Improving DVB-S2 PLHEADER Detection at Low SNR

Reliable detection of the DVB-S2 Physical Layer Header (PLHEADER) is one of the first challenges faced by a DVB-S2 receiver. Before the receiver can decode the payload, it must determine where a physical-layer frame begins, establish carrier phase and frequency sufficiently accurately, and decode the signalling contained in the PLS code.

This becomes particularly interesting at very low SNR.

I have been comparing two approaches:

1. A traditional differential SOF/PLS detector.
2. A new coherent, maximum-likelihood-oriented PLHEADER detector designed to exploit substantially more of the information contained in the DVB-S2 header.

Both methods are trying to answer essentially the same question:

At what symbol does the DVB-S2 PLHEADER begin, and what carrier frequency and phase best explain the received 90-symbol header?

The difference is how much of the available information they use to answer that question.

The DVB-S2 PLHEADER

Every DVB-S2 physical-layer frame begins with a 90-symbol PLHEADER:

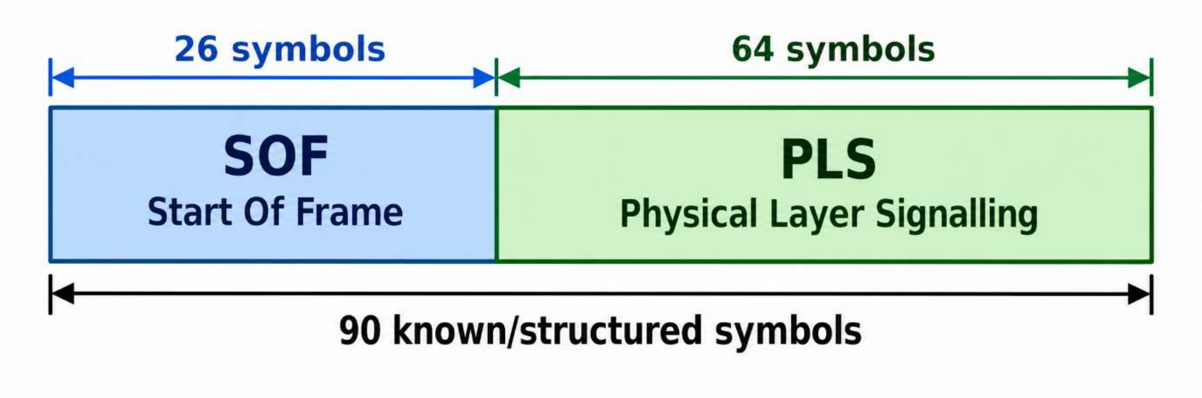

The first 26 symbols are the **Start Of Frame (SOF)**.

The following 64 symbols contain the **Physical Layer Signalling Code (PLSCODE)**, which identifies parameters including the MODCOD, frame type and pilot configuration.

The important observation is that the receiver has a great deal of prior knowledge about these 90 symbols.

The SOF sequence is known exactly.

Once a PLS hypothesis is chosen, the corresponding 64-symbol PLS sequence is also known.

This means that frame detection can potentially be treated as a problem of finding the known 90-symbol waveform hidden in noise.

Traditional Differential Detection

The traditional approach is elegant and computationally inexpensive.

Before symbol timing has been fully recovered, several fractional-symbol sampling positions are tested. For each sampling phase, interpolated symbols are generated.

The detector then forms differential products between adjacent symbols:

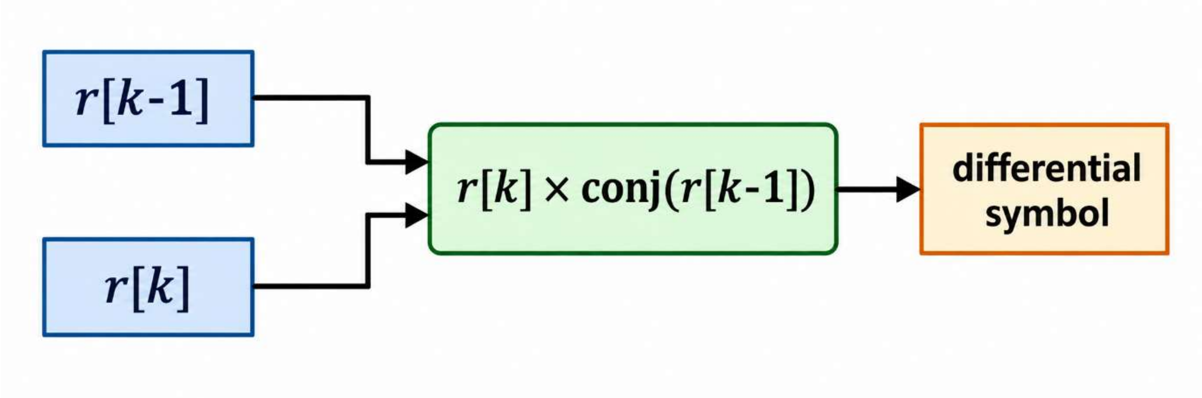

If the received symbols are approximately

r[k] = s[k] · exp(j(ϕ + kΔϕ)) + n[k]

where

s[k] = transmitted symbol

ϕ = unknown carrier phase

Δϕ = carrier rotation per symbol

n[k] = noise

then, ignoring noise,

r[k] · conj(r[k-1])

contains the differential phase between consecutive transmitted symbols plus the carrier rotation per symbol.

The unknown common carrier phase ϕ disappears.

That is the great attraction of differential detection.

Why differential detection works so well

At initial acquisition, the receiver does not yet know carrier phase accurately.

A coherent correlation such as

Σ r[k] conj(s[k])

can lose correlation energy if the carrier rotates significantly during the sequence.

Differential detection largely removes this problem.

It is therefore:

- simple;
- robust to unknown carrier phase;
- tolerant of carrier-frequency error;
- computationally inexpensive;
- suitable for searching many candidate positions.

The traditional implementation consequently uses differential correlation to locate the PLHEADER and obtain a rough carrier-frequency estimate.

It is a very sensible acquisition technique.

The Cost of Differential Detection

The advantage of differential detection is also its weakness.

Consider again:

d[k] = r[k] · conj(r[k-1])

Both r[k] and r[k-1] contain noise.

The operation therefore does not merely remove the unknown phase. It also combines noise from two received symbols.

Expanding the expression illustrates the problem:

r[k]   = s[k]   + n[k]

r[k-1]   = s[k-1]   + n[k-1]

and therefore

r[k] conj(r[k-1])

contains the desired signal product, but also signal/noise cross terms and a noise/noise term.

At high SNR this is not particularly troublesome.

At very low SNR it matters.

The differential representation is consequently not preserving all of the information that would be available to an ideal coherent detector.

This suggests a different approach:

Use differential or coarse correlation to solve the acquisition uncertainty, then use coherent processing once that uncertainty has been reduced.

That is the philosophy behind the new detector.

The New PLHEADER Detector

The new DvbS2FrameSynchronizer divides acquisition into stages.

Conceptually:

