DIVISOR OPERATOR

LINAS VEPSTAS

ABSTRACT. The "divisor operator" is defined as an infinite-dimensional matrix operator, encoding Dirichlet convolution in a linear algebra setting. The finite-dimensional variant is known as the Redheffer matrix. As a matrix operator, it naturally acts on the Banach space ℓ_1 of summable sequences. On this space, it is not a bounded operator. It's point spectrum consists of all completely multiplicative arithmetic series (that are ℓ_1-summable).

INTRODUCTION

This reviews various low-grade ideas surrounding something that I want to call the "divisor operator". It is a variation on a "Redheffer matrix".[2] It's catnip for the following reason: some of it's eigenvalues are parameterized by the Riemann zeta function. One might vaguely hope that exploration in this relatively novel direction can shed insight on the Riemann hypothesis. In practice, it's more of an exercise bridging across several well-established mathematical disciplines.

DIRICHLET CONVOLUTION

Define $f(n)$ and $g(n)$ arithmetic sequences, that is, functions from the natural numbers $\mathbb{N}$ to the complex numbers $\mathbb{C}$. Define Dirichlet convolution as usual:

$$(f*g)(n)=\sum_{d|n}f(d)\,g\left(\frac{n}{d}\right)$$

This is commutative, associative, Abelian, so its a form of "addition". There is a concept of zero. The zero is

$$\varepsilon(n)=\begin{cases}1 & \text{if } n=1\\ 0 & \text{otherwise}\end{cases}$$

so that $(f*\varepsilon)=(\varepsilon*f)=f$ for any f. It is conventional to define the one-unit as

$$\mathit{1}(n)=1 \quad \text{for all } n$$

Define the Mobius μ arithmetic sequence as minus-one, i.e.

$$\mu*\mathit{1}=\varepsilon$$

As such, Dirichlet convolution, together with ε, $\mathit{1}$ and μ can be used to create a model of arithmetic. The word "model" is used here in the model-theoretic sense. A model of arithmetic is a mapping from the axioms of arithmetic (together with the logical connectives and other parts of first-order logic, *e.g.* the tautologies) to the space of countably-infinite sequences. As a model, it contains "nothing more" than arithmetic; yet, because the space of countably-infinite sequences is much richer than the set of natural numbers, it suggests ways in which "ordinary arithmetic" can be extended in unusual ways. The reader is encouraged to loosely envision Cohen's "forcing" as an inspiration.

Date: January 2019.

To be precise, there are several models: the structure $(*,\varepsilon,\mathit{1})$ provides a model for the non-negative integers under addition$(+,0,1)$. The structure $(*,\mu,\varepsilon,\mathit{1})$ is an extension that provides a model for all of the integers $\mathbb{Z}=(+,-1,0,1)$ where $\mathbb{Z}$ is understood to stand for the first-order language $\mathscr{L}_{\mathbb{Z}}$ of expressions obtained from these operators and constants, together with equality as a relation. As usual, "less than" $<$ provides a total order on the integers.

Conjecture. *The relation $<$ on the integers embeds into the lexicographic ordering of strings.*

Proof. The above is surely a lemma not a conjecture, but I'm super-lazy right now to super-triple-check. Right? "Homework." □

As a matrix operator. The one-unit can also be written as an infinite-dimensional matrix. Picking the letter C for "convolution", write

$$(f*\mathit{1})(n)=\sum_{d|n}f(d)=\sum_{k=1}^{\infty}C_{nk}f(k)$$

with C_{nk} having matrix entries

$$C_{nk}=\begin{cases}1 & \text{if } k \text{ divides } n\\ 0 & \text{otherwise}\end{cases}$$

The visually explicit lower-triangular form is

$$C=\begin{bmatrix}1 &&&&&&\\ 1&1&&&&&\\ 1&0&1&&&&\\ 1&1&0&1&&&\\ 1&0&0&0&1&&\\ 1&1&1&0&0&1&\\ \vdots&&&&&&\ddots\end{bmatrix}$$

That is, the k'th column is just a repeating pattern with a one in it for every multiple of k. This can be emphasized by using the unit sequence $\mathit{1}=\mathit{1}(n)$ in the columns, and writing:

$$C_{nk}=\begin{cases}\mathit{1}\left(\frac{n}{k}\right) & \text{if } k \text{ divides } n\\ 0 & \text{otherwise}\end{cases}$$

This helps make clear that only multiples of k show up, and there are zeros otherwise. It also helps make clear that C can be decomposed as a sum of "diagonal rays", each "ray" progressing along $n/k=const$. The two most obvious rays are the diagonal: $n/k=1$ and the first column: $n/k=\infty$. The next obvious ray is $n/k=2$ which is just saying "down two, over one, place a 1 in that spot" can be iterated. The rays of non-zero entries are illustrated below. All matrix entries not on a ray are zero.

