On Linear Maps

by
K. S. Sarkaria

(From a talk given at
Brigham Young University,
Provo, Utah, U.S.A., on
September 26, 2006.) [1]

§ 1. We are by now blasé about *linear maps* f: V → V. Yes, their utility is undisputed—to cite just one famous instance, the dimension of the kernel minus that of the image of an exterior derivative, a linear map with $f^2 = 0$, rather magically yields the 'number of i-dimensional holes, $i \geq 0$' of the underlying smooth manifold—but, of and by themselves, you might well ask, are these homely objects worth a whole talk?

It is the possible *infinite dimensionality* of V that imparts interest to the topics I'll be considering, but I am not going to start with a topology on V and focus on maps f continuous with respect to it, that is, this is not going to be functional analysis. It will be pure linear algebra. Nevertheless, some topology and (co)homology will enter rather naturally (see §§ 9-12) before this topologist of sorts is done doing this algebra.

A researcher rash enough to venture into a new area is bound to rediscover as much, if not more, than what he discovers. My fate was no different (see Appendix 1)[2]. However the new light which such an interloper brings can be revealing. You are going to find this algebra talk uncommonly pictorial, and I feel that these *pictures*—definition below—constitute perhaps the right way of looking at this bit of algebra.

§ 2. I was led to them a year and half ago starting with an easy exercise: *does* ker(f) *always have an* f-*invariant complement*? Here, by 'complement' of ker(f) we mean a vector subspace C such that $\ker(f) \oplus C = V$, and by 'f-invariant' that $f(C) \subseteq C$. The answer to the exercise is *no*. For example, for any nonzero f with $f^2 = 0$, because f maps any complement of ker(f) into ker(f). More generally, for a nonzero f having some iterate $f^n = 0$ where n is least—for example, for the map $f(x_1, x_2, \ldots, x_n) = (x_2, x_3, \ldots, x_n, 0)$ of n-space—$\ker(f^{n-1})$ does not have an f-invariant complement.

At this point, Dinesh Khurana had asked: *does* $K_f = \bigcup_n \ker(f^n)$ *always have an* f-*invariant complement*? This new question was more challenging, the answer however was still *no*. Figure 1 shows the counterexample that I found after a few weeks. Here,

[1] January 10, 2007: this write-up may also be construed to be an outline—not in its final form, e.g., the mentioned Appendices are lacking—of an intended joint book with Khurana and Madahar.

[2] This history is intricate, for example, there is a largely forgotten—and not easily decipherable, for it has many gaps and mistakes, some quite serious—1917 book of F. Levi, which deserves priority for many ideas and results now routinely but wrongly ascribed to others: in November, 2006, I realized that "my" own counterexample, Figure 1, was also implicit in it!

and in all other pictures, the set of nodes forms a basis, and the arrows show how the linear map acts on these basis elements: each node has *just one or no arrow* issuing out of it, in the former case it is imaged to the next node, in the latter to zero. Many other pictures, and concomitant informal but self-explanatory terminology, will be used profusely from here on, however I'll postpone the further discussion of this particular picture for quite some time (till § 9).

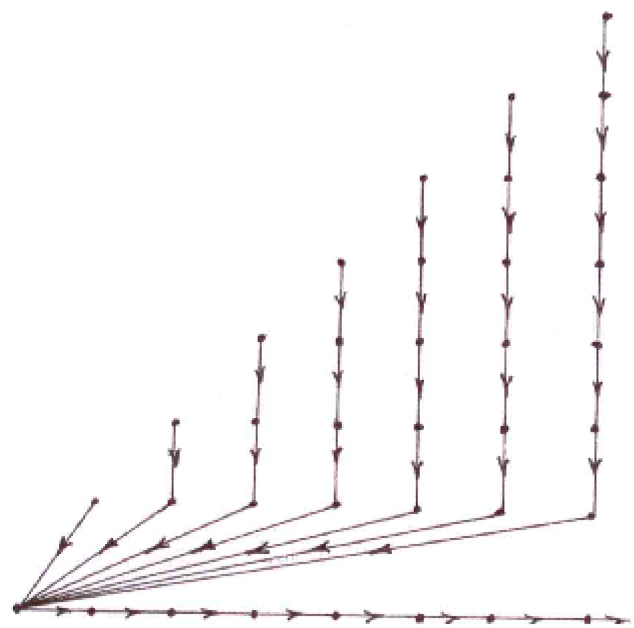

Figure 1

§ 3. Let's first assure ourselves that Dinesh's question has no finite dimensional counterexample. Note that if a term of either of the monotonic sequences,

$$0 \subseteq \ker f \subseteq \ldots \subseteq \ker f^n \subseteq \ker f^{n+1} \subseteq \ldots,$$

$$V \supseteq \operatorname{im} f \supseteq \ldots \supseteq \operatorname{im} f^n \supseteq \operatorname{im} f^{n+1} \supseteq \ldots,$$

is *stable,* that is, equal to the next term, then by applying f^{-1} or f repeatedly we see that it is equal to all subsequent terms. Also, it is always true that if $\ker f^n$ is stable, then $\ker f^n \cap \operatorname{im} f^n = 0$, because, if there were a nonzero $f^n(v)$ in this intersection, then $f^{2n}(v) = 0$ contradicts the assumed stability. Now invoke the finite dimensionality of V. It implies that the sum of $\dim(\ker f^i)$, a non-decreasing function of i, and $\dim(\operatorname{im} f^i)$, a non-increasing function of i, equals the constant dim(V). So, in this finite dimensional case, $\ker f^n$ is stable iff $\operatorname{im} f^n = 0$ is stable iff $\ker f^n \oplus \operatorname{im} f^n = V$, that is, now $K_f \oplus I_f = V$ where I_f denotes $\bigcap_n f^n(V)$. In fact this stable $\ker f^n$ and $\operatorname{im} f^n$ are the *unique* f-invariants complements of each other, for, if C is f-invariant with $\ker f^n \cap C = 0$, respectively $C \cap \operatorname{im} f^n = 0$, then $C \subseteq \operatorname{im} f^n$, respectively $C \subseteq \ker f^n$. This extra point can be useful.

For example, it implies that K_f has a unique f-invariant complement even under the much weaker hypothesis that each *orbit* $\{v, fv, f^2v, \dots\}$ spans a finite-dimensional vector subspace W. For, the intersection of K_f with W equals $K_{f|W}$ where $f|W$ denotes the restriction of f to W, and has in W the unique f-invariant complement $I_{f|W}$. On account of this uniqueness, the union $\bigcup_W I_{f|W}$ is itself a vector subspace, the asserted unique f-invariant complement C of K_f in V.

A natural question arises: is there a more convenient description of this unique C? The obvious suspect I_f is now easily ruled out, for example, the left infinite string

$$\dots\dots \to\bullet\to\bullet\to\bullet\to\bullet\to\bullet$$

shows we might even have $K_f = I_f$. The key is to observe that our orbital hypothesis is the same as saying that f is *locally algebraic,* i.e. that, for each v we can find a nonzero *polynomial* $c_0 + c_1f + c_2f^2 + ..$ of f which kills v, and then exploit—arguments below—the fact that if a bunch of polynomials kill v, then their h.c.f. also kills v.

