

Patch-based Thermodynamic Property Models for the Subcritical Region

Laughman, Christopher R.; Qiao, Hongtao

TR2021-053 June 04, 2021

Abstract

Model-based design approaches for vapor-compression cycles depend heavily upon refrigerant property representations that are fast, accurate, and consistent. We describe an approach based upon B-spline interpolants that describes properties such as density, temperature, and specific entropy as the intersection of multiple surfaces, which are referred to as “patches.” When combined with a transformation of thermodynamic coordinates, this approach can calculate the density over a domain with a maximum absolute percentage error less than 3×10^{-5} and a speedup over REFPROP of greater than 100x.

Purdue Air-Conditioning and Refrigeration Conference

© 2021 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Patch-based Thermodynamic Property Models for the Subcritical Region

Christopher LAUGHMAN*, Hongtao QIAO

Mitsubishi Electric Research Laboratories
Cambridge, MA, USA
{laughman, qiao}@merl.com

* Corresponding Author

ABSTRACT

Model-based design approaches for vapor-compression cycles depend heavily upon refrigerant property representations that are fast, accurate, and consistent. We describe an approach based upon B-spline interpolants that describes properties such as density, temperature, and specific entropy as the intersection of multiple surfaces, which are referred to as "patches." When combined with a transformation of thermodynamic coordinates, this approach can calculate the density over a domain with a maximum absolute percentage error less than 3×10^{-5} and a speedup over REFPROP of greater than 100x.

1. INTRODUCTION

Representations of thermophysical refrigerant properties form the core of any vapor compression cycle model because they provide the algebraic connections between the thermodynamic variables used to describe system behavior, such as pressure, temperature, density, and specific enthalpy. The precise nonlinear behavior of these properties over the domain of operation can qualify or disqualify a refrigerant for a particular application and can impact system-level design specifications (McLinden & Huber, 2020). Accurate representations of the thermophysical behavior of these fluids are thus essential to the model-based design of complex vapor compression cycles required by next-generation heating, air-conditioning, and refrigeration applications.

The calculation of these properties tends to be computationally expensive, particularly in the context of physics-based cycle models; it is common that up to 90% of the function calls in finite volume-based multiphase heat exchanger models are dedicated to the calculation of refrigerant properties (Aute & Radermacher, 2014). The underlying phase equilibrium computations required typically involve multiple layers of iteration because the boundaries between phase regions are defined by the equality of the pressure, temperature, and chemical potentials in the phase space. Moreover, while state of the art methods to characterize fluid properties are typically formulated in terms of the Helmholtz energy, which is expressed as a function of the fluid temperature and density, the fact that many cycle computations are more naturally expressed in terms of pressure and specific enthalpy often necessitates further changes of variables and thereby accompanies additional levels of iteration. Because any practical model-based design process is sensitive to the computation time required for its constituent tools, the refrigerant property models must be fast to prevent simulation time from limiting the utility of these approaches.

Consistency between property variables is also an important characteristic of these models for two related reasons. First, an engineering-based modeling approach must ensure the correct mathematical solution of the underlying equations without depending on the order of these equations. For example, $h(P_0, \rho(P_0, h_0^*)) - h_0^*$ (where $h(P, \rho)$ and $\rho(P, h)$ represent the computation of specific enthalpy and density from the related quantities) should be less than or equal to the numerical tolerance of the computation. In addition, the fact that some variables are related only through their thermodynamic derivatives imposes significant constraints on the relation between the variables and their derivatives. Consider the following standard formulation of mass conservation for heat exchanger models, e.g.,

$$\frac{d(\rho V)}{dt} = \dot{m}_{in} - \dot{m}_{out}. \quad (1)$$

Solving this equation when pressure and specific enthalpy are integrated by the differential equation solver requires an

additional equation to express the change of variables, which can be written as

$$\frac{d\rho}{dt} = \left(\frac{\partial\rho}{\partial p}\right)_h \frac{dp}{dt} + \left(\frac{\partial\rho}{\partial h}\right)_p \frac{dh}{dt}. \quad (2)$$

