

# Fair Tree

## Fairshare Algorithm for Slurm

Ryan Cox and Levi Morrison  
Brigham Young University



# Outline

- Introduction to job prioritization
- BYU's setup
- Issues with existing fairshare algorithms
- Fair Tree
- Fairshare=parent on accounts
- Appendix

# Job Prioritization

- Job Priority can include different components:
  - Age
  - JobSize
  - Partition
  - QOS
  - **Fairshare**

# Fairshare Factor

- Usage and Shares are the two components of Fairshare Factor

# Fairshare Factor

- Usage and Shares are the two components of Fairshare Factor
- Raw Shares are assigned to each association by an admin
  - “Shares” is Raw Shares when normalized to 0.0 .. 1.0
  - Similar to slices of a pie
  - Represents the part of the system that is “yours”

# Fairshare Factor

- Usage and Shares are the two components of Fairshare Factor
- Raw Shares are assigned to each association by an admin
  - “Shares” is Raw Shares when normalized to 0.0 .. 1.0
  - Similar to slices of a pie
  - Represents the part of the system that is “yours”
- “Usage” is a value between 0.0 and 1.0 that represents your proportional usage of the system

# Fairshare Equation

Basic premise of fairshare:

If  $\text{Shares} == \text{Usage}$ , you have hit your “fairshare target”

# Fairshare Equation

$$\text{Fairshare Factor} = 2^{-\text{Usage}/\text{Shares}}$$

*or*


$$F = 2^{-U/S}$$

- Result is between 0.0 and 1.0
- When Usage increases, F decreases
- When Usage decreases, F increases



# BYU's Setup

- Each professor has a Slurm account
  - Has *account coordinator* status for that account
  - BYU doesn't charge \$\$\$ for usage
  - Accounts are treated equally (Shares are the same)
- Students and collaborators are users in a professor's account
  - Users are treated equally (Shares are the same, unless account coordinator changes them)



# Our Definition of “Fair”

- **If accounts A and B are siblings and A has a higher fairshare factor than B, all children of A will have higher fairshare factors than all children of B**
  - True for all sibling associations at all levels

# Disclaimer

- We will discuss BYU's use case and how to support it
  - It may also apply to many other use cases
- We are building on the hard work of others
- We identified some issues with existing algorithms but our work would not have been possible without the other algorithms

# New Fairshare Algorithm

We wrote a new fairshare algorithm.

**Why?**

# The Problem

Account with 35% usage had  
higher priority than accounts with  
only 10% usage


# Fairshare Equation

$$\text{Fairshare Factor} = 2^{-\text{Usage}/\text{Shares}}$$


*or*

$$F = 2^{-U/S}$$

# Fairshare 3D Graph




# Fairshare Graph







# Fairshare Graph Zoomed




# Actual Fairshare Graph



# Actual Fairshare Zoomed



# Actual Fairshare More Zoomed



# Dampening Factor

- Added in 14.03
- $2^{-((\text{Usage}/\text{Shares})/\text{Damp})}$ 
  - Damp is a positive integer (slurm.conf)
- Great idea
- Isn't a complete solution
  - Small values may still be too small
  - Requires manual guess and check
- Can be replaced by linear interpolation
  - Ticket-Based can use this instead of MAX()
  - See Appendix

# Problems With Shares Calculation

- Normalized Shares equation (the “S” in  $2^-(U/S)$ ) for all current algorithms:

$$S = (\text{RawShares}_{\text{user}} / \text{RawShares}_{\text{siblings}})^* \\ (\text{RawShares}_{\text{account}} / \text{RawShares}_{\text{sibling-accounts}})^* \\ (\text{RawShares}_{\text{parent}} / \text{RawShares}_{\text{parent-siblings}})^* \dots$$

# Problems With Shares Calculation

- Normalized Shares equation (the “S” in  $2^-(U/S)$ ) for all current algorithms:

$$S = (\text{RawShares}_{\text{user}} / \text{RawShares}_{\text{siblings}})^* \quad 1/4^* \\ (\text{RawShares}_{\text{account}} / \text{RawShares}_{\text{sibling-accounts}})^* \quad 1/10^* \\ (\text{RawShares}_{\text{parent}} / \text{RawShares}_{\text{parent-siblings}})^* \dots \quad 1/20^* \dots$$

$$1/4^* 1/10^* 1/20^* = 1/800 = 0.00125$$

# Problems With Shares Calculation

- Normalized Shares equation (the “S” in  $2^-(U/S)$ ) for all current algorithms:

$$S = (\text{RawShares}_{\text{user}} / \text{RawShares}_{\text{siblings}})^* \quad 1/4^* \\ (\text{RawShares}_{\text{account}} / \text{RawShares}_{\text{sibling-accounts}})^* \quad 1/10^* \\ (\text{RawShares}_{\text{parent}} / \text{RawShares}_{\text{parent-siblings}})^* \dots \quad 1/20^* \dots$$

$$1/4^* 1/10^* 1/20^* = 1/800 = 0.00125$$

- Assume sibling associations have the same Raw Shares, tree has 2 levels:



# Problems With Shares Calculation

- Normalized Shares equation (the “S” in  $2^-(U/S)$ ) for all current algorithms:

$$S = (\text{RawShares}_{\text{user}} / \text{RawShares}_{\text{siblings}})^* \quad 1/4^* \\ (\text{RawShares}_{\text{account}} / \text{RawShares}_{\text{sibling-accounts}})^* \quad 1/10^* \\ (\text{RawShares}_{\text{parent}} / \text{RawShares}_{\text{parent-siblings}})^* \dots \quad 1/20^* \dots$$

$$1/4^* 1/10^* 1/20^* = 1/800 = 0.00125$$

- Assume sibling associations have the same Raw Shares, tree has 2 levels:

Equivalent to:

$$S = (1 / \text{number of users in account})^* \text{constant}$$

# Problems With Shares Calculation

- Normalized Shares equation (the “S” in  $2^{-(U/S)}$ ) for all current algorithms:

$$S = (\text{RawShares}_{\text{user}} / \text{RawShares}_{\text{siblings}})^{1/4} * \\ (\text{RawShares}_{\text{account}} / \text{RawShares}_{\text{sibling-accounts}})^{1/10} * \\ (\text{RawShares}_{\text{parent}} / \text{RawShares}_{\text{parent-siblings}})^{1/20} * \dots$$

$$1/4 * 1/10 * 1/20 = 1/800 = 0.00125$$


- Assume sibling associations have the same Raw Shares, tree has 2 levels:

Equivalent to:

$$S = (1 / \text{number of users in account}) * \text{constant}$$

- See appendix for demonstration of equivalence

# Problems With Shares Calculation



$$S = 1/2 = 0.5$$

$$S = 1/10 = 0.1$$

Problem applies to more complicated scenarios but is harder to model

# Usage Effective

- Usage Effective (the U in  $2^U - (U/S)$ ) =  
 $UA_{child} + ((UE_{parent} - UA_{child}) * S_{child} / S_{all\_siblings})$

# Usage Effective

- Usage Effective (the U in  $2^U - (U/S)$ ) =  
 $UA_{child} + ((UE_{parent} - UA_{child}) * S_{child} / S_{all\_siblings})$
- Assuming two levels in the tree,  $2^U - (U/S)$  expands to:

$$2 - \left( \frac{\frac{U_{User}}{U_{Total}} + \left[ \frac{U_{Acct}}{U_{Total}} + \left( 1 - \frac{U_{Acct}}{U_{Total}} \right) * \frac{S_{Acct}}{S_{Acct} + S_{ASiblings}} - \frac{U_{User}}{U_{Total}} \right] * \frac{S_{User}}{S_{User} + S_{USiblings}}}{\frac{S_{User}}{S_{User} + S_{USiblings}} * \frac{S_{Acct}}{S_{Acct} + S_{ASiblings}}} \right)$$

# Usage Effective

- Usage Effective (the U in  $2^U - (U/S)$ ) =  
 $UA_{child} + ((UE_{parent} - UA_{child}) * S_{child} / S_{all\_siblings})$
- Assuming two levels in the tree,  $2^U - (U/S)$  expands to:

$$2 - \left( \frac{\frac{U_{User}}{U_{Total}} + \left[ \frac{U_{Acct}}{U_{Total}} + \left( 1 - \frac{U_{Acct}}{U_{Total}} \right) * \frac{S_{Acct}}{S_{Acct} + S_{ASiblings}} - \frac{U_{User}}{U_{Total}} \right] * \frac{S_{User}}{S_{User} + S_{USiblings}}}{\frac{S_{User}}{S_{User} + S_{USiblings}} * \frac{S_{Acct}}{S_{Acct} + S_{ASiblings}}} \right)$$

- Has some unpleasant side effects that we consider bugs
  - (Have 20 minutes and a whiteboard? Let's talk!)

# Other Algorithms



- Ticket-Based

- Depends on queue state
  - Hard to explain to users
- $T = T_{\text{parent}} * S * F / \text{SUM}(S * F)_{\text{active\_siblings}}$ 
  - Similar issues to Norm Shares equation because  $(S * F / \text{SUM}(S * F)_{\text{active\_siblings}})$  depends on the number of active users in an account

- Depth Oblivious

- Very complex math
  - Hard to explain to users and admins
  - Hard to evaluate its fairness
- Many opportunities for floating point precision issues


# An Example






# Timeshare on a Spaceship

- The Beatles and Elvis split the cost of a spaceship and crew
- One passenger spaceship
- Elvis pays 50%
- Beatles pay 50%
  - Each member pays 25% of the Beatles' 50% (equals 1/8 of the total)
  - **If a band member isn't waiting to use the spaceship, other band members can take a ride**



Accounts



Users

- Lennon (Account=Beatles)
- McCartney (Account=Beatles)
- Harrison (Account=Beatles)
- Starr (Account=Beatles)
- Elvis (Account=Elvis)


How well does Slurm  
perform in this scenario?

# What actually happens?


**Simulated using the following account hierarchy:**

Account	User	Norm Shares
-----		
root		0.000000
beatles		0.500000
beatles	harrison	0.125000
beatles	lennon	0.125000
beatles	mccartney	0.125000
beatles	starr	0.125000
elvis		0.500000
elvis	elvis	0.500000


# Traditional Multifactor




4 Beatles vs Elvis



3 Beatles vs Elvis




2 Beatles vs Elvis




1 Beatle vs Elvis




# Ticket-Based




4 Beatles vs Elvis



3 Beatles vs Elvis




2 Beatles vs Elvis




1 Beatle vs Elvis




# Depth Oblivious




4 Beatles vs Elvis



3 Beatles vs Elvis




2 Beatles vs Elvis




1 Beatle vs Elvis




# Fair Tree




4 Beatles vs Elvis



3 Beatles vs Elvis



2 Beatles vs Elvis



1 Beatle vs Elvis




# Goals of Fair Tree

- **If accounts A and B are siblings and A has a higher fairshare factor than B, all children of A will have higher fairshare factors than all children of B**
  - True for all sibling associations at all levels
- Eliminate problems due to floating point precision loss
- sshare and debug support



# Trees

- Similar idea to binary search tree:



# Fairshare Calculation

- Level Fairshare =  $S / U$ 
  - $S = \text{RawShares}_{\text{self}} / \text{RawShares}_{\text{self+siblings}}$
  - $U = \text{RawUsage}_{\text{self}} / \text{RawUsage}_{\text{self+siblings}}$
- Range:  $[0, \text{infinity}]$ 
  - if  $U == S$  then  $LF = 1.0$
- Replaces  $2^{\lfloor U/S \rfloor}$
- **Only used for sorting**

# Tree Traversal

- Create “rooted plane tree”
- Traversal function, starting at root:
  - Calculate Level Fairshare for each child (S/U)
  - Sort children by Level Fairshare from highest to lowest

# Tree Traversal

- Create “rooted plane tree”
- Traversal function, starting at root:
  - Calculate Level Fairshare for each child (S/U)
  - Sort children by Level Fairshare from highest to lowest
  - Visit the children in order
    - If account, recurse with account as new root
    - If user, assign a final fairshare factor based on ranking

