Dimensionally Correct Is Not Functionally Aligned: Fastening Shift, Boresight Offset, and the Self-Test That Closes the Loop

Architecture, Computational Methods, and Field Evidence from MileSoft ADAS Radar Assembly

MileSoft Engineering Research Group

MileSoft Software Technologies

https://milesoft.net/

Corresponding authors: *Deepak Daryani* (deepakdaryani@milesoft.net), *Tanuj Daryani* (tanujdaryani@milesoft.net)

Working paper — Version 1.0 — 24 August 2026

Permanent link: https://milesoft.net/research/products/adas-radar-assembly

Abstract

A radar module can be mounted perfectly and point somewhere else. Vision measures where the housing sits; the boresight is where the beam actually goes, and the offset between them lives inside the part where no dimensional measurement can reach it. That gap is the reason this station exists and the reason its fourth check is not optional.

This paper presents the architecture and computational methods of MileSoft ADAS Radar Assembly, a station performing four checks in a fixed order: fiducial vision alignment, closed-loop torque fastening, post-fastening vision re-verification, and a sealed functional self-test in which the module measures its own alignment against a station fiducial and reports it back.

Three results follow. The order is load-bearing rather than procedural: fastening perturbs the position alignment established, so a station verifying only before the screws are driven certifies a state that no longer exists by the time the unit leaves. The error budget composes placement, fastening shift, internal boresight offset and measurement uncertainty, and only the last check observes their sum — which is what makes a functional test categorically different from a more precise dimensional one. And because sub-degree misalignment passes visual inspection, the escape latency is set by how long a vehicle takes to complain rather than by how long a plant takes to notice, so the population at risk is detection latency multiplied by production rate.

A Tier-1 automotive deployment reported a 92% reduction in customer containment incidents, and that figure is attributed in its own source to this station. The same deployment's zero wrong-part picks and five-day ramp-up are attributed to the companion kitting module and are not claimed here. The product page carries no screenshots, so every figure is an authored schematic.

Keywords: ADAS radar assembly; Radar module alignment; Boresight calibration; Vision guided assembly; Closed loop torque; Functional self test; Containment incident prevention; Automotive functional safety; Per unit digital pedigree; Sub-degree alignment tolerance; Fixture release interlock; Guard banding measurement uncertainty

1 Introduction

Radar modules feeding adaptive cruise control, lane keeping and automatic emergency braking are unforgiving. The product's material puts the failure mode plainly: sub-degree misalignment in the assembly fixture surfaces months later as a customer complaint and a containment letter that costs more than the line is worth.

Two properties make this defect class distinctive, and they compound. It passes visual inspection, so nothing at the plant notices. And the function it degrades is one the vehicle exercises rarely and the driver notices only when it matters, so the feedback arrives slowly.

A station that prevents it therefore has to do something a dimensional fixture cannot: confirm that the module is pointing where the design intends, rather than that it is sitting where the design intends. Those are different statements and the gap between them is inside the part.

1.1 Sitting correctly and pointing correctly

A vision system measures the housing against a fiducial. It can be arbitrarily precise about where the housing is and it observes nothing about where the antenna beam goes.

The offset between the two — the boresight offset — is a property of the individual module, arising from assembly tolerances inside it and from the calibration it carries. A module with a boresight offset sits perfectly in a perfect fixture and points somewhere else.

That is why the product's material specifies a sealed self-test in which the radar boots a routine against a fixed station fiducial and reports its own measured alignment. It is not a more accurate dimensional check. It is the only check of a different kind.

> A station that verifies position to a tighter tolerance is buying precision on three terms of a four-term budget. The fourth term does not get smaller.

1.2 Contributions

1. The four checks as an ordered sequence, with fastening shown to invalidate the check before it (Section 4.1).
2. An alignment error budget composing placement, fastening shift, boresight offset and measurement uncertainty (Section 4.2).

3. The functional self-test as the only observation of the composite, and the guard band that measurement uncertainty forces (Section 4.3).
4. Escape latency and the population at risk, connecting station prevention to containment breadth (Section 4.4).
5. An eight-dimension capability reference framework for alignment and validation stations (Section 6.3).

2 Background and Related Work

This station sits at the intersection of functional safety, fastening practice and measurement science, and each supplies an obligation the others do not.

2.1 What a safety-related item obliges of production

A radar feeding automatic emergency braking is a safety-related item, and ISO 26262 governs it. Part 1 supplies the vocabulary; Part 7 covers production, operation, service and decommissioning — the obligations that fall on the plant rather than on the design house.

The relevant consequence is that the production process itself carries requirements. It is not sufficient for a design to be safe; the process must be capable of producing it, must detect when it does not, and must be able to demonstrate both afterwards.

