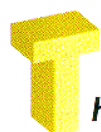


# Aluminum electrolysis process simulation

Marc Dupuis



**The Hall-Héroult aluminum electrolysis process is very complex. It involves many different physical and chemical phenomena, which are not completely understood and often interact with each other. Yet, we need to be able to represent those complex phenomena in**

comprehensive mathematical models to avoid designing cells by trial and error. This has been achieved; today's modern smelter technology has been designed with the support of these mathematical models [1].

In general, we can fit Hall-Héroult mathematical models into three broad categories:

- Thermal-electric models, which are generally associated with the problem of cell heat balance.
- Magneto-hydro-dynamic (MHD) models, which are generally associated with the problem of cell stability.
- Stress models, which are generally associated with cell shell deformation and cathode heaving issues.

All those aspects of cell design are linked in different ways:

- MHD is affected by the ledge profile, mostly dictated by the cell heat balance design.
- Local ledge profile is affected by the metal recirculation pattern, mostly dictated by the busbar's magnetic design.
- Shell deformation is strongly influenced by the shell thermal gradient, controlled by the cell heat balance design.
- Steel shell structural elements like cradles and stiffeners influence the MHD design through their magnetic shielding property.
- Global shell deformation affects the local metal pad height, which in turn affects both the cell heat balance and cell stability.

Unfortunately, at the present time, even taken individually, those models require tremendous computer resources, preventing us from merging them into a big "multiphysics" unified model [2]. Some initial model developments have, however, been made.

## Thermal-electric models

This article will concentrate on the thermal-electric aspect of the Hall-Héroult cell modelling and design, keeping in mind that this aspect of the design will influence the other aspects, as other aspects of the cell design will influence the cell energy balance.

Unfortunately, even at this level, limited computer resources force us to divide the modelling of the thermal-electric aspect of the process in two categories:

- Steady-state models are used to compute in detail and with good accuracy the global "average" heat balance.
- Dynamic models are used to compute the cell thermal response to discrete process perturbations (anode effect (AE), anode change, metal tapping, etc.).

Steady-state models are typically 2D thermal and often 3D thermal-electric finite element models. There is no reason why the three-dimensionality of the cell design (anode studs and collector bar) and the indissociable thermal-electrical coupling should be reduced into a 2D thermal model. Today, a full 3D thermal-electric finite element half-anode model or cathode side slice model only takes about ten minutes to solve on a fast PC.

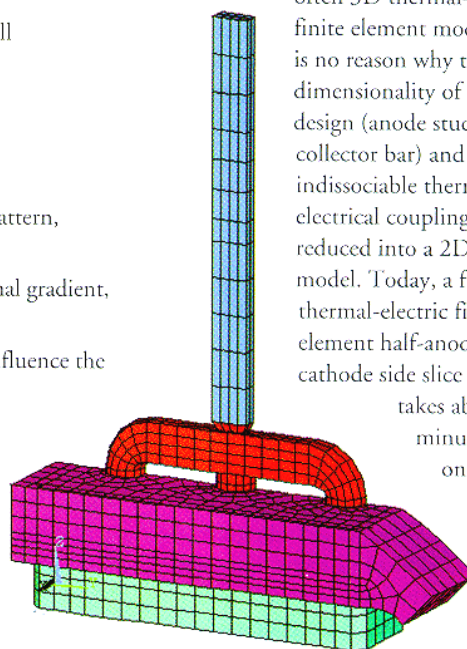


Figure 1. 3D half anode model mesh.

On the other hand, if ten minutes is reasonable to solve one steady-state solution, it is much too long to get only one time step in a dynamic model run. That is why dynamic models are still confined to 0D or 1D geometric representations typical of the second category of thermal-electric models.

### 3D steady-state thermal-electric models

The goal of 3D steady-state thermal-electric models is to determine as accurately as possible:

- The global cell heat dissipation when the cell is operating under well-established stable conditions.
- The corresponding ledge profile.
- The corresponding cathode lining drop.
- The corresponding anode drop.

