Graphing Elementary Riemann Surfaces

Robert M. Corless and David J. Jeffrey
Department of Applied Mathematics
University of Western Ontario
London, Ontario, CANADA N6A 5B7
Rob.Corless@uwo.ca, David.Jeffrey@uwo.ca

Published in SIGSAM Bulletin, Vol 32(1), no. 123, March 1998, pp 11–17

1 It's so simple!

This paper discusses one of the prettiest pieces of elementary mathematics or computer algebra, that we have ever had the pleasure to learn. The tricks that we discuss here are certainly "well-known" (that is, in the literature), but we didn't know them until recently, and none of our immediate colleagues knew them either. Therefore we believe that it is useful to publicize them further. We hope that you find these ideas as pleasant and useful as we do.

We show how to use a computer algebra system (or even a purely numerical graphing package) to graph the Riemann surfaces of various elementary functions. We first noticed the technique in Cleve Moler's MATLAB programs `cplxroot`, `cplxgrid`, and `cplxmap`, which are part of the MATLAB 5.1 DEMO package (in plots of complex functions). The command `type cplxroot` in MATLAB shows you the following.

```
function cplxroot(n,m)
%CPLXROOT Riemann  surface for the n-th root.
%   CPLXROOT(n) renders the Riemann surface for the
%                n-th root.
%   CPLXROOT, by itself, renders the Riemann surface
%              for the cube root.
%   CPLXROOT(n,m) uses an m-by-m grid.
%              Default m = 20.

%   C. B. Moler, 8-17-89, 7-20-91.
%   Copyright (c) 1984-97 by The MathWorks, Inc.
%   $Revision: 5.2 $  $Date: 1997/04/08 05:31:36 $

% Use polar coordinates, (r,theta).
% Cover the unit disc n times.

if nargin < 1, n = 3; end
if nargin < 2, m = 20; end
r = (0:m)'/m;
theta = pi*(-n*m:n*m)/m;
z = r * exp(i*theta);
s = r.^(1/n) * exp(i*theta/n);

surf(real(z),imag(z),real(s),imag(s));
```

This piece of code contains all the ideas that we shall explore in this paper, namely

1. if there is a 1–1 correspondence between the 3-d plot of $(x, y, \Im f(x+iy))$ and the Riemann surface of f, then we can plot the Riemann surface with little intellectual effort; alternatively, we can exploit any correspondence between $(x, y, \Re f(x+iy))$ and the Riemann surface;

2. the plots can be performed most easily with a parametric representation of the surface;

3. the 3-d plot can be coloured with the variable not used in the plotting (that is, the other one of $\Re f(x+iy)$ and $\Im f(x+iy)$) to give a true 4-dimensional plot of the complex-valued function.

So, if Cleve's piece of code contains all these ideas, why should we write this paper, and why should you read it? The main reason is that there is something to say about the 1–1 correspondence business, and this is not discussed anywhere that we have seen. Without discussing this, it is not obvious (at least to us) that the code really does produce a graph of the Riemann surface: *a priori*, why should a complex plot of the function give you the Riemann surface? The notions are not the same!

Incidentally, all the elementary complex variables texts that we have seen define Riemann surfaces only by example, in the usual cut-and-paste fashion. We believe the ideas of this paper can be used to flesh out those examples.

After the publication of the last issue of the Bulletin, where it was mentioned that this present article was coming, Michael Trott sent us a copy of his paper [6], and the Mathematica notebook that it was generated from. We recommend his paper to you: he produces many beautiful graphs of Riemann surfaces, using many variants of these ideas, including numerical computation. More articles by Trott on the subject are forthcoming.

However, we believe that this present article has something to contribute: namely, the emphasis on the 1–1 correspondence proof necessary to know that you are really looking at a faithful representation of a Riemann surface.

2 The logarithm

We begin with the Riemann surface for the logarithm function, as the simplest possible example. We will plot $w = \ln z$, where $w = u + iv$ and $z = x + iy$. If we plot the 3-d surface (x, y, u) we do *not* get a representation of a Riemann surface.