This is why the per-unit signed report the product describes is a requirement rather than a convenience. A safety argument that cannot be evidenced per unit is a safety argument about a design, and the vehicle contains a particular part.

2.2 Why fastening is a perturbation and not a fixation

The intuitive model of mounting is that alignment sets the position and fastening preserves it. In practice a tightening sequence applies torque about an axis, draws a joint down unevenly as successive fasteners take load, and moves the part.

VDI/VDE 2862 classifies joints by the consequence of failure, and a mounting whose misalignment degrades a braking function sits at the top of that classification. Its implication here is that the fastening must be captured — torque and angle per fastener — and not merely performed.

But capture alone does not address displacement. A perfectly executed tightening to specification can still move the module, because correct torque and unchanged position are independent facts. That is the argument for the third check, and Section 4.1 states it formally.

2.3 A pass verdict inherits its measurement chain

ISO 5725 separates trueness from precision; AIAG's measurement systems analysis supplies the repeatability and reproducibility study that quantifies a gauge. Both apply to the station's own instruments, which is easy to forget when the instrument is a camera and the output is a number on a screen.

The consequence is that acceptance limits cannot equal specification limits. A unit measured at the edge of tolerance by an instrument with its own uncertainty may be outside it, and Section 4.3 gives the guard band that follows.

ISO 22514 then governs how the resulting distribution may be characterised — and specifically, whether capability or performance indices may legitimately be reported, which depends on whether the process is in statistical control rather than on which is more flattering.

3 System Overview

The station is a vision-aligned mounting fixture, a multi-vendor torque controller integration, a functional self-test against a fixed station fiducial, and a per-unit report generator, sitting behind a release interlock that holds the unit until every check has passed.

> Every figure in this paper is an authored schematic. This product's page carries no screenshots of any kind, and no image belonging to another product is used to stand in for one.

3.1 The sequence, and why each step is where it is

The product's material names four things the station must do and is explicit that skipping any of them leaves a failure mode open. The ordering is what this paper adds.

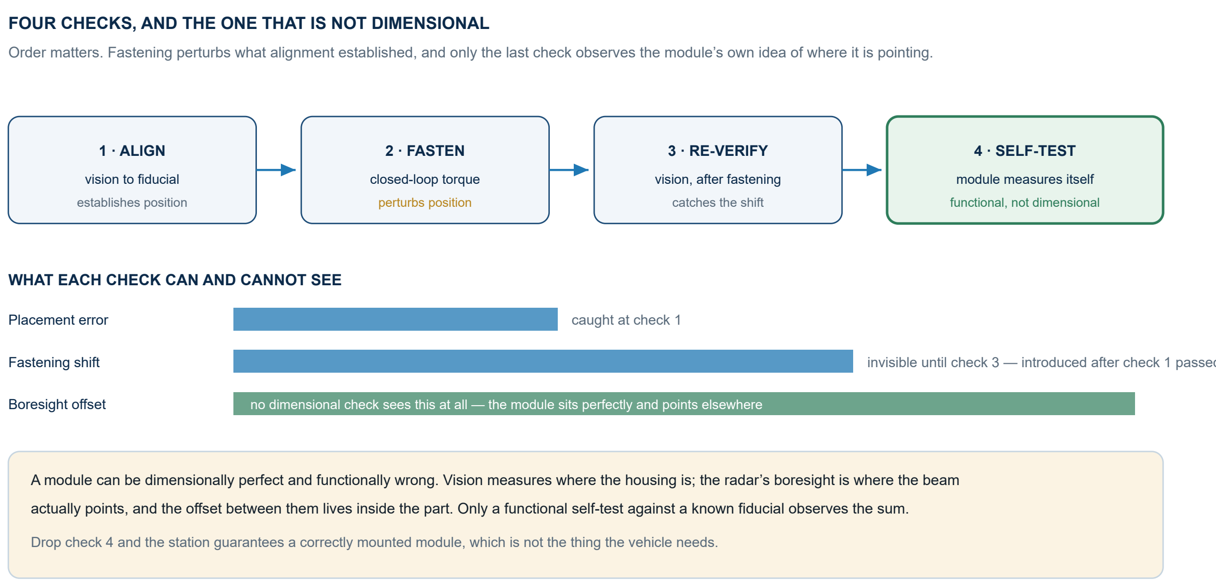

Figure 1. *Schematic.* Schematic (not product UI): the four checks and their coverage. Placement error is caught at the first; fastening shift is introduced after it and caught only at the third; and boresight offset is invisible to every dimensional check, observed only by the module measuring itself at the fourth.