The property models should ensure that the integral of these density derivatives is equivalent to the density itself, as inconsistency between these quantities can lead to numerical artifacts which cause nonphysical refrigerant charge dynamics (Laughman & Qiao, 2017). Particular care must be paid to these details because the density derivatives are not continuous across the saturation lines, as can be seen by computing the derivative of density with respect to specific enthalpy (Qiao, 2014), first for the single-phase region,

$$\left(\frac{\partial\rho}{\partial h}\right)_p = \frac{(\partial\rho/\partial T)_p}{(\partial h/\partial T)_p} = -\frac{\beta\rho}{c_p}, \quad (3)$$

where

$$\beta = -\frac{1}{\rho} \left(\frac{\partial\rho}{\partial T}\right)_p \quad c_p = \left(\frac{\partial h}{\partial T}\right)_p, \quad (4)$$

and then for the two-phase region,

$$\frac{\partial\rho}{\partial h} = \frac{\partial\rho}{\partial v} \frac{\partial v}{\partial h} = -\rho^2 \frac{v_g - v_f}{h_g - h_f}. \quad (5)$$

These discontinuous derivatives must be correctly described by the property model to generate accurate cycle simulations.

The importance of accurately, quickly, and consistently computing refrigerant properties has long motivated exploration and innovation in creating corresponding models. While there is extensive literature on the underlying physics and computational mathematics of fundamental refrigerant property models, Miyagawa & Hill (2001) represent a prime example of previous work in developing property models that are both accurate and computationally efficient via table-based Taylor series expansions. Further refinements of this method are proposed by Kunick (2018), which demonstrates significant performance improvements via the use of spline-based approaches, to which the work in this paper is deeply indebted. Similar methods are described and applied specifically to vapor compression cycles by Li et al. (2018). Wang et al. (2014) also describes spline-based methods, and uses extrapolation approaches similar to those explored in this paper. A related method is used by Aute & Rademacher (2014), who describe the application of Chebyshev polynomials to approximate refrigerant properties for similar applications. Finally, this paper builds on the authors' previous work (Laughman & Van, 2018), which served as both an initial proof-of-concept and highlighted shortcomings that will be addressed here.

The accurate description of refrigerant property behavior near the saturation lines represents the main focus of this work. The requirement of accurately describing discontinuous derivatives motivated the use of multi-surface representations of the refrigerant properties, where one surface characterizes the single-phase behavior and a second surface characterizes the two-phase behavior, and the overall properties are computed from the intersection of these surfaces, which are also referred to as "patches." Furthermore, we build these individual surfaces from B-splines because they enable the surface derivatives to be computed directly from the surface coefficients, rather than characterizing the $d\rho/dP$ and $d\rho/dh$ surfaces independent from the density surface itself. Finally, poor numerical behavior around the saturation lines related to their curvature in (P, h) coordinates prompted us to implement a coordinate change to pressure and thermodynamic quality x so that one of the coordinate axes is parallel to the saturation lines. While this specific coordinate transformation can only be applied in the subcritical region, the output of this transformation can be blended together with similarly efficient property representations around the critical point and in the supercritical region.

This approach to modeling refrigerant properties will be described and demonstrated over the course of this paper. Section 2 describes the method for formulating the refrigerant properties using the patch-based B-spline approach, as well as on the implementation of these methods in the Julia language (Bezanson et al., 2017). The B-spline and patch-based approaches can be applied to both standard (P, h) and (P, x) coordinates, so we refer to both h and x as "mixture coordinates" in this section as an indication of their generality. Section 3 presents the results of these studies, and we draw some conclusions in Section 4 and discuss next steps for this work.

2. METHOD & IMPLEMENTATION

The present method for modeling refrigerant properties can be described via three related ideas: the use of B-splines to fit property surface data as a function of two independent variables, the use of patches to accurately describe the discontinuities in the underlying property surface, and the use of a coordinate transformation between (P, h) and (P, x) to avoid numerical artifacts caused by the curvature of the saturation lines. We describe the details of each of these ideas in turn, after which we briefly review the most salient details of the method's implementation.

2.1 B-Splines

In general, splines are piecewise polynomial functions which map $\mathbb{R} \rightarrow \mathbb{R}^n$ that constitute interpolants on scalar or vector data, referred to as control points, and which are defined on a specified set of intervals of \mathbb{R} , referred to as a knot vector. While there are many different spline formulations, B-splines are well suited to the problem of describing refrigerant properties because they are numerically well-conditioned and because of the simple computational relationship between the spline surface and its derivatives.

