Automated bolting and head handling for multi-flange pressure-vessel closures: a simulation-based design evaluation with a delayed-coking case study

Leon Sandler

Independent Researcher, Northbrook, Illinois, USA

E-mail: sandler.leon@gmail.com · ORCID: 0009-0007-4584-808X

Abstract

Large bolted closures on process vessels are still assembled largely by hand, often in hazardous areas, using multipass tightening sequences whose duration and preload outcome are rarely quantified. An automated architecture for such closures is evaluated in this paper: orbital pneumatic torque runners on the flange, a bolt-management carousel, a counterbalanced head manipulator and an interlocking sequence controller, applied to a 92-bolt, three-flange delayed-coking drum. A Monte Carlo model built on the actual star-pattern tool travel shows that one torque head per flange needs a median of 8.0 h to unbolt and rebolt the drum and, within the investigated parameter bounds, cannot meet a 2.5 h target. Removing bottlenecks in turn gives 2.3 h with four synchronized heads and 1.8 h with a bolt-transfer channel per head, the preferred simulated configuration, which meets 2.5 h in 98% of samples; one fewer star pass would give 1.6 h but requires qualification. This ranking held across all 20,000 simulated samples under uniform, triangular and correlated input models. An exactly optimized tool path adds little, and the complete operation remains at approximately 3.4 h, governed by head handling and leak testing. Specifying the ±2% torque accuracy barely affects the preload scatter; under an assumed 3% angle-based preload-estimate uncertainty, the torque–angle retorque cuts the bolts outside ±10% of the target from 32% to 5%. Exhaustive verification revealed no safety violation but exposed deadlocks after interrupted sequences, which a resume rule removed. All the results are predictions; an experimental falsification plan is given.

Keywords: automated bolting; bolted flange joint; torque–angle control; Monte Carlo simulation; interlock verification; hazardous-area automation

1 Introduction

Bolted flange closures are the most frequently disturbed pressure boundaries in process plants, and their assembly is a manufacturing operation in all but name: a sequence of positioning, fastening and verification steps whose quality depends on the order in which fasteners are loaded and on how the applied torque is converted to preload. Guidance for pressure-boundary joints prescribes multipass tightening sequences for exactly this reason [1] because tightening one bolt changes the load in its neighbors through elastic interaction [2, 3]. When such joints are opened every few hours, the assembly procedure becomes a recurring production task, and the case for automating it is at once a quality, throughput and safety question.

Delayed-coking drums are an extreme example. Each drum is a large pressure vessel subject to severe thermal cycling and a range of damage mechanisms [4], and its heads are unbolted and rebolted on every coking cycle unless automatic unheading devices have been fitted. Bottom-head flange leaks are a common problem, and the industry trend is to install automatic bottom unheading devices, and the joints that remain require periodic retightening and careful assembly practices [5]. Thermal transients can also distort the drum flange until the bolt holes no longer align [6]. In the unit considered here, four drums on a 48 h fill cycle give an opening or closing operation every 12 h, and each operation involves 92 bolts across three flanged connections, handled manually at height in a potentially flammable, hot and dusty atmosphere.

An automated architecture for this operation was first proposed by the author in response to an industrial open-innovation challenge, as a pneumatic and automated bolt handling system (PABHS) with four subsystems: orbital torque runners, a bolt carousel, a head manipulator and a sequence controller. That proposal stated performance targets — bolting in 1.5–2.5 h against a reported 4–6 h manually, ±2% torque accuracy and no personnel in the hazard zone — without the models needed to establish whether the architecture could meet them. In this paper, those models are supplied. It treats the proposal as a design hypothesis and asks, for each target, what the architecture would have to achieve and under what conditions it would fail.

The contributions are as follows: (i) a cycle-time model for an automated multipass flange bolting that computes carriage travel from the actual tightening sequence and propagates step-time uncertainty by Monte Carlo simulation, with a rank-correlation sensitivity analysis; (ii) an admissibility analysis providing the per-bolt-pass time budget implied by a cycle-time target; (iii) a throughput-optimization study that removes the architecture's bottlenecks one at a time—torque heads, bolt-transfer channels, tightening passes, an exactly optimized tool path under a load-spreading constraint, and pipelined servicing—and reports bolting and complete-operation times separately, with a robustness check against the assumed input distributions, a design feasibility map and a servicing-time envelope for the complete operation; (iv) a preload analysis separating the contributions of tool accuracy, nut-factor scatter and torque–angle verification; and (v) exhaustive verification of the sequence and interlock logic, including mutation tests and a liveness check that exposes a deadlock in the logic as first specified; and (vi) a force budget for pneumatic head handling, a scored design FMEA and a hazard-exposure comparison. All the models are released as open code, and the paper ends with a validation plan whose acceptance criteria can falsify its predictions.

2 Related work

2.1 Flange assembly and tightening sequences

The legacy tightening pattern for pressure-boundary flanges loads bolts in a cross (star) order over several passes at increasing torque, followed by circular check passes until the nuts no longer rotate [1]. Alternative patterns that reduce assembly time or improve preload uniformity were added to that guidance with supporting laboratory evidence [7], and experimental comparisons of tightening sequences revealed that the choice of sequence matters for the uniformity achieved [8]. The underlying mechanism is elastic interaction: the bolt loads change as the neighboring bolts are tightened. Experimentally determined interaction coefficients allowed a one-pass, tension-controlled bolt-up to reach ±2% bolt-stress uniformity [2], and analytical models of the interaction have been used to optimize the loads applied during each pass [3, 9]. These studies optimize the sequence for joint quality; however, none consider the time taken by an automated tool to execute it, which depends on the path along which the tool must travel between bolts.

2.2 Torque control, preload scatter and torque–angle monitoring

Torque control infers preload from the applied torque through the nut factor, and the torque–tension relation is highly sensitive to thread and bearing friction [10]. In gasketed pipe flanges tightened by torque control, the scatter of axial bolt force is substantial and affects sealing performance [11], and the torque coefficient and bolt preload are closely correlated [12], and finite-element studies have shown that preload scatter due to elastic interaction is difficult to eliminate under torque control [13]. The accuracy of the assembly tool and the tightening strategy also affect the clamp–force scatter [14]. Monitoring the angle of the turn alongside the torque provides a second, friction-independent indication of the clamp force, and torque–angle formulations of the tightening process have been proposed and evaluated [15].

2.3 Robotic fastening and assembly data

Robotic fastening research has addressed bolt insertion and tightening with tactile in-hand localization [16], vision-based identification and pose estimation of nuts for a tightening robot [17], an adjustable spanner that engages bolts of different sizes automatically [18], and online vision inspection of screw-fastening quality [19]. Digital records of assembly data and process traceability have been developed for complex products using digital twin approaches [20]. These works concern small fasteners, structured factory settings or single tightening operations; large multipass flange closures in hazardous areas have received little attention.

2.4 Heavy-component handling

Gravity compensation of manipulators may use counterweights, springs or an auxiliary active force [21]. For pneumatic actuation, the available force is bounded by supply pressure and by standard cylinder sizes, which for the common ISO series reach a 320 mm bore [22]; therefore, whether a multitonne head can be handled pneumatically is a question of how much of the load a passive counterbalance carries.

2.5 Automation in hazardous process environments

Reviews of robotics in onshore and offshore oil and gas operations describe systems in which robots act while skilled operators retain cognitive decisions [23, 24]. Equipment in explosive atmospheres must follow the relevant protection concepts and certification requirements [25–27], and interlocks that protect against process hazards fall within functional safety practices for the process industry [28]. Model checking has been applied to verify PLC software [29], failure mode and effects analysis (FMEA) is the standard means of ranking design risks, with well-known limitations of the risk priority number [30], and discrete-event simulation is widely used to evaluate manufacturing system designs before they are built [31].

2.6 Research gap

In the literature, flange tightening for joint quality is optimized, robotic tightening for small fasteners is developed, and PLC logic is verified in general. What has not been done is to evaluate an integrated automated architecture for large multipass flange closures quantitatively: how long it takes as a function of its configuration, what preload accuracy its control strategy can actually deliver, whether its interlocking logic is both safe and free of deadlock, and whether its head-handling concept is physically feasible. This paper addresses those four questions for one industrial case.

3 Case study, requirements and architecture

3.1 Industrial case

The case is a four-drum delayed-coking unit. Each drum is approximately 28 m tall and 6 m in diameter. Three bolted connections are opened and closed on every cycle: the top head (32 bolts, M48–M64), the bottom head (48 bolts, M48–M64) and the feed inlet (12 bolts, M36–M48). Opening comprises unbolting, bolt storage, head removal and gasket removal; closing comprises face cleaning, gasket installation, head placement, bolt installation, multipass tightening, verification and a steam leak test. Hydraulic decoking, which occurs between opening and closing, is outside the scope. Flange drawings and head masses were not available; where the analysis needs them, it treats them as swept parameters rather than assumed point values.

3.2 Requirements

Requirement	Target	Source
Bolts per operation	92 (32 + 48 + 12)	Case data
Bolt sizes	M36–M64	Case data
Automated operations	Breakout, removal, storage, re-presentation, installation, multipass tightening	Design intent
Tightening procedure	Star passes at increasing torque + circular check passes	[1]
Record per bolt	Final torque and angle, pass/fail, time stamp	Design intent
Bolting time (92 bolts)	1.5–2.5 h (reported manual: 4–6 h)	Source proposal
Complete operation (open + close, excluding decoking)	Not stated in the source; evaluated here against the same 2.5 h	This paper
Personnel in hazard zone	None during routine operation	Design intent
Tool motion	Only with gas clearance confirmed	Design intent
Head motion	Only with no bolt engaged	Design intent
CLOSED status	Only after all bolts verified, leak test passed and operator sign-off	Design intent
Emergency stop	From any state, local and remote	Design intent

Table 1. Functional requirements. 'Source proposal' indicates the targets whose achievability is tested in this paper.