Since f is one-one on C, its v's can be killed by polynomials having nonzero constant terms, and since such polynomials are one-one on K_f this subspace contains no such nonzero v. Thus C coincides with the invariant subspace consisting of all v's killed by polynomials not divisible by f. But why stop here? Why not not split off a further K_φ from C where φ is a *prime polynomial* other than f, and so on? This indeed works like a charm and gives the canonical direct sum decomposition,

$$V = \oplus_\varphi K_\varphi ,$$

where φ runs over all prime polynomials with highest degree coefficient 1. The components $v_\varphi \in K_\varphi$ of any $v \in V$ are given thus. Take the polynomial μ of minimum degree with highest degree coefficient 1 which kills v. If a prime polynomial φ occurs k times in μ then v_φ is the element of K_φ mapped by the monomorphism μ/φ^k of K_φ to $(\mu/\varphi^k)(v)$. For almost all primes φ one has $k = 0$, then $(\mu/\varphi^k)(v)$ and so v_φ is zero. That $v - \Sigma_\varphi v_\varphi = 0$ holds follows by noting that the left side is killed by each polynomial μ/φ^k, and so must also be killed by their highest common factor, which is 1.

Dinesh's question has a positive answer (see Appendix 2) under some other finiteness hypotheses also, notably if ker f^n or im f^n is stable, and the proofs are akin to those above. These special cases can mislead one into thinking that the answer might be always yes. Indeed, interest in Dinesh's question was stimulated by an emailed misproof to this effect by a well-known U.S. algebraist, and even a year after the above counterexample had been circulated, another well-known U.S. algebraist was still toying with a misproof. On the other hand, counterexamples were not hard to find, once the pictorial approach being pushed in this talk had been adopted.

§ 4. The crowning glory of our undergraduate linear algebra courses is Jordan's similarity classification of finite dimensional complex linear maps. The first part of its proof is the above canonical direct sum decomposition. Due to the finite dimensionality of V only finitely many φ's occur in this formula, and since coefficients are algebraically closed (see Appendix 3 for Jordan theorem for non algebraically closed coefficients) these are all of degree one—$\varphi = f - \lambda$, where λ runs over the *eigenvalues* of f—and so f on each summand K_φ is known once we know $\varphi: K_\varphi \to K_\varphi$.

The job is thus to classify *nilpotent* maps, i.e., those with some power zero. Here pictures enter naturally, for all Jordan's theorem says is that, *any finite dimensional nilpotent map admits a basis on which its action gives us a disjoint union of strings,* for example, Figure 2 shows such a Jordan basis of a 23-dimensional nilpotent map.

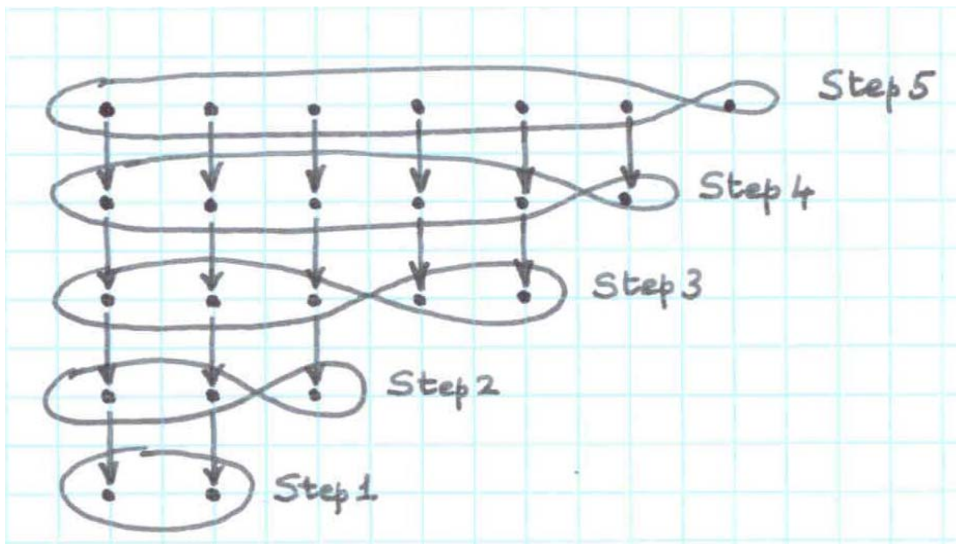

Figure 2

The correlation with the usual formulation is obvious. If say Figure 2 depicts $\varphi = f - \lambda: K_\varphi \to K_\varphi$ then what we are saying is that the matrix of f: $K_\varphi \to K_\varphi$ with respect to this basis of nodes is the direct sum of 7 Jordan matrices—λ's on main diagonal, 1's on the diagonal above it, and 0's in all other places—because each of the 7 strings contributes a Jordan matrix of the same size as its length.

This visualization also suggests an easy *going-up proof*. The required basis is constructed in steps using the fact that f induces a surjection of (im f^i)/(im f^{i+1}) onto (im f^{i+1})/(im f^{i+2}) for all i. If $f^n = 0$, n least (one has n = 5 in Figure 2) in Step 1 we simply choose any basis of im f^{n-1}. In Step 2 we use this surjection with i+2 = n to first pull this basis to the same number of elements of im f^{n-2} which stay linearly independent mod im f^{n-1}, and then augment this linearly independent set to obtain a basis of im f^{n-2}. In the next Step 3, we'll use the surjection with i+2 = n-1 to similarly lift this basis and then augment the lifted linearly independent set to a basis of im f^{n-3}, etc.

Alternatively one can use the fact that f induces an injection of (ker f^i)/ (ker f^{i-1}) into (ker f^{i-1}) / (ker f^{i-2}) to give a *going-down proof* indicated in Figure 3, which is Figure 2 redrawn so as to make the bottom nodes flush. Now the requisite basis is constructed by starting with elements representing any basis of (ker f^n)/(ker f^{n-1}), then we inject and augment to get elements representing a basis of (ker f^{n-1})/(ker f^{n-2}), etc.

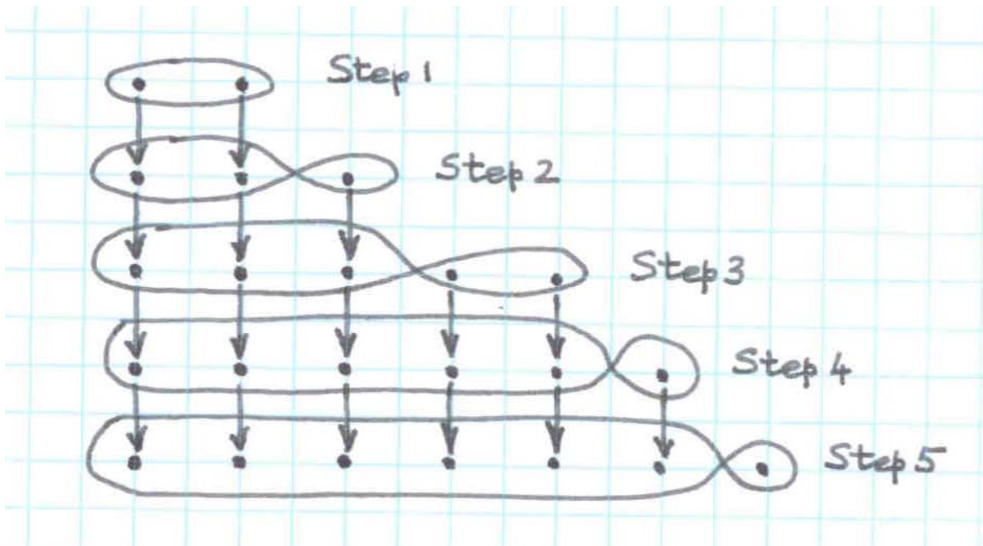

Figure 3

Noting how these pictures correlate with Young diagrams we also see at once that the number of similarity classes of nilpotent maps of a finite dimensional vector space V equals the number of *partitions* of dim(V).