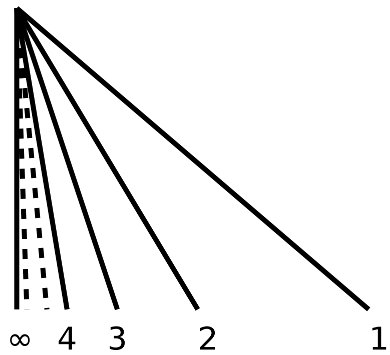

It can be more convenient to work with the transpose, which is upper-triangular, defined as

$$D = C^{\top}$$

The letter D stands for "divisor", as it has obvious relationships to the divisor function. One reason that D is nicer is that it maps finite sequences to finite sequences of the same length. That is, if $f(n) = 0$ for all $n \geq L$ so that f is a finite sequence of length L, then Df is also a sequence of length L.

Much of what follows will be concerned with infinite sequences, including the space of summable sequences; the interplay with finite sequences (which are manifestly summable) provides a nice counterpoint.

Mobius mu as Inverse. Visual representations provide a very useful and convenient means of thinking about the meaning of formulas. So, write D and D^{-1} bigger:

$$D = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & \cdots \\ & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & \\ & & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & \\ & & & 1 & 0 & 0 & 0 & 1 & 0 & 0 & \\ & & & & 1 & 0 & 0 & 0 & 0 & 1 & \\ & & & & & 1 & 0 & 0 & 0 & 0 & \\ & & & & & & 1 & 0 & 0 & 0 & \\ & & & & & & & 1 & 0 & 0 & \\ & & & & & & & & 1 & 0 & \\ & & & & & & & & & 1 & \\ & & & & & & & & & & \ddots \end{bmatrix}$$

and

$$D^{-1} = \begin{bmatrix} 1 & -1 & -1 & 0 & -1 & 1 & -1 & 0 & 0 & 1 & \cdots \\ & 1 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & -1 & \\ & & 1 & 0 & 0 & -1 & 0 & 0 & -1 & 0 & \\ & & & 1 & 0 & 0 & 0 & -1 & 0 & 0 & \\ & & & & 1 & 0 & 0 & 0 & 0 & -1 & \\ & & & & & 1 & 0 & 0 & 0 & 0 & \\ & & & & & & 1 & 0 & 0 & 0 & \\ & & & & & & & 1 & 0 & 0 & \\ & & & & & & & & 1 & 0 & \\ & & & & & & & & & 1 & \\ & & & & & & & & & & \ddots \end{bmatrix}$$

The matrix entries of D^{-1} are given by the Mobius μ function

$$\left[D^{-1}\right]_{nk}=\begin{cases}\mu\left(\frac{k}{n}\right) & \text{if } n \text{ divides } k\\ 0 & \text{otherwise}\end{cases}$$

Just as before, each row is just the Mobius μ with zeros interspersed. (I'm getting a feeling of *deja vu*, did I not write this same text ten years ago? Where is it? It dead-ended back then...)

Note that D^{-1} is both a left and right inverse, because Dirichlet convolution is commutative. That is, $D^{-1}D=DD^{-1}=I$ with I the identity matrix.

Just as the structure $(*,\mu,\varepsilon,\mathit{1})$ extracted from Dirichlet convolution provides a model of the structure $(+,-1,0,1)$ *i.e.* a model of arithmetic, so also the structure $\left(\cdot,D^{-1},I,D\right)$ also provides a model of arithmetic. In what follows, there will be some tension and confusion between these two models. The extension of relation $<$ on the integers to a partial order on infinite-dimensional operators is a bit less obvious.

Classical identities and fun facts. Clearly, all of the classical, well-known identities involving Dirichlet series, Lambert series, and so on can be be re-written in slightly unrecognizable form, using C and D instead of the classical notation. This includes all of the identities that can be found in Apostol.[1] For reference and orientation, a few are given here. The arrow $\to$ is used to denote the interpretation of a classical formula.

