BEAM-BASED CHARACTERISATION OF BPM ELECTRONICS THERMAL SENSITIVITY AFTER TWO DECADES OF OPERATION*

D. J. Scott †, SLAC, Menlo Park, CA, USA

Abstract

Beam-based measurements of the thermal sensitivity of beam position monitor (BPM) electronics were performed in SPEAR a third-generation storage ring after more than twenty years of routine user operations. Controlled building-temperature excursions were applied to two equipment buildings containing key BPM front-end and digitiser racks. Using orbit- and charge-normalised BPM signals and independent temperature logging, we performed lag-aware regression to estimate effective position-versus-temperature coefficients for each BPM, and compared several alternative temperature-driver models (global, per-building and hybrid). The method was applied to two measurement campaigns, months apart, with different ambient conditions. The results show clearly distinguishable building-level responses and reproducible patterns within subsets of BPMs, but also highlight strong correlations between temperature, beam conditions and slowly varying lattice effects.

INTRODUCTION

SPEAR BPM pickup buttons are installed in the accelerator tunnel, while the associated front-end and digitiser electronics are housed in surface equipment buildings, mainly Building 132 and Building 116. Building 132 was deliberately temperature-cycled using the HVAC system. Building 116 was not deliberately cycled and provides a comparison population. In this paper, thermal sensitivity is used in a beam-based sense, denoting an effective fitted coefficient between an electronics-building temperature driver and the BPM orbit in µm/°C. It should not be interpreted as a unique intrinsic electronics constant for each BPM, since it depends on the selected temperature driver, fitted lag, beam conditions and any remaining correlated orbit motion. Two fast orbit feedback (FOFB) off measurements are used for the primary analysis. Experiment 1 is the clean controlled measurement. Experiment 2 is a repeat measurement under different ambient and machine conditions, but with stronger coherent beam motion.

MEASUREMENTS

For each measurement, rack-temperature readbacks from the BPM electronics buildings were reduced to mean-centred temperature drivers (shown in Fig. 1). Experiment 1 gave the cleanest Building 132 excitation and the lowest apparent background structure. Experiment 2 repeated the measurement under different ambient and machine conditions, but showed larger coherent orbit structure.

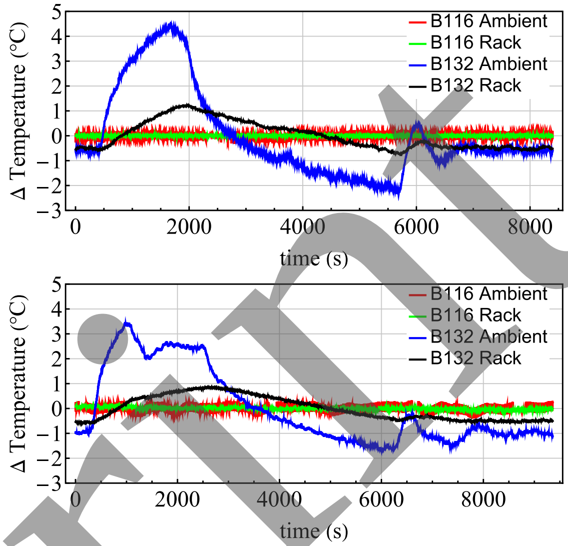

Figure 1: Mean-centred rack-temperature driver for Experiments 1 (top) and 2 (bottom).