§ 5. Now lets drop the finite dimensional hypothesis, but still insist on V = K_f, i.e., the map f is *locally nilpotent*, each element v killed by an $f^{n(v)}$, power depending on v. Can we still find a basis—we know, thanks to the axiom of choice, that bases exist—on which action of f gives disjoint but possibly infinitely long strings?

The answer to this question is also *no*. For example, the linear map depicted by Figure 4 (note that this is Figure 1 minus the horizantal string) cannot admit a basis of disjoint strings. For if it did, and all strings were finite, then we would have I_f = 0, while if there were an infinite string, then I_f would be infinite dimensional; but I_f , the span of the bottom node, is one-dimensional.

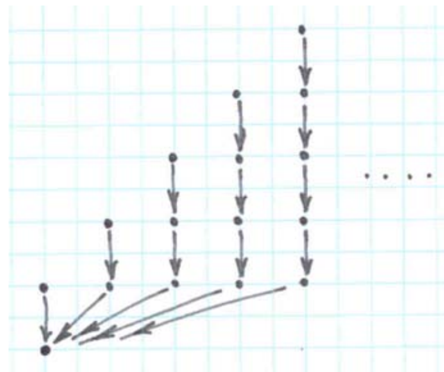

Figure 4

This prompts a modified question: maybe all locally nilpotent maps f are *graphic* admitting at least some picture as basis, i.e., a basis S such that the set $S \cup \{0\}$ is mapped by f into itself? We'll see towards the end—in § 12—that the answer is 'no' and that the problem of characterizing linear maps which are graphic is deep and very interesting.

§ 6. In the finite dimensional case this weaker statement carries almost the full force of Jordan's theorem, for, from a tree basis S we can very easily get another which is a disjoint union of strings. All we have to do is *strip* some branches off our tree(s)! Yet, a modicum of care is needed, the next picture illustrates this.

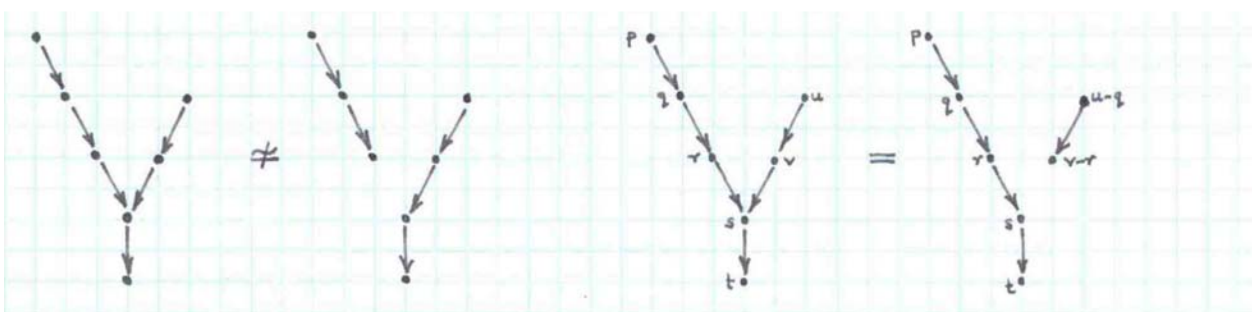

Figure 5a

Figure 5b

It is *invalid to strip off a longer branch*, for example, the maps of Figure 5a are dissimilar, only the 5^{th} power of the first is zero, while the 4^{th} power of the second is zero. Otherwise stripping is valid, in fact as the pictures of Figure 5b depict, the map remains the same if the nodes of the stripped branch are taken to be the differences of the old nodes and equidistant-from-fork nodes of an equal or longer branch still remaining at the fork (in matricial language a 'row transformation' has been done).

In other words, we insist from here on that stripping not alter the *height,* i.e. the largest α such that $f^{\alpha}V$ contains it, of any node. (For example, in Figure 2 the seven top most nodes have height 0, the six in the next lower level have height 1, … , the bottom most two have height 5, and none has height 6.) Therefore, *two finite trees* S *and* S' *are similar if and only if there is a height preserving bijection between their nodes.* For, by stripping, any tree basis can be made a Jordan basis, and if $h_\alpha = \dim(f^\alpha V/f^{\alpha+1}V)$ denotes the number of height α nodes, the *number* k_α *of height* α *strings* is given by $k_\alpha = h_\alpha - h_{\alpha+1}$. (In Figure 2 the decreasing sequence of h's was 7, 6, 5, 3, 2, 0, its successive differences 1, 1, 2, 1, 2 give us the sequence of k's, and adding the k's successively in reverse order to 0 we can recover the h's.) The k's are also invariants, for it is easy to check

$$k_\alpha = \dim(\ker f \cap f^\alpha V/ \ker f \cap f^{\alpha+1} V).$$

In the infinite dimensional case below, $h_\alpha - h_{\alpha+1}$ (difference of cardinals) seldom makes sense, so above formula will be our definition of k_α. Moreover, α will run over all ordinals, so we'll now recall these, and then define $f^\alpha V$ for any ordinal number α.

§ 7. In primary school, the finite ordinals α are often introduced as strings of beads [α]—see Figure 6, the initial bead will represent 0—the *next ordinal* being [α+1] = [α]→ •, the notation signifying 'lay a new bead after the last bead'.

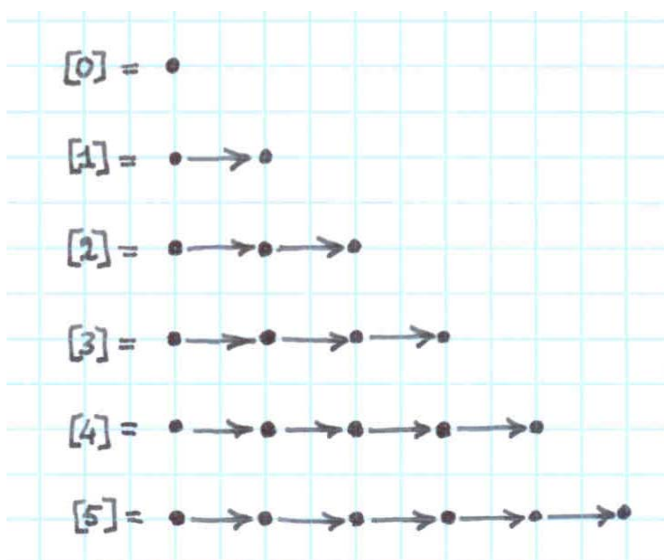

Figure 6

The set of finite ordinals cannot be augmented further by this construction. Whenever so stuck, we 'lay a new bead simultaneously after all the last beads' to construct the *next limit ordinal*. So, Figure 4 is the beads' picture [ω] of the first infinite ordinal, [ω+3] = [ω]→ •→ •→ •, and 2ω+2 and ω^2 are portrayed in Figure 7 below.

