


Sage Model Notes

GM-PTR-TwoStage-DPregulated.scfn

D. Gedeon


2 September 2021 (revised 27 October 2024)

A two-stage GM pulse-tube cryocooler with a pressure regulated compressor component providing the cold-head pressure boundary condition. The model schematic looks like this:



This is an extension of the single-stage GM pulse-tube cryocooler model, documented in **GM-PTR-SingleStage-DPregulated.pdf**. It might be a good idea to become familiar with that model before tackling this one.


The two-stage model is organized into submodels with the root-level model looking like this:



The root model includes the two *bypass valves* (aka double-inlet valves) which connect the *compressor* output to the secondary rejector exits of the two stages, as shown in the above schematic. The DC flow setpoint of bypass valve 2 is not zero in order to improve the 4K cooling power in the second stage. More on that later. Implementing an equivalent DC flow control in actual hardware may require an additional valve between the warm end of the second stage pulse tube or reservoir directly to the compressor suction pressure.

Compressor

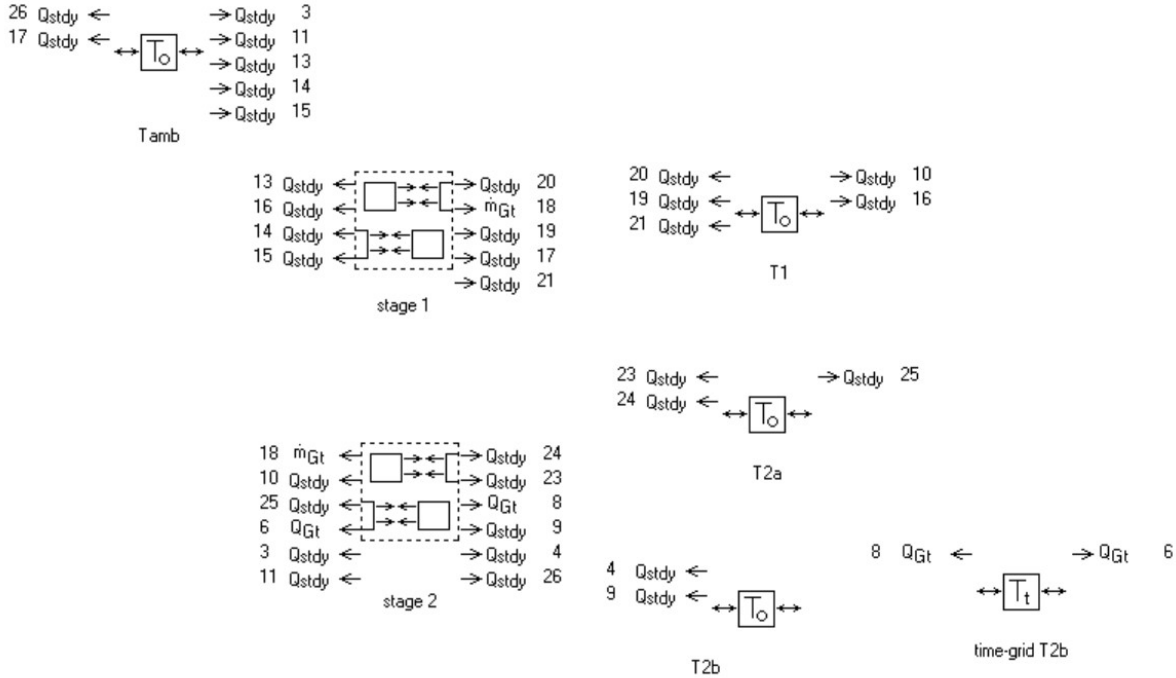
The model components that simulate the compressor are identical to those in the single-stage version of this model:



The *pressure-source* and *low P buffer* provide the low-side suction pressure. The *DP regulated compressor* boosts the pressure according to a pre-determined waveform. The *aftercooler* removes the adiabatic heat of compression. In the actual hardware any aftercooling is provided upstream of the rotary valve.

