CHEBYSHEV INTERPOLATION FOR FUNCTIONS WITH ENDPOINT SINGULARITIES VIA EXPONENTIAL AND DOUBLE-EXPONENTIAL TRANSFORMS

MARK RICHARDSON*

Abstract. We present five theorems concerning the asymptotic convergence rates of Chebyshev interpolation applied to functions transplanted to either a semi-infinite or an infinite interval under exponential or double-exponential transformations. This strategy is useful for approximating and computing with functions that are analytic apart from endpoint singularities. The use of Chebyshev polynomials instead of the more commonly used cardinal sinc or Fourier interpolants is important because it enables one to apply maps to semi-infinite intervals for functions which have only a single endpoint singularity. In such cases, this leads to significantly improved convergence rates.

Key words. endpoint singularity, conformal map, exponential transform, double-exponential transform, semi-infinite interval, infinite interval, Chebyshev interpolation, sinc function, Chebfun

AMS subject classifications. 41A05, 41A17, 41A25, 65D05

1. Introduction. Polynomial interpolation in Chebyshev points is known to be a robust and rapidly convergent technique for approximating smooth functions. The connection between Chebyshev expansions and Fourier cosine series provides a link to several fast evaluation and summation algorithms based upon the FFT, and for this reason, Chebyshev interpolants are often favoured in practice over other spectral expansion bases. An example of this is Chebfun, which uses Chebyshev interpolants to routinely compute integrals, derivatives, roots, and solutions of differential equations with high speed and accuracy close to machine-precision [24].

A failing of Chebyshev interpolants, however, is their inability to cope efficiently with functions that are singular, that is, functions which are not analytic at some point in the domain of interest. Some strategies do exist for dealing with simple instances of singularity. For example, in Chebfun, the function $|x|$ on $[-1,1]$ is not approximated by a global polynomial, but rather by two individual linear pieces. This approach is satisfactory in many cases, however, the difficulty is more serious when a function is encountered that cannot be divided into analytic pieces, such as $\sqrt{1-x}$.

A recent publication by Richardson and Trefethen [14] built upon the work of Stenger and others [2, 10, 16, 17] to construct an adaptive Chebfun-like system using mapped cardinal sinc interpolants to approximate such functions. At the heart of the method is a map which transplants functions f defined on $x \in [0,1]$ to functions $\mathcal{F}$ on $s \in (-\infty, \infty)$, mapping in the process any endpoint singularities to infinity. If the transplanted function $\mathcal{F}$ is analytic in some infinite strip about the real axis, then the sinc interpolant provides a rapidly convergent approximation. Indeed, it is well known that under reasonable hypotheses, this technique converges at a rate $C^{-\sqrt{n}}$ as $n \to \infty$, where n is the number of sample points.

Almost all of the literature concerning the approximation of functions with endpoint singularities via variable transformation has so far concentrated on the use of maps to $(-\infty, \infty)$. This implicitly treats the situation where the problem function has possible endpoint singularities at both ends of the domain. We shall argue in this paper that in many practical situations, it is preferable to instead use a map to a semi-infinite interval such as $(-\infty, 0]$. To motivate this point, consider approximating

*Oxford University Mathematical Institute, 24–29 St Giles', Oxford OX1 3LB, UK (mark.richardson@maths.ox.ac.uk, http://people.maths.ox.ac.uk/richardsonm/).

the function $f(x) = \sqrt{x} - x$ on $[0, 1]$. Since the only singularity is at $x = 0$, the use of a transformation to an infinite interval is wasteful. It is in fact much better is to use a transformation from $[0, 1]$ to $(-\infty, 0]$; then no "waste" is introduced on the right-hand side, since the point $x = 1$ is not mapped to infinity. However, since these transformations effect exponential decay in only one direction, sinc functions (and for that matter, Fourier series) cease to be a viable approximation basis. Chebyshev interpolants, on the other hand, will cope perfectly well with this decay asymmetry.

Related to this is a conclusion of [4, 14] that the use of sinc functions is not in itself a necessary component for the success of variable-transformation methods. With suitable modifications, other basis sets such as Fourier series and Chebyshev polynomials can and should be considered. In Sincfun, the problem functions $f(x)$ on $[0, 1]$ were assumed to take zero endpoint values (general functions were reduced to this case by subtracting a linear function). This resulted in exponential decay of the function $\mathcal{F}(s)$ on $(-\infty, \infty)$. Then, given a fixed precision ε, it was possible to determine in advance an interval $[-L, L]$, $L > 0$, such that $|\mathcal{F}(s)| < \varepsilon$ for all $|s| > L$. It is a simple task to approximate $\mathcal{F}$ on this finite interval, and any reasonable interpolation scheme will result in exponentially convergent approximations. Sinc functions were used in [14], but Fourier series could equally well have been used, since the exponential decay of $\mathcal{F}$ makes it a good approximation to regard it as periodic on $[-L, L]$. In particular, the associated Gibbs overshoot is negligible for large L.

Chebyshev polynomials could have been used too, though their use would incur the usual asymptotic loss of a factor of $\pi/2$ in the number of sample points [7, 26]. Crucially, for both Fourier series and Chebyshev polynomials, the FFT enables fast summation and evaluation; this is not the case for the sinc basis.

There are two classes of map that we shall discuss in this paper: *exponential transforms* and *double-exponential transforms*. In each case, we consider maps from the problem interval $[0, 1]$ to both $(-\infty, 0]$ and $(-\infty, \infty)$, giving a total of four transformations. As we shall see, exponential maps result in transplanted functions which decay exponentially, while double-exponential maps result in a transplanted functions which decay double-exponentially. In general, double-exponential transformations result in much more rapidly convergent approximations than exponential transforms, though there are caveats to this; see e.g. [21, 22]. Moreover, due to a result by Sugihara [18], at least in the infinite interval setting, we know that double-exponential formulas are optimal in the sense that the convergence rate cannot be improved further. There is therefore little point in considering, for example, triple-exponential transformations.

Table 1 summarises four of the five convergence results that will be proven. The fifth result is a special case of the situation involving the exponential map to $(-\infty, 0]$ for which C^{-n} geometric convergence may be achieved in certain circumstances. Naturally, it will be necessary to make certain assumptions in order to justify these results and the associated function classes will be defined in due course.

	$(-\infty,\ 0\ ]$	$(-\infty, \infty)$
exponential	$C^{-n^{2/3}}$	$C^{-\sqrt{n}}$
double-exponential	$C^{-n/\log n}$	$C^{-n/\log n}$

TABLE 1.1
Summary of convergence rates for Chebyshev interpolation applied to functions on $[0, 1]$ *transplanted to either* $(-\infty,\ 0\ ]$ *or* $(-\infty, \infty)$. *The constant* $C > 1$ *is different in each case.*

The original motivation for this paper is an early work of Boyd [3] in which Chebyshev interpolation combined with a strategy termed *domain truncation* was used to approximate functions defined on infinite and semi-infinite intervals. Drawing upon the steepest descent techniques developed by Elliott [6], estimates of asymptotic Chebyshev coefficients for several representative example functions were determined. In this paper, we link this domain truncation technique to the problem of approximating functions on the problem domain $[0, 1]$ by introduction of a variable transformation and provide general theorems describing the asymptotic convergence behaviour of this strategy for large classes of functions. In [3], an emphasis was placed upon the difference between convergence rates for entire functions and for functions that are merely analytic. We demonstrate that it is not enough for a function to be entire; to achieve the improved C^{-n} convergence rate, it must grow at most algebraically as $|x| \to \infty$ in the complex plane. Concerning double-exponential maps to a semi-infinite interval, the only other investigation that we are aware of is by Hamada in [8]. That article proposes, in essence, the method which we shall analyse in Theorem 3.3, but for an application related to solving differential equations. We are not aware of any previous study which discusses convergence rates for this transformation.

On the infinite interval, we have two results that are familiar from what one might call "classic sinc", and double-exponential quadrature theory. The $C^{-\sqrt{n}}$ convergence rate for exponential maps was proven for cardinal sinc functions in [16, 17] and for Fourier series in [25]. The last of these papers concerned the problem of using Fourier spectral methods to approximate the derivatives of an exponentially decaying function on the real line. Though the connection to our problem of approximating functions on $[0, 1]$ is not explicit, the arguments presented there lead us in the right direction toward proving the equivalent results in the Chebyshev case. The $C^{-n/\log n}$ result for double-exponential transformations to the infinite interval is also familiar and was originally demonstrated in the context of trapezoidal-rule quadrature in [20]. The closely related result pertaining to sinc interpolation was discussed in [11], and the Fourier equivalent in [25]. Both the $C^{-\sqrt{n}}$ and $C^{-n/\log n}$ results are widely known in these contexts, but as far as we are aware, have not yet been explicitly demonstrated for Chebyshev interpolants.

