Mathematical Foundations and Derivation of a Reduced - Order Model for Radial Piston Hydraulic Motor Optimisation

Abstract

This document presents the complete mathematical development of the reduced - order physics model used to analyse and optimise a radial piston hydraulic motor. Unlike the accompanying report, which presents only the governing equations required to describe the system, this document derives each equation from first principles and discusses the physical assumptions underlying every modelling decision.

The model couples geometry, kinematics, fluid dynamics, structural mechanics, and spectral analysis into a single computational framework suitable for design optimisation. Beginning with the geometric description of the cam mechanism, the derivation proceeds through chamber volume evolution, compressible hydraulic dynamics, piston mechanics, and harmonic analysis before concluding with the optimisation methodology and surrogate modelling framework.

The objective is not merely to present a collection of governing equations, but to demonstrate how each component of the model emerges naturally from the underlying laws of mechanics and fluid dynamics. Particular emphasis is placed on identifying the approximations introduced to reduce computational complexity while retaining the dominant physical behaviour governing torque ripple, pressure pulsation, and overall efficiency.

Contents

1. Introduction
2. Modelling Philosophy
3. Coordinate Systems
4. Cam Geometry
5. Piston Kinematics
6. Chamber Volume Derivation
7. Conservation of Mass
8. Fluid Compressibility
9. Hydraulic Pressure Dynamics
10. Valve Plate Flow
11. Leakage Modelling
12. Mechanical Dynamics
13. Equivalent Hydraulic Stiffness
14. Frequency - Domain Representation
15. Harmonic Metrics
16. Optimisation Formulation
17. Gaussian Process Surrogate
18. Bayesian Optimisation
19. Complete System Model

Introduction

Motivation

Hydraulic radial piston motors convert hydraulic energy into mechanical rotation through the interaction of pressurised fluid, reciprocating pistons, and a cam profile. Although the physical mechanism appears relatively straightforward, the resulting system exhibits strongly coupled nonlinear behaviour.

The instantaneous output torque depends not only upon the chamber pressure, but also upon the piston position, cam curvature, valve timing, leakage, fluid compressibility, and inertial effects. Each of these mechanisms influences the others. Changing the cam profile modifies piston motion, which alters chamber volume, which changes pressure dynamics, which changes the transmitted force and ultimately the output torque.

Consequently, designing a hydraulic motor cannot be reduced to independently optimising each subsystem. The entire machine must instead be considered as a coupled dynamic system.

Traditional analytical models often simplify the motor by assuming incompressible flow, ideal valve timing, or quasi-static pressure distributions. Although these assumptions allow closed-form solutions, they remove many of the physical mechanisms responsible for undesirable behaviour such as:

- Torque ripple
- Pressure pulsation
- Vibration
- Acoustic noise
- Reduced volumetric efficiency

Conversely, high-fidelity computational fluid dynamics (CFD) or finite element simulations can capture these effects with excellent accuracy, but their computational cost is prohibitive when thousands of candidate designs must be evaluated during optimisation.

This work therefore adopts an intermediate modelling philosophy. The objective is to develop a reduced-order model that retains the dominant physical mechanisms governing performance while remaining sufficiently inexpensive for repeated numerical evaluation.

Modelling Objectives

The mathematical model should satisfy four principal requirements: First, every equation should possess a clear physical interpretation derived directly from established conservation laws. Second, the model should remain computationally efficient enough to evaluate thousands of candidate geometries during optimisation. Third, the governing equations should preserve the nonlinear coupling between:

- Geometry
- Hydraulics
- Mechanics
- Spectral response

Finally, the model should naturally produce physically meaningful optimisation objectives such as torque ripple, pressure ripple, and efficiency.

Overall System Architecture

Rather than considering the hydraulic motor as a collection of independent subsystems, it is useful to view the complete model as a sequence of nonlinear mappings. Each stage transforms information into a new physical quantity:

$$x \rightarrow C(\theta) \rightarrow s(\theta) \rightarrow V(\theta) \rightarrow P(\theta) \rightarrow F(\theta) \rightarrow \tau(\theta)$$

Where:

- x: Design variables
- C: Cam geometry
- s: Piston displacement
- V: Chamber volume
- P: Chamber pressure
- F: Hydraulic force
- τ: Output torque

Each arrow represents the solution of a separate physical model. The optimisation process therefore operates not directly upon the torque, but upon the original design vector.

Governing Physical Principles

The complete model is constructed from only four fundamental physical principles:

Geometry: The cam profile determines the piston trajectory. Newton's Second Law: Hydraulic pressure generates forces that accelerate the piston. Conservation of Mass: Fluid entering or leaving the chamber changes the chamber pressure. Fluid Constitutive Behaviour: The compressibility of hydraulic oil links pressure changes to changes in chamber volume.

Everything that follows in this document is ultimately derived from these four principles.

Modelling Assumptions

Every mathematical model represents an approximation to reality. The assumptions introduced here reduce computational complexity while retaining the dominant physical behaviour.

- Isothermal Fluid: Temperature variations are neglected. Fluid density and viscosity are therefore assumed constant throughout the operating cycle. This assumption is appropriate when thermal time constants are significantly larger than the shaft rotation period.
- Weak Compressibility: Hydraulic oil is modelled as slightly compressible through a constant bulk modulus, β. This captures pressure-wave dynamics while avoiding the complexity of a full compressible Navier-Stokes formulation.
- Rigid Bodies: The pistons, housing, and cam are assumed rigid. Elastic deformation is neglected. Consequently, all compliance is attributed to the hydraulic fluid.
- Constant Rotational Speed: The shaft rotates at constant angular velocity, ω. Transient drivetrain dynamics are ignored.
- Lumped Parameter Chambers: Each piston chamber is represented as a spatially uniform control volume. Pressure is therefore assumed constant throughout an individual chamber. This approximation transforms the governing partial differential equations into ordinary differential equations.
- Quasi-Steady Orifice Flow: Flow through the valve plate is represented using the steady-state orifice equation. Dynamic jet formation and turbulence are incorporated empirically through the discharge coefficient.

Document Roadmap

The derivation proceeds in the same order as information flows through the physical system.

The geometry is developed first because every subsequent subsystem depends upon piston motion. Once the piston trajectory has been established, chamber volume can be determined. Volume evolution leads directly to conservation of mass, from which the pressure dynamics are derived. Pressure then generates hydraulic forces acting upon the piston. Finally, the resulting periodic signals are transformed into the frequency domain, enabling quantitative measures of torque ripple and pressure pulsation suitable for optimisation.

The optimisation framework presented in the final sections therefore emerges naturally as the final layer built upon the preceding physical derivations rather than existing as an independent mathematical construct.

Modelling Philosophy

Before deriving the governing equations, it is important to establish the modelling philosophy adopted throughout this work. Mathematical models exist on a spectrum ranging from highly simplified analytical descriptions to extremely detailed numerical simulations. The choice of model depends not only upon the desired accuracy but also upon the intended application.

The objective of this work is not to produce the most physically complete representation of a radial piston hydraulic motor. Instead, the objective is to construct a model that captures the dominant physical mechanisms governing performance while remaining sufficiently computationally efficient for large-scale optimisation.

Unlike validation studies, where a single design may justify several days of computational analysis, optimisation requires the repeated evaluation of thousands of candidate geometries. Consequently, the computational cost of the model becomes a design constraint in its own right.

This section establishes the rationale behind the adopted modelling strategy and introduces the mathematical structure upon which the remainder of this document is built.

Levels of Physical Modelling

Engineering systems may generally be represented using one of four modelling approaches.

Empirical Models

Empirical models, also known as surrogate models, are constructed directly from experimental observations without explicitly representing the underlying physical processes. A typical empirical relationship may be written as $y = f(x)$, where the mapping $f(x)$ is obtained through regression or curve fitting. While these models are computationally inexpensive, their predictive capability is limited to operating conditions similar to those used during calibration. Extrapolation beyond the training data often produces physically unrealistic behaviour. Consequently, empirical models are unsuitable for optimisation over large design spaces where many candidate solutions may lie far outside previously observed data.

High-Fidelity Numerical Models

At the opposite end of the modelling spectrum are high - fidelity numerical simulations. Examples include Computational Fluid Dynamics (CFD), Finite Element Analysis (FEA), Fluid-Structure Interaction (FSI), Finite Element Method Magnetics (FEMM), Reaction-Diffusion Modelling (RDM), and Multi-body dynamics. These methods solve discretised forms of the governing partial differential equations. For example, fluid flow may be represented through the compressible Navier–Stokes equations:

$$\frac{\partial \rho}{\partial t} + \nabla(\rho u) = 0$$

together with the corresponding momentum equations. Such approaches can accurately capture turbulent flow, cavitation, structural deformation, transient pressure waves, and thermal effects. However, the computational expense increases rapidly with model fidelity. A single CFD simulation may require several hours or even days of computation. Since optimisation typically requires tens of thousands of evaluations, direct use of high-fidelity simulations becomes computationally infeasible.

Analytical Models

Analytical models occupy the opposite extreme. These models intentionally simplify the governing equations until closed-form solutions become possible. Typical assumptions include incompressible fluid, negligible leakage, rigid mechanics, ideal valve timing, and steady-state operation. Although such assumptions allow elegant mathematical analysis, many important dynamic phenomena disappear from the resulting equations. In particular, pressure ripple, torque ripple, compressibility effects, and harmonic interactions cannot be represented adequately.

Reduced - Order Physics Models

Between these two extremes lies the reduced-order physics model. Rather than attempting to simulate every physical mechanism, reduced-order modelling seeks to retain only those dynamics that dominate system behaviour. The present work adopts this philosophy. The model preserves nonlinear chamber volume evolution, fluid compressibility, valve timing, leakage, piston inertia, and harmonic response, while neglecting thermal effects, elastic deformation, local turbulent structures, and detailed internal flow fields. This balance provides a model capable of predicting the dominant dynamic behaviour while remaining computationally efficient.

Hierarchy of Physical Models

The mathematical model developed in this work can be viewed as a hierarchy of coupled physical subsystems. Each subsystem receives information from the preceding stage before generating new state variables that are passed downstream. Conceptually, the flow is as follows:

$\mathbf{x} \rightarrow g(\theta)$	Cam Geometry
$\rightarrow s(\theta)$	Piston Kinematics
$\rightarrow V(\theta)$	Chamber Volume
$\rightarrow P(t)$	Pressure Dynamics
$\rightarrow F_h(t)$	Hydraulic Forces
$\rightarrow T(t)$	Output Torque
$\rightarrow R(T)$	Ripple Metrics
$\rightarrow J(\mathbf{x})$	Objective Function

Although represented sequentially, these quantities remain strongly coupled through nonlinear relationships. For example, changing the cam geometry alters piston motion. Piston motion changes chamber volume. Volume evolution modifies chamber pressure. Pressure generates hydraulic force. Hydraulic force determines output torque. Consequently, optimisation of the cam profile indirectly modifies every downstream subsystem.

Mathematical Structure of the Model

The complete mathematical model may be regarded as a nonlinear operator, $M : x \rightarrow y$, where:

$x = \{design\ variables\}$,

and $y = \{\tau(\theta), P(\theta), \eta, R_\tau, R_P\}$.

Rather than existing as a single equation, the operator is composed of several coupled mappings:

$$M = S \circ F \circ H \circ G$$

Where:

- G represents geometry
- H represents hydraulics
- F represents mechanics
- S represents spectral analysis

This decomposition is important because each subsystem may be independently improved without requiring modification of the remainder of the model.

State Variables

To avoid ambiguity, it is useful to distinguish between design variables, state variables, and performance variables.

- Design Variables: These are quantities directly controlled during optimisation. Examples include cam control points, piston diameter, valve timing, port geometry, leakage clearance. Collect these into the vector x.
- State Variables: State variables evolve according to the governing equations. Examples include chamber pressure, piston displacement, piston velocity, and chamber volume. These quantities cannot be chosen directly. Instead, $s = M(x)$.
- Performance Variables: Performance variables are derived from the state variables. Examples include torque ripple, pressure ripple, efficiency, and harmonic energy. These are the quantities used by the optimisation algorithm.

Computational Requirements

Since Bayesian optimisation repeatedly evaluates the physics model, computational efficiency is critical. Suppose one simulation requires $T_s = 15\ s$. If optimisation requires 20,000 design evaluations, the total computational time becomes $300{,}000\ s \approx 83\ hours$. Doubling the computational complexity therefore doubles optimisation time. For this reason, every modelling assumption must be justified not only physically but computationally. The reduced-order formulation developed throughout this document is specifically designed to minimise computational cost while preserving the dominant dynamic behaviour influencing optimisation.

Model Limitations

Every mathematical model possesses a finite range of validity. The present formulation should therefore be interpreted as a dynamic systems model rather than a high-fidelity fluid simulation. Specifically, the model does not resolve cavitation bubble dynamics, turbulent eddies, structural elasticity, thermal expansion, surface wear, lubrication films, or microscopic leakage paths. Instead, these effects are represented through equivalent parameters whenever their influence significantly affects system-level behaviour. This philosophy is consistent with classical lumped-parameter modelling used throughout hydraulic system analysis.

Summary

This section established the modelling philosophy underlying the remainder of this document. Rather than pursuing maximum physical fidelity, the objective is to retain only those mechanisms required to predict the dominant dynamic characteristics influencing optimisation. The resulting framework occupies the middle ground between purely analytical models and computationally intensive numerical simulations.

Having established this philosophy, the following section begins the mathematical derivation by defining the coordinate systems used to describe the geometry and motion of the radial piston mechanism.

Coordinate Systems and Geometric Representation

The motion of a radial piston hydraulic motor is fundamentally geometric. Every subsequent physical quantity – including chamber volume, hydraulic pressure, piston force, and output torque – is ultimately determined by the relative position of the pistons with respect to the cam profile. Consequently, the first step in constructing the mathematical model is to establish a consistent geometric framework.

Rather than treating the motor as an assembly of individual components, the mechanism is described using a small number of coordinate systems related through simple transformations. This approach allows the geometry, kinematics, and dynamics to be expressed in a compact and mathematically consistent manner.

Throughout this document the shaft is assumed to rotate with constant angular velocity, $\omega = \frac{d\theta}{dt}$, where t denotes time and θ denotes the angular position of the shaft. Because the motion is periodic, θ becomes the natural independent variable for much of the analysis. All state variables are therefore expressed as periodic functions of shaft angle rather than time wherever possible.

The Global Reference Frame

Consider the centre of the hydraulic motor to coincide with the origin of a fixed Cartesian coordinate system $F_G = \{x, y\}$. The positive x-axis defines the zero angular reference, and all angular quantities are measured counterclockwise according to the standard mathematical convention. Any point within the mechanism may therefore be represented by the position vector:

$$r = [x, y]^T$$

Although Cartesian coordinates provide a convenient global description, the geometry of a radial piston motor is inherently circular. Polar coordinates therefore provide a more natural representation.

Polar Coordinate Representation

The position vector may alternatively be expressed in polar form:

$$r = r\,[cos\theta, sin\theta]^T$$

Where r denotes the radial distance from the shaft centre and θ denotes the angular position. The transformation between Cartesian and polar coordinates follows directly from elementary trigonometry:

$$x = r\,cos\theta \; y = r\,sin\theta$$

With inverse relations:

$$r = \sqrt{x^2 + y^2}\theta = tan^{-1}\left(\frac{y}{x}\right)$$

Throughout the remainder of this document the radial coordinate r will generally correspond to the instantaneous cam radius.

The Rotating Shaft Frame

While the housing remains stationary, the cylinder block rotates with the shaft. It is therefore convenient to define a second coordinate frame, F_R, which rotates with angular velocity ω. The orientation of the rotating frame relative to the global frame is described by the rotation matrix:

$$R(\theta) = [\,cos\theta - sin\theta\,][\,sin\theta\; cos\theta\,]$$

Any vector expressed in the rotating frame may therefore be transformed into the stationary frame through:

$$r_G = R(\theta) r_R$$

Conversely:

$$r_R = R^{-1}(\theta) r_G$$

Where $R^{-1} = R^T$, since rotation matrices are orthogonal. This property significantly simplifies later derivations involving piston motion and force transmission.

Angular Position of Individual Pistons

Consider a motor containing N_p equally spaced pistons. Since the pistons are uniformly distributed around the cylinder block, the angular separation between neighbouring pistons is:

$$\Delta\theta = \frac{2\pi}{N_p}$$

If piston $i = 0$ is taken as the reference piston, the angular position of piston i becomes:

$$\theta_i = \theta + \frac{2\pi i}{N_p}$$

Where $i = 0, 1, \ldots, N_p - 1$.

As the shaft rotates, every piston experiences identical motion shifted only by a constant phase angle. This observation is particularly important because it allows the behaviour of the entire motor to be obtained from the motion of a single piston together with an appropriate phase shift. Instead of solving N_p completely independent problems, the model exploits this rotational symmetry:

$$s_{i(\theta)} = s\left(\theta + \frac{2\pi i}{N_p}\right)$$

The same relationship later applies to chamber volume, chamber pressure, hydraulic force, piston velocity, and piston acceleration. This symmetry substantially reduces computational cost.

Periodicity of the Mechanism

Unlike many mechanical systems, the radial piston motor possesses exact geometric periodicity. After one complete shaft revolution, $\theta \rightarrow \theta + 2\pi$, the physical configuration of the mechanism repeats. Consequently, every state variable satisfies $f(\theta) = f(\theta + 2\pi)$. Examples include $s(\theta)$, $V(\theta)$, $P(\theta)$, and $\tau(\theta)$.

This periodicity is one of the most important mathematical properties of the model. It permits the use of Fourier analysis in later sections and ultimately enables ripple to be characterised directly through harmonic amplitudes.

Design Variables as Geometric Parameters

The optimisation algorithm does not directly modify pressure or torque. Instead, it alters a collection of geometric parameters that define the motor. Collect these quantities into the design vector:

$$x = [x_1, x_2, \ldots, x_n]^T$$

Typical components include cam control point locations, piston diameter, dead volume, valve timing, port dimensions, leakage coefficients, and manufacturing tolerances. Every subsequent equation developed throughout this document is therefore a mapping from:

$$Geometry \rightarrow Kinematics \rightarrow Hydraulics \rightarrow Mechanics$$

Understanding this flow of information is fundamental to understanding the optimisation framework developed later.

Why the Cam Profile Is Treated as a Continuous Function

The instantaneous piston position depends upon the radial distance from the shaft centre to the cam surface. One could represent this profile using discrete measurements, $\{r_1, r_2, \ldots, r_n\}$, however such a representation introduces discontinuities in velocity and acceleration unless interpolation is performed. Since optimisation requires repeated differentiation of the geometry, a continuously differentiable representation is preferable. For this reason, the cam profile is treated as a smooth periodic function, $C(\theta)$, whose mathematical construction will be developed in the following section. The use of a smooth parametric representation ensures that velocity, acceleration and jerk all exist and remain continuous. This property is essential for both dynamic performance and numerical optimisation.

