

Harmony-Seeking Computations: a Science of Non-Classical Dynamics based on the Progressive Evolution of the Larger Whole

CHRISTOPHER ALEXANDER

University of California, Berkeley, and University of Cambridge

In this paper, I am laying out a proposal for a new form of computation, which focuses on the harmony reached in a system. This kind of computation resembles certain recent results in chaos theory and complexity theory. However, the orientation of harmony-seeking computation is towards finding harmonious configurations. It helps to create things, above all, in real world situations: buildings, towns, agriculture, and ecology.

This way of thinking about computation, though mathematical at base, is closer to intuition and artistic feeling than the processes we typically describe as “computations”. It is also useful, potentially, in a great variety of tasks we face in building and taking care of the surface of the Earth, and quite different in character since it is value-oriented, not value-free. It has everything to do with beauty. The harmony that is sought in these computations is indeed what we otherwise call “beauty”. But the result of harmony-seeking computations are not merely pretty or artistic. In most cases, they are also functionally and technically better than their inharmonious counterparts.

Harmony-seeking examples are taken from farming, art, architecture, embryology, physics, astrophysics, drawing, crystallography, meteorology, dynamics of living systems, and ecology.

1. BACKGROUND

1.1. Overview

The scientific importance of harmony-seeking computations arises in three ways.

First, there are a number of natural phenomena, especially those occurring in complex systems, that are approximated by present-day explanations, but which have annoying mismatches with reality. In a few cases that I describe in this paper, for example *Acetabularia* (§5.3), cosmology of large-scale structure in the universe (§5.7), group formations in bird flight (§7.2), cloud formations (§7.3), simulation of tree growth (§7.4), and snow crystal formation (§8.3), there is a possibility that the discipline of harmony-seeking computation can contribute something to solving puzzles and unanswered questions.

Second, there are cases where the conscious use of harmony-seeking computations strongly improves on current practice of design, planning, and ecology, in fields where it is our business to make beautiful and harmonious structures, but where, too often, contemporary architects and planners using current techniques fail to do this, for example, wind turbines (negative, §5.2), Tokyo apartment building (positive and negative, §5.4), Boston housing (positive, §10.1), drawing by Nakano (positive, §5.5), tree bole seat (positive, §3.4).

Third, there are other cases where study of traditional building processes, and the structures they generate, give insight into the nature of the harmony-seeking computations themselves, for example, St Mark's Square (§5.6), hayricks (§5.1), the Parthenon (§8.2).

This paper extends results first presented in the four-volume work *The Nature of Order (NOO)* [1]– [4]. The results that underpin this paper are the following. There is a structure, visible in any given part of the world, which we may call *the wholeness*. The wholeness is an abstract mathematical structure, existing in space. It captures what we may loosely consider as the global structural character of a given configuration, in itself and in relation to the world around it. The wholeness is a structure that exists at many levels of scale, and covers the interrelationships of the configurations at different scales. The primary entities of which the wholeness structure is built are *centers*, which become activated in the space as a result of the configuration as a whole. Centers have different levels of *strength* or coherence. The coherence of a configuration is caused by relationships among centers. There are 15 kinds of relationships among centers that increase or intensify the strength of any given center.

The 15 properties, summarized in §1.3, define the way that configurations within a configuration help each other. Within this scheme, unfolding of new configurations is a natural process, and can be understood and followed (see especially §10.1). We thus have a basis for making

computations about unfolding. These are somewhat similar to the bifurcations that have been observed and analyzed in complex non-linear systems [5], but they are much richer and more complex than the theory of bifurcations can at present contemplate. Unfolding occurs as a result of *wholeness-extending* transformations¹. These are combinations and sequences of 15 possible spatial transformations (based on the 15 properties) that determine how coherent centers may be built from one another. A computational theory of wholeness-extending transformations does not yet exist, but my aim here is to show how unfolding is built from these transformations, and how the outline of a new (computational) theory of unfolding can be established.

1.2. Wholeness: an illustrative example

To make more explicit what the term “wholeness” means, I describe the phenomenon of wholeness as it appeared in a series of psychological experiments. In studies performed at the Harvard Center for Cognitive Studies in the early 1960s [6], my co-workers and I discovered a curious phenomenon. When confronted with simple black and white strips, and asked to arrange them according to their relative similarities, different perceivers chose one of two different strategies, illustrated in figure 1.

¹ Throughout the four volumes of *NOO* [1]–[4], I call these transformations that move a system forward in time *structure-preserving*. This term is sometimes used in mathematics to refer to transformations that preserve some particular structural aspect of a given system, but by convention the particular aspect may be arbitrarily chosen. As a result of many discussions with colleagues, I have concluded that “structure-preserving” is not quite the right name here. The transformations are structure-confirming, structure-enhancing, structure-extending, structure-strengthening, structure-sensitive. Mere “preserving” sounds rather too static, too simplistic, and too arbitrary for the more active unfolding. The transformation preserves the whole, which is not arbitrary, but an objectively present feature of reality. Ability to see this feature of reality is dependent on training, on experience, on the observer’s learned ability to see the whole. We might better call these transformations *wholeness-extending*.



Figure 1a. Sequential-digital: Reading the individual strips left to right (from [6])




Figure 1b. Figural-holistic: Seeing the individual strips as wholes or as gestalten (from [6])

In figure 1a we see five columns, organized and defined by a left-right reading of the patterns, somewhat akin to the Dewey decimal system of organizing books in a library or a system of ordering binary numbers according to their left-right reading. The right hand ends of the patterns play virtually no part in placing it on the board: they are random with respect to black and white. The left hand end in each group starts with the same figure, two black squares in the first column, one black square in the second column, and so on. The *gestalt*, or overall pattern of each strip as a whole, plays no part in the arrangement of the strips on the board.

In figure 1b, we see various groups, and what ties together the patterns in each group, is the *overall* kind of pattern they have. One group has steady alteration of black and white squares. Another has a single chunk of black, with chunks of white on either side. Here the basis of classification is just the opposite from what is going on in figure 1a. The patterns are grouped according to their individual overall configurations, *as wholes*.

Of the two, the figural arrangements are more interesting. To a scientist, it is likely that this more morphological way of looking at the configurations is more pertinent in the study of physical systems such as the pattern of emerging galaxies, or the morphological development of embryos. As such we may consider the figurally-based classification as deeper and more interesting. The sequential basis for classifying patterns might interest a computer scientist concerned with binary codes, but is scientifically less interesting, and is functionally less likely to be important in the behavior of real systems.

However, our experiments revealed a remarkable fact. Our experimental subjects were women students at Radcliffe college, and therefore considered intelligent, analytical, and selected for good performance in academia.

About 80% of the students chose the first, sequential arrangement. Only about 20% used the second, figural arrangement.

The form of classification used by each student was spontaneous, and not “led” in any way by the experimenters. The experimental instructions asked only that they arrange the strips on the gray board supplied, and place those that look similar near each other. Each student made a unique arrangement; no two were exactly the same. But they fell into these two sharply different *kinds* of arrangements.

We were startled to find that these students should spontaneously create arrangements based not on the inherent figural configurations in the material they were looking at, but instead on an arbitrary classification scheme that was external to the nature of the material. This was especially startling in light of the fact that either scientific training or artistic training would typically look towards the figural aspect of the strips as their more salient feature.

A further interesting fact emerged. Through our experiments we had invented, in effect, a way of calibrating a perceiver’s way of looking at the strips, by objectifying their mode of perception in the arrangement that perceiver created, and using that arrangement as a definition of how that person was seeing the world. We began a second series of experiments in which we gave people various pre-experiment tasks, selected, we thought, in such a way that they would tend to induce “figural” perception in the subject. We used more than a dozen different tasks, including “creative play” in which subjects were encouraged to make enjoyable or beautiful arrangements with the strips, various memory tasks, perception under difficult conditions, ability to describe the strips verbally, and so on. After training the subject, we asked them to arrange all 35 strips on the gray board that grouped patterns which “looked alike to them”. We used subjects who had never seen the material before, so there was no residual effect from earlier experience.

To our surprise only one training technique had a marked effect on the way the subjects perceived the strips. In this technique we showed each subject a single strip, flashed on a screen, and allowed them to study it and absorb its character. We then put on the screen, for half a second, an overall array in which all 35 strips were tightly but randomly packed in a rectangle. We offered the subject 5¢ if, in the half second available, they could find the single strip that they had studied previously. It is fiendishly difficult to do. Half a second is long enough to study only a few strips before the whole array disappears. After repeated failures, the subjects changed their strategy.

The only way to find the targeted strip was to relax: instead of an earnest, focused type of perception, one had to lean back mentally, and try to take in the whole board in a single glance. Although this required great concentration, in order to make it work the subject had to approach the task in a fashion that was inherently unfocussed and relaxed. When people learned this technique (which they discovered spontaneously: we never explained it to anyone) they were then able to locate the targeted strip in the array each time, and so earned their 5¢ for each display. And then, after *this* training, without any prompting, they tended to arrange the strips more in the style of figure 1b.

The experiment suggests that there is such a thing as “seeing the whole”. Indeed, we began to recognize that the “whole” or “wholeness” that is seen is, apparently, an objectively present structure. It is not always easy to have the ability to see it, and students trained in analytical thinking are less likely than others to be able to see it at all: we found in other experiments [unpublished] that a group of 6 year old children, and a group of adults with learning difficulties, had much *greater* facility for seeing wholeness than the college students. To teach someone to see the wholeness, one had to break down the analytical habits of a lifetime of education [6].

The ability to see the wholeness is a prerequisite for the material contained in this paper. It is also, I believe, a prerequisite for seeing many of the key phenomena that are now an essential part of physics and biology, just as it is a prerequisite for seeing art with understanding.

1.3. Fifteen Properties of Wholeness

In *NOO* [1]–[4], I once again took up the study of wholeness that had first occupied me in those Harvard experiments 40 years earlier. Among other fundamental aspects of wholeness I focus on certain structural features which appear to be underpinning the wholeness structure, as it appears in the geometry of physical things. In *NOO* I identify 15 structural features that appear again and again in coherent systems, and appear to play a major role in establishing the wholeness of these systems. The 15 are summarized briefly here.

1. **STRONG CENTERS.** Wholeness is composed of centers, and centers arise from wholeness. A given wholeness is coherent to the degree that the

centers within it are coherent.² Centers are recursive in structure. Each center that exists acts to strengthen other centers, larger and smaller.

2. LEVELS OF SCALE. When a configuration contains centers, these centers are associated with centers at a range of sizes that occur at well-marked levels of scale. The scale jumps between levels are small: in coherent systems the centers of different sizes are often in size-ratios of 2 to 1, 3 to 1 and 4 to 1. If the jumps are larger – for example 10 to 1 or 100 to 1 – without intermediate levels, the coherence tends to fall apart. This means that in coherent structures, the ladder or hierarchy of levels has evenly spaced rungs, and is continuous and smooth.
3. THICK BOUNDARIES.³ Strong centers typically, though not always, have thick boundaries. The thick boundary may exist in 1-, 2- or 3-D, and is made up of smaller centers that have the LEVELS OF SCALE relation to the larger center being surrounded. Thick boundaries typically form a transition zone of interaction, allowing physical, chemical, or biological processes to occur without contaminating their centers. The boundary is often only one LEVEL OF SCALE smaller than the thing it surrounds, thus may be equal to the radius of the thing surrounded. Boundaries help to form the “field of force” that creates and intensifies a center; they surround, enclose, separate and connect.
4. ALTERNATING REPETITION. When repetition of similar centers occurs in a coherent system, the centers typically alternate with a second system of centers, thus forming a double system of centers with a beat or rhythmic alternation, from the positive space between the repetitions. Centers intensify other centers by their repeating rhythm; when a second system of centers repeats, in parallel, it intensifies the first system by providing a kind of counterpoint, or opposing beat.
5. POSITIVE SPACE. In coherent systems, there is no “background”, or figure and ground. Instead, every bit of space is coherent, well shaped;

² Coherence of a center, and strength of a center are two ways of referring to the same quality that centers have in greater or lesser degree. The degree of this quality that is present in a center, indeed in any whole, can be measured by empirical methods described extensively in [1].