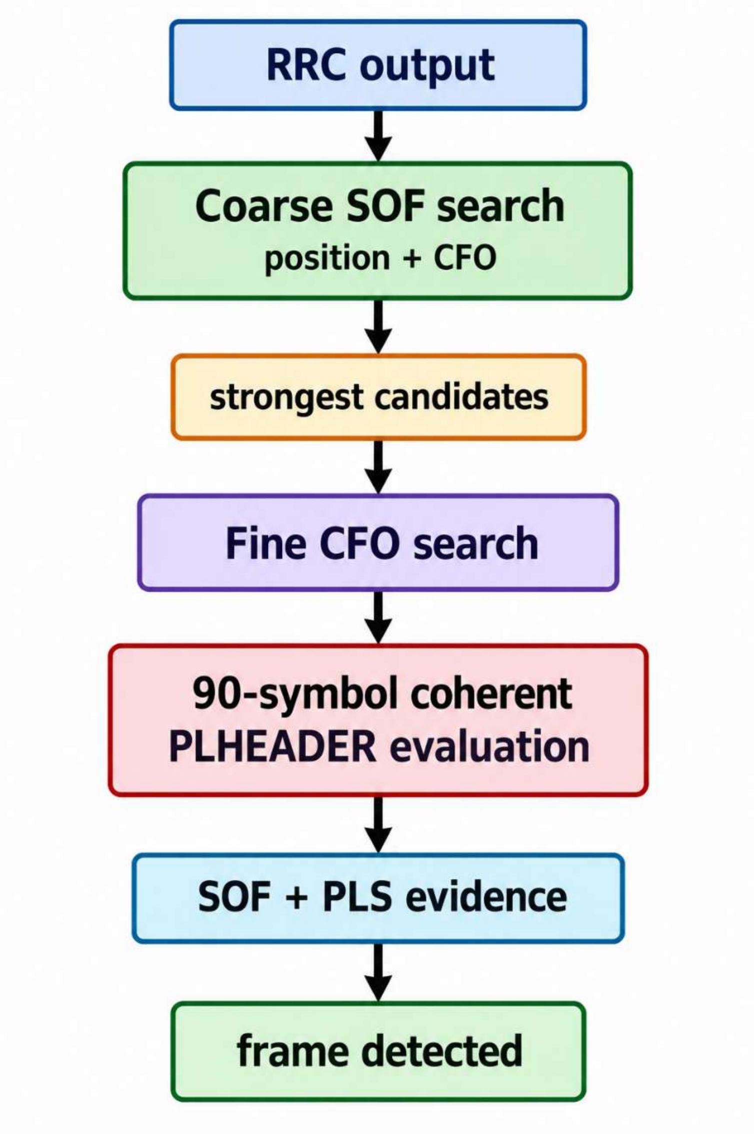

The important distinction is that **SOF detection is no longer the final decision**.

It is primarily a mechanism for finding promising hypotheses.

The expensive and sensitive test is performed only on those candidates.

Stage 1 — Searching the SOF

For initial acquisition, the detector searches the supplied symbol buffer for possible SOF locations.

For every candidate position it evaluates a CFO-aware correlation against the known 26-symbol SOF.

For a hypothesised carrier-frequency offset f, the phase advance per symbol is

$\Delta\phi = 2\pi f / Rs$

where Rs is the symbol rate.

Each received symbol can then be derotated according to that hypothesis before correlation with the known SOF.

Conceptually:

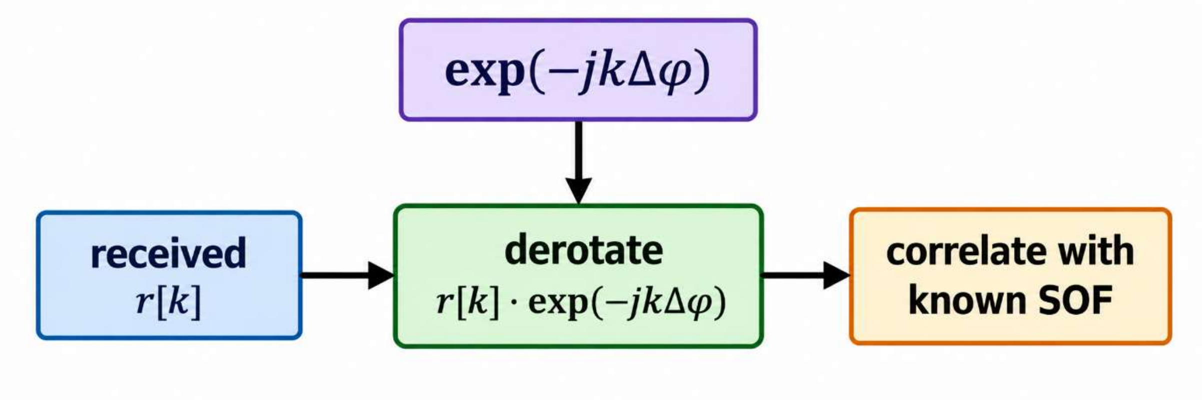

The detector therefore searches a two-dimensional space:

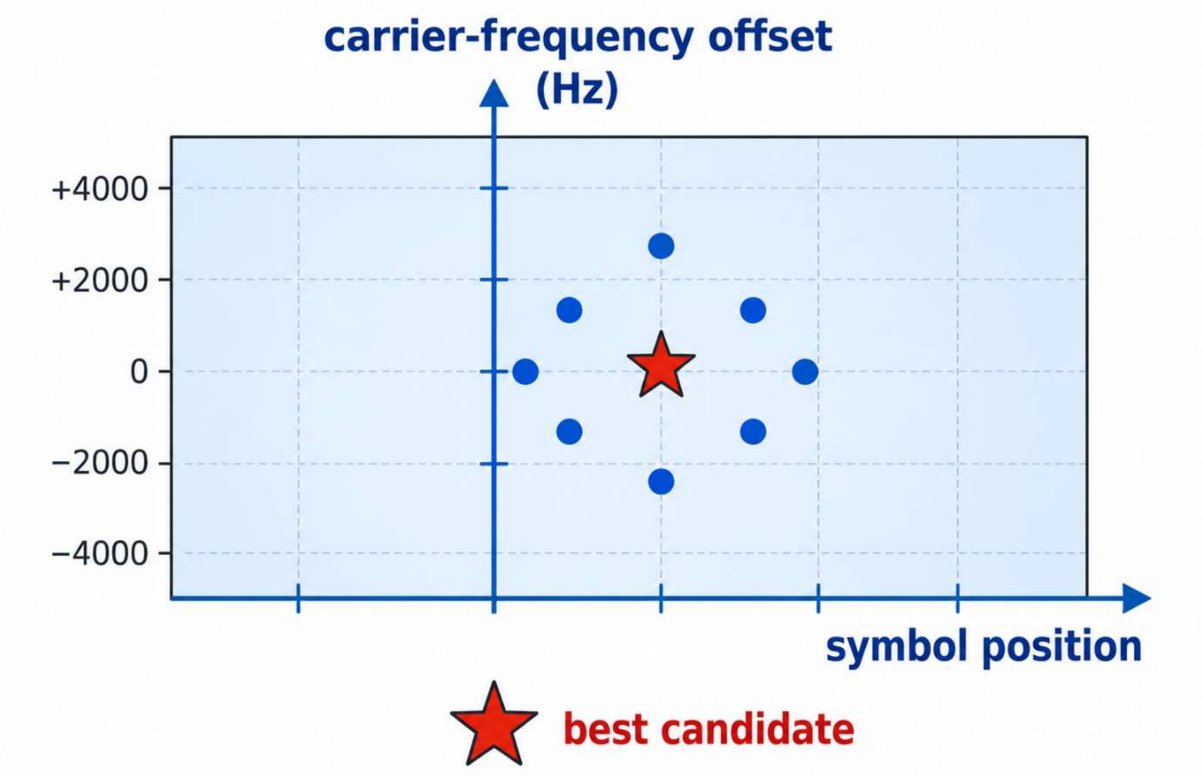

For our current application the acquisition range can, for example, be configured for approximately **±30 kHz at 333 ksym/s**.

The coarse CFO step does not need to provide the final frequency estimate. Its purpose is to retain enough coherent SOF energy to identify promising frame positions.

Normalised SOF Correlation

Simply measuring correlation amplitude is undesirable because the result would depend heavily on received signal amplitude.

Instead, the detector normalises correlation energy against received signal power.

In simplified form:

$$\text{metric} = \frac{\left|\sum r[k]\,\mathrm{conj}(s[k])\right|^2}{N \sum |r[k]|^2}$$

where N = 26 for the SOF.

This makes the detector much less dependent on absolute IQ amplitude.

Consequently, whether the average received power is 0.4, 1.0 or some other reasonable value is much less important than it would be with a raw unnormalised correlation.

This is useful in a practical SDR receiver where AGC level and signal amplitude may vary.

Candidate Selection Rather Than Immediate Detection

A particularly important design decision is that a moderately strong SOF correlation is **not immediately declared to be a DVB-S2 frame**.

Instead, the strongest SOF results are retained as candidates.

