EMO 2013

A Real-World Application of a Many-Objective Optimisation Complexity Reduction Process

Robert J. Lygoe, Mark Cary, and Peter J. Fleming

22-March-2013

Contents

Introduction

- Stricter exhaust emissions legislation => minimising emissions during start-up before catalyst reaches operating temperature.
- Gasoline Direct Injection controls - many degrees of freedom - significant time/effort needed to develop a quick, stable and reliable engine start and run-up whilst minimising fuel and HC emissions.
- Start profile should be robust to background variation/noise – can be formulated as noise sensitivity objectives which may conflict with performance objectives – increases dimensionality of optimisation.
- Extends complexity reduction strategy in previous paper (Lygoe et al, 2010) with several enhancements to a high-dimensional multi-objective optimisation: 10 objectives + 1 constraint.

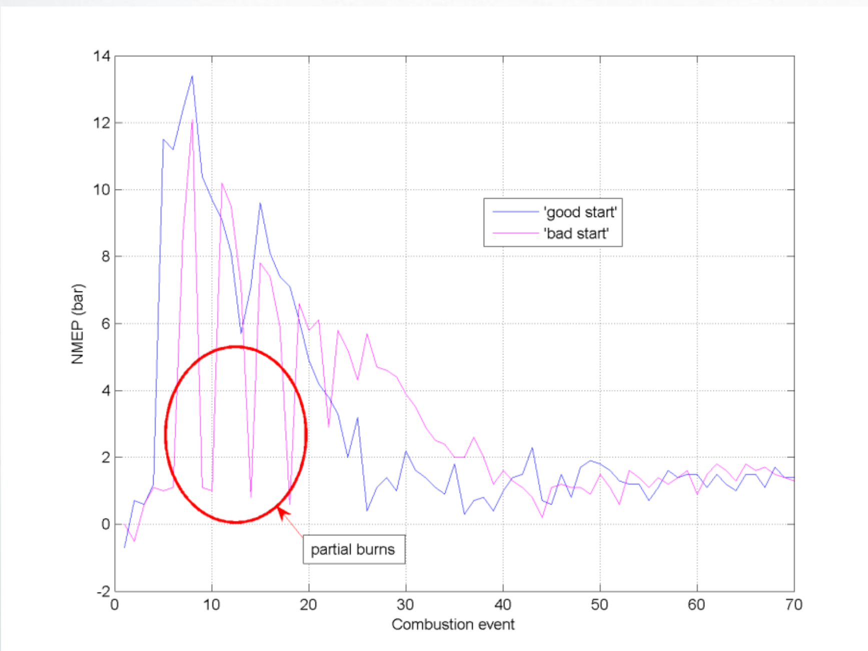

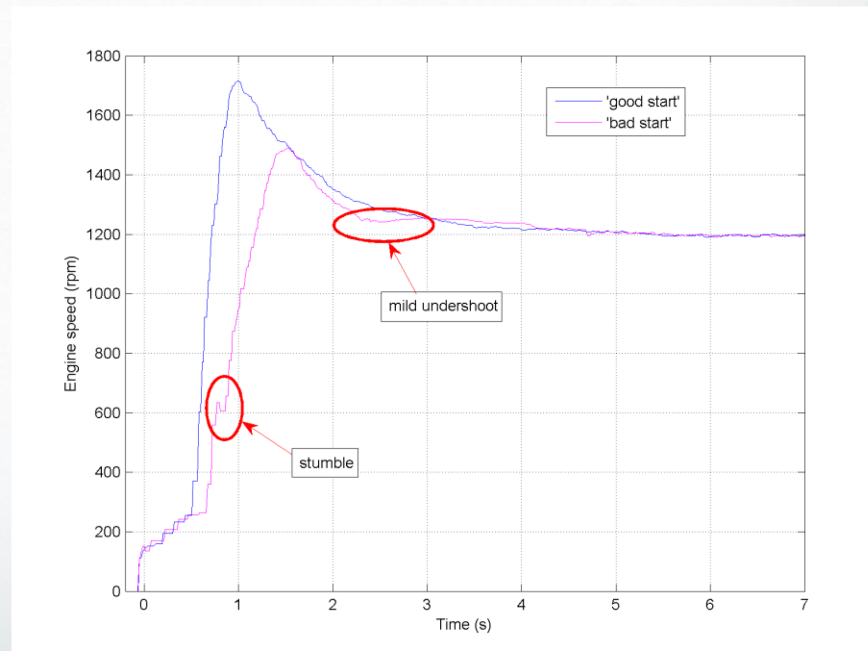

Cold start profiles for a good and bad start

Go Further

Background

- Proposed objective reduction process based on principle: local objective harmony (positive correlation) may exist for many-objective problems.
- Local objective dependency can be revealed by clustering the Pareto-optimal front – suitable algorithm: k*-Means algorithm - correct number and location of clusters verified by simulation.
- Several dimension reduction methods are available to identify and quantify objective dependency.
  - Linear methods:
    - Factor Analysis assumes some random data error exists – n/a to models used for objective functions.
    - Principal Components Analysis (PCA) - no such assumption, but cited as being unsuitable for non-linear data such as Pareto-optimal fronts.