³ In *NOO*, the property is called simply BOUNDARIES. I now use the term THICK BOUNDARIES, since it more clearly distinguishes the boundaries that work most effectively to create coherence, and also eliminates confusion with the common mathematical use of boundary as a dimensionless interface.

and the space between coherent bits of space are also coherent and well-shaped. Thus every bit of space swells outward, is substantial in itself, and is never the leftover from an adjacent shape – like ripening corn, each kernel swelling until it meets the others, each one having its own positive shape caused by its growth as a cell from the inside. The positiveness of space is difficult to pin down exactly, but it is like a weak kind of convexity, or quasi-convexity. In systems where the space is positive, the principal elements of space are nearly all quasi-convex, and the pieces of space between these elements are also quasi-convex.

6. GOOD SHAPE. This describes a particular, coherent quality of the shapes that occur in or around a coherent center. This kind of “good” shape is somewhat unusual, and is marked by the fact that the shape itself is made up from multiple coherent centers which together form the shape, and of other coherent centers which together form the shape of the space around the shape.
7. LOCAL SYMMETRIES. A local symmetry is a symmetry of a localised region of space that is not possessed by the space beyond . Strong centers often have strong local symmetries within them, and local parts of space with strong symmetries are typically strong centers. This feature binds together smaller centers within the whole, further creating coherence.
8. DEEP INTERLOCK AND AMBIGUITY. This occurs where coherent centers are “hooked” into their surroundings, making it difficult to disentangle the center from its surroundings. Often there are ambiguous zones which belong both to the center and to its surroundings, again making it difficult to disentangle the two. (See, for example [7]).
9. CONTRAST. Every center relies to some degree on the contrast of discernible opposites, and on its differentiation from the ground where it occurs. It is intensified when the ground, against which it is contrasted, is clarified and is itself made of centers and POSITIVE SPACE. The essence of this feature is that this differentiation arises from the degree or sharpness of contrast that is attained between adjacent centers. Note, though, that too much contrast is likely to be harmful, and must then be offset by NOT-SEPARATENESS, below.
10. GRADIENTS. Centers are generated and strengthened by gradients of size, shape, or quality. Thus any quality among a system of centers that

varies systematically produces a gradient, and this gradient, by pointing to a particular center, helps to build that center and to intensify its coherence.

11. ROUGHNESS. In coherent structures we usually find a rough arrangement and repetition of centers rather than exact repetition in shape, spacing and/or size. Thus apparently similar centers are different according to context, allowing each part to be adapted to the geometric constraints around it, modifying details of the repeating structure as it needs to be. Texture and imperfections are generated, and in part create the possibility of uniqueness and life.
12. ECHOES. Within coherent configurations there are often deep underlying similarities or family resemblances among the elements. These similarities are often characterized by typical angles, and typical curves, so that they generate what appear to be deeply related structures, sometimes so deep that everything seems to be related.
13. THE VOID. In the most profound centers that have perfect wholeness, there is often at the heart of the structure a void that is large, undifferentiated, like water, infinite in depth, surrounded by and contrasted with the clutter of the structure and fabric all around it.
14. SIMPLICITY AND INNER CALM. Essential to the completion of a coherent whole is a quality of simplicity. Every structural feature that is unnecessary has been removed, so that the remaining structure has slowness, majesty, quietness.
15. NON-SEPARATENESS. This describes the connectedness of each center to the world beyond it. When a whole is a living center, we experience it as being at one with the world around it. When not-separateness exists, visible strands of continuity of line, angle, shape, and form, connect the inside of a living center with the parts of the world beyond that center; it is impossible to draw a line separating the two.

These 15 properties are thoroughly explained in [1]. In [2], these structural features are also shown to form a basis for the wholeness-extending transformations that create life and coherence as configurations unfold.

2. THE SUBTLE PROCESS OF ADAPTATION

The core topic of this paper is geometric adaptation. In many real world systems, both in purely natural systems, and in those places where people form communities with animals, plants, and other people, the central observable is a close-knit adaptation of the system elements, usually arising over time, and expressed in the intricate geometry of the system.

This close-knit geometric adaptation has not been a major focus of scientific study, because it eludes simple algorithmic formulations. Studies of co-evolution and ecological evolution have moved in this direction, but they rarely concentrate on the geometry of the evolving system. That is not because this adaptation is too complex to be modeled, but rather because the elements of such are so simple, and so rooted in common sense, that they nearly elude the algorithmic and algebraic formulations that we view, wrongly, as more sophisticated.

For example, if a farmer places a row of fence posts, then runs a top rail, braces it here and there where it seems needed, allows it to relate in natural ways to declivities in the ground, or to nearby trees, this is supremely ordinary; it is characterized entirely by common sense, and by the farmer's ability to pay attention to the situation of each post, each rock, each bit of soil, each slope – and do it right. This oh-so-simple process eludes algorithmic formulation, because algorithmic formulation is not well tailored to this task, and at best only partially helpful in allowing us to grasp what is really going on.

That is not to say that a sensitive fence-building process is trivial or unimportant. On the contrary, the character of this minute, step-by-step adaptation is vitally important in the world, and we have been ignoring it, in recent decades, at our peril. But we do not have a theoretical model that emulates this process. As a result, the ability to perform adaptations in the real world according to such a process has been worn away and destroyed by other processes that are largely bureaucratic – often too bureaucratic – and, in their essence, algorithmic. The planners, building officials, construction companies, and engineers who have redefined everyday processes during the last 100 years, working in a broad context of algorithmic thinking, have, without explicitly intending to, destroyed a far more subtle process.

Until that subtle process is acknowledged, and then redefined in modern terms, it will not have the status it requires to play an effective role. The deep adaptation that nourishes the physical world *requires* this kind of adaptation. We can think of this adaptation process as a highly sophisticated

computation, performed on real sticks and stones, producing deep and subtle results.

The progress of an evolving (unfolding) natural landscape, or the development of an embryo, have similar qualities. As cell division progresses, new cells take shape within the context of the surrounding cells, and, at the same time, adapt, so that they are shaped by these surroundings, and simultaneously play an active role in shaping their surroundings. Again this process has not been modeled geometrically. The prevailing assumption is that what is going on in the real world is a nearly random aggregation of simple mechanical processes, with no special coordination or behavior as a whole. Within that view, there is little to be learnt from studying the process; it is just number crunching, without new insight.

But this view is, I believe, mistaken. The cells' progressive adaptations, as each part rubs up to its neighbors, shapes them, and is shaped by them, coordinates the whole. Some profound coordination of the whole is occurring. There are strong reasons to think that this aggregation of apparently random events is, instead, a highly organized wholeness-extending process, in which the process in the large progressively pays attention to the whole, reflects the whole, and extends and makes more beautiful the whole.

That is what makes the wholeness-extending process worth studying. It is a kind of computation, unfamiliar to conventional mathematics, but a computation nonetheless, and one that reaches profound results. By observing this kind of computation going on, and then, hopefully understanding it well enough to simulate it, we may lead to a new era of our ability to think.

In this paper, I rely heavily on examples. The subject of harmony-seeking computations is difficult, and one builds a sense of its feasibility by considering many examples from different spheres, and slowly grasping the general propositions that underlie all of them. Harmony-seeking computations occur in many parts of nature, and can also occur in human creative processes. The constant awareness of these different spheres, and the process of comparing them, is what makes this kind of computation interesting.

3. HARMONY-SEEKING COMPUTATIONS

3.1. Relation of a Given Computation to the Larger Whole beyond it

In the examples I give, an essential feature of the process described is that the new configuration that is created does something to participate in, strengthen, extend, and enlarge the force or “presence”, of some larger configuration in the world around it. In many cases, this larger configuration is not itself changed by the “computation”, but it nevertheless enters into the computation as a crucial input.

If a lane exists, connecting two villages, and the road between them is then asphalted, the new, wider, harder road will primarily affect the *whole* configuration of these two villages, a tract of land much larger than the lane itself. Indeed, the new road may (according to its position in the region) also affect and strengthen an even larger part of the region’s economic network. We recognize this as a fundamentally useful act (or, if undertaken without consideration of the wholeness, a potentially destructive act).

A played musical note is effective according to how it develops a melodic line, a rhythm, or a harmony. If it does all three simultaneously, it is the most effective. Once again, we recognize the usefulness of this act.

These are simple cases. Other cases involve creation of highly organized and complex artistic wholes. These have so far eluded science. Even the theories of bifurcation, symmetry breaking, chaos, and generative algorithms have not yet deeply plumbed the meaning, or the origin, of true complexity.

All this is hardly more than common sense. Yet the fact remains that this kind of adaptive process (the farmer building his fence with respect for the land) does not currently have an acknowledged part in theories of algorithms, in developmental biology, in architecture, or even in system theory. It is not part of the mental models in our current toolkit.

3.2. The Essence of Harmony-Seeking Computation

In essence, harmony-seeking computation creates new configurations, unknown configurations, and *good* ones, by taking off from a known configuration, extending, and enhancing, and preserving its wholeness, but without (necessarily) requiring human creativity. The process *itself* is creative.


Consider a given configuration with certain features visible, which define whatever wholeness we see in the configuration. In addition, in every configuration there are traces, hints, of dim structures not yet fully developed, existing in a latent form, “between the lines” of the configuration. What happens in harmony-seeking computation is that some

process “notices” these latent structures, and enhances them, develops them. Sometimes what develops is relatively modest with respect to the size of the entire configuration. At other times, large structures may also be latent in a configuration, and *these* are enhanced. If the whole-seeking computation identifies this latent larger whole, and strengthens it, so that what was before only barely visible now becomes strong and easily visible, the configuration seems, to an untrained eye, to have gone in a new direction all by itself. It is this process that is the essence of harmony-seeking computation.

Once we add knowledge and perception of this type of transformation to our toolkit, we will want to program our own wholeness-extending computations. Even were the task of discovering latent structures too difficult, one could imagine a computational system that examines possible moves forward from a given state of a system, tests them, and evaluates the extent to which each of these moves participates in and adds to the larger wholes in the world around it, then chooses the most harmony creating and the most successfully whole-creating. Such a process of computing, if it can be attained, would be enormously powerful, and powerful in its implications.

3.3. Example: a natural system: embryogenesis

Consider an example of embryogenesis, a growing mouse foot (figure 2). Each stage contains within it some structure that is defined, and some that is, for the time being, a chemical gradient, then a vague and fuzzy mass of jelly, both latent structures which anticipate the shape of the next step. Shortly afterwards, the embryo consolidates and solidifies what was merely latent only hours before.



Figures 2a, 2b, 2c, 2d. Step-by-step development of a mouse foot, showing its growth over four days: 11.5, 13, 14, and 15 days. (From [8], fig 10.5)

What are some of the transformations that constitute the wholeness-preserving character of these moves? The form starts with a barely visible axis from the point of attachment to the embryo’s body to the tip of the structure (2a). In figure 2b we see the emergence of a STRONG CENTER at the tip itself, forming a thick BOUNDARY (in one dimension) at the end of the arm, together with a dark fork that has emerged within the jelly. This center

is then accentuated by the appearance of a GRADIENT leading to the fingers, and this gradient is then embodied in figures 2c and 2d by LEVELS OF SCALE, CONTRAST, and LOCAL SYMMETRIES, finally finding expression in the GOOD SHAPE of the whole.

3.4. Example: an engineered system: a bench around a tree




Figure 3. Building a delightful seat around a tree. (a) staking out the boundary (b) weaving the basket (c) the top of the seat finished with turf as a cushion. (From [9])

Consider the example of a growing willow tree, and the act of the landowner who chooses to build a bench around the base of the trunk. The bench places a BOUNDARY around the tree. Compared to the length of the trunk, this is a relatively small step. But the bench starts from the cylinder of three trunks, putting a small ring-shaped structure around it (the seat). We may say that the possibility of this seat was inherent in the previous structure, that it was latent there, and that the bench builder simply made explicit and more solid the structure that was already present mathematically in a weak and latent form, at the base of the willow tree.

