An Exact Algorithm for the Linear Tape Scheduling Problem

Valentin Honoré, **Bertrand Simon**, Frédéric Suter

IN2P3 Computing Center – CNRS, Villeurbanne – Lyon, France

JCAD – October 2022

Tape usage today

(or 11 football fields)

$\approx$ 20TB on 1000s $\times$ 1km read at 10m/s – 100s MB/s

https://commons.wikimedia.org/wiki/File:LTO2-cart-wo-top-shell.jpg
https://commons.wikimedia.org/wiki/File:Usain_Bolt_Rio_100m_final_2016i-cr.jpg

Primordial for HTC (High Throughput Computing)

e.g., CCIN2P3 CERN ECMWF Forecast User Guide (100s PB)

also: media companies, cloud archive...

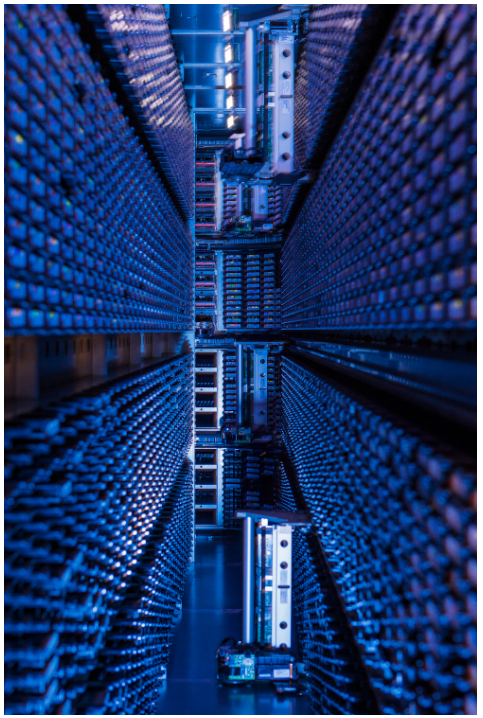

☺ Impressive technology improvements
density: + 30% / year (vs HDD: + 8%)

☹ High latency (mount, load, position $\rightarrow$ few mn)
Adapted for Write Once Read Many

