



Optimization strategies for single-stage, multi-stage and continuous ADRs



Peter Shirron

NASA/Goddard Space Flight Center, Greenbelt, MD 20771, United States

ARTICLE INFO

Article history:

Available online 18 April 2014

Keywords:

Adiabatic demagnetization refrigeration
Magnetic refrigeration
Magnetocaloric effect

ABSTRACT

Adiabatic Demagnetization Refrigerators (ADR) have many advantages that are prompting a resurgence in their use in spaceflight and laboratory applications. They are solid-state coolers capable of very high efficiency and very wide operating range. However, their low energy storage density translates to larger mass for a given cooling capacity than is possible with other refrigeration techniques. The interplay between refrigerant mass and other parameters such as magnetic field and heat transfer points in multi-stage ADRs gives rise to a wide parameter space for optimization. This paper first presents optimization strategies for single ADR stages, focusing primarily on obtaining the largest cooling capacity per stage mass, then discusses the optimization of multi-stage and continuous ADRs in the context of the coordinated heat transfer that must occur between stages. The goal for the latter is usually to obtain the largest cooling power per mass or volume, but there can also be many secondary objectives, such as limiting instantaneous heat rejection rates and producing intermediate temperatures for cooling of other instrument components.

Published by Elsevier Ltd.

1. Introduction

Compared to other low-temperature refrigeration technologies, ADRs tend to have the highest thermodynamic efficiency, in large part due to the reversible nature of the magnetocaloric effect [1] used in the cycle. However, their need for large magnetic fields and the relatively low density of paramagnetic refrigerants also means they tend to have the lowest cooling capacity per unit mass. Consequently, ADR design must focus both on minimizing heat loads to the cold stage, and ensuring that the stage achieves the maximum possible cooling capacity for its mass. Complicating the process is that fact that a wide variety of paramagnetic refrigerants can be used, and that the cooling capacity depends not only on the amount of refrigerant, but also on the magnetic field used to drive the cycle. The interdependence of these parameters means that many design solutions are possible that achieve the same cooling capacity, but with potentially large variations in total stage mass. In general, though, the choice of refrigerant is guided by relatively straightforward selection rules [2,3], and, as will be seen, it is possible to develop analytic expressions for the optimal magnetic field as a function of the ADR's operating temperature range, which can then be used to compute the refrigerant mass

needed to obtain the desired cooling capacity. When applied to single ADR stages, these yield designs for which the cooling capacity per unit mass is maximized.

However, this does not necessarily achieve an optimal design from a system perspective. For ADRs operating over a limited temperature range (ratio of high end and low end temperature less than about 20), a single stage is likely the lowest mass solution. When larger operating range or notably higher cooling capacities are needed, system mass likely can be lowered by using multi-stage arrangements [4] including ADRs that cool continuously [5]. The reduced operating range of each stage and the corresponding reduction in conducted heat loads, plus the ability to use refrigerants tailored to each stage's operating range, can result in significant mass savings and higher overall efficiency. The latter can also translate to lower peak heat rejection rates, which can be a design driver in systems that use cryocoolers as heat sinks.

Multi-staging, though, introduces additional operational parameters, including heat transfer points, hold temperatures, and recycling sequences. Collectively these determine the operating range and heat loads for each stage, and, therefore, stage mass. The number of free parameters can be daunting, but in practice it is possible to develop “rules of thumb” for the parameters that govern the interactions between stages, and separately to develop expressions for, for example, the optimal magnetic field of a stage once the operating range is established. System designs can then

E-mail address: peter.shirron@nasa.gov

be built up in a systematic way from a much smaller number of fixed parameters.

This paper presents those strategies and guidelines for optimizing ADR systems, starting first with individual ADR stages, and progressing to two-stage and then multi-stage (including continuous) ADRs. For single stages, where the operating range and required cooling capacity are fixed, optimization mainly involves determining the correct balance between magnetic field and refrigerant mass. For multi-stage systems, it also involves determining the operating range for each stage that minimizes the parasitic entropy generation within the ADR. As such, the optimization process not only identifies the lowest mass configuration, but also maximizes system efficiency, which has the further benefit of potentially relaxing requirements for the cryogenic system acting as the ADR's heat sink.

2. Optimization of a single ADR stage

An ADR stage consists of relatively few components: a superconducting magnet (producing a magnetic field, B , at maximum current, I_{max}), a salt pill containing the magnetic refrigerant, a heat switch to thermally link the salt pill to the heat sink, and a suspension system that mechanically supports the salt pill within the magnet bore. For sensitive instruments, there may also be a need for magnetic shielding, either passive (ferromagnetic or superconducting) or active (bucking coils). The stage is shown schematically in Fig. 1.

A single-stage ADR is typically designed to operate between a fixed heat sink temperature, T_{sink} , and a lower operating point (or operating range), T_{low} . Fig. 2 is a plot of the entropy versus temperature for a stage using chrome potassium alum (CPA) salt, operating at 50 mK using a 1.3 K heat sink, and a 2 T magnetic field. The figure assumes that the ADR stage has already been cycled to low temperature; if not, the cycle will start with an isothermal magnetization from the point on the zero-field entropy curve that intersects the line B–C.

At point A, the stage is cold but has no remaining cooling capacity, and must be recycled. The salt pill is thermally isolated, and is warmed (A–B) by adiabatic magnetization until its temperature reaches that of the heat sink. The heat switch is then closed, and magnetization continues isothermally (B–C). The entropy of the salt decreases as heat is transferred to the sink until full field is reached. The heat switch is opened, and the stage is adiabatically demagnetized (C–D) until the desired operating point is reached. The field is then slowly decreased to maintain temperature as heat is absorbed. When the field reaches zero, the cooling capacity is exhausted, and the stage is back at its starting point.

In Fig. 2, the entropy difference between points D and A (ΔS) is the entropy capacity of the stage. The curves show specific entropy

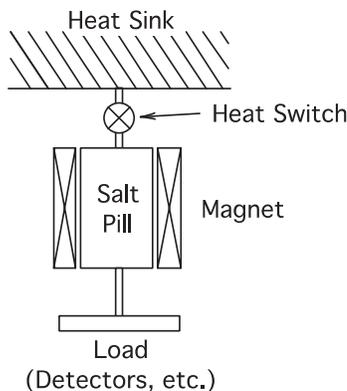


Fig. 1. Schematic of a single-stage ADR.

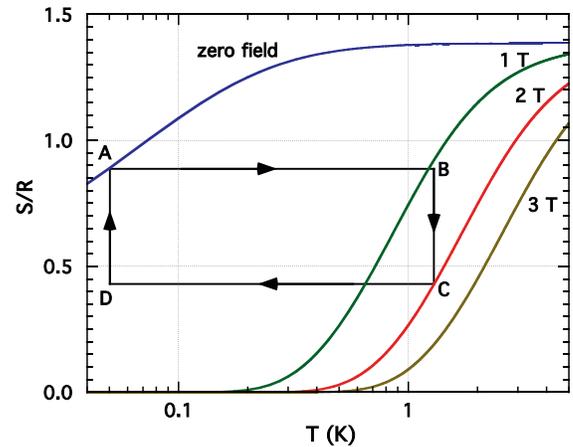


Fig. 2. Ideal cycle for a single ADR stage using CPA salt, operating at 50 mK using a 1.3 K heat sink, and a 2 T magnetic field.