Cold Head

The cold-head submodel is organized into *stage 1* and *stage 2* submodels with a number of temperature sources for anchoring the components within:




Temperature source T_1 anchors the cold temperature of *stage 1*, currently fixed at 50 K. Temperature source T_{2a} anchors the temperature of an intermediate heat exchanger in *stage 2*, between two regenerators. This intermediate heat exchanger is an optional part of the model used as a way to improve model convergence and control regenerator temperature boundary conditions. In actual hardware such a heat exchanger could be used to anchor some intermediate heat source (e.g. radiation shield or electrical leads), or not implemented, in which case it could be removed from the model. Temperature sources T_{2b} and *time-grid* T_{2b} anchor the cold heat exchanger of *stage 2*, currently fixed at 4 K. The time-grid heat source anchors time-varying temperatures at the cold ends of the regenerator and pulse tube.

Net cooling powers are available in the *cold head* user-defined variables Qlift1, Qlift2a and Qlift2b. These include the heat flows absorbed by the helium in the acceptor heat exchangers, less the conduction losses down the regenerator and pulse-tube canisters.

Stage 2 Submodel

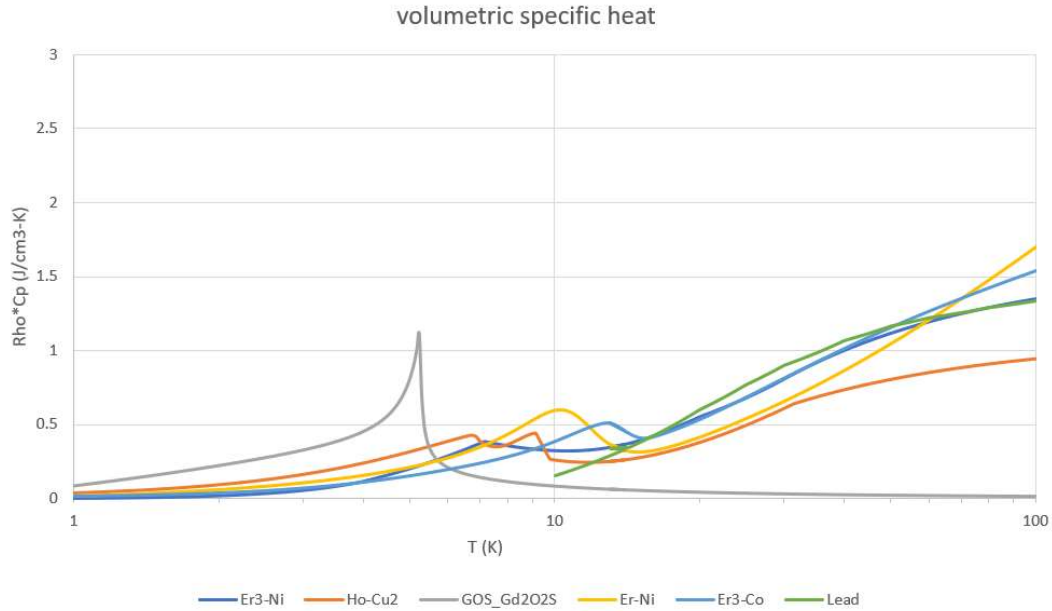
The components in the stage 1 submodel have already been documented in **GM-PTR-SingleStage-DPregulated.pdf**. The components in the *stage 2* submodel are slightly different as shown here:



Discharge flow from the first stage acceptor heat exchanger enters at the negative (left) boundary of the first regenerator (*regen Lead*) and proceeds through an intermediate heat exchanger (*accept 2a*), the second regenerator (*regen HoCu2*), the cold heat exchanger (*accept 2b*), pulse tube (*ptube 2*), secondary warm heat exchanger (*sec rej 2*), tuning valve (*orifice 2*) and into the reservoir (*reservoir 2*). During the suction part of the cycle the flows are reversed. The bypass inlet (*bypass inlet 2*) serves to redirect the flow from the secondary warm heat exchanger to or from the second-stage bypass valve (*bypass valve 2*) at the root level. The components are arranged in rows, with colder components at the bottom. The "2" suffix in all the component names designates that this is the second stage.