This paper is structured as follows. In Section 2, we make some key definitions and establish several notational conventions that will be needed throughout. In Section 3, we provide convergence proofs regarding the results on the semi-infinite interval. In Section 4, we do the the same for the infinite interval. In Section 5, we present numerical experiments which verify the results of Sections 3 and 4. Finally, in Section 6, we discuss our findings and make some concluding remarks.

2. Definitions and notation. Throughout this paper, we shall consider functions $f(x)$ which are continuous on $[0, 1]$, and analytic on at least $(0, 1)$. Functions defined on more general intervals $[a, b]$ may be trivially rescaled by use of a linear transformation, so it will suffice for us to consider only the $[0, 1]$ case.

The four conformal maps we shall analyse consist of two exponential transforms and two double-exponential transforms and are defined as follows:

$$\begin{aligned}
&\varphi_{\text{E}}(x) = \log(x), && \varphi_{\text{E}}^{-1}(s) = e^s;\\
&\psi_{\text{E}}(x) = \log(x/(1-x)), && \psi_{\text{E}}^{-1}(s) = e^s/(1+e^s);\\
&\varphi_{\text{DE}}(x) = -\log(1-\varphi_{\text{E}}(x)), && \varphi_{\text{DE}}^{-1}(s) = \varphi_{\text{E}}^{-1}(1-e^{-s});\\
&\psi_{\text{DE}}(x) = \sinh^{-1}(\psi_{\text{E}}(x)/\pi), && \psi_{\text{DE}}^{-1}(s) = \psi_{\text{E}}^{-1}(\pi\sinh(s)).
\end{aligned}$$

The action of these maps upon the interval $[0,1]$ can be summarised by

$$\varphi_{\text{E}}\,,\,\varphi_{\text{DE}}\;:\;[0,1]\mapsto(-\infty,\;0\,];$$
$$\psi_{\text{E}}\,,\,\psi_{\text{DE}}\;:\;[0,1]\mapsto(-\infty,\infty).$$

The following mnemonic may be helpful for remembering which symbol is which: E and DE, of course, stand for exponential and double-exponential. As for φ and ψ, we see that the symbol φ has only one endpoint tip, whereas ψ has two. These correspond to maps onto intervals with one or two points at infinity, respectively.

Letting $\varphi=\varphi_{\text{E}},\varphi_{\text{DE}}$ and $\psi=\psi_{\text{E}},\psi_{\text{DE}}$ as appropriate, we define

$$\mathcal{F}(s)=\begin{cases} f(\varphi^{-1}(s)), & s\in(-\infty,\;0\,];\\ f(\psi^{-1}(s)), & s\in(-\infty,\infty).\end{cases}$$

Thus, $\mathcal{F}(s)$ denotes transplanted functions in the s-variable. We shall see later that under reasonable continuity assumptions, $\mathcal{F}(s)$ approaches its limiting values either exponentially or double-exponentially quickly as s tends towards $-\infty$ or $\pm\infty$.

This motivates a strategy of truncating the semi-infinite and infinite intervals to $[-L,0]$ and $[-L,L]$, respectively, and approximating the transplanted functions with Chebyshev interpolants. We can accomplish this by applying a further transformation which scales $[-L,0]$ and $[-L,L]$ to the interval $[-1,1]$. We therefore define

$$F_L(y)=\begin{cases} f(\varphi^{-1}(L(y-1)/2)), & y\in(-\infty,\;1\,];\\ f(\psi^{-1}(Ly)), & y\in(-\infty,\infty).\end{cases}$$

These are the functions to which Chebyshev interpolation will be applied on $[-1,1]$. See Figure 2.1 for a schematic illustration of this process.

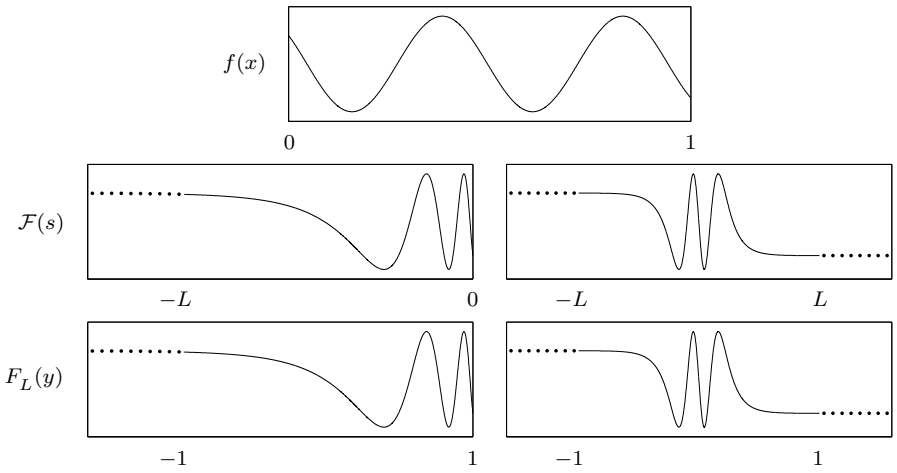

FIG. 2.1. *A function $f(x)$ on $[0,1]$ (top); its transplants $\mathcal{F}(s)$ on both $(-\infty,0]$ (middle, left), and $(-\infty,\infty)$ (middle, right); and scaled versions of each, $F_L(y)$, corresponding to mapping $[-L,0]$ and $[-L,L]$ to $[-1,1]$ for some arbitrary domain-size parameter L (bottom). Chebyshev interpolation will be applied to the functions $F_L(y)$ on $[-1,1]$, as indicated by the solid part of the curve.*

It now becomes clear that there are two sources of error to be considered. The first is the *interpolation error*, corresponding to sampling $F_L(y)$ at a finite number of grid points on $[-1,1]$; the second is the *domain error*, which we define as the maximum of

the two quantities $|F_L(-1) - f(0)|$ and $|F_L(1) - f(1)|$. The overall convergence rates of our approximations will be obtained by determining the optimal balance of these two sources of error. For a related discussion, see e.g. [2, Ch. 17].

It remains to specify an approximation to the function $f(x)$ on $[0, 1]$. By construction, for any $L > 0$, the three functions $f(x)$, $\mathcal{F}(s)$ and $F_L(y)$ have a point-wise correspondence across their respective domains. However, the Chebyshev interpolant $P_n(y)$ only provides an approximation to $F_L(y)$ on $[-1, 1]$, which corresponds to only a sub-interval of $[0, 1]$ in the x-variable. For the maps $\varphi_{\text{E}}, \varphi_{\text{DE}}$ to $(-\infty, 0]$, we denote this subinterval by $[x_L, 0]$, and for the maps $\psi_{\text{E}}, \psi_{\text{DE}}$ to $(-\infty, \infty)$, by $[x_L, 1-x_L]$. The number x_L is positive and, depending on the transform, decreases either exponentially or double-exponentially to zero as $L \to \infty$. Outside of these intervals, we define our approximation to be equal to the limiting endpoint values $f(0)$ or $f(1)$.

Our approximation to $f(x)$ on $[0, 1]$ is denoted by $p_L^{(n)}$. In the $(-\infty, 0]$ case, letting φ denote either φ_{E} or φ_{DE}, we define this to be

$$p_L^{(n)}(x) = \begin{cases} f(0), & x \in [0, x_L); \\ P_n(2\varphi(x)/L + 1), & x \in [x_L, 1]; \end{cases} \tag{2.1}$$

where $x_L = \varphi^{-1}(-L)$. Similarly, in the $(-\infty, \infty)$ case, letting ψ denote either ψ_{E} or ψ_{DE}, with $x_L = \psi^{-1}(-L) = \psi^{-1}(L)$, we define

$$p_L^{(n)}(x) = \begin{cases} f(0), & x \in [0, x_L); \\ P_n(\psi(x)/L), & x \in [x_L, 1 - x_L]; \\ f(1), & x \in (1 - x_L, 1]. \end{cases} \tag{2.2}$$

The error in these approximations will be measured by examining the quantity $||f - p_L^{(n)}||$, where $||\cdot||$ is the L_∞ norm. Since $p_L^{(n)}$ is defined piecewise, the total error is the maximum of the errors over each piece. The approximation error corresponding to (2.1) is therefore

$$||f - p_L^{(n)}|| = \max\left\{||F_L - P_n||_{y\in[-1,1]}, ||f - f(0)||_{x\in[0,x_L)}\right\}, \tag{2.3}$$

and the error corresponding to (2.2) is

$$\begin{aligned} ||f - p_L^{(n)}|| = \max\big\{&||F_L - P_n||_{y\in[-1,1]}, \\ &||f - f(0)||_{x\in[0,x_L)}, ||f - f(1)||_{x\in(1-x_L,1]}\big\}. \end{aligned} \tag{2.4}$$

In each of (2.3) and (2.4), the first term on the right-hand side represents the interpolation error, and the remaining term or terms the domain error.