(in units of the gas constant, R) which is multiplied by the refrigerant mass to get ΔS . The corresponding cooling capacity at low temperature is $T_{low}\Delta S$, and the heat rejected each cycle is $T_{sink}\Delta S$.

The target cooling capacity is based on the heat loads that must be absorbed at low temperature, and the duration (or hold time). The heat loads must include both internal (or parasitic) sources, such as the heat switch, salt pill suspension and sensor wiring, and external sources: detectors, wiring, amplifiers, etc. It is not uncommon for the parasitic loads to represent the majority of the heat budget, and it should be understood that they tend to scale with the size of the salt pill, and so the design process is usually iterative. The heat switch, in particular, has to be sized to reject heat to the sink within what is usually a prescribed period of time, and the open-state conduction will typically be a function of the size needed to achieve a suitable closed-state conductance.

2.1. Selection of refrigerant

Paramagnetic materials used as refrigerants are characterized by a spin quantum number, J , a magnetic ion density, and an ordering temperature. Magnetic ordering occurs below the temperature where the energy of nearest neighbor interactions is comparable to the thermal energy. The interaction competes with the effect of an applied magnetic field, reducing the latter's ability to affect the entropy and produce cooling. Effectively the ordering temperature is a lower limit to the useful range of a paramagnetic refrigerant.

The ordering temperature is positively correlated with J and the magnetic ion density. Consequently, the choice of refrigerant for an ADR stage is driven by the lowest temperature at which the stage needs to operate. Lower temperature requires the use of less dense, lower J -value materials, with correspondingly lower refrigeration capacity.

For reviews of low temperature refrigerants, their cooling capabilities and selection guidelines, see [2,3]. Here we note that the goal is typically to obtain the highest cooling capacity per volume in the salt pill, although some consideration must be given to the chemical properties and, hence, how the material must be packaged.

Low temperature materials (for use below, say, 100 mK) are almost exclusively hydrated salts [6]. In order to make thermal connection between the salt and other components, the salt is grown from solution inside a hermetic container, onto a thermal bus. The bus is constructed from thermally conductive [7–10] material, such as gold or copper wires. Some salts, like ferric ammonium alum (FAA), are highly corrosive to copper, whereas CPA and CMN are not. This leaves gold as the only practical choice for FAA. Depending on the extent of the bus needed, its cost or

labor-intensive fabrication time may be factors which bias the design away from a thermodynamically superior material to one that can be more easily (or inexpensively) utilized. In any event, the volume occupied by the salt pill container and thermal bus reduce the amount of salt that can be accommodated in the magnet bore, and this inefficiency must be weighed against the benefit of materials (if any exist) that do not require such structures.

Higher temperature materials, such as gadolinium gallium garnet (GGG), tend to be available as single-crystals, and stable in any environment. Most, in the temperature range of use, have high enough thermal conductance and low enough boundary resistance to not require an integral thermal bus. Whether one is used or not depends on the specific material's properties and the heat fluxes that must be supported. That is, a thermal bus may be advantageous if the thermal gradients within the salt pill reduce the realizable cooling capacity by more than the reduction in salt volume from its inclusion would.

2.2. Optimization of magnetic field and refrigerant mass

The cooling capacity of an ADR stage operating is determined by both the amount of salt used and the peak magnetic field used to drive the cycle. Since a reduction in one can be compensated for by an increase in the other, and since the magnet (and shield) will typically dominate the ADR's mass, there are potentially significant mass savings to be gained by correctly balancing the salt mass and magnetic field.

The process involves developing a relationship between salt mass and peak field that yields the desired cooling capacity, then identifying the combination having the lowest mass. As shown in Fig. 2, cooling capacity is proportional to the difference between the entropy of the salt at zero field and the operating temperature, $S(0, T_{low})$, and at peak field and the sink temperature, $S(B, T_{sink})$.

Entropies can be calculated using the function:

$$s(B, T)/R = x \coth(x) - (2J + 1)x \coth((2J + 1)x) + \ln\left(\frac{\sinh((2J + 1)x)}{\sinh(x)}\right) \quad (1)$$

$$x \equiv \mu_B g B_{eff} / 2k_B T = (0.336 \text{ K/T}) g B_{eff} / T$$

(where g is the Landé g -factor, μ_B is the Bohr magneton, and k_B is Boltzmann's constant), using the internal field approach to modeling the zero-field entropy [4]:

$$B_{eff} = \sqrt{B^2 + b^2} \quad (2)$$

where b is the background field generated by each ion's nearest neighbors.

Assuming the operating and sink temperatures are fixed, the entropy capacity is then only a function of the peak field. As an example, we consider the case of a CPA stage operating between 1.3 K and 50 mK. Fig. 3 shows the specific entropy of CPA at 1.3 K as a function of peak magnetic field.

If the required cooling capacity is ΔQ , the required salt mass (in moles) as a function of magnetic field is

$$n_{salt}(B) = \frac{\Delta Q / T_{low}}{S(B, T_{sink}) - S(0, T_{low})} \quad (3)$$

The result is shown in Fig. 4 for the example case of CPA operating at 50 mK under a $1 \mu\text{W}$ load for 24 h.

Translating the salt mass into a magnet bore volume, and from there a magnet mass, requires some details of the salt pill design, such as the aspect ratio (length to diameter), the fraction of the volume that will be devoted to the thermal bus and hermetic structure, and the gap between the salt pill and magnet bore. For many

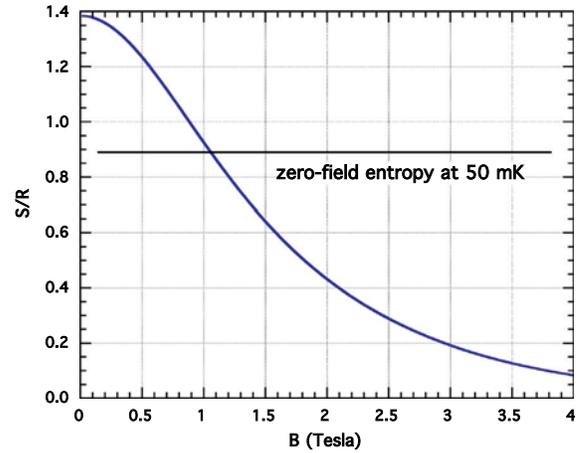


Fig. 3. Specific entropy of CPA at 1.3 K. The zero-field entropy for CPA at 50 mK is indicated.

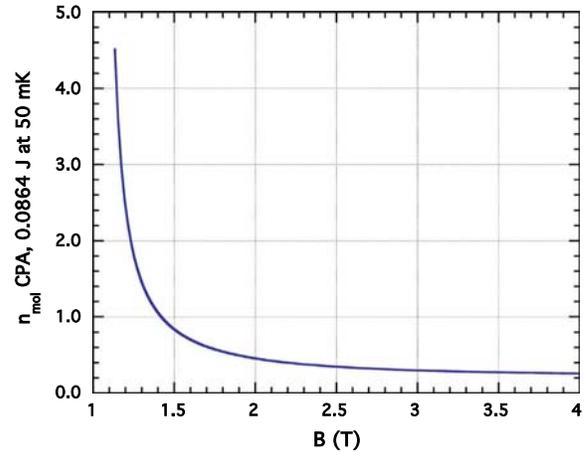


Fig. 4. Moles of CPA required to absorb $1 \mu\text{W}$ for 24 h at 50 mK, when demagnetized from 1.3 K and magnetic field, B . Ideal (Carnot) efficiency is assumed.

of these, general guidelines will suffice. The salt typically occupies at least 90% of the salt pill volume, and a gap of 0.75 mm between the salt pill and magnet bore is adequate, given the rigidity of typical suspension systems.