ANALYSIS METHOD

For each BPM and plane, lag scans were used to allow for delayed electronics response to the building-temperature change. The reported coefficient was obtained from a current-adjusted linear model:

$$x_i(t) = a_i + \beta_{T,i} T_b(t - \tau_b) + \beta_{I,i} I(t) + \epsilon_i(t)$$

where x_i is the measured BPM position for BPM i in the selected plane, T_b is the rack-temperature driver for the relevant electronics building, I is the stored beam current, and τ_b is the fitted building-temperature lag. The reported coefficient is $\beta_{T,i}$, in µm/°C. The current term reduces false attribution of current-dependent BPM response to temperature. The regression was tested using both ordinary least squares (OLS) and robust Huber fitting. Huber fitting reduces the influence of transient outliers and poorly modelled points, and was used as a robustness check. The Huber fits gave the same qualitative building-level conclusions as OLS. OLS is used for the reported coefficients because the fitted slopes, residuals, confidence intervals and R^2 values are simpler to interpret. The R^2 value is the fraction of variance in the fitted BPM position trace explained by the selected regression model. It is used as a fit-quality indicator, not as proof that all explained motion is purely thermal. Error bars on individual BPM coefficient plots show the 95% confidence interval for the fitted temperature coefficient. They do not include model dependence or repeatability between measurements.

---

* Work supported by DoE Contract DE-AC02-76SF00515

† dscott@slac.stanford.edu

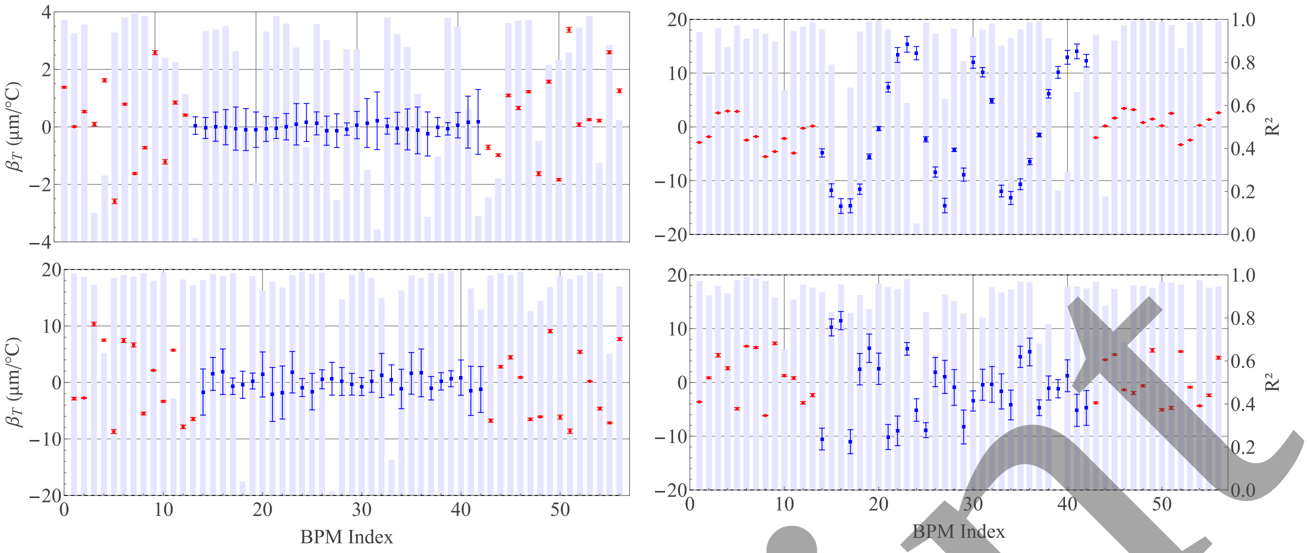

Figure 2: Experiment 1 horizontal temperature-response coefficients. Building 132 BPMs show the larger fitted response, while Building 116 BPMs remain close to zero. Error bars show fit uncertainty; pale bars show R².

RESULTS

The clearest result is obtained from Experiment 1. In the horizontal plane, the fitted coefficients separate strongly by electronics building, as shown in Fig. 2: the deliberately cycled Building 132 BPMs show the larger response, while the non-cycled Building 116 BPMs remain close to zero. This building-level separation is the strongest evidence that the fitted response is linked to the applied electronics-temperature change rather than only to global orbit drift. The coefficient summary is given in Table 1.