# Traversal in the Fair Tree Algorithm for Slurm

# Fair Tree: Traversal




**= visit association**



**= calculate level fairshare**




**= sort siblings**

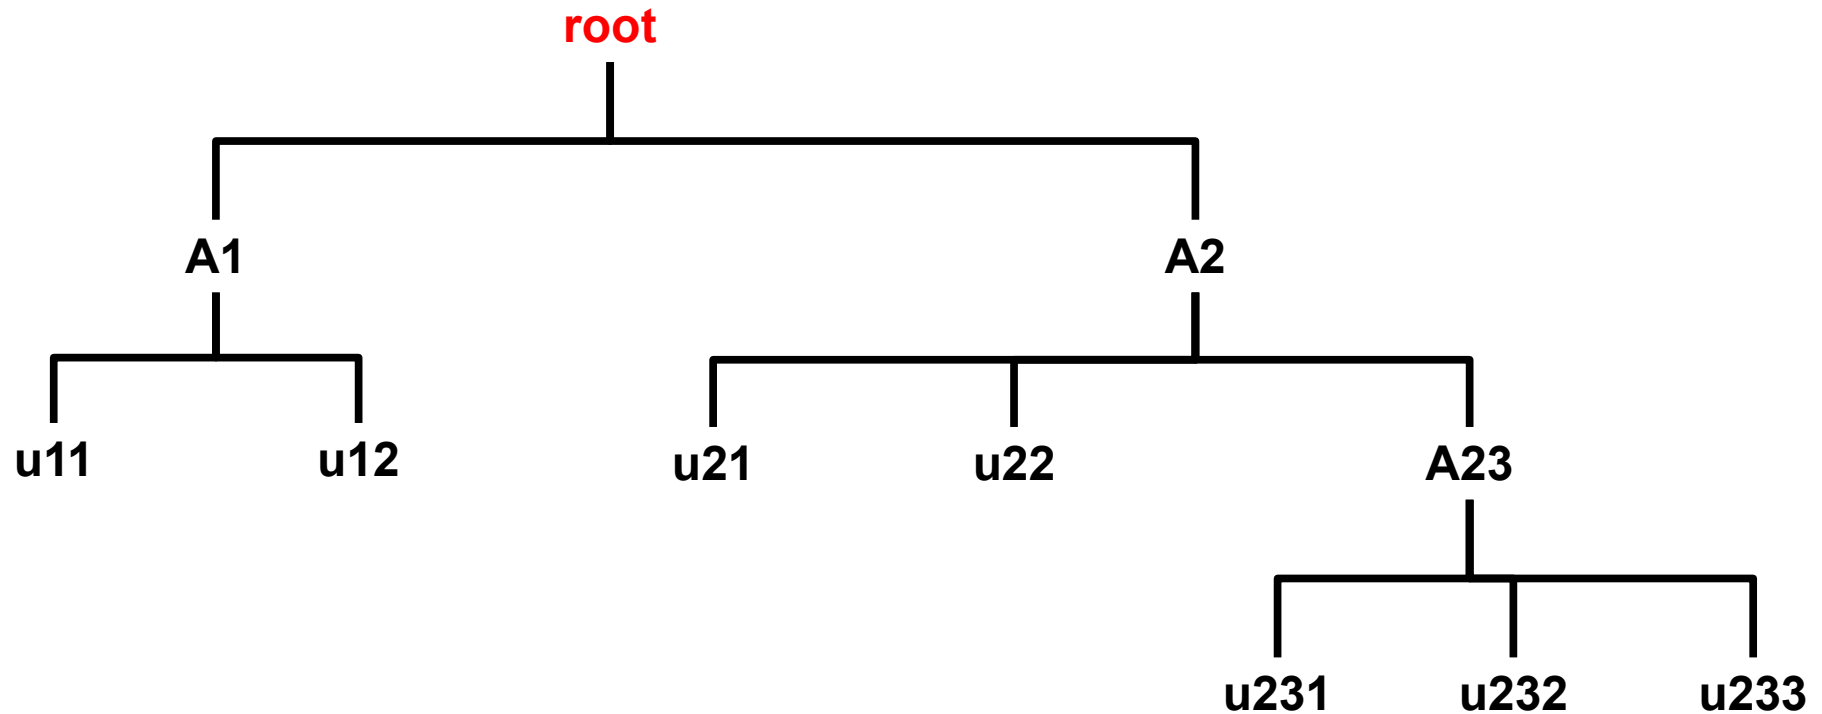


**= calculate final fairshare for user**

# Fair Tree: Traversal




# Fair Tree: Traversal



*Visit association*




# Fair Tree: Traversal




*Calculate level fairshare*

# Fair Tree: Traversal




*Calculate level fairshare*

# Fair Tree: Traversal




*Sort by level fairshare*

# Fair Tree: Traversal




*Sort by level fairshare*

# Fair Tree: Traversal




*Visit association*

# Fair Tree: Traversal




*Calculate level fairshare*

# Fair Tree: Traversal



*Calculate level fairshare*


# Fair Tree: Traversal



*Sort by level fairshare*




# Fair Tree: Traversal



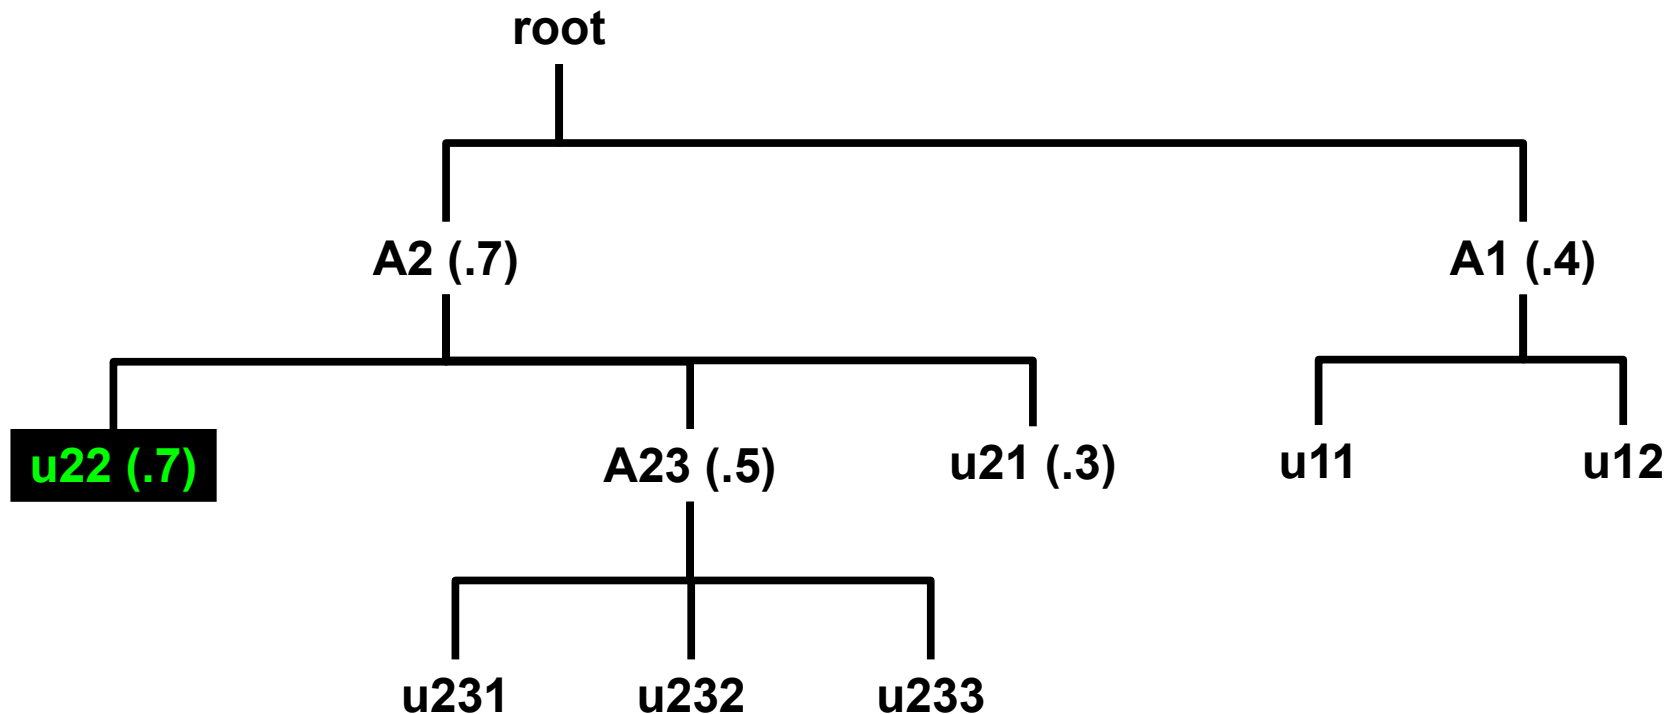
*Sort by level fairshare*

# Fair Tree: Traversal




*Visit association*

# Fair Tree: Traversal




*Final fairshare value = rank-- / user\_count*

# Fair Tree: Traversal




*Visit association*

# Fair Tree: Traversal



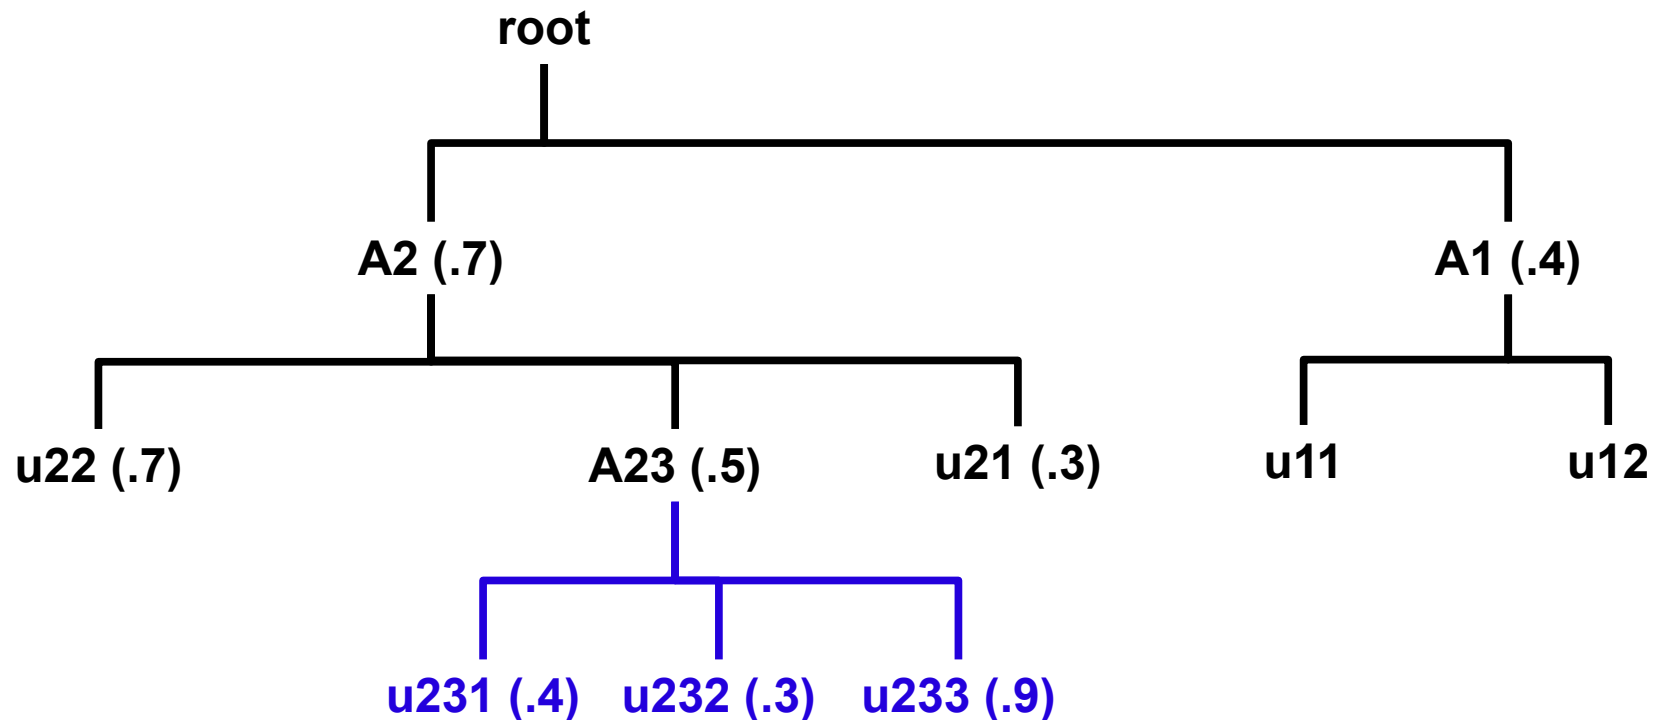
*Calculate level fairshare*

# Fair Tree: Traversal



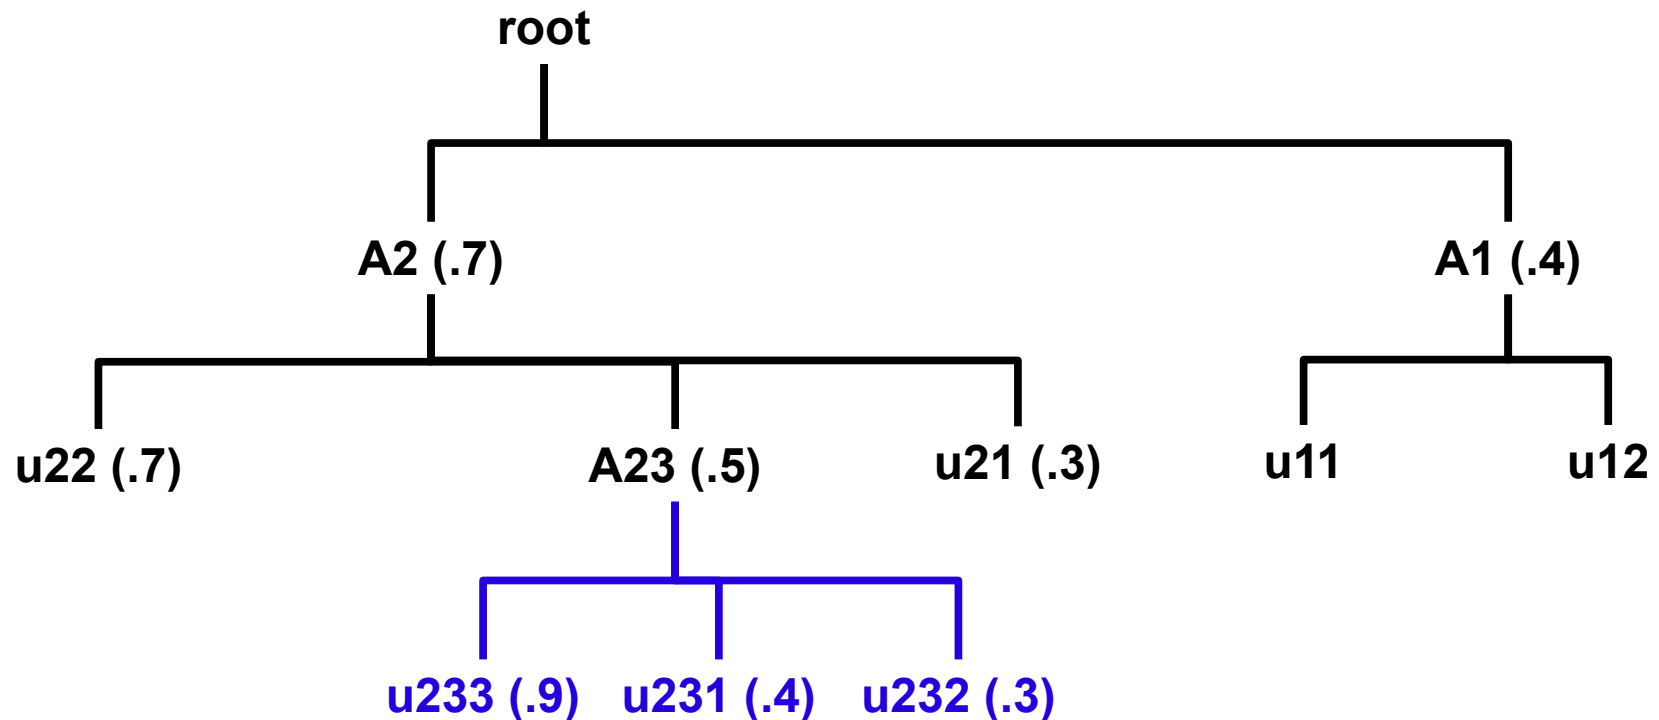
*Calculate level fairshare*

# Fair Tree: Traversal



*Sort by level fairshare*


# Fair Tree: Traversal



*Sort by level fairshare*




# Fair Tree: Traversal




*Visit association*

# Fair Tree: Traversal




*Final fairshare value = rank-- / user\_count*

# Fair Tree: Traversal




*Visit association*

# Fair Tree: Traversal




*Final fairshare value = rank-- / user\_count*

# Fair Tree: Traversal




*Visit association*

# Fair Tree: Traversal




*Final fairshare value = rank-- / user\_count*

# Fair Tree: Traversal



*Visit association*


# Fair Tree: Traversal



*Final fairshare value = rank-- / user\_count*




# Fair Tree: Traversal



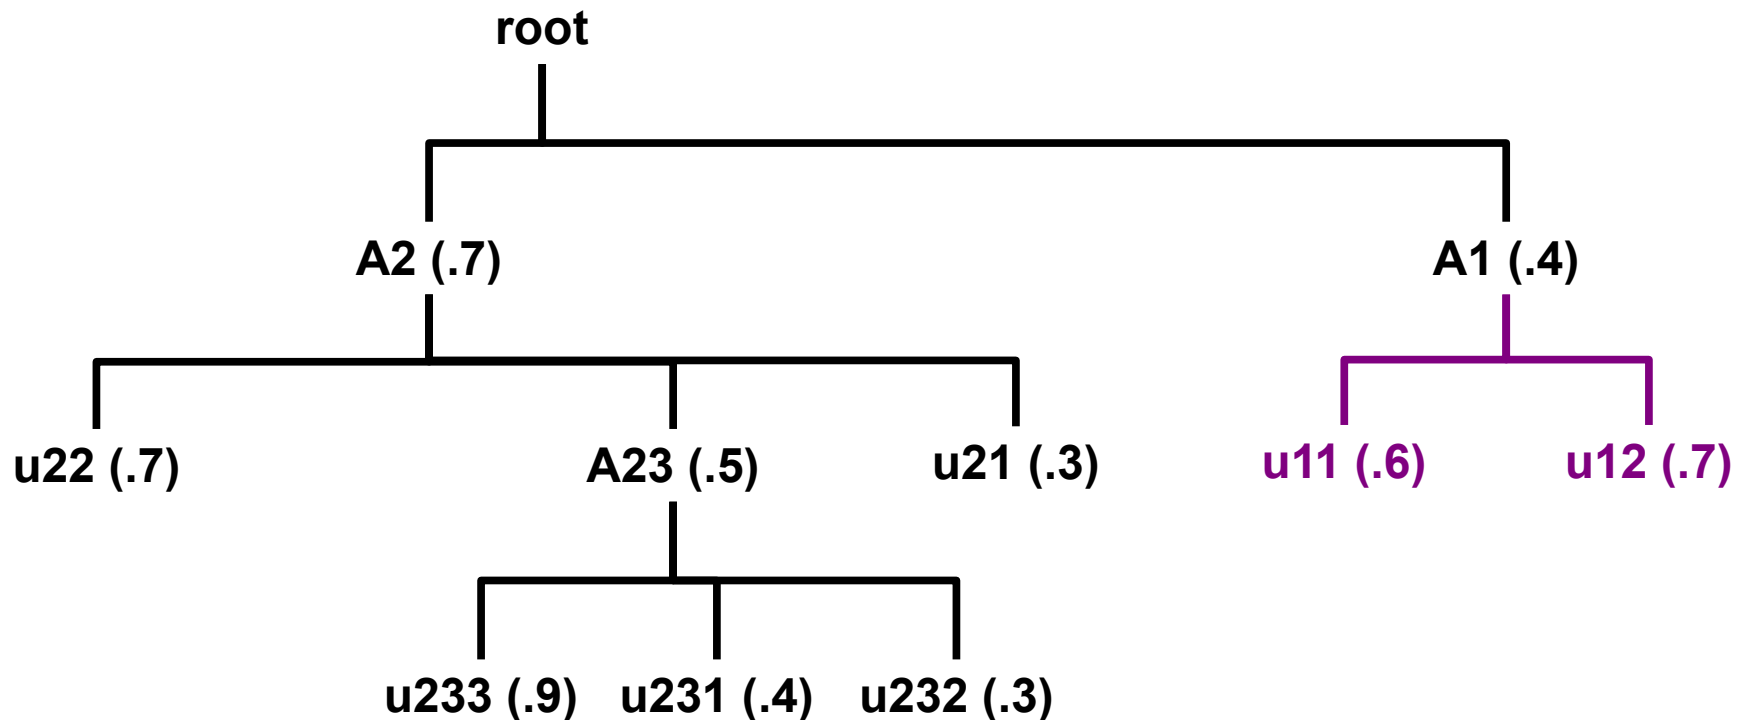
*Visit association*

# Fair Tree: Traversal



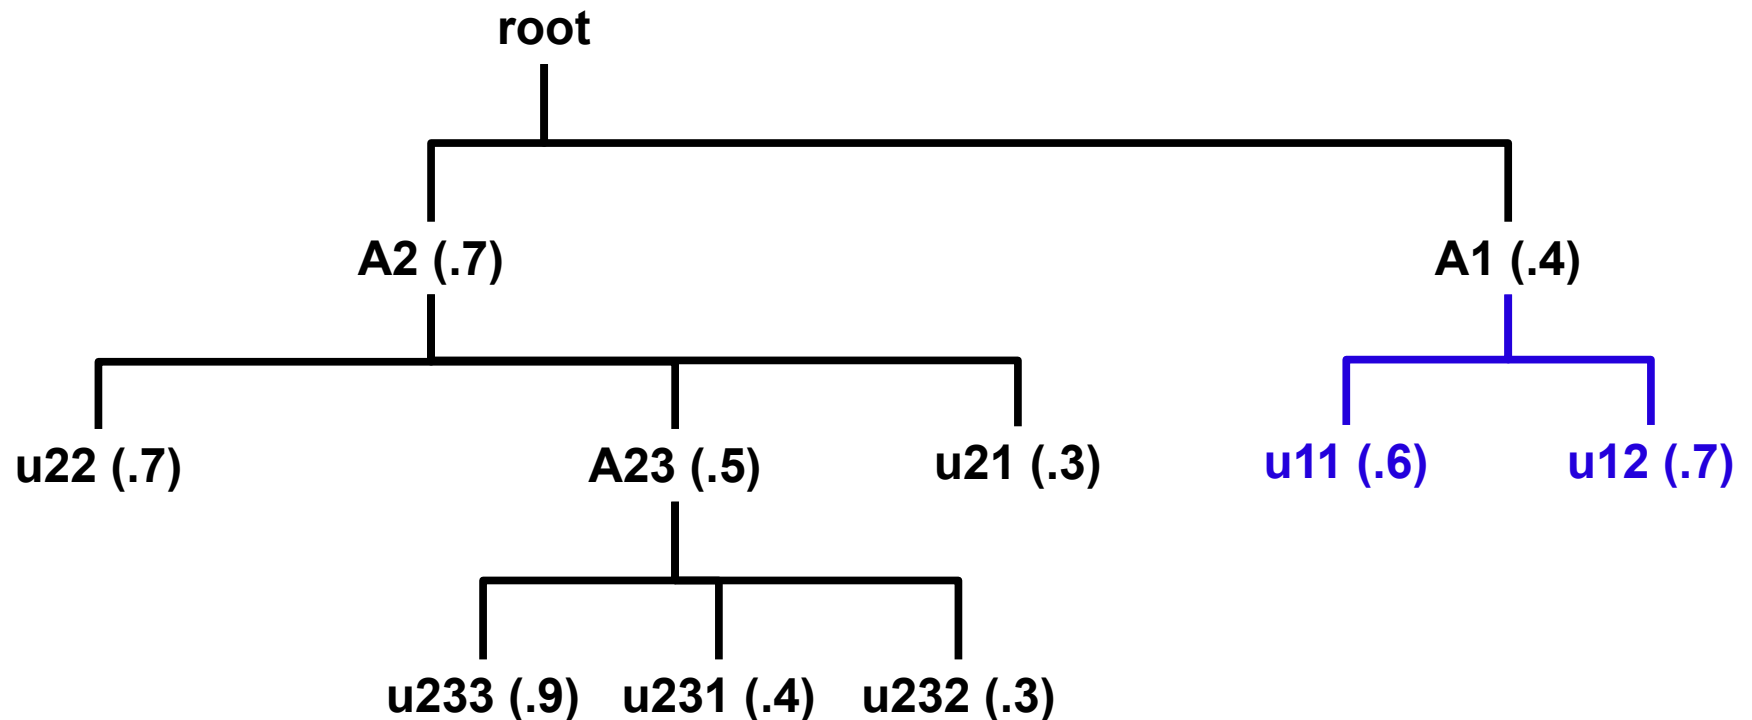
*Calculate level fairshare*

# Fair Tree: Traversal



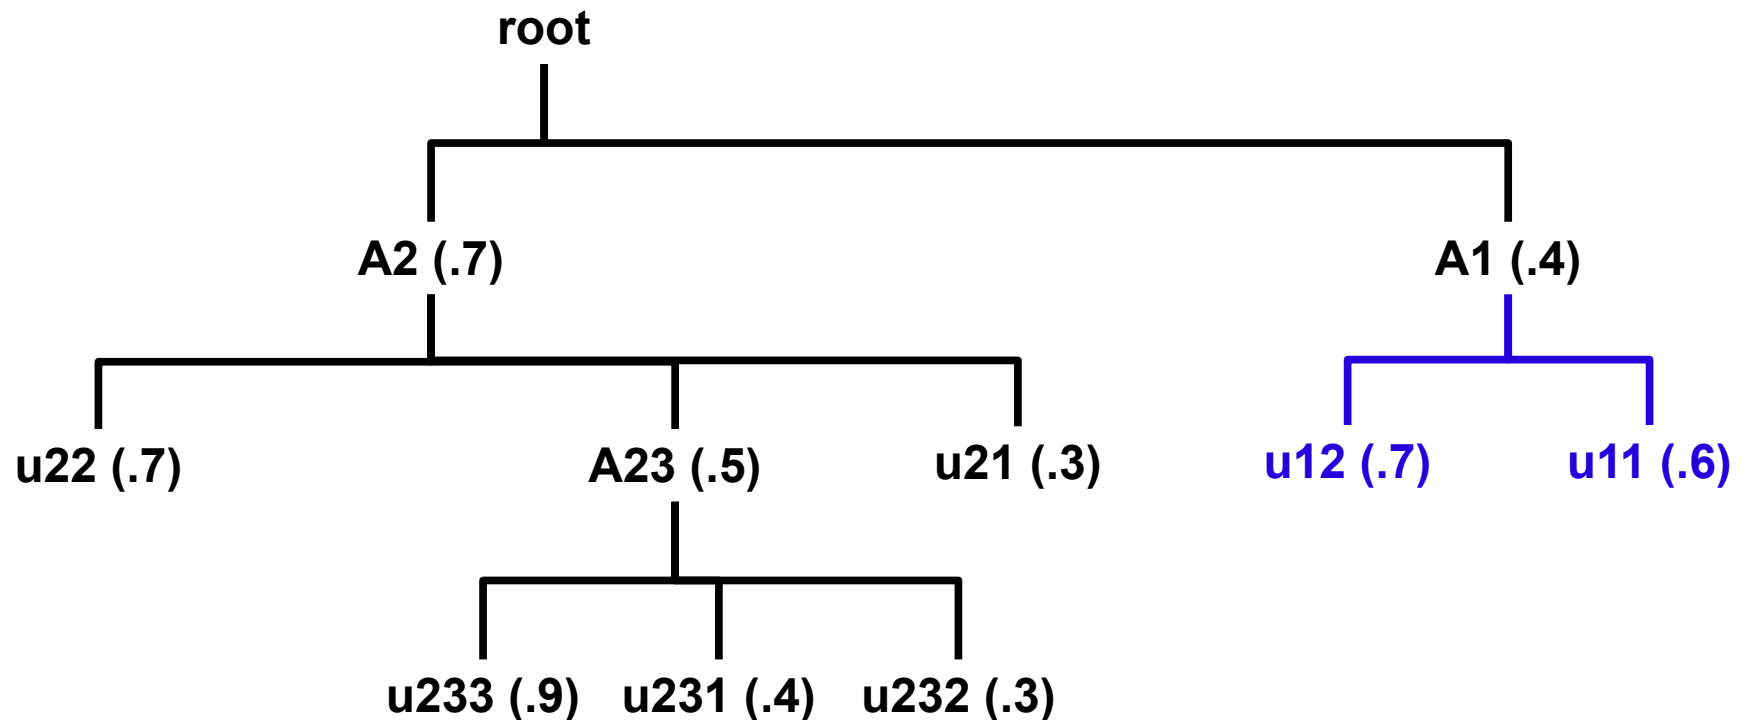
*Calculate level fairshare*

# Fair Tree: Traversal



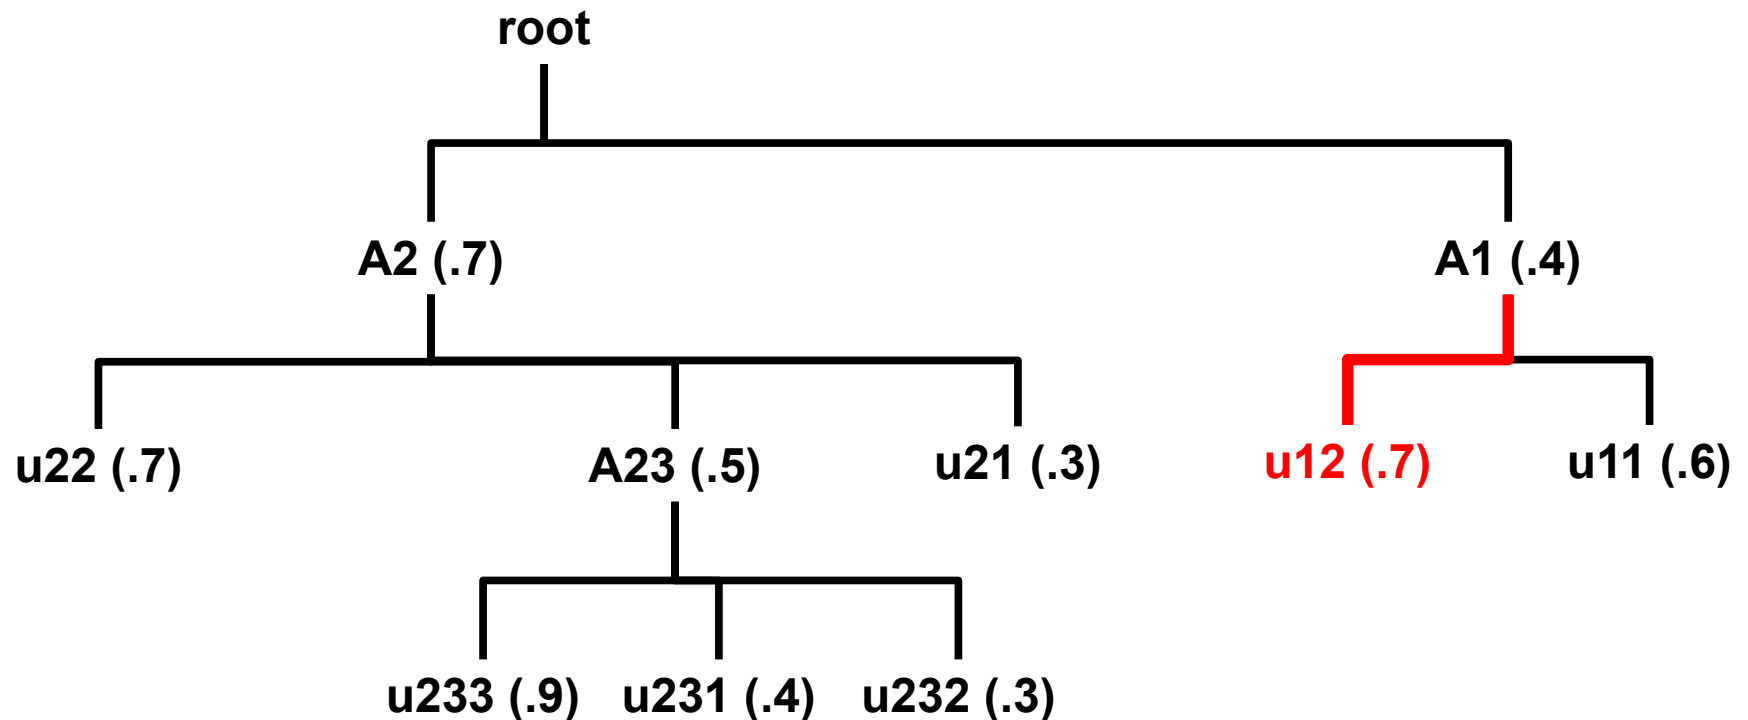
*Sort by level fairshare*

# Fair Tree: Traversal



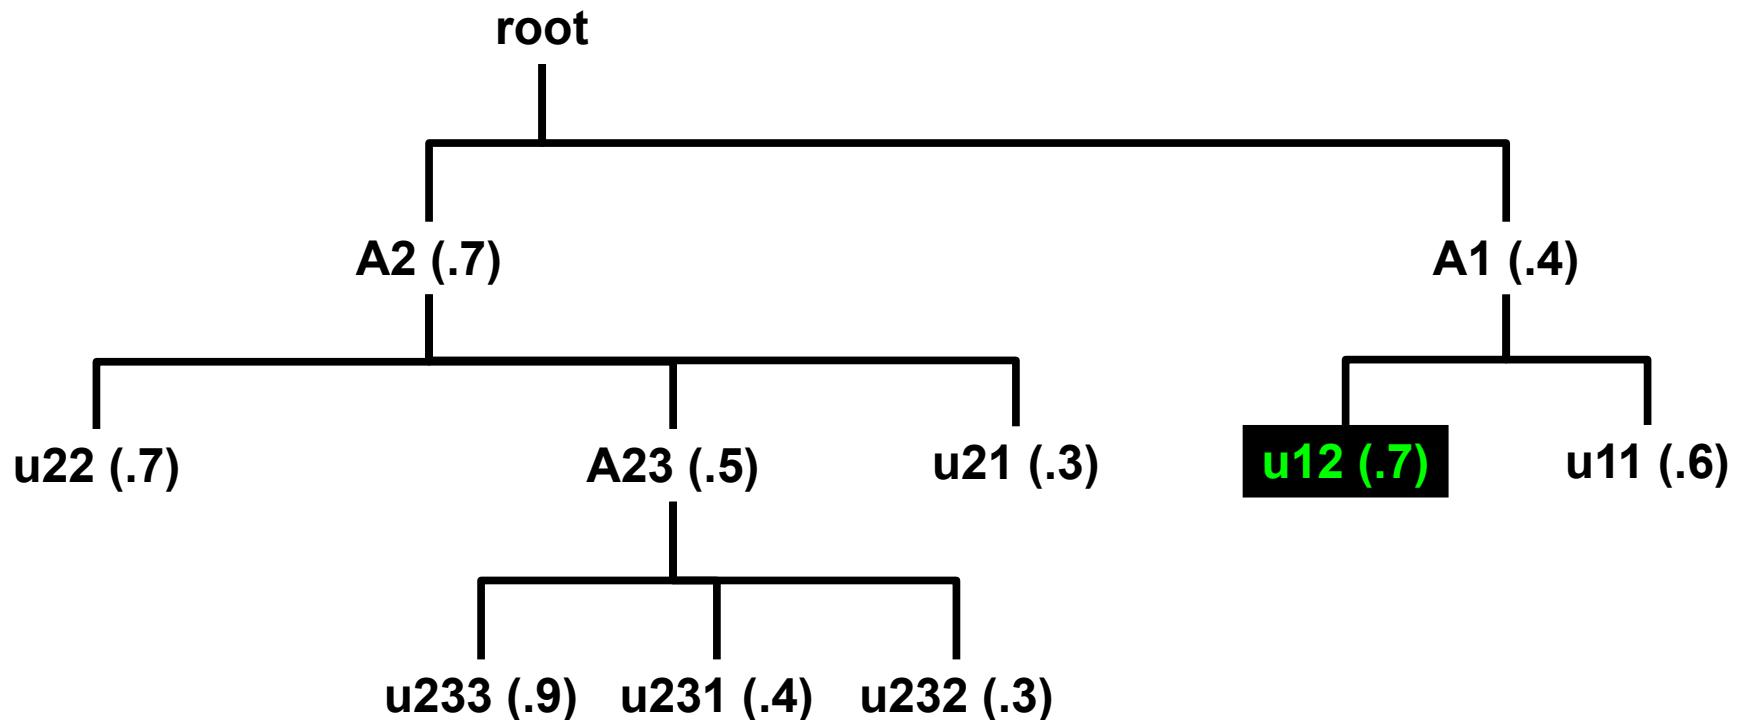
*Sort by level fairshare*

# Fair Tree: Traversal



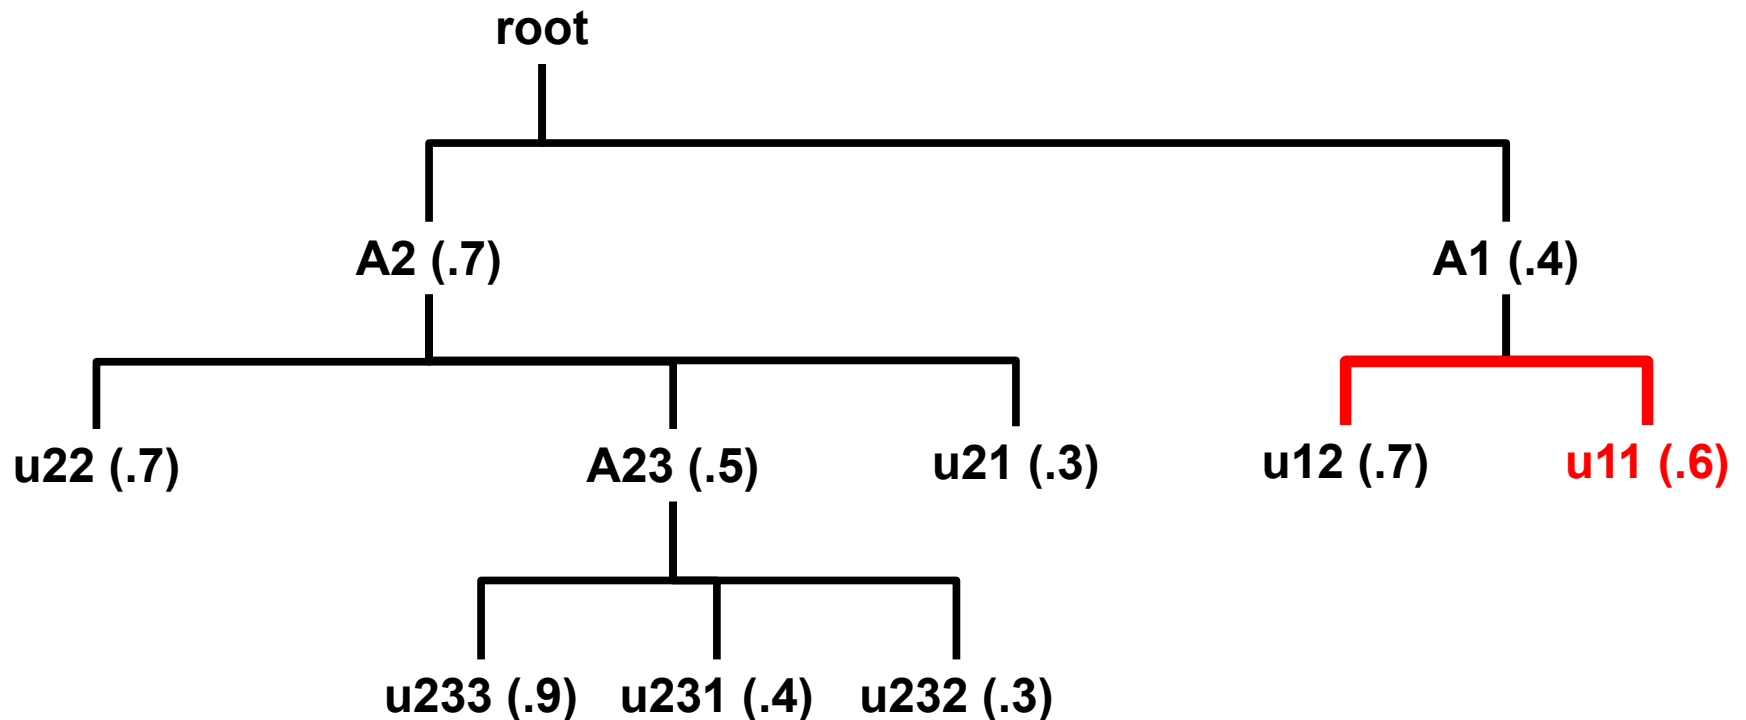
*Visit association*

# Fair Tree: Traversal



*Final fairshare value = rank-- / user\_count*

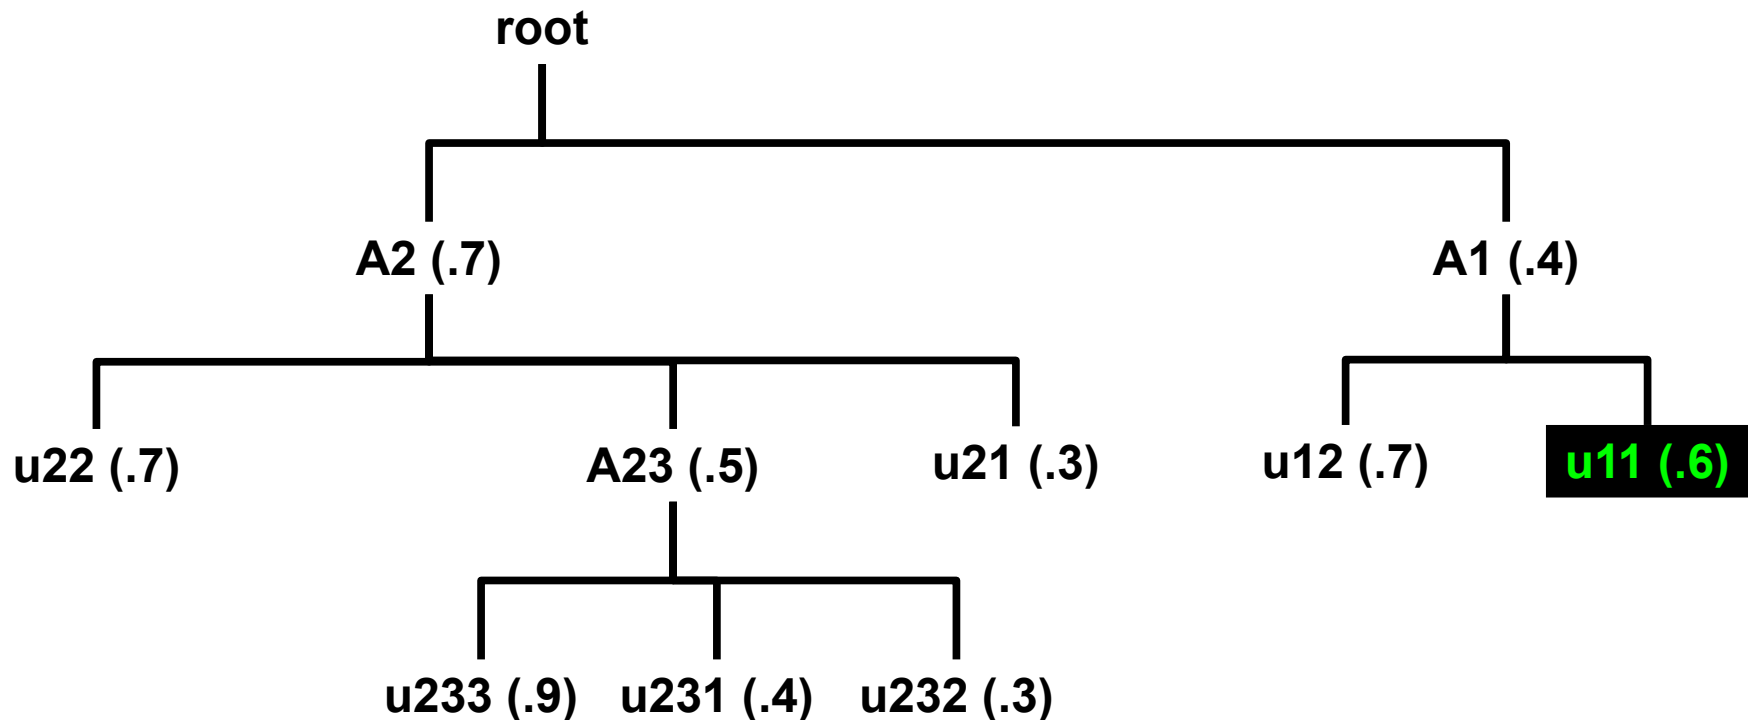
# Fair Tree: Traversal



*Visit association*




# Fair Tree: Traversal



*Final fairshare value = rank-- / user\_count*

# Fair Tree: Traversal Complete




Traversal Complete!