The issue of aspect ratio is more complex, as it is a key parameter affecting the thermal conductance of the salt pill's thermal bus [7,10], and is one of several parameters (along with magnet winding length, wire diameter and magnetic shield configuration) that affect the spatial distribution of the magnetic field in the volume of the salt, and hence the salt's overall cooling capacity.

For simple solenoids, the field decreases rapidly toward the end of the windings. For a semi-infinite solenoid, the field at the open end is half the central field. For smaller aspect ratio magnets, the same winding density produces a lower central field, but the field is more uniform along the length. Adding ferromagnetic shielding will both increase the central field, and further improve axial uniformity. Superconducting shields and bucking coils, on the other hand, tend to decrease the central field and increase gradients. Their advantage is that the shield components are less massive than ferromagnetic shielding. But as the magnet size must be increased to compensate for the lower average field, there is generally no net mass advantage.

As a result, it has become common to use ferromagnetic shields to provide basic control of fringing fields in ADRs. Some basic design rules were developed by Hagmann and Richards [14]. The

shield is sized so that when the flux returned outside of the windings is concentrated into the cross section of the shield, the shield just reaches saturation. The flux can be calculated from the central field, and the inner and outer diameters of the windings. This would seem to give a simple means of estimating shield mass for an ADR stage, and that is true if the shield is a simple cylindrical shield. However, the mass of a shield can be reduced significantly by contouring the shape to reflect the field profile along the axis and at the ends. An optimization of this sort requires specialized magnetic field modeling software, such as MAGNETO [15], to compute fields for asymmetric geometries.

Consequently, it is beyond the scope of a general optimization analysis to include magnetic shields. We note, though, that for a given central field, ferromagnetic shield thickness – and mass – is minimized when the magnet's outer diameter is minimized. Thus when optimizing aspect ratio and magnetic field, some weight could be given to the combination of parameters that yields the smallest magnet diameter. In fact, when the optimal aspect ratio and magnetic field are based on minimum stage mass, the resulting magnet diameter is the minimum.

2.3. Optimization of aspect ratio

To demonstrate the effect of varying aspect ratio, we choose a point design that is representative of low temperature ADRs in terms of operating temperature and cooling capacity, in this case a salt pill containing 0.5 mol of CPA, and magnet with an average field of 2 T. For a given aspect ratio, one can compute the salt pill's length and diameter, and the magnet's length and inner diameter. (It is assumed that the two lengths are equal.) From these, the outer diameter of the windings needed to achieve the required field (2 T) is computed. Mathematica™ [11] is used to perform numerical integration of Maxwell's equations for a loop current over the extent of the magnet's windings and the salt pill volume. Magnet mass is computed using an average density of 5 g/cm³.

One final input to the calculation is the engineering current density, j_E , of the superconducting wire, which itself is dependent on the maximum field that is produced in the windings. In the production of ADRs for space applications, small diameter magnet wire is typically used in order to limit the current required. The peak current affects both the power draw of the control electronics, and static conduction and ohmic dissipation in the dewar's current leads. The latter are often one of the major heat loads on the cryogenic system. Optimization of the current therefore involves trades between the spacecraft power and thermal systems, the control electronics, the instrument, and the dewar, as well as such considerations as the commercial availability of small diameter wire.

Magnets for space ADRs have tended to use currents in the range of 2–5 A. The ADRs built for XRS [12] on Astro-E, -E2 and for SXS on Astro-H [13] all used 0.1 mm diameter multifilamentary NbTi wire, and a maximum current of 2 A to produce fields in the range of 2–3 T. The corresponding j_E is about 2.0×10^{-4} A/cm². Representative critical current values (at 4.2 K) for 15S40 NbTi wire supplied by Supercon, Inc. [16], used in the Astro-H/SXS magnets, is shown in Fig. 5.

Note that at 2 A, the wire is superconducting at fields greater than 7 T. While laboratory magnets are routinely operated to within a few % of their critical current, for space instruments, some margin is usually added (the statistics here are very limited). For SXS on Astro-H, the required margin on current and field was 25%. In reality both the 2 T and 3 T magnets could sustain currents of in excess of 2.7 A before quenching.

Magnet mass versus aspect ratio is shown in Fig. 6. While the mass has a minimum at about 2.75, the minimum is rather broad and shallow, and suggests that the aspect ratio can be adjusted as needed for other concerns (such as meeting diameter or height

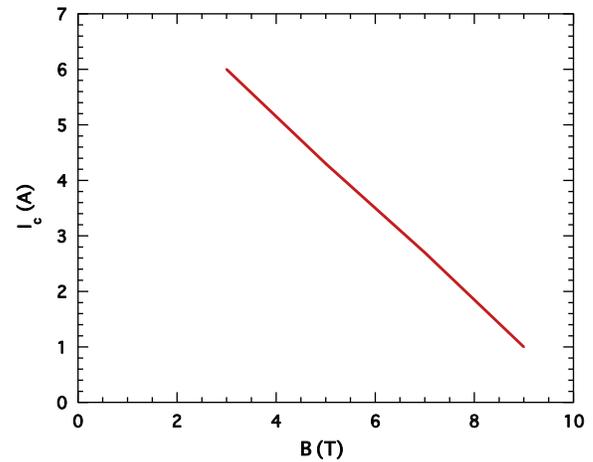


Fig. 5. Critical current versus magnetic field for multifilamentary NbTi wire (15S40) supplied by Supercon, Inc.

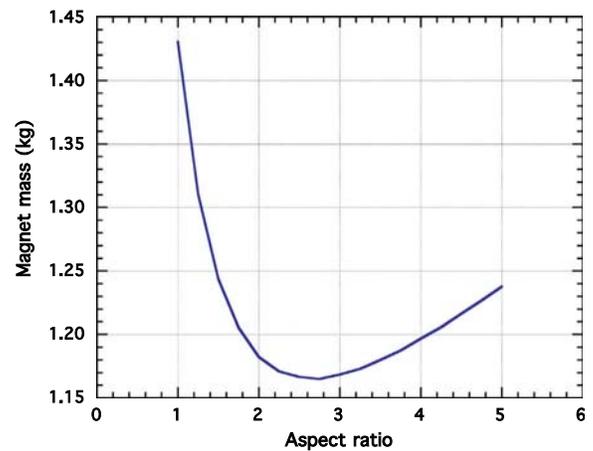


Fig. 6. Mass of a 2 T (average field) magnet as a function of salt pill aspect ratio for a 0.5 mol CPA salt pill.

constraints, or optimizing the thermal bus) without significant penalty in performance.