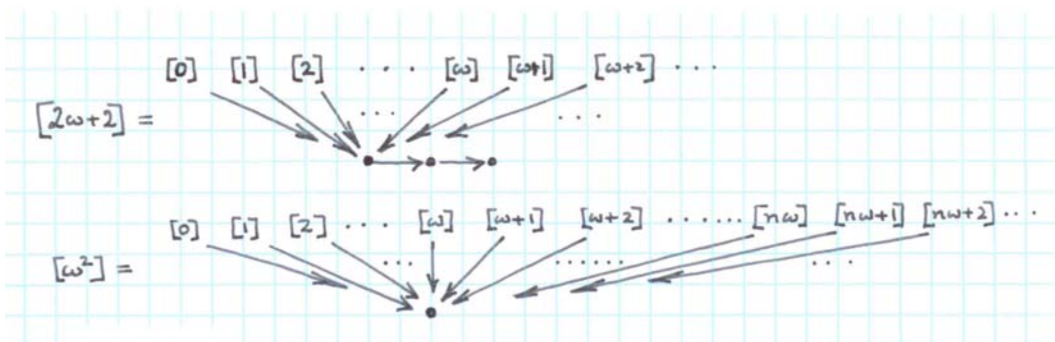

Figure 7

Coming back to a linear map f: $V \to V$, note that the increasing sequence ker f^n starts with the subspace 0, and each subsequent term is obtained by applying f^{-1} to the previous term, while the decreasing sequence im f^n starts with V, and each subsequent term is obtained by applying f to the previous term. If we apply f^{-1} to the union K_f of the increasing sequence we only get back K_f, however, applying f to the intersection I_f of the decreasing sequence we can get something strictly smaller. So we'll put $f^{\omega}(V) = I_f$, and the subspaces obtained by applying f repeatedly to I_f are denoted $f^{\omega+1}(V)$, $f^{\omega+2}(V)$, $f^{\omega+3}(V)$, $f^{\omega+4}(V)$, …, after this comes the intersection $f^{2\omega}(V)$ of all these, etc.

This transfinite sequence $f^{\alpha}(V)$, α any ordinal, ceases to be strictly decreasing at some ordinal $\alpha = \lambda$, for, by choosing an element in each term which is not in the next we obtain a linearly independent subset of V, and such a set has cardinality at most dim(V). We use $\lambda = \lambda(f)$ to denote the least such ordinal and call it the *length* of f.

We already have in hand rather canonical examples showing that length can be arbitrarily big. For the locally nilpotent linear map f_α: $V_\alpha \to V_\alpha$ (over any field of coefficients) defined by the beads' picture [α] of the ordinal, $(f_\alpha)^{\alpha}(V_\alpha)$ is the one-dimensional span of the last bead, so $\lambda(f_\alpha) = \alpha + 1$; and, if we knock off the last bead from [α], the new picture defines another such linear map having length exactly α .

The *stable part*, i.e., the restriction to $f^{\lambda(f)}(V)$, of f is surjective; conversely, given any α, and any surjective linear map g: $W \to W$, here is a construction of an f with length α and stable part g. Make nonzero elements of W—or only some chosen w's such that the union of their orbits $\{g^i v : i \geq 0\}$ spans W—final beads of disjoint copies of [α]. Treating the other beads of these *attached* disjoint pictures as linearly independent, let V be the linear span of W and these, and let f: $V \to V$ be the linear extension of g which acts on these other beads as per the arrows of their pictures.

For the locally nilpotent case $V = K_f$, the restriction $f^{\lambda}V \to f^{\lambda}V$ is graphic, in fact it is easily seen that any disjoint union of $s_{-\infty} = \dim(\ker f \cap f^{\lambda}V)$ left infinite strings $\ldots\bullet\to\bullet\to\bullet$ forms a basis whenever their last nodes form a basis of $\ker f \cap f^{\lambda}V$. Moreover, this stable part has an f-invariant complement. Zorn's lemma tells us that there is a maximal f-invariant subspace C of V with $f^{\lambda}V \cap C = 0$. We assert $f^{\lambda}V + C = V$. If the left side were smaller, then, using local nilpotence, we can find a $v \in V$ outside it, whose image $f(v) = (u,c)$ is in it. We can also arrange $u = 0$. Otherwise, we just replace v by v-u'where $u' \in f^{\lambda}V$ is such that $f(u') = u$. But then the bigger f-invariant subspace C' generated by C and v also satisfies $f^{\lambda}V \cap C' = 0$, which is impossible. Thus f is known, once we know $s_{-\infty}$ and the structure of the locally nilpotent map induced by f in the quotient $V/f^{\lambda}(V) \approx C$. This last map is *reduced*, that is, its stable part is zero, so we may now consider this subcase only.

§ 8. The stripping of §6 generalizes to show that *if a reduced locally nilpotent map has a graphic basis, it also has one in which forks occur only at nodes whose heights are limit ordinals,* so finally the finitely many other nodes of any string that runs between nodes of such heights must have consecutive heights. Assume only partial stripping has been done, that is, there exists still a consecutive heights string from a limit ordinal height α node v_{-n} with last fork before one at a higher limit ordinal at the vertex $v = v_0$ of height $\alpha+n$, $\ldots \to v_{-n} \to v_{-n+1} \to \ldots \to v_{-1} \to v_0 = v \to \ldots$. Take any another branch coming to this fork, say $\ldots \to u_{-2} \to u_{-1} \to u_0 = v$. We strip this off to the disconnected terminating tree $\ldots \to (u_{-2} - v_{-2}) \to (u_{-1} - v_{-1})$. Here u_{-t} signifies any of the nodes of this branch which is t arrows behind v. For $t \leq n$ its height is no more than that of v_{-t}, for $t > n$ we take care to choose higher height v_{-t}'s, also t arrows behind v, on suitably chosen corresponding branches ending at the height α node v_{-n}. The unchanged nodes and these form a new basis, in which the same map has the indicated new picture.

In such a *normalized* picture let $S_{\alpha,n}$ – here α is 0 or a limit ordinal and n is a finite ordinal – denote the set of all strings of n arrows issuing out of nodes of height α ; the last node of such a string might be hanging free as a terminus, or else might have an arrow going out to a node whose height is a bigger limit ordinal than α (this height may be bigger than $\alpha+\omega$). Since the cardinality of each $S_{\alpha,n}$ is obviously equal to the invariant $k_{\alpha+n}$, it follows that two normalized pictures arising from the same map, or even from two *similar* – i.e. admitting a linear isomorphism commuting with the two – maps necessarily have equally big sets $S_{\alpha,n}$. Before checking that this necessary condition is also sufficient for similarity, we pause to look at a few simple examples.

1. In case the length of the (reduced locally nilpotent and graphic) map is at most ω, there is just one normalized picture: cardinality k_n sets $S_{0,n}$ of strings of length n hanging free, à la classical Jordan, and the map has length ω iff the n's are unbounded.

2. For (the map defined by) $[\omega]$ the invariants are $k_0 = 1$, $k_1 = 1$, $k_2 = 1$, $k_3 = 1$, $k_4 = 1$, $k_5 = 1$, …; $k_\omega = 1$ (always $k_\alpha = 0$ for $\alpha \geq \lambda$, the length). Another normalized picture of the same map is obtained if we detach or strip off the string [n] – to see this, let nodes of detached strip be original nodes minus 'equidistant from the height ω node' nodes of a longer string, say [n+1] – indeed, we can likewise detach any number of [n]'s, provided we take care that arbitraily long [n]'s will remain undetached.

3. Now lay a bead simultaneously after the last beads of all the even finite ordinals [n], the length is still $\omega+1$, but now $k_0 = 1$, $k_1 = 0$, $k_2 = 1$, $k_3 = 0$, $k_4 = 1$, $k_5 = 0$, …; $k_\omega = 1$, so this map is not similar to the last example. And both these maps are dissimilar

from that obtained by laying a bead simultaneously after all the odd [n], for then $k_0 = 0$, $k_1 = 1$, $k_2 = 0$, $k_3 = 1$, $k_4 = 0$, $k_5 = 1$, …; $k_\omega = 1$. The disjoint union—Figure 8—of the last two pictures, has once again the same length, but is in yet another similarity class for now $k_0 = 1$, $k_1 = 1$, $k_2 = 1$, $k_3 = 1$, $k_4 = 1$, $k_5 = 1$, …; $k_\omega = 2$. We get exactly the same invariants as these, if instead of even and odd, we use some other partition of the finite ordinals into two infinite subsets. To check that the map is similar detach as above, from a component having arbitrarily long odd [n]'s, all its even [n]'s, then attach these to the other component, now detach from this all the odd [n]'s and attach to the first component.

