

Mass Conservation in Vapor Compression Cycles: A Method for Ensuring Consistency with Redundant Dynamic States

Daniel Andersson¹ John Batteh²
Matthis Thorade³ Lixiang Li⁴

¹Modelon AB, Sweden, {daniel.andersson}@modelon.com

^{2,4}Modelon Inc., USA, {john.batteh, lixiang.li}@modelon.com

³Modelon Deutschland GmbH, Germany, {matthis.thorade}@modelon.com

Abstract

This paper describes a method to resolve a potential inconsistency when employing redundant dynamic thermofluid states for modeling of vapor compression cycles. Following a brief introduction regarding the motivation and use of redundant thermofluid states, a series of test models ranging from simple component models to complex system models are developed to illustrate the potential inconsistency with Air Conditioning Library. Based on observations of the simulation results from these test models, a method for ensuring consistency is proposed and implemented. The method is then demonstrated on the test suite and evaluated for effectiveness, robustness, and computational efficiency.

Keywords: thermofluid modeling, vapor compression cycles, two phase systems, thermodynamics, state selection

1 Introduction

State selection is a key element of thermofluid modeling. Selecting appropriate dynamic states, or independent thermodynamic variables, is critical for ensuring robust, accurate, and computationally efficient models (Tummescheit 2002). In particular, the following considerations are often made when considering state choices for thermofluid models:

- Thermodynamic considerations to allow full phase identification over the operating regime for the fluid (i.e. single phase, multi-phase, etc.)
- State compatibility with medium model implementation (i.e. choosing states such that medium model can be evaluated explicitly from states as much as possible to avoid unnecessary nonlinear systems)
- Formulation for conservation laws explicitly in state variables (either manually or by Modelica compiler)

- Initialization structure to avoid unnecessary nonlinear systems
- Initialization ease (i.e. typical values specified by the user to describe initial state of system including the use of intensive or extensive variables)

Given the importance of state selection and the tight integration with the medium model implementation, there has been significant focus on state selection and on the development of compatible media packages in the design of the Modelica.Fluid package (Franke 2009) and in various commercial thermofluid libraries. For example, Modelica.Fluid indicates that it is compatible with medium models that have T, (p,T), (p,h), (T,X), (p,T,X) or (p,h,X) as independent variables while other state variables are possible but would require initialization changes [T is temperature, p is pressure, h is specific enthalpy, and X is a mass fraction vector]. For two-phase systems, the standard state choice is (p,h) or (p,h,X) since temperature cannot always be used to uniquely identify the fluid state in the two-phase region.

Mass conservation as denoted by the system charge in vapor compression cycles is a challenge in modeling two-phase systems due to the sharp nonlinearities in thermodynamic properties, especially at the saturated liquid phase boundary. While charge conservation can prove to be numerically challenging, it is an important part of the modeling since variations in charge can produce unphysical responses from the system model, both in steady state and dynamically. Variations in pressures, flowrates, temperatures, heat flow, and key control values such as superheat and subcool due to modeling-induced charge variations can cause not only unphysical results but also computational degradation due to the additional unphysical dynamics. Since a key use case for dynamic models of vapor compression cycles is to develop the system and optimize system charge for operation, these sorts of charge issues are especially problematic from an engineering standpoint.

Excellent work in this area has shown that the use of redundant thermofluid states can greatly improve mass conservation in vapor compression cycle modeling (Laughman 2015). This paper is a follow-up to that work and describes a method to resolve a potential inconsistency that can result when employing the redundant dynamic state approach. Following a brief introduction of the redundant state approach, the results of the potential inconsistency are shown via a series of test models ranging from simple component tests to complex system models using Air Conditioning Library. Based on observations of the simulation results from these test models, a method for ensuring consistency is proposed and implemented. The method is then demonstrated on the test suite and evaluated for effectiveness, robustness, and computational efficiency.