Summary

This section established the geometric framework used throughout the remainder of the mathematical model. By defining both stationary and rotating coordinate systems, introducing piston phase relationships, and exploiting the inherent periodicity of the mechanism, the geometry of the radial piston motor has been reduced to a compact mathematical description suitable for further analysis.

The next section constructs the cam profile itself. Beginning with the mathematical theory of B-splines, the derivation will show how a smooth periodic cam geometry can be represented using a finite set of control points while maintaining the continuity required for dynamic simulation and gradient-based optimisation.

Constructing the Cam Profile

The cam profile is the single most influential design variable in a radial piston hydraulic motor. Unlike many geometric parameters, such as piston diameter or chamber volume, the cam directly governs the trajectory of every piston throughout the operating cycle. Consequently, its shape simultaneously influences chamber volume evolution, hydraulic pressure generation, force transmission, torque production, vibration, and dynamic loading.

From a mathematical perspective, the optimisation algorithm does not directly optimise pressure or torque. Instead, it searches for an improved cam geometry. Every subsequent physical quantity is therefore a consequence of the chosen profile. The objective of this section is to develop a mathematical representation of the cam that is sufficiently flexible to describe complex geometries while remaining smooth, differentiable, computationally efficient, and suitable for numerical optimisation.

Desired Properties of the Cam Representation

Before selecting a mathematical representation, it is useful to identify the characteristics that an ideal cam description should possess.

- First, the profile must be closed. Since the cam forms a continuous mechanical surface, its start and end points must coincide exactly after one revolution.
- Second, the profile must be smooth. Sudden changes in curvature produce impulsive contact forces that increase vibration, stress concentration, and wear.
- Third, the profile should require relatively few parameters. Excessively large parameter spaces significantly increase optimisation time and make convergence more difficult.
- Finally, the representation should permit analytical differentiation. Since piston velocity, acceleration, and jerk are all obtained by differentiating the cam geometry, the chosen mathematical form should possess continuous derivatives of sufficiently high order.

These requirements immediately eliminate several common geometric representations. For example, piecewise linear interpolation introduces discontinuous slopes, while high-order global polynomials suffer from Runge's phenomenon and exhibit undesirable oscillatory behaviour. Spline functions provide an attractive compromise between these two extremes.

Why Splines?

A spline is a piecewise polynomial curve constructed by joining low-order polynomial segments together while enforcing continuity at their boundaries. Historically, the term spline originates from the flexible wooden strips used by shipbuilders and draftsmen to produce smooth curves through a series of fixed points. The mathematical spline performs an analogous role: rather than fitting a single polynomial through every point, the curve is assembled from many local polynomials whose transitions are carefully constrained.

This local construction provides several important advantages. A modification to one section of the curve influences only neighbouring regions rather than the entire geometry. The resulting representation is numerically stable. Derivatives are straightforward to compute. Smoothness may be controlled explicitly through the polynomial degree. For these reasons, spline functions have become the standard geometric representation in computer-aided design (CAD), computer graphics, and numerical optimisation.

Polynomial Curves

Consider first the simplest polynomial representation:

$$r(\theta) = a_0 + a_1\theta + a_2\theta^2 + \ldots + a_n\theta^n$$

Although this equation appears attractive, it possesses several significant shortcomings. Increasing the polynomial order introduces increasingly oscillatory behaviour between interpolation points. Furthermore, every coefficient affects the entire curve. Consequently, modifying a single design variable changes the geometry globally. This behaviour is undesirable during optimisation because small parameter changes may produce unexpectedly large geometric changes. The solution is to abandon global polynomials in favour of piecewise polynomials.

Piecewise Polynomial Representation

Suppose the domain $0 \leq \theta < 2\pi$ is divided into a sequence of intervals: $[\theta_0, \theta_1], [\theta_1, \theta_2], \ldots, [\theta_{\{m-1\}}, \theta_m]$. Rather than describing the entire cam using one polynomial, each interval possesses its own cubic polynomial:

$$r_i(\theta) = a_i + b_i\theta + c_i\theta^2 + d_i\theta^3$$

At first glance this appears to introduce discontinuities where neighbouring segments meet. The key idea behind spline theory is to enforce continuity conditions at every junction.

Continuity Requirements

Suppose two neighbouring polynomial segments meet at $\theta = \theta_i$. The simplest requirement is positional continuity, $r_i(\theta_i) = r_{i+1}(\theta_i)$. This is known as C^0 continuity. However, position alone is insufficient. If the first derivatives differ, $\frac{dr_i}{d\theta} \neq \frac{dr_{i+1}}{d\theta}$, the curve contains a sharp corner. To eliminate these discontinuities, slope continuity is imposed, $\frac{dr_i}{d\theta} = \frac{dr_{i+1}}{d\theta}$. This produces C^1 continuity.

For mechanical systems this is still not enough. Since piston acceleration depends upon the second derivative of the cam profile, discontinuities in curvature generate impulsive forces. Consequently, the second derivatives must also agree, $\frac{d^2r_i}{d\theta^2} = \frac{d^2r_{i+1}}{d\theta^2}$, producing C^2 continuity. A cubic spline is the lowest - order polynomial capable of satisfying all three continuity conditions simultaneously. This is one of the principal reasons cubic splines are almost universally adopted in mechanical cam design.

B-Spline Basis Functions

Instead of defining every polynomial segment explicitly, B-splines express the curve as a weighted combination of basis functions. The cam profile is written as:

$$C(\theta) = \sum_{i=0}^{n} N_{\mathrm{i},\,\mathrm{p}}(\theta) P_{\mathrm{i}}$$

Where:

- P_{i} are the control points
- $N_{\mathrm{i},\,\mathrm{p}}$ are the B-spline basis functions
- p denotes the spline degree

Unlike interpolation methods, the control points do not generally lie on the curve itself. Instead, they act as geometric attractors that influence the surrounding region of the curve. This distinction is fundamental. Moving one control point modifies only a local portion of the geometry, providing the locality required for efficient optimisation.

Why a Cubic B-Spline?

Throughout this work the cam profile is represented using a cubic B-spline, $p = 3$. This choice reflects a balance between smoothness and computational efficiency. A linear spline would introduce discontinuous velocities. A quadratic spline would produce discontinuous accelerations. Higher-order splines increase computational cost while providing relatively little additional benefit for the present application. The cubic spline represents the lowest - order spline possessing continuous position, velocity, and acceleration throughout the complete cam profile. Since the piston dynamics depend directly upon the first and second derivatives of the cam geometry, this level of continuity is sufficient for accurate dynamic simulation.

Periodic Boundary Conditions

Unlike an open curve, the cam forms a closed mechanical surface. Consequently, the spline must also be periodic. Mathematically this requires $C(0) = C(2\pi)$, together with equality of the first and second derivatives:

$$C'(0) = C'(2\pi) C''(0) = C''(2\pi)$$

These conditions ensure that the cam closes without a geometric discontinuity. In practice, periodicity is achieved by wrapping the first few control points onto the end of the knot vector. The resulting spline possesses seamless continuity across the $0 \to 2\pi$ boundary, making it indistinguishable from a single continuous curve. This periodic construction is essential for a rotating mechanism. Any discontinuity at the closure point would generate impulsive contact forces once per revolution, producing unrealistic pressure spikes and severe numerical instability.

The Cam as an Optimisation Variable

Having established the mathematical representation of the cam, the optimisation problem may now be viewed as the search for an optimal set of control points, $P = \{P_0, P_1, \ldots, P_{\mathrm{n}}\}$. Rather than directly modifying thousands of points along the cam surface, the optimisation algorithm adjusts only the relatively small set of control points. The B-spline basis functions then interpolate these into a smooth continuous profile. This dramatically reduces the dimensionality of the optimisation problem while preserving the ability to generate highly complex cam geometries. The cam profile therefore becomes a compact parametric description from which all subsequent geometric and dynamic quantities are derived.

Summary

This section has established the mathematical representation of the cam profile used throughout the remainder of this work. Beginning with the requirements imposed by the physical mechanism, it has been shown that a periodic cubic B-spline provides an excellent balance between geometric flexibility, numerical stability, local control, and smoothness. By representing the cam as a weighted combination of locally supported basis functions, the optimisation problem is reduced to the selection of a finite number of control points while maintaining continuous position, velocity, and acceleration.

The following section develops the mathematical properties of the B-spline basis functions themselves before deriving the piston kinematics from the resulting cam geometry.

The Mathematics of Cubic Periodic B-Splines

The previous section established that the cam profile is represented using a periodic cubic B-spline. While this choice was motivated from an engineering perspective, it is now necessary to develop the mathematical theory underlying this representation.

Having established the engineering motivation for representing the cam profile as a periodic cubic B-spline, the mathematical construction of the spline can now be developed. Rather than viewing the cam simply as a smooth curve, it is now treated as a weighted combination of polynomial basis functions possessing compact support. This section derives these basis functions from first principles, establishes their principal mathematical properties, and demonstrates why they provide an efficient and numerically robust parameterisation for optimisation.

Piecewise Polynomial Spaces

Consider the interval $0 \leq \theta < 2\pi$. Rather than approximating the cam profile using a single polynomial over the entire interval, the domain is subdivided into a sequence of smaller intervals: $[t_0, t_1), [t_1, t_2), \ldots, [t_\mathrm{m}, t_{\mathrm{m}+1})$. These intervals are known as knot spans, and the collection $T = \{t_0, t_1, \ldots, t_\mathrm{m}\}$ is called the knot vector.

Within each knot span, the curve is represented by a polynomial of fixed degree. For a cubic spline, $p = 3$, each interval contains an independent third-order polynomial:

$$P_\mathrm{i}(\theta) = a_\mathrm{i} + b_\mathrm{i}\theta + c_\mathrm{i}\theta^2 + d_\mathrm{i}\theta^3$$

The complete spline is therefore assembled from many local polynomial pieces rather than one global polynomial. This seemingly minor change has profound implications for optimisation.

Local Basis Functions

Rather than expressing the spline in terms of the polynomial coefficients of each interval, B-splines employ a set of basis functions that span the underlying piecewise polynomial space. The curve is therefore written as:

$$C(\theta) = \sum_{i=0}^{n} N_{\mathrm{i},\,\mathrm{p}}(\theta) P_\mathrm{i}$$

Where:

- P_i are control points
- $N_{\mathrm{i},\,\mathrm{p}}$ are basis functions
- p is the spline degree

Each basis function is itself a polynomial. The remarkable property of these basis functions is that each one is non-zero over only a small portion of the total domain. This property is called local support. Consequently,

moving one control point alters only a small region of the curve while leaving the remainder completely unchanged.

Constructing the Basis Functions

The B-spline basis functions are defined recursively using the Cox–de Boor recursion relation. The recursion begins with piecewise constant basis functions:

$$N_{i,0}(\theta) = 1, if\ t_i \leq \theta < t_{i+1} N_{i,0}(\theta) = 0, otherwise$$

These functions simply identify whether the parameter lies inside a particular knot span. Higher - order basis functions are then generated recursively:

$$N_{i,p}(\theta) = \left[\frac{\theta - t_i}{t_{i+p} - t_i}\right] N_{i,p-1}(\theta) + \left[\frac{t_{i+p+1} - \theta}{t_{i+p+1} - t_{i+1}}\right] N_{i+1,p-1}(\theta)$$

This equation is known as the Cox–de Boor recursion formula. Although it appears algebraically complicated, its geometric interpretation is surprisingly intuitive. Each higher-order basis function is simply a weighted interpolation between two lower-order basis functions. The weights vary continuously across each knot span. Repeated application of this recursion progressively smooths the piecewise constant functions until a continuously differentiable cubic basis is obtained.

Why Cubic Splines are Smooth

Every recursion step increases both the polynomial degree and the continuity of the resulting curve. For degree p, the spline possesses C^{p-1} continuity provided no knot is repeated. Consequently:

- Linear: C^0 continuity
- Quadratic: C^1 continuity
- Cubic: C^2 continuity

Since piston acceleration depends upon the second derivative of the cam profile, continuous acceleration requires at least C^2 continuity. Thus, cubic B-splines possess continuous position, first derivative, and second derivative throughout the parameter domain.

Partition of Unity

One of the defining properties of B-spline basis functions is that they always satisfy:

$$\sum_i N_{i,p}(\theta) = 1$$

for every value of θ. This is known as the partition of unity property. The consequence is extremely important. Suppose every control point is translated by the same vector, $P_i \rightarrow P_i + d$. Then:

$$C(\theta) = \sum_i N_i (P_i + d) = \sum_i N_i P_i + d \sum_i N_i$$

Since the basis functions sum to one, $C(\theta) = C(\theta) + d$. Therefore, translating every control point simply translates the entire curve. The curve neither stretches nor distorts. This desirable behaviour follows entirely from the partition of unity property.

Local Support

Unlike ordinary polynomial interpolation, each cubic basis function is non-zero over only four consecutive knot spans. Everywhere else, $N_{i,3}(\theta) = 0$. Consequently, moving one control point affects only four neighbouring spline segments. Consequently, each basis function has compact support over only four consecutive knot

spans. This locality is one of the defining mathematical characteristics of cubic B-splines and plays a central role in their numerical efficiency.

Convex Hull Property

Another important consequence of positive basis functions is that the spline always lies inside the convex hull of its neighbouring control points. Mathematically, $N_i(\theta) \geq 0$, combined with $\Sigma_i N_i = 1$, implies that every point on the spline is a convex combination of nearby control points. Therefore, the spline cannot exhibit uncontrolled oscillations between control points. This behaviour stands in sharp contrast to high - order polynomial interpolation, which may oscillate dramatically despite passing exactly through every data point.

Periodic Splines

To represent a closed parametric curve, periodic boundary conditions are imposed upon the spline. Consequently, the spline itself must be periodic. Rather than terminating at $\theta = 2\pi$, the curve wraps continuously back onto itself. Mathematically:

- $C(0) = C(2\pi)$
- $C'(0) = C'(2\pi)$
- $C''(0) = C''(2\pi)$

For cubic splines, periodicity is commonly enforced by wrapping the first three control points and constructing a compatible periodic knot vector. The resulting basis functions remain continuously differentiable across the closure point, ensuring that the spline is mathematically periodic.

Differentiating the Cam Profile

One of the principal advantages of B-splines is that differentiation is straightforward. Differentiating $C(\theta) = \Sigma_i N_i(\theta) P_i$ gives:

$$\frac{dC}{d\theta} = \sum_i P_i \left(\frac{dN_i}{d\theta}\right)$$

Likewise:

$$\frac{d^2C}{d\theta^2} = \sum_i P_i \frac{\left(\frac{d^2N_i}{d\theta^2}\right) d^3C}{d\theta^3} = \sum_i P_i \left(\frac{d^3N_i}{d\theta^3}\right)$$

Notice that the control points themselves remain constant. Only the basis functions are differentiated. This observation greatly simplifies numerical implementation because the derivatives of the basis functions may be precomputed once and reused throughout every optimisation iteration.

Summary

This section established the mathematical formulation of the periodic cubic B-spline used to parameterise the cam profile. Beginning with piecewise polynomial spaces, the Cox–de Boor recursion was introduced to construct the spline basis functions. The resulting representation possesses compact local support, partition of unity, convex hull preservation, and C^2continuity, while remaining analytically differentiable to arbitrary order. These properties make the spline particularly well suited to optimisation and provide the mathematical foundation required to derive the piston kinematics presented in the following section. Having established the mathematical representation of the cam, the following section derives the piston kinematics by relating the

Piston Kinematics

The previous sections established a smooth periodic representation of the cam profile using a cubic B-spline. While the cam geometry itself is important, it is not the quantity directly responsible for hydraulic energy conversion. Instead, the primary role of the cam is to prescribe the motion of the pistons.

As the cylinder block rotates, each piston remains in contact with the cam surface. Variations in cam radius force the pistons to move radially inward and outward, producing cyclic changes in chamber volume. These volume changes drive the hydraulic dynamics that ultimately generate torque.

The objective of this section is therefore to derive the piston displacement, velocity, acceleration, and jerk directly from the spline representation of the cam.

Physical Interpretation of the Cam

Consider a single piston located within the rotating cylinder block. As the shaft rotates, the piston axis remains fixed relative to the cylinder block while the cam remains fixed relative to the housing. Consequently, the piston experiences a varying radial distance between the shaft centre and the cam surface. This varying radius determines the instantaneous piston position.

If the cam radius increases, the piston is forced outward. If the cam radius decreases, the piston retracts inward. The piston motion is therefore completely determined by the radial geometry of the cam.

Radial Cam Representation

The spline derived in the previous section defines a closed planar curve $C(\theta)$. For the purposes of hydraulic analysis, only the radial distance from the shaft centre is required. The radial coordinate is therefore defined as:

$$r(\theta) = \left|\left|C(\theta)\right|\right|$$

Expanding the Euclidean norm:

$$r(\theta) = \sqrt{x^2(\theta) + y^2(\theta)}$$

This quantity represents the instantaneous distance from the shaft centre to the cam surface. Since the cam is periodic, $r(\theta + 2\pi) = r(\theta)$.

Reference Radius

The absolute cam radius is not itself physically important. What determines chamber volume is the displacement of the piston relative to a reference position. Define a reference radius, r_0. The piston displacement is then:

$$s(\theta) = r(\theta) - r_0$$

The choice of reference is somewhat arbitrary. Common choices include minimum cam radius, mean cam radius, or neutral piston position. Provided the same reference is used consistently throughout the model, the resulting dynamics remain unchanged.

Periodic Nature of Piston Motion

Since the cam profile is periodic, piston motion is also periodic. Therefore, $s(\theta + 2\pi) = s(\theta)$. This property is extremely important because it permits the later use of Fourier analysis. Unlike many mechanical systems, the forcing function is known to repeat exactly every revolution. The displacement may therefore be represented as a harmonic series:

$$s(\theta) = \Sigma\,(n = -\infty\ to\ \infty)\hat{s}_n\, e^{in\theta}$$

Although the Fourier representation will not be used until later sections, it is useful to recognise at this stage that the kinematic quantities are inherently periodic.

Velocity

The piston velocity is obtained by differentiating displacement with respect to time. Applying the chain rule:

$$\dot{s} = \frac{ds}{dt} = \left(\frac{ds}{d\theta}\right)\left(\frac{d\theta}{dt}\right)$$

Since $\frac{d\theta}{dt} = \omega$, the velocity becomes:

$$\dot{s} = \omega \left(\frac{ds}{d\theta}\right)$$

This equation converts derivatives with respect to shaft angle into derivatives with respect to time. Because the spline representation is differentiable, $\frac{ds}{d\theta}$ may be obtained analytically. No numerical differentiation is required.

Physical Interpretation of Velocity

The velocity determines the instantaneous volumetric flow generated by piston motion. When $\dot{s} > 0$, the piston moves outward and chamber volume increases. When $\dot{s} < 0$, the piston retracts and chamber volume decreases. Since volume flow is directly proportional to piston velocity, the hydraulic dynamics derived later are highly sensitive to the shape of $\dot{s}(\theta)$. Even relatively small changes in cam geometry can therefore produce substantial changes in pressure ripple.

