Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

When does a lower bound solve the distributed assembly flowshop? Optimality certificates, rates, a sharp threshold and 22 new optimality proofs

Felipe Provezano Coutinho

fpcoutinho@alumni.usp.br

Independent Researcher

Research Article

Keywords:

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11639999/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

When does a lower bound solve the distributed assembly flowshop? Optimality certificates, rates, a sharp threshold and 22 new optimality proofs

Felipe Provezano Coutinho[1*]

[1*]Independent Researcher, São Paulo, Brazil.

Corresponding author(s). E-mail(s): fpcoutinho@alumni.usp.br;

Abstract

In the distributed assembly permutation flowshop, several identical flowshop factories make the jobs of a set of products, and a single machine then assembles each product. We explain when and why the optimal makespan equals a lower bound, and how fast failure becomes unlikely on random instances. The answer depends on the load ratio between assembly work and factory work. If the assembly machine is the bottleneck, the weakest assembly bound is exact with probability tending to one, and the failure probability is exponentially small when the processing-time distribution has an atom at its minimum. If the factories are the bottleneck, every product-decomposition bound fails by a linear amount, and the right bound is a new covering end-chain bound. For integer processing times and at least two factories we prove a sharp threshold: the failure probability is exponentially small if as many products of minimal assembly time as there are factories fit into one minimal processing time, and of exact order $t^{-1/2}$ for t products otherwise, because machine-load ties then force idle assembly time. Without assembly, we prove that the factory bound of the distributed permutation flowshop is exact up to an exponentially small probability. The proofs yield polynomial optimality certificates, and two preregistered simulations confirm the main predictions. On the benchmark of Hatami et al., the certificates explain which instances are closed by bounds, prove optimality for 22 further large instances (two with new best values), and cut search time 5.5-fold when used as a stopping rule.

Keywords: Distributed assembly flowshop, Lower bounds, Optimality certificates, Probabilistic analysis, Benchmark instances, Vector summation

1 Introduction

In the distributed assembly permutation flowshop scheduling problem (DAPFSP), several identical factories, each a permutation flowshop, produce the components (jobs) of a set of products, and a single assembly machine then assembles each product once all its jobs are finished (Hatami et al. 2013). The objective is the makespan. The problem combines three classical difficulties: assigning jobs to factories, sequencing within each factory, and sequencing the assembly operations. It has become a standard test bed for metaheuristics, and since its introduction more than a dozen algorithms have been proposed for it and evaluated on the benchmark of Hatami et al. (2013) (e.g. Wang and Wang 2016; Lin and Zhang 2016; Lin et al. 2017; Pan et al. 2019; Ferone et al. 2020; Zhang et al. 2021; Song et al. 2023).

A recent benchmark study (Coutinho, unpublished manuscript) found that this benchmark is *almost closed*: simple lower bounds already equal the best known makespan on most instances, and an exact model solves many of the rest. But the closures are not evenly spread. On the large set they concentrate on instances with 100 or 200 jobs and on instances whose "start" product has few jobs, while instances with 500 jobs stay open far more often. This pattern raises a question that is not about algorithms at all: *when, and why, does the optimal makespan of a DAPFSP instance equal a simple lower bound?* If the answer is known, the bound becomes an optimality certificate, a stopping rule for search, and a way to tell which instances of a benchmark can discriminate between algorithms.

For the ordinary permutation flowshop the analogous question goes back to Taillard (1993), who observed that on random instances the machine-based lower bound is often attained and conjectured that this happens with probability tending to one as the number of jobs grows. A companion paper (Coutinho, manuscript under review) proves this conjecture and gives the rates. The present paper does the same for a richer problem, in which two new forces appear: the assembly machine, which can be the bottleneck, and the distribution of the work over several factories, which turns load *ties* between machines into a genuine obstruction.

1.1 Main results

The results are organised by the *load ratio* of a random instance model: the expected assembly work divided by the expected work per factory machine. When the ratio exceeds one the assembly machine is the bottleneck; when it is below one the factories are. For every value of the load ratio we determine whether the optimum equals a lower bound and how fast the probability of failure vanishes as the number of products grows. Table 4 in Section 6 summarises the answers.

(1) *A sharp threshold on the factory-dominant side.* For integer processing times, at least two factories and a load ratio below one, the answer depends only on whether as many products of minimal assembly time as there are factories can be assembled within one minimal processing time. If they can, a new lower bound, the covering end-chain bound, equals the optimum except on an event of exponentially small probability, whatever the machine loads, exact ties included (Theorem 21).

If they cannot, two factory machines with equal loads force the optimum above that bound and above the best combined bound, and the failure probability is exactly of the order of the inverse square root of the number of products (Proposition 17, Theorem 19). The boundary case falls on the exponential side (Corollary 22). A preregistered simulation with 1.72 million instances recovers the predicted exponent: -0.494 ± 0.016 against $-1/2$.

(2) *Exactness on the assembly-dominant side.* If the product that can finish first has no more jobs than there are factories, a constant number of fast "ramp" products exists, and the assembly load exceeds the busiest factory machine's share, then the optimum equals the weakest assembly bound, and a polynomial algorithm certifies it (Theorem 6). For random instances with a load ratio above one this happens with probability tending to one, exponentially fast when the processing-time distribution has an atom at its minimum (Theorem 9).

(3) *Taillard's law for the distributed permutation flowshop.* Without the assembly stage, the problem is the distributed permutation flowshop of Naderi and Ruiz (2010). For integer processing times we prove that the factory lower bound fails with a probability that decays exponentially in the number of jobs, with matching upper and lower exponential bounds (Theorem 16). To our knowledge this is the first result of this kind for any distributed flowshop.

(4) *What fails, and how fast.* When the factories are the bottleneck, every bound that decomposes the problem by products (and so ignores the factory side) is off by an amount linear in the number of products; at a load ratio of exactly one its expected gap grows like the square root (Proposition 11). The combined bound stays within a constant of the optimum at every load ratio (Theorem 10). If every product has at least as many jobs as there are factories, a start obstruction keeps the weakest assembly bound inexact with probability bounded away from zero (Proposition 12). For processing times with a density, the failure probability of that bound decays polynomially, with an exponent set by the number of machines (Proposition 13, Theorem 14).

(5) *Two new lower bounds and two certificates.* The covering end-chain bound (Proposition 3) and the two-machine end bound (Proposition 4) are valid for every schedule and computable in polynomial time. The start–ramp–Steinitz schedule (Theorem 6) and its mirror image, in which fast "down-ramps" give the assembly machine credit at the end (Theorem 8), certify optimality whenever they meet a bound.

(6) *Consequences for the benchmark.* The certificates explain, by construction, the split of the Hatami et al. (2013) benchmark into closed and open instances by product size. Combined with exact single-product releases and short local search, they give 22 new proved optima on the large set (two with new best values), and a lower-bound stopping rule that cuts the search budget 5.5-fold (Section 4).

The common thread is a *credit* mechanism: the assembly machine is never idle if, at every moment, the work released to it so far covers the time elapsed. On the assembly side, credit comes from the products released early; on the factory side, from fast products released at the very end. Hariri and Potts describe this intuition informally

for the two-stage assembly problem (Hariri and Potts 1997, p. 553); Theorems 6, 8 and 21 turn it into exact inequalities. The tools are the Steinitz lemma, in the form developed for flowshops by Sevast'janov (Sevast'janov 1994a, 1995), a new exact lattice partition of the work over factories (Lemma 15), local limit theorems for the machine loads, and elementary concentration bounds.

1.2 Related work

The DAPFSP and the distributed permutation flowshop.

Hatami et al. (2013) introduced the DAPFSP, its mixed-integer model and its benchmark of 900 small and 810 large instances. Apart from that model, the work since then has been heuristic: estimation-of-distribution and memetic algorithms (Wang and Wang 2016; Zhang et al. 2021; Li et al. 2023), genetic algorithms (Li et al. 2015; Zhang et al. 2020), biogeography-based optimisation (Lin and Zhang 2016), hyper-heuristics (Lin et al. 2017; Song et al. 2023), iterated local search and greedy variants (Pan et al. 2019; Ferone et al. 2020; Ying and Lin 2023; Yu et al. 2023), variable-neighbourhood search (Liu et al. 2016) and population methods (Hamzadayı et al. 2021). These papers report makespans, not optimality proofs. For the distributed permutation flowshop, Gogos (2023) solves most of the Taillard-based instances of Naderi and Ruiz (2010) with constraint programming, using a head-plus-average-load-plus-tail lower bound; it makes no asymptotic claim. Bai et al. (2025) study a distributed permutation flowshop with heterogeneous factories and release dates and show that a dispatching rule is asymptotically optimal in the ratio sense (the makespan per job converges to the optimal one), with no statement that the optimum equals a bound and no rate.

Assembly flowshops.

The two-stage assembly problem was introduced and analysed by Lee et al. (1993) and Potts et al. (1995). Potts et al. (1995) prove an additive guarantee by compact vector summation, and Hariri and Potts (1997) give a branch-and-bound method with Johnson-type bounds. Koulamas and Kyparisis (2001) study the three-stage assembly flowshop, with ratio heuristics and an additive bound, and Koulamas and Kyparisis (2000) give asymptotically optimal linear-time algorithms for flexible flowshops. Theorem 10 extends this additive line to several factories and multi-job products. The polynomial cases of Lee et al. (1993) rely on job-wise dominance between assembly and component times; we show in Section 3 that job-wise dominance does not make the optimum equal the weakest assembly bound.

Probabilistic analysis of flowshops.

Ramudhin et al. (1996) show that, for the balanced two-machine flowshop with independent uniform times, Johnson's makespan equals a simple bound with probability tending to one. Kaminsky and Simchi-Levi (1998, 2001) and Xia et al. (2000) prove asymptotic optimality of simple rules for completion-time objectives in flowshops, in the ratio sense. Koryakin (2005) shows that every permutation schedule is asymptotically optimal in ratio for stochastic flow shops and single-factory assembly lines; the additive gap in his bounds grows like the square root of the number of jobs up

to logarithms, and he lists the rate of convergence as an open question. Our results answer that question for the DAPFSP and the distributed permutation flowshop: we study the event that the optimum equals a lower bound, not a ratio, and determine its probability up to constants in the exponent or the power. For the single-factory permutation flowshop, the corresponding exactness theorem and rates are those of Coutinho (manuscript under review), whose staircase-block and absorption lemmas we reuse inside each factory.

Vector summation.

The Steinitz lemma and its algorithmic versions (Grinberg and Sevast'yanov 1980; Sevast'janov 1994a; Sevast'yanov 1994b; Sevast'janov 1995) are the classical route to additive guarantees in flow shops and open shops. We use them in three new ways: to order whole products in time with an assembly coordinate, to split products over factories with prescribed loads (together with exact unit exchanges), and to keep every factory's assembly releases close to a straight line.

1.3 Organisation

Section 2 defines the problem, the lower bounds and the random model. Section 3 gives the deterministic certificates, and Section 4 applies them to the benchmark of Hatami et al. (2013): what they explain, the new optimality proofs, and what the bounds enable in practice. Section 5 treats random instances on the assembly side and the obstructions for general distributions. Section 6 treats integer processing times: the distributed permutation flowshop, the tie obstruction, the factory-side certificate, the rates at every load ratio and the threshold. Section 7 discusses open problems. All proofs are in the appendices.

2 Model, lower bounds and random instances

The problem.

There are n jobs, partitioned into t products $N_1, \dots, N_t$ with sizes $s_h = |N_h|$, and F identical factories, each a flowshop with machines $1, \dots, m$. Job j has processing time $p_{ij} \geq 0$ on machine i. A schedule assigns every job to a factory and, in each factory, processes its jobs in one permutation on all machines (a permutation flowshop). Following Hatami et al. (2013), the jobs of one product may be assigned to different factories. Product h is *released* to the assembly stage at $c_h = \max_{j \in N_h} C_j$, the latest completion time of its jobs on machine m. A single assembly machine then processes the products without preemption; product h needs $A_h \geq 0$ time units and cannot start before c_h. The makespan is the completion time of the last assembly, and OPT is its minimum over all schedules. With $A \equiv 0$ the problem is the distributed permutation flowshop (DPFSP) of Naderi and Ruiz (2010).

Notation.

$P_j = \sum_i p_{ij}$ is the total processing time of job j, $L_i = \sum_j p_{ij}$ the load of machine i, $x_{h,i} = \sum_{j \in N_h} p_{ij}$ the load of product h on machine i, and $A_{\mathrm{tot}} = \sum_h A_h$. The entries

lie in $[a, b]$, where a and b are their global minimum and maximum; $\bar{s} = \max_h s_h$, $\bar{A}$ bounds the assembly times and $\bar{x} = \max_{h,i} x_{h,i}$. For machine k, $\mathrm{hd}_k = \min_j \sum_{i<k} p_{ij}$ and $\mathrm{tl}_k = \min_j \sum_{i>k} p_{ij}$ are the smallest head and tail. $R_h = \max_{j \in N_h} P_j$ is the earliest time at which product h could be released, $T_0 = \min_h R_h$, and r_h^* is the optimal makespan of the DPFSP on N_h alone, with all F factories.

2.1 Lower bounds

The assembly stage is a single machine with release dates, so its optimal schedule is classical.

Lemma 1 (earliest release date) *For release times c and processing times A on one machine, the minimum makespan is*

$$\mathrm{ERD}(c) = \max_h \Big[c_h + \sum_{g:\, c_g \geq c_h} A_g \Big],$$

attained by every order of non-decreasing c. ERD *is non-decreasing in every c_h, and for a fixed factory schedule the best assembly schedule has makespan* $\mathrm{ERD}(c)$.

Since $c_h \geq r_h^* \geq R_h$ in every schedule, every vector $\underline{r}$ with $R_h \leq \underline{r}_h \leq r_h^*$ gives a valid bound $\mathrm{ERD}(\underline{r})$. The weakest of them is

$$\mathrm{LB}_0 = T_0 + A_{\mathrm{tot}} = \min_h R_h + \sum_h A_h,$$

and the strongest is the *product-decomposition bound* $\mathrm{LB}^* = \mathrm{ERD}(r^*)$. We call a lower bound X *admissible* if $\mathrm{LB}_0 \leq X \leq \mathrm{OPT}$. LB_0 is the first term of a branch-and-bound bound of Lee et al. (1993, eq. (8)) for the three-machine assembly flowshop; we use it for multi-job products spread over F factories.

Lemma 2 *If $s_h \leq F$, then $r_h^* = R_h = \max_{j \in N_h} P_j$.*

So the release of a product with at most F jobs is known exactly, and LB_0 is attained by the start product alone if $s_{h_1} \leq F$ for the product h_1 with $R_{h_1} = T_0$. The benchmark study (Coutinho, unpublished manuscript) uses two further bounds. Its closed-form bound $\mathrm{LB}_{\mathrm{cf}}$ is the maximum of a factory-load term, two assembly terms and a Johnson-type fluid term; it satisfies $\mathrm{LB}_{\mathrm{cf}} \geq \max(\mathrm{LB}_0, A_{\mathrm{tot}}, \max_k L_k/F)$. Its product-decomposition bound DB is $\mathrm{ERD}(\underline{r})$ with exact or CP-bounded single-product releases $\underline{r}_h \leq r_h^*$.

The bounds above look at products one at a time and ignore how the work is shared among factories. The next bound looks at the factory side. For $r \geq 1$ let

$$W(r) = \min \Big\{ \sum_{g \in \mathcal{C}} A_g : \ \mathcal{C} \text{ a set of products with } \sum_{g \in \mathcal{C}} s_g \geq r \Big\},$$

the least assembly work of products that together have at least r jobs ($W(r) = \infty$ if $r > n$). W is a covering knapsack, computable in $O(tF)$ time for $r \le F$.

Proposition 3 (covering end-chain bound)

$$\mathrm{LB}_{\mathrm{end}} = \max_{k} \min_{1 \le F' \le F} \Big\lceil \mathrm{hd}_k + \mathrm{tl}_k + \frac{L_k + \sum_{r=1}^{F'} W(r)}{F'} \Big\rceil \le \mathrm{OPT},$$

also over non-permutation factory schedules. When all data are integers, $\lceil \mathrm{LB}_{\mathrm{end}} \rceil$ *is valid as well, and we then write* $\mathrm{LB}_{\mathrm{end}}$ *for the rounded value.*

The idea is that the last job of each used factory leaves machine k no earlier than hd_k plus that factory's load on k, then needs at least tl_k more, and then its product must still be assembled. The products that end the F' used factories cover at least F' distinct jobs, so the assembly work after the r-th latest factory end is at least $W(r)$. Averaging over the factory ends gives the bound. The core term with $W \equiv 0$ and $F' = F$ is the DPFSP bound of Gogos (2023, Alg. 1); the covering term, the averaging over the factory ends and the minimum over the number of used factories are new.

Proposition 4 (two-machine end bound) *Let all data be integers. For machines* $k < l$ *let* $\mu_{kl} = \mathrm{hd}_k + (m - k - 1)a + \mathrm{hd}_l + \mathrm{tl}_l$, $\nu_k = \mathrm{hd}_k + \mathrm{tl}_k$ *and* $\pi_{kl} = \max_j (p_{kj} + p_{lj})$. *For* $1 \le F' \le \min(F, n)$, $0 \le F_2 \le \min(F', n - F')$ *and* $N = F' + F_2$, *put*

$$V_{kl}(F', F_2) = \frac{1}{N} \Big[F_2 \mu_{kl} + (F' - F_2) \frac{\nu_k + \nu_l - \pi_{kl}}{2} + L_k + L_l + \sum_{r=1}^{N} W(r) \Big].$$

Then $\mathrm{LB}_{\mathrm{tie}} = \lceil \max_{k<l} \min_{F', F_2} V_{kl}(F', F_2) \rceil \le \mathrm{OPT}$.

$\mathrm{LB}_{\mathrm{tie}}$ charges the last *two* completions of every factory with at least two jobs: machine k forces the last one and machine l the one before it. It is never needed when one machine dominates the loads, but it is exactly the bound that detects the tie obstruction of Section 6. We write $\mathrm{LB}_{\mathrm{comb}} = \max(\mathrm{LB}_{\mathrm{cf}}, \mathrm{LB}_{\mathrm{end}})$ for the combined bound; it is admissible.

2.2 Random instances

In the random model, products are i.i.d. A product has S jobs with $S \sim \nu$ on $\{1, \dots, \bar{s}\}$; given S, its jobs have i.i.d. entries with law μ on $[a, b]$, $a = \min \operatorname{supp} \mu < b = \max \operatorname{supp} \mu$; and its assembly time is $A \sim \kappa_S$ on $[0, \bar{A}]$, independent of the entries given S. The number of products t grows, and $n = \sum_h s_h \in [t, \bar{s}t]$. The *load ratio*

$$\beta = \frac{F\, \mathbb{E}[A]}{\mathbb{E}[S]\, \mathbb{E}p}$$

compares the expected assembly work with the expected work per factory machine. When $\beta > 1$ the assembly machine is the bottleneck (*assembly-dominant* instances); when $\beta < 1$ the factories are (*factory-dominant* instances). The benchmark of Hatami

et al. (2013) draws A_h uniformly from $[s_h, 99s_h]$ and entries from $U\{1..99\}$, so $\beta = F$ there. The model assumptions used for $\beta > 1$ are collected as model M (Assumption 32 in the appendices):

(M1) $F \geq 2$ and $\nu(1) > 0$: single-job products occur;
(M2) some α with $(F-1)\alpha > a$ has $\kappa_1([\alpha, \bar{A}]) > 0$: single-job products can have assembly times large enough that $F-1$ of them outlast a minimal processing time;
(M3) either $\mu(\{a\}) > 0$ (the *atom case*) or $\kappa_1(\{0\}) = 0$;
(M4) $\beta > 1$.

3 Deterministic exactness certificates

Both certificates in this section build an explicit schedule and compare it with a lower bound. They rest on one tool.

Lemma 5 (Steinitz lemma, polynomial form) *Let $v_1, \ldots, v_N \in \mathbb{R}^d$ with $\|v_l\|_\infty \leq 1$ and $\sum_l v_l = 0$. There is a permutation of $\{1, \ldots, N\}$ all of whose prefix sums have sup-norm at most d, and it can be found in $O(d^2N^2)$ time.*

This is a special case of Sevast'janov (1995, Lemma 1); the bound d for every norm goes back to Sevast'janov's work with the proof method of Grinberg and Sevast'yanov (1980), and Sevast'janov (1994a) surveys its use in scheduling. In our setting the vectors are centred product loads, with one coordinate per machine and one for the assembly time, so an order with small prefix sums releases products to the assembly machine at an almost constant rate.

3.1 The assembly side: when the weakest bound is exact

The schedule behind Theorem 6 has three parts.

- A *start set* H_0 of at most F jobs is placed at the first position of distinct factories, so its products are released at exactly R_h. It contains the product h_1 with $R_{h_1} = T_0$.
- A constant number of *ramp* levels follows. Each level holds F single-job products, one per factory, with small entries and assembly times $A \geq \alpha$ where $F\alpha$ exceeds a ramp's processing time. Each level therefore releases more assembly work than the time it takes, and the surplus builds a credit that absorbs later fluctuations.
- The other products (the *main set* H') are ordered by Lemma 5 on their centred machine loads and assembly times, and the order is cut into F consecutive blocks, one per factory.

The assembly machine is never idle after T_0 if, for every product h, the assembly work released before h covers the time from T_0 to h's release. This is a queue condition, and Lemma 1 turns it into the equality $\mathrm{ERD}(c) = T_0 + A_{\mathrm{tot}} = \mathrm{LB}_0$. The hypotheses

(H1)–(H5) of Assumption 28 in the appendices make the queue condition hold for the start set, for every ramp level, and for every main level up to a deviation that the Steinitz order keeps below a constant. In words:

(H1) the start set has at most F jobs and contains h_1;
(H2) within the start set, earlier releases bring enough assembly work to cover later ones;
(H3) the ramps are single jobs with entries at most $a+\delta$, and $F\alpha \geq a+\delta$;
(H4) the *load margin* $F\sum_{h\in H'} A_h \geq \max_i \sum_{h\in H'} x_{h,i}$: on the main set, the assembly machine is at least as loaded as the busiest factory machine's share;
(H5) the *credit* built by the start set and the ramps is at least the Steinitz deviation constant $m(4m+5)\bar{x}' + 2F(m+1)\bar{A}'$, where $\bar{x}'$ and $\bar{A}'$ are the largest product load and assembly time in H'.

Theorem 6 (exactness) *Under (H1)–(H5),* $\mathrm{OPT} = \mathrm{LB}_0 = X$ *for every admissible bound* X*, also over non-permutation factory schedules. The start–ramp–Steinitz schedule attains* LB_0 *and is computed in polynomial time.*

Proof idea The Steinitz order bounds every product's release by a line of slope $\max_i \bar{y}_i$ per level plus a constant K_1, where $\bar{y}_i$ is the mean product load on machine i; the ramps and the start set are released at explicit times. The queue condition at a main level q needs the assembly work released at earlier levels, at least $F[(q-1)\bar{a} - E_A]$, to cover $(q-1)\max_i \bar{y}_i + K_1$ plus the start and ramp phases. (H4) says $F\bar{a} \geq \max_i \bar{y}_i$, and (H5) covers the constants. The full proof is in Section A.2. □

When the start product has more than F jobs, Lemma 2 fails, but the same argument works with an optimal sub-schedule of that product.

Corollary 7 (large start product) *Replace* H_0 *by the single product* h_1*, scheduled by an optimal DPFSP sub-schedule of* N_{h_1}*, and* T_0 *by* $r^*_{h_1}$*. If (H3)–(H5) hold with* $C_0 = A_{h_1}$ *and* $T_1 = T_0 = r^*_{h_1}$*, then* $\mathrm{OPT} = r^*_{h_1} + A_{\mathrm{tot}}$*. This value equals* $\mathrm{ERD}(\underline{r})$ *for every vector* $\underline{r}$ *with* $R_h \leq \underline{r}_h \leq r^*_h$*,* $\underline{r}_{h_1} = r^*_{h_1}$ *and* $\underline{r}_h \geq r^*_{h_1}$ *for all* h*.*

Theorem 6 uses an aggregate margin, not a job-wise one. Lee et al. (1993) solve the three-machine assembly flowshop in polynomial time when every job satisfies the same dominance condition between its assembly and component times. Their algorithms are sequencing rules; they do not make the optimum equal a bound. Indeed, two single-job products with $P = A = 1$ and $P = A = 10$ satisfy job-wise dominance, yet $\mathrm{LB}_0 = 12$ and $\mathrm{OPT} = 21$.

3.2 The factory side: a mirror-image certificate

When the factories are the bottleneck, the schedule must end well rather than start well. Fix a machine r, the bottleneck. An *up-ramp* is a single-job product with entries

a before machine r and b on r; a *down-ramp* has b on r and a after it. The ramps–Steinitz–end (RSE) schedule gives each factory f a sequence $U_f \cdot \mathrm{Mid}_f \cdot \mathrm{Dn}_f$: a block of up-ramps, a middle of whole products in a Steinitz order, and a block of $\beta_\downarrow$ down-ramps with assembly times $A \le \alpha^*$. The factories end at staggered times, $f \mapsto M - (m-1)a - S_{\rho(f)}$, where S_ρ is the assembly work of the end products of the first ρ factories in some ranking.

The down-ramps have processing time b on machine r and assembly time at most α^*, with $F\alpha^* < b$. So each level of down-ramps releases less assembly work ($F\alpha^*$) than the time it takes (b), and the gap $\gamma = b - F\alpha^*$ per level lets the assembly machine catch up with any backlog from the middle. The hypotheses (T1)–(T5) of Assumption 42 in the appendices state this precisely: (T1) the start- and end-block conditions of Coutinho (manuscript under review) hold in every factory, (T2) the staggered targets, (T3) an assembly margin along every suffix of the middle, (T4) $F\alpha^* < b$, and (T5) enough down-ramp levels to absorb the backlog.

Theorem 8 (factory-dominant exactness certificate) *Under (T1)–(T5), the RSE schedule has makespan at most M. If $M \le X$ for a valid lower bound X (for example $M = \mathrm{LB}_{\mathrm{end}}$ when the end products realise the covering function W), then* $\mathrm{OPT} = M$.

The analogy with Theorem 6 is one of mechanism, not a reduction. Reversing time within a factory turns a flowshop into a flowshop, but reversing the whole DAPFSP turns the assembly machine into a first-stage disassembly machine with fan-out precedence, which is not a DAPFSP. Theorem 8 is deterministic; Section 6 shows that its hypotheses hold with probability tending to one for lattice laws when $\beta < 1$, with a margin that does not grow with t.

4 The certificates on the benchmark of Hatami et al.

This section is exploratory: none of its measurements was preregistered. It asks what the theory explains, and what it enables, on the standard benchmark. We write BS for the companion benchmark study , whose closed-form bound $\mathrm{LB}_{\mathrm{cf}}$, product-decomposition bound DB, greedy heuristic, iterated-greedy (IG) schedules and exact CP-SAT model we reuse. All runs used one thread at low priority on a shared laptop, so wall times are indicative.

We ran a polynomial certifier on all 1,710 instances of Hatami et al. (2013). It takes the best of the constructive schedules behind Theorem 6 (start–ramp–Steinitz), its large-start-product corollary, Theorem 8 and Theorem 10, and compares that makespan with $\mathrm{LB}_{\mathrm{comb}} = \max(\mathrm{LB}_{\mathrm{cf}}, \mathrm{LB}_{\mathrm{end}})$, where $\mathrm{LB}_{\mathrm{cf}}$ is the closed-form bound of the companion benchmark study . An instance is self-certified when the two coincide. Table 2 gives the counts.

The certifier closes no instance beyond the BS closed set. It certifies fewer instances than the greedy heuristic of BS (123 and 326 against 257 and 546 on the small and large sets), but it is complementary to it: the union adds 25 small and 16 large instances. $\mathrm{LB}_{\mathrm{end}}$ never exceeds $\mathrm{LB}_{\mathrm{cf}}$, and the end-side certificate never applies. Both

facts are predicted, because the benchmark is assembly-dominant (assembly times are $U[|N_h|, 99|N_h|]$, so $\beta = F$).

The informative result is the split by product size. The start–ramp certificate needs a start product with at most F jobs (Lemma 2). It closes 176, 117 and 2 large instances at $n = 100$, 200 and 500. The large-start-product certificate closes 0, 3 and 28 at the same n. This reproduces, by construction, the benchmark split that BS observed.

On four open large instances the large-start-product schedule improves the BS reference upper bound: from 25,022 to 25,000, from 27,900 to 27,899, from 24,512 to 24,496 and from 29,778 to 29,776. None of the four closes against $\mathrm{LB}_{\mathrm{comb}}$; all four close against the BS product-decomposition bound (next paragraph). The whole certifier uses about 6.5 CPU-seconds for the 1,710 instances.

The open large instances.