2 Redundant Thermofluid States

This section provides a high-level overview of the motivation for the use of redundant thermofluid states and also describes the potential inconsistency. For a full treatment of the problem, including the derivation of the systems of equations, the interested reader is referred to previous work (Laughman 2015).

Using the standard assumptions for two phase thermofluid systems (one dimensional flow, thermodynamic equilibrium, homogeneous two-phase flow), conservation of mass and energy for a fixed control volume result in ODEs in terms of the following derivatives:

$$\frac{dM}{dt} = V \frac{d\rho}{dt} \quad (1)$$

$$\frac{dU}{dt} = V \frac{d(\rho u)}{dt} \quad (2)$$

where V is the volume, M is the volume mass, ρ is the density, U is the total internal energy, and u is the specific internal energy. These thermodynamic derivatives are then related to the mass, enthalpy, and heat flow into and out of the thermodynamic volume to provide the complete set of ODEs that must be integrated. Using different thermodynamic relationships (Thorade 2013), these derivatives can be written explicitly in different state variables. As discussed previously, the standard choice for two phase thermofluid systems is (p, h) . With Equations (1) and (2) rewritten as ODEs in terms of p and h , p and h are integrated via the solver and thus are error controlled by the solver tolerance. Thermodynamic relationships can then be used to calculate other variables such as density $\rho(p, h)$.

In the standard approach with p and h as state variables, there is no explicit error control on the density $\rho(p, h)$. Figure 1 shows the p - h diagram for the common refrigerant R1234yf with isotherms for temperature and lines of constant specific volume $1/\rho$. Note that the sharp

change in density at the liquid phase boundary. The density isochor increases in slope at the liquid phase boundary, becoming nearly vertical as the pressures are reduced below approximately 6 bar. This nonlinear behavior is especially challenging for numerical integrators and necessitates high precision on the integration of p and h to ensure that the calculated density and thus the mass are accurate. The reasons for this behavior are both that in the low-pressure liquid region, density has a weak dependency on pressure, and therefore there is large sensitivity in the solving of pressure time-derivative from the density and internal energy derivatives, and secondly the density partial derivatives with respect to p and h involved in that solution are discontinuous at the saturated liquid line. In general, precision can be increased by tightening integrator tolerances. In heat pump operation at cold ambients, it is not uncommon for the vapor cycle to operate with the high-pressure side in the range of 4-8 bar. Depending on the level of subcooling in the system, the condenser outlet condition and the high-pressure part of the system up to the expansion valve could have state points that reside just at the liquid phase boundary. These volumes are critical for charge conservation, as the accumulated error in the integration of p and h can lead to large errors in density and thus changes in system charge.




Figure 1. p - h diagram for R1234yf

As proposed in previous work (Laughman 2015), adding density as a redundant state can greatly reduce or even eliminate charge conservation issues. With density as a state, density is also error-controlled by the integrator. The authors propose the following additional equation:

$$\frac{dh}{dt} = \left. \frac{\partial h}{\partial p} \right|_{\rho} \frac{dp}{dt} + \left. \frac{\partial h}{\partial \rho} \right|_{p} \frac{d\rho}{dt} \quad (3)$$

Equation (3) can be combined with the formulated equations from Equations (1) and (2) to provide the set of ODEs that can be integrated for the state variables p , h , and ρ . This approach has the benefit of keeping the standard state variables p and h and thus allowing medium

models explicit in p and h (such as the Spline-Based Table Look-Up (SBTL) medium implementations in Air Conditioning Library which have demonstrated significant improvements in computational speed, especially when combined with analytic Jacobians (Li 2018; Li 2020)) to continue to be used even with the redundant density state.

In a subsequent publication (Laughman 2017), the authors acknowledge the potential inconsistency in the redundant state approach:

While this approach in theory relaxes the constraint forcing the state variable being integrated to be equal to the computation of the specific enthalpy as a function of the other state variables, for example, $h(P, \rho)$, simulations presented later in the paper provide evidence that these deviations are small in practice.