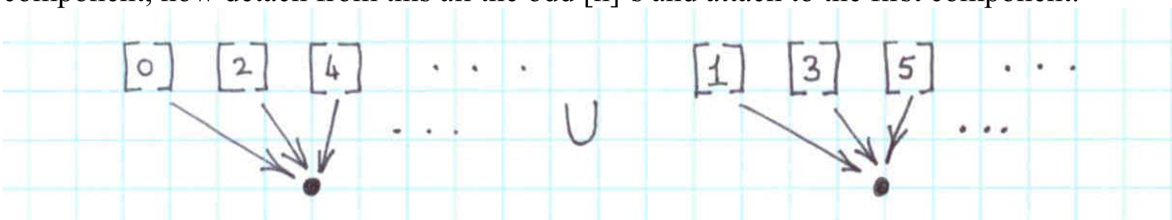

Figure 8

In general, by suitably detaching and attaching strings, *a normalized picture of a map can be replaced by any other having equally big sets* $S_{\alpha,n}$. Assume otherwise, and let β be the smallest height at which only some proper maximal subsets of nodes has been made 'good'. That is, some 'bad' node v remains with the totality of strings terminating v not of a desired 'kind', i.e. does not have the specified cardinalities of attached strings with last nodes of specified heights having as upper limit the height β of v. If modulo a subset of strings having a smaller upper limit, this 'kind' is already available on a single available (not already made 'good') node, detaching superfluous stuff from it we could have made that node good. So only subsets with smaller upper limits are available in various available nodes, but then these could have been safely detached from each of these, and then attached to a single node. It follows that there is no such 'bad' v..

Thus the similarity class of a graphic reduced locally nilpotent map is determined by the invariants k_α. Also we have an abundance of examples, any picture that terminates and does not have left infinite strings in it is an example, and we can draw hordes of these examples of any length we like. Still, the very nature of these pictures tells us that *the invariants satisfy some conditions*. The number of nodes of height α, a limit ordinal, is the sum of $k_{\alpha+n}$ over all finite n. The requirement of height tells us that, given a ordinal θ less than α we should be able to attach to each of these nodes a different string whose last node has height between θ and α, so the sum of all k_β with $\theta < \beta < \alpha$ must be at least as big as the sum before. It can be easily seen that these conditions are sufficient also, that is, we can easily build a normalized picture having the given k's once they hold.

Thus we have a satisfying and extensive generalization of Jordan's theorem, but an obvious lacuna remains. We have not delineated the exact force of the hypothesis 'graphic', that is, we have not characterized those nilpotent maps that are graphic. Towards the end I'll give an example of a non-graphic reduced locally nilpotent map, and state such a characterization: it is necessary and sufficient that a certain cohomological condition holds. This happens, for instance, if the dimension is countable.

Now, finally, we are ready to leave the locally nilpotent case, and return once more to our initial Figure 1. This, and a host of other counterexamples to Dinesh's question, will show that pictures are useful even now. However, we note that there are rather obvious linear maps f which are not graphic, e.g., multiplication by a nonzero scalar λ other than 1 (but now, being zero, the prime polynomial $f - \lambda$ is graphic). Thus, though we won't use these, there is room now for a more general notion, in which different arrows, even within the same component of the picture, can represent different prime polynomials of f. In the locally algebraic case this is avoidable, for the canonical direct sum decomposition allows us to separate the various prime polynomials.

§ 9. If $\alpha = \beta+1$, then $f^\alpha(V) = f(f^\beta(V))$, while for α a limit ordinal, $f^\alpha(V)$ was defined to be the intersection of all $f^\beta(V)$, $\beta < \alpha$. To see how K_f sits within V we need to look also at the analogous transfinite decreasing sequence of subspaces of the quotient linear map induced in V/K_f. For α a limit ordinal, the α th term of this sequence pulls back to the subspace of V which is obtained by intersecting all the subspace $K_f + f^\beta(V)$ as β runs over all ordinals less than α. This obviously contains all translates of K_f that intersect $f^\alpha(V)$, and in one important but special case it contains nothing else: *if* K_f *has an* f-*invariant complement then*

$$K_f + \bigcap_{\beta<\alpha} f^\beta(V) \quad = \quad \bigcap_{\beta<\alpha} (K_f + f^\beta(V)) \qquad\qquad (*)_\alpha$$

for all limit ordinals α. To see this check that, if $V = K_f \oplus C$, $f(C) \subseteq C$, then we have $f^\beta(V) = f^\beta(K_f) \oplus f^\beta(C)$ for all β, which implies that both sides of $(*)_\alpha$ are $K_f \oplus f^\alpha(C)$. The examples below will show that any $(*)_\alpha$ can fail. Thus the dimension of the right side of $(*)_\alpha$ mod its left side gives us a transfinte sequence of non-trivial invariants of the linear map f that measures how far its locally nilpotent part is from being a direct summand.

The first of these conditions $(*)_\omega$ *does not hold for* Figure 1. The span of the nodes on the its bottom most horizontal equals $\bigcap_n f^n V$, but only a codimension one subspace of the span of the nodes on any other horizantal line is in K_f, namely, that defined by the condition 'sum of coefficients is zero'. So the left side of $(*)_\omega$ is a proper subspace of V. On the other hand, the right side of $(*)_\omega$ equals V. Indeed, given any n, if we go sufficiently to the right on any horizantal its nodes—of course, no node of the picture is in K_f—are all in $f^n V$, so we have in fact $K_f + f^n(V) = V$ for all n.

In the notation of § 7, Figure 1 is the same as $[\omega] \to \bullet\!\to \bullet\!\to \bullet \ldots$, that is, it is obtained by laying a right infinite string after the last bead of ω. More generally, a similar calculation shows that, *for any limit ordinal* α*, the condition* $(*)_\alpha$ *fails for the linear map* $[\alpha] \to \bullet\!\to \bullet\!\to \bullet \ldots$, *while all other other conditions* $(*)_\beta$, $\beta\neq\alpha$, *hold.*

For each limit ordinal α let the α-*adic topology* on V be the one whose open sets are all translates of all the subspaces $f^\beta V$, $\beta < \alpha$.. In this topology the right side of $(*)_\alpha$ is the closure of K_f while the closure of the origin is $f^\alpha(V)$. Thus $(*)_\alpha$ is the same as saying K_f is *quasi-closed* in this topology, that is, it surjects to a closed subspace of the associated Haussdorff quotient topological space $V/ f^{\alpha}(V)$. Our counterexamples so far were non-Hausdorff, however there are Hausdorff ones too.

Draw in the first quadrant of the x-y plane a straight line having a positive slope m and let the nodes be all points on or below this line which have both coordinates non-negative integers; and, if its y-coordinate is positive, map a node to the one immediately below it, otherwise, to the one immediately to its right. The linear map corresponding to this picture—see Figure 9—has $f^{\omega}V = 0$, that is, its ω-adic topology is always Hausdorff, however $(*)_{\omega}$ *holds for this map if and only if* $m \leq 1$.