All final fairshare values were assigned

# Ranking

- Users are ranked as they are found
  - Ties are allowed (e.g. 8,8,8,5,4) if Level Fairshare is equal
  - See the appendix for details of tie handling
- Avoids precision loss

# sshare

- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

\*Note: Some sshare columns were cut for brevity in this presentation

# Debug

```
slurmctld: Fair Tree fairshare algorithm, starting at root:
slurmctld:      elvis (elvis):  1.11010830324909747294
slurmctld:      elvis (elvis):  1.000000000000000000000000
slurmctld:      beatles (beatles):  0.90976331360946745562
slurmctld:      mccartney (beatles):  4.56756756756756756785
slurmctld:      lennon (beatles):  1.65686274509803921568
slurmctld:      starr (beatles):  0.71610169491525423724
slurmctld:      harrison (beatles):  0.56146179401993355479
```

- Live view of depth-first traversal
- Users are printed in order of final fairshare factor
- Number shown is Level FS

# Easy to Try

- `slurm.conf`:
  - `PriorityFlags=FAIR_TREE`
- `scontrol reconfigure`
  - Wait for next iteration (`$PriorityCalcPeriod` minutes)
- Available starting in 14.11.0pre6

# Advantages

- Uses entry-level computer science (ordered tree)
- Uses entry-level math
  - $S/U$  is simple
  - $\text{LevelFS} < \text{LevelFS}$  is simple
- Precision loss issues *extremely* unlikely
- Handles unbalanced trees
- Demonstrably fair

# Possible Concerns

- A user can use all the account's shares if no one else is running
  - (Great!)
  - No algorithm can *prevent* this
  - User limits
    - Set by admin or account coordinator



# Possible Concerns

- A user can use all the account's shares if no one else is running
  - (Great!)
  - No algorithm can *prevent* this
  - User limits
    - Set by admin or account coordinator
- Tiny user in a huge group has trouble running
  - Optional QOS w/UsageFactor=20 and Priority=1000000000
    - User runs quickly but pays heavily
  - Secondary account with tight limits

# Possible Concerns

- A user can use all the account's shares if no one else is running
  - (Great!)
  - No algorithm can *prevent* this
  - User limits
    - Set by admin or account coordinator
- Tiny user in a huge group has trouble running
  - Optional QOS w/UsageFactor=20 and Priority=1000000000
    - User runs quickly but pays heavily
  - Secondary account with tight limits
- Strictness of hierarchical prioritization
  - Fairshare=parent on accounts
  - Fuzzy matching (see “[Future Development](#)”)

Fairshare=parent modifications  
(all algorithms)

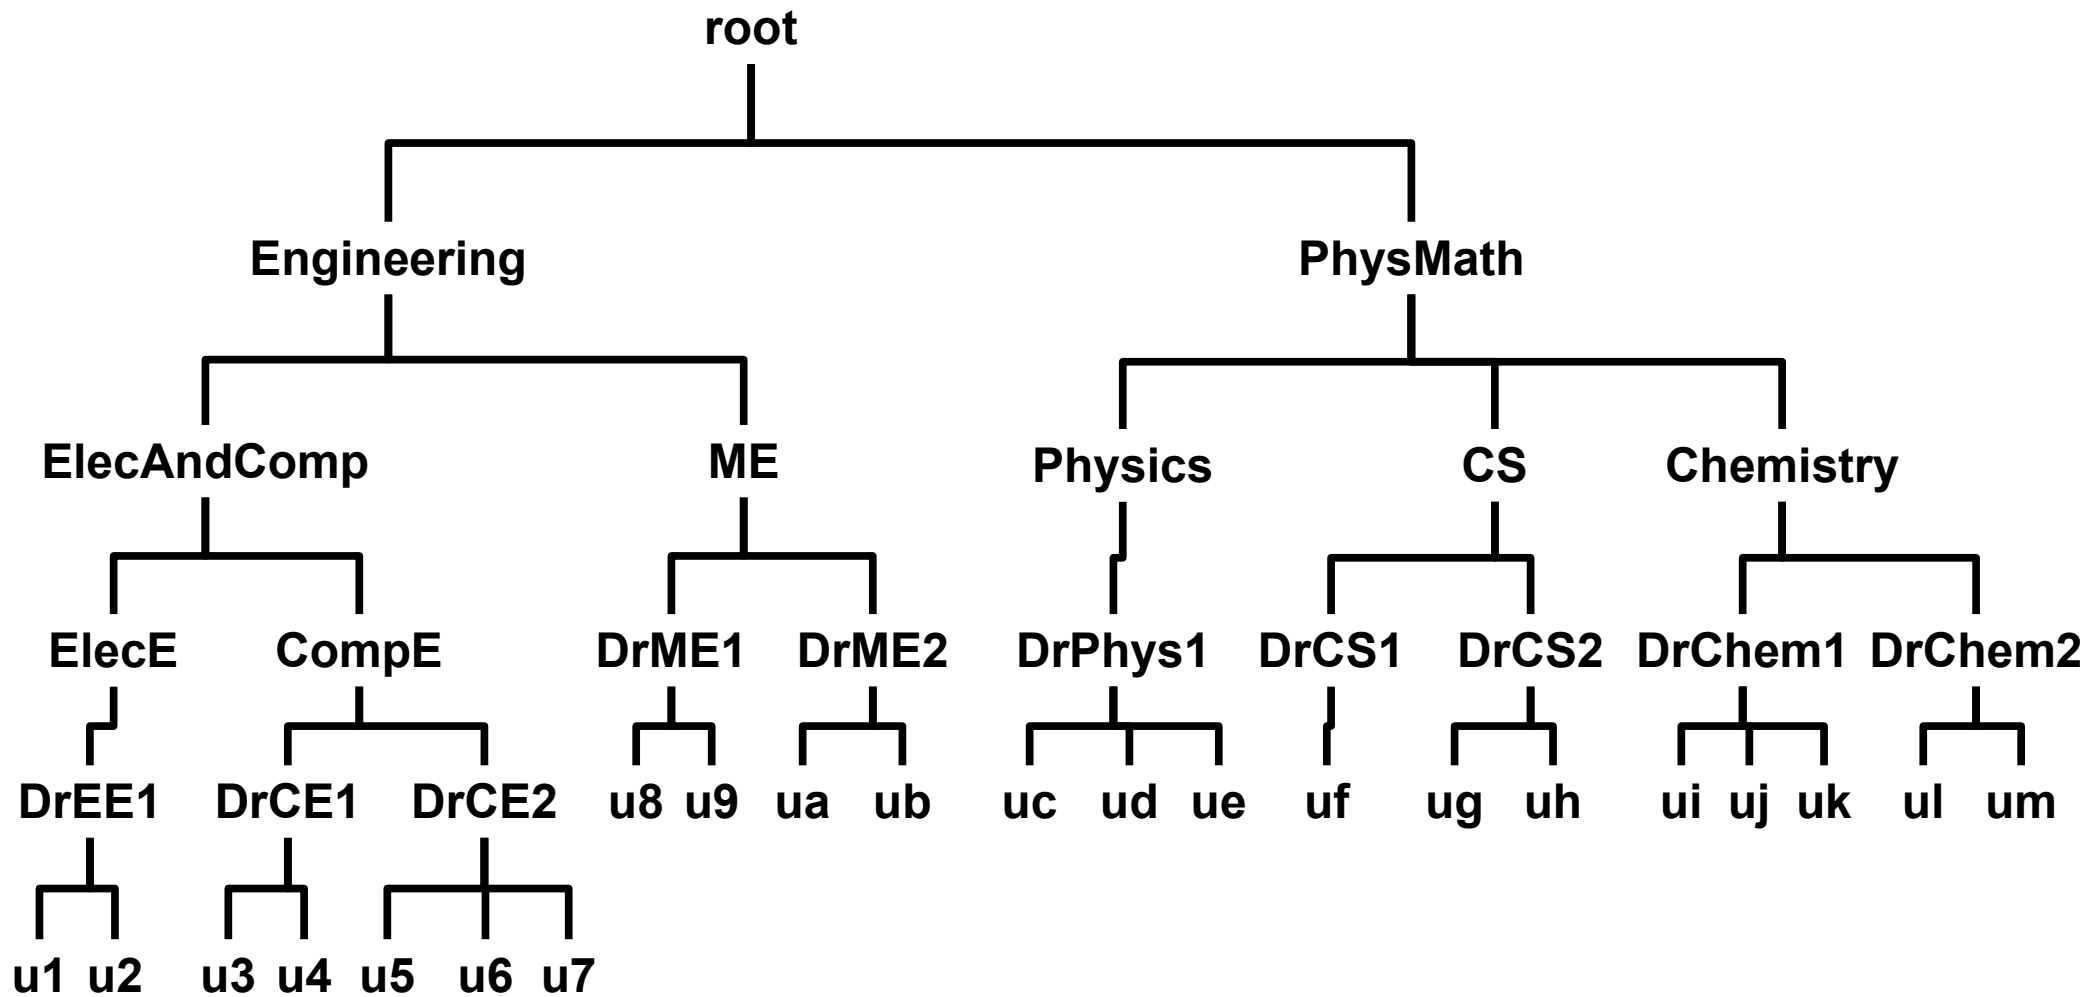
# Fairshare=parent

- Fairshare=parent behavior on an account was previously undefined
- It is now defined:

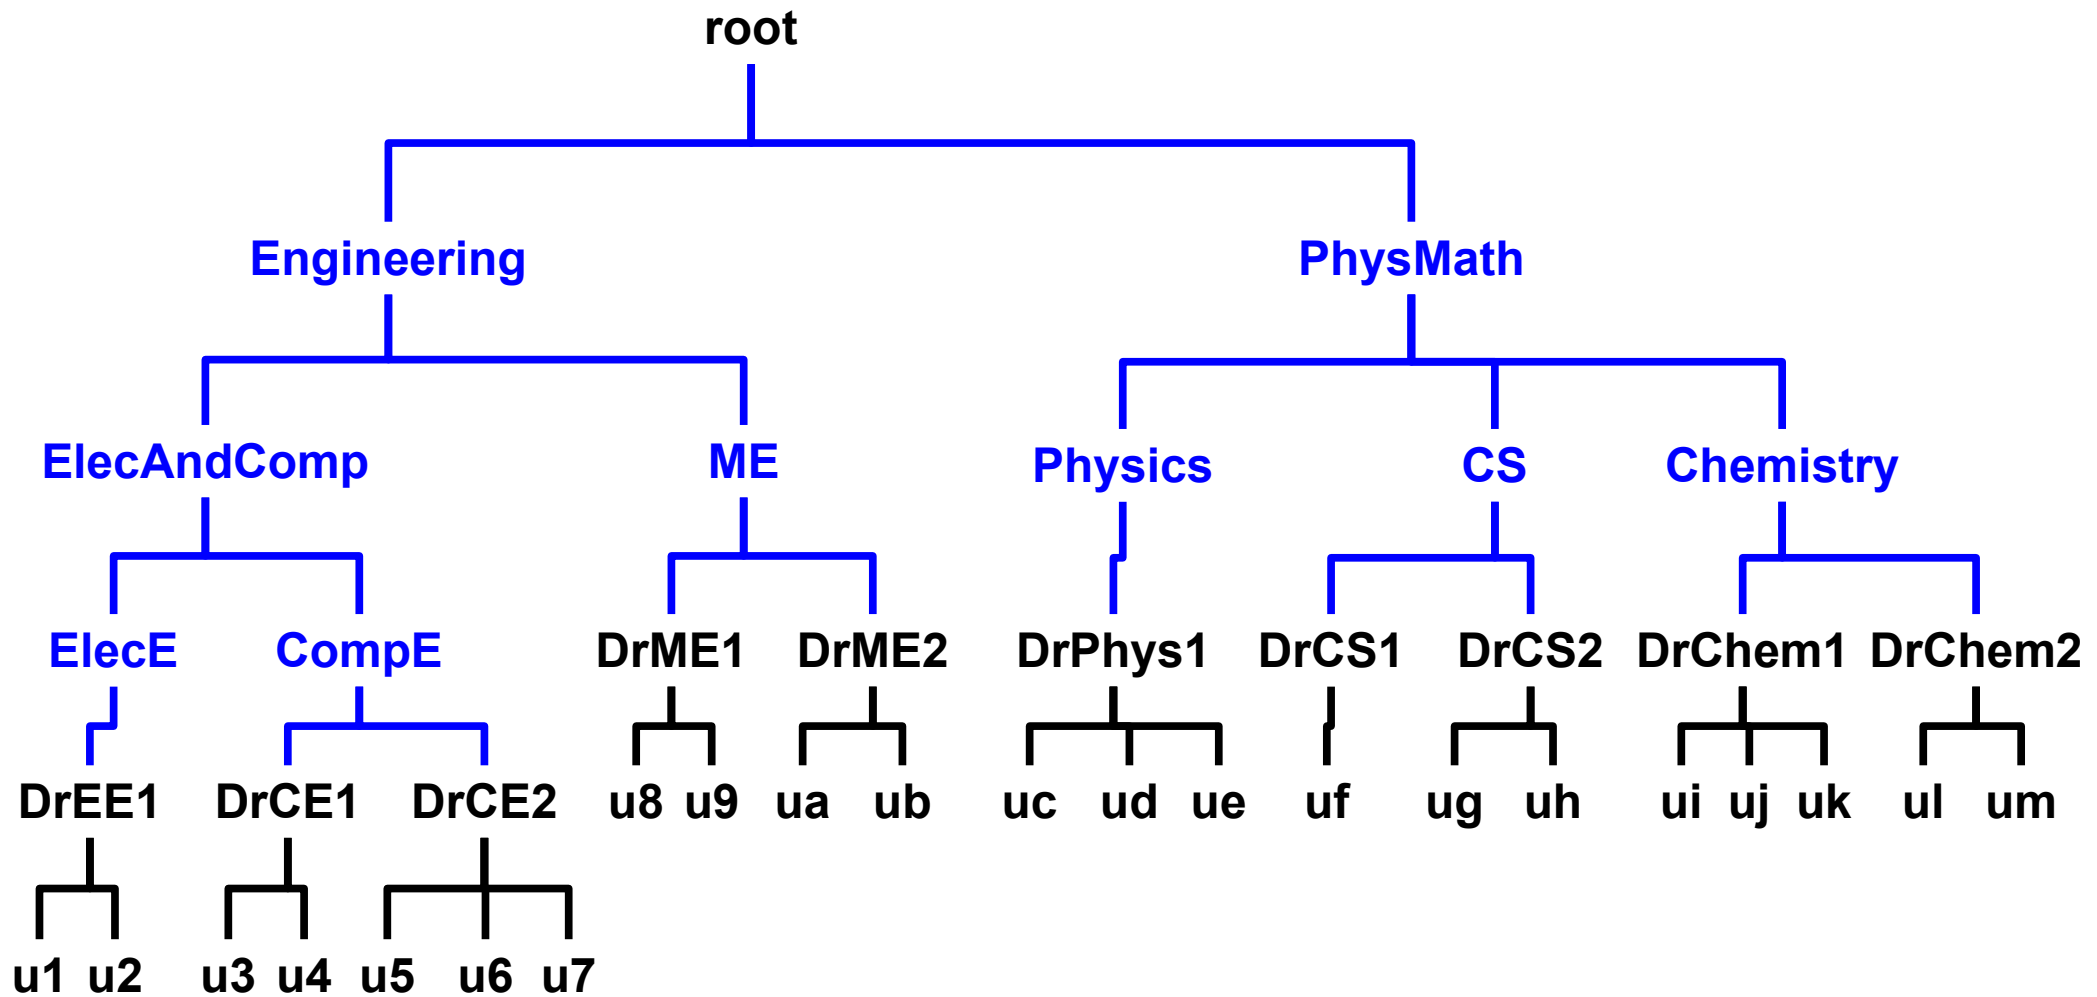
If Fairshare=parent is set on an account, that account's children will be effectively reparented for fairshare calculations to the first ancestor that is not Fairshare=parent.

- The behavior of limits is unchanged
- Available in 14.11 **for all algorithms**

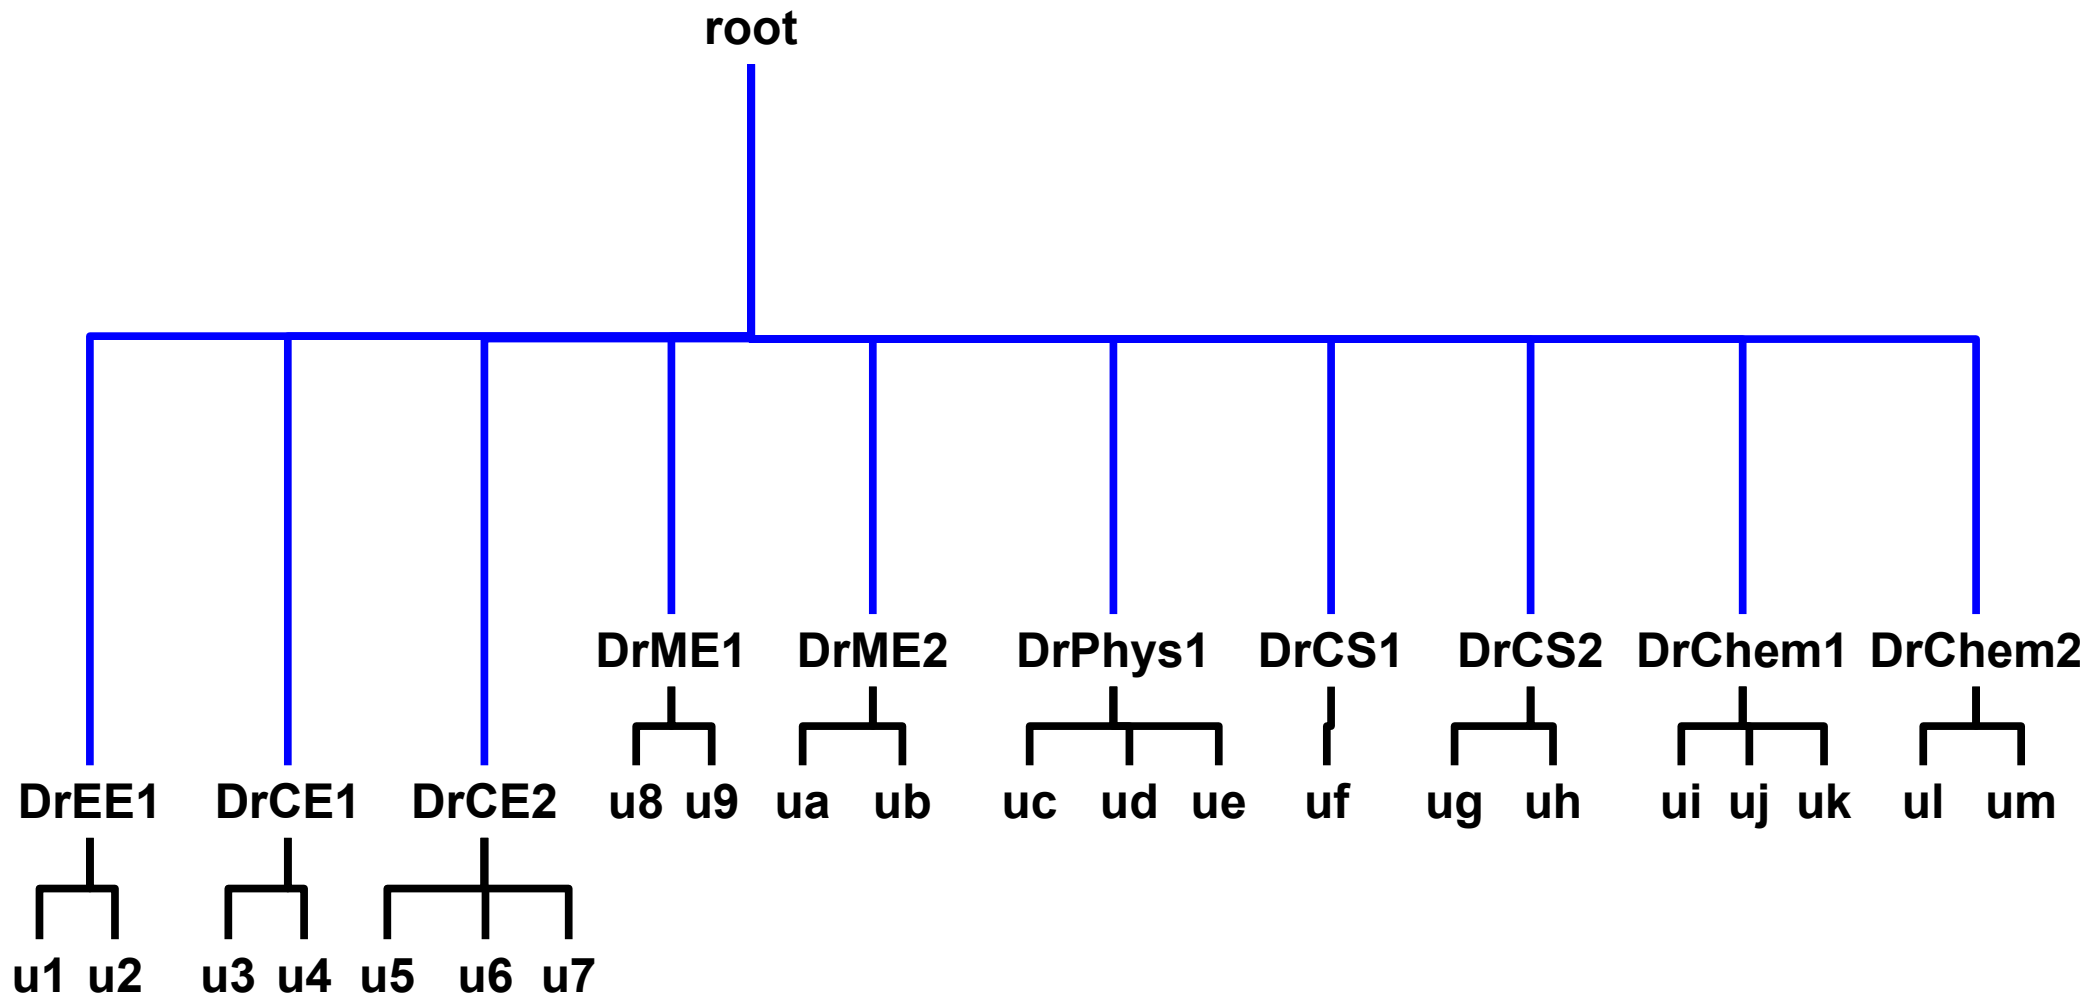
# Fairshare=parent



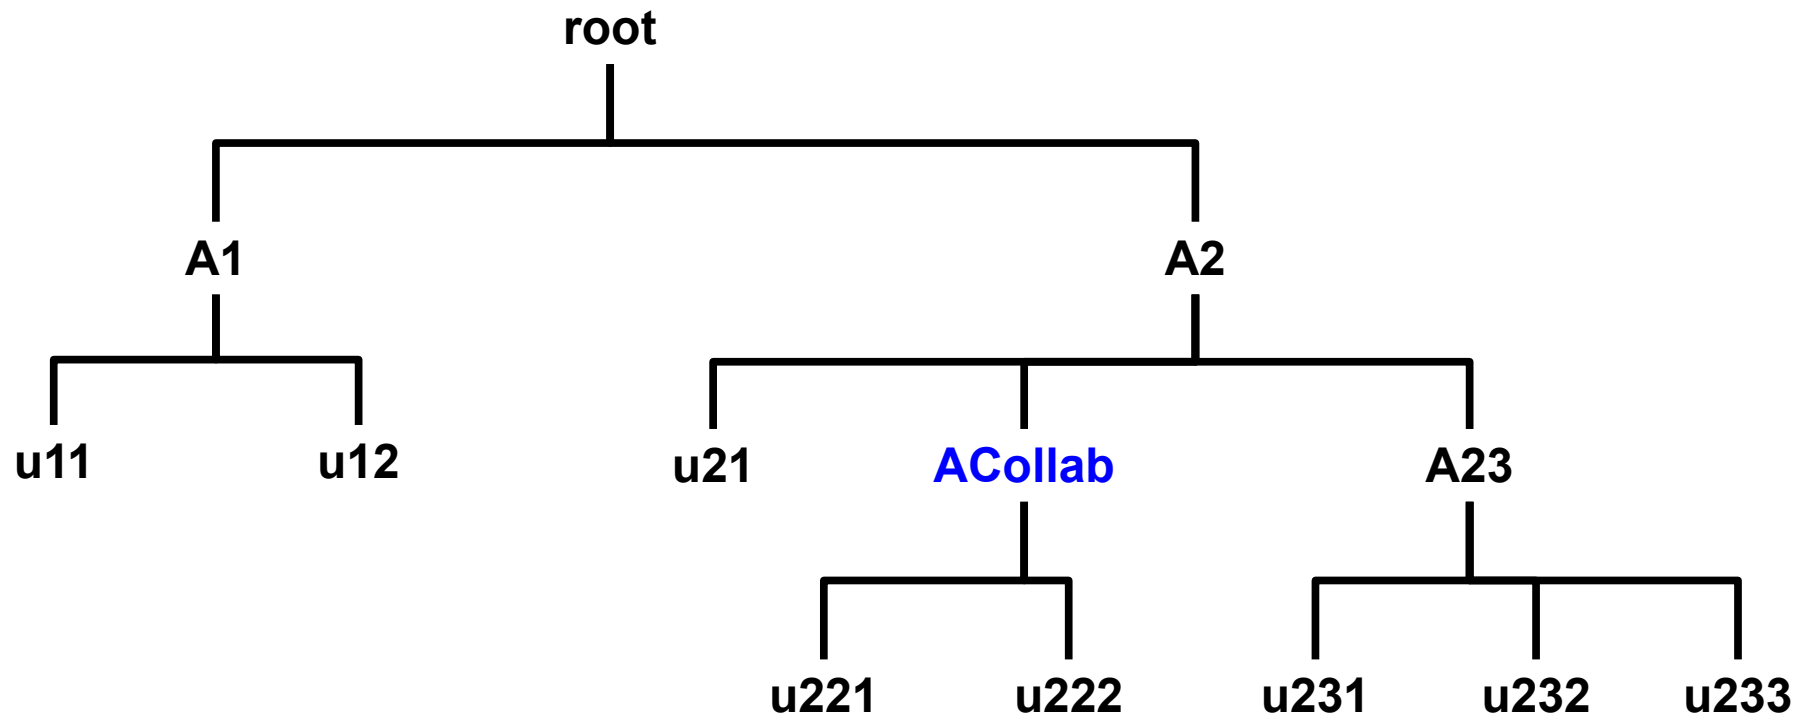
# Fairshare=parent



# Fairshare=parent




# Fairshare=parent





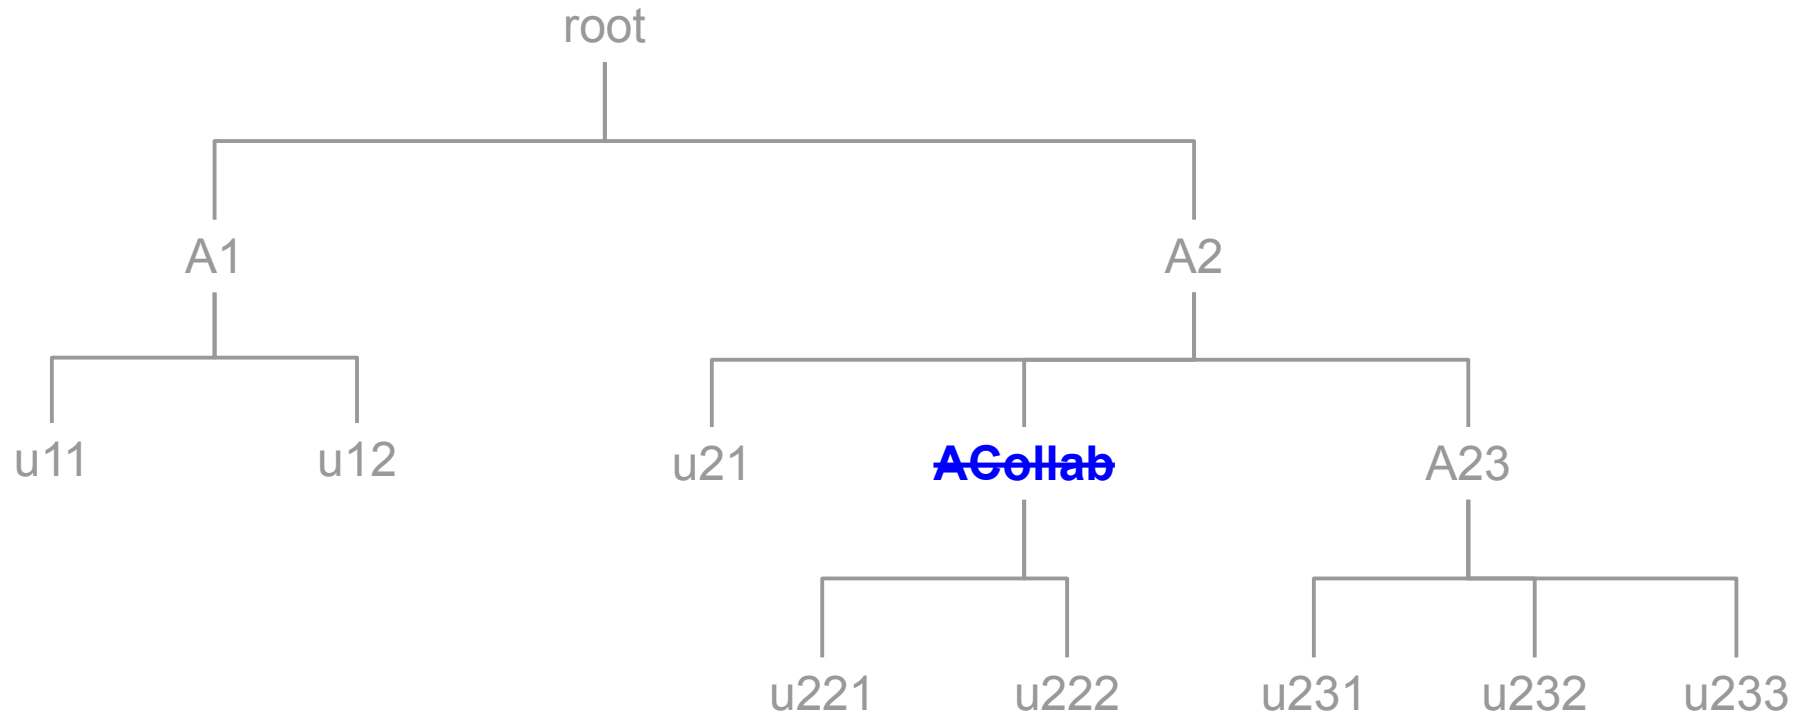
# Fairshare=parent



Example:


**ACollab** is the only association with Fairshare=parent

# Fairshare=parent




Example:  
**ACollab** effectively disappears

# Fairshare=parent




Example:  
**u221** and **u222** are reparented

# Fairshare=parent



The following are now **siblings** for purposes of fairshare calculations:  
u21, u221, u222, A23

# Fairshare=parent



The following are now **siblings** for purposes of fairshare calculations:  
u21, u221, u222, A23

Questions?

# Appendix

# Appendix Index

- More information
- Future Development
- Problems with Shares Calculation
- Linear Interpolation
- sshare example
- Tie Handling



# More Information

- Until 14.11 is released, Fair Tree documentation is at:
  - [https://fsl.byu.edu/documentation/slurm/fair\\_tree.php](https://fsl.byu.edu/documentation/slurm/fair_tree.php)
- Available in 14.11.0pre6

# Future Development

- Allow for less strictness
  - Merge accounts based on configurable delta
    - Allows reuse of tie handling. The comparison function needs a very minor change
  - LevelFSTieDelta=0.01,0.05,0.02
    - Values are the epsilon value at increasing depths
  - Either:
    - $\text{abs}((a-b)<\text{delta})$
    - $a > b * (1 - \text{delta}) \ \&\& \ a < b * (1 + \text{delta})$

# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms:

$$S = (S_{\text{user}} / S_{\text{siblings}}) *$$

$$(S_{\text{account}} / S_{\text{sibling-accounts}}) *$$

$$(S_{\text{parent}} / S_{\text{parent-siblings}}) * \dots$$

# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms:

$$S = (S_{\text{user}} / S_{\text{siblings}}) * \\ (S_{\text{account}} / S_{\text{sibling-accounts}}) * \\ (S_{\text{parent}} / S_{\text{parent-siblings}}) * \dots$$

Example:

- The tree is no deeper than: root->account->user
- Assume sibling associations are treated equally (same Shares)
- Problem still applies for more complicated scenarios

# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms:

$$S = (S_{\text{user}} / S_{\text{siblings}}) *$$

$$(S_{\text{account}} / S_{\text{sibling-accounts}}) *$$

~~$$(S_{\text{parent}} / S_{\text{parent-siblings}}) * \dots$$~~

- The tree is no deeper than: root->account->user
- Therefore:

$$S = (S_{\text{user}} / S_{\text{siblings}}) * (S_{\text{account}} / S_{\text{sibling-accounts}})$$

# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms:

$$S = (S_{\text{user}} / S_{\text{siblings}}) *$$

$$(S_{\text{account}} / S_{\text{sibling-accounts}}) *$$

~~$$(S_{\text{parent}} / S_{\text{parent-siblings}}) * \dots$$~~

- If each account is Shares=500
- $S_{\text{account}} / S_{\text{sibling-accounts}} = 500 / (500 * \text{count}(\text{accounts}))$
- $S_{\text{account}} / S_{\text{sibling-accounts}} = 1 / \text{count}(\text{accounts})$
- This is a constant for all accounts**

# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms:

$$S = (S_{\text{user}} / S_{\text{siblings}}) * \text{constant}$$

~~$$(S_{\text{account}} / S_{\text{sibling-accounts}}) *$$~~

~~$$(S_{\text{parent}} / S_{\text{parent-siblings}}) * \dots$$~~

- $S = (S_{\text{user}} / S_{\text{siblings}}) * \text{constant}$
- Since we are comparing... constants are thrown out:
- $S = (S_{\text{user}} / S_{\text{siblings}})$

# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms:

$$S = (S_{\text{user}} / S_{\text{siblings}}) * \text{constant}$$

---

$$(S_{\text{account}} / S_{\text{sibling-accounts}}) *$$

---

$$(S_{\text{parent}} / S_{\text{parent-siblings}}) * \dots$$

- If each user is Shares=100
- $S = 100 / (100 * \text{count}(\text{users\_in\_account}))$
- $S = 1 / \text{count}(\text{users\_in\_account})$




# Problems With Shares Calculation

- Normalized Shares equation for all current algorithms when accounts are treated equally and users in an account are treated equally:

$$S_{\text{user}} = 1 / \text{count}(\text{users\_in\_account}) * \text{constant}$$

**The problem still exists for other scenarios but is harder to model**

# Problems With Shares Calculation




$$S = 1/2 = 0.5$$

$$S = 1/10 = 0.1$$

# Linear Interpolation (lerp)

- Shares are between 0.0 .. 1.0
- Map them to the range 0.1 .. 1.0
  - Use linear interpolation (lerp)
- $\text{lerp}(\text{min}, \text{max}, f) = \text{min} + f * (\text{max} - \text{min})$ 
  - Has floating point issues
- $\text{lerp}(\text{min}, \text{max}, f) = \text{min} * (1.0L - f) + \text{max} * f$ 
  - Shouldn't have floating point issues
- $\text{lerp}(0.1, 1.0, 0.000001) = 0.1000009$

# Linear Interpolation (lerp)




# Unmodified vs lerp()

Usage	Shares	Fairshare (unmodified)	Fairshare (lerp)
0.15	.2	0.59460355750136062447097629046766	0.68981706017270247755391460176000
0.15	.1	0.35355339059327384187980669594253	0.57855511864135993283170603107557
0.15	.02	0.00552427172801990291201024163570	0.41431890600793676033685580817334
0.15	.01	0.00003051757812500000000000000000	0.38524635476490996013063453085046
0.15	.002	0.00000000000000000000000000000000	0.36011325130027945195512403819826
0.15	.001	0.00000000000000000000000000000000	0.35684750329912502426078085848715
0.15	.0002	0.00000000000000000000000000000000	0.35421449328863424725827649397480
0.15	.0001	0.00000000000000000000000000000000	0.35388408479113951296806533930450
0.15	.000000000002	0.00000000000000000000000000000000	0.35355339065944120049999645216676
0.15	.000000000001	0.00000000000000000000000000000000	0.35355339062635748816239289471497

- Fairshare factor is multiplied by PriorityWeightFairshare, making the problem much worse
- If lerp() is implemented, add an option (-L?) to sshare:
  - change Norm Shares → Lerp Shares

# sshare


- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

McCartney: Why do I have a lower fairshare factor than Elvis?

