

Analysis of Software Artifacts

Performance Analysis

Jonathan Aldrich

Asymptotic Performance Analysis

- How do we compare algorithm performance?
 - Abstract away low-level details
 - E.g. exact assembly code
 - Investigate scaling behavior
 - Which is better on really large inputs?
 - For small inputs, any algorithm will do
- Asymptotic performance
 - As input size grows, how does execution time grow?

Counting Operations

- Input size N
- Number of elements of array to be searched or sorted
- Dimensions of matrix to be multiplied
- Number of bits in a binary number

- Idea: count operations performed
- Ignore constant factors
 - An “operation” can be anything that is implemented in a fixed number of assembly operations
- Let $T(N)$ = number of operations for input size N
 - Could be average or maximum
 - May be complex to calculate
 - Say $T(N) = 2^N + 2N^2 + 10$
 - Do we really care about the constant 10?
 - If N is big all that matters is 2^N

Big-O Notation

Notation taken from
notes by Dan Ellard

- Order of algorithm $T(N) = O(f(N))$
- There exist positive constants c and n_0
- For all $N \geq n_0$ we have $T(N) \leq c f(N)$

- **Example**
- $2^{N+1} + 2N^2 + 10 = O(2^N)$
- Choose $c = 4$ and $n_0 = 5$
- $T(n_0) = 2^{5+1} + 2(5)^2 + 10 = 64 + 50 + 10 = 124$
- $c f(n_0) = 4 * 2^5 = 4 * 32 = 128$
- We can prove inductively that this is true for all $N \geq 5$

Big-O Rules

- If $T_1(N) = O(f_1(N))$ and $T_2(N) = O(f_2(N))$
 - $T_1(N) + T_2(N) = O(\max(f_1(N), f_2(N)))$
 - E.g. $N^2 + N = O(\max(N^2, N)) = O(N^2)$
 - $T_1(N) * T_2(N) = O(f_1(N) * f_2(N))$
 - E.g. $N * (N^2 + N) = O(N * N^2) = O(N^3)$
- Domination of functions
 - Exponentials dominate polynomials
 - Polynomials dominate logarithms
 - Logarithms dominate constants
- If a function has several terms, only consider the order of the dominant term
 - By this shortcut we know $2^{N+1} + 2N^2 + 10 = O(2^M)$
 - Caveat: $N * 2^{N+1} + 2N^2 + 10 = O(N * 2^M)$, **not** $O(2^M)$
 - Can ignore terms that are added but not terms that are multiplied

Check Your Understanding (1)

- $2N^2 + 3N + 5 = O(\quad)$
- $N + \log N = O(\quad)$
- $2 \log N + 100 = O(\quad)$
- $2^{N/2} + 100000N^{25} = O(\quad)$

How to Analyze an Algorithm

- **Straightline code**
 - Always $O(1)$!
- **Code with loops**
 - How many times might each loop execute?
 - Multiply bounds for nested loops
 - Add bounds for sequential loops
- **Recursive functions**
 - How many times is the function called?
 - How many operations per function, counting each recursive call as a single operation?
 - (the recursive call is accounted for separately)
 - Multiply the two numbers above
- **Sometimes better to count operations**
 - Bounds based on loops or recursion could be too loose

Check Your Understanding (2)

```
if (x == y)
    x := x + 200
else
    y := y * 8
    for (int i = 0; i < N; ++i)
        for (int j = 0; j < i; ++j)
            x := x + j
    for (int k = 0; k < 2*N; ++k)
        x := x - k
```

- $O()$
- $O()$

Check Your Understanding (2)

```
int foo(int N) {
    if (N == 0)
        return 1;
    else
        return foo(N-1) + foo(N-1);
}

• O( )

// constant-time hashtable operations
int lookup(int N) { ... }
void store(int N, int V) { ... }

int foo(int N) {
    if (N == 0)
        return 1;
    val = lookup(N);
    if (val == -1) {
        val = foo(N-1) + foo(N-1);
        store(N, val);
    }
    return val;
}

• O( )
```

Optimization

- Making code run fast is fun!
- But beware...

Premature optimization is the root of all evil. –Donald Knuth (CC2, p. 594)

Cost of Optimization

- Optimizing your program may reduce other quality attributes
 - Evolvability
 - Optimizations may increase coupling
 - Understandability
 - Optimizations obscure high-level design
 - On-Time Delivery
 - Optimizations may delay the project's delivery
 - Reliability
 - Optimizations may make your code less likely to be correct (because it is more complex)
- Don't optimize at all unless you know you need to!

Do you really need speed?

- TRW system
 - Original requirement: sub-second response time
 - Cost estimate: \$100 million
 - Revised requirement: 4 second response time is OK
 - Cost estimate \$30 million
- What level of performance can your users live with?
- Can you achieve the necessary performance by buying more hardware?
 - Maybe not if you are buying 10 million widgets; otherwise probably yes!

Optimization

- Your boss tells you, “make it run faster.”
What do you do?

Speedup Analysis Scenario

- Assumptions
 - Procedure x takes 75% of execution time
 - We can speed up x by a factor of 3
- How much faster will the program run?

Amdahl's Law

- P = % of program you can speed up
- S = speedup of that part of the program
- Maximum overall speedup:

$$\frac{1}{(1-P) + P/S}$$

time to run rest of program time to run sped-up program part

80/20 rule (Pareto principle)

- 20% of the program takes 80% of the time
- No point in optimizing anything until you know what 20% is causing the problem!
- Another way of stating Amdahl's law

Getting Good Performance

- **Algorithms**
 - #1 way to optimize your code
 - If you pick an $O(n^2)$ algorithm instead of $O(n)$, nothing else (typically) matters
- **Cache performance**
 - Penalty for a cache miss may be hundreds of cycles
 - Know your data layout and cache architecture
 - Access multiple words in the same cache line together
 - Try to access your data in groups that fit into the size of the cache
- **Parallelism**
 - Your computer likely has multiple processors – use them if you can
- **Caching**
 - Trade space for time, e.g. to avoid recomputing results
- **Doing less work**
 - Trade precision for time, e.g. by using approximations

What the compiler should do for you

- Inlining calls
 - E.g. accessor functions like `getX()`
 - C++ allows you to recommend inlining
 - Often an optimizing compiler is smarter than you are
- Recursion vs. loops
 - Functional compilers optimize away “tail recursion”
 - Where the function ends with “return recursive_call(...)”
 - For most other systems it doesn’t matter—write what is most natural
 - If you have to encode the stack with a data structure of your own it probably isn’t worth it, the machine stack is faster than yours
 - For inner loops it may make a difference—verify with a profiler before optimizing
 - Do the caller and callee make up substantial time?
- Low-level arithmetic
 - Good compilers can often optimize well
 - Exception: floating-point optimizations may change the result due to round-off errors.
 - Compilers won’t optimize automatically
 - You should only optimize if you know what you are doing

Don't assume you know the bottlenecks

- They may depend on the compiler, hardware, input, etc.
- You don't know until you measure

Profiler

- Dynamic analysis tool
- How much time is spent in each function
 - Direct and transitive
- How much memory is used by each data structure

How Profilers Work

- **Timing**
 - **Sampling**
 - Periodically interrupt program, look at stack
 - Ascribe 1 time unit to function at top of stack, 1 time unit (transitively) to all other functions on stack
 - Requires running for a while to get averages
 - **Measuring**
 - Measure time on entry and exit of each function
 - High overhead, requires very accurate clock
- **Memory**
 - Periodically, or on demand, capture snapshot of heap
 - Write size and type of each object to file

Profiling Principles

- Identify the hot spots
- Places where the program spends most of its time
- Make incremental improvements
- Re-run profiler after each change
- Keep only the changes that make improvements

Session Summary

- **Big-O Performance Analysis**
 - Determines how an algorithm's run-time scales as the input grows large
- **Optimization**
 - Don't optimize prematurely, or without measuring
 - Focus on algorithms, caching, concurrency, hot spots
- **Profiling tools**
 - Measure where execution time and memory is spent

Further Reading

- McConnell, Steve. *Code Complete*, Second Edition, ch. 25.
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*, Second Edition. Section 3.1: Asymptotic notation, pp.41–50.

FOR NEXT YEAR

- backup slide with rate monotonic theorem
- understand WHY the limit is what it is
- rate monotonic synonym for deadling monotonic (and NOT EDF, right?)

Analysis of Software Artifacts

Analyzing Real-Time Systems

Reference: **Burns and Wellings, *Real-Time Systems and Programming Languages***

Jonathan Aldrich

Real-Time Systems

- Definition
 - Any system in which the time at which output is produced is significant. This is usually because the input corresponds to some movement in the physical world, and the output has to relate to that same movement. The lag from input time to output time must be sufficiently small for acceptable timeliness – *Oxford Dictionary of Computing*
- Examples make up 99% of microprocessors
 - Industrial process control
 - Manufacturing
 - Communication – cell phones, etc.
 - Device controllers – thermostat, etc.

Hard and Soft Real Time

- **Hard**
 - Deadline must be met
- **Soft**
 - A task that may tolerate occasional missed deadlines
 - May be a limit on how often deadlines are missed, or how late a process may be
- **Firm**
 - May miss deadlines—but no benefit if result delivered after deadline
- **Real world**
 - Many systems have both hard and soft aspects
 - Many tasks are somewhere between hard and soft

Analysis of Real-Time Systems

- Many complex issues
 - Concurrency
 - All the problems we studied before
 - Priority inversion: can a low-priority task cause a high-priority one to miss its deadline?
 - Resource use
 - Especially memory resources
 - Safety criticality
- Most unique & critical: Timing
 - Will a task be completed by its deadline?
 - Two critical issues
 - How long does the task run in isolation?
 - When multiple tasks interact, will they each meet their deadline?

Timing and Memory Management

- Conventional malloc/free
- Heap fragmentation
 - can take unbounded space (in theory)

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----

Timing and Memory Management

- Conventional malloc/free
- Heap fragmentation
 - can take unbounded space (in theory)



Timing and Memory Management

- Conventional malloc/free
- Heap fragmentation
 - can take unbounded space (in theory)



Timing and Memory Management

- Conventional malloc/free
- Heap fragmentation
 - can take unbounded space (in theory)




Timing and Memory Management

- Conventional malloc/free
- Heap fragmentation
 - can take unbounded space (in theory)



Timing and Memory Management

- Conventional malloc/free
- Heap fragmentation
 - can take unbounded space (in theory)



Timing and Memory Management

- **Conventional malloc/free**
- **Heap fragmentation**
 - can take unbounded space (in theory)
 - example: 10x fragmentation overhead
- **malloc/free may not take bounded time**
 - may need to search a free list for a block of the right size
- **Use malloc/free designed for RT**
 - e.g. TLSF (www.ocera.org)

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----

Timing and Memory Management

- **Virtual memory**
 - Allow graceful performance degradation
 - pages swapped to disk when application uses more than available memory
 - **Page faults unpredictable**
 - substantial delay
 - typically can't be used in RT systems
- **Caches**
 - Speed up average case memory access
 - Make worst-case access hard to predict
 - Hard RT systems: don't use cache
 - pay huge cost in performance
 - Typical RT systems: measure worst-case and add a safety factor

Timing and Memory Management

- **Garbage collection (GC)**
 - Automatically free unreachable data
 - Eliminates many leak problems
 - Can still have storage “leaks”
 - If you reference memory you aren’t going to use again
 - e.g. adding data objects to a global table
 - System pauses to collect garbage periodically
 - Typically unacceptable for real-time systems
- **Real-time garbage collection**
 - State of the art: IBM Metronome GC
 - Guarantees at least X% utilization of every time interval of length Y using space Z
 - Typical parameters
 - Y = 1 ms, X = 45%
 - Z = 2 * application space usage
- **Conclusion: RT GC is feasible but must roughly double space, processor usage**

Memory Management Strategies

- **Static allocation**
 - pre-allocate all memory in static data blocks
 - pro: predictable memory usage
 - con: wastes a lot of space
- **Stack allocation**
 - allocate memory on stack as functions are called
 - pro: predictable memory usage
 - must analyze all call chains
 - must have no recursive functions (generally true of RT systems)
 - con: still wastes space when memory usage doesn't fit stack discipline
 - con: dangling pointers are possible
 - but static analysis can find these
- **Region-based allocation – used in Real Time Java**
 - allocate memory in regions
 - regions are freed when no longer needed
 - pro: efficient in time and memory usage
 - when memory is used in predictable chunks with a predictable lifetime
 - con: dangling pointers possible
 - but static analysis can find these
- **Real-Time GC**
 - pro: nicest programming model
 - con: constant % of wasted space and time (but in some cases the space constant may be smaller than other techniques)

Real-Time Process Model

- Fixed set of processes ($a-z$)
- Processes are periodic with period T
- Processes have deadlines D equal to period T
- Processes have fixed worst-case execution time C
- Processes are independent
- System/context switch overhead is ignored (assumed to be zero)
- Some of these assumptions will be relaxed later

Scheduling

- How to schedule the execution of processes so that each meets its deadline
- Multiple approaches
 - Cyclic Executive
 - Run processes in a fixed order
 - Fixed-Priority Scheduling
 - Order processes by priority
 - Always run highest priority process
 - Earliest-Deadline First Scheduling
 - Dynamically pick process closest to its deadline
- Complicating factors
 - Computing Worst-Case Execution Time
 - Sporadic processes
 - Interaction between processes and blocking

Computing Worst Case Execution Time

- Compile to machine code
- Divide code into basic blocks
 - straightline code with no jumps into or out of the middle
- Compute worst case time for each block
 - analytic: requires detailed model of processor and memory hierarchy
 - cache, pipelines, memory wait states, etc.
 - measurement: may not be worst case
 - may add engineering safety factor
- Collapse control flow graph
 - choice between blocks → choose max time
 - loops → use bound on maximum loop iterations
 - specialized knowledge may tighten bounds
 - taking if branch precludes taking a later else branch
 - branch may be taken only a limited number of times in loop
- Safer to measure basic blocks & combine with analysis than to measure worst case of entire program

Worst Case Execution Time Example

```
for i = 1 to 100 do
  if (i % 2)
    then A (cost 10)
    else B (cost 30)
  if (i % 10)
    then C (cost 100)
    else D (cost 10)
end

cost of test = 1
```

Cyclic Executive

[example from Burns & Wellings]

- Assume all deadlines are multiples of the *minor cycle time*
 - 25 ms in the example

<u>Process</u>	<u>T</u>	<u>C</u>
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

Cyclic Executive

[example from Burns & Wellings]

- Assume all deadlines are multiples of the *minor cycle time*
- 25 ms in the example
- Idea
- If a task must be completed within n cycles, do $1/n$ of the work each cycle

<u>Process</u>	<u>T</u>	<u>C</u>
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

Cyclic Executive

[example from Burns & Wellings]

- Assume all deadlines are multiples of the *minor cycle time*
- 25 ms in the example
- Idea
- If a task must be completed within n cycles, do $1/n$ of the work each cycle

<u>Process</u>	<u>T</u>	<u>C</u>
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

```
loop
wait()
a()      10.0
b()      8.0
half_of_c()  2.5
half_of_d()  2.0
quarter_of_e() 0.5
end loop // total 23 ms
```

Cyclic Executive

[example from Burns & Wellings]

- Assume all deadlines are multiples of the *minor cycle time*
- 25 ms in the example
- Idea
- If a task must be completed within n cycles, do $1/n$ of the work each cycle
 - benefit: if schedule is possible this one will work
 - cost: have to break up processes
 - expensive to add context-switches
 - bad for software engineering

