The Appearance of Big Integers in Exact Real Arithmetic based on Linear Fractional Transformations*

Reinhold Heckmann

FB 14 – Informatik, Universität des Saarlandes
Postfach 151150, D-66041 Saarbrücken, Germany

e-mail: heckmann@cs.uni-sb.de

Abstract. One possible approach to exact real arithmetic is to use linear fractional transformations to represent real numbers and computations on real numbers. In this paper, we show that the bit sizes of the (integer) parameters of nearly all transformations used in computations are proportional to the number of basic computational steps executed so far. Here, a basic step means consuming one digit of the argument(s) or producing one digit of the result.

1 Introduction

Linear Fractional Transformations (LFT's) provide an elegant approach to real number arithmetic [8, 16, 11, 14, 12, 6]. One-dimensional LFT's $x \mapsto \frac{ax+c}{bx+d}$ are used as digits and to implement basic functions, while two-dimensional LFT's $(x, y) \mapsto \frac{axy+cx+ey+g}{bxy+dx+fy+h}$ provide binary operations such as addition and multiplication, and can be combined to infinite expression trees denoting transcendental functions. In Section 2, we present the details of the LFT approach. This provides the background for understanding the results in the remainder of this paper.

LFT's can be modelled within linear algebra. If the four parameters of a one-dimensional LFT are written as a (2,2)-matrix (shortly called *matrix*), functional composition becomes matrix multiplication. Likewise, the eight parameters of a two-dimensional LFT can be written as a (2,4)-matrix (called *tensor*). We refer to matrices and tensors collectively as *transforms*. Basic computational steps such as consuming one digit of the argument(s) (*absorption*) or producing one digit of the result (*emission*) can be realised as variants of matrix multiplication applied to a transform and a digit matrix.

Usually, all the transforms used in real number arithmetic have integer components. Naively, one may think that these components become bigger by absorptions, and become smaller again by emissions. Technically, the components may decrease by *reduction*, i.e., division of all components of the transform by

* Most of the results in this paper were found during a visiting fellowship of the author at Imperial College, London. This visit was organised by Abbas Edalat and funded by EPSRC.

a common factor; as transforms denote rational functions, reduction does not affect their semantics.

Practical experiments have shown, however, that in most cases, the potential for reduction is negligible. The greatest common factor of the components of a transform is usually 1, and in nearly all of the remaining cases, it is just 2. In Sections 3 and 4, we show some upper and lower bounds for common factors. The full proof of the practically observed behaviour is obtained later (Corollary 12 in Section 6.4).

Practical experiments have also shown that in most cases, the bit size of the entries of a transform is roughly equal to the number of emitted digits. The main contribution of this paper is the formalisation (and of course proof) of these practical observations. First, we derive upper bounds for the sizes of the entries of a transform in Section 5. In Section 6, lower bounds for the determinant and the size of the biggest entry are obtained in the case of matrices. Tensors are handled in Section 7. Finally, we discuss these results and their impact on the complexity of real number computation.

2 Exact Real Arithmetic by Linear Fractional Transformations

In this section, we present the framework of exact real arithmetic by LFT's [8, 16, 11]. After a general introduction, we specialise to the version used by the group of Edalat and Potts at Imperial College [14, 12, 13, 15, 6].

2.1 From Digit Streams to Linear Fractional Transformations

There are many ways to represent real numbers as infinite objects [3, 2, 4, 5]. Here, we are only concerned with representations as infinite streams of "digits". These streams are evaluated incrementally; at any given time, only a finite prefix of the stream is known.

There are several different stream representations which can be grouped into two large families: variations of the familiar decimal representation [1, 3, 2, 5, 7, 11, 10], and continued fraction expansions [8, 16, 9].

For the first family, consider the usual decimal representation.[1] A number such as $0.142\cdots$ can be unravelled from left to right as follows:

$$0.142\cdots = \tfrac{1}{10}(1 + 0.42\cdots);\ 0.42\cdots = \tfrac{1}{10}(4 + 0.2\cdots);\ 0.2\cdots = \tfrac{1}{10}(2 + 0.\cdots)$$

Thus, every digit d corresponds to an affine map α_d with $\alpha_d(x) = \frac{1}{10}(d + x) = \frac{x+d}{10}$. A number of the form $0.\cdots$ can be any element of the closed interval $[0, 1]$, and so, a number of the form $0.142\cdots$ can be any element of the interval

[1] This representation is not suitable for practical purposes, as it lacks redundancy, and thus, most arithmetic functions are not computable. However, it provides a familiar example.

$(\alpha_1 \circ \alpha_4 \circ \alpha_2)[0,1] = [0.142, 0.143]$. In general, the infinite stream $0.d_1d_2d_3\cdots$ represents the unique real number in the intersection $\bigcap_{n=1}^{\infty}(\alpha_{d_1} \circ \cdots \circ \alpha_{d_n})[0,1]$.

In the classical continued fraction expansion, irrational numbers in the interval $[0,\infty]$ can be written as $a_0 + \frac{b_0}{a_1 + \frac{b_1}{a_2 + \cdots}}$ with natural numbers a_n and b_n. Every pair $p = (a,b)$ corresponds to the rational function ρ_p with $\rho_p(x) = a + \frac{b}{x} = \frac{ax+b}{x}$. Similar to the case above, an infinite continued fraction corresponds to the intersection $\bigcap_{n=1}^{\infty}(\rho_{p_1} \circ \cdots \circ \rho_{p_n})[0,\infty]$.

The formal similarity between the two approaches presented above leads to the following generalisation [8, 16, 14, 12, 13, 15, 6]: Real numbers in some *base interval* I are represented by infinite streams of digits. Digits are certain *Linear Fractional Transformations* (LFT's) $x \mapsto \frac{ax+c}{bx+d}$, parameterised by numbers a, b, c, d (in practical cases usually integers). The meaning of an infinite stream $\tau_1, \tau_2, \ldots$ of LFT's is the intersection $\bigcap_{n=1}^{\infty}(\tau_1 \circ \cdots \circ \tau_n)(I)$. This intersection is filtered (decreasing) if $\tau_n(I) \subseteq I$ holds for all digits τ_n.