    - Higher order methods (allow for non-Gaussian data) - computationally expensive or may rely on other methods.
  - Non-linear methods (may require non-linear transform or distributional assumption)
    - Multi-Dimensional Scaling - not able to project onto lower dimensions.
    - Self-Organising Maps - issues with subjectivity involved in hierarchical clustering, convergence and interpretation.
    - Vector Quantisation - target dimension must be specified a priori and no consideration given to objective harmony/conflict.
  - In summary, PCA was justified to identify any objective redundancy because:
    - It can quantify any local harmony for objective reduction if the Pareto-optimal front is first partitioned into groups of like-solutions.
    - It has widespread usage and is computationally efficient.

Process

Notes on process:

- Verification rules have been designed to justify the number and location of clusters.
- Objective reduction rules have been designed supplement the PCA to identify objective harmony and which objectives to retain.

Observations from previous (6-objective) application:

- There was only 1 stage of objective reduction. More objectives -> potentially more stages of objective reduction -> more lengthy application of rules. Compact rules suitable for automation, would be useful.
- Higher-dimensional problems -> larger populations (for effective search) -> more clusters -> more computational efficiency. Parallel computing can address this requirement.
- Increased objectives -> potentially increased Principal Components (PCs) -> finer gradation of variation represented by PCs -> potentially retain a different number of PCs -> a different degree of objective reduction.