To quantify the interpolation error, we shall require the notion of a *Bernstein ellipse* E_μ, which, for some $\mu > 0$, we define to be the open set in the complex plane bounded by the curve

$$\left\{\frac{1}{2}\left(e^{\mu}e^{i\theta} + e^{-\mu}e^{-i\theta}\right) : \theta \in [0, 2\pi]\right\}.$$

The degree n Chebyshev interpolant to $F_L(y)$ on $y \in [-1, 1]$ can be written

$$P_n(y) = \sum_{k=0}^{n} c_k T_k(y),$$

where T_k are the Chebyshev polynomials of the first kind. P_n corresponds to interpolation at the second-kind Chebyshev points $y_k = \cos(k\pi/n)$, and the c_k are aliased expansion coefficients corresponding to the function samples $F_L(y_k)$.

If F_L is analytic in a particular ellipse E_μ, and M_μ denotes the maximum absolute value of F_L in E_μ, then a bound for the error in $n+1$ point Chebyshev interpolation going back to Bernstein in 1912 is [23, Ch. 8]

$$||F_L - P_n||_{y\in[-1,1]} \leq \frac{4}{\mu} M_\mu e^{-\mu n}. \tag{2.5}$$

The domain error depends upon both the function f and the conformal map. In particular, if f satisfies, for some real $A, \alpha > 0$,

$$|f(x) - f_0| \leq A|x|^\alpha, \qquad x \in [0,1], \tag{2.6}$$

then it is straightforward to show that there exist $A_1, A_2 > 0$ such that

$$||f - f_0||_{x\in[0,x_L)} \leq \begin{cases} A_1 e^{-\alpha L} & \text{(for the map } \varphi_{\text{E}}\text{);} \\ A_2 e^{-\alpha e^L} & \text{(for the map } \varphi_{\text{DE}}\text{).} \end{cases} \tag{2.7}$$

Similarly, if f satisfies, for some real $B, \beta > 0$,

$$\begin{aligned} |f(x) - f_0| &\leq B|x|^\beta, \\ |f(x) - f_1| &\leq B|1-x|^\beta, \end{aligned} \qquad x \in [0,1], \tag{2.8}$$

then there exist $B_1, B_2 > 0$ such that

$$\begin{aligned} \max\left\{||f - f_0||_{x\in[0,x_L)}, ||f - f_1||_{x\in(1-x_L,1]}\right\} & \\ \leq \begin{cases} B_1 e^{-\beta L} & \text{(for the map } \psi_{\text{E}}\text{);} \\ B_2 e^{-\beta e^L} & \text{(for the map } \psi_{\text{DE}}\text{).} \end{cases} & \end{aligned} \tag{2.9}$$

We now have enough preparation to proceed to the theorems.

3. Maps to the semi-infinite interval. Typically in exponential and double-exponential mapping methods, one is concerned with a function analytic in an infinite strip about the real axis. Such strips are the natural setting for trapezoidal rule quadrature, and for sinc and Fourier interpolation. The vast majority of past literature has focussed on this setting; see [10, 11, 12, 14, 15, 16, 17, 19, 20]. Given some conformal map from the infinite interval to $[0,1]$, the image of an infinite strip under this transformation corresponds to a particular region of analyticity in the x-variable from which suitable definitions of function classes follow; see [21, 22].

Since in this paper we will be working with Chebyshev interpolants, we shall be concerned primarily with *ellipses*. Nevertheless, we shall see that for the maps ψ_{E} and ψ_{DE} to $(-\infty,\infty)$ considered in Section 4, in order to accommodate the limiting case $L \to \infty$, this still amounts to requiring analyticity in an infinite strip in the s-variable. On the other hand, for the maps φ_{E} and φ_{DE} to $(-\infty, 0]$, we shall require analyticity in a region which is essentially parabolic. We begin by describing this case.

Given some complex number $\zeta = \zeta_1 + i\zeta_2$, $\zeta_1 < 0$, $\zeta_2 \neq 0$, we define the quasi-parabolic region $\mathcal{P}_\zeta$ in the complex s-plane to be the the union of 1) the interior of a parabola symmetric about the real axis passing through both ζ and the origin, and 2) the interior of an ellipse with foci at $2\zeta_1, 0$ also passing through ζ. We consider this region to be an open set bounded by the parabola for $\text{Im}(s) < \zeta_1$, and by the ellipse for $\text{Im}(s) \geq \zeta_1$. See Figure 3.1.

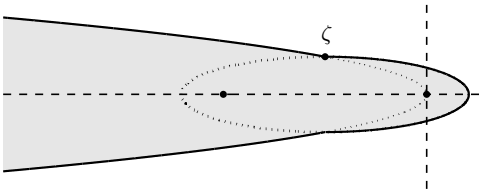

FIG. 3.1. *A visualisation of the quasi-parabolic region $\mathcal{P}_\zeta$ in the s-plane. The complex number $\zeta = \zeta_1 + i\zeta_2$ is indicated, and the two black dots at $s = 2Re(\zeta), 0$ are the foci of the ellipse.*

A formal description of this region is as follows:

$$\begin{aligned}
\mathcal{P}_\zeta^{(1)} &= \left\{ s \in \mathbb{C} : \zeta_2^2 \mathrm{Re}(s) < \zeta_1 \mathrm{Im}(s)^2 \right\}, \\
\mathcal{P}_\zeta^{(2)} &= \left\{ s \in \mathbb{C} : \zeta_2^2 (\mathrm{Re}(s) - \zeta_1)^2 + (\zeta_1^2 + \zeta_2^2) \mathrm{Im}(s)^2 < \zeta_2^2 (\zeta_1^2 + \zeta_2^2) \right\}, \\
\mathcal{P}_\zeta &= \mathcal{P}_\zeta^{(1)} \cup \mathcal{P}_\zeta^{(2)}.
\end{aligned}$$

The reason for requiring a region of analyticity such as this is in order to ensure that for all $L \geq 2|\zeta_1|$, the interior of every ellipse which passes through ζ and has foci at $-L, 0$ is contained within the quasi-parabolic region $\mathcal{P}_\zeta$. Incidentally, there is no special reason why the appended region needs to be elliptical: any sufficiently large region containing a fixed neighbourhood of the origin could have been used. We have chosen to use an ellipse for convenience.

The transforms $\varphi_{\mathrm{E}}^{-1}$ and $\varphi_{\mathrm{DE}}^{-1}$ are both analytic functions throughout $\mathcal{P}_\zeta$ which map it onto the following regions in the complex x-plane:

$$\begin{aligned}
\mathcal{Q}_\zeta &= \varphi_{\mathrm{E}}^{-1}(\mathcal{P}_\zeta), \\
\mathcal{R}_\zeta &= \varphi_{\mathrm{DE}}^{-1}(\mathcal{P}_\zeta).
\end{aligned}$$

These maps are not one-to-one. To make them one-to-one, one must view the images as infinitely sheeted Riemann surfaces wrapping around $x = 0$. It is in this sense, as Riemann surfaces, that we shall regard the regions $\mathcal{Q}_\zeta$ and $\mathcal{R}_\zeta$. See Figure 3.2.

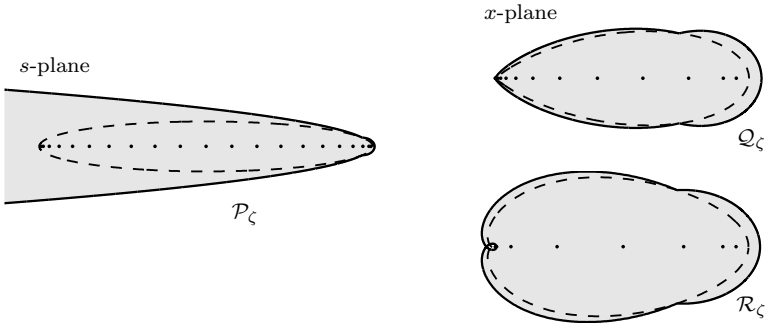

FIG. 3.2. *The functions φ_E^{-1} and φ_{DE}^{-1} map the quasi-parabolic region $\mathcal{P}_\zeta$ in the s-plane onto the regions $\mathcal{Q}_\zeta = \varphi_E^{-1}(\mathcal{P}_\zeta)$ and $\mathcal{R}_\zeta = \varphi_{DE}^{-1}(\mathcal{P}_\zeta)$ in the x-plane. $\mathcal{Q}_\zeta$ and $\mathcal{R}_\zeta$ are both infinitely sheeted Riemann surfaces which wrap around $x = 0$. We also show a sample ellipse (dashed line) and some interpolation nodes (dots) in the s plane, and the images of each in the regions $\mathcal{Q}_\zeta, \mathcal{R}_\zeta$.*

The following theorem is the more general of our two results concerning the transformation φ_{E}. It describes the convergence rate of a sequence of Chebyshev interpolants, given an assumption of analyticity in the region Q_ζ.