2.2 LFT's and Matrices

Every 2-2-matrix $A = \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right)$ of real numbers denotes an LFT $\langle A \rangle$, which is given by $\langle A \rangle(x) = \frac{ax+c}{bx+d}$. LFT's described by non-singular matrices, i.e., matrices A with determinant $\det A = ad - bc \neq 0$, are considered as endofunctions of $\mathbb{R}^\star = \mathbb{R} \cup \{\infty\}$, the one-point compactification of the real line. The value ∞ arises as $r/0$ with $r \neq 0$, and on the other hand, $\langle A \rangle(\infty)$ is defined to be a/b. For LFT's described by singular matrices, an additional 'number' $\perp$ (undefined) is needed which arises as $0/0$. The value of $\langle A \rangle(\perp)$ is defined to be $\perp$.

The mapping $A \mapsto \langle A \rangle$ is not one-to-one; for, $\langle A \rangle = \langle rA \rangle$ holds for all $r \neq 0$. We shall write $A \cong B$ if $\langle A \rangle = \langle B \rangle$, or equivalently $B = rA$ for some $r \neq 0$. Composition of LFT's can be expressed by matrix multiplication: $\langle A \rangle \circ \langle B \rangle = \langle A \cdot B \rangle$. The equivalence relation '$\cong$' is a congruence w.r.t. multiplication.
The determinant $\det A$ is a well-known property of a matrix A.

$$\det \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right) = ad - bc \qquad \det(A \cdot B) = \det A \cdot \det B \qquad \det(rA) = r^2 \det A \quad (1)$$

By the last equation, the determinant of a matrix is not invariant under equivalence '$\cong$', but its sign (1, 0, or -1) is, i.e., the sign of the determinant of A is a well-defined property of the LFT $\langle A \rangle$. LFT's with non-zero determinant (non-singular LFT's) are invertible; $\langle A \rangle^{-1}$ is given by $\langle A^{-1} \rangle$. Thus, non-singular LFT's form a group under composition.

A *rational LFT* is an LFT which can be represented by a matrix with rational entries, and therefore even by an integer matrix. As $\langle A \rangle = \langle kA \rangle$ for $k \neq 0$, there are infinitely many integer matrices denoting the same rational LFT. An integer matrix is called *k-reducible* if k is a common factor of its four components. Division of a k-reducible matrix by k is called *reduction by k*. A matrix is *in lowest terms* if there is no common factor other than 1 and -1. All integer matrices different from $\left(\begin{smallmatrix} 0 & 0 \\ 0 & 0 \end{smallmatrix}\right)$ are equivalent to an integer matrix in lowest terms.

To obtain an integer representation of $\langle A\rangle^{-1}$ for a non-singular integer matrix A, the *pseudo-inverse* A^* can be used. It is defined by

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix}^* = \begin{pmatrix} d & -c \\ -b & a \end{pmatrix} \tag{2}$$

Clearly, $\det(A^*) = \det A$ holds. The main property of the pseudo-inverse operation is

$$A \cdot A^* = A^* \cdot A = \det A \cdot \mathrm{E} \tag{3}$$

where $\mathrm{E} = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ is the identity matrix, and so, $A \cdot A^* = A^* \cdot A \cong \mathrm{E}$ if $\det A \neq 0$, whence $\langle A\rangle^{-1} = \langle A^*\rangle$.

2.3 The Signed Digit Approach

The group of Edalat and Potts at Imperial College [13, 6] represents the elements of $\mathbb{R}^\star = \mathbb{R} \cup \{\infty\}$ as infinite streams of matrices $S, D_1, D_2, \ldots$, standing for LFT's. The first matrix is a *sign matrix*, while the remaining ones are *digit matrices*. The base interval is $[0, \infty]$, and so, the meaning of the stream is

$$\bigcap_{n=1}^{\infty} \langle S \cdot D_1 \cdot \ldots \cdot D_n\rangle [0, \infty] \ . \tag{4}$$

The base interval $[0, \infty]$ was chosen because there is a simple check for the inclusion property [14]: for a non-singular matrix A, $\langle A\rangle([0, \infty]) \subseteq [0, \infty]$ holds iff all four entries of A are ≥ 0, or all are ≤ 0. Matrices with entries ≥ 0 are called *positive*. Digit matrices are positive, and so, the intersection (4) is filtered (decreasing).

The number set $\mathbb{R}^\star$ can be visualised as a circle. Intervals $[u, v]$ are counter-clockwise arcs from u to v, e.g., $[0, 1] = \{x \in \mathbb{R} \mid 0 \leq x \leq 1\}$, and $[1, 0] = \{x \in \mathbb{R} \mid 1 \leq x \text{ or } x \leq 0\} \cup \{\infty\}$.

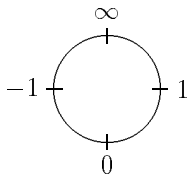

There are four possible sign matrices, corresponding to rotations by 0°, 90°, 180°, and 270°. They can be explicitly described as follows:

$$\begin{aligned}
S_+ &= \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right) & \langle S_+\rangle\,[0, \infty] &= [0, \infty] \\
S_\infty &= \left(\begin{smallmatrix} 1 & 1 \\ -1 & 1 \end{smallmatrix}\right) & \langle S_\infty\rangle [0, \infty] &= [1, -1] \\
S_- &= \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right) & \langle S_-\rangle\,[0, \infty] &= [\infty, 0] \\
S_0 &= \left(\begin{smallmatrix} 1 & -1 \\ 1 & 1 \end{smallmatrix}\right) & \langle S_0\rangle\ [0, \infty] &= [-1, 1]
\end{aligned}$$

S_0 and S_∞ are pseudo-inverse to each other; $S_0 \cdot S_\infty = S_\infty \cdot S_0 = 2\mathrm{E}$ holds.

There are many possible sets of digit matrices, one for every base $r > 1$. Edalat and Potts [6] discuss non-integer bases, but their implementation uses base $r = 2$. In this paper, we consider integer bases $r > 1$.

Fix an integer $r > 1$. Every real number in the interval $[-1, 1]$ has a representation as $\sum_{n=1}^{\infty} k_n r^{-n}$ with integer digits k_n satisfying $|k_n| < r$. (Digits may be negative [1].) As in Section 2.1, these digits correspond to affine maps $\alpha_k^r = \langle A_k^r \rangle$ with $A_k^r = \left(\begin{smallmatrix} 1 & k \\ 0 & r \end{smallmatrix}\right)$.