Following the terminology of Piegl & Tiller (1995), B-splines are defined over a general knot vector $U = \{u_0, \dots, u_m\}$, which is a nondecreasing sequence of real numbers, i.e., $u_i \leq u_{i+1}$, $i = 0, \dots, m-1$. In this case, the B-spline basis function of order $p+1$, denoted by $N_{i,p}(u)$, can be defined recursively by the Cox-de Boor formula as

$$N_{i,0}(u) = \begin{cases} 1 & \text{if } u_i \leq u < u_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

$$N_{i,p}(u) = \frac{u - u_i}{u_{i+p} - u_i} N_{i,p-1}(u) + \frac{u_{i+p+1} - u}{u_{i+p+1} - u_{i+1}} N_{i+1,p-1}(u). \quad (7)$$

The sparsity of these functions is evident from Equations 6- 7 as only the basis functions from $p+1$ intervals around a given point u are nonzero, so that the computation of the interpolant has a fixed cost regardless of the size of the data. Cubic B-splines of degree 3 have the following properties of particular value:

1. The functions are C^2 continuous everywhere as long as there are no repeated knots.
2. The coefficients for all non-zero for a given point sum to 1, so that they are numerically well-conditioned.
3. The resulting interpolating curve passes through the control points.
4. The recursive definition of the basis functions is helpful for computation.

In addition, the derivatives of these basis functions are given by

$$N'_{i,p} = \frac{p}{u_{i+p} - u_i} N_{i,p-1}(u) - \frac{p}{u_{i+p+1} - u_{i+1}} N_{i+1,p-1}(u), \quad (8)$$

so that the derivatives of the basis functions can be computed directly from the basis functions themselves.

A p^{th} degree B-spline curve $\mathbf{C}(u)$ is generated by the inner product of the nonzero basis functions and the pertinent set of control points \mathbf{P}_i , e.g.,

$$\mathbf{C}(u) = \sum_{i=0}^n N_{i,p}(u) \mathbf{P}_i \quad a \leq u \leq b \quad (9)$$

where the knot vector is given by

$$U = \underbrace{\{a, \dots, a, u_{p+1}, \dots, u_{m-p-1}, b, \dots, b\}}_{p+1} \quad (10)$$

The derivative of this curve is given by

$$\mathbf{C}^{(k)}(u) = \sum_{i=0}^n N_{i,p}^{(k)}(u) \mathbf{P}_i, \quad (11)$$

where k represents the order of the derivative. The fact that the derivative of the spline is represented by the same control points multiplied by the derivatives of the basis functions improves the accuracy and consistency of the derivative curves and conserves memory. These curves were used to characterize the saturation lines, such as the liquid and vapor saturated specific enthalpies, as a function of the pressure.

B-spline surfaces are defined via a net of control points in each of the coordinate axes and multiplying the control points by a set of univariate spline coefficients for each axis, as given by

$$\mathbf{S}(u, v) = \sum_{i=0}^n \sum_{j=0}^m N_{i,p}(u) N_{j,q}(v) \mathbf{P}_{i,j} \quad (12)$$

with the knot vectors in each coordinate direction defined analogously to the univariate curve. The derivatives of this surface are defined as

$$\frac{\partial^{k+l}}{\partial^k u \partial^l v} \mathbf{S}(u, v) = \sum_{i=0}^n \sum_{j=0}^m N_{i,p}^{(k)} N_{j,q}^{(l)} \mathbf{P}_{i,j}, \quad (13)$$

similar to the derivatives of the curve.

Given the mathematical structure of these basis functions for both curves and surfaces, their practical use for interpolation requires the specification of knot vectors for both coordinate axes and the calculation of the control points based on the reference data. While a uniformly spaced knot vector along the pressure axis was sufficient, the large variations in the density derivative motivated the use of increased knot density along the mixture coordinate axis in the vicinity of the liquid saturation curve. This knot vector was constructed with smooth changes in the spacing by defining a smooth curve that represented the derivative of the knot vector, which was then integrated and sampled to determine the location of the knots, because large changes in the knot spacing can cause oscillations in the interpolation output.