THEOREM 3.1. *Let f satisfy the condition* (2.6) *and be analytic in the region $\mathcal{Q}_\zeta$ for some ζ, where it satisfies $|f(x)| \leq M$ for some M. Let $p_L^{(n)}$ be the approximation defined by* (2.1) *corresponding to the exponential transform $\varphi = \varphi_{_E}$. If L is defined for any $c > 0$ by $L = cn^{2/3}$, then for all sufficiently large n,*

$$||f - p_L^{(n)}|| \leq C^{-n^{2/3}} \qquad \textit{for some } C > 1.$$

Proof. Since $f(x)$ is analytic in $\mathcal{Q}_\zeta$, the transplanted function $\mathcal{F}(s)$ is analytic in the quasi-parabolic region P_ζ, an example of which is given in Figure 3.2. Consequently, given any $L > 0$, if we define $p = 2\zeta_1/L + 1$, $q = 2\zeta_2/L$ as the real and imaginary parts of the image of ζ in the y-variable, then the function $F_L(y) = \mathcal{F}(L(y-1)/2)$ is analytic inside a Bernstein ellipse E_μ whose parameter μ is given by[1]

$$\mu = \sinh^{-1}\sqrt{\frac{1}{2}\left(p^2 + q^2 - 1 + \sqrt{(1-p^2-q^2)^2 + 4q^2}\right)}. \tag{3.1}$$

It can be shown that with μ defined by (3.1), there exists a constant $\eta > 0$ such that $\mu > \eta/\sqrt{L}$ for all $L > 0$. Thus, from (2.3), (2.5) and (2.7), we have

$$||f - p_L^{(n)}|| \leq \max\left\{\frac{4\sqrt{L}}{\eta} M_\mu e^{-\eta n/\sqrt{L}}, A_1 e^{-\alpha L}\right\}.$$

Now for any $L \geq 2|\zeta_1|$, with μ defined as above, the ellipse $L(E_\mu - 1)/2$ in the s-plane is entirely contained within the region $\mathcal{P}_\zeta$. Thus M_μ may be bounded by M, the maximum absolute value of f in $\mathcal{Q}_\zeta$, which exists by assumption. We then see that for any $c > 0$, the choice $L = cn^{2/3}$ gives

$$||f - p_L^{(n)}|| \leq \max\left\{\frac{4M\sqrt{c}}{\eta} n^{1/3}, A_1\right\} e^{-\delta n^{2/3}},$$

where $\delta = \min(\eta/\sqrt{c}, \alpha c)$. This is valid for all $n \geq (2|\zeta_1|/c)^{3/2}$. The result now follows with a suitable C chosen to absorb the bracketed terms. ∎

The proof of Theorem 3.1 is methodologically very similar to the proofs of Theorems 3.3, 4.1 & 4.2. The general pattern followed in these four theorems is to fix a particular region of analyticity in the s-plane and then deduce how rapidly the Bernstein ellipses in the y-variable must shrink as $L \to \infty$. For example, in the above theorem, we showed that $\mu = \mathcal{O}(L^{-1/2})$, which implies that the choice $L = \mathcal{O}(n^{2/3})$ balances the error contributions. Theorem 3.2, however, is fundamentally different to the other theorems in that we shall fix the size of the Bernstein ellipse in the y-variable at the outset and allow the corresponding regions of analyticity in the s- and x-variables to grow without bound as $L \to \infty$. A Bernstein ellipse of fixed size is of course necessary if one wishes to obtain true C^{-n} geometric convergence.

Before proceeding, it will be helpful to consider some additional properties of the transformation $\varphi_{\text{E}}^{-1}(s) = e^s$. For any integer k, this conformal map transplants the

[1] This is one possible formula for computing the elliptic coordinate parameter μ in terms of the real and imaginary parts of some complex number $p + iq$ which is assumed to lie on the ellipse's boundary. The elliptic coordinate system is defined to be: $p = \cosh(\mu)\cos(\theta)$, $q = \sinh(\mu)\sin(\theta)$, for $\mu > 0$, $\theta \in [0, 2\pi]$.

infinite strip $(2k-1)\pi < |\mathrm{Im}(s)| < (2k+1)\pi$ in the complex s-plane onto the slit plane $\mathbb{C}\backslash(-\infty, 0]$ in the complex x-plane. Corresponding to each k are individual sheets of a Riemann surface in the x-plane, each separated by a branch cut along the negative real axis. See Figure 3.3 for an illustration.

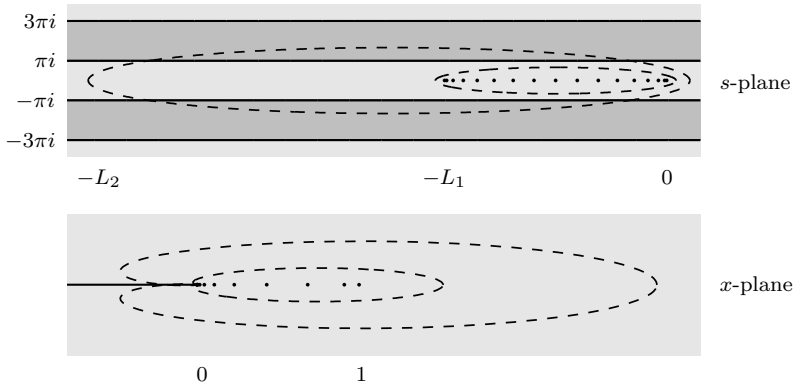

FIG. 3.3. *For each $k \in \mathbb{Z}$, φ_E^{-1} maps the infinite strips $(2k-1)\pi < |Im(s)| < (2k+1)\pi$ in the s-plane onto sheets of a Riemann surface in the x-plane separated by the branch cut $(-\infty, 0]$. In the bottom panel is the image of the principal strip $-\pi < |Im(s)| < \pi$. We also show two sample ellipses $L_1(E_\mu - 1)/2$ and $L_2(E_\mu - 1)/2$ in the s-plane and their images in the x-plane. The first ellipse remains within the principal strip and does not cross the branch cut in the x-plane; the opposite is true for the second ellipse. Sample Chebyshev interpolation nodes corresponding to the first ellipse are also shown, along with their images.*

We see from Figure 3.3 that in the limit $L \to \infty$, we shall require analyticity of the function $\mathcal{F}(s)$ in the entire complex s-plane. However, since only infinite strips of width 2π are mapped in one-to-one fashion onto the x-plane, we must require that the problem function $f(x)$ is analytic on the entire Riemann surface in the x-plane, including on the branch cuts. This condition ensures that ellipses in the s-plane can cross over onto adjacent strips without difficulty.

Allowing these regions of analyticity in the s- and x-variables to grow without bound comes with an associated cost. This is in having to account for the quantity M_μ in (2.5), which will now grow rapidly as L increases. It turns out that the required condition upon the function $\mathcal{F}$ is that it should be a function of *exponential type*, i.e., satisfy $|\mathcal{F}(s)| = \mathcal{O}(\exp|s|^\rho)$ for some $\rho > 0$ as $|s| \to \infty$. If this is the case, then one can prove geometric convergence of the approximations $p_L^{(n)}$. However, most problem functions f in the x-variable do not correspond to functions of exponential type in the s-variable. One can see that this is the case by recalling that for the map $\varphi_{\mathrm{E}}^{-1}$, we have $\mathcal{F}(s) = f(e^s)$; thus, only functions which grow *algebraically* in the x-plane will correspond to functions of exponential type in the s-variable.

We state these concepts precisely in the following theorem:

THEOREM 3.2. *Let f satisfy the condition* (2.6) *and be analytically continuable along any curve in the complex plane disjoint from $x = 0$. Suppose it satisfies $|f(x)| = \mathcal{O}(|x|^\gamma)$ as $|x| \to \infty$ for some $\gamma > 0$ uniformly on all such curves. Let $p_L^{(n)}$ be the approximation defined by* (2.1) *corresponding to the exponential transform $\varphi = \varphi_E$. If L is defined for any $c > 0$ by $L = cn$, then*

$$||f - p_L^{(n)}|| \le C^{-n} \qquad \textit{for some } C > 1.$$

Proof. From (2.3), (2.5) and (2.7), we have

$$||f - p_L^{(n)}|| \leq \max\left\{ \frac{4M_\mu}{\mu} e^{-\mu n}, A_1 e^{-\alpha L} \right\}. \tag{3.2}$$

The quantity M_μ is defined as the maximum absolute value of the function $F_L(y) = f(\varphi_{\text{E}}^{-1}(L(y-1)/2))$ on the ellipse E_μ. Our assumption on the growth rate of f implies $M_\mu \leq W \exp\left(L\gamma\left(e^\mu + e^{-\mu} - 2\right)/4\right)$ for some $\gamma, W > 0$. Thus with $L = cn$,

$$||f - p_L^{(n)}|| \leq \max\left\{ \frac{4W}{\mu} e^{-w(\mu)n}, A_1 e^{-\alpha c n} \right\},$$

where $w(\mu) = \mu - \frac{c\gamma}{4}(e^\mu + e^{-\mu} - 2)$. Using a Taylor series argument, it is straightforward to show that for all $c, \gamma > 0$ this function takes positive values in some interval $(0, \tilde{\mu})$, and that w attains a maximum at some point $\mu_m \in (0, \tilde{\mu})$. (For the purposes of proving this theorem, it is not crucial to compute the maximum μ, as any value in $(0, \tilde{\mu})$ will suffice; we do so in order to obtain as strong a bound as possible.) Using this value μ_m, we then have

$$||f - p_L^{(n)}|| \leq \max\left\{ \frac{4W}{\mu_m}, A_1 \right\} e^{-\delta n},$$

where $\delta = \min(w(\mu_m), \alpha c)$. The result now follows with a suitable choice of C. □

It is quite remarkable that true geometric convergence is attainable under a domain-truncation strategy. However, we reiterate that this result only applies to a relatively small class of functions: see Section 5 for examples. We note again that Boyd was aware of this result in a less general form [3].