Since the base interval is not $[-1, 1]$, but $[0, \infty]$, the maps α_k^r have to be transformed into that interval. This can be done by composition with the maps $\langle S_\infty \rangle$ and $\langle S_0 \rangle$, which are mutually inverse bijections between $[-1, 1]$ and $[0, \infty]$. Thus, the actual digit matrices are

$$D_k^r = S_\infty \cdot A_k^r \cdot S_0 = \begin{pmatrix} r+k+1 & r+k-1 \\ r-k-1 & r-k+1 \end{pmatrix} \quad . \tag{5}$$

Since the two entries in the top row differ by 2, these matrices are either in lowest terms or 2-reducible. The latter case occurs iff the parities of r and k are different. In this case, reduction by 2 may be performed. Hence, we distinguish between *unreduced digits* D_k^r and *reduced digits* $\tilde{D}_k^r = \frac{1}{2} D_k^r$. Table 1 illustrates the case $r = 2$. In the column "lowest terms", the first and third matrix ($k \neq 0$) are reduced, while the second matrix ($k = 0$) is unreduced.

Table 1. Digit matrices for base 2

k	A_k^2	D_k^2	lowest terms	$\langle D_k^2 \rangle([0, \infty])$
-1	$\left(\begin{smallmatrix} 1 & -1 \\ 0 & 2 \end{smallmatrix}\right)$	$\left(\begin{smallmatrix} 2 & 0 \\ 2 & 4 \end{smallmatrix}\right)$	$\left(\begin{smallmatrix} 1 & 0 \\ 1 & 2 \end{smallmatrix}\right)$	$[0, 1]$
0	$\left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right)$	$\left(\begin{smallmatrix} 3 & 1 \\ 1 & 3 \end{smallmatrix}\right)$	$\left(\begin{smallmatrix} 3 & 1 \\ 1 & 3 \end{smallmatrix}\right)$	$[\frac{1}{3}, 3]$
1	$\left(\begin{smallmatrix} 1 & 1 \\ 0 & 2 \end{smallmatrix}\right)$	$\left(\begin{smallmatrix} 4 & 2 \\ 0 & 2 \end{smallmatrix}\right)$	$\left(\begin{smallmatrix} 2 & 1 \\ 0 & 1 \end{smallmatrix}\right)$	$[1, \infty]$

2.4 Computation by LFT's

LFT's can not only be used to represent real numbers, but also to perform computations with real numbers. For the sake of simplicity, we only present computations within the interval $[0, \infty]$ where real numbers can be represented by a stream of digit matrices without a leading sign matrix.

Using suitable LFT's $x \mapsto \frac{ax+c}{bx+d}$, basic functions such as $x \mapsto x + 1$, $x \mapsto 2x$, and $x \mapsto \frac{1}{x}$ can be easily expressed. Recall that an LFT maps $[0, \infty]$ into itself iff it can be represented by a positive matrix (all components ≥ 0).

Given a positive matrix M, the actual computation of $\langle M \rangle(x)$ is performed by a sequence of *absorptions* and *emissions*. Absorption means that M consumes the first digit D of x, thereby becoming $M \cdot D$, which is positive again. It corresponds to the equality

$$M \cdot (D_1 \cdot D_2 \cdot \ldots) \;=\; (M \cdot D_1) \cdot (D_2 \cdot \ldots) \quad . \tag{6}$$

Emission means that M produces one further digit D of the result, thereby becoming $D^* \cdot M$. It corresponds to the equivalence

$$(D_1 \cdot \ldots \cdot D_n) \cdot M \;\cong\; (D_1 \cdot \ldots \cdot D_n \cdot D) \cdot (D^* \cdot M) \;. \tag{7}$$

Emission of a digit D is allowed only if $D^* \cdot M$ is positive. Therefore, a possible strategy for the computation of $\langle M \rangle(x)$ is as follows: emit digits until no further emission is possible, then absorb one digit of x, again emit digits until no longer possible, etc.

2.5 Tensors

To compute sums, products, etc., *two-dimensional LFT's* are employed. They are characterised by 8 parameters, and thus can be represented by 2-4-matrices, so called *tensors*. A tensor $T = \left(\begin{smallmatrix} a & c & e & g \\ b & d & f & h \end{smallmatrix}\right)$ denotes the function $\langle T \rangle : \mathbb{R}^\star{}_\perp \times \mathbb{R}^\star{}_\perp \to \mathbb{R}^\star{}_\perp$ given by $\langle T \rangle(x,y) = \frac{axy+cx+ey+g}{bxy+dx+fy+h}$. For integer tensors, the notions of reducible, reduction, and lowest terms can be defined analogous to the case of matrices. Likewise for positivity: a two-dimensional LFT maps $[0,\infty]^2$ to $[0,\infty]_\perp$ iff it can be represented by a positive tensor, i.e., a tensor with components ≥ 0. Because of these analogies, we refer to matrices and tensors collectively as *transforms*.

It is easy to represent addition, subtraction, multiplication, and division by suitable integer tensors [8, 16, 14, 12, 13]. Tensors may also be used to represent transcendental functions, e.g., $\arctan x = \langle T_0 \rangle(x, \langle T_1 \rangle(x, \langle T_2 \rangle(x, \ldots)))$ where $T_n = \left(\begin{smallmatrix} 0 & 1 & 0 & 0 \\ (n+1)^2 & 0 & 0 & 2n+1 \end{smallmatrix}\right)$. It remains to show how to actually compute $\langle T \rangle(x,y)$ for a given positive integer tensor T [12, 13].

Emissions can be done as in the one-dimensional case: in emitting a digit D, tensor T is replaced by $D^* \cdot T$, which is a tensor again. Emission of D is only allowed if $D^* \cdot T$ is positive.

Since digits can be absorbed from both arguments, there are two kinds of *absorptions*: absorption of a digit D from the left argument transforms T into $T \cdot L(D)$, while absorption from the right argument yields $T \cdot R(D)$. Here, $L(D)$ means $D \otimes \mathrm{E}$, and $R(D)$ means $\mathrm{E} \otimes D$. An explicit definition of these operations looks as follows:

$$L\begin{pmatrix} a & c \\ b & d \end{pmatrix} = \begin{pmatrix} a & 0 & c & 0 \\ 0 & a & 0 & c \\ b & 0 & d & 0 \\ 0 & b & 0 & d \end{pmatrix} \qquad R\begin{pmatrix} a & c \\ b & d \end{pmatrix} = \begin{pmatrix} a & c & 0 & 0 \\ b & d & 0 & 0 \\ 0 & 0 & a & c \\ 0 & 0 & b & d \end{pmatrix} \tag{8}$$

They satisfy the following equations:

$$L(A \cdot B) = L(A) \cdot L(B) \qquad R(A \cdot B) = R(A) \cdot R(B) \tag{9}$$

$$L(\mathrm{E}) = R(\mathrm{E}) = \mathrm{E}_4 \qquad L(A) \cdot R(B) = R(B) \cdot L(A) \tag{10}$$

where E_4 denotes the identity 4-4-matrix.