Inter-Regenerator Temperature

Temperature T2a anchors the heat exchanger between the lead and HoCu2 regenerators. The model fixes this temperature at 12 K, where the volumetric heat capacities of lead and HoCu2 are roughly equal (see plot below). T2a below 12 K will progressively destabilize the lead regenerator due to diminishing lead heat capacity. T2a above 12 K reduces performance of the HoCu2 regenerator relative to a lead regenerator.




4 K Issues

The temperature at the cold-end of the second stage is 4 K in this model, which raises a number of issues with helium properties and heat capacities of solid materials.

The non-ideal properties of helium become an issue. The operating pressure range is chosen to avoid the helium (helium-4) entering a two-phase state, which is a problem in a periodic flow machine because the relative vapor fraction (quality) is indeterminate. In a Sage model, two-phase helium would likely lead to convergence issues.

To avoid the two-phase region there are two options, operate below the equilibrium vapor pressure for the minimum temperature or above the critical pressure (see P-V plot below). Since the equilibrium vapor pressure at 4 K is rather low (low power density) this model operates above the critical pressure. The minimum time-average pressure (*pressure source* anchoring the compressor *low P buffer*) is set at 0.5 MPa (72 psia). This is well above the critical pressure of 0.23 MPa. Some pressure margin is a good idea to avoid convergence issues when the solver is adjusting temperatures and pressures to home in on the solution.




P-V plot from NIST Refprop software

Another problem is that at high pressures, helium loses its compressibility near 4 K, meaning that the volume change for a given temperature change is smaller than for an ideal gas. The non-dimensional form of compressibility is formulated in thermodynamic textbooks as

$$T\beta = \frac{T}{v} \left(\frac{\partial v}{\partial T} \right)_P$$

where v is specific volume ($1/\rho$).

The following plot shows $T\beta$ for helium 4, as a function of temperature, at various pressures.