The simplest "fun fact" is probably that the columns of D and rows of C sum to the divisor function $d(n)$:

$$d=\mathit{1}*\mathit{1}\to d(n)=\sum_{k=1}^{\infty}C_{nk}=\sum_{k=1}^{\infty}D_{kn}\to\vec{d}=C\vec{\mathit{1}}=\vec{\mathit{1}}^{\top}D$$

That is, the divisor function $d(n)$ is a model for the integer 2. The vector notation just emphasizes the linear-algebra aspect of the equations, so that $\vec{d}$ is the vector whose vector components are $d_n=d(n)$. Simplifying further,

$$\mathit{1}=\mathit{1}*\varepsilon\to 1=\sum_{k=1}^{\infty}C_{nk}\delta_{k1}=C_{n1}\to\vec{\mathit{1}}=C\vec{\varepsilon}$$

and so the representation of 2 becomes more clear:

$$d=\mathit{1}*\mathit{1}*\varepsilon\to\vec{d}=CC\vec{\varepsilon}=C^2\vec{\varepsilon}$$

Divisor functions for general powers are obtained by defining the sequence

$$N(n)=n$$

and more generally

$$N^{\alpha}(n)=n^{\alpha}$$

with $N^0=\mathit{1}$. Then the divisor function

$$\sigma_{\alpha}(n)=\sum_{d|n}d^{\alpha}$$

becomes

$$\sigma_{\alpha}=\mathit{1}*N^{\alpha}\to\sigma_{\alpha}(n)=\sum_{k=1}^{\infty}C_{nk}k^{\alpha}=\sum_{k=1}^{\infty}k^{\alpha}D_{kn}\to\vec{\sigma}_{\alpha}=C\vec{N^{\alpha}}=\vec{N^{\alpha}}^{\top}D$$

The Euler totient function φ is given by

$$\varphi=\mu*N\to\varphi\left(n\right)=\sum_{k=1}^{\infty}kD_{kn}^{-1}\to\vec{\varphi}=\vec{N}^{\top}D^{-1}$$

In what follows, most of the work will concern the operator D, and so all of these classical identities are statements about left-vectors of D, as opposed to the more natural use of right-vectors in linear algebra.

There is a deep and broad collection of further such relations; the Wikipedia article "Redheffer matrix"[2] provides a good entry point.

Affine space. Although the structure $(*,\mu,\varepsilon,\mathit{1})$ provides a model of arithmetic, the same structure can be based on any completely multiplicative function f. More precisely, the arithmetic acts on such functions:

$$\varepsilon*f=f$$

and so $\mathit{1}*f$ and $\mathit{2}*f$ and $\mathit{3}*f$ are also completely multiplicative functions, where a short-hand is introduced: $\mathit{2}=\mathit{1}*\mathit{1}$ and $\mathit{3}=\mathit{2}*\mathit{1}$and so on. Similarly writing $-\mathit{1}=\mu$ and $-\mathit{2}=\mu*\mu$ and so on, and finally writing $+$ for $*$, one has objects

$$\cdots,\ -2+f,\ -1+f,\ f,\ f+1,\ f+2,\ \cdots$$

all of which are completely multiplicative. Effectively, completely multiplicative functions form an affine space over these "integers".

Questions

It is tempting to try to view operator matrix C as a raising operator of some sort: that is, a shift. After all, it generates the integers, so this should be natural. This is not the case, and can be seen in two ways. One way would be to note that raising and lowering operators come in non-commuting pairs. In this case, the "lowering" operator C^{-1} commutes with C, that is, $C^{-1}C-CC^{-1}=0$, and so the prize property of such operator pairs is lost. Another way to see this is that raising operators have shift states (*aka* "coherent states"), given as eigenvectors. But C has no (right-)eigenvectors. That is, there are no explicit vectors $\vec{v}$ such that $C\vec{v}=\lambda\vec{v}$. In classical notation, there are no multiplicative functions f such that $\mathit{1}*f=\lambda f$ for some scalar λ. In operator language, C does not have a point spectrum. It may still have a continuous or a residual spectrum; this question is explored below. So although C carries one from one integer to the next, it is not a shift.