Finally in this section, we consider the double-exponential transform φ_{DE}.

THEOREM 3.3. *Let f satisfy the condition* (2.6) *and be analytic in the region $\mathcal{R}_\zeta$ for some ζ, where it satisfies $|f(x)| \leq M$ for some M. Let $p_L^{(n)}$ be the approximation defined by* (2.1) *corresponding to the double-exponential transform $\varphi = \varphi_{DE}$. If L is defined for any $c > 0$ by $L = \log(cn)$, then for all sufficiently large n,*

$$||f - p_L^{(n)}|| \leq C^{-n/\log n} \qquad \textit{for some } C > 1.$$

Proof. Since $f(x)$ is analytic on R_ζ, $\mathcal{F}(s)$ is analytic in $\mathcal{P}_\zeta$, as in Figure 3.2. If we again define μ according to (3.1), then using (2.3), (2.5) and (2.7), we have

$$||f - p_L^{(n)}|| \leq \max\left\{ \frac{4\sqrt{L}}{\eta} M_\mu e^{-\eta n/\sqrt{L}}, A_2 e^{-\alpha e^L} \right\}.$$

As before, for any $L \geq 2|\zeta_1|$, M_μ may be bounded by the maximum absolute value of f in $\mathcal{R}_\zeta$. Choosing $L = \log(cn)$ for any $c > 0$ then gives

$$||f - p_L^{(n)}|| \leq \max\left\{ \frac{4\sqrt{\log(cn)}}{\eta} M e^{-\eta n/\sqrt{\log(cn)}}, A_1 e^{-\alpha c n} \right\},$$

which is valid for all $n \geq 1/c$. If $c \leq 1$, then $e^{-\eta n/\sqrt{\log(cn)}} \leq e^{-\eta n/\log n}$; if $c > 1$, then there is a constant $C_1 > 1$ such that $e^{-\eta n/\sqrt{\log(cn)}} \leq C_1^{-n/\log n}$ for all $n \geq 1/c$. In either case, the result follows with a suitable choice of C. □

4. Maps to the infinite interval. Next, by considering the transformations ψ_{E} and ψ_{DE}, we verify that the $C^{-\sqrt{n}}$ and $C^{-n/\log n}$ results from classic sinc and double-exponential quadrature theory also hold in the Chebyshev interpolation setting.

In the proofs of the following two theorems, we shall see that in the limit $L \to \infty$, the region of analyticity that bounds the ellipses in the s-variable is now required to be an infinite strip, rather than a (quasi-)parabola, as was the case for the maps φ_{E} and φ_{DE}. However, there is also a more subtle difference associated with this setting: regardless of whether or not f has any singularities, the transplanted function $\mathcal{F}(s)$ always does have singularities, due to those present in the transplanting maps ψ_{E}^{-1} and ψ_{DE}^{-1}. Consequently, in contrast to the results of the previous section, there can be no class of functions for which, say, geometric convergence is attainable.

The regions of analyticity associated with the transformations ψ_{E} and ψ_{DE} are precisely those which appear in classic sinc and double-exponential theory, and are defined as follows. Given some $0 < d \leq \pi$, let $\mathcal{S}_d$ be the infinite strip of half-width d in the complex s-plane described by

$$\mathcal{S}_d = \{\, s \in \mathbb{C} \;:\; |\text{Im}(s)| < d \,\}.$$

The transforms ψ_{E}^{-1} and ψ_{DE}^{-1} are both analytic functions throughout $\mathcal{S}_d$ which map it onto the following regions in the complex x-plane:

$$\mathcal{T}_d = \psi_{\text{E}}^{-1}(\mathcal{S}_d),$$
$$\mathcal{U}_d = \psi_{\text{DE}}^{-1}(\mathcal{S}_d).$$

There is a one-to-one correspondence between $\mathcal{S}_d$ and $\mathcal{T}_d$, but not between $\mathcal{S}_d$ and $\mathcal{U}_d$. As before, a one-to-one correspondence exists in the latter case if $\mathcal{U}_d$ is regarded as an infinitely sheeted Riemann surface wrapping around both $x = 0$ and $x = 1$. $\mathcal{T}_d$ is simply a lens-shaped region formed by two circular arcs meeting at $x = 0, 1$ with half-angle d. See Figure 4.1.

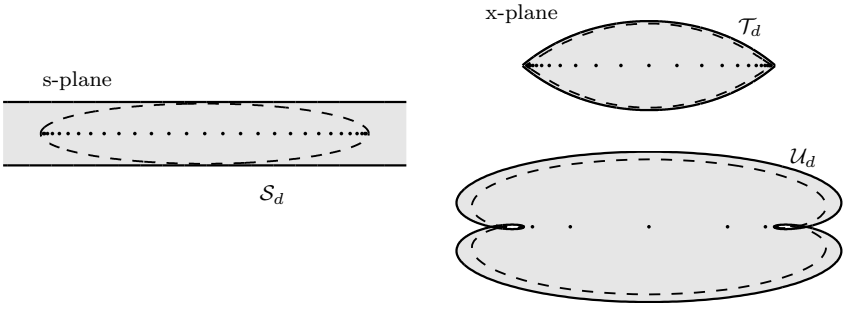

FIG. 4.1. *The functions ψ_E^{-1} and ψ_{DE}^{-1} map $\mathcal{S}_d$ onto $\mathcal{T}_d = \psi_E^{-1}(\mathcal{S}_d)$ and $\mathcal{U}_d = \psi_{DE}^{-1}(\mathcal{S}_d)$. $\mathcal{T}_d$ is a lens-shaped region formed by two circular arcs meeting with half-angle d. $\mathcal{U}_d$ is an infinitely sheeted Riemann surface wrapping around $x = 0, 1$. We also show a sample ellipse (dashed line) and some interpolation nodes (dots) in the s plane and the images of each in the regions $\mathcal{T}_d$,$\mathcal{U}_d$.*

We begin with the $C^{-\sqrt{n}}$ result concerning the exponential transform φ_{E}.

THEOREM 4.1. *Let f satisfy the condition* (2.8) *and be analytic in the region $\mathcal{T}_d$ for some $d > 0$, where it satisfies $|f(x)| \leq M$ for some M. Let $p_L^{(n)}$ be the*

approximation defined by (2.2) *corresponding to the exponential transform* $\psi = \psi_{E}$. *If* L *is defined for any* $c > 0$ *by* $L = c\sqrt{n}$, *then*

$$||f - p_L^{(n)}|| \leq C^{-\sqrt{n}} \qquad \textit{for some } C > 1.$$

Proof. From (2.4), (2.5) and (2.9), we have

$$||f - p_L^{(n)}|| \leq \max\left\{\frac{4}{\mu} M_\mu e^{-\mu n}, B_1 e^{-\beta L}\right\}.$$

Due to the assumed analyticity of $\mathcal{F}(s)$ in the strip $\mathcal{S}_d$, for any $L > 0$, the function $F_L(y) = \mathcal{F}(Ly)$ is analytic inside the Bernstein ellipse E_μ with semi-minor axis length d/L. The corresponding ellipse parameter is

$$\mu = \log\left(d/L + \sqrt{d^2/L^2 + 1}\right), \tag{4.1}$$

and if we define $\eta = \log(d + \sqrt{d^2+1})$, then $\mu > \eta/L$ for all $L > 0$.