Overview of a tape

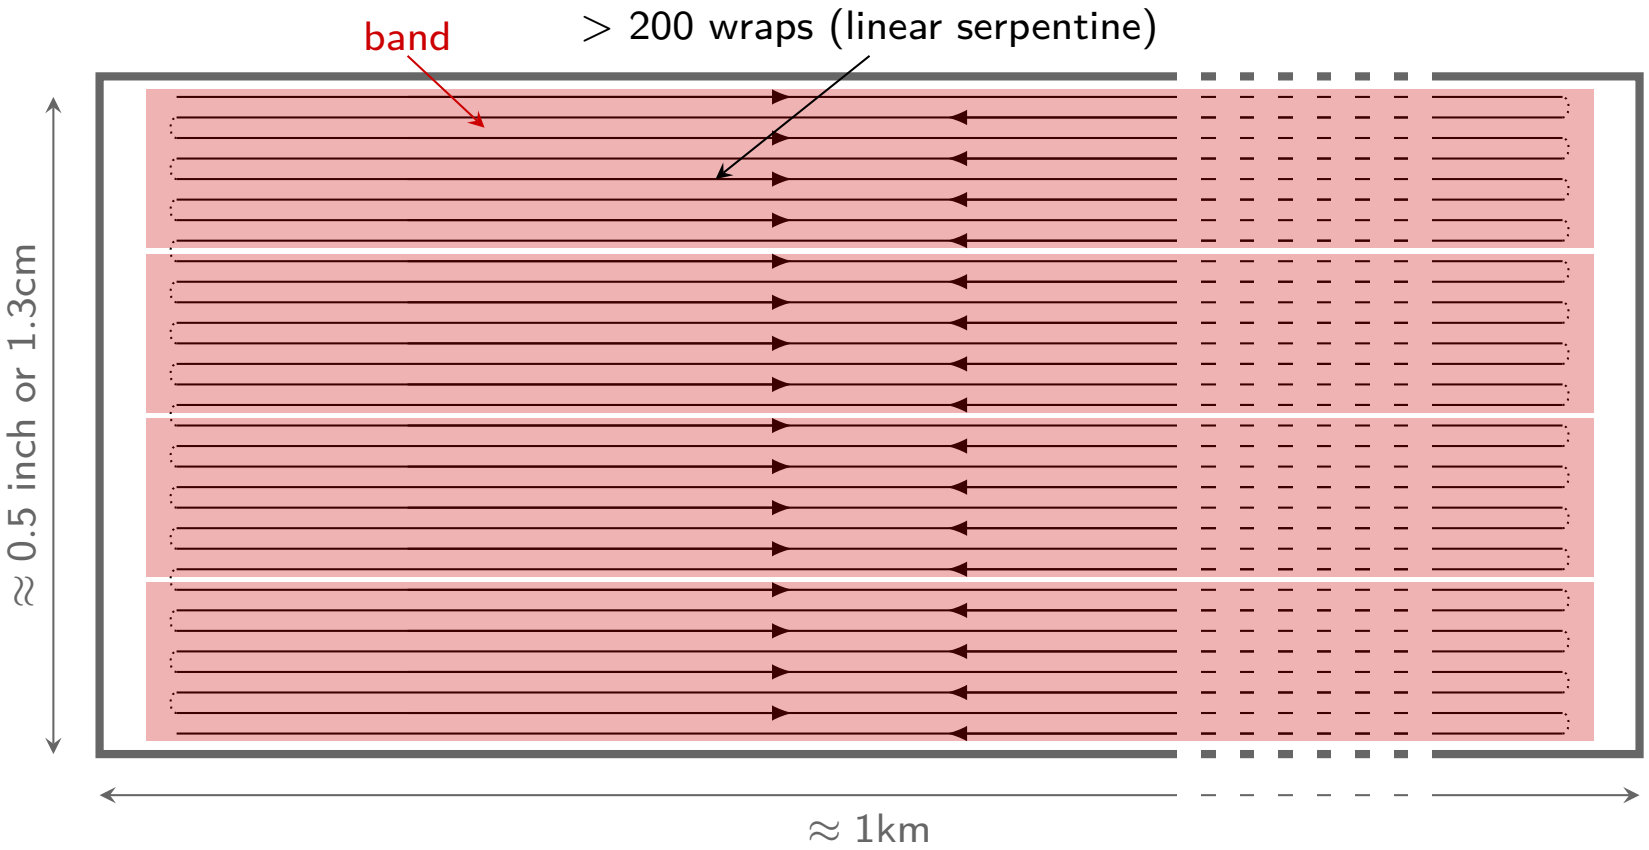

wrap = dozens of tracks read / written simultaneously by parallel heads

Overview of our tape model

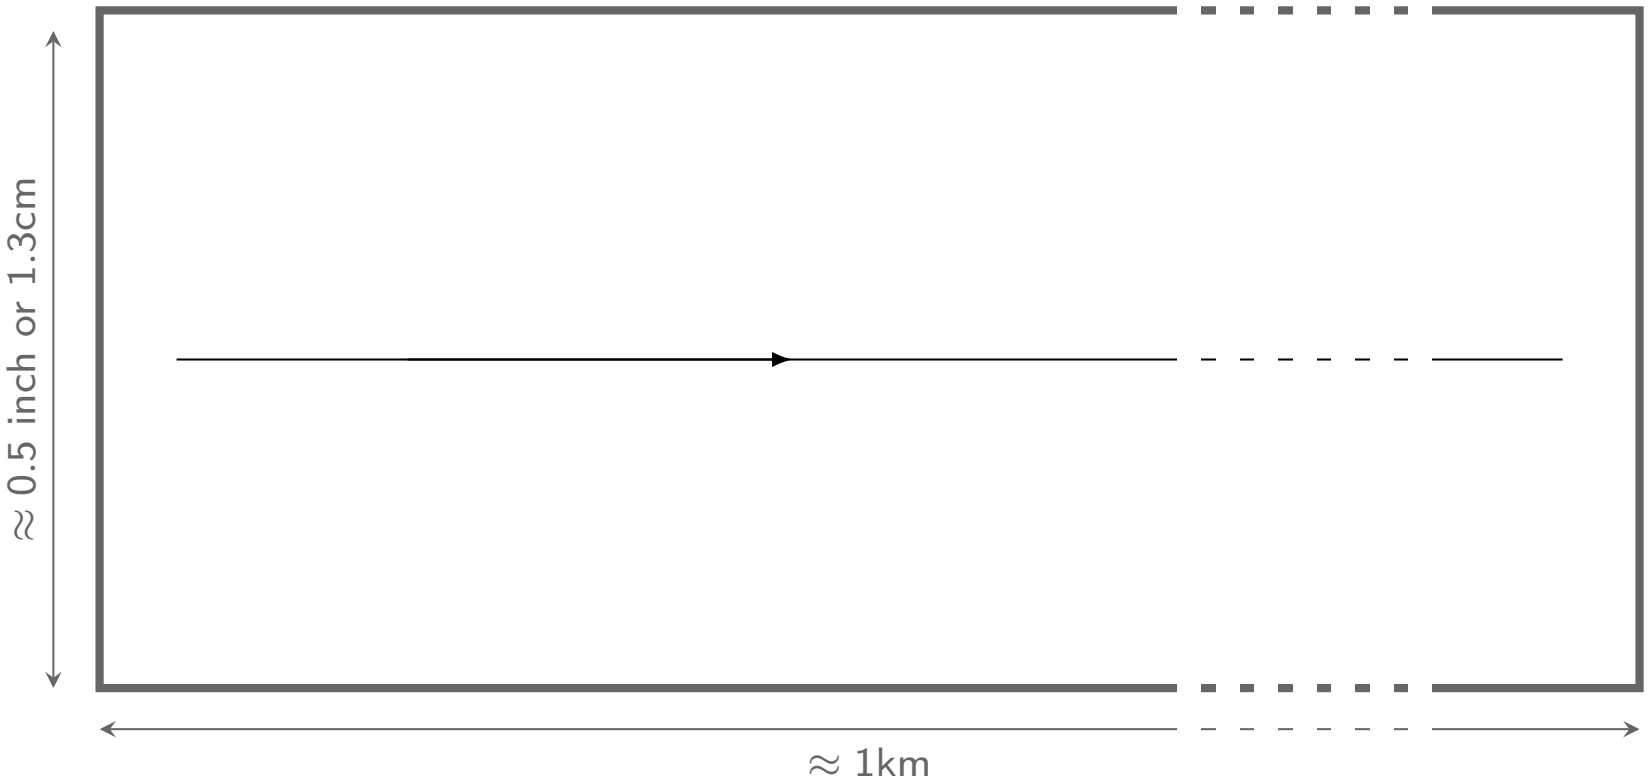

Linear Tape Scheduling Problem [CardonhaReal'16]

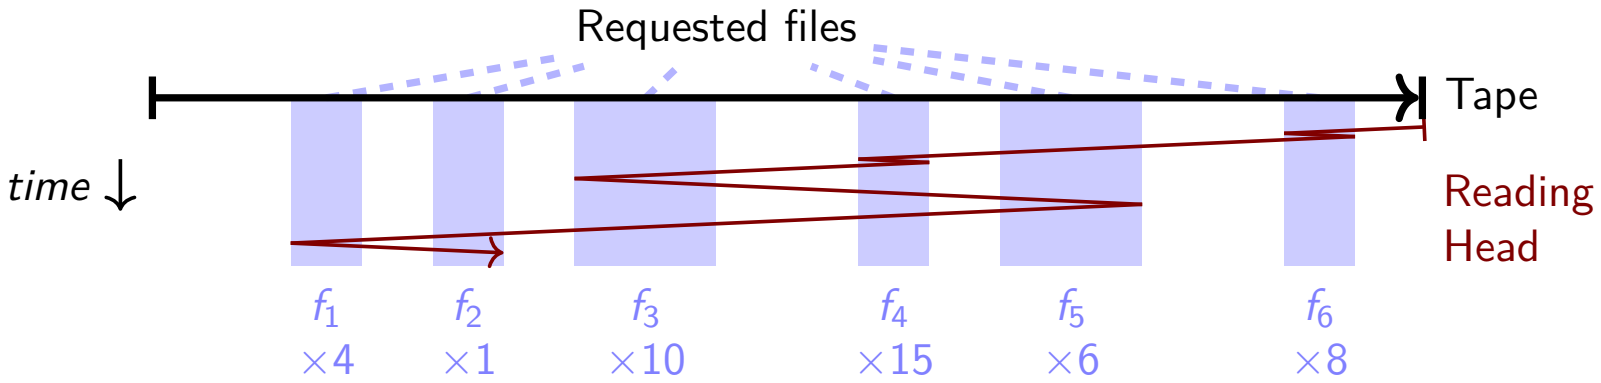

Assumptions:

- files are read left-to-right
- start on the right
- constant speed

Input:

- tape of n_f consecutive files
- n file requests (44 here)
- n_{req} distinct files requested (6)

Objective: average service time

Motivation: lack of fundamental theoretical results, models local files

Linear Tape Scheduling Problem [CardonhaReal'16]

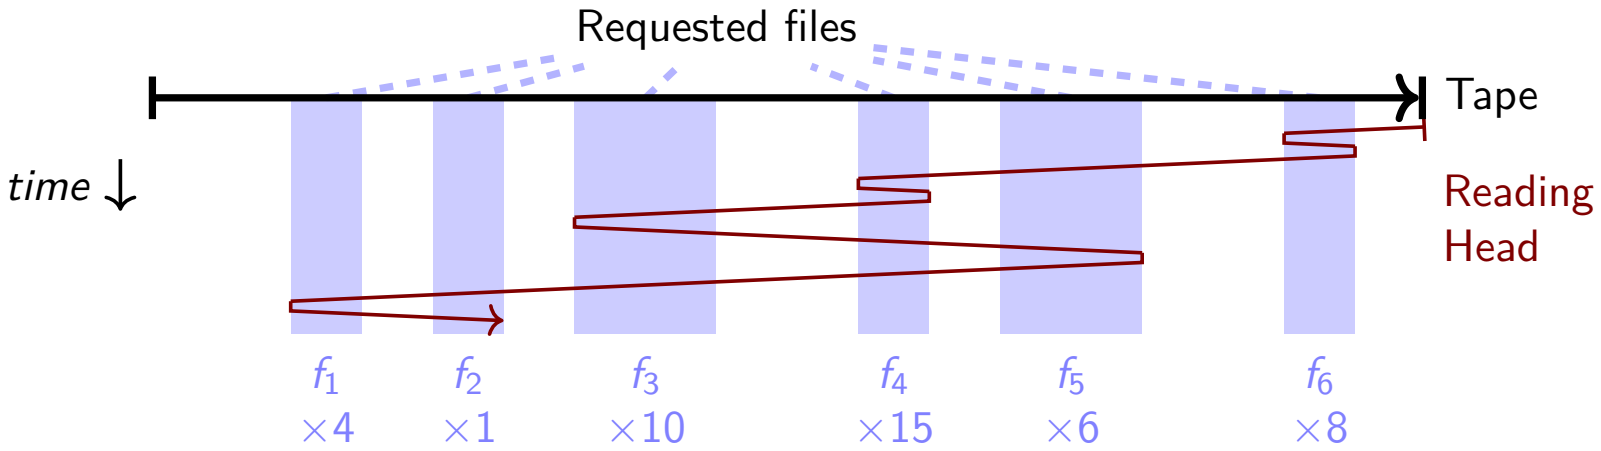

Assumptions:

- files are read left-to-right
- start on the right
- constant speed
- [new] U-turn penalty U

Input:

- tape of n_f consecutive files
- n file requests (44 here)
- n_{req} distinct files requested (6)

Objective: average service time

Motivation: lack of fundamental theoretical results, models local files

Related to the Linear Tape Scheduling Problem

Travelling Salesperson Problem (TSP)

- super-famous NP-hard problem
- recent $(1.5 - 10^{-36})$-approximation [KarlinKleinGaran'21]
- ☹ minimizes makespan, trivial on the real line

Minimum Latency Problem / TRP (Repair) - variant

- ☺ minimize average service time $\in P$ on the real line
- delays to repair a node: complexity open

Dial-a-ride variant on the real line

- $\approx$ LTSP but with overlapping files in both directions $\longrightarrow$ NP-hard

Tapes except LTSP: 2 specific experimental papers in the 90's

Structural results

Any optimal solution

- after reaching $\ell(f_1)$, go straight to the rightmost unread request
- can be described by a set of **detours** done before

Definition (Detours)

A solution includes the **detour** (a,b) with $a \leq b$ if:

- the 1st time the head reaches $\ell(a)$, go straight to $r(b)$, back to $\ell(a)$

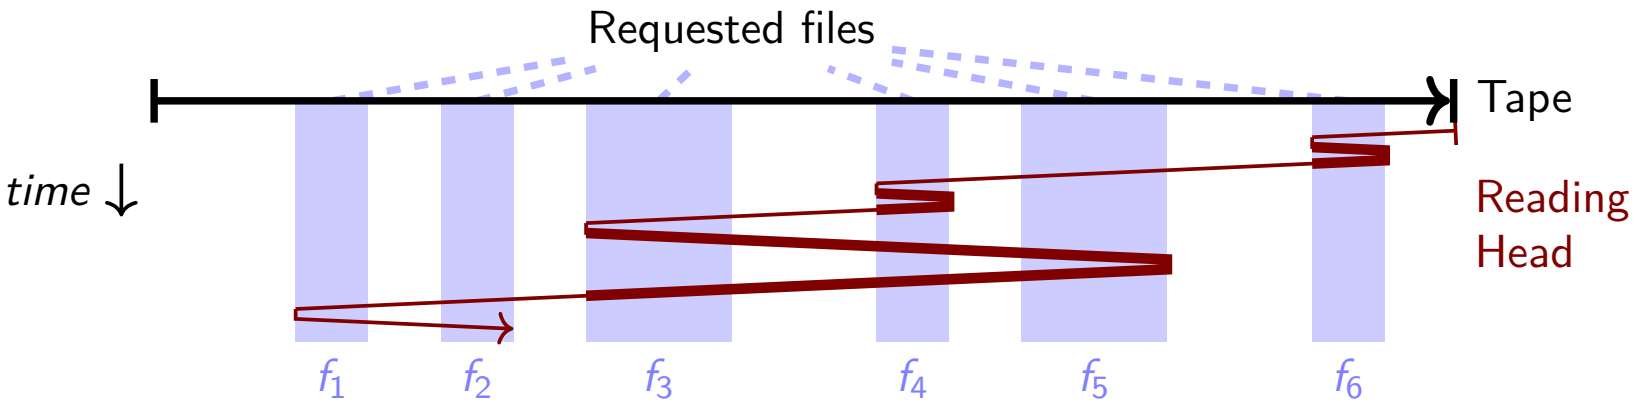

detours: $[(f_6, f_6), (f_4, f_4), (f_3, f_5)]$

Structural results

Any optimal solution

- after reaching $\ell(f_1)$, go straight to the rightmost unread request
- can be described by a set of **detours** done before