3.3 Architecture

The architecture (Fig. 1) has four subsystems. (A) An orbital bolt runner is a circular rail fixed around each flange carrying a carriage with one or more pneumatic torque multipliers, each with a socket, reaction arm, and torque and angle transducers; with several heads, they are spaced equally and act simultaneously. (B) A bolt-management carousel captures each bolt as it is removed, stores it in an indexed slot and presents it for reinstallation in the programmed order. (C) A counterbalanced articulated pneumatic manipulator engages, lifts, swings clear and locks each head. (D) A sequence and interlock controller enforces the tightening protocol, records torque and angle for every bolt, and exchanges start, stop and status signals with the plant control system. Pneumatic actuation is chosen because compressed air is available as a plant infrastructure and avoids electrical ignition sources at the tool; the hazardous-area certification of the integrated system is discussed in Section 6.

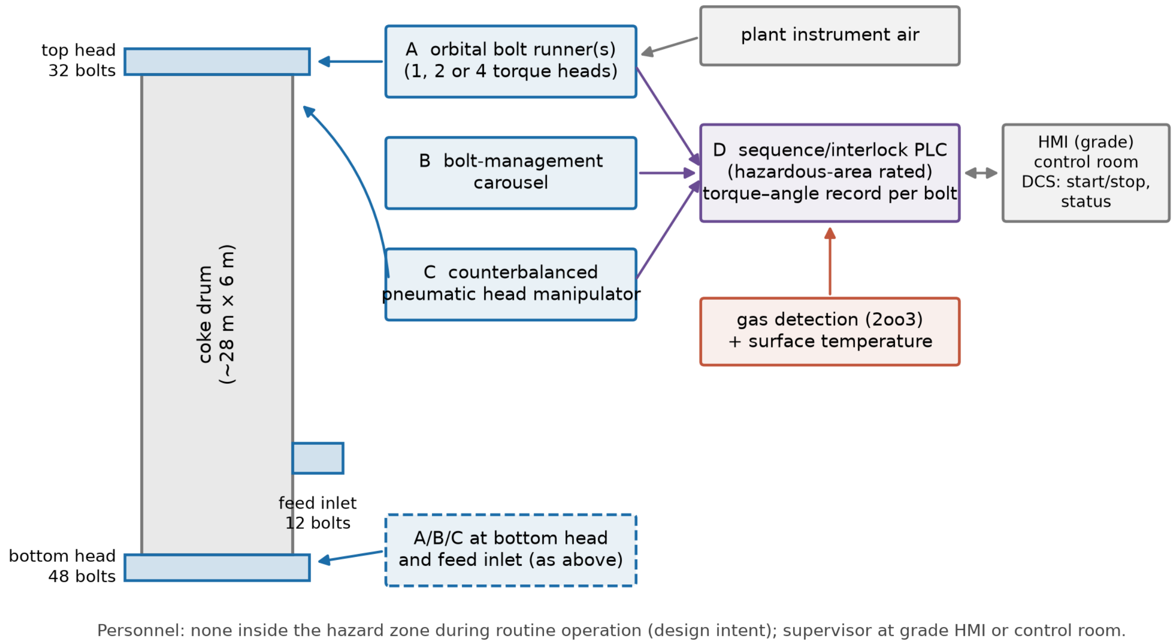

Fig. 1 Architecture of the automated bolting and head-handling system for a three-flange drum. Subsystems A–C are replicated at each flange (C at the two heads only); D is common.

4 Methods

4.1 Geometry

Bolt i of a flange with N bolts sits at angle

$$\theta_i = 2\pi i/N, \quad i = 0, 1, \ldots, N - 1 \qquad (1)$$

on a bolt circle of diameter D. Because D was not available, it is bracketed by a circumferential spacing rule s = k·d, where d is the nominal bolt diameter and k is between 2.5 and 5.0 (wrench clearance sets the lower end, gasket-load uniformity the upper):

$$D = N\,k\,d/\pi \qquad (2)$$

giving D from 1.83 to 4.89 m for the bottom head, 1.22 to 3.26 m for the top head and 0.34 to 0.92 m for the feed inlet. The spacing factor k is sampled in the Monte Carlo analysis, so its effect appears in the sensitivity results rather than being fixed.

4.2 Tightening sequence and rail travel

The legacy star pattern for N divisible by four is generated as groups of four bolts at 90° spacing, with the groups visited in bit-reversed order, which reproduces the cross-pattern family of the legacy sequence [1]. A circular pass visits the bolts in angular order. For a pass visiting stations $j_1, j_2, \ldots$ the carriage travels the shortest way around the rail between consecutive stations,

$$L = (\pi D/N)\, \Sigma_k\, \delta(j_k, j_{k+1}), \quad \delta(a, b) = \min((b - a) \bmod M,\ M - (b - a) \bmod M) \qquad (3)$$

where $M = N/r$. With r equally spaced heads acting together, bolts $j, j + N/r, \ldots$ are tightened at one station, so a pass visits M stations, and the carriage reaches the next station by moving whichever head is nearest, i.e., the shortest way on a ring of M positions (for $r = 1$, $M = N$). The 48-bolt star sequence and the resulting travel per pass are shown in Fig. 2: a single head covers 56 bolt-circle diameters per star pass on the bottom head, against 3.1 for a circular pass and 2.7 for four heads.

Much of that travel might be regarded as wasted motion, so the tool path is also optimized. Let the quartets $g = 0, \ldots, q - 1$ ($q = N/4$) be visited in order π; each quartet is worked as $g, g + 2q, g + q, g + 3q$ so that opposite bolts are loaded consecutively as the legacy procedure requires. The optimized order solves

$$\min \text{ over } \pi \text{ of } \Sigma_k\, c(\pi_k, \pi_{k+1}) \quad \text{subject to} \quad \delta\_q(\pi_k, \pi_{k+1}) \geq s\_min \text{ for all } k \qquad (4)$$

where c is the carriage travel from the last bolt of one quartet to the first of the next plus the travel within the next quartet and δ_q is the separation of two quartets on the ring of q positions. The load-spreading constraint requires every consecutive pair of quartets to be at least as far apart as the closest consecutive pair in the legacy sequence (s_min = 2 for the top and bottom heads), so the optimized path never concentrates load more locally than the legacy path does at any step. The problem is a constrained shortest Hamiltonian path on at most $q = 12$ nodes and is solved exactly by dynamic programming (Held–Karp). Whether a sequence that satisfies this geometric constraint also preserves preload uniformity under elastic interaction must be established experimentally [3, 9].

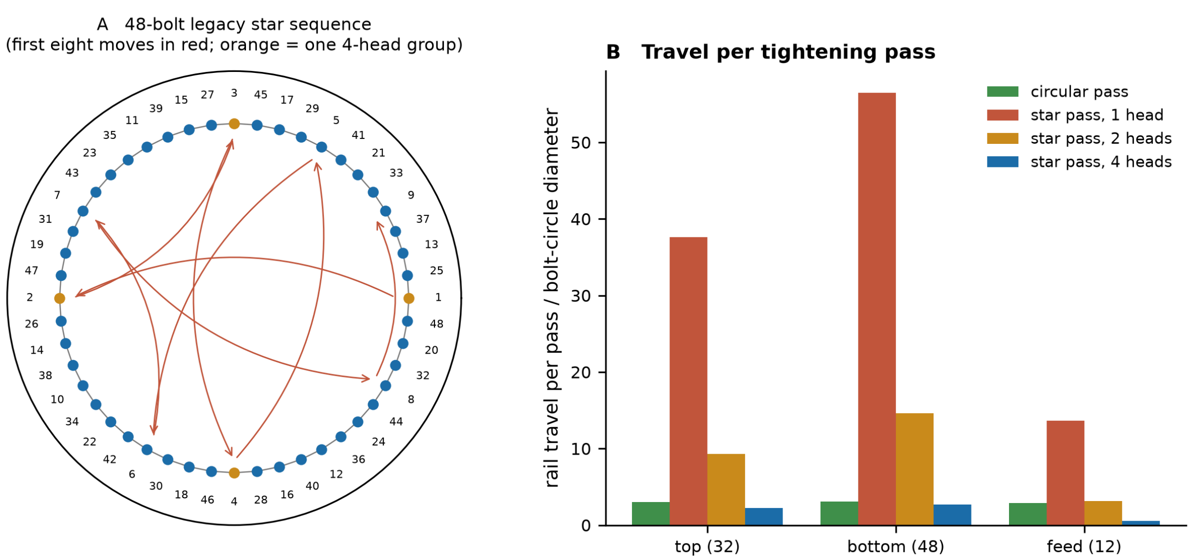

Fig. 2 (a) Legacy star sequence for the 48-bolt bottom head; numbers indicate the order; red arrows indicate the first eight carriage moves, and orange marks one group of four bolts tightened together by a four-head carriage. (b) Rail travel per pass, in terms of the bolt-circle diameter, for each flange and configuration.

4.3 Cycle-time model

One pass over a flange takes

$$t\_pass = (N/r)(t\_align + t\_engage + t\_tool + t\_release + t\_verify + t\_retry + t\_review) + L/v + m \cdot t\_move \quad (5)$$

where t_tool is the torque, breakout or run-on/run-off time of the pass, v is the carriage speed, m is the number of indexing moves and t_move a per-move acceleration and settling overhead. A station is complete only when all r heads are complete, so the expected retry and operator-review delays per station are

$$t\_retry = [1 - (1 - p\_retry)^r](t\_engage + t\_tool), \quad t\_review = [1 - (1 - p\_op)^r]\, t\_op \quad (6)$$

Run-on and run-off take the time to turn the nut through its engagement length, $t\_spin = 60 \cdot (1.5d/P)/n\_spin$, where P is the ISO coarse pitch [32] and n_spin is the spin speed in rev/min. Opening is one staged breakout pass in star order, a circular run-off pass and transfer of every bolt to the carousel; closing is a transfer from the carousel, a circular run-on pass, three star passes at increasing torque and one to three circular check passes, following the legacy procedure [1]. Bolt transfer between the flange and the carousel takes $N \cdot t\_transfer/c$, where c is the number of transfer channels working in parallel: c = 1 for a single carousel serving one bolt at a time and c = r when each torque head has its own feed point. The flanges are either worked in parallel, each by its own runner, or in series.