<u>Process</u>	<u>T</u>	<u>C</u>
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

```
loop
wait()
a()      10.0
b()      8.0
half_of_c()  2.5
half_of_d()  2.0
quarter_of_e() 0.5
end loop // total 23 ms
```

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

<u>Process</u>	<u>T</u>	<u>C</u>
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

Process	T	C
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

loop

wait()

wait()

wait()

wait()

end

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

Process	T	C
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

loop

wait()
a() 10

wait()
a() 10

wait()
a() 10

wait()
a() 10

end

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

loop		Process	T	C
wait()	10	a	25	10
a)	8	b	25	8
b)		c	50	5
		d	50	4
		e	100	2
wait()	10			
a)	8			
b)				
wait()	10			
a)	8			
b)				

end

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

Process	T	C
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

loop

wait()	
a()	10
b()	8
c()	5
wait()	
a()	10
b()	8
wait()	
a()	10
b()	8
c()	5
wait()	
a()	10
b()	8

end

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

Process	T	C
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

loop

wait()	
a()	10
b()	8
c()	5
wait()	
a()	10
b()	8
d()	4
wait()	
a()	10
b()	8
c()	5
wait()	
a()	10
b()	8
d()	4

end

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*

Process	T	C
a	25	10
b	25	8
c	50	5
d	50	4
e	100	2

loop	
wait()	
a()	10
b()	8
c()	5
wait()	
a()	10
b()	8
d()	4
e()	2
wait()	
a()	10
b()	8
c()	5
wait()	
a()	10
b()	8
d()	4

end

Cyclic Executive, Revised

[example from Burns & Wellings]

- Improved idea
 - If several tasks must be split over several minor cycles, divide them among the minor cycles
 - Repeat this *major cycle*
- This is the bin-packing problem
 - NP complete, therefore solved heuristically
- If it works, use it
 - Simple, effective
- But sometimes inapplicable
 - Deadlines are not a multiple of any reasonable minor cycle time
- Very long processes
 - have to be broken up

	loop	Process	T	C
	wait()	a	25	10
	a)	b	25	8
	b)	c	50	5
	c)	d	50	4
	wait()	e	100	2
	a)			
	b)			
	d)			
	e)			
	wait()			
	a)			10
	b)			8
	c)			5
	wait()			
	a)			10
	b)			8
	d)			4

end

Fixed Priority Scheduling (FPS)

- Assign each process a fixed priority
 - Based on time requirements, not importance
- Always run the highest priority active process
 - Processes released on schedule, according to period T
 - Run highest priority process first
 - When it completes run next priority process
 - New high-priority processes may preempt running processes of lower priority
 - non-preemptive schemes are possible but are more complicated to analyze

Deadline Monotonic Priorities

[example from Burns & Wellings]

- Shortest deadline first

- If $D_i < D_j$ then $P_i > P_j$

<u>Process</u>	<u>T/D</u>	<u>P</u>
a	25	5
b	60	3
c	42	4
d	105	1
e	75	2

- Optimality theorem: *if any process set can be scheduled using preemptive priority-based scheduling, deadline monotonic priority ordering will schedule it. (Leung and Whitehead, 1982)*

Timeline Schedulability Test

- How do you know if a set of processes is schedulable given a set of priorities?
 - i.e. will all processes meet their deadlines?
 - Theorem about deadline monotonic priorities only says that approach will find a schedule *if one exists*...but it may not exist!

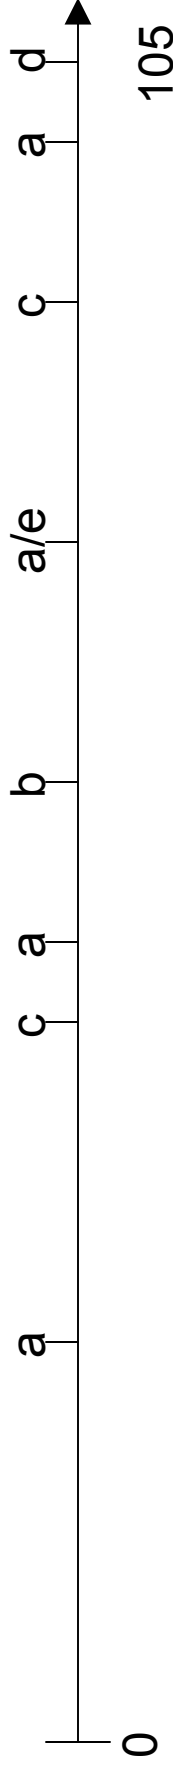
Timeline Schedulability Test

- How do you know if a set of processes is schedulable given a set of priorities?
 - i.e. will all processes meet their deadlines?
 - Theorem about deadline monotonic priorities only says that approach will find a schedule *if one exists*...but it may not exist!
- Timeline test (for deadline $D = \text{period } T$)
 - Assume all processes are released together
 - This is the maximum possible load on the processor
 - Known as the critical instant
 - Track process execution until all processes have passed their deadline
 - Theorem: *If a set of processes are tracked from the critical instant until the process with the longest period is complete, and all processes meet their deadlines, then all deadlines will be met thereafter (Liu and Layland, 1973)*
 - assumes deadline $D = T$

Timeline Schedulability Example

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

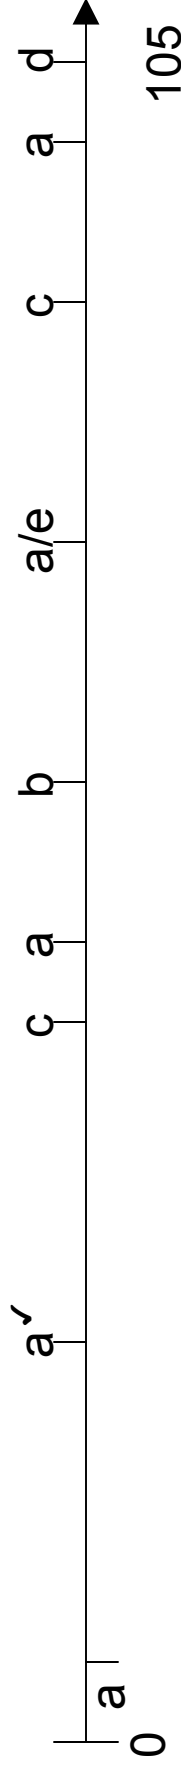


Process Execution

Timeline Schedulability Example

Process	<i>T/D</i>	<i>C</i>	<i>P</i>
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

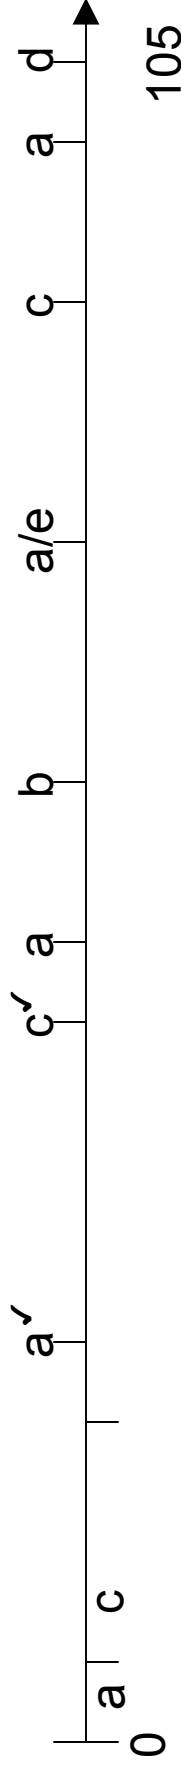


Process Execution

Timeline Schedulability Example

Process	<i>T/D</i>	<i>C</i>	<i>P</i>
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

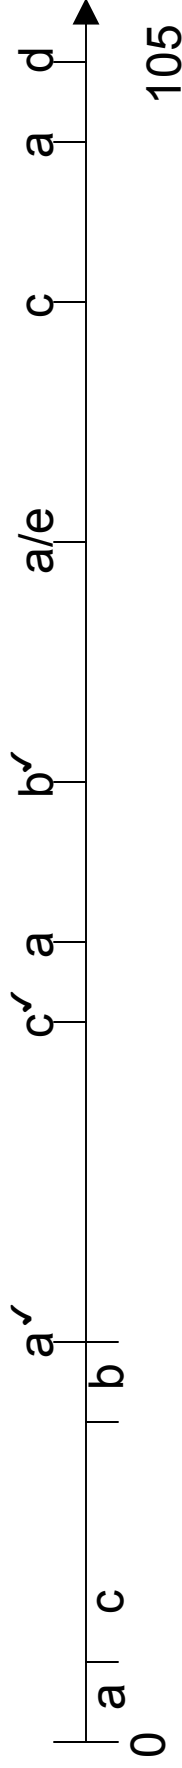


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

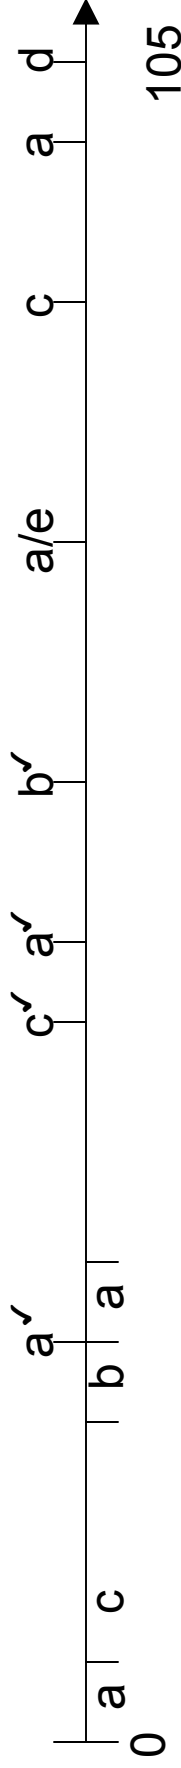


Process Execution

Timeline Schedulability Example

Process	<i>T/D</i>	<i>C</i>	<i>P</i>
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

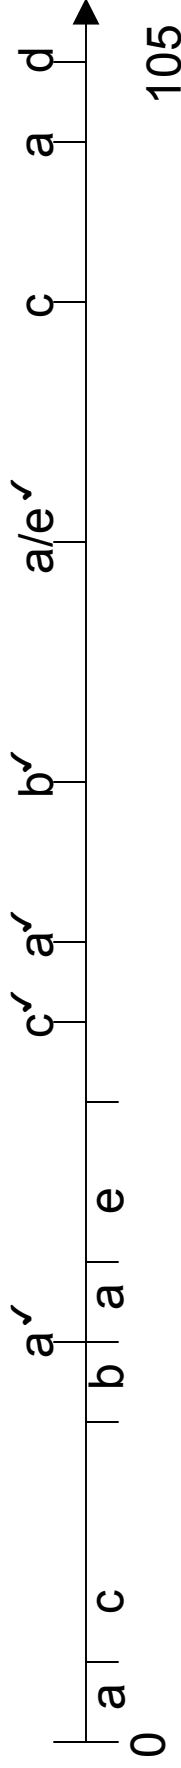


Process Execution

Timeline Schedulability Example

Process	<i>T/D</i>	<i>C</i>	<i>P</i>
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

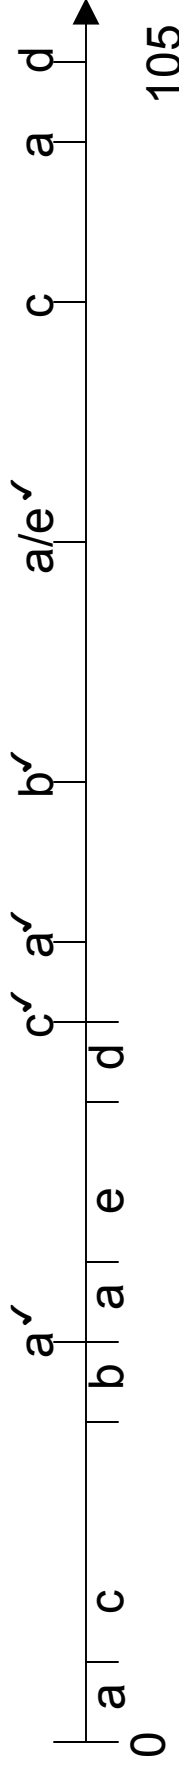


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

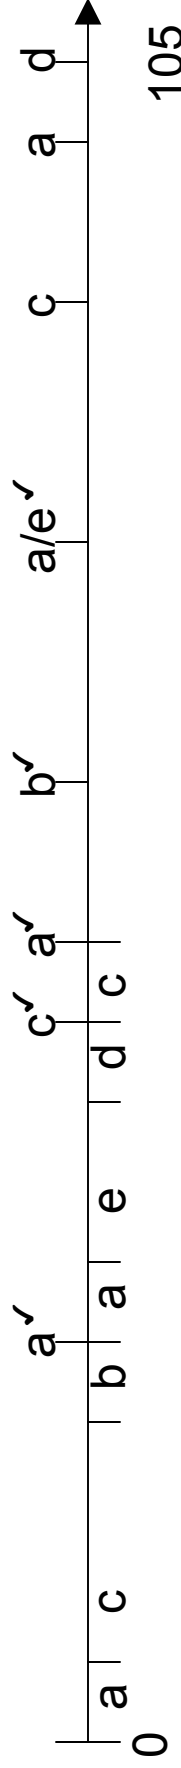


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

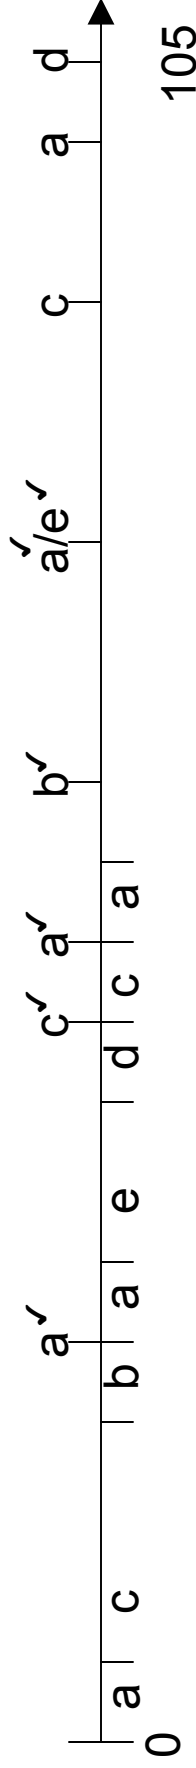


Process Execution

Timeline Schedulability Example

Process	<i>T/D</i>	<i>C</i>	<i>P</i>
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