Reading the bars gives the design rule. Each check covers the errors introduced before it and none of those introduced after, so a check is worth exactly the errors that can still occur downstream of it — which is why the last one carries the guarantee.

3.2 The fixture that will not let go

The mounting fixture does not release the unit until alignment is within tolerance and the self-test has passed. That is a control function in Shingo's sense — the next step is impossible rather than discouraged — and it is the same architectural choice the companion kitting paper analyses at a bin.

The failure direction of that interlock is a safety question rather than a convenience one. A release circuit whose failure mode is to open allows an unverified safety-related part onto a vehicle, so ISO 13849-1 governs its design and the performance level it must achieve.

There is an operational cost worth acknowledging. A fixture that holds a failing unit stops the station, and a station that stops stops the line — so the reaction path for a failed check has to be designed alongside the check itself, which IATF 16949 requires in any case.

3.3 Fastening on the same controller backbone

The station integrates torque controllers from the major vendors over the same protocol backbone the traceability module uses, so torque and angle per fastener are captured rather than assumed.

Sharing the backbone is a real architectural benefit rather than a convenience: the fastening record on a radar mounting has the same shape as every other critical joint on the line, so it lands in one audit trail and one capability analysis rather than in a station-specific log.

3.4 What ships with the unit

Every unit exits with a signed report carrying serial or vehicle identification, alignment values in azimuth, elevation and roll, torque per fastener, station identity, operator identity and timestamp, in a form exportable to the customer's traceability specification.

Three properties of that record decide whether it is worth anything in a dispute. The alignment values must be the self-test's measurements rather than the fixture's targets. The operator must be authenticated rather than inherited from a station login. And the record must be produced by the station rather than assembled afterwards, because a report generated later is a report that could have been generated differently.

> The value of a per-unit pedigree is proportional to how much of it was measured rather than asserted. A report reproducing the specification the unit was supposed to meet is a copy of the drawing.

3.5 Calibrating the station itself

The product's material notes calibration validation of the station on a schedule an auditor will recognise, which is easy to read past and is load-bearing.

The fiducial the self-test measures against is the reference for every alignment value the station has ever produced. If it drifts, every unit passed since the last validation is suspect — and by Section 4.4 that population is the validation interval multiplied by the production rate.

This makes the calibration interval a containment-breadth decision rather than a maintenance one, and it should be set with that in view rather than by convention.

4 Computational Methods

Four computations carry the paper: why the order matters, what the error budget contains, what a pass verdict requires given uncertainty, and what a slow escape costs.

4.1 Fastening invalidates the check before it

Let theta be the module's angular position relative to the design intent, measured about the axis that matters.

the position at the first check is the placement error; the position after fastening adds a shift; so a unit inside tolerance at check 1 can be outside it at check 3 (shift)

Delta(fasten) is the angular displacement introduced by the tightening sequence and tau the alignment tolerance. The inequality is strictly positive whenever the shift has any variance at all, which it does.

That inequality is the entire argument for the third check and it holds regardless of how good the fastening process is. A perfectly capable tightening to specification still has a displacement distribution, and a unit placed near the tolerance edge will occasionally be pushed across it.

The design implication is not to tighten more carefully but to measure afterwards. A station verifying only before fastening is certifying theta(1) and shipping theta(3), and the two are different random variables.

> This is the check most often removed on cost grounds, because the first vision check already exists and re-running it looks redundant. It is not redundant; it is the only check that observes the second term.

4.2 The four terms, and which instruments see which

The alignment a vehicle actually experiences composes several independent contributions.

effective alignment is placement plus fastening shift plus the module's internal boresight offset plus measurement error; the combined uncertainty is the root sum of squares of the individual contributions (budget)

delta(bore) is the offset between the housing datum and the antenna boresight, a property of the individual module, and epsilon the measurement error of whatever observed it.

Table 1. Which check observes which term. The third column is the one that decides the station's architecture.

Term	Origin	Observed by
theta(place)	Fixture and operator placement	Vision, check 1
Delta(fasten)	The tightening sequence	Vision, check 3 — and nothing earlier
delta(bore)	Inside the module	The functional self-test only
epsilon	The station's own instruments	Nothing — it is absorbed by guard banding

The third row is why a functional check is categorically different from a better dimensional one. Boresight offset has no external geometric signature; a vision system of unlimited resolution measuring a perfectly manufactured housing reports zero error on a module whose beam is off by half a degree.

It also explains why upstream variant control matters to this station. A close-cousin variant has a different nominal boresight relationship, so a wrong-variant module can pass every dimensional check and fail functionally — which is the coupling to the kitting station the product's material describes.