BS closes 703 of the 810 large instances, using its closed-form bound and a product-decomposition bound DB (the ERD bound with exact or CP-bounded single-product releases). That leaves 107 open. We ran on all instances its closed-form bound leaves open, a superset of the 107. On each, we solved the start product exactly (CP-SAT, 10 s, optimal on 155 of 157) and rebuilt the large-start-product schedule from the exact sub-schedule. We then ran 30 s of local search (relocate/swap moves with ERD assembly) from that schedule and, separately, from the BS 60 s iterated-greedy schedule. Every schedule was re-evaluated with the BS evaluator. Every single-product release that a closure needs (33 sub-problems with 8–21 jobs) was re-solved with an independent CP-SAT interval model, and by complete enumeration where $|N_h| \leq 9$. All 33 agree with the BS values, so all 20 closures below stay proved.

On 20 of the 107 open instances our schedules reach DB, so these are new proved optima. On 7 of them the certificate schedule is optimal without any search. Local search from the certificate closes 8 more, and local search from the IG schedule closes 5. Nineteen of the 20 have $n = 500$, the large-start-product regime of the corollary, and the four upper bounds above are among them. With these, 723 of the 810 large instances are solved; two more closures from the stopping experiment below raise this to 725. In the terms of BS, a closure is *certified* when a schedule attains the bound and *published-value-only* when only a published makespan does. Of the 725, 713 are certified (later searches of BS certify 691 of its 703, and our runs certify all 22 new ones) and 12 are published-value-only, resting on values of the 2013 sheet or Ferone et al. (2020). The current count of BS, 735, has 36 published-value-only closures: these 12, 12 of our 22 (which our runs certify), and 12 instances where only a value from the later tables equals DB (see below). We do not count those last 12 as closed. Our upper bound is also below that of BS on 38 of the 107.

Table 1 lists the new optima with the best per-instance value reported for each instance. Besides the 2013 sheet and Ferone et al. (2020), four later studies report makespans for individual instances of Hatami's large set: Lin and Zhang (2016), Lin et al. (2017), Zhang et al. (2021) and Song et al. (2023). The first two label instances as in the 2013 results sheet, which is offset by one from the file names for replicates 2–10. Read with that offset, their "best known" columns match the 2013 sheet on 91/91

and 92/92 rows; read without it, 35 and 33 of their values would fall below a valid lower bound. Zhang et al. and Song et al. use the file names.

For 18 of the 20 instances the optimal value had already been reported by one of these studies; what is new there is the proof of optimality. Two values are new: I_500_5_8_30_7 (26,522) and I_500_5_8_30_9 (24,875). No reported value is below any of the 20 optima. Of the other 18 improvements over the BS upper bound, two are below every reported value (I_500_10_6_50_7, 25,585, and I_500_5_6_30_3, 31,888), four tie the best reported value and twelve are above it.

On 14 open instances a reported value equals DB. For two of them we now have a schedule of that value (see the stopping experiment below); for the other 12, if the reported values are feasible, they are optimal, which would give 737 solved instances. The papers publish no schedules, so we do not count them. Five of Zhang et al.'s 123 large values lie below an elementary bound (the ERD bound with releases $r_h = \max_{j \in N_h} \sum_k p_{jk}$, which is below DB); four are even below $\min_h r_h + \sum_h a_h$. We therefore treat reported values as reported, not verified.

Other studies are not comparable:

- Hamzadayı et al. (2021) report new bests on small instances only.
- Pan et al. (2019), Ying and Lin (2023) and Yu et al. (2023) solve different variants or objectives (factory-level assembly, flexible assembly with setups, total flowtime).

Wang and Wang (2016), Li et al. (2015) and Liu et al. (2016) report only aggregate deviations, so their values cannot be checked per instance. Zhang et al. (2020) also report only average deviations, and Li et al. (2023) use their own unreleased instances, so neither is comparable per instance. "Best known" above means the best published per-instance value.

What the bounds enable.

- *Early stopping.* We ran the BS iterated greedy on all 810 large instances with a 60 s cap. Each run stopped as soon as the makespan reached a given lower bound. A fixed 60 s budget costs 48,600 CPU-s. Stopping at the BS closed-form bound, as BS does, would cost 10,528 s (computed from the same runs). Stopping at the combined bound $\max(\mathrm{LB_{cf}}, \mathrm{DB}, \mathrm{LB_{end}})$ costs 8,806 s: 5.5 times less than the fixed cap, and 16 % less than the BS rule. The saving falls where the gap between the bounds lies. On the 50 instances that only DB closes, the cost drops from 3,000 to 1,373 s (34 reach the bound, median 2.5 s). On the 653 instances closed by the BS bound, 640 reach it with a median of 0.007 s. On the 107 open instances, few runs reach the bound: 5 of the 20 closed in Table 1 and 2 of the other 87. Those two are new closures: I_500_20_4_30_4 at 26,791 and I_500_5_4_40_7 at 28,538. Their DB uses CP-bounded releases, so we re-solved the three single-product releases they need with three independent models (a CP-SAT interval relaxation, a positional CP-SAT model, and HiGHS with presolve off), plus enumeration for the 7-job product. All three were confirmed. Both values had already been reported by Lin and Zhang (2016), Lin et al. (2017) and Zhang et al. (2021); what is new is the

Table 1 New proved optima on the Hatami et al. (2013) large set. OPT = DB: the optimum, equal to the product-decomposition bound DB of BS. 2013/Ferone: min of the 2013 results sheet and Ferone et al. (2020). Best later report: lowest value reported by Lin and Zhang (2016), LZ16, Lin et al. (2017), LWL17, Zhang et al. (2021), Z21 or Song et al. (2023), S23, mapped to file names; – if none of them lists the instance. $^{=}$: equal to our optimum, so only the optimality proof is new. BS UB: min of published values and a 60 s IG. ^{b}: DB is marked inexact because some non-binding products are not solved to optimality. Every release that the closure needs was re-solved to optimality with a second CP-SAT formulation (an interval model), and agrees. Route: certificate schedule alone (cert), local search from it (LS-cert) or from the IG schedule of BS (LS-IG).

Instance	2013/Ferone	Best later report	BS UB	OPT = DB	Route
I_200_5_4_50_4	9,390	9,389$^{=}$ (S23)	9,390	9,389	LS-IG
I_500_10_4_30_8	26,670	26,644$^{=}$ (LZ16, Z21)	26,653	26,644^{b}	LS-IG
I_500_10_4_50_1	28,887	28,837$^{=}$ (LZ16, LWL17, Z21)	28,843	28,837	LS-cert
I_500_10_6_30_10	27,992	27,984$^{=}$ (LZ16, S23)	27,992	27,984^{b}	LS-cert
I_500_10_6_30_2	25,022	24,997$^{=}$ (LZ16, Z21)	25,022	24,997^{b}	cert
I_500_10_6_30_3	28,516	28,502$^{=}$ (S23)	28,515	28,502^{b}	LS-cert
I_500_10_8_30_4	27,900	27,896$^{=}$ (S23)	27,900	27,896^{b}	cert
I_500_10_8_40_10	24,055	24,043$^{=}$ (S23)	24,050	24,043	LS-IG
I_500_20_4_40_3	24,585	24,583$^{=}$ (LZ16, LWL17, Z21)	24,585	24,583^{b}	LS-IG
I_500_20_4_40_8	21,272	21,267$^{=}$ (LZ16, LWL17, Z21)	21,272	21,267^{b}	LS-IG
I_500_20_6_40_2	25,589	25,578$^{=}$ (S23)	25,589	25,578	LS-cert
I_500_20_8_30_7	22,579	22,568$^{=}$ (LZ16, LWL17, Z21)	22,576	22,568^{b}	LS-cert
I_500_5_4_30_1	27,845	27,841$^{=}$ (LZ16, LWL17, Z21)	27,845	27,841^{b}	LS-cert
I_500_5_4_30_3	24,811	24,793$^{=}$ (S23)	24,803	24,793^{b}	LS-cert
I_500_5_4_50_2	27,780	27,776$^{=}$ (LZ16, S23)	27,780	27,776	LS-cert
I_500_5_4_50_6	24,512	24,491$^{=}$ (S23)	24,512	24,491	cert
I_500_5_8_30_3	29,877	29,875$^{=}$ (Z21)	29,877	29,875^{b}	cert
I_500_5_8_30_4	29,782	29,772$^{=}$ (LZ16, LWL17, Z21)	29,778	29,772^{b}	cert
I_500_5_8_30_7	26,526	–	26,526	26,522	cert
I_500_5_8_30_9	24,883	24,879 (S23)	24,883	24,875^{b}	cert

proof of optimality. Running the certificate first, and calling the search only when the certificate fails, costs 8,502 s.

- *The certificate as an exact algorithm.* On the instances it certifies, the Theorem 6 certifier is a polynomial exact solver: it needs a median of 3.4 ms (at most 12 ms) per instance. On a fixed random sample of 30 certified instances with $n \in \{100, 200\}$, the BS exact CP-SAT model (one worker) proved optimality on only 10 within 60 s, with a median of 27 s when it did. On 19 it found a feasible schedule only, and on 1 it found none.
- *An a priori rule.* Two conditions of Theorem 6 can be read off the data before any search: a start product with at most F jobs (Lemma 2; the BS class E) and an assembly load above the busiest factory machine's share, $\rho = \sum_h A_h / \left(\max_k \sum_j p_{jk}/F\right) > 1$. Neither threshold is fitted. We applied the rule "E and $\rho > 1 \Rightarrow$ the lower bound closes" to all 3,114 instances whose status is known (Hatami's sets, the BS load-controlled sets and our factory-dominant sets). It has precision 0.86, recall 0.52 and balanced accuracy 0.70, against a base rate of 0.58. Class E alone has precision 0.64 and balanced accuracy 0.55. The $\rho > 1$ condition cuts false positives from 558 to 152. On Hatami's large set ($\beta = F$) the ρ condition removes no class-E instance, and both rules have precision 0.96

Table 2 Self-certified instances on the Hatami et al. (2013) benchmark. BS closed: the reference upper bound of BS (proven optima on the small set; min of published values and a 60 s IG on the large set) equals $\mathrm{LB}_{\mathrm{cf}}$. BS greedy: the heuristic of BS reaches $\mathrm{LB}_{\mathrm{cf}}$. Theorem 6, Corollary 7: instances certified by each. Ours (any): by any of our certificates. Union: by ours or by BS greedy. Theorem 8 and the block schedule certify none.

Set	Class	N	BS closed	BS greedy	Thm 6	Cor. 7	ours (any)	union
small	$F = 2$	300	37	17	5	4	9	21
small	$F = 3$	300	134	90	21	23	44	99
small	$F = 4$	300	192	150	47	23	70	162
small	$n = 8$	180	111	91	42	4	46	93
small	$n = 12$	180	93	75	22	11	33	80
small	$n = 16$	180	73	45	7	14	21	51
small	$n = 20$	180	51	32	2	10	12	36
small	$n = 24$	180	35	14	0	11	11	22
small	all	900	363	257	73	50	123	282
large	$F = 4$	270	177	123	41	5	46	127
large	$F = 6$	270	223	196	101	7	108	200
large	$F = 8$	270	253	227	153	19	172	235
large	$n = 100$	270	262	243	176	0	176	243
large	$n = 200$	270	252	217	117	3	120	219
large	$n = 500$	270	139	86	2	28	30	100
large	all	810	653	546	295	31	326	562

and recall 0.76. On the 971 instances with $\beta \leq 1$ the rule is close to the majority classifier (balanced accuracy 0.51), consistent with Proposition 11: there the decomposition bounds are off by $\Theta(t)$, and the rare closures need the end-side bound.

Factory-dominant instances.

The Hatami benchmark has $\beta = F$, so neither $\mathrm{LB}_{\mathrm{end}}$ nor Theorem 8 can act there. We therefore generated factory-dominant instances with the BS load-controlled rule ($A_h \sim U[\ell|N_h|, u|N_h|]\cdot\beta/F$, otherwise as Hatami et al.). We used $\beta = 0.25$ on $U\{1..99\}$ and $\beta = 0.5$ on the atom-rich lattice $U\{1..9\}$, which favours the exact ramp classes of Theorem 8, and added the $\beta \in \{0.5, 1\}$ sets of BS (Table 3).

The result is negative and explanatory. The certifiers add no closure beyond the BS machinery. The hypotheses of Theorem 8 hold on none of the 936 random instances, so the theorem is a certificate for structured instances, consistent with our negative random-model evidence for probabilistic factory-dominant exactness, which remains open. $\mathrm{LB}_{\mathrm{end}}$ improves $\mathrm{LB}_{\mathrm{cf}}$ on 66 of the 216 large instances, but only by 0.04–1 %.

The gap in this regime is on the bound side. On the small sets our local search reaches the proven optimum on 454 of the 515 solved instances, while the bound sits 9–11 % below the optimum at $\beta \leq 0.5$. This is what Proposition 11 predicts for decomposition-type bounds when $\beta \leq 1$. On the large factory-dominant sets, 5 s of local search from our block schedules beats the BS upper bound on 213 of 216 instances.

Table 3 Factory-dominant instances ($\beta \le 1$). BS closed: the bound of BS equals its upper bound or the optimum. Solved: optimum known (exact CP-SAT model of BS; 300 s on the BS sets, 20 s on ours). Theorem 8: instances where its hypotheses hold. Ours closed: our schedule meets the bound. Gap: median relative gap from the bound to the best upper bound.

Set	N	BS closed	solved	$\mathrm{LB}_{\mathrm{end}} > \mathrm{LB}_{\mathrm{cf}}$	Thm 8	ours closed	LB→UB gap
small, $\beta = 0.25$	180	12	98	7	0	12	9.0 %
small, $\beta = 0.5$, $U\{1..9\}$	180	19	125	0	0	19	9.1 %
small, $\beta = 0.5$ (BS)	180	13	128	0	0	13	8.6 %
small, $\beta = 1$ (BS)	180	47	164	0	0	44	5.6 %
large, $\beta = 0.25$	54	0	0	17	0	0	12.8 %
large, $\beta = 0.5$, $U\{1..9\}$	54	0	0	7	0	0	18.5 %
large, $\beta = 0.5$ (BS)	54	0	0	29	0	0	13.8 %
large, $\beta = 1$ (BS)	54	1	0	6	0	0	10.2 %

5 Random instances: the assembly side and general laws

5.1 Exactness for $\beta > 1$

When $\beta > 1$, the margin (H4) holds with a margin linear in t, and the start set and the ramps of Theorem 6 are found among the many single-job products guaranteed by (M1)–(M2).

Theorem 9 ($\beta > 1$) *Under model M,* $\Pr[\mathrm{OPT} = X$ *for every admissible* $X] \to 1$ *as* $t \to \infty$*, and the start–ramp–Steinitz certificate exists with the same probability. In the atom case* $\mu(\{a\}) > 0$*, the failure probability is at most*

$$\Pr\big(\mathrm{Bin}(t, q_*) < F\big) + \Pr\big(\mathrm{Bin}(t, q_R) < F(R_\ell + 1)\big) + m \exp\Big(-\frac{2(t\mu_\xi - c_0)^2}{t(F\bar{A} + \bar{s}b)^2}\Big) = e^{-\Omega(t)},$$

with explicit constants $q_*, q_R > 0$*,* R_ℓ*,* c_0 *and* $\mu_\xi = \mathbb{E}[S]\,\mathbb{E}p\,(\beta - 1)$ *given in the appendices, equation* (A1).

In the atom case, F single-job products with all entries equal to a form the start set, and their number is binomial. Without an atom, the start set needs $F-1$ products whose total processing time lies within the first one's assembly time of the minimum ma, and the probability of finding them depends on how much mass μ puts near a and κ_1 near 0 (Theorem 14 below). In expectation the bound is also exact: $\mathbb{E}[\mathrm{OPT} - X] \to 0$ for every admissible X (Corollary 34).

5.2 An additive guarantee at every load ratio

The Steinitz blocks alone, without start set or ramps, give a schedule within a constant of a simple bound.

Theorem 10 (additive guarantee) *For every instance,*

$$\mathrm{OPT} \le \max\Big(A_{\mathrm{tot}},\ \max_i \frac{L_i}{F}\Big) + m(4m + 5)\bar{x} + 2F(m + 1)\bar{A}.$$

Hence $\mathrm{OPT} - X \le m(4m+5)\bar{x} + 2F(m+1)\bar{A}$ *for every lower bound* $X \ge \max(A_{\mathrm{tot}}, \max_i L_i/F)$*, in particular for* $X = \mathrm{LB}_{\mathrm{cf}}$ *and* $X = \mathrm{LB}_{\mathrm{comb}}$.

The constant does not depend on n when products have bounded size. This extends to F factories and multi-job products the additive guarantees of Potts et al. (1995, Thm. 7) for the two-stage assembly problem and of Koulamas and Kyparisis (2001, Thm. 5) for the three-stage assembly flowshop, both also based on vector summation.

5.3 When the assembly bounds fail

Theorem 10 says that the *combined* bound is always close. Bounds that look only at the assembly side are not.

Proposition 11 ($\beta \le 1$) *In the random model of Section 2.2:*

(a) if $\beta < 1$*, there are* $c, c' > 0$ *with* $\Pr[\mathrm{OPT} - \mathrm{LB}^* < ct] \le e^{-c't}$ *for all large* t*; so the product-decomposition bound, and every admissible bound below it, is off by* $\Theta(t)$ *with probability* $1 - e^{-\Omega(t)}$*;*

(b) if $\beta = 1$ *and* $\mathrm{Var}(x_{h,1}/F - A_h) > 0$*, then* $\mathbb{E}[\mathrm{OPT} - \mathrm{LB}^*] = \Theta(\sqrt{t})$*,* $\mathrm{OPT} - \mathrm{LB}^* = O_P(\sqrt{t})$ *and* $\liminf_t \Pr[\mathrm{OPT} - \mathrm{LB}^* \ge c\sqrt{t}] > 0$ *for some* $c > 0$.

The reason is that $\mathrm{LB}^* \le \bar{s}mb + A_{\mathrm{tot}}$ (every single product can be finished within the sum of its entries), while $\mathrm{OPT} \ge \max_i L_i/F$. When the factories are the bottleneck, the decomposition bounds cannot see it. So, for $\beta \le 1$, any exactness result must use a bound that sees the factory side, such as $\mathrm{LB}_{\mathrm{end}}$; this is what Section 6 does.

The next two results show that even for $\beta > 1$ the weakest bound LB_0 can fail, and identify why.

Proposition 12 (start obstruction) *Let* $t \ge 2$*, let* h_1 *be the unique minimiser of* R *with* $s_{h_1} \ge F$*, and put* $P_{\min} = \min_{j \in N_{h_1}} P_j$*,* $D = T_0 - P_{\min}$ *and* $T_0' = \min_{h \ne h_1} R_h$*. Then*

$$\mathrm{OPT} \ge \mathrm{LB}_0 + \max\{0,\ \min(T_0' - T_0,\ a - D - A_{h_1})\}.$$

In the random model with every product having at least F *jobs,* μ *without atoms,* $a > 0$ *and* $\kappa_{\min} := \min_{s \in \mathrm{supp}\,\nu} \kappa_s([0, a/2)) > 0$*, we get* $\liminf_t \Pr[\mathrm{OPT} > \mathrm{LB}_0] \ge \kappa_{\min}$*, for every* β.

If the start product has at least F jobs, either two of them share a factory, or every factory starts with one of them; in both cases some job is delayed by at least one processing time a, and if h_1's assembly is short, the assembly machine idles. The proposition does not show that single-job products are necessary: the regime in which every product has at least two but some have fewer than F jobs is open.

Proposition 13 (gap obstruction) *Let* $t \ge 2$*, let* h_1 *be the unique minimiser of* R*, and let* $R_{(2)}$ *be the second smallest release. Then* $\mathrm{OPT} \ge \mathrm{LB}_0 + (R_{(2)} - T_0 - A_{h_1})$*. Suppose that* $\nu(1) > 0$*, that* $c_H x \le \mu([a, a+x]) \le C_H x$ *for* $0 \le x \le x_H$*, and that* $\kappa_1([0,u)) \ge \kappa_0 u$ *for small* u*. Then* $\Pr[\mathrm{OPT} > \mathrm{LB}_0] \ge c\, t^{-1/m}$ *for some* $c > 0$ *and all large* t*, for every* β.

Theorem 14 (rate for continuous laws) *Under (M1), (M2) and (M4), if* $\mu([a, a+x]) \geq c_H x$ *for* $x \leq x_H$ *and* $\kappa_1([0, u)) \leq C_\kappa u$ *for all* u*, then* $\Pr[\mathrm{OPT} > \mathrm{LB}_0] = O(t^{-1/m})$*. Together with Proposition 13,* $\Pr[\mathrm{OPT} > \mathrm{LB}_0] = \Theta(t^{-1/m})$ *when* $\beta > 1$*,* μ *has a density bounded above and below near* a*, and* κ_1 *has a density bounded above and below near* 0*.*

So the rate on the assembly side is set by the lower tail of the processing times: an atom at a gives an exponentially small failure probability, a density gives $t^{-1/m}$. When κ_1 vanishes near 0, as in the benchmark generator ($A \geq s$), the failure probability is $O(1/t)$; a matching lower bound is open there.

5.4 A preregistered test

We tested three predictions of this section in a simulation whose protocol, tests and minimum effects were fixed before any instance was generated: eight tests, Holm-corrected at level 0.05, on 1,320 instances with $F = 4$. The upper bound came from the certificate of Theorem 6, the block schedule of Theorem 10, the greedy heuristic of the benchmark study, 20 s of local search and, for $n \leq 150$, 60 s of CP-SAT. The lower bound was $\max(\mathrm{LB}_{\mathrm{cf}}, \mathrm{DB}, \mathrm{LB}_{\mathrm{end}})$, together with Proposition 12.

- $\beta > 1$ *(Theorem 9).* At $t = 800$ the bound solved 36/40, 40/40 and 40/40 instances for $\beta = 1.5$, 2 and 4 (Holm $p \leq 0.005$; all above the preregistered 0.85).
- $\beta = 1$ *(Proposition 11).* The gap to the decomposition bound grows like $t^{0.57 \pm 0.08}$ (theory: $t^{1/2}$; Holm $p = 10^{-5}$). The preregistered claim that the gap to the combined bound is flat (slope below 0.2, minimum effect 0.10) was *not* confirmed: the slope was 0.14 ± 0.06 (Holm $p = 0.16$). That gap is flat for $m = 3$ (median 19 at $t = 25$ and at $t = 800$). For $m = 5$ it rises at $t \geq 400$, where the upper bound came from a fixed 20 s local search on about 1,600 jobs; the relative gap still falls from 5.6 % to 0.46 %.
- *Start obstruction (Proposition 12).* When every product has at least F jobs, $\mathrm{OPT} > \mathrm{LB}_0$ was certified on 55–63 % of the instances at $t = 400$ (Holm $p < 10^{-17}$). As t grows, the decomposition bound certifies this less often (from 22–25 to 14 of 40 instances), and the bound of Proposition 12 more often (from 12–13 to 16–21 of 40).

6 Lattice laws: the factory side at every load ratio

On the factory side the decomposition bounds fail (Proposition 11), and the question becomes whether $\mathrm{LB}_{\mathrm{end}}$, or the combined bound, is exact. For integer processing times we can answer it at every β. In this section the entries have an integer-valued law μ that satisfies

(L) $a = \min \operatorname{supp} \mu < b = \max \operatorname{supp} \mu$, and $\operatorname{supp} \mu$ contains two consecutive integers $s, s+1$.

Every $U\{a..b\}$ with $b \geq a+1$ qualifies, including Taillard's $U\{1..99\}$. *Lattice model* M is the random model of Section 2.2 with μ satisfying (L) and $a \geq 1$, $\nu(1) > 0$, and integer assembly times: κ_s is supported on the integers in $[s\alpha_0, \bar{A}]$ for an integer $\alpha_0 \geq 0$, and $\kappa_1(\{\alpha_0\}) > 0$. So α_0 is the smallest assembly time per job, and it is attained by single-job products. Throughout, $\mathrm{LB}_{\mathrm{end}}$ denotes the rounded-up bound of Proposition 3.

6.1 Exact partition over factories

Inside one factory, the staircase blocks of Coutinho (manuscript under review) make the makespan equal to the largest machine load plus $(m-1)a$, for every load vector, once the instance contains enough *exact staircase jobs*: front jobs of type $r = 2, \ldots, m$ (entries a before machine r and b from r on) and end jobs of type $r = 1, \ldots, m-1$ (b up to machine r and a after it). To reach the factory bound, every factory must also receive almost exactly a $1/F$ share of every machine's load. The next lemma does this exactly, using the two consecutive support points of (L).

Lemma 15 (exact lattice partition) *Let each of F groups consist of a common fixed part, a Steinitz block (Lemma 5) of a core set of d-dimensional vectors, and a common reservoir that holds, for each coordinate k in a set K, E items of vector τ_k^0 and E items of vector $\tau_k^1 = \tau_k^0 + e_k$. Let the targets $t_{f,k}$ satisfy $\sum_f t_{f,k} = L_k$ and $|t_{f,k} - L_k/F| \leq \theta$. If $E \geq (2d+1)\bar{Z}_k + \theta$, where $\bar{Z}_k$ bounds coordinate k of the core vectors, then at most $|K|FE$ exchanges of reservoir items make every group's coordinate-k load equal to $t_{f,k}$, for every $k \in K$. The other coordinates, the fixed parts and the blocks do not change.*

The Steinitz blocks bring every factory's loads within a constant of L_k/F, and exchanging a τ_k^1 for a τ_k^0 moves exactly one unit of machine-k load between two factories without touching the other machines. With $\tau_k^0 = a\mathbf{1} + (s-a)e_k$, both vectors are in the support of the job law, so the exchanges use jobs that random instances contain in linear numbers.

6.2 Taillard's law for the distributed permutation flowshop

With $A \equiv 0$, the factory bound is $\mathrm{LB}_D = \max_k[\mathrm{hd}_k + \lceil L_k/F \rceil + \mathrm{tl}_k]$, the bound of Gogos (2023, Alg. 1) rounded up. It is valid for every schedule.

Theorem 16 (Taillard's law for the DPFSP)

(a) Let the entries be integers in $[a, b]$ and $E = 2m(b-a) + b + 1$. Suppose the instance contains Fm jobs of each of the $2(m-1)$ staircase types and, for every k, FE jobs each of τ_k^0 and τ_k^1, with demands added when vectors coincide. Then $\mathrm{OPT} = (m-1)a + \max_k\lceil L_k/F \rceil = \mathrm{LB}_D = X$ for every $X \in [\mathrm{LB}_D, \mathrm{OPT}]$, also over non-permutation schedules, and an optimal schedule is computed in polynomial time.

(b) Let the entries be i.i.d. μ satisfying (L), with $m \geq 2$ and $F \geq 1$. Put $q_L = \min\{\min_{1 \leq s' \leq m-1} \mu(a)^{s'}\mu(b)^{m-s'},\ \mu(a)^{m-1}\min(\mu(s), \mu(s+1))\}$ and $D = 2mF(m-1+E)$. Then $\Pr[\mathrm{OPT} \neq X] \leq (4m-2)\Pr[\mathrm{Bin}(n, q_L) < D] \leq (4m-2)e^{-nq_L/8}$ for $n \geq 2D/q_L$ and every $X \in [\mathrm{LB}_D, \mathrm{OPT}]$.

(c) With $\rho_0 = \min(\mu(a), \mu(b), \mu(s), \mu(s+1))$ *and* $n \geq F(mb+1)+2$*, the permutation optimum satisfies* $\Pr[\mathrm{OPT} > \mathrm{LB}_D] \geq \rho_0^{mn}$*. So* $-\frac{1}{n}\log \Pr[\mathrm{OPT} \neq \mathrm{LB}_D]$ *stays between two positive constants.*

Part (a) puts m staircase jobs of each type in every factory and uses Lemma 15 to give every factory the loads $\lceil L_k/F \rceil$ or one less. Part (b) is a union bound with a Chernoff tail, and part (c) exhibits an instance of all-constant jobs with equal loads divisible by F, on which the job of entries b costs one unit. For $F = 1$ the theorem reduces to the lattice case of the single-factory result of Coutinho (manuscript under review). For $U\{1..B\}$ the scale of n needed is B^m, far beyond benchmark sizes when $B = 99$; the theorem fixes the order of the failure probability, not its size at $n \leq 500$.

6.3 The end side: load ties

With an assembly stage the picture changes. Suppose two machines $k < l$ carry the same, maximal load L. The last job of every factory must finish on machine k and on machine l, and in a permutation schedule the job before it must already have finished on l. Charging both completions of every factory gives $\mathrm{LB}_{\mathrm{tie}}$ (Proposition 4). On a tie, $\mathrm{LB}_{\mathrm{tie}}$ is larger than $\mathrm{LB}_{\mathrm{end}}$ exactly when the end products of the factories cannot be assembled fast enough, which depends on whether $F\alpha_0$ exceeds a.