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc(exp(w));
```

$$z := e^u \cos(v) + I\,e^u \sin(v)$$

```
> x := evalc(Re(z)):
> y := evalc(Re(z)):


> plot3d([x,y,u],
        u=-6..1,
        v=0..10*Pi,
        orientation=[-22,48],
        labels=["x ","y ","u "],
        view=[-1..1,-1..1,-4..1],
        grid = [50,50],
        style=PATCHNOGRID,
        axes=FRAME,
        colour=v);
```

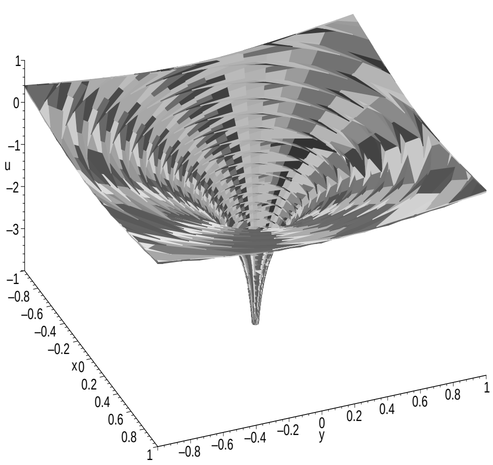

Figure 1: Not the Riemann surface for $\ln z$.

Instead, we get the single funnel seen in Figure 1; moreover, we have asked Maple to colour the surface with the value of v, but Maple's plotting algorithms get confused and we see "stripes". The plot in Figure 1 is somewhat different from the Maple session version, but still bad. The stripes probably occur because although the real part $u = \ln(x^2 + y^2)/2$ is always the same, no matter what v is, rounding errors may cause the plotting software to get confused. In any event, this shows that plotting (x, y, u) does *not* give us the Riemann surface for the logarithm function. So why should plotting (x, y, v) instead do so?

Carathéodory [1] defines the Riemann surface of logarithm by considering the map $w \to \exp w =: z$ and piecing the sheets together. We do the same thing by plotting $z = x + iy$ computed from an array of (u, v) values where $w = u + iv$. But, since this is very similar to what failed earlier, we shall have to explain its success.

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc(exp(w));
```

$$z := e^u \cos(v) + I\,e^u \sin(v)$$

```
> x := evalc(Re(z)):
> y := evalc(Re(z)):


> plot3d([x,y,v],
        u=-6..1,
        v=-3*Pi..3*Pi,
        orientation=[-56,72],
        labels=["x","y","v "],
        view=[-1..1,-1..1,-3*Pi..3*Pi],
        grid = [50,50],
        style=PATCHNOGRID,
        axes=FRAME,
        colour=u);
```

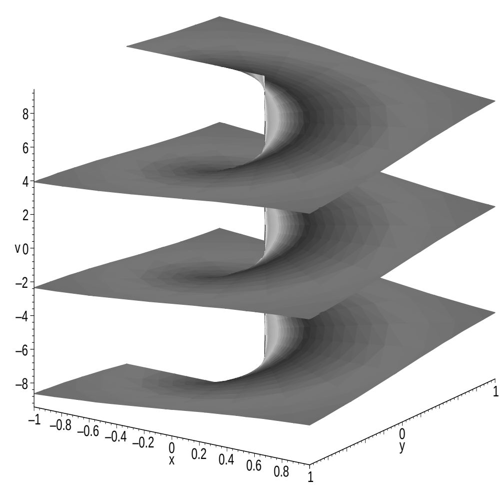

Figure 2: The Riemann surface of $\ln z$.

There are three subtleties in this short piece of code: the first is that we have changed to Cartesian coordinates, so we need not do the "θ/n trick" that Cleve uses in his MATLAB code; the second is that we write the equation backwards, putting $z = \exp(w) = \exp(u + iv) = \exp(u)(\cos v + i \sin v)$, which gives us a parametric representation for the surface (u and v are our parameters); and the third is the real trick.

