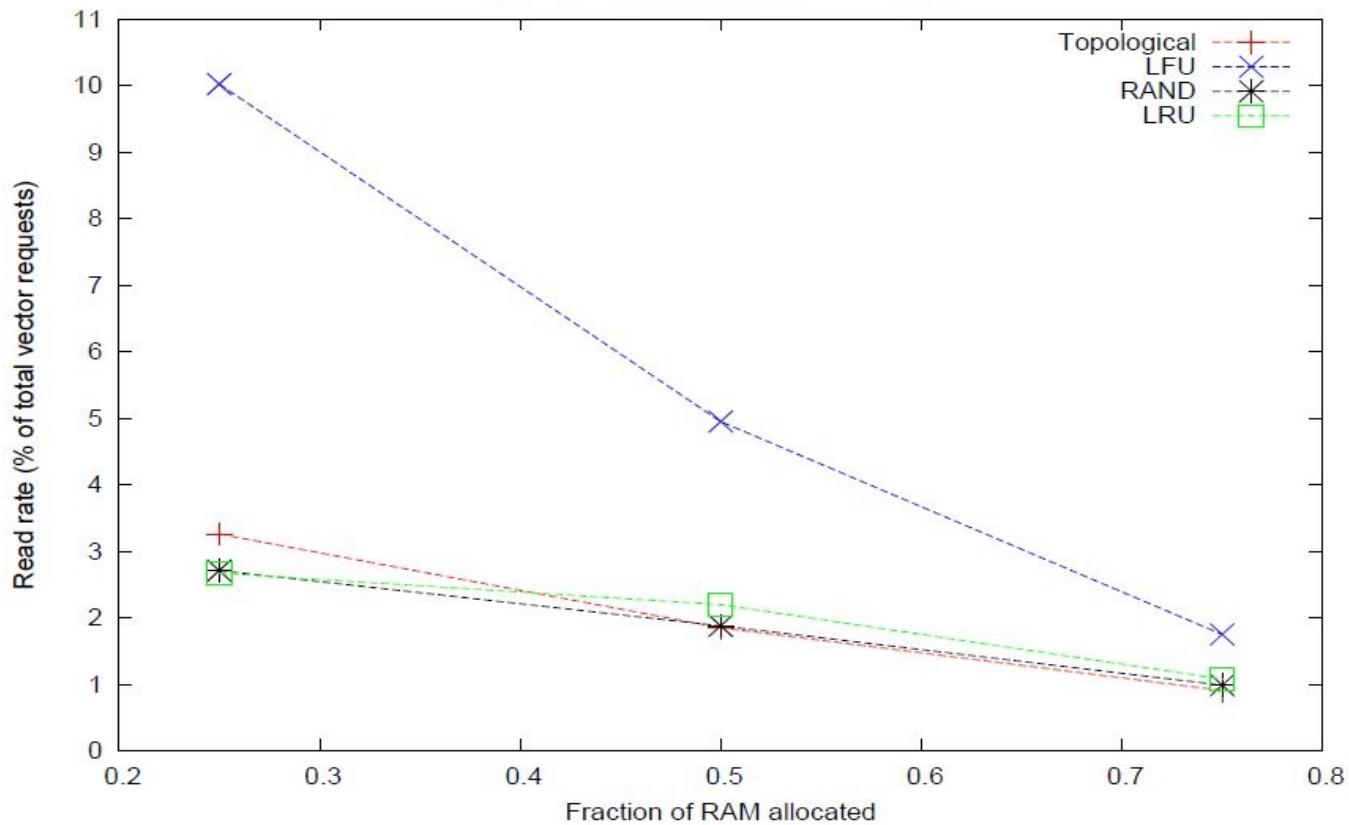
DISK (~100  
MB/s)



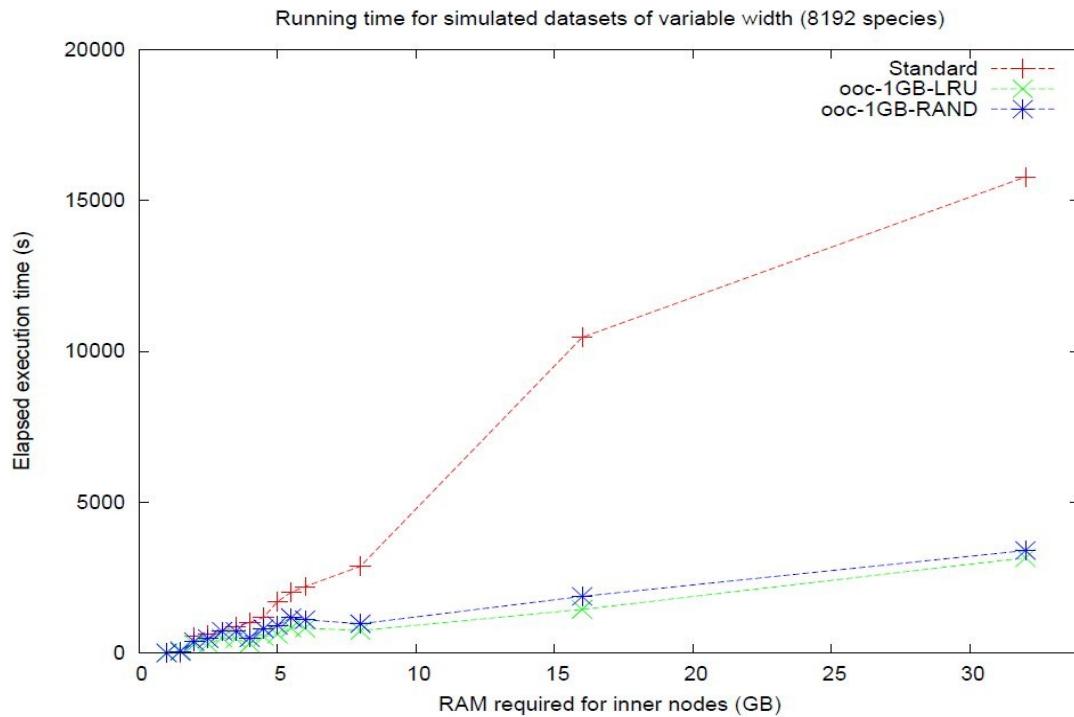
Binary  
file(s)

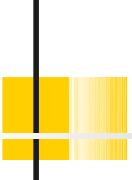
# Miss Rate

Read rate for dataset with 1288 species



# Performance

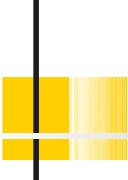




# Summary: Out-of-core

---

- Performance substantially better than for paging
- ... but still disappointing, not practical
- Works for dense datasets!
- **Future Work:** test with solid state disks



# Recomputing Vectors instead of storing them

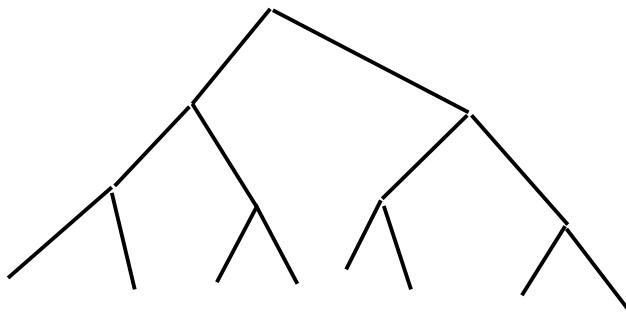
---

- Out of core: high data locality of vectors
- Proof: we only need  **$\log(\#\text{taxa}) + 2$**  vectors in memory to compute the likelihood on any unrooted binary tree
- Performance does not depend on gappyness!
- Can be combined with the orthogonal SEV approach