Proposition 17 (tie obstruction)

(i) Let all data be integers. Suppose that some job has all entries equal to a*; that machines* $k < l$ *have* $L_k = L_l = L = \max_i L_i$*; that* $A_h \geq s_h\alpha_0$ *for an integer* α_0*, with at least* $2F$ *single-job products of assembly time* α_0*; that* $F\alpha_0 > a$ *and* $\lceil x \rceil - x < (F\alpha_0 - a)/2$ *for* $x = L/F + \alpha_0(F+1)/2$ *(this holds whenever* $F\alpha_0 \geq a+2$*, and whenever* F *divides* $L + \alpha_0 F(F+1)/2$*); and that* $L \geq F^2 b + F(2F-1)[(m-1)a + \alpha_0(2F+1)/2 + 1]$*. Then* $\mathrm{OPT} \geq \mathrm{LB}_{\mathrm{tie}} \geq \mathrm{LB}_{\mathrm{end}} + 1$*. If moreover* $FA_{\mathrm{tot}} \leq L - F\bar{s}mb$ *and, for every machine* k*, at least* $F\bar{A}$ *single-job products have* $A = \alpha_0$ *and* $p_k \geq F\alpha_0 + 1$*, then also* $\mathrm{OPT} \geq \mathrm{LB}_{\mathrm{comb}} + 1$*.*

(ii) Under lattice model M *with* $m \geq 2$*,* $a < F\alpha_0 < b$ *and* $\beta \leq 1$*, there is* $c > 0$ *with* $\Pr[\mathrm{OPT} > X] \geq c\,t^{-1/2}$ *for* $X \in \{\mathrm{LB}_{\mathrm{end}}, \mathrm{LB}_{\mathrm{comb}}\}$ *and all large* t*. For* $X = \mathrm{LB}_{\mathrm{end}}$ *the restriction on* β *is not needed.*

Part (i) is deterministic and needs neither $F \geq 2$ nor a particular β. Part (ii) follows from a local limit theorem for the vector of machine loads: an exact tie $L_1 = L_2 = \max_i L_i$, with L in the residue class modulo F that makes x an integer, has probability of order $t^{-1/2}$, and the other conditions of (i) fail only with probability $e^{-\Omega(t)}$. At $\beta = 1$ the assembly condition $FA_{\mathrm{tot}} \leq L - F\bar{s}mb$ holds with probability bounded away from zero, independently of the tie given the product sizes. The obstruction is driven by the assembly stage: with $A \equiv 0$ it disappears, in agreement with Theorem 16.

6.4 The factory side with a constant margin

Theorem 8 needs staggered factory ends and a bottleneck margin. With Lemma 15, both can be arranged with a margin that does not grow with t.

Theorem 18 (factory-side certificate) *Let the data be integers in $[a, b]$, fix a machine r and an integer $\alpha_0 \geq 0$ with $F\alpha_0 < b$, and let $s, s+1 \in [a, b]$. There are constants Δ^*, c_1, $\beta_\uparrow$, $\beta_\downarrow$ and E_T, depending only on $(m, F, a, b, s, \bar{s}, \bar{A}, \alpha_0)$ and r, such that the following holds. Suppose that $A_h \geq s_h\alpha_0$ for every product, and that the instance contains, as single-job products, $F\beta_\uparrow$ copies of $v_r = a\mathbf{1} + (b-a)e_r$, $F\beta_\downarrow$ copies of v_r with $A = \alpha_0$, and FE_T copies each of τ_r^0 and τ_r^1. Suppose also that $L_r - L_i \geq \Delta^*$ for $i \neq r$, that $FA_{\mathrm{tot}} \leq L_r - c_1$, and that $L_r \geq \alpha_0 F(F-1)/2$. Then*

$$\mathrm{OPT} = M := (m-1)a + \left\lceil (L_r + \alpha_0 F(F+1)/2)/F \right\rceil = \mathrm{LB}_{\mathrm{end}} = X$$

for every valid lower bound $X \geq \mathrm{LB}_{\mathrm{end}}$, and the schedule of Theorem 8, built with Lemma 15, attains it in polynomial time.

Lemma 15 puts load $M - (m-1)a - f\alpha_0$ on machine r of factory f, which is the staggering that Theorem 8 needs; the middle of every factory is reordered with its own means, so the ramp conditions hold factory by factory. The constants are explicit but large: for $m = 2$, $\bar{s} = 1$, $U\{1..3\}$ ($a = 1$, $b = 3$, $s = 1$), bottleneck $r = m = 2$, $F = 2$, $\alpha_0 = 1$ and $\bar{A} = 4$, $\Delta^* = 584$ and $c_1 = 1{,}816$. A constructive check with exactly these constants attained M on every instance built.

6.5 Rates at every load ratio

Combining the two certificates (Theorems 9 and 18) with the obstructions gives the rate at every β. Let $\mathcal{X}$ be the set of valid lower bounds $X \geq \max(\mathrm{LB}_0, \mathrm{LB}_{\mathrm{end}})$; it contains $\mathrm{LB}_{\mathrm{comb}}$.

Theorem 19 (every load ratio, lattice laws) *Assume lattice model M with $F \geq 2$, $F\alpha_0 < b$, and (M2) for $\beta \geq 1$. Then:*

(a) for $\beta < 1$, $\Pr[\mathrm{OPT} \neq X] \leq C_1 t^{-1/2}$ for every valid $X \geq \mathrm{LB}_{\mathrm{end}}$;

(b) for $\beta = 1$, $\Pr[\mathrm{OPT} \neq X] \leq C_2 t^{-1/2}$ for every $X \in \mathcal{X}$;

(c) for $\beta > 1$, $\Pr[\mathrm{OPT} \neq X] = e^{-\Omega(t)}$ for every admissible X;

(d) if $m \geq 2$ and $F\alpha_0 > a$, the rate is $\Theta(t^{-1/2})$ for $X \in \{\mathrm{LB}_{\mathrm{end}}, \mathrm{LB}_{\mathrm{comb}}\}$ when $\beta < 1$, and for $X = \mathrm{LB}_{\mathrm{comb}}$ when $\beta = 1$;

(e) if $m \geq 2$, $F\alpha_0 \leq a$ and $\beta < 1$, then $\Pr[\mathrm{OPT} \neq X] \leq e^{-\Omega(t)}$ for every valid $X \geq \mathrm{LB}_{\mathrm{end}}$.

For (a) and (b), the factory-side certificate applies when $FA_{\mathrm{tot}} \leq L_{r^*} - c_1$ at the most loaded machine r^*, and the assembly-side certificate when $FA_{\mathrm{tot}} - L_{r^*} \geq c_0$. Both need a load margin $L_{r^*} - L_i \geq \Delta^*$, which fails with probability $O(t^{-1/2})$ by the Berry–Esseen theorem. The window between the two certificates has constant width; for $\beta < 1$ it has probability $e^{-\Omega(t)}$, and for $\beta = 1$ probability $O(t^{-1/2})$. Part (d) adds the lower bound of Proposition 17. Part (e) is the subject of the next subsection: below and on the threshold, ties do not matter, and the margin event is not needed at all.

6.6 The sharp threshold: harmless ties

Proposition 17 needs $F\alpha_0 > a$. When $F\alpha_0 \le a$, ties are harmless and the factory side behaves like the distributed flowshop. The certificate uses staircase blocks in every factory, an assembly coordinate in the Steinitz order of the middle, and a short *buffer* of single-job products of vector $w = a\mathbf{1} + (b-a)e_m$ (front type m) between the middle and the end block.

Lemma 20 (staircase blocks with an assembly coordinate) *Let $m \ge 2$, let the entries lie in $[a, b]$ with $g = b - a > 0$, and let every product have at most $\bar{s}$ jobs, kept consecutive in one factory. Suppose a factory holds c exact front jobs of each type, c exact end jobs of each type, and a buffer of $K_B \ge 0$ jobs of vector w, all single-job products, with $c \ge (m+1)(\bar{s}+1) + K_B$; let Mid be its other products, with at least K_B items after padding. Pad* Mid $\cup$ Buf *to equal machine loads with single-entry jobs placed in Mid, and order Mid by Lemma 5 in dimension m, over the centred machine differences $x_{h,l} - x_{h,l+1}$ and the centred coordinate $y_h = A_h - x_{h,m}/F$. Sequence the factory as Front · Mid · Buf · End, with the front and end jobs in decreasing type order. Then:*

(a) the factory makespan is $\max_k L_k + (m-1)a$;

(b) if $\sum_{h \in \text{Mid}} y_h \le 0$, every suffix Z of whole products of Mid has $\sum_{h \in Z} y_h \le m(\bar{A} + \bar{s}b/F)$.

The buffer is not centred, so it shifts the last machine difference by up to $K_B g$ inside Mid; the staircase count c absorbs this shift, by the absorption lemma of Coutinho (manuscript under review). The padding must cover the buffer as well: if Mid alone is padded and the buffer appended, the buffer raises machine m above the factory's largest load (Remark after the proof of Theorem 21 in the appendices).

Theorem 21 (harmless ties on and below the threshold)

(a) Let $m \ge 2$, let the data be integers with entries in $[a, b]$, $a \ge 1$, and fix an integer $\alpha_0 \ge 0$ with $F\alpha_0 \le a$ and $s, s+1 \in [a, b]$. There are constants c, E, K_B and c_2, depending only on $(m, F, a, b, \bar{s}, \bar{A}, \alpha_0)$, such that the following holds. Suppose that $A_h \ge s_h \alpha_0$ for every product, and that the instance contains, as single-job products, Fc front jobs of each type, Fc end jobs of each type with $A = \alpha_0$, FK_B jobs of vector w with $A = \alpha_0$, and FE each of τ_k^0, τ_k^1 for every k, with demands added when vectors coincide. Suppose also that $FA_{\text{tot}} \le L_m - c_2$ and $L = \max_k L_k \ge \alpha_0 F(F-1)/2$. Then

$$\text{OPT} = M := (m-1)a + \lceil (L + \alpha_0 F(F+1)/2)/F \rceil = \text{LB}_{\text{end}} = X$$

for every valid lower bound $X \ge \text{LB}_{\text{end}}$, whatever the load vector, exact ties included. The certificate is computed in polynomial time. If $F\alpha_0 < a$, one may take $K_B = 0$.

(b) Under lattice model M with $m \ge 2$, $F\alpha_0 \le a$ and $\beta < 1$, $\Pr[\text{OPT} \ne X \text{ for some valid } X \ge \text{LB}_{\text{end}}] \le e^{-\Omega(t)}$.

Proof idea Lemma 15, with every machine targeted, gives factory f load at most $T_f = M - (m-1)a - f\alpha_0$ on every machine. These targets are within $F\alpha_0/2 + 1$ of L_k/F for every k, so no load margin is needed. Lemma 20 then gives factory completion times at most $M - f\alpha_0$. Machine m processes the last jobs of factory f one after another, so the

Table 4 Rate of $\Pr[\mathrm{OPT} \neq X]$ by load ratio for lattice laws, under the hypotheses of Theorem 19; X is the combined bound $\mathrm{LB}_{\mathrm{comb}}$, and the DPFSP row is for the permutation optimum and LB_D. †: the lower bound is proved for $F\alpha_0 > a$; for $F\alpha_0 \le a$ at $\beta = 1$ it is open.

Regime	Rate	Mechanism
DPFSP ($A \equiv 0$)	$e^{-\Theta(n)}$	lattice partition and staircase blocks (Thm. 16)
$\beta > 1$	$e^{-\Omega(t)}$	assembly-side certificate (Thm. 9)
$\beta = 1$	$\Theta(t^{-1/2})^\dagger$	window between the two certificates; load ties
$\beta < 1$, $F\alpha_0 > a$	$\Theta(t^{-1/2})$	load ties among factory machines (Prop. 17)
$\beta < 1$, $F\alpha_0 \le a$	$\le e^{-\Omega(t)}$	staircase blocks, end slots and buffer (Thm. 21)

j-th end job from the end is released by $M - f\alpha_0 - (j-1)a$. Give it the assembly slot $[M - f\alpha_0 - (j-1)a,\ M - (f-1)\alpha_0 - (j-1)a]$. For fixed j the F slots tile a window of length $F\alpha_0 \le a$, so the slots are disjoint and the end jobs alone never delay the assembly machine. The middle products are released before the end and buffer jobs and may exceed their straight-line share of assembly work by a constant (Lemma 20(b)); this excess must be absorbed by idle time of the assembly machine within the end region. When $F\alpha_0 < a$, each round of end jobs leaves a gap $a - F\alpha_0 > 0$. When $F\alpha_0 = a$, the end rounds leave no gap, but the buffer jobs have last-machine entry b, so each buffer round leaves $b - F\alpha_0 > 0$. A constant number of rounds absorbs the excess. The condition $FA_{\mathrm{tot}} \le L_m - c_2$ covers the case in which the counted work reaches the front jobs. Part (b) follows from binomial counts and Hoeffding's inequality, since $\mathbb{E}[FA_h - x_{h,m}] < 0$ for $\beta < 1$. □

The certificate keeps every product in one factory, so the theorem also holds when products may not be split. Its constants are large; it fixes the order of the failure probability, not its size at benchmark scale.

Corollary 22 (dichotomy) *Under lattice model M with $m \ge 2$, $F \ge 2$, $F\alpha_0 < b$ and $\beta < 1$:*

- *if $F\alpha_0 \le a$, then $\Pr[\mathrm{OPT} \neq \mathrm{LB}_{\mathrm{end}}] \le e^{-\Omega(t)}$;*
- *if $F\alpha_0 > a$, then $\Pr[\mathrm{OPT} > X] = \Theta(t^{-1/2})$ for $X \in \{\mathrm{LB}_{\mathrm{end}}, \mathrm{LB}_{\mathrm{comb}}\}$.*

The hypothesis $F\alpha_0 < b$ is implied by $\beta < 1$, since $A_h \ge s_h\alpha_0$ gives $\beta \ge F\alpha_0/\mathbb{E}p$.

So the assembly stage changes the order of the failure probability exactly when F products of minimal assembly time cannot be assembled within one minimal processing time. The distributed flowshop ($\alpha_0 = 0 < a$) lies on the exponential side, in agreement with Theorem 16. The threshold has a simple operational reading: on the last machine the factories release their final products at intervals of at least a, and the assembly machine keeps up with F parallel release streams if and only if $F\alpha_0 \le a$. Above the threshold, a tie between two machines forces two slow end products per factory, and no rearrangement of the end can absorb them.

6.7 A preregistered test of the tie obstruction

We tested Proposition 17(ii) in a second simulation, whose protocol, code and analysis were frozen by hash before any instance was generated. The cell was $m = 2$, $F = 3$, entries $U\{1..9\}$, sizes uniform on the multiset $\{1, 1, 2\}$, and $A_h = s_h \cdot (1$ with probability 0.8, else 2), so that $F\alpha_0 = a + 2$; the realised β was 0.72. We used $t \in \{400, 1{,}600, 6{,}400, 25{,}600, 102{,}400\}$, with 1.72 million instances in total, and two confirmatory tests, Holm-corrected at level 0.05.

- *Rate.* The frequency of an exact tie $L_1 = L_2$ on which $\mathrm{LB}_{\mathrm{tie}} > \mathrm{LB}_{\mathrm{comb}}$ decays like $t^{-0.494\pm0.016}$; the theory predicts $t^{-1/2}$. The preregistered equivalence test for a slope in $[-0.75, -0.25]$ is confirmed ($p < 10^{-15}$). The tie frequency times $\sqrt{t}$ stays between 0.088 and 0.096.
- *Certification.* $\mathrm{LB}_{\mathrm{tie}}$ exceeds $\mathrm{LB}_{\mathrm{comb}}$ on 562 of the 562 ties at $t \geq 25{,}600$ ($p = 7 \times 10^{-170}$), and on all 1,307 ties at every t.

Three further results are exploratory. In random subsamples of 1,000 instances, certified gaps without an exact tie fall from 31 to 0 as t grows. In a contrast cell with $F\alpha_0 < a$ ($F = 2$, $U\{3..9\}$, $\beta \approx 0.40$), ties are as frequent, but none of 782 is certified, as Theorem 21 predicts. In the DPFSP ($m = 3$, $F = 2$, $U\{1..3\}$), the counts of Theorem 16(a) never hold at $n \leq 1{,}200$; at $n = 4{,}800$ they hold on 100 of 100 instances, and the construction attains LB_D on all of them. The whole simulation used about one CPU-hour.

The certificate of Theorem 21 was also checked constructively on planted exact ties: it attained $\mathrm{LB}_{\mathrm{end}}$ on all 226 instances below the threshold ($m \in \{2, 3\}$, $F \in \{2, 3\}$, $\alpha_0 \in \{0, 1, 2\}$, up to 11,421 jobs; 106 of them built with the theorem's own constants), and, with a buffer of 20 jobs per factory, on all 75 instances on the threshold $F\alpha_0 = a$ with $\beta \leq 0.94$. Without the buffer it failed on 30 of these 75. On the bound side, exact CP-SAT runs on small instances found $\mathrm{LB}_{\mathrm{tie}} > \mathrm{LB}_{\mathrm{comb}}$ on ties only above the threshold.

7 Discussion

What the results say.

On random instances a simple lower bound solves the DAPFSP with probability tending to one, in every regime we analyse except one: when every product has at least F jobs, the start obstruction keeps LB_0 from being exact. (On the factory side, non-lattice laws are not analysed.) What changes with the load ratio is which bound is exact and how fast. When the assembly machine is the bottleneck ($\beta > 1$), the weakest assembly bound LB_0 is exact, and the failure probability is set by the lower tail of the processing times: exponentially small with an atom at the minimum, $\Theta(t^{-1/m})$ with a density. When the factories are the bottleneck ($\beta < 1$), every product-decomposition bound fails by $\Theta(t)$, and the right bound is the covering end-chain bound $\mathrm{LB}_{\mathrm{end}}$. For integer data and $F \geq 2$ its failure probability is exponentially small if $F\alpha_0 \leq a$ and exactly of order $t^{-1/2}$ if $F\alpha_0 > a$, the slow rate being caused by ties between machine

loads. At $\beta = 1$ the two certificates leave a window of constant width, and the rate is $O(t^{-1/2})$.

What the results enable.

The certificates are polynomial and need no search. On the benchmark they act as an exact algorithm on the instances they certify (median 3.4 ms against a CP-SAT model that proves optimality within 60 s on a third of a sample), as a stopping rule that cuts a fixed search budget 5.5-fold, and as an a priori rule that predicts closure from two quantities read off the data. They also explain the closed/open split of the benchmark by product size, which a purely computational study can only observe. Combined with exact single-product releases and short local search, they prove optimality for 22 large instances of Hatami et al. (2013) that no published study had closed, two of them with new best values, and raise the number of solved large instances from 703 to 725 of 810. For generators of new benchmarks the message is direct: on random instances with large n the optimum is a lower bound with overwhelming probability unless the instance sits near $\beta = 1$, has products of at least F jobs, or has machine-load ties with $F\alpha_0 > a$. Instances meant to discriminate between algorithms should be generated in those regimes, which the load-controlled generator of Coutinho (unpublished manuscript) makes possible.

Scope.

The asymptotic results are statements about orders, not about benchmark-size constants. The constants of Theorems 16, 18 and 21 are explicit but large, as in the single-factory case, and for $U\{1..99\}$ the guarantees become effective only far beyond $n = 500$. The two preregistered simulations show that the predicted behaviour (the exponent $-1/2$ for ties, the $\sqrt{t}$ growth of the decomposition gap at $\beta = 1$, and near-certain exactness for $\beta > 1$) is already visible between a few hundred and 10^5 products. The lattice results assume integer data and a law with two consecutive support points, which covers every discrete uniform law used in practice. Product sizes are bounded; under the size law of Hatami et al. (2013) (one plus a multinomial) the additive constant of Theorem 10 grows, but only as $o(\sqrt{n})$.

Open problems.

(1) *The lower bound at $\beta = 1$ for $F\alpha_0 \le a$.* The upper bound $O(t^{-1/2})$ holds, but ties are harmless there, and the window between the certificates may be the only obstruction.

(2) *Non-lattice laws on the factory side.* The exact partition of Lemma 15 uses unit exchanges. A tolerant version, as in the bridge lemmas of Coutinho (manuscript under review), should give the continuous analogue, presumably with a polynomial rate.

(3) *A hierarchy of tie bounds.* Adding $\mathrm{LB}_{\mathrm{tie}}$ to the combined bound removes the two-machine tie obstruction. Three machines tied in load presumably give an obstruction of order t^{-1}, which would suggest a hierarchy of bounds, each improving the rate by a factor $t^{-1/2}$.

(4) *Intermediate product sizes.* When no product is a single job but some have fewer than F jobs ($\nu(1) = 0 < \nu(\{2, \dots, F-1\})$), neither Theorem 9 nor Proposition 12 applies.
(5) *Assembly times bounded away from zero.* When κ_1 vanishes near 0, as in the benchmark generator, the failure probability of LB_0 for continuous laws is $O(1/t)$; a matching lower bound is open.

Use of large language models. The author used AI assistants (Anthropic's Claude models, Claude Opus 5.5 among them) for literature search, drafting and editing of text, writing experiment code, exploring proof strategies and candidate lemmas, and adversarial checking of proofs and numerics. The rules used to steer them were: every claim must be proved in full and checked line by line by independent reading passes; numeric claims are verified by independent recomputation; confirmatory experiments are preregistered; and credit to prior work is checked against the original papers. The author verified, and takes full responsibility for, all content.

Declarations

- **Funding.** No funding was received for this work.
- **Competing interests.** The author has no competing interests to declare.
- **Ethics approval and consent to participate.** Not applicable.
- **Consent for publication.** Not applicable.
- **Data availability.** The benchmark instances of Hatami et al. (2013) are public. The generated instances, the computed bounds, the optimal schedules of the 22 newly certified instances and the preregistration records are available from the author on reasonable request and will be deposited in a public repository before publication.
- **Materials availability.** Not applicable.
- **Code availability.** The code for the certificates, the bounds and the simulations is available from the author on reasonable request and will be deposited together with the data.
- **Author contribution.** F. P. Coutinho conceived the study, proved the results, wrote the code, ran the experiments and wrote the manuscript.

Appendix A Proofs for Sections 2, 3 and 5

A.1 Basic lemmas

Throughout, a and b denote the global minimum and maximum processing-time entries, $P_j = \sum_i p_{ij}$, $L_i = \sum_j p_{ij}$, $x_{h,i} = \sum_{j \in N_h} p_{ij}$, $R_h = \max_{j \in N_h} P_j$, $T_0 = \min_h R_h$, $\mathrm{LB}_0 = T_0 + A_{\mathrm{tot}}$, and $c_h = \max_{j \in N_h} C_j$ is the completion time of product h in the factories. Unless stated otherwise, the jobs of one product may be assigned to different factories, as in Hatami et al. (2013). A bound X is *admissible* if $\mathrm{LB}_0 \le X \le \mathrm{OPT}$.

Proof of Lemma 1 Recall the claim: for releases c and processing times A on one machine, the optimal makespan is $\mathrm{ERD}(c) = \max_h \big[c_h + \sum_{g:\, c_g \geq c_h} A_g\big]$; it is attained by every order of non-decreasing c; ERD is non-decreasing in every c_h; and for a fixed factory schedule the best assembly schedule has makespan $\mathrm{ERD}(c)$.

Lower bound. In every feasible sequence, every product g with $c_g \geq c_h$ starts at or after c_h, and these products are processed one at a time. Hence the makespan is at least $c_h + \sum_{g:\, c_g \geq c_h} A_g$ for every h.

Upper bound. Take a sequence of non-decreasing c and its last busy period. Let h be the first product of that period, so that h starts at c_h. Every product of the period follows h in the order, so $c_g \geq c_h$ for each of them. Hence the makespan equals c_h plus the work of the period, which is at most $c_h + \sum_{g:\, c_g \geq c_h} A_g$. (If a product tied with h came before h in the period and had $A > 0$, it would end after c_h, contradicting that h starts at c_h; tied products with $A = 0$ do not affect the sum.)

Monotonicity. An assembly sequence that is feasible for releases $\hat{c} \geq c$ is also feasible for c, with no larger makespan.

Last claim. For a fixed factory schedule, the assembly machine faces exactly the single-machine problem $1|r_h|C_{\max}$ with $r = c$. □

Corollary 23 (Validity of release-based bounds) *Let r_h^* be the optimal makespan of the distributed permutation flowshop on N_h alone, using all F factories. Every vector $\underline{r}$ with $R_h \leq \underline{r}_h \leq r_h^*$ for all h gives $\mathrm{ERD}(\underline{r}) \in [\mathrm{LB}_0, \mathrm{OPT}]$. In particular $\mathrm{LB}^* = \mathrm{ERD}(r^*)$, the closed-form bound $\mathrm{LB}_A = \mathrm{ERD}(\underline{r})$ of Coutinho (unpublished manuscript) (whose first term is $\max_j P_j$), and $\max(\mathrm{LB}_A, \mathrm{LB}_F, \mathrm{LB}_J)$ are admissible. LB_0 is valid also for non-permutation factory schedules.*

Proof In any schedule, deleting all jobs outside N_h leaves a feasible schedule of N_h alone, so $c_h \geq r_h^* \geq \underline{r}_h$. By the monotonicity in Lemma 1, $\mathrm{ERD}(\underline{r}) \leq \mathrm{ERD}(c) \leq$ makespan. Moreover $\mathrm{ERD}(\underline{r}) \geq \min_h \underline{r}_h + A_{\mathrm{tot}} \geq T_0 + A_{\mathrm{tot}} = \mathrm{LB}_0$, taking h a minimiser of $\underline{r}$ in the formula of Lemma 1. Each of $\mathrm{LB}_A, \mathrm{LB}_F, \mathrm{LB}_J$ is a valid lower bound, and $\mathrm{LB}_A \geq \mathrm{LB}_0$, so their maximum is admissible. The argument $c_h \geq R_h$ uses only that each job j completes no earlier than P_j, which holds also for non-permutation factory schedules. □

Proof of Lemma 2 Let $s_h \leq F$. Put one job of N_h in each of s_h distinct, otherwise empty factories. Each job j then finishes at P_j, so product h is complete at $\max_{j \in N_h} P_j = R_h$. No schedule finishes product h before $\max_{j \in N_h} P_j$. Hence $r_h^* = R_h$. □

Remark 24 For $s_h > F$, r_h^* is the optimum of a distributed permutation flowshop instance, which is NP-hard already for $F = 1$ and $m = 3$, and in general $r_h^* > R_h$.

The Steinitz lemma.

Lemma 5 is the polynomial form of the Steinitz lemma due to Sevast'janov (Sevast'janov 1994a); see also Sevast'janov (1995, Lemma 1). It states: let $v_1, \ldots, v_N \in \mathbb{R}^d$ with $\|v_l\|_\infty \leq 1$ and $\sum_l v_l = 0$; then some order of the vectors has every prefix sum of sup-norm at most d, and such an order can be found in $O(d^2 N^2)$ time. We use it exactly as in Coutinho (manuscript under review).

We use the path formula for permutation flowshops: in a factory with job sequence π, the last-machine completion time $C_m(k)$ of the job at position k is the maximum, over monotone lattice paths from cell $(1,1)$ to cell (m,k), of the sum of the entries of the visited cells.

Lemma 25 (Column groups) *Split a factory sequence into consecutive groups $G_1, \dots, G_r$. For a job at position k in G_r, its last-machine completion time is at most the sum, over the groups before G_r, of each group's makespan when processed alone from an empty factory, plus the completion time of position k inside G_r, also processed alone.*

Proof By the path formula, $C_m(k)$ is the maximum over monotone lattice paths from $(1,1)$ to (m,k) of the sum of the visited cells. Such a path splits into monotone sub-paths, one per group, and every visited cell belongs to exactly one group. Each sub-path in a group before G_r is at most the longest path of that group from row 1 to row m, because entries are non-negative and extending a sub-path upwards (to row 1) or downwards (to row m) only adds cells. The same argument bounds the sub-path in G_r by the completion time of position k inside G_r processed alone. □

Lemma 26 (Start and ramp group) *Fix a factory. Optionally, a start job j_f comes first, with $P_{j_f} \le T_1$ and all its entries $\ge a$. Then come l ramp jobs, all of whose entries are $\le \tilde{a} := a + \delta$, with $\delta \ge 0$. Assume $(m-1)a \le T_1$. Then every monotone path inside the group formed by the start column and the first l ramp columns has value at most $T_1 + l\tilde{a} + (m-1)\delta$, whatever its first and last rows.*

Proof *Case with a start job.* Suppose the path leaves the start column at row i. Its start part is at most $\sum_{i' \le i} p_{i',j_f} = P_{j_f} - \sum_{i' > i} p_{i',j_f} \le T_1 - (m-i)a$. Its ramp part visits at most $l + (m-i)$ cells, each of value at most $\tilde{a}$. The total is at most

$$T_1 - (m-i)a + (l+m-i)(a+\delta) = T_1 + l\tilde{a} + (m-i)\delta \le T_1 + l\tilde{a} + (m-1)\delta.$$

Case without a start job. The path visits at most $l + m - 1$ cells, so its value is at most $(l+m-1)\tilde{a} = l\tilde{a} + (m-1)a + (m-1)\delta \le T_1 + l\tilde{a} + (m-1)\delta$. □

Lemma 27 (Product block) *A factory processes products $g_1, \dots, g_K$ in this order from an empty state, with the jobs of each product contiguous. Let $X_i(q) = \sum_{l \le q} x_{g_l,i}$, let $\bar{x}$ bound every $x_{g_l,i}$, and suppose $|X_i(q) - q\bar{y}_i| \le E_x$ for every q and i, where $\bar{y}_i \ge 0$. Then every job of g_q completes on the last machine by $(q-1)\max_i \bar{y}_i + 2mE_x + m\bar{x}$.*

Proof Take a chain $1 = k_0 \le k_1 \le \cdots \le k_m = k$, where k is a position of g_q, describing a monotone path that runs along machine i over positions $k_{i-1}, \dots, k_i$. Let q_i be the product index of position k_i, so that $q_0 = 1$ and $q_m = q$. Write $Q_i(k)$ for the prefix load of machine i by position. Since products are contiguous, $Q_i(k_i) \le X_i(q_i)$ and $Q_i(k_{i-1}-1) \ge X_i(q_{i-1}-1)$. Hence the path value $\sum_i \big[Q_i(k_i) - Q_i(k_{i-1}-1)\big]$ is at most

$$\sum_{i=1}^{m} \big[X_i(q_i) - X_i(q_{i-1})\big] + \sum_{i=1}^{m} x_{g_{q_{i-1}},i} \le \sum_{i=1}^{m} \big[(q_i - q_{i-1})\bar{y}_i + 2E_x\big] + m\bar{x}$$

$$\leq (q-1)\max_i \bar{y}_i + 2mE_x + m\bar{x}.$$

The path formula completes the proof. □

A.2 Deterministic exactness

Data.