The real trick is to prove that, given (x, y, v), we can solve for u uniquely. If we can do this, then the picture is in 1–1 correspondence with the Riemann surface for the logarithm. In this case, this is trivial: we have that $u = \ln(x^2 + y^2)/2$, thus uniquely determining the logarithm given v. Because each point on this surface is associated with one and only one point on the Riemann surface for the logarithm, this surface *is* a representation of the Riemann surface.

This was too easy: let's look at a more interesting example.

3 The Lambert W function

Our main interest here is plotting the Riemann surface for the Lambert W function [2, 3, 4]. The first attempts at this were done by hand some years ago, following the 'piece together the cut planes' approach commonly discussed in textbooks. A Maple program to graph the sheets three-dimensionally produced moderately pleasing results (an Axiom version was better because, at the time, Axiom's graphics were considerably superior to Maple's). The following few lines of Maple code make all that work redundant (though it is gratifying to see that the pictures are qualitatively the same).

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc(w*exp(w));
```

$$z := u\,e^{u}\cos(v) - v\,e^{u}\sin(v) + I\,(v\,e^{u}\cos(v) + u\,e^{u}\sin(v))$$

```
> x := evalc(Re(z)):
> y := evalc(Im(z)):
> plot3d([x,y,v], u=-6..1,
        v=-5..5, axes=FRAME,
        orientation=[-110,73],
        labels=["x","y","v "],
        style=PATCHNOGRID,
        colour=u,
        view=[-1..1,-1..1,-5..5],
        grid=[50,50]);
```

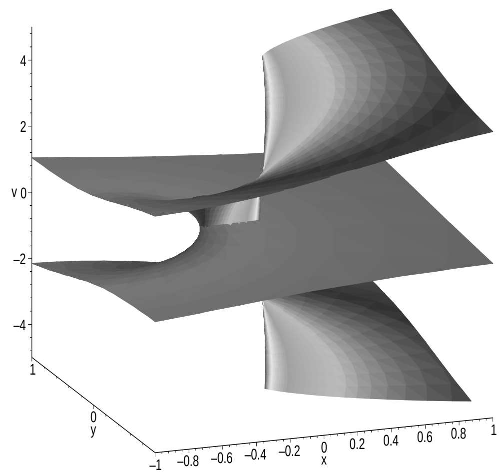

Figure 3: The Riemann surface of $W(z)$.

3.1 1–1 correspondence proof

Given x, y, and v, we have to solve for u. We have

$$(u + iv)e^{u+iv} = x + iy\;, \tag{1}$$

which gives

$$\begin{aligned} ue^u + ive^u &= (x + iy)e^{-iv} \\ &= (x + iy)(\cos v - i \sin v)\;. \end{aligned} \tag{2}$$

Therefore

$$\begin{aligned} ue^u + ive^u &= (x \cos v + y \sin v) \\ &\quad + i(y \cos v - x \sin v)\;. \end{aligned} \tag{3}$$

If $v \neq 0$, and moreover $y \cos v - x \sin v \neq 0$, then by dividing the real part by the imaginary part we have the following equation defining u in terms of x, y, and v:

$$u = \frac{v(x \cos v + y \sin v)}{y \cos v - x \sin v}\;. \tag{4}$$

Moreover this solution is unique. Investigation of the exceptional conditions $v = 0$ or $y \cos v - x \sin v = 0$ leads to $u \exp u = x$, which has two solutions if and only if $-1/e \leq x < 0$, in the case $v = 0$, and to the singular condition $u = -\infty$ and $x = y = 0$.

This is precisely what we observe in the graph: two sheets intersect only if $-1/e \leq x < 0$ (note the colours are different and hence the corresponding sheets on the Riemann surface do not "really" intersect), and all sheets except the central one which contains $v = 0$ have a singularity at the origin. This is as good a representation of the Riemann surface for the Lambert W function as can be produced in three dimensions.

Remark. The above static representation of the Riemann surface is really nowhere near as intelligible as the live Maple plot (OpenGL), which can be rotated to give a good sense of what it is really like.

4 The Arcsin function

Thus emboldened, we look at the Riemann surface for Arcsin. One new twist is that we plot (x, y, u), because the sin function has real periods.

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc(sin(w));
```