Definition (Detours)

A solution includes the **detour** (a,b) with $a \leq b$ if:

- the 1st time the head reaches $\ell(a)$, go straight to $r(b)$, back to $\ell(a)$

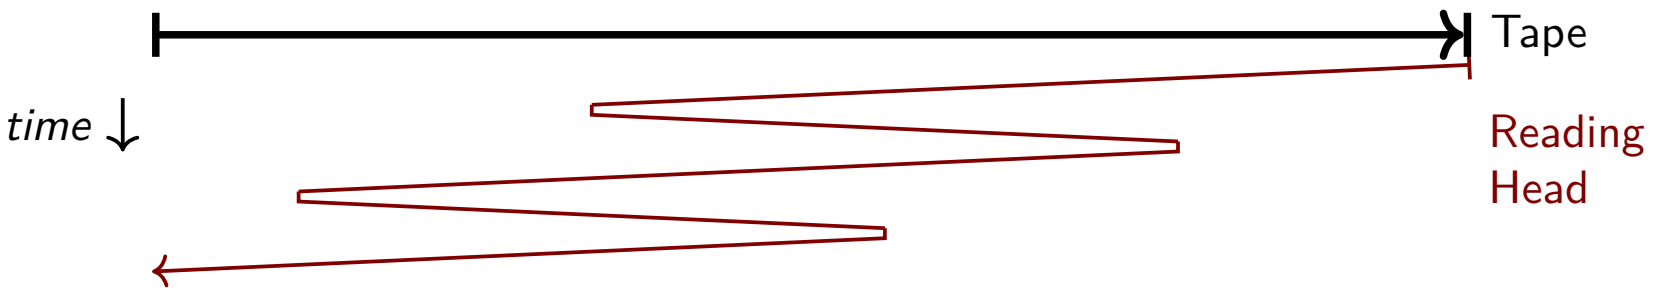

Lemma: detours never partially overlap (*strictly laminar*)

Structural results

Any optimal solution

- after reaching $\ell(f_1)$, go straight to the rightmost unread request
- can be described by a set of **detours** done before

Definition (Detours)

A solution includes the **detour** (a,b) with $a \leq b$ if:

- the 1st time the head reaches $\ell(a)$, go straight to $r(b)$, back to $\ell(a)$

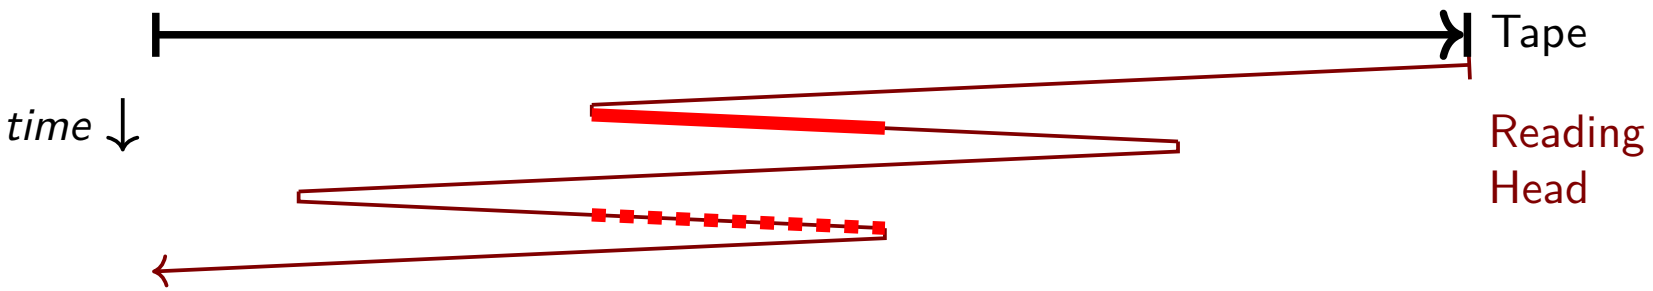

Lemma: detours never partially overlap (*strictly laminar*)

Naive algorithms

NoDetour: go to the leftmost request, then to the rightmost

can be arbitrarily bad (place urgent requests on the right)

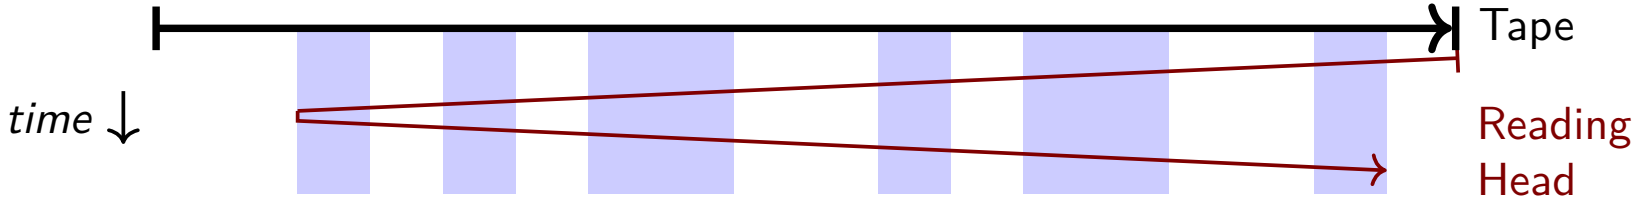

GS (Greedy Schedule): do all atomic detours, *i.e.*, $\{(f_i, f_i)\}_{\forall i}$

Lemma [CardonhaReal'16] : **GS** is a 3-approximation if $U = 0$

Proof: does ≤ 3 times the optimal distance before reading each request

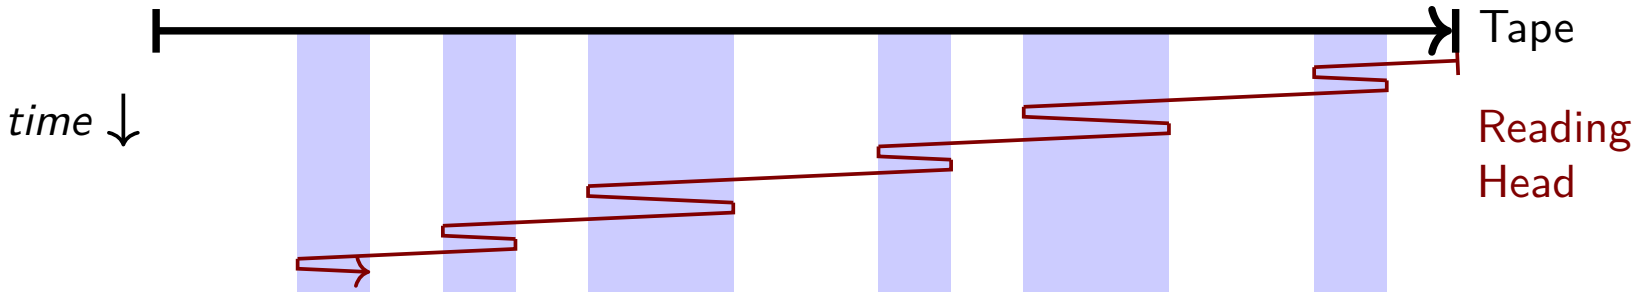

Heuristic improvements [CardonhaCiréReal'18]