Two cycle times are distinguished. The bolting time is the time to unbolt and rebolt all 92 bolts, the quantity to which the source proposal's 1.5–2.5 h target referred:

$$T\_bolting = \max\_f T\_open,f + \max\_f T\_close,f \quad (7)$$

The complete operation adds atmosphere confirmation, head handling, gasket and face work and the steam leak test. In a phase-synchronized schedule, every flange finishes a phase before the next begins; in a pipelined schedule, each flange runs its own chain, and only the leak test waits for all of them:

$$T\_complete = t\_atmos + \max\_f(T\_open,f + t\_head,f) + \max\_f(t\_gasket + t\_head,f + T\_close,f) + t\_leak \quad (8)$$

where t_head,f = 0 for the feed inlet. Hydraulic decoking, which separates opening from closing, is excluded from both.

Parameter	Range (uniform)	Unit
Carriage speed	0.05 – 0.2	m/s

Indexing move overhead	1 – 4	s
Alignment to nut	4 – 15	s
Socket engagement	3 – 10	s
Torque pass, one bolt	10 – 40	s
Breakout, one bolt	10 – 35	s
Reaction release and retract	2 – 6	s
Torque–angle check	1 – 3	s
Run-on/run-off speed	20 – 60	rev/min
Carousel transfer, one bolt	10 – 35	s
P(automatic retry) per bolt-pass	0.005 – 0.03	–
P(operator review) per bolt-pass	0.0005 – 0.005	–
Operator review duration	300 – 1200	s
Circular check passes	1 – 3	–
Head handling (each way)	900 – 1800	s
Gasket and face work, per flange	900 – 2400	s
Atmosphere confirmation	120 – 600	s
Steam leak test	1200 – 2700	s

Table 2. Assumed step-time ranges. None is a measurement; the ranges are the inputs whose influence the sensitivity analysis ranks, and the validation plan (Section 7) specifies how each is to be measured.

4.4 Monte Carlo simulation and sensitivity

All inputs in Table 2, the spacing factor k and the position of each bolt size within its range are sampled independently 20000 times with a fixed seed (20260925). Six configurations are compared: one head with flanges in parallel, in series, and with circular passes only; two heads; four heads; and four heads with two instead of three star passes, as an example of a reduced-pass procedure of the kind permitted as an alternative to the legacy sequence [7]. Sensitivity is measured by the Spearman rank correlation between each sampled input and the total bolting time. A best case, with every step at the fast end of its range, bounds what each configuration could achieve. A second study then removes the bottlenecks of the four-head design one at a time on the same samples: (A) four heads and a single carousel; (B) one transfer channel per head; (C) two star passes; (D) the optimized tool path; (E) pipelined servicing. Three uncertainty models are compared to test whether the conclusions depend on the distributional assumptions (Section 5.6). Spearman rank correlation is used as a computationally transparent screening measure of sensitivity; variance-based (Sobol) analysis is deferred until empirical parameter distributions are available from the validation program, since apportioning variance among assumed distributions would add precision without adding information. Discrete-event simulations of this kind are standard for evaluating manufacturing designs before they are built [31].

4.5 Preload model

The torque needed for a target preload F in a bolt of diameter d is

$$T = K F d \qquad (9)$$

where K is the nut factor. The target is taken as half the specified minimum yield strength of SA-193 B7 bolting, 724 MPa [33], acting on the tensile stress area [34]

$$A_s = (\pi/4)(d - 0.9382P)^2 \qquad (10)$$

The achieved preload under torque control is $F = T(1 + \varepsilon_{tool})/(K d)$, with a tool error ε_{tool} uniform within the specified accuracy and K normally distributed with a coefficient of variation CV_K. To first order

$$CV_F \approx (CV_{tool}^2 + CV_K^2)^{1/2} \qquad (11)$$

Therefore, the benefit of a tighter torque specification depends on its size relative to that of CV_K. Torque–angle monitoring is modeled as an independent estimate of preload from the angle turned after the snug, with error CV_θ representing joint and gasket stiffness uncertainty; bolts whose estimate lies outside ±10% of the target are retorqued to the angle-corrected value. Each case uses 200,000 samples.

4.6 Head manipulator force budget

For a head of mass M_h lifted with safety factor SF = 1.5, the pneumatic force needed when a passive counterbalance carries a fraction c of the weight is

$$F_p = SF(1 - c)\, M_h\, g \qquad (12)$$

and the bore of a single cylinder at supply pressure p = 0.6 MPa and force efficiency η = 0.9 is $b = [4F_p/(\pi p \eta)]^{1/2}$. The minimum counterbalance fraction for the largest standard bore (320 mm) [22] follows by setting b = 320 mm. Because the head mass is not known, M_h is swept from 1 to 20 t.

4.7 Control logic and its verification

The controller is written as an explicit transition function over 13 modes (Fig. 3) and a state comprising the number of engaged bolts, the number verified, the leak test result and the phase to resume after an interruption. Nine Boolean inputs (gas clearance, emergency stop, torque pass, angle pass, head aligned, leak pass, start, confirm, and reset) are applied in every combination from every reachable state, by breadth-first enumeration, and seven safety invariants are checked on every transition: I1 tool motion begins only with gas clearance; I2 the head moves only with no bolt engaged; I3 the sequence advances by at most one bolt and only on a torque and angle pass; I4 CLOSED follows only a passed leak test with all bolts verified; I5 CLOSED requires operator sign-off; I6 an emergency stop forces the safe state from any state; and I7 loss of gas clearance halts tool motion. A liveness check then identifies reachable states from which no input sequence can reach a completed state (OPEN or CLOSED). To test that the invariants have discriminating power, each interlock is removed in turn, and the check is repeated; a check that no such mutation fails would test nothing. This is explicit-state model checking of the logic as specified [29]; it does not verify a PLC implementation.

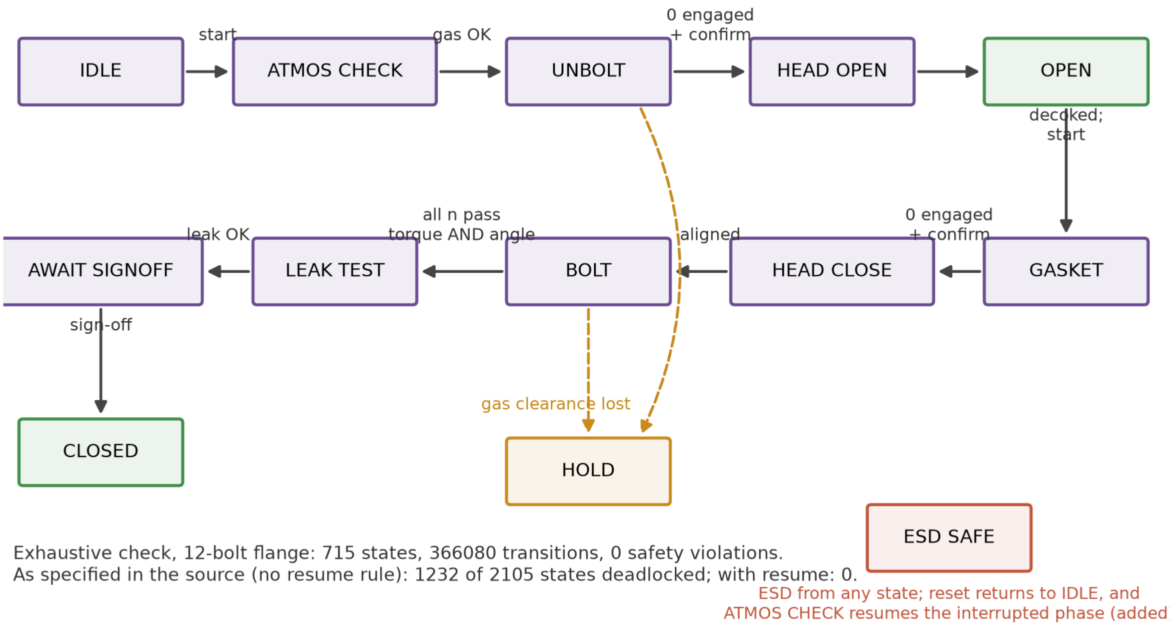

Fig. 3 Sequence and interlock state machine. Dashed transitions are interlock responses; the resume rule on reset (bottom right) was added after the liveness check found deadlocks in the logic as first specified.

4.8 FMEA, automation coverage and exposure

Twelve failure modes are scored for severity S, occurrence O (with the stated control) and detection D on 1–10 scales and are ranked by RPN = S·O·D [30]. The scores are the author's engineering judgment, not field data. Automation coverage is

$$A = \Sigma\, w_i a_i / \Sigma\, w_i \qquad (13)$$

with a_i = 1 for fully automatic, 0.5 for operator-initiated and 0 for manual operations, weighted by modeled task duration w_i. Hazard zone exposure is defined as person-hours per operation.

Implementation note: The models were written in Python with the assistance of a generative AI tool (see Declarations); the author specified, reviewed and checked every model and result.

5 Results

5.1 Torque requirement

Bolt	A_s (mm²)	Target preload (kN)	Torque, K = 0.16 (kN·m)	Torque, K = 0.20 (kN·m)
M36	817	296	1.7	2.1
M42	1121	406	2.7	3.4
M48	1473	533	4.1	5.1
M56	2030	735	6.6	8.2

M64	2676	969	9.9	12.4

Table 3. Target preload at half the specified minimum yield strength of SA-193 B7 and the corresponding torque for the two nut factors.

Table 3 gives the torques the runner must deliver: 9.9–12.4 kN·m for the M64 bolts of the heads and 1.7–2.1 kN·m for M36 on the feed inlet. These are within the range of torque multipliers, and they set the reaction load that the carriage and rail must resist at every station.

5.2 Cycle time

Configuration	Opening (h)	Closing (h)	Bolting total (h)	P(≤ 2.5 h)	Best case (h)
1 head, parallel flanges	2.12	5.84	7.96 [6.09–10.51]	0.00	2.96
1 head, serial flanges	3.87	10.60	14.49 [11.12–18.79]	0.00	5.42
1 head, circular passes only	1.70	4.54	6.26 [4.79–8.06]	0.00	2.42
2 heads, parallel flanges	1.14	2.84	3.99 [3.10–5.06]	0.00	1.48
4 heads, parallel flanges	0.72	1.56	2.28 [1.74–2.90]	0.73	0.84
4 heads, 2 star passes	0.72	1.34	2.06 [1.55–2.64]	0.90	0.76

Table 4. Modeled bolting time for all 92 bolts (median, with 5th–95th percentiles), probability of meeting the 2.5 h target, and the best case with every step at its fast limit. Opening and closing exclude head handling, gasket work and leak tests.