The calculation of the control points is straightforward and can be accomplished via least squares methods given a set of reference data, as the interpolant is a linear combination of basis functions. The set of nonzero basis functions is first determined via the knot vector for a specific data point, and the control points are then calculated using least squares methods. This is performed sequentially for each coordinate axis in the case of the density surface fit. As the condition number of the basis function matrix can be quite large, Tikhonov regularization was used to improve the numerical stability of the solution of the control points by adding a constant $\lambda = 10^{-5}$ to the diagonal of the basis function matrix.

The computation of the reference property data from which these interpolants are built is obtained from REFPROP (Lemmon et al., 2018), and can normally be accomplished via standard REFPROP function calls when constructing models for pure fluids. However, the iterations used by REFPROP for complex fluid mixtures sometimes do not converge around the saturation lines close to the critical point. In this case, filtering or other data cleaning measures must be taken to ensure that the property surfaces do not contain artifacts related to nonconvergence of these iterative routines.

2.2 Patch-Based Approach

B-splines can accurately describe thermodynamic property surfaces because these surfaces are everywhere smooth except at the saturation lines, but these edges require special treatment to ensure correct behavior in their vicinity. While the most conventional approach for representing saturation lines smoothes over the derivative discontinuities, we seek to capture the discontinuities in the property models because of their potential importance in the overall system dynamics. B-splines can potentially be used to describe surfaces with edges by using repeated knots, but these were not found to be effective under standard parameterizations due to numerical artifacts in the surface derivatives.

Because the saturation lines represent the only points of discontinuity, we can construct independent smooth approximations of the single and two-phase surfaces and combine them via a nonlinear function. We call these separate representations "patches", in reference to the analogous computer graphics terminology. Given that the surface in each of these phase regions can essentially be described by a quadratic form in near the saturation lines, we use quadratic extrapolation to extend each region across the saturation lines so that the derivatives for each patch are continuous in the vicinity of these lines. The discontinuity is then captured by switching between the surfaces.

For example, the data for the surface fit for the two-phase region (the two-phase "patch" $\rho_{\varphi 2}$) includes all of the points in the two-phase region, as well as the saturation lines, and then uses quadratic extrapolation to extend the surface into the portion of the single-phase region adjacent to the saturation lines. The same process is followed to construct the single-phase patch $\rho_{\varphi 1}$, which is continuous throughout the entire single-phase region and extends into the two-phase




Figure 1: Patches ρ_{φ_1} and ρ_{φ_2} for R32 at 1 MPa compared against reference data from REFPROP.




Figure 2: Patches ρ_{φ_1} and ρ_{φ_2} for R32 in the neighborhood of the liquid saturation curve at low pressure.

region. The resulting property surface is thus described by the intersection of these two patches that cover the entire domain of interest.

The thermodynamics of the property under study and the coordinate system can be used to design simple nonlinear functions that take in the interpolated output from both patches and return the correct property value. When the density surface is approximated using specific enthalpy as a mixture variable, the fact that $d\rho/dh|_P$ has a jump in magnitude as the enthalpy increases across both saturation lines enables the use of the simple nonlinear function $\rho = \min(\rho_{\varphi_1}, \rho_{\varphi_2})$ to calculate the correct value. Under a coordinate transformation from specific enthalpy to thermodynamic quality x , the value of x can be directly used in an `if` statement to select the appropriate surface. Moreover, the derivative discontinuities can be accurately reproduced because information about the specific patch used at any coordinate in the thermodynamic space also provides information about the value of the derivative at that point.

Figure 1 illustrates this construction for the density of the refrigerant difluoromethane (R32) at 1 MPa over a range of specific enthalpies. The ρ_{φ_2} surface is built from the points in the two-phase region by using quadratic extrapolation to extend the two-phase density surface into the single phase regions, with an additional limit imposed to ensure that the magnitude of the density stays within reasonable bounds. The resulting two-phase surface is denoted by the dashed line. Similarly, the ρ_{φ_1} surface is constructed by using quadratic extrapolation to extend the single-phase surfaces from the liquid and vapor regions into the two-phase region. These two extrapolations are then connected to form a contiguous surface that is represented by the dotted line. The minimum of these two surfaces can be seen to match the reference data obtained from REFPROP, which is marked by the circles.