# Tree Shapes

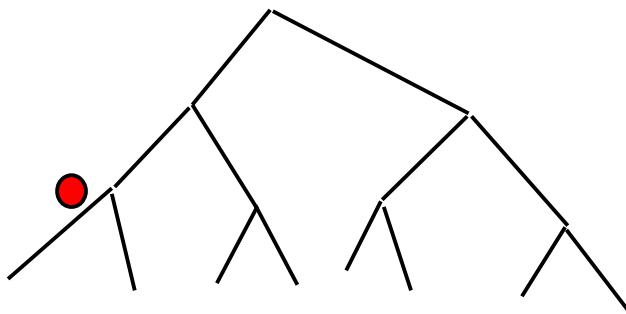
Virtual Root





# Tree Shapes

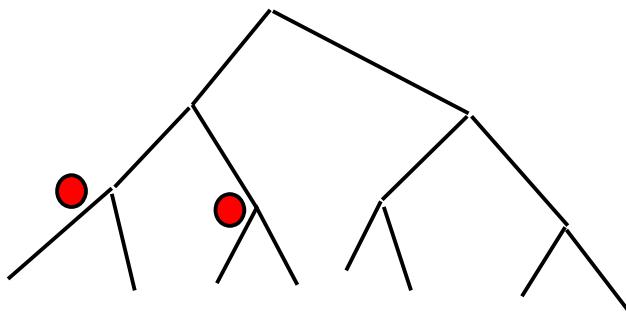
Virtual Root





# Tree Shapes

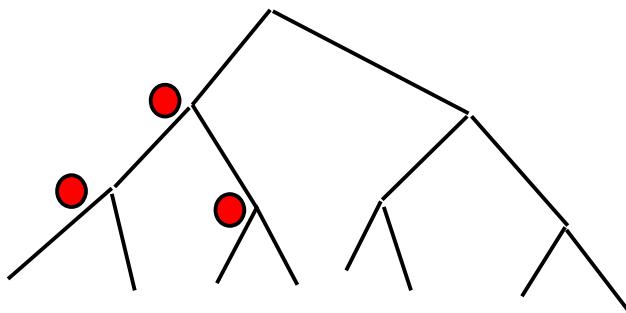
Virtual Root





# Tree Shapes

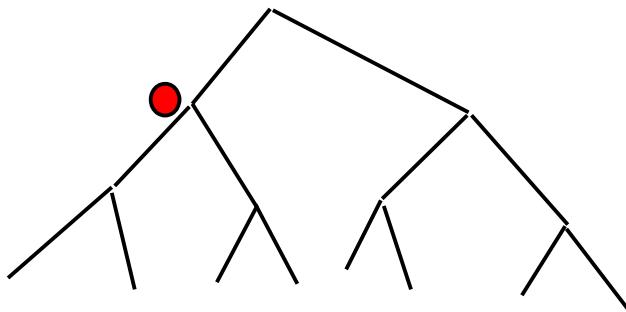
Virtual Root





# Tree Shapes

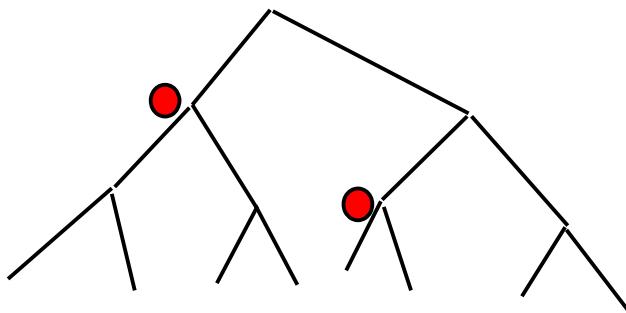
Virtual Root





# Tree Shapes

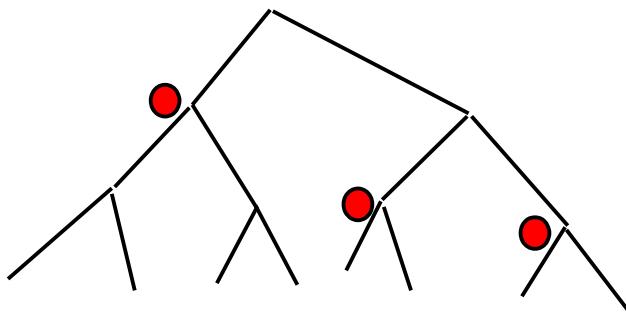
Virtual Root





# Tree Shapes

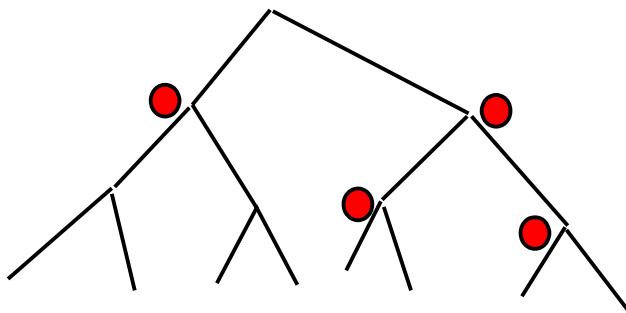
Virtual Root





# Tree Shapes

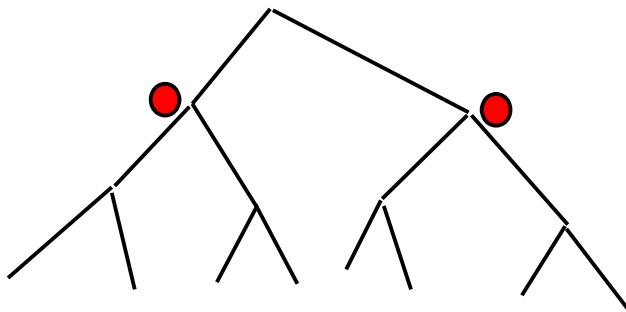
Virtual Root





# Tree Shapes

Virtual Root

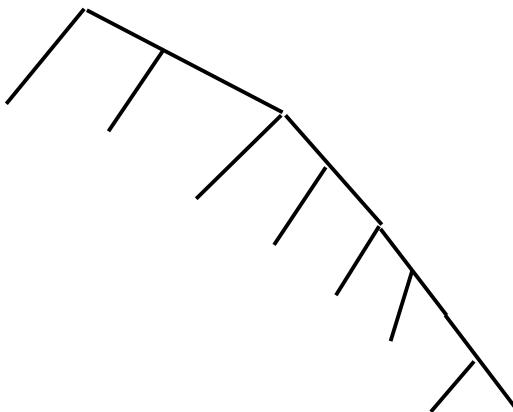


Max number of  
required vectors: 4