My standard approach to achieve the above goal is to develop two independent models: a half-anode model, and a cathode side slice model.

When considering heat transfer by convection, the “enhanced conductivity” of bath and metal are orders of magnitude greater than those of any lining materials. So, it is a good modelling practice to consider the bath and metal to be isotherms and to remove them from the model domain.

Having done that, the only coupling between the anode and the cathode parts of the cell is located in the cell’s side channel crust. For practical reasons, I prefer to completely separate the anode part from the cathode part by cutting the crust along an adiabatic line perpendicular to the crust surface.

By having a decoupled half-anode model and cathode side slice model, we can solve them independently. But to do so, we must specify the cell operating temperature and the cell superheat as boundary conditions (i.e., as input to the model). Under such conditions, the models will calculate the global cell heat loss. If the calculated cell loss does not correspond to the separately computed cell internal heat, it means that the cell steady-state conditions do not correspond to the assumed operating cell temperature and cell superheat for that cell design.

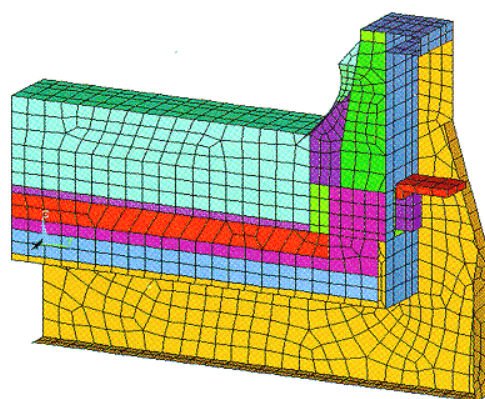


Figure 2. 3D cathode side slice model mesh.

This leaves the designer with many options for the next step in the design loop:

- Adjust the cell internal heat; e.g., reduce the anode to cathode distance (ACD).
- Change the assumed cell temperature and/or superheat.
- Change the cathode lining design or liquid height.
- Change the anode design or insulation cover thickness.
- A combination of the above.

The most important thing to remember with this approach is that the user specifies the operating temperature and superheat. The model then calculates the corresponding cell heat loss that may or may not match the cell internal heat generation. In the “classic 2D full cell slice” model approach [3], the user specifies the cell internal heat and the liquidus temperature and the model

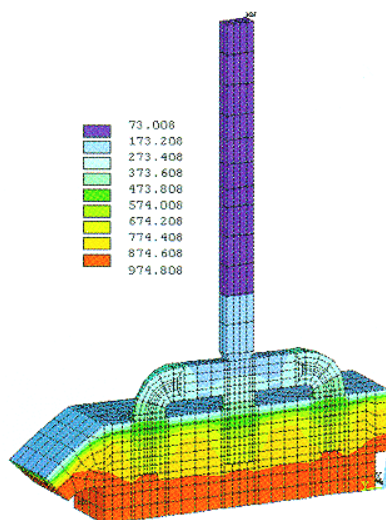


Figure 3. 3D half anode isotherms.

calculates the corresponding steady-state temperature. But, the steady-state temperature may turn out to be unreasonable. This is yet another example of the “chicken and egg” issue and should not be a source of concern for a modeler using my approach.

### Model validation issues

I cannot stress strongly enough that developing a cell model without validating it is a waste of time. Furthermore, using an unvalidated cell model can be very misleading (garbage in, garbage out). So, if you are planning to develop a mathematical model of your cell, be ready to spend the time required to validate it properly.

The best way I know to validate a 3D thermal-electric cell model is to conduct a coordinated “thermal blitz” campaign. Essentially, a “thermal blitz” consists in measuring enough heat fluxes very rapidly around a cell (in stable operating conditions) to be able to integrate them along all the external surfaces in order to obtain an experimental cell heat loss “snap shot” (see *Tables 1 and 2* on pg. 31).