2.4. Optimization of magnetic field

Using 2.75 for the salt pill's aspect ratio, Fig. 7a presents the estimated mass of both the salt pill and magnet for an ADR stage using CPA that is sized for a cooling capacity of 1 μW for 24 h at 50 mK, and assuming the salt pill is demagnetized from $T_{high} = 1.3$ K. Fig. 7b shows the same analysis for GGG sized for 20 μW for 24 h at 1 K using a 4.25 K heat sink. The heat loads chosen represent the same entropy absorption for both stages, but in practice the cooling capacity does not appreciably affect the optimal field value.

At low fields, a very large salt mass is needed due to the rapidly diminishing entropy capacity. At high fields, the salt's entropy is already suppressed close to zero, and additional field does not appreciably increase the cooling capacity.

A fuller mapping of the optimal magnetic field provides a useful tool for optimizing stages for any combination of operating parameters. Fig. 8 shows the results for CPA over the range of $T_{low} = 50$ –300 mK, and $T_{high} = 0.3$ –2.6 K. It turns out the two dependences are separable, as the function:

$$B_{optimal_{CPA}} = 0.572 T_{low}^{-0.35} T_{high}^{0.75} \text{ (in Tesla)} \quad (4)$$

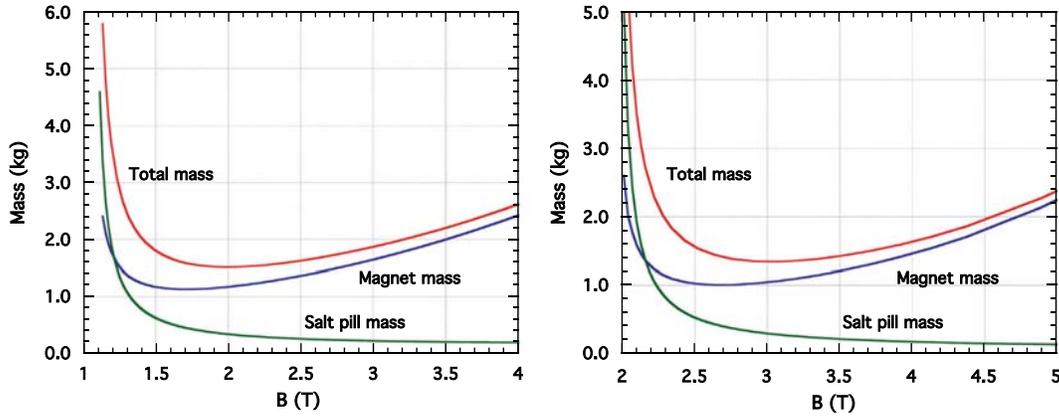


Fig. 7. Component masses for (a) a CPA stage designed to operate at 50 mK from a 1.3 K heat sink, and (b) a GGG stage designed to operate at 1 K from a 4.2 K heat sink, as a function of magnetic field.

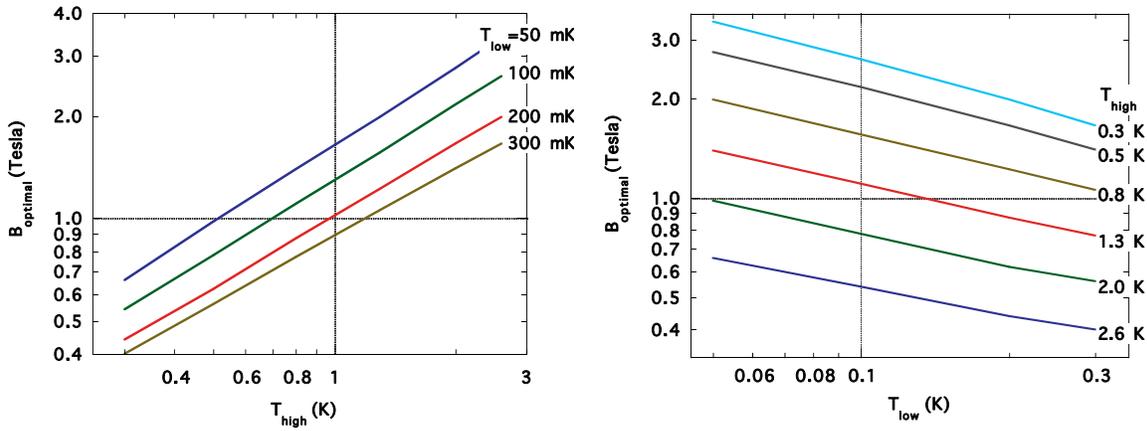


Fig. 8. Optimal magnetic field for a CPA stage as a function of demagnetization temperature and operating temperature.

The largest deviations are at the high end of the range for sink temperature, where the optimal field is overestimated by less than 3%.

A similar analysis for FAA yields:

$$B_{optimal_{FAA}} = 0.355 T_{low}^{-0.45} T_{high}^{0.72} \text{ (in Tesla)} \quad (5)$$

As an example, consider the ADR flow with the XRS instrument on Astro-E and Astro-E2, for which $T_{high} \sim 1.5$ K (using a 1.3 K helium bath as a sink), and $T_{low} = 0.06$ K. From Eq. (5), the optimal field is 1.69 T. In practice, since the optimal design point is at a very shallow minimum, somewhat larger magnetic fields (2 T for XRS) are often chosen to provide margin on operating conditions on orbit.

3. Optimization of two-stage ADRs

Two-stage ADRs that are configured in parallel are essentially independent stages that use the same heat sink, and are optimized separately as for single stages. When the stages are arranged in series, as in Fig. 9, the upper stage (S2) has two roles: to absorb heat from the lower stage (S1) in order to regenerate its entropy capacity, and to absorb heat loads during the hold time (and thereby reduce the heat load on S1). The transfer point between stages, T_i , and S2's hold temperature, T_2 , are parameters that need to be optimized.

Recycling involves two separate heat transfers, from S1 to S2, and from S2 to the heat sink, not necessarily in that order. In fact, there are many possible sequences, including

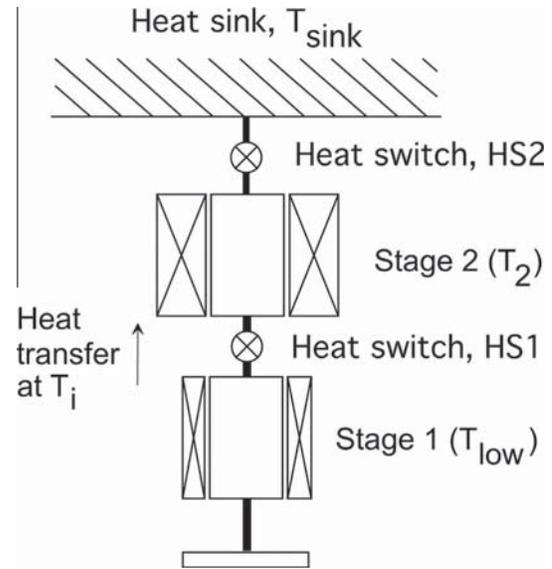


Fig. 9. 2-stage ADR coupled in series. The two stages are stationed at T_2 and T_{low} during the hold time, and heat is transferred between stages at T_i .

1. Recycle S1; recycle S2; demagnetize S1 and S2.
2. Recycle S2; recycle S1; demagnetize S1 and S2.
3. Recycle S2; recycle S1; recycle S2; demagnetize S1 and S2.