I mean this literally, not metaphorically. The ring-shaped structure that later finds embodiment and physical form in the seat, is already present before the seat is built, in the ring-shaped system of symmetries and sub-symmetries of the space around the tree, because they are induced by the presence of the trunk and its roughly cylindrical shape. The ALTERNATING REPETITION in the basket weave, then arises through a natural wholeness-extending transformation from the simple (but structurally weaker) homogeneous cylinder first contemplated. This statement contains the crux of my argument, and in this statement – when it is generalized – lies the mathematical kernel of what I have to say in this paper.

In each of the two examples – the mouse foot, and the bench around the willow tree – the steps build on the structure that was there before. They do not destroy it or interfere with it, but rather enhance it, elaborate it, deepen it. As a result, what arises has wholeness, coherence, and beauty. *That is the trick, in a nutshell.* By continuously preserving and enhancing and extending the wholeness of the existing structure, a beautiful thing arises, naturally. Yet, because each whole is unique, and the idiosyncrasies lying latent in it are also unique, the new whole that springs from this process is unpredictable, original, and creative.⁴

To perform this trick we must ask rather more precisely what it *means* to take steps that intensify the latent structure that is present in a configuration. Can we define this idea precisely enough, so that we can start doing it on purpose?

4. STRUCTURE OF WHOLENESS

The idea that every geometric configuration might have its own “wholeness” seems odd at first. However, there is a definable mathematical structure underlying any configuration that captures, or “is”, the structure of the whole configuration, its *global* structure. I call this structure *the wholeness*. It is not uniquely defined (mainly because the structure is so difficult and complex to describe precisely), but it is definite, and various attempts to represent this structure approximately, if they are any good, all have something in common. There is a deeper, global “something” there (for example, see the discussion of four Matisse self portraits shown in Alexander [1] page 97).

I first began to recognize this possibility in the experiments described in §1.2, to find out how different perceivers see configurations [6]. These experiments were inspired by the early 20th-century gestalt psychologists –

⁴ Susan Stepney raised the question: What if the existing structure is ugly? Can it be made beautiful by this process (like sand in an oyster resulting in a pearl), or should one not start from there? The answer is, in varying degrees, both. In nature, structure generally grows from very small disruptions in a simple and almost pure homogeneity. The beauty is there from the beginning, and continues as the system evolves. Human created ugliness always arises from moves that are non-wholeness-extending. Among human-made urban landscapes, where ugliness is common, it is tempting to tear down and start again. However as an architect, I myself usually resist this temptation, and leave much of the ugly part, pruning out only the worst. When wholeness-extending transformations are applied to this softened ugly place, it soon brings the system back to order, and leads to more wholesome results.

Wertheimer [10], Koehler [11], Koffka [12] and others – who wrote extensively about the perception and cognition of wholeness and of wholes in the world around us. As explained in §1.2, it was clear that the figural perception was somehow more real, while the analytical perception was arbitrary, and made according to arbitrary invented schemes. The figural quality was inherent in the material. The ability to perceive figurally was therefore of enormous importance. In broad terms, we may say that the wholeness or global configuration is a structure of the main features of any configuration – in particular *when it is seen as a whole* [13]. During the same series of experiments, we made another significant discovery. The relative coherence of different configurations could be measured by counting the number of locally symmetric sub-configurations it contains [14]. This measure gives results in strong agreement with other experimental results from cognition, memory, perception and so on. Thus a configuration and its structure do depend, in some considerable degree, on the nesting of LOCAL SYMMETRIES and on the structure of the system of locally symmetric subsets that occur in it. The fact that a simple measure (the number of locally symmetric subsets) to some extent characterizes the different configurations, lends further weight to the idea that for a given configuration, *the wholeness* is something real.

5. EXAMPLES OF HARMONY-SEEKING COMPUTATIONS IN DIFFERENT DISCIPLINES

5.1. Example: hayricks in a field

Figure 4 shows a photograph of two modest hayricks placed in a rolling meadow in Romania. The placement is done in such a way as to complement the land, to enrich the land. It is humble, self effacing, there is no ego visible, only concern that the land and its harmony should be enhanced.

Several readers of drafts of this paper have commented that they believe they could equally well explain the placement in terms of the pure practicality of drainage, wind protection, a flattish area on which to build the rick, a roughly equal spacing of the ricks in the meadow, so that each has a similar sized catch basin of the available growing hay, and to reduce the work of carrying while forming the rick. It is true that all these practical arguments are sound, and essential to the success of the rick-building venture. Attention to these functional issues is necessary. But is not sufficient.

In order to create a coherent object, the inner coherence of the physical thing and its geometry, must also play a foundational role. That means that the unfolding of the geometry from the wholeness of the place, must also have a paramount role. Without having this as an explicit focus of attention, a pure practical solution without consideration for this harmony, will fall flat, and most often create disharmony, not harmony.



Figure 4. Hay ricks in Romania [photograph by Radu B. Chindris].

During my work with traditional builders and craftspeople all over the world I have found that these so-called “humble” artisans, are always conscious of the harmony of the whole, and are always making sure that they are guided by it. I know of no exceptions. That makes me quite certain that we would be mistaken in believing that this Romanian field and its harmony could be successfully explained by pure 20th-century functionalism. It is the explicit search for harmony that underlies traditional practical techniques, and it is this element which makes the results so good.

How is one to characterize what is happening here? What is it, that we see, and recognize as so supremely successful? Of course, the placement is indeed primarily practical. The hay ricks are placed on almost-level spots, but ones which have a slight slope so that they can drain. But there are many subtle harmonies, also visible. The ricks are on the same scale as the trees on the ridge; the trees gently surround and enclose the place where the

ricks are; the rolling quality of the land, with its dips and hollows, and the hedges and trees as they are placed, is respected. The two hayricks are placed with considerable care, on slight outcrops of land, so that each hayrick is not only kept dry and free from rot, but so that it actively enhances and emphasizes the slight bumps in the land. In addition, the actual structure of the hayricks themselves (each rick a tree of light timber, with hay piled up on it), mirrors and echoes the character of the grass land and of the surrounding vegetation.




Figure 5. Impressionistic depiction of the wholeness structure in that field

Figure 5 is a sketch that, in broad-brush terms, approximates the wholeness present in that field. The wholeness includes the sinuous curves, the kinds of shapes which are present in the field, the dips and hollows, the surrounding trees, the trees along the ridge, and the fallen down enclosing fence below. I believe that the people who tended the land, placed the trees forming a hedge on the ridge, allowed the hollows in the land; those who built and placed these hayricks were consciously performing harmony-seeking computations every time they worked the land. There are ECHOES of shape and size between land and hayricks, ECHOES of a certain kind of curve, LOCAL SYMMETRIES in the ricks themselves, the placing of the ricks emphasizes naturally occurring STRONG CENTERS that are generated by shelves and flattened places, bounded so that the hayricks nestled in the land, are subdued and congruent, and inside the structure which exists. The hayricks are kindly to the land, and they are placed with enormous care. They follow the wholeness of the land itself. The ricks fit into the land

because of the way they function and because of what they are made of. In order to function well, they have to sit on a locally flat surface that is uphill from where water collects, and downhill from the crest of the hills so that the effort it takes to stack the hay is minimized. They are made of hay, the material produced by the place, so their scale, color and form are the same as their surroundings. Nothing foreign is introduced into the place for the sake of function alone. Form and function together are part of the wholeness of the place. The wholeness is a harmonious structural and geometric quality within the whole.

5.2. Example: giant wind turbines on the Danish coast

In order to accentuate the harmony that was consciously generated in the Romanian field, let us compare it with the much cruder results typical of 20th- and 21st-century thinking about engineering and design. The 20th-century thinking was, usually, more or less algorithmic. When the computations about subtle and delicate wholes are rationalized according to crude dollar-based ways of thinking, they usually go badly wrong, especially as far as respect for the background global structure of the landscape is concerned. The subtlety and depth of the existing global structure is not easily recognized by technological thinking.



Figure 6. Two wind turbines on the flat north coast of Denmark nr Nibe. (From [15])

For example, in figure 6, the land has a global wholeness that is made of layered, nested flat plates, a two dimensional system of plates. The ecologist who placed the two giant turbines was evidently unconcerned about creating harmony with this structure. Cover up the two towers on this photograph, so that you can see the global structure of the land as it was before the towers were placed there. It is hard to see it, because the towers are now so obtrusive, and so alien to the kind of structure that is there.

When an ecologist becomes too concerned with one aspect of sustainability (the importance of wind energy, as in this case), and plants 300 foot high windmills on an ocean front, it destroys the landscape. This is no longer a wholeness-extending transformation, but rather a well-meaning but ill-judged algorithmic computation of some inappropriate kind about cost benefit. In this case, the wholeness and the structure of the place have been severely damaged. Sometimes, the place, its landscape and its internal adaptations – often the result of centuries of patient work – are irredeemably destroyed.




Figure 7a. The wholeness before construction of the turbines



Figure 7b. The wholeness after construction of the turbines, which destroys virtually all the structure that was previously there.

In the coastal landscape where the Danish turbines have been placed, the structure that was there before (figure 7a) is flat like a disc, a great flat disc, the size of the bay, flatted at the edge where the waves meet the sand and the grassland just beyond it. This enormous gentle discus-shape, perhaps half a mile across, is suddenly violated by the two vertical structures, sprouting turbine blades dumped on it (figure 7b). The two massive vertical things have no relation to the structure of the subtle disc that lay in the land before. The slash of the new straight road is at odds with the curves of the land. And, in the foreground there are two teeny little sheds that look lost and

forlorn, too. The building site is not even harmonious with itself. Altogether, the structure has been violated.

Compare what is happening in the hayrick example, and in the turbine example. In the case of the hayricks, the process of placing them has left the previously existing structure alone, indeed the process enhances it, strengthens it, increase its harmony. But in the case of the turbines, their shape and placement are at odds with the previously existing structure, they damage it, they destroy it.

What is the “it” that is being left alone in the first case and is being destroyed in the second case? It is the *wholeness*, that system of centers existing in the geometry of each place that gives it its character, shape, organization, rhythm, feeling, and global configuration and relation to the land around it. One may describe the relation of the action to the wholeness like this. In the first case, the injection is friendly, helpful: what is there is enhanced, complemented, and continued. In the second case, the injection is violent, disrespectful, and structurally at war with the land that was there before.

My judgment about these two injections is not a romantic yearning for an idyllic past, but a structural judgment about something that is objectively present in one case, and missing in the other. It is also true that wholeness-enhancing processes allow growth and extension to continue while wholeness-destroying processes freeze things and inhibit improvement. One injection respects what is there, the other does not.

These structural judgments are about phenomena that are essentially mathematical in character. It is not just a question of feeling, but also a question of structural congruence, or the lack of it, that exists between an existing structure and an injected structure which has been brought into it. We need to understand more about this kind of structural respect.

In order to widen our understanding of the idea of harmony-seeking or wholeness-extending computations, we now examine a further series of examples of such computations, many of them occurring in the natural world.

5.3. Example: Evolution of the whorled cap of *Acetabularia* – 6 steps




Figure 8. Whorls forming at the tips of developing *Acetabularia* (from [16]). The bumps in the caps develop as shown at the right hand end of Fig.11a.

This example is taken from Brian Goodwin's work on the morphogenesis of *Acetabularia*, a single-celled alga which takes the form of a stem with a cupped whorl at the end (figure 8). Goodwin has studied the growth of this form, and the series of transformations that generate the form of the completed alga [17] [18].




Figure 9. Brian Goodwin's five diagram sequence (from [17] fig 4.6)

The key morphogenetic sequence that Goodwin describes is illustrated in his five-diagram sequence (figure 9), showing the stem bulging and then making the whorl. It goes as follows:

1 to 2. A hemisphere, formed by the bulging upward growth of the cell – turns into → a hemisphere with a small centrally symmetric swelling on top (an ogee curve).

2 to 3. The hemisphere with a small centrally symmetric swelling on top — *turns into* → an elongation of the swelling forming a neck.

3 to 4. The hemisphere with the elongated narrow neck — *turns into* → a neck with a flattened top, upper end.