At the same time, a detailed cell voltage measurement is carried out to establish the cell’s internal heat. If the cell was truly in a stable operating condition and the measurements were perfectly carried out (which is not an easy task), the measured cell heat loss should match the cell internal heat within 5 percent. If this is the case, then the measured heat balance data can be used to validate the cell model.

At the end of the thermal blitz, a series of side ledge profiles should also be measured. By imposing the average measured ledge profile on the validated cathode model, we obtain from the model solution the “experimental” heat transfer coefficients at the fluids/ledge interfaces. Those heat transfer coefficients are then used to predict the ledge profile thickness of the “retrofitted” design.

Obviously, the behavior of the cathode thermal-electric model relies heavily on the lining material properties data defined in it. Having a reliable material properties database is one of the keys of successful model development. It is especially important not to blindly use data provided by the lining material’s manufacturers, since the properties that should be used in the model are not the “new” material properties but the “used” lining material properties, after the

lining materials have been affected by sodium diffusion under cell operating conditions. That is why, typically, the final model calibration involves adjusting/degrading lining materials properties until the model globally reproduces the measured heat balance within 5 percent.

On the electric side, the cell models should also match the measured cathode lining drop and the average mid-life anode drop. This usually translates into the adjustment of the model's electrical contact resistance between cast iron and carbon.

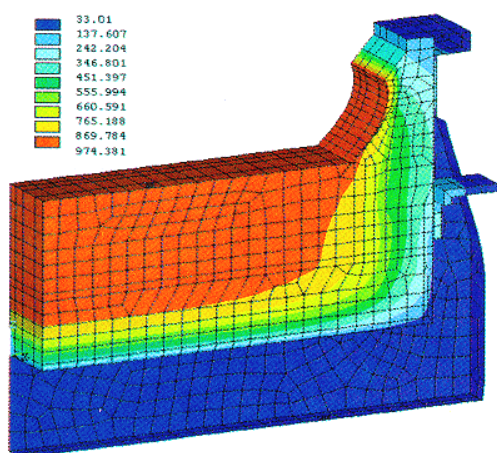


Figure 4. 3D cathode side slice isotherms.

An extra experimental measurement can be conducted to measure at least one such contact resistance to set up an instrumented anode. An instrumented anode setup consists of installing thermocouples in and around studs. Those thermocouples, normally used to measure temperature only, are used at the same time as voltage probes. With this setup, it is possible to measure the contact resistance both between the steel stud and the cast iron, and between the cast iron and the anode carbon.

Obviously, the contact resistance is one of the key parameters in both the anode and cathode models; having a good pressure contact between cast iron and carbon parts is one of the important features of a successful cell thermal-electric design. Yet, it is widely considered as being only a mechanical issue.

### Example of 3D model usage: Sensitivity analysis

In a recent presentation at the CQRDA symposium [4], I presented a 3D model application of an imaginary cell retrofit project, where I redesigned a 300 kA cell operating at 92.9 percent current efficiency (CE) and 13.7 kWh/kg into a 265 kA cell operating at 96 percent CE and 11.9 kWh/kg. Although this was a good demonstration of the power of modelling tools, the fact that I lumped 17 design changes into one design loop was not an ideal way to demonstrate what was going on in the model and how the model's results should be analyzed.

For this article, I will use a parametric sensitivity study.

The base case model that I am going to use is in the public domain; I created it for demonstration purposes, inspired by information published by VAW in the *Journal of Metals* in February 1994 [5]. I basically used all the available information and made an "educated guess" to fill the gaps.

The models that resulted from that "guessing" exercise are presented in Figures 1 and 2. By solving them, you get the results shown in Figures 3 to 6. Although fringe plots of the isotherms and equipotentials are quite useful to look at (especially

in the debugging phase of model development), the key results are located in the heat balance tables, which are produced automatically by the models.