Though clearly true in the simulations presented in the cited publication (Laughman 2017), the authors of this work have observed these inconsistencies and their manifestation in system charge issues as described in the following section.

3 Simulations with Inconsistency

As outlined in Section 2, there is a potential inconsistency that can result from the redundant state approach. This section describes a series of models that demonstrate the impact of this inconsistency between the dynamic states. These models are built using Air Conditioning Library (Modelon 2023) and simulated with Dymola (Dassault Systemes 2023).

3.1 Single Volume Test

Since it can be difficult to put models in the state to demonstrate the inconsistency, it is valuable to have simple models that can be easily manipulated, either via initialization and/or via dynamics, and whose correct results are clearly understood. Figure 2 shows a model of a fixed volume with a trapezoidal heat input. The model is initialized such that the volume state is right at the liquid phase boundary (denoted by the blue circle on the p - h diagram) as this state is extremely sensitive due to the sharp changes in density at the liquid phase boundary. The heat input is manipulated such that the volume state repeatedly just enters the two-phase region and then leaves again (into the liquid region) as heat is added and the pressure increases. Since there is no mass flow into or out of the volume, the correct physical result is that the mass is constant, and the fluid state should move along a constant density line. Since we repeatedly add and remove the same amount of heat, all thermodynamic states should also return to their initial values between each full heating and cooling cycle.




Figure 2. Fixed volume test with heat input

In Air Conditioning Library, users can select different state choices at the top level of the model via the systemACL record. The option for the redundant (p, ρ, h) states is enabled for these simulations. Simulations are run with the SBTL 1234yf medium model in Air Conditioning Library with a relative tolerance of $1e-7$. Figure 3 shows simulation results from the fixed volume test. The top plot shows the refrigerant mass in the control volume which is constant as expected when the density is an integrated state. However, the second plot shows that the density state in the volume is not equal to the density $\rho(p, h)$. Thus, there is an inconsistency between the integrated states ρ, p , and h . The third and fourth plots show the pressure and specific enthalpy, indicating that they do not return to their initial values after many heat injections and extraction cycles as they should.

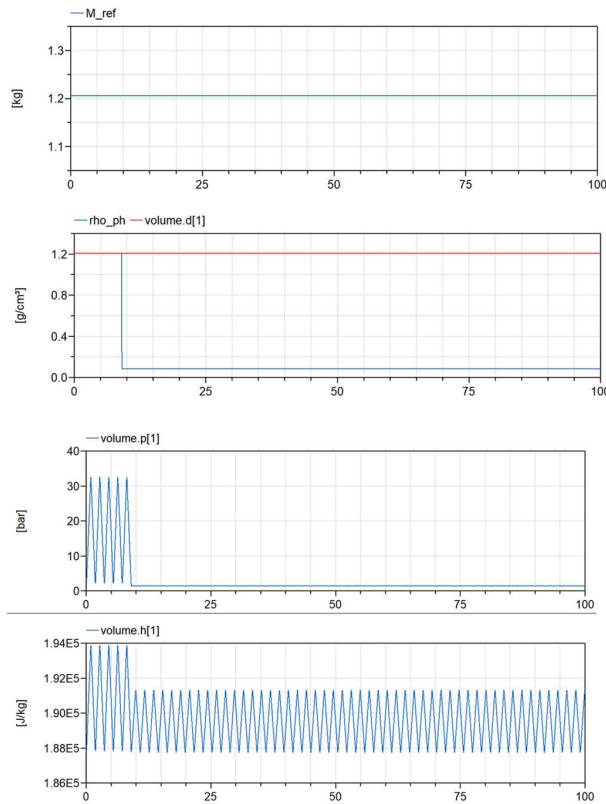


Figure 3. Simulation results from fixed volume test. System mass (top), inconsistent values for density from integrated density state and $\rho(p, h)$ (second figure), volume pressure (third figure), and volume enthalpy (bottom).