Right absorption can be easily expressed with block matrices. Observe $R(A) = \left(\begin{smallmatrix} A & 0 \\ 0 & A \end{smallmatrix}\right)$ where the four entries are matrices. Likewise, a tensor can be written as a row $(T^{\mathrm{L}}, T^{\mathrm{R}})$ of two matrices, and so

$$(T^{\mathrm{L}}, T^{\mathrm{R}}) \cdot R(A) \;=\; (T^{\mathrm{L}}A,\; T^{\mathrm{R}}A) \;\; . \tag{11}$$

Left and right absorption are closely connected. Let $T^{\times}$ be T with the two middle columns exchanged. Then

$$(T \cdot L(D))^{\times} = T^{\times} \cdot R(D) \qquad (T \cdot R(D))^{\times} = T^{\times} \cdot L(D) \;\; . \tag{12}$$

Later, we shall see that D-emissions and D-absorptions have many properties in common. Thus, we introduce a common name: a *D-transaction* at a transform is either a D-emission or a D-absorption.

3 Small Factors

After a transaction at a transform in lowest terms, the entries of the result may have a non-trivial common factor. The most drastic example is $D^* \cdot D = \det D \cdot \mathrm{E}$ for a digit matrix D. Yet apart from this, practical experience shows that common factors are usually quite small. The goal of this section is to find bounds for such factors. We start off with a property involving determinants.

Proposition 1. *Let A be a matrix, and let B be a transform in lowest terms. Then every common factor of the entries of $A \cdot B$ divides* $\det A$.

Proof. Let g be a common factor of $A \cdot B$, i.e., $A \cdot B = gC$ for some transform C. We may compute:

$$g \cdot (A^* \cdot C) = A^* \cdot gC = A^* \cdot A \cdot B \overset{(3)}{=} (\det A \cdot \mathrm{E}) \cdot B = (\det A) \cdot B \;\; .$$

Hence, g divides $(\det A) \cdot B$. Since B is in lowest terms, g must divide $\det A$. □

For matrices, there is a dual statement with an analogous proof so that we obtain:

Theorem 2. *Let A and B be matrices in lowest terms. Then every common factor of $A \cdot B$ divides both* $\det A$ ***and*** $\det B$.

There is a similar statement for the two versions of multiplying a tensor and a matrix:

Proposition 3. *Let T be a tensor in lowest terms, and M an arbitrary matrix. Then every common factor of $T \cdot L(M)$ or $T \cdot R(M)$ divides* $\det M$.

Proof. We consider the L case; the other one is analogous. If $T \cdot L(M) = gC$ for some tensor C, then

$$\begin{aligned} g \cdot (C \cdot L(M^*)) = T \cdot L(M) \cdot L(M^*) &\overset{(9)}{=} T \cdot L(M \cdot M^*) \\ &\overset{(3)}{=} T \cdot L(\det M \cdot \mathrm{E}) \overset{(10)}{=} T \cdot (\det M \cdot \mathrm{E}_4) = (\det M) \cdot T \end{aligned}$$

Since T is in lowest terms, g divides $\det M$. □

Now, consider a transform T in lowest terms. Let T' be the result of a D-absorption at T, i.e., $T' = T \cdot D$ if T is a matrix, or $T' \in \{T \cdot L(D),\ T \cdot R(D)\}$ if T is a tensor. By Theorem 2 and Proposition 3, any common factor of T' divides $\det D$. If T' is the result of a D-emission at T, i.e., $T' = D^* \cdot T$, then by Prop. 1 any common factor of T' divides $\det D^* = \det D$. Summarising, we obtain:

Theorem 4. *Let T be a transform in lowest terms, and D a digit matrix. After a D-transaction at T, any common factor of the result divides* $\det D$.

How big is $\det D$? Recall the definition of the digit matrices for base r from Section 2.3. As $A_k^r = \left(\begin{smallmatrix} 1 & k \\ 0 & r \end{smallmatrix}\right)$, $\det A_k^r$ is r. Since $\det S_0 = \det S_\infty = 2$, we have $\det D_k^r = \det(S_\infty A_k^r S_0) = 4r$. Therefore, we obtain $\det \tilde{D}_k^r = r$ for reduced digits $\tilde{D}_k^r = \frac{1}{2} D_k^r$.

Corollary 5. *Let T be a transform in lowest terms, and D a digit matrix for base r. After a D-transaction at T, any common factor of the result divides $4r$ if D is unreduced, and even divides r if D is reduced.*

Specialising to the case $r = 2$, we see that any common factor of the result divides 2 in case of a transaction with a non-zero digit ($k \neq 0$), and divides 8 in case of $k = 0$.

Corollary 12 in Section 6.4 shows that in many cases, the result of Corollary 5 can be strengthened from $4r$ (r) to 2 (1), ruling out most reductions.

4 Possibilities for Reductions

In the last section, we have seen that there is not much potential for reductions. Here, we show a result of opposite flavour: certain reductions are always possible.

Consider unreduced digit matrices $D_k^r = S_\infty A_k^r S_0$. We have already mentioned that some of them are in lowest terms, while others are 2-reducible; higher reducibilities do not occur. Multiplying two digit matrices yields:

$$D_k^r D_{k'}^{r'} = S_\infty A_k^r S_0 S_\infty A_{k'}^{r'} S_0 = 2 S_\infty A_k^r A_{k'}^{r'} S_0 = 2 D_{kr'+k'}^{rr'} \tag{13}$$

Here, the second equality is due to $S_0 S_\infty = 2\mathrm{E}$, and the third due to

$$A_k^r \cdot A_{k'}^{r'} = \begin{pmatrix} 1 & k \\ 0 & r \end{pmatrix} \cdot \begin{pmatrix} 1 & k' \\ 0 & r' \end{pmatrix} = \begin{pmatrix} 1 & k' + kr' \\ 0 & rr' \end{pmatrix} \tag{14}$$

together with the estimation $|kr' + k'| \leq (r-1)r' + (r'-1) = rr' - 1$.
Iterating (13) leads to

$$D^r_{k_1} \cdot \ldots \cdot D^r_{k_n} = 2^{n-1} D^{r^n}_k \text{ where } k = \sum_{i=1}^{n} k_i r^{n-i} \quad . \tag{15}$$

Hence, we obtain:

1. The product of n digit matrices is always 2^{n-1}-reducible.
2. After 2^{n-1}-reduction, the result is again a digit matrix, and so it is either in lowest terms or 2-reducible.