4 to 5. The flattened neck on top of the hemisphere, grows a ring of small projections (a whorl) that ultimately turn into the umbrella shape visible in the completed algae.

5 to 6. The whorl formed around the top of the cylinder grows outward — *turns into* → a full umbrella-like ring of small projections that ultimately turn into the umbrella shaped cap visible in the completed algae.

As Goodwin says, this growth sequence is generated by natural physical phenomena. It has relatively little to do with DNA or genetic guidance. What is happening is a naturally occurring and inevitable progression of morphological transformations that arise directly from the geometry and dynamics of the form itself, to generate the next geometry [18].

It is significant that a prominent biologist has begun taking the emphasis of morphogenesis away from the influence of DNA, and has begun to see morphogenesis as an autonomous process in the geometry itself.

There is an interesting sequel to this story. When Brian first showed me the diagrammatic sequence illustrated above, I asked him if the progression from round-ended neck to flat-topped neck was correct. Why do you ask? he said. I said that, looking at this purely from the point of view of wholeness-extending transformations, I would have expected something different. It seemed to me very unlikely that a round topped structure could transform into a flat-topped structure. There is nothing like the flat top latent in the configuration. There is indeed a *ring* latent in a hemispherical configuration, but it is further down the curve. If you have a rounded hill, or hemispherical end of a neck-like structure, there is, inherent in that configuration, a latent structure something like a shoulder, which is incipient near the top of the hill and just below it (see figure 10b).

In figure 10a, we see such a structure on the prehistoric Wittenham Clumps near Oxford. It evidently was, for the builders who made these mounds, a natural way to continue the smooth structure of a rounded hill. This point of inflexion in the curve of the hill is a natural development from the smooth curve.



Figure 10a. One of the Wittenham clumps near Oxford

Figure 10b. A diagram of the structure created on the mound by the prehistoric embankment. It is a wholeness-extending transformation. The smoothness of surface, the volume, possibly the area of the curved surface, are all preserved. It is generated by a continuous and smooth process. It preserves the feeling and morphological character of the whole.

Figure 10c. A cut. It is a transformation of a kind, but it does not extend or preserve the wholeness of the curved volume. Rather, it disrupts it, and alters its feeling and morphological character very substantially.

If you are in doubt that it is “natural”, consider two possible transformations of a hemispherical hill, shown in figure 10. Figure 10b introduces a band below the top, with two points of inflection in the curve, as shown below. Figure 10c just flattens the top by squashing it down into a flat surface. Ask yourself which of these leaves the hill alone more, and preserves the global structure of the hill better. I think that you will agree that it is figure 10b. It is in this sense that the transformation in figure 10b is more natural than that in figure 10c.

So I asked Brian, referring to his *Acetabularia* diagrams: if such a configuration were latent, and did then transform to create such a ring-like

shoulder, then it would be that *shoulder*, rather than the top, that would be the most likely place where whorls and other irregularities might naturally form. So, are you sure that the real morphogenesis of *Acetabularia* doesn't go like that, rather than the way your diagram portrays it? As shown in figure 11, for instance, and the photograph in figure 8.




Figure 11. My analysis of the developmental sequence, based on W-E-transformations

Brian laughed, and told me: you are absolutely right: that *is* what actually happens – I just got the diagram wrong.


This story sheds light on the way harmony-seeking computations may be able to help observational biology, theoretical biology, and structural modelling and simulation. We come back to this later. The point for now is that it is possible to make accurate statements about what structures are latent at a given phase of morphogenesis; and it is therefore possible to say, objectively, how things are likely to unfold.

Look again at the photograph of the *Acetabularia* in figure 8, and look inside the caps: there is indeed a remaining *hillock*, just as my whole-seeking computation predicted. At the time I made the prediction, I had never seen enough detail of *Acetabularia* to know about this kind of cross-section. My guess arose solely from wholeness considerations, and from thought about what would be wholeness-extending.

5.4. Example: Two possible plans for a five-story apartment building in Tokyo

Figure 12a shows a plan for this apartment building in Tokyo, done by a student. Figure 12b is the plan of the five-story apartment building that he and I subsequently built on the site (the Emoto Building, Komagome, Tokyo [3], figure 12c). Though typical of many architectural projects around 1987 (the era of the project), the student's plan lacks the wholeness-preserving

quality of the final plan. The layout process for the final design has about fifty wholeness-extending steps (described in detail in [3]).



Figures 12a and 12b. Two possible plans for the apartment building. The lower one is the one we built. [(Upper drawing by Hajo Neis when a student. Lower drawing by Christopher Alexander.)]



Figure 12c. The completed Emoto apartment building in Tokyo [Photograph by Hajo Neis]

The plan as built complements and intensifies the wholeness of the site. This is clear from the ground plan alone. However, to emphasize what this really means, and to illustrate the deeper structures that are involved, I describe some specific features of the global structure that existed in that place, and that were solidified and complemented by the building placed there. Some of the structural features of the original site are: the center formed by the Y-configuration of the fork; the curving nature of the two branches; the space between the forks, and its specific V-shape; the southern spot hit by the sun.

The new building, when injected into this site, does the following things that help to accentuate and strengthen these structural features of the place:

1. The sharp end is given a snub nose, creating a place in front of it (it became the entry to a shop).
2. The front gardens of ground floor apartments are placed to form a boundary between the building and the street.
3. Since the boundary is of uniform width, the space of the street is maintained and continues its positive shape.
4. The building is wrapped round the sunny south-facing spot on the site, making sunny space for all the apartment owners.
5. On upper floors an inner layer of galleries and terraces form a boundary to this place.

As these various features are consolidated, they form a coherent structure, in which each feature supports and helps the others. Marked among this coherence is the presence of positive space *throughout*, even in a geometrically complex configuration; the fact that there are levels of scale within the structure; the focus of major centers to the structure, each with its own strength and beauty.

The entire process of building and injecting this structure into the configuration of the site – designing and planning and building the entire five story-complex – I view *altogether* as a complex computation. It takes the initial configuration (the site) as its starting point, then transforms the configuration, by about 50 steps, finally ending with the configuration illustrated. Although the work of an architect of this kind would conventionally be viewed as an artistic act, hence as an arbitrary and “private” process, I believe this to be seriously in error. In the process illustrated here, each step is based on the configuration that is there at the moment before, and the latent structures that exist, and then makes one step forward that specifically, and carefully, strengthens or supports, *and evolves naturally out of*, the structure that was there before.

5.5. Example: An ornament drawn by Hiro Nakano – 8 steps

Figure 13 shows a much simpler example, the evolution of a sketch, that follows the same process. This example is a result of a classroom exercise, where the students were given the initial row of dots.




Figure 13. Eight wholeness-extending transformations creating an ornament [from an original drawing by Hiro Nakano, now lost, kindly recreated from a copy, by Hajo Neis. CES archives.]

- Start with a row of even spaced dots.
- Make a circle around alternate dots. (ALTERNATING REPETITION)
- Connect the alternating dots and circles with a line (NOT SEPARATENESS)
- Intensify these new centers with an additional circle, and also make a much smaller circle around every other dot. (STRONG CENTERS, ALTERNATING REPETITION)
- Take the bigger circles and subdivide each into eight sectors, four plain, and four that are fan-shaped alternating with the plain ones. Give each fan-shaped sector a special shape to emphasize it and give it more detail. (GOOD SHAPE, LEVELS OF SCALE, ALTERNATING REPETITION, THICK BOUNDARIES).
- Add a strong dark, shaded band, to connect the larger circles with the smaller ones they alternate with more strongly. (CONTRAST, NOT SEPARATENESS, LOCAL SYMMETRIES).
- Add two additional lines, making borders to frame the dark band, and thus creating a still stronger sense of unity. (THICK BOUNDARIES, NOT SEPARATENESS).
- Darken the central band much more, and blacken the flower shaped quadrants of the inner large circles. (CONTRAST).

Through this sequence of seven transformations, we see the growth of a coherent and sophisticated structure, from simple moves. The names given at the end of each step are references to the 15 properties and wholeness-extending transformations, which play a fundamental role in the creation of this drawing.

The example is instructive because, although simple, it is complex enough to be interesting. We can see, for the first time among these examples, how recursive use of these transformations on the emerging centers, and on the centers that the transformations themselves generate, is potentially a powerful process, perhaps one capable of generating all complex configurations in their totality.

However, there is still much innovative work needed to take this informally described harmony-seeking computation process into an effective and operational form to cope with the broad range of problems we may wish to apply it to.

5.6. Example: Historical evolution of St Mark's Square – 10 cycles


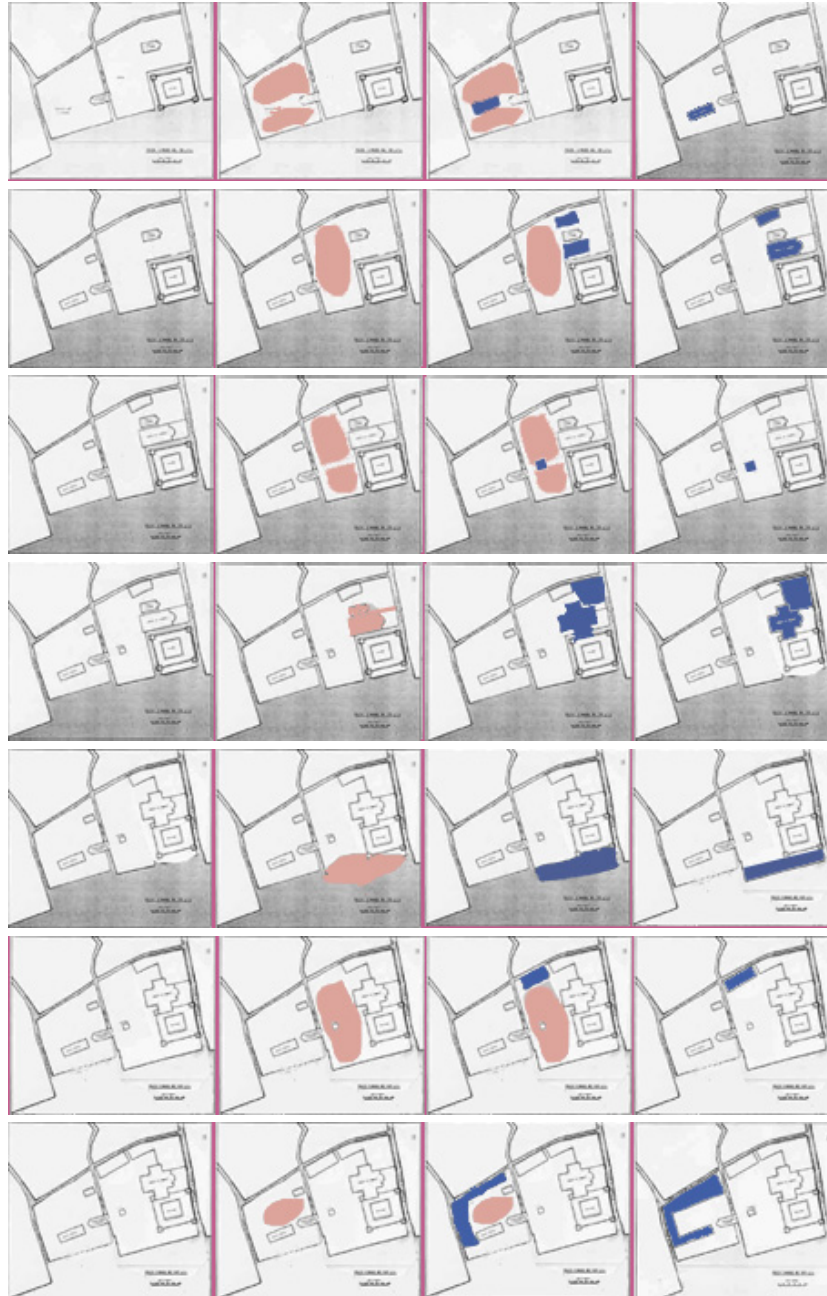
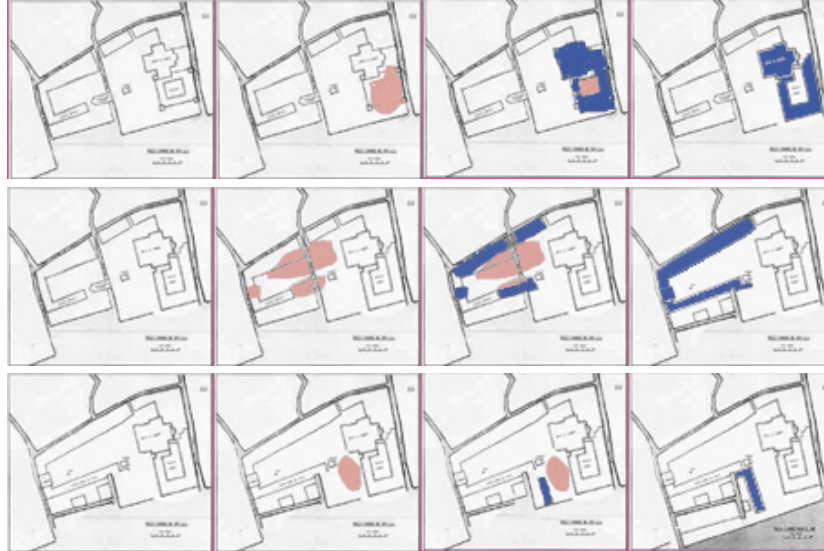