3.2 Two Volumes and Flow Model Test

One issue that was reported a few times by users of Air Conditioning Library when using redundant states (p, d, h) is the appearance of negative density values. Negative density is clearly an inaccurate result and often followed by simulation crashes. A hypothesis is that these negative values are caused by inconsistent redundant states along these lines:

1. Consider a simple system of two control volumes connected by an orifice model.
2. An inconsistency in pressure in one of the volumes will induce a mass flow rate from one volume to the other.
3. In a standard (p, h) state model a numerical inaccuracy in pressure would cause a change in density too, as density is just an algebraic function of p & h , and the mass flow would proceed until the pressures of the two volumes are equal again.
4. In a redundant state model though, it can be imagined that a low-density volume has a sudden inaccurate increase in pressure, driving a mass flow out of the volume. In the redundant state model, that inaccuracy does not cause the density

to increase, but the mass flow out of the volume will cause the density state to decrease from the accurate value. Therefore, it is postulated that inconsistent redundant states can indirectly cause (locally) incorrect densities via the flow models (momentum balances) of the system although density is under solver error control, and in some situations reach negative values if sufficiently high flow rates out from an already low-density volume is induced by the inconsistency.

To verify the above idea and potential solutions, the authors attempted, and successfully reproduced negative density in a relatively simple model, shown in Figure 4. This model includes the same fixed volume with repeated forced heat injection and rejection and is further connected via an orifice model to another adiabatic, fixed volume, and that is in turn also connected via another orifice model to a fixed pressure reservoir. The two connected volumes allow representation of the effect of mass flow rate induced by inaccurate pressure appearing in the system, and the fixed pressure reservoir prevents the pressure level from going far too high or low. Additionally, the heat flow rate to the first volume is controlled, such that in the event of numerical inaccuracy causing the volume to deviate from the vicinity of the sensitive saturated liquid line, heating will be adjusted to bring the volume state back there again. The idea is to trigger as much state inconsistency as possible and let it accumulate over many heat cycles.




Figure 4. Negative density reproducer model

Simulating this model indeed resulted in build-up of state inconsistency, eventually reaching negative density in the adiabatic volume, followed by a simulation crash.

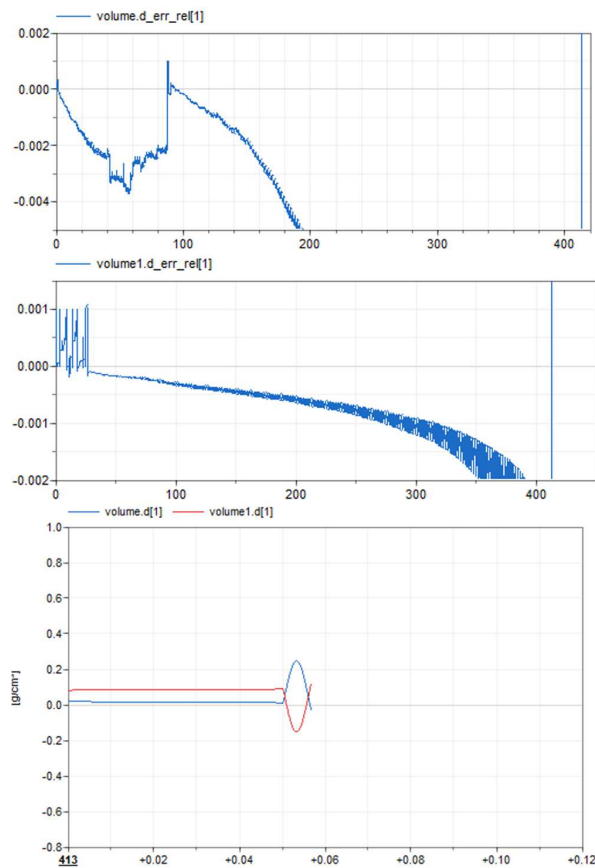


Figure 5. Negative density reproducer results. Relative error in density state and density from p & h states in non-adiabatic volume (top), adiabatic volume (middle). Densities in both volumes at the end of simulation (bottom).