The effect on the overall surfaces can be seen clearly in Figure 2 for a low-pressure region around the liquid saturation line. The liquid saturation line is represented by a smooth curve, but that the density derivatives will manifest the correct discontinuous jump as the choice of these patches switches across the saturation line. Similar functions can be constructed for other properties of interest, such as temperature or specific entropy.

2.3 Coordinate Transformations

Though (P, h) represents the most natural coordinate system for common calculations in vapor-compression cycles, the resulting curvature of the saturation lines can pose challenges when developing approximations of thermodynamic properties. Of particular concern is the fact that $P(h)$ is not single-valued along the vapor saturation line for a range of specific enthalpies up to the maximum specific enthalpy on this curve, which is also referred to as the maxcondenthalp (MCH) point. The vanishing length of an isobaric path in the two-phase region as the MCH point is approached from lower specific enthalpies often gives rise to numerical artifacts and oscillations in the property surfaces in this region that are large enough to cause problems during integration in differential equation solvers.

This problem may be avoided by transforming the coordinate system from (P, h) into one in which the saturation lines have no curvature, and instead run parallel to one of the coordinate axes. One such transformation maps the specific




Figure 3: $d\rho/dh|_P$ in the (P, h) coordinate system.




Figure 4: $d\rho/dh|_P$ in the (P, x) coordinate system.

enthalpy h to the thermodynamic quality x , which is defined as

$$x(P, h) = \frac{h - h_f(P)}{h_g(P) - h_f(P)}, \quad (14)$$

Note that the thermodynamic quality is distinct from conventional definitions of static quality, though they are qualitatively similar; while the static quality has a minimum of 0 and a maximum of 1, the thermodynamic quality can go above and below zero with pressure dependent limits on the on upper and lower ranges of this coordinate transformation. In the case of R32 with the standard default REFPROP specific enthalpy reference, the lower limit of x corresponding to $h = 0$ at 300 kPa is -0.432, while the lower limit of x corresponding to $h = 0$ at 5.5 MPa is -4.63. This coordinate transformation cannot be used close to the critical point because $h_f(P_{crit}) = h_g(P_{crit}) = h_{crit}$, rendering the transformation singular. Nevertheless, the accuracy of this transformation is quite beneficial for common subcritical cycle models.

The effect of this coordinate transformation can be seen by comparing the $d\rho/dh|_P$ surface in the (P, h) coordinate system (Figure 3) to the same surface in the (P, x) coordinate system (Figure 4), where the black lines represent contours of constant magnitude. Whereas the curvature of the saturation lines is evident in Figure 3, the straight iso-quality lines at $x = 0$ and $x = 1$ in Figure 4 simplifies the extrapolation process used to generate the surface patches corresponding to the single-phase and two-phase regions and eliminates the corresponding numerical artifacts. These coordinate axes are also helpful because they facilitate the use of a more regular grid, reducing the need for computationally expensive searching to find node locations.

Because the mass, momentum, and energy conservations equations are usually expressed in (P, h) coordinates, cycle models must first compute the coordinate transform to (P, x) , calculate the thermodynamic properties in this space, and then transform the resulting values back into (P, h) coordinates. Such a conversion is straightforward for the fundamental thermodynamic variables (ρ, T, u, s), but is more complex for their derivatives. As illustrated in Equation 2, the derivatives $\partial\rho/\partial h|_P$ and $\partial\rho/\partial P|_h$, are often used in expressing mass conservation, but the surface derivatives available from the B-splines are instead $\partial\rho/\partial x|_P$ and $\partial\rho/\partial P|_x$. The relation between these terms can be derived by calculating the differentials of the variables of interest and matching the corresponding terms. The differential of the specific volume $v = v(P, x)$ is given by

$$dv = \left(\frac{\partial v}{\partial P}\right)_x dP + \left(\frac{\partial v}{\partial x}\right)_P dx, \quad (15)$$




Figure 5: Absolute percentage error along liquid (upper) and vapor (lower) saturation lines.