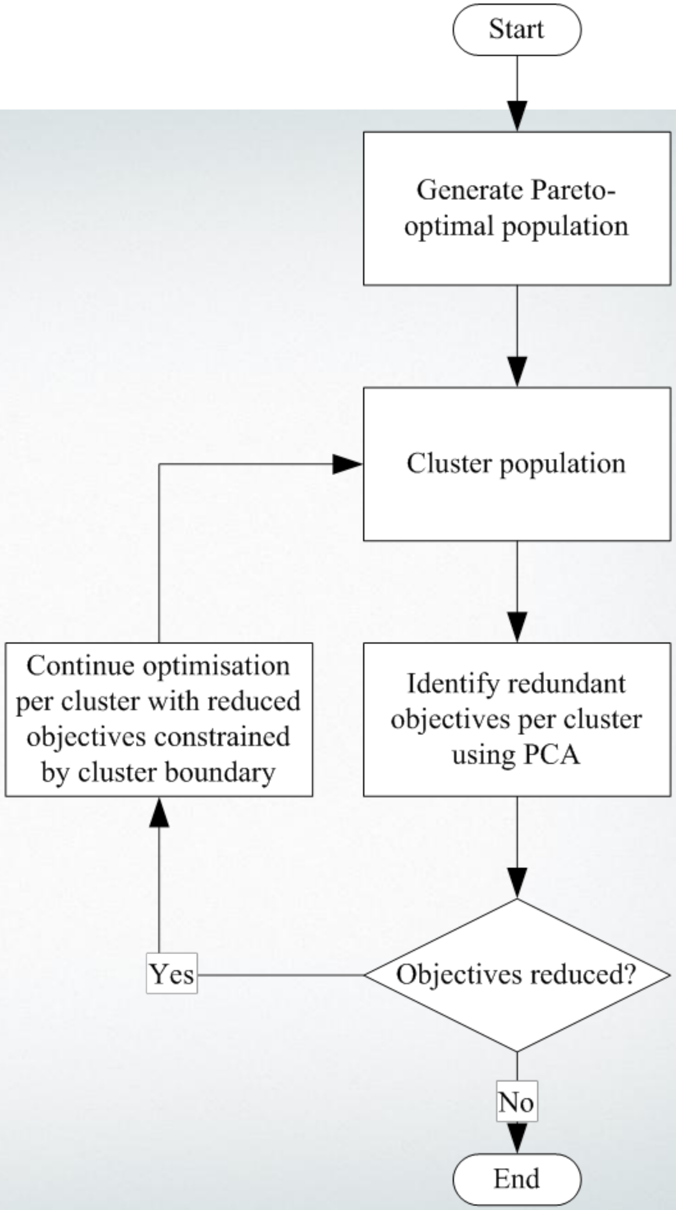

Go Further

Enhancements

- The previously published complexity reduction strategy was extended with several significant enhancements to support many-objective problems:
  - **Sensitivity objectives**
    - To enable the the optimiser simultaneously searches for solutions, which are optimal for performance and which minimise the sensitivity to background noise.
  - **Variations on thresholds for reducing the number of objectives**
    - Varying the threshold used for selecting Principal Components may affect the number of objectives retained using the objective reduction rules. This can provide flexibility in the dimension reduction process.
  - **Parallel computing methods**
    - A parallel MOEA to evaluate large populations distributed across a cluster of processors.
    - Batch processing in parallel to accelerate the clustering task.

- Fundamental concept: exploit any local harmony to allow various degrees of complexity reduction in several local domains of the Pareto-optimal population.
- The resultant sequence of optimisations, clustering and objective reduction processes enables the decision maker, working in conjunction with an experienced calibration engineer, to propose potential solutions.
- These results, developed systematically using the methods described, are shown to out-perform the existing calibration developed using empirical approaches.

Cold Start Engine Calibration Optimisation

Implementation of enhancements:

- To include both sensitivity (along with engine response) objectives in the optimisation problem, the Direct Derivatives (forward finite difference) method was selected.
  - computationally efficient, easy-to-implement approach, easy-to-integrate into the optimisation and sufficiently general for engine calibration optimisation problems.
- Historically, parallel or distributed computing has been an important initiative in solving time-consuming real-world optimisation problems. Three broad paradigms:
  - Master-Slave model (for expensive objective functions)
  - Island model (for distributed computing)
  - Diffusion model (massively parallel computer)
- A compute cluster was available and the execution speed of the objective functions used was very fast (i.e. msecs) -> most suitable approach was the island-based pMOEA -> parallel NSGAII.
- A validation test was carried out on the six-objective problem (Lygoe et al. 2010) resulting in good agreement between the Pareto-optimal populations from the serial and the parallel NSGAII (see next slide).
- A significant speed-up was achieved (~x80) reducing execution time from approximately 21h to 15min – achieved through parallelisation (20 processors used) and the resulting efficiency of cache speed-up arising from the parallel configuration.