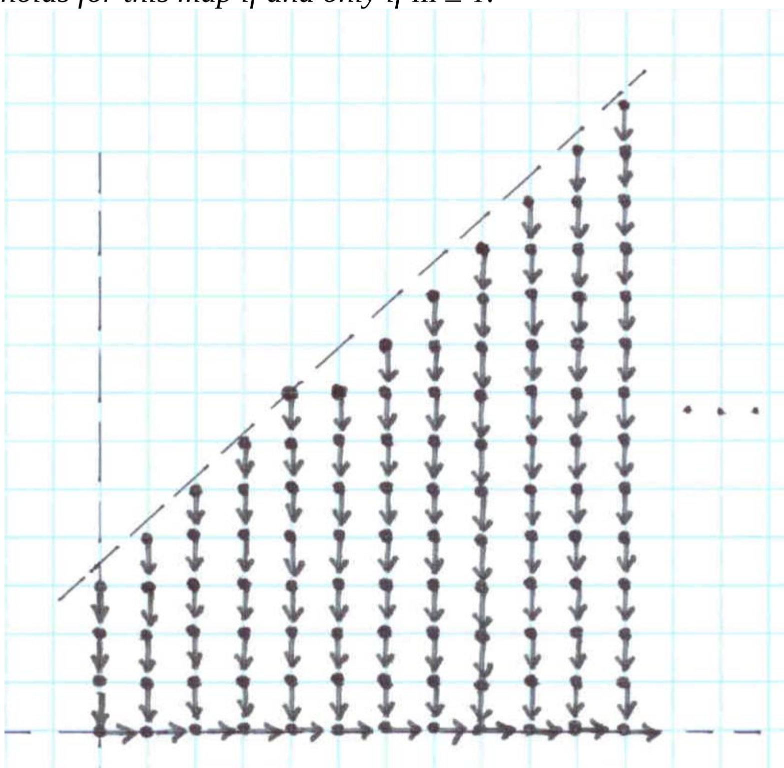

Figure 9

The left side of $(*)_{\omega}$ is now only K_f, and since it contains no node, it is certainly smaller than V. On the other hand K_f does contain a codimension one subspace of the span of all nodes on any slope 1 line, namely, that defined by the condition 'sum of coefficients is zero'. If m is bigger than 1, then for any n, all nodes sufficiently to the right on any such line are in f^nV, which shows that $K_f + f^n(V) = V$ for all n, so the right side of $(*)_{\omega}$ equals V. By sliding it down if need be, we can assume that the slope m 'roof' of our picture has at least one node (if m is irrational it has only one node, otherwise infinitely many). If $m \leq 1$ it is easily seen that the span C of the orbit of such a node is an f-invariant complement of K_f, à fortiori, $(*)_{\omega}$ must hold.

For $m \leq 1$ we can go further: some obvious stripping shows that the map is given by a disjoint union of strings (of which just one, the aforementioned orbit, is infinite). This is quite like what we saw before for any locally nilpotent map with $f^{\omega}V = 0$, namely, that it can be represented by a disjoint union of strings. However, for m > 1, our example shows that this need not be always so if the map is not locally nilpotent.

§ 10. *If the linear map* f *is graphic then the conditions* $(*)_\alpha$ *guarantee that the locally nilpotent part is a direct summand.* The argument is sketched below..

We analyze each *component*, that is, a maximal subset such that any two nodes have a common successor, of a given pictorial basis separately. There are five mutually disjoint possibilities. **(a)** If the component has an arrowless node or a *loop* then this is necessarily unique, and is its 'sink', that is, everything finally drains into it.

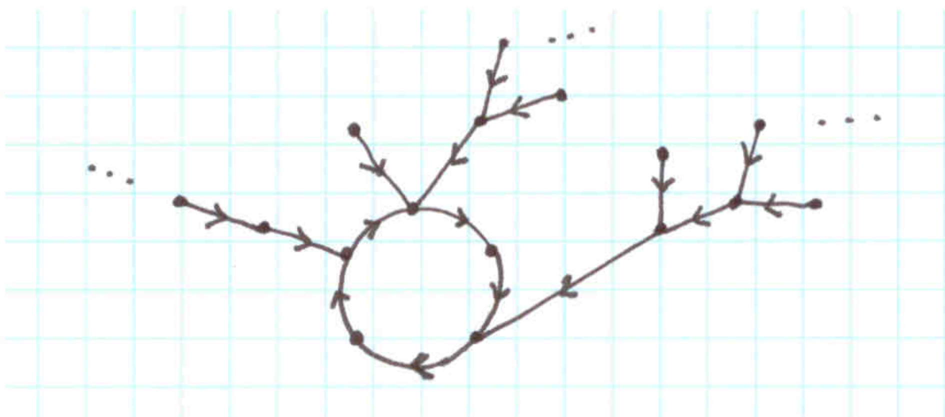

Figure 10

The component is thus locally nilpotent, or, if it sinks into a loop with $t \geq 1$ arrows, its locally nilpotent part has codimension t, with as an f-invariant complement the span of the loop; indeed, we can strip everything else off the loop as a basis for the component's locally nilpotent part. **(b)** The component contains a doubly infinite string of nodes … → •→ •→ •→ •→ … Now everything strips off from this as the locally nilpotent part, and has the span of this doubly infinite string as an invariant complement. **(c)** The least ordinal bigger than the heights of its nodes is ω, and moreover, the height of any node of its right infinite strings •→ •→ … is eventually only one more than the height of the previous node (note that two such strings of the component eventually coincide, so we need to check this condition for only one string). By generalizing slightly the argument given before for the $m \leq 1$ case of Figure 9 it follows that the locally nilpotent part strips off a suitably chosen right infinite strip whose span serves as its invariant complement. **(d)** If instead the height of the nodes of such a string frequently exceeds that of the previous node by two or more then, by imitating the argument given before for the $m > 1$ case of Fig. 9, it follows that $(*)_\omega$ fails for the component, and so also for the map as a whole. **(e)** The remaining possibility is that the least ordinal bigger than the heights of the nodes of the component is a limit ordinal α bigger than ω. If $\alpha = \beta + \omega$, $(*)_\beta$ or $(*)_\alpha$ fails, depending on whether the height of a node of a right infinite string issuing out of any height β node is, or is not, eventually only one more than that of the previous node. If α is not of the type $\beta + \omega$, then $(*)_\alpha$ fails.

Thus t*he conditions* $(*)_\alpha$ *preclude the possibility that any component is of type* (d) *or* (e), and for the allowed possibilities (a), (b) or (c) we have explicitly found an invariant complement for the component's locally nilpotent part, their direct sum now gives us the required invariant complement C of K_f .

Indeed, the proof yields *a classification of all graphic linear maps satisfying* $(*)_\alpha$. The invariant complement $C \approx V/K_f$ we constructed has as basis a disjoint union of $s_{\pm\infty}$ doubly infinite strings, $s_{+\infty}$ right invariant strings, and c_t disjoint loops with t arrows, $t \geq 1$. The rest of the given graphic basis was stripped off to a basis of K_f. If this has $s_{-\infty}$ disjoint left infinite strings … → •→ •→ • we can further strip off from these a graphic basis of an invariant complement - a reduced nilpotent map - of their span in K_f. As in § 8 we now normalize this basis eliminating forks at non limit ordinal heights, and for each limit ordinal α, denote by $k_{\alpha+n}$ the number of n arrow strings issuing out from height α nodes. *The similarity classes are in bijective correspondence with the sets of invariants* $\{ s_{-\infty}, s_{\pm\infty}, s_{+\infty} ;\ k_\alpha ;\ c_t \}$. The left and doubly infinite strings, and the loops, form a basis of the stable part $f^\lambda V$, with the left infinite strings alone spanning the locally nilpotent part of $f^\lambda V$, while the loops alone span the intersection of C with the locally algebraic part, which must coincide by §3 with the direct sum $\oplus\ \{K_\varphi : \varphi \neq f\}$. Thus, for such linear maps, this direct sum is exhausted by the *periodic points* of V, namely those whose orbits are finite, in particular, all *eigenvalues of such an* f *are necessarily roots of unity.*

It seems likely that one can drop the conditions $(*)_\alpha$ altogether and even classify all graphic linear maps. Indeed, if the locally nilpotent part of a component of type (d) or (e) is known to be graphic, then the component would be classified by the k_α's of this part, and the unique *germ at infinity of the height sequences* of its right infinite strings. It may even be possible to push further still, and classify all linear maps that are graphic only in the general sense alluded to towards the end of § 8.