Figure 14. Plan of St Mark's Square. (From [19], p.104)

Figure 14 shows the plan of St. Mark's Square, Venice, one of the most famous and beautiful public places in the world. This highly complex structure grew steadily, over a period of about a thousand years. What happened there may be understood as a series of harmony-seeking computations, carried out at intervals, with intense and deliberate care. The series of growth cycles shown in figure 15 is historically accurate, although there is no historical evidence to suggest the nature of the activity undertaken at each stage. However, I am fairly certain that it must have gone roughly as I suggest: the fact that it is possible to define a coherent and rather simple paradigm for each cycle of activity, and that in their essentials the cycles are all the same, suggests that what happened was indeed something along the lines of the underlying computation I propose. (Historical data and additional analysis of the evolution of St. Mark's Square are provided in [2] [3]).





Figures 15. Ten cycles (four frames per cycle) of wholeness-extending transformations in the history of the 1000-year development of St Mark's Square

In every cycle, there are three entities at work:

1. *The configuration in the large*, roughly the size of the whole of the frame. The focus of attention shifts from cycle to cycle, and is usually about half the area of the entire frame.
2. *The particular latent center* that most calls for elaboration and development. The focus of the latent center, shown in pink, is usually slightly smaller than the latent center itself.
3. *One or more smaller centers*, shown in blue. These blue centers are placed in such a way as to strengthen both the previously weak latent center, and the large configuration that is helped as a structure by the elaboration and articulation of the pink.

As a computation, the procedure goes like this. Find the latent center that is most salient: the center that seems most likely, if strengthened in the next step, to strengthen the wholeness of the larger configuration. Act locally, so that this latent center is strengthened, and globally, so that this strengthening also helps to strengthen the largest whole. Repeat this cycle ten times over a period of about 1000 years, (roughly once per century), and the result is St Mark's Square as we know it today. This rule explains (or generates) the ten

actual cycles of construction and improvement that occurred around St. Mark's from 600AD to 1600AD.

5.7. Example: formation of giant voids in the universe: a very large example of a generated wholeness

In recent years, attention of cosmologists has focused on the structures of galaxies and galactic clusters that are now called filaments. Attention focuses on the voids that seem to be encircled by these filaments. The filaments are typically hundreds of millions of light years long, and tens of million of light years thick. The voids are often hundreds of millions of light years in diameter.




Figure 16. Intergalactic voids and filaments (from [20])

Figure 16 is one of the early pictures of the distribution of galaxies and galaxy clusters that for the first time clearly showed that these voids and filaments existed [20]. The voids appear to be convex, and quasi-spherical. To me, that implies the presence of some active principle pushing outwards. In terms of the language of wholeness, the existence of THE VOID (a general structure existing at all scales, and not specific to astrophysics and cosmology), and of THICK BOUNDARIES, and of the scale ratio given by LEVELS OF SCALE, is to be expected in all wholeness, since these are three of the 15 main ways in which coherent structure typically occurs. Further, the astrophysical voids seem to be close packed, rather like soap bubbles, and,

like soap bubbles, where the bubbles meet, there is a film where the boundary material is concentrated. In just this way, the galaxies are spread out in linear filaments.




Figure 17. Simulation of the 3D structure, showing roughly how the ring-like filaments lie in space, on the surface of roughly spherical voids (from [21])

Filaments are typically 200 to 500 million light-years in length. The so called Great Wall is roughly 1500 million light-years long, 600 million light-years across (band) and 15 million light-years thick. The voids are 30 to 500 million light-years in diameter. The actual physical size of one of these voids, (typically about 100 million light-years in diameter), is almost unimaginably huge. A jet plane flying at 600 mph would take something like 80,000 times the age of our universe to cross this void. I say this to emphasize the huge size of the voids that we are talking about, and in particular to draw attention to the fact that if the void is that huge, then the ratio of ring thickness to ring diameter could (hypothetically) have an enormous range of possible values.




Figure 18. Screen shot of a simulation by George Lake and Tom Quinn (from [22])

Yet that is not what we observe. One of the structure's most unexpected features is that the ratio of filament thickness to void diameter is about 1 to 10, roughly one order of magnitude. (Figure 18 shows this structure more clearly.) This is interesting because the filaments could have been much thinner or much fatter, compared with the diameter of the voids. There is no *a priori* reason for the ratio to be in the range of 1:10 (one order of magnitude). It could just as easily be 2, or 3, or 6 orders of magnitude. Why is the ratio of filament thickness to void diameter on the order of one order of magnitude?

5.8. Harmony-seeking, but not teleology

Stated differently, what reason could there be that a circular filament lying on the surface of one of these giant voids should have roughly the same geometric proportion as a wedding ring or an elastic band? In the absence of any current theory to give us a predictive model of the relative size (unless one fits the parameters just to get the result we observe), there is one unusual but possible theory. It is worth speculating that the universe itself – its action – is (among other things) making harmony-seeking computations, and that the results of these computations lead to a system of voids and filaments and define the proportion of ring diameter to filament

thickness. This could explain why the proportion of thickness to diameter so closely parallels our own cognitive and intuitive sense about the natural proportion of an intuitively sensible and natural ring.

What exactly do I mean by this? I mean to propose that there is a certain class of mathematical structures, in space itself, that has features that, for spatial and computational reasons alone, gradually appear in space, not because of something that is caused by forces, but *because of the geometry of space itself*. In short, certain kinds of ring thickness to ring diameter are more likely to occur simply because of their geometry. For some essentially geometric reason, the universe chooses to go in this direction.

If this hypothesis holds up, it would provide us with a process that looks *like* teleology, *without being teleology at all*. It would indicate that there is, inherent in space itself, a motion caused by the whole that can be understood in mathematical terms that are quite simple. Even very small local modifications of space cause a reshuffling, or reorganization, of the entire system of larger local symmetries, sometimes extending far beyond the point where the disturbance occurred. A new global configuration introduces new latent centers, sometimes far from the original disturbance. Harmony-seeking computations then strengthen the latent centers that were at first only dimly present, and we see how space itself, by virtue of its geometric organization, can move towards new configurations.

This process is *like* teleology, because it provides the system with a constantly moving goal that forms new structures all the time. It does not come merely from the dynamic trajectories of the past. At the same time, it is not the anthropomorphic teleology that scientists very properly reject. Rather the system's next move comes from the current state of system, *including its latent structures*, not from some future or external hand. It is an internally-fueled trajectory, caused by the phenomenon of the geometrical whole.

Try the following experiment when brushing your teeth: spit the frothy toothpaste into a bowl of gently swirling water. The toothpaste then moves in ring-like swirls on the surface of the water, and once again the roughly circular swirling ring-shapes that are formed have a diameter to thickness ratio in the range of about one order of magnitude. I do not believe the ring ratio merely fits a cognitive niche. Rather, once again, I suspect that there are, hidden in the geometry of space itself, reasons why ring-like structures with this kind of ratio occur. In current jargon, rings of this particular ratio might be viewed as attractors in some phase space. However, the discovery of geometric attractors in the solutions to systems of dynamic equations is,

in my view, only one particular manifestation of the far more general harmony-seeking computations that occur naturally in three-dimensional space. The tendency for such structures to occur might be shared over a wide range of circumstances, and the voids and ring-filaments in the cosmos might form in thickness to diameter proportions similar to those found in chains of foam in the bathtub, not for dynamic reasons having to do with physics, but for mathematical reasons having to do with the structure of space itself.

If we accept that the harmony-seeking computation is based on transformations that are (for whatever reason) largely congruent with human cognition and mentality, then we can consider such a process as explanatory even in the most outlandish structures.

This supposes that there may be mathematical reasons for the generation of circles or ring-like filaments whose thickness is just about one order of magnitude smaller than their diameter. There is, at the moment, no physical theory that would predict this, but, from the point of view of harmony-seeking computation, the cognitive simplicity of this ratio for a ring-structure could suggest that it might have to be so, for mathematical or quasi-mathematical reasons.

Is this strange? Yes indeed, to our ears it is very strange. Yet it is not more strange than Pauli's exclusion principle⁵ ? It seemed strange when first put forward, but is now widely regarded as one of the cornerstones of quantum theory. A purely geometric harmony-seeking computation might explain at least one aspect of the structure better than any current physical theory. In the absence of better computations, some reason needs to be found why such processes, congruent with cognition, might occur in structures at the scale of the universe and at the scale of human beings, and why the ring-ratios – in both cases – tend to be roughly similar.

⁵ Pauli's exclusion principle (see, for example [23]) states that no two fermions of the same kind (for example, no two electrons, or two protons, or two neutrons) can be in the same quantum state at the same time. What this amounts to in atoms is that the electrons can form only certain mathematically symmetrical patterns, and results in the chemical properties of atoms. This is a quasi-geometric principle, consistent with the antisymmetric nature of a fermion's wavefunction, that has no known "force-based" explanation.

6. WHOLENESS-EXTENDING TRANSFORMATIONS

6.1. Viewing the previous examples as computations

The previous examples all suggest the possibility of some kind of harmony-seeking computations at work in a variety of systems. In all cases, there is some kind of wholeness-extending process at work, mainly geometric in nature, sometimes intellectually graspable, sometimes not so clear. Apparently, as a result of this wholeness-extending process, at each step the generated configuration moves a given configuration towards a more harmonious, more coherent, state than it was in before.

We may regard *all* these examples – physical processes, simulations, and paper and pencil exercises – as *computations*, each of which takes its system forward, step by step, from an existing state W_1 to $W_2, W_3, \dots W_i, W_{i+1}$, increasing structure and harmony and wholeness in the system as it goes. These steps are carefully calculated, commonplace in many spheres of human experience and in many natural systems. Some of them are done by animal instinct; some by human instinct; some by adherence to a certain tribal sensibility; some by concern for ecology. Many occur naturally (without human intervention) as part of the behavior of some complex physical system.

I use the word *computation* with a variety of related meanings:

- *Intuitive computations*, made by craftsmen: hayricks; the bench around the tree; the construction of St Mark's Square.
- *Conscious computations*, made by a designer, according to a consciously used scheme that calls for certain actions to be performed in sequence: the Emoto apartment building in Tokyo; the ornamental drawing by Nakano.
- *Natural computations*, possessing explanatory power for straightforward scientific problems that have not yet been solved: the development of the whorl tips of *Acetabularia*; geese flocking to form a V-shape (later).
- *As yet half-formulated harmony-seeking computations*, which may be implicit in the behavior of complex systems, and would shed important light on the nature of the phenomena: the formation of giant voids in the unfolding of the cosmos; the formation of blue sky and clouds with their complementary positive space (later); the growth of tree limbs to create positive space between the limbs (later).