Schematically these sequences are shown in Fig. 10. Each recycling sequence begins at the end of a hold time, and each recycling operation follows that of a single stage. For simplicity, the heat transfer point between S1 and S2 (T_i) is the same as S2's hold temperature (T_2). In all sequences, S1 and S2 absorb heat during the hold time. In sequence 1, S2 then absorbs entropy from S1, and finally rejects all entropy absorbed during the cycle to the sink. In sequence 2, the order of heat rejection is simply reversed. In sequence 3, the first recycling of S2 rejects the entropy it absorbed during the hold time, and the second recycling rejects the entropy absorbed from S1. By recycling S2 multiple times, the stage can be smaller, sized for the *larger* of the two entropy loads, not the sum. The trade-off is a somewhat longer recycle time. Note, though, that each recycle of S2 is shorter than for sequences 1 and 2 since less heat is transferred each time.

The choice between sequences 1 and 2 essentially depends on whether T_i is greater than or less than T_2 . The largest entropy change for the cycle will be obtained when S2 completes its cycle at the higher of the two temperatures. To illustrate the point, Fig. 11 shows the entropy diagram for the upper stage of a 2-stage ADR designed for a 1.8 K heat sink, using $T_i = 0.8$ K and $T_2 = 0.5$ K. Using sequence 1, after S2 is recycled, it is cooled to 0.5 K for the hold period. At the end of the hold, the temperature is raised to 0.8 K where heat is absorbed from S1, then S2 is recycled again. The entropy capacity of the stage is much larger than if the sequence were reversed.

3.1. Optimization of T_i and T_2

Depending on how the ADR and the components it cools are configured, the heat loads on S1 will be dependent either on the heat sink temperature, or on the intermediate temperature, T_2 . The load (detectors, amplifiers, etc.) and salt pill suspension generally have direct links to the heat sink, but some designs will include a heat intercept to S2. In the latter case, S2 will absorb the majority of the heat conducted through those components, and the smaller heat load reaching S1 will be dependent on T_2 . The conducted load from heat switch HS1 will necessarily depend on T_2 .

Entropy absorbed by S1 during the hold time is transferred to S2, before rejecting it to the sink. The heat load on S1 therefore affects the required cooling capacity of both stages. The required entropy capacity of S1 is

$$\Delta S_1 = \frac{\dot{Q}_{1,sink}(T_{sink}) + \dot{Q}_{S2}(T_2)}{T_{low}} \Delta t_{hold} / \varepsilon_1 \quad (6)$$

where $\dot{Q}_{1,sink}(T_{sink})$ is the conduction to S1 through components linked to the heat sink, $\dot{Q}_{S2}(T_2)$ is the conduction through

components linked to S2, Δt_{hold} is the hold time, and ε_1 is the heat absorption efficiency (the “usable” cooling capacity as a fraction of ideal). For stages operating below 200 mK, the efficiency tends to be in the range of 70–80%, where warmer stages can achieve efficiencies of 80–90%.

The required entropy capacity for S2 is then

$$\Delta S_2 = \left(\Delta S_1 / \varepsilon_{1 \rightarrow 2} + \frac{\dot{Q}_{2,sink}(T_{sink}) - \dot{Q}_{S2}(T_2)}{T_2} \Delta t_{hold} \right) / \varepsilon_2 \quad (7)$$

where $\dot{Q}_{2,sink}(T_{sink})$ is the conduction to S2 through components linked to the heat sink, ε_2 is the efficiency of stage 2, and $\varepsilon_{1 \rightarrow 2}$ is the efficiency of the entropy transfer from S1 to S2. $(1 - \varepsilon_{1 \rightarrow 2})$ is essentially the ratio of the temperature gradient across HS1 divided by the average temperature of S1 and S2 as heat is transferred. Typically this will be 90%, although larger or smaller gradients can be used depending on the heat switch conductance.

The above expression assumes that either recycling sequence 1 or 2 is being used, for which the capacity of S2 is based on entropy absorption from S1 and during the hold time. When sequence 3 is used, the expression should include only the larger of the two entropy terms.

The entropy capacity of both stages can also be expressed in terms of their cycle parameters:

$$\Delta S_1 = n_{salt_1} (s(0, T_{low}) - s(B_1, T_i)) \quad (8)$$

$$\Delta S_2 = n_{salt_2} (s(0, T_i) - s(B_2, T_{sink})) \quad (9)$$

where n_{salt_1} and n_{salt_2} are the salt masses used in S1 and S2, and B_1 and B_2 are the peak fields. In the expression for ΔS_2 it is assumed that T_i is larger than T_2 . If not, the zero-field entropy term should be replaced with $s(0, T_2)$.

The effects of varying T_i and T_2 can now be seen. Decreasing T_2 lowers the heat load on S1, and the subsequent heat transfer to S2, but it increases the entropy generation associated with the heat load on S2 during the hold time. Below the optimal value of T_2 , the increase in salt mass and magnetic field for S2 needed to accommodate the elevated entropy generation rate more than offsets the reduction in salt mass and magnetic field for S1 and S2 associated with the lower heat load on S1.

Decreasing T_i increases the entropy capacity of S1, but decreases the entropy capacity of S2. The optimal value of T_i is the point where the combined mass of the salt pills and magnets for S1 and S2 is a minimum.

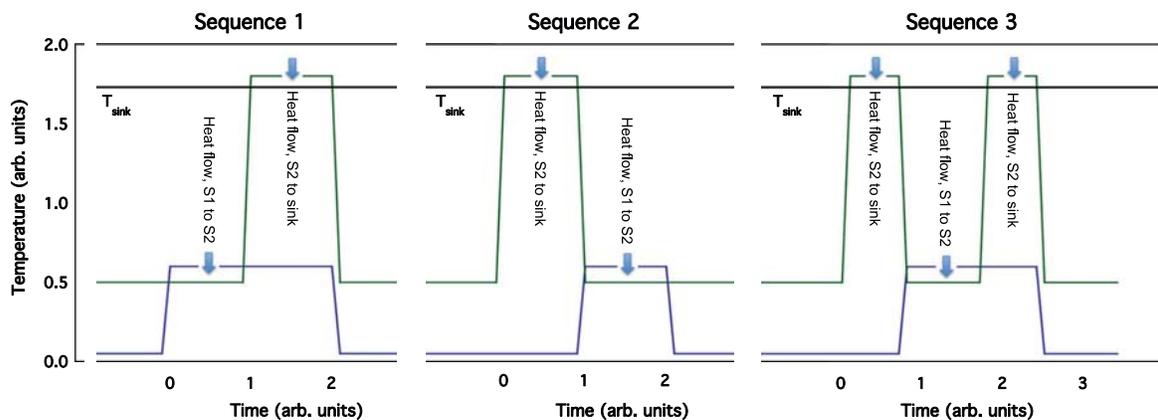


Fig. 10. Typical recycling sequences for a 2-stage ADR. Recycling begins at the end of the hold time (time equal zero).