Table 1: Median Fitted Thermal-Sensitivity Summary by Experiment, Plane and Electronics Building

Exp.	Plane	Building	β_T μm/°C
1	X	116	1.03
1	X	132	6.15
1	Y	116	0.08
1	Y	132	1.09
2	X	116	4.75
2	X	132	4.22
2	Y	116	10.67
2	Y	132	2.45

For Experiment 1, the Building 132 horizontal median absolute coefficient was 6.15 μm/°C, compared with 1.03 μm/°C for Building 116. The vertical response was weaker, with median absolute coefficients of 1.09 μm/°C for Building 132 and 0.08 μm/°C for Building 116. Experiment 2 provides a repeat measurement under different ambient and machine conditions. The Building 132 horizontal response remained at the few-μm/°C scale, with a median absolute coefficient of 4.22 μm/°C. However, Experiment 2 also contained stronger background structure. In particular, the Building 116 vertical distribution was much larger than in Experiment 1. Since Building 116 was kept at nominal, this is not interpreted as a clean Building 116 electronics thermal response. It indicates that slow beam, lattice, current-dependent or environmental terms can remain correlated with the temperature drivers in short beam-based measurements.

RESIDUAL MODES AND MODEL LIMITS

The two FOFB-off measurements were not equally clean. Experiment 1 produced the clearest driven thermal excitation and the strongest building-level separation. Experiment 2 repeated the Building 132 horizontal response scale, but also showed larger comparison-building structure, especially in the vertical plane. This difference is visible in the PCA explained-variance comparison in Fig. 3, where the Experiment 2 residual data retain stronger coherent structure after the current-adjusted fit.

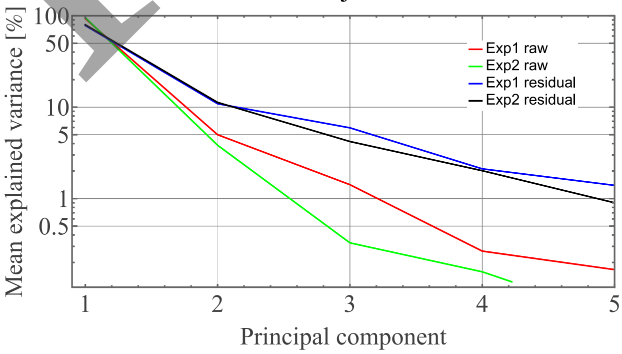

Figure 3: PCA comparison for Experiments 1 and 2, before and after removal of the fitted thermal component. Experiment 1 shows a cleaner reduction of the dominant coherent mode, while Experiment 2 retains stronger coherent beam motion.

We used PCA as a diagnostic check on the vertical BPM data, not as a replacement for the regression model. The purpose was to test whether the fitted β_T pattern looked like a local electronics response, or whether it was aligned with a coherent orbit pattern already present in the raw BPM data. In the PCA, each mode is described by a time-dependent amplitude and a spatial pattern across the BPMs.

The spatial pattern identifies which BPMs participate in the mode and with what sign. A beam or orbit mode would therefore be expected to appear as a coherent pattern over many BPMs, rather than as isolated excursions at a few devices. For the Experiment 2 vertical data, the mode-correlation check in Fig. 4 shows that PC2 is the PCA mode most strongly associated with the fitted β_T pattern. This means that the fitted vertical temperature-coefficient pattern has a similar BPM-by-BPM shape to a coherent spatial mode in the raw vertical BPM data. We then made a direct test by isolating only the PC2 contribution to the BPM data. This PC2-only signal is formed from the product of the PC2 time score and the PC2 spatial loading pattern. We fitted this reconstructed PC2-only BPM signal with the same current-adjusted temperature model used for the reported coefficients. The resulting PC2-implied β_T pattern closely reproduced the observed Experiment 2 vertical β_T pattern, as shown in Fig. 5.