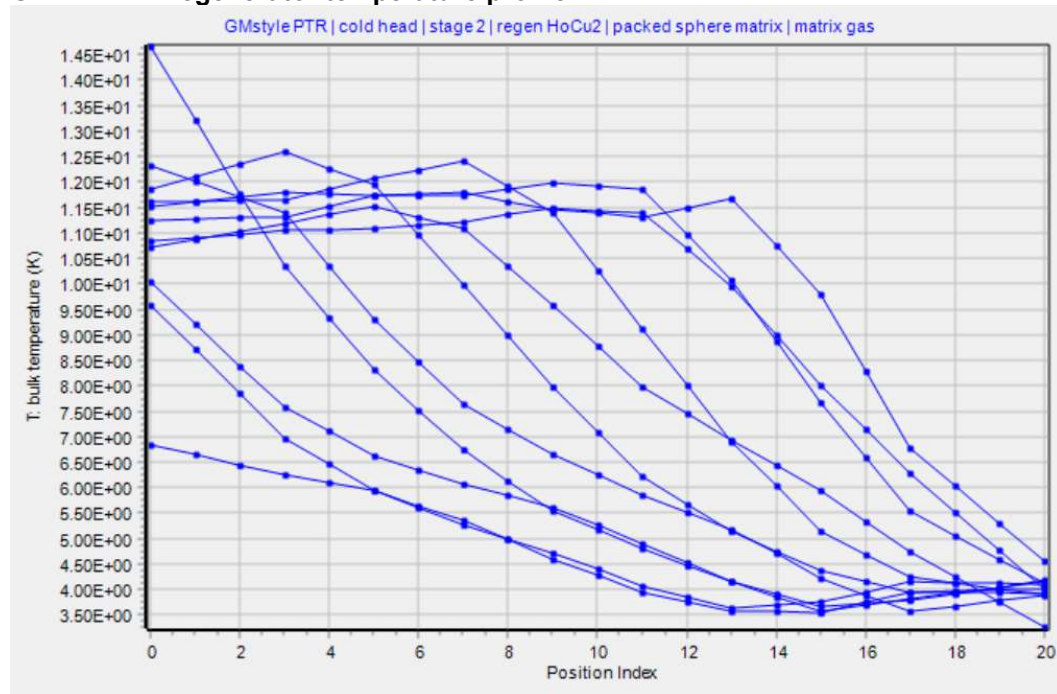


T*Beta plot from NIST Refprop software

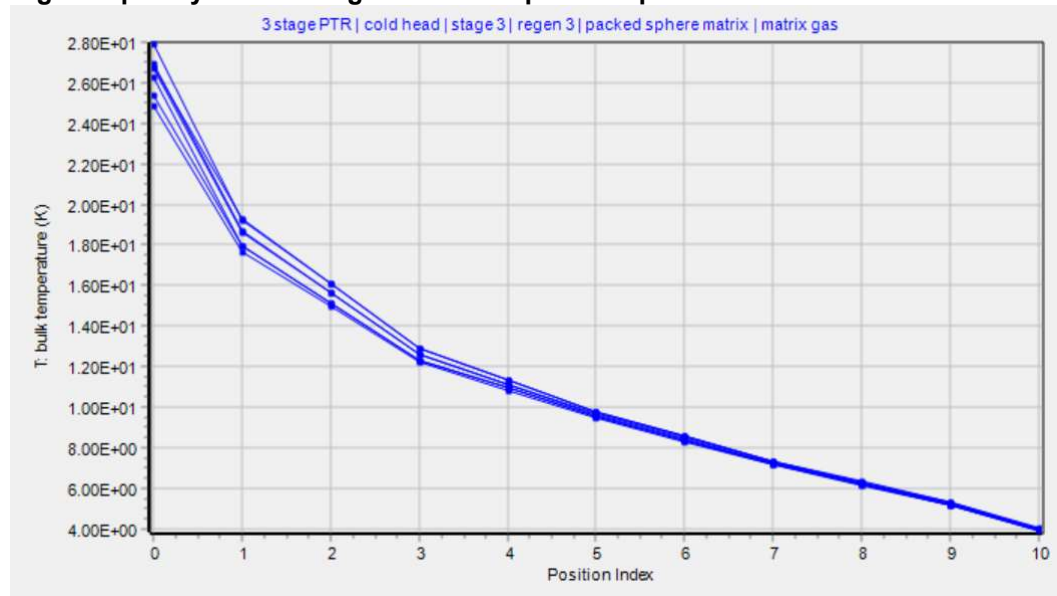
At low and intermediate pressures the compressibility can be higher than for an ideal gas. The documentation for the high-frequency three-stage pulse-tube cryocooler (*HiFreqPTR-ThreeStage.scfn*) discusses the implications on helium enthalpy of $T\beta$ being higher or lower than 1. The pressure-dependent component of enthalpy flow tends to produce concave temperature profiles in regenerators.

Low solid heat capacities of regenerator matrix materials are another issue — the most important issue from a model convergence perspective. This was not nearly so critical in the 4 K high-frequency pulse-tube model *HiFreqPTR-ThreeStage.scfn* because pressure swings and resulting adiabatic temperature swings in the helium were much lower. In this model the longitudinal temperature profiles in the second-stage regenerators show extremely large time variations, superimposed on effects of pressure-dependent enthalpy flows. The plots below show regenerator helium temperature profiles from the 4 K stage of present model, compared to the 4 K stage of the high-frequency pulse-tube model. In the present model the low heat capacity of the regenerator matrix cannot limit the gas temperature swings as it can in a higher temperature (first stage) regenerator.