$$z := \sin(u)\cosh(v) + I\cos(u)\sinh(v)$$

```
> x := evalc(Re(z)):
> y := evalc(Re(z)):
> B := 3:

> plot3d([x,y,u],
        u=-Pi..Pi,
        v=-B..B,
        orientation=[-127,52],
        labels=["x","y","u "],
        view=[-B..B,-B..B,-Pi..Pi],
        grid=[50,50],
        axes=FRAME,
        color=v,
        style=PATCHNOGRID);
```

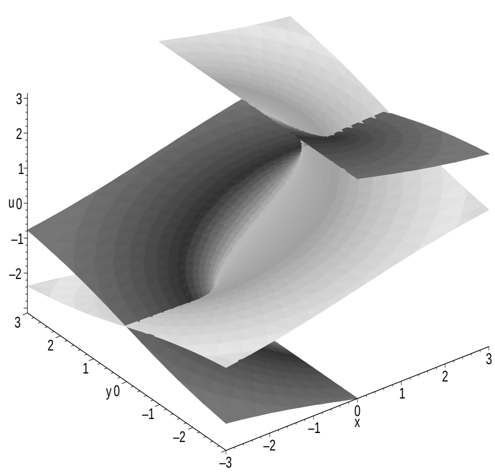

Figure 4: The Riemann surface for $w = \arcsin z$.

4.1 1–1 correspondence proof

This is trivial here: $\sinh v \cos u = y$ implies that $v = \sinh^{-1}(y/\cos u)$, uniquely whenever $\cos u \neq 0$. Again, this is what we observe in the graph: the only apparent intersections occur when $u = (2k+1)\pi/2$ for some k, and $y = 0$. The other equation ($x = \sin v \cosh u$) then gives two possible values of u, from the two real branches of $\cosh^{-1}$.

5 The Arctan function

The surface for Arctan is a bit more complicated, and we see that Maple introduces spurious flat sheets across singularities in the surface (much as the old 2-d plots would "connect the dots" across a singularity, one expects). To combat this, we plot using the POINT style. We see twin helices, wrapping around the logarithmic singularities at $\pm i$ in opposite directions.

```
> w := u +I*v;
```

$$w := u + I\,v$$

```
> z := evalc(tan(w));
```

$$z := \frac{\sin(u)\cos(u)}{\cos(u)^2 + \sinh(v)^2} + \frac{I\sinh(v)\cosh(v)}{\cos(u)^2 + \sinh(v)^2}$$

```
> x := evalc(Re(z)):
> y := evalc(Im(z)):

> plot3d([x,y,u],
        u=-Pi..Pi,
        v=-2..2,
        view=[-4..4, -4..4, -Pi..Pi],
        grid=[200,200],
        style=POINT,
        axes=FRAME,
        labels=["x","y","u "],
        orientation=[50,86],
        color=v);
```

5.1 1–1 correspondence proof

If we divide y by x (provided $x \neq 0$) we get $\sinh(2v)/\sin(2u) = y/x$, so $v = \sinh^{-1}(y\sin(2u)/x)/2$, uniquely. The only intersections possible would have $x = 0$ and $\sin 2u = 0$, and the y-equation then would give $y = \sinh v \cosh v/(t^2 + \cosh^2 v)$ which has a unique solution no matter whether $t = \cos u$ is zero or one, its only possibilities. Thus, this plot is exactly the Riemann surface for arctan.

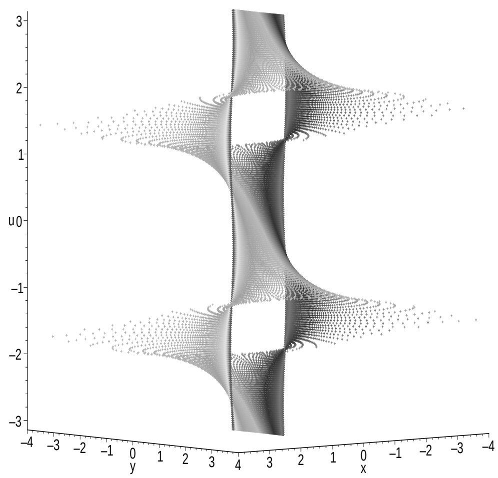

Figure 5: The Riemann surface for $w = \text{Arctan}\, z$.