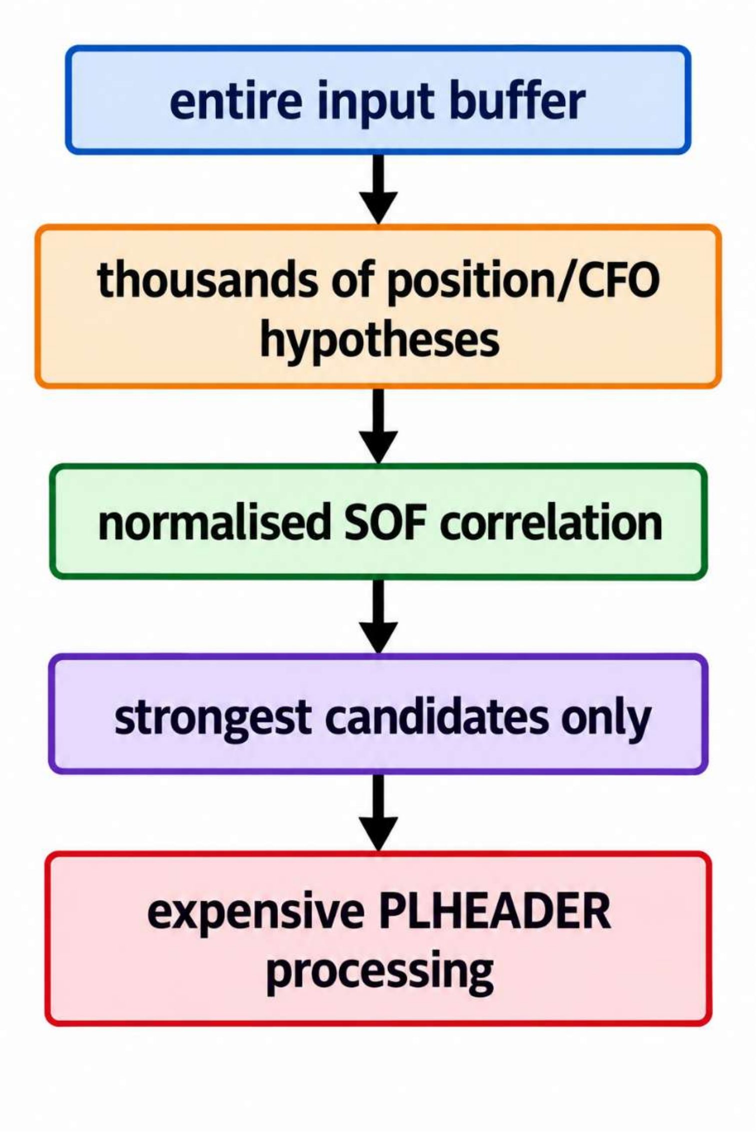

This permits the initial SOF threshold to be relatively permissive.

That matters at low SNR.

If the SOF detector itself has to make the final yes/no decision, its threshold must be high enough to suppress false detections.

But if a much stronger 90-symbol test follows it, the SOF stage only needs to avoid discarding the real frame.

This is a classic detection-system trade-of

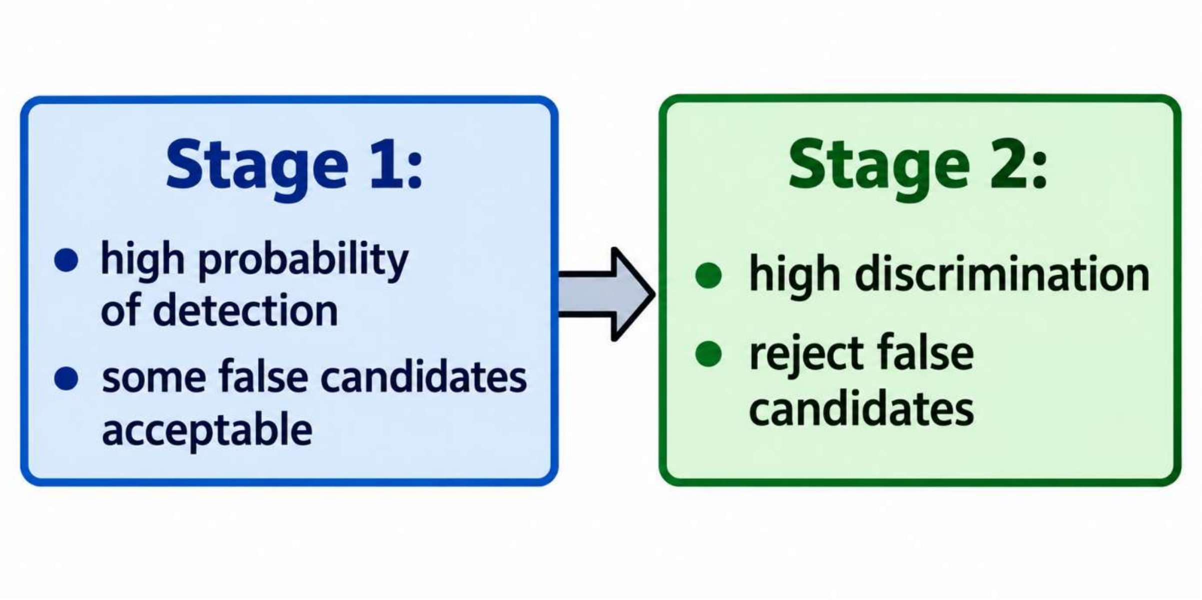

That allows the first stage to be more sensitive.

Stage 2 — Fine Carrier-Frequency Search

The coarse CFO search is deliberately not the final frequency measurement.

Once a promising SOF position has been identified, the synchronizer searches a much smaller CFO interval with finer resolution.

Instead of searching:

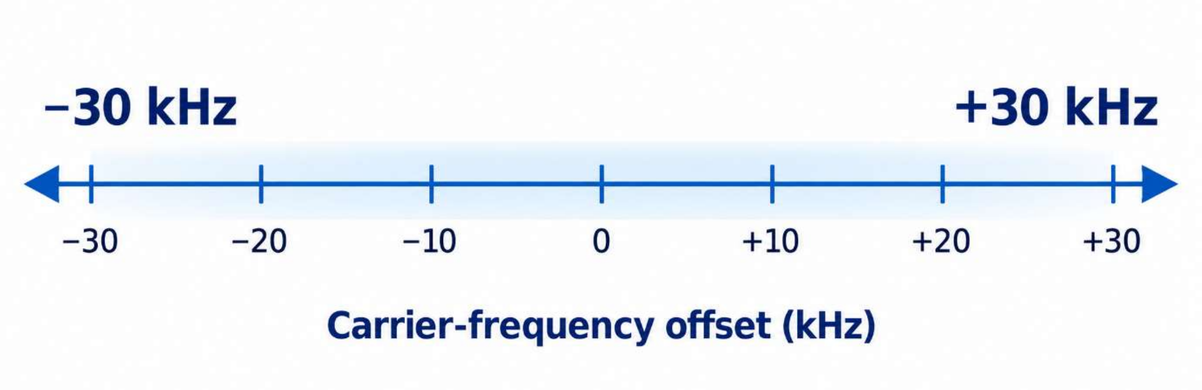

at high resolution everywhere, it can first identify a neighbourhood:

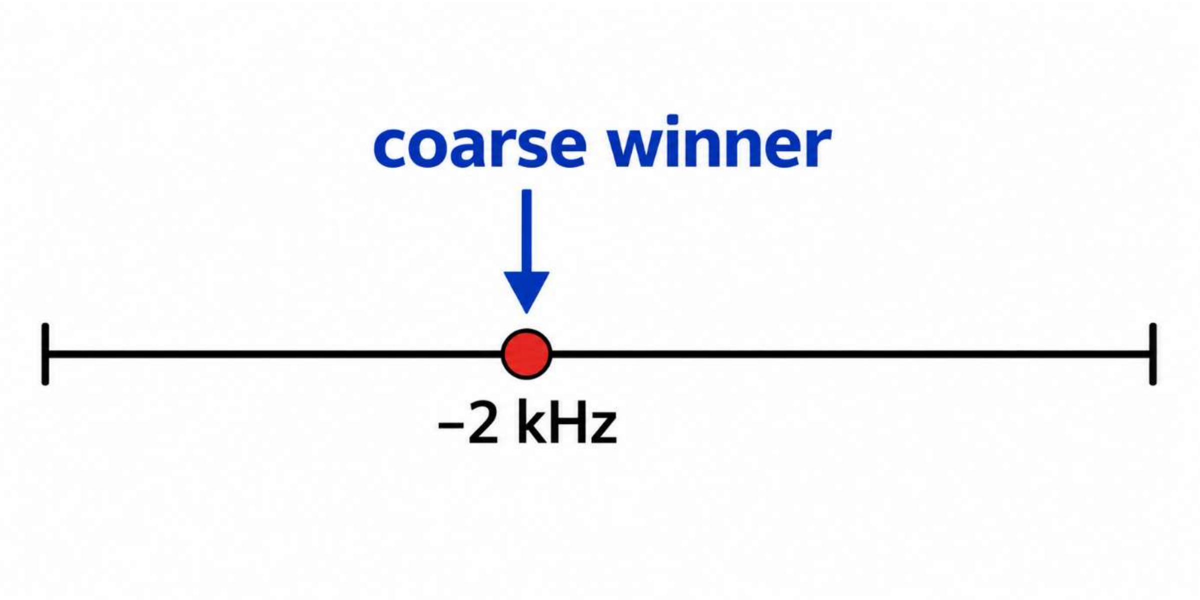

and then inspect a much narrower region:

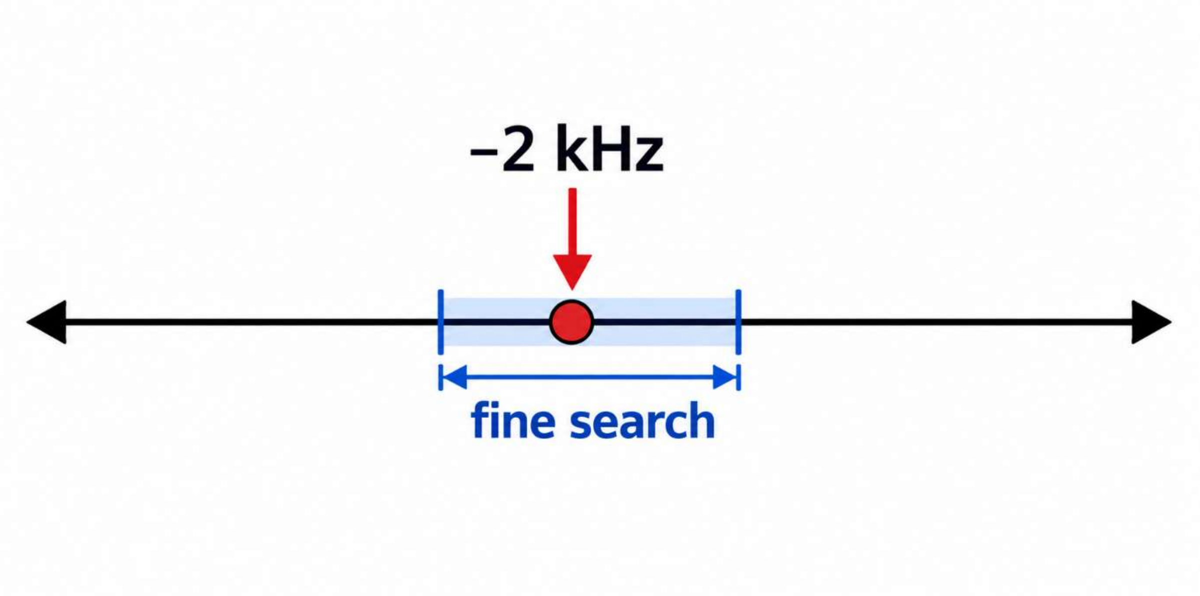

This substantially reduces computational cost.

It also separates two fundamentally different tasks:

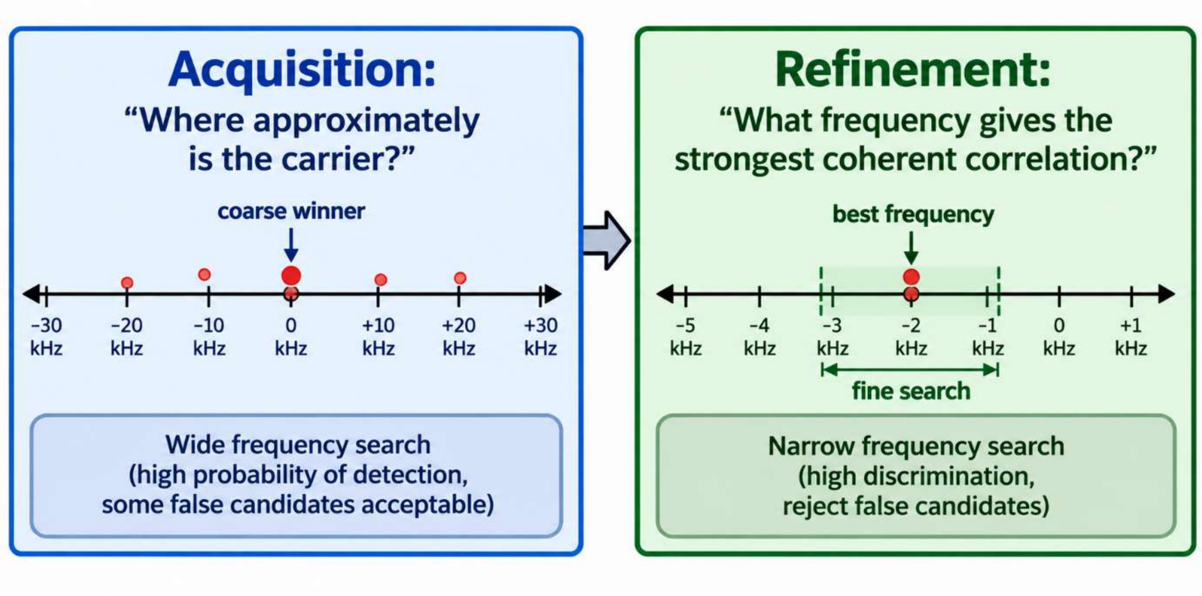

Stage 3 — Use the Entire 90-Symbol PLHEADER

This is where the new detector gains much of its potential sensitivity.

A traditional SOF detector primarily has **26 symbols** available for frame detection.

The new algorithm can eventually evaluate:

26 SOF symbols

+

64 PLS symbols

=

90 PLHEADER symbols

The PLS symbols cannot simply be treated as one universally known sequence because their contents depend on the transmitted signalling.

Instead, during initial acquisition the receiver evaluates possible PLS hypotheses.

For each hypothesis it asks:

If this were the transmitted PLS codeword, how well would the received 64-symbol sequence match it?

The strongest hypothesis becomes the candidate PLS result.

Coherent PLS Processing

After CFO correction and phase estimation from the SOF, the PLS symbols can be coherently projected according to their π/2-BPSK structure.

For alternating PLS symbols, the useful soft projection is equivalent to combinations such as

(I + Q) / √2

and

(Q - I) / √2

depending on symbol parity.

The receiver therefore retains a **soft value** rather than immediately converting each PLS symbol into a hard zero or one.

That distinction is important.

Consider two received values:

+0.95, +0.03

A hard detector sees:

1, 1

But clearly the first decision is much more reliable.

A soft detector retains that information.

When operating near the decoding threshold, preserving reliability information can make a substantial difference.

Maximum-Likelihood Interpretation

The new detector can be understood as an approximate maximum-likelihood search.

For a candidate position p, CFO f, phase ϕ, and PLS hypothesis h, we want to know how well the received samples fit the waveform predicted by those parameters.

Conceptually:

$$\Lambda = \left| \sum_{k=0}^{89} r[p+k] \, \mathrm{conj}(s_h[k]) \exp(-jk\Delta\phi) \right|$$

with suitable normalisation and handling of the SOF/PLS structure.

The detector is effectively trying to find:

(position, CFO, phase, PLS)

that best explains the received PLHEADER.

That is fundamentally different from merely asking:

"Does this 26-symbol differential pattern look like the SOF?"

The extra information is what gives the algorithm the opportunity to operate at lower SNR.

PLS Margin

Another useful measurement is not merely the strength of the winning PLS hypothesis, but how much better it is than the alternatives.

Suppose the two strongest hypotheses produce:

Best PLS correlation = 0.72

Second-best correlation = 0.31

That is persuasive.

But:

Best PLS correlation = 0.72

Second-best correlation = 0.70

is much less convincing.

The detector therefore uses the separation between the strongest and second strongest hypotheses as additional evidence.

Conceptually:

PLS margin = best hypothesis - second-best hypothesis

This is particularly valuable when noise happens to produce a high correlation for an incorrect candidate.

Joint SOF + PLS Decision

Rather than treating SOF and PLS as completely independent yes/no tests, the detector combines their evidence.

Conceptually:

jointMetric = Wsof × SOF_metric + Wpls × PLS_metric

This allows a somewhat weak SOF candidate to survive if the following PLS sequence provides extremely strong evidence.