4.3 What a pass verdict costs in tolerance

The station's own measurement has uncertainty, so accepting everything measured inside specification accepts some units that are outside it.

the acceptance limit is the specification limit tightened by a multiple of the combined uncertainty; raising that multiple reduces escapes and increases false rejects (guard)

k is the guard-band factor, chosen against the relative cost of an escape and a false reject. It is a business decision expressed as a number, not a technical constant.

Choosing k for a safety-related part is not symmetric. A false reject costs a rework at the station; an escape costs a containment programme whose breadth Section 4.4 computes. The cost ratio runs to several orders of magnitude, which argues for a k that would be indefensibly conservative on an ordinary part.

This is also where the measurement-systems obligation bites. u(tot) cannot be assumed; it has to be established by a repeatability and reproducibility study on the station's own instruments, and a guard band computed from an unmeasured uncertainty is a number with no content.

> A station reporting pass and fail without a stated uncertainty and guard band is reporting a verdict whose false-accept rate nobody has computed. For a braking-related part that is the substantive gap, not a documentation one.

4.4 What a silent defect costs, and why it is not per unit

The defining property of this failure mode is that it passes inspection. That sets the detection latency outside the plant's control and makes the cost superlinear in it.

the population at risk is production rate times the interval from onset to detection; total cost is that population times the per-unit containment cost plus a fixed programme cost; and the detection time is not a property the plant controls (escape)

r is the production rate, t(onset) when the process began producing misaligned units, t(detect) when the customer raised it, and c(contain) the per-unit cost of finding, recalling and replacing.

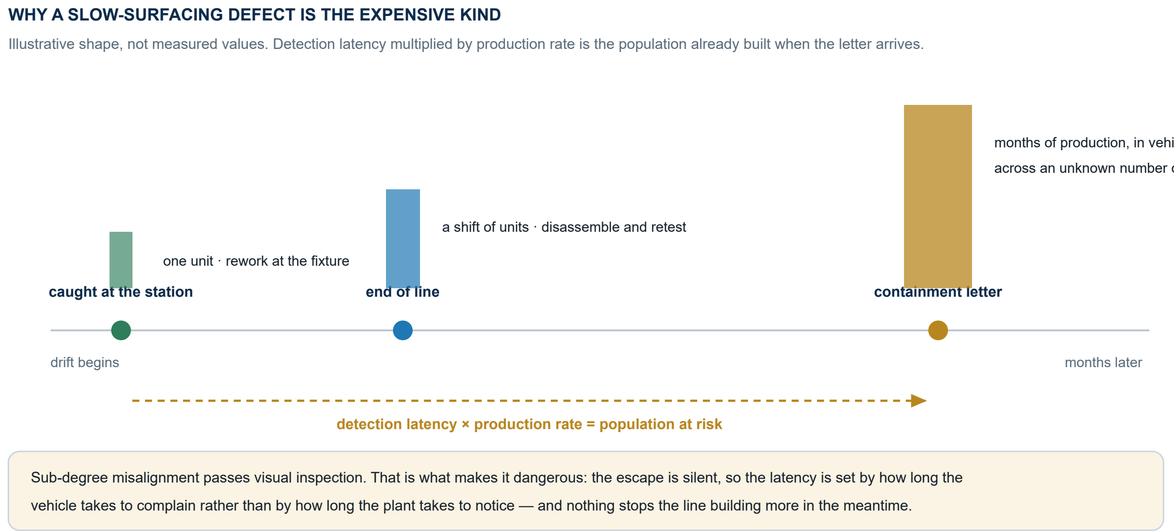

Figure 2. *Schematic.* Schematic (illustrative shape, not measured values): the same drift caught at three points. At the station it is one unit and a rework; at end of line a shift's production; at the customer, months of output already in vehicles. Because sub-degree misalignment passes visual inspection, the latency is set by how long a vehicle takes to complain.

Two consequences follow that a per-unit cost model misses entirely. The cost of an escape is not the cost of one defective part; it is the cost of every part built while nobody knew, so the same defect is a thousand times more expensive in a high-rate plant than a low-rate one. And because t(detect) is outside the plant's control, the only lever available is t(onset) — which means detecting the drift rather than the defect.

That reframes what the station's statistical reporting is for. Capability tracking on the alignment distribution detects a process drifting toward the tolerance edge before it crosses, which reduces t(detect) to something the plant does control — and it is the only mechanism in this paper that acts on the multiplier rather than the per-unit term.

The connection to containment breadth is direct. Once an escape has occurred, the affected population is bounded by whatever genealogy the plant can compute, which is the subject of the companion material-traceability paper. Prevention at the station and containment breadth afterwards are two ends of the same argument.

5 Reported Outcomes and Field Evidence

This station carries one attributed deployment result and a set of design statements. They are separated here because the design statements are checkable and the result is not reproducible from what is published.