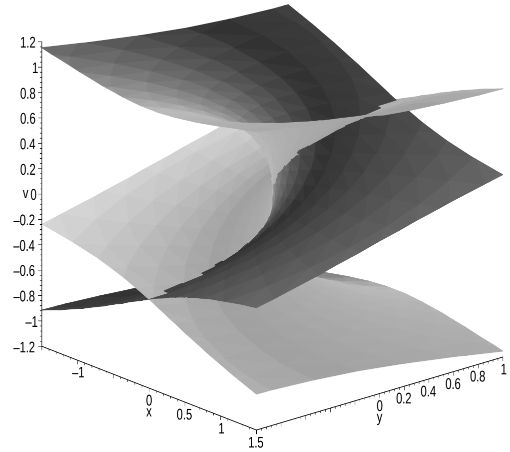

Figure 6: The Riemann surface for $w = z^{1/3}$.

6 n th roots

Now we look at the Riemann surface for nth roots; more specifically we will look at $w = z^{1/3}$.

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc(w^3);
```

$$z := u^3 - 3\,u\,v^2 + I\,(3\,u^2\,v - v^3)$$

```
> x := evalc(Re(z)):
> y := evalc(Re(z)):


> plot3d([x,y,v],
        u=-2..2,
        v=-2..2,
        view=[-1.5..1.5,-1..1,-1.2..1.2],
        grid=[50,50],
        orientation=[-41,70],
        style=PATCHNOGRID,
        color=u,
       labels=["x","y","v"]);
```

6.1 Proof of 1–1 correspondence

Considering first the case $v = 0$, we have $u = x^{1/3}$ uniquely (using the real cube root). On the other hand we have that $u^2 = (y + v^3)/(3v)$ if $v \neq 0$, and since $x = u(u^2 - 3v^2)$ we have

$$x = u\left(\frac{y + v^3}{3v}\right).$$

If $x \neq 0$, then we must have

$$u = \frac{3xv}{y + v^3}$$

uniquely. This leaves the case $x = 0$. If $x = 0$, then either $u = 0$, which gives $y = -v^3$ and in any event is unique since $v \neq 0$ by assumption, or else $y = 8v^3$ which leads to $u^2 = 3v^2$ which has two possible solutions. This agrees with the graph; the only apparent intersections we see have $x = 0$ and as usual the colours are different on each intersecting sheet. Thus this graph is as good a representation of the Riemann surface as can be found.

7 More on fractional powers

We look at a simple example, $w = z^{2/3}$. It turns out that the Cartesian coordinate approach used heretofore doesn't work very well, and indeed Maple will force polar coordinates on us (in disguise) anyway. This is shown by the code fragment below, which gives an incomplete plot (not shown) because the $\arctan(v, u)$ which shows up uses the principal branch.

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc(w^p);
```

$$z := e^{(1/2\,p\ln(u^2+v^2))}\cos(p\arctan(v,\,u)) + I\,e^{(1/2\,p\ln(u^2+v^2))}\sin(p\arctan(v,\,u))$$

```
>  x := evalc(Re(z)):
>  y := evalc(Im(z)):

> plot3d(subs(p=3/2,[x,y,v]),
        u=-2..2,
        v=-2*Pi..2*Pi,
        grid=[50,50],
        colour=u,
        axes=FRAME,
        orientation=[150,40],
        style=PATCHNOGRID,
        labels=["x","y","v "]);
```