That is precisely what we want near the sensitivity limit.

At very low SNR the real SOF may not always be the single strongest 26-symbol correlation in a large search buffer.

But a random noise peak is unlikely to be followed by a 64-symbol sequence that also convincingly matches a valid DVB-S2 PLS codeword.

So the complete header provides much stronger discrimination.

Acquisition Versus Tracking

There is another major optimisation once the first frame has been acquired.

For the first frame we do not initially know:

- position
- CFO
- PLS

so a relatively broad search is necessary.

Once the first frame has been decoded, the situation changes dramatically.

We now know:

- previous frame position
- frame length
- previous CFO
- PLS index

The next PLHEADER position can therefore be predicted:

nextPosition = previous.position + previous.frameSymbols

Instead of searching an entire buffer, the receiver only needs to inspect a small region around that prediction.

---

findFirstFrame() and findFrameNear()

This leads naturally to two different operations.

Initial acquisition

findFirstFrame()

performs the expensive global search.

Conceptually:

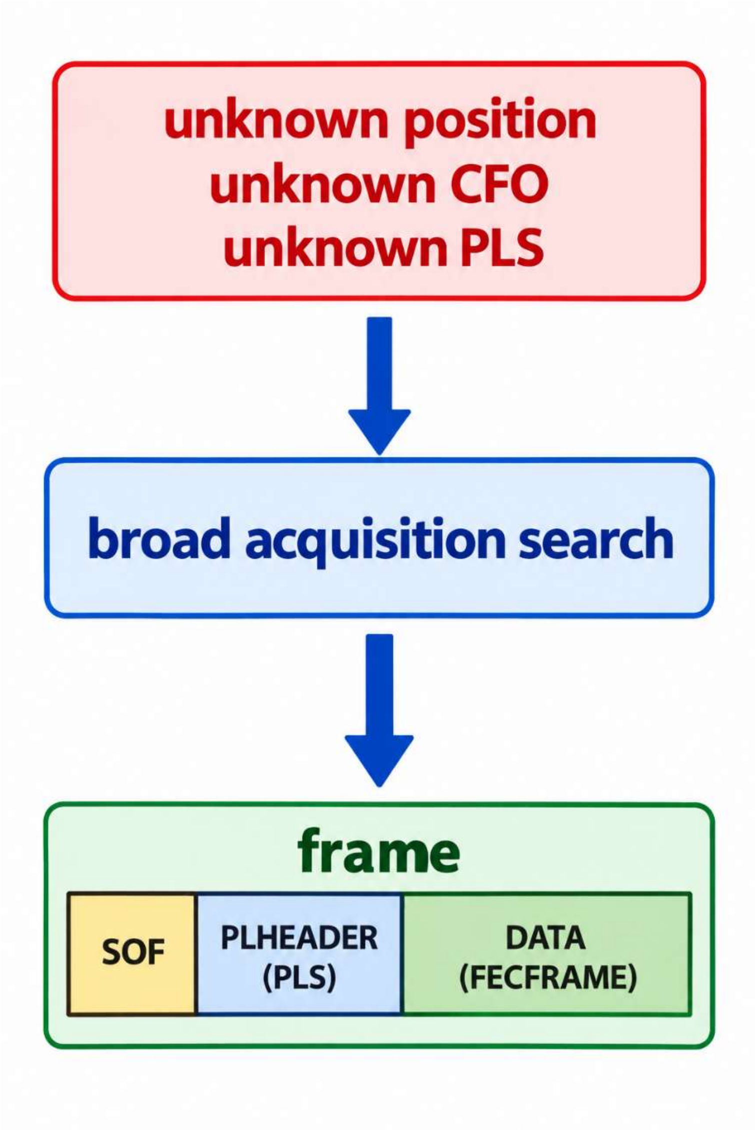

Local refinement and tracking

findFrameNear()

starts with prior knowledge:

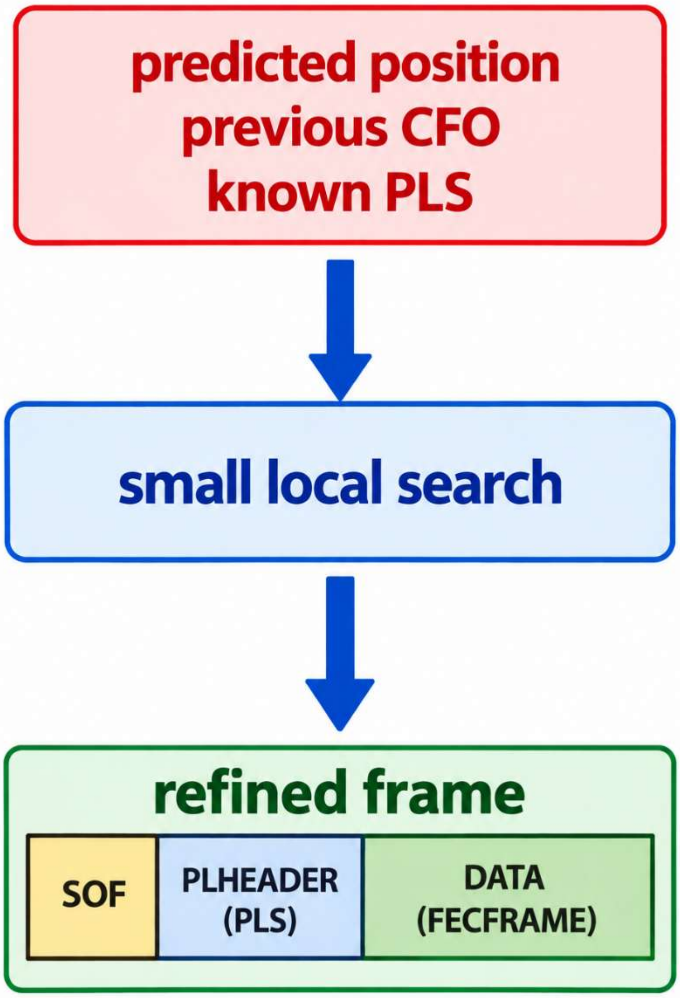

In my receiver I can therefore use:

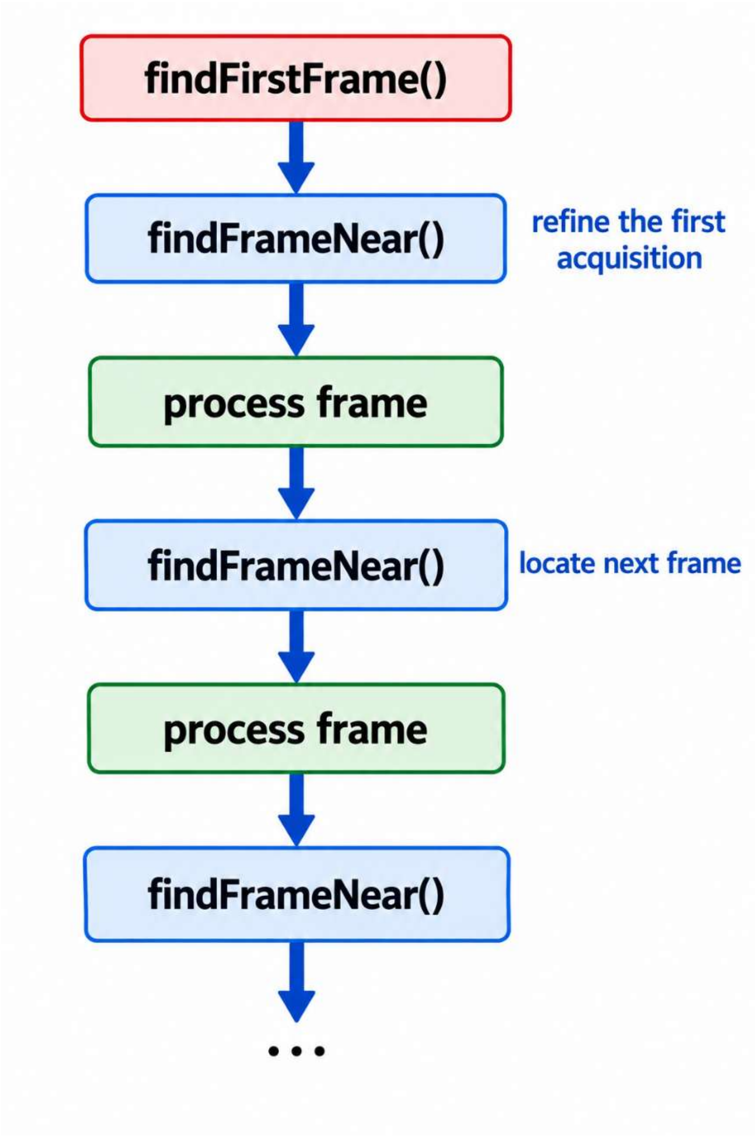

This is much more efficient than repeatedly performing complete acquisition.