# Tree Shapes

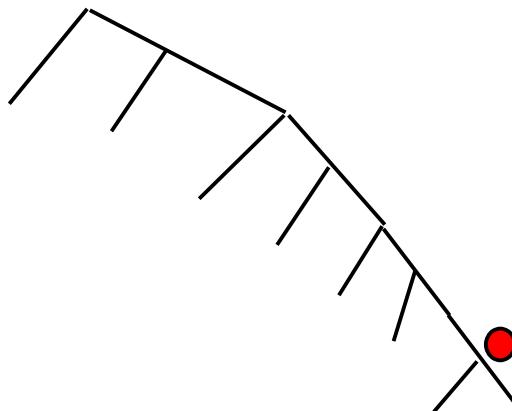
Virtual Root





# Tree Shapes

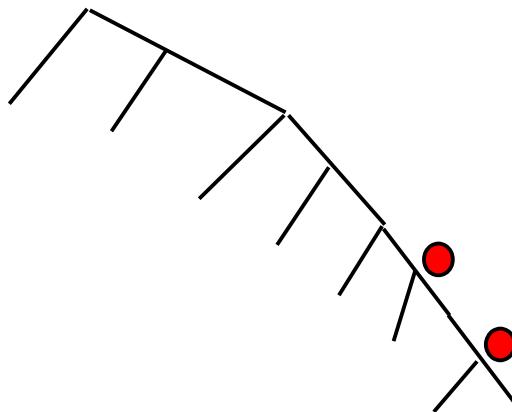
Virtual Root





# Tree Shapes

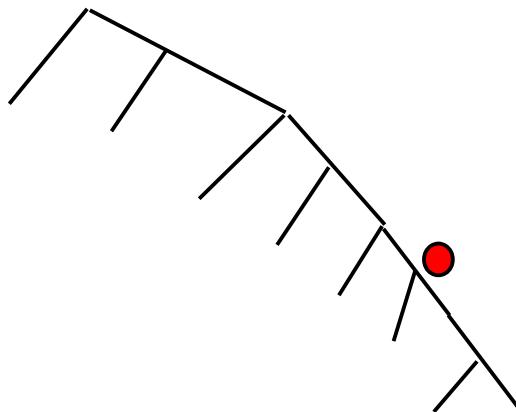
Virtual Root





# Tree Shapes

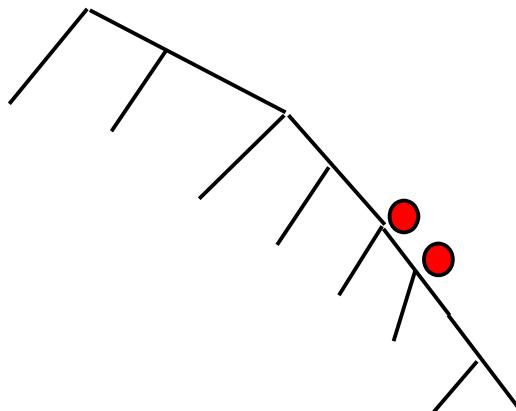
Virtual Root





# Tree Shapes

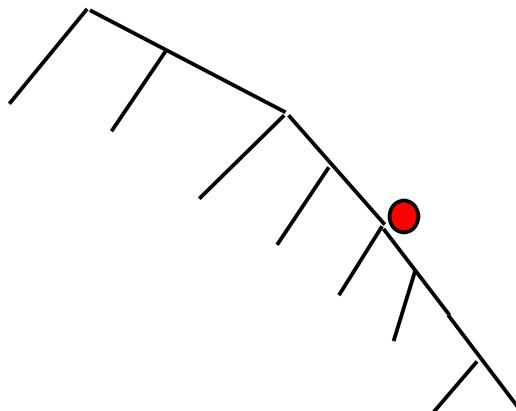
Virtual Root





# Tree Shapes

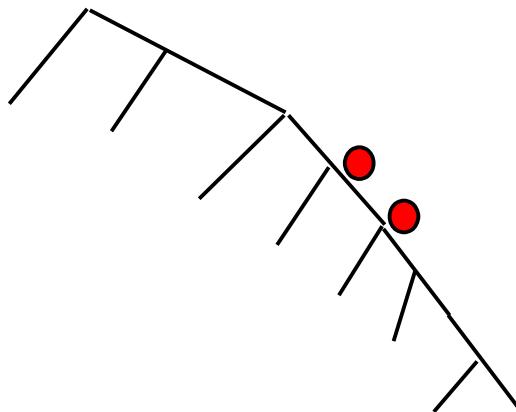
Virtual Root





# Tree Shapes

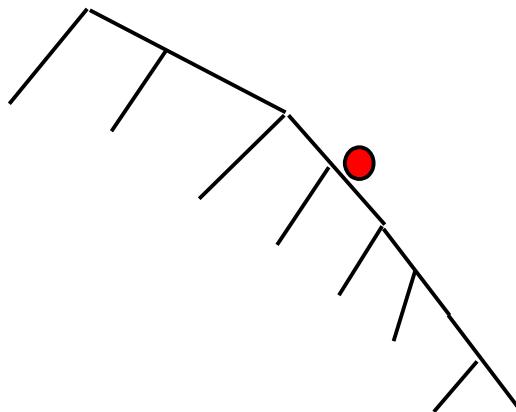
Virtual Root





# Tree Shapes

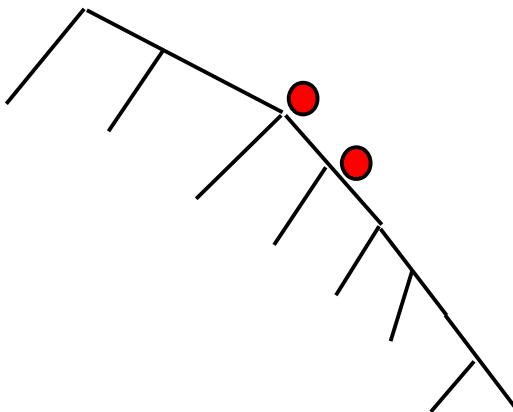
Virtual Root





# Tree Shapes

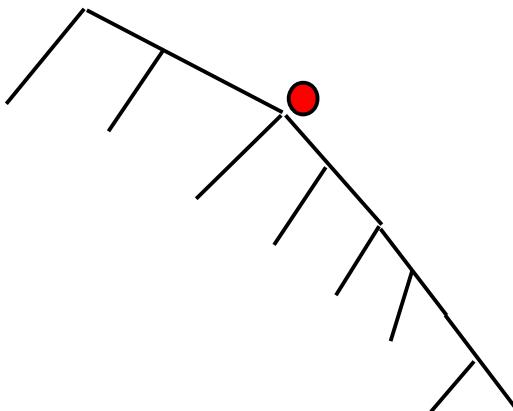
Virtual Root





# Tree Shapes

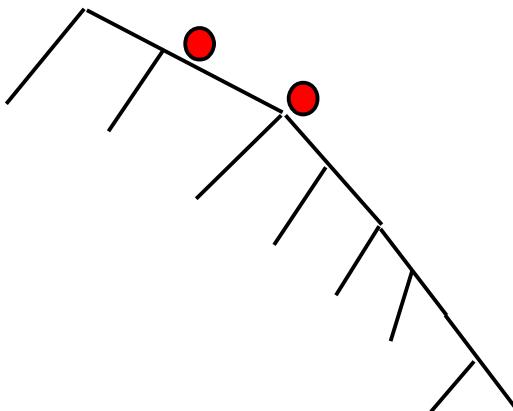