The plot in Figure 7 uses polar coordinates explicitly, giving us a good parameterization of the Riemann surface. It turns out that we can use either the real part of w or the imaginary part of w as the height.

```
>  w := r*cos(theta) + I*r*sin(theta);
```

$$w := r\cos(\theta) + I\,r\sin(\theta)$$

```
>  x := r^p*cos(p*theta);
```

$$x := r^p\cos(p\,\theta)$$

```
>  y := r^p*sin(p*theta);
```

$$y := r^p\sin(p\,\theta)$$

```
> plot3d(subs(p=3/2,[x,y,evalc(Im(w))]),
        r=0..1.5, theta=-2*Pi..2*Pi,
        grid=[50,50],
        style=PATCHNOGRID,
        view=[-1..1,-1..1,-2..2],
        colour=evalc(Re(w)),
        orientation=[36,66],
        labels=["x","y","v "],
        axes=FRAME);
```

7.1 Proof of 1–1 correspondence

Here we do not wish to identify r and θ uniquely, but rather, given x, y, and either of $r\sin\theta$ or $r\cos\theta$, find $r\cos\theta$ or $r\sin\theta$ respectively. Thus the equations that we have to solve are

$$r^p\cos p\theta = x \qquad (5)$$
$$r^p\sin p\theta = y\,. \qquad (6)$$

This gives immediately that $r = (x^2+y^2)^{1/(2p)}$, the unique positive root. Then knowing $r\sin\theta$ gives us $r\cos\theta = \pm r\sqrt{1-\sin^2\theta}$, and all that remains to be determined is the sign (similarly if $r\cos\theta$ is known and $r\sin\theta$ desired). This shows that at most two sheets can intersect in any one place.

Let us now look at the case $p = 3/2$, which corresponds to $w = z^{2/3}$. Thus we know $\sin 3\theta/2$ and $\cos 3\theta/2$

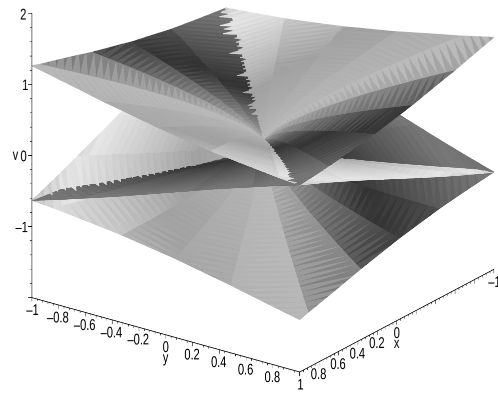

Figure 7: The Riemann surface for $w = z^{2/3}$.

and wish to identify $\cos\theta$ given $\sin\theta$. The following trig identities give us what we want.

$$\begin{aligned}
\sin\tfrac{3}{2}\theta &= \sin\theta\cos\tfrac{1}{2}\theta + \cos\theta\sin\tfrac{1}{2}\theta \\
\cos\tfrac{3}{2}\theta &= \cos 2\theta\cos\tfrac{1}{2}\theta + \sin 2\theta\sin\tfrac{1}{2}\theta \\
&= (1-2\sin^2\theta)\cos\tfrac{1}{2}\theta \\
&\quad +2\sin\theta\cos\theta\sin\tfrac{1}{2}\theta\,,
\end{aligned}$$

which gives

$$\cos\tfrac{3}{2}\theta - 2\sin\theta\sin\tfrac{3}{2}\theta = (1-4\sin^2\theta)\cos\tfrac{1}{2}\theta\,.$$

This determines $\cos(\theta/2)$ uniquely unless $\sin\theta = \pm 1/2$; and once we know $\cos(\theta/2)$ we can find $\cos\theta = 2\cos^2(\theta/2) - 1$ uniquely. After this calculation, you might be puzzled because the graph clearly shows intersections at 45 degrees, not at 30 degrees. However, everything is correct: we plot $x = r^p\cos p\theta$, $y = r^p\sin p\theta$, which means if $p = 3/2$ that the intersections occur at $3/2 * 30 = 45$ degrees. Thus our calculation is in agreement with the graph, and the graph really shows the Riemann surface for the $2/3$ power of z.