θ_{sat}	REFPROP	Spline
ρ_f	7.055 s	569.8 μ s
ρ_g	7.368 s	567.4 μ s
θ_{surf}	REFPROP	Spline
ρ @ 500 kPa	3.002 s	758.1 μ s
ρ @ 2.5 MPa	3.084 s	797.1 μ s
ρ @ 5.0 MPa	3.279 s	796.4 μ s

Table 1: Computational time comparison between REFPROP and B-splines.

while the definition of x in Equation 14 enables us to write its differential as a function of P and h , yielding

$$dx = \frac{1}{h_g - h_f} \left(dh - \left[\frac{dh_f}{dP} + x \left(\frac{dh_g}{dP} - \frac{dh_f}{dP} \right) \right] dP \right). \quad (16)$$

Thus, combining expressions and collecting terms, we have

$$dv = \left(\frac{dv}{dP} \Big|_x - \frac{1}{(h_g - h_f)} \frac{dv}{dx} \Big|_P \left[\frac{dh_f}{dP} + x \left(\frac{dh_g}{dP} - \frac{dh_f}{dP} \right) \right] \right) dP + \frac{dv}{dx} \Big|_P \frac{1}{(h_g - h_f)} dh, \quad (17)$$

and this can be transformed back into expressions for the relevant density derivatives, e.g.,

$$\frac{d\rho}{dP} \Big|_h = \left(\frac{d\rho}{dP} \Big|_x - \frac{1}{(h_g - h_f)} \frac{d\rho}{dx} \Big|_P \left[\frac{dh_f}{dP} + x \left(\frac{dh_g}{dP} - \frac{dh_f}{dP} \right) \right] \right), \quad (18)$$

$$\frac{d\rho}{dh} \Big|_P = \frac{1}{(h_g - h_f)} \frac{d\rho}{dx} \Big|_P. \quad (19)$$

Similar transformations can be implemented for other variables, such as temperature, specific internal energy, or specific entropy, to facilitate the incorporation of these property models into existing cycle models.

2.4 Implementation

Version 1.5.2 of the Julia language was used to implement these property models due to its high-performance numerical computing capabilities and its extensive support of metaprogramming, which was used to create an interface to existing REFPROP libraries. The advanced capabilities of its ecosystem of differential equation solvers¹ also suggests significant potential for applications to challenging problems in cycle simulation.

The interface between Julia and REFPROP was of particular importance, due to the need for reference data in the spline computation process. Property values were generated in REFPROP by using the Python-based ctREFPROP wrapper², which instantiates a connection to the REFPROP dll in Windows and interfaces to Julia via the PyCall package³. This flow of data facilitated rapid prototyping for these property codes, but imposed a computational penalty on the REFPROP algorithm performance as compared to the speed of the underlying Fortran implementation.

¹<https://sciml.ai/>

²<https://github.com/usnistgov/REFPROP-wrappers>

³<https://github.com/JuliaPy/PyCall.jl>




Figure 6: Absolute percentage error of density over the validation domain.




Figure 7: Absolute percentage error of $d\rho/dh|_P$ over the validation domain.

3. RESULTS

We evaluated the performance of this method for approximating thermodynamic properties by implementing a set of models for the refrigerant R32. This included approximations of the liquid and vapor saturation curves, as well as of the surface properties. We first described the saturation lines as a function of pressure by using REFPROP to sample the density, temperature, and specific enthalpy along the saturation lines from 300 kPa to 5.6 MPa every 5 kPa, and then built the corresponding 1-D B-splines. The performance of this approximation was then compared to the performance of REFPROP by evaluating both of these functions at every 1 kPa for this set of outputs. We also compared the derivatives of density along the liquid saturation line ($d\rho_f/dP$) and along the vapor saturation line ($d\rho_g/dP$) as well as the derivative of temperature with respect to pressure (dT/dP) along the saturation line. These derivatives were computed from the spline coefficients, rather than from a separate fit to derivative data obtained from REFPROP.

The absolute value of the percentage error (APE) between the output of REFPROP and the B-spline approximation is illustrated in Figure 5, while the maximum of the APE as well as the root mean squared error (RMSE) for the quantities under consideration are given in the upper portion of Table 2. These errors are generally quite small, with no value of the maximum APE greater than 10^{-8} for any property values or greater than 2×10^{-5} for their derivatives. This performance will likely be quite adequate for many applications, as the errors are on the order than the square root of machine epsilon for double precision arithmetic.