FGS (Filtered): remove detrimental atomic detours in $O(n_{req}^2)$

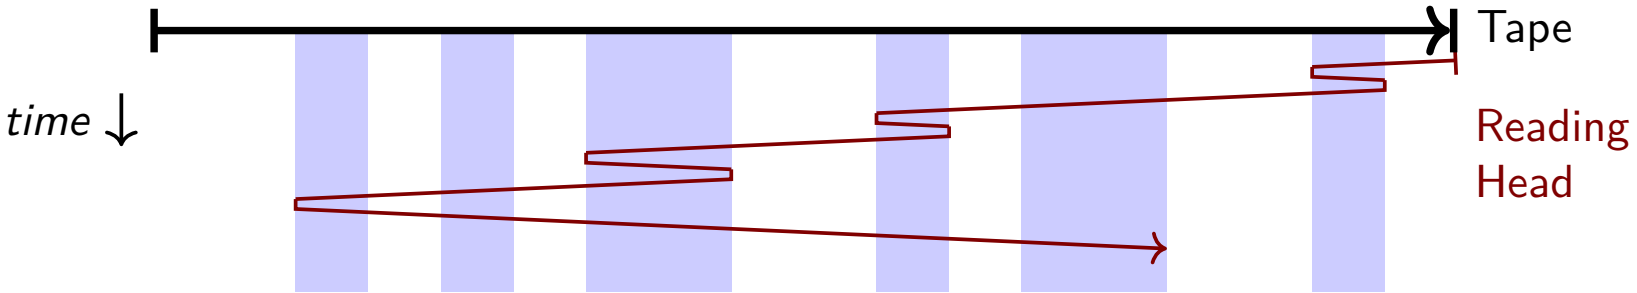

NFGS (Non-atomic): greedily add long detours if currently beneficial. Make one pass from left to right. Complexity in $O(n_{req}^3)$

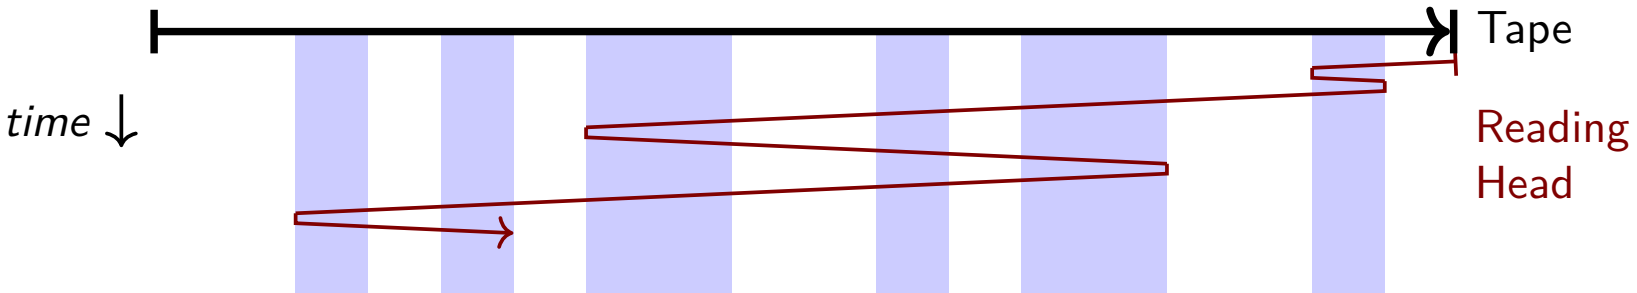

Dynamic Program DP: overview

Each cell: three parameters $T[a, b, n_{skip}]$

- compute the best strategy from $r(\mathbf{b})$ to $\ell(\mathbf{a})$ assuming:

1. there is a detour $(\mathbf{a}, f)$ for some $f \geq \mathbf{b}$,
2. there is no detour (f_1, f_2) such that $\mathbf{a} < f_1 < \mathbf{b} < f_2$,
3. when reaching $r(\mathbf{b})$, exactly n_{skip} requests have been skipped.

$\Rightarrow$ value $\approx$ cost contribution from 'first $r(b)$' to '$r(b)$ after reading a'

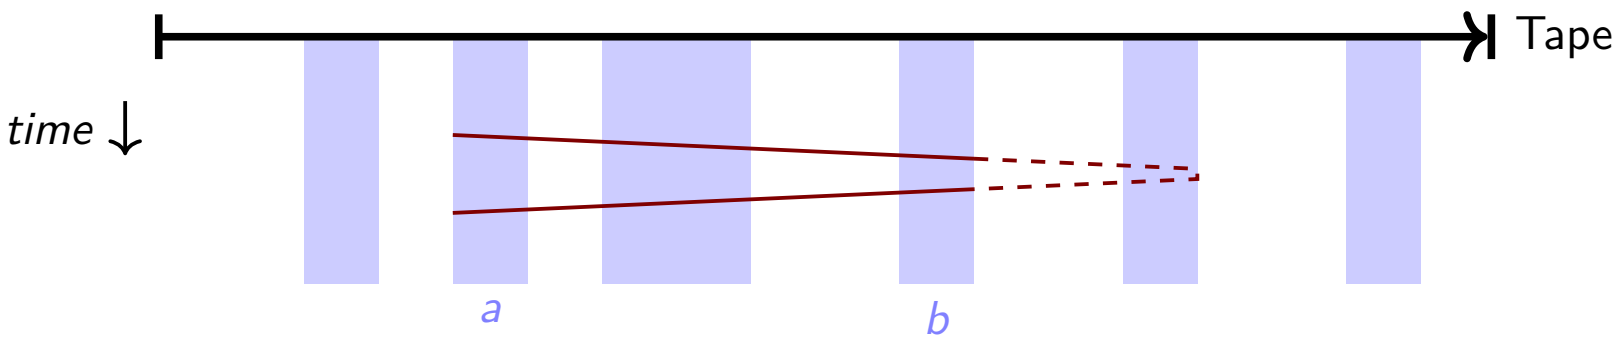

Subtleties: $\forall$ request on f, do not count the cost $m \to \ell(f) \to r(f)$
if f is read after b, remove one U (counted before)