The result of applying n_1 absorptions and n_2 emissions of unreduced digits to a matrix M has form $A^*_2 \cdot M \cdot A_1$ where A_i is a product of n_i digit matrices. Thus, the result has a common factor of $2^{n_1-1} \cdot 2^{n_2-1} = 2^{n_1+n_2-2}$. For a tensor T, we obtain a result of the form $A^*_3 \cdot T \cdot L(A_2) \cdot R(A_1)$, and thus a common factor of $2^{n_1+n_2+n_3-3}$.

Theorem 6. *Let T_0 be some initial transform, and T_n the result of applying n transactions with unreduced digits to T_0. Then T_n is at least 2^{n-2}-reducible in case of matrices, and at least 2^{n-3}-reducible in case of tensors.*

5 An Upper Bound for the Entries

Next, we derive an exponential upper bound for the entries of a transform after n transactions. An estimate for the entries is the maximum of their absolute values: $\left\|\left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right)\right\| = \max(|a|, |b|, |c|, |d|)$ for matrices, and analogously for tensors, and vectors $\left(\begin{smallmatrix} a \\ b \end{smallmatrix}\right)$. Let us consider how this norm is affected by emissions.

Recall the definition of the digit matrices for base r (Equation (5) in Section 2.3):

$$D^r_k = \begin{pmatrix} r+k+1 & r+k-1 \\ r-k-1 & r-k+1 \end{pmatrix} \quad . \tag{16}$$

Consider the product of $(D^r_k)^*$ with a vector $\left(\begin{smallmatrix} u \\ v \end{smallmatrix}\right)$:

$$(D^r_k)^* \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} 1-k+r & 1-k-r \\ 1+k-r & 1+k+r \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} (1-k)(u+v) + r(u-v) \\ (1+k)(u+v) - r(u-v) \end{pmatrix} \tag{17}$$

Using $|k| < r$, we obtain

$$\left\|(D^r_k)^* \begin{pmatrix} u \\ v \end{pmatrix}\right\| \;\leq\; (1 + |k| + r)(|u| + |v|) \;\leq\; 2r \left\|\begin{pmatrix} u \\ v \end{pmatrix}\right\| \tag{18}$$

Since the norm of a transform is the maximum of the norms of its column vectors, we obtain $\|(D^r_k)^* \cdot T\| \leq 2r\|T\|$ — for unreduced digits. For reduced digits, the right hand side is $r\|T\|$.

Now, let us study absorption. For the absorption of a digit into a matrix, it suffices to consider products $(u,\, v) \cdot D_k^r$ of a row vector and a digit matrix.

$$(u,\, v) \begin{pmatrix} r+k+1 & r+k-1 \\ r-k-1 & r-k+1 \end{pmatrix} = (r(u+v)+(k+1)(u-v),\ r(u+v)+(k-1)(u-v))$$

By an estimation as above, we obtain $\|M \cdot D_k^r\| \leq 2r\|M\|$ for matrices M. By (11), the block formula for right absorption into a tensor, an analogous result holds for $\|T \cdot R(D_k^r)\|$, and by (12), the formula connecting left and right absorption, the same holds for $\|T \cdot L(D_k^r)\|$. Summarising, we obtain:

Proposition 7. *Let T be a transform, D a digit matrix for base r, and T' the result of a D-transaction at T. Then $\|T'\| \leq 2r\|T\|$ if D is unreduced, and $\|T'\| \leq r\|T\|$ if D is reduced.*

By induction, we see that after n transactions, $\|T'\| \leq (2r)^n\|T\|$ holds if unreduced digits are used. Applying all the reductions that are possible by Theorem 6, we obtain:

Theorem 8. *Let T_0 be some initial transform, and T_n the result of applying n transactions in base r to T_0, and all possible reductions. Then $\|T_n\| \leq 4r^n\|T_0\|$ in case of matrices, and $\|T_n\| \leq 8r^n\|T_0\|$ in case of tensors.*

In the moment, there is some hope that further reductions may lead to a much smaller increase. Unfortunately, we shall soon see that this does not work; in most cases, an exponential increase is guaranteed.

6 Big Numbers in Matrices

In this section, we derive lower bounds for the entries of a matrix after n transactions and all possible reductions. This is done by observing how the determinant and another quantity, the column difference, are changed by transactions and reductions, and by deriving a reduction invariant from this.

6.1 Determinant

Determinants are easy because of $\det(A \cdot B) = \det A \cdot \det B$. The determinants of the digit matrices and their pseudo-inverses are calculated in Section 3 just before Corollary 5. In the following list, let M be a matrix, and let M' be the result of applying a transaction to M.

- Transaction with an unreduced digit: $\det M' = 4r \det M$,
- Transaction with a reduced digit: $\det M' = r \det M$,
- Reduction by k: $\det M' = \frac{1}{k^2} \det M$.

These facts allow the derivation of an upper bound for the determinant after n transactions. Working with unreduced digits gives a factor of $(4r)^n$, and performing all reductions admitted by Theorem 6 gives a factor of $2^{-2(n-2)}$. Together, we get the following:

Theorem 9. *Let M_0 be some initial matrix, and M_n the result of applying n transactions in base r to M_0, and all possible reductions. Then* $|\det M_n| \leq 16r^n |\det M_0|$.

6.2 Column Difference

Consider again the explicit formulae for digit matrices of base r and their inverses (Equation (5) in Section 2.3):

$$D_k^r = \begin{pmatrix} r+k+1 & r+k-1 \\ r-k-1 & r-k+1 \end{pmatrix} \qquad (D_k^r)^* = \begin{pmatrix} 1-k+r & 1-k-r \\ 1+k-r & 1+k+r \end{pmatrix} \tag{19}$$

It is easy to see that in both cases the difference of the two column sums is 0. This motivates the definition of the *column difference* $\operatorname{cd}\left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right) = (a+b)-(c+d)$ of a matrix. Thus, $\operatorname{cd} D_k^r = \operatorname{cd}(D_k^r)^* = 0$. In general, $\operatorname{cd} A^* = -\operatorname{cd} A$ holds.