8 A final example: the "Dilbert Λ function"

In [5], Bill Gosper suggests that instead of using the Lambert W function, we could instead use the function $\Lambda(z)$ which satisfies

$$\Lambda e^{\Lambda^2} = z\,. \qquad (7)$$

The chief advantage of this function is that it is single-real-valued for real z. See [5] for details. Here we shall look at its Riemann surface.

```
> w := u + I*v;
```

$$w := u + I\,v$$

```
> z := evalc( w*exp(w^2) ) ;
```

$$z := u\,\%1\cos(2\,u\,v) - v\,\%1\sin(2\,u\,v) + I\,(v\,\%1\cos(2\,u\,v) + u\,\%1\sin(2\,u\,v))$$

$$\%1 := e^{(u^2-v^2)}$$

```
> x := evalc(Re(z)):
> y := evalc(Re(z)):

> B := 2:

> plot3d([x,y,v],
        u=-2..2,  v=-B..B,
        axes=FRAME,
        grid=[40,40],
        view=[-1.5..1.5,-1..1,-B..B],
        style=PATCHNOGRID,
        orientation=[151,77],
        labels=["x","y","v "],
        color=u);
```

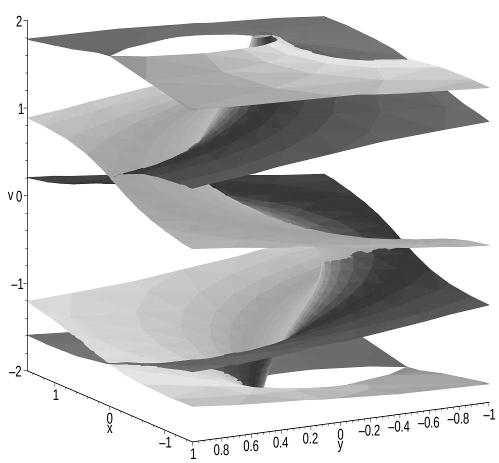

Figure 8: The Riemann surface for the function $\Lambda(z)$ which solves $\lambda\exp(\lambda^2) = z$.

We leave the proof of 1–1 correspondence to the reader. This Riemann surface plot shows that it would be natural to index the sheets for the domain of $\Lambda(z)$ so that $\Lambda_0(z)$ has branch cuts along the imaginary axis from $\pm i/\sqrt{2e}$ out to infinity, which means that the real axis does not cross a cut; $\Lambda_1(z)$ and $\Lambda_{-1}(z)$ are then symmetric, with cuts to $i/\sqrt{2e}$ and 0 from $\pm i\infty$ along the imaginary axes and to 0 and $-i/\sqrt{2e}$ from $\mp i\infty$ respectively; finally the cuts on the domains of $\Lambda_k(z)$ for $|k| \geq 2$ go from $\pm i\infty$ to 0. Counter-clockwise continuity is always possible.

Acknowledgements

Cleve Moler took time out to answer several questions on the phone, for which we are grateful. Michael Trott kindly provided a copy of the Mathematica Journal containing his paper.

References

[1] C. Carathéodory. *Theory of Functions of a Complex Variable.* Chelsea, 1964.

[2] Robert M. Corless, Gaston H. Gonnet, David E. G. Hare, David J. Jeffrey, and Donald E. Knuth. "On the Lambert W function". *Advances in Computational Mathematics*, 5:329–359, 1996.

[3] Robert M. Corless and David J. Jeffrey. "The unwinding number". SIGSAM *Bulletin*, 30(2):28–35, June 1996.

[4] Robert M. Corless, David J. Jeffrey, and Donald E. Knuth. A sequence of series for the Lambert W function. In Wolfgang W. Küchlin, editor, *Proceedings of ISSAC '97, Maui*, pages 197–204, 1997.

[5] Bill Gosper. The solutions of ye^{y^2} and ye^y. SIGSAM *Bulletin, Communications in Computer Algebra*, 31(2):8–10, 1998.

[6] Michal Trott. Visualization of Riemann surfaces of algebraic functions. *Mathematica in Education and Research*, 6(4):15–36, 1997.