Virtual Root

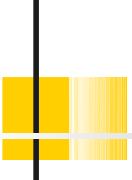




# Tree Shapes

Virtual Root

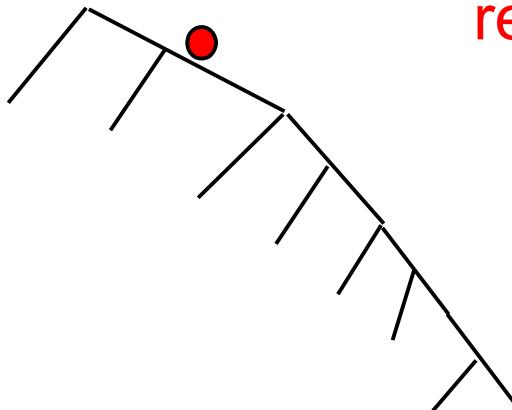




# Tree Shapes

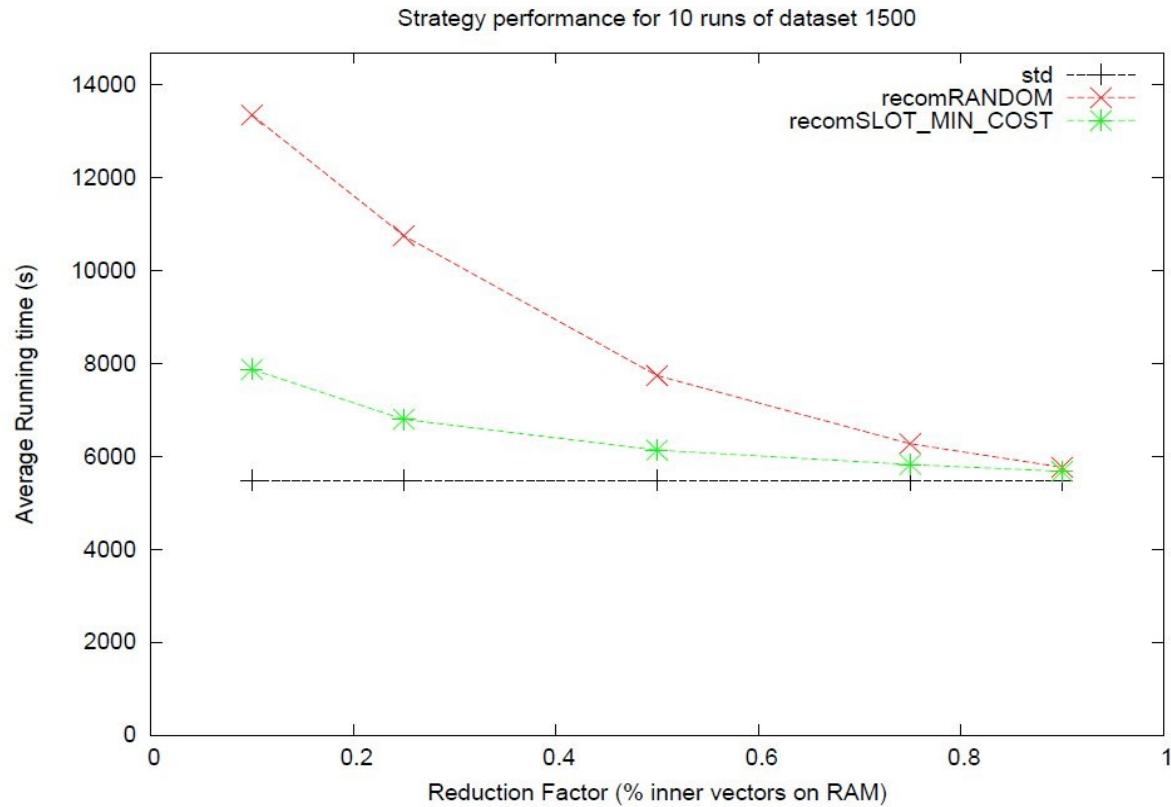
---

Virtual Root



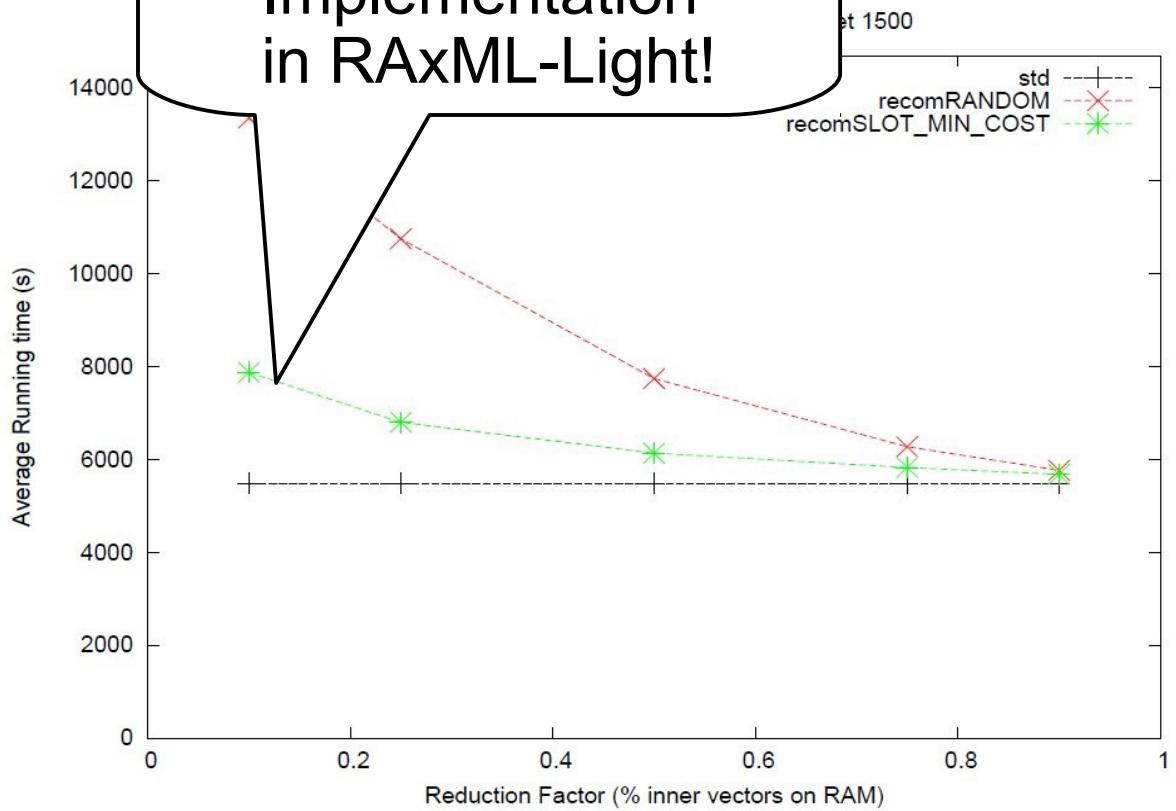
Max number of  
required vectors: 2

# Results

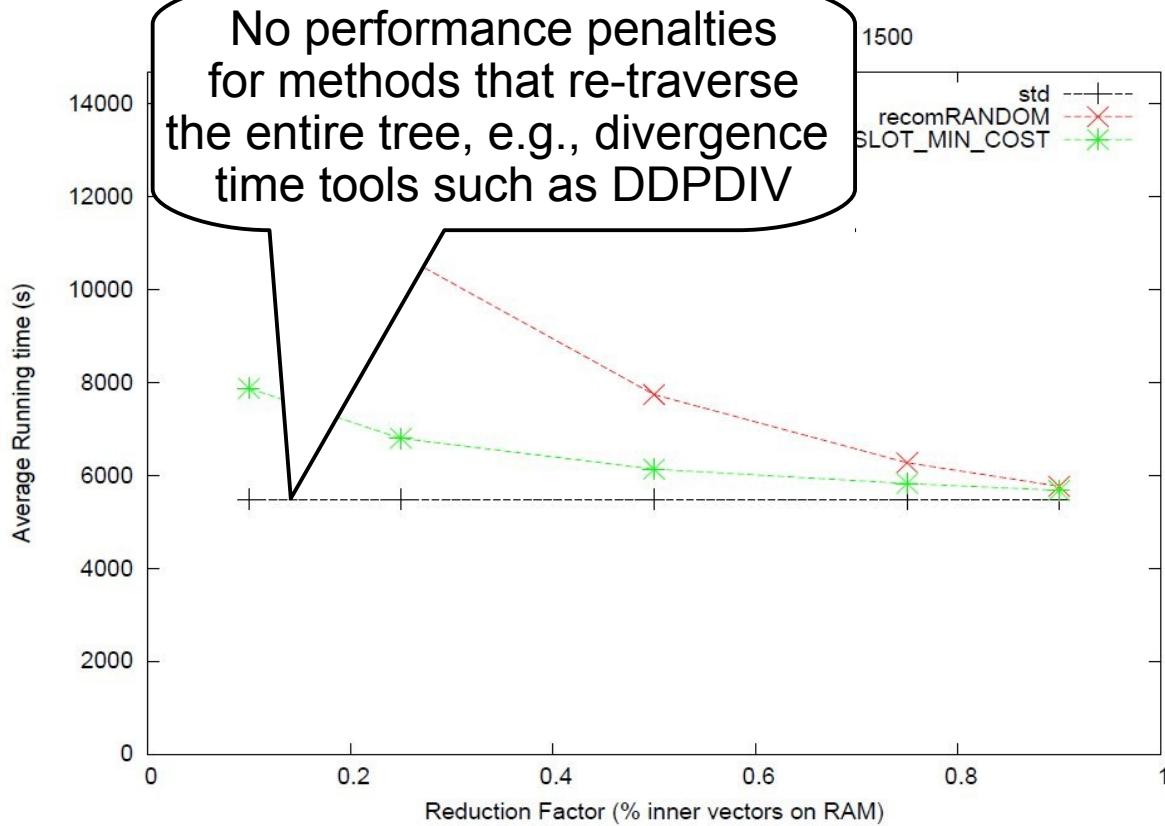


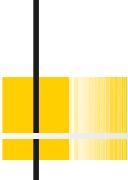
# Results

Production-Level  
Implementation  
in RAxML-Light!



# Results

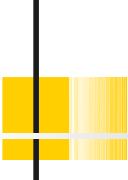




# Outline

---

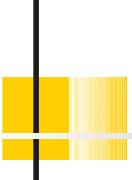
- Computing the Likelihood
- Optimizing & Parallelizing Likelihood Computations
- Saving Memory in Likelihood Computations
- HPC population genetics