§ 11. Though components are prominent in the above arguments, their number is, as such, not an invariant of the graphic map, because stripping can increase it. This anomaly is removed by using an *augmented picture*: add a new node 0, and arrows to this new node from previously arrowless nodes (a very natural thing to do really, for it is $S \cup \{0\}$, and not S, that is always mapped into itself by f). The *number of components of an augmented basis is an invariant of the graphic map*, now stripping merely enlarges the component containing 0. In particular, for a graphic linear map satisfying the conditions $(*)_\alpha$ this is one more than $s_{\pm\infty} + s_{+\infty} + \Sigma_t\, c_t$, the component not counted in this sum being that containing 0, i.e., the locally nilpotent part. So $s_{\pm\infty} + s_{+\infty} + \Sigma_t\, c_t = \beta_0$, the *zeroth Betti number*—definition below—of the one-dimensional simplicial complex determined by the arrows of the augmented picture, modulo its distinguished vertex 0.

Using coefficients from the underlying field—other coefficients are also useful—the vector space of *zero-dimensional chains* C_0 , the span of all the nodes mod that of the node 0, coincides with V. On the other hand, the vector space of *one-dimensional chains* C_1, the span of all the arrows (s , fs), coincides with the *graph of* f, that is, with the subspace of $V \oplus V$ consisting of all ordered pairs (v , fv). The boundary of each arrow is its initial node minus its final node, that is, the *boundary operator* ∂: $C_1 \rightarrow C_0$ is given by $\partial(v, fv) = v - fv$. If we identify C_1 with V under the isomorphism that takes v to (v , fv) we can identify $\partial = 1 - f : V \rightarrow V$. Thus the *zero-dimensional homology* mod 0 of the augmented picture is given by $H_0 \approx V/\text{im}(1\text{-}f)$, its dimension is the aforementioned zeroth Betti number β_0 . The *one-dimensional homology* is given by $H_1 \approx \ker(1\text{-}f)$, the subspace consisting of all the *fixed points* of the linear map. Its dimension is the *first Betti number* β_1 of the picture, so dimension of the fixed subspace $\beta_1 = \Sigma_t c_t$, the number of loops in any pictorial basis; indeed, by taking for each loop the sum of its nodes one obtains a rather canonical basis of the fixed subspace.[3] Note $\beta_0 = \beta_1 + s_{\pm\infty} + s_{+\infty}$, which suggests—even though the difference of cardinals $\beta_0 - \beta_1$ is not always defined—that $s_{\pm\infty} + s_{+\infty}$ ought to be called the *Euler characteristic* of the picture. This is an example of the *index of a linear map*, that is, dimension of cokernel minus dimension of kernel, a very important notion from functional and global analysis.

Since $f^n(S \cup \{0\}) \subseteq S \cup \{0\}$, the same nodes span also a picture of f^n, whose invariants can be worked out in terms of the invariants of the picture of f. The first Betti numbers $\beta_1(f^n)$ turns out to be linear combinations of the invariants c_t , and conversely, the c_t's are linear combinations of the first Betti numbers of the iterates of f. In fact, *essentially all the similarity invariants of* f *are homological in nature*, however some more computations, which we won't pursue here, are needed to fully justify this remark; instead, I will now conclude with some words about cochains and cohomology.

§ 12. The vector space C^0 of *zero-dimensional cochains*, i.e., functions c from the nodes of a given picture into the field F of coefficients, coincides with the dual V* as each c extends uniquely to a functional on V, and the dual map $f^*: V^* \rightarrow V^*$ is given by 'pulling back' each c under the arrows of the picture, i.e., $(f^*c)(s) = c(fs)$. It is natural to consider analogous questions for this map too, for example, when exactly does the locally nilpotent part of f* have an invariant complement in V* ?

The *support* of a cochain c is the set supp(c) of all nodes on which it is nonzero. Since $\ker(f^*)^n$ consists of all cochains c such that all nodes of supp(c) have height less than n, *the locally nilpotent part of* f* *is contained in the invariant subspace of cochains vanishing on nodes of infinite height*. This easily reduces the above question to the case when the picture is merely a disjoint union of finite or right infinite strings.

For just one string with n arrows, the nth pull-back of any c is zero. For just one right infinite string, the locally nilpotent part of the dual map consists of all cochains that are eventually zero, and, since f* is now onto, this has an invariant complement. From this, however, one cannot conclude that the same is true for the dual map of a disjoint union of such strings, because its locally nilpotent part can be much smaller than the direct product of the locally nilpotent parts of the dual maps of the strings.

---

[3] We recall that 'the number of i-dimensional holes, $i \geq 0$' alluded to in the very beginning of this talk is also, more formally, called the i*th Betti number*, and is defined similarly using the boundary operator ∂ of any simplicial complex triangulating the smooth manifold.