GM PTR 4 K regenerator temperature profile

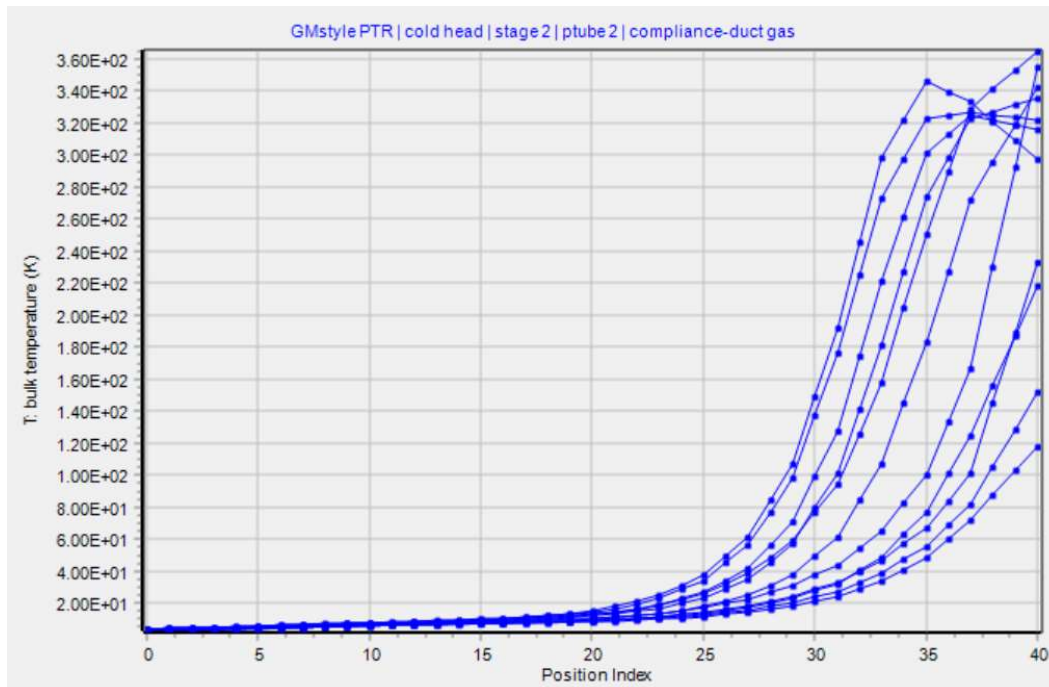


High-frequency PTR 4 K regenerator temperature profile



Low solid heat capacity is not an issue per-se in the pulse-tubes of a GM cooler because the heat transfer between the helium and solid wall is relatively low in the first place, so the helium responds adiabatically, for the most part. The extreme pressure variations are a problem though because they adiabatic temperature variations tend to be relatively large, as illustrated below.

GM PTR pulse-tube temperature profile



The large temperature swings in the second-stage regenerator and pulse tube can lead to convergence issues.

DC flow effects The 4 K cooling power is increased significantly if the time-average mass flow rate through the stage 2 components is slightly positive (from the warm to cold end of the regenerator). In the Sage model such a DC flow increases the net enthalpy flow down the regenerator slightly (a loss) but increases the enthalpy flow up the pulse tube significantly (a gain). The result is a net increase in cooling power of the 4 K heat exchanger (*accept 2b*). The reason for the increased pulse-tube enthalpy flow seems to be that the thermal component of net enthalpy flow in the negative direction (thermal loss) is diminished because the temperature gradient is flattened toward the cold end as a result of the DC flow, while the net PV power flow in the positive direction is not affected. For the regenerator, the reason that the enthalpy flow does not increase much may be that the temperature gradient does not increase significantly as a result of the DC flow, possibly because of the effects of pressure-dependent enthalpy active in the interior part of the regenerator.