The total heat dissipated adds up to 620 kW, of which 38 percent is dissipated by the anode panel and the remaining 62 percent is dissipated by the cathode shell, cradles and collector bars.

A third table is produced to verify that the total heat dissipated does correspond to the cell internal heat generated. At 5 cm ACD, the cell internal heat turns out to be 622 kW. Calculation of the cell internal heat is based on public-domain equations:

- Bath resistivity using Wang's equation. [6]
- Bath voltage using Haupin's equation. [7]
- Electrolysis voltage using Haupin's equation. [7]
- Equivalent voltage to make metal using Haupin's equation. [7]
- Cell current efficiency using Solli's equation. [8]

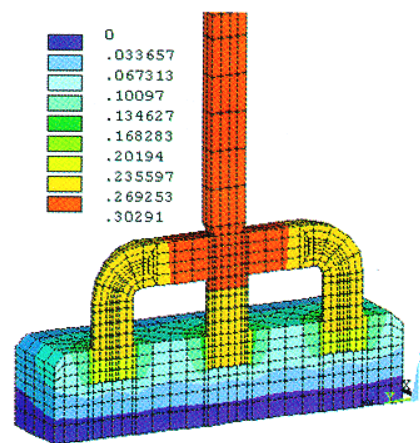


Figure 5. 3D half anode equipotentials.

## Simple 1D steady-state model for Monte Carlo analysis

One of the drawbacks of using only a 3D modelling approach in obtaining the "optimum steady-state design" is that the optimum steady-state solution could very well be on the edge of practical limits. For example, the ACD can be on the edge of the MHD stability limit, or the ledge thickness can be very thin. Yet, you don't know how sensitive that design will be to process

perturbations, which could amplify into serious operation problems.

When designing "on the edge" it is a good engineering practice to allow for a safety margin. One efficient way to estimate the size of the safety margin is to carry out a Monte Carlo analysis. In one type of Monte Carlo analysis, you may want to assess the risk involved in using mathematical models

of a given accuracy to predict a cell behavior in real life. You replace the value computed by the 3D models for key results like anode voltage drop, anode panel heat dissipation, etc. by a probability function that accounts for the model's assumed accuracy.

I personally like to use a Poisson distribution for the probability function of the model predictions because I use the model

As mentioned before, it is up to the user to ensure that the total heat dissipated matches the internal heat generation. It is critical that this is achieved before the model results can be considered to correspond to the cell steady-state conditions.

In the sensitivity test case study, the alumina cover is decreased from the 16 cm setup in the base case model to 13 cm. The first step is to rerun the anode model with that change. We can note at this point that having developed the model by using parameters, it was quite easy to change the base case model into the new configuration. The results of the reduced alumina cover indicate that the anode panel heat loss is increased from 234 kW to 247 kW.

In the second step, the cell operating temperature and the corresponding cell superheat must be decreased to reduce the cathode heat loss of 13 kW in order to maintain the total heat loss equal to the internal heat. It is up to the user to "guess" how much the cell operating temperature should be decreased. Yet, it is straightforward to make a good "educated guess" by assuming that only the heat loss through the ledge will be affected by the change of the operating temperature and that it will be proportional to the change of the cell superheat. At 20°C of cell superheat, the heat going through the ledge in the model is  $766 + 1472 = 2238 \text{ W}$ , corresponding to 209 kW for the total cathode.

From the above assumption the new cell eutectic superheat should be:

$$\Delta T = 20^\circ \times \frac{(209 - 13)}{209} = 18.75^\circ \text{C}$$

The cathode model is rerun after decreasing the cell operating temperature from 975°C to 973.75°C.

The total cathode heat loss is reduced from 385 kW to 374 kW for a new total cell heat loss dissipation of 622 kW, still in thermal equilibrium.