Fix a start set H_0, a ramp set G of $F \cdot R_\ell$ products ($R_\ell \geq 0$ levels), and $\delta \geq 0$. The main set is $H' = H \setminus (H_0 \cup G)$, with $t' = |H'|$. For the main set let $\bar{x}' = \max_{h\in H',i} x_{h,i}$, $\bar{A}' = \max_{h\in H'} A_h$, $\bar{a} = \sum_{h\in H'} A_h/t'$ and $\bar{y}_i = \sum_{h\in H'} x_{h,i}/t'$. Put $T_1 = \max_{h\in H_0} R_h$, $C_0 = \sum_{h\in H_0} A_h$, $\tilde{a} = a + \delta$ and $\alpha = \min_{g\in G} A_g$. The constants are $K_1 = m(4m+5)\bar{x}'$ and $E_A = 2(m+1)\bar{A}'$.

Assumption 28 (Hypotheses of Theorem 6)

(H1) $\sum_{h\in H_0} s_h \leq F$, *and some* $h_1 \in H_0$ *has* $R_{h_1} = T_0$.

(H2) For every $h \in H_0$: $\sum_{g\in H_0:\, R_g < R_h} A_g \geq R_h - T_0$.

(H3) Every $g \in G$ *is a single job with all entries* $\leq \tilde{a}$. *If* $R_\ell \geq 1$, *also* $\tilde{a} > 0$, $F\alpha \geq \tilde{a}$, *and* $C_0 \geq (T_1 - T_0) + \tilde{a} + (m-1)\delta$.

(H4) (Margin) $F\sum_{h\in H'} A_h \geq \max_i \sum_{h\in H'} x_{h,i}$.

(H5) (Credit) $C_0 - (T_1 - T_0) - (m-1)\delta + R_\ell(F\alpha - \tilde{a}) \geq K_1 + F \cdot E_A$.

Definition 29 (Start–ramp–Steinitz (SRS) schedule)

1. Put the jobs of H_0 at position 1, in distinct factories (possible by (H1)).
2. In each factory, add R_ℓ ramp jobs; level $l \in \{1, \ldots, R_\ell\}$ holds F ramps, one per factory.
3. Order H' by Lemma 5 applied to the vectors $w_g = \big((x_{g,i} - \bar{y}_i)/\bar{x}'\big)_{i\leq m}$ together with the coordinate $(A_g - \bar{a})/\bar{A}'$. Here $d = m+1$, each vector lies in $[-1,1]^d$, and the vectors sum to 0.
4. Cut this order into F consecutive blocks whose lengths differ by at most one. Block f follows the ramps in factory f, with the jobs of each product contiguous.

Proof of Theorem 6 Block deviations. Every block prefix is the difference of two global prefixes of the Steinitz order. Each global prefix deviates from its mean by at most $d = m+1$ in normalised units, so, within each block, $|X_i(q) - q\bar{y}_i| \leq 2(m+1)\bar{x}' =: E_x$ and $\big|\sum_{l\leq q} A_{g_l} - q\bar{a}\big| \leq 2(m+1)\bar{A}' = E_A$. Note that $2mE_x + m\bar{x}' = K_1$.

Upper envelope $\hat{c} \geq c$. We define $\hat{c}$ and show $c \leq \hat{c}$.

- Products in H_0: $\hat{c}_h = R_h$ exactly, since every job of h is at position 1 of an empty factory.
- Ramps of level l: $\hat{c} = T_1 + l\tilde{a} + (m-1)\delta$, by Lemma 26. Its assumption $(m-1)a \leq T_1$ holds because $T_1 \geq T_0 \geq ma$.
- Main product at level q of a block: $\hat{c} = T_1 + R_\ell\tilde{a} + (m-1)\delta + (q-1)\max_i \bar{y}_i + K_1$, by Lemma 25 (with the start-and-ramp group and the block as groups), Lemma 26 and Lemma 27.

Queue condition. We show

$$\text{(Q)} \qquad \sum_{g:\,\hat{c}_g<\hat{c}_h} A_g \geq \hat{c}_h - T_0 \quad \text{for every } h.$$

By Lemma 1, (Q) gives $\mathrm{ERD}(\hat{c}) = \max_h \left[\hat{c}_h + A_{\mathrm{tot}} - \sum_{g:\,\hat{c}_g<\hat{c}_h} A_g\right] \leq T_0 + A_{\mathrm{tot}}$. We check (Q) in three cases.

- $h \in H_0$. This is exactly (H2), because products outside H_0 released before h can only add credit.
- Ramp level l. The products released earlier include all of H_0 and the $F(l-1)$ ramps of lower levels, so the credit is at least $C_0 + (l-1)F\alpha$. The requirement is $T_1 - T_0 + l\tilde{a} + (m-1)\delta$. It follows from (H3), because $(l-1)F\alpha \geq (l-1)\tilde{a}$.
- Main level q. The products released earlier include all of H_0, all ramps and, in every block, the first $q-1$ products. Every block has at least $q-1$ products, because $q-1 \leq n_{\max} - 1 \leq n_{\min}$, where $n_{\max}$ and $n_{\min}$ are the largest and smallest block lengths. The credit is therefore at least $C_0 + FR_\ell\alpha + F\left[(q-1)\bar{a} - E_A\right]$, and the requirement is $T_1 - T_0 + R_\ell\tilde{a} + (m-1)\delta + (q-1)\max_i \bar{y}_i + K_1$. (H4) says $F\bar{a} \geq \max_i \bar{y}_i$, so (H5) suffices. At $q = 1$ the term $-F \cdot E_A$ is absent.

Conclusion. By Lemma 1, the SRS makespan is $\mathrm{ERD}(c) \leq \mathrm{ERD}(\hat{c}) \leq \mathrm{LB}_0 \leq X \leq \mathrm{OPT} \leq$ makespan, for every admissible X. Since LB_0 is valid for every schedule (Corollary 23), the equality holds also over non-permutation factory schedules. The SRS schedule is computed in $O(m^2 t'^2 + nm)$ time plus the linear programs of Lemma 5, so in polynomial time, and its makespan compared with LB_0 is the certificate.

Degenerate cases. If $\bar{x}' = 0$ or $\bar{A}' = 0$, drop that coordinate. If $H' = \emptyset$, (H4) and (H5) are void. □

Remark 30 (One-factory-per-product variant) Some variants force all jobs of a product into one factory. Under that variant, Theorem 6 holds if H_0 is restricted to single-job products; Lemma 2 fails for $s_h \geq 2$. Theorem 9 holds, because its H_0 and ramps are single-job products and the SRS schedule keeps every H' product inside one block. Theorem 10 holds. Proposition 12 holds and only gets stronger. Theorem 8 and Proposition 3 hold, since the schedule of Definition 41 keeps products contiguous in one factory and the bound does not use splitting.

Remark 31 (Relation to element-wise dominance) The polynomial cases of Lee et al. (1993) for the three-machine assembly flowshop (one job per product, two dedicated first-stage machines, one assembly machine, no factories) are sequencing algorithms under job-wise dominance; none states that OPT equals a lower bound, and job-wise dominance does not give $\mathrm{OPT} = \mathrm{LB}_0$. For example, with $F = m = 1$, two single-job products with $P = A = 1$ and $P = A = 10$ satisfy their condition $p_{i2} \geq \max(p_{ia}, p_{ib})$, yet $\mathrm{LB}_0 = 1 + 11 = 12$ while $\mathrm{OPT} = 21$. The aggregate margin (H4), the start conditions (H1)–(H2) and the ramp credit (H3), (H5) are therefore not implied by those cases.

Proof of Corollary 7 Replace H_0 by $\{h_1\}$, scheduled by an optimal sub-schedule of N_{h_1} on the F factories, so that $T_0 := r^*_{h_1}$; put $C_0 = A_{h_1}$ and $T_1 = T_0$. Assume $r^*_{h_1} \leq \underline{r}_h$ for all h. The requirement of Lemma 26 that the last job of each start group has all entries $\geq a$ holds automatically, because a is the global minimum entry.

The proof of Lemma 26 holds with the start group in place of the start column: a path that leaves the start group at row i has value at most $C_i^{\mathrm{grp}}(\mathrm{last}) \le C_m^{\mathrm{grp}}(\mathrm{last}) - (m-i)a \le T_1 - (m-i)a$, where C^{grp} are the completion times within the start group. (H1) and (H2) hold trivially for the single start product. The rest of the proof of Theorem 6 is unchanged, and (H3)–(H5) give a makespan of at most $r^*_{h_1} + A_{\mathrm{tot}}$.

Conversely, for every admissible ERD($\underline{r}$) with $\underline{r}_{h_1} = r^*_{h_1}$ and $\underline{r}_h \ge r^*_{h_1}$ for all h, Lemma 1 gives $\mathrm{ERD}(\underline{r}) \ge \min_h \underline{r}_h + A_{\mathrm{tot}} = r^*_{h_1} + A_{\mathrm{tot}}$, and Corollary 23 gives $\mathrm{ERD}(\underline{r}) \le \mathrm{OPT}$. Hence $\mathrm{OPT} = r^*_{h_1} + A_{\mathrm{tot}} = \mathrm{ERD}(\underline{r})$. Without the condition $\underline{r}_h \ge r^*_{h_1}$ for all h, ERD($\underline{r}$) can be smaller than OPT; for example, $\underline{r} = r^*$ qualifies if and only if h_1 minimises r^*. □

A.3 Probabilistic exactness for $\beta > 1$

Assumption 32 (Model M) *Products are i.i.d. The size S of a product follows a law ν on $\{1, \dots, \bar{s}\}$, with $\bar{s}$ fixed. Given the size, the S jobs have i.i.d. entries with law μ on $[a, b]$, where $a = \min \operatorname{supp} \mu < b = \max \operatorname{supp} \mu$. The assembly time is $A_h \sim \kappa_S$ on $[0, \bar{A}]$, independent of the entries given the size. Let $\mathbb{E}p$ be the mean of μ; the load ratio is $\beta := F\,\mathbb{E}[A]/(\mathbb{E}[S]\,\mathbb{E}p)$. The assumptions are:*

(M1) $F \ge 2$ *and* $\nu(1) > 0$.

(M2) Some α *with* $(F-1)\alpha > a$ *has* $\kappa_1([\alpha, \bar{A}]) > 0$.

(M3) Either $\mu(\{a\}) = q_a > 0$ *(the atom case), or* $\kappa_1(\{0\}) = 0$.

(M4) $\beta > 1$.

Proof of Theorem 9 We prove that $\Pr[\mathrm{OPT} = X \text{ for every admissible } X] \to 1$ as $t \to \infty$ and that the SRS certificate exists with the same probability, and, in the atom case, the quantitative bound

$$\begin{aligned}\Pr[\text{failure}] &\le \Pr\big(\mathrm{Bin}(t, q_*) < F\big) + \Pr\big(\mathrm{Bin}(t, q_R) < F(R_\ell + 1)\big) \\ &\quad + m \exp\Big(-\frac{2(t\mu_\xi - c_0)^2}{t(F\bar{A} + \bar{s}b)^2}\Big) = e^{-\Omega(t)}, \end{aligned} \tag{A1}$$

where $q_* = \nu(1)\, q_a^m\, \kappa_1([\alpha, \bar{A}])$, $q_R = \nu(1)\, \mu([a, a+\delta])^m\, \kappa_1([\alpha, \bar{A}])$, $\mu_\xi = \mathbb{E}[S]\,\mathbb{E}p\,(\beta - 1)$, $c_0 = F^2(R_\ell + 1)\bar{A}$, $\gamma = F\alpha - a - \delta > 0$ and $R_\ell = \lceil \big(m(4m+5)\bar{s}b + 2F(m+1)\bar{A}\big)/\gamma \rceil$.

Choice of δ. By (M2), pick $\delta > 0$ with $(F-1)\alpha > a + m\delta + \delta$. Then $\gamma = F\alpha - \tilde{a} > 0$, with $\tilde{a} = a + \delta$.

Events and failure terms.

- E_{margin}: $F \cdot A_{\mathrm{tot}} - \max_i L_i \ge c_0$. For each machine i, $\sum_h (FA_h - x_{h,i})$ is a sum of t i.i.d. terms in $[-\bar{s}b, F\bar{A}]$ with mean $\mu_\xi > 0$ (by (M4)). Hoeffding's inequality (Hoeffding 1963) and a union bound over the m machines give the third term of (A1).
- E_{ramp}: at least $F(R_\ell + 1)$ products have size 1, all entries $\le a + \delta$ and $A \ge \alpha$. Their number is $\mathrm{Bin}(t, q_R)$ with $q_R > 0$, because $a \in \operatorname{supp} \mu$. This gives the second term.
- Atom case: at least F products belong to the class Q_a of products of size 1 with all entries equal to a and $A \ge \alpha$; their number is $\mathrm{Bin}(t, q_*)$, which gives the first term. Then $T_0 = ma$. Take $h_1 \in Q_a$ and $F - 1$ further members of Q_a as H_0. All of them have $R = T_0$, so (H2) holds trivially, and $T_1 = T_0$.
- General case ($\kappa_1(\{0\}) = 0$): let $\Phi(x) = \Pr(P_j \le x)$, which is > 0 for every $x > ma$, and choose $\varepsilon_t \downarrow 0$ with $\varepsilon_t \le \delta$ and $t\,\Phi(ma + \varepsilon_t) \to \infty$. The events are:

(i) Some size-1 product attains T_0. Its complement requires a product of size ≥ 2 two of whose jobs both lie strictly below all n_1 totals of the size-1 products. By exchangeability this has probability at most $2/\big((n_1+1)(n_1+2)\big)$ per product, and $n_1 \sim \mathrm{Bin}(t, \nu(1))$, so the total failure probability is $O(1/t)$.

(ii) $A_{h_1} \geq \varepsilon_t$. Choose h_1 (a size-1 product attaining T_0) by a rule that uses only entries and sizes; then $A_{h_1} \sim \kappa_1$, and the failure probability is $\kappa_1([0, \varepsilon_t)) \to 0$.

(iii) At least F products have size 1, $P \leq ma + \varepsilon_t$ and $A \geq \alpha$. This is a binomial count with mean $t\,\nu(1)\,\Phi(ma+\varepsilon_t)\,\kappa_1([\alpha, \bar{A}]) \to \infty$, so it fails with probability $\to 0$. Take as H_0 the product h_1 and $F-1$ of these products other than h_1.

Hypotheses of Theorem 6 on these events. The constant $a = \min \operatorname{supp} \mu$ bounds every start entry from below. Take the ramps from the E_{ramp} class, disjoint from H_0; since $|H_0| = F$, at least FR_ℓ remain.

- (H1) holds: H_0 consists of F single-job products and contains h_1.
- (H2): in the atom case it is trivial; in the general case every extra product has $R - T_0 \leq \varepsilon_t \leq A_{h_1}$.
- (H3): ramps are single jobs with entries $\leq \tilde{a}$, $\tilde{a} > 0$, $F\alpha \geq \tilde{a}$ (as $\gamma > 0$), and $C_0 \geq (F-1)\alpha > a + m\delta + \varepsilon_t \geq (T_1 - T_0) + \tilde{a} + (m-1)\delta$, using $T_1 - T_0 \leq \varepsilon_t \leq \delta$ (and $T_1 = T_0$ in the atom case).
- (H4): $F \cdot A(H') \geq F \cdot A_{\mathrm{tot}} - F \cdot |H_0 \cup G| \cdot \bar{A} \geq F \cdot A_{\mathrm{tot}} - c_0 \geq \max_i L_i \geq \max_i X_i(H')$ on E_{margin}, because $|H_0 \cup G| \leq F(R_\ell + 1)$.
- (H5): by (H3) the left side is at least $\tilde{a} + R_\ell \gamma \geq K_1 + F \cdot E_A$, since $\bar{x}' \leq \bar{s}b$ and $\bar{A}' \leq \bar{A}$, and by the choice of R_ℓ.

Theorem 6 then gives $\mathrm{OPT} = X$ for every admissible X, certified by the SRS schedule, on the intersection of the events. □

Remark 33 (The generator of Hatami et al. (2013)) In the benchmark of Hatami et al. (2013), $\mu = U\{1, \ldots, 99\}$ (an atom at $a = 1$) and $A_h \sim U\{s_h, \ldots, 99 s_h\}$, which is the law $\kappa_s = \mathrm{round}(U[s, 99s] \cdot \beta/F)$ at $\beta/F = 1$. These instances therefore have model ratio $\beta = F \geq 2$ (assembly-dominant), and $\kappa_1 \geq 1$ with an atom at 1, so (M2) and (M3) hold. The regime is exponential, but $q_* \approx \nu(1)\, 99^{-m}$, so (A1) becomes informative only for $t \gg 99^m$. The benchmark sizes are 1 plus a multinomial rather than i.i.d. bounded; that extension is open.

Corollary 34 (Expectation) *Under Assumption 32,* $\mathbb{E}[\mathrm{OPT} - X] \to 0$ *for every admissible* X.

Proof By Theorem 10, $\mathrm{OPT} - X \leq \mathrm{OPT} - \mathrm{LB}_0 \leq C + (\max_i L_i/F - A_{\mathrm{tot}})^+$, with a deterministic constant C (recall $\mathrm{LB}_0 \geq A_{\mathrm{tot}}$). On the success event of Theorem 9, $\mathrm{OPT} = X$. Hence $\mathbb{E}[\mathrm{OPT} - X] \leq C \cdot \Pr(\mathrm{fail}) + \mathbb{E}\big[(\max_i L_i/F - A_{\mathrm{tot}})^+\big]$. Both terms vanish: the first by Theorem 9, and the second is exponentially small for $\beta > 1$ by Hoeffding's inequality and a union bound over machines. □

The rate outside the atom case is the subject of Section A.9.

A.4 Additive guarantee

Proof of Theorem 10 Apply the SRS schedule with $H_0 = G = \emptyset$ ("Steinitz blocks only"): all products are ordered by Lemma 5 as in Definition 29 and cut into F blocks, with $\bar{x}$, $\bar{A}$, $\bar{a} = A_{\mathrm{tot}}/t$ and $\bar{y}_i = L_i/t$ computed over all products, $K_1 = m(4m+5)\bar{x}$ and $E_A = 2(m+1)\bar{A}$. By Lemma 27 and the block deviations in the proof of Theorem 6, a product at level q of a block has $c \le \hat{c}(q) := (q-1)\max_i \bar{y}_i + K_1$. The products with $\hat{c} \ge \hat{c}(q)$ are those of levels $\ge q$. For $q \ge 2$ their assembly work is at most $A_{\mathrm{tot}} - F(q-1)\bar{a} + F \cdot E_A$; for $q = 1$ it is A_{tot}. Lemma 1 then gives

$$M \le \mathrm{ERD}(\hat{c}) \le K_1 + F \cdot E_A + A_{\mathrm{tot}} + \max\bigl(0, (n_{\max} - 1)(\max_i \bar{y}_i - F\bar{a})\bigr),$$

where $n_{\max} = \lceil t/F \rceil$ is the largest block length. Finally $n_{\max} - 1 \le t/F$, and $(t/F)(\max_i \bar{y}_i - F\bar{a}) = \max_i L_i/F - A_{\mathrm{tot}}$. Hence

$$M \le \max\bigl(A_{\mathrm{tot}}, \max_i L_i/F\bigr) + m(4m+5)\bar{x} + 2F(m+1)\bar{A},$$

and $\mathrm{OPT} - X \le C(m, F, \bar{x}, \bar{A}) := m(4m+5)\bar{x} + 2F(m+1)\bar{A}$ for every lower bound $X \ge \max(A_{\mathrm{tot}}, \max_i L_i/F)$. □

Remark 35 (Bounds covered) The bound $\max(\mathrm{LB}_A, \mathrm{LB}_F, \mathrm{LB}_J)$ of Coutinho (unpublished manuscript) is at least $\max(A_{\mathrm{tot}}, \max_i L_i/F)$: its factory-load term $\mathrm{LB}_F = \max\bigl(\max_k[\min_j \mathrm{hd}_{jk} + \lceil L_k/F \rceil + \min_j \mathrm{tl}_{jk}], \max_j P_j\bigr) + \min_h A_h \ge \max_k L_k/F$, and its ERD term is at least $\min_h r_h + A_{\mathrm{tot}} \ge A_{\mathrm{tot}}$. So Theorem 10 applies to it. The constant is free of n only for bounded $\bar{s}$; under the size law of Hatami et al. (2013) (1 plus a multinomial) it is $o(\sqrt{n})$.

Remark 36 (Relation to prior additive guarantees) For $Am||C_{\max}$ (one job per product, m dedicated component machines, no factories), Potts et al. (1995, Thm. 7) obtain $C_{\max} - \mathrm{OPT} \le (m + 1/m)p^*$ in $O(m^2n^2)$ time, by padding the loads to a common P^* and applying compact vector summation. For $AF(m, 1, 1)//C_{\max}$, Koulamas and Kyparisis (2001, Thm. 5) obtain $C_{\max} \le P^* + 2(m + 1 + 1/(m+1))p^*$ with the same padding and Sevast'janov's theorem, and note asymptotic optimality in ratio. For the three-machine assembly flowshop, Lee et al. (1993, Thm. 4, Cor. 4) show that every non-delay permutation schedule has $C \le T_{\max} + T_2 \le C^* + \min\{T_{\max}, T_2\} \le 2C^*$, a ratio-type ancestor whose error term is a whole machine load. Theorem 10 extends this additive line to the distributed assembly flowshop, where jobs are assigned to F factories and grouped into products; the constant scales with the largest product load $\bar{x}$ and with $F\bar{A}$, not with p^*. The method (vector summation) is theirs.

A.5 Start obstruction

Recall the deterministic hypotheses of Proposition 12: $t \ge 2$; h_1 is the unique minimiser of R; $|N_{h_1}| \ge F$; a is the minimum entry; $P_{\min} = \min_{j \in N_{h_1}} P_j$, $D = T_0 - P_{\min}$ and $T_0' = \min_{h \ne h_1} R_h$. The claim is $\mathrm{OPT} \ge \mathrm{LB}_0 + \max\{0, \min(T_0' - T_0,\, a - D - A_{h_1})\}$. The probabilistic part assumes model M (Assumption 32) with $\nu(\{1, \dots, F-1\}) = 0$, μ without atoms, $a > 0$, and $\kappa_{\min} := \min_{s \in \mathrm{supp}\,\nu} \kappa_s([0, a/2)) > 0$, and claims $\liminf_t \Pr[\mathrm{OPT} > \mathrm{LB}_0] \ge \kappa_{\min}$.

Proof of Proposition 12 Deterministic part. Fix a schedule. *Key claim: either $c_{h_1} \ge P_{\min} + a$, or every job outside N_{h_1} completes at or after $P_{\min} + a$.*

- If one factory holds two jobs of h_1, the later of them completes at or after $P_{\min} + a$ (it completes at least $p_{mj} \geq a$ after the earlier one, which completes at or after $P_{\min}$).
- Otherwise, since $|N_{h_1}| \geq F$, each factory holds exactly one job j_f of h_1. A job j' placed before j_f forces $C(j_f) \geq p_{1,j'} + P_{j_f} \geq P_{\min} + a$, so $c_{h_1} \geq P_{\min} + a$.
- If no job is placed before any j_f, every other job j' is placed after some j_f and completes at or after $C(j_f) + p_{m,j'} \geq P_{\min} + a$.

Case analysis on the first assembled product.

- If it is not h_1, it starts at or after T_0', so the makespan is at least $T_0' + A_{\text{tot}}$.
- If it is h_1 and $c_{h_1} \geq P_{\min} + a$, the makespan is at least $P_{\min} + a + A_{\text{tot}}$.
- Otherwise, by the key claim, the second assembled product (which exists since $t \geq 2$) starts at or after $P_{\min} + a$, so the makespan is at least $P_{\min} + a + A_{\text{tot}} - A_{h_1}$.

Since $P_{\min} + a - T_0 = a - D$, every schedule has makespan at least $\text{LB}_0 + \min(T_0' - T_0,\, a - D - A_{h_1})$, and also at least LB_0. This proves the deterministic part. The hypothesis $t \geq 2$ is necessary: for $t = 1$ the claim fails on small instances.

Probabilistic part. Every product has at least F jobs. $T_0 \to ma$ in probability, since for every $\varepsilon > 0$, with probability tending to one some product has all of its jobs with $P \leq ma + \varepsilon$. The R_h are almost surely distinct (μ has no atom), so h_1 is the unique minimiser and $T_0' > T_0$. Moreover $P_{\min} \geq ma$, so $D \leq T_0 - ma$, and hence $a - D - A_{h_1} > 0$ on the event $\{T_0 - ma < a/2,\ A_{h_1} < a/2\}$. On this event the deterministic part gives $\text{OPT} > \text{LB}_0$. Since $A_{h_1} \sim \kappa_{s_{h_1}}$ independently of the entries given the sizes, $\Pr(A_{h_1} < a/2 \mid \text{sizes, entries}) \geq \kappa_{\min}$, and $\Pr(T_0 - ma < a/2) \to 1$. Hence $\liminf_t \Pr[\text{OPT} > \text{LB}_0] \geq \kappa_{\min}$. □

Remark 37 (Scope) If no product has fewer than F jobs, and assembly times can fall below $a/2$ ($\kappa_{\min} > 0$), then LB_0 fails with probability bounded away from 0, for every β (atomless μ, $a > 0$). This does not show that $\nu(1) > 0$ in (M1) is necessary. The intermediate regime $\nu(1) = 0 < \nu(\{2, \dots, F-1\})$ is covered neither by Theorem 9 (which needs a single-job h_1) nor by Proposition 12, and it is open. That the credit conditions matter is shown only through $\kappa_s([0, a/2)) > 0$.

A.6 The regime $\beta \leq 1$ and the decomposition bounds

Proof of Proposition 11 Deterministic part. For every $r \leq r^*$, Lemma 1 gives $\text{ERD}(r) \leq \max_h r_h + A_{\text{tot}}$, and $r_h^* \leq \bar{s}mb$, since scheduling the at most $\bar{s}$ jobs of h in one factory finishes them by the sum of their entries. Hence $\text{ERD}(r) \leq \bar{s}mb + A_{\text{tot}}$. Also $\text{OPT} \geq \max_i L_i/F$, since some factory carries at least L_i/F on machine i. Hence $\text{OPT} - \text{LB}^* \geq \max_i L_i/F - A_{\text{tot}} - \bar{s}mb$.

$\beta < 1$. Let $c = \mathbb{E}[S]\,\mathbb{E}p\,(1-\beta)/(2F)$. The sum $\sum_h (x_{h,1}/F - A_h)$ is a sum of t i.i.d. bounded terms with mean $t\,\mathbb{E}[S]\,\mathbb{E}p\,(1-\beta)/F = 2ct$. Hoeffding's inequality gives, for some $c' > 0$, $\Pr\big(L_1/F - A_{\text{tot}} < \tfrac{3}{2}ct\big) \leq e^{-c't}$; on the complement, $\text{OPT} - \text{LB}^* \geq \tfrac{3}{2}ct - \bar{s}mb \geq ct$ for $t \geq 2\bar{s}mb/c$. Hence $\Pr[\text{OPT} - \text{LB}^* < ct] \leq e^{-c't}$ for all large t, so $\Pr[\text{OPT} = \text{LB}^*] \to 0$ exponentially and the gap is $\Theta(t)$ (it is $O(t)$ by Theorem 10).