In these cases we might also talk about the system behavior, and the rules according to which system behavior develops. It is perhaps odd to call them *computations*: but if we wish to simulate the system behavior

that leads to harmonious results, we may need to use such a computation as the explanation of the physical system's behavior.

In a generic sense they are all computations. But in nearly all these cases they are not recognizable as computations in the way we conventionally understand calculations or algorithms. There are reasons to believe that a fundamental kind of process is at work in all these cases. We need to examine these computations, and understand or extract the underlying way they work – all of them together. If we succeed in getting the gist of this type of computation, we may find a powerful new way of computing that is guided by emerging harmony, and by a motion towards harmony.

I view these examples as computations because they all embody some kind of single process that has the capacity, starting with an arbitrary configuration, to reach a better and more harmonious configuration. Although the idea of a wholeness-extending transformation is, for the moment, still somewhat loosely defined, there is considerable similarity between the examples – in the way that they work and in what they accomplish.

To learn to see these examples as embodiments of a single type of computation, we must be specific about what this deep seated type of computation might be, how it might work, and how we can put it to work to give better organization, better harmony, better deep adaptation and better coherence for some real-world configuration.

6.2. Experimental Confirmation

My own experiments in cognition show strong agreement among observers as to comparative judgments regarding which steps are wholeness-extending and which steps are not (see [2] pp 51-84, especially pp 59-61). A step is a harmony-seeking step if it extends and enhances the wholeness structure. When asked, under experimental conditions, to compare different possible steps, starting from a given configuration, in terms of the degree to which the steps are wholeness-extending or not, people largely agree, based on their intuitions, or on their own cognition. This strongly suggests that the quality of “being wholeness-extending” is to some degree objective, that the phenomenon reflects an objective underlying physical reality.

If we ask people taking part in such an experiment [2] to make (in their own chosen terms) any kind of structural diagram in which they try to capture the wholeness of the thing they are looking at, the agreement among observers improves. When people make such diagrams before judging which

steps from it are wholeness-extending and which are not, the level of agreement between observers goes up strongly. Even though their diagrams are not alike, the people who have performed such a private diagram-making process then agree more strongly with one another on what steps from that structure are wholeness-extending. This occurs even though the making of a diagram that describes the structure is a private matter, and people are given no special instructions about how to do it, and the different diagrams people make are dissimilar. After simply *attempting* to represent the structure, they can judge more reliably what is wholeness-enhancing and what is not. This, once more, is evidence that the phenomenon of wholeness-extending is real and objective as a phenomenon, with objective structural content.

6.3. The wholeness-extending postulate: always helping a larger whole to form

How may we formulate, mathematically, the character of the wholeness-extending steps that occur in the examples?

In each case there is a whole W , and within the whole a latent center L that is being modified, transformed, or reshaped, by a certain step. This latent center is the focus of the transformation, and the latent center sets the boundary of the geometrical and physical transformations that are then undertaken.

The larger whole W is often an order of magnitude bigger than L . The wholeness-extending transformation improves the structure of W , and to do so modifies L relation to its whole context. Thus the output from this step is a modification geometrically within L , but it is a function of both L and W . In addition, there is a sense in which L is being *fitted* to W , it is being made more congruent with W , adapted to W , harmonious with W . Further, in modifying L , new centers N_i are created *within* and around L .

To undertake this transformation in such a way that it is indeed a wholeness-extending transformation, the various N_i are generated by 15 generic types of transformation acting on L . These are *the* 15 principal center-creating and wholeness-extending transformations. The N_i can then become the bases for further transformations.

6.4. The wholeness-extending transformations of St Mark's Square

Now look again at figure 15. In each cycle, the pink area is a latent center L . The next building is built inside L . The larger context of the whole St. Mark's area is W . However, there is now a subtlety. In a particular step, we know what L is, because we are looking back in time and see what the step was. But the people who did the step did not have this perspective.

They could not know what L was, until examination of the context and the larger whole revealed it to them. (And there is not always a unique answer. Sometimes L is obvious, but sometimes there are several equally good choices. Other choices will generate other St Mark's Squares. Nevertheless, the total number of harmonious results from this process, even taking into account the variety of possible choices, is a tiny fraction of all possible moves, most of which are *less* harmonious.) A further subtlety relates to the context W , which is not so general as the whole St. Mark's area. It is, rather, a particular area within St. Mark's square, where the latent center L has been identified as in need of improvement, or presenting itself for elaboration and strengthening. This latent center L , which plays a crucial role in the wholeness-extending transformations, *and* the detailed effects, are created by the 15 transformations acting together. It is the immediate local context of L that then gives rise to the step that transforms L .

Alternatively we might say that each latent center L , once identified, is the entity to be transformed, and the transformation is to be done in such a way as to strengthen L in its ability to help and make cohesive the larger context W . The N_i of the transformation are the various buildings and partial buildings generated by the action of the transformations.



Figure 19 A second look at one of the wholeness-extending cycles in St Mark's Square [repeated from fig 15]

In figure 19, the pink ellipse is L , a latent center formed by the three buildings around it. To enhance and strengthen L , the blue building mass is built, thus forming a stronger rectangular space by enclosure, and by establishing continuity with the buildings on the right of the latent configuration. In the process two of the original buildings are absorbed into the new one, which continues and extends the wholeness.

7. HARMONY-SEEKING RATHER THAN “EMERGENCE”

In recent years, complexity theory has begun to see emergence of complex structure as a result of coupled local atomic events generating larger wholes through interaction. Such bottom-up processes are said to demonstrate “emergence”. It is used to explain how birds fly in flocks [24], how ants accomplish complex tasks together [25], or how slime mold (apparently an uncoordinated aggregate of cells) is able to move coherently as a whole [26]. In most cases, the explanation is, in some form or other, that the individual agents mimic the behavior or action of their nearest neighbors, or base their next state on the state of their neighbors. In many cases this ultra-simple rule of action, plus some refinements, explains considerably coherent movements of the larger body of cells [27]. In other more complicated cases, the atomic actions of individual cells are coupled, and the coupling helps more complex forms of behavior to occur in the group [28]. In all these cases, the resultant aggregate seems to be acting as a whole.

However, I do not think that this bottom-up emergence, caused only by local interactions, can adequately account for the world’s ability to generate highly-ordered, complex phenomena. When wholeness, or harmony, exists in a part of the world, what is under discussion is always the relationship of a given system to the larger world beyond that system. So the issue is not merely whether a group of elements act together. What is important is that when the elements are grouped together to form a system, the resulting system then either does or does not act in such a way as to sustain, or improve, the coherence and health of its own component systems, and also of some still larger system around it: some part of the world *both inside and outside and beyond* the group.

The phenomenon of bottom-up emergence is a *two-fold* relationship between a set of elements and the group they form. The *harmony* phenomenon is a *three-fold* relationship between a set of elements, the group they form, and the helpfulness of the emerging group, through adaptation, to an even larger portion of the world beyond the group.




Figure 20a. Emergence, a two level relationship

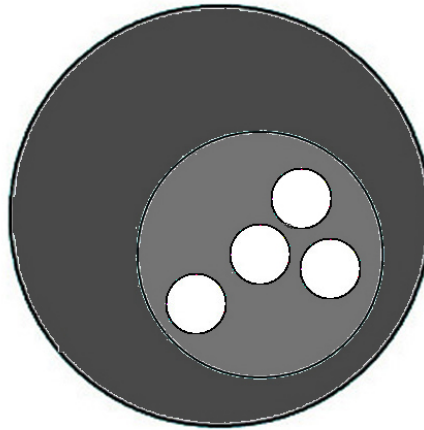


Figure 20b. Harmony, a three level relationship

In figure 20a we have a group of small elements (white circles), and because of their interactive coupling they act as a whole, represented by the middle-sized grey circle. In figure 20b we also have a group of coupled elements that act as a whole because of the coupling, but in addition, *this middle-sized whole acts to support or help the coherence of a still larger system*, represented by the outermost black circle. The first system is interested only in itself: it does nothing to help the world beyond it, and does not contribute harmony to the world beyond it. The second system is not only embedded in the outermost ring: it *helps* that ring, and establishes harmony between the middle ring of rings and the largest ring, so that the middle ring is well adapted to the largest, and the largest benefits from the existence of the middle ring.

I now compare the doctrine of emergence from local, bottom-up interactions, with the richer notion that in many cases there is a *whole-based, harmony-seeking process that originates in wholeness*, works by continually strengthening latent centers that appear in that wholeness, and thereby operates on the larger wholeness that exists – all in such a way as to generate a new and deeper wholeness.

7.1. What does it mean for a system to *help* a larger system?

In the example of St Mark's Square, in each cycle the process (through human agency) identifies a latent center in the larger configuration. This latent center, if strengthened, would improve the coherence of the whole.

The larger area immediately around that latent center is then enhanced and made more coherent by the transformation of the latent center.

Abstractly we may express this concept by looking again figure 20b. The middle, grey system is a latent center, a weak system that has the capacity to enhance the coherence of the black system beyond it: and it is able to do so by creation and modification of the small white systems – they are built and shaped to create more coherence in the grey system. However, the key point is that the enhancement will not take place, unless the grey system then *also* helps to enhance the larger system beyond it, the black one.

This three-level structure is a matter of common experience both in nature and in everyday affairs. In a family, for example, the individual members of the family do what they can to help the family be coherent, and one of the most effective ways this can be made to happen is through a situation where the family as a whole makes a positive contribution to the village. The individuals become stronger, and the family becomes stronger, as a result of the help they give to the village which contains them. We observe a similar phenomenon in the edge ecology of a river. It may be said that if the two lower levels do not reach out to help the coherence of a larger system, their own coherence is likely to falter.

In general, the way that smaller systems are modified and arranged to help the coherence of larger systems, in the fashion described above, is too easily lost in too much focus on bottom-up “emergence”. Emergence, in the form that is currently popular, is thought of as a two-level phenomenon: the emergent property and the level below. The appearance of “helping” in the world, depending as it does on transformations which increase coherence in the large, is crucial, and embraces a higher level: the wider environmental aspects of a larger system.

7.2. Example: Flying Geese and the V-formation

Boid simulations, due to Reynolds [29] and others, provide models in which the particles fly about, avoid each other, all go in the same direction, flock, play: altogether, simulate some of the behavior of birds in a fairly realistic fashion. When used to model bird flock behavior, there is a surprising realism in some of these models. However, simulations of this type do not do as well when it comes to generating more complex formations, such as the characteristic V-formations of migrating Canada geese (figure 21).



Figure 21. The V-formation of Canada geese. It is most interesting to look at the spaces between adjacent birds. In each case there is a rectangle of empty space formed by the left wing of one bird, the right wing of the other bird, and the two bodies. You may see these rectangles clearly in this photograph and in Figure 23b. They play a crucial role in generating the overall coherence of the V-formation.

When flying in V-formation, Canada geese can increase their flying range by up to 70%, because they use less energy, as a result of the vortex interactions [30][31][32]. When aircraft arrange themselves in similar V-formations, similar energy savings are experienced [33]. But that does not explain *how* the V-shape comes about in the case of migrating geese.

Let us consider what real geese may be doing to achieve the V-formation. Reynolds' three basic boids rules (as most frequently given) are these [34]:

1. Collision Avoidance: If you are about to crash into another bird, change direction or reverse direction.
2. Velocity Matching: Attempt to match velocity with nearby flockmates
3. Flock Centering: Fly in a direction that moves you closer to the centroid of the nearby birds.

This system of three rules generates *general* flocking configurations, to an extent that seems amazing the first time you see it. But these three rules, alone, certainly do not create a persistent, stable, *V-formation*. To do that we need to add two further rules.

4. Try to fly along the sweet spot, or line of maximum uplift, or the vortex cylinder of the bird in front (which the geese can do because they can sense the change in pressure and the easier ride in that position).




Figure 22a. In a bird, as in an airplane, there are two vortex cylinders behind the wingtips, each rolling towards the center. The zones of maximum uplift are visible along the outer edges of the drawing, along a line some distance outside the vortex axis. The same maxima of uplift are also shown in Figure 22c, as dotted lines. (From [35]). If you look at the cross section in Figure 22a, you see that the line of maximum uplift is parallel to the vortex cylinder, but outwards from it, and it is this line that gives the second bird it maximum uplift.