Now, since f is bounded in $\mathcal{T}_d$, we may infer that $M_\mu \leq M$ for all $L > 0$. Thus, setting $L = c\sqrt{n}$ for any $c > 0$, we have

$$\begin{aligned} ||f - p_L^{(n)}|| &\leq \max\left\{\frac{4M}{\eta} c\sqrt{n} e^{-\eta\sqrt{n}/c}, B_1 e^{-\beta c\sqrt{n}}\right\}, \\ &\leq \max\left\{\frac{4M}{\eta} c\sqrt{n}, B_1\right\} e^{-\delta\sqrt{n}}, \end{aligned}$$

where $\delta = \min(\eta/c, \beta c)$. The result follows with a suitable choice of C. □

In the final theorem, we demonstrate the $C^{-n/\log n}$ convergence result associated with the double-exponential map ψ_{DE}.

THEOREM 4.2. *Let* f *satisfy* (2.8) *and be analytic in the region* $\mathcal{U}_d$ *for some* $d > 0$, *where it satisfies* $|f(x)| \leq M$ *for some* M. *Let* $p_L^{(n)}$ *be the approximation defined by* (2.2) *corresponding to the double-exponential transform* $\psi = \psi_{DE}$. *If* L *is defined for any* $c > 0$ *by* $L = \log(cn)$, *then for all sufficiently large* n,

$$||f - p_L^{(n)}|| \leq C^{-n/\log n} \qquad \textit{for some } C > 1.$$

Proof. From (2.4), (2.5) and (2.9), we have

$$||f - p_L^{(n)}|| \leq \max\left\{\frac{4}{\mu} M_\mu e^{-\mu n}, B_2 e^{-\beta e^L}\right\}.$$

As in Theorem 4.1, for any $L > 0$, $F_L(y)$ is analytic inside the Bernstein ellipse E_μ, where μ is given by (4.1), and again we have $\mu > \eta/L$ where $\eta = \log(d + \sqrt{d^2+1})$. Also, for all $L > 0$, we may again bound M_μ by M.

Setting $L = \log(cn)$ for any $c > 0$ then yields

$$||f - p_L^{(n)}|| \leq \max\left\{\frac{4M}{\eta}\log(cn) e^{-\eta n/\log(cn)}, B_2 e^{-\beta cn}\right\}.$$

If $c \leq 1$, then $e^{-\eta n/\log(cn)} \leq e^{-\eta n/\log n}$; if $c > 1$, then there is a constant $C_1 > 1$ such that $e^{-\eta n/\log(cn)} \leq C_1^{-n/\log n}$. In either case, the result follows with a suitable choice of C. □

5. Numerical experiments. We now consider some examples of functions defined on $[0,1]$ which may be treated by the above techniques. Consider for instance

$$f_1(x) = x,\ \ f_2(x) = \sqrt{x},\ \ f_3(x) = 1 + x^{1/4},\ \ f_4(x) = x\log x.$$

Each of these functions is free from singularities away from $x = 0$, and grows only algebraically in the complex plane. The conditions of Theorem 3.2 are met in each case, and therefore using the transform φ_{E} and setting $L = cn$ will lead to geometric convergence.

In contrast to this, the following functions, for one reason or another, do not meet the conditions of Theorem 3.2:

$$f_5(x) = \sin(x),\ \ f_6(x) = \sqrt{x}\cos x,\ \ f_7(x) = \frac{1 + x^{1/4}}{x^2 - x + 1},\ \ f_8(x) = \frac{x\log x}{1 + x}.$$

The reasons are that although f_5 and f_6 do not have any singularities away from $x = 0$, they grow faster than algebraically in the complex plane. By contrast, f_7 and f_8 grow only algebraically, but have singularities away from $x = 0$. Functions f_5, f_6, f_7, f_8 do all meet the conditions of Theorems 3.1 and 3.3, however. Thus by using the transforms φ_{E} or φ_{DE}, and setting $L = cn^{2/3}$ or $L = \log(cn)$ as appropriate, we may expect to obtain either $C^{-n^{2/3}}$ or $C^{-n/\log n}$ convergence, respectively.

Finally, consider

$$f_9(x) = x^{1/3}(1-x)^{2/3} + x,\ \ f_{10}(x) = \sqrt{x - x^2}\tanh(3x - 2).$$

Since these functions have singularities at both endpoints, they cannot satisfy the conditions of any of Theorems 3.1, 3.2, or 3.3. They do, however, satisfy the conditions of Theorems 4.1 and 4.2. Thus, by using the transforms ψ_{E} or ψ_{DE}, and setting $L = c\sqrt{n}$ or $L = \log(cn)$ as appropriate, we may expect to achieve either $C^{-\sqrt{n}}$ or $C^{-n/\log n}$ convergence, respectively. We note also that as soon we are forced to work with the maps ψ_{E} and ψ_{DE} rather than φ_{E} and φ_{DE}, issues such as growth rates in the complex plane or the presence of additional singularities become factors that are irrelevant with regards to the overall convergence rate.

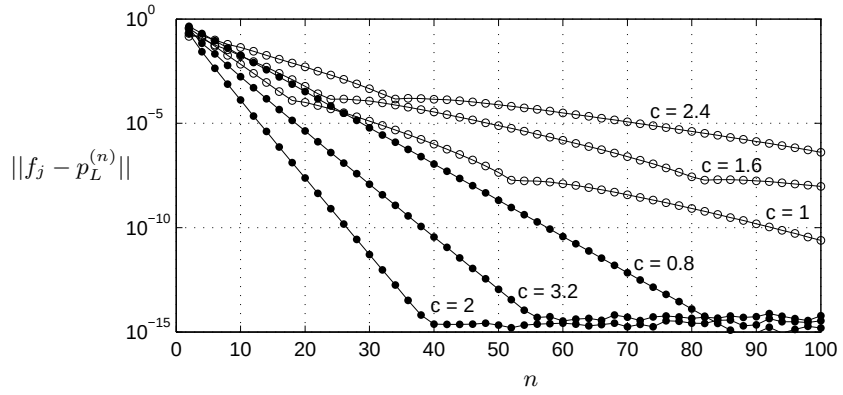

FIG. 5.1. *Convergence of the approximations $p_L^{(n)}$ defined by (2.1) corresponding to the exponential transformation φ_{E} for $f_2(x) = \sqrt{x}$ (dots), and $f_6(x) = \sqrt{x}\cos x$ (circles). In each case L is set according to $L = cn$, for some sample values of c, which are indicated in the plot. For f_2, we observe geometric convergence as predicted by Theorem 3.2.*

There are of course also many examples of functions which are not suited to any of these classes, for example $\sin(1/x)$ and $e^{\sqrt{x}}/x$. We note, however, that often, functions which blow up only algebraically at the endpoints (such as the latter of these two examples) may be treated in our framework after multiplication of a suitable algebraic factor.

Let us now provide some numerical confirmation of the theoretical convergence rates. We consider first the functions f_2 and f_6, which are ostensibly similar, yet exhibit markedly different convergence behaviour under the transformation φ_{E}. Setting $L = cn$ for any $c > 0$, we observe geometric convergence for f_2, but, due to the presence of the cosine, sub-geometric convergence for f_6; see Figure 5.1.

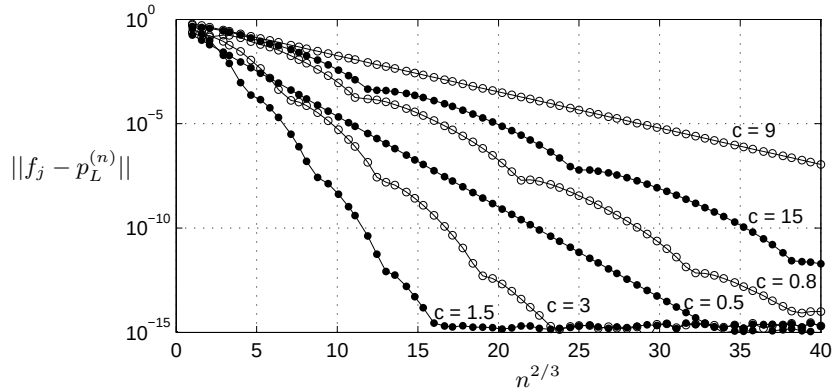

FIG. 5.2. *Convergence of the approximations $p_L^{(n)}$ defined by* (2.1) *corresponding to the exponential transformation φ_E for $f_2(x) = \sqrt{x}$ (dots), and $f_6(x) = \sqrt{x}\cos x$ (circles). L is set according to $L = cn^{2/3}$, for some sample values of c, which are indicated in the plot. For each function, we observe $C^{-n^{2/3}}$ convergence as predicted by Theorem 3.1.*