# HPC Pop. Gen.

---

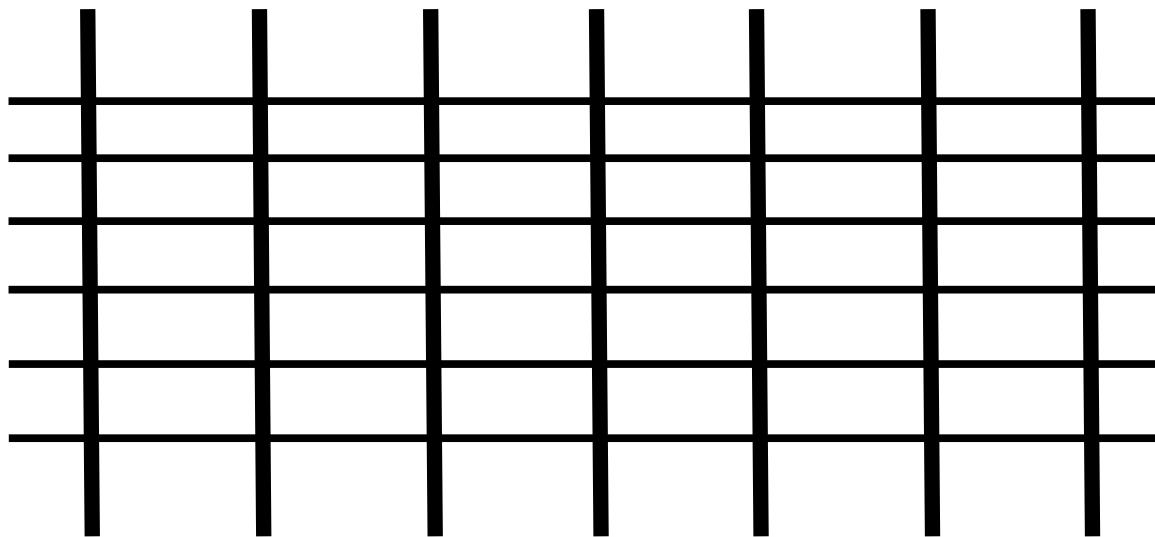
- Omega statistics code
  - Dynamic programming algorithm
  - Up to two orders of magnitude:
    - faster than existing codes
    - less memory
  - Multi-grain parallelization
- Forward in time simulator
  - Algorithmic engineering
  - New data structures
  - Vector instructions
  - 1-2 orders of magnitude faster than SFS
  - Work in progress: Parallelization



# OmegaPlus: Selective Sweep Detection using the $\omega$ Statistic

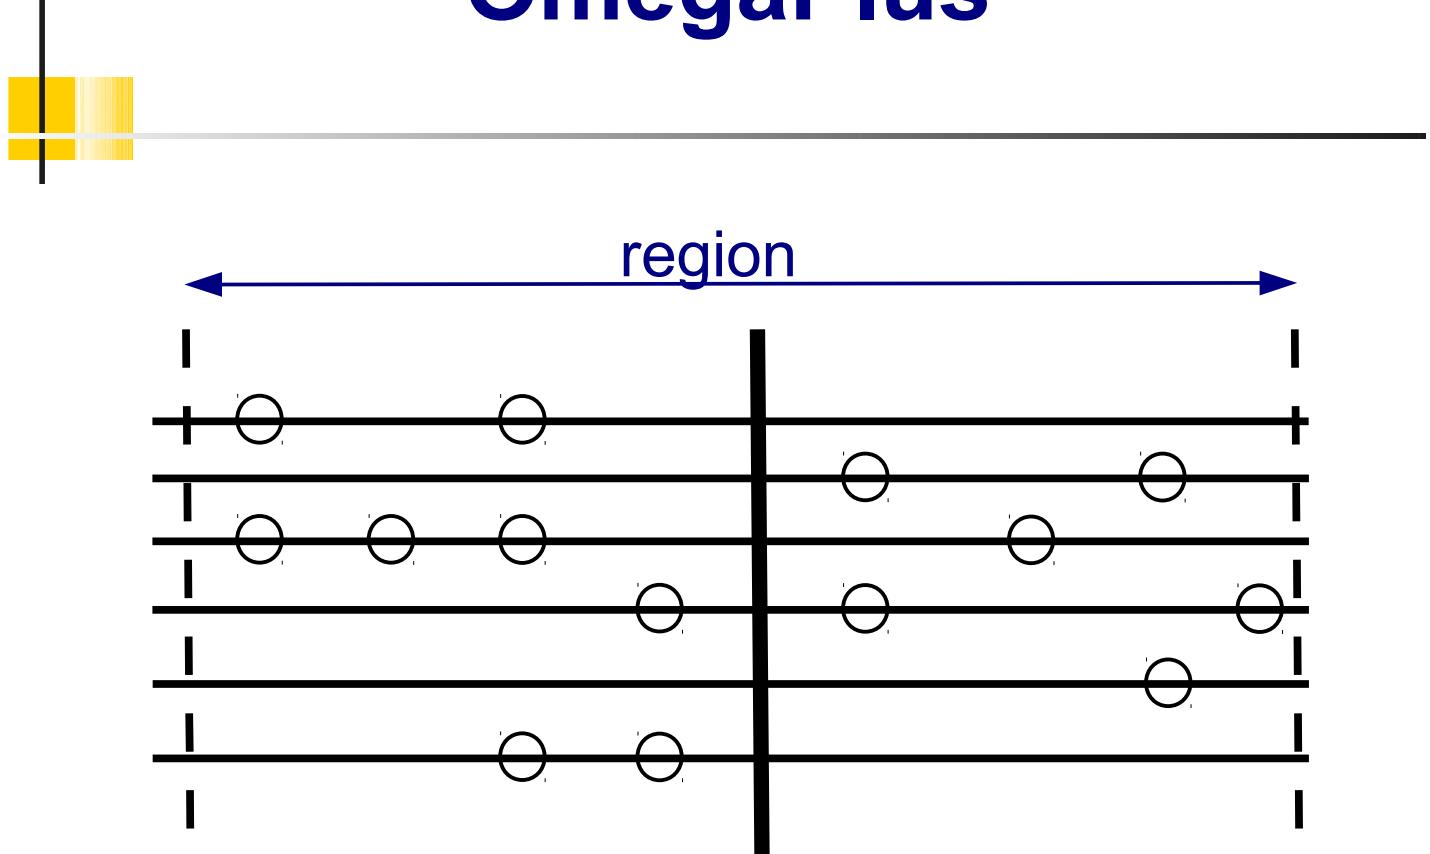
---

OmegaPlus computes the  $\omega$  statistic at N equidistant positions in the alignment.



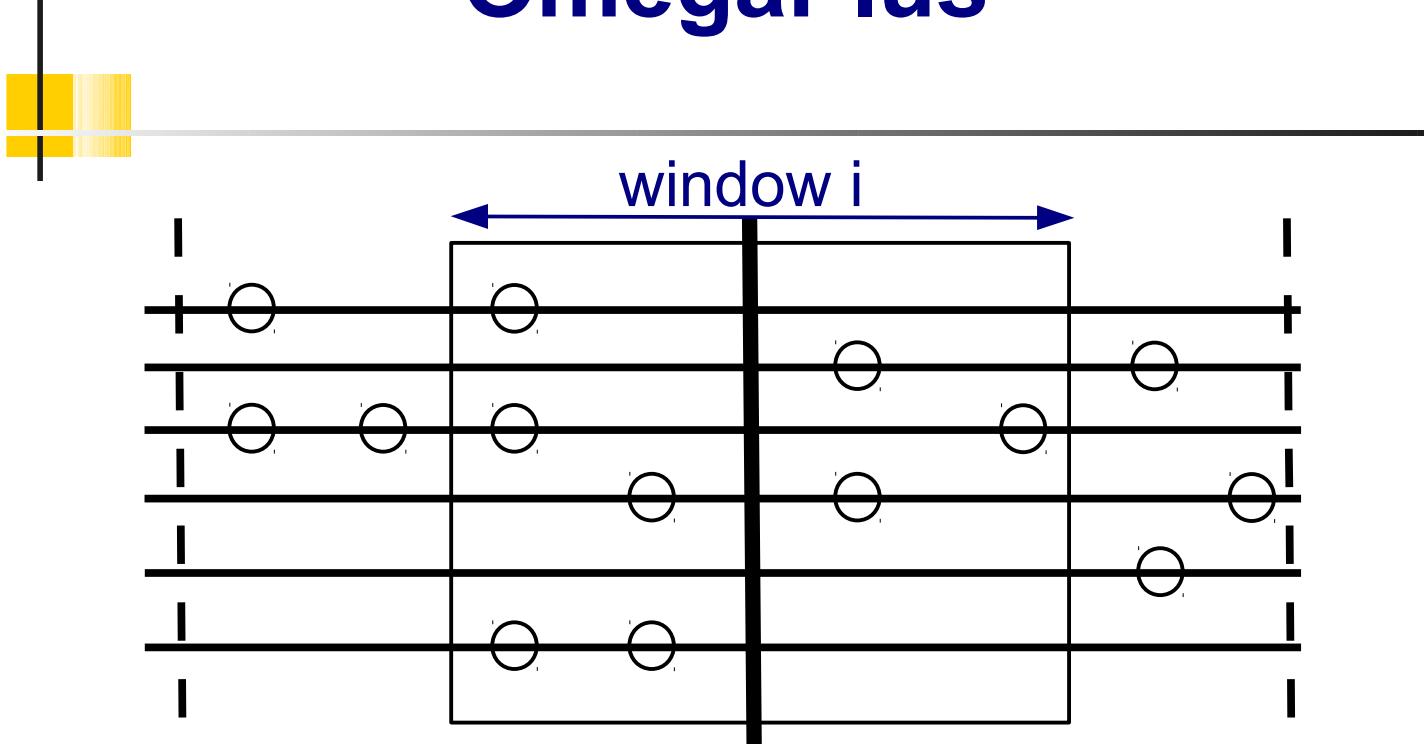
compute the  $\omega$  statistic at this position

# OmegaPlus



For each position, a maximum user-defined region that the sweep may have affected is analyzed.