Known PLS Makes Tracking More Sensitive

Once the first frame has established the PLS index, there is no reason to perform another 128-way PLS search for every following frame if the transmission parameters are expected to remain unchanged.

The receiver already knows which 64-symbol PLS sequence it expects.

Tracking therefore becomes a **known-sequence detection problem**.

That is much easier than acquisition.

Instead of:

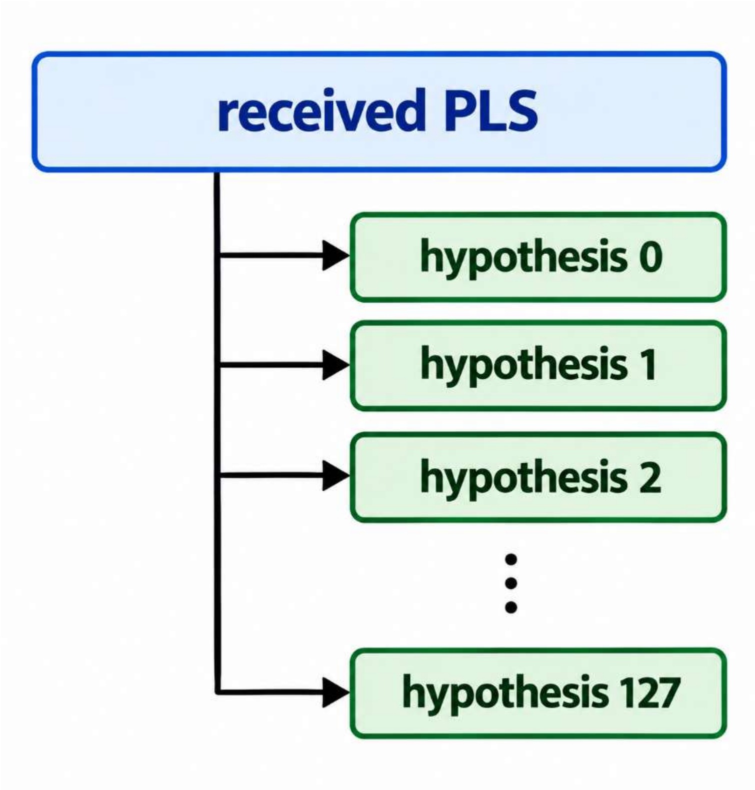
received PLS
hypothesis 0
hypothesis 1
hypothesis 2
hypothesis 127

we have:

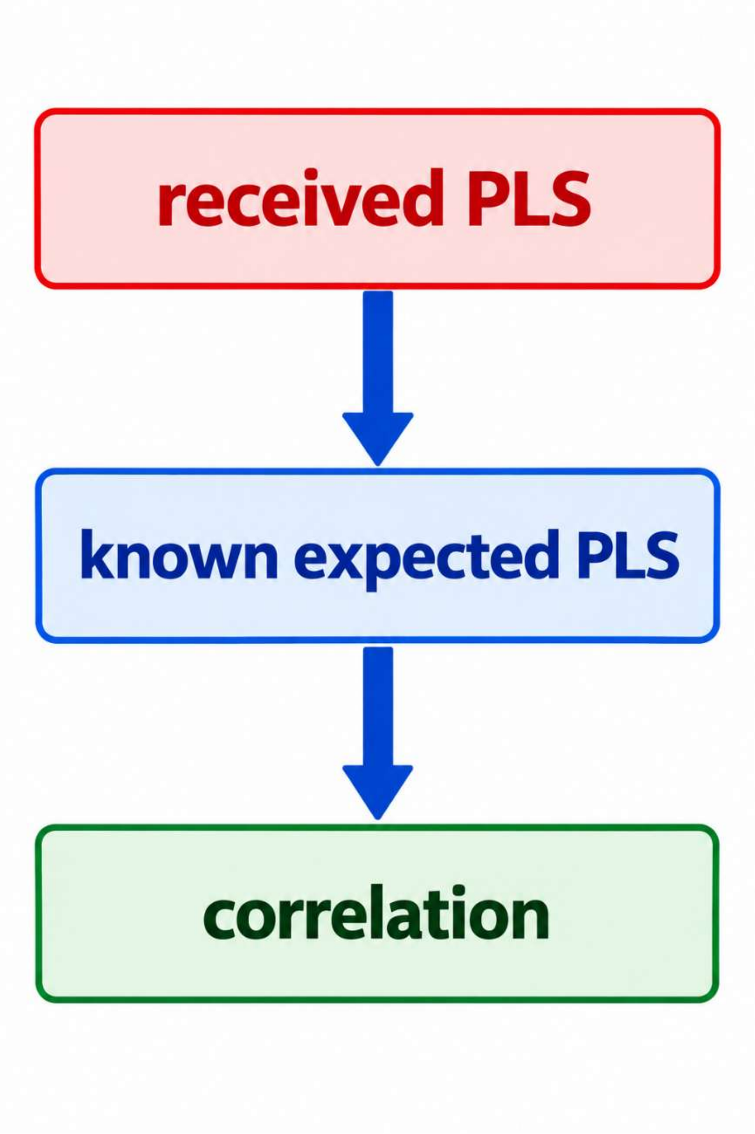

This removes an enormous amount of ambiguity.

It also explains why a well-designed DVB-S2 receiver should distinguish between **acquisition sensitivity** and **tracking sensitivity**.

Tracking should normally work below the SNR at which completely blind acquisition becomes reliable.

Improving the CFO Estimate During Tracking

The complete known PLHEADER provides another useful opportunity.

Once the PLS is known, all 90 PLHEADER symbols are known.

Let the received symbol be

r[k]

and the expected PLHEADER symbol be

$s[k]$.

Remove the known modulation:

$z[k] = r[k]\, \mathrm{conj}(s[k])$

Ideally this leaves:

$z[k] \approx A \exp(j(\phi + k\Delta\phi))$

Now calculate adjacent differential products:

$z[k]\, \mathrm{conj}(z[k-1])$

The unknown common phase cancels, leaving approximately:

$A^2 \exp(j\Delta\phi)$

Accumulating this over the PLHEADER gives an estimate of the phase rotation per symbol:

$\Delta\phi = \mathrm{atan2}(\Sigma\ \mathrm{imag}, \Sigma\ \mathrm{real})$

and therefore:

$$\mathrm{CFO} = \frac{\Delta\phi\, R_s}{2\pi}$$

This is an attractive tracking estimator because it uses **89 differential phase measurements across the known 90-symbol PLHEADER**.

It also provides a continuous frequency estimate rather than restricting the answer to a coarse frequency-search grid.

Why Not Use the Differential Method Everywhere?

There is an interesting apparent contradiction here.

I have argued that differential detection loses sensitivity, yet differential phase is useful for estimating CFO.

The distinction is what we are asking it to do.

Using differential detection as the **primary frame detector** throws away information that coherent integration could otherwise exploit.

Using differential phase **after the header has already been identified and its modulation is known** is a different problem.