When we compare the ledge profile of the base case (Figure 7) vs. the ledge profile of the sensitivity test case (Figure 8), the main conclusion we can draw from the analysis is that when we remove 3 cm of alumina cover, the cell reacts by increasing the ledge thickness by an average of 0.74 cm.

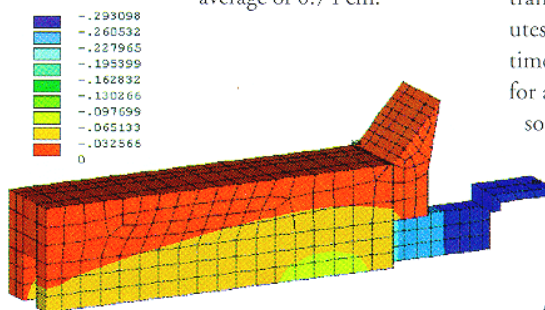


Figure 6. 3D cathode side slice equipotentials.

## The need for a simple 1D thermal model

Even with a fast PC, the above example of sensitivity study will require about 30 minutes: 10 minutes of CPU time for the anode model, 10 minutes of CPU time for the cathode model and 10 minutes for model change preparation and model results review.

This is far too long to wait for testing "what if" scenarios like the following (from the above example):

- What if we increase the ACD (instead of letting the cell cool down)?
- What if we decrease the metal level?
- What if we increase the cell amperage?

In the preliminary brainstorming phase of a retrofit project, you don't want to have to wait 30 minutes each time you want an answer to one of those "what if" questions.

If you want to study the impact of multiple changes in a Monte Carlo analysis, it is obviously out of the question to have to wait 10 minutes per calculation cycle, since we will need to conduct thousands of such calculations (see sidebar).

Additionally, the same is true for the transient analysis. You cannot wait 10 minutes per time step if you need thousands of time steps to compute the system response for a long enough period of time with a reasonably small time step.

For all those reasons, it is more convenient to develop a simple 1D thermal model of the reduction cell. The simple 1D thermal model that I have developed is based on the following assumptions:

- The heat produced in the cell can escape by four different paths (through the anode panel, through the cathode block panel, through the ledge at the bath level and through the ledge at the metal level).
- The global thermal resistance of the first two paths is constant.
- The global thermal resistance of the last two paths varies in order to maintain heat flux through the ledge proportional to the cell's superheat.

Using the above assumptions, it is possible to reduce the 3D finite element thermal model containing thousands of nodes (differential equations) into a system of four 1D equations:

$$Q_{\text{ANODE}} = R_{\text{ANODE}} * (T_{\text{OP}} - T_{\text{AIR}})$$

$$Q_{\text{CATHODE}} = R_{\text{CATHODE}} * (T_{\text{OP}} - T_{\text{AIR}})$$

predictions as the left side limit of the Poisson distribution. This way, my distribution assumes that the model prediction is the best outcome expected with a significant trailing edge for the less favorable outcomes. The distributions are based on the same Poisson distribution characterized by a mean of 1, appropriately translated, scaled and (if needed) mirrored.

At each of the 15,000 cycles, the results of the random evaluation of the four input parameters are used to solve the steady-state

temperature of the cell with the 1D model. Even with a fast 1D model, it took a few minutes of CPU time to get the results!

The output temperature distribution is neither a Poisson nor a normal distribution, although a normal distribution characterized by a mean of 975.67 °C and a standard deviation of 0.645 °C is a fair approximation. We can obviously compute from the same analysis the distributions of other indirect output variables like the cell internal heat, the cell superheat or the ledge thick-

ness at the bath and metal level.