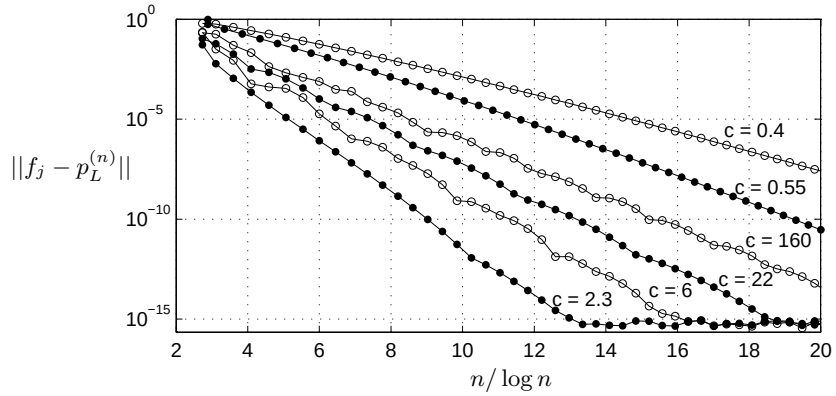

FIG. 5.3. *Convergence of the approximations $p_L^{(n)}$ defined by* (2.1) *corresponding to the double-exponential transformation φ_{DE} for f_2 (dots), and f_6 (circles). L is set according to $L = \log(cn)$, for some sample values of c, which are indicated in the plot. For each function, we observe $C^{-n/\log n}$ convergence as predicted by Theorem 3.3.*

Figure 5.1 provides an illustration of the importance of growth rates in the complex plane for the map ψ_{E}. We considered for illustration the functions f_2 and f_6, but could equally well have compared any of f_1, f_2, f_3, f_4 to any of f_5, f_6, f_7, f_8 and observed similar effects. Particularly interesting is the function f_8, which has a sin-

gularity (a simple pole) lying on the branch cut $(-\infty, 0]$. Due to the mapping effects of φ_{E} described in Figure 3.3, this is enough to ensure that ellipses in the s-variable remain bounded by a point lying on the boundary of the principal infinite strip. In contrast to this, we see in Figure 5.2 that setting $L = cn^{2/3}$ leads to $C^{-n^{2/3}}$ convergence for both f_2 and f_6. The same functions are considered again in Figure 5.3, though this time with the double-exponential transform ψ_{DE}. As expected, setting $L = \log(cn)$, leads to $C^{-n/\log n}$ convergence.

In Figures 5.4 and 5.5, we approximate the functions f_4 and f_9 under the transformations ψ_{E} and ψ_{DE}. Due to the singularity present at $x = 1$ in f_9, we have used Maple to obtain the necessary data down to the level of around 10^{-15}. Barycentric interpolation [1] seems to be the most straightforward way of implementing this, since one never need worry about computing Chebyshev expansion coefficients. All results in the other figures were obtained using Chebfun [24], which automates the interpolation process, but shares the limitations of double-precision arithmetic.

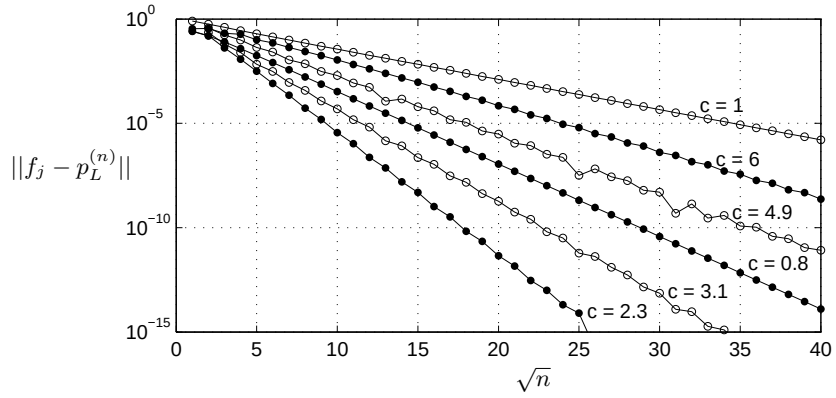

FIG. 5.4. *Convergence of the approximations $p_L^{(n)}$ defined by* (2.2) *corresponding to the exponential transformation ψ_E for $f_4(x) = x\log(x)$ (dots), and $f_9(x) = x^{1/3}(1-x)^{2/3} + x$ (circles). L is set according to $L = c\sqrt{n}$ for some sample values of c which are indicated in the plot. For each function, we observe $C^{-\sqrt{n}}$ convergence as predicted by Theorem 4.1.*

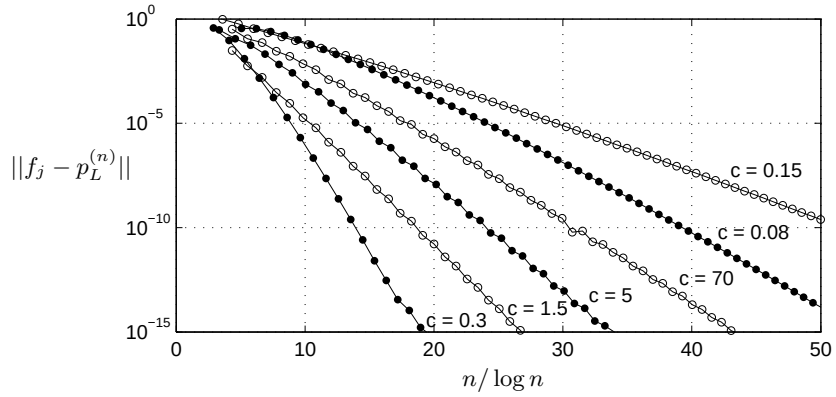

FIG. 5.5. *Convergence of the approximations $p_L^{(n)}$ defined by* (2.2) *corresponding to the dobule-exponential transformation ψ_{DE} for f_4 (dots), and f_9 (circles). L is set according to $L = \log(cn)$ for some sample values of c which are indicated in the plot. For each function, we observe $C^{-n/\log n}$ convergence as predicted by Theorem 4.2.*

We also give a brief experimental illustration of the differences between what one might call the *resolution properties* of the four maps. In the sense that we intend, resolution power is the notion of the average number of points-per-wavelength (ppw) required to resolve high frequency oscillations. For example, Fourier series are known to require at least 2 ppw before exponential convergence sets in, and Chebyshev series π ppw [7]. Table 2 shows the results of a simple resolution experiment involving the function $\sin(Mx)$. Without drawing rigorous conclusions, it is apparent that in this respect, Chebyshev interpolation under the maps to $(-\infty, 0]$ is much superior to Chebyshev interpolation under the maps to $(-\infty, \infty)$. This is still true even after one accounts for the $\pi/2$ factor associated with using Fourier or sinc interpolants in the $(-\infty, \infty)$ case. A forthcoming publication by Richardson and Trefethen and will conduct a detailed analysis in this regard.

	φ_{E}		φ_{DE}		ψ_{E}		ψ_{DE}	
M	n	ppw	n	ppw	n	ppw	n	ppw
1	84	–	46	–	443	–	99	–
10	144	90.5	64	40.2	724	455	168	106
100	501	31.5	204	12.8	2 073	130	499	31.4
1 000	3 334	20.9	1 295	8.14	12 967	81.5	3 097	19.5
10 000	29 730	18.7	11 582	7.28	118 322	74.3	27 458	17.3

TABLE 5.1

An indication of the resolution power of the four maps. For each, the function $\sin(Mx)$ *has been approximated to machine precision by an* $n+1$*-point Chebyshev interpolant to* $\mathcal{F}(s)$ *on either* $[-L, 0]$ *or* $[-L, L]$*. The parameter* L *is determined in each case such that* $|\mathcal{F}(s)|$ *is negligible outside of these intervals. ppw stands for "points-per-wavelength", which we define to be* $2\pi n/M$*. We note that* L *grows with* M*, albeit slowly; thus it is not clear if these ppw estimates converge as* $M \to \infty$*. We leave a theoretical analysis of resolution issues to a future work.*

Finally, we suggest how the techniques described in this paper can used for practical computation. Utlilising the black-box interpolation capabilities of Chebfun, we merely need to define the transplanted function $F_L(y)$ for some L, which we compute in advance such that $F_L(-1) \approx 10^{-16}$, and then rely on Chebfun to adaptively compute the correct number of interpolation points to use. This *a priori* domain determination technique also was utilised in the Sincfun software system described in [14]. The Matlab code is extremely simple; here is an example involving φ_{DE}:

```
f6 = @(x) sqrt(x).*cos(x); L = 4.3;                    % function; domain
phiIDE = @(s) exp(1-exp(-s));                          % inverse transform
Pn = chebfun(@(y) f6(phiIDE(L*(y-1)/2)));              % chebfun on [-1,1]
```

It is then a straightforward matter to evaluate such a representation:

```
phiDE = @(x) -log(1-log(x));                           % transform
pnL   = @(x) Pn(2*phiDE(x)/L+1);                       % representation
```

Of course, one must account for the endpoint values separately in line with (2.1) since the chebfun `Pn` only provides an approximation for $x \in [x_L, 1]$. Similar code can easily be written to deal with situations involving the other three maps.