Acceleration

Differentiating velocity once more gives the piston acceleration:

$$\ddot{s} = \frac{d}{dt}\left(\omega \frac{ds}{d\theta}\right)$$

Assuming constant rotational speed ($\dot{\omega} = 0$), this gives:

$$\ddot{s} = \omega^2 \left(\frac{d^2 s}{d\theta^2}\right)$$

This expression reveals an important scaling law. Acceleration increases with the square of rotational speed. Doubling shaft speed therefore quadruples inertial loading.

Physical Interpretation of Acceleration

Acceleration determines the inertial force acting on the piston. Through Newton's Second Law, $F_I = m_p \ddot{s}$, where m_p is the piston mass. Consequently, acceleration directly influences contact forces, vibration, stress, and dynamic torque ripple. The second derivative of the cam profile therefore plays a major role in the overall dynamic behaviour of the motor.

Jerk

The third derivative of displacement is known as jerk, $j = \frac{d^3 s}{dt^3}$. Applying the chain rule once more gives:

$$j = \omega^3 \left(\frac{d^3 s}{d\theta^3}\right)$$

Jerk measures the rate at which acceleration changes. Although jerk does not appear explicitly in the hydraulic equations, it provides an extremely useful measure of smoothness. Profiles with large jerk values tend to generate vibration, acoustic noise, dynamic impact loading, and numerical stiffness. For this reason jerk is later included as a regularisation term within the optimisation objective function.

Phase Shift Between Pistons

The previous coordinate system section established that the pistons are evenly distributed around the cylinder block. For a motor containing N_p pistons, the displacement of piston i becomes:

$$s_{\mathrm{i}}(\theta) = s\left(\theta + \frac{2\pi i}{N_p}\right)$$

Differentiation immediately yields:

$$\dot{s}_{\mathrm{i}}(\theta) = \omega \frac{d}{d\theta}\left[s\left(\theta + \frac{2\pi i}{N_p}\right)\right]$$

and similarly for acceleration and jerk. Thus, once the motion of a single piston is known, the motion of every other piston follows through a simple phase shift. This symmetry substantially reduces computational effort because the complete kinematic state of the motor may be reconstructed from a single reference piston.

Kinematic Constraints

Not every mathematically valid spline corresponds to a physically feasible cam. Several constraints must therefore be imposed during optimisation.

Stroke Constraint: The piston stroke is defined as $S = s_{max} - s_{min}$. This quantity must remain compatible with the available cylinder length.

Velocity Constraint: Excessive piston velocity increases hydraulic losses and may induce cavitation. Therefore,

$$|\dot{s}| \leq \dot{s}_{max}$$

Acceleration Constraint: Large accelerations generate excessive inertial loading. Therefore,

$$|\ddot{s}| \leq \ddot{s}_{max}.$$

Jerk Constraint: Excessive jerk produces undesirable dynamic behaviour. Consequently,

$$|j| \leq j_{max}.$$

These constraints ensure that the optimisation algorithm searches only within physically meaningful regions of the design space.

Summary

This section established the kinematic relationship between the cam geometry and piston motion. Beginning with the radial representation of the periodic spline, expressions were derived for piston displacement, velocity, acceleration, and jerk. Through repeated application of the chain rule, all temporal derivatives were expressed directly in terms of angular derivatives of the cam profile. The resulting equations reveal how rotational speed amplifies inertial effects and provide the foundation for the hydraulic volume calculations developed in the following section.

Having determined the piston trajectory, the next step is to establish how this motion alters chamber volume. This volume evolution forms the geometric link between piston kinematics and hydraulic pressure generation.

Chamber Volume and Control Volume Formulation

The previous section established the kinematic relationship between the cam profile and piston motion. While the piston displacement completely describes the mechanical motion of the reciprocating assembly, the hydraulic behaviour depends not on the displacement itself but on the volume enclosed within the piston chamber.

As the piston moves, the chamber continuously expands and contracts. During expansion, fluid must enter the chamber to occupy the newly created volume. During contraction, fluid is expelled through the valve plate. If the incoming and outgoing flow cannot exactly match the volumetric change imposed by the piston motion, the pressure within the chamber changes accordingly.

The objective of this section is therefore to derive the chamber volume as a function of piston displacement and establish the control volume used throughout the remainder of the hydraulic analysis.

The Hydraulic Control Volume

A fundamental principle of fluid mechanics is that conservation laws are applied not to individual fluid particles but to a control volume. For the present analysis, each piston chamber is treated as an independent control volume bounded by:

- The piston face
- The cylinder walls
- The valve plate
- The surrounding hydraulic fluid

Although fluid continuously enters and leaves this region, the control volume itself moves with the piston. Within this control volume we seek to determine the evolution of chamber volume, fluid mass, pressure, inflow, and outflow.

The simplifying assumption adopted throughout this work is that pressure is spatially uniform within each chamber. This converts the governing equations from partial differential equations into ordinary differential equations while retaining the dominant compressibility effects.

Geometry of the Chamber

Consider a piston of constant circular cross-sectional area A_p. If the piston moves through a differential displacement ds, the corresponding differential change in chamber volume is simply $dV = A_p ds$. This result follows directly from elementary geometry. Since the piston area remains constant, every infinitesimal piston movement sweeps out a cylindrical volume whose base area is A_p and whose height is ds. Integrating both sides gives:

$$V = A_p s + C$$

where C is an integration constant.

Dead Volume

The integration constant possesses an important physical interpretation. Even when the piston reaches its minimum displacement, a finite quantity of fluid remains trapped inside the chamber. This residual volume arises from:

- Machining clearances
- Valve plate recesses
- Sealing grooves

- Manufacturing tolerances
- Geometric limitations preventing complete piston closure

This quantity is known as the dead volume, V_{dead}. Substituting this physical interpretation into the previous equation gives the chamber volume:

$$V(\theta) = V_{dead} + A_p s(\theta)$$

This equation forms one of the central geometric relationships of the entire hydraulic model.

Physical Interpretation

The chamber volume consists of two independent contributions. The first is a constant offset, V_dead, which represents the minimum attainable chamber volume. The second is the variable swept volume, A_p s(θ), generated directly by piston motion. As the cam rotates, the piston periodically increases and decreases the chamber volume, causing the fluid to undergo alternating expansion and compression. This cyclic volume variation is ultimately responsible for the periodic pressure oscillations analysed later.

Chamber Volume Rate

Hydraulic pressure depends not only upon chamber volume, but also upon the rate at which the volume changes. Differentiating the chamber volume with respect to time gives:

$$\frac{dV}{dt} = A_p \left(\frac{ds}{dt}\right)$$

or equivalently, since the piston area is constant:

$$\dot{V} = A_p \dot{s}$$

Substituting the kinematic relationship derived previously, $\dot{s} = \omega \left(\frac{ds}{d\theta}\right)$, gives:

$$\dot{V} = A_p \omega \left(\frac{ds}{d\theta}\right)$$

This equation establishes the direct relationship between cam geometry and volumetric flow generation.

Swept Volume

The total volume displaced during one complete piston cycle is known as the swept volume. If s_{max} and s_{min} denote the maximum and minimum piston positions, then:

$$V_{swept} = A_p(s_{max} - s_{min})$$

The swept volume determines the theoretical fluid displacement of a single piston during one complete reciprocating cycle. For a motor containing N_p pistons, the total theoretical displacement per revolution becomes:

$$V_D = N_p A_p(s_{max} - s_{min})$$

This quantity is commonly referred to as the geometric displacement of the hydraulic motor. In the absence of leakage and compressibility, it represents the ideal fluid volume transferred during one shaft revolution.

Chamber Volume Constraints

The chamber volume must remain strictly positive throughout the operating cycle. Therefore, $V(\theta) > 0$ for every θ. Substituting the chamber volume equation:

$$V_{dead} + A_p s(\theta) > 0$$

This constraint places a lower bound on the permissible dead volume during optimisation. Violation of this condition would imply a physically impossible negative fluid volume.

Influence of Dead Volume

Although increasing the dead volume guarantees positive chamber volume, it also reduces hydraulic stiffness. This relationship becomes evident later through the compressibility equation, where pressure dynamics depend upon $\frac{\beta}{V}$. Large dead volumes therefore produce slower pressure response, reduced volumetric efficiency, and increased compressibility effects.

Conversely, extremely small dead volumes improve stiffness but increase sensitivity to manufacturing tolerances and may promote cavitation. The optimisation process must therefore balance these competing effects.

Coupling with the Cam Geometry

At this stage, the chain of physical relationships developed throughout the document can be summarised as:

$$x \rightarrow C(\theta) \rightarrow s(\theta) \rightarrow V(\theta)$$

The optimisation variables define the cam geometry. The cam geometry determines piston displacement. The piston displacement determines chamber volume. The chamber volume will now become the primary state variable governing the hydraulic dynamics.

Summary

This section derived the chamber volume directly from the geometry of the piston and established the control volume used throughout the hydraulic analysis. Beginning with the differential relationship between piston displacement and swept volume, the chamber volume equation was obtained by integration and interpreted in terms of dead volume and swept volume. Expressions for the chamber volume rate and theoretical motor displacement were subsequently derived, providing the geometric foundation required for the conservation of mass.

Having established the evolution of chamber volume, the next section introduces the principle of mass conservation and derives the governing pressure equation for the compressible hydraulic fluid. This marks the transition from purely geometric modelling to fluid dynamics.

Conservation of Mass Within the Hydraulic Chamber

The previous section established the geometric relationship between piston motion and chamber volume. However, a changing chamber volume alone does not determine the hydraulic pressure. Pressure is generated by the interaction between the moving boundary of the chamber and the finite compressibility of the hydraulic fluid.

The fundamental principle governing this interaction is conservation of mass. Fluid cannot be created or destroyed within the hydraulic chamber. Therefore, any change in the mass contained inside the chamber must be caused by a difference between the mass entering and leaving through the valve ports.

This section develops the governing mass balance for a single piston chamber. The resulting equation forms the foundation upon which the pressure dynamics are derived.

Control Volume Mass Balance

Consider a hydraulic chamber represented as a moving control volume. The total mass of fluid contained within the chamber is $m = \rho V$, where ρ is the fluid density and V is the instantaneous chamber volume. Applying conservation of mass gives:

$$\frac{dm}{dt} = \dot{m}_{in} - \dot{m}_{out}$$

The right - hand side represents the net mass flow entering the chamber. Define $\dot{m}_{net} = \dot{m}_{in} - \dot{m}_{out}$. Therefore:

$$\frac{dm}{dt} = \dot{m}_{net}$$

Expanding the Mass Derivative

Since the fluid mass is given by $m = \rho V$, differentiation gives:

$$\frac{dm}{dt} = \frac{d}{dt}(\rho V)$$

Applying the product rule:

$$\frac{dm}{dt} = V\left(\frac{d\rho}{dt}\right) + \rho\left(\frac{dV}{dt}\right)$$

Therefore, the conservation equation becomes:

$$V\,\dot{\rho} + \rho\,\dot{V} = \dot{m}_{net}$$

This equation represents the complete mass balance for the hydraulic chamber. The first term describes changes in fluid density. The second term describes changes in the physical chamber volume.

Conversion Between Mass Flow and Volumetric Flow

Hydraulic systems are conventionally described using volumetric flow rather than mass flow. The relationship between mass flow and volumetric flow is $\dot{m} = \rho Q$, where Q is volumetric flow rate. Therefore, $\dot{m}_{net} = \rho\, Q_{net}$. Substituting gives:

$$V\,\dot{\rho} + \rho\,\dot{V} = \rho\, Q_{net}$$

Dividing by density:

$$\left(\frac{V}{\rho}\right)\dot{\rho} + \dot{V} = Q_{net}$$

or equivalently:

$$\dot{V} = Q_{net} - \left(\frac{V}{\rho}\right)\dot{\rho}$$

This equation states that the physical chamber expansion rate is supplied by two mechanisms:

1. External fluid entering or leaving the chamber.
2. Compression or expansion of the fluid itself.

The second effect is often neglected when modelling incompressible fluids. However, it is the dominant mechanism responsible for pressure dynamics in hydraulic systems.

The Incompressible Limit

Before introducing compressibility, it is useful to examine the limiting case. For an incompressible fluid, $\rho = constant$, therefore $\dot{\rho} = 0$. The mass equation reduces to:

$$V = Q_{net}$$

This result is physically intuitive. If the fluid density cannot change, any increase in chamber volume must be exactly matched by fluid entering the chamber. Likewise, if the chamber volume decreases, fluid must leave at the same rate. However, real hydraulic fluids are not perfectly incompressible. At high pressures, even small density variations generate significant pressure changes.

Introducing Fluid Compressibility

The relationship between pressure and density is described through the bulk modulus of the fluid. The bulk modulus is defined as $\beta = \rho \left(\frac{dP}{d\rho}\right)$. Equivalently:

$$\frac{d\rho}{dt} = \left(\frac{\rho}{\beta}\right)\left(\frac{dP}{dt}\right)$$

This equation describes how rapidly fluid density changes in response to pressure changes. For hydraulic oil, β is typically very large, meaning the fluid undergoes only a small volume change even under substantial pressure. However, because hydraulic systems operate at high pressures, this small deformation cannot be neglected.

Substituting the Compressibility Relationship

Returning to the mass balance: $V\,\dot{\rho} + \rho\,\dot{V} = \rho\,Q_{net}$

Substitute $\dot{\rho} = \left(\frac{\rho}{\beta}\right)\dot{\mathrm{P}}$. This gives:

$$V\left(\left(\frac{\rho}{\beta}\right)\dot{\mathrm{P}}\right) + \rho\,\dot{V} = \rho\,Q_{net}$$

Dividing through by ρ:

$$\left(\frac{V}{\beta}\right)\dot{\mathrm{P}} + \dot{V} = Q_{net}$$

Rearranging:

$$\dot{\mathrm{P}} = \left(\frac{\beta}{V}\right)(Q_{net} - \dot{V})$$

This is the fundamental pressure evolution equation for a compressible hydraulic chamber.

Physical Interpretation

The pressure equation can be interpreted directly from the balance of volume flow. The quantity Q_{net} represents the external fluid supplied to the chamber. The quantity $\dot{V}$ represents the volume demand created by piston motion. Therefore, the difference, $Q_{net} - \dot{V}$, represents the volumetric mismatch.

If $Q_{net} > \dot{V}$, more fluid enters than the piston requires. The additional fluid compresses, increasing pressure: $\dot{\mathrm{P}} > 0$. Conversely, if $Q_{net} < V$, the piston expands the chamber faster than fluid can enter. Pressure decreases: $\dot{\mathrm{P}} < 0$.

Effective Hydraulic Capacitance

The pressure equation may be rewritten as:

$$\dot{P} = \left(\frac{1}{C_h}\right)(Q_{net} - \dot{V})$$

where:

$$C_h = \frac{V}{\beta}$$

is the hydraulic capacitance. This quantity describes how much volume change is required to produce a given pressure change. A large chamber volume gives $C_h \uparrow$, meaning the system is more compliant and pressure changes more slowly. A large bulk modulus gives $C_h \downarrow$, meaning the fluid behaves more stiffly and pressure responds rapidly.

Incorporating Piston Motion

From the previous section, $\dot{V} = A_p\dot{s}$. Substituting into the pressure equation gives:

$$\dot{P} = \left(\frac{\beta}{V(\theta)}\right)(Q_{net} - A_p\dot{s})$$

This equation represents the fundamental coupling between mechanical motion and hydraulic behaviour. The cam geometry determines $s(\theta)$, which determines $\dot{s}(\theta)$, which directly affects the chamber pressure.

Summary

This section derived the hydraulic chamber mass balance from first principles. Beginning with conservation of mass, the relationship between chamber volume change, fluid flow, and density variation was established. Introducing the fluid bulk modulus transformed this relationship into the governing pressure evolution equation.

The final result, $\dot{P} = \left(\frac{\beta}{V(\theta)}\right)(Q_{net} - A_p\dot{s})$, forms the central hydraulic dynamic equation of the model. It reveals the fundamental mechanism responsible for pressure ripple: any mismatch between the flow supplied through the valve ports and the instantaneous volumetric demand created by piston motion appears directly as a pressure fluctuation.

The following section develops the valve port flow model required to determine Q_{net}, completing the hydraulic subsystem.

Valve Plate Flow and Port Timing

The previous section derived the pressure evolution equation by applying conservation of mass to the hydraulic chamber. However, the pressure equation contains one unknown quantity: Q_{net}. This term represents the net volumetric flow entering or leaving the chamber through the valve plate.

In a radial piston hydraulic motor, this flow is not externally controlled by an independent valve. Instead, it is determined by the relative position between the rotating cylinder block and the stationary valve plate. As the shaft rotates, each piston chamber periodically connects to the high-pressure supply port, the low-pressure return port, or neither port during transition regions. The resulting flow therefore depends on both the pressure difference and the geometric overlap between the chamber port and the valve plate opening.

This section develops the mathematical representation of this process.

Flow Through an Orifice

The fundamental relationship between pressure difference and hydraulic flow is obtained from Bernoulli's equation. Consider an ideal incompressible fluid flowing through a restriction. The conservation of energy equation between upstream and downstream locations is:

$$P_1 + \frac{1}{2}\rho v_1^2 = P_2 + \frac{1}{2}\rho v_2^2$$

Assuming the upstream velocity is small compared with the velocity through the restriction ($v_1 \approx 0$), this gives:

$$P_1 - P_2 = ½\rho v_2^2$$

Solving for velocity:

$$v_2 = \sqrt{\frac{2(P_1 - P_2)}{\rho}}$$

The volumetric flow rate is velocity multiplied by area: Q = Av. Therefore:

$$Q = A\sqrt{\frac{2(P_1 - P_2)}{\rho}}$$

Real Orifice Flow

The previous derivation assumes an ideal restriction. Real hydraulic ports contain losses caused by viscous effects, flow separation, turbulence, and edge geometry. These effects are represented using the discharge coefficient, C_d. The practical flow relationship becomes:

$$Q = C_d A\sqrt{\frac{2\Delta P}{\rho}}$$

where $\Delta P = |P_1 - P_2|$. The discharge coefficient therefore represents the difference between ideal energy conversion and the actual behaviour of the restriction. For a well - designed hydraulic port, C_d is approximately constant over the operating range. This assumption is adopted throughout this model.

Valve Plate Geometry

The valve plate determines when each piston chamber connects to the supply and return ports. Consider a single chamber rotating relative to the stationary valve plate. The relative angular position is θ. The valve plate contains regions corresponding to inlet connection, outlet connection, and closed transition regions. Therefore, the hydraulic connection state is a periodic function, $\alpha(\theta)$. This function represents the fraction of the port opening available to the chamber at a given shaft position.