Let us compute the column difference of the product of $A = \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right)$ and $B = \left(\begin{smallmatrix} a' & c' \\ b' & d' \end{smallmatrix}\right)$:

$$\begin{aligned} \operatorname{cd}(A \cdot B) &= \operatorname{cd} \begin{pmatrix} aa'+cb' & ac'+cd' \\ ba'+db' & bc'+dd' \end{pmatrix} \\ &= (a+b)a' + (c+d)b' - (a+b)c' - (c+d)d' \\ &= (a+b)(a'-c') - (c+d)(d'-b') \end{aligned}$$

If $B = D_k^r$, then $a'-c' = d'-b' = 2$, and so, $\operatorname{cd}(A \cdot D_k^r) = 2 \operatorname{cd} A$. If $A = (D_k^r)^*$, then $a+b = c+d = 2$, and so, $\operatorname{cd}((D_k^r)^* \cdot B) = 2 \operatorname{cd} B$. If reduced digits are used instead, the factor 2 disappears. Thus, we obtain:

- Transaction with an unreduced digit: $\operatorname{cd} M' = 2 \operatorname{cd} M$,
- Transaction with a reduced digit: $\operatorname{cd} M' = \operatorname{cd} M$,
- Reduction by k: $\operatorname{cd} M' = \frac{1}{k} \operatorname{cd} M$.

Hence, the properties of having zero or non-zero column difference are transaction invariants.

6.3 The Quotient

Let M be a matrix with $\operatorname{cd} M \neq 0$. For such a matrix, the quotient $\operatorname{qcd} M = \frac{\det M}{(\operatorname{cd} M)^2}$ is a well-defined rational number. By a transaction with an unreduced digit, this quotient is multiplied by $\frac{4r}{2^2} = r$; by a transaction with a reduced digit, the factor is $\frac{r}{1^2} = r$; and a k-reduction yields a factor of $\frac{1/k^2}{(1/k)^2} = 1$. Thus, the quotient qcd is invariant under reductions, and is multiplied by r in every transaction.

Lemma 10. *Let M_0 be some initial matrix with* $\operatorname{cd} M_0 \neq 0$, *and M_n the result of applying n transactions in base r to M_0, and all possible reductions. Then* $\operatorname{qcd} M_n = r^n \operatorname{qcd} M_0$.

6.4 Big Determinant

The equation in Lemma 10 can be turned into an integer equation by multiplying with the denominators:

$$\det M_n \cdot (\mathrm{cd}\, M_0)^2 = r^n \cdot \det M_0 \cdot (\mathrm{cd}\, M_n)^2 \tag{20}$$

If $\mathrm{cd}\, M_0 \neq 0$, then $\mathrm{cd}\, M_n \neq 0$, too. As an integer, $(\mathrm{cd}\, M_n)^2$ is at least 1. Hence, we obtain:

$$|\det M_n| \cdot (\mathrm{cd}\, M_0)^2 \geq r^n \cdot |\det M_0| \tag{21}$$

This gives a lower bound for the determinant; an upper bound was provided by Theorem 9.

Theorem 11. *Let M_0 be some initial matrix with $\mathrm{cd}\, M_0 \neq 0$, and M_n the result of applying n transactions in base r to M_0, and all possible reductions. Then*

$$\frac{|\det M_0|}{(\mathrm{cd}\, M_0)^2} \cdot r^n \;\leq\; |\det M_n| \;\leq\; 16|\det M_0| \cdot r^n \quad .$$

The upper bound was obtained by working with unreduced digits and performing the 2^{n-1}-reduction guaranteed by Theorem 6. In case of $\det M_0 \neq 0$, the quotient of upper bound over lower bound shows that only a constant number of further reductions is possible; they combine to a factor of at most $4\,\mathrm{cd}\, M_0$. This implies the promised strengthening of Corollary 5:

Corollary 12. *When working with a matrix with non-zero determinant and column difference, the average maximal reducibility is 2 after a transaction with an unreduced digit, and 1 after a transaction with a reduced digit.*

6.5 Law of Big Numbers for Matrices

A lower bound for the determinant of a matrix M can be turned into a lower bound for the norm $\|M\|$ using the inequality $\|M\| \geq \sqrt{\frac{1}{2}|\det M|}$, which follows from the definition of the determinant as $\det \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right) = ad - bc$. Thus, we obtain together with Theorem 8:

Theorem 13. *Let M_0 be some initial matrix with $\mathrm{cd}\, M_0 \neq 0$, and M_n the result of applying n transactions in base r to M_0, and all possible reductions. Then*

$$\sqrt{\frac{|\det M_0|}{2(\mathrm{cd}\, M_0)^2}} \cdot \left(\sqrt{r}\right)^n \;\leq\; \|M_n\| \;\leq\; 4\|M_0\| \cdot r^n \quad .$$

Thus, if in addition $\det M_0 \neq 0$, *even if all possible reductions are performed, the entries of the matrix are bound to grow exponentially in the number of transactions.*

It sounds a bit more optimistically to speak of the bit sizes of the entries instead of the entries themselves. The bit size of a number m is $\log m$.

Theorem 14 (Law of big numbers).
Let M be a matrix with non-zero determinant and non-zero column difference. After n transactions at M, at least one entry of the result has bit size $\Omega(n)$, even if all possible reductions are performed.

The law of big numbers means that the usage of big integers is unavoidable in exact real arithmetic, at least in the signed digit approach of Edalat's group. It applies even in the simplest cases. For instance, doubling of an unsigned real is effected by the matrix $\left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ that has determinant 2 and column difference 1, halfing by $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right)$ with determinant 2 and column difference -1, and addition of 1 by the matrix $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ with determinant 1 and column difference -1.

The law of big numbers does not apply to matrices with zero column difference. The simplest example is the identity matrix $\mathrm{E} = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$. According to (3), after a D-absorption, a subsequent D-emission, and a reduction by $\det D$, the identity matrix is recovered. Repeating this cycle, we see that there are arbitrarily long sequences of transactions at the identity matrix which do not lead to entries bigger than $4r$. It is an open problem whether such a fixed bound can be found for any matrix with column difference 0.