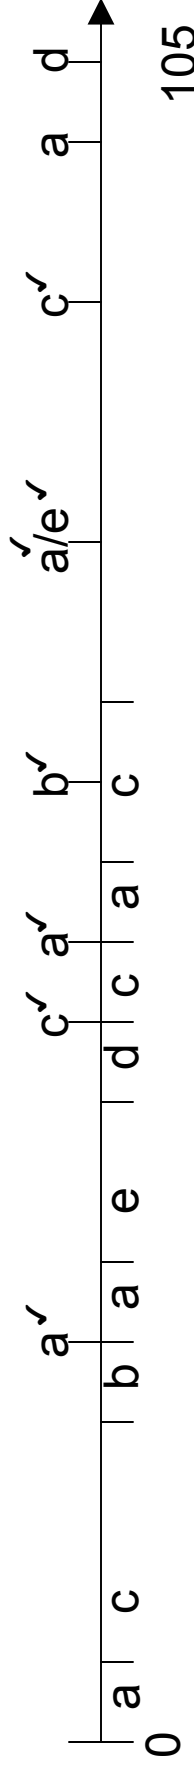


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

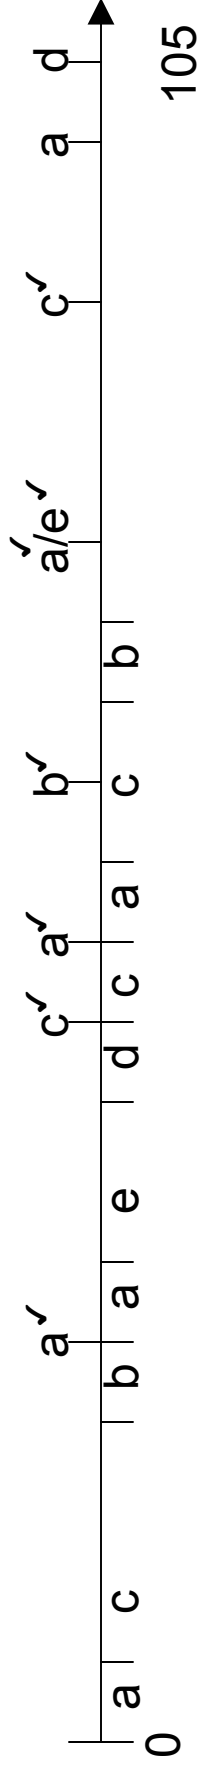


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

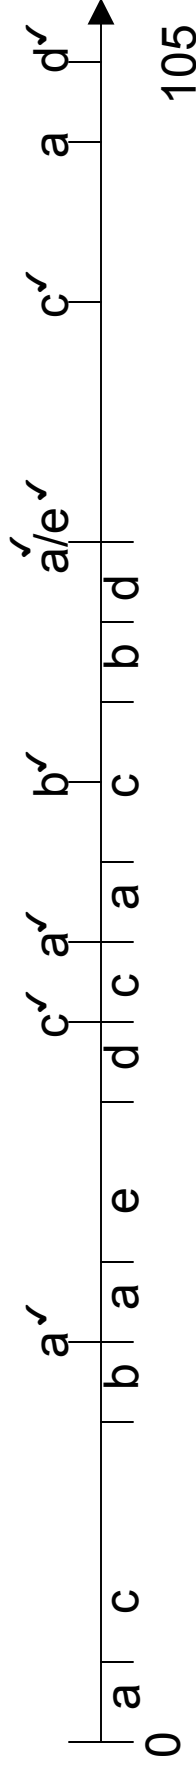


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

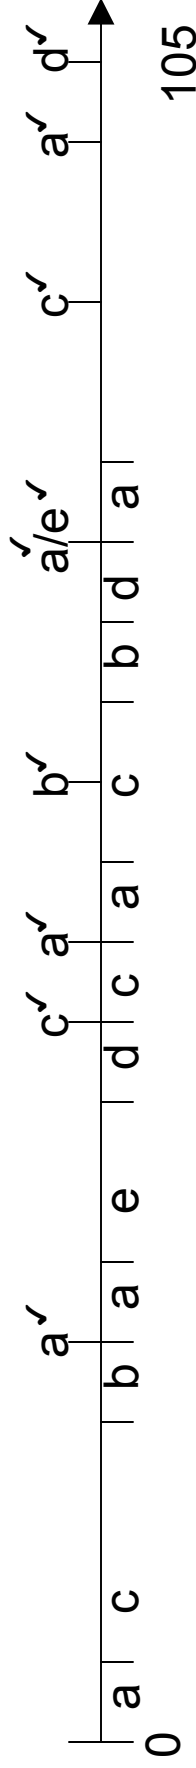


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

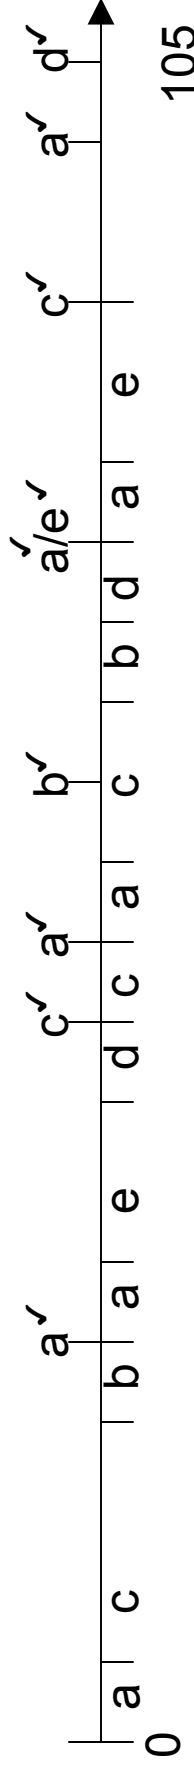


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines

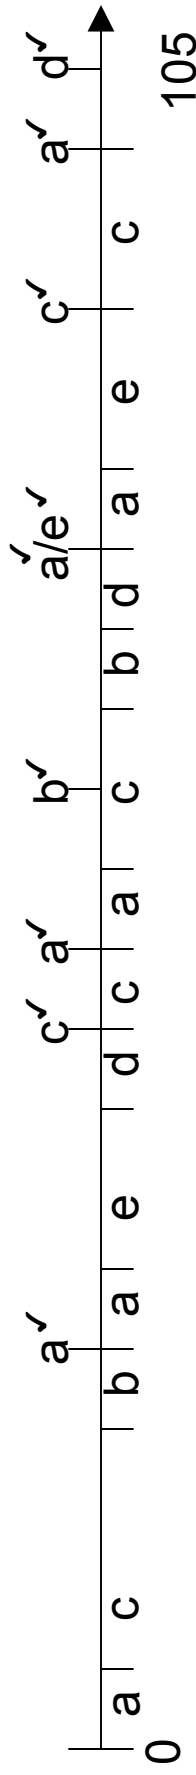


Process Execution

Timeline Schedulability Example

Process	T/D	C	P
a	25	5	5
b	60	5	3
c	45	15	4
d	105	10	1
e	75	10	2

Process Deadlines



Process Execution

Earliest Deadline First (EDF) Scheduling

- Always run the process with the earliest deadline
- Theorem: *As long as the sum of the utilizations of the processes is less than or equal to 100%, EDF can schedule a process set. (Liu and Layland, 1973)*
- Utilization of process $i = C_i/T_i$

Example: EDF vs. FPS

[example from Burns & Wellings]

Process	<i>T/D</i>	<i>C</i>	<i>P</i>	<i>Util</i>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

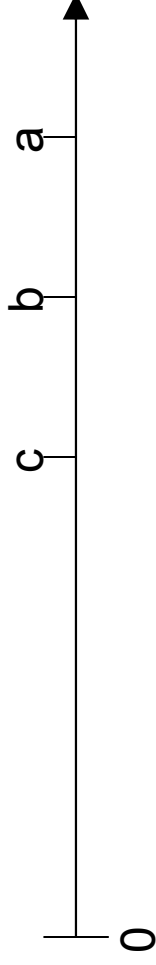
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

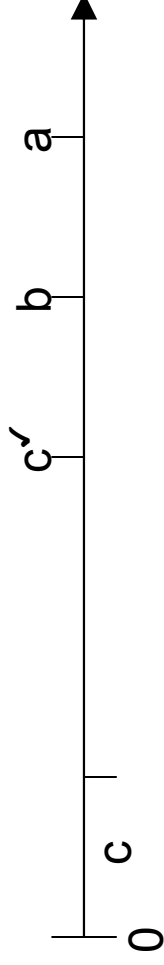
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



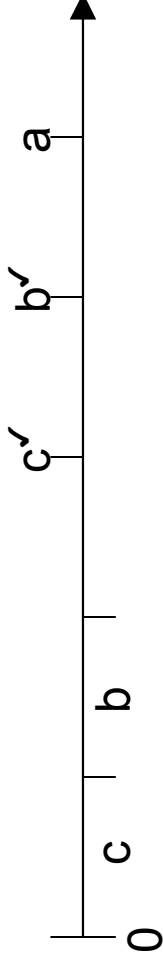
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

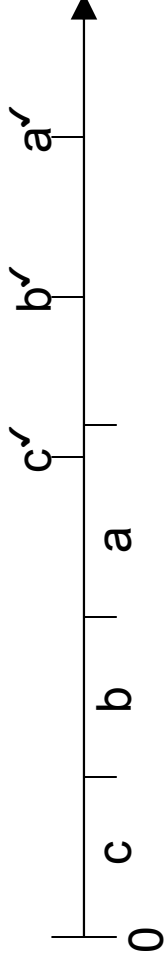
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

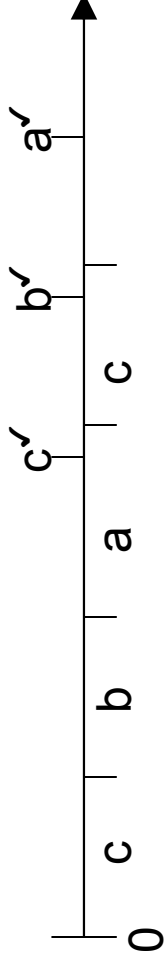
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

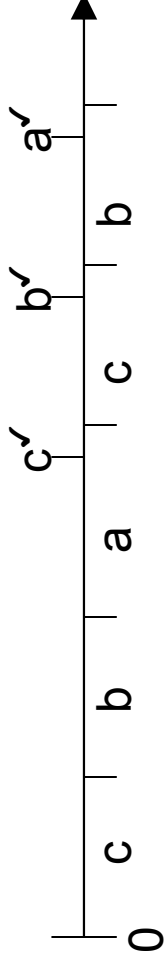
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

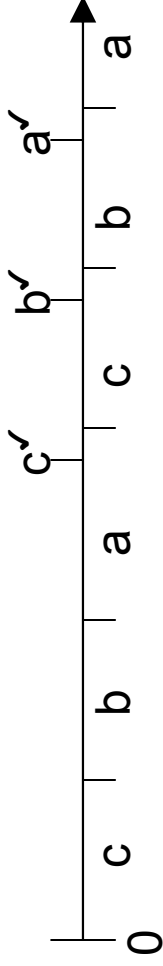
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

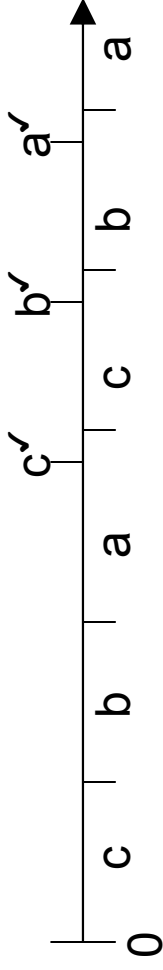
Example: EDF vs. FPS

[example from Burns & Wellings]

<u>Process</u>	<u>T/D</u>	<u>C</u>	<u>P</u>	<u>Util</u>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

EDF Process Deadlines



EDF Process Execution

All deadlines are met

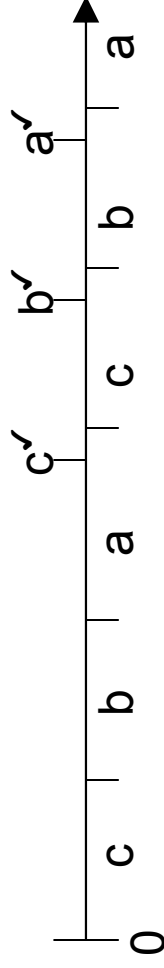
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

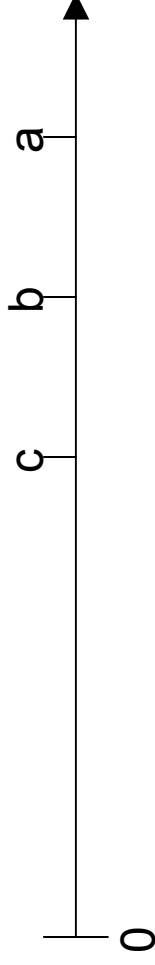
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

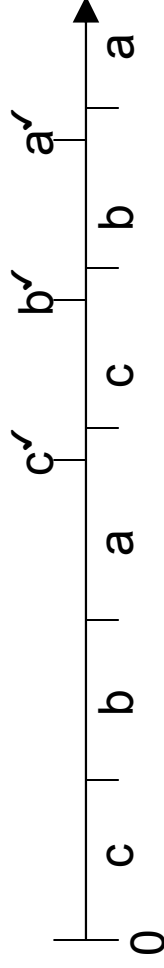
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

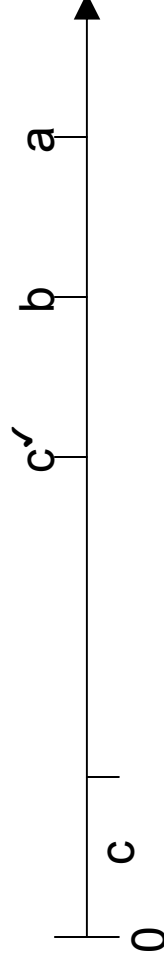
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

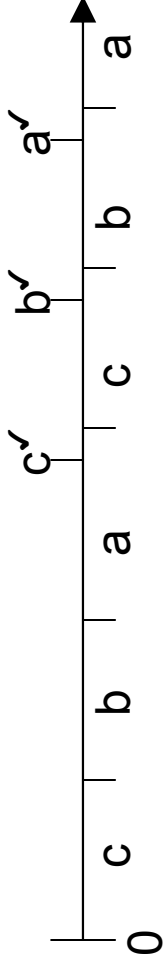
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

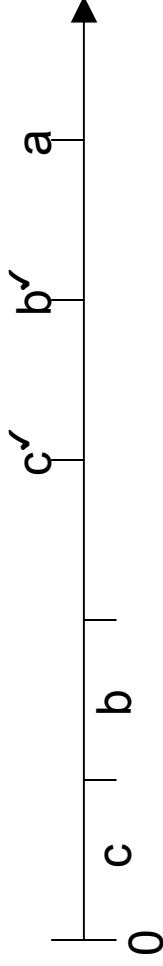
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

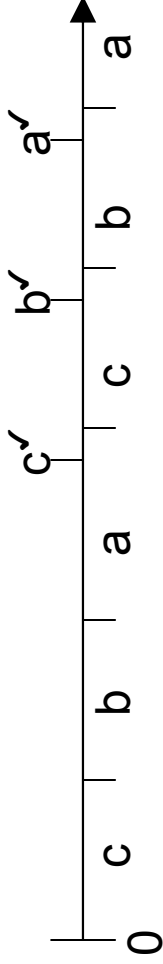
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	<i>T/D</i>	<i>C</i>	<i>P</i>	<i>Util</i>
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

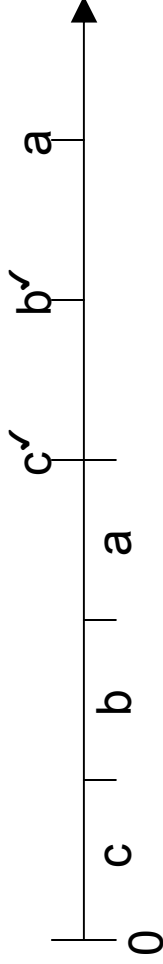
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

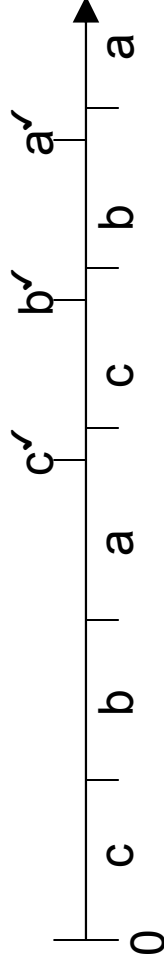
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