Effective Port Area

The geometric port area is represented by A_{port}. However, the complete area is not always available. During partial overlap only a fraction of the port is open. Therefore, the effective flow area becomes:

$$A(\theta) = A_{port}\alpha(\theta)$$

where $0 \leq \alpha(\theta) \leq 1$. The overlap function acts as a geometric switching function between fully closed and fully open states.

Port Overlap Function

For a simple idealised valve plate, the overlap function may be represented as a piecewise function. For example:

- $\alpha(\theta) = 0$ (closed region)
- $\alpha(\theta) = 1$ (fully open region)
- $0 < \alpha(\theta) < 1$ (transition region)

However, a discontinuous switching function introduces unrealistic flow impulses. Real valve ports have finite opening and closing times. Therefore, a smooth approximation is preferred.

Linear Window Approximation

The overlap region may be approximated using a triangular window function. For a port centred at angular position θ_p, with angular width $\Delta\theta_p$, the contribution of that port is:

$$\alpha_{p(\theta)} = \max\left(0, 1 - \left[\frac{|\theta - \theta_p|}{\left(\frac{\Delta\theta_p}{2}\right)}\right]\right)$$

This produces zero flow outside the port region, maximum overlap at the centre, and linear opening and closing transitions. For multiple ports, the complete overlap function becomes:

$$\alpha(\theta) = \sum_{p=1}^{N_{ports}} \max\left(0, 1 - \left[\frac{|\theta - \theta_p|}{\left(\frac{\Delta\theta_p}{2}\right)}\right]\right)$$

This representation preserves the periodic nature of the valve plate while remaining computationally inexpensive.

Port Timing

Each port is positioned according to:

$$\theta_p = \theta_{shift} + \frac{2\pi\rho}{N_{ports}}$$

The shift angle, θ_{shift}, controls the timing relationship between piston motion and hydraulic connection. Changing this parameter modifies pressure rise timing, torque production, pressure ripple, and volumetric efficiency. Therefore, valve timing is included as an optimisation parameter.

Chamber Flow Equation

The pressure difference across the valve plate determines the flow direction. For a supply connection ($P_s > P$), the flow enters the chamber:

$$Q_{in} = C_d A(\theta)\sqrt{\frac{2(P_s - P)}{\rho}}$$

For a return connection ($P > P_t$), fluid leaves the chamber:

$$Q_{out} = C_d A(\theta)\sqrt{\frac{2(P - P_t)}{\rho}}$$

The net chamber flow becomes:

$$Q_{net} = Q_{in} - Q_{out}$$

Substituting the port area relationship:

$$Q_{net} = C_d A_{port} \alpha(\theta) \sqrt{\frac{2\Delta P}{\rho}}$$

with the sign determined by the pressure difference.

Incorporating Leakage

A perfectly sealed hydraulic chamber does not exist. Small clearances between the piston and cylinder, and the valve plate and cylinder block, allow internal leakage. A first - order approximation models leakage as proportional to pressure difference:

$$Q_{leak} = C_l(P - P_t)$$

where C_l is the leakage coefficient. The net hydraulic flow therefore becomes:

$$Q_{net} = Q_{port} - Q_{leak}$$

or explicitly:

$$Q_{net} = C_d A(\theta) \sqrt{\frac{2\Delta P}{\rho}} - C_l(P - P_t)$$

Complete Hydraulic Pressure Equation

Substituting the flow model into the pressure equation derived previously, $\dot{P} = \left(\frac{\beta}{V}\right)(Q_{net} - A_p \dot{s})$, gives:

$$\dot{P} = \left(\frac{\beta}{V(\theta)}\right)\left[C_d A(\theta) \sqrt{\frac{2\Delta P}{\rho}} - C_l(P - P_t) - A_p \dot{s} \right]$$

This equation represents the complete nonlinear hydraulic subsystem. The pressure evolution is determined by three competing effects:

1. Valve plate flow supply: $C_d A(\theta) \sqrt{\frac{2\Delta P}{\rho}}$
2. Leakage losses: $C_l(P - P_t)$
3. Volumetric demand from piston motion: $A_p \dot{s}$

Model Interpretation

The equation above demonstrates the fundamental coupling within the hydraulic motor. The cam determines piston velocity, ṡ(θ). The piston velocity determines the chamber volume demand. The valve plate determines available flow. The difference between supplied flow and required flow generates pressure. Pressure then creates hydraulic force on the piston. Therefore, the complete causal chain is:

$$C(\theta) \rightarrow s(\theta) \rightarrow \dot{s}(\theta) \rightarrow P(\theta) \rightarrow F(\theta) \rightarrow \tau(\theta)$$

Summary

This section derived the valve flow model required to complete the hydraulic subsystem. Starting from Bernoulli's equation, the orifice flow relationship was obtained and modified using a discharge coefficient to account for real hydraulic losses. The valve plate geometry was then represented through a smooth overlap function, allowing the changing hydraulic connection state to be incorporated directly into the model.

The resulting pressure equation now contains all major hydraulic effects: compressibility, port timing, valve geometry, leakage, and piston - induced volume change.

The next section introduces the mechanical force balance, where chamber pressure is converted into piston force and ultimately rotational torque.

Piston Force Transmission and Mechanical Dynamics

The previous sections developed the complete hydraulic subsystem. The cam geometry determines piston motion, the piston motion determines chamber volume variation, and the valve plate determines the hydraulic flow entering or leaving the chamber. These interactions produce a time-varying chamber pressure.

However, pressure alone does not produce useful mechanical output. The hydraulic pressure must generate a force on the piston, which is then transmitted through the cam mechanism to produce shaft torque. The piston therefore represents the coupling point between the hydraulic and mechanical domains.

This section develops the mechanical equations governing piston motion and derives the equivalent dynamic behaviour of the hydraulic actuator.

Hydraulic Force Generation

The pressure acting on the piston face produces a force according to Pascal's law. For a piston with effective area A_p, the hydraulic force is:

$$F_h = A_p P$$

where P is the instantaneous chamber pressure relative to the return pressure. More generally, $F_h = A_p(P - P_t)$. For simplicity, the return pressure is incorporated into the pressure reference throughout this model. The hydraulic force acts along the piston axis and attempts to accelerate the piston according to the instantaneous pressure difference.

Free Body Description of the Piston

The piston experiences several competing forces. The primary forces are:

- Hydraulic force: $F_h = A_p P$, which drives piston motion.
- Inertial force: $F_I = m_p \ddot{s}$, which opposes changes in piston velocity.
- Friction force: F_f, which resists motion due to mechanical losses.
- Contact force: F_c, which transmits the piston force into the cam surface.

The complete force balance may therefore be written as:

$$F_h - F_f - F_I - F_c = 0$$

Newton's Second Law

Applying Newton's second law directly to the piston gives:

$$\sum F = m_p \ddot{s}$$

Taking the positive radial direction as the direction of piston extension:

$$A_p P - F_f - F_c = m_p \ddot{s}$$

Rearranging:

$$m_p \ddot{s} = A_p P - F_f - F_c$$

This equation represents the mechanical dynamics of the piston.

Friction Modelling

Mechanical friction within hydraulic pistons is complex. The true friction force depends on lubrication conditions, velocity, temperature, surface roughness, and pressure. A complete friction model may include Coulomb friction, viscous friction, and Stribeck effects. For optimisation purposes, a reduced-order representation is preferred. The present model uses linear viscous damping:

$$F_f = c_m \dot{s}$$

where c_m is the equivalent viscous damping coefficient. Substituting this into the force balance:

$$m_p \ddot{s} = A_p P - c_m \dot{s} - F_c$$

Cam Reaction Force

The piston does not directly create torque. Instead, the piston pushes against the cam surface. The cam converts radial piston force into tangential shaft torque. Consider the cam pressure angle, $\varphi(\theta)$, defined as the angle between the piston force direction and the normal direction of the cam surface. The tangential component of force is:

$$F_t = F_c sin\varphi$$

The instantaneous torque contribution of a single piston is therefore:

$$\tau_i = r(\theta) F_c sin\varphi(\theta)$$

The total motor torque is obtained by summing the contribution of every piston:

$$\tau = \sum_{i=0}^{n} \tau_i$$

This is the mechanical origin of torque ripple. Even if hydraulic pressure were perfectly constant, variations in cam geometry would still produce torque variation through the changing mechanical leverage.

Simplified Dynamic Model

For many system - level analyses, the detailed cam reaction force is not required. Instead, the piston dynamics may be expressed in terms of the hydraulic force required to generate the prescribed displacement. The mechanical equation becomes:

$$m_p \ddot{s} + c_m \dot{s} = A_p P$$

This represents the piston as a damped hydraulic actuator. The equation reveals that piston acceleration depends upon the balance between hydraulic driving force, viscous losses, and inertial effects.

Hydraulic Compressibility as a Mechanical Spring

The most important dynamic effect introduced by fluid compressibility is hydraulic stiffness. Consider a piston displaced by a small amount, Δs. The corresponding chamber volume change is $\Delta V = A_p \Delta s$. The bulk modulus relationship is:

$$\beta = -\frac{\Delta P}{\left(\frac{\Delta V}{V}\right)}$$

Rearranging:

$$\Delta P = -\left(\frac{\beta}{V}\right)\Delta V$$

Substituting:

$$\Delta P = -\left(\frac{\beta A_p}{V}\right)\Delta s$$

The corresponding hydraulic force change is $\Delta F_h = A_p \Delta P$. Therefore:

$$\Delta F_h = -\left(\frac{\beta A_p^2}{V}\right)\Delta s$$

Comparing this with Hooke's law ($F = -kx$), the equivalent hydraulic stiffness becomes:

$$k_f = \frac{\beta A_p^2}{V}$$

Physical Meaning of Hydraulic Stiffness

The hydraulic fluid behaves as a mechanical spring. When the piston compresses the chamber, fluid pressure increases, and the fluid pushes back against the piston. When the piston retracts, pressure decreases, and the fluid assists piston motion.

The stiffness is increased by an increase in β, because a less compressible fluid resists deformation. The stiffness is increased by an increase in A_p, because a larger piston converts pressure into greater force. The stiffness is reduced by an increase in V, because a larger fluid volume provides greater compliance.

Equivalent Mass - Spring - Damper System

Including hydraulic stiffness, the piston dynamics become:

$$m_p \ddot{s} + c_m \dot{s} + k_f s = F_{hyd}$$

where F_{hyd} represents the externally applied hydraulic excitation. This is mathematically identical to a classical second - order mechanical oscillator. The hydraulic system therefore possesses natural dynamic characteristics despite containing no conventional mechanical spring.

Natural Frequency

For a mass-spring system, the natural frequency is:

$$\omega_n = \sqrt{\frac{k_f}{m_p}}$$

Substituting the hydraulic stiffness:

$$\omega_n = \sqrt{\frac{\beta\, A_p^2}{V\, m_p}}$$

This equation provides important physical insight. The hydraulic natural frequency increases with fluid stiffness and piston area and decreases with chamber volume and piston mass.

Damping Ratio

The damping ratio is defined as:

$$\zeta = \frac{c_m}{2\sqrt{m_p k_f}}$$

Substituting the hydraulic stiffness:

$$\zeta = \frac{c_m}{2\sqrt{m_p\left(\frac{\beta A_p^2}{V}\right)}}$$

The damping ratio determines the dynamic response of the piston - fluid system. For $\zeta < 1$, the system is underdamped and exhibits oscillatory behaviour. For $\zeta = 1$, the system is critically damped. For $\zeta > 1$, the response is overdamped.

Connection to Torque Ripple

The mechanical dynamics introduce another source of periodic variation. Torque ripple arises from the interaction between:

1. Pressure variation, $P(\theta)$
2. Cam geometry, $C(\theta)$
3. Piston inertia, $m_p\ddot{s}$

The output torque may therefore be interpreted as a combined hydraulic - mechanical response:

$$\tau(\theta) = T\big(P(\theta), s(\theta), \dot{s}(\theta), \ddot{s}(\theta)\big)$$

This relationship motivates the frequency - domain analysis developed in later sections, where torque ripple is quantified through harmonic energy.

Summary

This section derived the mechanical subsystem connecting hydraulic pressure to output torque. Starting from Newton's second law, the piston force balance was developed, and frictional losses were introduced through a viscous damping approximation. The compressibility of the hydraulic fluid was then shown to produce an equivalent mechanical spring, yielding a classical mass - spring - damper representation.

The resulting equation, $m_p\ddot{s} + c_m\dot{s} + k_f s = F_{hyd}$, demonstrates that a hydraulic motor is not merely a fluid power device but a coupled dynamic system containing both hydraulic and mechanical energy storage.

The following section develops the frequency - domain representation of this coupled system, allowing periodic pressure and torque variations to be analysed through their harmonic content.

Frequency - Domain Representation of Periodic Hydraulic Dynamics

The previous sections developed the coupled mechanical and hydraulic equations governing the motor behaviour. These equations describe the evolution of piston position, chamber volume, pressure, and torque in the time domain.

However, the operating cycle of a hydraulic motor is inherently periodic. Every shaft revolution produces a repeating sequence of piston extension, piston retraction, chamber filling, chamber compression, pressure generation, and torque production.

This periodic behaviour means that the resulting signals are not arbitrary transient responses. Instead, they are composed of a finite set of repeating harmonic components. Representing these signals in the frequency

domain provides several advantages: torque ripple can be quantified directly, dominant excitation frequencies can be identified, optimisation objectives can target specific harmonic content, and computational cost is reduced compared with full transient simulation.

This section develops the Fourier representation of the hydraulic motor dynamics and derives the relationship between cam geometry, pressure harmonics, and torque ripple.

Periodic Signal Representation

Consider a generic periodic motor quantity, $x(\theta)$, where θ is shaft angle. Because the motor repeats every revolution, $x(\theta + 2\pi) = x(\theta)$. Any periodic function satisfying this condition may be represented using a Fourier series. The complex Fourier representation is:

$$x(\theta) = \sum_{-\infty}^{\infty} \hat{x}_n e^{\mathrm{in}\theta}$$

where $\hat{x}_n$ is the complex amplitude of the nth harmonic. The coefficient is defined as:

$$\hat{x}_n = \frac{1}{2\pi} \int_0^{2\pi} x(\theta) e^{-\mathrm{in}\theta} \, d\theta$$

Physical Interpretation of Harmonics

The Fourier series separates the motor behaviour into independent rotational frequencies. The zero - order term, $\hat{x}_0$, represents the average value. The first harmonic, $\hat{x}_1$, represents one oscillation per revolution. The second harmonic, $\hat{x}_2$, represents two oscillations per revolution.

In general, $\hat{x}_n$ represents an excitation occurring at $n\omega$ in physical time. Therefore, low-order harmonics correspond to large-scale variations in motor output, while higher-order harmonics correspond to finer geometric details.

Why Harmonics Determine Torque Ripple

The shaft torque generated by a hydraulic motor is periodic. Therefore:

$$\tau(\theta) = \sum_{n} \hat{\tau}_n e^{\mathrm{in}\theta}$$

The average torque is $\hat{\tau}_0$. All remaining components represent unwanted oscillations ($n \geq 1$). The torque ripple magnitude is therefore related to the energy contained within these harmonics. Using Parseval's theorem, the total signal energy is $E = \sum_n |\hat{x}_n|^2$. Therefore, the harmonic energy associated with ripple is:

$$E_{ripple} = \sum_{n=1}^{N} |\hat{x}_n|^2$$

where N represents the highest harmonic considered physically significant.

Transforming the Hydraulic Equation

The nonlinear hydraulic equation previously derived is:

$$\dot{\mathrm{P}} = \left(\frac{\beta}{V(\theta)}\right)\left(Q_{net} - A_p \dot{\mathrm{s}}\right)$$

For frequency - domain analysis, the system is locally linearised around a nominal operating point. Let $P = P_0 + \tilde{P}$, and $s = s_0 + \tilde{s}$. The perturbation pressure is $\tilde{P}$. The linearised hydraulic equation becomes:

$$\dot{\tilde{P}} + \left(\frac{\beta C_f}{V}\right)\tilde{P} = -\left(\frac{\beta A_p}{V}\right)\dot{\tilde{s}}$$

This equation represents the small-signal hydraulic response.

Fourier Transform of Pressure Dynamics

Represent the pressure perturbation as $\tilde{P}(\theta) = \sum_n \hat{P}_n e^{in\theta}$. Likewise, $\tilde{s}(\theta) = \sum_n \hat{s}_n e^{in\theta}$. The time derivative becomes $\frac{d}{dt} = \omega\left(\frac{d}{d\theta}\right)$. Therefore, $\frac{d}{dt}(e^{in\theta}) = in\omega e^{in\theta}$.

Substituting into the pressure equation gives:

$$in\omega\,\hat{P}_n + \left(\frac{\beta C_f}{V}\right)\hat{P}_n = -\left(\frac{\beta A_p}{V}\right)(in\omega)\hat{s}_n$$

Factoring the pressure term:

$$\left(in\omega + \frac{\beta C_f}{V}\right)\hat{P}_n = -\left(\frac{\beta A_p}{V}\right)in\omega\hat{s}_n$$

Therefore:

$$\hat{P}_n = -\frac{\left(\frac{\beta A_p}{V}\right)in\omega\hat{s}_n}{in\omega + \frac{\beta C_f}{V}}$$

Hydraulic Transfer Function

The previous expression may be written as $\hat{P}_n = G_{h(n)}\hat{s}_n$, where:

$$G_{h(n)} = -\frac{\left(\frac{\beta A_p}{V}\right)in\omega}{in\omega + \frac{\beta C_f}{V}}$$

is the hydraulic transfer function. This function describes how each cam - induced displacement harmonic is converted into a pressure harmonic.

Low - Frequency Behaviour

For low harmonic frequencies ($n\omega \ll \frac{\beta C_f}{V}$), the denominator is dominated by hydraulic damping. Therefore, $G_{h(n)} \approx -\left(\frac{A_p}{C_f}\right)in\omega$. Pressure variation is therefore relatively small. The fluid has sufficient time to flow through the valve and leakage paths.

High - Frequency Behaviour

At high frequencies ($n\omega \gg \frac{\beta C_f}{V}$), the denominator becomes dominated by inertial fluid stiffness. The transfer function approaches $G_{h(n)} \approx -\frac{\beta A_p}{V}$. The fluid cannot respond quickly enough through the ports. Instead, volume changes are accommodated by compression of the hydraulic fluid. This produces large pressure oscillations.

Pressure Ripple Metric

The pressure ripple can now be quantified directly from the harmonic coefficients. The RMS harmonic energy is:

$$E_P = \sum_{n=1}^{N} \left|\hat{P}_n\right|^2$$

Normalising against the average pressure gives:

$$R_P = \frac{\sqrt{\sum_{n=1}^{N} \left|\hat{P}_n\right|^2}}{\left|\hat{P}_0\right|}$$

This metric represents the relative magnitude of pressure oscillation.