The matrix operator D. The transpose of C, the matrix D has one obvious eigenvector: $D\varepsilon=\varepsilon$. There are not any others that are of finite length. Any triangular matrix has all eigenvalues on the diagonal; for D they are all 1. Thus, D has a huge algebraic multiplicity; we conclude it is not diagonalizable (for if it was, the eigenvalues would lie on the diagonal – as they are all 1, it would be the identity matrix.)

Some obvious questions about D are:

- Can it be brought into something resembling Jordan normal form? The next section shows that the answer is no, at least not in the naive sense: there are no Jordan blocks. There are finite sequences of vectors that resemble Jordan chains: iterated upon, they eventually terminate. The sequences correspond exactly to the factorization of integers; each vector in the sequence corresponds to the removal of the least prime factor. Unfortunately, these sequences are not uniquely determined; nor can they be made orthogonal. They cannot be projected out, as they seem to "mix". Without a clear projection, there does not seem to be any clean or elegant way of writing classes of these sequences (or classes of subspaces spanned by the vectors). Thus, there are no simple or immediately insightful results that instantly appear.

- Perhaps it makes more sense to ask about the Frobenius normal form? A curious idea, but also one that depends on producing a chain of vectors, which does not seem possible.

One can also pose questions from classical operator theory. If the matrix D is treated as an operator, acting on some kind of vector space (choosing Banach ℓ_1 space seems natural), what is it's spectrum? The answer here is much simpler to arrive at: D has an immense point spectrum, the set of all completely multiplicative arithmetic functions (that are ℓ_1-summable). The spectrum is continuous (not surprising, as it was not diagonalizable), and the operator D itself is not a bounded operator (on the space of ℓ_1-summable sequences).

Jordan Decomposition

To find the Jordan decomposition, we look for the generalized eigenvectors. The first one is already $D\varepsilon = \varepsilon$. It's useful to define $\varepsilon_1 = \varepsilon$. The second one must be $(D - \lambda I)\,\varepsilon_2 = \varepsilon_1$ with $\lambda = 1$ and I the identity matrix. These are easier to visually read off if we write

$$D - I = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & \cdots \\ & 0 & & 1 & & 1 & & 1 & & 1 & \\ & & 0 & & & 1 & & & 1 & & \\ & & & 0 & & & & 1 & & & \\ & & & & 0 & & & & & 1 & \\ & & & & & 0 & & & & & \\ & & & & & & 0 & & & & \\ & & & & & & & 0 & & & \\ & & & & & & & & 0 & & \\ & & & & & & & & & 0 & \\ & & & & & & & & & & \ddots \end{bmatrix}$$

Thus, $\varepsilon_2 = (a, 1, 0, \cdots)$ fills the bill, for any a. It spans a one-dimensional affine space. In the following, ε_2 will be called a 'generalized eigenvector' even though any vector from the affine space will do; the generalized vectors are actually classes of vectors.

The next obvious one is $\varepsilon_3 = (b, c, 1-c, 0, \cdots)$ for any b and c. Clearly $(D-I)\,\varepsilon_3 = \varepsilon_1$. So ε_2 and ε_3 are necessarily in distinct Jordan chains. Next comes $\varepsilon_4 = (d, e, a-e-1, 1, 0, \cdots)$ which obeys $(D-I)\,\varepsilon_4 = \varepsilon_2$, so its part of an existing Jordan chain. This was written so that the a in ε_4 is the same a as in ε_2. It spans a two-dimensional affine space (*i.e.* holding a constant). Then what?

Clearly, for any prime number p one has $\varepsilon_p = (a, 0, \cdots, 0, 1, 0, \cdots)$ starts a new chain: $(D-I)\,\varepsilon_p = \varepsilon_1$. The next interesting composite is $6 = 1 \cdot 2 \cdot 3$. So

$$\varepsilon_6 = (a, b, c, -1, -c, 1, 0, \cdots)$$

must necessarily join the chain of it's greatest divisor: $(D-I)\,\varepsilon_6 = \varepsilon_3$. There is no other way: so ε_n is necessarily a vector of length n, with a 1 in the last position. As this multiplies through, this 1 in the last position will necessarily hit the row with the greatest divisor in it, and this row will necessarily have only one non-zero entry in it. The greatest divisor of n is n/d where d is the least prime factor of n. Repeating this process walks through each of the prime factors of n, with multiplicity. Thus, we have proven:

Lemma. *The length of the Jordan chain of ε_n is given by the number of prime factors $\Omega(n)$, so that*

$$(D-I)^{\Omega(n)}\varepsilon_n = \varepsilon_1$$

Writing the factorization of n as

$$n = p_1^{a_1} p_2^{a_2} \cdots p_k^{a_k}$$

with $p_i < p_j$ when $i < j$, the corresponding Jordan chain is given by the sequence

$$\begin{array}{c} p_1^{a_1} p_2^{a_2} \cdots p_k^{a_k} \\ p_1^{a_1-1} p_2^{a_2} \cdots p_k^{a_k} \\ p_1^{a_1-2} p_2^{a_2} \cdots p_k^{a_k} \\ \vdots \\ p_2^{a_2} \cdots p_k^{a_k} \\ p_2^{a_2-1} \cdots p_k^{a_k} \\ p_2^{a_2-2} \cdots p_k^{a_k} \\ \vdots \\ p_k \end{array}$$

which repeatedly peels off the least prime factor of n.

How can one characterize the affine spaces indicated by these ε_n? One can proceed constructively. Let e_n be the standard-notation basis vector: i.e. all zero's except for a 1 in the n'th location. The first vector is $\varepsilon_1 = e_1$, which spans a one-dimensional space. Other representative elements of the affine space can be decomposed as

$$\varepsilon_n = \sum_{k=1}^{n} a_{nk} e_k$$

with various constraints applied to the a_{nk} to be discovered.

Proceeding as before, $\varepsilon_2 = (a, 1, 0, \cdots) = a\varepsilon_1 + e_2$ and since $(D-I)\varepsilon_1 = 0$ one may as well take ε_2 to be orthogonal to ε_1, and normalized, so that $\varepsilon_2 = e_2$. The second affine space is generated by ε_2 and spanned by ε_1.

Next is $\varepsilon_3 = (b, c, 1-c, 0, \cdots)$. Asking for orthogonality to ε_1 sets $b = 0$. At this point, it seems reasonable to ask that ε_3 be orthogonal to ε_2, so as to avoid mixing of these subspaces. This constrains $c = 0$ so that the third space is generated by $\varepsilon_3 = e_3$ and spanned by ε_1.

As before, one may write $\varepsilon_4 = (c, d, a-d-1, 1, 0, \cdots)$ and asking that this be orthogonal to ε_1 and ε_2 implies that $c = d = 0$ so that $\varepsilon_4 = (0, 0, a-1, 1, 0, \cdots)$. Insisting on the Jordan chain condition $(D-I)\varepsilon_4 = \varepsilon_2$ implies that $a = 0$, so that $\varepsilon_4 = (0, 0, -1, 1, 0, \cdots)$. This means that ε_4 is not orthogonal to ε_3. At this point, it seems that there is no reason to ask for this; 3 is not a prime factor of 4. The corresponding affine space is spanned by the Jordan chain ε_1, ε_2 and ε_4.

Continuing as before, we had seen that $\varepsilon_6 = (a, b, c, -1, -c, 1, 0, \cdots)$ and asking for orthogonality to the prime factors ε_1, ε_2 and ε_3 means that $a = b = c = 0$ and so $\varepsilon_6 = (0, 0, 0, -1, 0, 1, 0, \cdots)$. The Jordan chain is ε_1, ε_3 and ε_6. Next is $\varepsilon_8 = (0, 0, 0, 0, 0, -1, 0, 1, \cdots)$, which is uniquely determined by $(D-I)\varepsilon_8 = \varepsilon_4$ and the demand that it be orthogonal to the rest of it's Jordan chain.

Next is ε_9. This is not well-constrained by the previous pattern of constraints. The vector $\varepsilon_9 = (0, a, 0, 0, b, 0, -1-a-b, 0, 1, \cdots)$ satisfies $(D-I)\varepsilon_9 = \varepsilon_3$ and is orthogonal to the rest of it's Jordan chain (and to it's own prime factors). It is not constrained in the

2,5 and 7 positions, effectively because 2, 5 and 7 are all relatively prime to 9, and none appear in the Jordan chain. To form a constraint, one can insist on orthogonality to two of these, but one cannot insist on orthogonality to all three (unless one instead discards the constraint that $(D-I)\,\varepsilon_9=\varepsilon_3$.)