The cell design team may not be happy with the idea that the real cell might end up having to dissipate 630 kW instead of the 620 kW predicted by the 3D models. Yet, the only alternative to building up a safety margin of 10 kW in the design is to develop and use more accurate 3D models. At least, with the help of the 1D model and the Monte Carlo analysis, the level of risk associated with using 3D models of a given accuracy can be estimated fairly well. 🌈

$$Q_{FRBATH} = h_{FRB} * A_{BATH/LEDGE} * (T_{OP} - T_{MLT})$$

$$= U_{GLB} * A_{BATH/LEDGE} * (T_{OP} - T_{AIR})$$

$$Q_{FRMETAL} = h_{FRM} * A_{METAL/LEDGE} * (T_{OP} - T_{MLT})$$

$$= U_{GLM} * A_{METAL/LEDGE} * (T_{OP} - T_{AIR})$$

$$U_{GLB} = \frac{1}{\left(\frac{1}{U_{FIXB}} + \frac{1}{U_{FRB}}\right)}$$

$$U_{GLM} = \frac{1}{\left(\frac{1}{U_{FIXM}} + \frac{1}{U_{FRM}}\right)}$$

- Where  $R_{ANODE}$  and  $R_{CATHODE}$ , the anode and cathode panels' global thermal resistance, can be determined from the 3D model results.
- Where  $A_{BATH/LEDGE}$  and  $A_{METAL/LEDGE}$ , the bath-to-ledge and metal-to-ledge interface surface areas, can be determined from the 3D model results and set proportional to the bath and metal levels, respectively.
- Where  $U_{FIXB}$  and  $U_{FIXM}$ , the fix thermal resistance behind the ledge at bath and metal level, can be determined from the 3D model results.
- Where  $U_{FRB}$  and  $U_{FRM}$ , the ledge thermal resistance at bath and metal level, are functions of the ledge thickness and thermal conductivity.

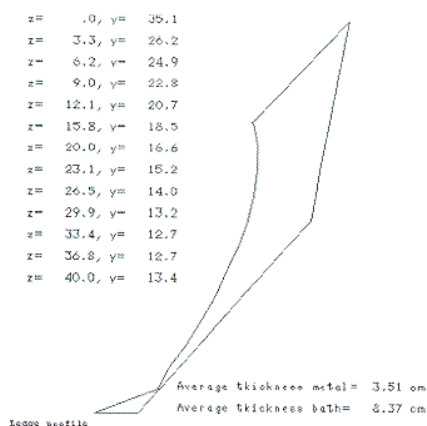


Figure 7. Ledge profile at 20°C of superheat.

For the base case model presented above, the following numbers are determined:

$$U_{GLB} = 10.53 \text{ W/m}^2\text{°C}$$

$$U_{GLM} = 15.16 \text{ W/m}^2\text{°C}$$

$$U_{FIXB} = 29.43 \text{ W/m}^2\text{°C}$$

$$U_{FIXM} = 21.70 \text{ W/m}^2\text{°C}$$

$$U_{FRB} = 16.39 \text{ W/m}^2\text{°C}$$

$$U_{FRM} = 50.26 \text{ W/m}^2\text{°C}$$

$$L_{BF} = 8.26 \text{ cm}$$

$$L_{MF} = 3.70 \text{ cm}$$

Obviously, the simple 1D model reproduces very well the 3D base case model results, since it was built from it!

## Sensitivity analysis

The real test is to compare the behavior of the simplified thermal model against the predictions of the 3D models for the same sensitivity test case presented before.

To do so, the thermal model that predicts the global heat dissipation is combined with the set of equations that predict the cell internal heat to come up with a set of equations that predict the thermal unbalance of the cell.

Having done that, it is quite straightforward to apply the Newton-Raphson convergence scheme to find the root ( $f(x)=0$ ) of the cell thermal unbalance function. That root corresponds to the cell steady-state conditions. What is very convenient is that many parameters appearing in the function can be used as the root search variable [9].

Considering that the simple 1D model cannot "exactly" match the 3D models results, the base case 1D model results are presented in *Tables 9 and 10* (see this magazine's web site, [www.analysismag.com](http://www.analysismag.com), for these and other tables not shown here) using  $T_{LIQUID}$  as the convergence variable for the Newton-Raphson scheme.