Torque Harmonics

The torque signal is similarly represented as:

$$\tau(\theta) = \sum_{n} \hat{\tau}_n e^{\mathrm{in\theta}}$$

Torque is generated from hydraulic force: $F_h = A_p P$. Therefore, pressure harmonics directly contribute to torque harmonics. A simplified harmonic relationship is:

$$\hat{\tau}_n = \Gamma_n A_p \hat{P}_n$$

where Γ_n represents the geometric transformation from piston force to shaft torque. This term incorporates cam radius, pressure angle, piston phase relationships, and piston count.

Effect of Multiple Pistons

For a motor with N_p pistons, the total torque is $\tau = \sum_{i=0}^{N_p-1} \tau_i$. Each piston contributes a phase-shifted harmonic:

$$\hat{\tau}_{\{n,i\}} = \hat{\tau}_{\{n,0\}} e^{in\frac{2\pi i}{N_p}}$$

Therefore:

$$\hat{\tau}n = \hat{\tau}_{\{n,0\}} \sum_{i=0}^{N_p-1} e^{in\frac{2\pi i}{N_p}}$$

The summation term determines which harmonics reinforce and which cancel. This explains why piston count is a fundamental design parameter affecting torque smoothness.

Torque Ripple Metric

The equivalent torque ripple metric is:

$$R_\tau = \frac{\sqrt{\sum_{n-1}^{N} |\hat{\tau}_n|^2}}{|\hat{\tau}_0|}$$

Unlike average torque, which represents useful power output, the harmonic terms represent undesirable periodic excitation.

Spectral Objective Formulation

Since both pressure and torque ripple are fundamentally harmonic phenomena, the optimisation objective may be written directly in terms of spectral energy. A combined objective is defined as:

$$J_{spec} = \sum_{n=1}^{N} w_P \left|\hat{P}_n\right|^2 + w_\tau |\hat{\tau}_n|^2$$

where w_P and w_τ are weighting coefficients. This formulation allows the optimisation algorithm to reduce the physically important oscillatory components rather than only minimising a time - domain error.

Summary

This section developed the frequency - domain representation of the hydraulic motor model. By exploiting the periodic nature of cam - driven piston motion, the system was transformed from a time - domain dynamic model into a harmonic representation.

Fourier decomposition revealed that pressure ripple and torque ripple are fundamentally the result of harmonic energy contained within the periodic response. The linearised hydraulic dynamics were transformed into a harmonic transfer function describing how cam displacement generates pressure oscillations.

The resulting spectral formulation provides the foundation for the optimisation objective. Rather than simply reducing geometric variation, the optimisation algorithm can now directly minimise the harmonic components responsible for vibration, noise, and torque irregularity.

The following section combines the physical model into a complete deterministic objective function and introduces the constraints required to ensure physically feasible designs.

Deterministic Objective Function and Design Constraints

The previous sections developed the complete reduced - order physics model of the hydraulic motor. The model now provides a deterministic mapping between the design variables and the resulting motor performance. The complete physical chain may be expressed as:

$$x \rightarrow C(\theta) \rightarrow s(\theta) \rightarrow V(\theta) \rightarrow P(\theta) \rightarrow \tau(\theta)$$

where the design vector x contains all variables defining the motor geometry and operating characteristics.

The purpose of optimisation is to identify the design that provides the best compromise between competing objectives: low torque ripple, low pressure pulsation, high efficiency, smooth manufacturable geometry, and acceptable mechanical loading. This requires converting the physical model outputs into a single scalar objective function while enforcing physical feasibility constraints.

Design Vector Definition

The optimisation variables define the degrees of freedom available to the designer. The general design vector is written as:

$$x = \{C(\theta), A_{port}, \theta_{shift}, C_l, c_m, \ldots\}$$

In practice, the cam profile is not optimised directly as a continuous function. Instead, the cubic B-spline representation introduced previously is used. Therefore, $C(\theta) = C(\theta, p)$, where $p = \{p_1, p_2, \ldots, p_n\}$ contains the spline control parameters. The final optimisation vector therefore becomes:

$$x = \{p_1, p_2, \ldots, p_n, A_{port}, \theta_{shift}, C_l, \ldots\}$$

This parameterisation provides a compact design space while maintaining smooth geometry.

Deterministic Physics Evaluation

For any candidate design vector, the physics model evaluates the following sequence:

- Geometry generation: $x \rightarrow C(\theta)$. The cubic B-spline generates the periodic cam profile.
- Kinematic calculation: $C(\theta) \rightarrow s(\theta)$. The cam profile determines piston displacement. Velocity and acceleration are obtained analytically: $\dot{s} = \omega \left(\frac{ds}{d\theta}\right)$ and $\ddot{s} = \omega^2 \left(\frac{d^2s}{d\theta^2}\right)$.
- Hydraulic simulation: The chamber dynamics are solved: $\dot{P} = \left(\frac{\beta}{V}\right)\left(Q_{net} - A_p\dot{s}\right)$ to determine the pressure waveform.
- Torque calculation: The mechanical transformation produces $\tau(\theta)$.
- Spectral analysis: Fourier decomposition gives $\hat{P}_n$ and $\hat{\tau}_n$. These harmonic coefficients form the basis of the optimisation objective.

Objective Function Formulation

The optimisation problem is formulated as:

$$min_x J(x)$$

where the objective function represents total undesirable system behaviour. The objective combines several physical terms.

Torque Ripple Contribution

Torque ripple is quantified using:

$$R_\tau = \frac{\sqrt{\Sigma_n |\hat{\tau}_n|^2}}{|\hat{\tau}_0|}$$

The corresponding objective contribution is:

$$J_\tau = w_\tau R_\tau^2$$

Squaring provides two advantages: large ripple values are penalised more strongly, and the objective remains smooth for optimisation.

Pressure Ripple Contribution

Similarly,

$$R_P = \frac{\sqrt{\Sigma_n \left|\hat{P}_n\right|^2}}{\left|\hat{P}_0\right|}$$

The pressure contribution becomes:

$$J_P = w_P R_P^2$$

Reducing pressure ripple improves acoustic performance, valve loading, and structural fatigue behaviour.

Efficiency Contribution

The hydraulic efficiency is defined as $\eta = \frac{P_{out}}{P_{in}}$. Equivalently, $\eta = \frac{P_{hyd}}{P_{hyd} + P_{loss}}$. The loss components are represented as $P_{loss} = P_{fric} + P_{leak} + P_{visc}$. Since optimisation minimises the objective, inefficiency is represented as:

$$J_\eta = w_\eta(1 - \eta)$$

Cam Smoothness Regularisation

A mathematically valid cam profile may still be unsuitable for manufacturing or high - speed operation. To discourage aggressive geometric changes, jerk is penalised. The smoothness objective is:

$$J_{jerk} = \int_0^2 \pi \left(\frac{d^3 s}{d\theta^3}\right)^2 d\theta$$

This term discourages profiles with excessive curvature variation and high dynamic excitation.

Combined Physical Objective

The complete deterministic objective is therefore:

$$J_{phys} = w_1 R_\tau^2 + w_2 R_P^2 + w_3(1 - \eta) + w_4 J_{spec} + w_5 J_{jerk}$$

where the weighting coefficients determine the relative importance of each design goal. This formulation converts the multi - objective engineering problem into a scalar optimisation problem while preserving the underlying physical meaning.

Constraint Formulation

Optimisation must not only improve performance; the resulting design must also remain physically feasible. The constrained optimisation problem becomes:

$$\min_{\mathrm{x}} J_{phys}(x)$$

subject to:

$$g_i(x) \leq 0$$

where each constraint represents a physical limitation.

Cavitation Constraint

The chamber pressure must remain above the vapour pressure of the hydraulic fluid. Therefore, $P(\theta) \geq P_{vap}$. The constraint function is:

$$g_1(x) = P_{vap} - \min P(\theta)$$

A positive value indicates cavitation risk.

Mechanical Stress Constraint

The cam and piston assembly must remain below allowable stress. The maximum Hertzian/contact stress is represented as σ_H. The constraint becomes:

$$g_2(x) = \sigma_H - \sigma_{\max}$$

Dynamic Stability Constraint

The equivalent hydraulic spring system introduces a natural frequency and damping ratio. The damping ratio must remain above a minimum acceptable value ($\zeta \geq \zeta_{min}$). Therefore:

$$g_3(x) = \zeta_{min} - \zeta$$

Geometric Constraints

The cam geometry must also satisfy manufacturing requirements. Examples include:

- Stroke limit: $s_{min} \leq s(\theta) \leq s_{max}$
- Maximum velocity: $|\dot{s}| \leq \dot{s}_{max}$
- Maximum acceleration: $|\ddot{s}| \leq \ddot{s}_{max}$
- Periodicity: $C(0) = C(2\pi)$

Optimisation Landscape

The resulting objective landscape is highly nonlinear due to nonlinear valve flow, pressure - dependent leakage, harmonic interactions, geometric coupling, and manufacturing constraints. Additionally, the objective does not provide a convenient analytical gradient because every evaluation requires a complete physics simulation. Therefore, conventional gradient - based optimisation methods are poorly suited. A derivative - free optimisation framework is required.

Optimisation Strategy Overview

The final optimisation approach combines global exploration with local refinement. The process consists of two stages.

- **Global Search Using Gaussian Process Surrogate:** The expensive physics model is first approximated using a Gaussian Process surrogate. The surrogate learns the relationship $x \rightarrow J_{phys}$ from previously evaluated designs. The GP provides both a predicted objective value and an uncertainty estimate. This allows efficient exploration of promising regions without evaluating every candidate using the full physics model.
- **Local Refinement Using Nelder-Mead:** Once a promising region has been identified, a derivative-free local optimiser is used. The Nelder-Mead algorithm refines the design by evolving a simplex of candidate solutions. Unlike gradient-based approaches, Nelder-Mead requires only objective evaluations. This makes it suitable for the final refinement stage where the objective may contain penalties, gradients may be unreliable, and local geometric interactions dominate.

Complete Optimisation Loop

The final computational workflow is therefore:

$$x \rightarrow Physics\ Model \rightarrow J_{phys} \rightarrow Gaussian\ Process\ Search \rightarrow Nelder-Mead\ Refinement \rightarrow x_{opt}$$

The final design is validated using the full deterministic physics model.

Summary

This section converted the physical motor model into a constrained optimisation problem. The objective function combines harmonic ripple metrics, efficiency, and cam smoothness into a single physically meaningful cost function. Constraints ensure that all generated designs remain mechanically and hydraulically feasible.

Because each objective evaluation requires a complete nonlinear simulation, the optimisation problem is computationally expensive and unsuitable for direct brute - force search. The adopted solution therefore combines a Gaussian Process surrogate for efficient global exploration with Nelder - Mead derivative - free refinement for local optimisation.

The next section develops the Gaussian Process surrogate model and explains how it reduces the number of expensive physics evaluations required during optimisation.

Optimisation Framework

The deterministic physics model provides an accurate representation of the hydraulic motor behaviour; however, evaluating the objective function requires solving the coupled nonlinear hydraulic and mechanical equations for every candidate design. The optimisation problem therefore becomes computationally expensive. A direct search across the complete design space would require an impractical number of physics evaluations, particularly as the design vector contains multiple coupled geometric and hydraulic parameters.

To overcome this limitation, the optimisation framework uses a surrogate - assisted approach. The objective function is approximated using a Gaussian Process model, allowing the optimisation algorithm to explore the design space efficiently while minimising the number of expensive physics simulations.

Surrogate-Assisted Optimisation Strategy

The optimisation process consists of two interacting models:

Deterministic Physics Model: The full physics simulation evaluates $J_{phys}(x)$ for selected designs. This provides high-fidelity evaluations but is computationally expensive.

Gaussian Process Surrogate: The surrogate model approximates the objective landscape:

$$J(x) \sim GP\big(m(x), k(x, x')\big)$$

where $m(x)$ is the predicted mean objective value, and $k(x, x')$ describes the correlation between designs. The surrogate provides an estimate of both expected performance and uncertainty in the prediction.

Bayesian Optimisation Search

The purpose of the surrogate is not to replace the physics model. Instead, it determines where the next expensive physics evaluation should occur. At each iteration, the Gaussian Process predicts the objective landscape. The next candidate design is selected using an acquisition function that balances exploitation of known good regions and exploration of uncertain regions. The optimisation criterion is therefore:

$$x_{\{k+1\}} = argmin_x[\,\mu(x) - \kappa\sigma(x)]$$

where:

- $\mu(x)$ is the predicted objective value
- $\sigma(x)$ is the prediction uncertainty
- κ controls the exploration-exploitation balance

Physics-Informed Surrogate Improvement

Directly learning the complete nonlinear physics response requires a large number of training samples. However, much of the system behaviour is already captured by the analytical model. Therefore, instead of learning the complete objective, the surrogate may learn only the nonlinear correction term:

$$f(x) = J_{sim}(x) - J_{linearised}(x)$$

This allows the Gaussian Process to focus on the complex residual behaviour rather than rediscovering known physics. The expected advantages are improved sample efficiency, reduced training requirements, and better generalisation across the design space.

Gaussian Process Optimisation Loop

The complete optimisation cycle is therefore:

1. Generate initial design samples.
2. Evaluate each design using the deterministic physics model.
3. Train the Gaussian Process surrogate.
4. Select the next candidate using the acquisition function.
5. Evaluate the candidate using the physics model.
6. Update the surrogate model.
7. Repeat until convergence.

The computational workflow is:

$$x \rightarrow Physics\ Simulation \rightarrow J(x) \rightarrow GP\ Model \rightarrow Acquisition\ Function \rightarrow x_{new}$$

Final Local Refinement Stage

Implementation Note: *The original optimisation formulation is based on Gaussian Process assisted Bayesian optimisation. During implementation, an additional local refinement stage was introduced after the global surrogate search. This refinement step does not replace the Bayesian optimisation framework. Instead, it is used only after a promising region of the design space has been identified. The purpose is to perform a local search around the best candidate found by the GP stage.*

The surrogate-assisted search identifies promising regions of the design space but does not guarantee that the local optimum within that region has been reached. Therefore, following the global optimisation stage, a derivative-free local refinement step is applied. The final implementation uses the Nelder–Mead simplex algorithm.

Nelder–Mead was selected because the objective function contains nonlinear hydraulic behaviour, penalty terms from constraints, harmonic interactions, and spline parameter coupling. These characteristics make analytical gradients unreliable. The method requires only objective evaluations and therefore integrates naturally with the existing physics simulation.

The final optimisation sequence is therefore:

$$GP\ Global\ Search \rightarrow Nelder-Mead\ Local\ Refinement \rightarrow Final\ Physics\ Validation$$

This hybrid approach combines the global exploration capability of Bayesian optimisation with the local convergence properties of derivative - free simplex optimisation.

Summary

The optimisation framework transforms the expensive nonlinear hydraulic motor simulation into a practical design exploration method. A Gaussian Process surrogate model is used to approximate the objective landscape and guide the selection of promising candidate designs through Bayesian optimisation. The surrogate does not replace the physical model; instead, it reduces the number of required high - fidelity simulations. Following the global search stage, Nelder–Mead refinement is applied to further improve promising solutions without requiring objective gradients.

Gaussian Process Surrogate Modelling

The previous section established the optimisation framework and demonstrated why surrogate - assisted optimisation is required for the hydraulic motor design problem. Although the reduced - order physics model is computationally efficient compared with a full computational fluid dynamics simulation, each evaluation still requires generation of the cam geometry, piston kinematic calculation, hydraulic pressure simulation, torque calculation, Fourier decomposition, and objective evaluation. When this process is repeated thousands of times during optimisation, the total computational cost becomes significant.

A Gaussian Process (GP) surrogate model is therefore introduced to approximate the deterministic objective function. The purpose of the surrogate is not to replace the physical model, but to provide a statistical approximation of the design landscape that can guide future evaluations. The GP provides two critical quantities: $\mu(x)$, the predicted objective value, and $\sigma^2(x)$, the uncertainty associated with that prediction. This uncertainty allows the optimisation algorithm to intelligently balance searching known good regions and exploring poorly understood regions.

Objective Function Approximation

The deterministic optimisation problem is defined as $J_{phys}(x)$, where x is the motor design vector. The Gaussian Process approximates the relationship:

$$J_{phys}(x) \approx f(x)$$

where $f(x)$ is treated as a random function. Unlike conventional regression methods, a Gaussian Process does not assume a fixed parametric form for the function. Instead, it defines a probability distribution over possible functions.

Definition of a Gaussian Process

A Gaussian Process is defined as a collection of random variables where any finite subset follows a multivariate Gaussian distribution. It is written as:

$$f(x) \sim GP\big(m(x), k(x, x')\big)$$

where:

- $m(x) = E[f(x)]$ is the mean function
- $k(x, x') = E\big[\big(f(x) - m(x)\big)\big(f(x') - m(x')\big)\big]$ is the covariance function

The mean describes the expected objective behaviour, while the covariance defines how strongly two design points influence each other.

Training Data

The Gaussian Process is trained using previously evaluated designs. The available dataset is:

$$D = \{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)\}$$

where $y_i = J_{phys}(x_i)$. The input matrix is $X = [x_1, x_2, \ldots, x_n]$ and the corresponding objective values are $y = [y_1, y_2, \ldots, y_n]^T$.

Covariance Function

The covariance function determines the smoothness and correlation structure assumed by the surrogate. A common choice for engineering optimisation problems is the squared exponential kernel:

$$k(x, x') = \sigma_f^2 exp\big(-½ (x - x')^T L^{-2} (x - x')\big)$$

where:

- σ_f^2 is the signal variance
- L contains characteristic length scales for each design variable

The length scales determine how rapidly the objective changes with each parameter. A small length scale indicates strong sensitivity. A large length scale indicates a smoother response.

Observation Noise

Although the deterministic physics model is theoretically noise - free, numerical effects may introduce small variations. These may arise from finite Fourier resolution, solver tolerance, numerical integration error, and constraint penalties. Therefore, a noise term is included:

$$y_i = f(x_i) + \varepsilon, where\ \varepsilon \sim N(0, \sigma_n^2)$$

The covariance matrix becomes:

$$K_n = K + \sigma_n^2 I$$

where $K_{ij} = k(x_i, x_j)$.

Posterior Prediction

After training, the GP predicts the objective at a new design point x_*. The covariance between the new point and training points is:

$$k_* = [k(x_*, x_1), \ldots, k(x_*, x_n)]^T$$

The covariance at the new point is $k_{**} = k(x_*, x_*)$. The posterior mean is:

$$\mu(x_*) = k_*^T (K + \sigma_n^2 I)^{-1} y_*$$

This represents the predicted objective value.

Posterior Uncertainty

The prediction variance is:

$$\sigma^2(x_*) = k_* - k_*^T (K + \sigma_n^2 I)^{-1} k_{**}$$

This uncertainty is fundamental to Bayesian optimisation. Near previously evaluated designs, $\sigma^2(x) \to 0$. In unexplored regions, $\sigma^2(x)$ becomes larger. Therefore, the GP naturally identifies regions where additional simulation data would be valuable.