Convergence Tricks


Compared to higher temperature cryocooler models, numerical convergence for this model is much more of a problem. Solving takes longer and sometimes terminates in an out-of-range helium temperature exception. Below are some suggestions for dealing with such problems.

Monitor Solution Progress with Plot Windows

It can take a long time for the pressure wave imposed by the *DP regulated compressor* component to build up from initial conditions. Initialization sets the pressure to P_{norm} everywhere, where P_{norm} is a root-model input. To get an idea what is going on it is helpful to open plot windows for various key components and click the refresh button from time to time during the solution process to see how the solution is progressing. Key components might be the *DP regulated compressor* itself and the gas domains of the second-stage cold regenerator and pulse tube. As long as the plots show progress in converging to a solution there is hope. If the plots get stuck along with the convergence error tolerance, then it may be time to stop the solver, adjust the model and try again. Possibly reinitializing the solution first.

Monitor Step Reductions

Pay attention to the step-limit information in the status dialog that appears when solving or optimizing. If the step limit factor is less than 1 the information in the box below can help you figure out where in the model the solver is having difficulty. For example, the screen shot below suggests there is some difficulty at node (35, 6) (position, time indices) in the stage 2 regenerator:



Such problems will often clear up by themselves. But if the step limit factor is too small (less than $1.0E-2$) for too many iterations you may want to plot the solution (see below) and see if you can spot anything strange going on. In a 4K model you may see a temperature decreasing to zero or otherwise unhinged from reality. Adding thermal connectors at endpoints often helps.

Diagnostic Dialog

Showing the solver diagnostic dialog can sometimes spot problem areas (Options -> Sage -> Solver, check show diagnostic dialog). Clicking on high points in the plot of " $|F|$ values > tol" can give you an idea where the problem lies.

Increased Cold-End Temperature

Especially after model reinitialization it may be helpful to increase the cold-end temperature above 4 K temporarily, then lower it again after the model has successfully converged at a higher temperature.

Rough Tuning

There are four valves to adjust in this model, two bypass valves and two tuning valves. When they are properly adjusted the pressure variations will be roughly in phase with the mass flow rate variations in the regenerators. If *bypass valve 2* allows too much flow into the stage 2 pulse tube it is possible to shift the mass flow rate by 180 degrees which effectively kills the cooling power. You can visualize the phasing between pressure and mass flow rate with plot windows or by looking at the first harmonic phases of tabular outputs FP_{mean} and $FRhoUA_{mean}$ in the regenerator gas domain. For initial convergence it may help to temporarily reduce the DC flow input $DCRhoUA$ of bypass valve 2, since it can skew the temperature distributions one way or the other in the stage 2 regenerator and pulse tube. You may be able to use the optimizer to adjust the valves (optimize lengths) subject to the constraints that $FP_{mean}.Arg.1 = FRhoUA.Arg.1$ in the stage 1 regenerator and stage 2 cold regenerator. After that you can re-define the optimization in order to optimize cooling power.


UpwindFrac

There is an obscure input UpwindFrac, available in all gas domains, with a default value of 0.01. As explained in the Sage User's Guide the purpose of UpwindFrac is to smooth out a jagged longitudinal temperature profile at some cost in reduced accuracy. If this model fails to converge you can often see zig-zag temperature swings in the temperature profiles of the second-stage cold regenerator and pulse tube. To stabilize the temperatures in the present model the UpwindFrac values are set to 0.8 in the cold regenerator and 0.7 in the pulse tube. These values stabilize the solution but do not seem to affect the 4 K cooling power appreciably.


UpwindFrac affects only the interpolated values of density in the solution grid, and therefore also temperature, which is a function of density. Interpolated values are those at the boundaries of computational cells (every other node starting with node 0 at the negative endpoint boundary), where density is not actually solved. UpwindFrac determines the relative weighting of central and upwind interpolation. Central interpolation amounts to the average of cell-center values on either side of a given reference point. Upwind interpolation amounts to a linear extrapolation of cell-center values in the upwind direction from a given reference point. The upwind direction is tied to the local gas mass flow rate. The UpwindFrac value is limited to 1, which amounts to 100% upwind extrapolation.