7 Big Numbers in Tensors

In this section, we derive analogues of the results of the previous section for tensors. The proceeding is similar, but a major obstacle is that tensors do not have determinants. Fortunately, a suitable substitute can be found.

7.1 Double Column Difference

We start by introducing an analogue to the column difference of a matrix. For a tensor T, the *double column difference* $\operatorname{dcd} T$ is defined by

$$\operatorname{dcd} \begin{pmatrix} a & c & e & g \\ b & d & f & h \end{pmatrix} = (a+b) - (c+d) - (e+f) + (g+h) \quad . \tag{22}$$

Writing a tensor T as a row $(T^{\mathrm{L}}, T^{\mathrm{R}})$ of two matrices, the double column difference can be reduced to the column differences of the two matrices: $\operatorname{dcd}(T^{\mathrm{L}}, T^{\mathrm{R}}) = \operatorname{cd} T^{\mathrm{L}} - \operatorname{cd} T^{\mathrm{R}}$. Hence, by (11) and the properties of cd, we obtain for all digit matrices D

$$\operatorname{dcd}((T^{\mathrm{L}}, T^{\mathrm{R}}) \cdot R(D)) = \operatorname{cd}(T^{\mathrm{L}} \cdot D) - \operatorname{cd}(T^{\mathrm{R}} \cdot D) = 2 \operatorname{dcd}(T^{\mathrm{L}}, T^{\mathrm{R}}) \quad .$$

By $(T \cdot R(D))^{\times} = T^{\times} \cdot L(D)$ (12) and $\operatorname{dcd}(T^{\times}) = \operatorname{dcd} T$, we obtain the corresponding formula $\operatorname{dcd}(T \cdot L(D)) = 2 \operatorname{dcd} T$.
We still have to derive a formula for emission. Recall (17)

$$(D_k^r)^* \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} (1-k)(u+v) + r(u-v) \\ (1+k)(u+v) - r(u-v) \end{pmatrix} \tag{23}$$

which implies

$$(D_k^r)^* \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} u' \\ v' \end{pmatrix} \implies u' + v' = 2(u+v) \ . \tag{24}$$

From this, $\operatorname{dcd}(D^* \cdot T) = 2 \operatorname{dcd} T$ follows for all digit matrices D. Therefore, dcd for tensors behaves exactly as cd for matrices:

- Transaction with an unreduced digit: $\operatorname{dcd} T' = 2 \operatorname{dcd} T$,
- Transaction with a reduced digit: $\operatorname{dcd} T' = \operatorname{dcd} T$,
- Reduction by k: $\operatorname{dcd} T' = \frac{1}{k} \operatorname{dcd} T$.

Again, the properties of having zero or non-zero double column difference are transaction invariants.

7.2 Column Determinant

A suitable substitute for the determinant of a matrix is the *column determinant* $\operatorname{cdet} T$ of a tensor T, defined by

$$\operatorname{cdet} \begin{pmatrix} a & c & e & g \\ b & d & f & h \end{pmatrix} = (a+b)(g+h) - (c+d)(e+f) \ . \tag{25}$$

Because of (24), $\operatorname{cdet}(D^* \cdot T) = 4 \operatorname{cdet} T$ holds for all tensors T and digit matrices D. Note that in contrast to the determinant of matrices, the factor is not $\det D^* = 4r$, but only 4. On the other side, the column determinant is multiplicative w.r.t. absorptions; for any tensor T and matrix M,

$$\operatorname{cdet}(T \cdot L(M)) = \operatorname{cdet}(T \cdot R(M)) = \operatorname{cdet} T \cdot \det M \tag{26}$$

holds. Here, the first equality follows from (12) and $\operatorname{cdet}(T^\times) = \operatorname{cdet} T$, while the proof of the second equality is a straightforward, but tedious exercise in algebraic manipulations.
Summarising and specialising to the case of digit matrices, we obtain:

- Emission of an unreduced digit: $\operatorname{cdet} T' = 4 \operatorname{cdet} T$,
- Emission of a reduced digit: $\operatorname{cdet} T' = \operatorname{cdet} T$,
- Absorption of an unreduced digit: $\operatorname{cdet} T' = 4r \operatorname{cdet} T$,
- Absorption of a reduced digit: $\operatorname{cdet} T' = r \operatorname{cdet} T$,
- Reduction by k: $\operatorname{cdet} T' = \frac{1}{k^2} \operatorname{cdet} T$.

In contrast to matrices, emissions and absorptions behave differently.

7.3 The Quotient

For a tensor T with $\operatorname{dcd} T \neq 0$, we consider the quotient $\operatorname{qdcd} T = \frac{\operatorname{cdet} T}{(\operatorname{dcd} T)^2}$. This quotient is invariant under reductions and also invariant under emissions. Every absorption yields a factor of r.

Lemma 15. *Let T_0 be some initial tensor with $\operatorname{dcd} T_0 \neq 0$, and T_n the result of applying n absorptions, any number of emissions, and all possible reductions to T_0. Then $\operatorname{qdcd} T_n = r^n \operatorname{qdcd} T_0$.*

As in the case of matrices, a lower bound for the column determinant follows:

Theorem 16. *Let T_0 be some initial tensor with $\operatorname{dcd} T_0 \neq 0$, and T_n the result of applying n absorptions, any number of emissions, and all possible reductions to T_0. Then*

$$\operatorname{cdet} T_n \;\geq\; \frac{|\operatorname{cdet} T_0|}{(\operatorname{dcd} T_0)^2} \cdot r^n \quad .$$

7.4 Law of Big Numbers for Tensors

For tensors T, $\|T\| \geq \frac{1}{2}\sqrt{\frac{1}{2}|\operatorname{cdet} T|} = \sqrt{\frac{1}{8}|\operatorname{cdet} T|}$ holds. Thus, we obtain together with Theorem 8:

Theorem 17. *Let T_0 be some initial tensor with $\operatorname{dcd} T_0 \neq 0$, and T_n the result of applying n absorptions, any number of emissions, and all possible reductions to T_0. Then*

$$\sqrt{\frac{|\operatorname{cdet} T_0|}{8(\operatorname{dcd} T_0)^2}} \cdot \left(\sqrt{r}\right)^n \;\leq\; \|T_n\| \;\leq\; 8\|T_0\| \cdot r^n \quad .$$

Theorem 18 (Law of big numbers for tensors).
Let T be a tensor with non-zero column determinant and non-zero double column difference. After n absorptions and any number of emissions at T, at least one entry of the result has bit size $\Omega(n)$, even if all possible reductions are performed.