$\beta = 1$, *with* $\text{Var}(x_{h,1}/F - A_h) = \sigma^2 > 0$. The sum $\Sigma_t = \sum_h (x_{h,1}/F - A_h)$ has mean 0 and variance $t\sigma^2$, and its terms are bounded. By the Berry–Esseen theorem (Shevtsova 2011), $\Pr(\Sigma_t > \bar{s}mb) \to 1/2$ and $\liminf_t \Pr(\Sigma_t \geq c\sqrt{t} + \bar{s}mb) > 0$ for some $c > 0$. Since $\text{OPT} - \text{LB}^* \geq \Sigma_t - \bar{s}mb$ and $\text{OPT} - \text{LB}^* \geq 0$, this gives $\liminf \Pr[\text{OPT} > \text{LB}^*] \geq 1/2$, $\liminf_t \Pr[\text{OPT} - \text{LB}^* \geq c\sqrt{t}] > 0$, and $\mathbb{E}[\text{OPT} - \text{LB}^*] \geq \mathbb{E}[(\Sigma_t - \bar{s}mb)^+] \geq c\,\sigma\sqrt{t}$ for large t. Conversely, $\text{LB}^* \geq A_{\text{tot}}$, so Theorem 10 gives $\text{OPT} - \text{LB}^* \leq C + (\max_i L_i/F - A_{\text{tot}})^+ =$

$O_P(\sqrt{t})$, and its expectation is $O(\sqrt{t})$ by the same central-limit bounds applied to each machine. Precisely: $\mathbb{E}[\mathrm{OPT} - \mathrm{LB}^*] = \Theta(\sqrt{t})$, $\mathrm{OPT} - \mathrm{LB}^* = O_P(\sqrt{t})$, and $\liminf_t \Pr[\mathrm{OPT} - \mathrm{LB}^* \geq c\sqrt{t}] > 0$. This is not a statement that the gap is at least $c\sqrt{t}$ with high probability, and $\sigma^2 > 0$ is needed. □

Remark 38 (Combined bound) For the combined bound, Theorem 10 gives $\mathrm{OPT} - \max(\mathrm{LB}_A, \mathrm{LB}_F) \leq C$ at every β. A $\sqrt{t}$ gap at $\beta = 1$ is therefore a property of the decomposition bound only, not of the combined bound.

A.7 The covering end-chain bound

Recall the definition. For $r = 0, \ldots, F$ let $W(r) = \min\{\sum_{g\in\mathcal{C}} A_g : \mathcal{C}$ a set of products with $\sum_{g\in\mathcal{C}} s_g \geq r\}$, with $W(0) = 0$ and $W(r) = \infty$ when r exceeds the total number of jobs (no cover exists; such F' terms drop out of the minimum below, so effectively $F' \leq \min(F, n)$). This is a covering knapsack, computed by dynamic programming in $O(tF)$ time. For each machine k let $\mathrm{hd}_k = \min_j \sum_{i<k} p_{ij}$ and $\mathrm{tl}_k = \min_j \sum_{i>k} p_{ij}$, and

$$\mathrm{LB}_{\mathrm{end}} = \max_k \min_{1\leq F'\leq F} \Big[\mathrm{hd}_k + \mathrm{tl}_k + \frac{L_k + \sum_{r=1}^{F'} W(r)}{F'}\Big],$$

rounded up for integer data.

Proof of Proposition 3 We show $\mathrm{LB}_{\mathrm{end}} \leq \mathrm{OPT}$, also over non-permutation factory schedules and under the one-factory-per-product variant. Fix a schedule and a machine k. Let F' be the number of factories that process at least one job. In such a factory f, let j_f be the last job processed on machine k, and let $\mathrm{load}_k(f)$ be the machine-k load of f. Then $C_{j_f} \geq \mathrm{hd}_k + \mathrm{load}_k(f) + \mathrm{tl}_k$: the first job processed on machine k cannot start there before hd_k, and j_f still needs at least tl_k after machine k. Order these factories so that $C_{j_{f_1}} \geq \cdots \geq C_{j_{f_{F'}}}$.

Fix $r \leq F'$. The products containing $j_{f_1}, \ldots, j_{f_r}$ are all assembled at or after $C_{j_{f_r}}$. These r jobs are distinct, so the sizes of these products sum to at least r, and their assembly takes at least $W(r)$. Hence makespan $\geq C_{j_{f_r}} + W(r) \geq \mathrm{hd}_k + \mathrm{load}_k(f_r) + \mathrm{tl}_k + W(r)$. Averaging over $r = 1, \ldots, F'$ gives

$$\text{makespan} \geq \mathrm{hd}_k + \mathrm{tl}_k + \frac{\sum_f \mathrm{load}_k(f) + \sum_{r\leq F'} W(r)}{F'},$$

and $\sum_f \mathrm{load}_k(f) = L_k$. Taking the minimum over the unknown F' and the maximum over k gives the claim; with integer data the makespan is an integer, so the bound may be rounded up. The argument uses neither the permutation property nor splitting of products. □

Remark 39 $W(r) \geq S_{\lceil r/\bar{s}\rceil}$, the sum of the $\lceil r/\bar{s}\rceil$ smallest assembly times, so $\mathrm{LB}_{\mathrm{end}}$ dominates the simpler version with W replaced by these sums. With $F' = F$ and only the $r = 1$ term, $\mathrm{LB}_{\mathrm{end}}$ is the factory-load bound of Coutinho (unpublished manuscript) up to the rounding of L_k/F. Because of the minimum over F', $\mathrm{LB}_{\mathrm{end}}$ does not dominate the bound of Coutinho (unpublished manuscript); the two are combined by taking their maximum.

Remark 40 (Credit) For the plain distributed permutation flowshop, Gogos (2023, Sec. 4.4, Alg. 1) gives LB $= \max_i[\min \text{ head} + L_i/F + \min \text{ tail}]$, which is our core term with $W \equiv 0$ and $F' = F$. In the three-machine assembly flowshop, Lee et al. (1993, Thm. 2, Cor. 1) show that a job whose assembly time is at most every remaining processing time can be put last (adjacent interchange); this is the sequencing intuition behind the end side, but it is a dominance rule, not a bound. The averaging-with-tails argument is classical for parallel machines with heads and tails, $P|r_j, q_j|C_{\max}$. New here are the covering term $W(r)$, which accounts for products that can end several factories at once, the averaging over the r latest factory ends, and the minimum over the unknown number F' of used factories.

A.8 Factory-dominant exactness certificate

Machines are numbered $1, \ldots, m$; $a < b$ are the global minimum and maximum entries. Up-ramps and down-ramps for a machine r are as in Coutinho (manuscript under review): an *up-ramp* is a single-job product with $p_i = a$ for $i < r$ and $p_r = b$; a *down-ramp* is a single-job product with $p_r = b$ and $p_i = a$ for $i > r$.

Definition 41 (RSE schedule) Fix a machine r. Each factory f gets a sequence $\pi_f = U_f \cdot \mathrm{Mid}_f \cdot \mathrm{Dn}_f$ of n_f jobs:

- U_f holds $\beta_\uparrow$ up-ramps; we require $\beta_\uparrow \geq 1$ when $r \geq 2$, and U_f is empty if $r = 1$;
- Mid_f holds the other products, with each product's jobs contiguous;
- Dn_f holds $\beta_\downarrow \geq 1$ down-ramps with $A \leq \alpha^*$; its last product is the end product e_f.

ρ is a bijection from factories to ranks $1, \ldots, F$, and $S_\rho = \sum_{j \leq \rho} A_{e(j)}$, where $e(j)$ is the end product of the rank-j factory. M is a number (the target).

In factory f, $Q_i(l)$ is the prefix load of machine i up to position l of π_f, $L_i(f) = Q_i(n_f)$, $D_i = Q_i - Q_r$, $\Delta_i = L_r(f) - L_i(f)$, and $\bar{D}_i(l) = \sum_{l' > n_f - l}(p_{i,\pi_f(l')} - p_{r,\pi_f(l')}) = -\Delta_i - D_i(n_f - l)$ for $i > r$. Following the bottleneck decomposition of Coutinho (manuscript under review), let $\theta_k = C_{r-1}(k) - Q_r(k-1)$ (with $C_0 \equiv 0$) and $\omega_k = R_{r+1}(k) - (L_r(f) - Q_r(k))$, where $R_{r+1}(k)$ is the longest path value from cell $(r+1, k)$ to (m, n_f); then $C_f = L_r(f) + \max_{k \leq k'}(\theta_k + \omega_{k'})$.

Assumption 42 (Hypotheses of Theorem 8)

(T1) In every factory f the start and end block conditions of Coutinho (manuscript under review) hold, with constants $E_\uparrow, E_\downarrow \geq 0$ that are the same in every factory (for example, the maximum over factories):

(U1) $D_i(l) \leq -W_\uparrow$ *for all* $i < r$ *and* $l > \beta_\uparrow$;
(U2) $D_i(l) - D_i(s) \leq E_\uparrow$ *for all* $i < r$ *and* $\beta_\uparrow \leq s < l \leq n_f$;
(D1) $\bar{D}_i(l) \leq -W_\downarrow$ *for all* $i > r$ *and* $l > \beta_\downarrow$;
(D2) $\bar{D}_i(l) - \bar{D}_i(s) \leq E_\downarrow$ *for all* $i > r$ *and* $\beta_\downarrow \leq s < l \leq n_f$;

where $W_\uparrow = (r-2)^+ E_\uparrow + (r-1)(b-a)$ *and* $W_\downarrow = (m-r-1)^+ E_\downarrow + (m-r)(b-a)$. *(U1)–(U2) are void if* $r = 1$, *and (D1)–(D2) are void if* $r = m$.

(T2) (Targets) $L_r(f) \leq M - (m-1)a - S_{\rho(f)}$ *for every* f.

(T3) (Assembly margin along suffixes) For every factory f and every product-aligned suffix Z of $U_f \cdot \mathrm{Mid}_f$, $A(Z) \le x_r(Z)/F + E_A$.

(T4) $F\alpha^* < b$. Put $\gamma = b - F\alpha^*$, $\Lambda_0 = (m-r)(b + E_\downarrow - a)$, $k_0 = \lceil \Lambda_0/(b-a) \rceil$, and let $\bar{x}_r$ be the largest machine-r load of a product in $U \cup \mathrm{Mid}$.

(T5) (a) $\beta_\downarrow \gamma \ge \Lambda_0 + \bar{x}_r + F \cdot E_A$, and (b) $(\beta_\downarrow - k_0)\gamma \ge \Lambda_0 + F\alpha^*$.

Proof of Theorem 8 We show that the RSE makespan is at most M. Hence, if $M \le X$ for some valid lower bound X (for example $M = \lceil \mathrm{LB}_{\mathrm{end}} \rceil$ with integer data; $\mathrm{LB}_{\mathrm{end}}$ need not be $\ge \mathrm{LB}_0$, so admissibility is not required), then $\mathrm{OPT} = M = X$, certified in polynomial time by "makespan = bound".

Step 1 (factory makespans). By the bottleneck decomposition and the start and end block lemmas of Coutinho (manuscript under review), (T1) gives $\theta_k \le (r-1)a$ and $\omega_k \le (m-r)a$ for every k, so $C_f \le L_r(f) + (r-1)a + (m-r)a = (m-1)a + L_r(f)$. Conversely, $C_f \ge (r-1)a + L_r(f) + (m-r)a$, since a is the minimum entry. So equality holds. Machine r therefore has no idle time after $(r-1)a$, and $C_r(l) = (r-1)a + Q_r(l)$ for every position l. (This identity also follows directly from the start block lemma: $C_r(l) = Q_r(l) + \max_{k \le l} \theta_k$.)

Step 2 (prefix envelope). For $r < m$, every machine $i > r$ and every position l:

$$C_i(l) \le (r-1)a + Q_r(l) + (i-r)(b + E_\downarrow).$$

A critical path to (i,l) leaves row $r-1$ at column k_1, which contributes $\theta_{k_1} \le (r-1)a$ (start block lemma). It runs along row r to column k_2, and then through rows $r+1, \dots, i$ to column l. Write $Q_s = Q_r + D_s$ and telescope the Q_r part: it contributes $Q_r(l) - Q_r(k_2 - 1)$, plus at most $i - r - 1$ breakpoint terms $p_r \le b$. The D_s increments are each at most $E_\downarrow$: (D2) is equivalent to $D_s(v) - D_s(u) \le E_\downarrow$ for $u < v \le n_f - \beta_\downarrow$, and inside Dn_f the increments $a - b$ are ≤ 0. Adding the breakpoint term $p_r(k_2) \le b$, where the path leaves row r, gives the bound. For $r = m$ the bound reads $C_m(l) \le (m-1)a + Q_m(l)$, which is Step 1.

Step 3 (the end block becomes exact). Put $\lambda_i(l) = C_i(l) - C_r(l) - (i-r)a$ for $i \ge r$. For a down-ramp at position l, with machine r busy (Step 1), so that $C_r(l) = C_r(l-1) + b$, the recursion $C_i(l) = \max(C_{i-1}(l), C_i(l-1)) + a$ for $i > r$ gives

$$\lambda_i(l) = \max\big(\lambda_{i-1}(l),\ \lambda_i(l-1) - (b-a)\big), \qquad \lambda_r \equiv 0.$$

By Step 2, $\max_i \lambda_i \le \Lambda_0$ at the start of Dn_f, and $\lambda_i \ge 0$. Each down-ramp lowers $\max_i \lambda_i$ by $b - a$ until it reaches 0, so after k_0 down-ramps $\lambda \equiv 0$. Hence the last $\beta_\downarrow - k_0$ jobs of factory f complete exactly at $C_f - kb$, for $k = 0, \dots, \beta_\downarrow - k_0 - 1$.

Step 4 (release envelopes). Let product g of factory f have its last job at position l_g, and let $\sigma_g = L_r(f) - Q_r(l_g)$ be the machine-r load after it. Define

$$\hat{c}_g = \begin{cases} M - S_{\rho(f)} - \sigma_g & \text{for the last } \beta_\downarrow - k_0 \text{ ramps of } f \text{ ("exact")}, \\ M - S_{\rho(f)} - \sigma_g + \Lambda_0 & \text{for every other product.} \end{cases}$$

Steps 1–3 and (T2) give $c_g \le \hat{c}_g$: by Step 2, $c_g \le (m-1)a + Q_r(l_g) + \Lambda_0 = (m-1)a + L_r(f) - \sigma_g + \Lambda_0$, which is at most $M - S_{\rho(f)} - \sigma_g + \Lambda_0$ by (T2), and Step 3 removes the Λ_0 for the exact ramps. By Lemma 1, makespan $= \mathrm{ERD}(c) \le \mathrm{ERD}(\hat{c})$. It remains to show that, for every h, with $\tau = M - \hat{c}_h$,

$$\sum_{g:\, \hat{c}_g \ge M - \tau} A_g \le \tau.$$

$\tau \ge 0$ for every h. For exact h, $\tau = S_\rho + kb \ge 0$. For non-exact h, $\sigma_h \ge (\beta_\downarrow - k_0)b \ge (\beta_\downarrow - k_0)\gamma \ge \Lambda_0$ by (T5b), so $\tau \ge S_\rho + \sigma_h - \Lambda_0 \ge 0$.

Case A: $\tau < (\beta_\downarrow - k_0)b - \Lambda_0$. Every non-exact g has $\sigma_g \geq (\beta_\downarrow - k_0)b$, so $\hat{c}_g < M - \tau$. Only exact ramps count, and h is one of them: $h = (\rho, k)$, the ramp with k exact ramps after it in the rank-ρ factory, with $\hat{c}_h = M - S_\rho - kb$.

- $k = 0$ (h is an end product): the ramps counted are the end products $e(\rho')$ with $S_{\rho'} \leq S_\rho$. Their sum is at most $S_\rho = \tau$.
- $k \geq 1$: every end product counts, because $S_F \leq F\alpha^* < b$. A non-end ramp (ρ', k') counts only if $k' \leq k$. For $\rho' > \rho$ with $S_{\rho'} > S_\rho$, only $k' \leq k - 1$ counts, because $0 < S_{\rho'} - S_\rho < b$. Let $N_> = \#\{\rho' > \rho : \ S_{\rho'} > S_\rho\}$. Then $S_F - S_\rho \leq \alpha^* N_>$, and the total is at most $S_F + \alpha^*(Fk - N_>) \leq S_\rho + \alpha^* Fk \leq S_\rho + kb = \tau$.

Case B: $\tau \geq (\beta_\downarrow - k_0)b - \Lambda_0$. Put $y = \tau + \Lambda_0$. Every counted g has $\sigma_g \leq y$.

- B1, $y < \beta_\downarrow b$: only ramps count, because Mid and U products have $\sigma \geq \beta_\downarrow b$. Each factory contributes at most $\lfloor y/b \rfloor + 1$ ramps, each with $A \leq \alpha^*$, so the total is at most $F\alpha^*(y/b + 1)$. This is at most $y - \Lambda_0$ if and only if $y\gamma/b \geq \Lambda_0 + F\alpha^*$, which holds because $y \geq (\beta_\downarrow - k_0)b$ and (T5b).
- B2, $y \geq \beta_\downarrow b$: in each factory all $\beta_\downarrow$ ramps count, plus a product-aligned suffix Z of $U_f \cdot \mathrm{Mid}_f$ with $x_r(Z) \leq y - \beta_\downarrow b + \bar{x}_r$. By (T3) the total is at most $F\beta_\downarrow \alpha^* + (y - \beta_\downarrow b + \bar{x}_r) + F \cdot E_A$, which is at most $y - \Lambda_0$ by (T5a).

Conclusion. In all cases the counted work is at most τ, so $\mathrm{ERD}(\hat{c}) \leq M$. If $M \leq X \leq \mathrm{OPT}$, then $\mathrm{OPT} = M$. □

Corollary 43 (When the certificate meets $\mathrm{LB}_{\mathrm{end}}$) *Suppose that an up-ramp exists or $r = 1$, so that $\mathrm{hd}_r = (r-1)a$; that a down-ramp exists, so that $\mathrm{tl}_r = (m-r)a$; and that the ends realise the cover, $W(\rho) = S_\rho$ for $\rho \leq F$. Then summing (T2) over f gives $M \geq (m-1)a + (L_r + \sum_\rho S_\rho)/F$, which is the $F' = F$ term of $\mathrm{LB}_{\mathrm{end}}$ at machine r. (This step needs only $W(\rho) = S_\rho$; note that $W(\rho) \leq S_\rho$ always holds, because the first ρ end products form a cover.) If moreover the minimum over F' in $\mathrm{LB}_{\mathrm{end}}$ is attained at $F' = F$ and the maximum over k is attained at $k = r$, that term equals $\mathrm{LB}_{\mathrm{end}}$, and (T2) with $M = \lceil \mathrm{LB}_{\mathrm{end}} \rceil$ is an exact staggered partition of machine r's load, with total slack less than F.*

Remark 44 (Mechanism and scope) In the factory-dominant regime the end of the schedule is the binding side. Three things make the bound exact: every factory's bottleneck runs without idle time; the F factory end times are staggered by the assembly times of the cheapest covering end products (Proposition 3); and a block of fast down-ramps, each with $A \leq \alpha^* < b/F$, gives the assembly machine credit at rate γ per level, which absorbs the backlog. The analogy with Theorem 6 is one of mechanism, not a reduction: reversing the whole problem turns the assembly machine into a first-stage disassembly with fan-out precedence, which is not a distributed assembly flowshop. The proof uses only the exact within-factory reversal in the end block lemma of Coutinho (manuscript under review), and handles the assembly side directly by the ERD envelope. The term $W(r)$ mirrors the factory-count condition $|N_{h_0}| \leq F$ on the bound side only: in the bound, one product of s jobs may end s factories at once, while the certificate uses single-job end products, one per factory. When the cheapest cover uses a multi-job product, $W(\rho) < S_\rho$, Corollary 43 does not apply, and the RSE schedule cannot reach $\lceil \mathrm{LB}_{\mathrm{end}} \rceil$; a certificate with split end products is open.

A.9 The continuous-law rate for LB_0

Let $R_{(2)}$ denote the second smallest release R_h. Proposition 13 assumes, in its probabilistic part, model M (Assumption 32) with $\nu(1) > 0$ and

$(\mathrm{H}\mu)$ $c_H x \le \mu([a, a+x]) \le C_H x$ for $0 \le x \le x_H$;
$(\mathrm{H}\kappa^-)$ $\kappa_1([0,u)) \ge \kappa_0 u$ for $u \le u_1$.

Theorem 14 assumes (M1), (M2) and (M4), and

$(\mathrm{H}\mu^+)$ $\mu([a, a+x]) \ge c_H x$ for $x \le x_H$;
$(\mathrm{H}\kappa^+)$ $\kappa_1([0,u)) \le C_\kappa u$ for all u.

Proof of Proposition 13 *Deterministic part.* Let $t \ge 2$ and let h_1 be the unique minimiser of R. Apply Lemma 1 at $h = \arg\min_{g \ne h_1} R_g$: every product except h_1 has $R \ge R_{(2)}$, so

$$\mathrm{ERD}(R) \ge R_{(2)} + A_{\mathrm{tot}} - A_{h_1} = \mathrm{LB}_0 + \big(R_{(2)} - T_0 - A_{h_1}\big).$$

Corollary 23 (with $\underline{r} = R$) gives $\mathrm{ERD}(R) \le \mathrm{OPT}$.

Probabilistic part. Let $x_t = (2t)^{-1/m}/(2C_H)$, and let t satisfy $2x_t \le \min(x_H, u_1)$. We show $\Pr[\mathrm{OPT} > \mathrm{LB}_0] \ge c_1\kappa_0 x_t = \Omega(t^{-1/m})$ with $c_1 = \nu(1)\big(c_H/(2mC_H)\big)^m/4$.

- Let $E_h = \{s_h = 1,\ P_h \le ma + x_t\} \cap \{R_g > ma + 2x_t \text{ for all } g \ne h\}$. These events are disjoint.
- $\Pr(s_h = 1,\ P_h \le ma + x) \ge \nu(1)\,\mu([a, a + x/m])^m \ge \nu(1)(c_H x/m)^m$ for $x/m \le x_H$.
- $\Pr(R_g \le ma + 2x) \le \mu([a, a+2x])^m \le (2C_H x)^m$ for $2x \le x_H$, because $R_g \ge P_j$ for each job j of g, and $P_j \le ma + 2x$ forces every entry of j into $[a, a+2x]$.
- With $t(2C_H x_t)^m = 1/2$, independence across products gives $\Pr(E_h) \ge \nu(1)(c_H x_t/m)^m(1 - \frac{1}{2t})^{t-1} \ge \nu(1)(c_H x_t/m)^m/2$. Summing over the t disjoint events gives $\Pr(\bigcup_h E_h) \ge c_1$.
- On E_h: $h_1 = h$, it is a single job, and $R_{(2)} - T_0 \ge x_t$. A_h is independent of the entries and sizes given $s_h = 1$, so $\Pr(A_h < x_t \mid E_h) \ge \kappa_0 x_t$ (as $x_t \le u_1$).
- On $E_h \cap \{A_h < x_t\}$ the deterministic part gives $\mathrm{OPT} - \mathrm{LB}_0 \ge R_{(2)} - T_0 - A_{h_1} > 0$.

The proof uses only $\nu(1) > 0$, $(\mathrm{H}\mu)$ and $(\mathrm{H}\kappa^-)$, and holds for every β. □

Remark 45 On this event every ERD-type bound ($\mathrm{ERD}(R)$, LB_A, LB^*) already exceeds LB_0. The obstruction defeats LB_0 only, and is invisible to the ERD-type bounds.

Proof of Theorem 14 We prove

$$\Pr[\mathrm{OPT} > \mathrm{LB}_0] \le C_\kappa\big[(2F)^{1/m} + 8^{1/m}\Gamma(1 + 1/m)\big](tc)^{-1/m} + O(1/t) + e^{-\Omega(t)},$$

where $c = \nu(1)\,\kappa_1([\alpha, \bar{A}])\,(c_H/m)^m$. Together with Proposition 13, this gives $\Pr[\mathrm{OPT} > \mathrm{LB}_0] = \Theta(t^{-1/m})$ under two-sided $(\mathrm{H}\mu)$ and $(\mathrm{H}\kappa^-)$, $(\mathrm{H}\kappa^+)$. We rerun the general case of the proof of Theorem 9, with a random start window in place of ε_t; $(\mathrm{H}\kappa^+)$ implies $\kappa_1(\{0\}) = 0$.

Start set. Let h_1 be the smallest-index single-job product attaining T_0. The event that none exists is event (i) of that proof, with failure probability $O(1/t)$. The extras are $F - 1$

single-job products $g \neq h_1$ with $P_g \leq ma + A_{h_1}$ and $A_g \geq \alpha$. The earlier cap $\min(A_{h_1}, \delta)$ is not needed: start jobs need not lie in the ramp class, because T_1 enters Theorem 6 only through $C_0 - (T_1 - T_0)$ and through the condition $(m-1)a \leq T_1$ of Lemma 26; and extras outside the ramp class do not reduce the ramp count.

Hypotheses of Theorem 6. Ramps are taken from the E_{ramp} class minus H_0, so at least FR_ℓ remain, using the count $F(R_\ell + 1)$ of Theorem 9. (H2) holds, since $R_g - T_0 \leq A_{h_1}$ for every extra g (as $T_0 \geq ma$). (H3) and (H5) hold as in Theorem 9, since $C_0 - (T_1 - T_0) \geq A_{h_1} + (F-1)\alpha - A_{h_1} = (F-1)\alpha > a + m\delta$. (H4) is as before.

Counting the extras (coupling). Let A' equal A, except that A'_{h_1} is replaced by an independent copy. Since h_1 is a function of sizes and entries only, the triples $(s_g, \text{entries}_g, A'_g)$ are i.i.d. across products, and A_{h_1} is independent of (sizes, entries, A'), with law κ_1 on $\{s_{h_1} = 1\}$. Let

$$N'(u) = \#\{g : \ s_g = 1, \ P_g \leq ma + \min(u, mx_H), \ A'_g \geq \alpha\}.$$

Then $N'(u) \sim \text{Bin}\big(t, \nu(1)\,\Phi(ma + \min(u, mx_H))\,\kappa_1([\alpha, \bar{A}])\big)$, with $\Phi(x) = \Pr(P_j \leq x)$, and the number of extras is at least $N'(A_{h_1}) - 1$. By $(\text{H}\mu^+)$, $\Phi(ma + v) \geq \mu([a, a + v/m])^m \geq (c_H v/m)^m$ for $v \leq mx_H$; the cap at mx_H is needed for $(\text{H}\mu^+)$ (if one keeps a δ-cap, $\delta \leq mx_H$ is required). Hence

$$\begin{aligned}\Pr(\text{fewer than } F-1 \text{ extras} \mid A_{h_1} = u) &\leq \Pr(N'(u) < F) \\ &\leq f(u) := \Pr\big(\text{Bin}(t, c\min(u, mx_H)^m) < F\big),\end{aligned}$$

which is non-increasing in u. By the layer-cake formula and $(\text{H}\kappa^+)$, since $\{f > s\} \cap [0, mx_H)$ is an initial interval for every s,

$$\int f\, d\kappa_1 \leq C_\kappa \int_0^{mx_H} f(u)\, du + f(mx_H).$$

By the Chernoff bound $\Pr(\text{Bin}(t,p) < tp/2) \leq e^{-tp/8}$, $f(u) \leq \exp(-tcu^m/8)$ once $tcu^m \geq 2F$. Splitting the integral at $u_0 = (2F/(tc))^{1/m}$ gives

$$\int_0^{mx_H} f(u)\, du \leq u_0 + \int_0^\infty e^{-tcu^m/8}\, du = \big[(2F)^{1/m} + 8^{1/m}\Gamma(1 + 1/m)\big](tc)^{-1/m},$$

which is the first term; $f(mx_H)$ is exponentially small.

Other events. E_{margin} and E_{ramp} fail with probability $e^{-\Omega(t)}$, as in Theorem 9. □

Remark 46 (Regimes) With an atom of μ at a the rate is exponential (Theorem 9). With no atom at a and κ_1 with positive density at 0 it is $\Theta(t^{-1/m})$. When κ_1 vanishes near 0 (for example $A \geq s$, as in Hatami et al. 2013), the failure probability is $O(1/t)$, dominated by the event that a product of size ≥ 2 attains T_0; a matching lower bound in that case is open.

Appendix B Proofs for the lattice results of Section 6

B.1 Lattice partition with unit transfers

Throughout this part of the appendices, a and b denote the global minimum and maximum entries, $\bar{s} = \max_h s_h$, $\bar{A}$ is an upper bound on every assembly time, $A_{\text{tot}} = \sum_h A_h$, and e_k is the k-th unit vector. Items i carry vectors $v_i \in \mathbb{Z}^m_{\geq 0}$ and possibly other attributes (assembly times). There are groups $f = 1, \ldots, F$, and $K \subseteq \{1, \ldots, m\}$ is the set of coordinates to be hit exactly.

Lemma 47 (unit transfers) *Each group G_f consists of a fixed part Φ_f, a core block B_f and a reservoir R_f. For each $k \in K$ and each f, R_f contains E designated items of type τ_k^0 and E designated items of type τ_k^1, with $v(\tau_k^1) = v(\tau_k^0) + e_k$; copies designated for different k are distinct items, even when their vectors coincide. Let $t_f \in \mathbb{Z}^m$ be targets with $\sum_f t_{f,k} = \sum_f v_k(G_f)$ for $k \in K$, and suppose $|v_k(G_f) - t_{f,k}| \le E$ for all f and all $k \in K$. Then exchanging reservoir items between groups gives groups G'_f with*

- *$v_k(G'_f) = t_{f,k}$ for every f and every $k \in K$;*
- *$v_{k'}(G'_f) = v_{k'}(G_f)$ for every $k' \notin K$;*
- *Φ_f and B_f unchanged.*

The number of exchanges is $\sum_{k\in K} \sum_f (v_k(G_f) - t_{f,k})^+ \le |K|FE$.

Proof Treat the coordinates of K one at a time. Fix k and put $x_f = v_k(G_f) - t_{f,k}$, so that $\sum_f x_f = 0$ and $|x_f| \le E$.