Training Wheels

The wild temperature variations in the 2nd stage regenerator and pulse tube components can overwhelm Sage's solver algorithm. If all else fails you can temporarily reconnect their gas domains from rigorous surfaces to isothermal surfaces:



Or anchor the rigorous surface to an independent line heat source.

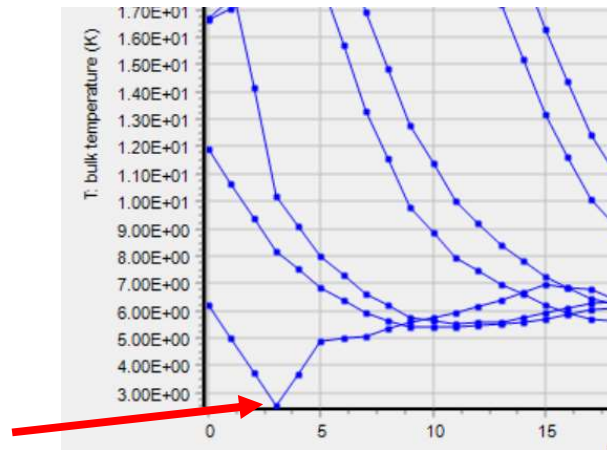


This will anchor the gas domains according to the temperature distributions specified by the Tinit input of the parent component, making convergence much more likely. After convergence do not forget to remove the training wheels to avoid the nonsensical physics

produced by the infinite heat capacity and heat flows produced by isothermal components.

Increased Thermal Conduction

When the regenerator matrix or pulse-tube wall heat capacity and thermal conductivity are too low, the gas temperature is weakly determined, and can drift erratically especially in non-ideal fluid states where pressure-dependent enthalpy flows produce local heating or cooling effects. For example, the spike at 3 K below:



Temporarily increasing the gas or wall thermal conductivity by increasing the `Kmult` input can iron out such spikes, allowing the solver to converge. Subsequent model adjustments may eliminate the reason for the temperature drift, allowing you to reset `Kmult` to the default value of 1.

NCell reduction

Another way to deal with local temperature drifting is by reducing the number of computational cells in the component. As `NCell` decreases there is an increasing amount of numerical thermal diffusion in the solution which has a similar effect to increasing the thermal conductivity. There is a trade-off here because decreasing `NCell` reduces solution accuracy. Once the model converges reliably some numerical experimentation with the value of `NCell` may be in order.

Optimization

This model implements a simple optimization to maximize the 4 K cooling power subject to 2 kW compressor power input. The 4 K cooling power (`Qlift2b`), along with the cooling powers at intermediate temperatures are defined in the *cold head* submodel as

<code>Qlift1</code>	net lift stage 1	1.000E+01
<code>Q1</code>		
<code>Qlift2a</code>	net lift stage 2a	2.694E-05
<code>Q2a</code>		
<code>Qlift2b</code>	net lift stage 2b	9.767E-01
<code>Q2bStdY + Q2bGt</code>		

The intermediate-temperature cooling powers are subject to the constraints

```
Qlift1 = 10
Qlift2a >= 0
```

The compressor power is defined in the *DP regulated compressor* component of the compressor submodel as

```
Wcompressor      adiabatic compressor power      -2.000E+03
FWc.Mean
```

FWc is a built-in Fourier series output variable.

The optimized variables are:

- *Stage 1* and *stage 2* lengths (*Lstage1*, *Lstage2*)
- Regenerator and pulse-tube diameters for both stages (*Dreg* and *Dptb*)
- Lengths of the two bypass valves and tuning orifi (*Length*), and
- DC mass-flow rate for the stage 2 bypass valve (*DCRhoUA*)