Figure 22b. The cylinders of vortices seen as two outer parallel trails emanating from the wing tips. (The inner trails come from the engines.) (From [36])




Figure 22c. Here each bird flies in the outer wake of the bird ahead of it, since this is the spot that takes the least energy – just like riding a bicycle behind a bus, to catch a ride from the uplift in the airflow.

The optimum uplift positions are in the outer wake of the bird ahead, off-center from the bird's axis [37]. A sophisticated algorithm would have to model the aerodynamics to identify the position of this sweet spot. A simpler (and more easily constructed) algorithm would merely need to identify the spot outside the wake (figure 22c), and give positioning in that area a high probability.

5. Finally, if you are the leader at the tip of the V, you have the smallest energy advantage in the flock. So, after a while, drop away, and let another bird take your place. This causes the V to regenerate continuously, and lets each bird carry an equal share of work.

When we use all five of these rules, I believe the dynamics will now generate a stable V that regenerates itself continuously.

We may easily see how this structure of events is a harmony-seeking computation, not merely an example of bottom-up emergence, since it has the three-level structure postulated earlier (figure 20b). The new rules of action require that the birds compute in a way that makes a local center (the bird's own body) work to create a second level element (the positive space shown in figure 23b). These second level elements are strung together to form the third level diagonal structures of which the V-formation is formed), thus demonstrating the action of the three-level structure.

The algorithm needs to be yet more subtle. Vortices form during flight, as shown in figure 23a. The shaded space between adjacent birds (figure 23b) is a positive shape, composed by one bird's left wing, the other bird's right wing, and the two bodies. Two birds can discover, through experience, that there is an optimum ride when they are in a position in which the left wingtip of one is approximately in line with the right wingtip of the one ahead of it, generating the rectangle in space. At that spacing the vortices cancel, minimising turbulence [38]. This relies on the positive and definite shape of the space between the birds, which strengthens the stability of the dynamic configurations.




Figure 23a: the vortices formed in a V formation (from [38])




Figure 23b: position of the virtual rectangle of positive space between two birds.

Now we see how the three levels work together. Elements at level 1 (the birds) form chunks of positive space between adjacent birds. These rectangular chunks of positive space are the elements of level 2. The birds experience the best functioning when they are forming these chunks of positive space between them. The chunks of positive space then join up to form a coherent diagonal (elements of level 3), and these are the diagonals we see as the V-formation. This can only be generated by using a computation that looks at the whole, and the emerging structure of the V-configuration as a whole.

This situation cannot in principle be modelled properly without a harmony-seeking transformation based on the structure of the whole V as a guiding field, because the computation needed to generate the V is a computation that explicitly relates the individual to the whole. See figures 22 and 23.

The current enthusiasm for “emergence” as evidenced in bottom-up boid-like computer games and simulations does little to unravel the real and more subtle problems of the whole and its organization, which hinge (almost always) on the way that small parts *work to help the coherence of a larger whole*, and the ways that the larger whole then also shapes and modifies the action of the smaller parts.

The real issue is that the emergence that is being attributed to the birds (when they are viewed as mechanisms) is not as dumb or mechanistic as some emergence enthusiasts like to think. To produce the V-formation, the birds themselves have to perform a harmony-seeking computation, in the way that they act to relate themselves to the larger whole, and to help that larger whole. This computation might be performed geometrically, by using the diagram illustrated above which offsets the bird behind by a certain ratio

of the wake angle. Or, equally well, the computation might be performed physically by relating the bird's position to the sensed pressure. These two computations, though different, are nearly equivalent. In both cases, the fact that matters is that the birds themselves perform this harmony-seeking computation by relating themselves to the larger whole in a certain way.

7.3. Example: clouds and the positive space that arises between them

In figures 24b, 25, and 26 we see various examples of cloudy skies. Look at the shape of the blue sky *between* the clouds. Notice (with something of the eye of a painter) how the blue patches are well-shaped, different from the childlike cotton wool cloud formations we perhaps carry in our minds as a picture of the way clouds are, where the clouds are objects and the sky is background (figure 24a).



Figure 24a. Imaginary structure: cotton wool clouds. The clouds are thought of as objects and the sky as background.

Figure 24b. The real thing: a very different structure. Both sky and clouds are positive in shape, co-existent and interlocking.



Figure 25. Clouds and the sky between them, as they appear in nature



Figure 26. The blue systems are clearly visible here as having their own shape and character

The blue sky is made of definite and positive shapes (POSITIVE SPACE). For a painter this notion is straightforward: a picture cannot be good unless *all* the spaces and components and fragments, have their own positive shape. But this positive space is tricky to formulate in mathematical terms. The idea of POSITIVE SPACE is something like convexity. In mathematics, a convex body is one that has the property that, for any straight line that connects two points inside the body, all the points in between, along that

line, also lie inside the body. Positive space is coherent in shape, is formed of positive, somewhat convex lumps that have definite and recognizable shape, but is less tightly constrained than mathematical convexity – hence *quasi-convex*.

There is no *a priori* reason to expect POSITIVE SPACE (seen in projection as blue sky) to appear between the clouds. Yet, clouds blowing and changing minute by minute maintain the positive space of the blue, as well as the positive space of the clouds themselves, at each instant. Mathematically, this may be expressed by a simple idea: the blue sky is not a mere background for the white clouds; the space that we see as empty (or blue in projection) is itself a three-dimensional system, a dynamic object, that has its own shape and behavior.

So how may we explain what nature is doing in these cases? Somehow the positiveness of space appears between the clouds. The system seems to have a disposition to have this positive space appear. What is causing it, and how does it work?




Figure 27a. Configuration arising from the Swift Hohenberg equation




Figure 27b. Configuration arising from the Ginsburg Landau equation

The easiest way to explain it is to consider both the white bits, *and* the blue bits, as quasi-convex bodies in 3-D space. It is a co-evolving system in which this loose packing of differently sized blue and white quasi-convex bodies is maintaining itself. Such a changing dynamic packing of irregular cells could be a complex 3D analog of Taylor vortices or Benard convection. Figures 27a and 27b show two simulations from the laboratory of Professor Michael Cross at Caltech, demonstrating stages of Benard-like formations in

a medium, one under the impetus of the Swift-Hohenberg equation, and one under the impetus of the complex Ginsburg-Landau equation [39]. They are more regular than clouds, but one can see how the positive space between clouds could come about as a result of more complex interactions stemming from these kinds of effects, and possibly driven, in addition, by some iterative rule similar to the POSITIVE SPACE transformation.

Although this is not (in this form) part of the accepted physics of cloud systems, it is what one would expect from a harmony-seeking computation, which has, as one of its wholeness-extending transformations, the continuous maintenance of POSITIVE SPACE at every step. This reflects the fact that the space we see as blue sky between clouds, has its own rules of formation that are active “empty-space” systems maintaining their own coherence dynamically as rolls or vortices, and that shape the space of the cloud systems as much as being shaped by them.

7.4. Example: the branching shapes of real trees and the positive space between the branches

The nature of POSITIVE SPACE as an active principle, which plays a major role in harmony-seeking computations, may be understood further by looking at the growth of trees. If we look at a tree we see volumes of positive space between the branches. It is visible, for example, in figure 28. So, apparently, the empty space between the branches must, like the blue sky, have a driving force, for it to be a definite “thing”. How exactly might this be explained? The space between the branches has no obvious energy to push and create its positiveness.

Prusinkiewicz has made great progress by simulating tree growth through the use of L-systems, and has recently developed “environmentally-sensitive” simulations, which create more realistic structure that develops, in part, according to neighbour conditions [40] (figure 29). But the idea of empty space as positive does not yet have an accepted mathematical formulation in L-system theory, which focuses on the growth of branches, twigs and leaves. The L-system simulations that I have seen lack the positive character of the space between the branches in figure 28.



Figure 28. A sketch of mine showing a real tree as it typically is in winter, and strongly showing the positive space generated between the tree's branches




Figure 29. An environmental L-system simulation of tree growth in two adjacent, interacting trees (from [40]). The spaces between branches still lack the intensity as centers caused by the action of each volume of space as a system in its own right.

One can readily account for the positive space phenomenon as it occurs in trees. The twigs and leaves in a growing tree are phototropic: they look for the light. (This *is* simulated in environmental L-systems.) The tree also needs air and breeze, and the leaves and twigs and branches modify their growth to allow the passage of air, and to avoid crowding of leaves that deny them light. Thus, the “empty” volumes of space that nestle among the leaves are highly sophisticated systems, as necessary to the tree as the complementary organization of limbs and branches that provides structural strength and the flow of sap. As in all systems of positive space, the one system complements the other. Both are necessary to the tree’s living character. In this case, observation of the harmony-seeking calculations performed by the action of the tree teaches us more about the tree as a system, and could allow us, if we chose to do it, to write more convincing algorithms of a tree’s morphogenesis.

Environmental L-Systems currently model the tree’s environment in a manner very different from the way they model the growing tree itself, typically using a conventional simulation rather than a growth grammar. I wonder if an extended new type of L-system might be constructed that allows two parallel and complementary growth grammars, (1) one for the branch and leaf system and (2) another for the system of space volumes between the branches, the two being coupled to evolve and unfold in parallel, and in a way that makes each the inverse or geometric complement of the other? I believe such a model is likely to provide a more accurate picture of the real structure that is generated in plants and trees.

8. WHOLENESS-EXTENDING TRANSFORMATIONS AND SYMMETRY BREAKING

8.1. The possibility that wholeness-extending transformations are deep generalizations of symmetry-breaking

As part of an analysis of the evolution of natural phenomena and configurations, Stewart and Golubitsky were perhaps the first to lay emphasis on the phenomenon of symmetry-breaking [41]. The regular spacing of dewdrops on a spider's web is a case that they discuss in detail (figure 30). The water on a thread of the web starts out as a uniformly coated cylinder of water, made roughly uniform in thickness by surface tension. When the surface tension starts to break up the continuity of the coating, the configuration falls apart, but still repeats at least some of the symmetries present in the continuous cylinder, since there is no reason for all *those* symmetries to be removed as well.




Figure 30. Dewdrops on a spider's web [from Fotosearch..com]

This idea is similar to that of a wholeness-extending transformation. We have a structure: in this case, the infinite translational symmetries of the cylinder along the thread, and the rotational symmetries of the cylinder

around the thread. As this system moves to a fragmented version (caused by the action of the surface tension, or by jiggling of the thread), the simplest end product is a configuration that destroys the fewest possible number of symmetries. Golubitsky and Stewart's work on symmetries in the equation systems of bifurcation theory continue such ideas [5].

This one way of preserving structure is somewhat limited, if compared with all the possible ways of transforming a system while preserving and enhancing structure. Thus symmetry-breaking alone, as an interpretation of what is going on in the simple dewdrop example, is geometrically too limited. Each dewdrop also takes on local spherical symmetries and axial symmetries normal to the thread, and symmetries parallel to the thread but not aligned with it, that are not present in the symmetry scheme of the infinite cylinder.

In my view, the symmetry-breaking idea is not yet sufficiently developed to be useful as a general theory explaining complex configurations, or to account for the harmony-seeking phenomenon. "The" wholeness consists of the entire system of overlapping local symmetries at a wide variety of scales in a configuration [1, pp. 79-142 and 446-452]. We need to have a view where the underlying structure of all these symmetries, working together, is preserved. And further, the centers in a given wholeness are not all LOCAL SYMMETRIES. Other centers are formed by GRADIENTS, ECHOES, THICK BOUNDARIES, DEEP INTERLOCK, POSITIVE SPACE, NOT SEPARATENESS, and so on. These other properties, and the entire structure, too, have to be preserved when a harmony-seeking computation starts with a currently existing structure, and finds its way to a stronger structure that is latent in this overall configuration, and can be brought out by a few transformations.