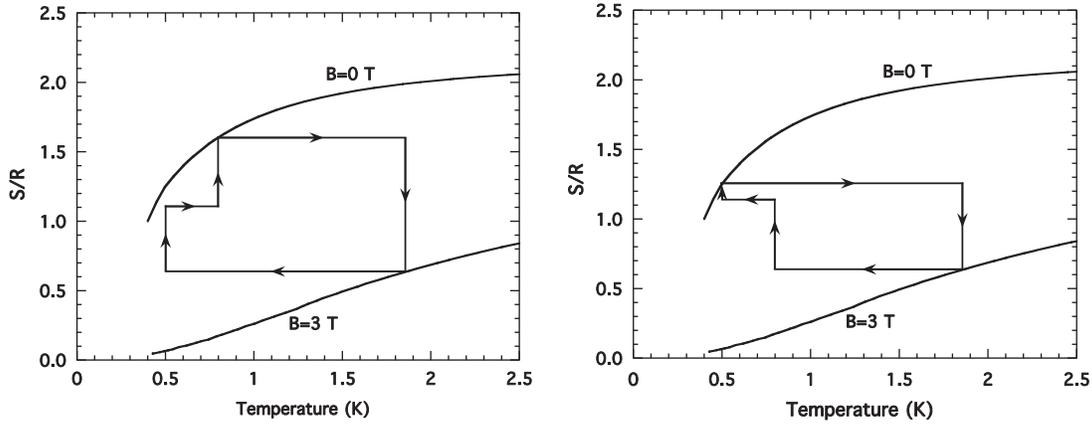


Fig. 11. Entropy diagram for the upper stage of a 2-stage series ADR, depicting sequence 1 (left) and sequence 2 (right) for the case when T_i is greater than T_2 . The diagram is based on the properties of gadolinium lithium fluoride (GLF) at representative fields.

4. Optimization of multi-stage and continuous ADRs

In the preceding sections, it has been assumed that the number of ADR stages is fixed, and the parameters affecting their operation are varied in order to obtain the highest cooling capacity per unit mass. But for an ADR using stages connected in series between two temperatures, the number of stages can also be considered an adjustable parameter. Multi-staging can reduce the overall mass by (1) reducing the magnetic field requirements for each stage, and (2) reducing parasitic heat loads and associated entropy generation.

A simple case to examine is the use of a one or two ADR stages to span an increasing temperature range. Using 50 mK as a “standard” operating temperature, and a fixed cooling capacity of 0.0864 J (1 μ W absorbed over 24 h), Fig. 12 compares the mass of single- and two-stage ADRs as a function of T_{sink} . CPA is used for all stages. For the two-stage ADR, the hold temperature, T_2 , and the intermediate temperature, T_i , are both equal to the logarithmic average of 50 mK and T_{sink} . (In this case, one can use either recycling sequence 1 or 2, as they are equivalent.) The optimal magnetic fields are computed from Eq. (4).

For fields above ~ 2 T, the single-stage mass has a temperature dependence of $\sim T_{sink}^{4.5}$. The two-stage ADR uses much lower magnetic fields in each stage, with correspondingly lower magnet masses. As a result, it is advantageous to transition to two stages when the ratio of temperatures (T_{sink}/T_{low}) exceeds ~ 30 , at which point each stage has an operating range of about a factor of ~ 5.5 in temperature.

The advantage of multi-staging is even more pronounced when the effect on cooling capacity is considered. Components, such as heat switches, that link to an intermediate stage rather than directly to the heat sink will impose lower heat loads on the lower stage by a factor of

$$\frac{\dot{S}_2}{\dot{S}_1} = \frac{\int_{T_{low}}^{T_i} k(T) dT}{\int_{T_{low}}^{T_{sink}} k(T) dT} \approx \left(\frac{T_i}{T_{sink}}\right)^{m+1} \quad (10)$$

where thermal conductance is taken to be of the form $k \sim T^m$, and $T_{low}^m \ll T_i^m, T_{sink}^m$. If $T_i/T_{sink} \sim 1/5$ and $m = 1$, the reduction is a factor of 25.

The total heat load will, in general, be reduced by a smaller fraction due to contributions that are not affected by the intermediate stage, and a full optimization needs to be performed for each stage as outlined in the previous section. However, the combined benefits of smaller magnetic fields and lower heat loads suggest that an ADR’s performance can be improved if, by adding a stage, the temperature span (T_{high}/T_{low}) of each stage is reduced into the range of 5–6.

4.1. Continuous ADRs

The physical arrangement of stages in a series CADR [5] is the same as for single-shot, multi-stage ADRs, and each stage operates on the same cycle: giving off heat during an isothermal magnetization phase at “high” temperature with its heat switch closed, and

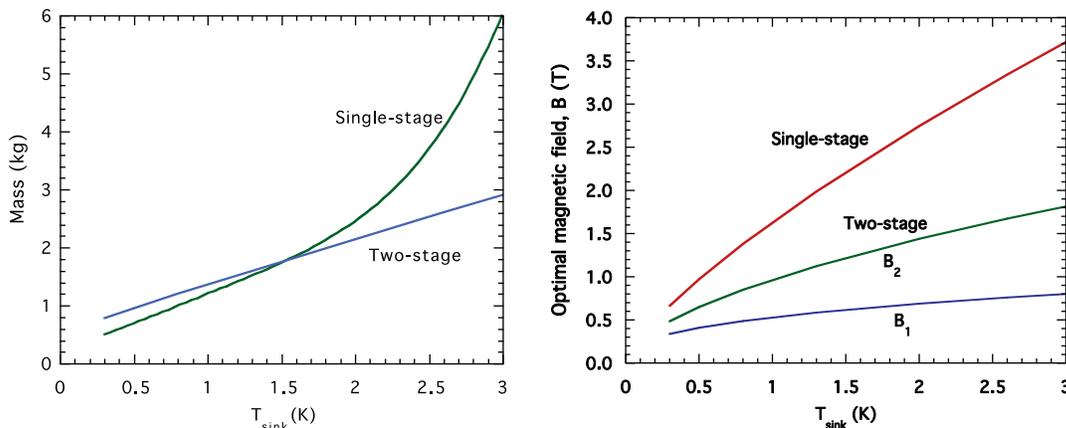


Fig. 12. Comparison (left) of the masses of a single-stage and a two-stage ADR spanning 50 mK to T_{sink} , and the optimal magnetic fields (right) for both ADRs. For the 2-stage ADR, $T_2 = T_i = \sqrt{(50 \text{ mK})T_{sink}}$, and recycling sequence 1 is used.

absorbing heat at “low” temperature during an isothermal demagnetization phase with its heat switch open. If there is a lower stage, the heat absorption is done with that stage’s heat switch closed in order to regenerate its cooling capacity.

The first notable difference between CADRs and single-shot ADRs involves the control of the lowest stage in the chain (the “continuous” stage). Its low and high temperatures are the same, and it alternates between isothermal magnetization and isothermal demagnetization as needed to maintain a constant temperature. The second difference is the duty cycle of each stage. In a single-shot ADR, the cascade of heat to the sink is performed over a small fraction of the total cycle, so that the majority of time is spent in the isothermal demagnetization phase. In a CADR, the cycling is continually repeated, with each stage spending about half the time rejecting heat, and half the time absorbing heat.

This has an impact on the requirements for the heat switches, and consequently, for the operating range of each stage. Heat switches are designed to have high enough thermal conductance in the closed state to recycle the ADR within an allocated time, and low thermal conduction in the open state to minimize parasitic heat loads. The two requirements tend to drive the design in opposite directions – larger structures needed to obtain higher “on” conductance generally have higher heat leak in the “off” state.