Next is ε_{10}. Since it is relatively prime to both 3 and 7, one has that $\varepsilon_{10}=(0,0,a,0,0,0,-1-a,0,0,1,\cdots)$. Clearly, the strategy of constructing Jordan chains in this fashion – asking each element to be orthogonal to the last – leaves them under-constrained. Asking for the chains to all be orthogonal is an over-constraint.

I don't see any particular way out. Since D has nothing but 1's along the diagonal, it would be absurd to envision an infinite-dimensional Jordan block. Despite this, one might hope to still have some kind of pseudo-blocks, feeding into one-another based on the factorization of each integer. How to make this explicit is unclear; its not obvious that this dream of pseudo-locks is even possible.

SPECTRA

Let $\vec{s}=(s_n)$ be the vector given by vector components $s_n=n^{-s}$. Let $\vec{\zeta}=D\vec{s}$, then the vector components of $\vec{\zeta}$ are $\zeta_k=k^{-s}\zeta(s)$. This can be easily seen by inspecting the visual form of the D matrix. The top row corresponds to $k=1$ and clearly sums all terms. The second row, for $k=2$, sums only even terms, and so on.

This demonstrates that $\vec{\zeta}$ is just a constant multiple of $\vec{s}$, the constant of proportionality given by $\zeta(s)$. In formulas: $\vec{\zeta}=D\vec{s}=\zeta(s)\,\vec{s}$. This is a diagonalization, and eigenvector-eigenvalue equation. We've written down at least some of the spectrum of D.

A similar identity holds for any completely multiplicative sequence. That is, if $f(mn)=f(m)\,f(n)$ for all integers m,n then writing $g=Df$ one has that

$$g(1)=\sum_{n=1}^{\infty}f(n)\quad\text{and}\quad g(k)=f(k)\,g(1)$$

which is obviously a formal series when the summation does not numerically converge. This too has the form of an eigenvalue equation:

$$Df=g(1)\,f$$

and provides a spectrum whenever $g(1)$ is actually summable and convergent.

Note that a completely multiplicative sequence is determined entirely by it's values on the prime numbers: that is, by the values $f(p)$ for p prime. It would appear that we've demonstrated a lemma:

Lemma 1. *Let M be the space of completely multiplicative sequences. Let ℓ_1 be the Banach space of summable sequences. The the space $M\cap\ell_1$ lies within the eigenspace of D.*

I suspect the converse is true, so that $M\cap\ell_1$ is exactly the eigenspace of D, but am too lazy to pursue a proof right now.

Spectral Theory. The goal of this section is to describe the spectrum of D. It's a slow walk through the standard vocabulary and definitions of spectral decompositions, mostly posing questions but leaving them unanswered.

Lets assume (for now) that D is a bounded operator on the Banach space ℓ_1. Actually, it isn't bounded, as shown in the next section. Reviewing common definitions for bounded operators:

- The set of numbers λ such that $(D-\lambda I):\ell_1\to\ell_1$ is one-to-one and onto is called the resolvent set, denoted by $\rho(D)$. It seems that perhaps the resolvent is the empty set in this case (or perhaps, per Picard theorem, contains a single point. Or two.).
- The complement of the resolvent set is the spectrum, denoted as $\sigma(D)=\mathbb{C}\backslash\rho(\sigma)$.

What kind of spectra are there?

- The point spectrum of D are all those values $\lambda\in\sigma(D)$ for which $D-\lambda I$ is not one-to-one. That is, the kernel of $D-\lambda I$ is non-trivial; *i.e.* there is a vector, other than the zero vector, that is in the kernel of $D-\lambda I$.
- The continuous spectrum of D are all those values $\lambda\in\sigma(D)$ for which $D-\lambda I$ is one-to-one but not onto, and the range of $D-\lambda I$ is dense in ℓ_1. Examples include the "almost-eigenvectors", i.e. vector-sequences $\vec{v}_n$ that converge to ever-smaller $(D-\lambda I)\vec{v}_n\to 0$, while at the same time not converging themselves: $|\vec{v}_n|\to\infty$. Such λ's belong to the continuous spectrum, not the point spectrum.
- The residual spectrum of D are all those values $\lambda\in\sigma(D)$ for which $D-\lambda I$ is one-to-one but not onto, and the range of $D-\lambda I$ is not dense in ℓ_1.