The comparison of the computation time in Julia for the full set of 5300 saturation data points is provided in Table 1, which demonstrates that the B-spline approach is quite favorable in this respect. These values represent the time required to obtain property values in Julia, and do not accurately represent the use of REFPROP in other environments, which may be faster. Nevertheless, the significant increase in speed of the spline-based methods, even in light of such considerations, is apparent.

The properties over the (P, h) domain were constructed in an analogous manner. The same pressure points were used in the surface construction, while nonuniformly spaced specific enthalpy values between 100 kJ/kg and 700 kJ/kg were used to define the mixture coordinate data. These values were translated to the corresponding pressure-dependent axis limits in the thermodynamic quality, and the spacing over most of the single-phase domain was selected to be 0.01. Additional samples were required in the two-phase region adjacent to the liquid saturation line due to the large magnitude of the density derivatives, and these were spaced according to the approach described at the end of Section 2.1, with a minimum spacing of 0.0002 in the interval $x = [-0.1, 0.1]$ and an increased spacing outside of this interval. This resulted in a data grid of 1060×1564 samples from which the B-spline surface patches were constructed. This sampling density could potentially be reduced, but such optimizations were not explored in the context of this work.

θ_{sat}	max(APE)	RMSE	θ_{sat}	max(APE)	RMSE
T	1.123e-09	6.387e-09 K	dT/dP	3.984e-06	5.815e-12 K/Pa
h_f	2.827e-09	1.415e-05 J/kg	h_g	1.294e-09	1.406e-05 J/kg
ρ_f	3.710e-09	5.132e-08 kg/m ³	ρ_g	7.182e-09	4.827e-08 kg/m ³
$d\rho_f/dP$	1.012e-05	4.695e-11 kg/(m ³ Pa)	$d\rho_g/dP$	1.112e-05	4.418e-11 kg/(m ³ Pa)
θ_{surf}	max(APE)	RMSE	θ_{surf}	max(APE)	RMSE
T	2.097e-05	6.445e-03 K	$d\rho/dP _h$	1.547e-03	1.171e-08 kg/(m ³ Pa)
ρ	2.060e-05	1.228e-02 kg/m ³	$d\rho/dh _p$	4.036e-04	1.526e-07 kg ² /(m ³ J)
s	2.503e-05	3.393e-02 J/(kg · K)			

Table 2: Accuracy of refrigerant property approximations over saturation lines (upper portion) and over (P, h) validation domain (lower portion).

The accuracy of these property models was then evaluated over a more finely sampled grid in which the pressure vector was sampled at every 2 kPa, and the specific enthalpy was sampled every 288.3 J/kg (15 J/mol). This resulted in a validation grid of 2666×2133 points over which the density, temperature, and specific entropy were sampled, as well as the density derivatives.

Figure 6 illustrates the small APE for the density over the domain of interest, Figure 7 illustrates the same error metric for the derivative $d\rho/dh|_p$, and the numerical values of these errors are provided in the lower portion of Table 2. It is again of particular note that the $d\rho/dh|_p$ surface was computed directly from the coefficients of the density surface, rather than from a separate set of reference data. This provides an indication of the excellent consistency of this data-efficient approach. It is also evident from Figure 6 that the mean absolute error over the domain is $2e-5$, while the deviations in these errors are concentrated around the saturation lines. The magnitude of the errors are connected to the regularization used in the spline coefficient computation, which was chosen to trade off accuracy over the majority of the domain against numerical artifacts in the neighborhood of the saturation lines. These remaining deviations are also related to the nonuniform spacing of the thermodynamic quality vector; while a uniformly spaced vector could be used with the minimum sample spacing to maintain accuracy close to the liquid saturation line, such a vector was judged to be impractical because it would be approximately 50x larger than that which is currently used.

The computational time required to interpolate the refrigerant surface properties was assessed by recording the amount of time needed to compute all values for the validation specific enthalpy vector, as described earlier in this section, for three different pressures. The results of these experiments are provided in Table 1, and again demonstrate a significant speedup over the equivalent REFPROP functions. We expect that the time required by REFPROP will increase with increasing pressure due to the change in the curvature of the saturation lines and the concomitant increase in the number of iterations required for the flash calculations, but that the spline calculations will require essentially constant time. Note that the additional time required to evaluate the splines on the surface can be attributed to the fact that interpolation is required in two directions for the surface, rather than one for the saturation lines.