While some $x_f > 0$, pick f with $x_f > 0$ and h with $x_h < 0$. Move one designated τ_k^1 from f to h and one designated τ_k^0 from h to f. Then $v(G_f)$ changes by $-e_k$ and $v(G_h)$ by $+e_k$, so x_f falls by 1 and x_h rises by 1. No other coordinate changes, so the coordinates fixed earlier stay fixed.

A group with $x_f > 0$ only ever sends; its excess stays ≥ 0, and it sends exactly $x_f \le E$ times, so its E designated copies of τ_k^1 suffice. Symmetrically, a receiver uses at most $|x_h| \le E$ designated copies of τ_k^0. Items designated for coordinate k are never used for another coordinate. The process stops when every $x_f = 0$. □

Lemma 48 (Steinitz blocks) *Let C be a set of N items with vectors $z_i \in \prod_{c=1}^d [0, \bar{Z}_c] \subset \mathbb{R}^d$. Some order of C, cut into F consecutive blocks $B_1, \dots, B_F$ of sizes $\lfloor N/F \rfloor$ or $\lceil N/F \rceil$, satisfies*

$$|z_c(B_f) - z_c(C)/F| \le (2d+1)\bar{Z}_c \qquad \text{for every } f \text{ and } c.$$

The order is computable in polynomial time.

Proof Put $\bar{z} = z(C)/N$ and $u_i = \big((z_{i,c} - \bar{z}_c)/\bar{Z}_c\big)_c$ (a coordinate with $\bar{Z}_c = 0$ is dropped). Since $z_{i,c}$ and $\bar{z}_c$ both lie in $[0, \bar{Z}_c]$, $\|u_i\|_\infty \le 1$, and $\sum_i u_i = 0$. Lemma 5 gives, in polynomial time, an order whose prefix sums are at most d in sup-norm. A block sum is a difference of two prefix sums, so $|z_c(B_f) - N_f \bar{z}_c| \le 2d\bar{Z}_c$, where $N_f = |B_f|$. Also $|N_f \bar{z}_c - z_c(C)/F| = |N_f - N/F|\, \bar{z}_c \le \bar{Z}_c$. □

Proof of Lemma 15 The statement of Lemma 15 is the following special case of Lemma 47: the fixed parts Φ_f and the reservoirs R_f, as multisets of vectors, are the same for all f, the blocks B_f are the Steinitz blocks of Lemma 48 of a core set C, and the targets satisfy $\sum_f t_{f,k} = L_k$ and $|t_{f,k} - L_k/F| \le \theta$ for $k \in K$, where L_k is the total load.

The common part contributes $v_k(\Phi \cup R)$ to every group and $F\, v_k(\Phi \cup R)$ to L_k. Hence

$$v_k(G_f) - L_k/F = z_k(B_f) - z_k(C)/F,$$

which is at most $(2d+1)\bar{Z}_k$ in absolute value by Lemma 48. So $|v_k(G_f) - t_{f,k}| \le (2d+1)\bar{Z}_k + \theta$, and Lemma 47 applies with any integer $E \ge (2d+1)\bar{Z}_k + \theta$, in particular $E = \lceil (2d+1)\bar{Z}_k + \theta \rceil$. Also $\sum_f v_k(G_f) = L_k = \sum_f t_{f,k}$, as Lemma 47 requires. □

Remark 49 Lemma 48 is the Steinitz-block device already used for the deterministic exactness results, based on the lemma of Sevast'janov (1994a) (Lemma 5). The reservoir-exchange step of Lemma 47 is an elementary exchange argument; we found no source for it and claim nothing beyond its use here. Exchanges change the non-vector attributes of the groups (assembly times); Lemma 47 claims nothing about them, and every later use bounds them separately.

B.2 The distributed permutation flowshop

In this section all $A_h = 0$ ($\beta = 0$), so the model is the DPFSP of Naderi and Ruiz (2010). Recall $\mathrm{hd}_k = \min_j \sum_{i<k} p_{ij}$, $\mathrm{tl}_k = \min_j \sum_{i>k} p_{ij}$ and

$$\mathrm{LB}_D = \max_k \big[\mathrm{hd}_k + \lceil L_k/F \rceil + \mathrm{tl}_k\big].$$

For integer data this is the bound of Gogos (2023) (Alg. 1, $\max_k[\min \text{head} + L_k/F + \min \text{tail}]$ rounded up; Gogos (2023) also takes the maximum with $\max_j P_j$.

Validity for every schedule, non-permutation included. In each factory, machine k starts its first operation no earlier than hd_k and then works for its load; its last job still needs at least tl_k afterwards. Some factory has machine-k load at least $\lceil L_k/F \rceil$ (integer loads, pigeonhole). So $\mathrm{OPT} \geq \mathrm{LB}_D$, and also $\mathrm{OPT} \geq \lceil \mathrm{Gogos} \rceil$, the rounded-up bound of Gogos (2023).

Vector types. The $2(m-1)$ exact staircase types of Coutinho (manuscript under review) are the front types $r = 2, \ldots, m$ (a on machines $i < r$, b on machines $i \geq r$) and the end types $r = 1, \ldots, m-1$ (b on machines $i \leq r$, a on machines $i > r$). For each k, $\tau_k^0 = a\mathbf{1} + (s-a)e_k$ and $\tau_k^1 = \tau_k^0 + e_k$, where $s, s+1 \in [a, b]$. The demand of a vector type is Fm per staircase type plus FE per τ role; when vectors coincide, their demands add up.

Proof of Theorem 16(a) Let $E = 2m(b-a) + b + 1$ and suppose that every distinct vector occurs in the instance at least as often as the summed demand of the types equal to it.

Value of LB_D. A front job of type m gives $\mathrm{hd}_k = (k-1)a$, and an end job of type 1 gives $\mathrm{tl}_k = (m-k)a$. So $\mathrm{LB}_D = (m-1)a + T$ with $T := \max_k \lceil L_k/F \rceil$.

Targets. For each k put $e_k = FT - L_k \geq 0$ and $t_{f,k} = T - \lfloor e_k/F \rfloor - [f \leq e_k \bmod F]$. Then $\sum_f t_{f,k} = L_k$, $t_{f,k} \leq T$, and $|t_{f,k} - L_k/F| < 1$.

Groups. The fixed part Φ_f consists of m jobs of each staircase type. The reservoir R_f consists of E designated τ_k^0 and E designated τ_k^1 for every k. The core C is all other jobs. Cut C into Steinitz blocks by Lemma 48 with $d = m$, using the centred vectors $z - a\mathbf{1} \in [0, b-a]^m$, which gives $\bar{Z} = b - a$. In the original coordinates the deviation is then at most $2m(b-a) + |N_f - N/F| \max_c \bar{z}_c \leq 2m(b-a) + b$. With $\theta < 1$, Lemma 15 applies with $E = 2m(b-a) + b + 1$ and $K = \{1, \ldots, m\}$. The result is groups G'_f with $v(G'_f) = t_f \leq T\mathbf{1}$ that still contain m jobs of every staircase type.

Per factory. Factory f receives G'_f. Its entries lie in $[a, b]$, and it holds m exact staircase jobs of each type, so the staircase theorem of Coutinho (manuscript under review) (exact atoms, at least m staircase jobs per type, every load vector) gives

$$C_f = \max_k L_k(G'_f) + (m-1)a \leq T + (m-1)a = \mathrm{LB}_D.$$

Conclusion. Hence $\mathrm{LB}_D \leq X \leq \mathrm{OPT} \leq \max_f C_f \leq \mathrm{LB}_D$ for every $X \in [\mathrm{LB}_D, \mathrm{OPT}]$, in particular for $X = \lceil \mathrm{Gogos} \rceil$. Since LB_D is valid for every schedule, the non-permutation optimum is the same.

Running time. Lemma 5 runs in polynomial time, the exchanges take $O(mFE)$ steps, and the staircase schedule of Coutinho (manuscript under review) is computed in polynomial time. □

Proof of Theorem 16(b) Put

$$q_L = \min\Big\{ \min_{1\le s'\le m-1} \mu(a)^{s'}\mu(b)^{m-s'},\ \mu(a)^{m-1}\min(\mu(s),\mu(s+1)) \Big\} > 0$$

and $D = 2mF(m-1+E)$. A front type r has $r-1$ entries a and an end type r has $m-r$ entries a, so s' ranges over $1,\dots,m-1$, and each of the at most $4m-2$ relevant vectors occurs in a given job with probability at least q_L. The total demand is $2(m-1)\cdot Fm + 2m\cdot FE = D$, so each vector's demand is at most D. Since $\Pr(\mathrm{Bin}(n,p) < D)$ is non-increasing in p, part (a), a union bound over the vectors and the Chernoff bound $\Pr(\mathrm{Bin}(n,q) \le nq/2) \le e^{-nq/8}$ give, for every n and every $X \in [\mathrm{LB}_D, \mathrm{OPT}]$,

$$\Pr[\mathrm{OPT} \ne X] \le (4m-2)\Pr(\mathrm{Bin}(n,q_L) < D) \le (4m-2)e^{-nq_L/8},$$

the second form for $n \ge 2D/q_L$. □

Remark 50 (Taillard's lattice) For $U\{1..B\}$ with $B \ge 2$: $q_L = B^{-m}$, $E = 2m(B-1)+B+1$ and $D = 2mF(m-1+E)$. So $\Pr[\mathrm{OPT} \ne X] \le (4m-2)\exp(-n/(8B^m))$ for $n \ge 2D\,B^m$. For $B = 99$ this is vacuous at benchmark sizes ($n \le 500$), as in the permutation-flowshop case of Coutinho (manuscript under review).

Proof of Theorem 16(c) Assume (L), $m \ge 2$ and $F \ge 1$; put $\rho_0 = \min(\mu(a),\mu(b),\mu(s),\mu(s+1))$ and let $n \ge \max(F+2,\, Fmb+F+2) = F(mb+1)+2$. Here OPT is the permutation optimum, which is the DPFSP definition. We show $\Pr[\mathrm{OPT} > X] \ge \rho_0^{mn}$ for $X \in \{\mathrm{LB}_D, \lceil \mathrm{Gogos} \rceil\}$.

The event. Choose $y \in \{0,\dots,F-1\}$ such that F divides $b+a+(n-2)(s+1)-y$ (possible since $n-2 \ge F-1 \ge y$). The event is: job 1 is $b\mathbf{1}$, job 2 is $a\mathbf{1}$, jobs $3,\dots,y+2$ are $s\mathbf{1}$, and the remaining $n-2-y$ jobs are $(s+1)\mathbf{1}$. Each of its mn entries takes its prescribed value with probability at least ρ_0, so the event has probability at least ρ_0^{mn}.

Bounds on the event. Every machine has the same load $L = b+a+ys+(n-2-y)(s+1) = b+a+(n-2)(s+1)-y$, and $F \mid L$. Job 2 gives $\mathrm{hd}_k = (k-1)a$ and $\mathrm{tl}_k = (m-k)a$, so $\mathrm{LB}_D = (m-1)a + L/F$. Also $L \ge n-2-y \ge n-F-1 \ge Fmb$, so $\max_j P_j = mb \le (m-1)a + L/F$, and therefore $\lceil \mathrm{Gogos} \rceil = \mathrm{LB}_D$.

Lower bound on OPT. Every job is a constant vector $c_j\mathbf{1}$. In a factory, consider the path that runs along machine 1 to the column of a job j^*, down that column, and along machine m. Its value is $L(f)+(m-1)c_{j^*}$. So $C_f \ge L(f)+(m-1)a$ for every non-empty factory, and $C_f \ge L(f)+(m-1)b$ for the factory holding job 1. Averaging over the $F' \le F$ non-empty factories gives

$$\mathrm{OPT} \ge (m-1)a + \frac{L+(m-1)(b-a)}{F'} \ge (m-1)a + \frac{L+(m-1)(b-a)}{F}.$$

Since OPT is an integer, $F \mid L$ and $(m-1)(b-a) \ge 1$, $\mathrm{OPT} \ge (m-1)a + L/F + 1 = \mathrm{LB}_D + 1$.

Hence $-\frac{1}{n}\log\Pr[\mathrm{OPT} \ne \mathrm{LB}_D]$ stays between two positive constants for the permutation optimum: the upper constant from this event, the lower one from part (b). The lower bound uses the permutation path formula; for the non-permutation optimum it is not proved when $m \ge 4$. The upper bound of part (b) covers both. □

B.3 The end side: two machines tied in load

Throughout, the data are integers (entries and assembly times) and factory schedules are permutation schedules, which is the DAPFSP definition. W is the covering function of Proposition 3, extended to $r \le 2F$: $W(r) = \min\{\sum_{g\in G} A_g : \sum_{g\in G} s_g \ge r\}$ for $r \le \min(2F, n)$, and $W(r) = \infty$ for $r > n$.

Proof of Proposition 4 Fix a schedule. Let F' be the number of factories used and F_2 the number of those with at least two jobs. In factory f let z_f be the last job and, if $n_f \ge 2$, y_f the second-to-last.

(a) *Factories with $n_f \ge 2$.* Machine k starts its first operation at or after hd_k and then works $L_k(f)$, so $C_k(z_f) \ge \mathrm{hd}_k + L_k(f)$. Since $l > k$ and there are $m-k-1$ other machines after k,

$$C_{z_f} \ge \mathrm{hd}_k + L_k(f) + \sum_{i>k} p_{iz_f} \ge \mathrm{hd}_k + L_k(f) + p_{lz_f} + (m-k-1)a.$$

On machine l, the first $n_f - 1$ jobs carry load $L_l(f) - p_{lz_f}$, and in a permutation schedule y_f is the last of them, so $C_{y_f} \ge \mathrm{hd}_l + L_l(f) - p_{lz_f} + \mathrm{tl}_l$. Adding, $C_{z_f} + C_{y_f} \ge \mu_{kl} + L_k(f) + L_l(f)$.

(b) *Factories with $n_f = 1$.* $C_{z_f} = P_{z_f} \ge \mathrm{hd}_k + p_{kz_f} + \mathrm{tl}_k$, and likewise for l. So $C_{z_f} \ge (\nu_k + \nu_l + p_{kz_f} + p_{lz_f})/2$.

(c) *Sum of the events.* Put $F_1 = F' - F_2$. Using $\sum_f (L_k(f) + L_l(f)) = L_k + L_l$ and $p_{kz} + p_{lz} \le \pi_{kl}$ (in (b), $(\nu_k + \nu_l + p_k + p_l)/2 = p_k + p_l + (\nu_k + \nu_l - (p_k + p_l))/2 \ge p_k + p_l + (\nu_k + \nu_l - \pi_{kl})/2$), the $N = F' + F_2$ event times $\{C_{z_f}\} \cup \{C_{y_f}\}$ sum to at least

$$F_2\mu_{kl} + L_k + L_l + F_1(\nu_k + \nu_l - \pi_{kl})/2.$$

(d) *Covering argument, as in Proposition 3.* Sort the events in decreasing order, $E_{(1)} \ge \cdots \ge E_{(N)}$. For each r, the r latest events belong to r distinct jobs. Their products cover at least r jobs, are all assembled at or after $E_{(r)}$, and need total assembly time at least $W(r)$. So the makespan is at least $E_{(r)} + W(r)$. Averaging over $r = 1, \ldots, N$ gives makespan $\ge V_{kl}(F', F_2)$. Take the minimum over the unknown (F', F_2) (with $1 \le F' \le \min(F, n)$ and $0 \le F_2 \le \min(F', n - F')$), the maximum over (k, l), and round up (integer data). □

Remark 51 The bound charges the last two completions of every factory with at least two jobs: machine k forces the last one and machine l the one before it, and the covering argument then charges up to $2F$ assembly releases. *Intuition only (not a proved property of optimal schedules):* when both machines bind in every factory, the last two jobs of a factory finish close together, and the assembly machine cannot absorb $2F$ releases at cost α_0 each.

We prove Proposition 17(i) under the following conditions, for machines $k < l$ and an integer $\alpha_0 \ge 0$; put

$$x := L/F + \alpha_0(F+1)/2, \qquad \delta := (F\alpha_0 - a)/2.$$

(1) $\mathrm{hd}_i = (i-1)a$ and $\mathrm{tl}_i = (m-i)a$ for every i (for example, an all-a job exists);
(2) $L_k = L_l = L := \max_i L_i$;
(3) $A_h \ge s_h\alpha_0$ for every product, and at least $2F$ single-job products have $A = \alpha_0$, so that $W(r) = r\alpha_0$ for $r \le 2F$;
(4′) $F\alpha_0 > a$ and $\lceil x \rceil - x < \delta$;

(5) $L \geq L_* := F^2 b + F(2F-1)\big[(m-1)a + \alpha_0(2F+1)/2 + 1\big]$;
(6a) $FA_{\text{tot}} \leq L - F\bar{s}mb$;
(6b) for every machine i there are at least $F\bar{A}$ single-job products with $A = \alpha_0$ and $p_i \geq F\alpha_0 + 1$.

Condition (4′) holds whenever $x \in \mathbb{Z}$, that is, whenever F divides $L + \alpha_0 F(F+1)/2$ (note that $\alpha_0 F(F+1)/2 \in \mathbb{Z}$), because then $\lceil x \rceil - x = 0 < \delta$. It holds always when $\delta \geq 1$, that is, when $F\alpha_0 \geq a + 2$ (call this condition (4)), because $\lceil x \rceil - x < 1$. Condition (6b) needs $F\alpha_0 + 1 \leq b$.

Let LB_{cf} denote the closed-form bound of Coutinho (unpublished manuscript), the ceiling of the maximum of four terms: a *factory-load* term $\max\big(\max_i[\text{hd}_i + \lceil L_i/F \rceil + \text{tl}_i], \max_j P_j\big) + \min_h A_h$; two *assembly* terms (a trivial one and one from the earliest-release-date rule), each at most $\max_h r_h + A_{\text{tot}}$, where the release bound r_h of product h is at most the sum of its entries; and a *Johnson fluid* term $\max_i[\text{hd}_i + J_i + \text{tl}_i]$, where J_i is Johnson's two-machine makespan for the jobs $(x_{h,i}/F, A_h)$. Let $\text{LB}_{\text{comb}} = \max(\text{LB}_{\text{cf}}, \text{LB}_{\text{end}})$.

Proof of Proposition 17(i) Assume (1), (2), (3), (4′) and (5). Condition (3) provides $n \geq 2F$, so the pair $(F', F_2) = (F, F)$ is admissible in Proposition 4.

The value of LB_{end}. By (1), $\text{hd}_i + \text{tl}_i = (m-1)a$ for every i, so by (2) the maximum over machines in LB_{end} is attained at k. By (3), the F' term is $(m-1)a + L/F' + \alpha_0(F'+1)/2$. For $F' < F$ it exceeds the $F' = F$ term by

$$(F - F')\Big[\frac{L}{FF'} - \frac{\alpha_0}{2}\Big] \geq \frac{L}{F(F-1)} - \frac{\alpha_0 F}{2} \geq 0,$$

by (5). So $\text{LB}_{\text{end}} = (m-1)a + \lceil x \rceil$.

LB_{tie} *at* $(F', F_2) = (F, F)$. Here $N = 2F$ and $\mu_{kl} = (k-1)a + (m-k-1)a + (l-1)a + (m-l)a = (2m-3)a$, and $\sum_{r=1}^{2F} W(r) = \alpha_0 F(2F+1)$. So

$$V_{kl}(F, F) = (m-1)a - \frac{a}{2} + \frac{L}{F} + \frac{\alpha_0(2F+1)}{2} = (m-1)a + x + \delta.$$

Every other (F', F_2). Here $N \leq 2F - 1$, μ_{kl} and ν are ≥ 0, $\pi_{kl} \leq 2b$ and $\sum W \geq 0$. So the numerator of V_{kl} is at least $2L - F'b \geq 2L - Fb \geq 0$, and

$$V_{kl}(F', F_2) \geq \frac{2L - Fb}{2F - 1} = \frac{L}{F} + \frac{L - F^2 b}{F(2F-1)} \geq (m-1)a + x + 1,$$

by (5). (Condition (5) is sufficient but conservative: this step needs only $\alpha_0(F+1)/2$ where (5) has $\alpha_0(2F+1)/2$.)

Hence $\min_{F', F_2} V_{kl} \geq (m-1)a + x + \min(\delta, 1)$. By (4′), $\lceil x \rceil - x < \delta$, and always $\lceil x \rceil - x < 1$; so $x + \min(\delta, 1) > \lceil x \rceil$. Since $(m-1)a$ is an integer,

$$\text{OPT} \geq \text{LB}_{\text{tie}} \geq \big\lceil (m-1)a + x + \min(\delta, 1) \big\rceil \geq (m-1)a + \lceil x \rceil + 1 = \text{LB}_{\text{end}} + 1.$$

Under (4), $\delta \geq 1$ and this reads $\text{LB}_{\text{tie}} \geq \lceil (m-1)a + x + 1 \rceil = \text{LB}_{\text{end}} + 1$ directly.

The bound LB_{cf}. Assume (6a) and (6b) as well. We show $\text{LB}_{\text{cf}} \leq \text{LB}_{\text{end}}$, term by term.

- *Factory load.* It is at most $(m-1)a + \lceil L/F \rceil + \alpha_0 \leq (m-1)a + \lceil x \rceil$, because $\max_j P_j \leq \bar{s}mb \leq L/F$ by (6a) (and $FA_{\text{tot}} \geq 0$), $\min_h A_h = \alpha_0$ by (3), α_0 is an integer, and $(F+1)/2 \geq 1$. Condition (5) does *not* imply $\max_j P_j \leq (m-1)a + L/F$: for $m = 5$, $F = 2$, $a = 1$, $b = 99$, $\alpha_0 = 2$ one has $L_*/F + (m-1)a = 232 < mb = 495$.
- *Assembly terms.* Both are at most $\max_h r_h + A_{\text{tot}} \leq \bar{s}mb + A_{\text{tot}} \leq L/F \leq x$, by (6a).

- *Johnson fluid.* Write, in Johnson order, $J_i = \max_u \big[L_i/F + A_u + \sum_{v>u}(A_v - x_{v,i}/F)\big]$.
  - For u in the second Johnson group ($A_u < x_{u,i}/F$) with $A_u = \alpha_0$, the later products are all in the second group, so the later terms are ≤ 0 and the value is at most $L_i/F + \alpha_0$.
  - For u in the second group with $A_u > \alpha_0$: the second group is sorted by decreasing A, and the products of (6b) have $A = \alpha_0 < A_u$ and $x_i/F \ge \alpha_0 + 1/F > A$, so all of them come after u, each contributing at most $-1/F$. The value is at most $L_i/F + \bar{A} - F\bar{A}/F = L_i/F$.
  - For u in the first group ($x_{v,i}/F \le A_v$), the value equals $\sum_{v<u} x_{v,i}/F + x_{u,i}/F + \sum_{v\ge u} A_v \le A_{\mathrm{tot}} + x_{u,i}/F \le A_{\mathrm{tot}} + \bar{s}\bar{b} \le L/F$, by (6a).

  So the Johnson term is at most $(m-1)a + L/F + \alpha_0 \le (m-1)a + x$.

Hence $\mathrm{LB}_{\mathrm{cf}} \le \mathrm{LB}_{\mathrm{end}}$, so $\mathrm{LB}_{\mathrm{comb}} = \mathrm{LB}_{\mathrm{end}}$ and $\mathrm{OPT} \ge \mathrm{LB}_{\mathrm{comb}} + 1$. □

Remark 52 Part (i) has no premise $F \ge 2$. At $F = 1$ the $(1,1)$ term exceeds $\mathrm{LB}_{\mathrm{end}}$ by $(\alpha_0 - a)/2$, so the obstruction is driven by the assembly stage, not by the number of factories. With $A \equiv 0$ (the DPFSP) it disappears, consistent with Theorem 16. Part $(4')$ is a sufficient condition for this mechanism; whether other obstructions exist when $F\alpha_0 \le a$ is not addressed here.

Proof of Proposition 17(ii) for $\beta < 1$ Assume lattice model M (Section 6) with $m \ge 2$, $a < F\alpha_0 < b$ and $\beta < 1$. We first treat $F\alpha_0 \ge a + 2$, so that (4) and hence $(4')$ hold deterministically; the band $a < F\alpha_0 < a + 2$ is treated at the end. We show that there are $c > 0$ and t_0 with $\Pr[\mathrm{OPT} > X] \ge c\, t^{-1/2}$ for $X \in \{\mathrm{LB}_{\mathrm{end}}, \mathrm{LB}_{\mathrm{comb}}\}$ and $t \ge t_0$.

Conditioning. Condition on the sizes; then $n \in [t, \bar{s}t]$ jobs have i.i.d. entries with law μ, of variance $\sigma^2 > 0$.

Tie event. The vectors $\xi_j = (p_{1j} - p_{ij})_{i=2,\dots,m}$ are i.i.d. in $\mathbb{Z}^{m-1}$, with mean 0 and covariance $Q = \sigma^2(I + \mathbf{1}\mathbf{1}^\top)$, which is non-singular. Varying $p_{ij} \in \{s, s+1\}$ for a single i shows that $\operatorname{supp}\xi - \operatorname{supp}\xi$ contains every unit vector and hence generates $\mathbb{Z}^{m-1}$, so the walk $S_n = \sum_{j\le n} \xi_j$ is strongly aperiodic. The local limit theorem (Spitzer 1976, P7.9) gives

$$n^{(m-1)/2} \Pr(S_n = x) = (2\pi)^{-(m-1)/2} |Q|^{-1/2} e^{-x\cdot Q^{-1}x/(2n)} + o(1),$$

uniformly in x. The coordinates of S_n are $L_1 - L_i$, $i = 2, \dots, m$. The set

$$\Lambda_n = \{x : \ x_2 = 0,\ 0 \le x_i \le \sqrt{n} \text{ for } i \ge 3\}$$

(indexing by $i = 2, \dots, m$) lies inside $\{L_1 = L_2 = \max_i L_i\}$, has at least $n^{(m-2)/2}$ points, and on it $x \cdot Q^{-1}x \le \lambda_{\max}(Q^{-1})(m-2)n$, so the Gaussian factor is bounded below there. So $\Pr(\mathrm{Tie}) \ge c' n^{-1/2} \ge c'(\bar{s}t)^{-1/2}$ for $n \ge n_0$. For $m = 2$, $\Lambda_n = \{0\}$. On the tie event, (2) holds with $(k, l) = (1, 2)$.

Good event. It consists of the following, each of which fails with probability $e^{-\Omega(t)}$:

- an all-a job exists, giving (1);
- at least $2F$ single-job products with $A = \alpha_0$, giving (3) (the lower bound $A_h \ge s_h\alpha_0$ holds by the model), and, for each machine i, the $F\bar{A}$ products of (6b); each has positive probability per product, at least $\nu(1)\kappa_1(\{\alpha_0\})\mu(b)$ for (6b), since $F\alpha_0 + 1 \le b$ and b is an atom;
- (5): deterministic for $t \ge L_*/a$ when $a \ge 1$, since $L \ge na$; by Hoeffding's inequality (Hoeffding 1963), $L \ge n\mathbb{E}p/2$, when $a = 0$;

- (6a), by Hoeffding's inequality, since $L \geq L_1$ and $\mathbb{E}[FA_h - x_{h,1}] = -\mathbb{E}[S]\,\mathbb{E}p\,(1-\beta) < 0$.

Conclusion. $\Pr(\text{Tie} \cap \text{Good}) \geq \Pr(\text{Tie}) - \Pr(\text{Good}^c) \geq c'(\bar{s}t)^{-1/2} - e^{-\Omega(t)}$, and no independence is needed. On this event part (i) gives $\text{OPT} > \text{LB}_{\text{comb}} \geq \text{LB}_{\text{end}}$.

The band $a < F\alpha_0 < a+2$. Here $(4')$ is not automatic. The proof in Section B.5 applies verbatim with $\beta < 1$ in place of $\beta = 1$: the tie event is replaced by the residue-restricted local event Z_n^ρ, which makes x an integer and so gives $(4')$; the assembly event G_A has probability tending to 1 by the law of large numbers, since $\mathbb{E}\zeta_h < 0$; and every other step does not use β. That proof also covers $F\alpha_0 \geq a+2$. □

B.4 The factory side with a constant margin

We use the notation of Theorem 8 and its hypotheses (T1)–(T5) (42). Fix a factory and a machine r (machines are numbered from 1). Write $v_r = a\mathbf{1} + (b-a)e_r$, let $Q_i(t)$ be the prefix load of machine i over the first t jobs of the factory sequence, and let $D_i = Q_i - Q_r$. The conditions (U1), (U2) (for $i < r$, with parameters $(\beta_\uparrow, E_\uparrow, W_\uparrow)$) and (D1), (D2) (for $i > r$, with $(\beta_\downarrow, E_\downarrow, W_\downarrow)$) are those of the start-block and end-block lemmas of Coutinho (manuscript under review), as used in (T1).