We can therefore combine the strengths of both approaches:

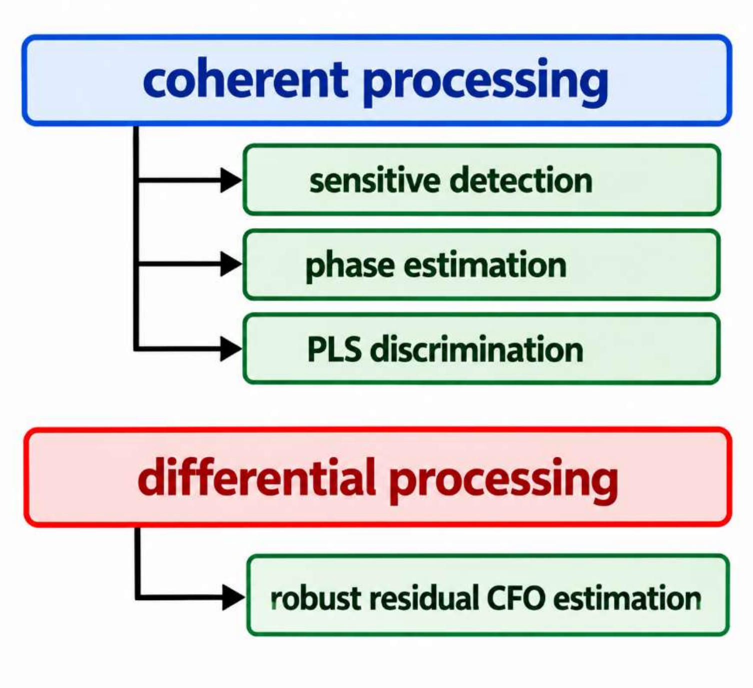

There is no requirement that one technique be used exclusively.

A good receiver can use whichever representation is best suited to each estimation problem.

Computational Cost

The obvious disadvantage of the new detector is computational complexity.

A simple differential correlator is inexpensive.

A detector which searches:

position × CFO × PLS hypothesis

can become extremely expensive.

The solution is to avoid doing the expensive work everywhere.

The implementation therefore uses a hierarchy:

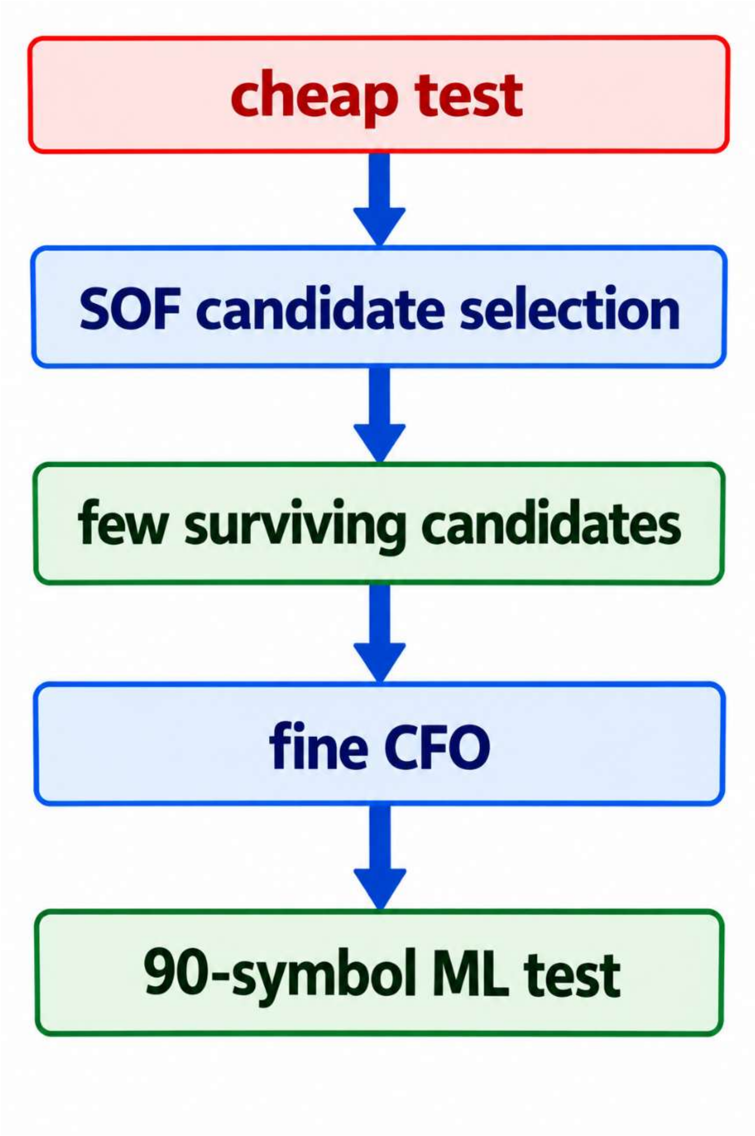

AVX2 is also useful during the broad SOF search.

Eight candidate positions can be evaluated in parallel, making the large initial scan substantially cheaper.

This is a good example of an important optimisation principle:

Optimise the part of the algorithm that is executed millions of times, not the part executed a handful of times.

The full PLS test is comparatively expensive, but it is performed only for a small number of promising candidates.

Traditional Detector Versus New Detector

The main differences can be summarised as follows:

Property	Traditional differential detector	New PLHEADER detector
Primary acquisition sequence	SOF/PLS differential pattern	SOF followed by full PLHEADER
Initial SOF length	26 symbols	26 symbols
Final evidence available	Primarily differential header correlation	Up to all 90 PLHEADER symbols
Unknown carrier phase	Naturally cancelled	Explicitly estimated
CFO handling	Differential estimate	Coarse search + refinement
PLS processing	Differential/hard-oriented	Coherent soft correlation
PLS alternatives	Code detection	ML-style hypothesis comparison
Reliability information	Limited	Soft metrics retained
False-candidate rejection	Header correlation	SOF + PLS + PLS margin
Tracking	Local differential detection	Known-PLS coherent tracking
SIMD optimisation	Possible	AVX2 coarse SOF search
Computational cost	Low	Higher
Expected low-SNR sensitivity	Good	Potentially significantly better

The final line is deliberately phrased as **expected** sensitivity.

Theoretical arguments tell us why coherent use of the complete PLHEADER should be better, but the real engineering question is:

How many decibels better is it on an actual DVB-S2 signal?

That has to be measured.

What Should Be Tested?

The most useful comparison is not simply whether each detector can find one particular frame.

A proper test should repeatedly generate or receive frames at progressively lower SNR and measure detection statistics.

For example:

SNR	Traditional	New detector
+3 dB	?	?
0 dB	?	?
-1 dB	?	?
-2 dB	?	?
-3 dB	?	?
-4 dB	?	?
-5 dB	?	?

For every SNR, I would measure at least:

- correct frame detections
- missed frames
- false detections
- position errors
- CFO estimation error
- PLS decoding errors

Ideally thousands of frames should be tested at each SNR.

That produces a much more meaningful result than saying:

"It seemed to work at −3 dB."

The useful result is something such as:

At SNR = X:

Traditional detector:

91.2% detection

0.3% false detection

New detector:

99.7% detection

0.02% false detection

The precise numbers remain to be established experimentally.

An Especially Interesting Experiment

One experiment should separate **initial acquisition** from **locked tracking**.

These are fundamentally different operating conditions.

I would therefore measure two curves:

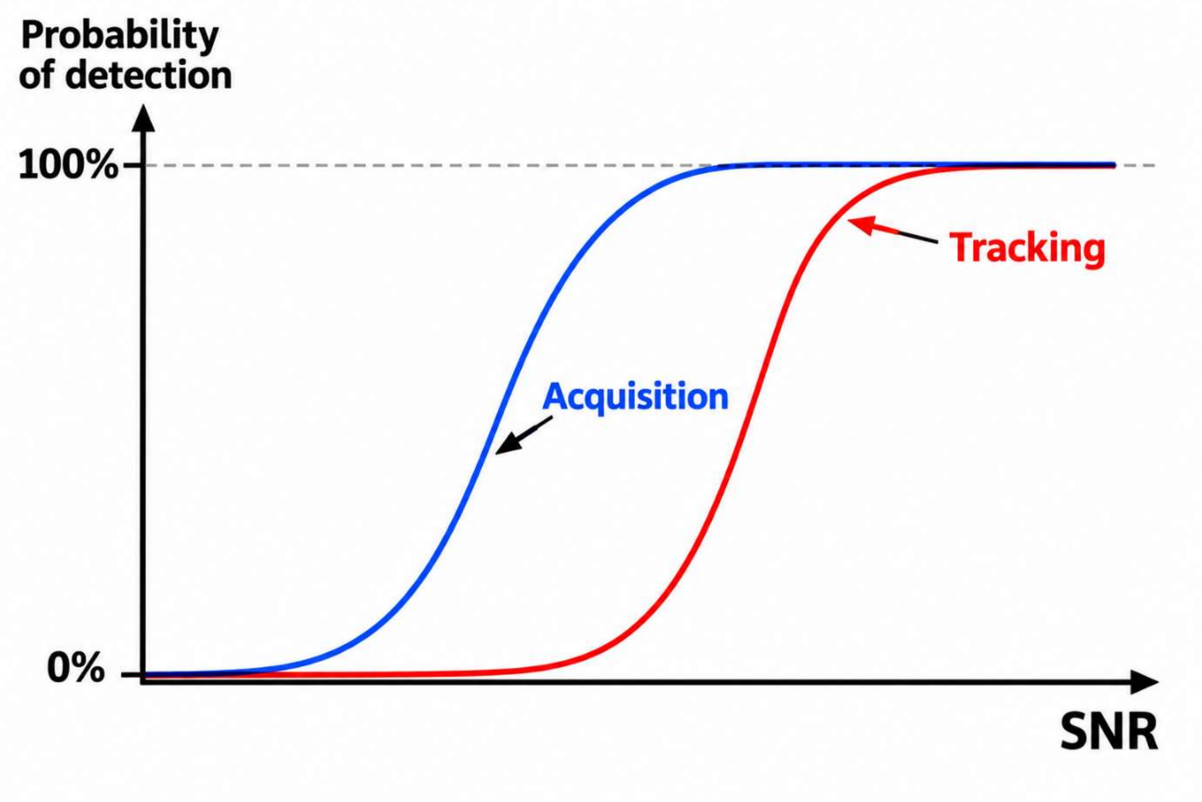

Once position, frame length and PLS are known, the receiver has far more information available.

I therefore expect the tracking detector to remain reliable below the point where a completely blind findFirstFrame() acquisition becomes unreliable.

That difference is itself useful information when designing a practical receiver.

Why This Matters

The DVB-S2 LDPC decoder can operate at remarkably low SNR, particularly for the lower code rates.

There is little benefit in having an extremely powerful FEC decoder if the receiver cannot reliably find the physical-layer frame at the same SNR.

The complete receiver chain is only as sensitive as its weakest acquisition stage:

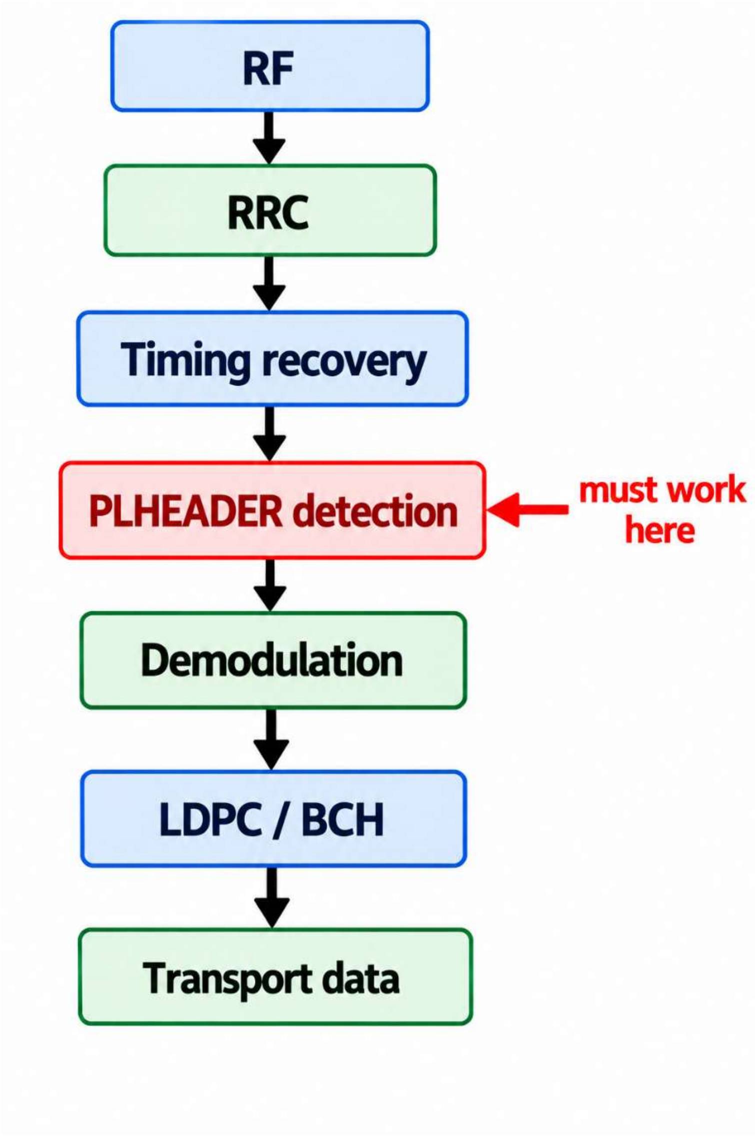

Improving PLHEADER detection can therefore translate directly into usable receiver sensitivity.

Conclusion

The traditional DVB-S2 differential PLHEADER detector is an excellent engineering solution.

It is simple, fast and naturally robust against unknown carrier phase and moderate frequency error.

Its weakness appears when we push the receiver towards the lowest possible SNR.

Differential processing inevitably sacrifices some information, while the DVB-S2 PLHEADER actually provides **90 symbols of highly structured information** that can be exploited.

The new synchronizer therefore takes a different approach:

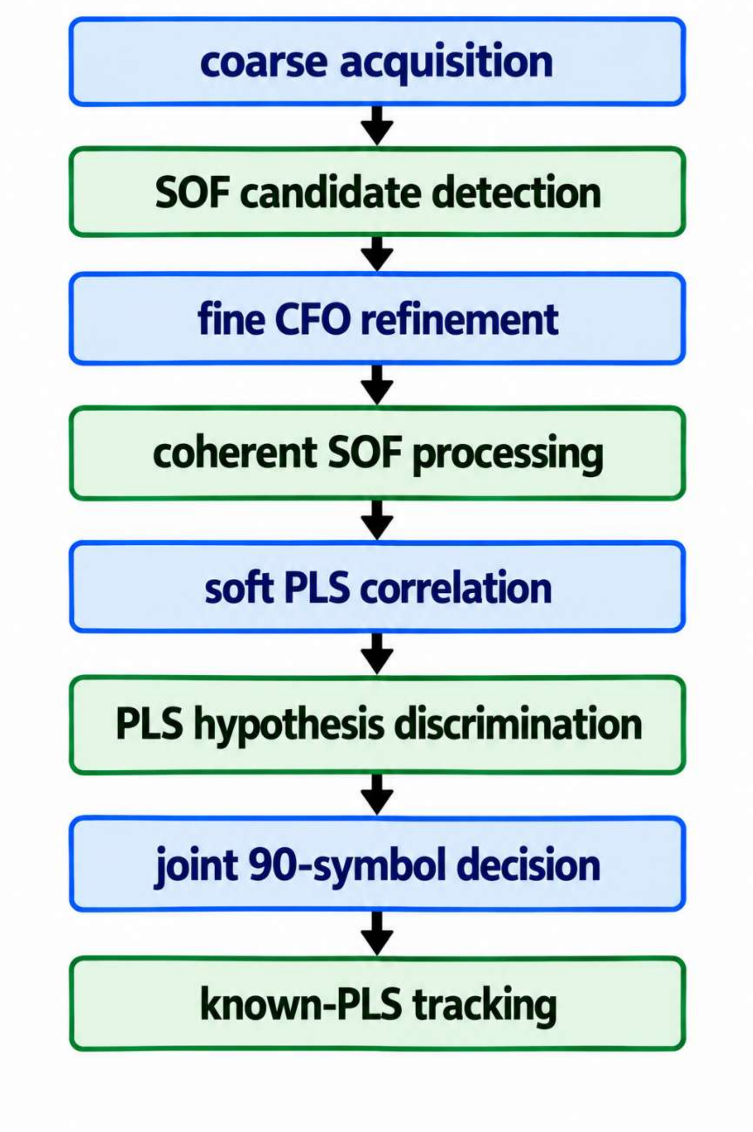

The objective is not to replace every differential operation with something more complicated.

It is to use each technique where it is strongest:

Differential processing for robustness and frequency estimation; coherent processing for sensitivity and discrimination.

The next step is the most interesting one: measuring the two algorithms against exactly the same DVB-S2 signal while progressively reducing the SNR.

That will tell us how much additional sensitivity the extra complexity actually buys.

And in SDR development, that measured result is ultimately what matters.