Cold Start Engine Calibration Optimisation

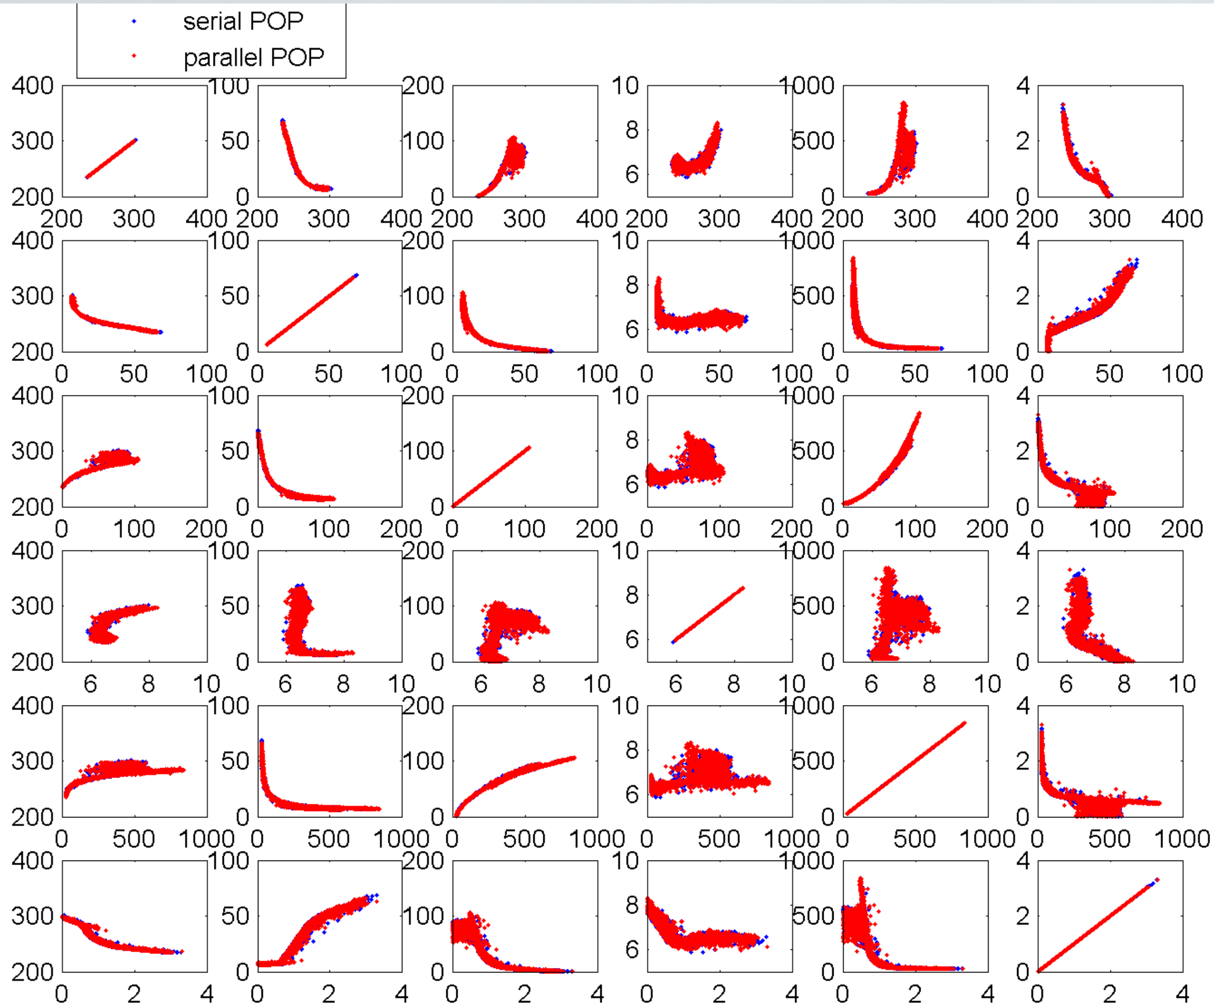

Pareto-optimal populations resulting from runs with serial and parallel versions of NSGAII. Objective axes not labelled for compactness and as they are not required.

Cold Start Engine Calibration Optimisation

Problem Formulation:

- Objective functions based on empirical engine models generated from experimental cold-start test data from a 2-litre in-line four cylinder turbocharged direct injection gasoline passenger car engine.
- The optimisation was formulated as a ten-objective, single constraint problem.
  - Minimise the objectives listed below subject to a constraint on the mild extrapolation of valid domain or boundary of the models.

Label	Description	Units
Obj1	Combustion variation metric for cycles 2-5	Bar
Obj2	Combustion variation metric for cycles 6-12	Bar
Obj3	Negative run-up combustion intensity for cycles 2-5	Bar
Obj4	Negative run-up combustion intensity for cycles 6-12	Bar
Obj5	Fuel quantity	Unitless
Obj6	Maximum engine speed flare after start	RPM
Obj7	Absolute value of sensitivity of combustion variation metric for cycles 2-5 to Fuel Pressure	Bar/MPa
Obj8	Absolute value of sensitivity of combustion variation metric for cycles 6-12 to Fuel Pressure	Bar/MPa
Obj9	Absolute value of sensitivity of run-up combustion intensity for cycles 2-5 to Fuel Pressure	Bar/MPa
Obj10	Absolute value of sensitivity of run-up combustion intensity for cycles 6-12 to Fuel Pressure	Bar/MPa

  - pMOEA parameters: 50,000 gens; pop. size: 20,000; external archive used and updated every 1 gen.; 2% migrants migrated every gen.; tournament selection; SBX crossover; polynomial mutation; took ~14h to run.