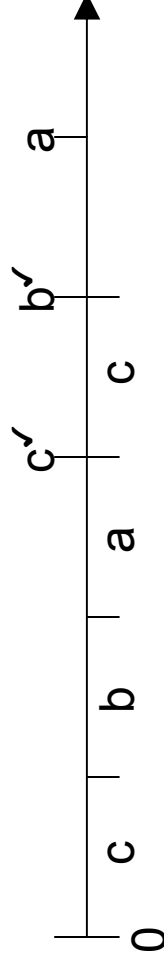
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

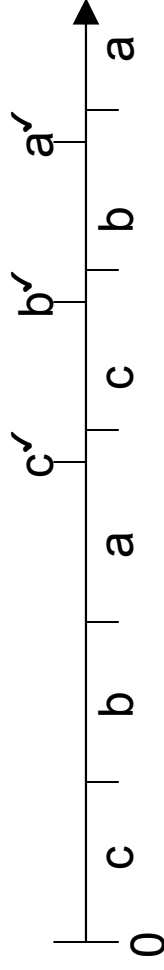
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

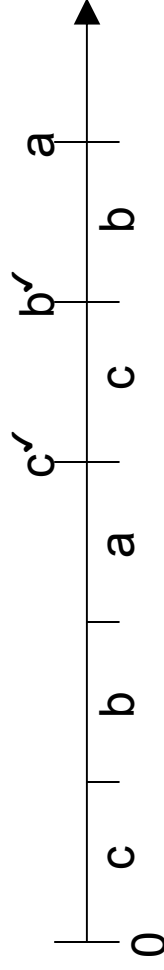
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

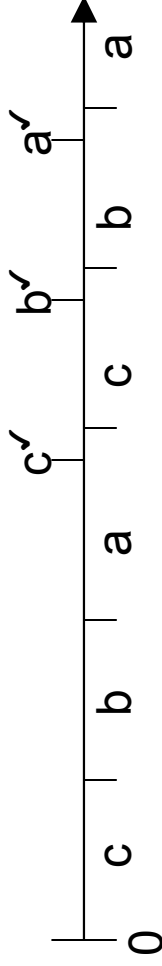
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

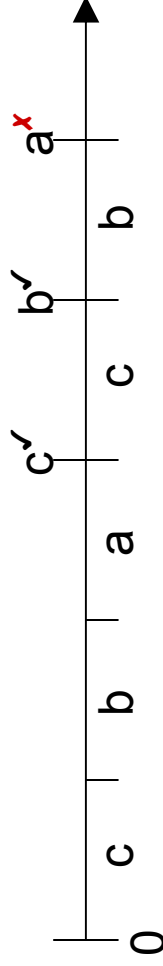
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



FPS Process Execution

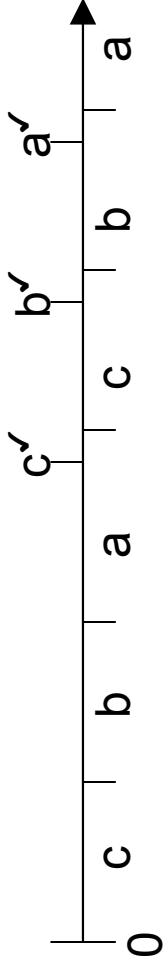
Example: EDF vs. FPS

[example from Burns & Wellings]

Process	T/D	C	P	Util
a	50	12	1	24%
b	40	10	2	25%
c	30	10	3	33%

Total utilization 82%

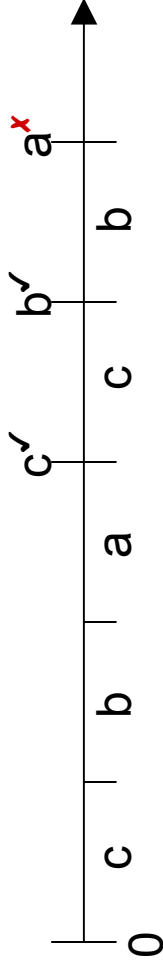
EDF Process Deadlines



All deadlines are met

EDF Process Execution

FPS Process Deadlines



Deadline for a is missed

FPS Process Execution

Discussion: EDF vs. FPS

[from Burns & Wellings]

- EDF can schedule process sets which FPS cannot
 - fixed priority is not flexible enough
 - EDF makes a dynamic choice based on deadlines
- Why not always use EDF?
 - FPS is easier to implement and more efficient
 - Schedule based on purely static priorities
 - Dynamic choices require complex run-time system and higher overhead
 - Incorporating processes without deadlines is easier in FPS
 - Just give them a priority; no need to set artificial deadline
 - Can use other factors (e.g. importance) for priority in FPS
 - EDF forces all scheduling decisions to be based on deadlines
 - FPS is more predictable in overload situations
 - Lower priority processes will miss deadlines first
 - In EDF cascade of processes may miss deadlines – domino effect

Sporadic Processes

- What about processes that do not arrive according to a fixed periodic schedule?
- Fixed priority scheduling
 - Set T to minimum inter-arrival interval
 - Sufficient to guarantee responsiveness
 - $D=T$ may be unrealistic
 - Allow $D < T$
 - Set priorities based on D

Soft Processes

- Sporadic processes
 - Worst case may be much worse than average-case
 - May be unrealistic (and unnecessary) to make all deadlines
- Distinguish hard and soft processes
- Scheduling rules [Burns & Wellings]
 - All processes should be schedulable using average execution times and average arrival rates
 - All hard real-time processes should be schedulable using worst-case execution times and worst-case arrival rates of all processes (including soft)

Servers

- Typically sporadic soft real time processes have low priority
- Means they may often miss their deadlines
 - Even when resources are available to meet their deadlines
- Deferrable Server (Lehoczky et al., 1987)
 - Schedule all hard processes
 - Add a server process
 - Highest priority
 - Some period T_s
 - Maximum capacity C_s such that system is still schedulable
 - When an sporadic process arrives
 - Run as part of server process until complete or server has used C_s capacity
 - After time T_s server can continue executing the sporadic process


Blocking and Priority Inversion

- Unrealistic assumption: processes are independent
- Typical case: processes access shared resources
- Resources protected by locks
- Priority inversion (Lauer and Satterwaite, 1979)
 - Low-priority process acquires a lock
 - High-priority process has to wait for low priority process to release the lock
 - Could wait a long time, especially if medium-priority processes preempt the low-priority process

Priority Inversion Example

[example from Burns & Wellings]


<u>Process</u>	<u>Priority</u>	<u>Execution</u>	<u>Release Time</u>
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

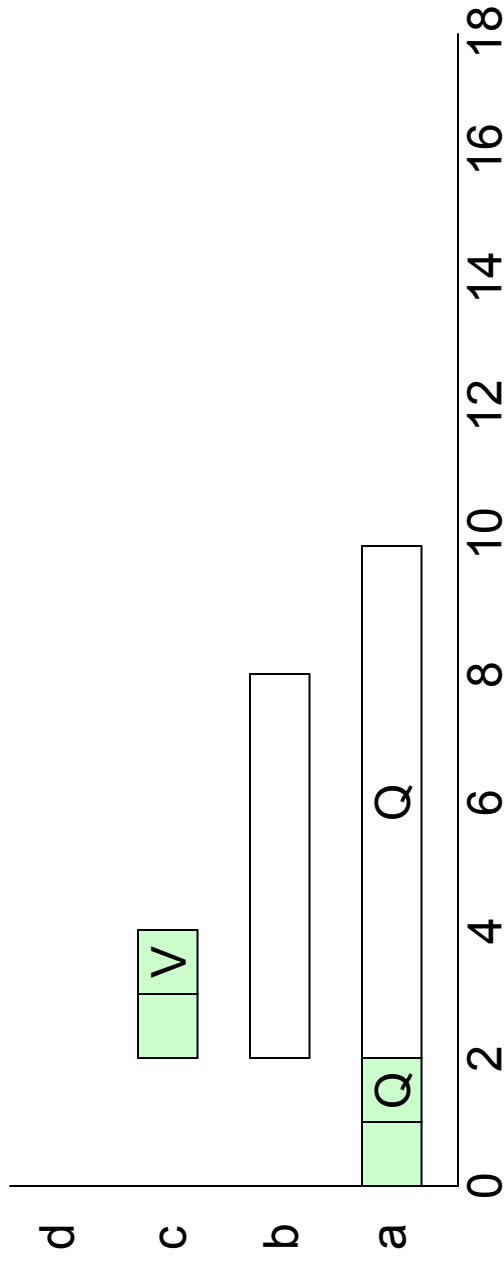
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

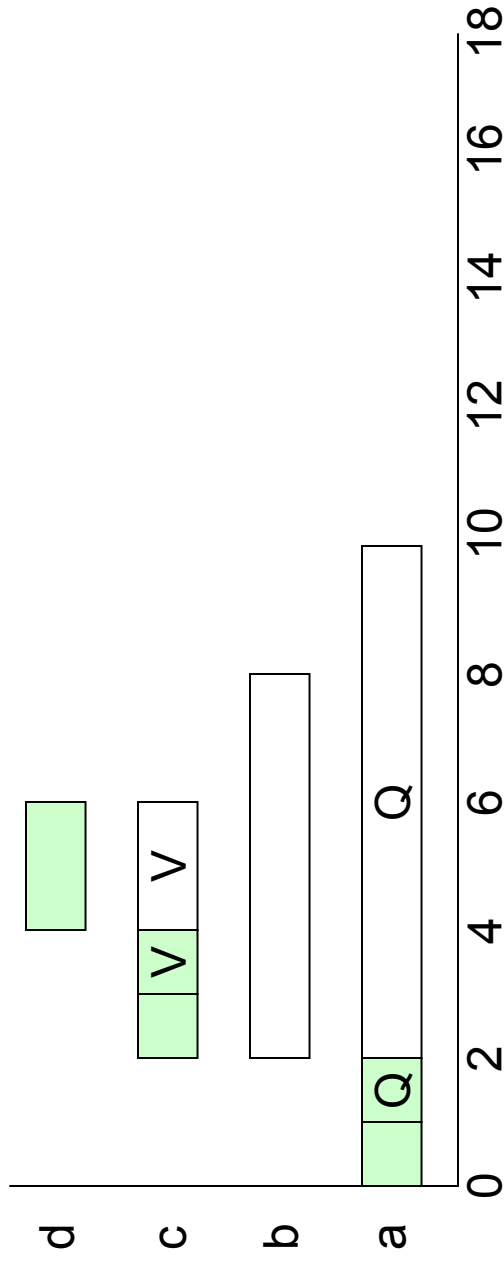
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

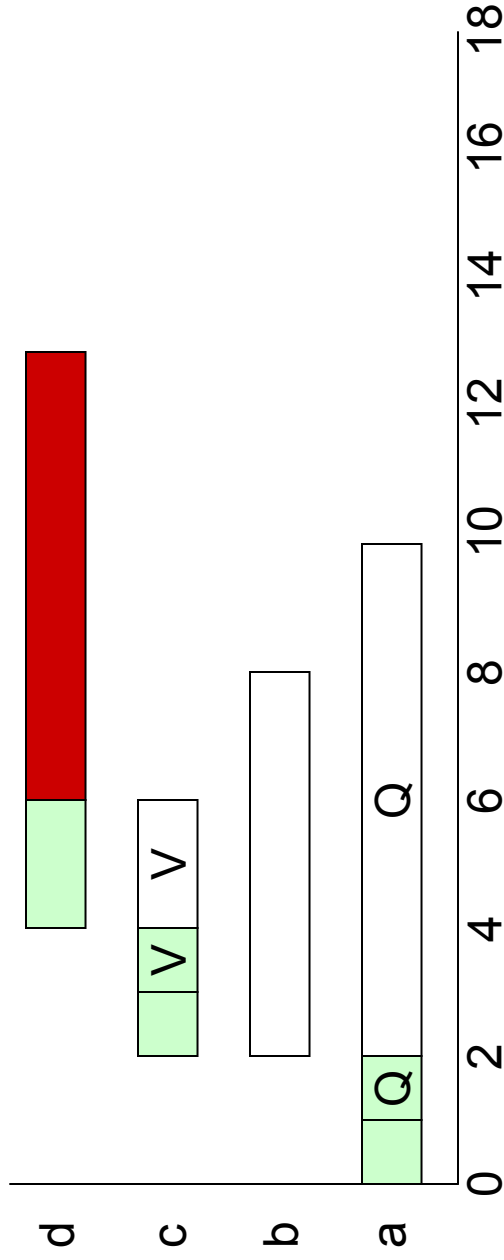
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

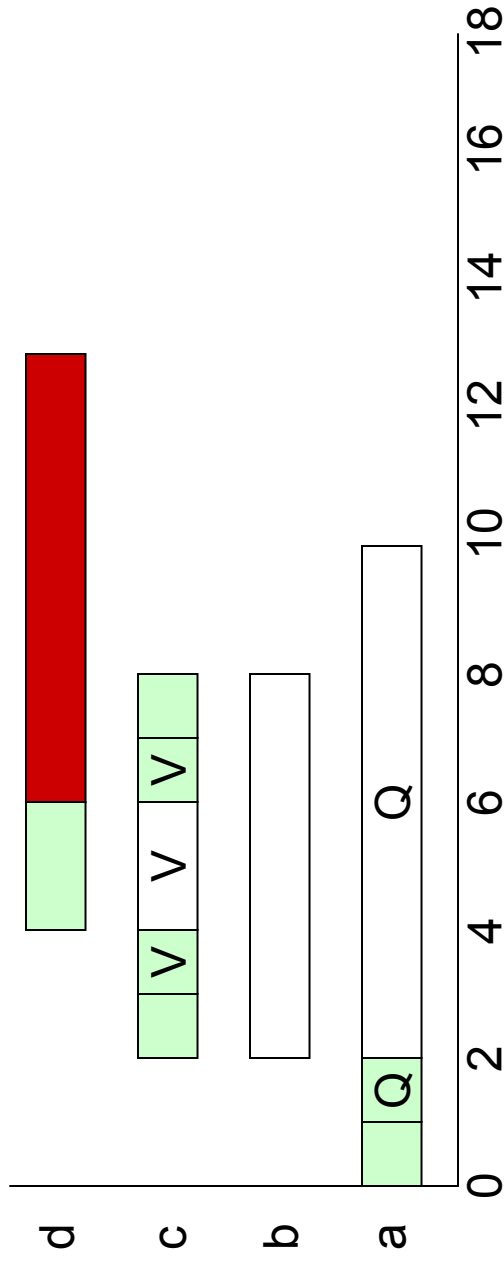
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

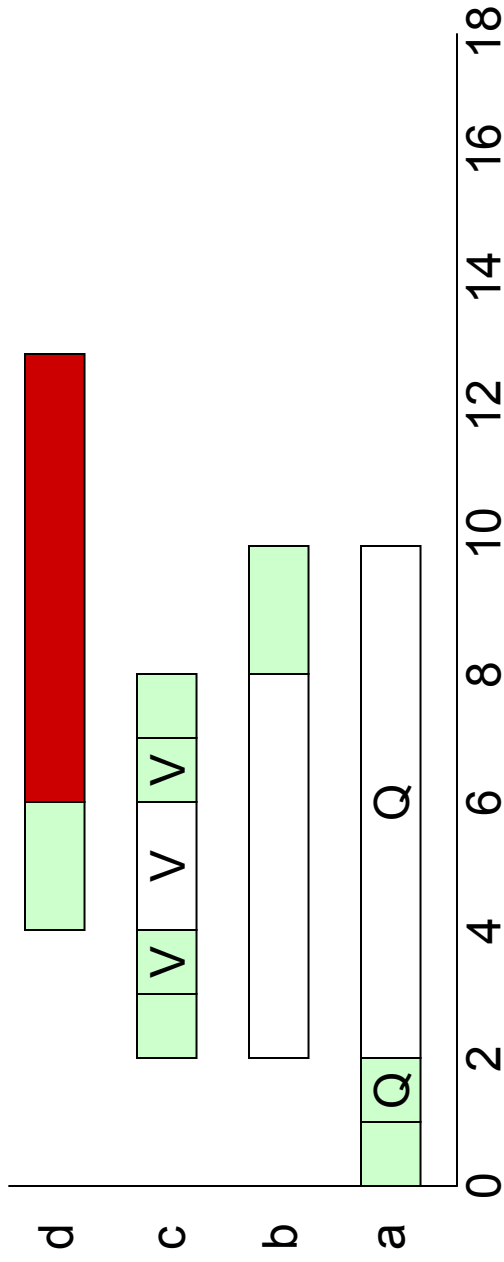
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