The linkage is not absolute, of course. For superconducting heat switches, the A/L of the superconductor dictates both the on and off state conductances [17]. But for gas-gap heat switches, the on conductance is proportional to the volume occupied by the conductive elements, while the off conduction is a function of the aspect ratio of that volume. Thus depending on the type of switch used, there may be design choices that allow higher on conductance without sacrificing off conduction.

The requirements for heat switch performance can be derived for different ADR configurations. Following an analysis by Kittel [18], we evaluate the entropy flow of a heat switch in both the on and off states. During the isothermal demagnetization phase (absorbing heat at T_{low} while the warm end of the heat switch is T_{high}), the entropy load from the heat switch is

$$\dot{S}_{inHS} = \left(\frac{\tau_{hold}}{\tau}\right) \frac{1}{T_{low}} \int_{T_{low}}^{T_{high}} k_{OFF}(T) dT \quad (11)$$

and the entropy flow out during the isothermal magnetization (rejecting heat at T_{reject}) is

$$\dot{S}_{out} = \left(1 - \frac{\tau_{hold}}{\tau}\right) \frac{1}{T_{reject}} \int_{T_{reject}}^{T_{reject} + \Delta T} k_{ON}(T) dT \quad (12)$$

where τ_{hold}/τ is the fraction of time that the stage is cold (also called the duty cycle), ΔT is the gradient across the switch when heat is being rejected, and $k_{ON}(T)$ and $k_{OFF}(T)$ are the on and off state conductances. At low temperature, the latter can be expressed as $k_{ON}(-T) = \alpha T^n$ and $k_{OFF}(T) = \beta T^m$.

In a practical system, the stage is designed to lift not only its own parasitic heat, but also heat from a load, which adds an entropy term \dot{Q}/T_{low} to \dot{S}_{in} . Assuming $T_{low}^m \ll T_{high}^m$ and $\Delta T \ll T_{reject}$, the entropies reduce to

$$\dot{S}_{in} = \frac{\dot{Q}}{T_{low}} + \left(\frac{\tau_{hold}}{\tau}\right) \frac{1}{T_{low}} \frac{\beta}{m+1} T_{high}^{m+1} \quad (13)$$

$$\dot{S}_{out} = \left(1 - \frac{\tau_{hold}}{\tau}\right) \alpha T_{reject}^{n-1} \Delta T \quad (14)$$

In an idealized ADR, the heat switch must have characteristics that yield $\dot{S}_{out} > \dot{S}_{in}$. In a practical system, the realizable cooling capacity is typically 70–90% of ideal, depending on temperature. A more realistic requirement is therefore $\dot{S}_{out} > \dot{S}_{in}/\varepsilon$, where ε is

an efficiency factor that accounts for gradients within the salt pill, eddy current losses, etc.

Applying these expressions to the continuous stage: $T_{reject} = T_{low}$, and $\tau_{hold}/\tau \sim 0.5$. The on and off state conductivities will depend on the particular switch technology used. For a stage that will operate below about 200 mK, the choice is limited to those that use metallic conductors: superconducting, magnetoresistive, or mechanical. For this discussion, a superconducting switch is assumed. The off thermal conductance is governed by the population of quasiparticles, giving a relatively strong temperature dependence. Lead (Pb) switches developed at NASA/GSFC [17] have $k_{OFF}(T) = 2.15 T^{3.21}$ (mW/K). In the on state, $k_{ON}(T) = 110 T$ (mW/K).

Inserting these values gives the requirement

$$\begin{aligned} \dot{S}_{out} &= 55 \Delta T (\text{mW/K}) > \dot{S}_{in}/\varepsilon \\ &= \frac{1}{\varepsilon T_{low}} \left(\dot{Q} + 0.256 T_{high}^{4.21} (\text{mW/K}) \right) \end{aligned} \quad (15)$$

For a CADR targeting $\dot{Q} = 5 \mu\text{W}$ of cooling power at $T_{low} = 50$ mK (using $\Delta T = 5$ mK), and using $\varepsilon = 0.7$, T_{high} is constrained to be < 0.38 K. At this temperature, the time average heat load on the continuous stage is $9.36 \mu\text{W}$.

If the size of the superconducting element is changed, for example to a larger A/L, both the on and off conductances change by the same proportion. The effect will be to make the fixed heat inputs (\dot{Q}) smaller relative to the heat rejection capability. On the surface this would provide leeway to increase T_{high} , but higher heat flows can lead to less efficient transfers, and such adjustments should be approached cautiously.

For the upper stages of the CADR, switching ratios are less of a driver for operating temperature, and there is greater flexibility in the choice of technologies. Gas-gap switches can be used at temperatures above about 0.3 K, with very high intrinsic on/off ratios. On state conductances on the order of $k_{ON}(T) \sim 25 T$ (mW/K) are easily achieved, with off conductances on the order of $k_{OFF}(-T) \sim 0.5 T^{2.4}$ ($\mu\text{W/K}$)

Using these parameters, along with $\tau_{hold}/\tau \sim 0.5$, Eqs. (13) and (14) reduce to

$$\dot{S}_{in} = \frac{\dot{Q}}{T_{low}} + 0.074 \frac{T_{high}^{3.4}}{T_{low}} (\mu\text{W/K}) \quad (16)$$

$$\dot{S}_{out} = 12.5 \Delta T (\text{mW/K}) \quad (17)$$

The direct heat load term is retained, and represents the entropy that each stage must absorb from the next lower stage. For the second stage of the CADR (connected to the continuous stage via the superconducting heat switch), this term is $9.38 \mu\text{W}$ divided by the efficiency of 0.7. The remaining parameters are: $T_{low} = 45$ mK, $T_{high} = 0.38$ K, and $\Delta T = 0.038$ K. The two entropy terms contributing to \dot{S}_{in} are 0.30 mW/K and 0.030 mW/K, the smaller attributable to the gas-gap heat switch. \dot{S}_{out} , in this case, is 0.48 mW/K – more than adequate for cascading heat to the next (3rd) stage.

Since heat switch performance is not the driver for the operating range of the 3rd and higher stages, their optimization will be driven by the individual and overall stage masses as described at the beginning of this section. The main result is that each stage is optimally designed to operate over about a factor of 5 in temperature. Table 1 summarizes the operating parameters for a 4-stage CADR [5] designed to operate between 50 mK and 5 K.

4.2. CADR cooling power

A final consideration for CADR design is the relationship between system parameters (heat switch characteristics, operating

Table 1
Operating Parameters for a 4-stage CADR operating between 50 mK and 5 K.

Stage	Refrigerant	B (T)	T_{high} (K)	T_{low} (K)	Mass (kg)
1	60 g CPA	0.1		0.05	0.5
2	100 g CPA	0.5	0.3	0.045	1.7
3	100 g CPA	1.5	1.3	0.25	2.4
4	65 g GLF	4.0	5	1.15	3.1

temperatures, salt pill size, etc.) and cooling power. Unlike single-shot ADRs where salt pill size is dictated by heat loads and required hold times, CADRs do not have a hold time requirement. Instead, the relevant time frame for computing cooling capacity is the cycle period. The cycle period, however, depends largely on the time required to transfer the cooling capacity of the continuous stage to the 2nd stage.