Physics-Informed Residual Learning

The deterministic hydraulic model already captures the dominant physical behaviour. Therefore, rather than learning the complete objective directly, the surrogate may model the difference between the full simulation and a simplified analytical approximation. Define the residual:

$$r(x) = J_{phys}(x) - J_{approx}(x)$$

The Gaussian Process is trained on $r(x)$ rather than $J_{phys}(x)$. The final prediction becomes:

$$J_{pred}(x) = J_{approx}(x) + r_{GP}(x)$$

This approach embeds known physics into the learning process. The GP is therefore responsible only for modelling the unknown nonlinear correction.

Advantages of Residual Surrogate Modelling

Compared with direct black-box learning, residual modelling provides several advantages:

- Reduced Data Requirement: The dominant system behaviour is already known. The GP does not need to learn basic hydraulic stiffness, piston volume relationships, or harmonic structure. It only learns the remaining nonlinear effects.

- Improved Generalisation: Because the prediction is constrained by physical knowledge, the surrogate is less likely to produce unrealistic behaviour in unexplored regions.
- Better Optimisation Efficiency: A more accurate surrogate means fewer physics evaluations, faster convergence, and improved exploration.

Acquisition Function

The Gaussian Process prediction is used to select the next candidate design. For minimisation problems, the lower confidence bound acquisition function is used:

$$a(x) = \mu(x) - \kappa\sigma(x)$$

The next design is selected by:

$$x_{next} = argmin_x a(x)$$

The first term encourages exploitation: μ(x) selects designs predicted to have low objective values. The second term encourages exploration: $-\kappa\sigma(x)$ favours regions where uncertainty is high.

Optimisation Convergence

As additional physics evaluations are collected, the GP prediction becomes more accurate, uncertainty decreases, and the acquisition function concentrates around promising regions. Eventually the optimisation converges when additional evaluations provide negligible improvement. The final design is then passed to the full physics model for validation.

Summary

This section developed the Gaussian Process surrogate used within the optimisation framework. The GP was formulated as a probabilistic approximation of the expensive deterministic objective function, providing both predicted performance and uncertainty. The posterior mean allows estimation of likely objective values, while the posterior variance enables intelligent exploration of unknown regions. By incorporating physics - informed residual learning, the surrogate focuses on modelling nonlinear behaviour not already captured by the analytical model. This improves sample efficiency and reduces the number of required high - fidelity simulations. The resulting surrogate model forms the global exploration component of the optimisation framework, with final candidate solutions subsequently refined using derivative - free local optimisation.

Complete Computational Modelling and Optimisation Workflow

The previous sections developed the mathematical foundations of the hydraulic motor model and the surrogate - assisted optimisation strategy. The final computational framework combines these components into a sequential design loop. The objective of the framework is to transform a set of geometric and hydraulic design variables into a predicted motor performance metric while minimising the number of expensive physics evaluations.

The complete model therefore represents a chain of coupled transformations:

$$x \rightarrow C(\theta) \rightarrow s(\theta) \rightarrow V(\theta) \rightarrow P(\theta) \rightarrow \tau(\theta) \rightarrow J(x)$$

where x is the design vector, $C(\theta)$ is the cam profile, $s(\theta)$ is piston displacement, $V(\theta)$ is chamber volume, $P(\theta)$ is chamber pressure, $\tau(\theta)$ is output torque, and $J(x)$ is the optimisation objective.

Design Parameterisation

The first stage of the computational process converts the optimisation variables into a physically meaningful cam geometry. The cam profile is represented using a periodic cubic B-spline:

$$C(\theta) = \sum_{i=0}^{n} N_{i,3}(\theta) P_i$$

where $N_{i,3}$ are cubic B-spline basis functions and P_i are control points. The periodic boundary condition ensures continuity: $C(0) = C(2\pi)$. The optimisation variables therefore modify the spline control parameters rather than directly modifying individual points along the cam surface. This provides smooth geometry, reduced optimisation dimension, guaranteed differentiability, and improved manufacturability.

Kinematic Evaluation

Once the cam profile has been generated, piston motion is calculated. The radial piston displacement is $s(\theta) = ||C(\theta)||$. The corresponding velocity is:

$$\dot{s} = \omega \left(\frac{ds}{d\theta}\right)$$

Acceleration is:

$$\ddot{s} = \omega^2 \left(\frac{d^2 s}{d\theta^2}\right)$$

Higher derivatives are used to evaluate dynamic smoothness:

$$\dddot{s} = \omega^3 \left(\frac{d^3 s}{d\theta^3}\right)$$

The resulting displacement waveform represents the mechanical excitation imposed on the hydraulic chamber.

Hydraulic Chamber Simulation

The piston displacement determines the chamber volume:

$$V(\theta) = V_{dead} + A_p s(\theta)$$

The pressure evolution is obtained from the nonlinear hydraulic model:

$$\dot{P} = \left(\frac{\beta}{V(\theta)}\right)\left(Q_{net} - A_p \dot{s}\right)$$

where $Q_{net} = Q_{port} - Q_{leak}$. The valve flow contribution is:

$$Q_{port} = C_d A(\theta) \sqrt{\frac{2\Delta P}{\rho}}$$

Leakage is $Q_{leak} = C_l(P - P_t)$. This produces the periodic chamber pressure waveform $P(\theta)$.

Mechanical Torque Calculation

The hydraulic pressure generates piston force: $F_h = A_p P$. This force is transmitted through the cam mechanism to generate shaft torque. For each piston:

$$\tau_i = r(\theta) F_i \sin(\varphi_i)$$

The total torque is:

$$\tau(\theta) = \sum_{i=0}^{n} \tau_{\mathrm{i}}(\theta)$$

The multiple piston contributions are phase shifted according to their angular positions. This phase relationship determines harmonic reinforcement and cancellation.

Spectral Analysis

The pressure and torque waveforms are transformed into the frequency domain. For a periodic quantity $x(\theta) = \sum_n \hat{x}_{\mathrm{n}}\, e^{\mathrm{in}\theta}$, the Fourier coefficients are calculated:

$$\hat{x}_{\mathrm{n}} = \frac{1}{2\pi}\int_0^2 \pi\, x(\theta) e^{-\mathrm{in}\theta}\, d\theta$$

The resulting harmonic coefficients are $\hat{P}_{\mathrm{n}}$ and $\hat{\tau}_{\mathrm{n}}$. These values quantify the undesirable periodic content of the system.

Objective Evaluation

The spectral information is combined with efficiency and smoothness metrics. The objective function is:

$$J_{phys} = w_1 R_\tau^2 + w_2 R_P^2 + w_3(1 - \eta) + w_4 J_{spec} + w_5 J_{jerk}$$

where R_τ is torque ripple, R_P is pressure ripple, η is efficiency, and J_{jerk} is the cam smoothness penalty. Constraint penalties are applied if the design violates cavitation limits, stress limits, geometric limits, or dynamic stability requirements.

Initial Sampling Strategy

Before surrogate optimisation begins, an initial set of designs is evaluated using the deterministic model. The initial dataset is:

$$D_0 = \{(x_1, J_1), (x_2, J_2), \ldots\}$$

These samples provide the first approximation of the objective landscape. Because no prior information about the optimum is assumed, the initial samples must provide good coverage of the design space.

Gaussian Process Search Loop

The initial dataset is used to train the Gaussian Process. The surrogate predicts $\mu(x)$ and $\sigma(x)$. The acquisition function selects the next candidate:

$$x_{new} = argmin[\, \mu(x) - \kappa\sigma(x)]$$

The selected candidate is then evaluated using the full physics model. The new result is added:

$$D_{\{k+1\}} = D_k \cup (x_{new}, J_{new})$$

The process repeats until convergence.

Local Derivative – Free Refinement

After the global search identifies a promising region, the solution is refined using the Nelder-Mead simplex algorithm. The optimisation problem remains $\min J_{phys}(x)$, but only a local region surrounding the best GP candidate is considered. Nelder-Mead updates a simplex through geometric operations: reflection, expansion, contraction, and shrink. The method requires only objective evaluations $J(x)$ and therefore avoids requiring

analytical gradients. This is advantageous because the objective contains nonlinear pressure dynamics, spectral calculations, constraint penalties, and coupled geometric effects.

Final Validation

The final candidate design is not accepted solely based on surrogate prediction. The optimised design is evaluated again using the complete deterministic physics model. The validation stage verifies torque ripple reduction, pressure ripple reduction, efficiency improvement, constraint satisfaction, and manufacturable cam geometry. The final output is therefore a physically validated design rather than simply the minimum predicted surrogate value.

Complete Optimisation Architecture

The complete computational process can be summarised as the following sequential workflow:

1. Input: Design vector x
2. Stage 1: Periodic Cubic B-Spline Cam Generator
3. Stage 2: Piston Kinematics
4. Stage 3: Nonlinear Hydraulic Solver
5. Stage 4: Mechanical Torque Model
6. Stage 5: Fourier Spectral Analysis
7. Stage 6: Objective Evaluation $J_{phys}(x)$
8. Stage 7: Gaussian Process Global Search
9. Stage 8: Nelder-Mead Local Refinement
10. Output: Optimal design vector $x_{optimal}$

Summary

This section combined the individual mathematical components into a complete computational framework for hydraulic motor optimisation. The model begins with a compact geometric representation of the cam profile and progressively transforms this geometry into piston motion, chamber pressure, output torque, and finally a scalar performance objective. The optimisation framework combines high - fidelity physics evaluation with surrogate - assisted global exploration and derivative - free local refinement. This allows the design space to be searched efficiently while maintaining physical accuracy. The resulting methodology provides a complete link between geometric design variables and system - level performance metrics, enabling automated optimisation of hydraulic motor characteristics such as torque ripple, pressure pulsation, and efficiency.

Linearisation of the Hydraulic Pressure Dynamics

The nonlinear pressure equation derived previously is:

$$\dot{P} = \left(\frac{\beta}{V(\theta)}\right)\left(Q_{net} - A_p\dot{s}\right)$$

where the net flow is determined by the valve port relationship, leakage, and pressure difference. This equation fully describes the nonlinear hydraulic subsystem. However, for frequency - domain analysis and harmonic response evaluation, a linearised representation is required. The objective of this section is to derive the small - signal pressure model around a nominal operating point.

Nominal Operating Point

Consider the hydraulic system operating around a steady periodic condition:

- $P = P_0 + \tilde{P}$

- $Q = Q_0 + \tilde{Q}$
- $s = s_0 + \tilde{s}$

where P_0 is the nominal chamber pressure, Q_0 is the nominal flow, s_0 is the nominal piston position, and quantities with a tilde represent small perturbations. The purpose of linearisation is to retain only first-order variations.

Linearisation of Chamber Volume

The chamber volume is $V(\theta) = V_{dead} + A_p s(\theta)$. Expanding around the operating point:

$$V = V_0 + \tilde{V}$$

where $\tilde{V} = A_p \tilde{s}$. For small perturbations, the volume term is approximated as constant: $V(\theta) \approx V_0$. Therefore:

$$\dot{P} = \left(\frac{\beta}{V_0}\right)\left(Q_{net} - A_p \dot{s}\right)$$

Linearisation of Valve Flow

The nonlinear port flow is:

$$Q_{port} = C_d A(\theta)\sqrt{\frac{2(P_s - P)}{\rho}}$$

Expanding around P_0:

$$Q_{port} \approx Q_0 + \left(\frac{\partial Q}{\partial P}\right)\tilde{P}$$

Define the pressure-flow coefficient $C_f = -\frac{\partial Q}{\partial P}$, therefore:

$$\tilde{Q}_{port} = -C_f \tilde{P}$$

The coefficient C_f represents the effective hydraulic damping created by the valve restriction.

Leakage Linearisation

The leakage relationship is $Q_{leak} = C_l(P - P_t)$. Perturbing around the operating point:

$$\tilde{Q}_{leak} = C_l \tilde{P}$$

The total flow perturbation is therefore $\tilde{Q}_{net} = \tilde{Q}_{port} - \tilde{Q}_{leak}$, giving:

$$\tilde{Q}_{net} = -(C_f + C_l)\tilde{P}$$

Define the combined hydraulic flow coefficient:

$$C_h = C_f + C_l$$

so that:

$$\tilde{Q}_{net} = -C_h \tilde{P}$$

Linearised Pressure Equation

Substituting into the pressure equation:

$$\tilde{\dot{P}} = \left(\frac{\beta}{V}\right)\left(-C_h \tilde{P} - A_p \tilde{\dot{s}}\right)$$

Rearranging:

$$\dot{\tilde{P}} + \left(\frac{\beta C_h}{V}\right)\tilde{P} = -\left(\frac{\beta A_p}{V}\right)\dot{\tilde{s}}$$

This is the linearised hydraulic subsystem. The two coefficients represent $\frac{\beta C_h}{V}$ (hydraulic damping) and $\frac{\beta A_p}{V}$ (mechanical-to-hydraulic coupling).

Frequency Domain Transformation

Represent pressure and displacement using Fourier components:

$$\tilde{P} = \sum_n \hat{P}_n e^{in\theta} \quad \tilde{s} = \sum_n \hat{s}_n e^{in\theta}$$

The time derivative relationship is $\frac{d}{dt} = \omega\left(\frac{d}{d\theta}\right)$, therefore $\frac{d}{dt}(e^{in\theta}) = in\omega e^{in\theta}$. Substituting into the pressure equation gives:

$$in\omega\hat{P}_n + \left(\frac{\beta C_h}{V}\right)\hat{P}_n = -\left(\frac{\beta A_p}{V}\right)in\omega\hat{s}_n$$

Collecting terms:

$$\left(in\omega + \frac{\beta C_h}{V}\right)\hat{P}_n = -\left(\frac{\beta A_p}{V}\right)in\omega\hat{s}_n$$

Therefore:

$$\hat{P}_n = \frac{-\left(\beta\frac{A_p}{V}\right)in\omega}{in\omega + \frac{\beta C_h}{V}}\hat{s}_n$$

Hydraulic Frequency Response Function

Define:

$$G_{h(n)} = \frac{-\left(\frac{\beta A_p}{V}\right)in\omega}{in\omega + \frac{\beta C_h}{V}}$$

giving:

$$\hat{P}_n = G_h(n)\hat{s}_n$$

This represents the harmonic transfer function between cam-induced piston motion and chamber pressure response.

Hydraulic Time Constant

The hydraulic response time is determined by:

$$\tau_h = \frac{V}{\beta C_h}$$

Therefore $G_h(n)$ can be expressed in terms of this time constant. The quantity τ_h describes how rapidly pressure responds to piston-induced volume changes.

Low and High Frequency Limits

Low Frequency: For $n\omega \ll \tau_h^{-1}$, the system follows the flow dynamics and $G_h(n) \approx -\left(\frac{A_p}{C_h}\right) in\omega$. Pressure fluctuations remain small because fluid can flow through the ports.

High Frequency: For $n\omega \gg \tau_h^{-1}$, the fluid behaves as a stiff spring and $G_h(n) \approx -\beta \frac{A_p}{V}$. The valve flow cannot respond quickly enough, and pressure variations are generated through fluid compression.

Summary

The nonlinear hydraulic equations were reduced to a linear frequency - domain model by perturbing the system around a nominal operating condition. Linearisation of the valve flow produced an effective hydraulic flow coefficient, which represents the damping introduced by the hydraulic restriction and leakage paths. The resulting harmonic transfer function, $\hat{P}_n = G\_h(n)\ \hat{s}_n$, provides the mathematical connection between cam geometry and pressure ripple. This relationship forms the basis for the spectral torque and pressure ripple metrics used in the optimisation objective.

Cam Force Transmission and Torque Generation

The previous sections derived the relationship between piston displacement and hydraulic pressure. However, pressure itself is not the final output of interest. The purpose of the hydraulic motor is to convert hydraulic energy into rotational mechanical power. The conversion occurs through the interaction between the piston and the cam ring. The hydraulic pressure generates a radial piston force, which acts against the cam surface. Due to the cam geometry, this radial force is transformed into a tangential force acting on the rotating assembly. This section derives the relationship between chamber pressure and output torque.

Hydraulic Force on the Piston

The pressure acting over the piston face produces a normal hydraulic force:

$$F_h = A_p P$$

where A_p is piston area and P is chamber pressure relative to return pressure. For a piston located at angular position θ_i, the hydraulic force acts along the piston axis. Therefore:

$$F_{\{h,i\}} = A_p P_i$$

for the i'th piston.

Cam Contact Geometry

The piston force is not directly converted into torque. The force acts normal to the piston axis, while useful torque is generated by the tangential component of the cam reaction force. Define $\varphi(\theta)$ as the cam pressure angle. This is the angle between the piston force direction and the normal direction of the cam surface. The cam reaction force is therefore resolved into a normal component and a tangential component. The tangential component is:

$$F_t = F_h sin\varphi$$

therefore:

$$F_t = A_p P \sin\left(\varphi(\theta)\right)$$

Torque Contribution of a Single Piston

Torque is defined as $\tau = r F_t$, where $r(\theta)$ is the instantaneous moment arm from the shaft centre. Substituting:

$$\tau_i = r(\theta_i) A_p P_i \sin(\varphi(\theta_i))$$

Therefore:

$$\tau_i(\theta) = A_p P_i(\theta) r(\theta_i) sin\varphi(\theta_i)$$

This represents the torque contribution from one piston.

Cam Geometry Relationship

The pressure angle is determined directly from the cam profile. For a radial cam, $r(\theta) = C(\theta)$ and the cam tangent direction is $t = \frac{dC}{d\theta}$. The surface normal is perpendicular to the tangent: $n \cdot t = 0$. The pressure angle is therefore:

$$\varphi = cos^{-1}\left(\frac{(n \cdot e_r)}{|n|}\right)$$

where e_r is the radial piston direction. Therefore, changes in cam geometry modify piston displacement, pressure angle, mechanical leverage, and torque waveform.

Multiple Piston Contributions

A radial piston motor contains multiple pistons distributed around the cylinder block. For N_p pistons, the angular position of piston i is:

$$\theta_i = \theta + \frac{2\pi i}{N_p}$$

where $i = 0, 1, \ldots, N_p - 1$. Each piston experiences a phase - shifted version of the same cam profile:

$$s_i(\theta) = s\left(\theta + \frac{2\pi i}{N_p}\right)$$

and therefore:

$$P_i(\theta) = P\left(\theta + \frac{2\pi i}{N_p}\right)$$

The total torque is obtained by summing all piston contributions:

$$\tau(\theta) = \sum_{i=0}^{N_p - 1} A_p P_i(\theta) r_i(\theta) \sin(\Phi_i(\theta))$$

Harmonic Representation of Torque

The periodic torque waveform may be expanded using a Fourier series:

$$\tau(\theta) = \sum_n \hat{\tau}_n e^{in\theta}$$

Each piston contributes $\hat{\tau}_{\{n,i\}}$ such that:

$$\hat{\tau}_n = \sum_i \hat{\tau}_{\{n,i\}}$$

The phase relationship between pistons determines harmonic cancellation. For equally spaced pistons:

$$\hat{\tau}_{\{n,i\}} = \hat{\tau}_{\{n,0\}} e^{in\frac{2\pi i}{N_p}}$$

therefore:

$$\hat{\tau}_n = \hat{\tau}_{\{n,0\}} \sum_{i=0}^{N_p-1} e^{in\frac{2\pi i}{N_p}}$$

Harmonic Cancellation Condition

The summation term:

$$S_n = \sum_{i=0}^{N_p-1} e^{in\frac{2\pi i}{N_p}}$$

Determines whether a harmonic reinforces or cancels. For certain harmonic orders, $S_n = 0$ and therefore $\hat{\tau}_n = 0$. This explains why piston count strongly influences torque smoothness. A suitable piston arrangement can suppress dominant harmonic components without changing the fundamental operating principle of the motor.