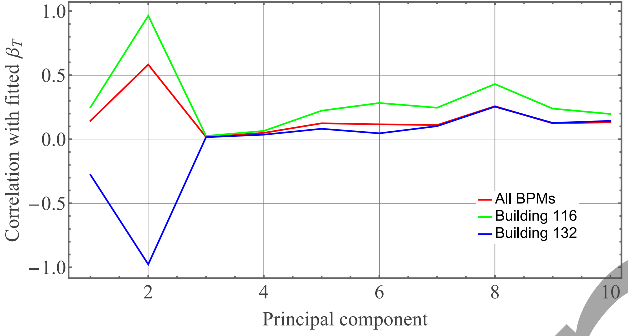

Figure 4: Experiment 2 vertical PCA diagnostic. PC2 shows the strongest building-dependent correlation with the fitted βT pattern, consistent with coherent beam/orbit contamination.

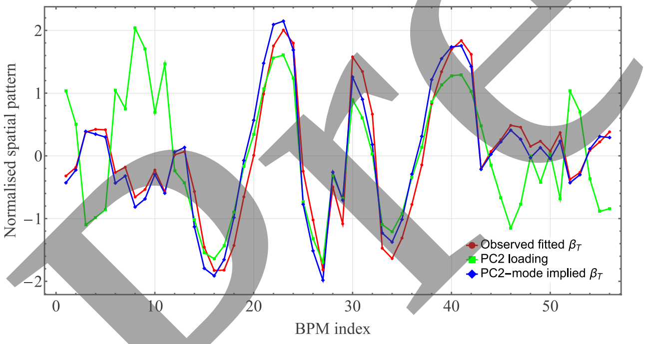

Figure 5: Experiment 2 vertical mode diagnostic. The observed fitted βT pattern is compared with the raw PC2 loading and the βT pattern implied by fitting the isolated PC2 contribution.

This shows that much of the anomalous fitted vertical response can be generated by the coherent PC2 mode alone. This indicates that the Experiment 2 vertical coefficients contain a coherent beam/orbit mode correlated with the temperature drivers. They should not be interpreted as a clean Building 116 electronics thermal response. The coefficients reported here are therefore effective beam-based response estimates. They resolve the scale and building-level pattern of the temperature-linked response in the cleanest measurement, but they do not uniquely separate BPM electronics drift from all correlated slow beam, lattice or environmental contributions.

DISCUSSION

The most robust result is the building-level separation in Experiment 1. The deliberately cycled Building 132 BPMs show the largest fitted response, especially in the horizontal plane, while the non-cycled Building 116 population remains close to zero. This supports a real temperature-linked BPM/electronics response at the population level. Experiment 2 confirms that the Building 132 horizontal response remains at the few-µm/°C scale, but also shows that short beam-based measurements can retain coherent orbit structure correlated with the temperature drivers. The Experiment 2 vertical coefficients, especially for Building 116, should therefore be treated as effective fitted coefficients rather than clean electronics sensitivities.

CONCLUSION

Controlled temperature cycling of the SPEAR BPM electronics environment produced a measurable beam-based response after more than two decades of operation. The clearest result is the Experiment 1 FOFB-off horizontal measurement: Building 132 was deliberately cycled and showed a larger fitted response, while the non-cycled Building 116 population remained close to zero. For Building 132 X, the median absolute coefficient was 6.15 µm/°C in Experiment 1 and 4.22 µm/°C in Experiment 2. The vertical response was weaker and less uniform. Experiment 2 also contained larger coherent background structure. A PCA diagnostic showed that the anomalous Experiment 2 vertical coefficient pattern was strongly associated with a raw-data vertical PCA mode. This supports the interpretation that this part of the result contains beam/orbit contamination rather than a clean electronics thermal response. The coefficients should therefore be treated as effective beam-based thermal-response estimates, not immutable per-BPM electronics constants. Future work should use longer datasets, repeated controlled excursions, explicit orbit-preparation checks and additional environmental signals to better separate BPM electronics response from other slow machine effects. The state of the FOFB could should also be checked for the 2$^{\text{nd}}$ experimental time.