5.1 The result attributed to this station

A Tier-1 automotive supplier building radar modules reported that off-spec alignment was slipping through quality control and surfacing at the customer, generating containment incidents and damaging the relationship.

After integrating this station with the manufacturing execution system — each radar booting through a sealed alignment routine and shipping a signed alignment report — the deployment reports a 92% reduction in customer containment incidents.

> Attribution, stated plainly. This result is attributed in its own source to this station. The same deployment's zero wrong-part picks and five-day operator ramp-up are attributed to the companion kitting module, and its throughput gain to the programme; neither is claimed here.

The published mechanism is the one Section 4.2 formalises: the sealed self-test is what converts a dimensional guarantee into a functional one, and the signed report is what makes the guarantee evidenceable per unit.

Three caveats belong with the figure. A 92% reduction in incidents is a count over an unstated period, and incidents are lumpy — one containment programme can span thousands of units, so a reduction in incident count is not proportional to a reduction in units affected. No baseline incident rate is published. And the same programme installed the kitting station upstream, which Section 4.2 notes prevents wrong-variant modules reaching the fixture — a mechanism that would also reduce containment incidents and is not separated out.

5.2 Design statements published for this station

Table 2. Statements published about this station, separated by kind. All of the design statements are checkable by demonstration.

Statement	What it asserts	Kind
Vision-aligned fixture	the fixture does not release until alignment is in tolerance	Design
Closed-loop torque	torque captured per fastener on a multi-vendor backbone	Design
Post-fastening re-verification	alignment re-checked after the screws are driven	Design
Sealed self-test	the module measures its own alignment against a station fiducial	Design
Per-unit signed report	identity, azimuth, elevation, roll, torque, station, operator, time	Design
Station calibration validation	on a schedule an external auditor will recognise	Design
Sub-degree tolerance	the regime standard fixtures cannot guarantee	Context

The last row is deliberately labelled context rather than claim. No numeric tolerance is published anywhere in the product's material, so this paper treats sub-degree as the order of magnitude of the problem and computes nothing from it.

5.3 Modelled rework avoided

The published return model prices the saving as rework avoided through automated alignment and validation. Its assumptions are printed here so a reader can substitute their own.

Table 3. Modelled scenario — not a deployment result. Assumptions are the published defaults of the return-on-investment model for this product.

Assumption	Value
Modules assembled per month	5,000 (60,000 per year)
Baseline alignment defect rate	3% (1,800 per year)
Reduction attributed to the station	70%
Modelled defects avoided	1,260 per year
Cost per defect	site-specific; the published default illustrates only

The model is conservative in a way worth naming, because Section 4.4 says the conservatism is structural. It prices defects at a per-unit rework cost, and the defect this station exists to prevent is precisely the one that does not get reworked — it escapes, and its cost is the population at risk rather than the unit.

> A site sizing this investment on rework cost alone is valuing the least expensive thing the station prevents. The containment exposure of Equation (escape) is the term that dominates, and it is the harder one to put in a spreadsheet.

6 Discussion

6.1 Functional checks are not precision upgrades

The most common way to get this station wrong is to treat the self-test as a redundant confirmation of what vision already established, and therefore as the obvious candidate for removal when cycle time is under pressure.

Section 4.2 shows why that reasoning fails. The self-test is not measuring the same quantity more precisely; it is measuring a different quantity — the composite that includes a term inside the part. Removing it does not reduce confidence in an estimate, it eliminates the only observation of a whole error source.

The general principle transfers. Whenever a station verifies a proxy for the property that matters, adding precision to the proxy has a ceiling set by the proxy's relationship to the property — and past that ceiling only a measurement of the property itself helps.

> Vision measures where the housing is. The vehicle cares where the beam points. Everything else in this paper follows from those being different sentences.

6.2 Detect the drift, not the defect

Equation (escape) makes the cost of this failure mode proportional to a latency the plant does not control. That looks like a dead end and is not, because the plant does control when it notices the process moving.

A capability analysis on the alignment distribution — reported per station, per shift, against a stated control state — detects a mean walking toward the tolerance edge while every individual unit is still passing. That converts an uncontrollable detection time into a controllable one, and it is the only lever in this paper that acts on the multiplier rather than the per-unit cost.

This is also where ISO 22514's control-state distinction earns its place. A capability index computed on a process that is drifting is not describing capability, and a station reporting a healthy index on an out-of-control process is providing exactly the false assurance that lets the latency run.

The practical recommendation is therefore narrower than continuous monitoring in general: track the alignment mean and spread per station, alarm on movement rather than on failure, and treat a station whose distribution has shifted as requiring investigation even when its yield is unchanged.