6. Discussion. Preliminary numerical experiments such as those conducted in the previous section provided the initial stimulus for this investigation. These experiments seemed to suggest that variable transform methods onto a semi-infinite interval were often preferable to those onto an infinite interval. The purpose of this paper has been to attempt to quantify this observation.

The majority of previous investigations concerning exponential or double exponential variable transform methods have been concerned with the process of transplanting the problem function to an infinite interval and subsequently approximating it there using some infinite interval basis set. If one assumes from the outset that the only sensible way to go about such approximation is to use sinc functions or Fourier series, then of course, the notion of transplanting functions to a semi-infinite interval does not arise. This is perhaps the reason that the discussion of transformations to a semi-infinite interval is conspicuously absent from the literature.

The strategy of using Chebyshev interpolants as a viable alternative to sinc functions was recommended in both [4] and [14], and we wish to reinforce that position here. In [14], the primary reason for citing Chebyshev polynomials as a good alternative was essentially ease of computation. Crucially, however, a further advantage of Chebyshev polynomials is their ability to handle semi-infinite interval maps. This is in contrast to sinc, or Fourier discretisations, which are most readily applicable in the infinite interval case. It should be noted that a cursory inspection of the literature (see e.g. [16, 21]) may suggest that the problem of dealing with functions defined on a semi-infinite interval has been effectively addressed. However, one will find that the strategy followed is to transplant again to an infinite interval. This, of course, does not really address the issue of introducing unnecessary waste into the computation on one side of the domain if there happens to be no endpoint singularity present there.

We have shown in Section 3 that one may expect genuinely superior asymptotic convergence rates in the case of the exponential map φ_{E}. For the double-exponential transform φ_{DE}, although the exponent in the asymptotic convergence rate is the same as for ψ_{DE}, the constants involved are better: numerical evidence provided in Table 2 suggests that the resolution properties associated with φ_{DE} are superior.

Our recommendation, then, is to use the map φ_{DE} wherever possible. In practice, this means that if one has a function which has only a single endpoint singularity, then the map φ_{DE} should usually be used. Functions with singularities at both endpoints must of course be dealt with with one of ψ_{E} or ψ_{DE} (at least, if one requires a single global representation), and similarly we recommend the use of the transformation ψ_{DE} wherever possible. The caveat "wherever possible" in both cases refers to the fact that functions can arise in practice which will not be analytic in the regions $\mathcal{R}_\zeta$ or U_ζ, for any ζ, corresponding to the maps φ_{DE} or ψ_{DE}. However, situations where this is the case occur infrequently. One can, if pressed, construct examples of functions that are not analytic in these regions: such functions tend to have an infinite number of singularities which cluster near the wrap-around points of the Riemann surface; Tanaka *et al.* [21] provided examples of this in the infinite interval case. The numerical experiments there also suggest that in such cases the convergence rate reverts back to $C^{-\sqrt{n}}$ anyway; so in practice, one probably need not worry too much about whether or not it is safe to use φ_{DE} or ψ_{DE}.

An additional consideration is that that in double precision arithmetic, we usually cannot approximate functions such as f_9 and f_{10} anyway, even using the maps $\psi_{\text{E}}, \psi_{\text{DE}}$. The reason is the inability of the floating point number system to cope well with singularities at $x = 1$; these matters were discussed in [14], and also previously in [17].

At $x = 0$, the situation is much better, since the density of floating point numbers there is much greater [9, 13]. Though one may not always have a choice, in floating point arithmetic it is always better to position the singularities at $x = 0$ if it is possible to do so. This is also the reason we chose the default problem interval to be $[0, 1]$, rather than, say, $[-1, 1]$. For these reasons, and in light of the demonstrated superiority of the semi-infinite interval maps, we believe that the infinite interval maps $\psi_{\text{E}}, \psi_{\text{DE}}$ should be used infrequently in practice. It only makes sense to use these maps if one is prepared to go beyond the usual floating-point arithmetic.

From a purist's perspective, if one does happen to be approximating on a truncated infinite interval after transplantation using ψ_{E} or ψ_{DE}, then Fourier series are preferential to Chebyshev polynomials, due to the $\pi/2$ saving; this point has also been made previously by Boyd [5]. However, given the availability and convenience of Chebfun, unless one has an excellent reason for using Fourier series, it hardly seems worth going to the extra effort of coding from scratch.

It is also worth noting that we have not provided any converse theorems, so we are not able to conclude that the bounds we have provided are best possible. However, the numerical evidence does suggest that our bounds are sensible. One way to improve the bounds associated with ψ_{E} and ψ_{DE} would be to modify the continuity assumption (2.8) in such a way that differing decay rates of $\mathcal{F}$ as $s \to \pm\infty$ are accounted for. Such steps are often taken in the sinc literature (see e.g. [10, 16, 17]), but since the modifications are trivial, yet cumbersome, and do not affect the asymptotic convergence rates, we have chosen not to do so here in order to expedite the discussion. We also mention once more the optimality result regarding ψ_{DE} due to Sugihara [18].

Acknowledgements. The author would like to thank Nick Trefethen, Alex Townsend and Pedro Gonnet for valuable discussions and input.

REFERENCES

[1] J.-P. Berrut and L. N. Trefethen, *Barycentric Lagrange interpolation*, SIAM Review 46 (2004), 501–517.
[2] J. P. Boyd, *Chebyshev and Fourier spectral methods*, Second Edition, Dover, 2001.
[3] J. P. Boyd, *The optimization of convergence of Chebyshev polynomial methods in an unbounded domain*, J. Comput. Phys. 45, (1982), 43–79.
[4] J. P. Boyd, *Polynomial series versus sinc expansions for functions with corner or endpoint singularities*, J. Comp. Phys. 64 (1986), 266–269.
[5] J. P. Boyd, *Chebyshev domain truncation is inferior to Fourier domain truncation for solving problems on an infinite interval,* J. Sci. Comp. 3 (1988), 109–120.
[6] D. Elliott, *The evaluation of the Chebyshev coefficients in the Chebyshev series expansion of a function*, Math. Comp. 18 (1964), 274–284.
[7] D. Gottlieb and S. A. Orszag, *Numerical analysis of spectral methods: theory and applications*, SIAM, 1977.
[8] S. Hamada, *Numerical solutions of Serrin's equations by transformation,* Publ. RIMS, Kyoto Univ., 43 (2007), 795–817.
[9] N. J. Higham, *Accuracy and stability of numerical algorithms,* Second Edition, SIAM, 2002.
[10] J. Lund and K. Bowers, *Sinc methods for quadrature and differential equations,* SIAM, 1992.
[11] M. Mori, *Discovery of the double exponential transformation and its developments,* Publ. RIMS, Kyoto Univ., 41 (2005), 897–935.
[12] M. Mori and M. Sugihara *The double-exponential transform in numerical analysis,* J. Comp. Appl. Math., 127 (2001), 287–296.
[13] M. L. Overton, *Numerical computing with IEEE floating point arithmetic,* SIAM, 2001.
[14] M. Richardson and L. N. Trefethen, *A sinc function analogue of Chebfun*, SIAM J. Sci. Comp.

33, (2011).
[15] F. Stenger, *Numerical methods based on Whittaker cardinal, or sinc functions*, SIAM Rev., 23 (1981), 165–224
[16] F. Stenger, *Numerical methods based on sinc and analytic functions,* Springer, 1993.
[17] F. Stenger, *Handbook of sinc numerical methods,* CRC Press, 2010.
[18] M. Sugihara, *Optimality of the formula—functional analysis approach,* Numer. Math. 75 (1997), 379–395.
[19] M. Sugihara and M. Takayasu, *Recent developments of the Sinc numerical methods,* J. Comp. App. Math. 164–165 (2004), 673–689.
[20] H. Takahasi and M. Mori, *Double exponential formulas for numerical integration,* Publ. RIMS, Kyoto Univ., 9 (1974), 721–741.
[21] K. Tanaka, M. Sugihara, K. Muroto, and M. Mori, *Function classes for integration formulas,* Numer. Math. 111 (2009), 631–655.
[22] K. Tanaka, M. Sugihara and K. Muroto, *Function classes for successful DE-sinc approximations,* Math. Comp. 78 (2009), 1553–1571.
[23] L. N. Trefethen, *Approximation Theory and Approximation Practice,* manuscript in preparation (2011), draft available at: `http://www.maths.ox.ac.uk/chebfun/ATAP`.
[24] L. N. Trefethen, et al., Chebfun software package, `http://www.maths.ox.ac.uk/chebfun/`.
[25] J. A. C. Weideman and A. Cloot, *Spectral methods and mappings for evolution equations on the infinite line*, Comput. Meths. Appl. Mech. Eng. 80 (1990) 467–481.
[26] J. A. C. Weideman and L. N. Trefethen, *The eigenvalues of second-order spectral differentiation matrices*, SIAM J. Numer. Anal. 25 (1988), 1279–1298.