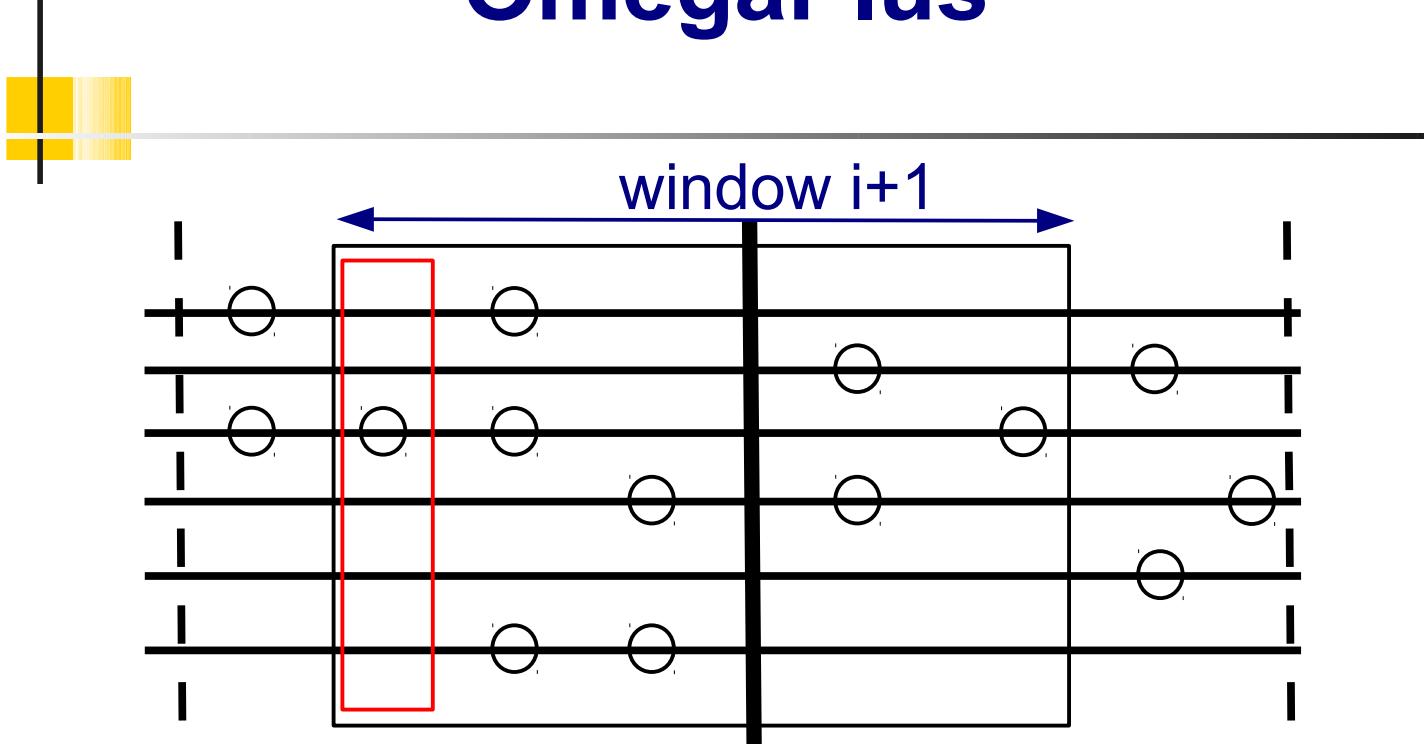
# OmegaPlus



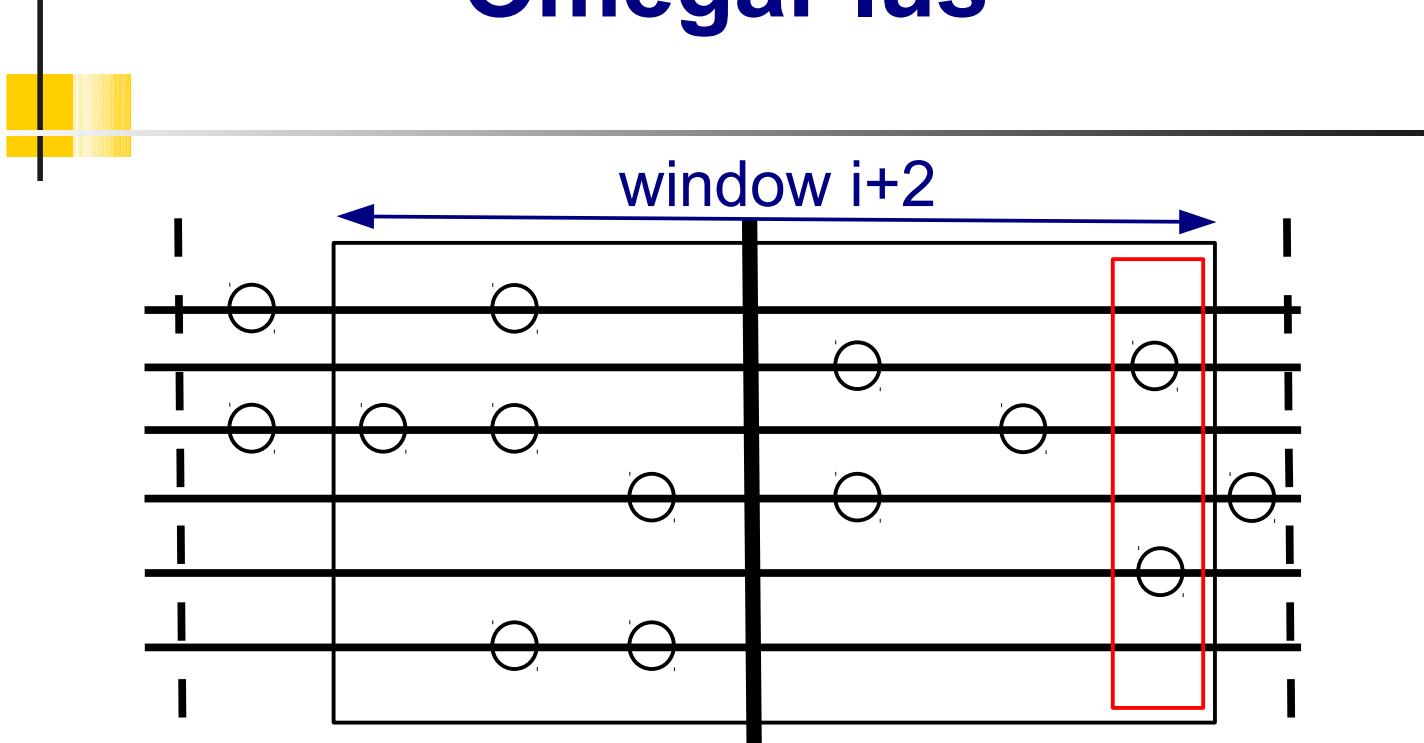
For each position, OmegaPlus computes  $\omega$  for all possible window sizes within the user-defined region.

Each window  $i+1$  contains one SNP more than the previous window  $i$ .

# OmegaPlus



# OmegaPlus



# OmegaPlus

The  $\omega$  statistic uses linkage-disequilibrium (LD) information to detect non-random associations of alleles at different alignment positions.