Cold Start Engine Calibration Optimisation

Dynamic dynamometer encapsulated test facility – can provide efficient, cost effective and realistic testing on a rig as opposed to building expensive prototype vehicles, which require specialised vehicle-based test facilities or testing in remote cold climate locations.

Go Further

Cold Start Engine Calibration Optimisation

Results:

Initial clustering of 18,552 Pareto-Optimal Population (POP), various learning rates, initial number of clusters, max. iterations = 5000 & convergence tol. = 0.1 -> reference solution of 4 converged clusters.

Reference solution randomly sub-sampled per cluster to generate smaller POPs of 10k, 5k, 2k & 1k and each clustered -> 10k POP smallest sub-sampled POP that provided acceptable agreement with the reference POP.

Cold Start calibration engineer specified that for Objective 6 (Peak Flare Speed), only solutions in the 1300-1500rpm range were of interest. Cluster 1 violated this limit and so was discarded.

The objective reduction rules with objective priorities were applied to each cluster to identify potential objective reduction. e.g. in cluster 2:

- 6 objectives were retained (Included):1, 3, 6, 7, 9 & 10.
- 4 (the remainder) were discarded (Excluded): 2, 4, 5 and 8.

Objective Reduction Stage 1 Flowchart

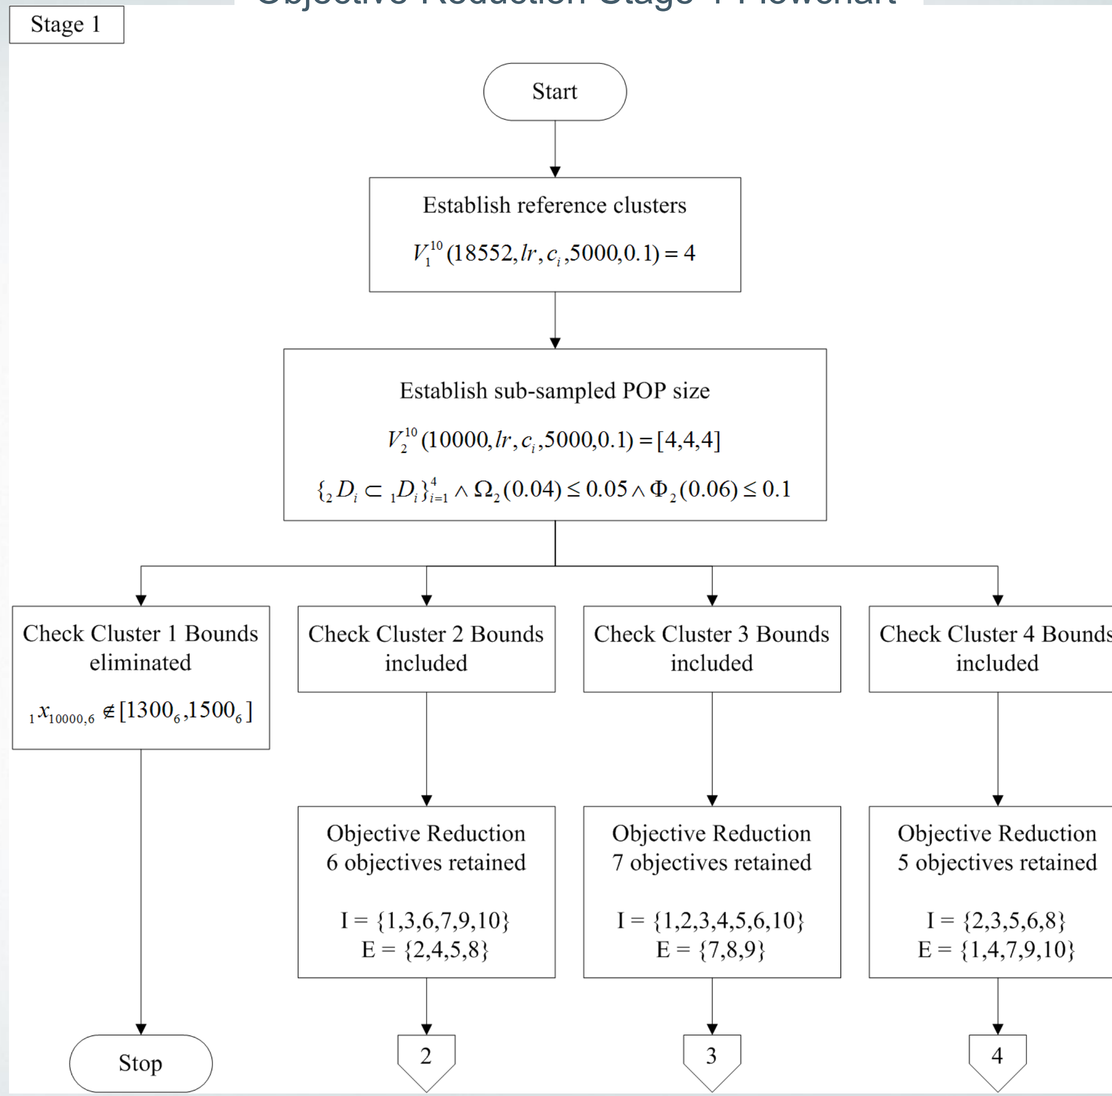

Go Further

Cold Start Engine Calibration Optimisation

The Effect of a Reduced Percentage of Variation on Objective Reduction

- Main purpose of PCA is to reduce the dimensionality of the problem by replacing the objectives by a smaller number of Principal Components (PCs) which account for most of the variation.