It would appear that $D-\lambda I$ has a point spectrum that is pretty much the entire complex plane. This is because $\zeta(s)$ has a pole at $s=1$ and so as $s\to 1$ it follows that $\lambda\to\infty$. Whether or not the point spectrum really is the entire complex plane or not requires answering the question: does $\zeta(s)$ attain every possible complex value, or not? (I do not know).

Another way of saying is: the spectral radius of D is infinite (unbounded) because $\zeta(s)$ is unbounded as $s\to 1$ (and the sequence sequence $s_n=n^{-s}$ is ℓ_1-summable as long as $\Re s>1$; that is $\vec{s}\in\ell_1$ whenever $\Re s>1$). Does $\zeta(s)$ attain every possible complex value in the domain $\Re s>1$?

Note that just because the the point spectrum might be the entire complex plane, this does not mean that the continuous and residual spectra are empty. These are distinct questions.

Operator norm. What's the operator norm of D? The operator norm is defined as

$$\limsup_{\vec{v}}\frac{|D\vec{v}|}{|\vec{v}|}$$

but this is clearly unbounded: taking $\vec{v}=\vec{s}$, one instantly has

$$\frac{|D\vec{s}|}{|\vec{s}|}=\frac{|\zeta(s)\vec{s}|}{|\vec{s}|}=|\zeta(s)|$$

which diverges as $1/s$ as $s\to 1$. A simple pole, as it were. It's operator norm is divergent in this vector-direction. From this we can only conclude that D is not a bounded operator.

Are there other problematic directions? It would appear that any Dirichlet character will result in a Dirichlet L-function having a pole at $s\to 1$. As vector directions, they are distinct. As there are a countable infinity of characters, it would appear that D is unbounded in a countable number of distinct directions (as characters are linearly independent of one-another).

Are there other directions? Let's try our luck. Consider the vector $\vec{q}=\left(1,q,q^2,\cdots\right)$ then $|\vec{q}|=1/(1-q)$. For the following exercise, it's pointless to imagine that q is a complex number; so instead, take $0\le q<1$ real. The norm of the divisor operator applied to this vector is

$$|D\vec{q}| = \sum_{n=1}^{\infty} \frac{q^{n-1}}{1-q^n}$$

which is a Lambert series. Since

$$\frac{1}{1-q^n} \leq \frac{1}{1-q}$$

one can write

$$\begin{aligned}\frac{|D\vec{q}|}{|\vec{q}|} &= (1-q)\sum_{n=1}^{\infty} \frac{q^{n-1}}{1-q^n} \\ &\leq \sum_{n=1}^{\infty} q^{n-1} \\ &= \frac{1}{1-q}\end{aligned}$$

which is useless as a bound; it is no bound at all. Can we do better? It would seem that the asymptotic expansion $q \to 1$ for the Lambert series is

$$\sum_{n=1}^{\infty} \frac{q^{n-1}}{1-q^n} = \frac{-\log(1-q)}{1-q} + \frac{\gamma}{1-q} - \frac{1}{2}\log(1-q) + \mathscr{O}(1-q)$$

which implies that

$$\begin{aligned}\frac{|D\vec{q}|}{|\vec{q}|} &= (1-q)\sum_{n=1}^{\infty} \frac{q^{n-1}}{1-q^n} \\ &= -\log(1-q) + \mathscr{O}(1)\end{aligned}$$

as $q \to 1$. So this is a distinct direction in which D is not bounded.

The asymptotic expansion closely resembles Stirling's approximation, but presence of the Euler–Mascheroni constant γ suggests that that the derivative of the gamma function, *i.e.* the digamma, appears in the derivation.

A speculative question presents itself, at this junction: is it possible to decompose D into two parts, a bounded part, and an unbounded part? That is, is there any way possible to write

$$D = D_{bounded} \oplus D_{divergent}$$

If so, are the two parts uniquely determined? What might they be?

Why even pose such a crazy question? Easy: the eigen-equation $\vec{\zeta} = D\vec{s} = \zeta(s)\vec{s}$ given earlier is meaningful only when $\Re s > 1$, and is otherwise blocked by a pole in $\zeta(s)$ at $s = 1$. Now, of course, $\zeta(s)$ can be analytically continued to the rest of the complex plane. Can one do something analogous for $D\vec{s}$? How would that work?

Lambert basis. For example: define $\vec{q}$ as a kind-of principle vector/direction, and then build an orthonormal basis around it. Perform a change-of-basis, i.e. apply the similarity transform; what does D look like then? Lets try this.