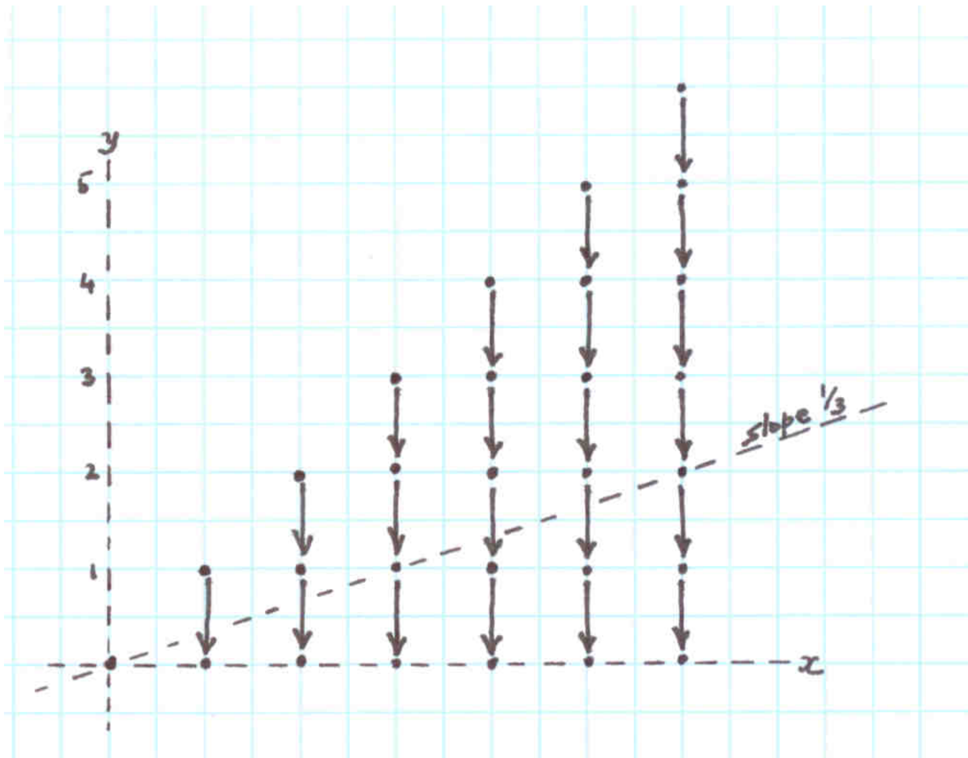

Figure 11

Indeed, for the disjoint union of the finite ordinals, i.e., *for Figure 10, the locally nilpotent part of the dual map does not have an invariant complement,* in fact, the necessary condition $(*)_\omega$, i.e., $K_{f^*} + \cap_n (f^*)^n V^* = \cap_n(K_{f^*} + (f^*)^n V^*)$ does not hold. The nth pull-backs comprising $(f^*)^n V^*$ are cochains supported on or above the horizantal y = n, so their intersection is zero. Thus the left side coincides with the locally nilpotent part, which consists of cochains supported above some line with slope 1. In particular, the left side does not contain a cochain c that is nonzero on a node iff the node is on a given line through the origin whose slope—1/3 in Fig. 11—is a positive rational less than 1. On the other hand, for any n, all points of this line sufficiently to the right are above y = n, which shows that c is in $K_{f^*} + (f^*)^n(V^*)$ for all n, that is, c is in the right side $(*)_\omega$.

A straightforward generalization of this argument now settles the probem: K_{f^*} *has an invariant complement iff* K_f *has an invariant complement, and moreover,* $s_{+\infty}$ *is finite and the* k_α*'s are all zero for* α *bigger than some finite ordinal.* The proof in fact classifies all pictures having such cochain maps—upto stripping it must be a disjoint union of some doubly or left infinite strings, loops, finitely many right infinite, and finite strings of a bounded length—and it seems reasonable to hope that one can even classify all linear maps that are *cographic*, that is, arise as dual maps of pictures.

Of course (just as for graphic maps) restrictions of cographic maps need not be cographic. In fact it is not hard to see that the restriction to K_{f^*} of the dual map f* of Figure 11 is not cographic. Moreover, *the restriction of the* f* *of Fig. 11 to* K_{f^*} *gives us a reduced locally nilpotent map which is not graphic.* To see this last, note that c is in ker(f*) iff it is supported on some nodes (t,t), and is in $\text{im}(f^*)^n V^*$ iff all such $t \geq n$. So k_α

$= \dim[(\ker f^*)\cap(f^*)^{\alpha} V^*/(\ker f^*)\cap(f^*)^{\alpha+1}V^*)]$ is zero unless $\alpha = n$, a finite ordinal, and then $k_n = 1$. But we know already that the only graphic reduced nilpotent map with these k_α's is the map f in the space of chains of Figure 11. This vector space is countable dimensional, while K_{f^*} is not, a contradiction.

In the above instance, we obtained identical k_α's for the chain and the cochain map of a picture. This is exceptional, the cohomological invariants of a picture, though related to its homological invariants, are usually different.

The vector space C^1 of *one-dimensional cochains*, i.e., functions from the arrows (s,fs) to the field F of coefficients, identifies also—each arrow being determined by its starting point—in an obvious way with V*. The *coboundary operator* δ: $C^0 \rightarrow C^1$, defined by $(\delta c)(s,fs) = c(s) - c(fs)$, thus identifies with the dual of $1 - f : V \rightarrow V$. The *zero-dimensional cohomology* of the augmented picture, relative to its distinguished node 0, is therefore $H^0 \approx \ker(1\text{-}f)^*$, the subspace Hom(V;F) of V* consisting of all *equivariant zero-dimensional cochains* c, $f^*c = c$, i.e., those that take a constant value on each component, this constant value being 0 on the component of the node 0. So $\beta^0 = \dim H^0$ equals $\beta_0 = \dim H_0$ iff this cardinal is finite, otherwise it is strictly bigger. Turning now to the *one-dimensional cohomology* it is $H^1 \approx V^*/\text{im}(1\text{-}f)^*$. Since functionals of $\text{im}(1\text{-}f)^*$ annihilate $\ker(1\text{-}f) \approx H_1$, by associating to each cohomology class $[c] \in H^1$ the restriction c: $\ker(1\text{-}f) \rightarrow F$, we obtain a well-defined surjection—but not necessarily a bijection, because, in the infinite dimensional case, the annihilator of ker(1-f) can be bigger than $\text{im}(1\text{-}f)^*$—of H^1 onto the dual of H_1. So $\beta^1 = \dim H^1$ equals $\beta_1 = \dim H_1$ iff this cardinal is finite, otherwise it too is strictly bigger.

Using only cochains vanishing on U, we can associate to each invariant subspace U of V, a *relative cohomology* H*(V,U), which is tied to H*(V) and H*(U) by the usual *exact cohomology sequence.* Using more generally the coboundary operator $\delta = \varepsilon - f^*$, in the vector space L(U,W) of all linear maps from U to W, we get *cohomologies* $H^*(V;W)$ *and* H*(V,U;W) for each fixed linear map ε: $W \rightarrow W$, instead of just the identity map of the one-dimensional vector space F. Note that one has now $H^0(V;W) \approx \text{Hom}(V;W)$, the vector space of *all equivariant linear maps from* V *to* W, and, using standard categorical language, $V \rightarrow H^1(V;W)$ coincides with the *first derived functor* of $V \rightarrow \text{Hom}(V;W)$. Finally note that equivariant maps induce obvious contravariant and covariant maps in these cohomologies; in particular, f induces a contravariant map in its own cohomologies $H^*(V;W)$ and H*(V,U;W), and the iterated images of this induced map, which are defined as before, are denoted by $f^{\alpha}H^*(V;W)$ and $f^{\alpha}H^*(V,U;W)$, where α is any ordinal.

Pictures give a covariant functor from the category of partially defined functions $S \rightarrow S$ *to the category of linear maps* $V \rightarrow V$. We have already computed some (co)homologies for linear maps occuring in the image of this functor, that is, for graphic linear maps. It is fairly routine to push these computations further and increase the number of these necessary (co)homological conditions. Moreover, *one should be able to write down a brief and explicit list of necessary cohomological conditions that probably "characterizes" graphicity.* We won't propose such a list here, but do want to mention a couple of known cases to explain the quotation marks used in this remark.

(1) If a reduced map f which is locally nilpotent, is known to be graphic, then a straightforward computation shows $f^{\alpha}H^{1}(V,f^{\alpha}V;W) = 0$ for all α and W, and *the converse* (i.e., that this condition guarantees that such a map is graphic) *has been proved*,

(2) If a reduced map f which is injective, is known to be graphic, then it is easily seen that $H^{1}(V;W) = 0$, for any right infinite string W, but now, *the converse cannot be proved, however we can safely assume it as a new axiom!*

This intrusion of proof theory is striking, but if one thinks about it, it was perhaps inevitable from the very beginning. Everything we have said flowed from *the basic assumption* that the identity maps are graphic, i.e., that any vector space V has a basis S, i.e., that the axiom of choice holds. It is known that this statement cannot be proved from the other axioms of the current logic in which we do mathematics, but can be—and in fact has been—safely incorporated into this logic as an additional new axiom, because it is known that it cannot introduce any new contradictions (and it is the fervent, but alas, unprovable belief of most mathematicians that no contradiction will be found in current logic). For the case (1) the sufficiency of the cohomological condition can be barely squeezed out of the basic assumption by a fairly delicate argument, however in case (2) we run out of gas, so to speak, and the same situation re-occurs, i.e., the sufficiency cannot be proved from the extant axioms, but can be safely assumed as a new axiom.