- For this study, a cumulative percentage of variation threshold is used where the number of PCs retained is the smallest number of PCs whose cumulative percentage of variation exceeds this threshold.
- The default value for this threshold (95%) was evaluated as well as 90% and 86% and the effect on objective reduction is shown in the table below.
- Some objective reduction was achieved using a threshold of 95%, but 86% was chosen as it gives significantly more reduction.
- Consequently, this stage culminated in three Clusters (2, 3 and 4) with six, seven and five objectives being retained, respectively.

Threshold for cumulative % of total variation	No. of objectives retained		
	Cluster 2	Cluster 3	Cluster 4
95	7	10	7
90	6	8	7
86	6	7	5

Go Further

Cold Start Engine Calibration Optimisation

Conclusions from the Objective Reduction Process:

- The results from the final objective reduction (3rd Stage) are displayed in the table below and show the number of objectives retained in each cluster at each stage.

Objective reduction Stage	No. of objectives retained			
	Cluster 2	Cluster 3	Cluster 4	
1st	6	7	5	
			Cluster 4_1	Cluster 4_2
2nd	4	5	4	4
3rd	4	4	4	3

- A Parallel Coordinates plot (next slide) of the Pareto-optimal populations resulting from the final objective reductions was reviewed with the Cold Start calibration engineer.
- Clusters 3 and 4_1 were discarded as they contained solutions considered to be comparatively inferior to those in clusters 2 and 4_2 (details in paper).