6.3 A capability reference framework for alignment and validation stations

Table 4. Capability reference framework. Each dimension separates a validation station from an instrumented fixture, and each is answerable by demonstration on a live station rather than by reading a specification.

Dimension	Question the system must answer by demonstration
D1 Post-fastening check	Is alignment re-measured after the screws are driven, or only before?
D2 Functional measurement	Does the module report its own alignment, or does the station infer it?
D3 Release interlock	Fail the self-test deliberately. Does the fixture let the unit go?
D4 Failure direction	If the release circuit fails, does the fixture hold or open?
D5 Stated uncertainty	What is the station's measurement uncertainty, and how was it established?
D6 Guard band	Is the acceptance limit tighter than the specification limit, and by how much?
D7 Measured pedigree	Are the values on the per-unit report measurements, or the targets?
D8 Drift alarm	Does the station alarm on a distribution shift, or only on a failing unit?

D8 is the dimension that acts on cost rather than on defects. Every other dimension improves the chance of catching a bad unit; only this one shortens the interval in which bad units are being built.

6.4 Generalisability

The proxy-versus-property argument generalises to every station verifying a geometric surrogate for a functional requirement: camera modules, lidar housings, antenna arrays, optical assemblies, wheel alignment. In each case the internal offset term exists and no external measurement observes it.

The escape-latency model generalises to any defect that passes inspection and surfaces in service — which is the defining class of expensive quality failures across industries, and the reason drift detection matters more than yield in all of them.

What does not generalise is the 92% figure. It is one supplier's incident count against an unpublished baseline, in a programme that also installed an upstream variant control with an overlapping mechanism.

7 Threats to Validity and Limitations

1. The 92% result counts incidents, not units. Containment incidents are lumpy — one programme can span thousands of units — so a reduction in incident count does not translate proportionally into units affected.
2. The result is partly confounded. The same programme installed the kitting station upstream, which prevents wrong-variant modules reaching the fixture; that mechanism would also reduce containment

incidents and is not separated out.

3. No baseline incident rate is published, and no period is stated, so no rate can be computed from the percentage.
4. No numeric tolerance is published. Sub-degree is treated in this paper as the order of magnitude of the problem, and nothing is computed from it.
5. No measurement uncertainty is published. Equation (guard) requires u(tot) and no figure for the station's vision or self-test uncertainty appears in the product's material.
6. The fastening-shift distribution is unpublished. Section 4.1 argues the shift has non-zero variance; its magnitude at a given joint is a site-specific quantity the paper does not supply.
7. The self-test's own reference can drift. Every alignment value the station has produced is relative to a station fiducial, so a drifted fiducial invalidates a population bounded by the calibration interval.
8. No figure in this paper is a product capture. This product's page carries no screenshots, so nothing here demonstrates the described station exists in the form modelled.

> The seventh limitation is the one that inverts the paper's own argument, and it should be stated rather than buried. The functional self-test removes the plant's dependence on a dimensional proxy and replaces it with a dependence on a reference artefact — and the calibration interval on that artefact is now a containment-breadth decision.

8 Future Work

- Publishing the station's measurement uncertainty and its guard band. Equation (guard) cannot be applied without u(tot), and a pass verdict whose false-accept rate has never been computed is the substantive gap for a braking-related part.
- Characterising the fastening-shift distribution per joint. Section 4.1 argues the shift has non-zero variance; measuring it would let the first check's tolerance be set so that the third check rarely fails.
- Drift alarms on the alignment distribution rather than on failing units, which is the only lever in this paper that shortens the interval during which bad units are built.
- Continuous validation of the station fiducial, so the population invalidated by a drifted reference is bounded by minutes rather than by a calibration interval.
- Reporting containment outcomes in units affected rather than incident count, so a reduction can be sized against the exposure Equation (escape) computes.

9 Conclusion

A radar module can be mounted perfectly and point somewhere else, and no dimensional measurement of any precision will say so. That single fact determines the architecture of this station.

Three results follow. The check order is load-bearing: fastening perturbs the position alignment established, so a station verifying only before the screws are driven certifies a state that has already changed, and the post-fastening re-verification is the step most often removed on cost grounds and the one that makes the first check mean anything. The error budget composes placement, fastening shift, boresight offset and measurement uncertainty, and the functional self-test is the only observation of the sum — which makes it categorically different from a more precise dimensional check rather than a redundant one. And because sub-degree misalignment passes visual inspection, the cost of an escape is not per unit: it is the production rate multiplied by a latency the plant does not control.

That last result points at the only lever available. The plant cannot shorten the time a vehicle takes to complain, but it can notice its own process moving before any unit fails — which is why capability tracking on the alignment distribution, alarmed on movement rather than on failure, is the intervention that acts on the multiplier rather than the unit.