7.5 Examples

The tensors that realise the four basic arithmetic operations satisfy the hypotheses of the law of big numbers:

Addition:	$\left(\begin{smallmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{smallmatrix}\right)$	$\operatorname{cdet} = -1$	$\operatorname{dcd} = -1$
Subtraction:	$\left(\begin{smallmatrix} 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{smallmatrix}\right)$	$\operatorname{cdet} = 1$	$\operatorname{dcd} = 1$
Multiplication:	$\left(\begin{smallmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{smallmatrix}\right)$	$\operatorname{cdet} = 1$	$\operatorname{dcd} = 2$
Division:	$\left(\begin{smallmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{smallmatrix}\right)$	$\operatorname{cdet} = -1$	$\operatorname{dcd} = -2$

Yet the tensor for the mean value operation is different:

Mean value:	$\left(\begin{smallmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{smallmatrix}\right)$	$\operatorname{cdet} = -1$	$\operatorname{dcd} = 0$

Does this mean that $\frac{1}{2}x$, which leads to big numbers as shown in Section 6.5, can be computed as $\frac{0+x}{2}$ avoiding big numbers? The answer is no, at least in the case $r = 2$. Let T^{R} be the matrix on the right hand side of the tensor T. The equations $(D^* \cdot T)^{\mathrm{R}} = D^* \cdot T^{\mathrm{R}}$ and $(T \cdot R(D))^{\mathrm{R}} = T^{\mathrm{R}} \cdot D$ hold for all tensors

T and digit matrices D. This means that the right half of $\left(\begin{smallmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{smallmatrix}\right)$ behaves exactly as the halfing matrix $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right)$ during emissions and absorptions from the right. Since the number 0 is represented by the infinite product $(\tilde{D}^2_{-1})^\omega$, and $(T \cdot L(\tilde{D}^2_{-1}))^{\mathrm{R}} = 2T^{\mathrm{R}}$, the correspondence is only changed by a common factor during absorptions from the left. Hence, after any number of transactions, the right half of the resulting tensor is a multiple of the matrix resulting from $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right)$ by the corresponding sequence of transactions. Thus, it has entries which are at least as big as the entries of the matrix, which are big by Theorem 14.

8 Discussion and Conclusion

The laws of big numbers as derived in this paper apply to unsigned reals only. For instance, halfing in the zero interval $[-1, 1]$ with base $r = 2$ means putting D^2_0 in front of the unsigned part of the argument, an operation possible without employing big integers.

Of course, our results crucially depend on the choice of the digit matrices. All digit matrices for all bases have zero column difference, and this fact is implicitly used in the derivations of the formulae for the cd and dcd values after transactions. A completely different choice of digit matrices, with non-zero column difference, may change everything. Also, the results may look different if irrational bases are used such as the golden ratio. However, we believe that big numbers cannot be avoided even in these cases, although we do not have a proof.

The appearance of big integers affects the complexity of real number arithmetic. Consider an LFT satisfying the hypotheses of the laws of big numbers. If it absorbs and emits digits one by one, then the nth transaction needs time $\Omega(n)$ since it involves integers of bit size $\Omega(n)$. Consequently, the computation of the first n digits of the result of the LFT needs time $\Omega(n^2)$. This time can only be reduced by replacing the one by one treatment of digits by algorithms absorbing and emitting many digits at once. Of course, the price for this reduction in time are much more involved algorithms.

References

1. A. Avizienis. Signed-digit number representations for fast parallel arithmetic. *IRE Transactions on Electronic Computers*, 10:389–400, 1961.
2. H.J. Boehm, R. Cartwright, M. Riggle, and M.J. O'Donell. Exact real arithmetic: A case study in higher order programming. In *ACM Symposium on Lisp and Functional Programming*, 1986.
3. H.J. Boehm and R. Cartwright. Exact real arithmetic: Formulating real numbers as functions. In D. Turner, editor, *Research Topics in Functional Programming*, pages 43–64. Addison-Wesley, 1990.
4. P. Di Gianantonio. *A Functional Approach to Real Number Computation.* PhD thesis, University of Pisa, 1993.

5. P. Di Gianantonio. Real number computability and domain theory. *Information and Computation*, 127(1):11–25, May 1996.
6. A. Edalat and P. Potts. A new representation for exact real numbers. In S. Brookes and M. Mislove, editors, *MFPS '97*, volume 6 of *Electronic Notes in Theoretical Computer Science*, 1997. URL: `http://www.elsevier.nl/locate/entcs/volume6.html`.
7. M. H. Escardó. PCF extended with real numbers. *Theoretical Computer Science*, 162(1):79–115, August 1996.
8. W. Gosper. Continued fraction arithmetic. Technical Report HAKMEM Item 101B, MIT Artificial Intelligence Memo 239, MIT, 1972.
9. P. Kornerup and D. W. Matula. Finite precision lexicographic continued fraction number systems. In *Proc. 7th IEEE Symposium on Computer Arithmetic*, pages 207–214. IEEE Computer Society Press, 1985.
10. V. Menissier-Morain. Arbitrary precision real arithmetic: Design and algorithms. *submitted to J. Symbolic Computation*, 1996.
11. A. Nielsen and P. Kornerup. MSB-first digit serial arithmetic. *J. of Univ. Comp. Scien.*, 1(7), 1995.
12. P. J. Potts and A. Edalat. Exact real arithmetic based on linear fractional transformations. Draft, Imperial College, available from `http://www-tfm.doc.ic.ac.uk/~pjp`, December 1996.
13. P. J. Potts and A. Edalat. Exact real computer arithmetic. Draft, Imperial College, available from `http://www-tfm.doc.ic.ac.uk/~pjp`, March 1997.
14. P. J. Potts. Computable real arithmetic using linear fractional transformations. Draft PhD Thesis, Imperial College, available from `http://www-tfm.doc.ic.ac.uk/~pjp`, June 1996.
15. P. Potts, A. Edalat, and M. Escardó. Semantics of exact real arithmetic. In *Twelfth Annual IEEE Symposium on Logic in Computer Science*. IEEE, 1997.
16. J. E. Vuillemin. Exact real computer arithmetic with continued fractions. *IEEE Transactions on Computers*, 39(8):1087–1105, 1990.