Dynamic Program DP: overview

Each cell: three parameters $T[a, b, n, \cdot]$

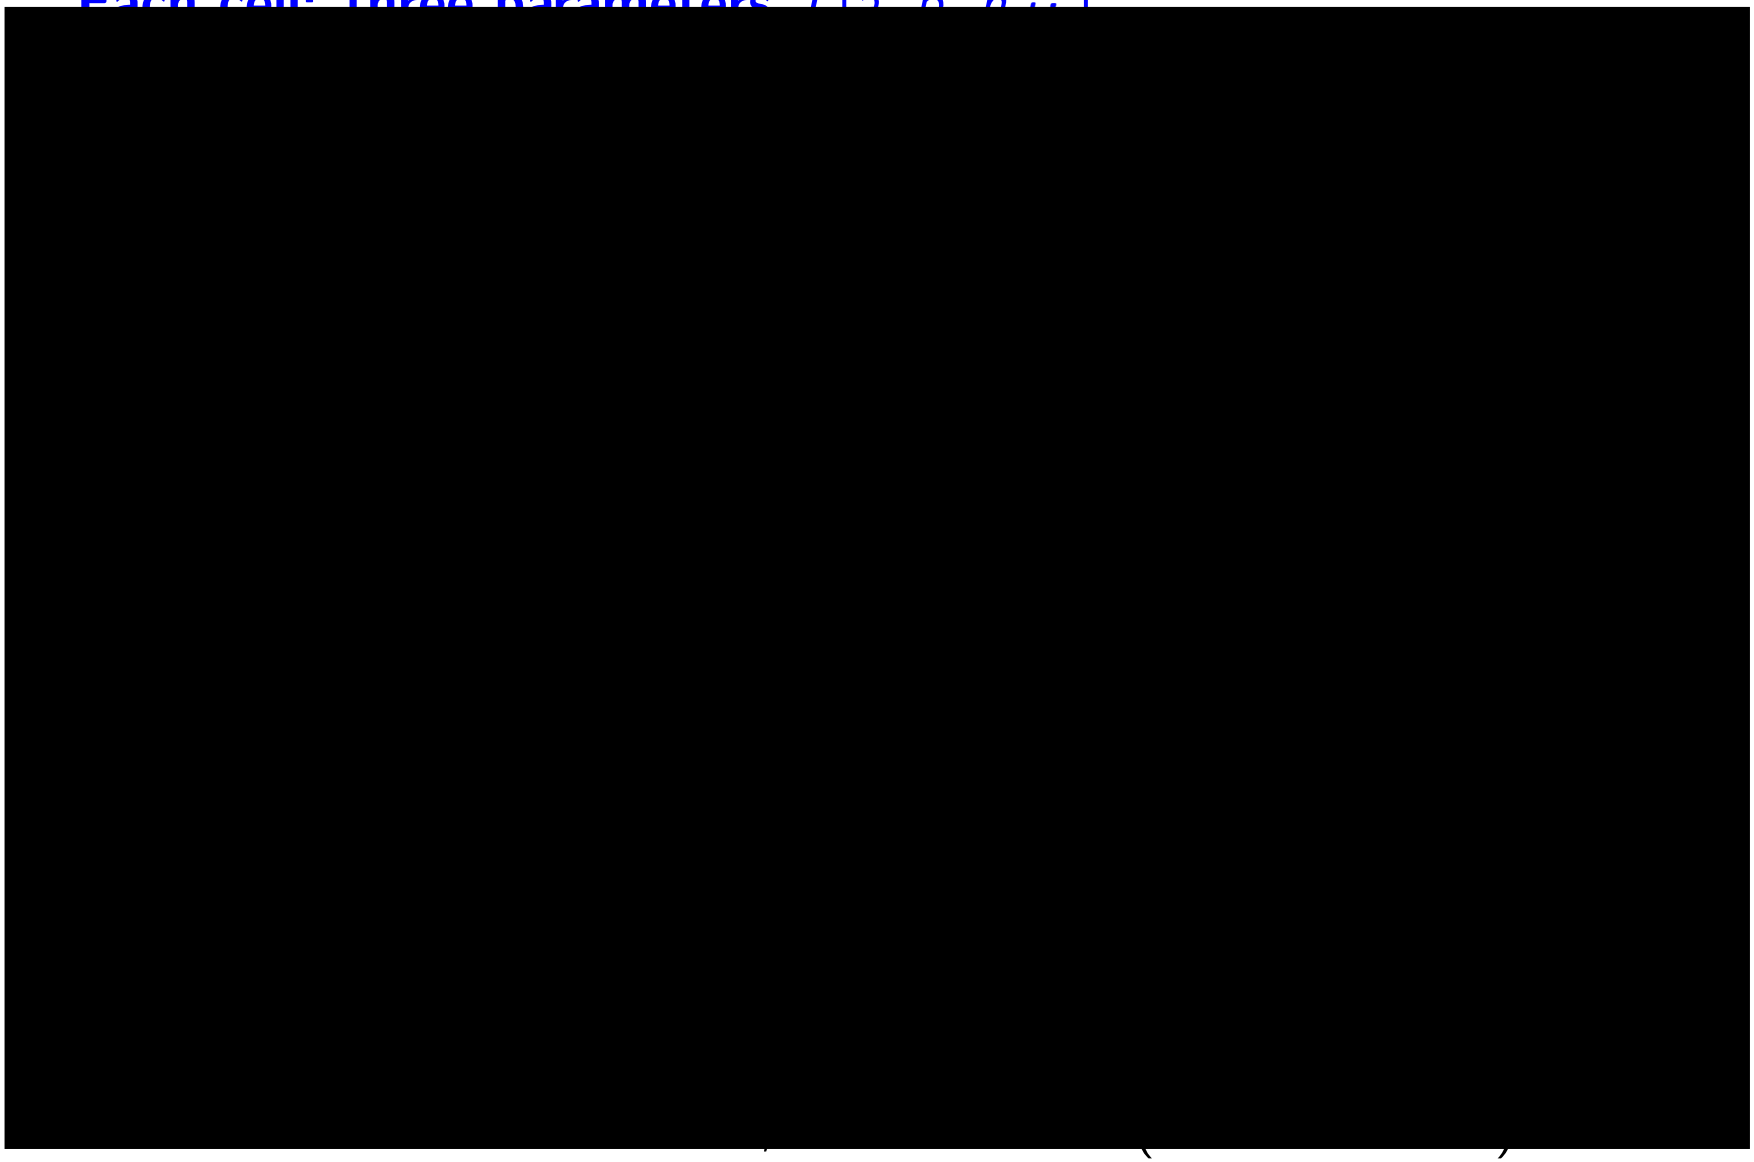

More dynamic programs

Theorem

DP *solves LTSP exactly in time* $O(n \cdot n_{req}^3)$.