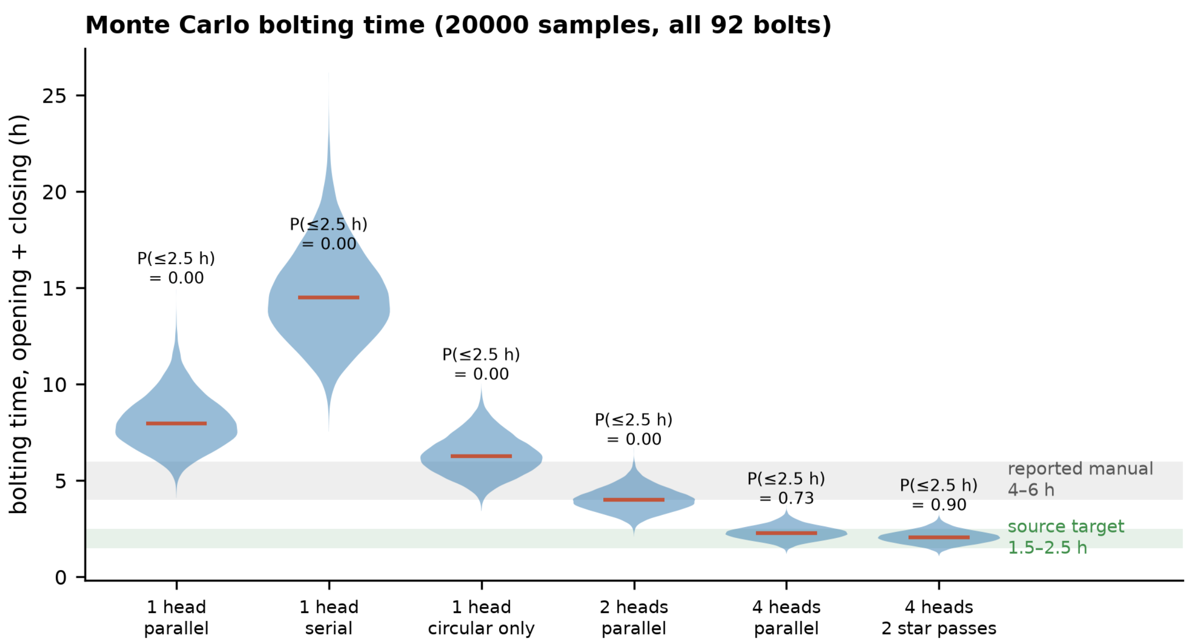

Fig. 4 Distribution of total bolting time by configuration, with the 1.5–2.5 h target of the source proposal and the reported 4–6 h manual duration.

With one head per flange and the flanges working in parallel, bolting takes 7.96 h (90% interval 6.09 –10.51 h), set by the bottom head. Working the flanges in series nearly doubles this to 14.5 h, and using only circular passes, which does not satisfy the legacy procedure, it still leaves 6.3 h. Most importantly, within the investigated parameter bounds, the single-head configuration cannot meet the 2.5 h target even in its best case of 2.96 h. A second head per flange halves the time to 4.0 h. Only four synchronized heads leave the median within the target at 2.28 h (90% interval 1.74–2.90 h), and even then, the target is met in 73% of samples; four heads with one fewer star pass increase this to 90%. Including head handling, gasket work and the leak test, the full opening and closing operation takes 9.8 h with one head and 4.1 h with four.

5.3 Admissibility of the cycle-time target

The legacy procedure applies each bolt approximately eight times per open–close operation (one breakout, one run-off, one run-on, three star passes and a median of two check passes), yielding 384 bolt passes on the bottom head and 736 in total. Meeting 2.5 h with flanges in parallel therefore allows 23 s per bolt pass on the bottom head and 12 s if the flanges share one runner in series. The modeled median is 75 s per bolt pass. The target is thus not a matter of faster tools: alignment, engagement, torque, release and verification of a large bolt cannot plausibly fit in approximately 20 s. It is achievable only by working several bolts at once, which is what a four-head carriage does, and which is also what a manual crew of four to six does.

5.4 Sensitivity

With one head (Fig. 5a), the time is driven by the carriage speed (ρ = -0.54), torque time per pass (0.43) and number of check passes (0.43): the single head spends much of the time traveling through the star pattern. With four heads (Fig. 5b), travel almost disappears, and the ranking changes: carousel transfer becomes the strongest driver (ρ = 0.55), followed by check passes (0.46) and torque time (0.42), and the probability of an operator review increases in importance because any one of four heads can stall a station. Carousel transfer accounts for 26% of the bottom-head bolting time in the four-head configuration. A design with four heads must therefore also increase carousel throughput, for example, with a feed point per head; otherwise, the carousel, not the torque tool, sets the pace.

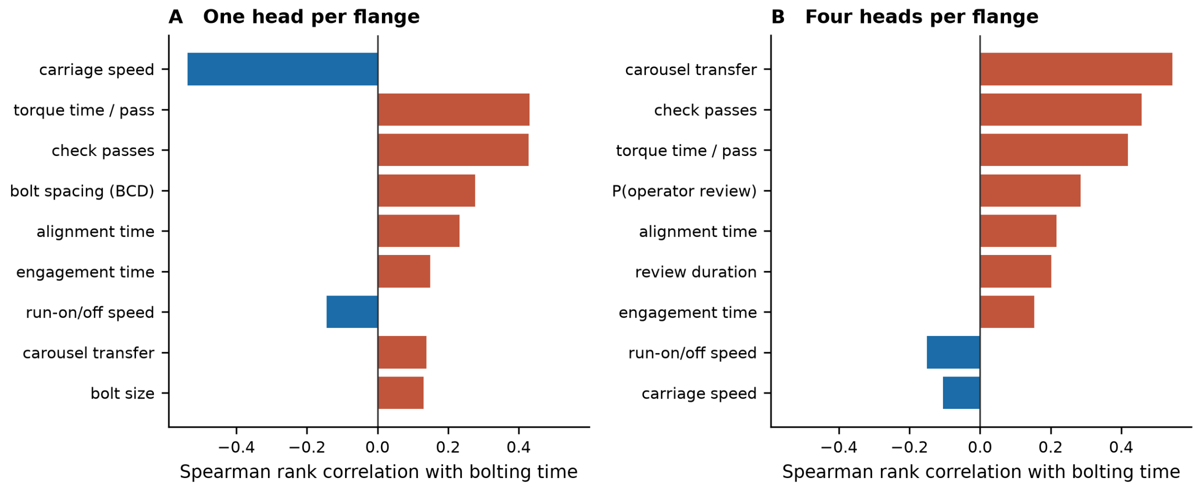

Fig. 5 Spearman rank correlation of each sampled input with total bolting time: (a) one head per flange and (b) four heads per flange. The nine strongest inputs are shown.

5.5 Throughput-optimized architecture

Configuration	Bolting (h)	P(bolting ≤ 2.5 h)	Complete operation (h)	P(complete ≤ 2.5 h)	P(complete ≤ 4 h)
One head (baseline)	7.96 [6.09–10.51]	0.00	9.81 [7.91–12.41]	0.00	0.00
A Four heads, single carousel	2.28 [1.74–2.90]	0.73	4.13 [3.47–4.86]	0.00	0.37
B + four transfer channels	1.82 [1.38–2.37]	0.98	3.68 [3.09–4.33]	0.00	0.80
C + two star passes	1.60 [1.21–2.10]	1.00	3.46 [2.91–4.08]	0.00	0.93
D + optimized tool path	1.57 [1.18–2.07]	1.00	3.43 [2.88–4.04]	0.00	0.94
E + pipelined servicing	1.57 [1.18–2.07]	1.00	3.43 [2.88–4.04]	0.00	0.94

Table 5. Removing the bottlenecks of the automated architecture one at a time (median and 5th–95th percentiles, same Monte Carlo samples throughout). Each row adds one change to the row above.

Table 5 and Fig. 6 remove the bottlenecks in turn. Giving each of the four heads its own bolt-transfer channel, for example, a carousel with four access stations, cuts bolting from 2.28 to 1.82 h and increases the probability of meeting 2.5 h from 0.73 to 0.98. This is a change to the machine alone, with the legacy tightening procedure unaltered, and it is the configuration the present analysis supports and recommends for experimental development: four heads and four transfer channels meet the bolting target without any change to how the joint is assembled. Removing one star pass brings bolting further to 1.60 h inside the target in all the simulated samples, but this is a change to the joint-assembly procedure rather than to the machine. It is an optional process improvement, admissible only if a reduced-pass procedure of the kind described as an alternative to the legacy sequence [7] is shown to give equivalent preload uniformity on the actual joint (Section 7); the architecture does not depend on it.

The exactly optimized tool path adds little, and it is reported for what it shows rather than as part of the preferred configuration. With four heads, it cuts the carriage travel per star pass on the bottom head from 2.68 to 1.51 bolt-circle diameters, but travel is by then a small part of the cycle and bolting falls only from 1.60 to 1.57 h. With one head, where travel matters, the optimum reduces travel per star pass only from 56.5 to 55.0 diameters and bolting from 7.96 to 7.91 h. The reason is structural: most of the single-head travel is incurred within each quartet because opposite bolts, half a turn apart, must be loaded consecutively, and no ordering of the quartets can remove that. The wasted motion of Fig. 2 is thus a consequence of serving a symmetric loading rule with one tool, and the remedy is more tools rather than a better path. Pipelining the servicing steps changes nothing (3.43 h in both D and E) because the bottom head, with the most bolts and a head to handle, is on the critical path in every phase, so the other flanges' chains finish within it.