# sshare


- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  

$$\text{assoc} / (\text{assoc} + \text{siblings})$$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

McCartney: Why do I have a lower fairshare factor than Elvis?

# sshare


- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

McCartney: Why do I have a lower fairshare factor than Elvis?



# sshare


- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

Answer: Compare the accounts' Level FS values

# sshare


- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  

$$\text{assoc} / (\text{assoc} + \text{siblings})$$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

Answer: Compare the accounts' Level FS values


# sshare

- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

McCartney: How is Level FS calculated?


# sshare

- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

Answer: Norm Shares / Effectv Usage

# sshare

- **Norm Shares** and **Effectv Usage** – reflect the Fair Tree approach  
 $\text{assoc} / (\text{assoc} + \text{siblings})$
-  **Level FS** (*with -l option*) – Fairshare compared to siblings
- **FairShare** – This is the final fairshare factor (i.e. it works like it should)

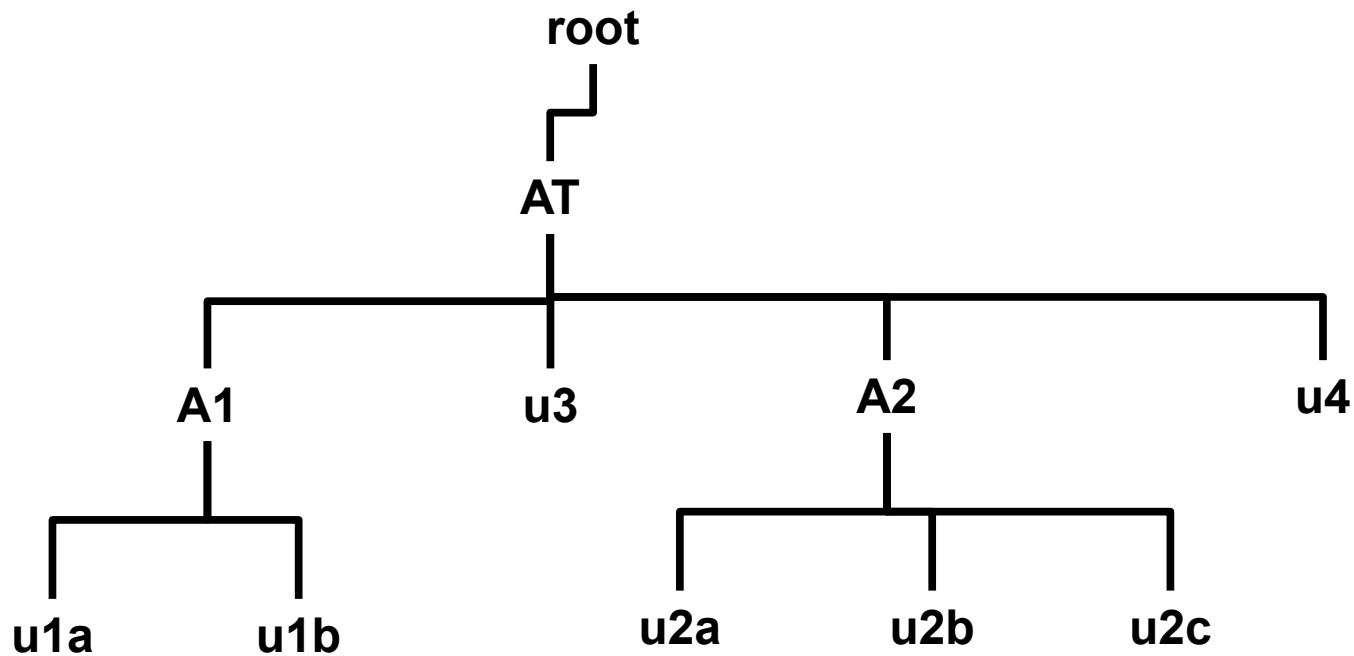
Account	User	Raw Shares	Norm Shares	Raw Usage	Norm Usage	Effectv Usage	FairShare	Level FS
root			0.000000	1230		1.000000		1.000000
beatles		500	0.500000	676	0.549593	0.549593		0.909763
beatles	harrison	25	0.250000	301	0.244715	0.445266	0.200000	0.561462
beatles	lennon	25	0.250000	102	0.082927	0.150888	0.600000	1.656863
beatles	mccartney	25	0.250000	37	0.030081	0.054734	0.800000	4.567568
beatles	starr	25	0.250000	236	0.191870	0.349112	0.400000	0.716102
elvis		500	0.500000	554	0.450407	0.450407		1.110108
elvis	elvis	1	1.000000	554	0.450407	1.000000	1.000000	1.000000

Answer: Norm Shares / Effectv Usage

# Ties

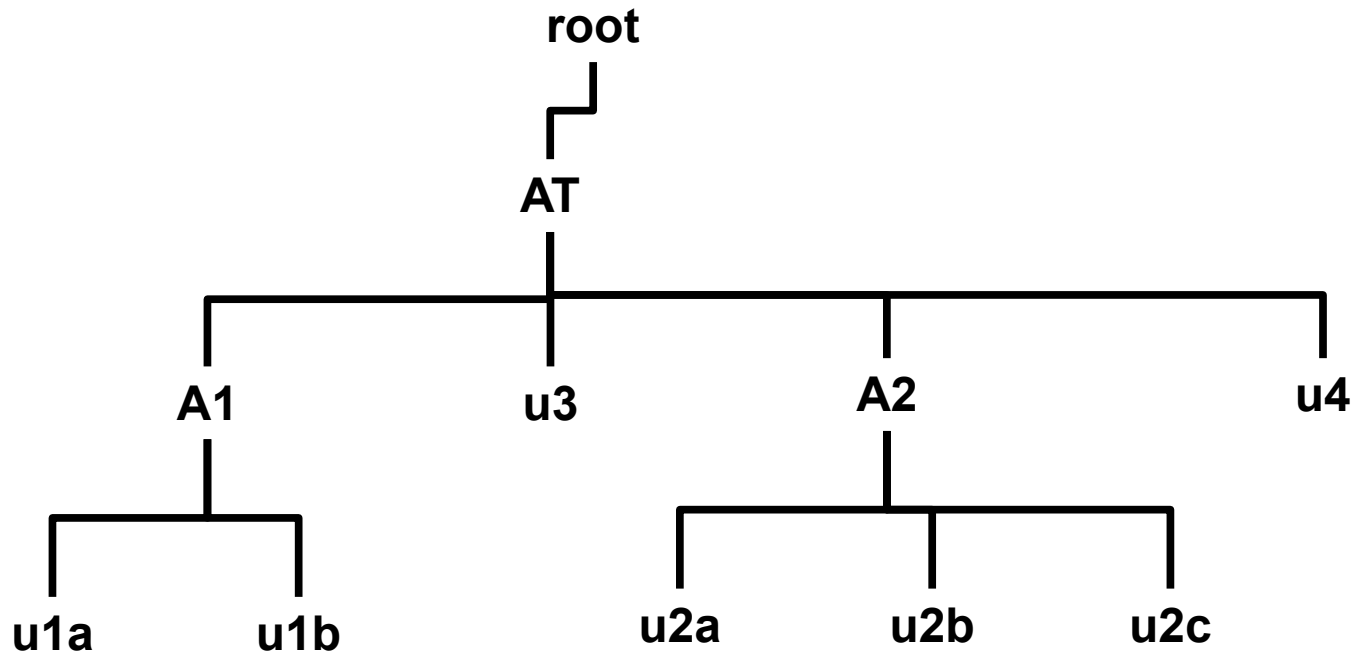
- Occur when sibling associations have the same level\_fs
  - Extremely rare except when raw usage == 0.0
  - Added complexity to the code
- Rules:
  - Sibling **users** receive the same rank
  - Sibling **accounts** have their children lists merged
  - A **user** with the same level\_fs as a sibling **account** will receive the same rank as the account's highest ranked user

# Fair Tree: Tie Handling



This is one portion of the tree. AT is pictured in its entirety.

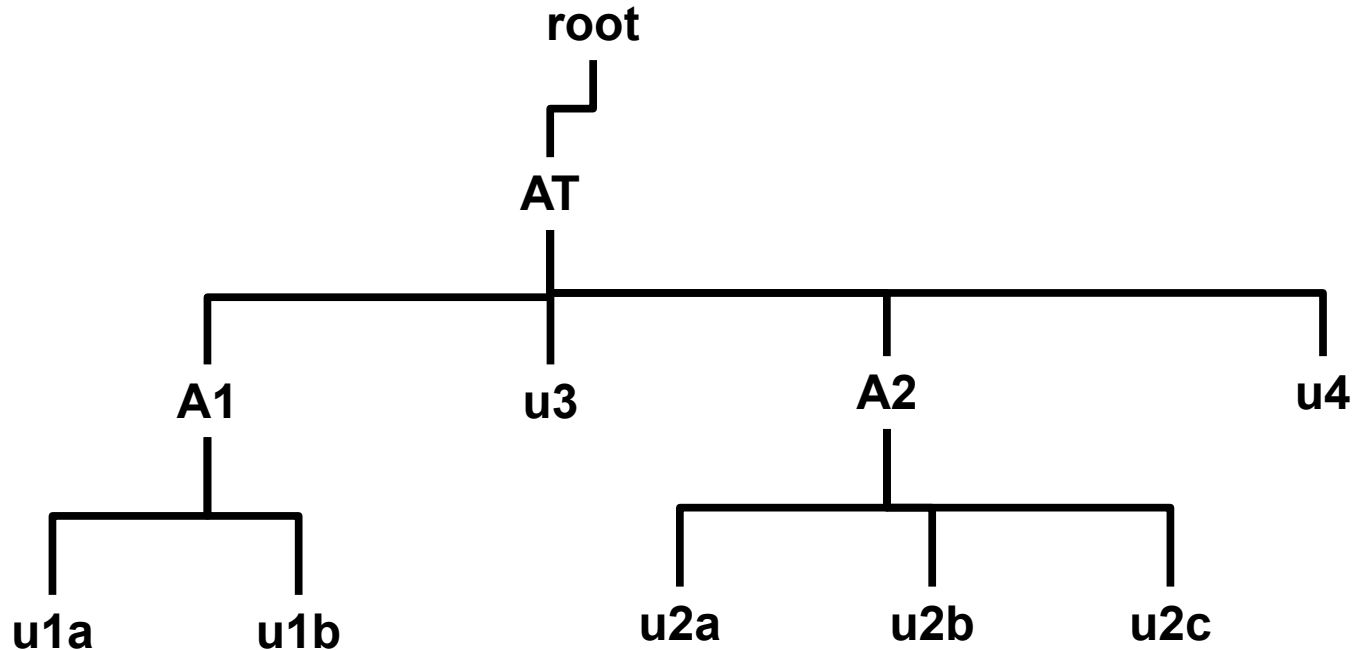
# Fair Tree: Tie Handling



There are 83 users (not all pictured).  
Some have been visited.



# Fair Tree: Tie Handling

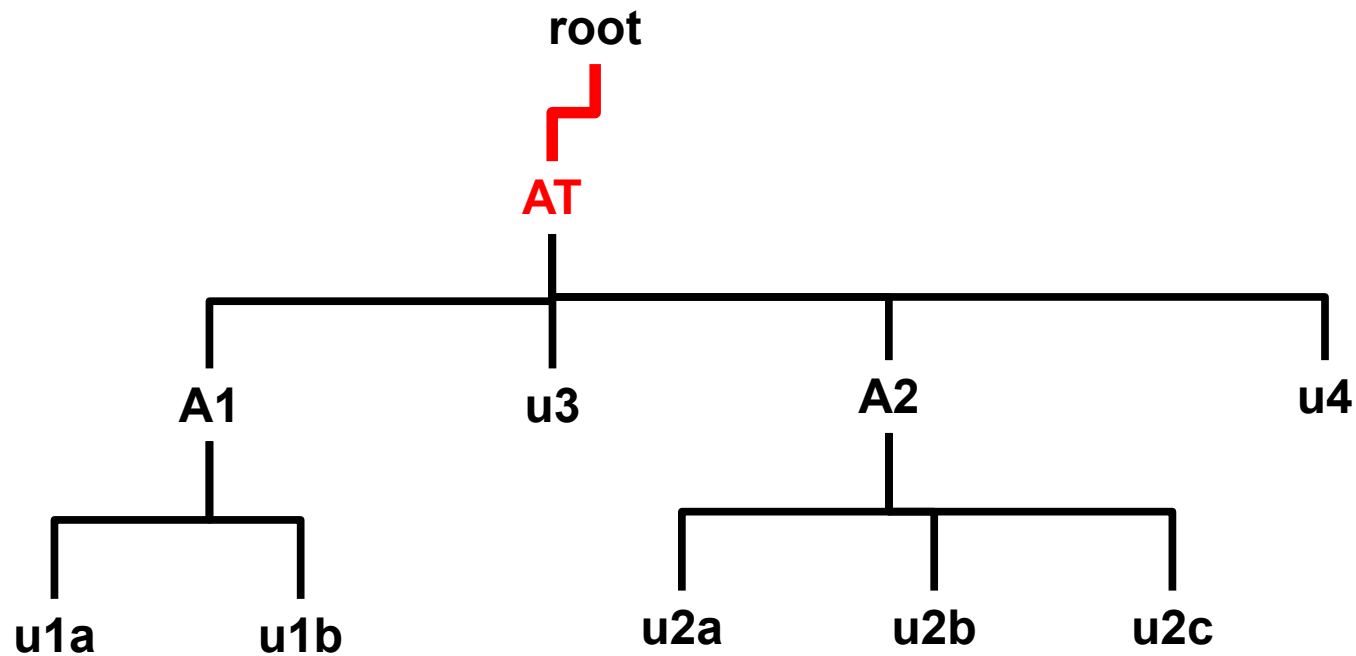


There are 83 users (not all pictured).  
Some have been visited.