LogDP(λ): **DP** restricted to detours spanning $\lambda \log n_{req}$ requested files

Reduced complexity in $O(\lambda^2 \cdot n_{req} \cdot n \cdot \log^2(n_{req}))$, *tested with* $\lambda \in \{1, 5\}$

SimpleDP: **DP** forbidding intertwined (*i.e.*, overlapping) detours

Reduced complexity in $O(n \cdot n_{req}^2)$, *better theoretical guarantees*

Simulations: overview

Dataset: 2 weeks at CC-IN2P3

- 169 tapes, $>$ 3M files
- focus on reading operations
- filtering steps, data processing (e.g., merge reads on aggregates)
- median data: 150 files requested, 3k requests, 50% file size variation

Code + dataset (with statistical descriptions) available online

Experimental methodology

- vary U
- median timings (seconds, on a compiled Python program):

FGS	NFGS	LogDP(1)	SimpleDP	LogDP(5)	DP
< 0.1	1	2	3	7	30

Simulation results, $U = 0$

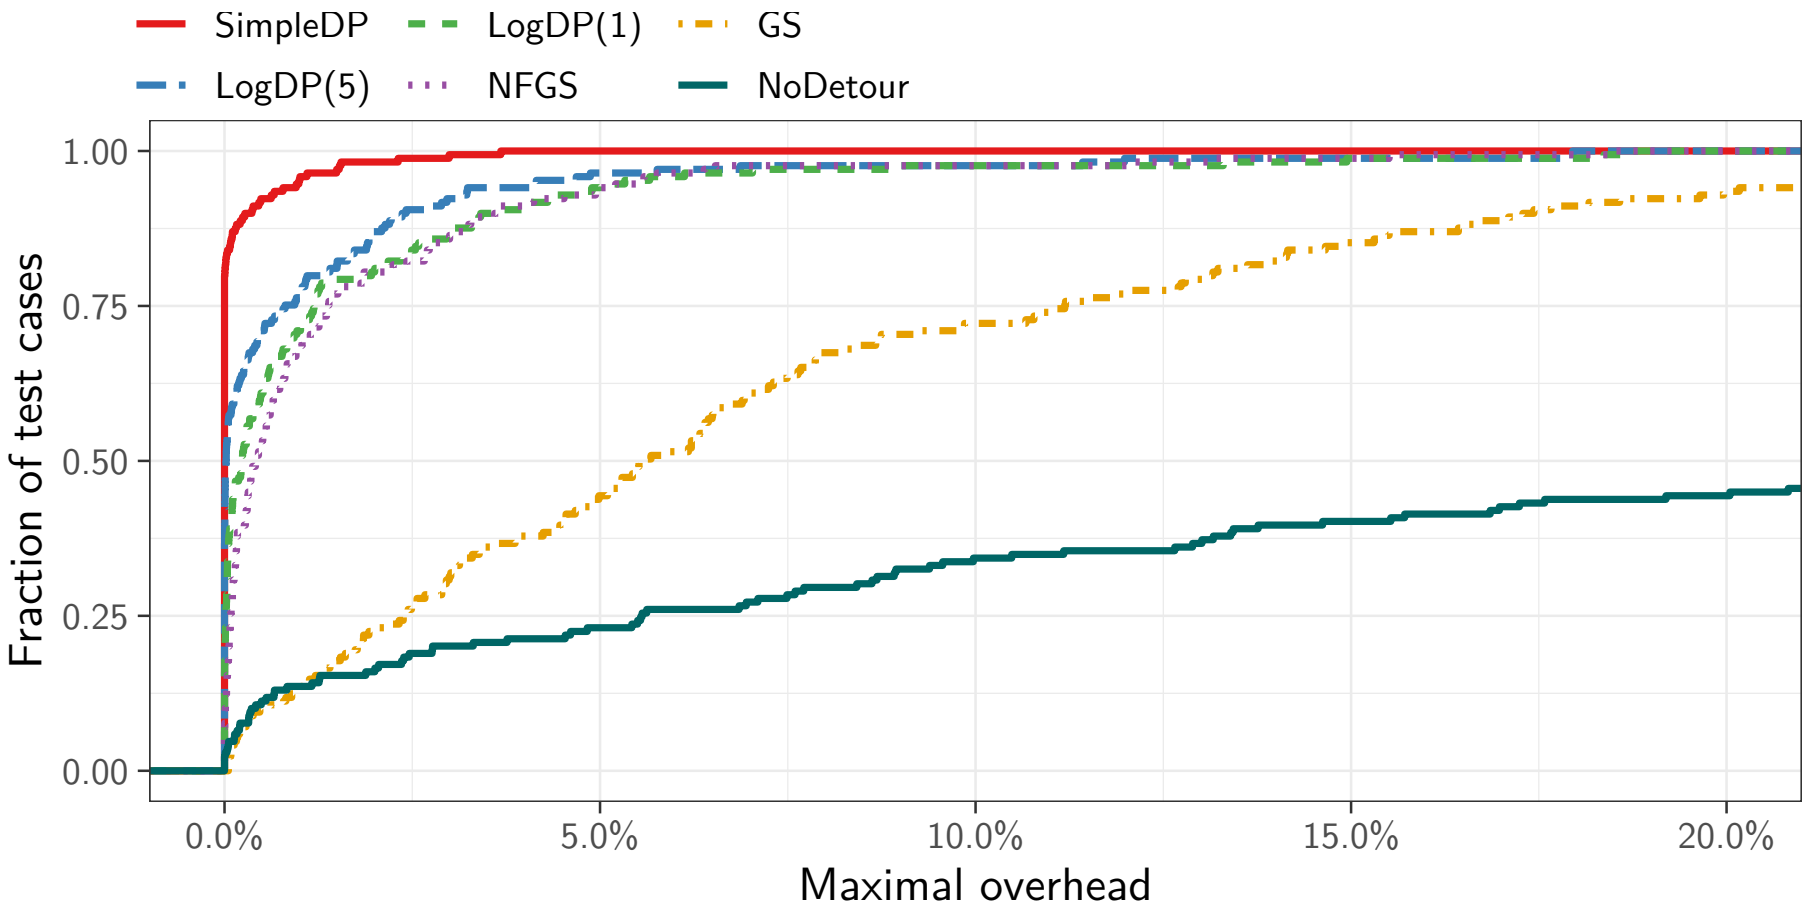

Performance profile: best is top-left (most instances with low overhead vs OPT)