The framework of Section 6.3 is offered as the durable contribution, and its second dimension is the question to open with: does the module report its own alignment, or does the station infer it?

Appendix A Nomenclature

Table 5. Symbols and abbreviations used in this paper.

Symbol / term	Meaning
theta(eff)	Effective alignment the vehicle experiences
theta(place)	Placement error established at the first vision check
Delta(fasten)	Angular displacement introduced by the tightening sequence
delta(bore)	Boresight offset — housing datum against antenna beam, inside the part
epsilon	Measurement error of whatever instrument observed the value
u(tot)	Combined measurement uncertainty, root sum of squares of contributions
tau(spec), tau(accept)	Specification limit, and the acceptance limit after guard banding
k	Guard-band factor — a business decision expressed as a number
r	Production rate
t(onset), t(detect)	When the process began producing misaligned units, and when it was found
N(risk)	Population at risk — production rate times detection latency
c(contain)	Per-unit cost of finding, recalling and replacing an escaped unit
Boresight	The direction the antenna beam actually points
Fiducial	A fixed reference feature the vision system or module measures against
ADAS	Advanced driver assistance systems

Appendix B Worked Numerical Examples

Appendix B.1 Why the third check catches units the first one passed

A tolerance of 0.30 degrees. Placement error is centred at zero with a standard deviation of 0.09 degrees. The tightening sequence introduces a shift centred at 0.04 degrees with a standard deviation of 0.07.

Applying Equation (shift): units passing check 1 have absolute placement error under 0.30. Take a unit measured at 0.26 — comfortably inside, and a routine pass. After fastening its position is 0.26 plus a draw from a distribution centred at 0.04 with spread 0.07, so it exceeds 0.30 with probability about 0.72.

Aggregating over the placement distribution, roughly 3.1% of units passing check 1 fail check 3. A station without the third check ships all of them, and none of them looks wrong at any point in the process.

Note that the fastening process here is entirely healthy — a mean shift of 0.04 on a 0.30 tolerance, tightly held. The escapes come from the interaction with placement, not from a fastening defect, which is why tightening more carefully does not address them.

Appendix B.2 The term no camera sees

The same station, now with a boresight offset distribution across incoming modules centred at zero with a standard deviation of 0.11 degrees, and a station measurement uncertainty of 0.04.

Applying Equation (budget): the dimensional terms contribute a spread of 0.114 degrees. Adding boresight gives 0.158, and with measurement uncertainty 0.163.

Boresight offset alone contributes nearly as much spread as placement and fastening combined. A station verifying only dimensionally certifies a 0.114 process and ships a 0.158 one — it is wrong about its own capability by 39%, in the optimistic direction.

Against a 0.30 tolerance the dimensional-only view suggests a comfortable 2.6 sigma; the functional reality is 1.8 sigma. The first supports a capability claim and the second does not, and no amount of vision resolution moves the second number.

Appendix B.3 What the latency multiplies

A line builds 5,000 modules a month. A fixture wears and the alignment mean walks past tolerance. The plant's yield does not change, because the drift is inside the guard band for the first six weeks and the units that do fail are attributed to normal variation.

The customer raises a containment letter four months after onset. Applying Equation (escape): the population at risk is 5,000 x 4 = 20,000 modules. At a per-unit containment cost of 45 — finding, recalling, replacing, plus field labour — the exposure is 900,000 plus the programme's fixed cost.

Compare a drift alarm on the alignment distribution firing at week two, when the mean has moved by one standard deviation and every unit is still passing. The population at risk becomes 5,000 x 0.5 = 2,500 modules, and most are still in the plant or in transit rather than in vehicles.

The alarm did not improve the station's ability to detect a bad unit; it detected a moving process. That distinction is worth 17,500 modules of exposure in this example, and it is the whole content of Section 6.2.

Data provenance

Every quantitative claim in this paper resolves to one of the entries below. Figures marked as modelled estimates are not deployment measurements; their assumptions are stated.