As a result, the size of a CADR's salt pills have far less impact on cooling power than the conductance of the heat switches. And in this regard, the continuous stage's heat switch is most important, since it has a direct effect on the entropy load that must be cascaded through the CADR. Its on conductance can be increased with a larger superconducting element, but this also raises the off conductance. An effectively larger switching ratio can be obtained by reducing the peak temperature of the 2nd stage. In the 4-stage CADR, T_{high} for the superconducting switch was limited to 0.3 K, reducing its average load on the continuous stage from 4.2 μ W to 1.6 μ W – nearly a 30% reduction in the total heat load.

CADR performance can also be improved by reducing the fraction of the cycle time spent in “overhead” operations, including ramping the upper stages over their temperature range, and opening and closing heat switches. Because the CADR magnets are small, the stages can easily be ramped between operating points within a minute (each way), and heat switch transition times can be limited to zero in the case of passive gas-gap heat switches, or to less than a minute for active versions. At these (de)magnetization rates, eddy currents [9] are a concern, and the salt pill's are specifically designed to minimize dissipation. Primarily this consists of minimizing closed loops of high purity copper used in the thermal buses, and locating bulk materials used for thermal interfaces outside the magnetic shields. A worst-case estimate is that each salt pill's cooling capacity must be increased by 5% to account for eddy current losses over a cycle.

In practice the overhead operations add 2–4 min to the CADR cycle. For the system described in [19], the resulting cycle time was \sim 20 min. Since the overhead operations are relatively fixed in duration, some improvement in performance can be realized by increasing the cooling capacity of each salt pill, and increasing the duty cycle for heat transfer.

5. Summary

ADRs are unique among ultra-low temperature refrigerators in that they can use a wide range of (magnetic) refrigerants, and their cooling capacity is determined both by the amount of refrigerant and the magnetic field used. In addition, there are many architectures that can be employed – from single-stage units, to multi-stage single-shot systems, to multi-stage continuous versions. There is, then, a very wide parameter space for optimizing their design and performance that can at first be daunting.

However, a given set of operating requirements can be translated into an optimized system by design guidelines outlined in this and other references [3,7,9,A,17]. These optimization rules developed in this paper, based on obtaining the lowest stage mass for a required cooling capacity, can be summarized as follows.

- The optimal refrigerant will have an ordering temperature slightly below the lowest desired operating temperature. CPA and FAA are typical for stages operating down to 30 mK and 50 mK; above 0.5 K, GLF and GGG become usable.
- The use of ferromagnetic shielding enhances a magnet's central field and reduces axial and radial field gradients; superconducting shielding and bucking coils have the opposite effect, and require scaling up the central field.
- The optimal salt pill aspect ratio is \sim 2.75, although the mass minimum for this parameter is broad and shallow. Different aspect ratios can be used to accommodate size constraints or to improve thermal bus conductance without significant mass penalty.
- For CPA and FAA, an expression for the optimal magnetic field is presented in terms of the high and low operating temperatures of the stage. From these, a stage operating between 50 mK and 1.5 K would optimally use a \sim 2 T field.
- ADR stage mass grows rapidly as the required field exceeds \sim 2 T, and/or the operating range exceeds a factor of \sim 30. Increasing the number of stages reduces mass by reducing the required magnetic fields and cooling capacities. The intermediate temperature of a 2-stage ADR should be chosen to minimize the entropy absorption during the hold time, and will be close to the logarithmic average of the low end and heat sink temperatures.
- For multi-stage ADRs, the recycling sequence has a significant affect on the mass of each stage. Optimization is based on the relative size of internal (parasitic) and external heat loads. Recycling upper stages multiple times can significantly reduce their required cooling capacity.
- A continuous ADR (CADR) has significantly lower mass than a single-shot ADR for the same cooling power. Performance is dictated by cycle period, and is far more driven by heat switch conductances than salt pill size.
- The operating range of each stage in a CADR is dictated by the on- and off-state conductances of the lowest temperature heat switches, and is optimal at $T_{high}/T_{low} \sim 5$.

References

- [1] Debye P. *Ann Phys* 1926;81:1154.
- [2] Wikus P, Burghart G, Figueroa-Feliciano E. *Cryogenics* 2011;51:555–8.
- [3] Wikus P, Canavan E, Heine ST, Matsumoto K, Numazawa T. Magnetocaloric materials and guidelines to their selection. *Cryogenics* 1994;34:213–9. This issue on magnetic refrigeration, this issue.
- [4] Shirron P. Applications of the magnetocaloric effect in single-stage, multi-stage and continuous adiabatic demagnetization refrigerators. *Cryogenics*, *Cryogenics special issue on magnetic refrigeration*, this issue.
- [5] Shirron PJ, Canavan ER, DiPirro MJ, Tuttle JG, Yeager CJ. A multi-stage continuous-duty adiabatic demagnetization refrigerator. *Adv Cryog Eng* 2000;45:1629–38.
- [6] Vilches OE, Wheatley JC. Measurements of the specific heat of three magnetic salts at low temperatures. *Phys Rev* 1966;148:509.
- [7] Hagmann C, Benford DJ, Richards PL. Paramagnetic salt pill design for magnetic refrigerators used in space applications. *Cryogenics* 1994;34:213–9.
- [8] Serlemitsos AT, SanSebastian M, Kunes E. Design of a spaceworthy adiabatic demagnetization refrigerator. *Cryogenics* 1992;32:117.
- [9] McCammon D, Shirron P. Salt pill design and fabrication for adiabatic demagnetization refrigerators. *Cryogenics* (this issue on magnetic refrigeration).
- [10] Wilson GW, Timbie PT. Construction techniques for adiabatic demagnetization refrigerators using ferric ammonium alum. *Cryogenics* 1999;39:319–22.
- [11] Mathematica™ is a registered trademark of Wolfram Research Inc., 100 Trade Center Drive, Champaign, IL 61820-7237.
- [12] Serlemitsos AT, SanSebastian M, Kunes E. Final design of the XRS/Astro-E ADR. In: Kittel P et al., editors. *Advances in cryogenic engineering* 43A. New York: Plenum; 1988. p. 957–63.
- [13] Shirron P, Kimball M, Wegel D, Miller F. ADR design for the soft X-ray spectrometer instrument on the Astro-H mission. *Cryogenics* 2010;50:494–9.
- [14] Hagmann C, Richards PL. Adiabatic demagnetization refrigerators for small laboratory experiments and space astronomy. *Cryogenics* 1995;35:305.

- [15] MAGNETO is a product of INTEGRATED Engineering Software, 220-1821 Wellington Avenue, Winnipeg, Manitoba, R3H 0G4 Canada.
- [16] Supercon Inc, 830 Boston Turnpike, Shrewsbury, MA 01545.
- [17] DiPirro M, Shirron P. Heat switches for adiabatic demagnetization refrigerators. *Cryogenics* (this issue on magnetic refrigeration).
- [18] Kittel P. Heat switch limitations on multi-stage magnetic refrigeration. *AIP Conf Proc* 2002;613:1167.
- [19] Shirron PJ, Canavan ER, DiPirro MJ, Francis J, Jackson M, Tuttle JG, et al. Development of a cryogen-free continuous ADR for the constellation-X mission. *Cryogenics* 2004;44:581–8.