In the base case, the anode panel heat loss is specified to be 234.35 kW at 975 °C. As a result, the model converges to an operating temperature of 974.05 °C. In the selected bath chemistry, this corresponds to a cell eutectic superheat of 19.47 °C. This gives a global cell dissipation of 620.4 kW. The corresponding ledge thickness is predicted to be 8.61 cm at bath level and 4.04 cm at metal level.

In the sensitivity run, the anode panel heat loss is specified to be 247.57 kW. As a result, the model converges to an operating temperature of 972.89 °C for a decrease in the cell superheat of 1.15 °C to 18.32 °C. The corresponding ledge thickness is predicted to be 9.45 cm at bath level and 4.88 cm at metal level for an average increase of the ledge thickness of 0.84 cm.

With this tool, in a matter of minutes it is possible to obtain answers to the "what if" questions asked previously like "What if I increase the ACD instead of letting the cell cool down?" By changing the convergence variable from  $T_{LIQUID}$  to ACD and solving again, we get the answer: 5.17 cm ACD will increase the internal heat to the 633.6 kW required to maintain the new cell thermal balance.

At this point, you may wonder since the 1D model is so fast and flexible, why you should bother to build 3D models? Well, remember that you need the 3D model results to build the 1D model; also, without the 3D models you would not have known that removing 3 cm of alumina cover will increase the anode panel heat dissipation by 13 kW!

## Conclusion

In conclusion, I hope I was able to convince you of the value of using process simulation models in your work. I strongly believe that they offer tremendous opportunities to improve the productivity of all of the smelters in operation today.

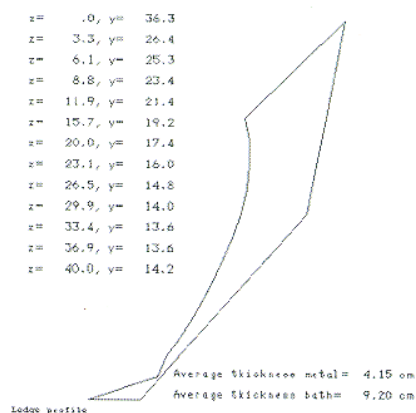


Figure 8. Ledge profile at 18.75 °C of superheat.

## References

1. *Histoire technique de la production d'aluminium*, Paul Morel, Presses Universitaire de Grenoble, 1992.
2. J.P. Antille, M. Givord, Y. Kraehenbuehl and R. Von Kaenel, "Effects of Current Increase on Aluminium Reduction Cells," *Light Metals* (1995), 315-321.
3. J.N. Bruggeman and D.J. Danka, "Two-Dimensional Thermal Modelling of the Hall-Hérault Cell," *Light Metals* (1990), 203-209.
4. M. Dupuis, "Les modèles thermiques," 2nd Symposium québécois sur le procédé d'électrolyse de l'aluminium, CQRDA (1997).
5. V.A. Kryukovski, G.A. Sirasutdinov, J. Klein and G. Peychal-Heiling, "International Cooperation and High Performance Reduction in Siberia," *JOM*, 46(2) (1994), 23-25.
6. X. Wang, R.D. Peterson and A.T. Tabereaux, "A Multiple Regression Equation for the Electrical Conductivity of Cryolitic Melts," *Light Metals* (1993), 247-255.
7. W. Haupin, "Cell Voltage Components" and "Heat Balance and Energy Consumption," CMP/PCPE Course on Aluminum Electrolysis (1992).
8. P.A. Solli, T. Haarberg, T. Eggen, E. Skybakmoen and A. Sterten, "A Laboratory Study of Current Efficiency in Cryolitic Melts," *Light Metals* (1994), 195-203.
9. M. Dupuis and I. Tabsh, "Using a Steady-State Model of an Aluminum Reduction Cell to Investigate the Impact of Design Changes," *Proceeding of the International Symposium on Light Metal*, CIM (1996), 417-429.