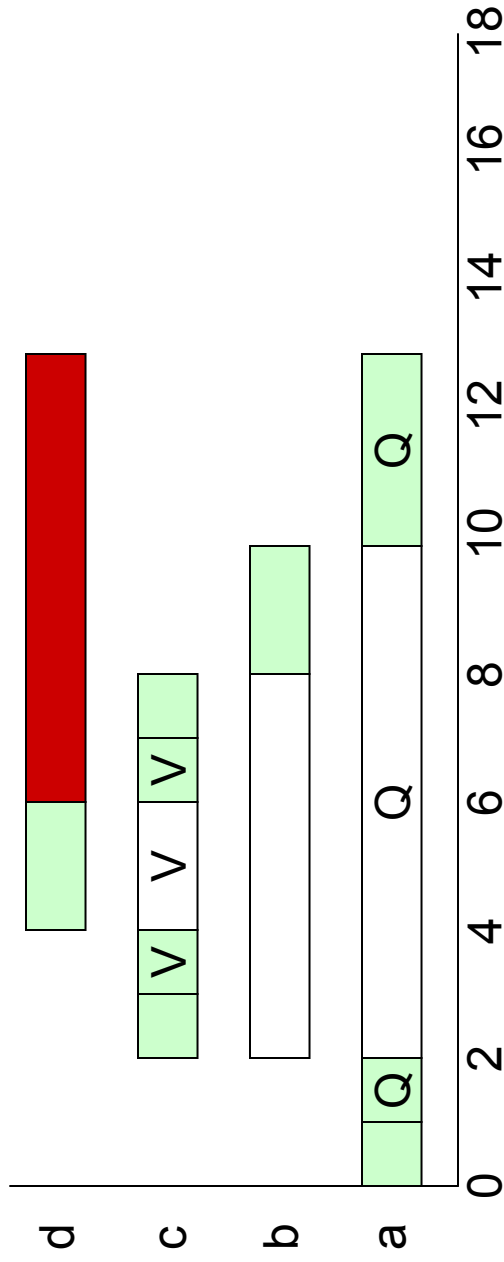
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

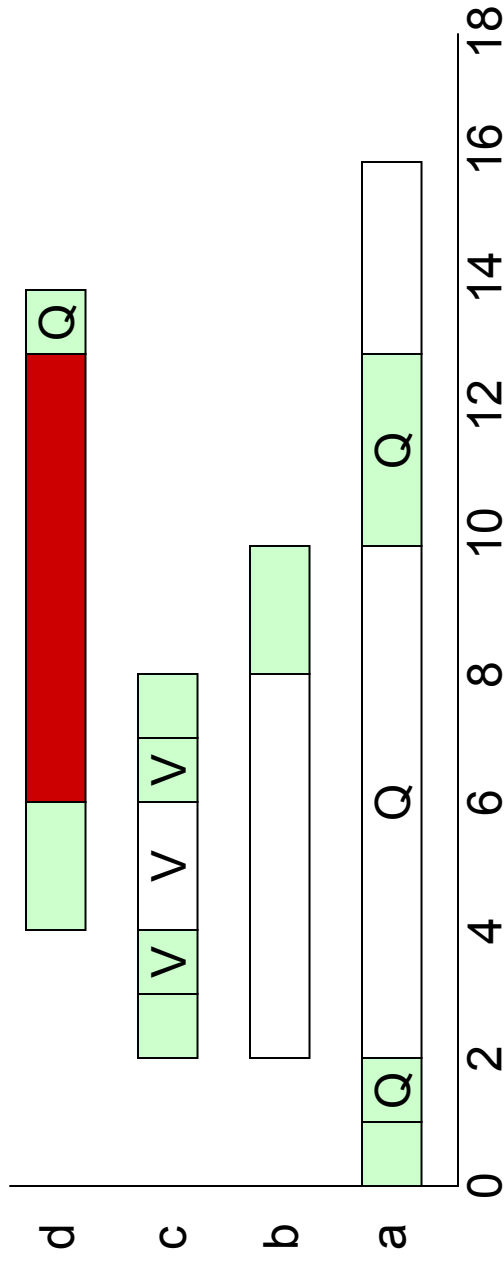
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

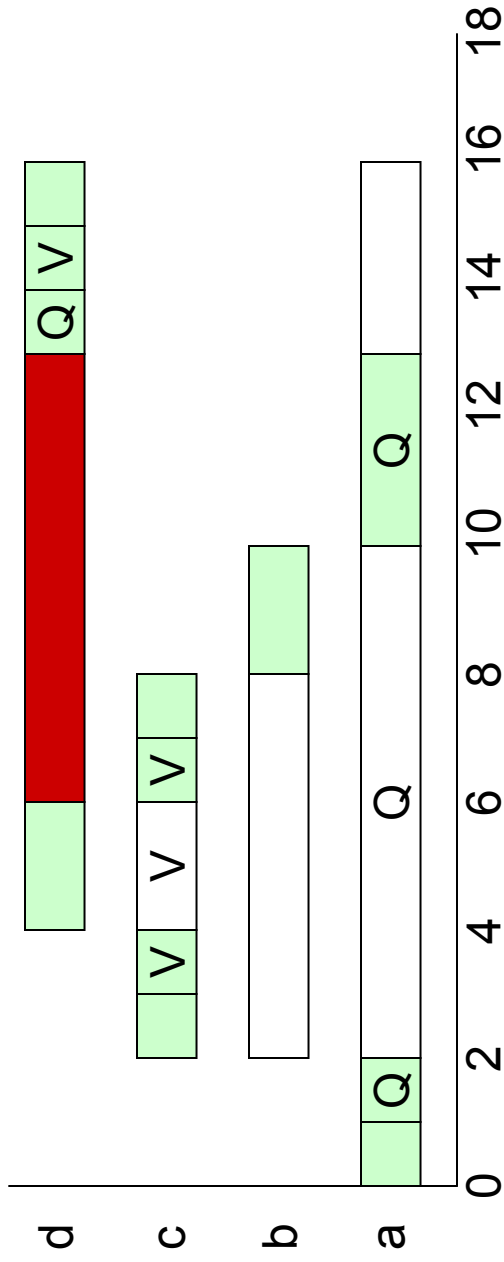
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

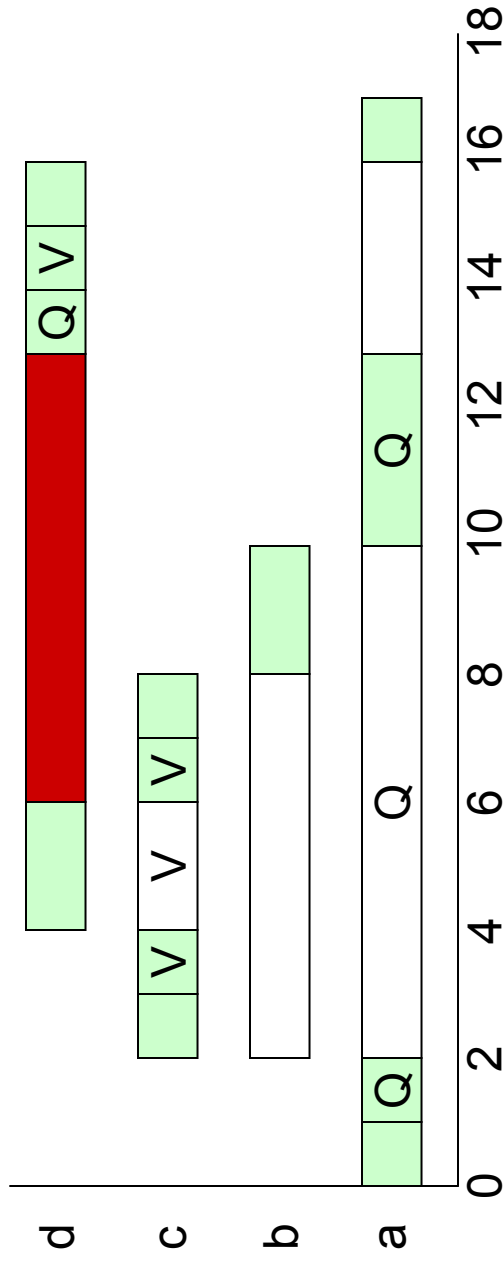
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

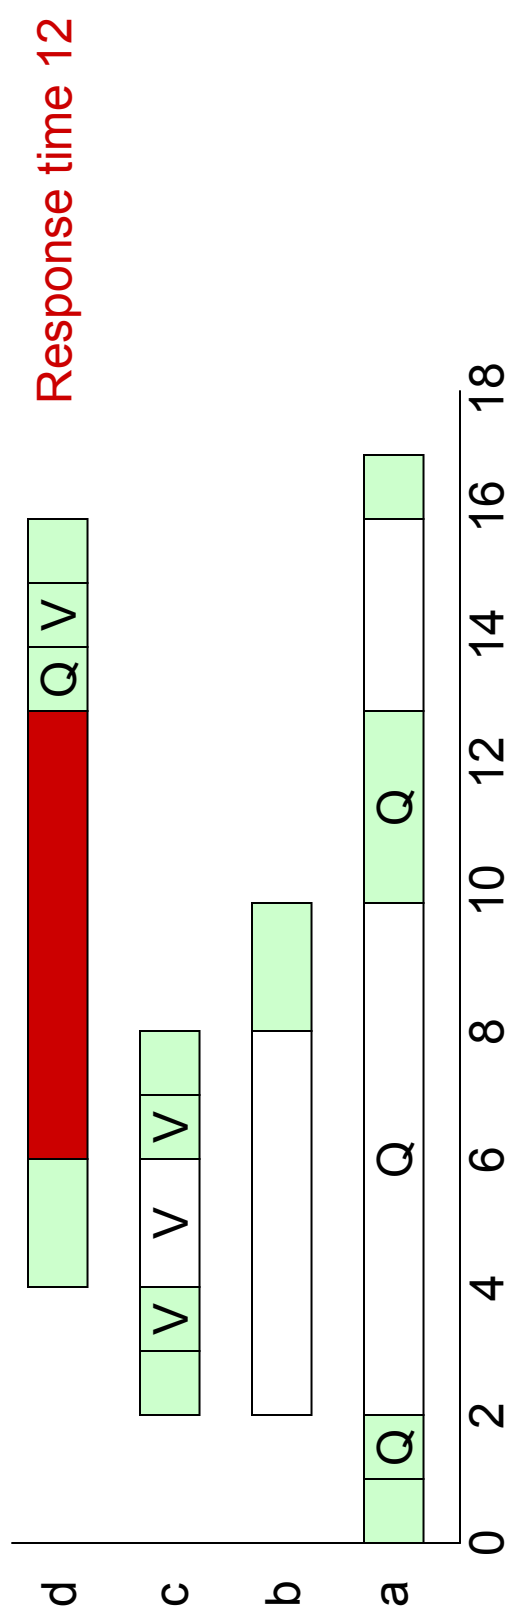
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inversion Example

[example from Burns & Wellings]

Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



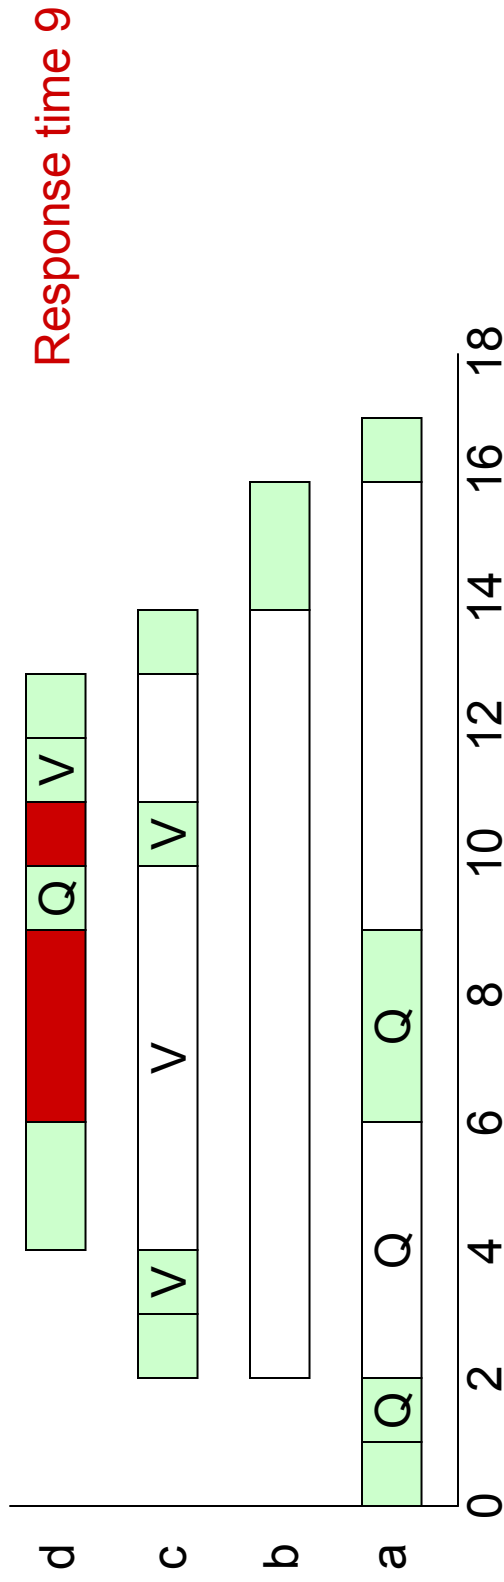
Priority Inheritance

- High priority process P is blocked on resource R
- R is currently held by low-priority process Q
- Priority Inheritance (Cornhill et al., 1987)
 - Boost priority of Q to that of P
 - Ensures that Q will finish quickly so that P can proceed

Priority Inheritance Example

[example from Burns & Wellings]


Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

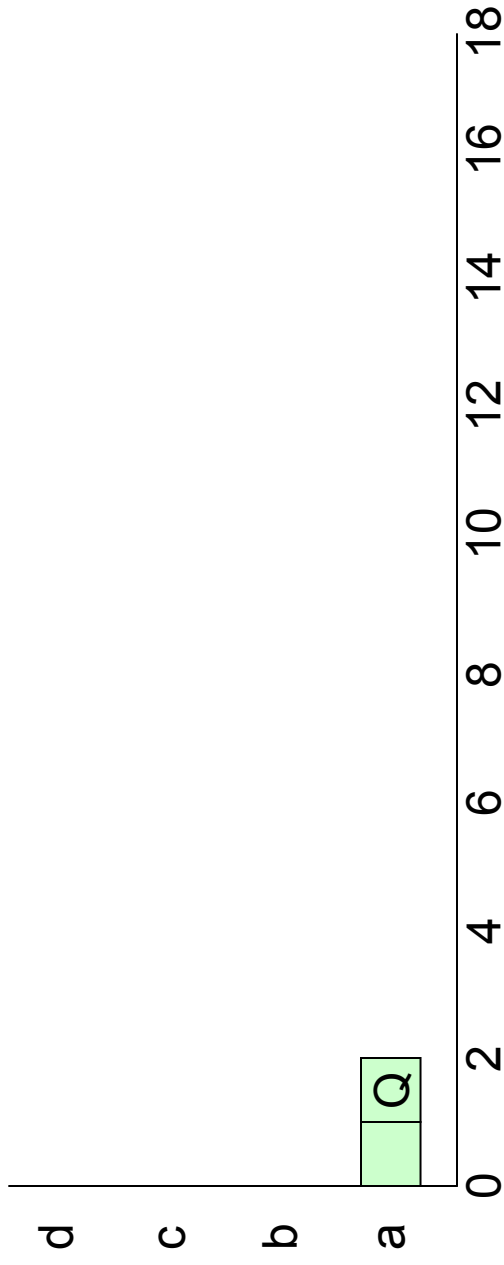
<u>Process</u>	<u>Priority</u>	<u>Execution</u>	<u>Release Time</u>
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

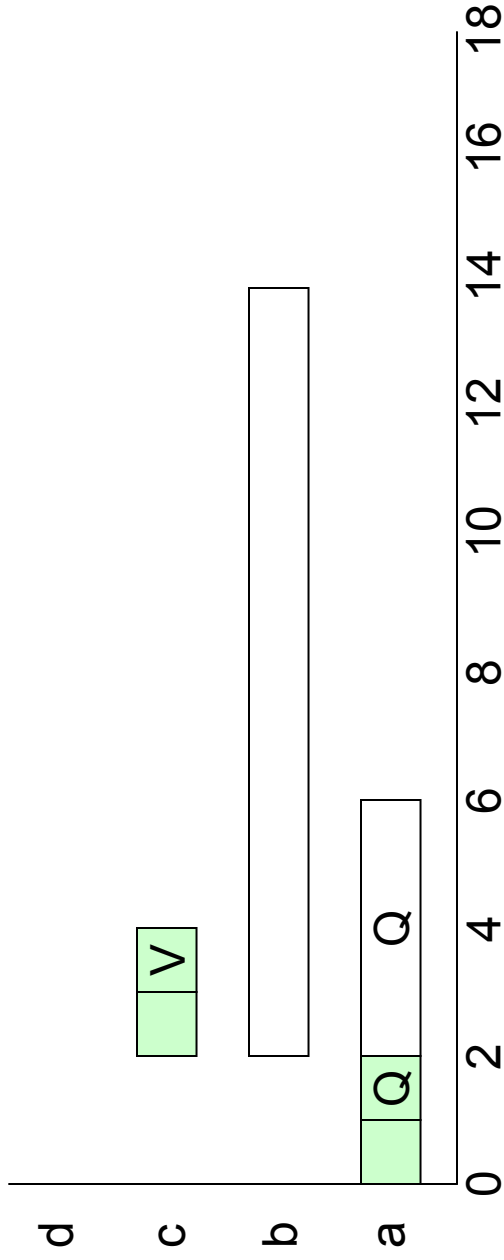
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

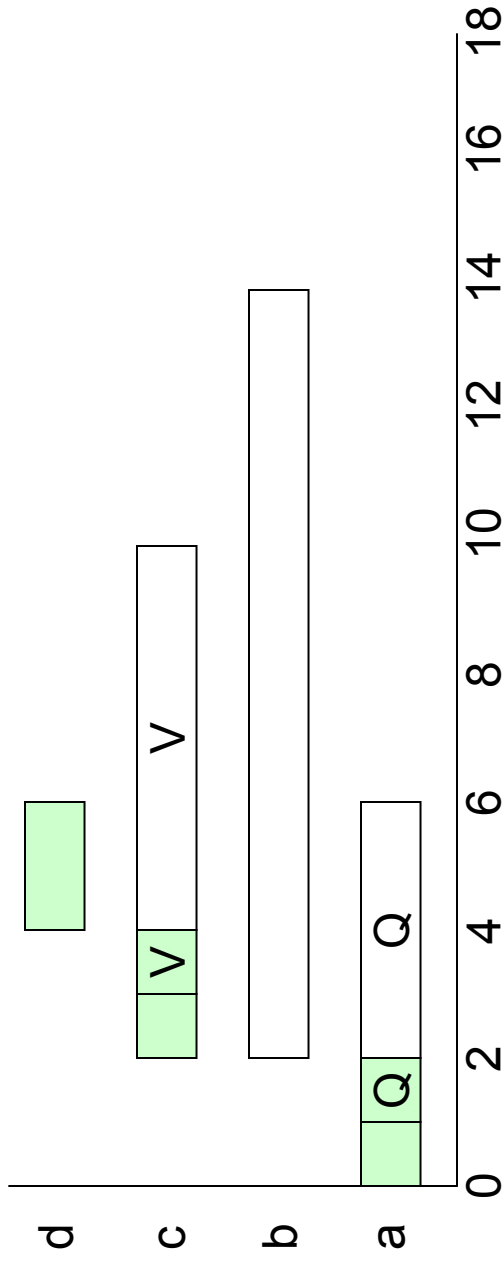
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

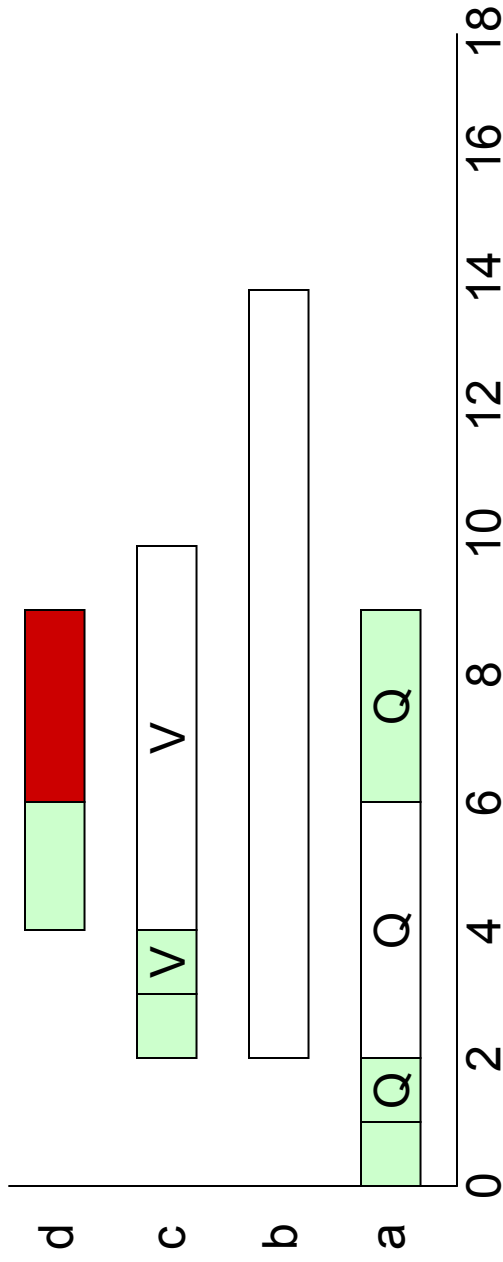
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

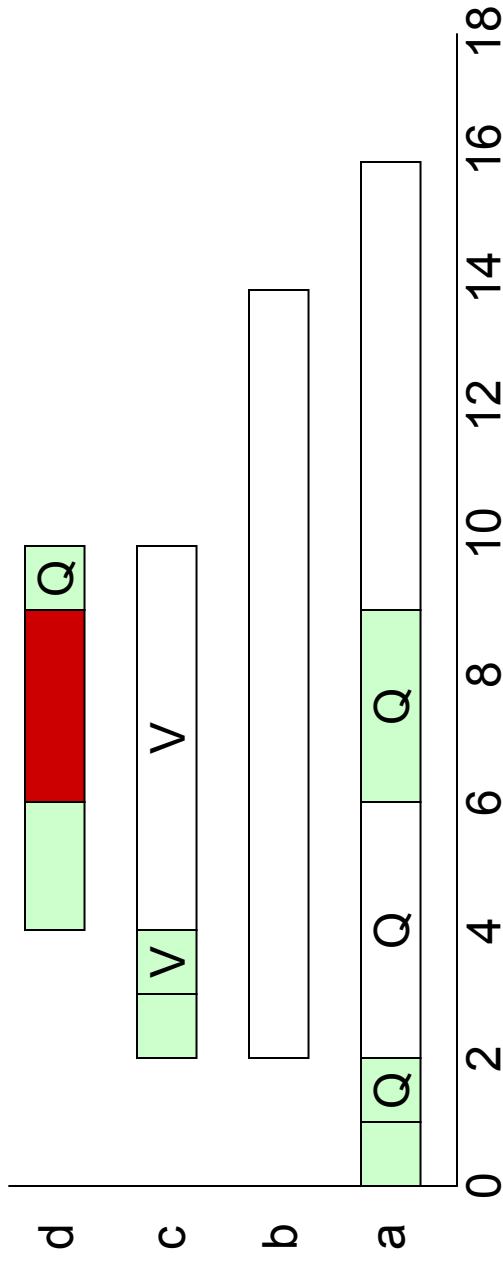
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

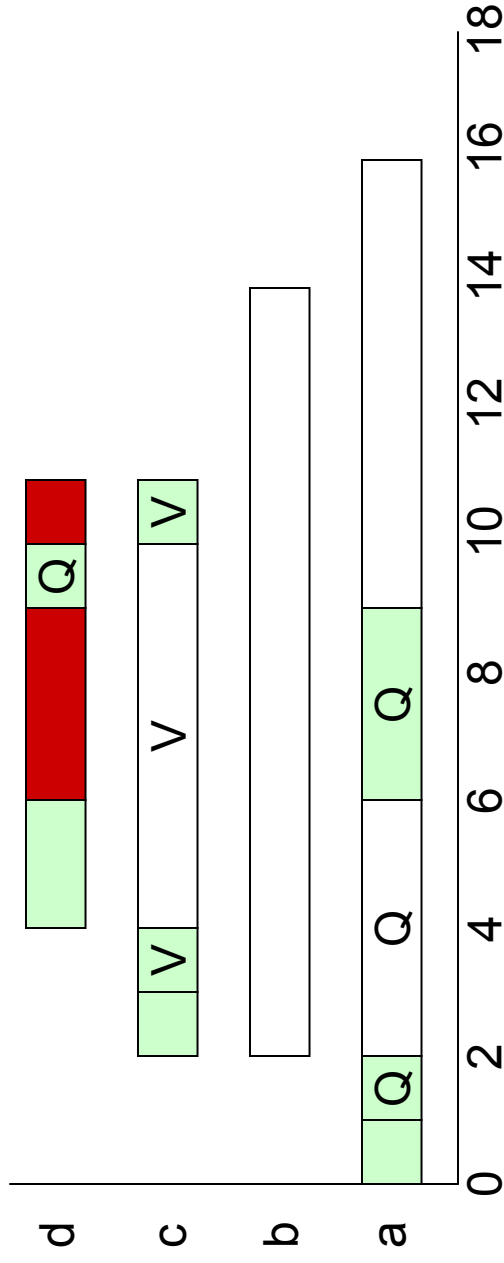
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

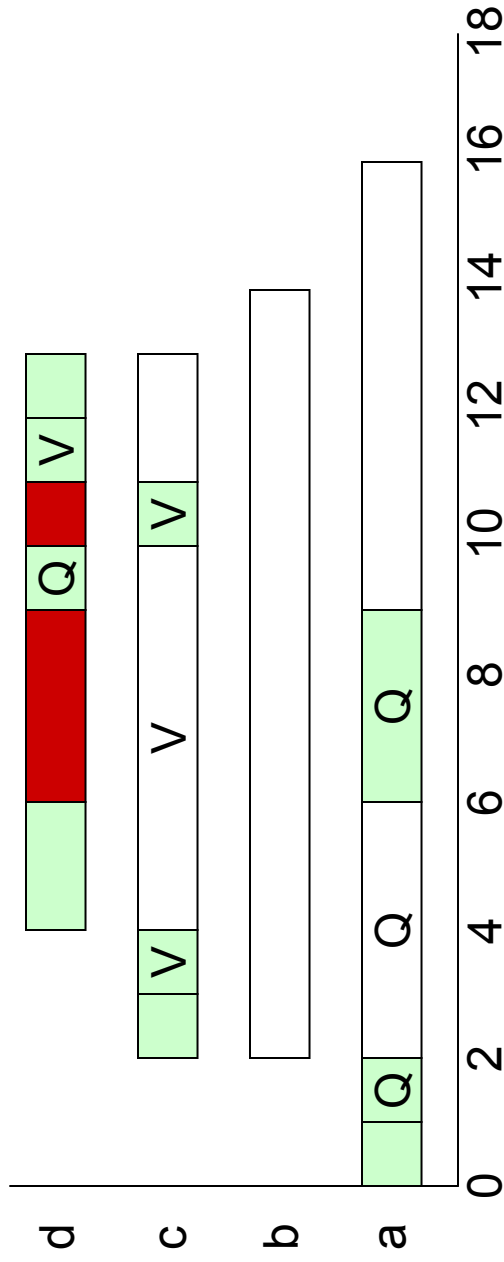
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

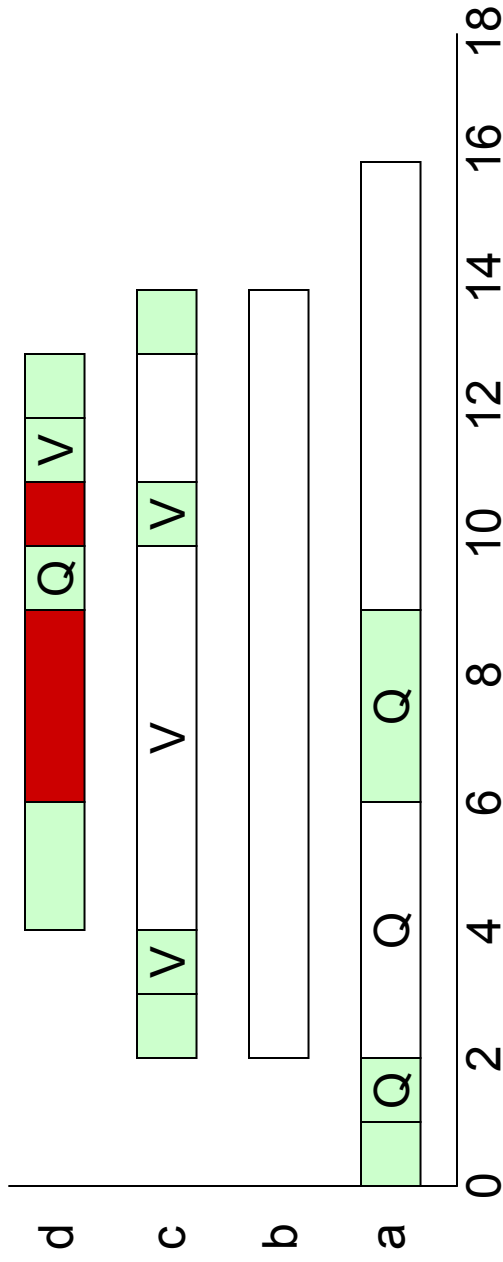
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