Figure 5 show the gradual build-up of state inconsistency, illustrated by the relative difference between the density state variable and density computed as a medium property function evaluation from pressure and enthalpy dynamic states in each volume. The total system mass cannot be used to evaluate consistency here because the model is not a closed system. Toward the end of the simulation, the density dynamic state in the adiabatic volume also shows a negative value for some time followed by sign changes in both density states and then simulation crash.

4 Method for Consistent Redundant Dynamic States

4.1 Method Overview

Given the possibility for inconsistency between the redundant states and the observation of this inconsistency in test models, a method for ensuring consistency between the redundant states is proposed. This method attempts to solve the redundant state inconsistency problem via computation of the density from the (p, h) dynamic states and comparison with the density dynamic state. This comparison is done for every control volume in the system

and continuously throughout the transient simulation. If the deviation exceeds a certain allowed relative error (which can be set via a model parameter), the pressure state is corrected such that the three states are consistent again. This correction is done using the Modelica `reinit` operator and thus happens momentarily in an event which is triggered when the inconsistency is large enough. For simplicity, the new pressure value is taken as a medium property function of pressure from density and temperature, where temperature is as function of pressure and enthalpy. This approach is a simple alternative because the medium package already has equation of state implemented as a function with these inputs.

4.2 Variants Tested

There are multiple possible variants to formulate the inconsistency criterion. The criterion can be specified as tolerance on any of the three state variables and compared to the same property computed from the other two. The authors currently have only tested a consistency criterion for density.

There are also different possible choices of what or which state(s) to re-initialize and how to compute the state to re-initialize to. It was tested to reinitialize enthalpy instead of pressure, but this approach was not at all equally successful. Pressure seems to be the state with most inaccuracy in the low-pressure liquid region and thus could explain why it works better to initialize it.

Lastly, a good value must be found for the maximum allowed relative error of checked state (density in this case). Too tight tolerance may give very frequent re-initializations, reducing simulation performance due to large number of events, and too loose tolerance may fail to keep the states sufficiently consistent to mitigate the original issues.

4.3 Further Improvement

In the initial testing of the `reinit` approach, the authors observed promising results for maximum relative error in the range of $1e-3 - 1e-4$, but for tighter tolerances, observed unexpected results trending towards those of the original redundant states model without consistency check and correction. After some investigation the authors found that this behavior was a result of the simple calculation of the corrected pressure as a function of ρ and T . Since the consistency check uses the medium density from pressure and enthalpy function and `reinit` used the pressure from density and temperature function, a small inconsistency between those two functions could cause the re-initialized pressure not to fulfill the consistency criterion. In this case the simulation proceeds with no further correction, unless by chance the states become consistent again. Accurate equation of state for two-phase fluid properties are only explicit for one causality, meaning that one of the two functions used in the detection and correction of inconsistent states will include

iteration. Therefore, they will only be consistent to a certain numerical tolerance that is part of the medium property model (and not related to ODE solver tolerance or the newly introduced maximum allowed density error parameter), and the reinit method needed to be improved to guarantee that after the reinit the three dynamic states are consistent according to the criterion introduced.

To achieve this response, the authors compute the potential new pressure as before but then the convergence criterion is evaluated. If the convergence criterion is fulfilled, the corrected pressure is taken as the new value for pressure state as before. If the criterion is not fulfilled, Brent iteration of pressure is employed until a new pressure is found such that density computed from p and h is within allowed tolerance. With this improvement to the method, the authors observed the expected trend of ever better consistency for smaller maximum allowed inconsistency, and computed absolute relative density error was never greater than the given allowed tolerance.

5 Simulations with Consistent Redundant States

This section provides results from the test suite from Section 3 using the method described in Section 4 to ensure consistency between the redundant dynamic states.