Define a sequence of vectors $\vec{q}_n$ of the form

$$\vec{q}_n = \left(1, 1-\frac{1}{n}, \left(1-\frac{1}{n}\right)^2, \cdots\right)$$

The length of each vector is

$$|\vec{q}_n| = \sum_{k=0}^{\infty}\left(1-\frac{1}{n}\right)^k = n$$

It's convenient to normalize to unit length: $\hat{q}_n = \vec{q}_n/n$ so that $|\hat{q}_n| = 1$. With this normalization, its clear that $\hat{q}_n \to 0$ strongly, and not weakly, because each individual component in $\hat{q}_n$ is going to zero. (TODO review formal definitions of strong/weak convergence; I might be mis-using the terms).

From the previous section,

$$|D\hat{q}_n| = \log n + \mathcal{O}(1)$$

The $\hat{q}_n$ are not mutually orthogonal; can we build some orthogonal basis that is weakly/strongly orthogonal to the limit $\hat{q}_n \to 0$ in some way? (Last time, when I tried to expand $\vec{s}$ in an orthonormal basis, I ended up with the topologists sine. Is this why that happened?)

Dirichlet characters. The Dirichlet characters mod k are completely multiplicative. Therefore, they too are eigenvectors. Given the character χ, define the vector $\vec{\chi} = \chi(n)n^{-s}$. The eigen-equation is

$$D\vec{\chi} = L(s,\chi)\vec{\chi}$$

with $L(s,\chi)$ the usual Dirichlet $L-$series

$$L(s,\chi) = \sum_{n=1}^{\infty}\frac{\chi(n)}{n^s}$$

Just like $\vec{\zeta}$, the character vectors $\vec{\chi} \in \ell_1$ whenever $1 < \Re s$ (since $\chi(n)$ is always a root of unity or zero). Unlike $\zeta(s)$, which has no zeros for $1 < \Re s$, the Dirichlet $L-$series have zeros in the strip $1 < \Re s < 1+\varepsilon$. (Right? Am I imagining this? WTF? I forget...) This implies that D has a non-trivial kernel. (ummm.... that can't be right, can it?)

An earlier section wrote down the matrix for D^{-1}, which seems to be "obviously" the inverse, by inspection. Since D has a kernel in ℓ_1, one concludes that D^{-1} must be unbounded on ℓ_1. I presume (without checking) that

$$\frac{1}{L(s,\chi)} = \sum_{n=1}^{\infty}\frac{(\mu*\chi)(n)}{n^s}$$

where $(\mu*\chi)$ is the Dirichlet convolution. (double check that the above is correct, I think it is; its a classical number-theory identity.). So while D had a divergence for the Lambert-series vector, it seems that D^{-1} has much stronger divergences at the XXX zeros. Hang on wtf. Where are the zeros of the L-series, again? This can't all be right.

Series Shifts

Interesting problems arise if one considers composing D with the left-shift and right-shift operators S and S^T, where the left shift is defined as $[Sf](n) = f(n+1)$ for any arithmetic function $f(n)$. Also intriguing are the ladder operators for the simple harmonic oscillator, A and A^T defined by $[Af](n) = \sqrt{n+1}f(n+1)$, or, in more conventional notation, $A|n\rangle = \sqrt{n}|n-1\rangle$ with $A|0\rangle = 0$. The natural questions arising are then to describe the product DS, and related constructions, such as the commutator $A^TD - DA$ or similar oddities, such as the Laplacian-like $2D - \left(A^TD + DA\right)$.

Numerical exploration suggests ... eigenvalues are stable functions of matrix dimensions. Some are near the real axis, some are near the imaginary axis. The eigenvalues seem bounded. (Not increasing as function of matrix size). So far, nothing hops out and says “look at me”.

Numerology

Identities worth not forgetting. The digamma:

$$\psi(z+1) = -\gamma - \sum_{k=1}^{\infty} \zeta(k+1)(-z)^k$$

Conclusion

None yet. This is yet another patented sprawl of randomized ideas masquerading as an essay. This essay may or may not be expanded and revised at some future date.

References

[1] Tom M. Apostol. *Introduction to Analytic Number Theory*. Springer; New York, 1976.
[2] Wikipedia. Redheffer matrix, 2019. https://en.wikipedia.org/wiki/Redheffer_matrix.