Metric	Reported value	Provenance
OEM containment incidents	92% reduction	**Operator-reported — Tier-1 Manufacturers (Pune & Nasik plant, India)** /case-studies/pokayoke
Checks the station performs before releasing a unit	fiducial vision alignment, closed-loop torque, post-fastening re-verification, sealed functional self-test	**Product specification** /products/adas-radar-assembly
Release condition on the mounting fixture	the fixture does not release until alignment is within tolerance and the self-test passes	**Product specification** /products/adas-radar-assembly
Fields on the signed per-unit report	serial or VIN, azimuth, elevation, roll, torque per fastener, station, operator, timestamp	**Product specification** /products/adas-radar-assembly
Alignment tolerance a standard fixture cannot guarantee	sub-degree	**Product specification** /products/adas-radar-assembly — Stated on the product blog and FAQs; no numeric tolerance is published, so this paper treats sub-degree as the order of magnitude rather than a figure.
Alignment defects avoided per year	1,260 of 1,800 (modelled)	**Modelled estimate** 5,000 modules assembled per month, 60,000 per year; Baseline alignment defect rate of 3%; 70% reduction attributed to the station; Model and defaults published in src/data/roiModels.ts

References

[1] International Organization for Standardization (2018). ISO 26262-1: Road vehicles - Functional safety - Part 1: Vocabulary. ISO, Geneva; the 2018 revision extends scope to all road vehicles except mopeds. https://www.iso.org/standard/68383.html

[2] International Organization for Standardization (2018). ISO 26262-7: Road vehicles - Functional safety - Part 7: Production, operation, service and decommissioning. ISO, Geneva.

[3] International Automotive Task Force (2016). Quality management system requirements for automotive production and relevant service parts organizations. IATF 16949:2016, first edition, 1 October 2016 (superseding ISO/TS 16949).

[4] Verein Deutscher Ingenieure / Verband der Elektrotechnik (2025). VDI/VDE 2862 Blatt 1: Minimum restrictions for application of fastening systems and tools - Applications in the automotive industry. VDI/VDE, Dusseldorf.

[5] International Organization for Standardization (2017). ISO 6789-1: Assembly tools for screws and nuts - Hand torque tools - Part 1: Requirements and methods for design conformance testing and quality conformance testing. ISO, Geneva. https://www.iso.org/standard/62549.html

[6] International Organization for Standardization (2017). ISO 6789-2: Assembly tools for screws and nuts - Hand torque tools - Part 2: Requirements for calibration and determination of measurement uncertainty. ISO, Geneva. https://www.iso.org/standard/62550.html

[7] International Organization for Standardization (2014). ISO 22514-1: Statistical methods in process management - Capability and performance - Part 1: General principles and concepts. ISO, Geneva. https://www.iso.org/standard/64135.html

[8] International Organization for Standardization (2017). ISO 22514-2: Statistical methods in process management - Capability and performance - Part 2: Process capability and performance of time-dependent process models. ISO, Geneva. https://www.iso.org/standard/71617.html

[9] International Organization for Standardization (2023). ISO 5725-1: Accuracy (trueness and precision) of measurement methods and results - Part 1: General principles and definitions. ISO, Geneva (superseding ISO 5725-1:1994).

[10] International Organization for Standardization (2019). ISO 5725-2: Accuracy (trueness and precision) of measurement methods and results - Part 2: Basic method for the determination of repeatability and reproducibility. ISO, Geneva.

[11] Automotive Industry Action Group (Chrysler, Ford, General Motors) (2010). Measurement Systems Analysis (MSA) Reference Manual. AIAG, 4th edition, MSA-4.

[12] International Organization for Standardization (2015). ISO 13849-1: Safety of machinery - Safety-related parts of control systems - Part 1: General principles for design. ISO, Geneva; a 2023 edition has since been published. https://www.iso.org/standard/69883.html

[13] Shingo, S. (translated by A. P. Dillon) (1986). Zero Quality Control: Source Inspection and the Poka-Yoke System. Productivity Press, Cambridge, MA. Originally published in Japanese, 1985.

[14] Nakajima, S. (1988). Introduction to TPM: Total Productive Maintenance. Productivity Press, Cambridge, MA.

[15] International Electrotechnical Commission / International Society of Automation (2025). Enterprise-control system integration - Part 1: Models and terminology. IEC 62264-1; ANSI/ISA-95.00.01-2025 (IEC 62264-1 Mod).

[16] International Society of Automation (2013). Enterprise-control system integration - Part 3: Activity models of manufacturing operations management. ANSI/ISA-95.00.03-2013 (IEC 62264-3 Modified).

[17] International Electrotechnical Commission (2025). IEC 62541-1: OPC unified architecture - Part 1: Overview and concepts. IEC, Geneva. https://webstore.iec.ch/en/publication/81513

[18] International Organization for Standardization (2014). ISO 22400-2: Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 2: Definitions and descriptions. ISO, Geneva; defines OEE as availability x effectiveness x quality ratio. https://www.iso.org/standard/54497.html

[19] International Organization for Standardization / International Electrotechnical Commission (2024). ISO/IEC 19987: Information technology - EPC Information Services (EPCIS) Standard, version 2.0. ISO/IEC, Geneva. https://www.iso.org/standard/85557.html

[20] Lee, J., Bagheri, B., and Kao, H.-A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18-23.