Mechanical Dynamic Contribution

The previous torque relationship assumes quasi-static force transmission. However, piston inertia introduces an additional force component. The piston force balance is:

$$m_p\ddot{s} + c_m\dot{s} = A_pP - F_c$$

where F_c is the cam contact force. Rearranging:

$$F_c = A_pP - m_p\ddot{s} - c_m\dot{s}$$

Therefore the torque contribution becomes:

$$\tau_i = r_i sin\varphi_i \left(A_pP_i - m_p\ddot{s}_i - c_m\dot{s}_i \right)$$

This equation shows that torque ripple is generated by both hydraulic pressure variation, $P(\theta)$, and mechanical excitation caused by piston acceleration, $\ddot{s}(\theta)$.

Torque Harmonic Relationship

Taking the Fourier transform:

$$\hat{\tau}_n = \Gamma_n \left(A_p\hat{P}_n - m_p(n\omega)^2\hat{s}_n - c_m(in\omega)\hat{s}_n \right)$$

where Γ_n represents the harmonic mechanical transmission coefficient. Therefore:

$$\hat{\tau}_n = \Gamma_n \left(A_p\hat{P}_n + m_p\omega^2n^2\hat{s}_n + c_min\omega\,\hat{s}_n \right)$$

This expression connects the hydraulic and mechanical frequency responses.

Interpretation

The torque harmonic contains three contributions:

- Hydraulic pressure contribution: $A_p\hat{P}_n$. Pressure oscillations directly create torque ripple.

- Inertial contribution: $m_p\omega^2 n^2 \hat{s}_n$. Higher harmonic components are amplified by acceleration effects.
- Damping contribution: $c_m in\omega\ \hat{s}_n$. Mechanical losses introduce phase-shifted torque components.

Summary

This section derived the mechanical conversion between hydraulic pressure and shaft torque. Starting from the hydraulic force acting on each piston, the cam pressure angle and moment arm were introduced to determine the tangential force responsible for torque generation. The contributions from multiple pistons were then combined, showing how piston count and phase relationships influence harmonic cancellation. The final harmonic torque relationship, $\hat{\tau}_n = \Gamma_n \left(A_p \hat{P}_n + m_p\omega^2 n^2 \hat{s}_n + c_m in\omega\ \hat{s}_n \right)$, provides the final coupling between the hydraulic and mechanical subsystems. Together with the hydraulic pressure transfer function, this equation forms the complete frequency - domain model used to evaluate torque ripple.

Complete Coupled Mathematical Model

The previous sections derived the individual components of the hydraulic motor model: periodic cam geometry, piston kinematics, chamber volume variation, nonlinear hydraulic pressure evolution, mechanical force transmission, and torque generation. These subsystems are strongly coupled. The cam profile determines piston motion, piston motion drives chamber volume variation, volume variation determines pressure evolution, and pressure generation produces output torque. The complete model can therefore be represented as a sequential nonlinear system:

$$x \rightarrow C(\theta) \rightarrow s(\theta) \rightarrow V(\theta) \rightarrow P(\theta) \rightarrow \tau(\theta)$$

where x represents the complete set of motor design variables.

Geometric Model

The cam surface is parameterised as a periodic cubic B-spline:

$$C(\theta) = \sum_{i=0}^{n} N_{i,3}(\theta) P_i$$

where $N_{i,3}$ are cubic B-spline basis functions and P_i are control points. Periodicity is enforced through $P_0 = P_n$ and $C(0) = C(2\pi)$. The radial displacement experienced by the piston is:

$$s(\theta) = \left|\left|C(\theta)\right|\right|$$

This geometric representation defines the complete piston excitation.

Kinematic Model

The piston velocity is obtained from:

$$\dot{s} = \omega \left(\frac{ds}{d\theta}\right)$$

and acceleration:

$$\ddot{s} = \omega^2 \left(\frac{d^2 s}{d\theta^2}\right)$$

The third derivative is:

$$\dddot{s} = \omega^3 \left(\frac{d^3 s}{d\theta^3}\right)$$

and is used as a measure of geometric smoothness.

Chamber Volume Model

The piston motion changes the hydraulic chamber volume. The chamber volume is:

$$V(\theta) = V_{dead} + A_p s(\theta)$$

Differentiating: $\dot{V} = A_p \dot{s}$. This term represents the volumetric displacement imposed on the fluid.

Nonlinear Hydraulic Model

The fundamental conservation equation is:

$$\dot{V} = Q_{in} - Q_{out} - Q_{leak}$$

Using the hydraulic bulk modulus relationship $\beta = -V\left(\frac{dP}{dV}\right)$ gives:

$$\dot{P} = \left(\frac{\beta}{V}\right)\left(Q_{in} - Q_{out} - Q_{leak} - A_p \dot{s}\right)$$

The valve port flow is:

$$Q_{port} = C_d A(\theta)\sqrt{\frac{2|P_s - P|}{\rho}}$$

where the effective port area is $A(\theta) = A_{port}\alpha(\theta)$. Leakage is represented as:

$$Q_{leak} = C_l(P - P_t)$$

Therefore, the complete hydraulic subsystem becomes:

$$\dot{P} = \left(\frac{\beta}{V}\right)\left(Q_{port} - C_l(P - P_t) - A_p \dot{s}\right)$$

This equation represents the nonlinear pressure dynamics.

Linearised Hydraulic Representation

For frequency - domain analysis, the system is linearised around a nominal operating point. Substituting the linearised flow relationship $\Delta Q_{net} = -C_h \Delta P$, where $C_h = C_f + C_l$, gives the reduced pressure dynamics:

$$\tilde{\dot{P}} + \left(\frac{\beta C_h}{V_{eq}}\right)\tilde{P} = -\left(\frac{\beta A_p}{V_{eq}}\right)\tilde{\dot{s}}$$

where V_{eq} is the equivalent chamber volume evaluated at the nominal operating condition. The corresponding harmonic transfer function derived previously is:

$$\hat{P}_n = \frac{-\left(\frac{\beta A_p}{V_{eq}}\right) in\omega}{in\omega + \frac{\beta C_h}{V_{eq}}}\hat{s}_n$$

This relationship provides the frequency-domain pressure response generated by periodic piston excitation. The linearised hydraulic model therefore provides the coupling $\hat{s}_n \rightarrow \hat{P}_n$, which is subsequently used in the torque generation model.

Mechanical Force Balance

The piston force balance is obtained from the interaction between hydraulic pressure, piston inertia, and mechanical damping:

$$m_p\ddot{s} + c_m\dot{s} = A_pP - F_c$$

where m_p is piston mass, c_m is the equivalent mechanical viscous damping coefficient, and F_c is the cam contact force. The hydraulic stiffness generated by fluid compressibility is:

$$k_f = \frac{\beta A_p^2}{V_{eq}}$$

where V_{eq} is the chamber volume evaluated about the nominal operating point. The approximate linear dynamic representation is therefore:

$$m_p\ddot{s} + c_m\dot{s} + k_f s = F_{hyd}$$

The corresponding natural frequency is:

$$\omega\_n = \sqrt{\frac{k_f}{m_p}}$$

and damping ratio:

$$\zeta = \frac{c_m}{2\sqrt{m_p k_f}}$$

This representation demonstrates the inherent hydraulic mass-spring-damper behaviour produced by fluid compressibility and piston inertia.

Hertzian Contact Stress Constraint

Contact Force Determination

The piston transmits hydraulic force to the cam surface through the contact interface. The normal contact force is obtained from the piston force balance:

$$F_{c(\theta)} = A_pP(\theta) - m_p\ddot{s}(\theta) - c_m\dot{s}(\theta)$$

The maximum contact load over one operating cycle is:

$$F_{max} = \max F_{c(\theta)}$$

Equivalent Material Properties

The contact between the piston and cam is approximated using Hertzian elastic contact theory. The equivalent elastic modulus is:

$$\frac{1}{E'} = \frac{(1-\nu_1^2)}{E_1} + \frac{(1-\nu_2^2)}{E_2}$$

where E_1, E_2 are Young's moduli and ν_1, ν_2 are Poisson ratios.

Equivalent Radius of Curvature

The local curvature of the cam determines the contact geometry. The equivalent radius is:

$$\frac{1}{R'} = \frac{1}{R_1} + \frac{1}{R_2}$$

where R_1 is piston radius and R_2 is local cam curvature radius.

Hertzian Contact Pressure

For a cylinder-on-cylinder contact approximation, the contact half-width is $b = \sqrt{\frac{4F_c R'}{\pi L E'}}$, where L is the effective contact length. The maximum Hertzian contact pressure is:

$$p_0 = \frac{2F_c}{\pi b L}$$

The Hertzian contact stress is approximated as $\sigma_H = p_0$.

Contact Stress Constraint

The mechanical feasibility constraint is:

$$g_{H(x)} = \sigma_H - \sigma_{max} \leq 0$$

where σ_{max} is the allowable contact stress determined from material fatigue requirements. This provides a computationally inexpensive approximation of cam/piston contact loading without requiring full finite-element contact analysis.

Torque Generation Model

The piston force is obtained from the mechanical force balance:

$$F_{\{p,i\}} = A_p P_{\mathrm{i}} - m_p \ddot{s}_{\mathrm{i}} - c_m \dot{s}_{\mathrm{i}}$$

The torque contribution of each piston is obtained using the principle of virtual work: $F_{\mathrm{i}} ds_{\mathrm{i}} = \tau_{\mathrm{i}} d\theta$, therefore:

$$\tau_{\mathrm{i}}(\theta) = F_{\{p,i\}}(\theta)\left(\frac{ds_{\mathrm{i}}}{d\theta}\right)$$

For a motor containing N_p pistons:

$$\tau(\theta) = \sum_{i=0}^{N_p-1} F_{\{p,i\}}(\theta)\left(\frac{ds_{\mathrm{i}}}{d\theta}\right)$$

This provides the instantaneous shaft torque generated by the interaction between piston force and cam geometry. The torque generation chain is therefore:

$$P(\theta) \rightarrow F_{p(\theta)} \rightarrow \tau(\theta)$$

Spectral Representation

The periodic torque and pressure responses are analysed using the harmonic formulation developed in the Spectral Representation section. The resulting coefficients $\hat{P}_{\mathrm{n}}$ and $\hat{\tau}_{\mathrm{n}}$ are used to evaluate the ripple metrics and optimisation objectives.

Complete Nonlinear Mapping

The entire hydraulic motor model can now be written as a sequence of transformations:

- $x \rightarrow C(\theta)$
- $C(\theta) \rightarrow s(\theta)$

- $s(\theta) \rightarrow V(\theta), \dot{s}(\theta)$
- $V(\theta), \dot{s}(\theta) \rightarrow P(\theta)$
- $P(\theta), s(\theta) \rightarrow \tau(\theta)$

Or compactly:

$$\tau(\theta) = M(x)$$

where M represents the complete coupled nonlinear motor model.

Energy Relationship

The hydraulic input power is $P_{hyd} = \Delta P Q$. The mechanical output power is $P_{out} = \tau\omega$. The instantaneous efficiency is therefore:

$$\eta = \frac{\tau\omega}{\Delta P\, Q}$$

Losses arise from $P_{loss} = P_{fric} + P_{leak} + P_{visc}$. Thus:

$$P_{hyd} = P_{out} + P_{loss}$$

This provides the energy consistency condition for the model.

Summary

The complete mathematical model represents the hydraulic motor as a chain of coupled physical transformations. The cam geometry defines piston motion, piston motion generates fluid compression, fluid compression produces pressure variation, and pressure variation produces mechanical torque. The resulting model captures the dominant mechanisms responsible for torque ripple, pressure pulsation, dynamic loading, and efficiency variation. The final representation is:

$$x \rightarrow C(\theta) \rightarrow s(\theta) \rightarrow P(\theta) \rightarrow \tau(\theta)$$

which forms the mathematical foundation for evaluating motor performance.

Spectral Representation of Periodic Motor Dynamics

The hydraulic motor operates as a periodic system because the piston motion is generated by a rotating cam profile. Every revolution of the shaft produces a repeated sequence of piston displacement, chamber volume variation, pressure generation, and torque production. Therefore, the resulting pressure and torque signals are periodic functions of shaft angle. Rather than analysing these signals only in the time domain, they can be represented in the frequency domain using Fourier analysis. This provides a direct measure of harmonic content, allowing torque ripple and pressure pulsation to be quantified independently of the average operating point.

Periodic Signal Representation

For any periodic motor quantity $x(\theta)$ with period 2π, the Fourier series representation is:

$$x(\theta) = \sum_{-\infty}^{\infty} \hat{x}_{\mathrm{n}}\, e^{\mathrm{in\theta}}$$

where $\hat{x}_{\mathrm{n}}$ is the complex Fourier coefficient of harmonic n. The coefficient is calculated as:

$$\hat{x}_{\mathrm{n}} = \frac{1}{2\pi} \int_{0}^{2} \pi\, x(\theta)\, e^{-\mathrm{in\theta}}\, d\theta$$

The zeroth harmonic represents the average value: $\hat{x}_0 = \frac{1}{2\pi}\int_0^2 \pi\, x(\theta)\, d\theta$, while higher-order terms represent oscillatory components.

Harmonic Decomposition of Pressure

The chamber pressure waveform is $P(\theta)$ and may be written as:

$$P(\theta) = P_0 + \sum_{n=1}^{\infty} \hat{P}_{\mathrm{n}}\, e^{\mathrm{in\theta}}$$

where P_0 is the mean chamber pressure. The oscillating component is $\tilde{P}(\theta) = \sum_{n=1}^{\infty} \hat{P}_{\mathrm{n}}\, e^{\mathrm{in\theta}}$. The pressure ripple is therefore determined by the magnitude of $\left|\hat{P}_{\mathrm{n}}\right|$.

Harmonic Decomposition of Torque

Similarly, shaft torque is $\tau(\theta)$ and is represented as:

$$\tau(\theta) = \tau_0 + \sum_{n=1}^{\infty} \hat{\tau}_{\mathrm{n}}\, e^{\mathrm{in\theta}}$$

where τ_0 is the average output torque. The oscillatory torque component is $\tilde{\tau}(\theta) = \sum_{n=1}^{\infty} \hat{\tau}_{\mathrm{n}}\, e^{\mathrm{in\theta}}$. The magnitude of each coefficient $|\hat{\tau}_{\mathrm{n}}|$ defines the contribution of each harmonic order to torque ripple.

Relationship Between Harmonic Order and Physical Excitation

The harmonic frequency associated with order n is:

$$\omega_n = n\omega$$

where ω is the shaft rotational speed. Therefore, higher harmonic orders correspond to higher excitation frequencies. The corresponding piston acceleration term is $\ddot{s}_{\mathrm{n}} = -(n\omega)^2 \hat{s}_{\mathrm{n}}$. This shows that mechanical inertia naturally amplifies higher-order geometric excitation.

Influence of Piston Arrangement

For a motor containing N_p equally spaced pistons, the total torque harmonic is the sum of phase-shifted piston contributions:

$$\hat{\tau}_{\mathrm{n}} = \sum_{i=0}^{N_p-1} \hat{\tau}_{\{n,i\}}$$

Each piston contribution has a phase shift $\theta_{\mathrm{i}} = \frac{2\pi i}{N_p}$, therefore:

$$\hat{\tau}_{\{n,i\}} = \hat{\tau}_{\{n,0\}}\, e^{in\theta_{\mathrm{i}}}$$

The total harmonic becomes:

$$\hat{\tau}_{\mathrm{n}} = \hat{\tau}_{\{n,0\}} \sum_{i=0}^{N_p-1} e^{in\frac{2\pi i}{N_p}}$$

The summation term acts as a harmonic filter. Certain harmonic orders reinforce ($|\hat{\tau}_n|$ becomes large) while others cancel ($|\hat{\tau}_n|$ approaches 0). This explains why piston count and phase arrangement strongly influence torque smoothness.

Parseval Relationship

The total energy contained within a periodic waveform is related to its Fourier coefficients through Parseval's theorem:

$$\frac{1}{2\pi}\int_0^{2\pi} |x(\theta)|^2 d\theta = \sum_{n=1}^{\infty} |\hat{x}_n|^2$$

Therefore, the total oscillatory energy is:

$$E_x = \sum_{n=1}^{\infty} |\hat{x}_n|^2$$

This provides the mathematical basis for using harmonic energy as a ripple metric.

Pressure Ripple Metric

The pressure ripple is defined relative to the mean pressure. The RMS harmonic energy is:

$$P_{ripple} = \sqrt{\sum_{n=0}^{N} |\hat{P}_n|^2}$$

Normalising by the average pressure gives:

$$R_P = \frac{\sqrt{\sum_{n=0}^{N} |\hat{P}_n|^2}}{|\hat{P}_0|}$$

This provides a dimensionless measure of pressure pulsation.

Torque Ripple Metric

Applying the same formulation to torque:

$$\tau_{ripple} = \sqrt{\sum_{n=0}^{N} |\hat{\tau}_n|^2}$$

Normalisation gives:

$$R_\tau = \frac{\sqrt{\sum_{(n=1)}^{N} |\hat{\tau}_n|^2}}{|\hat{\tau}_0|}$$

This quantity measures the magnitude of torque oscillation relative to the average produced torque.

Truncated Harmonic Representation

In practice, the Fourier series is truncated: $n = 1, \ldots, N$, where N is selected such that all significant harmonics are retained. Therefore:

$$x(\theta) \approx \hat{x}_0 + \sum_{n=1}^{N} \hat{x}_n e^{in\theta}$$

The truncation provides computational efficiency and the removal of negligible high-frequency numerical noise.

Unified Spectral Objective

Rather than treating pressure and torque ripple independently, the total harmonic energy of the motor can be represented as:

$$J_{spec} = \sum_{n=1}^{N} w\_P|\hat{P}_n|^2 + w\_\tau|\hat{\tau}_n|^2$$

where w_P and w_τ control the relative importance of pressure and torque smoothness. This formulation directly penalises harmonic energy rather than only peak-to-peak ripple.