Simulation results, $U =$ file

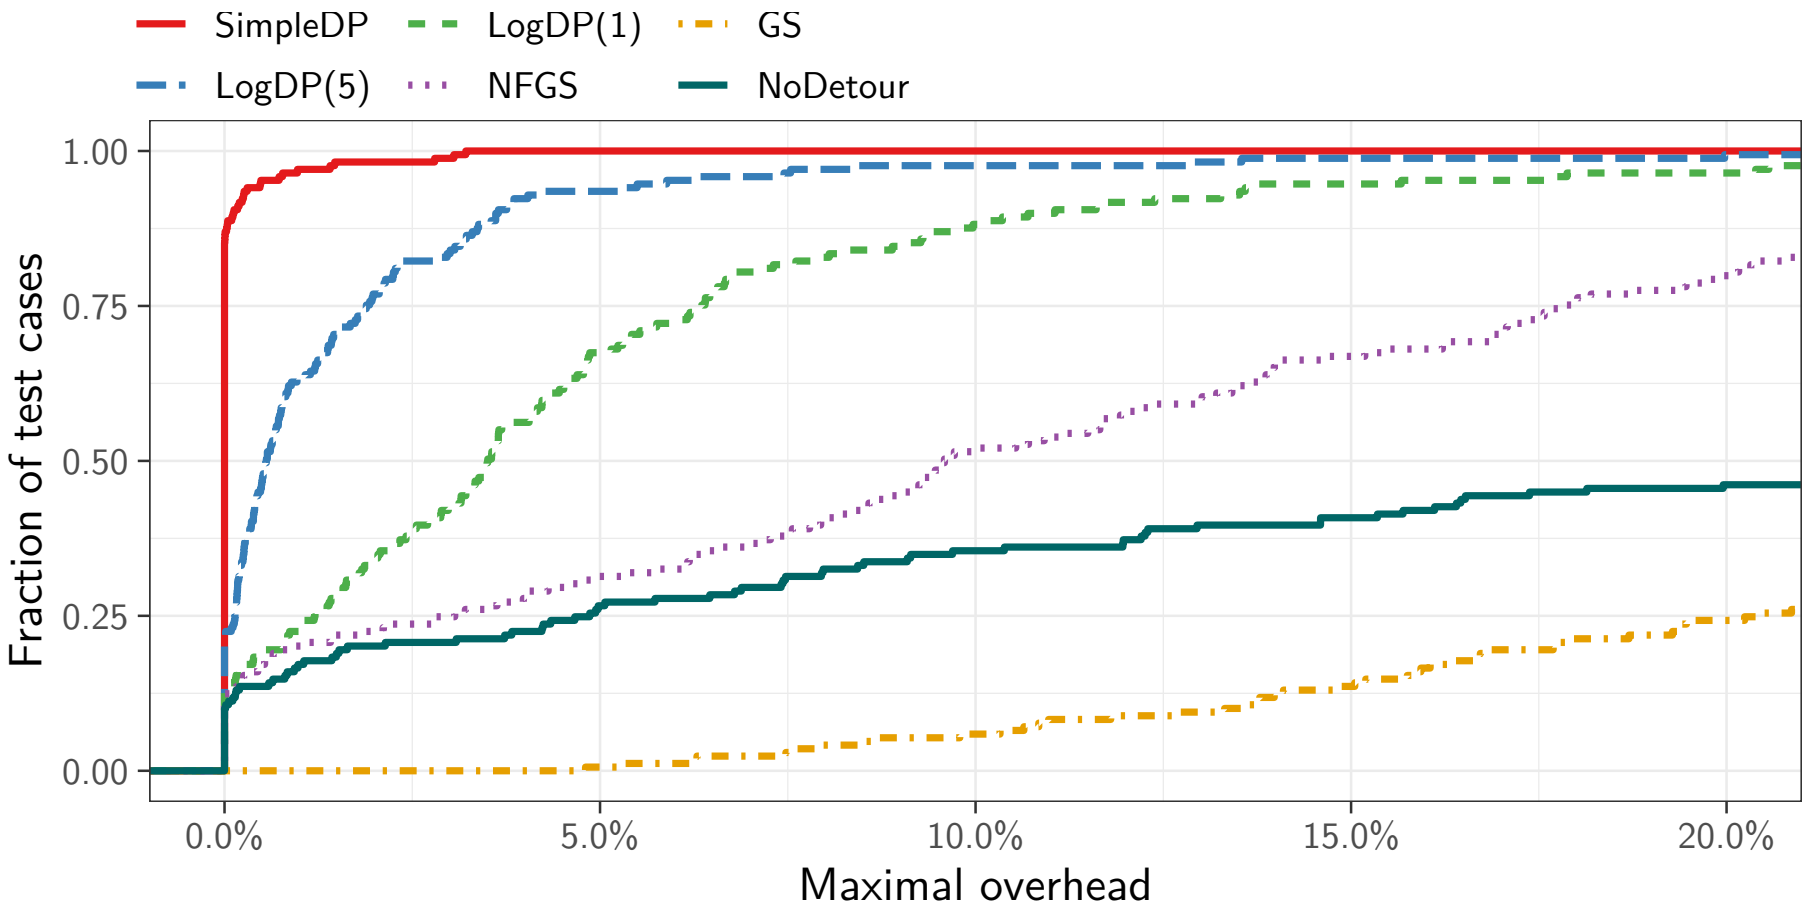

Performance profile: best is top-left (most instances with low overhead vs OPT)

Conclusion

https://commons.wikimedia.org/wiki/File:LTO2-cart-wo-top-shell.jpg

General: tapes are past & future

- tapes stay primordial in some fields but neglected by CS research
- fundamental problems are still open

On LTSP

- high-multiplicity variant remains open
- huge gap between theoretically studied models and practical heuristics

Perspectives on other tape-related topics

- multi-tape requests: optimize waiting queues
- optimize tape / disk storage ratio