Marc Dupuis is a Professional Engineer, working with GeniSim, Inc. in Quebec. He can be reached at [marc.dupuis@genisim.qc.ca](mailto:marc.dupuis@genisim.qc.ca).

## HEAT FLUX MEASUREMENTS FOR CELL HEAT BALANCE

date: 17-Aug-97  
cell: "VAW" 300

slice no: A2

### Shell Wall

Description	Flux	Temp
Wall above bath level	2000	150
Wall bath level	5500	230
Wall metal level	7500	250
Wall block level above bar	6000	235
Left collector bar	3000	190
Right collector bar	3000	190
Wall collector bar level	1500	90
Wall brick level	1000	60
Floor near centerline	500	50
Floor at quarter point	500	50
Floor near corner	500	50

### Cradle Web

Wall above bath level	1000	100
Wall bath level	2165	130
Wall metal level	2660	140
Wall block level above bar	955	125
Wall collector bar level	400	60
Wall brick level	155	50
Floor extension	0	0
In the corner	100	35
Wall extension wide section	0	0
Wall extension narrow section	0	0
Floor near centerline	100	35
Floor at quarter point	100	35
Floor near corner	100	35

### Cradle Flange

Wall above bath level	500	65
Wall bath level	1085	80
Wall metal level	1330	90
Wall block level above bar	475	40
Wall collector bar level	200	35
Wall brick level	50	30
Under the floor	50	30

Table 1. Thermal blitz input data sheet.

## HEAT BALANCE RESULTS

date: 17-Aug-97

Cell: "VAW" 300

Cathode Heat Losses	W / m2	kW	%
Shell side wall above bath level	2000	11.52	1.86
Shell side wall opposite to bath	5500	31.68	5.11
Shell side wall opposite to metal	7500	43.20	6.97
Shell side wall opposite to block above bar	6000	48.38	7.80
Shell side wall opposite to block between bars	1500	6.48	1.05
Collector bars to air	3000	17.28	2.79
Collector bars to flexible	60	9.68	
Shell side wall opposite to brick	1000	11.52	1.86
Shell floor close to corner	500	12.54	2.02
Shell floor quarter point region	500	10.44	1.68
Shell floor centerline region	500	8.34	1.34
Cradle above bath level	889	6.08	0.98
Cradle opposite to bath	1925	13.17	2.12
Cradle opposite to metal	2364	16.17	2.61
Cradle opposite to block above bar	848	8.12	1.31
Cradle opposite to block between bars	356	2.43	0.39
Cradle opposite to brick	132	1.80	0.29
Cradle corner	52	1.52	0.25
Cradle below floor close to corner	100	2.76	0.44
Cradle below floor quarter point region	100	2.76	0.44
Cradle below floor centerline region	100	2.76	0.44
Shell end wall opposite to metal	1500	2.61	0.42
Shell end wall opposite to block above bar	3000	7.31	1.18
Shell end wall opposite to block below top of bar	4000	6.96	1.12
Shell end wall opposite to brick	3000	10.44	1.68
Shell coverplate in the ends	500	1.52	0.25
Shell horizontal strip in the ends	1184	18.00	2.90
Shell vertical stiffeners in the ends	898	5.52	0.89
Shell horizontal stiffeners in the ends	100	0.45	0.07
<b>Total for the cathode part</b>		<b>371.76</b>	<b>59.95</b>
<b>Anode Heat Losses</b>			
Crust in side channels	1700	21.48	3.46
Crust above anodes	1800	81.91	13.21
Crust in center channel	1750	3.60	0.58
Studs	4000	27.14	4.38
Yoke	3640	83.87	13.53
Aluminum rod	822	30.31	4.89
<b>Total for the anode part</b>		<b>248.3</b>	<b>40.05</b>
<b>Total for the cell</b>		<b>620.1</b>	<b>100.00</b>

Table 2. Thermal blitz results.