Low-Order Harmonic Dominance

Although the Fourier series contains infinitely many harmonics, low-order components generally dominate the mechanical response. This occurs because:

1. Hydraulic filtering attenuates high frequencies. From the hydraulic transfer function $G_h(n) = \frac{-\left(\frac{\beta A_p}{V}\right) in\omega}{in\omega + \frac{\beta C_h}{V}}$, the pressure response approaches a bounded value as frequency increases.
2. Mechanical inertia suppresses high-frequency torque components. The inertial term $m_p(n\omega)^2\hat{s}_n$ is balanced by structural and hydraulic dynamics, reducing transmission of very high-frequency components.
3. Structural compliance and damping attenuate high-frequency excitation.

Therefore, the dominant contributors to perceived vibration and acoustic output are typically the first several harmonic orders.

Spectral Model Summary

The periodic motor behaviour can therefore be represented entirely in the frequency domain:

$$\hat{s}_n \rightarrow \hat{P}_n \rightarrow \hat{\tau}_n$$

with:

$$\hat{P}_n = G_h(n)\hat{s}_n$$

and:

$$\hat{\tau}_n = \Gamma_n\left(A_p\hat{P}_n + m_p\omega^2n^2\hat{s}_n + c_m in\omega\,\hat{s}_n\right)$$

This provides a complete harmonic representation of the coupled hydraulic-mechanical system. The resulting spectral quantities form the mathematical basis for evaluating motor smoothness and defining the performance objective.

Mathematical Formulation of the Design Objective

The previous sections developed the complete physical model of the hydraulic motor. The model provides the following outputs for any design vector:

$$x \rightarrow \{P(\theta), \tau(\theta), \eta, s(\theta)\}$$

However, optimisation requires these multiple physical outputs to be combined into a single scalar measure of performance. The purpose of the objective function is therefore to quantify the overall quality of a motor design while preserving the physical meaning of each performance requirement. The optimisation problem is formulated as:

$$\min_{\mathrm{x}} J(x)$$

subject to physical feasibility constraints.

Design Vector

The design vector contains all parameters that can be modified during optimisation. For a cam - based hydraulic motor:

$$x = \{p, A_{port}, \theta_{shift}, C_l, c_m, \ldots\}$$

where $p = \{p_1, p_2, \ldots, p_{\mathrm{n}}\}$ are the B-spline cam parameters. The design vector therefore defines cam geometry, hydraulic timing, flow characteristics, and mechanical damping properties.

Torque Ripple Objective

The torque ripple metric is obtained from the harmonic representation. The torque waveform is represented as:

$$\tau(\theta) = \tau_0 + \sum_{n=1}^{N} \hat{\tau}_{\mathrm{n}}\, e^{\mathrm{in}\theta}$$

where $\hat{\tau}_{\mathrm{n}}$ are the torque harmonic coefficients. The normalised torque ripple metric is therefore:

$$R_\tau = \frac{\sqrt{\sum_{n=1}^{N} |\hat{\tau}_{\mathrm{n}}|^2}}{|\hat{\tau}_0|}$$

This provides a dimensionless measure of oscillatory torque energy relative to the average output torque. The contribution of torque ripple to the design objective is:

$$J_\tau = w_\tau R_\tau^2$$

where w_τ is the weighting coefficient applied to torque smoothness.

Pressure Ripple Objective

The chamber pressure ripple metric follows directly from the Fourier representation. The pressure waveform is:

$$P(\theta) = P_0 + \sum_{n=1}^{N} \hat{P}_{\mathrm{n}}\, e^{\mathrm{in}\theta}$$

where $\hat{P}_{\mathrm{n}}$ are the pressure harmonic coefficients. The normalised pressure ripple is:

$$R_P = \frac{\sqrt{\sum_{n=1}^{N} |\hat{P}_{\mathrm{n}}|^2}}{|\hat{P}_0|}$$

This metric represents the magnitude of pressure pulsation relative to the mean chamber pressure. The corresponding objective contribution is:

$$J_P = w_P R_P^2$$

where w_P is the pressure ripple weighting coefficient.

Unified Spectral Objective

The individual pressure and torque ripple metrics quantify the relative magnitude of each output oscillation. However, the complete harmonic energy of the hydraulic motor can be represented using the combined spectral objective:

$$J_{spec} = \sum_{n=1}^{N} w_P \left|\hat{P}_{\mathrm{n}}\right|^2 + w_\tau \left|\hat{\tau}_{\mathrm{n}}\right|^2$$

This formulation directly penalises the harmonic energy contained within the hydraulic and mechanical outputs. The objective therefore considers both $\hat{P}_{\mathrm{n}}$ pressure excitation and $\hat{\tau}_{\mathrm{n}}$ shaft torque excitation.

Efficiency Objective

The hydraulic input power is $P_{in} = \Delta P Q$. The mechanical output power is $P_{out} = \tau\omega$. Therefore:

$$\eta = \frac{\tau\omega}{\Delta P\, Q}$$

The efficiency loss term is:

$$J_\eta = w_\eta (1 - \eta)$$

A perfectly efficient system gives $\eta = 1$ and therefore $J_\eta = 0$.

Cam Smoothness Objective

The cam profile must not only achieve low ripple but must also remain dynamically feasible. High jerk values produce increased vibration, higher contact forces, and manufacturing difficulty. The jerk penalty is defined as:

$$J_{jerk} = \int_0^2 \pi \left(\frac{d^3 s}{d\theta^3}\right)^2 d\theta$$

This term acts as a regularisation of the cam geometry.

Combined Objective Function

The individual performance terms are combined:

$$J_{phys} = w_1 R_\tau^2 + w_2 R_P^2 + w_3 (1 - \eta) + w_4 J_{spec} + w_5 J_{jerk}$$

where w_1, w_2, w_3, w_4, w_5 are weighting coefficients. Each coefficient determines the relative importance of the corresponding physical requirement.

Constraint Formulation

The objective function alone is insufficient because a mathematically optimal design may violate physical requirements. The constrained optimisation problem is therefore:

$$\min_{\mathrm{x}} J_{phys}(x)$$

subject to:

$$g_{\mathrm{i}}(x) \leq 0$$

where g_{i} are constraint functions.

Cavitation Constraint

The minimum chamber pressure must remain above the vapour pressure: $P(\theta) \geq P_{vap}$. The constraint is therefore:

$$g_1(x) = P_{vap} - min\big(P(\theta)\big)$$

The constraint is satisfied when $g_1 \leq 0$.

Geometric Stroke Constraint

The cam must maintain the required piston displacement range: $s_{min} \leq s(\theta) \leq s_{max}$. This can be represented as:

$$g_2(x) = \max s - s_{max} g_3(x) = s_{min} - \min s$$

Velocity Constraint

The maximum piston velocity must remain bounded: $|\dot{s}| \leq \dot{s}_{max}$, giving:

$$g_4(x) = max(|\dot{s}|) - \dot{s}_{max}$$

Acceleration Constraint

The acceleration limit is $|\ddot{s}| \leq \ddot{s}_{max}$, therefore:

$$g_5(x) = max(|\ddot{s}|) - \ddot{s}_{max}$$

Dynamic Stability Constraint

The equivalent hydraulic spring system gives $\omega\_n = \sqrt{\frac{k_f}{m_p}}$ and $\zeta = \frac{c_m}{2\sqrt{m_p k_f}}$. A minimum damping ratio is imposed: $\zeta \geq \zeta_{min}$, therefore:

$$g_6(x) = \zeta_{min} - \zeta$$

Penalty Formulation

For numerical optimisation, constraints may be incorporated into the objective using penalty functions. The penalised objective becomes:

$$J(x) = J_{phys}(x) + \lambda \sum_i \max\big(0, g_i(x)\big)^2$$

where λ is the penalty coefficient. This allows unconstrained optimisation algorithms to search only within physically meaningful regions.

Final Deterministic Optimisation Problem

The complete mathematical problem is:

$$\min_x w_1 R_\tau^2 + w_2 R_P^2 + w_3(1 - \eta) + w_4 J_{spec} + w_5 J_{jerk} + \lambda \sum_i \max(0, g_i)^2$$

The result is a single scalar objective that maps the complete hydraulic motor design space to a measurable performance value:

$$x \rightarrow J(x)$$

This deterministic objective forms the function evaluated by the surrogate-assisted optimisation framework.

Gaussian Process Surrogate Formulation

The deterministic objective function $J(x)$ provides an accurate evaluation of motor performance, but each evaluation requires the complete nonlinear physics model. To reduce the number of required simulations, the objective function is approximated using a Gaussian Process surrogate. The Gaussian Process does not attempt to replace the physical model. Instead, it constructs a probabilistic approximation of the objective landscape from previously evaluated designs. For a new design point, the surrogate provides $\mu(x)$, the predicted objective value, and $\sigma^2(x)$, the uncertainty of the prediction. These two quantities allow efficient exploration of the design space.

Gaussian Process Definition

A Gaussian Process is a distribution over functions. It is defined as:

$$f(x) \sim GP\big(m(x), k(x, x')\big)$$

where:

- $m(x) = E[f(x)]$ is the mean function
- $k(x, x')$ is the covariance function

The covariance determines the relationship between two points in the design space.

Objective Function Representation

The true objective is assumed to be generated by:

$$J(x) = f(x) + \varepsilon$$

where ε represents numerical uncertainty. The observation noise is modelled as:

$$\varepsilon \sim N(0, \sigma_n^2)$$

Therefore:

$$y_i = f(x_i) + \varepsilon_i$$

Training Dataset

The evaluated designs form the training set:

$$D = \{X, y\}$$

where X is the matrix of input design vectors and y is the corresponding vector of evaluated objective values.

Covariance Matrix Construction

The covariance matrix is defined element - wise as:

$$K_{ij} = k(x_i, x_j)$$

Including observation noise, the noisy covariance matrix becomes:

$$K_y = K + \sigma_n^2 I$$

Kernel Selection

The covariance function defines the assumed smoothness of the objective. A common engineering kernel is the squared exponential kernel:

$$k(x, x') = \sigma_f^2 exp\left(-\frac{1}{2}(x - x')^T L^{-2}(x - x')\right)$$

where σ_f^2 is the signal variance and L is the characteristic length - scale matrix. For d design variables, $L = diag(l_1, l_2, \ldots, l_d)$, where each l_i represents the sensitivity of the objective to variable x_i.

Joint Gaussian Distribution

For a new design point x_*, the GP prior gives the joint distribution between the training observations and the new point. The vector of covariances between the new point and training points is k_*, and the self-covariance of the new point is k_{**}.

Posterior Mean Derivation

Conditioning the joint Gaussian distribution gives the posterior mean:

$$\mu(x_*) = k_*^T(K + \sigma_n^2 I)^{-1} y_*$$

This represents the predicted objective value.

Posterior Variance Derivation

The conditional covariance is:

$$\sigma^2(x_*) = k_* - k_*^T(K + \sigma_n^2 I)^{-1} k_{**}$$

The uncertainty behaves as $\sigma(x_*) \to 0$ near known samples, while $\sigma(x_*)$ increases in unexplored regions.

Physics-Informed Residual Gaussian Process

The deterministic model already captures the dominant system physics. Therefore, rather than approximating $J(x)$ directly, the surrogate may model the residual:

$$r(x) = J_{full}(x) - J_{model}(x)$$

where J_{full} is obtained from the complete nonlinear simulation and J_{model} is the analytical approximation. The GP learns:

$$r(x) \sim GP(m_r, k_r)$$

The final prediction becomes:

$$J_{pred}(x) = J_{model}(x) + \mu_r(x)$$

with uncertainty:

$$\sigma_J^2(x) = \sigma_r^2(x)$$

Bayesian Optimisation Acquisition Function

The GP surrogate provides an estimate of the objective landscape. The next design point must balance selecting designs predicted to be optimal and exploring uncertain regions. For minimisation, the lower confidence bound acquisition function is:

$$a(x) = \mu(x) - \kappa\sigma(x)$$

where κ controls exploration.

Exploitation Term

The predicted objective, $\mu(x)$, encourages selection of designs with low expected cost. A small $\mu(x)$ therefore favours exploitation.

Exploration Term

The uncertainty term, $-\kappa\sigma(x)$, encourages evaluation of poorly sampled regions. Large uncertainty reduces $a(x)$, making those points attractive.

Candidate Selection

The next physics evaluation is chosen by:

$$x_{\{k+1\}} = argmin_x[\,\mu(x) - \kappa\sigma(x)]$$

The selected design is then evaluated as $J\left(x_{\{k+1\}}\right)$ and added to the training set. The GP is retrained:

$$D_{\{k+1\}} = D_k \cup \left(x_{\{k+1\}}, J_{\{k+1\}}\right)$$

and the process repeats.

Surrogate Optimisation Summary

The complete mathematical surrogate process is:

$$D \rightarrow GP \rightarrow \mu(x), \sigma(x) \rightarrow a(x) \rightarrow x_{new} \rightarrow J(x_{new})$$

The Gaussian Process therefore provides an adaptive approximation of the expensive hydraulic motor objective function while preserving the ability to quantify prediction uncertainty.

Optimisation Algorithm and Search Strategy

The deterministic motor model defines a mapping $x \rightarrow J(x)$, where $J(x)$ represents the physical performance of a candidate hydraulic motor design. However, the design space contains many interacting nonlinear effects: cam geometry, piston kinematics, hydraulic compressibility, port timing, torque harmonics, and efficiency losses. Direct optimisation of the full physics model is therefore computationally expensive.

The optimisation framework is formulated as a sequential model - based optimisation process, where a Gaussian Process surrogate is used to identify promising regions of the design space before performing expensive physics evaluations.

General Optimisation Formulation

The design problem is:

$$\min_{\mathrm{x}} J(x)$$

subject to:

$$g_{\mathrm{i}}(x) \leq 0$$

where x is the motor design vector and $J(x)$ is the deterministic physics objective.

Initial Design Sampling

Before the Gaussian Process can make predictions, an initial approximation of the design space must be generated. A set of initial samples is defined:

$$X_0 = \{x_1, x_2, \ldots, x_m\}$$

For each design x_i, the full deterministic model is evaluated as $y_i = J(x_i)$, forming the initial dataset:

$$D_0 = \{(X_0, Y_0)\}$$

The initial samples provide information about objective magnitude, design sensitivity, and feasible regions.

Gaussian Process Prediction

For a new design point x_*, the posterior mean prediction and associated uncertainty estimate are calculated using the formulations derived in the previous section. These two quantities define the surrogate prediction $J(x) \approx \mu(x)$ and uncertainty $\sigma(x)$ used by the acquisition function.

Surrogate - Based Search

The acquisition function is defined as:

$$a(x) = \mu(x) - \kappa\sigma(x)$$

where $\mu(x)$ represents exploitation of predicted low-cost regions and $\sigma(x)$ encourages exploration of uncertain regions. The next candidate design is selected:

$$x_{\{k+1\}} = argmin_x a(x)$$

and evaluated using the deterministic hydraulic motor model. The resulting objective value is then added to the training dataset, and the Gaussian Process is updated.

Physics Evaluation

The selected design is evaluated using the complete nonlinear motor model:

$$J_{\{k+1\}} = J\left(x_{\{k+1\}}\right)$$

The dataset is updated:

$$D_{\{k+1\}} = D_k \cup \left\{x_{\{k+1\}}, J_{\{k+1\}}\right\}$$

The Gaussian Process is then retrained. This sequential process continues until convergence criteria are reached.

Convergence Criteria

The optimisation may be terminated when one or more conditions are satisfied:

- Objective improvement: $\frac{|J_k - J_{\{k-1\}}|}{J_{\{k-1\}} < \varepsilon_J}$
- Design convergence: $\left|\left|x_k - x_{\{k-1\}}\right|\right| < \varepsilon_x$
- Iteration limit: $k > N_{max}$, where N_{max} is the maximum allowed number of physics evaluations.

Bayesian Optimisation Loop

The complete optimisation cycle is therefore:

1. Generate initial designs
2. Evaluate physics model
3. Train GP surrogate
4. Optimise acquisition function
5. Evaluate new design

6. Update GP
7. Repeat

This represents the original surrogate-assisted optimisation framework.

Local Refinement Using Nelder–Mead

The Nelder–Mead algorithm is a derivative-free optimisation method. For a design space with d variables, a simplex containing $d + 1$ vertices is constructed:

$$S = \{x_1, x_2, \ldots, x_{\{d+1\}}\}$$

Each vertex has an objective value $J(x_i)$. The vertices are ordered such that $J(x_1) \leq J(x_2) \leq \ldots \leq J(x_{\{d+1\}})$, where x_1 is the best point and $x_{\{d+1\}}$ is the worst point.

Simplex Centroid

The centroid of all points except the worst point is calculated:

$$x_c = \left(\frac{1}{d}\right)\sum_{i=1}^{d} x_i$$

Reflection Step

The worst point is reflected through the centroid:

$$x_r = x_c + \alpha(x_c - x_{\{d+1\}})$$

where α is the reflection coefficient. If $J(x_1) \leq J(x_r) < J(x_d)$, the reflected point replaces the worst point.

Expansion Step

If reflection produces a significantly improved point ($J(x_r) < J(x_1)$), an expansion step is attempted:

$$x_e = x_c + \gamma(x_r - x_c)$$

where γ is the expansion coefficient. The better of x_e and x_r is retained.

Contraction Step

If reflection fails to improve the simplex ($J(x_r) \geq J(x_d)$), a contraction is performed:

$$x_{con} = x_c + \rho(x_{\{d+1\}} - x_c)$$

where ρ is the contraction coefficient.

Shrink Step

If contraction fails, the simplex contracts around the best point:

$$x_i = x_1 + \sigma(x_i - x_1)$$

for $i = 2, \ldots, d + 1$.

Hybrid Optimisation Procedure

The final implemented optimisation procedure is therefore:

$$GP\ Bayesian\ Search \rightarrow Best\ Candidate\ Region \rightarrow Nelder - Mead\ Refinement$$

The GP stage provides global exploration, uncertainty-guided search, and reduced physics evaluations. The Nelder–Mead stage provides local improvement, derivative-free refinement, and convergence near the local optimum.

Complete Optimisation Algorithm

The complete procedure is:

1. Generate initial design samples
2. Evaluate nonlinear physics model
3. Construct GP surrogate
4. Optimise acquisition function
5. Evaluate selected design
6. Update GP
7. Repeat until convergence
8. Apply Nelder - Mead local refinement

The resulting design is the solution:

$$x = argmin_x J(x) *$$

within the explored design space.

Summary

The optimisation framework combines a physics - based motor model with a probabilistic surrogate search strategy. The Gaussian Process Bayesian optimisation stage reduces the number of expensive nonlinear simulations required by intelligently selecting promising designs. A final Nelder–Mead refinement stage may then be applied as an implementation enhancement to improve convergence around the best design identified by the global search. The complete workflow therefore combines:

$$Physics\ Model + Gaussian\ Process\ Search + Local\ Derivative - Free\ Refinement$$

to efficiently explore the nonlinear hydraulic motor design space.