rank = 64

# Fair Tree: Tie Handling



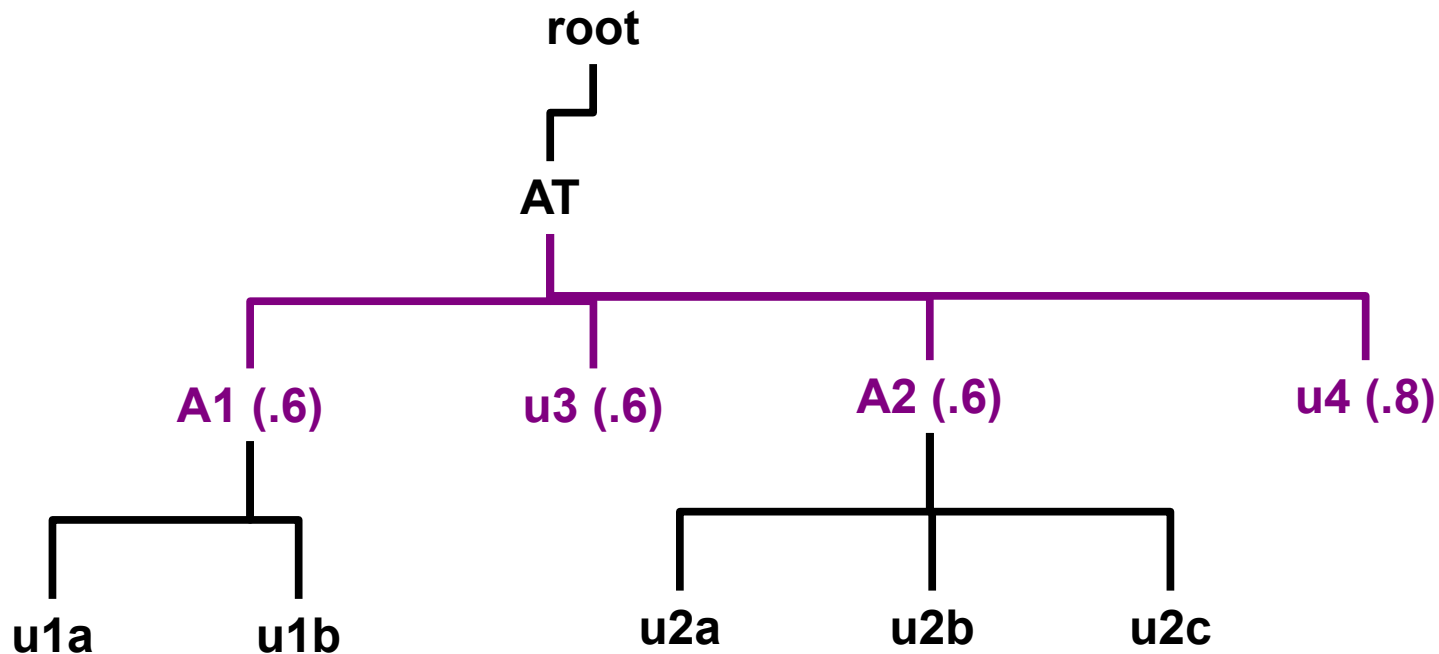
rank = 64

# Fair Tree: Tie Handling



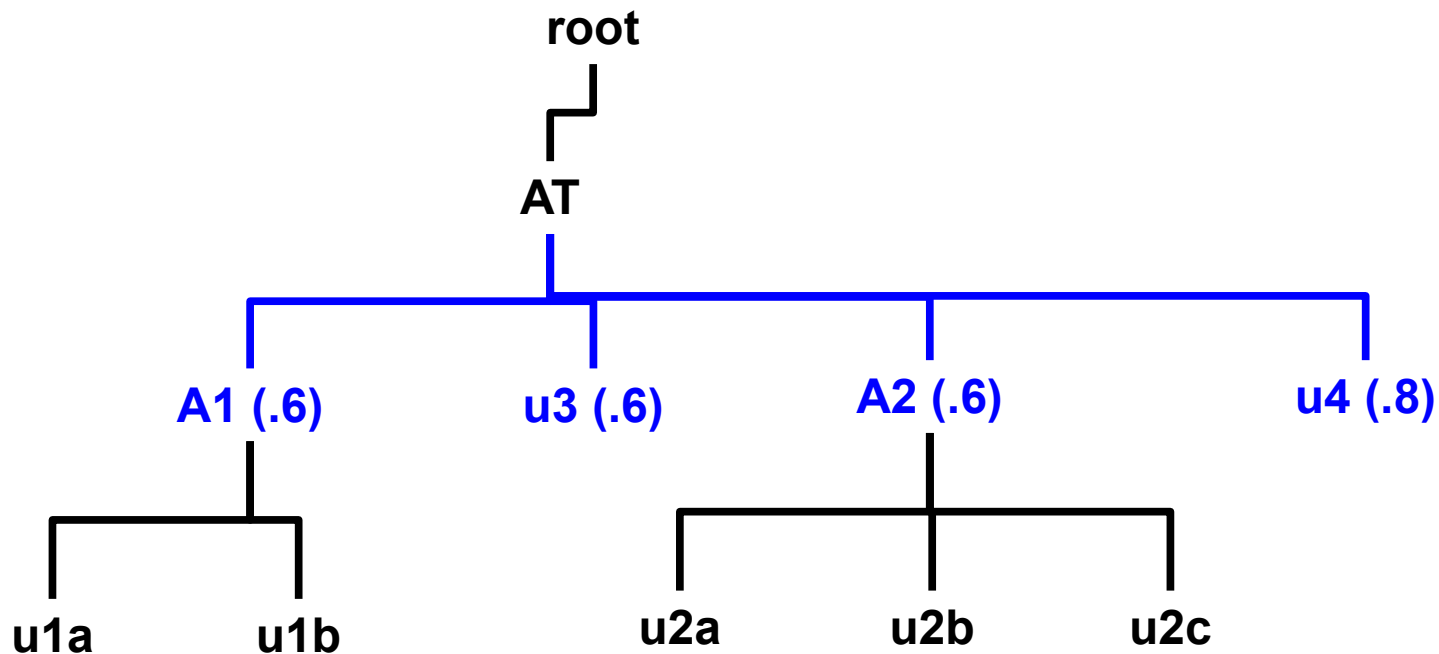
rank = 64

# Fair Tree: Tie Handling



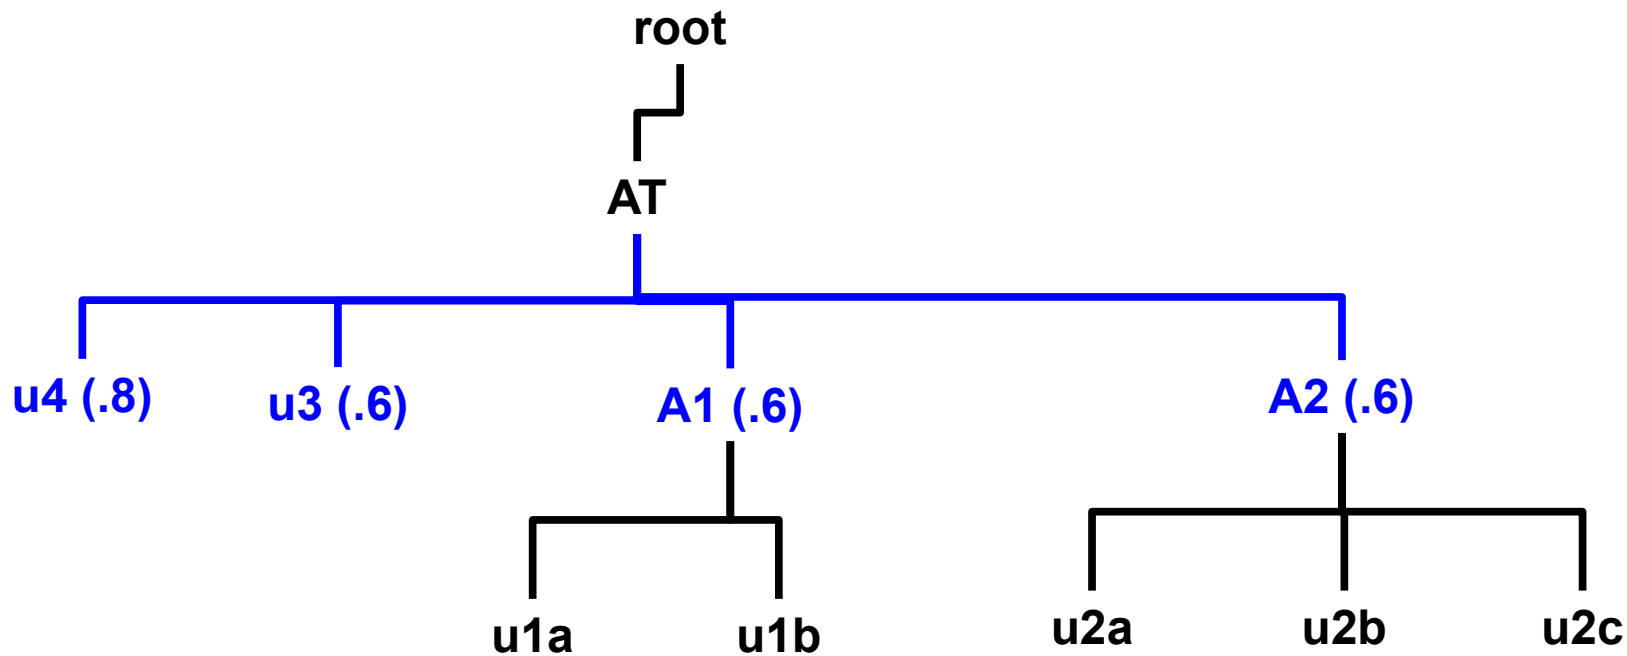
rank = 64

# Fair Tree: Tie Handling




rank = 64

# Fair Tree: Tie Handling



rank = 64


# Fair Tree: Tie Handling



u3, A1, and A2 are tied at .6

rank = 64

# Fair Tree: Tie Handling

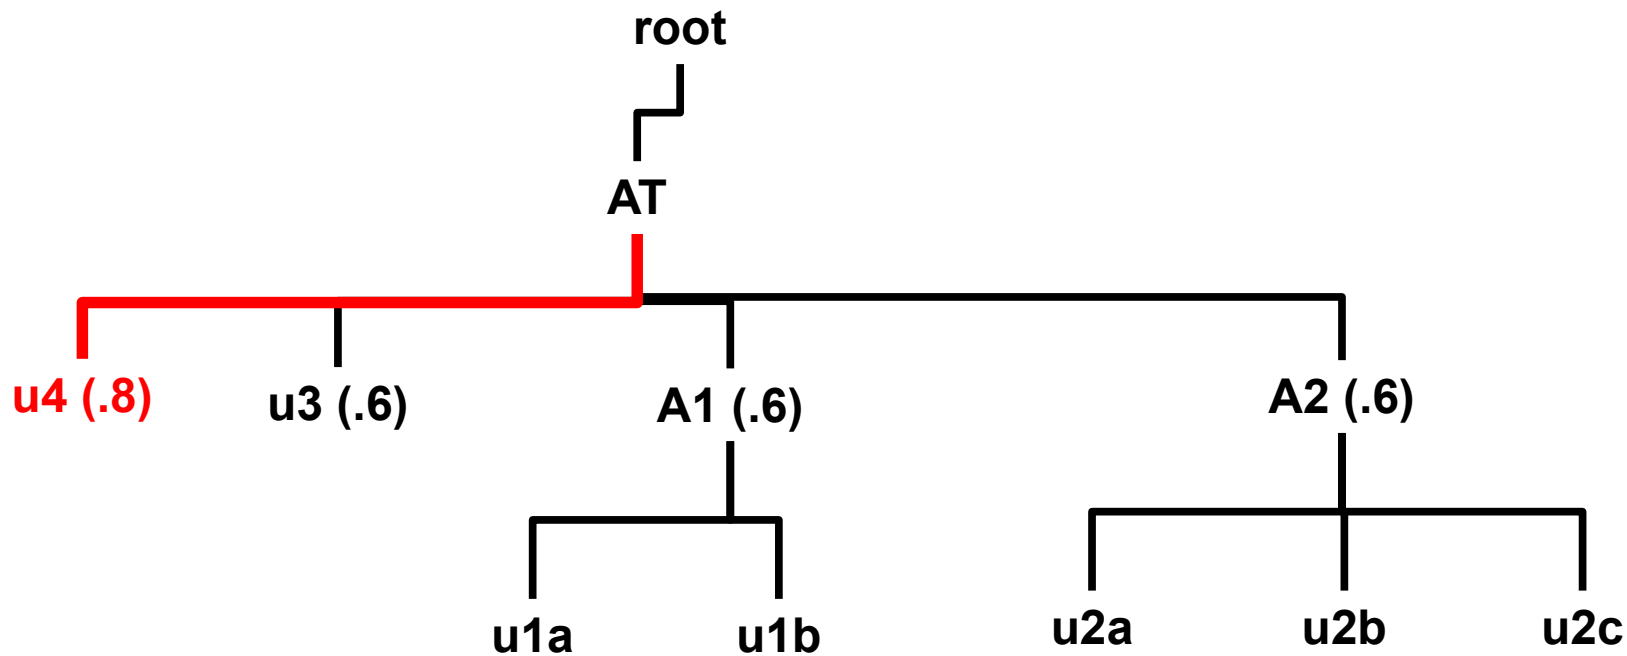


Tied users are placed to the left

rank = 64

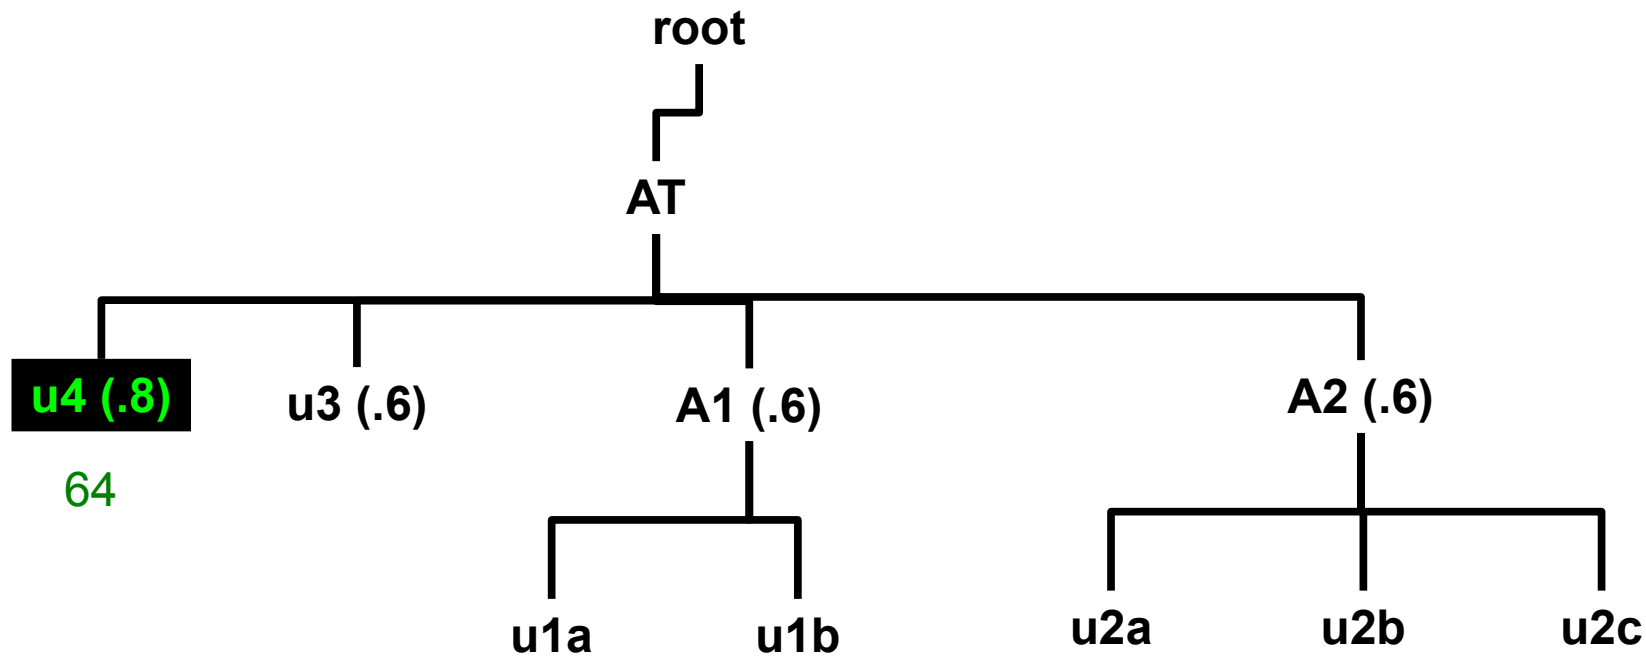


# Fair Tree: Tie Handling



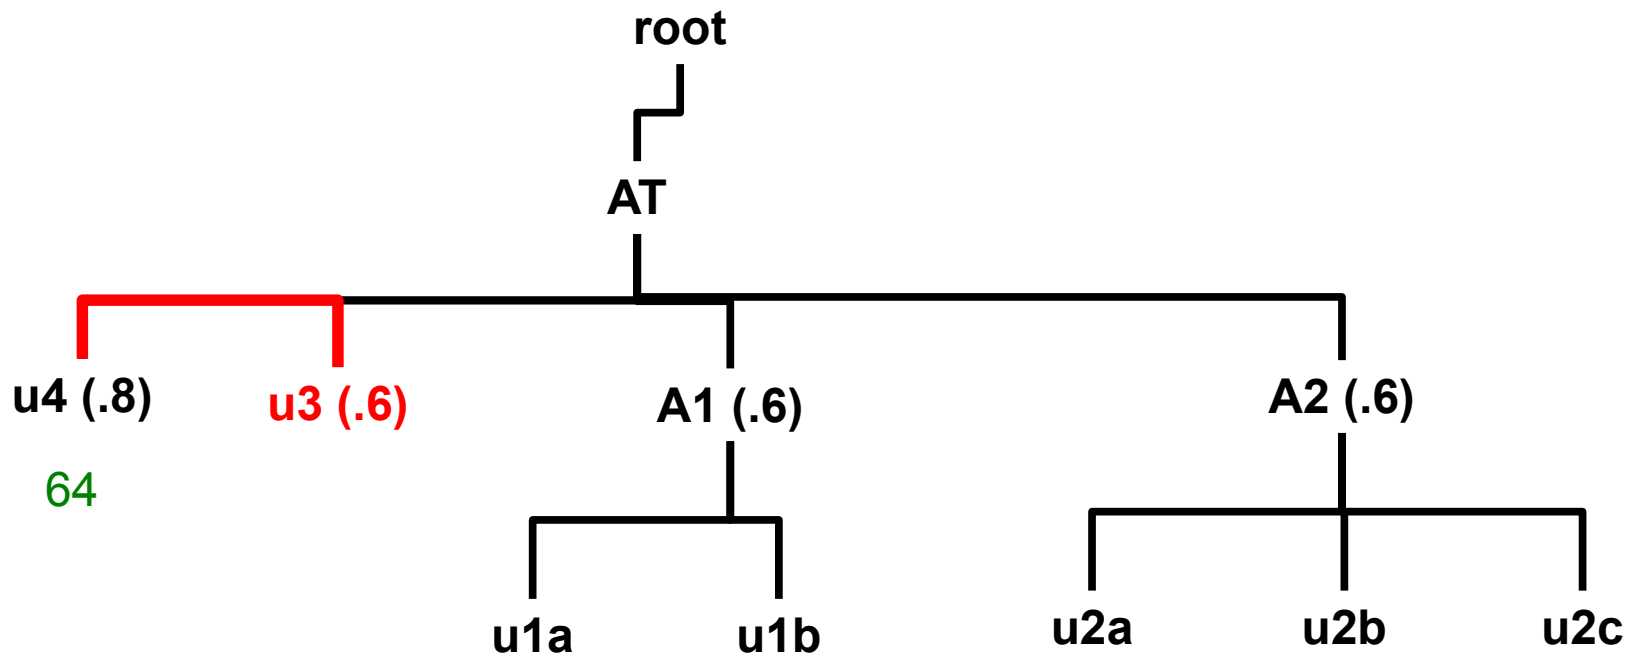
rank = 64

# Fair Tree: Tie Handling



rank = 64

# Fair Tree: Tie Handling




Next association **not** tied. Decrement rank.




rank = 63

# Fair Tree: Tie Handling



rank = 63


# Fair Tree: Tie Handling



Next association **is** tied. rank stays the same

rank = 63


# Fair Tree: Tie Handling



Next association **is** tied. rank stays the same

rank = 63


# Fair Tree: Tie Handling



Peek at next account to check for a tie

rank = 63

# Fair Tree: Tie Handling




$A1 \rightarrow \text{level\_fs} == A2 \rightarrow \text{level\_fs}$

rank = 63



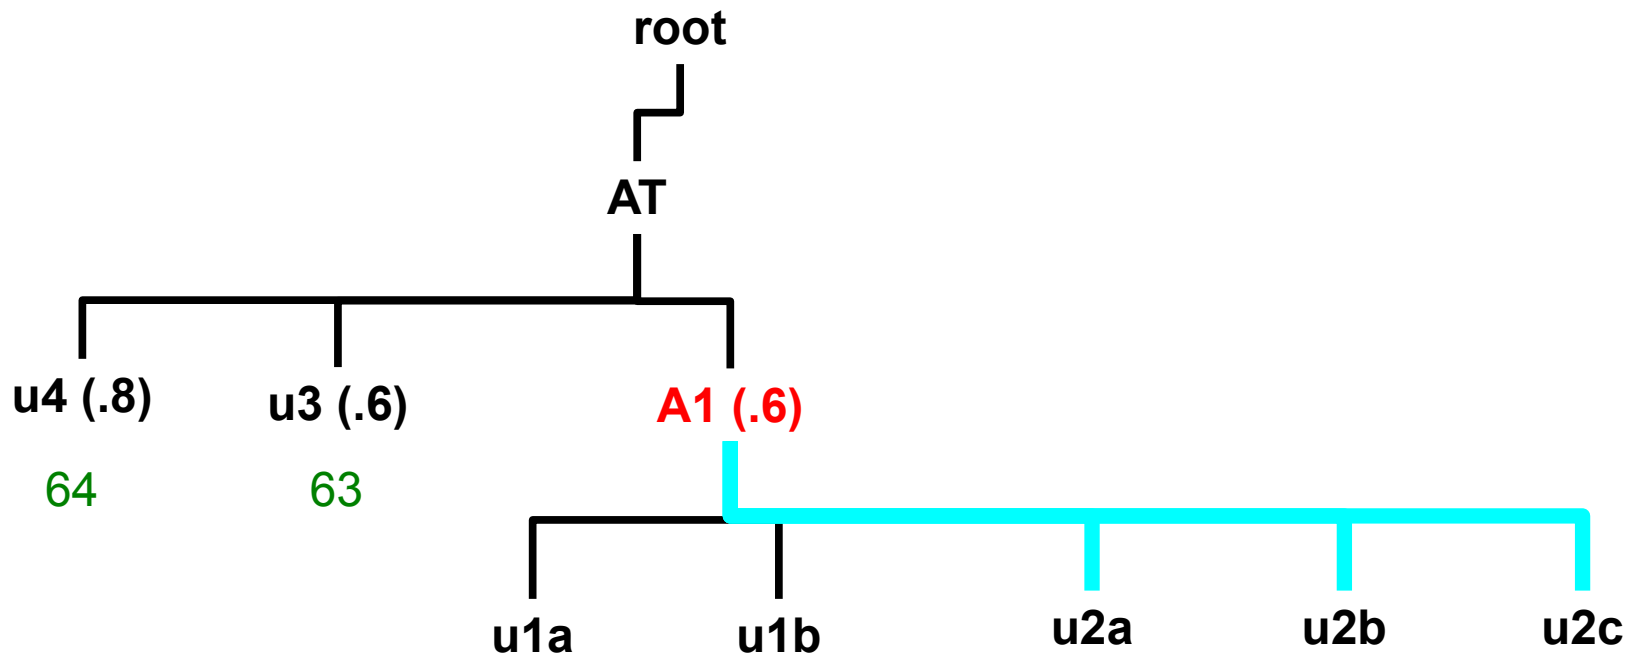
# Fair Tree: Tie Handling



Accounts are equal. Merge!

rank = 63

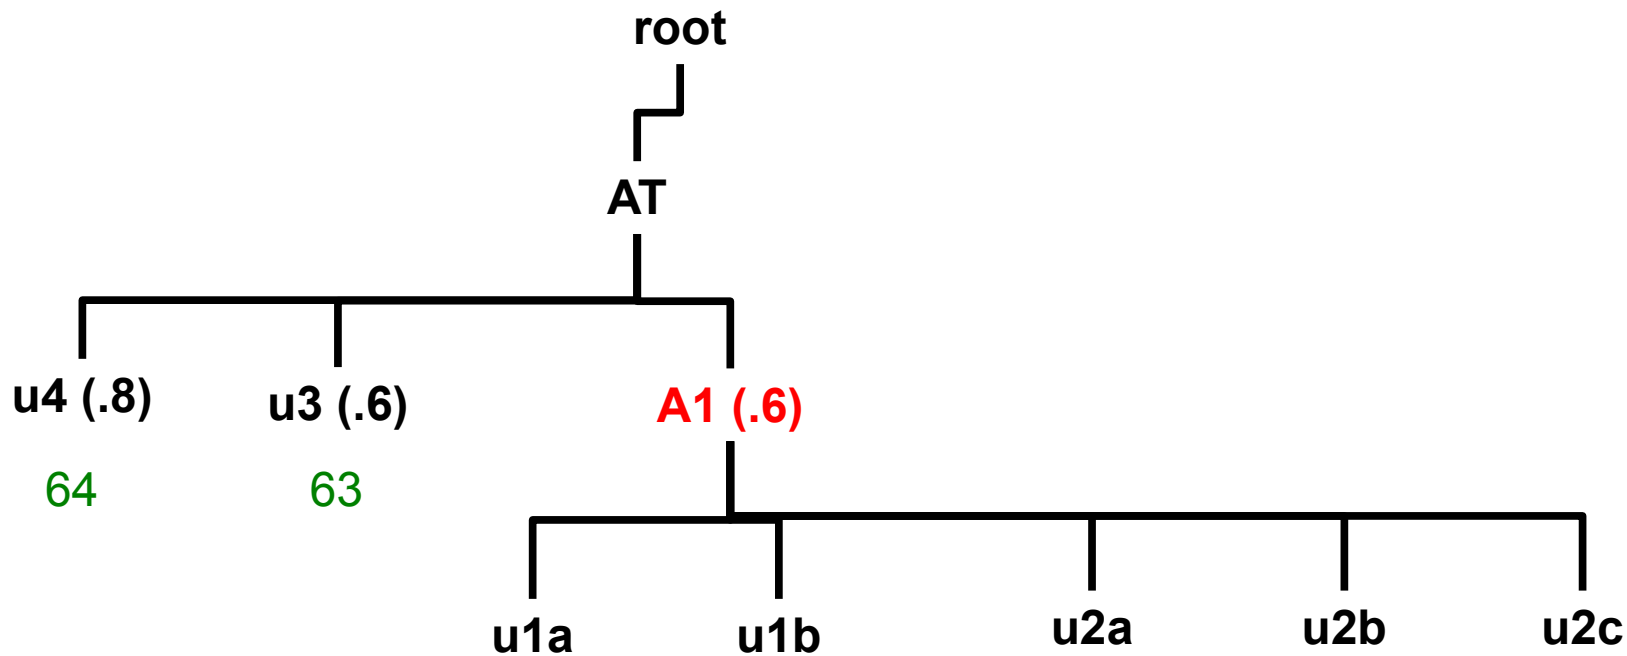
# Fair Tree: Tie Handling



Merge complete

rank = 63

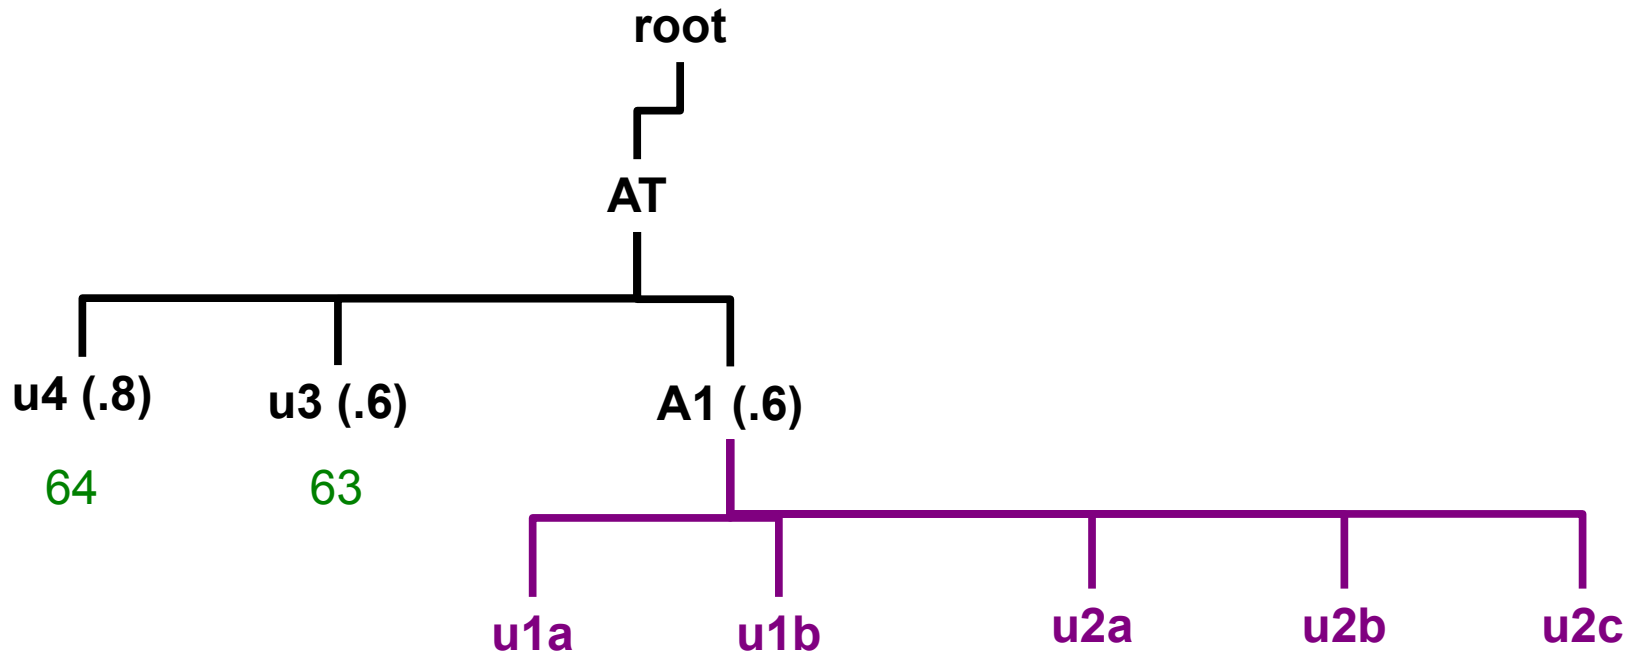
# Fair Tree: Tie Handling



Merge complete

rank = 63


# Fair Tree: Tie Handling



Calculate level\_fs using **actual siblings**, **not** adopted ones.  
The calculations pretend that the merge never happened.