The complete operation is different. Even in configuration E, it takes 3.43 h (90% interval 2.88–4.04 h) and does not meet 2.5 h for any of the 20000 samples. Of the median, 1.57 h is bolting, and 1.85 h is servicing: head handling in both directions (0.75 h), the leak test (0.54 h), gasket and face work (0.46 h) and atmosphere confirmation (0.10 h). The 2.5 h target is therefore achievable for bolting, which is what the source proposal stated, but not for the complete operation; once bolting is parallelised, head handling and the leak test become the next bottlenecks.

This observation can be turned into a performance requirement for the servicing subsystems. For a complete operation of at most 2.5 h, the servicing steps must satisfy

$$t\_atmos + 2t\_head + t\_gasket + t\_leak \leq 2.5\ h - T\_bolting \qquad (14)$$

With the preferred configuration B (median bolting of 1.82 h), the servicing budget is 0.68 h, or 41 min; with the optional reduced-pass configuration C (1.60 h), it is 0.90 h, or 54 min. The assumed servicing ranges are 1.85 h at the median and 1.12 h at best. With configuration C, a 2.5 h complete operation would need, for example, head handling of approximately 10 min each way, gasket and face work of approximately 10 min, a leak test of approximately 15 min and atmosphere confirmation under 5 min, approximately 50 min in all; with B, the same steps would have to fit in approximately 40 min. Either is roughly a halving of every servicing step. Equation (14) is the envelope that any faster head-handling or leak-test design would have to meet; no such design is evaluated here.

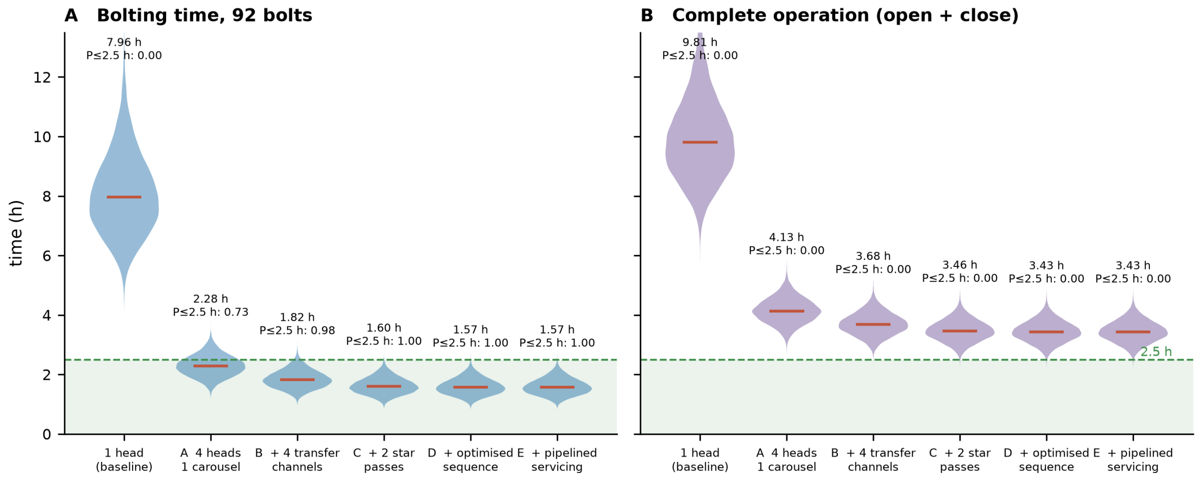

Fig. 6 (a) Bolting time and (b) complete-operation time when the bottlenecks are removed one at a time (Table 5). Shading indicates 2.5 h; labels indicate the median and the probability of meeting 2.5 h.

5.6 Robustness to the uncertainty model

Configuration	uniform, independent: median (h)/P(≤ 2.5 h)	triangular, independent: median (h)/P(≤ 2.5 h)	uniform, correlated (rho = 0.6): median (h)/P(≤ 2.5 h)
1 head	7.96/0.00	7.86/0.00	7.94/0.00
2 heads	3.99/0.00	3.96/0.00	3.98/0.01
A 4 heads	2.28/0.73	2.28/0.80	2.27/0.68
B + 4 channels	1.82/0.98	1.82/1.00	1.82/0.92
C + 2 star passes	1.60/1.00	1.60/1.00	1.60/0.98
Ranking 1 > 2 > A > B holds in	100.0% of samples	100.0% of samples	100.0% of samples

Table 6. Bolting time under three uncertainty models: the baseline independent uniform distributions, independent symmetric triangular distributions with the mode at the midpoint of each range, and uniform margins with a correlation of 0.6, through a Gaussian copula, among alignment, engagement and torque times, retry and review probabilities and bolt size.

The step-time ranges of Table 2 are assumptions, as are the uniform distributions and the independence with which they are sampled. It is plausible that a large, badly aligned bolt takes longer to align, engage and torque and is also likelier to need a retry or review so that those inputs move together. Table 6 repeats the analysis under two alternatives: triangular distributions centered on the nominal value and a pessimistic case in which those six inputs are positively correlated. The medians barely move because the ranges are unchanged; the tails do. In the correlated case, the probability that the preferred configuration B meets 2.5 h decreases from 0.98 to 0.92, and its 95th percentile increases from 2.37 to 2.60 h, whereas under triangular distributions, it increases to 0.997. What does not change is the architectural ranking: one head is slower than two, two slower than four, and four heads with a single carousel slower than four with four transfer channels across all 20,000 simulated samples under all three models. The paper's central claim is that ranking, not the particular value of 1.82 h.

The discrete configurations are generalized into a design feasibility map over the number of torque heads and transfer channels per flange for the legacy three-star-pass procedure and a two-pass alternative, as shown in Fig. 7. Heads are limited to 1, 2 and 4 because they must be equally spaced on all three flanges, and 32, 48 and 12 bolts are all divisible only by those. The feasible region, where bolting meets 2.5 h in at least 90% of samples, requires four heads under either procedure: with four heads, two transfer channels give 0.94 under the legacy procedure, and the two-pass procedure reaches 0.90, just short of the threshold, even with a single channel. With two heads, no combination exceeds 0.12. Adding channels without adding heads achieves almost nothing because transfer is a bottleneck only once torque application has been parallelized.

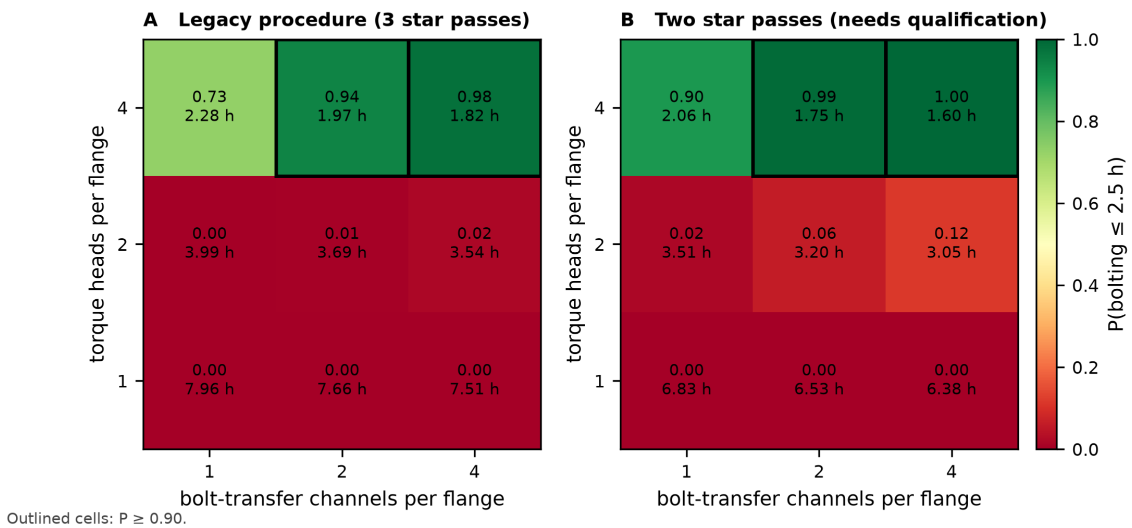

Fig. 7 Design feasibility map: probability that bolting meets 2.5 h (color and upper number) and median bolting time (lower number) against torque heads and bolt-transfer channels per flange for (a) the legacy three-star-pass procedure and (b) a two-pass alternative that would require qualification.

5.7 Preload

As shown in Fig. 8a, the ±2% torque accuracy for the system has almost no effect on preload scatter. With CV_K = 10%, the preload CV is 10.4% for a ±2% tool, 10.7% for ±5% and 11.8% for ±10%; 32% of the bolts fall outside ±10% of the target even with the ±2% tool. Preload scatter is set by the nut factor, as the torque–tension literature leads one to expect [10–12]. The angle channel is what changes the outcome (Fig. 8b): with an angle-based preload estimate of 3% uncertainty and a retorque of bolts outside ±10%, the fraction outside decreases from 32% to 5% (CV_K = 10%) and from 50% to 5% (CV_K = 15%), respectively, at the cost of retorquing 34% and 51% of the bolts, respectively. With a 5% angle-estimate uncertainty, the residual is approximately 10%; with 8%, the benefit largely disappears. The value of torque–angle monitoring therefore depends on how well joint stiffness is known, which the validation plan measures.