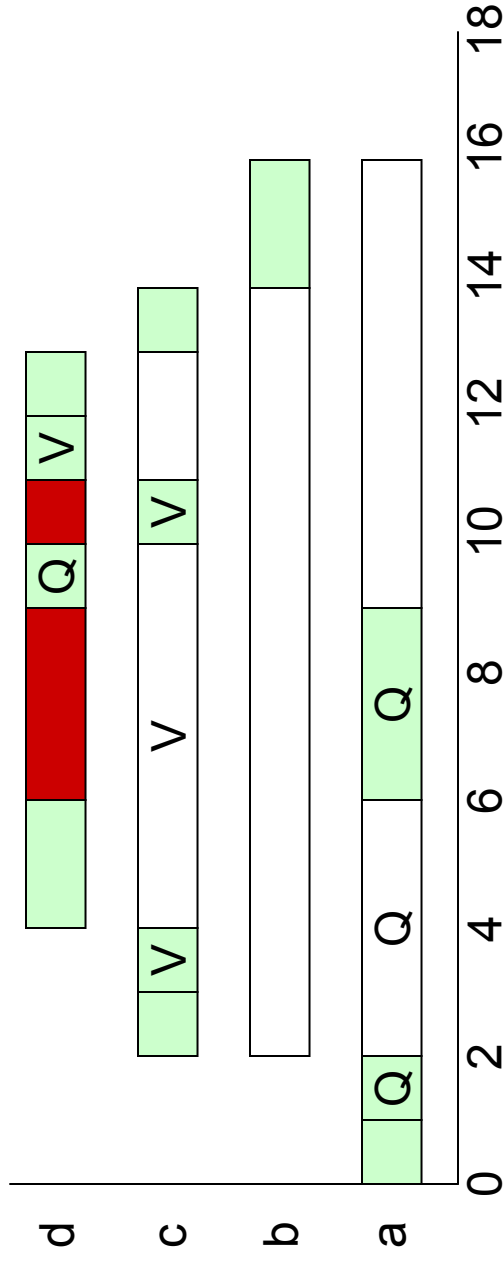
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

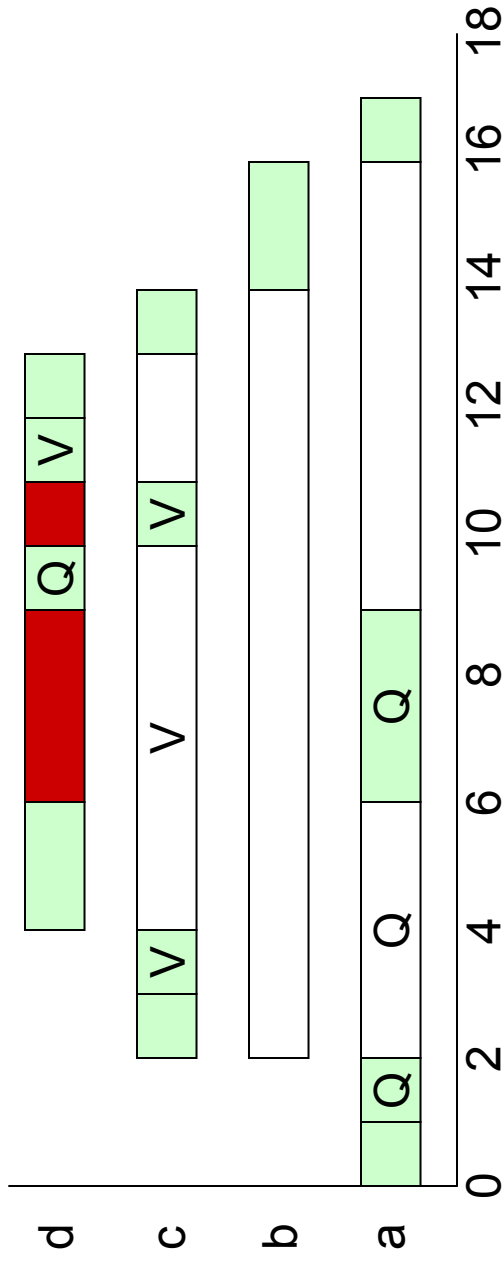
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

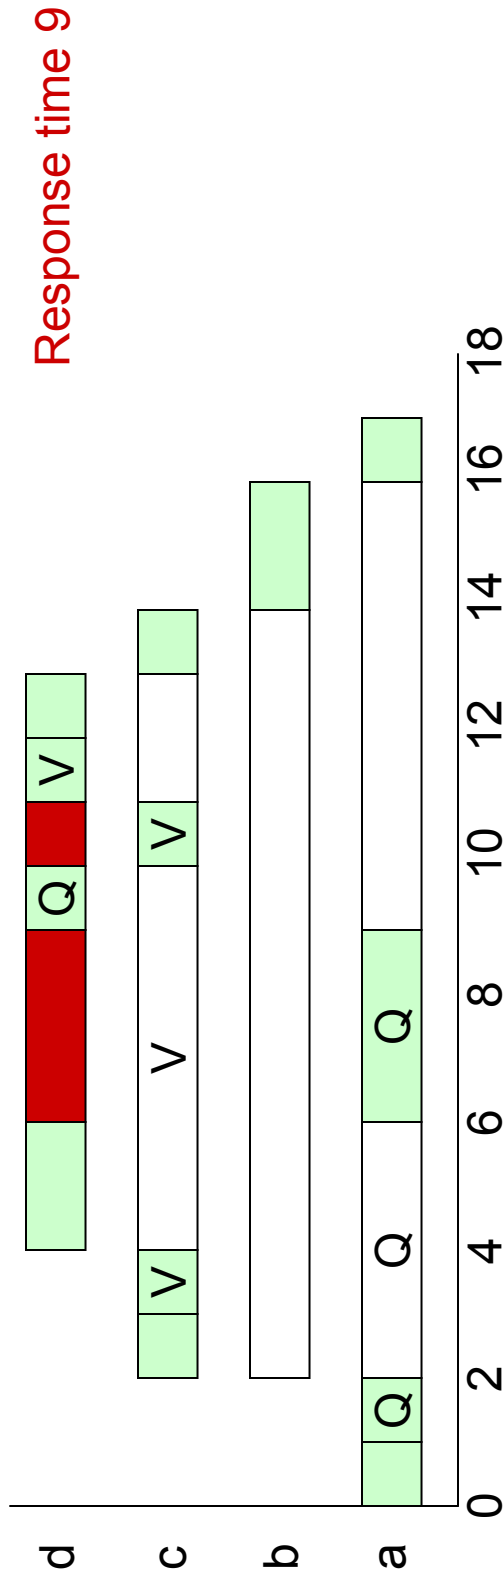
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Inheritance Example

[example from Burns & Wellings]

Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



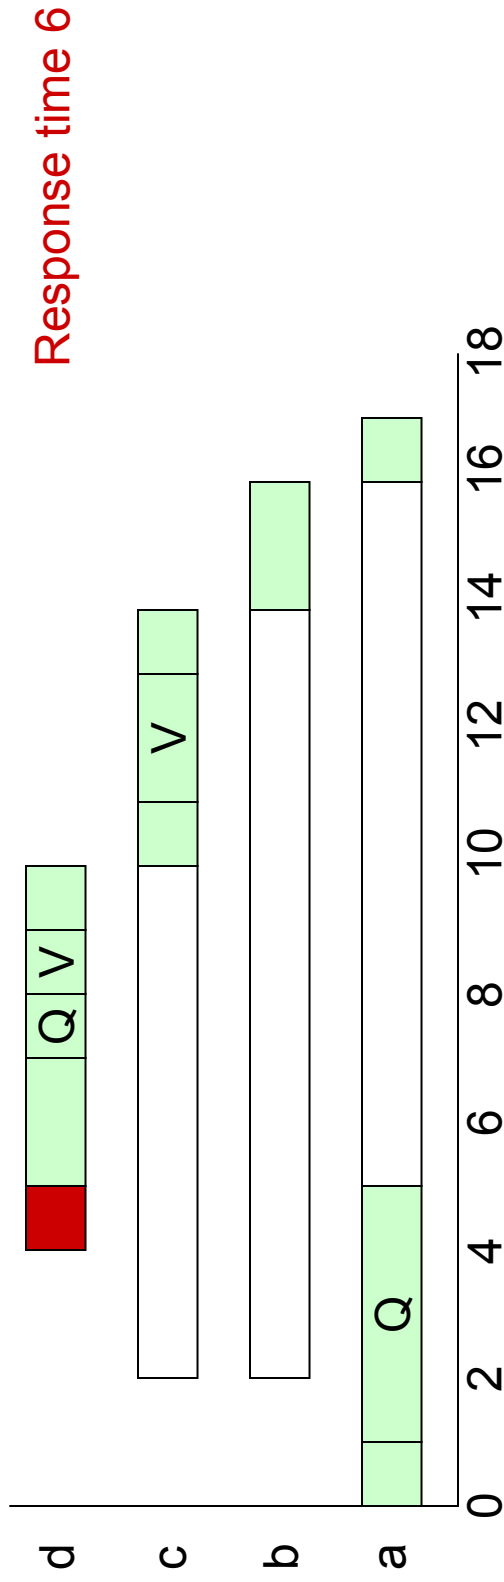
Priority Ceilings

- Can we do better than Priority Inheritance?
 - In the example, can we ensure process d is not blocked twice?
- Priority ceiling of resource R
 - The maximum priority of the processes that use R
- Dynamic priority of process P
 - The maximum of P 's static priority and the ceiling values of any resources it has locked
 - Intuition: if I've locked a resource R that some higher priority process Q might need, this rule ensures that no process with lower-priority than Q can possibly run and acquire some other resource R' , so that Q would have to wait for two processes to acquire R and R'

Priority Ceiling Example

[example from Burns & Wellings]


Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]


Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

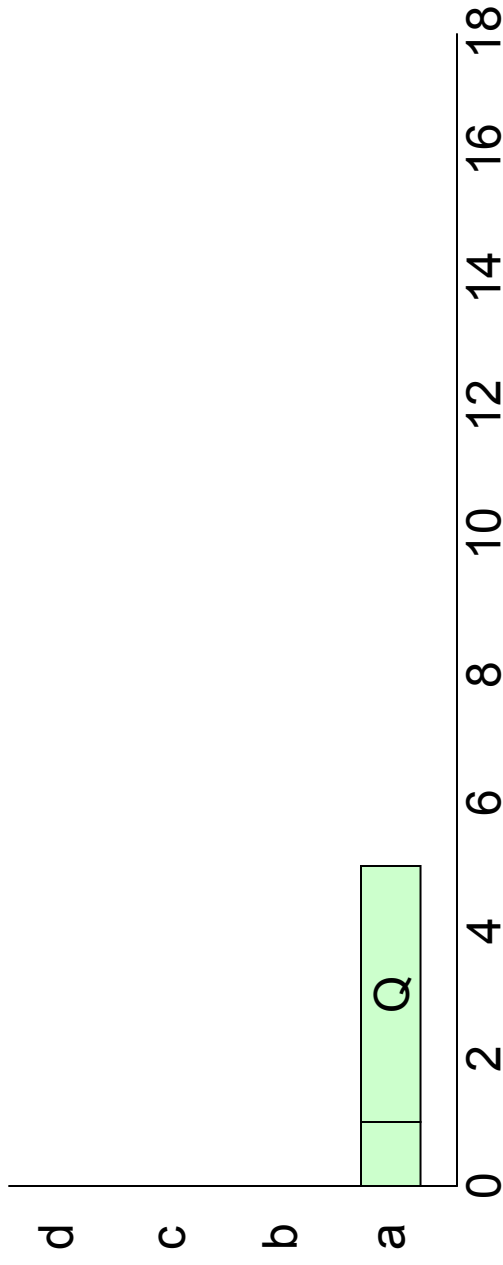
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

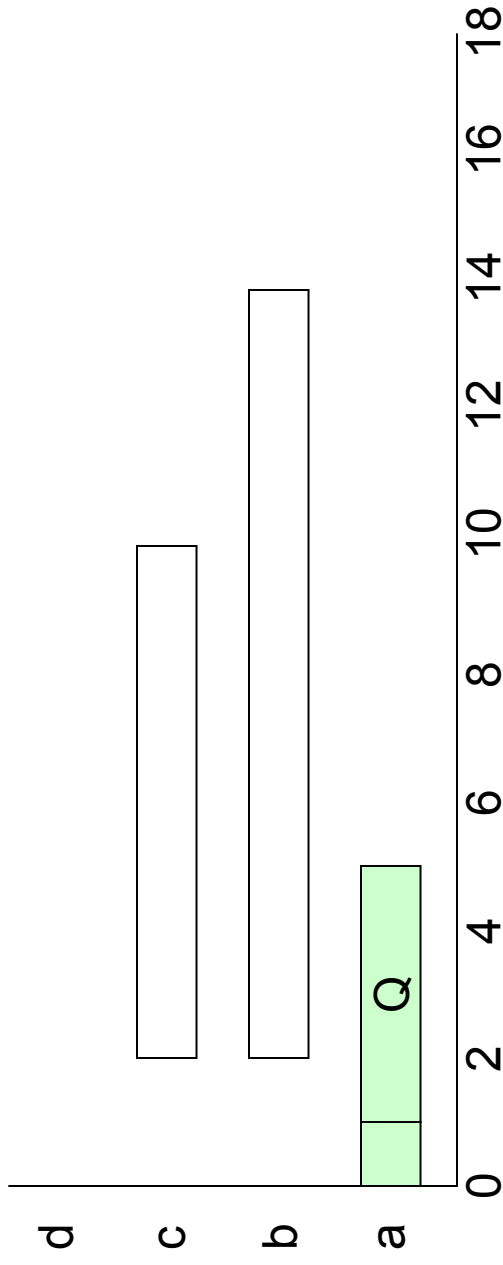
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

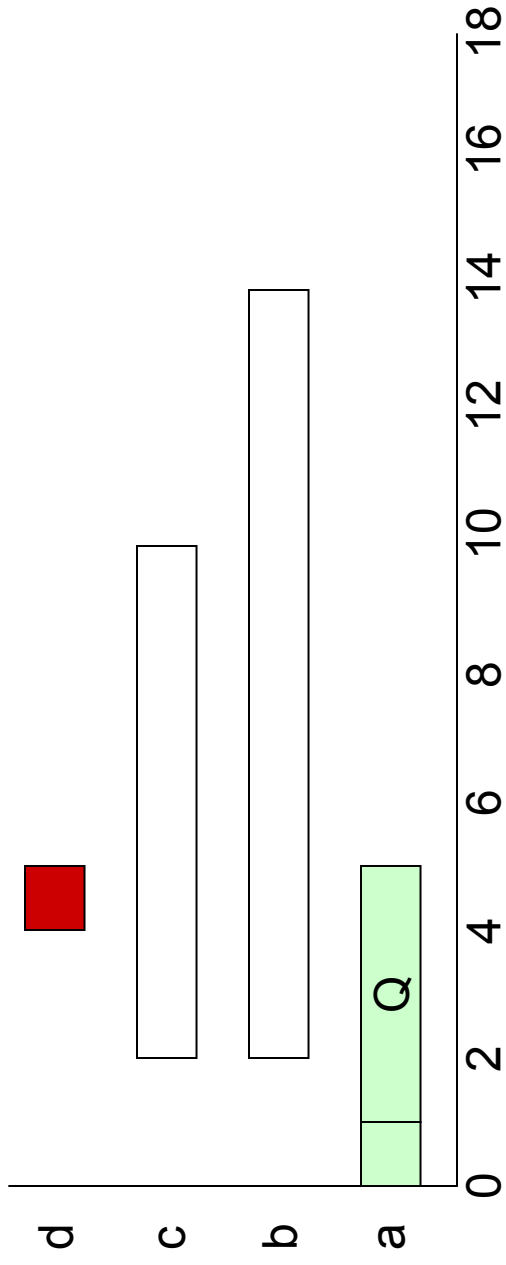
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