rank = 63

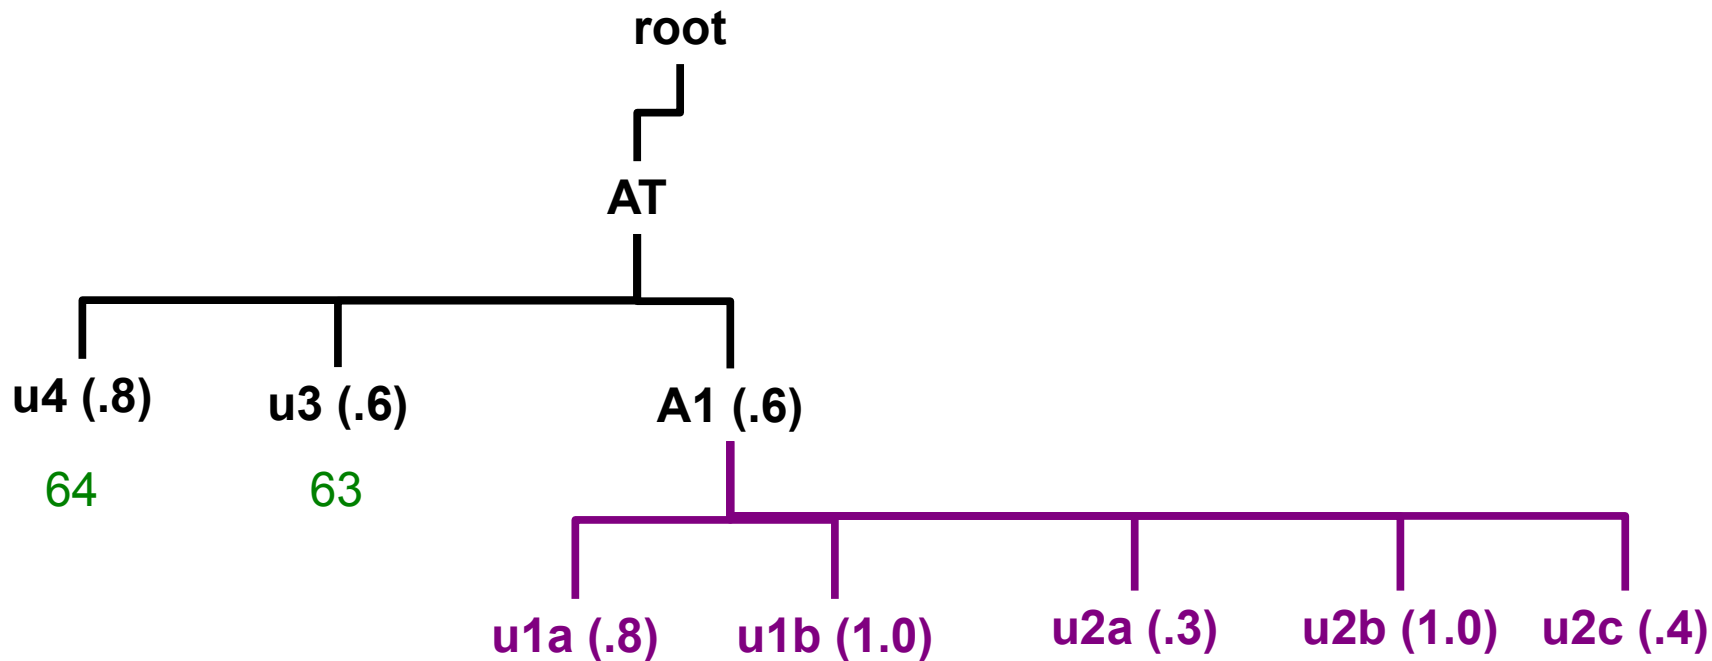
# Fair Tree: Tie Handling



Calculate level\_fs using **actual siblings**, **not** adopted ones.  
The calculations pretend that the merge never happened.

rank = 63

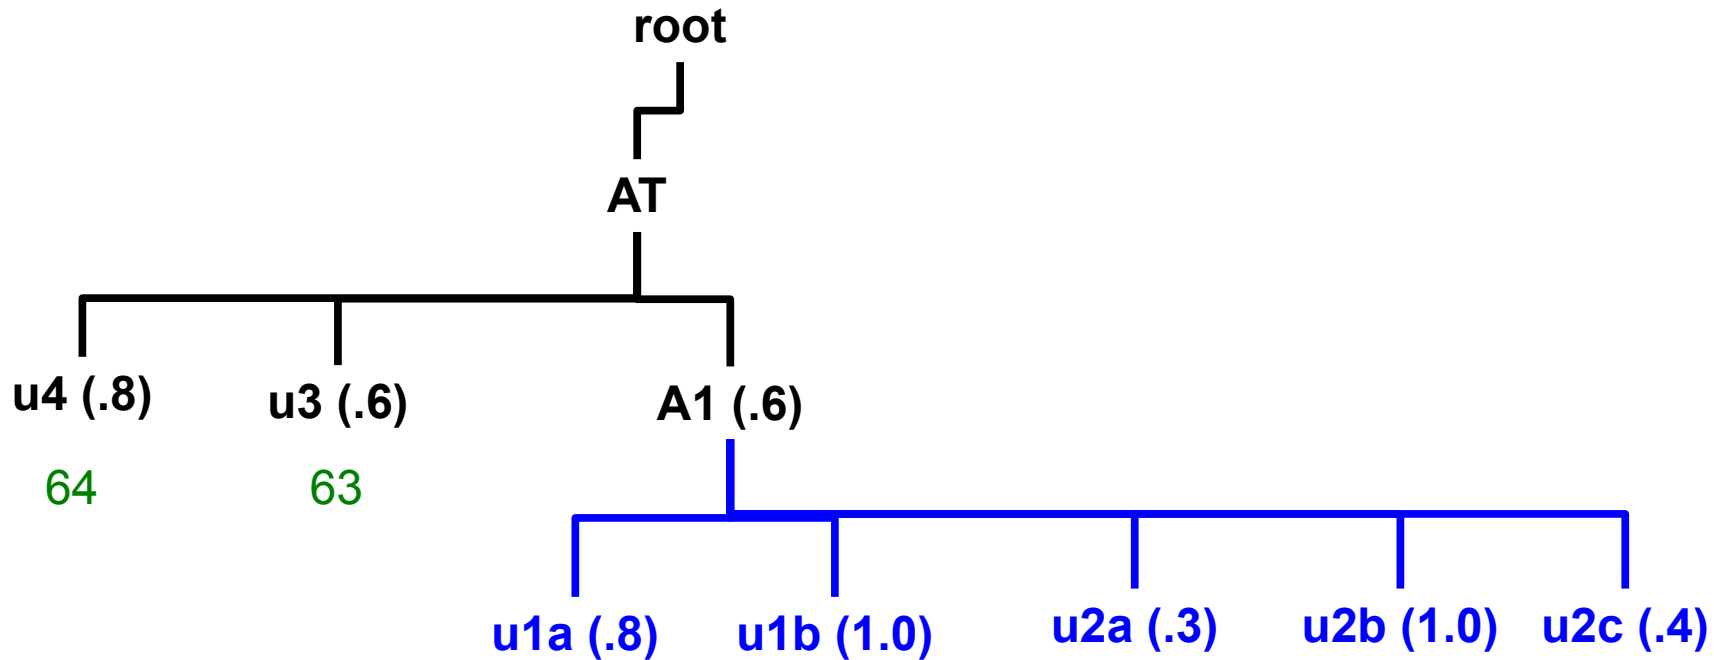
# Fair Tree: Tie Handling



Calculate level\_fs using **actual siblings**, **not** adopted ones.  
The calculations pretend that the merge never happened.

rank = 63

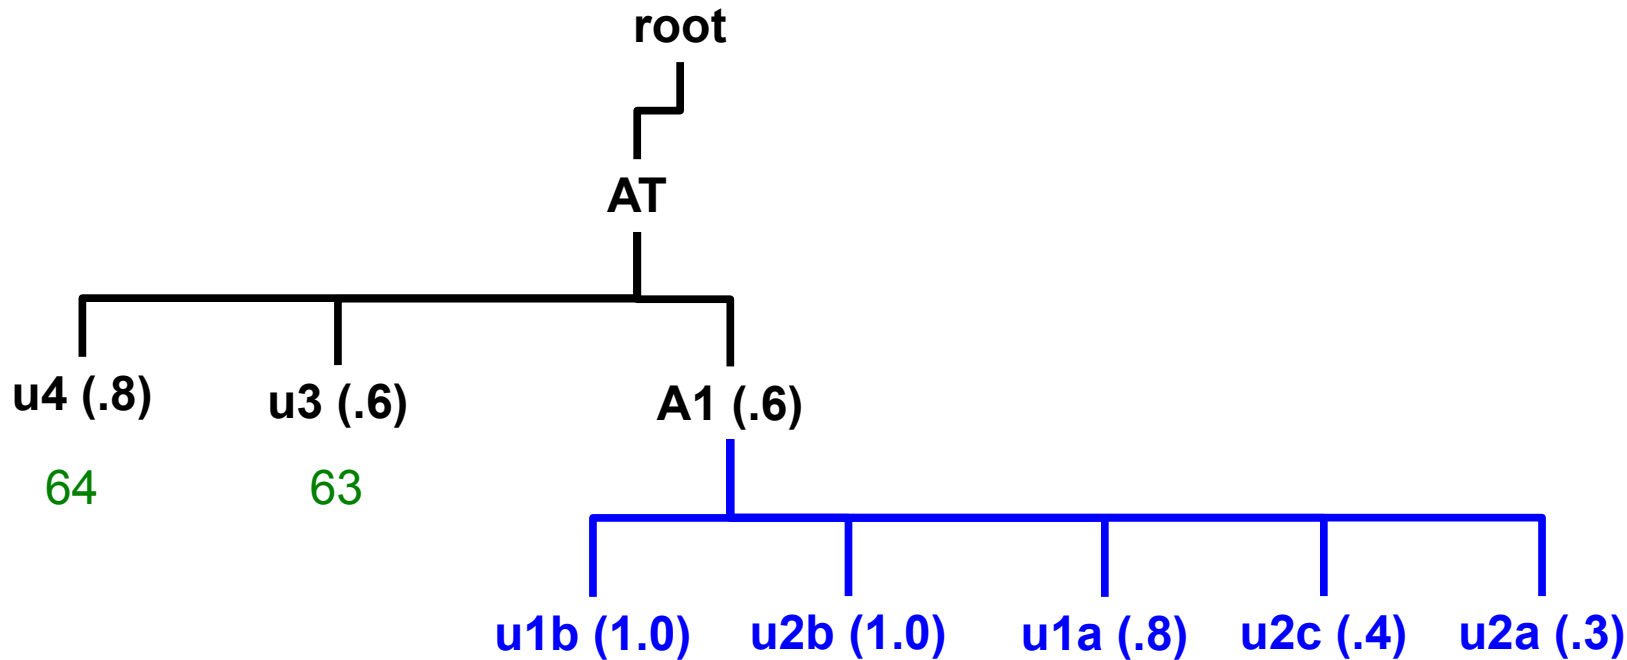
# Fair Tree: Tie Handling



Sorting uses the merged list

rank = 63

# Fair Tree: Tie Handling

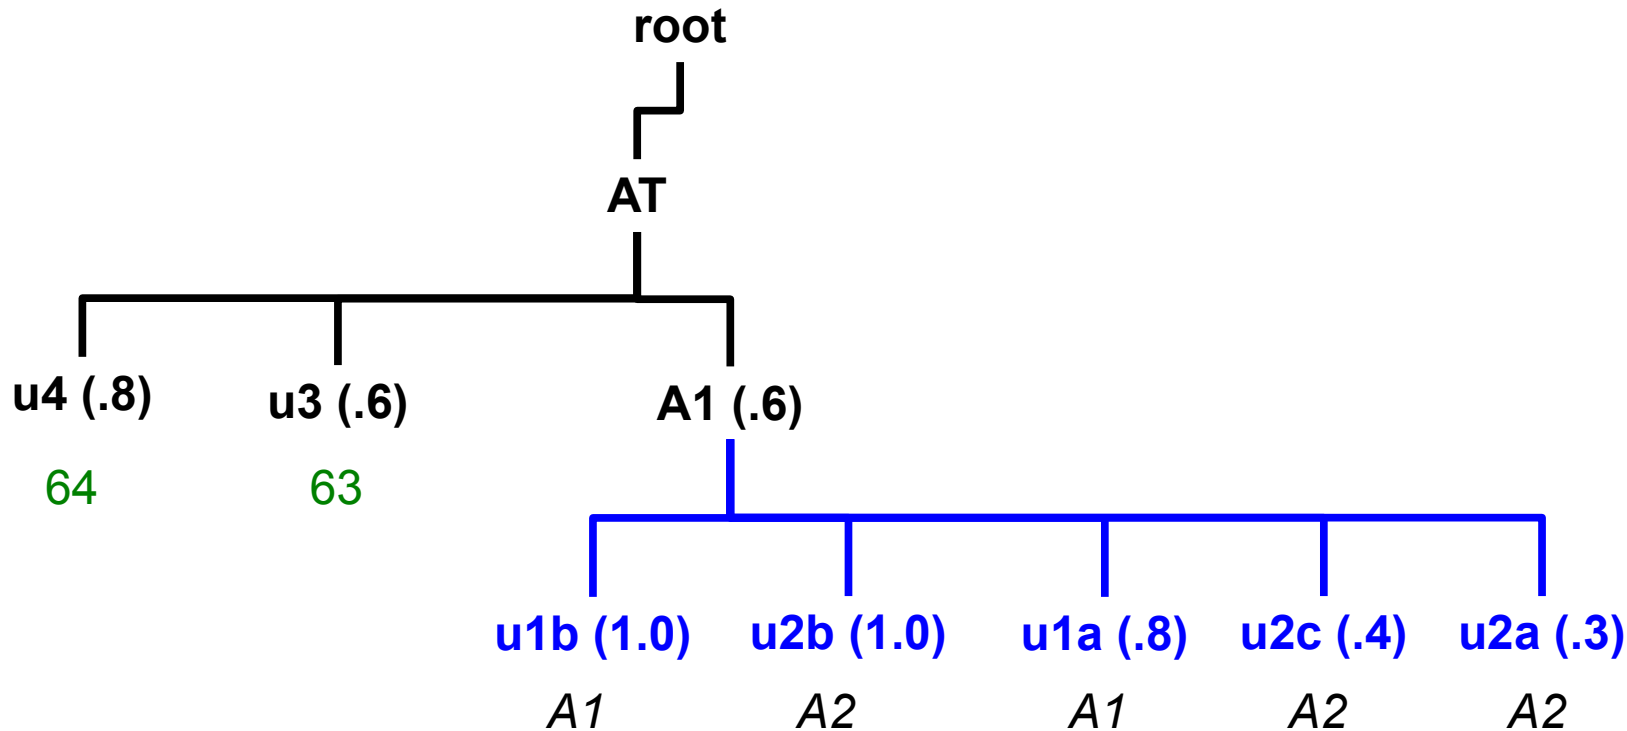


Sorting uses the merged list

rank = 63



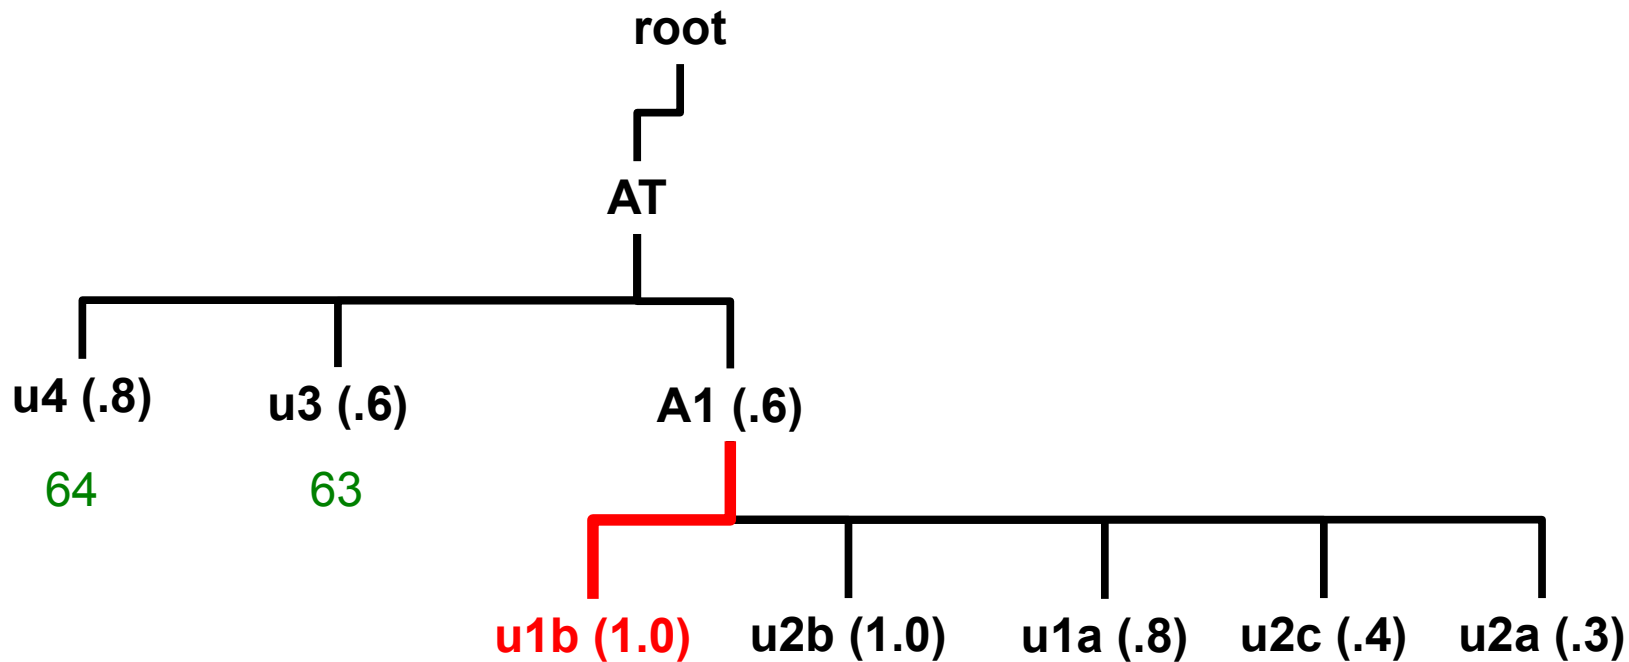
# Fair Tree: Tie Handling



Note that the merged users are mixed

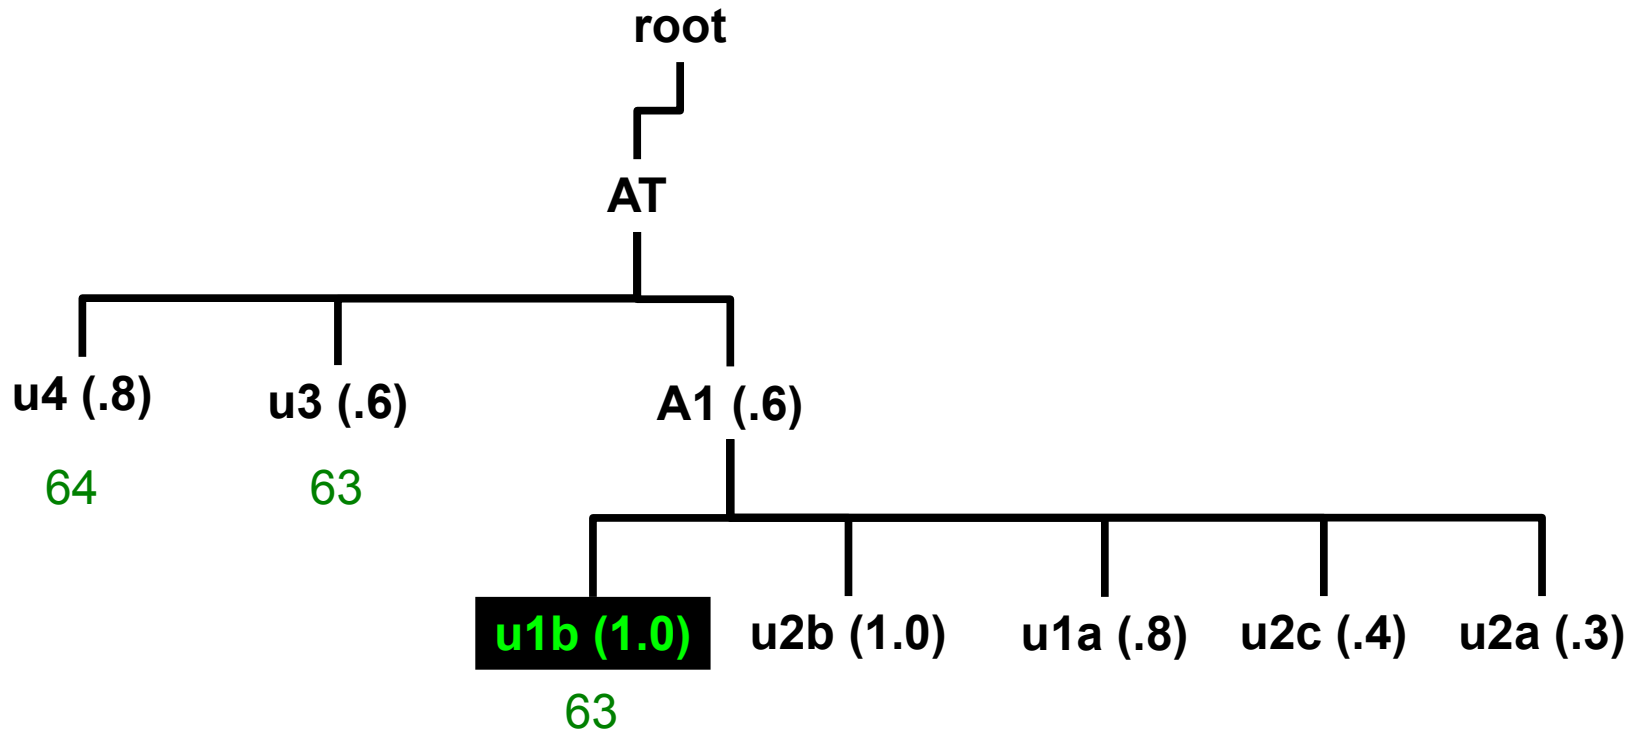
rank = 63

# Fair Tree: Tie Handling



rank = 63

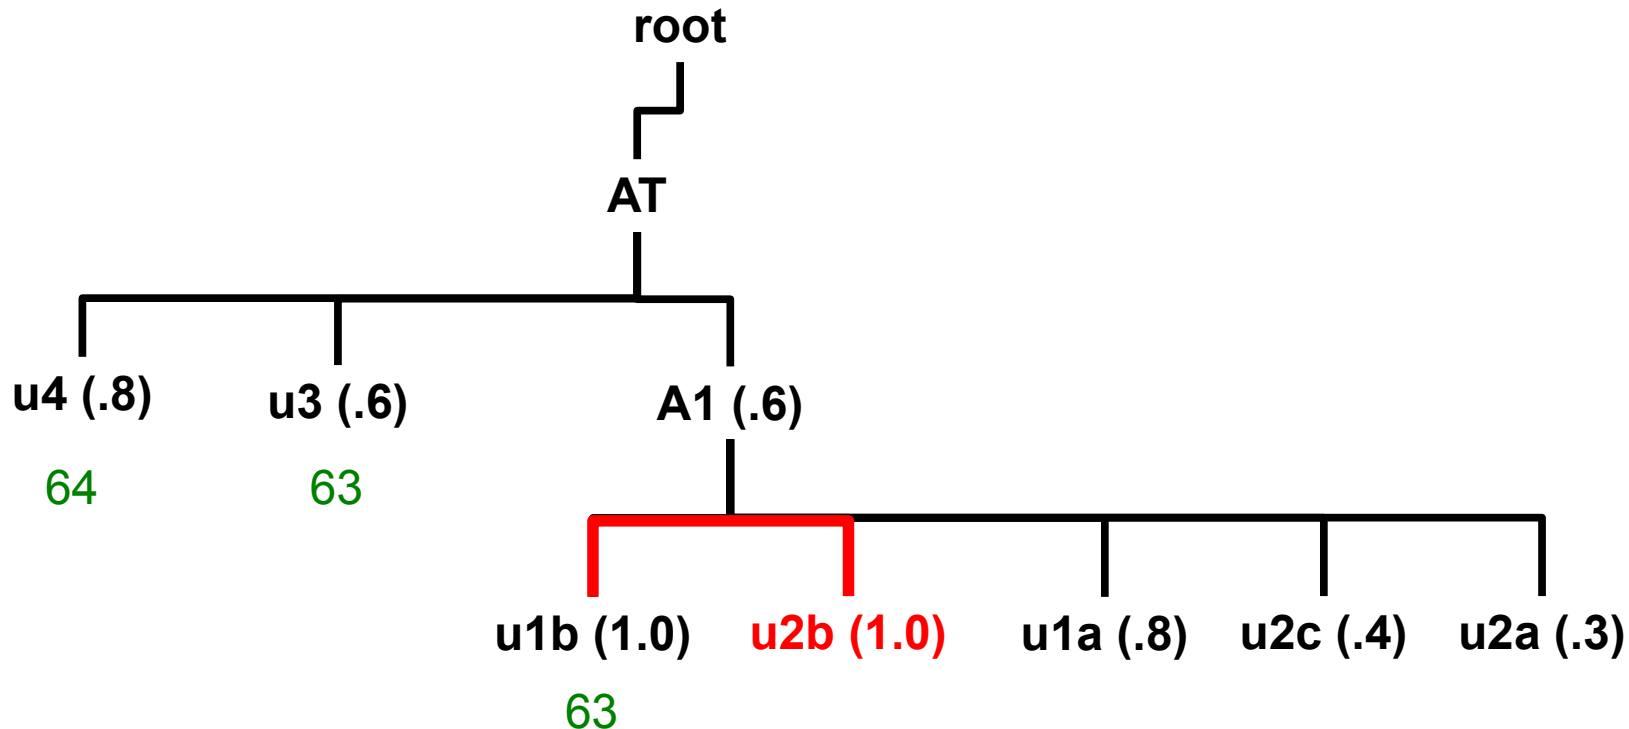
# Fair Tree: Tie Handling



Note that rank has **not** been decremented.  
u1b receives the same rank as u3 since A1 tied u3.