Go Further

Cold Start Engine Calibration Optimisation

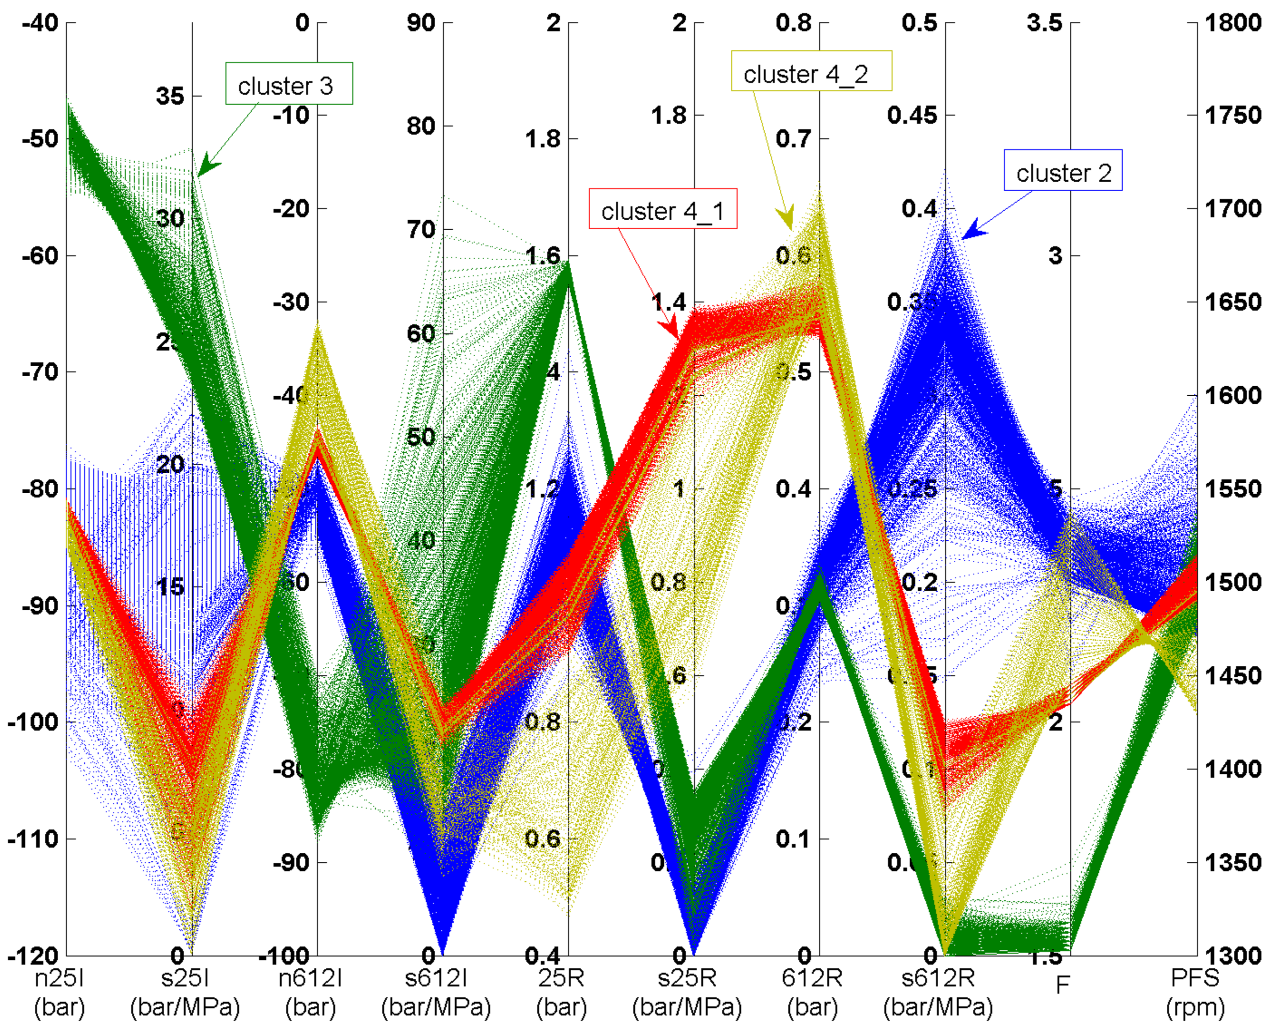

Parallel Coordinates plot of the final populations resulting from objective reduction in each cluster. The objectives have been re-ordered so that each sensitivity objective is adjacent to the objective to which it relates.

Go Further

Cold Start Engine Calibration Optimisation

Comparison of selected solutions:

- Parallel Coordinates plots of selected solutions vs. respective parent population vs. a recent calibration generated by the Cold Start calibration engineer using a manual, iterative tuning process.
- For both clusters, it can be seen that the calibration is inferior with respect to the population and the selected solutions. The exception is in Cluster 2, where the calibration is slightly better (smaller) than selected solution 271 with respect to Obj1 (n25I: combustion intensity for cycles 2-5).