5.1 Single Volume Test

The single volume test shown in Figure 2 was evaluated to understand the impact of the density relative error tolerance. Pressure correction at density relative error of $1e-3$, $1e-4$ and $1e-5$ was compared to the original results with uncorrected redundant states. Figure 6 shows huge error without any correction while the results are clearly staying within allowed error at different settings. Smaller allowed error can be seen increasing the frequency of corrections. Figure 7 shows density from (p, h) states. For the fixed volume in this test, a constant density is the correct result here. All corrected variants are much better than the uncorrected. For this particular model, CPU time does not vary much depending on tolerance here, and uncorrected is faster, likely because in such a simple model, inconsistency doesn't lead to expensive artificial transients.




Figure 6. Density relative error for uncorrected redundant states (top) and with maximum allowed error $1e-3$ (2nd plot), $1e-4$ (3rd plot) and densities at $1e-5$ (4th plot)




Figure 7. Density for uncorrected redundant states (top) and with maximum allowed error 1e-3 (2nd plot blue), 1e-4 (2nd plot red) and 1e-5 (2nd plot green). CPU times 3rd plot, event count 4th plot.

5.2 Two Volumes and Flow Model Test

The test shown in Figure 4 for negative density was evaluated to assess the original method and the improved method discussed in Section 4.3. Figure 8 shows the original method for state correction without Brent fallback. The relative error goes outside the maximum allowed error of 1e-3 at green circles.

Figure 9 shows results from the same model with the improved reinit and Brent iteration fallback. With this improved approach, the density error always stays within the tolerance. Figure 10 shows the volume densities. With the improved correction, volume densities are never close to zero. As can be observed, the corrections get very frequent toward the end. This result is likely due to unreasonable experiment setup with a high amount of heat injected.




Figure 8. State correction without Brent fallback. Density relative error in first heated volume (top) and second volume (bottom).




Figure 9. State correction with Brent fallback, Density relative error in first heated volume (top) and second volume (bottom).

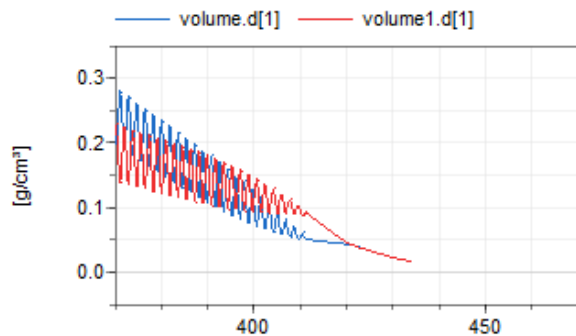


Figure 10. Volume densities

5.3 Complex System Models

This section illustrates results from complex vapor cycle system models. Figure 11 is an example system model from Air Conditioning Library using (p, ρ, h) states (AirConditioning.Examples.TXVCycleOnOff_pdh). Figure 12 compares results from the uncorrected redundant states with the method for ensuring state consistency. Note that this model does not have a large issue with mass conservation. However, employing the method for state consistency even with a 1e-3 tolerance for density inconsistency improves mass conservation significantly. The CPU time cost is low. The computed total system mass shown here is based on density from p & h function calculations. Observing the charge as per density states shows nearly perfect preservation.

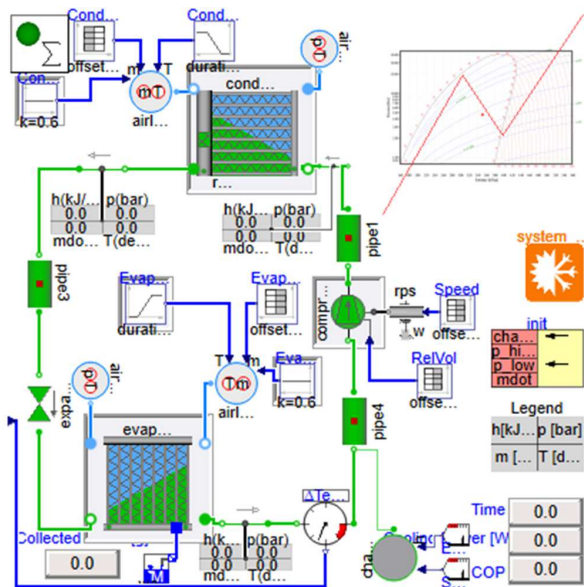


Figure 11. Air Conditioning Library example using (p, ρ, h) states




Figure 12. Air Conditioning Library example model with uncorrected redundant states (blue) and with maximum allowed error 1e-3 (red). System mass 1st plot, CPU times 2nd plot, event count 3rd plot.