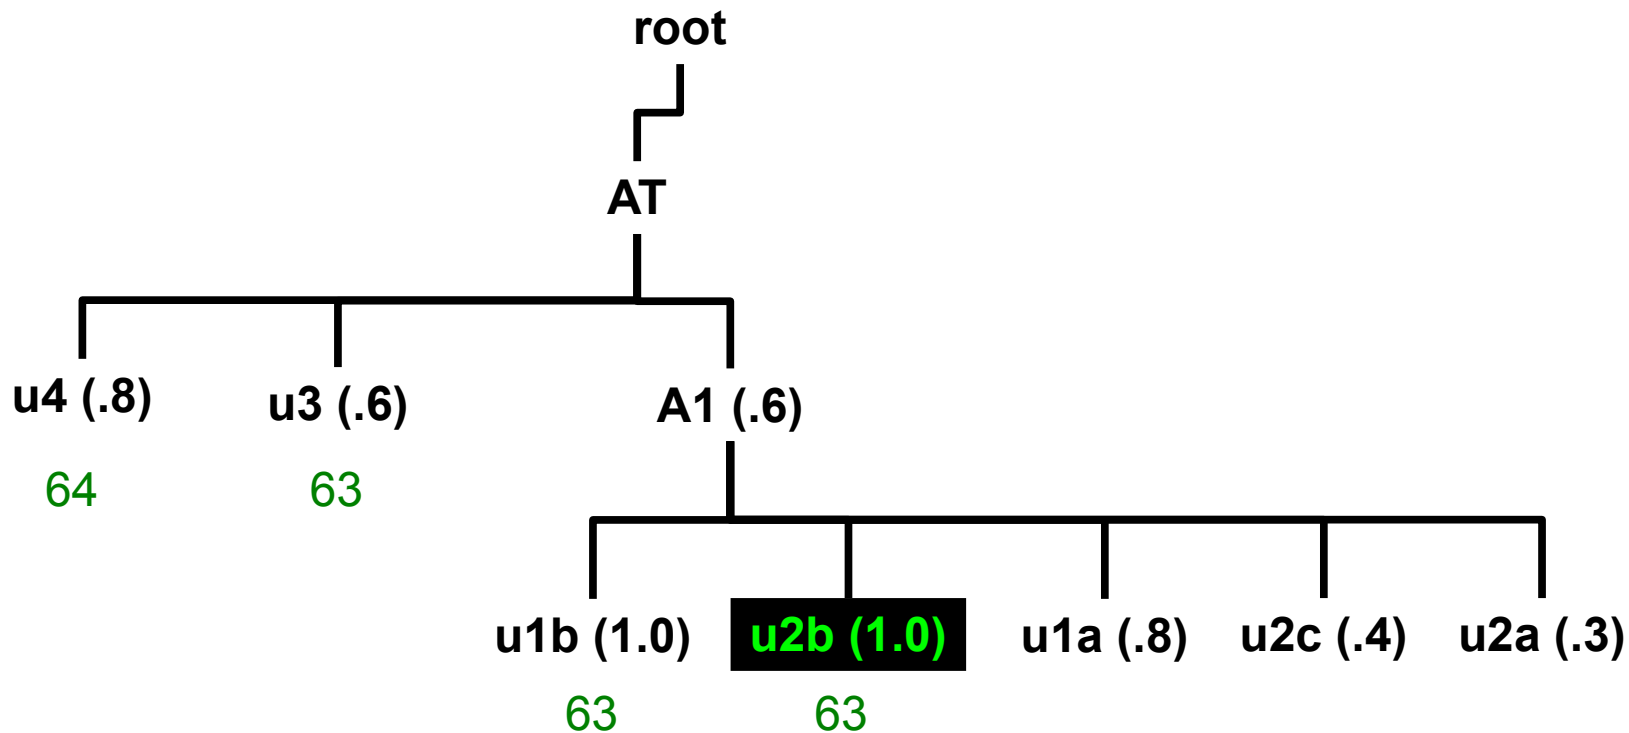
rank = 63

# Fair Tree: Tie Handling




rank = 63

# Fair Tree: Tie Handling



rank = 63

# Fair Tree: Tie Handling




Next association **not** tied. Decrement rank by 3 (3 ties)



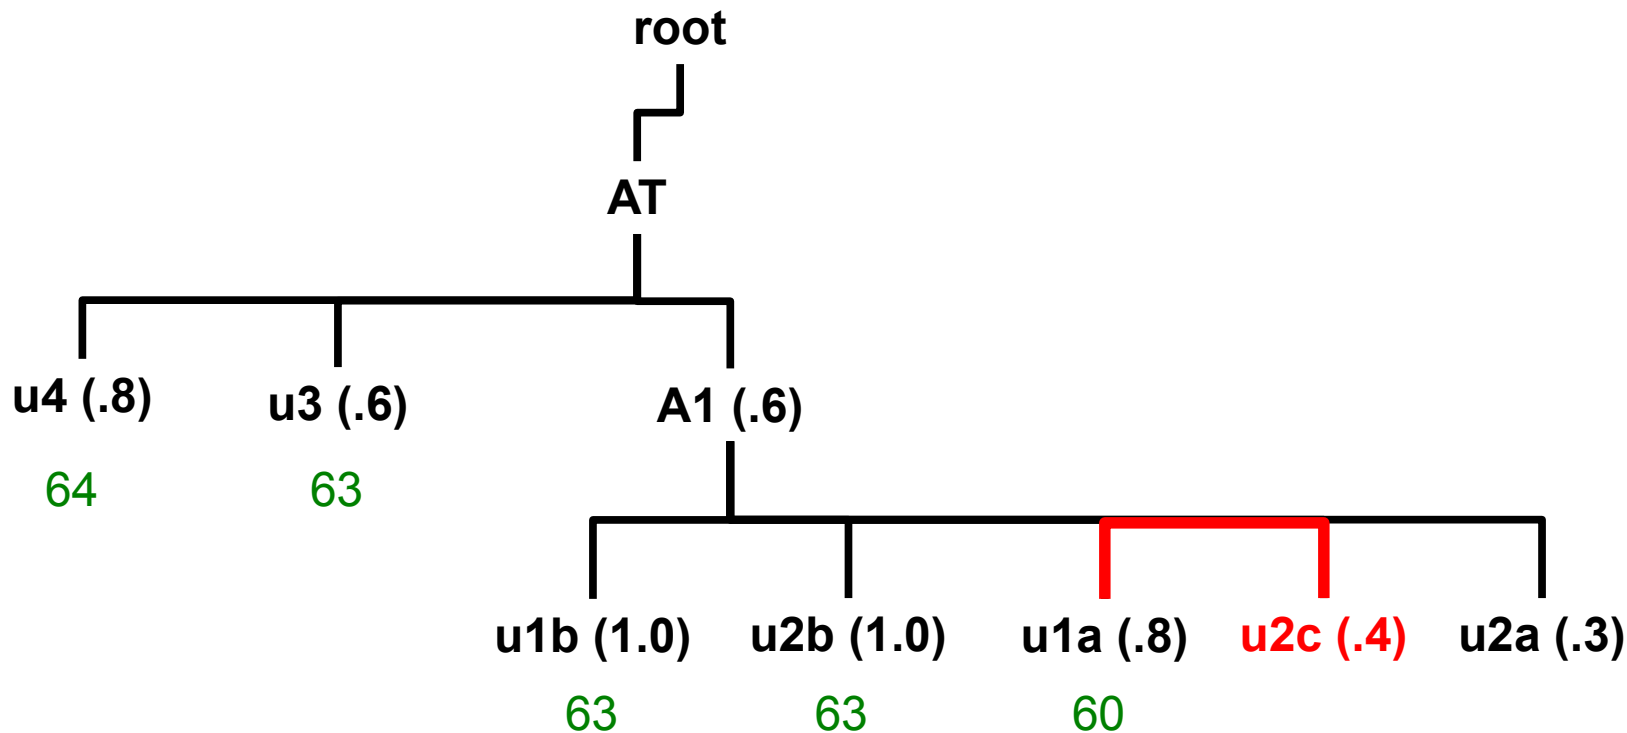
rank = 60

# Fair Tree: Tie Handling



rank = 60

# Fair Tree: Tie Handling



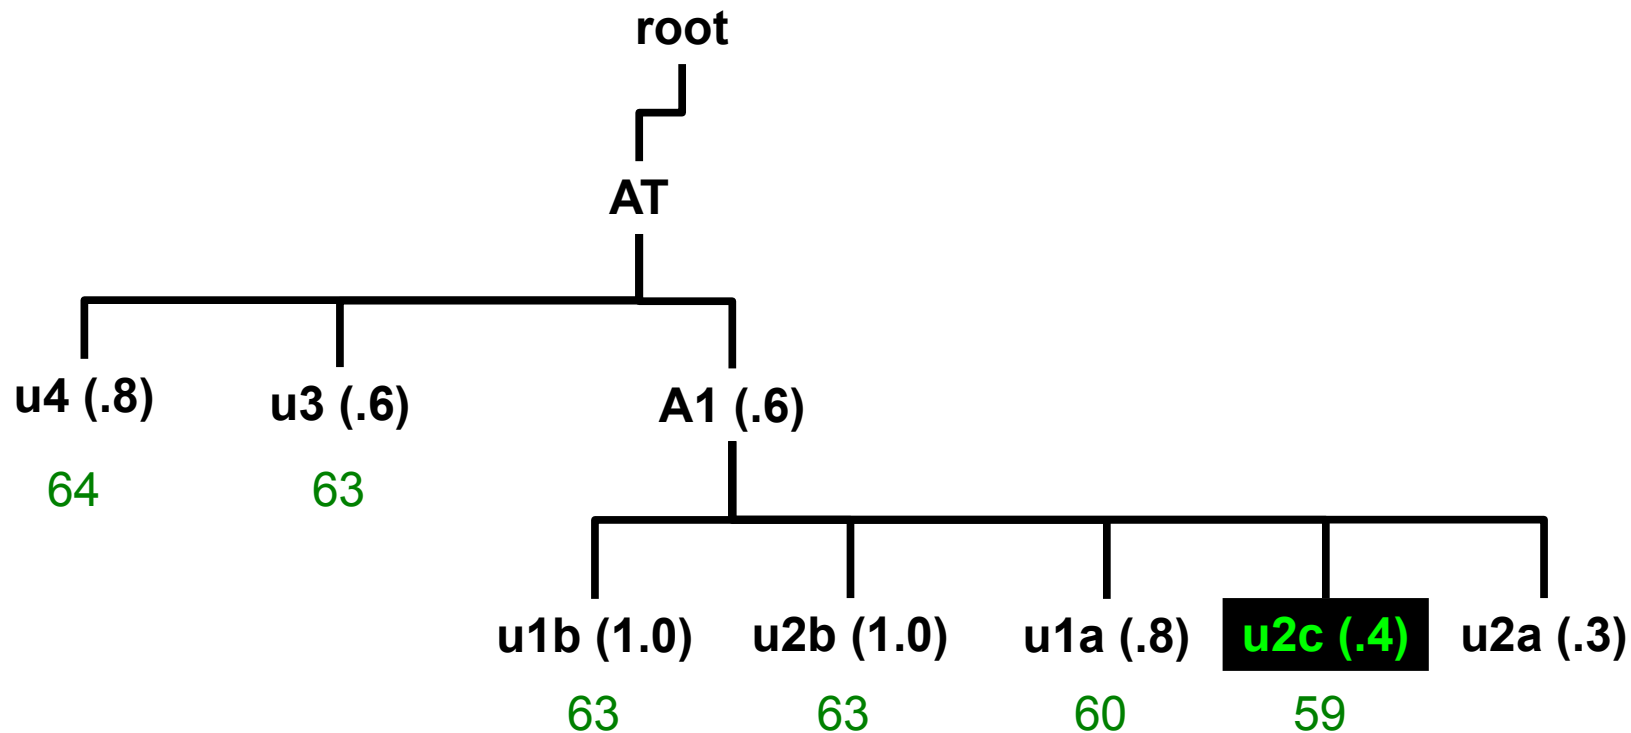
Next association **not** tied. Decrement rank



rank = 59

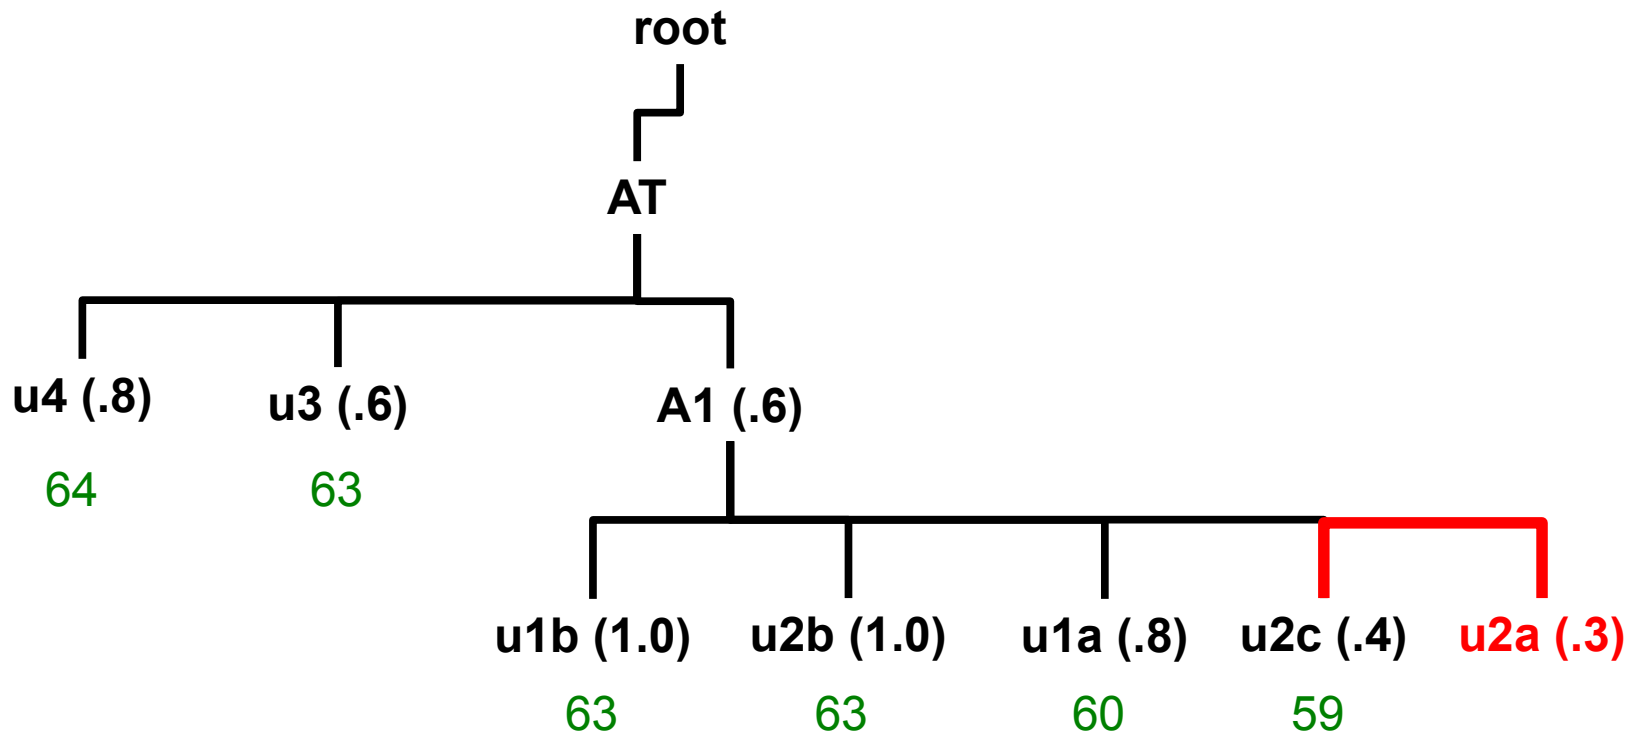


# Fair Tree: Tie Handling



rank = 59

# Fair Tree: Tie Handling

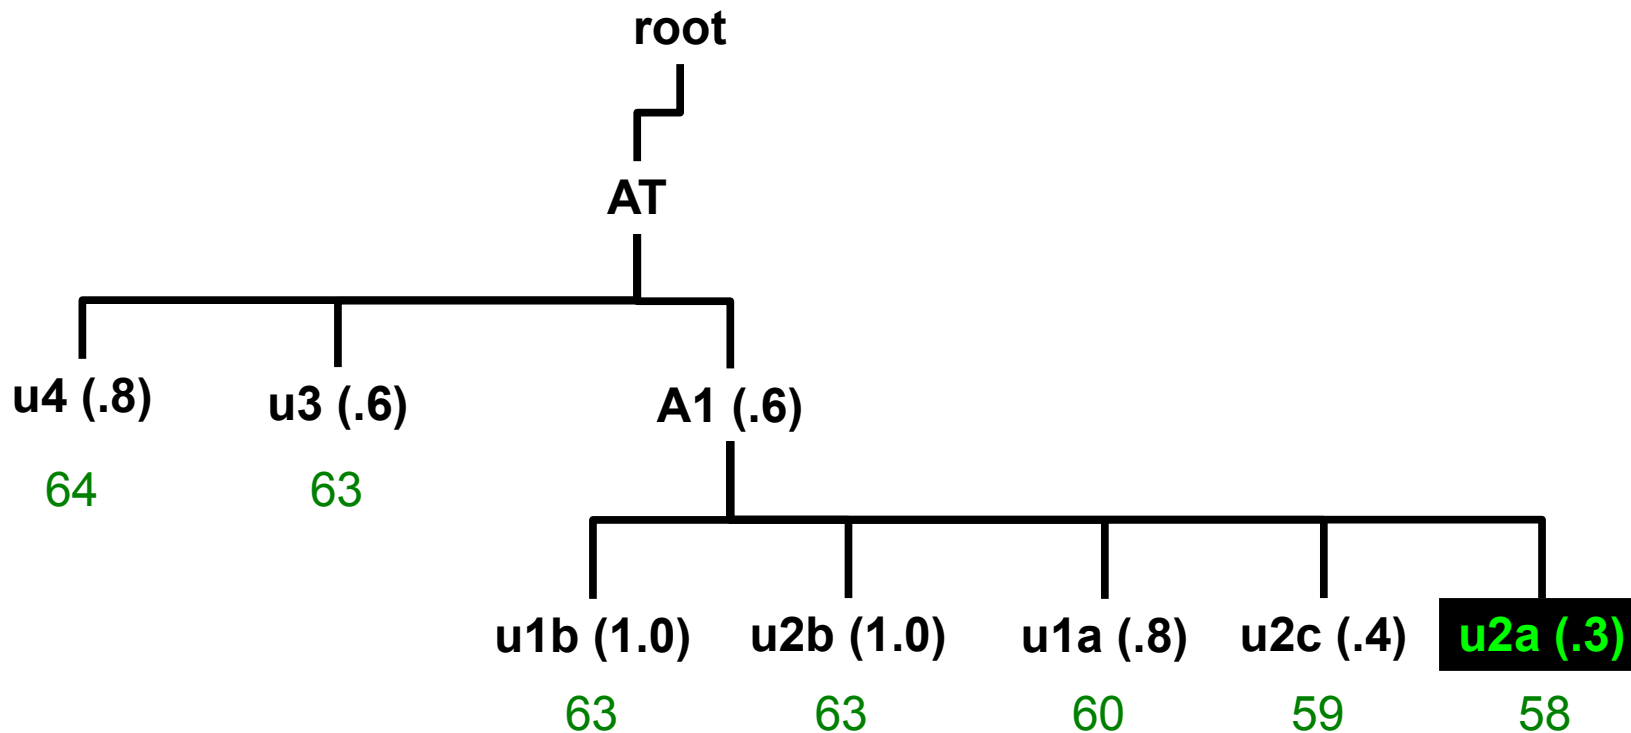


Next association **not** tied. Decrement rank



rank = 58


# Fair Tree: Tie Handling



Sorting uses the merged list


rank = 58

# Fair Tree: Tie Handling



u3, u1b, and u2b receive equal rankings

# Fair Tree: Tie Handling



- u3 tied A1, thus it tied A1's highest rank user, u1b
- u1b tied its adopted sibling user, u2b, since A1 and A2 were tied
- u3, u1b, and u2b receive the same fairshare factor