4. CONCLUSION & DISCUSSION

Patch-based B-spline interpolants based on a coordinate transformation from (P, h) to (P, x) can be seen to approximate thermodynamic properties in a accurate, consistent, and computationally efficient manner. The compact support of B-splines facilitates the high computational speed and accuracy of these models, while the patch-based approach is able to capture the derivative discontinuities well. Finally, the (P, x) coordinate transformation improves this method's numerical performance and eliminates artifacts that accompany approximations in other coordinate systems.

While the results presented here indicate the benefits of this approach, further work is required to complete a full-fledged set of property models. Perhaps most obvious is the current restriction to the subcritical domain; additional effort is required to augment these models so that they apply near the critical point and in the supercritical region. In addition, these techniques will likely require modification and extension to work on refrigerant mixtures. Finally, this work motivates further study into techniques that will further reduce deviations between reference data and the approximations, with the ultimate goal being a set of codes for which REFPROP's output and the approximated output are indistinguishable in dynamic cycle modeling applications.

NOMENCLATURE

P	pressure	[Pa]	t	time	[s]
T	temperature	[K]	x	thermodynamic quality	[-]
V	volume	[m ³]	θ	thermodynamic property	[-]
h	specific enthalpy	[J/kg]	ρ	density	[kg/m ³]
\dot{m}	mass flow rate	[kg/s]	v	specific volume	[m ³ /kg]
s	specific entropy	[J/(kg·K)]			

Subscript

f	liquid
g	vapor

REFERENCES

- Aute, V., & Radermacher, R. (2014). Standardized polynomials for fast evaluation of refrigerant thermophysical properties. In *International Refrigeration and Air-Conditioning Conference at Purdue*.
- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1), 65–98. Retrieved from <https://doi.org/10.1137/141000671>
- Kunick, M. (2018). *Fast calculation of thermophysical properties in extensive process simulations with the spline-based table look-up method (SBTL)* (No. 618). Fortschritt-Berichte VDI.
- Laughman, C., & Qiao, H. (2017). On the influence of state selection on mass conservation in dynamic vapor compression cycle models. *Math Comput Model Dyn Syst*, 23, 262-283.
- Laughman, C., & Van, C. (2018). Approximation of refrigerant properties for dynamic vapor compression cycle models. In *9th Vienna International Conference on Mathematical Modelling (MATHMOD)* (p. 625-630).
- Lemmon, E. W., Bell, I., Huber, M. L., & McLinden, M. O. (2018). *NIST Standard Reference Database 23: Reference Fluid Thermodynamic and Transport Properties-REFPROP, Version 10.0*, National Institute of Standards and Technology. Retrieved from <https://www.nist.gov/srd/refprop> doi: <https://doi.org/10.18434/T4/1502528>
- Li, L., Gohl, J., Batteh, J., Greiner, C., & Wang, K. (2018). Fast calculation of refrigerant properties in vapor compression cycles using spline-based table look-up method (SBTL). In *Proceedings of the American Modelica Conference 2018*.
- McLinden, M., & Huber, M. (2020). (R)Evolution of refrigerants. *Journal of Chemical and Engineering Data*, 65, 4176-4193. doi: 10.1021/acs.jced.0c00338
- Miyagawa, K., & Hill, P. (2001, Jul). Rapid and accurate calculation of water and steam properties using the tabular taylor series expansion method. *Journal of Engineering for Gas Turbines and Power*, 123, 707-712.
- Piegl, L., & Tiller, W. (1995). *The NURBS Book* (2nd ed.). Springer.
- Qiao, H. (2014). *Transient modeling of two-stage and variable refrigerant flow vapor compression systems with frosting and defrosting* (Doctoral dissertation, University of Maryland, College Park). doi: 10.13016/M24892
- Wang, X.-D., Wang, Z.-X., Duan, Y.-Y., An, B., & Lee, D.-J. (2014). Efficient evaluation of thermodynamic properties of water and steam on P-h surface. *Journal of the Taiwan Institute of Chemical Engineers*, 45, 372-379. doi: 10.1016/j.jtice.2013.06.016