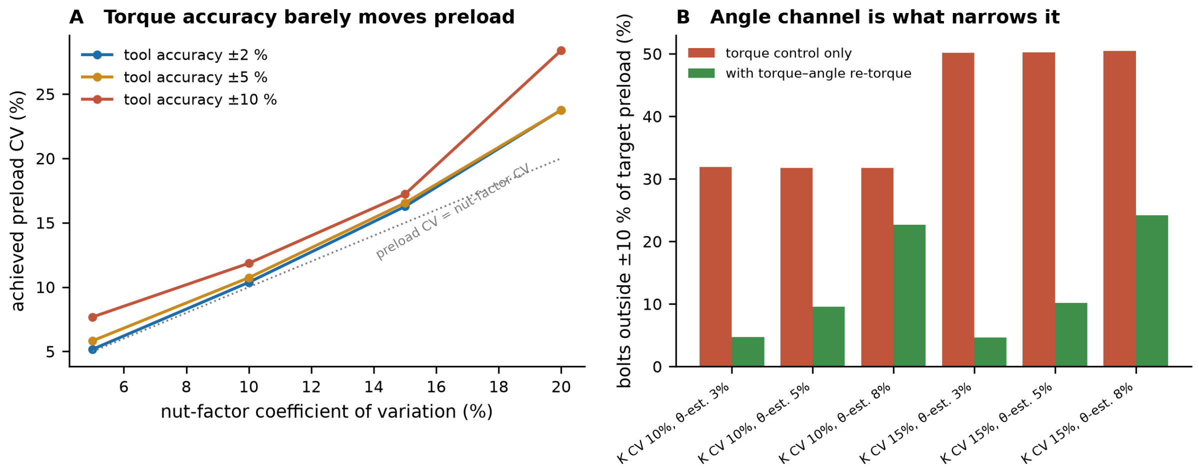

Fig. 8 (a) Achieved preload coefficient of variation against nut-factor variation for three tool accuracies. (b) Fractions of bolts outside ±10% of the target under torque control alone and with torque–angle retorque for two nut-factor scatters and three angle-estimate uncertainties.

5.8 Head handling

Without counterbalance, lifting a 5 t head at SF 1.5 from 6 bar air requires a 417 mm cylinder and a 10 t head 589 mm, beyond the largest standard bore. The minimum counterbalance fraction that keeps a single cylinder within 320 mm is 41% for 5 t, 70% for 10 t, 80% for 15 t and 85% for 20 t (Fig. 9). A pneumatic manipulator is therefore feasible only as a counterbalanced device in which air supplies balance correction and motion, not lift, consistent with passive gravity compensation [21]. Compressibility also makes a pneumatic axis compliant; thus, the final seating of the head must rely on mechanical guides and alignment pins, and holding must rely on a spring-applied mechanical lock rather than on air pressure.

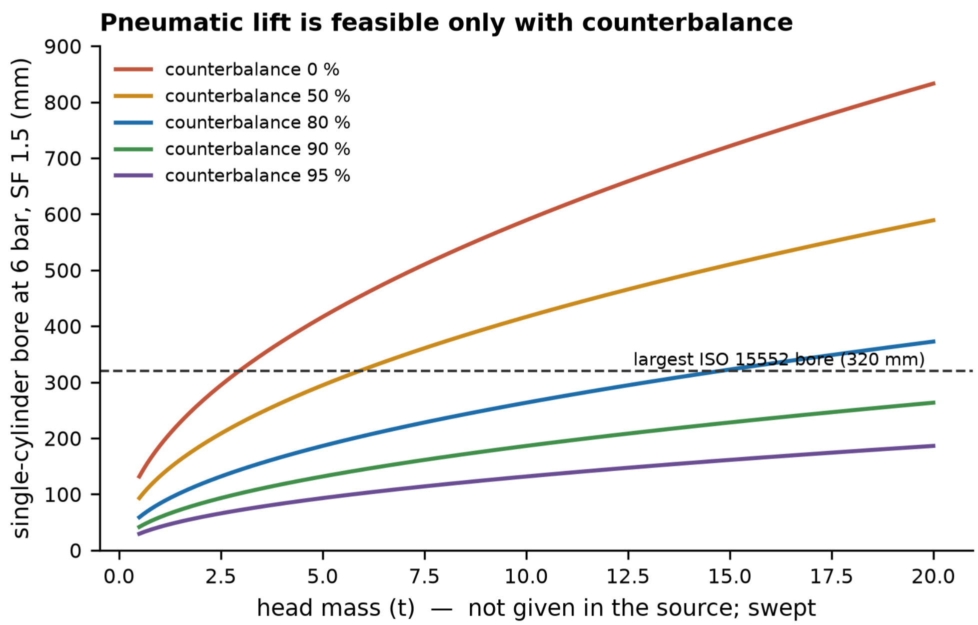

Fig. 9 Single-cylinder bore needed to lift a head against its unbalanced weight as a function of head mass and counterbalance fraction. The head mass was not available and was swept.

5.9 Control logic

With the full interlock set and the resume rule, exhaustive exploration reveals no safety violation: 715 reachable states and 366080 checked transitions for a 12-bolt flange (347 and 177664 for 4 bolts), with all 13 modes reachable and no deadlocked state. Every mutation is detected (Table 7); thus, each invariant is capable of failure. The liveness check, however, revealed a defect in the logic as first specified. Without an explicit rule for resuming an interrupted phase, 1232 of 2105 reachable states for a 12-bolt flange (144 of 385 for 4 bolts) can never reach OPEN or CLOSED: an emergency stop part-way through unbolting or tightening leaves a partially bolted joint from which the prose interlocks provide no way forward, so the operator would have to leave the automated system. Adding the rule that atmosphere confirmation routes to the interrupted phase — continue unbolting, resume tightening at the next unverified bolt, and repeat the leak test or await sign-off — removes every deadlock while preserving all seven invariants.

Interlock removed	**Invariant violated**	**Shortest counterexample (transitions)**
Emergency stop	I6 ESD forces the safe state from any state	1
Gas clearance before tool motion	I1 unbolting starts only with gas clearance	2
Hold on loss of gas clearance	I7 loss of gas clearance holds all tool motion	3
Leak test before CLOSED	I4 CLOSED only after leak test pass	18
No head motion with bolts engaged	I2 head moves only with no bolt engaged	3
Sequence advances only on torque AND angle pass	I3 sequence advances only on torque AND angle pass	13
Operator sign-off before CLOSED	I5 CLOSED only with operator sign-off	18

Table 7. Mutation test of the interlock logic (4-bolt flange). Each interlock removed in isolation is captured by the invariant it protects.

5.10 Failure modes

The highest-ranked failure mode (Fig. 10, Table 8) is an out-of-window preload that goes undetected (RPN 160), which follows directly from Section 5.7: torque control alone cannot see a nut-factor outlier, and detection depends on the angle channel. The second is the failure of gas detection to register residual hydrocarbons (RPN 150, severity 10), which is why the design specifies voted redundant detectors and a precycle functional test. The failure modes with the highest severity but lower RPN—head release with a bolt engaged and loss of air under load—are those the interlocks and the mechanical lock address by design; their low RPN reflects the control, not the absence of the hazard.

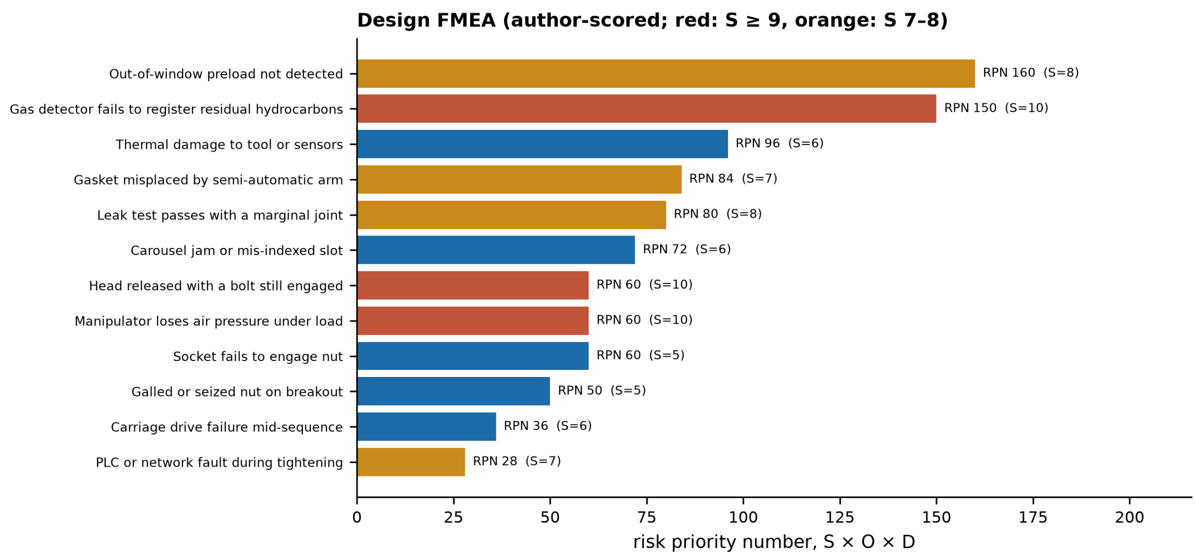

Fig. 10 Design FMEA ranked by risk priority number; color indicates severity. The scores are the author's judgement and are listed with their controls in Table 8.

Failure mode	Effect	S	O	D	RPN	Control
Out-of-window preload not detected	Gasket leak at restart	8	5	4	160	Torque-angle window per bolt; angle-corrected retorque
Gas detector fails to register residual hydrocarbons	Bolting starts in flammable atmosphere	10	3	5	150	Redundant detectors (2oo3 voting), bump test before each cycle
Thermal damage to tool or sensors	Tool failure, sensor drift	6	4	4	96	High-temperature rating; surface temperature interlock
Gasket misplaced by semiautomatic arm	Leak at steam test	7	4	3	84	Camera confirmation; steam test gate before CLOSED
Leak test passes with a marginal joint	Leak in service	8	2	5	80	Defined hold time and decay criterion; torque-angle record

						review
Carousel jam or mis-indexed slot	Wrong bolt returned or stall	6	4	3	72	Slot presence sensing; bolt ID; stop on mismatch
Head released with a bolt still engaged	Head/flange damage, dropped load	10	2	3	60	Carousel count must equal n AND per-station sensor; interlock
Manipulator loses air pressure under load	Uncontrolled head motion	10	3	2	60	Counterbalance carries static load; spring-applied mechanical lock
Socket fails to engage nut	Bolt skipped, sequence stall	5	6	2	60	Laser/contact alignment; engagement proof before torque
Galled or seized nut on breakout	Breakout torque exceeded, stall	5	5	2	50	Torque limit with alarm; operator review; spare stud set
Carriage drive failure mid-sequence	Sequence interrupted mid-pass	6	3	2	36	Resume-from-bolt state retention; manual procedure retained
PLC or network fault during tightening	Partially tightened joint	7	2	2	28	Safe-state on fault; state retention; ESD independent of PLC

Table 8. Design FMEA (1–10 scales; D = 10 is undetectable).

