



INDUSTRY INSIGHTS

Silvio Ruiu | Mechanical Engineer | Independent Consultant, www.cmblasterus.com — Italy

Building the Procedure

INTRODUCTION

The two cases in the first part describe processes that work once the procedure is already established. They do not say how you get there.

That is the part worth spelling out, because the procedure is where the whole benefit is won or lost, and because none of it is guesswork. What follows is the sequence I use.

START FROM THE PART, NOT FROM THE MACHINE

The first step is to look at the component and to ask the manufacturer where it typically breaks.

A part is not uniform. It is not loaded uniformly, and it has intrinsic weak points. There is no such thing as peening the part — there is peening the places that matter. Where the fatigue crack initiates determines where you will need intensity, where you will measure it, and where coverage actually has to be verified.

The manufacturer already has that information. Broken parts come back.

LOOK UP WHAT HAS BEEN DONE BEFORE

The second step is literature and archives: tests run on comparable parts, matched as far as possible on material and surface hardness, to arrive at a starting intensity that stands a reasonable chance of producing a usable result.

This step exists because of what comes later. Fatigue testing takes the time it takes. Starting from a plausible intensity rather than from nothing is how you avoid spending test campaigns on values that never had a chance. If there is no literature, prepare extra patience instead.

BE HONEST ABOUT VOLUMES

The third step is an honest assessment of volume. How many parts per hour do we actually have to process?

Nothing has been built yet, but that number decides the machine and, more importantly, the layout. The procedure will have to fit inside an existing production process, so what surrounds it counts as much as the machine itself.

My own preference, for peening, is to work in batches. A

batch gives far better control over how the process evolves, and therefore over its stability.

The reflex for high volume is a pass-thru machine; this is large, it is expensive, and everything has to happen in a single pass — which means nothing is allowed to go wrong. Two batch machines, with a small buffer at the infeed and another at the discharge, can deliver the same throughput far more stably. If one of them stops, or if the amperage and the strips say it is no longer performing, the other keeps running. You are down to half. Half is not stopped.

The same argument applies to correction. On a pass-thru, the only real way to recover intensity and coverage is to slow the pass-thru speed — and slowing it slows the whole line, because the machine is the line. With two batch machines and properly sized buffers, you correct one machine and, for a short period, you have halved capacity rather than lost it.

None of this is an argument against continuous flow. Automation has largely dissolved the conflict: a layout can present itself as a pass-thru — continuous infeed, continuous discharge, nobody handling anything — while internally running in batches. There are systems where loading, unloading and even the reading of Almen strips are handled by robots. Continuous throughput, discrete process. Both are available.

POSITION THE STRIPS AGAINST THE REAL THROW

Once the fragile points are identified, the strips have to be placed where those points sit relative to the throw of the machine you are using for the test. Keep that in mind, because it will matter when the process scales to the final layout.

Not in a conventional position. In the position that maps the throw onto that geometry. This is why one installation ends up with six strips, three per side, and another with four on a control tray. Those are not numbers — they are the map of where the media lands on that part in that machine.

These are also the positions you will be reading from for as long as the process runs. Get them wrong now and the work has to be redone later to have the reference values you need.

MEDIA IS THREE VARIABLES, NOT ONE

Media today is a varied and complex world, and there are three levers, not one: size, material — meaning specific weight — and the hardness of the media itself. All three converge on a single readable number, the Almen intensity.

Which is why enlarging the shot is the lazy way to go faster. It is one lever out of three, and it is the one that does the most damage to surface morphology.

I remain of the view that you should use the smallest media that will reach the required intensity. Throwing more kilos of fine media beats throwing fewer kilos of coarse media: the single impact does less morphological harm. Fine media typically costs more. The result is on another level in terms of coverage and consistency.

One condition attaches to this, and it is the real constraint. It requires a machine capable of running fine media properly. That is less common than it looks.

SAMPLE PROPERLY

With a starting target from the literature, prepare a small sampling. Not an enormous number of samples, but more than one, and drawn from different production lots. If there is more than one manufacturing site, take them from different sites. The purpose is to flatten out the statistics and the peculiarities each site carries.

This is the same problem the aerospace framework was built to solve — the same part behaving differently in different places. Here it is not normed. It is sampled.

PEEN, TEST, AND CORRECT IN BOTH DIRECTIONS

Peen the samples the way the process is meant to peen them, and send them to test. Then verify whether the improvement is what you expected, and correct. Higher if it falls short. But also lower: if you have tripled the fatigue life and doubling is enough, you reduce, and what you buy back is a faster process. The target is what the part needs, not the maximum obtainable.

ONE MORE THING ON CYCLE TIME

A note on what I mean by time here. I am talking about batch machines, where there is a cycle time rather than a pass-thru speed. As a very general order of magnitude, cycle times tend to sit somewhere around four to six minutes under ordinary conditions — and that figure is already 2T, the doubled time. With parameters this generic it is indicative, and there are cases that call for considerably more, or less.

Once the part is in the machine, the difference between

running the time T needed for full coverage and running 2T is marginal in practice. The part is already in there, already being worked.

What doubling the time buys is two things at once. Saturation is defined as the point where doubling the exposure no longer raises intensity by more than 10%, so running at 2T puts you past that point by definition. And coverage, which under SAE J2277 is extended beyond 100% precisely by extending exposure time, is at 200%. In industrial practice that is the reason for doing it: one decision, margin on both.

The cabin still has to be set up: parts loaded, parts unloaded. However good the handling systems are, that is time with the machine standing still. Extend the peening time and quite often it is still the standing time that exceeds the peening time, not the other way round.

The productivity lost to a doubled cycle is not nothing. But as a percentage it weighs less than the uncertainty of not doubling it.

Between literature and reasoning, there is a target value to aim at even at the start of a new process study. That is what changes the arithmetic. Say the sampling is three parts, enough to see whether the improvement is real: I can peen those three at the target intensity, using the parameters I worked out on the strips, and send them to fatigue testing — which takes as long as it takes.

Do it the other way and it becomes three samples at one minute, three more at two, three more at three. Nine tests instead of three.

The difference is that a target starting intensity is normally decided beforehand. Sailing towards the unknown is fine, but let us at least take a compass along.

SAME INSTRUMENT, BOTH PHASES

Using strips while building the procedure is what allows you to use strips to control it afterwards. Same instrument, both phases — which is what makes the control comparable to the construction, instead of being a different measurement of a different thing.

Skip that, and you are left rebuilding the reference values later, on a process that is already running. ●

ABOUT THE AUTHOR:

Silvio Ruiu is a surface engineering consultant based in Modena, Italy, with more than ten years of field experience. He works on wheel blast and shot peening installations, focusing on process setup, equipment efficiency and diagnostics.