The method for state consistency was also evaluated on a complex customer model (note that this model cannot be shown graphically as it is customer proprietary). This model uses R1234yf SBTL and has 505 (p, h) states. It was run using CVODE at a tolerance of 1e-7. The original results with the uncorrected redundant states are shown in Figure 13. The results show a significant mass defect and are thus problematic.

Figure 14 show results from the same model now set to use (p, ρ, h) states (591 states in total) and employing the method for correction of the redundant states with a tolerance for density of 1e-3. The figure shows results from the initial method for correction (blue) and improved with Brent (red). Both methods preserve mass much better than p, h (note reported total mass is computed from the p & h states). Note that the improved correction shows no CPU time penalty when compared with uncorrected (p, h) states.

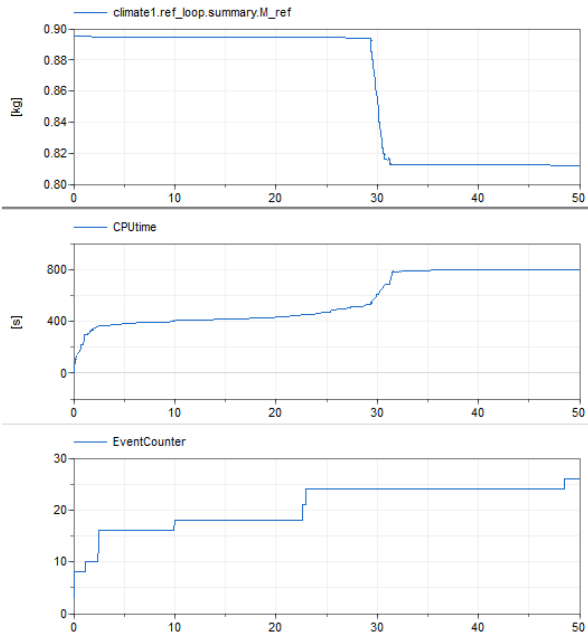


Figure 13. Customer vapor cycle system model with uncorrected redundant states. System mass 1st plot, CPU time 2nd plot, event count 3rd plot.

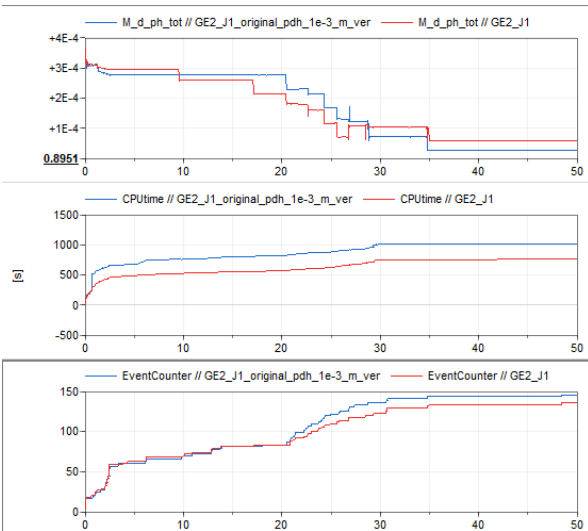


Figure 14. Customer vapor cycle system model with maximum allowed error 1e-3. System mass 1st plot, CPU time 2nd plot, event count 3rd plot.

Figure 15 shows results with a tightened density error tolerance of 1e-4. The initial reinit approach (blue) is compared with the improved approach with Brent (red). Here it can be observed that mass defects from correction are not successfully corrected to a consistent state with the initial method. The improved method is faster and more accurate. The CPU time is still roughly on par with uncorrected (p, h) states for the improved correction.