5.11 Automation coverage and exposure

Operation	**Modeled duration (h)**	**Automation level a_i**
Atmosphere check	0.10	1.0
Bolt loosening and removal	2.12	1.0
Head removal and replacement	0.75	0.5
Gasket removal, cleaning, installation	0.46	0.5
Bolt installation and tightening	5.84	1.0
Leak test	0.54	1.0

Table 9. Operations, modeled durations (one-head configuration, medians) and automation levels used in the coverage metric.

Weighted by the modeled duration, the automation coverage is A = 0.94 (Table 9). The reported manual procedure places four to six workers in the hazard zone for four to six hours, 16–36 person-hours per operation; with the automated system, the routine value is zero and increases to at most 0.46 person-hours if gasket cassette loading has to be performed at the flange rather than from outside the zone. The exposure reduction is the least sensitive benefit to the cycle-time modeling assumptions: it does not depend on the step times that dominate the cycle-time uncertainty, although it remains to be demonstrated in an integrated pilot.

6 Discussion

6.1 What the targets require

The analysis supports two of the three original targets and conditions the third. The architecture is designed to meet the exposure target, although achievement remains to be demonstrated in the integrated pilot. The torque-accuracy target is achievable but, on its own, does not deliver the preload quality it appears to promise; what delivers it is the angle channel. The cycle-time target is not met by the architecture as first described, with one runner per flange, anywhere within the investigated parameter bounds. The 8.0 h single-head result is not a failure of the concept but a finding about it: a single tool is the wrong architecture for a 92-bolt, multipass closure. The architecture the analysis supports has four synchronized heads per flange and one bolt-transfer channel per head, with the legacy tightening procedure unchanged, which results in bolting from 8.0 h to 1.8 h; a qualified two-star-pass procedure is an optional further step, not a requirement. Parallelization, not faster tools, is what matters: the admissibility analysis shows that no plausible per-bolt time allows one head to meet the target, and once torque application is parallelised every stage around it — bolt handling above all — must be parallelised too. This is a substantially larger and more expensive machine than the original proposal implied, and it changes the comparison with manual work: a single automated head is slower than a crew of four to six, and the automated system matches or improves on manual duration only when it too works several bolts at once. The target also has to be read as the source stated it, for bolting: the complete operation remains approximately 3.4 h, governed by head handling and the leak test, and shortening it is a separate design problem for those subsystems.

6.2 Sequence choice as a throughput lever

For a single head, the star pattern costs far more travel than a circular pass does (Fig. 2b), and circular-only tightening shortens the cycle, but the legacy procedure uses the star pattern because it limits the load redistribution caused by elastic interaction [2, 3]. Section 5.5 shows that optimizing the order of the tool path under a load-spreading constraint recovers only a few percent of the travel because the opposite-bolt rule fixes most of it. The penalty is removed for multihead carriages while the cross-pattern loading is maintained because each station tightens the bolts at 90° spacing together. The remaining sequence lever is the number of passes: with four heads and four transfer channels, removing one star pass shortens the closing by 17%. Alternative patterns with fewer passes [7] and sequences optimized using measured interaction coefficients [9] are the way to make that saving safely, and the manufacturing question becomes which reduced-pass procedure yields acceptable preload uniformity on this joint.

6.3 Flange condition and engagement

The engagement and alignment steps assume that nuts are where the rail expects them. Coke-drum bottom flanges can become distorted by thermal transients until the bolt holes no longer align [6], and the fouling of nuts by coke is likely. Both would raise the retry and operator-review probabilities, which the four-head sensitivity ranks among the stronger drivers. Vision-based nut localization and pose estimation

[17] and self-adjusting engagement tools [18] address part of this; a dimensional survey of the actual flanges before rail fabrication addresses the rest.

6.4 Relation to unheading devices

When automatic bottom unheading devices are installed, the bottom head is no longer unbolted on every cycle [5], and the per-cycle bolting problem addressed here does not arise at that flange; the drum-to-device joint is then opened only at maintenance intervals, when flange distortion can make reassembly difficult [6]. The architecture evaluated here is relevant to units that retain bolted heads, and its bolting subsystem could serve those maintenance reassemblies. Whether automated bolting or unheading retrofitting is preferable for a given unit is an economic and engineering decision outside the scope of this paper.

6.5 Hazardous-area compliance and functional safety

The source proposal described hazardous-area certification of the integrated system as achievable by selecting certified components. Component certification is necessary but not sufficient: the integrated assembly, including carriage drives, transducers, the carousel and the controller, must satisfy the applicable protection concepts as a whole [25–27]. The gas clearance and head motion interlocks are safety functions, and their integrity requirements belong to a functional safety assessment [28]. The exhaustive check in Section 5.9 shows that the logic is consistent and deadlock-free as specified; it does not establish the hardware fault tolerance or safety integrity level of its implementation.

6.6 Limitations

Every performance figure in this paper is a model prediction. The step-time ranges in Table 2 are assumptions, which are bounded but not measured; the bolt-circle diameters are bracketed by a spacing rule; and the head masses are swept. The distributions and their independence are also assumptions; Section 5.6 shows that they move the tails of the cycle-time distribution but not the ranking of architectures. Elastic interaction is represented only through the number of passes and is not modeled directly. The preload model treats the nut factor and angle estimation errors as independent and normally distributed. The FMEA scores are the author's judgement. The control-logic verification abstracts the controller to a finite-state model and checks it for flanges of 4, 8 and 12 bolts; its conclusions concern the logic, not an implementation. The manual baseline of 4–6 h is as reported and was not measured. Costs are not analyzed beyond the preliminary estimates noted in Section 8.

7 Experimental validation and falsification plan

The predictions above can be falsified by a staged test program. Table 10 lists the measurements, the rig on which each is made and the acceptance criterion that would confirm or refute the corresponding result.

Quantity	**Rig**	**Acceptance criterion**
Step times: align, engage, torque, release, verify, transfer	Full-size single-station rig with a representative M56–M64 stud and flange segment	Measured medians within the Table 2 ranges; if not, rerun the model with measured values
Carriage speed and indexing overhead	Rail segment, loaded carriage	As above
Per-bolt-pass time, four-head station	Four-head carriage on a full flange mock-up	≤ 94 s median to meet 2.5 h (Section 5.3)
Bolt transfer with four channels	Carousel with four access stations on the mock-up	Median transfer time within the Table 2 range per channel, with no interference between channels (configuration B)
Two-star-pass procedure and optimized tool path	Full-flange mock-up with instrumented studs and the service gasket	Final preload scatter and gasket-stress uniformity no worse than the legacy three-pass procedure; otherwise configuration B, not C–E, is the design
Nut-factor scatter CV_K	Instrumented studs (ultrasonic elongation) on the mock-up	Measured CV_K; preload prediction of Fig. 8a to be confirmed within ±2 points
Angle-estimate uncertainty CV_θ	As above, with gasket installed	CV_θ ≤ 5% for the torque–angle benefit of Fig. 8b to hold
Engagement success on distorted flange	Mock-up with imposed ovality and cupping	Retry probability within Table 2 range
Interlock logic	Controller hardware-in-the-loop	All seven invariants and the resume rule demonstrated; no deadlock after ESD in any phase
Head handling	Counterbalanced manipulator with dummy head at maximum mass	Lift, swing, lock and reseat without personnel in load path; lock holds on air loss

Table 10. Validation measurements and acceptance criteria. A pilot on one bottom head is proposed only after the single-station and mock-up tests have passed.

8 Industrial feasibility

Each subsystem uses established technology—pneumatic torque multipliers, rails and carriages, fastener magazines, counterbalanced manipulators and programmable controllers—so the development risk lies in integration and in the hazardous-area environment rather than in any single component. The preliminary cost estimates of the source proposal are approximately US$0.14–0.20 million for a single-flange pilot and US$0.74–1.1 million for all three flanges on four drums, with one runner assumed per flange. The results of Section 5 imply a four-head carriage and four bolt-transfer channels on each flange if the bolting target is to be met, so those estimates are likely to be low and should be revised after the single-station tests; shortening the complete operation further would require faster head handling and leak testing, which were not costed. Rather than substituting another point estimate, the revised cost per flange is structured as follows:

$$C = C\_fixed + r \cdot C\_head + c \cdot C\_channel + C\_controls + C\_certification \quad (15)$$

with r = 4 heads and c = 4 transfer channels in the preferred configuration; the coefficients—rail and structure, each torque head, each transfer channel, the controls and the hazardous-area certification of the integrated assembly—require vendor quotations and are not estimated here. A staged programme — single-station rig, full-flange mock-up, then a pilot on one bottom head — allows the project to stop cheaply if the measured step times, nut-factor scatter or engagement success fall outside the ranges on which the results depend.

9 Conclusions

An automated architecture for opening and closing a 92-bolt, three-flange pressure vessel was evaluated by simulation against the targets originally claimed for it. With one torque head per flange, bolting takes a median of 8.0 h and cannot meet a 2.5 h target within the investigated parameter bounds: a single tool is the wrong architecture for this closure. Parallelizing in stages brings bolting to 2.3 h with four synchronized heads and 1.8 h with a bolt-transfer channel per head. That machine-only configuration, with the legacy tightening procedure unchanged, is the configuration the present analysis supports and recommends for experimental development: it meets the bolting target in 98% of samples, and the ranking of architectures that leads to it being held across all 20,000 simulated samples under uniform, triangular and correlated input models. One fewer star pass would bring bolting to 1.6 h but is an optional process change that requires the qualification of preload uniformity. An exactly optimized tool path and pipelined servicing add little, the first because the opposite-bolt rule fixes most of the travel and the second because the bottom head is on the critical path in every phase. The complete operation remains at approximately 3.4 h, governed by head handling and the leak test, so the 2.5 h target holds for bolting but not for the whole operation; Eq. (14) gives the servicing-time envelope that a faster design would have to meet. Specifying ±2% torque accuracy does little for preload, for which nut-factor scatter is dominant, whereas torque–angle verification with retorque substantially narrows it to an extent that depends on how well joint stiffness is known. Explicit-state verification revealed no violation of the seven safety invariants but revealed widespread deadlock in the initial interlock specification after interrupted sequences; adding an explicit resume rule eliminated every deadlock while preserving all seven invariants. Pneumatic head handling requires a counterbalance to carry most of the load. The reduction in hazard zone exposure, from 16–36 person-hours to essentially none by design, is the benefit least sensitive to the cycle-time modeling assumptions, although it also awaits demonstration. These are predictions, and the validation plan specifies the measurements that would confirm or overturn them.