The LD at each site can be represented as a binary vector to calculate the squared correlation coefficient  $r_{ij}^2$  between two sites i and j.

$$r_{ij}^2 = \frac{(x_{11} - p_1 q_1)^2}{p_1 q_1 (1 - p_1) (1 - q_1)}$$

#1-1 pairs

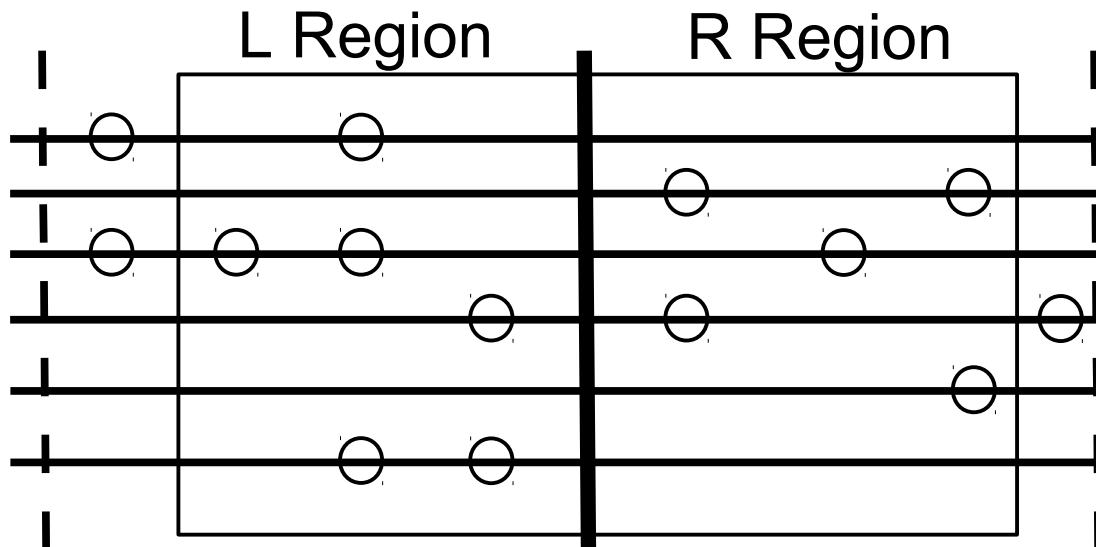
#1s at site i      #1s at site j

The diagram illustrates the components of the LD coefficient  $r_{ij}^2$ . It shows three circles: one at the top labeled  $(x_{11})$ , and two below it labeled  $(p_1)$  and  $(q_1)$ . Arrows point from the top circle to both the bottom circles. The label "#1-1 pairs" is positioned above the top circle, indicating the count of co-occurrences of alleles at the same site.

# OmegaPlus

Number of SNPs in L      Number of SNPs in R

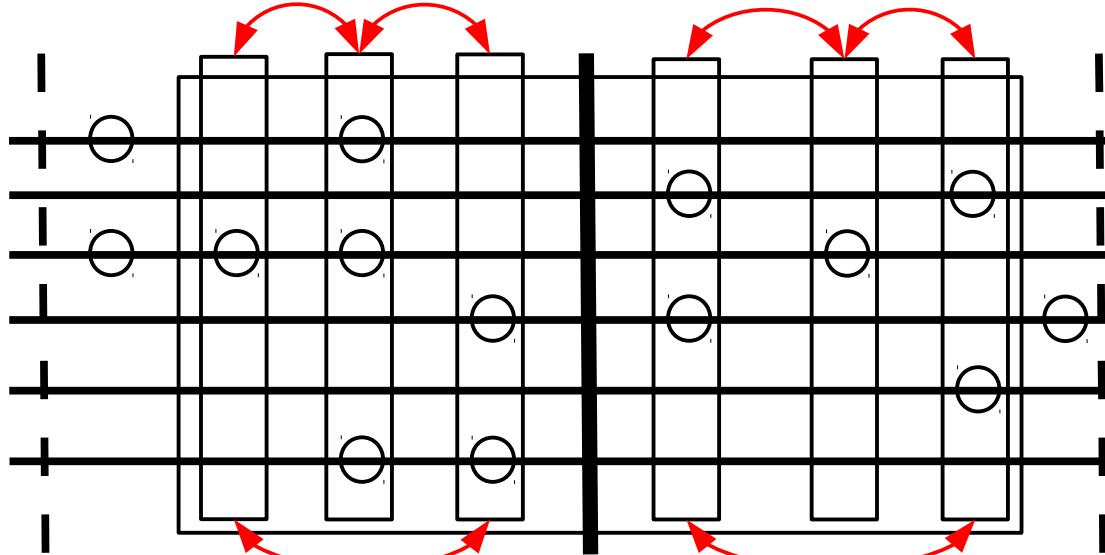
$$\omega = \frac{\binom{l}{2} + \binom{S-l}{2})^{-1} (\sum_{i,j \in L} r_{ij}^2 + \sum_{i,j \in R} r_{ij}^2)}{(l(S-l))^{-1} \sum_{i \in L, j \in R} r_{ij}^2}$$



# OmegaPlus

Selective sweep theory predicts LD increase at both sides of a selective sweep ...

$$\omega = \frac{\left(\binom{l}{2} + \binom{S-l}{2}\right)^{-1} \left( \sum_{i,j \in L} r_{ij}^2 + \sum_{i,j \in R} r_{ij}^2 \right)}{(l(S-l))^{-1} \sum_{i \in L, j \in R} r_{ij}^2}$$

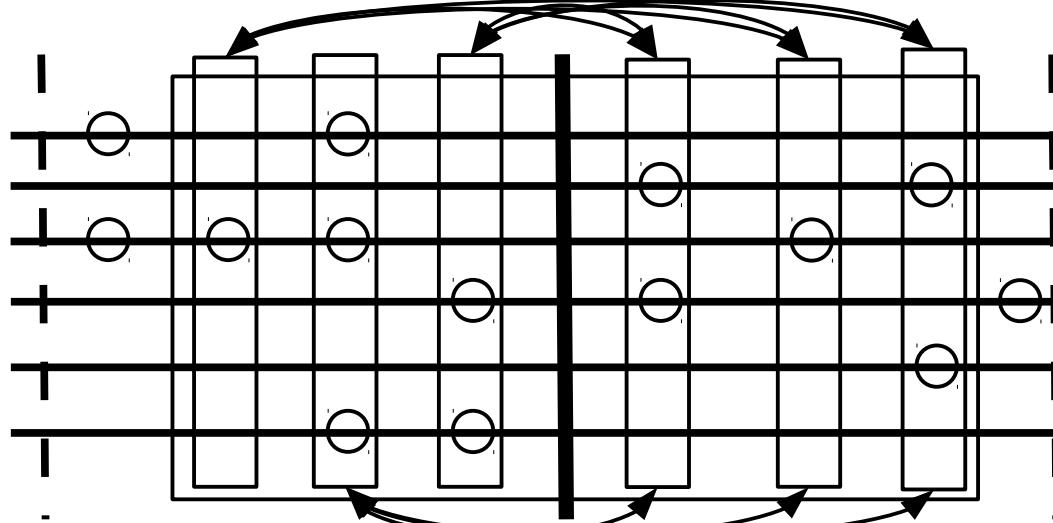


# OmegaPlus

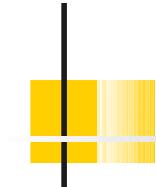


... but not across the selected site.

$$\omega = \frac{((l) + (S-l))^{-1} (\sum_{i,j \in L} r_{ij}^2 + \sum_{i,j \in R} r_{ij}^2)}{(l(S-l))^{-1} \sum_{i \in L, j \in R} r_{ij}^2}$$



# OmegaPlus



Site  $j$

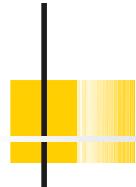
$$M_{i,j} = \begin{cases} 0 & 1 \leq i \leq S, j = i \\ r_{ij}^2 & 2 \leq i \leq S, j = i - 1 \\ M_{i,j+1} + M_{i-1,j} + \\ M_{i-1,j+1} + r_{ij} & 3 \leq i < S, i - 1 > j \geq 0 \end{cases}$$

i       $j$        $j+1$        $j+2$        $j+3$

$i+1$        $i+2$

Fast computation of sums for a region using dynamic programming.