Figure 15. Customer vapor cycle system model with maximum allowed error 1e-4. System mass 1st plot, CPU time 2nd plot, event count 3rd plot.

Figure 16 shows results with a further tightened density error tolerance of 1e-5. The initial reinit approach (blue) is compared with the improved approach with Brent (red). The initial method at this setting demonstrates worse preservation than only (p, h) as states. This response is probably because the correction is mostly dysfunctional due to running in an inconsistent state, and somehow the added density states introduce some additional inaccuracy when exerted to mass flows caused by pressure state inaccuracy. The improved correction performs well. Mass is very well preserved, and relative density error never exceeds 1e-5. There is a CPU time penalty versus only (p, h) as states of approximately 25%.




Figure 16. Customer vapor cycle system model with maximum allowed error 1e-5. System mass 1st plot, CPU time 2nd plot, event count 3rd plot.

6 Conclusions

This paper describes a method to resolve a potential inconsistency when employing redundant dynamic thermofluid states for modeling of vapor compression systems. While the impact of redundant thermofluid states is generally positive in terms of mass conservation and even CPU time, it is possible to observe inconsistencies in the separately integrated states as shown in this work. These inconsistencies could manifest not only as mass loss but also as incorrect pressure and enthalpy in the system as shown in the simulation results from both simple and complex models. These inconsistencies can certainly affect the engineering utility of these models if not resolved. A method is proposed and implemented to resolve any inconsistencies as detected during the transient simulation. This method has been tested on models that demonstrated inconsistencies and shows promising results in terms of consistent results and computational impact. These changes have been implemented in Air Conditioning Library 1.26 to ensure that users who employ the redundant state option will not only preserve charge but also will see consistency between the density, pressure, and enthalpy states.

Environmental earth sciences, 70(8), pp.3497 -3503. DOI: 10.1007/s12665-013-2394-z

Tummescheit, H. (2002). *Design and Implementation of Object-Oriented Model Libraries using Modelica*. PhD thesis, Lund Institute of Technology.

References

- Dassault Systemes (2023), Dymola, <https://www.3ds.com/products-services/catia/products/dymola>.
- Franke, R., F. Casella, M. Sielemann, K. Proelss, M. Otter, and M. Wetter. (2009). Standardization of thermo-fluid modeling in Modelica.Fluid. In *Proceedings of the 7th Modelica Conference*, Como.
- Laughman, C. and H. Qiao. (2015). Mass Conserving Models of Vapor Compression Cycles. In *Proceedings of the 11th International Modelica Conference*, Versailles, France. doi: 10.3384/ecp15118759.
- Laughman, C. and H. Qiao. (2017). On the influence of state selection on mass conservation in dynamic vapour compression cycle models. *Mathematical and Computer Modeling of Dynamical Systems*, Volume 23, Issue 3. <https://doi.org/10.1080/13873954.2017.1298625>
- Li, L., J. Gohl, J. Batteh, C. Greiner, and K. Wang. (2018). Fast Calculation of Refrigerant Properties in Vapor Compression Cycles Using Spline-Based Table Look-Up Method (SBTL). *Proceedings of the 1st American Modelica Conference*, p. 77 -84. DOI: <http://dx.doi.org/10.3384/ecp1815477>
- Li, L., J. Gohl, J. Batteh, C. Greiner, and K. Wang. (2020). Fast Simulations of Air Conditioning Systems Using Spline-Based Table Look-Up Method (SBTL) with Analytic Jacobians. *Proceedings of the American Modelica Conference 2020*, p. 64 -72. DOI: <https://doi.org/10.3384/ecp2016964>
- Modelon (2023). *Air Conditioning Library*. <https://modelon.com/library/air-conditioning-library/>.
- Thorade, M. and A. Saadat. (2013). Partial derivatives of thermodynamic state properties for dynamic simulation.