Declarations

Funding. This work received no external funding.

Competing interests. The author declares that there are no competing interests.

Ethics approval. Not applicable. No human participants or animals were involved.

Consent to participate/for publication. Not applicable.

Prior dissemination. An earlier, nonquantitative description of the architecture was submitted by the author in response to an industrial open-innovation challenge. That submission was not peer reviewed or published and contained none of the models or results reported here.

Data availability. No experimental data were generated. All the model inputs are stated in the paper.

Code availability. The models, the verification script and the figure generators are openly available at https://github.com/sandlerleon/pabhs-bolting-model and are archived at https://doi.org/10.5281/zenodo.22969548. Running pabhs_model.py, plc_verify.py and make_figures.py reproduces every number and figure.

Author contributions. L.S. conceived the study, developed the models, performed the analysis and wrote the manuscript.

Use of generative AI. During the preparation of this work, the author used Claude (Anthropic) to assist with the literature search, implementation of the models and drafting of the text. The author reviewed and edited all the content, verified every cited reference against its source, and takes full responsibility for the content of the article.

References

1. ASME (2022) PCC-1-2022 Guidelines for pressure boundary bolted flange joint assembly. American Society of Mechanical Engineers, New York
2. Bibel GD, Ezell RM (1992) An Improved Flange Bolt-Up Procedure Using Experimentally Determined Elastic Interaction Coefficients. Journal of Pressure Vessel Technology 114(4):439-443. https://doi.org/10.1115/1.2929252
3. Zhu L, Bouzid AH, Hong J, Zhang Z (2018) Elastic Interaction in Bolted Flange Joints: An Analytical Model to Predict and Optimize Bolt Load. Journal of Pressure Vessel Technology 140(4):041202. https://doi.org/10.1115/1.4040421
4. Ma X, Xie Y, Ge Z (2023) Identification of Typical Damage Mechanism of Coke Drum in Delayed Coking Unit. In: Proceedings of the ASME 2023 Pressure Vessels & Piping Conference, Volume 1: Codes & Standards. American Society of Mechanical Engineers. https://doi.org/10.1115/pvp2023-107648
5. Berry A, Brown W, Seijas A, Cook S (2019) Coke Drum Bottom Head Flange Design Optimisation. In: Proceedings of the ASME 2019 Pressure Vessels & Piping Conference, Volume 3: Design and Analysis. American Society of Mechanical Engineers. https://doi.org/10.1115/pvp2019-93136
6. Horstmeyer M, Corcoran A, Seijas A, Berry A, Brown W (2025) Coke Drum Bottom Head Flange Distortion, Root Cause and Solutions. In: Proceedings of the ASME 2025 Pressure Vessels & Piping Conference, Volume 3: Design & Analysis. American Society of Mechanical Engineers. https://doi.org/10.1115/pvp2025-153394
7. Brown W, Waterland J, Lay D (2010) Background on the New ASME PCC-1:2010 Appendix F "Alternatives to Legacy Tightening Sequence/Pattern". In: ASME 2010 Pressure Vessels and Piping Conference: Volume 3. ASMEDC. https://doi.org/10.1115/pvp2010-25772
8. Kumakura S, Saito K (2003) Tightening Sequence for Bolted Flange Joint Assembly. In: Proceedings of the ASME 2003 Pressure Vessels and Piping Conference, Analysis of Bolted Joints. ASMEDC. https://doi.org/10.1115/pvp2003-1867
9. Zhu L, Bouzid AH, Hong J (2019) A Novel Methodology to Optimize the Tightening Sequence in Bolted Flange Joints. In: Proceedings of the ASME 2019 Pressure Vessels & Piping Conference, Volume 2: Computer Technology and Bolted Joints. American Society of Mechanical Engineers. https://doi.org/10.1115/pvp2019-93062

10. Nassar SA, Ganeshmurthy S, Ranganathan RM, Barber GC (2007) Effect of Tightening Speed on the Torque-Tension and Wear Pattern in Bolted Connections. Journal of Pressure Vessel Technology 129(3):426-440. https://doi.org/10.1115/1.2749290
11. Sawa T, Matsumoto M, Nagata S (2003) Effects of Scatter in Bolt Preload of Pipe Flange Connections With Gaskets on Sealing Performance. In: Proceedings of the ASME 2003 Pressure Vessels and Piping Conference, Analysis of Bolted Joints. ASMEDC. https://doi.org/10.1115/pvp2003-1875
12. Omiya Y, Fujii M (2016) Correlation on Scatter Between Torque Coefficient and Bolt Preload in Pipe Flange Connection. In: Proceedings of the ASME 2016 Pressure Vessels and Piping Conference, Volume 2: Computer Technology and Bolted Joints. American Society of Mechanical Engineers. https://doi.org/10.1115/pvp2016-63720
13. Abid M, Khan YM, Nash DH (2013) Assembly of Gasketed Bolted Flange Joints Using Torque Control of Preload Method: FEA Approach. In: Proceedings of the ASME 2013 Pressure Vessels and Piping Conference, Volume 3: Design and Analysis. American Society of Mechanical Engineers. https://doi.org/10.1115/pvp2013-97863
14. Persson EV, Kumar M, Friberg C, Dressler N (2021) Clamp Force Accuracy in Threaded Fastener Joints Using Different Torque Control Tightening Strategies. SAE Technical Paper 2021-01-5073. SAE International. https://doi.org/10.4271/2021-01-5073
15. Nassar SA, Yang X (2008) Torque-Angle Formulation of Threaded Fastener Tightening. Journal of Mechanical Design 130(2):024501. https://doi.org/10.1115/1.2821388
16. Nozu K, Shimonomura K (2018) Robotic bolt insertion and tightening based on in-hand object localization and force sensing. In: 2018 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM). IEEE. https://doi.org/10.1109/aim.2018.8452338
17. Yibang Z, Xiaoyong W, Lanzhu Z (2023) Visual identification and pose estimation algorithms of nut tightening robot system. The International Journal of Advanced Manufacturing Technology 127(11-12):5307-5326. https://doi.org/10.1007/s00170-023-11597-6
18. Ali MAH, Alshameri MA (2019) An intelligent adjustable spanner for automated engagement with multi-diameter bolts/nuts during tightening/loosening process using vision system and fuzzy logic. The International Journal of Advanced Manufacturing Technology 101(9-12):2795-2813. https://doi.org/10.1007/s00170-018-3133-0
19. Martinez P, Al-Hussein M, Ahmad R (2020) Intelligent vision-based online inspection system of screw-fastening operations in light-gauge steel frame manufacturing. The International Journal of Advanced Manufacturing Technology 109(3-4):645-657. https://doi.org/10.1007/s00170-020-05695-y
20. Zhuang C, Gong J, Liu J (2021) Digital twin-based assembly data management and process traceability for complex products. Journal of Manufacturing Systems 58:118-131. https://doi.org/10.1016/j.jmsy.2020.05.011
21. Arakelian V (2016) Gravity compensation in robotics. Advanced Robotics 30(2):79-96. https://doi.org/10.1080/01691864.2015.1090334
22. ISO (2018) ISO 15552:2018 Pneumatic fluid power. Cylinders with detachable mountings, 1 000 kPa (10 bar) series, bores from 32 mm to 320 mm. Basic, mounting and accessories dimensions. International Organization for Standardization, Geneva
23. Shukla A, Karki H (2016) Application of robotics in onshore oil and gas industry—A review Part I. Robotics and Autonomous Systems 75:490-507. https://doi.org/10.1016/j.robot.2015.09.012
24. Shukla A, Karki H (2016) Application of robotics in offshore oil and gas industry— A review Part II. Robotics and Autonomous Systems 75:508-524. https://doi.org/10.1016/j.robot.2015.09.013
25. Kumar Singh A (2023) Explosion-Proof Equipment in Hazardous Area. Springer Nature Singapore. https://doi.org/10.1007/978-981-99-2516-2
26. IEC (2017) IEC 60079-0:2017 Explosive atmospheres. Part 0: Equipment. General requirements. International Electrotechnical Commission, Geneva
27. European Parliament and Council (2014) Directive 2014/34/EU on the harmonisation of the laws of the Member States relating to equipment and protective systems intended for use in potentially explosive atmospheres. Official Journal of the European Union L 96:309-356

28. IEC (2016) IEC 61511-1:2016 Functional safety. Safety instrumented systems for the process industry sector. Part 1: Framework, definitions, system, hardware and application programming requirements. International Electrotechnical Commission, Geneva
29. Ovatman T, Aral A, Polat D, Ünver AO (2016) An overview of model checking practices on verification of PLC software. Software & Systems Modeling 15(4):937-960. https://doi.org/10.1007/s10270-014-0448-7
30. Liu HC, Liu L, Liu N (2013) Risk evaluation approaches in failure mode and effects analysis: A literature review. Expert Systems with Applications 40(2):828-838. https://doi.org/10.1016/j.eswa.2012.08.010
31. Negahban A, Smith JS (2014) Simulation for manufacturing system design and operation: Literature review and analysis. Journal of Manufacturing Systems 33(2):241-261. https://doi.org/10.1016/j.jmsy.2013.12.007
32. ISO (1998) ISO 261:1998 ISO general purpose metric screw threads. General plan. International Organization for Standardization, Geneva
33. ASTM International (2023) ASTM A193/A193M-23 Standard specification for alloy-steel and stainless steel bolting for high temperature or high pressure service and other special purpose applications. ASTM International, West Conshohocken
34. ISO (2013) ISO 898-1:2013 Mechanical properties of fasteners made of carbon steel and alloy steel. Part 1: Bolts, screws and studs with specified property classes. International Organization for Standardization, Geneva