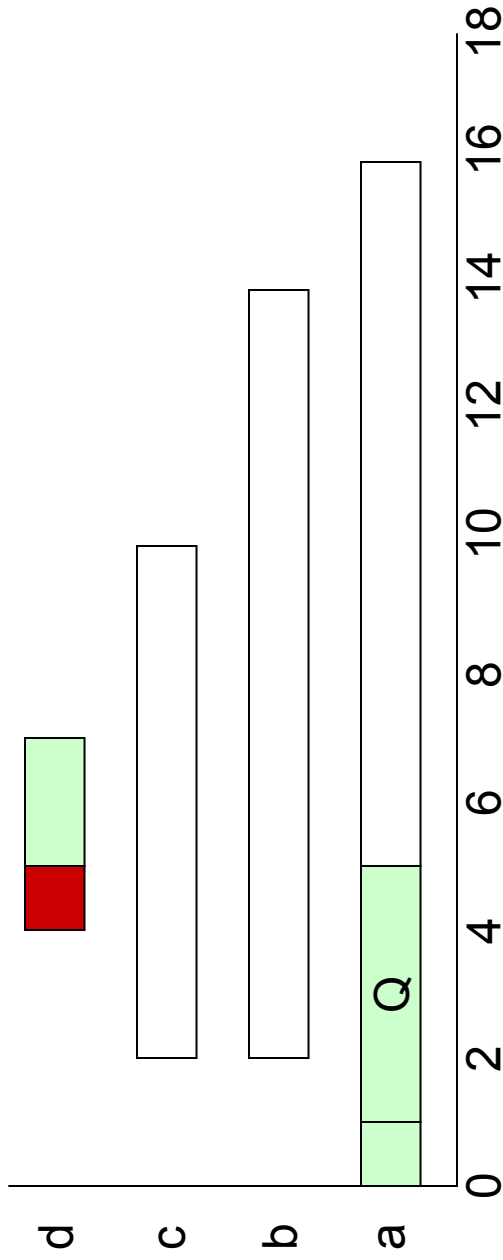
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

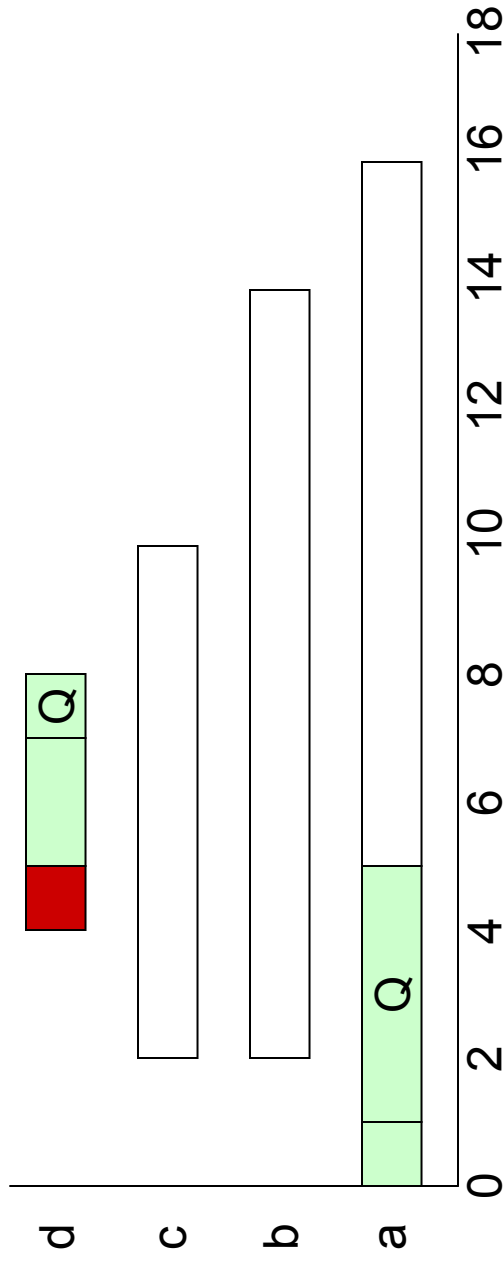
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

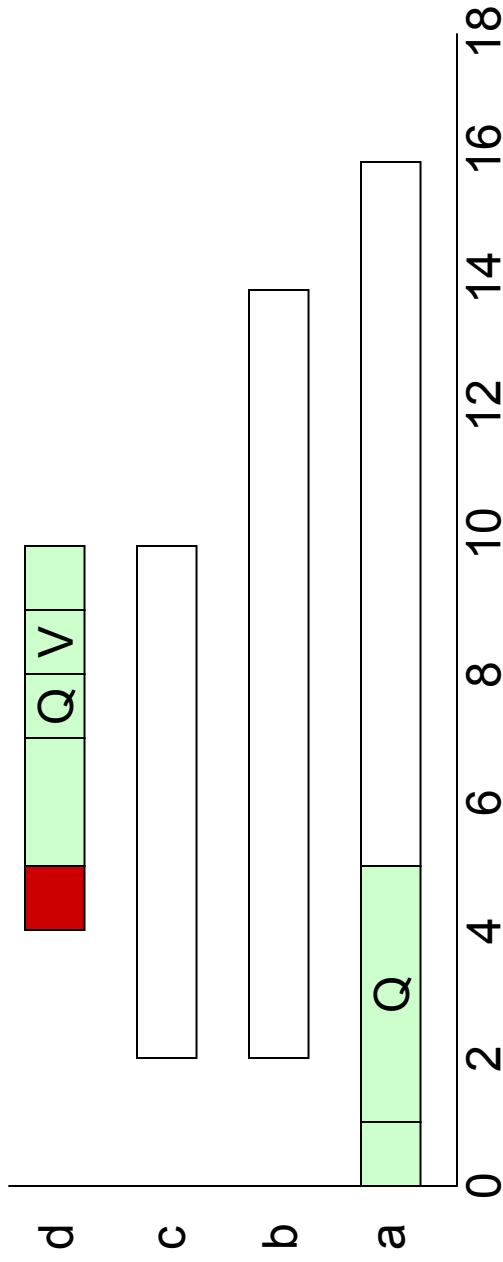
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

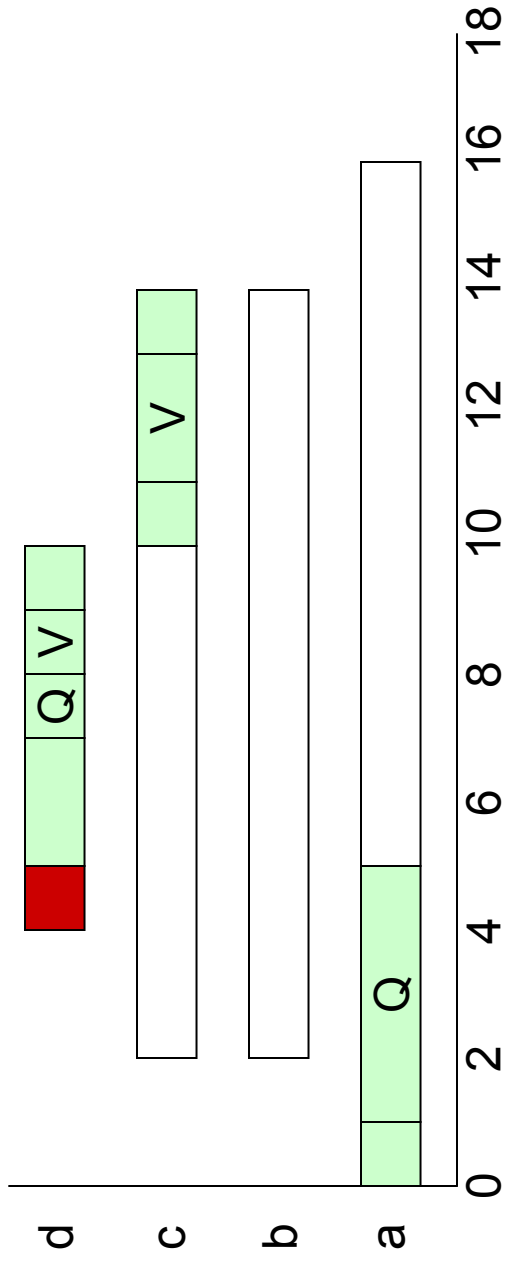
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

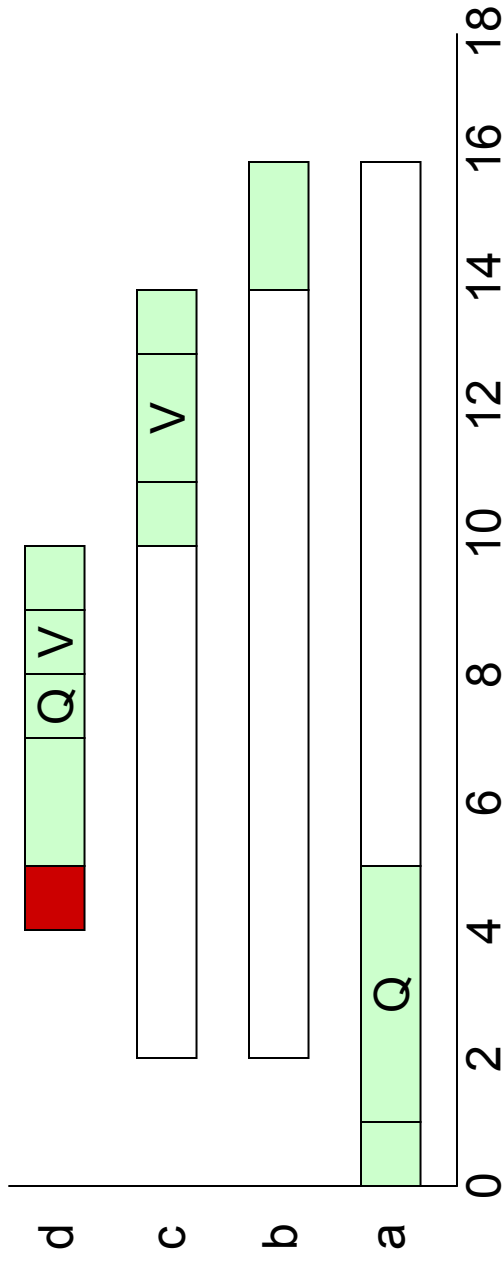
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

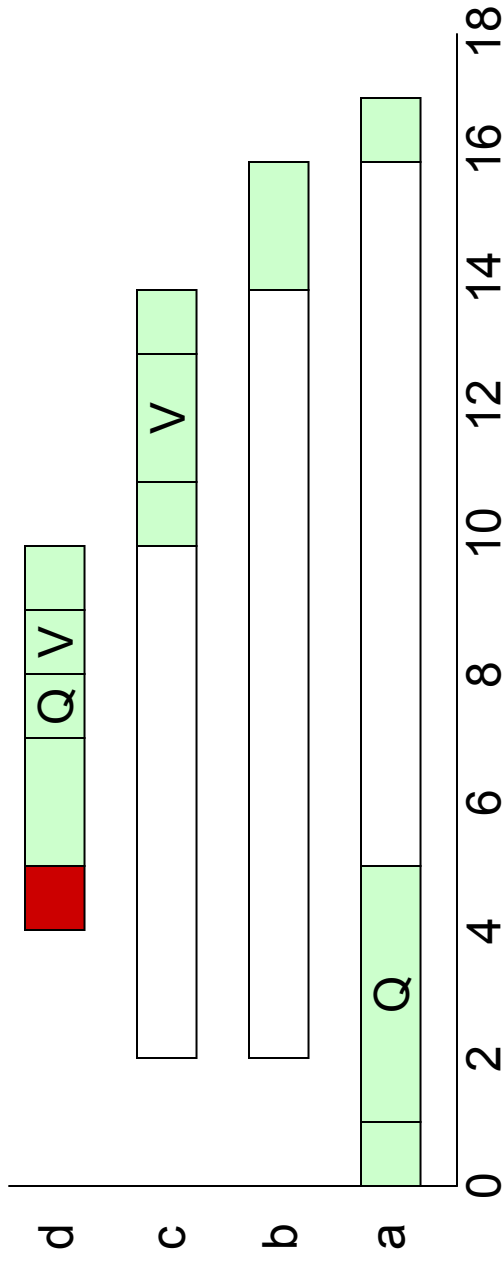
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

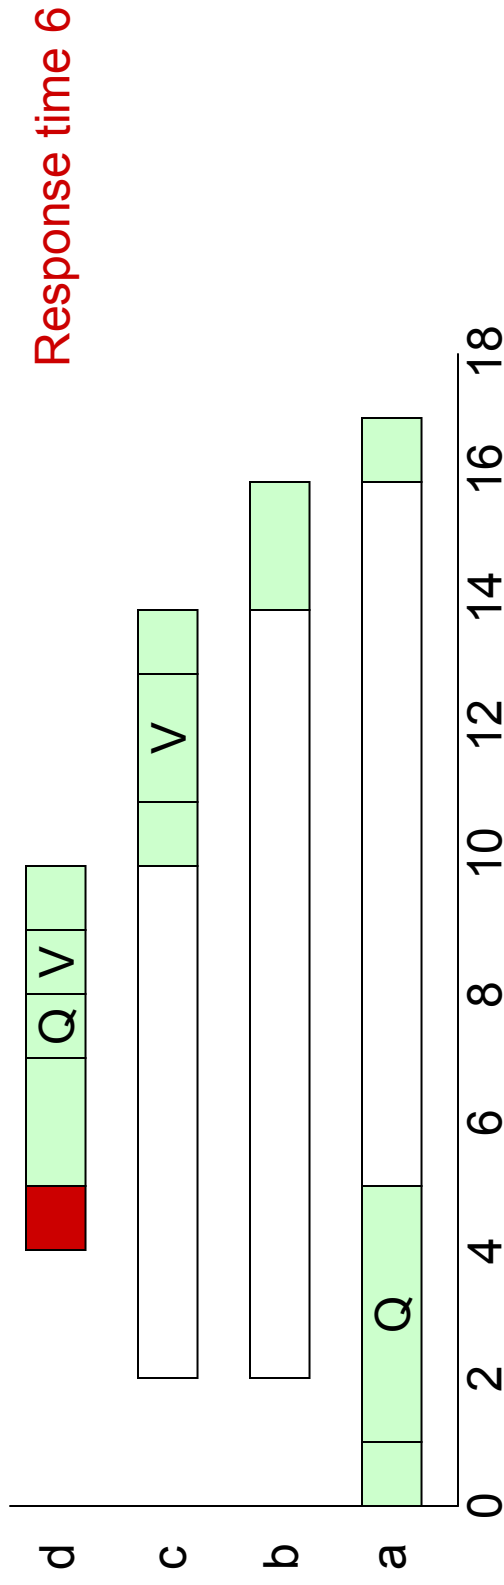
Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Priority Ceiling Example

[example from Burns & Wellings]

Process	Priority	Execution	Release Time
a	1	EQQQE	0
b	2	EE	2
c	3	EVVE	2
d	4	EEQVE	4



Benefits of Priority Ceilings

[Burns & Wellings]

- A high priority process can be blocked at most once by lower-priority processes
- Deadlocks are prevented
 - on single-processor systems
- Transitive blocking is prevented
- Mutually exclusive access to resources is ensured

Summary: Analysis of Real-Time Systems

- Resource use is challenging
 - Memory typically allocated statically or on stack
 - RT Java: regions
 - Future: real-time garbage collection
- Scheduling
 - Simple case: cyclic executive
 - General case: fixed priority scheduling
 - Assign priorities according to deadlines
 - Timelines used to check feasibility
 - Use priority inheritance to avoid priority inversion problems