# OmegaPlus

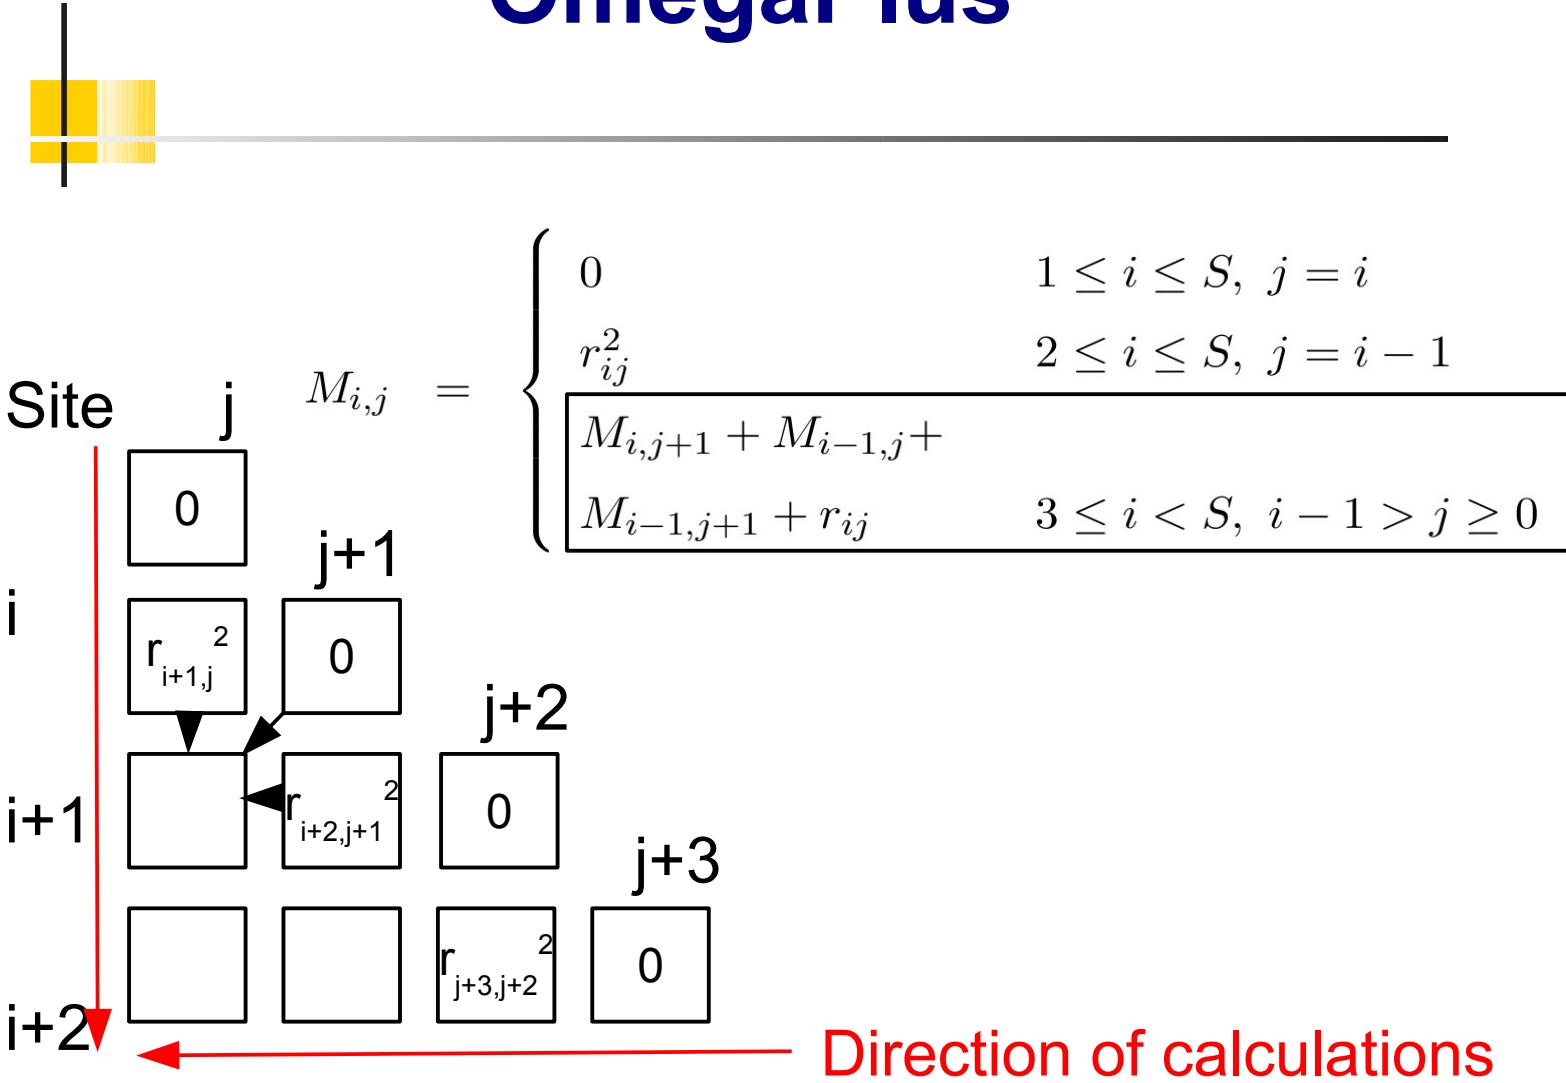


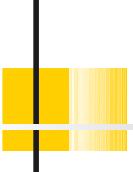
Site  $j$

$$M_{i,j} = \begin{cases} 0 & 1 \leq i \leq S, j = i \\ r_{ij}^2 & 2 \leq i \leq S, j = i - 1 \\ M_{i,j+1} + M_{i-1,j} + \\ M_{i-1,j+1} + r_{ij} & 3 \leq i < S, i - 1 > j \geq 0 \end{cases}$$

$0$	$j+1$	$j+2$	$j+3$
$i$	$r_{i+1,j}^2$	$0$	
$i+1$		$r_{i+2,j+1}^2$	$0$
$i+2$			$r_{j+3,j+2}^2$

# OmegaPlus





# Forward-in-Time Simulation (FiTS)

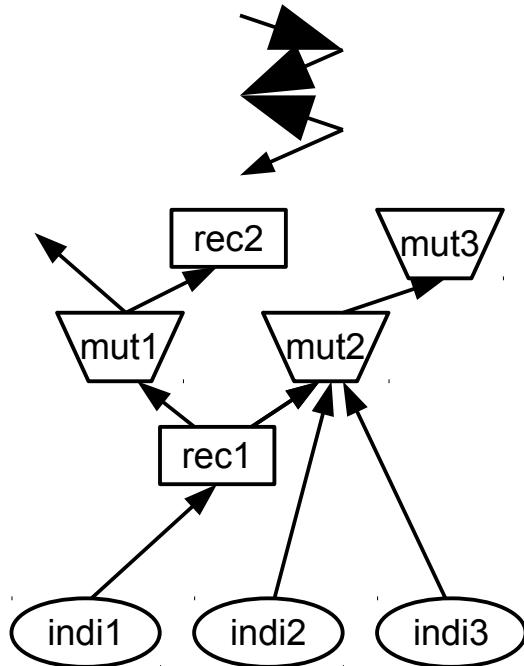
---

- Most accurate method for simulating pop. Gen. datasets
- Works for small population sizes
- Straight-forward method to simulate evolution of multiple non-neutral mutations
- Combines various evolutionary forces under complex demographic scenarios
- **Problem:** orders of magnitude slower than backward simulation

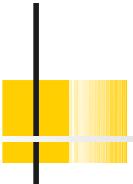
# Toward Whole-Genome FiTS



```
0: n: 000001001000s: fitness: 1.01748
1: n: 000001001101s: fitness: 1.02114
2: n: 000011001010s: fitness: 1.02083
3: n: 00001001010s: fitness: 1.02286
4: n: 000001001010s: fitness: 1.02286
```



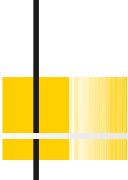
- Simulate non-neutral mutations forward in time
- Keep track of ancestry
- Progressively create ancestral graph
- Add neutral mutations in the end “backward-like”



# Towards Whole-Genome FiTS

---

- Advantage: yields complete history for all simulated sequences in the population
- Uses:
  - high-quality random number generators
  - efficient data structures
  - SIMD/SSE3 instructions
- At present: 1-2 orders of magnitude faster than SFS code
- Planned:
  - parallelization on shared-memory machines
  - trait-based fitness function



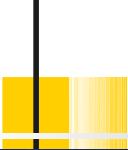
# Some other projects

---

- Protein model assignment in partitioned datasets with joint branch length estimates
  - Is this NP-hard?
  - Can we design good heuristics?
- RAxML goes Bayesian
- Phylogenetic short read identification and alignment in metagenomics
- Discrete algorithms on tree sets → building consensus trees and identifying rogue taxa
- Species delimitation in metagenomic samples
- Novel protein substitution models

# Tonight





# Thank you for your Attention !

---



Psiloritis peak, Crete, Greece, April 28, 2012