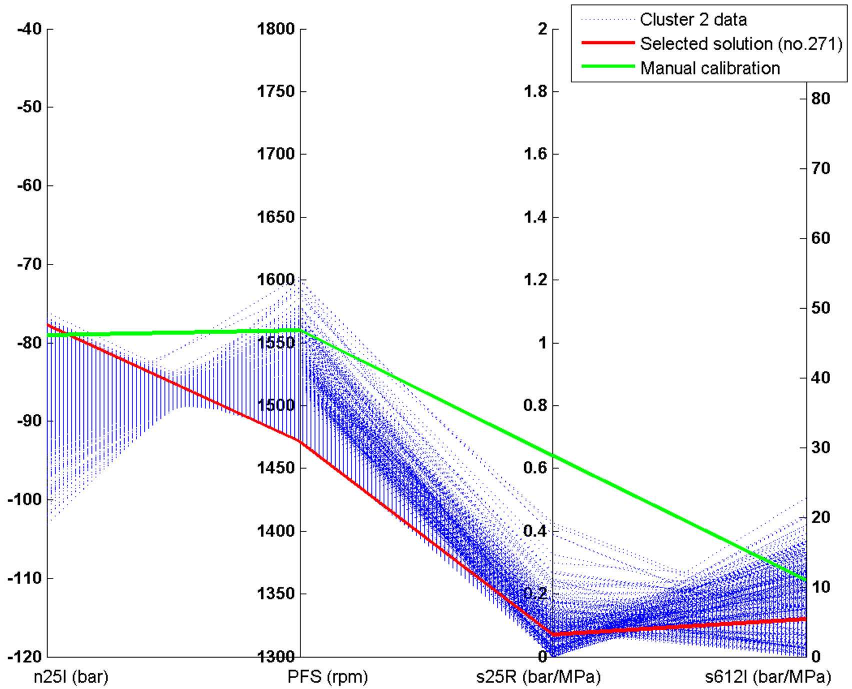

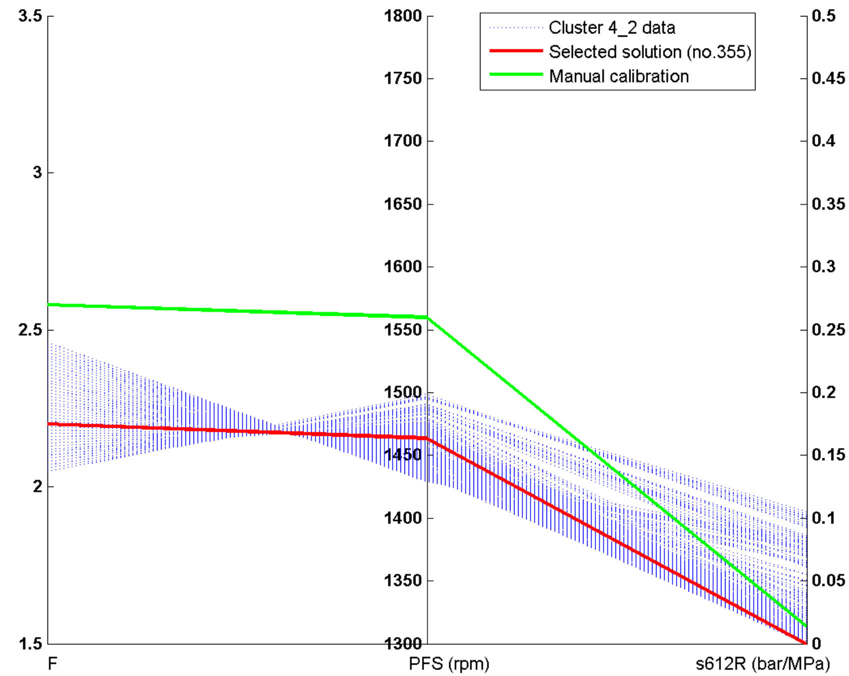

Go Further

Conclusions and Future Work

Conclusions:

- The Multi-Objective Optimisation Decision-Making (MOODM) process (Lygoe et al, 2010) has been extended and applied to a real-world automotive engine problem of increased complexity: a ten-objective engine cold start optimisation.
- The process enhancements comprised :
  - Defining and embedding sensitivity objectives into the optimisation to yield a robust calibration.
  - Application of parallel computing to make the process efficient.
  - The use of varying PC thresholds to explore the potential to achieve greater objective reduction.
- In general, the preferred solutions that resulted from this MOODM process compared favourably to those generated from a manual tuning calibration approach.

Future Work includes:

- Revisit the objective reduction process where there is evidence of independence between objectives. In such scenarios, an increased number of lower dimension optimisations may result.
- The two main stages of the search (evolution of an initial population for clustering and further evolution of populations within clusters) have involved many thousands of objective function evaluations. Is this justified or would fewer objective function evaluations would suffice?
- Implement mathematical notation for the Clustering verification and Objective reduction Rules in software. This will make a high-dimensional multi-stage objective reduction process more efficient, less error-prone and potentially fully automated including documentation of results at each stage.