Lemma 53 (block conditions with a product-level Steinitz middle) *Fix integers $\beta_\uparrow \geq [r \geq 2]$ and $\beta_\downarrow \geq 1$. The factory sequence is $U \cdot \text{Mid} \cdot \text{Dn}$ with n jobs, where*

- *U consists of $\beta_\uparrow$ single-job products of vector v_r;*
- *Dn consists of $\beta_\downarrow$ single-job products of vector v_r;*
- *Mid consists of N products of size at most $\bar{s}$, each contiguous, in the Steinitz order (Lemma 5, $d = m$) of the centred vectors $(w_g - \bar{w},\ y_g - \bar{y})$. Here $w_g = (x_{g,i} - x_{g,r})_{i \neq r}$, scaled by $2\bar{s}(b-a)$; $y_g = A_g - x_{g,r}/F$, scaled by $2Y$ with $Y = \bar{A} + \bar{s}b/F$; and $\bar{w} = w(\text{Mid})/N$, $\bar{y} = y(\text{Mid})/N$.*

Put $K = (2m+1)\bar{s}(b-a)$, $E = 2K$ and $W = (m-2)E + (m-1)(b-a)$. Assume

(a) $x_r(\text{Mid}) \geq x_i(\text{Mid})$ for every $i \neq r$;

(b) $(b-a)\beta_\uparrow \geq W + K$ when $r \geq 2$, and $(b-a)\beta_\downarrow \geq W + K$ when $r \leq m-1$.

Then (U1), (U2) hold for $i < r$ with $(\beta_\uparrow, E, W)$, and (D1), (D2) hold for $i > r$ with $(\beta_\downarrow, E, W)$. If moreover $y(\text{Mid}) \leq 0$, then every product-aligned suffix Z of $U \cdot \text{Mid}$ has

$$A(Z) - x_r(Z)/F \leq E_A := 2mY + \beta_\uparrow \bar{A},$$

which is (T3).

Proof We adapt Steps 2–4 of the proof of Theorem A of Coutinho (manuscript under review) from jobs to products. Since $W \geq (r-2)^+E + (r-1)(b-a)$ and $W \geq (m-r-1)^+E + (m-r)(b-a)$, and (U1), (D1) only get stronger with W, it suffices to prove them with W.

Line control. Let $q(t)$ be the number of complete Mid products by position t, and $\ell_i(t) = q(t)\bar{w}_i$. The Steinitz order bounds the prefix deviation at product boundaries by $m \cdot 2\bar{s}(b-a)$; inside a product, D_i moves by at most $\bar{s}(b-a)$. So $|D_i(t) - D_i(\beta_\uparrow) - \ell_i(t)| \leq K$ on Mid. By (a), $\bar{w}_i = (x_i - x_r)(\text{Mid})/N \leq 0$, so ℓ_i is non-increasing.

(U1). For $i < r$ and t in Mid: $D_i(t) \leq D_i(\beta_\uparrow) + K = -(b-a)\beta_\uparrow + K \leq -W$, by (b). In Dn the increments of D_i are $a - b \leq 0$.

(U2). For $s < t$ in Mid: $D_i(t) - D_i(s) \le \ell_i(t) - \ell_i(s) + 2K \le E$. The case $s = \beta_\uparrow$ is included ($q = 0$, deviation at most K). If t lies in Dn, its increments are ≤ 0, so $D_i(t) - D_i(s) \le D_i(n - \beta_\downarrow) - D_i(s) \le E$, or ≤ 0 if s is in Dn too.

(D1), (D2) for $i > r$. By Step 4 of Coutinho (manuscript under review), these are equivalent to $D_i(u) \ge -\Delta_i + W$ for $u < n - \beta_\downarrow$, and $D_i(v) - D_i(u) \le E$ for $u < v \le n - \beta_\downarrow$, where $\Delta_i = L_r(f) - L_i(f)$ is the load difference of the factory.

- The ramps have $p_r - p_i = b - a$, so $\Delta_i = (b-a)(\beta_\uparrow + \beta_\downarrow) + (x_r - x_i)(\text{Mid}) \ge (b-a)(\beta_\uparrow + \beta_\downarrow)$.
- $D_i(n - \beta_\downarrow) = D_i(\beta_\uparrow) + T_i$ exactly, where $T_i = N\bar{w}_i = \ell_i(\text{end})$. So, for u in Mid, $D_i(u) \ge D_i(n - \beta_\downarrow) - K = -\Delta_i + (b-a)\beta_\downarrow - K \ge -\Delta_i + W$.
- For u in U (including $u = 0$), $D_i(u) = -u(b-a) \ge -(b-a)\beta_\uparrow \ge -\Delta_i + (b-a)\beta_\downarrow \ge -\Delta_i + W$.
- Increments: ≤ 0 inside U; $\le 2K = E$ inside Mid; and from U into Mid, $D_i(v) - D_i(u) \le (\ell_i(v) + K) - (b-a)(\beta_\uparrow - u) \le K$.

(T3). For $Z \subseteq$ Mid starting after q products, the Steinitz bound on the y coordinate gives

$$y(Z) = y(\text{Mid}) - y(\text{prefix}_q) \le y(\text{Mid}) - (q\bar{y} - 2mY) = (N - q)\bar{y} + 2mY \le 2mY,$$

since $\bar{y} = y(\text{Mid})/N \le 0$. Each up-ramp in Z adds $y \le \bar{A}$, and there are at most $\beta_\uparrow$ of them. □

Constants of Theorem 18.

Let the data be integers in $[a, b]$ with $a < b$ the minimum and maximum entries, fix a machine r, an integer $\alpha_0 \ge 0$ with $\gamma := b - F\alpha_0 > 0$, and s with $s, s+1 \in [a, b]$. Here $\bar{A}$ is any upper bound on every A_h, including the up-ramps and the reservoir items, which may carry any assembly time; in the probabilistic use, $\bar{A} = \sup \operatorname{supp} \kappa$. With K, E, W, Y as in Lemma 53:

$$\begin{aligned}
\beta_\uparrow &:= \lceil (W+K)/(b-a) \rceil \cdot [r \ge 2], \qquad E_A := 2mY + \beta_\uparrow \bar{A}, \\
\Lambda_0 &:= (m-r)(b+E-a), \qquad k_0 := \lceil \Lambda_0/(b-a) \rceil, \\
\beta_\downarrow &:= \max\left\{ \left\lceil \frac{W+K}{b-a} \right\rceil,\ k_0 + \left\lceil \frac{\Lambda_0 + F\alpha_0}{\gamma} \right\rceil,\ \left\lceil \frac{\Lambda_0 + \bar{s}b + FE_A}{\gamma} \right\rceil \right\}, \\
\mathrm{dev}_x &:= (2m+3)\bar{s}b, \qquad \mathrm{dev}_A := (2m+3)\bar{A}, \\
E_T &:= \lceil \mathrm{dev}_x + F\alpha_0 + 1 \rceil, \\
\Delta^* &:= F(\beta_\uparrow + \beta_\downarrow)(b-a) + FE_T(2s+1-2a) + F(2\,\mathrm{dev}_x + E_T), \\
Y_1 &:= \left[(\beta_\uparrow + \beta_\downarrow)b + 2E_T b + \mathrm{dev}_x \right]/F + \mathrm{dev}_A + 2E_T\bar{A}, \\
c_1 &:= F^2 Y_1.
\end{aligned}$$

With $\tau_r^0 = a\mathbf{1} + (s-a)e_r$ and $\tau_r^1 = \tau_r^0 + e_r$, the hypotheses of Theorem 18 read:

(F1) $A_h \ge s_h \alpha_0$ for every product;
(F2) the instance contains, as distinct single-job products, $F\beta_\uparrow$ of vector v_r (any A), $F\beta_\downarrow$ of vector v_r with $A = \alpha_0$, and FE_T each of τ_r^0 and τ_r^1 (any A); demands add up when vectors coincide;
(F3) $L_r - L_i \ge \Delta^*$ for every $i \ne r$;

(F4) $FA_{\text{tot}} \le L_r - c_1$;
(F5) $L_r \ge \alpha_0 F(F-1)/2$.

Proof of Theorem 18 Let $M := (m-1)a + \lceil (L_r + \alpha_0 F(F+1)/2)/F \rceil$.

$M \le \text{LB}_{\text{end}}$. The v_r ramps give $\text{hd}_r + \text{tl}_r = (m-1)a$. By (F1) and the at least F products with $A = \alpha_0$, $W(\rho) = \rho\alpha_0$ for $\rho \le F$. So the F' term of LB_{end} at machine r is $(m-1)a + L_r/F' + \alpha_0(F'+1)/2$, which by (F5) is non-increasing in F' (its decrease from F' to $F'+1 \le F$ is equivalent to $L_r \ge \alpha_0 F'(F'+1)/2$). Hence $\text{LB}_{\text{end}} \ge \lceil \min_{F'} \text{term} \rceil = M$.

Partition. Factory f gets up-ramps U_f ($\beta_\uparrow$ copies of v_r), down-ramps Dn_f ($\beta_\downarrow$ copies of v_r with $A = \alpha_0$), a reservoir R_f (E_T designated τ_r^0 and E_T designated τ_r^1), and a Steinitz block B_f of the core C of all other products, by Lemma 48 applied to the vectors (x_h, A_h) with $d = m+1$, $\bar{Z} = \bar{s}b$ for the x coordinates and $\bar{Z} = \bar{A}$ for A. The deviations are therefore at most dev_x and dev_A.

Targets are set on coordinate r only: $T_f = M - (m-1)a - f\alpha_0$. Since $\alpha_0 F(F+1)/2$ is an integer, the ceiling in M gives $\sum_f T_f - L_r \in [0, F)$. Put $t_f = T_f - [f \le \sum_{f'} T_{f'} - L_r]$, so that $\sum_f t_f = L_r$ and $|t_f - L_r/F| \le (F-1)\alpha_0/2 + 1 \le F\alpha_0 + 1 =: \theta$. The ramps and reservoirs are the same vectors in every factory, so Lemma 15 with $K = \{r\}$ and $E = E_T = \lceil (2(m+1)+1)\bar{s}b + \theta \rceil$ gives $x_r(G'_f) = t_f \le M - (m-1)a - S_f$ with $S_f = f\alpha_0$. This is (T2) with $\rho(f) = f$.

Reordering. After the exchanges, each $\text{Mid}_f = B_f \cup R'_f$ is reordered within its factory by Lemma 5 on Mid_f's own centred vectors $(w_g - \bar{w},\ y_g - \bar{y})$, with Mid_f's own means, as Lemma 53 requires. The global order of C served only to form the blocks. Reordering changes none of the factory loads used in (T2), in margin (a) or in $y(\text{Mid}_f)$.

Margin (a) of Lemma 53 in every factory. Let $i \ne r$. Then $(x_r - x_i)(B_f) \ge (x_r - x_i)(C)/F - 2\,\text{dev}_x$. Every reservoir item, before or after the exchanges, has $x_r - x_i \in \{s - a, s + 1 - a\} \ge 0$, so the reservoir part of Mid_f contributes ≥ 0 item by item. The ramps contribute $b - a$ each and the reservoirs $E_T(s-a) + E_T(s+1-a)$ per factory, so

$$(x_r - x_i)(C) = L_r - L_i - F(\beta_\uparrow + \beta_\downarrow)(b-a) - FE_T(2s+1-2a) \ge F(2\,\text{dev}_x + E_T),$$

by (F3). Hence $(x_r - x_i)(\text{Mid}_f) \ge E_T \ge 0$.

$y(\text{Mid}_f) \le 0$. Exchanges move only reservoir items, and $y(R'_f) \le 2E_T\bar{A}$. Next, $y(B_f) \le A(C)/F + \text{dev}_A - x_r(C)/F^2 + \text{dev}_x/F$, with $A(C) \le A_{\text{tot}}$ and $x_r(C) \ge L_r - F(\beta_\uparrow + \beta_\downarrow)b - 2FE_Tb$ (the r-entries of the ramps and reservoir items are at most b). So

$$y(\text{Mid}_f) \le \frac{FA_{\text{tot}} - L_r}{F^2} + Y_1 \le 0,$$

by (F4).

The hypotheses of Theorem 8.

- (T1) and (T3): Lemma 53, with $E_\uparrow = E_\downarrow = E$ and $W_\uparrow, W_\downarrow \le W$; condition (b) of the lemma holds by the first term in the definitions of $\beta_\uparrow$ and $\beta_\downarrow$.
- (T2): shown above, with $S_\rho = \rho\alpha_0$.
- (T4): $F\alpha^* = F\alpha_0 < b$, since all down-ramps have $A = \alpha_0$; then γ, Λ_0 (with $E_\downarrow = E$) and k_0 are as in (T4).
- (T5a): $\beta_\downarrow\gamma \ge \Lambda_0 + \bar{x}_r + FE_A$, since $\bar{x}_r \le \bar{s}b$ (third term of $\beta_\downarrow$).
- (T5b): $(\beta_\downarrow - k_0)\gamma \ge \Lambda_0 + F\alpha_0$ (second term of $\beta_\downarrow$).

Conclusion. Theorem 8 gives, for the schedule of Theorem 8 built as above, makespan $\le M \le \text{LB}_{\text{end}} \le X \le \text{OPT} \le$ makespan, for every valid lower bound $X \ge \text{LB}_{\text{end}}$ (LB_{comb} included). So $\text{OPT} = M = \text{LB}_{\text{end}} = X$. Every step (Lemma 5, the exchanges, the schedule) is polynomial. □

Remark 54 (size of the constants) The constants are astronomic, as in Theorem A of Coutinho (manuscript under review). For $m = 2$, $\bar{s} = 1$, $U\{1..3\}$ ($a = 1$, $b = 3$, $s = 1$), $F = 2$, $\alpha_0 = 1$ and $\bar{A} = 4$: $K = 10$, $E = 20$, $W = 2$, $Y = 5.5$, $\gamma = 1$, $\mathrm{dev}_x = 21$, $\mathrm{dev}_A = 28$ and $E_T = 24$. At $r = m = 2$ (the larger case), $\beta_\uparrow = 6$, $E_A = 46$, $\Lambda_0 = k_0 = 0$, $\beta_\downarrow = 95$, $\Delta^* = 584$, $Y_1 = 454$ and $c_1 = 1{,}816$. At $r = 1$, $\beta_\uparrow = 0$, $E_A = 22$, $\Lambda_0 = 22$, $k_0 = 11$, $\beta_\downarrow = 69$, $\Delta^* = 456$ and $c_1 = 1{,}624$. A constructive check that builds the certificate with exactly these constants attained $M = \mathrm{LB}_{\mathrm{end}}$ on every built instance.

Proof of Theorem 19(a)–(d).

Assume lattice model M with $F \geq 2$, and (M2) (32) for $\beta \geq 1$. All constants of Theorem 18 are now indexed by the machine r; put $\Delta^*_{\max} = \max_r \Delta^*(r)$. Conditionally on the sizes, $n \geq t$ and the entries are i.i.d. μ with variance σ^2. By the load-margin lemma of Coutinho (manuscript under review) (Berry–Esseen, lattice laws included),

$$\Pr(\text{no } \Delta^*_{\max}\text{-margin}) \leq \binom{m}{2} \frac{\Delta^*_{\max}/\sqrt{\pi} + 0.672\,(b-a)}{\sigma\sqrt{t}}.$$

Proof of Theorem 19(a), (b) Let $r^* = \arg\max_i L_i$ (smallest index).

Events.

- E_M: $L_{r^*} - L_i \geq \Delta^*_{\max}$ for every $i \neq r^*$. By the display above, $\Pr(E_M^c) = O(t^{-1/2})$.
- E_C: the counts (F2) for every r, plus the atom-case counts of Theorem 9. These are binomial with positive success probabilities ($\nu(1)$, $\mu(a)$, $\mu(b)$, $\mu(s)$, $\mu(s+1)$, $\kappa_1(\{\alpha_0\})$), so $\Pr(E_C^c) = e^{-\Omega(t)}$.
- E_L: (F5) and the size conditions of the same type (such as L_* in Theorem 9); probability $1 - e^{-\Omega(t)}$.

The two certificates. On $E_M \cap E_C \cap E_L$:

- if $FA_{\mathrm{tot}} \leq L_{r^*} - c_1(r^*)$, Theorem 18 at $r = r^*$ gives $\mathrm{OPT} = \mathrm{LB}_{\mathrm{end}}$, so $X = \mathrm{OPT}$ for every valid $X \geq \mathrm{LB}_{\mathrm{end}}$;
- if $FA_{\mathrm{tot}} - L_{r^*} \geq c_0$, where c_0 is the constant of the load-margin event in the proof of Theorem 9, that proof (with its margin event replaced by this event; atom case) gives $\mathrm{OPT} = \mathrm{LB}_0$, so $X = \mathrm{OPT}$ for every admissible X.

The window. What remains is

$$\mathcal{W}_t := \{FA_{\mathrm{tot}} - L_{r^*} \in (-c_1(r^*), c_0)\} \subseteq \bigcup_{r=1}^{m} \Big\{ \sum_h \zeta_{h,r} \in (-c_1(r), c_0) \Big\}, \qquad \zeta_{h,r} = FA_h - x_{h,r},$$

where, for each r, the $\zeta_{h,r}$ are i.i.d. over the t products and bounded.

(a) $\beta < 1$. Then $\mathbb{E}\zeta_{\cdot,r} = -\mathbb{E}[S]\,\mathbb{E}p\,(1-\beta) < 0$, so Hoeffding's inequality gives $\Pr(\mathcal{W}_t) \leq m\,e^{-\Omega(t)}$. Hence $\Pr[\mathrm{OPT} \neq X] \leq \Pr(E_M^c) + e^{-\Omega(t)} = O(t^{-1/2})$ for every valid $X \geq \mathrm{LB}_{\mathrm{end}}$ ($\mathrm{LB}_{\mathrm{end}}$ itself included, which need not be $\geq \mathrm{LB}_0$), in particular for every $X \in \mathcal{X}$. On the good events of the proof, $\mathrm{OPT} = \mathrm{LB}_{\mathrm{end}}$.

(b) $\beta = 1$. Then $\mathbb{E}\zeta_{\cdot,r} = 0$. A_h and $x_{h,r}$ both depend on the size S, so A is not independent of the entries unconditionally; but given S they are independent, so

$$\mathrm{Var}\,\zeta_{\cdot,r} \geq \mathbb{E}[\mathrm{Var}(\zeta_{\cdot,r} \mid S)] \geq \mathbb{E}[\mathrm{Var}(x_{h,r} \mid S)] = \mathbb{E}[S]\sigma^2 > 0.$$

Write $\tau_r^2 = \operatorname{Var}\zeta_{\cdot,r}$. Since ζ is bounded, $\mathbb{E}|\zeta|^3 < \infty$. The Berry–Esseen inequality (with constant C_{BE}; Shevtsova 2011) and the union over r give

$$\Pr(\mathcal{W}_t) \le \sum_{r=1}^{m} \Big[\frac{c_0 + c_1(r)}{\sqrt{2\pi}\,\tau_r \sqrt{t}} + \frac{2C_{\mathrm{BE}}\mathbb{E}|\zeta_{\cdot,r}|^3}{\tau_r^3 \sqrt{t}} \Big] = O(t^{-1/2}).$$

Both ends of the window are fixed constants ($c_1(r)$ depends only on $(m, F, a, b, s, \bar{s}, \bar{A}, \alpha_0)$ and r; c_0 only on the constants of Theorem 9), not on the realised A_{tot}, so there is no circularity. Hence $\Pr[\mathrm{OPT} \neq X] \le \Pr(E_M^c) + \Pr(\mathcal{W}_t) + e^{-\Omega(t)} = O(t^{-1/2})$ for every $X \in \mathcal{X}$; both certificates are needed here, which is why $X \ge \mathrm{LB}_0$ is required. □

Proof of Theorem 19(c), (d) (c) For $\beta > 1$, (L) gives an atom at the lower end a of μ, and (M2) is assumed, so the atom case of Theorem 9 gives $\Pr[\mathrm{OPT} \neq X] = e^{-\Omega(t)}$ for every admissible X.

(d) Let also $F\alpha_0 > a$. Theorem 19 assumes $F\alpha_0 < b$; for $\beta < 1$ this is implied, since $A_h \ge s_h\alpha_0$ gives $\beta \ge F\alpha_0/\mathbb{E}p$, hence $F\alpha_0 < \mathbb{E}p < b$. For $\beta < 1$, Proposition 17(ii) (Section B.3; for $a < F\alpha_0 < a+2$ via the residue-class event of Section B.5) gives $\Pr[\mathrm{OPT} > X] \ge c\,t^{-1/2}$ for $X \in \{\mathrm{LB}_{\mathrm{end}}, \mathrm{LB}_{\mathrm{comb}}\}$, and part (a) gives the matching $O(t^{-1/2})$ for both, since both are valid lower bounds $\ge \mathrm{LB}_{\mathrm{end}}$. For $\beta = 1$, Proposition 17(ii) (Section B.5) gives the lower bound for $X = \mathrm{LB}_{\mathrm{comb}}$, and part (b) the upper bound, since $\mathrm{LB}_{\mathrm{comb}} \in \mathcal{X}$. So the rate is $\Theta(t^{-1/2})$ in these cases. For $\mathrm{LB}_{\mathrm{end}}$ alone at $\beta = 1$ there is no matching upper bound, since $\mathrm{LB}_{\mathrm{end}}$ need not be $\ge \mathrm{LB}_0$. □

B.5 The lower bound at $\beta = 1$

Proof of Proposition 17(ii) at $\beta = 1$ Assume lattice model M with $\beta = 1$: μ satisfies (L); $m \ge 2$; $\nu(1) > 0$; κ_s is integer-valued on $[s\alpha_0, \bar{A}]$ with $\kappa_1(\{\alpha_0\}) > 0$ and α_0 an integer; and $a < F\alpha_0 < b$. We show that there are $c > 0$ and t_0 with $\Pr[\mathrm{OPT} > \mathrm{LB}_{\mathrm{comb}}] \ge c\,t^{-1/2}$, hence also $\Pr[\mathrm{OPT} > \mathrm{LB}_{\mathrm{end}}] \ge c\,t^{-1/2}$, for $t \ge t_0$. Put $K' = F\bar{s}mb$ and $\bar{p} = \mathbb{E}p$; note $\bar{p} > 0$ since $b > a \ge 0$.

Entries (local event). Condition on the sizes, so $n = \sum_h s_h \in [t, \bar{s}t]$. The job vectors $p_j \in \mathbb{Z}^m$ are i.i.d., with mean $\bar{p}\mathbf{1}$ and covariance $\sigma^2 I$, non-singular since $a < b$. Their support differences contain every e_k (vary one coordinate between s and $s+1$), so $\mathbb{Z}^m$ is their minimal lattice. The multivariate lattice local limit theorem (Bhattacharya and Rao 1976, Thm. 22.1, with second moments) gives

$$\sup_{z \in \mathbb{Z}^m} \Big| n^{m/2} \Pr\Big(\sum_j p_j = z\Big) - \varphi_{\sigma^2 I}\big((z - n\bar{p}\mathbf{1})/\sqrt{n}\big) \Big| \to 0,$$

where $\varphi_{\sigma^2 I}$ is the centred normal density with covariance $\sigma^2 I$. On $|z - n\bar{p}\mathbf{1}| \le 3\sqrt{mn}$ the density is at least $(2\pi\sigma^2)^{-m/2} e^{-9m/(2\sigma^2)} > 0$. Hence $n^{m/2} \Pr(\sum_j p_j = z) \ge c_0 > 0$ uniformly over $|z - n\bar{p}\mathbf{1}| \le 3\sqrt{mn}$ and $n \ge n_0$.

Let

$$Z_n := \big\{ z \in \mathbb{Z}^m : \; z_1 = z_2 \in n\bar{p} + [\sqrt{n}, 2\sqrt{n}],\; z_1 - \sqrt{n} \le z_i \le z_1 \text{ for } i \ge 3 \big\}.$$

Every $z \in Z_n$ has coordinate deviations in $[0, 2\sqrt{n}]$, so it lies inside that ball, and Z_n has at least $c\,n^{(m-1)/2}$ points. On $\{\sum_j p_j \in Z_n\}$, $L_1 = L_2 = \max_i L_i \ge n\bar{p} + \sqrt{n}$, which is condition (2) with $(k, l) = (1, 2)$.

If $F\alpha_0 \geq a+2$, condition (4), and hence (4′), holds deterministically, and we use Z_n. If $a < F\alpha_0 < a+2$, we use instead

$$Z_n^\rho := \{z \in Z_n : \ z_1 \equiv \rho \pmod F\}, \qquad \rho \equiv -\alpha_0 F(F+1)/2 \pmod F.$$

On Z_n^ρ, $L = z_1$ and $x = (L + \alpha_0 F(F+1)/2)/F \in \mathbb{Z}$, so $\lceil x \rceil - x = 0 < (F\alpha_0 - a)/2$ and (4′) holds. Z_n^ρ has at least $\lfloor \sqrt{n}/F \rfloor \cdot \lfloor \sqrt{n} \rfloor^{m-2} \geq c\, n^{(m-1)/2}/F$ points inside the ball. In either case, writing Z for the event $\{\sum_j p_j \in Z_n\}$ (respectively Z_n^ρ),

$$\Pr(Z \mid \text{sizes}) \geq c' n^{-1/2} \geq c'(\bar{s}t)^{-1/2},$$

with c' including the factor $1/F$ in the second case.

Assembly (CLT event). The variables $\zeta_h := FA_h - s_h\bar{p}$ are i.i.d. over products, bounded, and have $\mathbb{E}\zeta = F\,\mathbb{E}A - \mathbb{E}S \cdot \bar{p} = 0$ at $\beta = 1$. Let $G_A := \{\sum_h \zeta_h \leq \sqrt{t}/2\}$. By the central limit theorem, $\Pr(G_A) \to \Phi(1/(2\tau)) > 1/2$ if $\tau^2 = \operatorname{Var}\zeta > 0$, and $\Pr(G_A) = 1$ if $\zeta \equiv 0$. So $\Pr(G_A) \geq 1/2$ for $t \geq t_1$. Given the sizes, G_A depends only on the assembly times, which are independent of the entries. Since the bound on $\Pr(Z \mid \text{sizes})$ is uniform over the sizes,

$$\Pr(Z \cap G_A) = \mathbb{E}\big[\Pr(Z \mid \text{sizes})\ \Pr(G_A \mid \text{sizes})\big] \geq c'(\bar{s}t)^{-1/2} \Pr(G_A) \geq \tfrac{c'}{2}(\bar{s}t)^{-1/2}.$$

(6a) on the intersection. Since $FA_{\text{tot}} = \sum_h(\zeta_h + s_h\bar{p}) = n\bar{p} + \sum_h \zeta_h$ and $n \geq t$,

$$FA_{\text{tot}} \leq n\bar{p} + \sqrt{t}/2 \leq L_1 - \sqrt{n} + \sqrt{t}/2 \leq L_1 - \sqrt{t}/2 \leq L_1 - K'$$

once $\sqrt{t} \geq 2K'$.

The other conditions of Proposition 17(i).

- (1): an all-a job exists;
- (3): at least $2F$ single-job products with $A = \alpha_0$ (the lower bound $A_h \geq s_h\alpha_0$ holds by the model);
- (6b): for each k, at least $F\bar{A}$ single-job products with $A = \alpha_0$ and $p_k \geq F\alpha_0 + 1$ (possible since $F\alpha_0 + 1 \leq b$ and b is an atom);
- (5): deterministic on Z_n, since $L = L_1 \geq n\bar{p} + \sqrt{n} \geq t\bar{p}$ with $\bar{p} > 0$, so it holds for $t \geq L_*/\bar{p}$. (The bound $L \geq na$ is not used; it fails when $a = 0$.)

Each of (1), (3), (6b) fails with probability $e^{-\Omega(t)}$ (binomial counts with positive success probability); (5) costs nothing. No independence from Z is needed: $\Pr(Z \cap G_A \cap \text{Good}) \geq \Pr(Z \cap G_A) - \Pr(\text{Good}^c)$. Integer data hold in model M, and $n \geq 2F$ for $t \geq 2F$, so Propositions 3 and 4 apply with W defined up to $2F$.

Conclusion. On $Z \cap G_A \cap \text{Good}$, Proposition 17(i) gives $\text{OPT} \geq \text{LB}_{\text{comb}} + 1$. So $\Pr[\text{OPT} > \text{LB}_{\text{comb}}] \geq \frac{c'}{2}(\bar{s}t)^{-1/2} - e^{-\Omega(t)}$. □

Remark 55 (scope)

- The same proof covers $\beta < 1$, where $\mathbb{E}\zeta < 0$ and G_A has probability tending to 1; so Proposition 17(ii) holds for $\beta \leq 1$. It fails for $\beta > 1$, where $FA_{\text{tot}} - L_1 \asymp +t$ and Theorem 9 gives exponential exactness.
- For $X = \text{LB}_{\text{end}}$ alone, conditions (6a), (6b) and $F\alpha_0 < b$ are not needed, since part (i) without them already gives $\text{OPT} \geq \text{LB}_{\text{end}} + 1$; this part of the argument uses no restriction on β.
- With Theorem 19(b), which additionally needs $F \geq 2$ and (M2) (these are not implied here: for example $F = 2$, $a = 2$, $\alpha_0 = 2$, $\kappa_1 = \delta_2$ violates (M2)), the rate at $\beta = 1$ is $\Theta(t^{-1/2})$ for $X = \text{LB}_{\text{comb}}$. For LB_{end} alone there is no matching upper bound at $\beta = 1$, since LB_{end} need not be $\geq \text{LB}_0$.

- The joint lattice of $(L_1 - L_i,\ FA_{\mathrm{tot}} - L_1)$, which depends on κ, is avoided: the local limit theorem is used for the entries alone, and the assembly times enter only through a CLT event that is conditionally independent of them.

Remark 56 (an elementary route to the local bound) In model M the m coordinates of p_j are independent, so $\Pr(\sum_j p_j = z) = \prod_i \Pr(L_i = z_i)$, and the one-dimensional lattice local limit theorem (Gnedenko and Kolmogorov 1954), with span 1 since $s, s+1 \in \operatorname{supp}\mu$, gives the same uniform bound $n^{m/2}\Pr(\sum_j p_j = z) \geq c_0$ on the ball. This avoids the multivariate theorem altogether. An exact computation of $n^{1/2}\Pr(Z_n)$ for $U\{1..9\}$ and $m = 2, \ldots, 5$ shows convergence to positive constants, so the order $n^{-1/2}$ is attained.

Appendix C Proofs for the threshold (Section 6)

C.1 Harmless ties on and below the threshold

Throughout this section $m \geq 2$, the data are integers, the entries lie in $[a, b]$ with $g = b - a > 0$, every product has at most $\bar{s}$ jobs and assembly time at most $\bar{A}$, and $\alpha_0 \geq 0$ is an integer with $A_h \geq s_h\alpha_0$ for every product. We write $Y = \bar{A} + \bar{s}b/F$ and $y_h = A_h - x_{h,m}/F$. A *front job of type* r $(r = 2, \ldots, m)$ has entries a on machines $i < r$ and b on machines $i \geq r$; an *end job of type* r $(r = 1, \ldots, m-1)$ has b on machines $i \leq r$ and a on machines $i > r$. The buffer vector is $w = a\mathbf{1} + g\,e_m$, which is the front type m.

We use four lemmas of the companion analysis (Coutinho, manuscript under review), stated there for a single flowshop in units shifted by $-a$ (entries in $[0, g]$). For a sequence S and machines $i \leq i'$, $C_{i\to i'}(S)$ is the largest value of a path from (machine i, first job of S) to (machine i', last job of S). Let $\mathcal{F}$ be c front jobs of each type in decreasing type order, $\mathcal{E}$ be c end jobs of each type in decreasing type order, and σ an order of the other jobs $\mathcal{M}$, with $\Lambda_{\mathcal{M}} = \max_l L_l(\mathcal{M})$.

- *Absorption* (exact atoms): $\mathcal{F}\sigma\mathcal{E}$ has makespan $\max_l L_l$ if $C_{i\to i'}(\sigma) \leq \Lambda_{\mathcal{M}} + c(i' - i)g$ for all $i < i'$.
- *Padding*: (i) deleting jobs from an order does not increase any $C_{i\to i'}$; (ii) adding $\lceil(\Lambda_{\mathcal{M}} - L_l(\mathcal{M}))/g\rceil$ jobs with a single nonzero entry $(\leq g)$ on each machine l equalises all loads at $\Lambda_{\mathcal{M}}$.
- *One-sided summation*: if $\mathcal{M}'$ has equal loads Λ and an order of $\mathcal{M}'$ has prefix differences $D_l = P_l - P_{l+1}$ with $\min\{D_l(k-1), D_l(k)\} \leq \beta_l$ for all $l \leq m-1$ and all k, then $C_{i\to i'} \leq \Lambda + \sum_{l=i}^{i'-1}(g + \beta_l)$.

Proof of Lemma 20 Work in shifted units. Let $\mathcal{M} = \mathrm{Mid} \cup \mathrm{Buf}$, and pad $\mathcal{M}$ to equal loads $\Lambda = \Lambda_{\mathcal{M}}$; the padding items are placed in Mid. Let N be the number of items of Mid (real products and padding items), with $N \geq K_B$ by assumption.

Sums. Write $v_h = (x_{h,l} - x_{h,l+1})_{l \leq m-1}$ for an item h. After padding, the difference vectors of $\mathcal{M}$ sum to 0. A buffer job has shifted vector $g\,e_m$, so it changes only $x_{m-1} - x_m$, by $-g$. Hence $S := \sum_{h\in\mathrm{Mid}} v_h = K_B g\,e_{m-1}$.

Order of Mid. Apply Lemma 5 with $d = m$ to the items of Mid with the vectors

$$u_h = \Big(\frac{v_h - S/N}{\rho},\ \frac{y_h - \bar{y}}{Y}\Big), \qquad \rho = (\bar{s}+1)g,$$

where $\bar{y}$ is the mean of y over the items of Mid (a padding item has $y = 0$). A real product has $|v_{h,l}| \le \bar{s}g$, a padding item $|v_{h,l}| \le g$, and $|S/N| \le K_B g/N \le g$; so $\|u_h\|_\infty \le 1$, and $\sum_h u_h = 0$. Every prefix of k items therefore satisfies $|P_k - (k/N)S| \le m\rho$ in each difference coordinate and $|y(\text{prefix}_k) - k\bar{y}| \le mY$.

Prefix differences. At item boundaries in Mid, $|D_l| \le m\rho + K_B g$. Inside a product, D_l moves by at most $\bar{s}g$ from its value at the preceding boundary. So $|D_l| \le ((m+1)\bar{s} + m + K_B)g =: \beta^*$ at every job position in Mid. At the end of Mid the prefix difference is exactly S, so $D_{m-1} = K_B g$ and $D_l = 0$ for $l < m-1$. Inside Buf, D_{m-1} decreases by g per job down to 0, and the other D_l stay at 0. Hence $|D_l| \le \beta^*$ at every position of Mid $\cdot$ Buf.

Makespan. One-sided summation with $\beta_l = \beta^*$ gives $C_{i\to i'} \le \Lambda + (i'-i)(g+\beta^*) = \Lambda + (i'-i)((m+1)(\bar{s}+1) + K_B)g \le \Lambda + c(i'-i)g$ for the padded order. Padding (i) gives the same for the restriction to $\mathcal{M}$. Absorption then gives makespan $\max_l L_l$ in shifted units. Undoing the shift, which lowers every path by $(n_f + m - 1)a$ for a factory with n_f jobs, the factory makespan is $\max_k L_k + (m-1)a$, which proves (a).

Suffixes, part (b). Let Z be a product-aligned suffix of Mid, and Z' the corresponding suffix of the item order, padding items included, with k items before it. Then $y(Z) = y(Z') = y(\text{Mid}) - y(\text{prefix}_k) \le N\bar{y} - (k\bar{y} - mY) = (N-k)\bar{y} + mY$. If $y(\text{Mid}) \le 0$, then $\bar{y} \le 0$ and $y(Z) \le mY$. □

Lemma 57 (slots) *Let F factories have completion times $C_f \le M - f\alpha_0$ $(f = 1, \dots, F)$. Suppose the last K jobs of each factory are single-job products with $A = \alpha_0$ and $p_m \ge a$, and the K_B jobs before them are single-job products with $A = \alpha_0$ and $p_m \ge b$. If $F\alpha_0 \le a$, then:*

- *(i) the j-th of the last K jobs from the end of factory f completes by $\hat{c} = M - f\alpha_0 - (j-1)a$, and the i-th buffer job from the end of the buffer by $\hat{c} = M - f\alpha_0 - Ka - (i-1)b$;*
- *(ii) the intervals $[\hat{c}, \hat{c} + \alpha_0]$ of all these jobs are pairwise disjoint and lie in $[\hat{c}_{\min}, M]$, so $\sum_{\text{such jobs with } \hat{c} \ge M-\tau} \alpha_0 \le \tau$ for every $\tau \ge 0$;*
- *(iii) the total assembly time of all $F(K+K_B)$ such jobs is $F(K+K_B)\alpha_0 = (Ka + K_B b) - K(a - F\alpha_0) - K_B(b - F\alpha_0)$.*

Proof (i) Machine m processes a factory's jobs in sequence, so a job completes by C_f minus the machine-m entries of the jobs after it. (ii) For a fixed end position j, the F intervals are consecutive, of length α_0, and fill $[M - (j-1)a - F\alpha_0,\ M - (j-1)a]$, a window of length $F\alpha_0 \le a$. For a fixed buffer position i they fill $[M - Ka - (i-1)b - F\alpha_0,\ M - Ka - (i-1)b]$, of length $F\alpha_0 \le a < b$. Consecutive windows start at least a (end part) or b (buffer part) apart, so all windows are disjoint. A counted interval lies in $[M - \tau, M]$. (iii) is arithmetic. □

Proof of Theorem 21(a) Put $Y = \bar{A} + \bar{s}b/F$, $E_A = mY$, $\theta = F\alpha_0/2 + 1$ and

$$K_E = \begin{cases} 0, & F\alpha_0 = a, \\ \lceil (\bar{s}b + FE_A)/(a - F\alpha_0) \rceil, & F\alpha_0 < a, \end{cases}$$

$$K_B = \Big\lceil \frac{\bar{s}b + FE_A}{b - F\alpha_0} \Big\rceil \qquad (K_B = 0 \text{ is allowed if } F\alpha_0 < a),$$

$$c = \max\{(m+1)(\bar{s}+1) + K_B,\ \lceil K_E/(m-1)\rceil\},$$
$$E = \max\{\lceil (2m+3)\bar{s}b + \theta\rceil,\ K_B\},$$
$$c_2 = F^2\Big[(2m+3)Y + 2mE\bar{A} + \big(2(m-1)c + 2mE + K_B\big)\frac{b}{F} + \big(c(m-1) + K_B\big)\bar{A}\Big].$$

Any $K_E, K_B \geq 0$ with $K_E(a - F\alpha_0) + K_B(b - F\alpha_0) \geq \bar{s}b + FE_A$ works in the argument below; the stated choice is one such pair. Let $M = (m-1)a + \lceil (L + \alpha_0 F(F+1)/2)/F\rceil$.

Step 0: $\mathrm{LB}_{\mathrm{end}} \geq M$. Every entry is at least a, a front job of type m has $\sum_{i<k} p_i = (k-1)a$ and an end job of type 1 has $\sum_{i>k} p_i = (m-k)a$; so $\mathrm{hd}_k + \mathrm{tl}_k = (m-1)a$ for every k. By $A_h \geq s_h\alpha_0$, $W(r) \geq r\alpha_0$. At $k^* = \arg\max_k L_k$ the F' term of Proposition 3 is therefore at least $\phi(F') := (m-1)a + L/F' + \alpha_0(F'+1)/2$. For $1 \leq F' \leq F-1$, $\phi(F') - \phi(F'+1) = L/(F'(F'+1)) - \alpha_0/2 \geq 0$ because $L \geq \alpha_0 F(F-1)/2 \geq \alpha_0 F'(F'+1)/2$. So the minimum over F' is at least $\phi(F)$, and, the data being integers, $\mathrm{LB}_{\mathrm{end}} \geq \lceil \phi(F)\rceil = M$ (for integer data $\mathrm{LB}_{\mathrm{end}}$ denotes the rounded-up bound; Section 6).

Step 1: partition. Factory f receives as its fixed part c of each front type, c of each end type (with $A = \alpha_0$), K_B buffer jobs (with $A = \alpha_0$), and E each of τ_k^0 and τ_k^1 for every k. The remaining products form the core C and are split into Steinitz blocks B_f (Lemma 48 on (x_h, A_h), $d = m+1$, $\bar{Z} = \bar{s}b$ for machine coordinates and $\bar{A}$ for A). Put $T_f = M - (m-1)a - f\alpha_0$, so $\sum_f T_f = FM - F(m-1)a - \alpha_0 F(F+1)/2 \geq L \geq L_k$ for every k. For each k, write $e_k = \sum_f T_f - L_k \geq 0$ and set $t_{f,k} = T_f - \lfloor e_k/F\rfloor - [f \leq e_k \bmod F]$. Then $\sum_f t_{f,k} = L_k$, $t_{f,k} \leq T_f$, and $|t_{f,k} - L_k/F| \leq |T_f - \sum_{f'} T_{f'}/F| + 1 \leq \theta$ for every k, whatever L_k is. Lemma 15 with $K = \{1,\dots,m\}$ and $E \geq (2m+3)\bar{s}b + \theta$ gives exact loads $t_{f,k}$. The Steinitz order of Lemma 48 only forms the blocks. After the exchanges, the middle Mid_f of factory f consists of B_f and the reservoir items R'_f ($2mE \geq K_B$ items), and it is reordered within the factory by Lemma 20, with its own padding and means. This changes no factory load.

Step 2: factory completions. Factory f is sequenced as $\mathrm{Front} \cdot \mathrm{Mid}_f \cdot \mathrm{Buf}_f \cdot \mathrm{End}$. By Lemma 20(a), $C_f = \max_k t_{f,k} + (m-1)a \leq M - f\alpha_0$.

Step 3: releases. Product h in factory f is complete by $\hat{c}_h := M - f\alpha_0 - x_m(\text{jobs after } h \text{ in } f)$, since machine m processes the jobs after h before C_f. By the monotonicity in Lemma 1 it suffices to show $\mathrm{ERD}(\hat{c}) \leq M$, i.e. $D(\tau) := \sum_{h:\,\hat{c}_h \geq M-\tau} A_h \leq \tau$ for every $\tau \geq 0$.

Step 4: assembly. The counted set of factory f is a suffix of its sequence: part of its $K_f := c(m-1) \geq K_E$ end jobs; then its K_B buffer jobs; then possibly a suffix Z'_f of Mid_f; then possibly front jobs. Let J be the set of factories whose counted set reaches into Mid_f.

(a) $J = \emptyset$. Only end and buffer jobs are counted, and Lemma 57(ii) gives $D(\tau) \leq \tau$.

(b) $J \neq \emptyset$. Some product of a middle is counted, so $\tau \geq K_f a + K_B b$, and all $F(K_f + K_B)$ end and buffer jobs are counted (their $\hat{c}$ exceed that of any middle product of any factory by Lemma 57(i) and $F\alpha_0 \leq a$). They contribute $F(K_f + K_B)\alpha_0$. For $f \in J$ with no front job counted, $x_m(Z'_f) \leq \bar{s}b + \tau - f\alpha_0 - K_f a - K_B b$ (the last counted product may straddle $M - \tau$ by at most $\bar{s}b$), and Lemma 20(b) gives $A(Z'_f) \leq x_m(Z'_f)/F + E_A$, provided $y(\mathrm{Mid}_f) \leq 0$, shown in (c). Summing over $f \in J$ ($|J| \leq F$),

$$\begin{aligned} D(\tau) &\leq F(K_f + K_B)\alpha_0 + (\tau - K_f a - K_B b) + \bar{s}b + FE_A \\ &= \tau - K_f(a - F\alpha_0) - K_B(b - F\alpha_0) + \bar{s}b + FE_A \leq \tau, \end{aligned}$$

by the choice of $K_E \leq K_f$ and K_B. When $F\alpha_0 = a$ the buffer term alone does the work.

(c) Front jobs counted for some $f \in J$ (a sub-case of (b); each factory of J is bounded by (b) or by (c)). Then all of Mid_f and Buf_f are counted. The contribution of Mid_f and of the front part is at most $A(\mathrm{Mid}_f) + c(m-1)\bar{A} = x_m(\mathrm{Mid}_f)/F + y(\mathrm{Mid}_f) + c(m-1)\bar{A}$. We bound $y(\mathrm{Mid}_f)$. The blocks satisfy $A(B_f) \leq A(C)/F + (2m+3)\bar{A}$ and $x_m(B_f) \geq$

$x_m(C)/F - (2m+3)\bar{s}b$ (Lemma 48, $d = m+1$); the reservoir adds at most $2mE\bar{A}$ to A and nothing negative to x_m; and $A(C) \le A_{\text{tot}}$, $x_m(C) \ge L_m - F(2(m-1)c + 2mE + K_B)b$. Hence

$$\begin{aligned} y(\text{Mid}_f) &\le \frac{FA_{\text{tot}} - L_m}{F^2} + (2m+3)Y + 2mE\bar{A} + \frac{(2(m-1)c + 2mE + K_B)b}{F} \\ &\le -\big(c(m-1) + K_B\big)\bar{A} \end{aligned}$$

by $FA_{\text{tot}} \le L_m - c_2$. So the contribution of this factory's middle and front part is at most $x_m(\text{Mid}_f)/F \le (\tau - f\alpha_0 - K_f a - K_B b)/F$, and the bound of (b) holds with no $\bar{s}b + E_A$ term for that factory. This also shows $y(\text{Mid}_f) \le 0$ for every factory, as (b) requires, since the same bound applies to every f.

Conclusion. The ERD assembly finishes by M. With Step 0, $M \le \text{LB}_{\text{end}} \le X \le \text{OPT} \le M$ for every valid lower bound $X \ge \text{LB}_{\text{end}}$. The certificate (blocks, exchanges, Steinitz orders, padding) is computed in polynomial time. □

Proof of Theorem 21(b) Assume lattice model M (Section 6) with $m \ge 2$, $F\alpha_0 \le a$ and $\beta < 1$, and let the constants be those of part (a).

- The required counts (front, end and buffer types with the stated assembly times, and τ_k^0, τ_k^1) are binomial over the t products, with success probabilities at least $\nu(1)\kappa_1(\{\alpha_0\}) \min(\mu(a), \mu(b), \mu(s), \mu(s+1))^m > 0$, against constant demands; each fails with probability $e^{-\Omega(t)}$ (Chernoff).
- $FA_{\text{tot}} \le L_m - c_2$: $\sum_h (FA_h - x_{h,m})$ is a sum of t i.i.d. bounded terms with mean $-t\,\mathbb{E}[S]\,\mathbb{E}p\,(1-\beta) < 0$, so Hoeffding's inequality (Hoeffding 1963) gives failure probability $e^{-\Omega(t)}$.
- $L \ge \alpha_0 F(F-1)/2$: $L \ge na \ge t$ for $a \ge 1$, so it holds for $t \ge \alpha_0 F^2$.
- $A_h \ge s_h \alpha_0$ holds by the model.

A union bound gives $\Pr[\text{OPT} \ne X \text{ for some valid } X \ge \text{LB}_{\text{end}}] \le e^{-\Omega(t)}$. □

Remark 58 (why the buffer is placed inside the padding) If Mid_f alone were padded to equal loads and the buffer appended afterwards, padding would raise machine m of the middle to Λ_{Mid} and the buffer would add $K_B g$ on top of it, which can exceed the factory's maximum load. A constructive check of this variant (planted exact ties, $F\alpha_0 = a$, $K_B = 20$, six parameter cells, 90 instances) exceeded the factory completion targets on every instance, while the construction of Lemma 20, with the padding computed for Mid ∪ Buf, attained LB_{end} on all 75 boundary instances with $\beta \le 0.94$.

Proof of Corollary 22 If $F\alpha_0 \le a$, Theorem 21(b) gives $\Pr[\text{OPT} \ne \text{LB}_{\text{end}}] \le e^{-\Omega(t)}$. If $F\alpha_0 > a$ (and $F\alpha_0 < b$), Proposition 17(ii) gives $\Pr[\text{OPT} > X] \ge c\,t^{-1/2}$ for $X \in \{\text{LB}_{\text{end}}, \text{LB}_{\text{comb}}\}$, and Theorem 19(a) gives $\Pr[\text{OPT} \ne X] = O(t^{-1/2})$ for both. □

Proof of Theorem 19(e) This is Theorem 21(b): every valid $X \ge \text{LB}_{\text{end}}$ equals OPT on the good events. □

Remark 59 (one factory per product) The certificate of Theorem 21 keeps every product whole inside one factory, and LB_{end} and LB_{tie} are valid under that restriction. So Theorem 21, Corollary 22 and Proposition 17 also hold in the variant where products may not be split across factories.

References

Bai, D., Liu, T., Zhang, Y., Chu, F., Qin, H., Gao, L., Su, Y., Huang, M.: Scheduling a distributed permutation flowshop with uniform machines and release dates. IEEE Transactions on Automation Science and Engineering **22**, 215–227 (2025) https://doi.org/10.1109/TASE.2023.3349167

Bhattacharya, R.N., Rao, R.R.: Normal Approximation and Asymptotic Expansions. Wiley, New York (1976)

Ferone, D., Hatami, S., González-Neira, E.M., Juan, A.A., Festa, P.: A biased-randomized iterated local search for the distributed assembly permutation flow-shop problem. International Transactions in Operational Research **27**(3), 1368–1391 (2020) https://doi.org/10.1111/itor.12719

Gnedenko, B.V., Kolmogorov, A.N.: Limit Distributions for Sums of Independent Random Variables. Addison-Wesley, Cambridge, MA (1954). Translated by K. L. Chung

Gogos, C.: Solving the distributed permutation flow-shop scheduling problem using constrained programming. Applied Sciences **13**(23), 12562 (2023) https://doi.org/10.3390/app132312562

Grinberg, V.S., Sevast'yanov, S.V.: Value of the Steinitz constant. Functional Analysis and Its Applications **14**(2), 125–126 (1980) https://doi.org/10.1007/BF01086559

Hamzadayı, A., Arvas, M.A., Elmi, A.: Distributed assembly permutation flow shop problem; Single seekers society algorithm. Journal of Manufacturing Systems **61**, 613–631 (2021) https://doi.org/10.1016/j.jmsy.2021.10.012

Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association **58**(301), 13–30 (1963) https://doi.org/10.1080/01621459.1963.10500830

Hariri, A.M.A., Potts, C.N.: A branch and bound algorithm for the two-stage assembly scheduling problem. European Journal of Operational Research **103**(3), 547–556 (1997) https://doi.org/10.1016/S0377-2217(96)00312-8

Hatami, S., Ruiz, R., Andrés-Romano, C.: The distributed assembly permutation flowshop scheduling problem. International Journal of Production Research **51**(17), 5292–5308 (2013) https://doi.org/10.1080/00207543.2013.807955

Koulamas, C., Kyparisis, G.J.: Asymptotically optimal linear time algorithms for two-stage and three-stage flexible flow shops. Naval Research Logistics **47**(3), 259–268 (2000) https://doi.org/10.1002/(SICI)1520-6750(200004)47:3⟨259::AID-NAV5⟩3.0.CO;2-K

Koulamas, C., Kyparisis, G.J.: The three-stage assembly flowshop scheduling problem.

Computers & Operations Research **28**(7), 689–704 (2001) https://doi.org/10.1016/S0305-0548(00)00004-6

Koryakin, R.: On asymptotic optimality of permutation schedules in stochastic flow shops and assembly lines. In: Fleuren, H., Hertog, D., Kort, P. (eds.) Operations Research Proceedings 2004, pp. 200–206. Springer, Berlin, Heidelberg (2005). https://doi.org/10.1007/3-540-27679-3_25

Kaminsky, P., Simchi-Levi, D.: Probabilistic analysis and practical algorithms for the flow shop weighted completion time problem. Operations Research **46**(6), 872–882 (1998) https://doi.org/10.1287/opre.46.6.872

Kaminsky, P., Simchi-Levi, D.: The asymptotic optimality of the SPT rule for the flow shop mean completion time problem. Operations Research **49**(2), 293–304 (2001) https://doi.org/10.1287/opre.49.2.293.13536

Lee, C.-Y., Cheng, T.C.E., Lin, B.M.T.: Minimizing the makespan in the 3-machine assembly-type flowshop scheduling problem. Management Science **39**(5), 616–625 (1993) https://doi.org/10.1287/mnsc.39.5.616

Lin, J., Wang, Z.-J., Li, X.: A backtracking search hyper-heuristic for the distributed assembly flow-shop scheduling problem. Swarm and Evolutionary Computation **36**, 124–135 (2017) https://doi.org/10.1016/j.swevo.2017.04.007

Liu, B., Wang, K., Zhang, R.: Variable neighborhood based memetic algorithm for distributed assembly permutation flowshop. In: 2016 IEEE Congress on Evolutionary Computation (CEC), pp. 1682–1686. IEEE, Vancouver, Canada (2016). https://doi.org/10.1109/CEC.2016.7743990

Lin, J., Zhang, S.: An effective hybrid biogeography-based optimization algorithm for the distributed assembly permutation flow-shop scheduling problem. Computers & Industrial Engineering **97**, 128–136 (2016) https://doi.org/10.1016/j.cie.2016.05.005

Li, X., Zhang, Z.-Q., Hu, R., Qian, B., Li, K.: Hyper-heuristic three-dimensional estimation of distribution algorithm for distributed assembly permutation flowshop scheduling problem. In: Huang, D.-S., Premaratne, P., Jin, B., Qu, B., Jo, K.-H., Hussain, A. (eds.) Advanced Intelligent Computing Technology and Applications (ICIC 2023). Lecture Notes in Computer Science, pp. 386–396. Springer, Singapore (2023). https://doi.org/10.1007/978-981-99-4755-3_34

Li, X., Zhang, X., Yin, M., Wang, J.: A genetic algorithm for the distributed assembly permutation flowshop scheduling problem. In: 2015 IEEE Congress on Evolutionary Computation (CEC), pp. 3096–3101. IEEE, Sendai, Japan (2015). https://doi.org/10.1109/CEC.2015.7257275

Naderi, B., Ruiz, R.: The distributed permutation flowshop scheduling problem. Computers & Operations Research **37**(4), 754–768 (2010) https://doi.org/10.1016/j.cor.2009.06.019

Pan, Q.-K., Gao, L., Li, X.-Y., Framinan, J.M.: Effective constructive heuristics and meta-heuristics for the distributed assembly permutation flowshop scheduling problem. Applied Soft Computing **81**, 105492 (2019) https://doi.org/10.1016/j.asoc.2019.105492

Potts, C.N., Sevast'janov, S.V., Strusevich, V.A., Van Wassenhove, L.N., Zwaneveld, C.M.: The two-stage assembly scheduling problem: complexity and approximation. Operations Research **43**(2), 346–355 (1995) https://doi.org/10.1287/opre.43.2.346

Ramudhin, A., Bartholdi, J.J., Calvin, J.M., Vande Vate, J.H., Weiss, G.: A probabilistic analysis of two-machine flowshops. Operations Research **44**(6), 899–908 (1996) https://doi.org/10.1287/opre.44.6.899

Sevast'janov, S.V.: On some geometric methods in scheduling theory: a survey. Discrete Applied Mathematics **55**(1), 59–82 (1994) https://doi.org/10.1016/0166-218X(94)90036-1

Sevast'yanov, S.V.: Nonstrict vector summation in scheduling problems. Sibirskiĭ Zhurnal Issledovaniya Operatsiĭ **1**(2), 67–99 (1994). In Russian. English translation in A. D. Korshunov (ed.), *Discrete Analysis and Operations Research*, Kluwer, Dordrecht, 1996, pp. 257–287

Sevast'janov, S.V.: Vector summation in Banach space and polynomial algorithms for flow shops and open shops. Mathematics of Operations Research **20**(1), 90–103 (1995) https://doi.org/10.1287/moor.20.1.90

Shevtsova, I.: On the absolute constants in the Berry–Esseen type inequalities for identically distributed summands. arXiv:1111.6554 [math.PR] (2011)

Spitzer, F.: Principles of Random Walk, 2nd edn. Graduate Texts in Mathematics, vol. 34. Springer, New York (1976). https://doi.org/10.1007/978-1-4757-4229-9

Song, H.-B., Yang, Y.-H., Lin, J., Ye, J.-X.: An effective hyper heuristic-based memetic algorithm for the distributed assembly permutation flow-shop scheduling problem. Applied Soft Computing **135**, 110022 (2023) https://doi.org/10.1016/j.asoc.2023.110022

Taillard, E.: Benchmarks for basic scheduling problems. European Journal of Operational Research **64**(2), 278–285 (1993) https://doi.org/10.1016/0377-2217(93)90182-M

Wang, S.-Y., Wang, L.: An estimation of distribution algorithm-based memetic algorithm for the distributed assembly permutation flow-shop scheduling problem. IEEE

Transactions on Systems, Man, and Cybernetics: Systems **46**(1), 139–149 (2016) https://doi.org/10.1109/TSMC.2015.2416127

Xia, C.H., Shanthikumar, J.G., Glynn, P.W.: On the asymptotic optimality of the SPT rule for the flow shop average completion time problem. Operations Research **48**(4), 615–622 (2000) https://doi.org/10.1287/opre.48.4.615.12423

Yu, H., Gao, K.-Z., Ma, Z.-F., Pan, Y.-X.: Improved meta-heuristics with Q-learning for solving distributed assembly permutation flowshop scheduling problems. Swarm and Evolutionary Computation **80**, 101335 (2023) https://doi.org/10.1016/j.swevo.2023.101335

Ying, K.-C., Lin, S.-W.: Reinforcement learning iterated greedy algorithm for distributed assembly permutation flowshop scheduling problems. Journal of Ambient Intelligence and Humanized Computing **14**(8), 11123–11138 (2023) https://doi.org/10.1007/s12652-022-04392-w

Zhang, X., Li, X., Yin, M.: An enhanced genetic algorithm for the distributed assembly permutation flowshop scheduling problem. International Journal of Bio-Inspired Computation **15**(2), 113 (2020) https://doi.org/10.1504/IJBIC.2020.106443

Zhang, Z.-Q., Qian, B., Hu, R., Jin, H.-P., Wang, L.: A matrix-cube-based estimation of distribution algorithm for the distributed assembly permutation flow-shop scheduling problem. Swarm and Evolutionary Computation **60**, 100785 (2021) https://doi.org/10.1016/j.swevo.2020.100785