In summary: a wholeness-extending transformation is a more complex and richer version of the phenomenon whose simplest cases have been called "symmetry breaking" or "symmetry reduction". This nomenclature is rather over-simplified, and does not do justice to the real potential complexity of the underlying phenomenon. A wholeness-extending transformation moves a complex configuration forward, retaining as much of its wholeness structure as possible, and drawing new structure from the latencies within the wholeness itself. In so doing the configuration becomes richer and more complex in unforeseeable ways that benefit the larger whole.

8.2. Local Symmetry Production

Even in the well-known cases of symmetry reduction, what we see, if we look closely, is symmetry *elaboration*. The infinite Euclidean 3D-space

does indeed lose some global symmetries. But at the same time, LOCAL SYMMETRIES are strengthened, or generated where none existed before. This is a better picture of how harmony-seeking computations work. The failure to see it, comes, I think, from the fact that in recent years the symmetry structure of the plane and of the 3D-continuum, have been viewed too simply through the symmetries *of the whole*. It has not been sufficiently clear that there are, hidden in the plane, or in the continuum, an infinite number of *smaller local sets*, all over the place, some of which are themselves symmetrical, while others are not. There is an infinity of systems of smaller and smaller symmetries that occurs in these nested systems of sets. (Many physical examples of the presence, and density, of local symmetries in a structure are described in [1]). What happens when a whole evolves under harmony-seeking computations, is that many of these smaller symmetries are strengthened or generated where the latencies are strongest, thus gradually generating a hierarchical nesting of local symmetries at different levels of scale.



Figure 31 Salingaros [42] studied 25 famous buildings by counting the number and density of local symmetries, according to a modified version of the measure proposed in *NOO*, adding further elaborations to estimate the density of symmetries. Among the buildings he studied, the Parthenon had a density of local symmetries (including both the small symmetries and the large ones) in the upper part of the range. [photograph from www.sapdesignguild.org]

We may see the result of such a process in many famous buildings, for example, the Parthenon (figure 31). The Parthenon is, of course,

symmetrical in the large. It is these large symmetries that most people see when they first look at it. But when we look for locally symmetric sets, large, medium, and small, we find an astonishing number of them: the columns, the spaces between columns, the flutes of the column, the capitals, the metopes, the triglyphs, abaci, stylobates, entablature, guttae, steps, etc, which are all locally symmetrical. The building is deceptively simple, yet the enormous density of locally symmetric sets within the configuration is among the highest of all those measured by Salingaros among a considerable range of famous buildings [42]. In general his findings are that the buildings with the most profound qualities are those for which the internal density of local symmetries is highest.




Figure 32a. Experimentally derived rank order of coherence of 35 different patterned strips

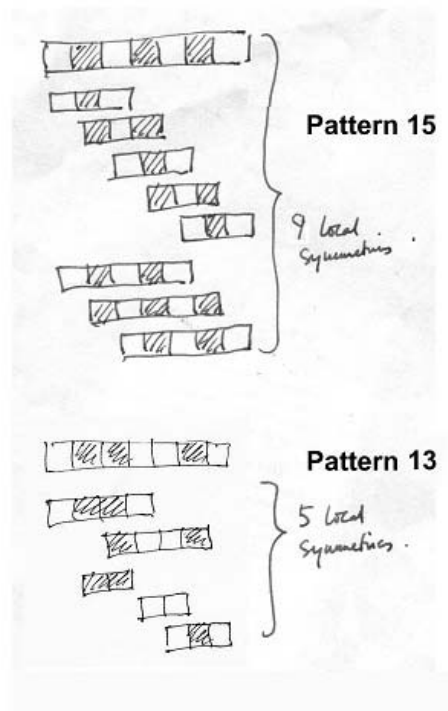


Figure 32b. Enumeration of local “subset” symmetries: nine local symmetries in pattern 15; five symmetries in pattern 13

High density of local symmetries, sometimes overlapping, is also visible in my cognitive experiments on black and white strips [14]. Figure 32a shows the experimentally derived rank order of coherence of 35 patterned strips, as measured by a variety of independent cognitive and perceptual tasks. Those at the top are (by a variety of experimental measures) found to be cognitively most coherent; those at the bottom least coherent. This rank order is predicted almost exactly by counting the number of local subsymmetries in each strip. The strips at the top (such as pattern 15) have nine local symmetries, those near the bottom have five, and the others lie in between (figure 32b).

The presence of these local subset symmetries, and the number of them, appear to cause what is seen as coherence. This structure is not a product of symmetry breaking, since the local symmetries of an asymmetrically placed symmetric subset is not a subgroup of the larger symmetry group. Rather, it needs to be described in a different way, as an example of multiple symmetry production. In organic development, local symmetries (limb-buds for example) are created, and this is one of the most important phenomena in organic development.

In order to understand this local symmetry as source of harmony, it is necessary to see the symmetries as part of a system of overlapping nested sets in space, each of which may take on local symmetries or not, within its own local frame. A harmony-seeking computation in this context is a type of computation that injects as many overlapping local symmetries as possible into a finite framework.

Salingaros [42] has shown that such compressed systems of local symmetries are present precisely in the acknowledged great buildings, especially of ancient society. He shows correlations between the local symmetry density measure, and the intuitive and accepted evaluations of many great historic buildings, especially when compared with more raucous modern examples.

8.3. Example: snow crystals

Another illustration of the deficiencies of algorithmic computations in morphological matters lies in attempts to simulate snow crystal formation using Diffusion Limited Aggregation (DLA) and cellular automata (CA) methods. The results of these simulations do not accomplish the extraordinary variety and beauty seen in real snow crystals. Bentley photographed some 5,000 snow crystals [43], and his precise and exquisite photographs show us the kinds of structures that a successful computation

must be able to create (see figures 33 and 34 for examples). In particular, we see LOCAL SYMMETRIES, DEEP INTERLOCK, POSITIVE SPACE, and GOOD SHAPE in these snow crystals, at levels that are common in works of art. Such features show up, only weakly, in the DLA or CA simulations.

The prevailing theory of snow crystal growth (described in [44], with simulations on the website [45]), is that the crystal grows outward starting from a small hexagonal plate. The six arms grow similarly, maintaining the six-fold symmetry, because the conditions of temperature and spatial constraint, although changing, are roughly the same on each arm at any given moment in the crystal's growth.



Figure 33. A snowflake with two three-fold symmetries (axis- and face-oriented) appearing within the six-fold scheme (from [43])

However, if you examine figure 33, you see that in addition to the six-fold symmetries, there are other symmetries. Most noticeable is an alternating pattern: three arms have one pattern, and the other three alternating with them have another. This may be seen in the moth-like figures that occur on the three arms that are 120 degrees apart.

Figure 34a, showing a simple plate crystal, has another aberration. Again, there is relatively uniform six-fold symmetry, centered on the six vertex axes as usual. However, if we look at the outer edges of the hexagon, along the three faces pointed to by the blue arrows, we see that these sides have a different symmetry: a shape that has two ogive S-curves (rather like a valance on the top of an 18th-century cupboard) just along the outer edge of

the hexagon. This new symmetry is three-fold, and is based on three axes that are edge-centered, not vertex-centered.




Figure 34a. Three-fold symmetries (axial and face-oriented) appearing within the six-fold scheme (from [43])

Figure 34b. A two-fold symmetry appearing within the six-fold scheme (from [43])

In figure 34b, there is a configuration where a figure like a pair of chromosomes appears within a conventional six sided plate, but shows up only twice, 180 degrees apart, thus introducing yet another symmetry.

Further information has been obtained by Ukichiro Nakaya, who observed and photographed thousands of snow crystals, and also experimented by growing them under controlled conditions [46].

The non-six-fold phenomena do not seem to be explained by the assumption [44] that the crystal's symmetrical growth is caused by uniform temperature variations on all six arms at the same time. Instead, there must be other mechanisms that generate a variety of alternating three-fold symmetry and two-fold symmetry. These various symmetries need to be the focus of our investigation, not only the overall growth. Some larger

wholeness-extending transformation is entering into the computation. The three axes of symmetry now coming forward are not a subsystem of the original six growth axes, but new ones that are interlaced among them. I believe (but have not yet shown how) the phenomenon can be explained by strongly invoking the LOCAL SYMMETRIES transformation, but not as a part of the simple kind of growth mechanism described in [44]. It seems that only a model based on a nested, multi-level system of symmetries and axes of symmetry can ultimately provide an adequate explanation. Such a model may be a step towards a capability of modeling the general character of highly ordered structure at many levels simultaneously.

9. WHOLENESS-EXTENDING TRANSFORMATIONS AND COMPLEXITY THEORY

9.1. A small modification needed in the work of the modern masters and Complexity Theory

What I have said in the forgoing, can be simply summarized. Consider the modern masters: Benoit Mandelbrot, Brian Goodwin, Ian Stewart, Przemyslaw Prusinkiewicz, Deborah Gordon, Craig Reynolds, *et al.* We must, and do, take off our hats to them. But if we examine the contributions made by complexity theory, we see that while making enormous strides, it has fallen slightly short in one all-important respect. It tries to explain complex emergence as a product of coupled interactions among local events. Yet, when one examines in detail what is actually going on, and what is emerging, it turns out that there is, in every case, some aspect of the emerging whole that cannot be properly explained by this approach. The aberrations are small, but not small enough to be overlooked. Although complexity theory seeks to explain the whole that emerges through analytical bottom-up means, it is the whole, the real whole in the world that makes us marvel in the first place, that is propelled by a second process, which is both whole-driven and whole-oriented.

I reiterate: this whole-seeking or harmony-seeking process is not teleological, not goal-seeking (see also §5.8). Instead, it comes about through a computation performed on the structure that exists, that brings to fruition a larger, unanticipated new structure of wholeness. The existence of such a computation, in virtually every creative process in nature and in art, perhaps embodies the real creativity of the universe at work.

I believe there is a realistic chance that this creative emergence, where the whole inspires the emergent structure and gives it direction, can, at least in part, be encompassed by a computational formula. This hinges on the

construction of a new mathematical formalism in which nested systems of symmetries and centers are acted on by the transformations I have described, and where the future (the $(t+1)$ -trajectory in phase space) is determined by action on the latent structures already present at time t , where these latent structures, too, are couched and described in terms of symmetries and centers.

10. USING HARMONY-SEEKING COMPUTATIONS IN BUILDING DESIGN

In natural systems we may expect (in the absence of truly extraordinary conditions) that when things are “okay”, harmony-seeking computations are occurring. This also happened in traditional societies when people made buildings. But in human-created buildings of the modern era, this no longer happens automatically. We now find it necessary to consciously ensure that harmony-seeking methods are used to generate plans and building designs.


10.1. Example: a housing project for 200 apartments in Boston, generated by Harmony-Seeking Computations

As an example, I show the first half dozen steps for a housing project. The design was performed by consciously following a sequence of harmony-seeking computations, in the fashion sketched in this paper. The process starts with an abandoned industrial site (figure 36) in Somerville, just north of Boston, Massachusetts.




Figure 36. The pre-development industrial site: it is 5.5 acres, triangular in shape, and lies between a railroad line, a bike path, and an existing neighborhood.


Figure 36. The following six diagrams show the first six steps of the harmony seeking process.




Step 1. The first step in the computation identifies latent centers in the site: the bike path, along the south side, the point of contact where Warwick avenue comes in from the west, and the entry from Lowell Street between two railroad bridges, in the south east corner.



Step 2. Starting from these latent centers, the computation identifies the connection with the surrounding neighborhood at the two endpoints, and generates a smooth curved line, aligned in a position that keeps roughly equal areas on either side, leading to stairs at the east end, where there is a 20 foot rise to Lowell street.




Step 3. To establish this line as a major spine or promenade, the next step designates it more clearly as a walking promenade, possibly tree-lined, with public open spaces at the two ends, thus giving it strong presence and identity in the neighborhood.



Step 4. The pedestrian spine now has, as its most important latent centers, the areas on either side of it. To place 200 dwellings in this limited area, and to do it in such a way that people have their own gardens that their neighbors can also enjoy, leads to houses laid in strips around squarish open spaces. The gardens should, typically, have a diameter on the order of 50 to 120 feet; the buildings in strips should have a depth of some 20-25 feet; the building height would then be 2 and 3 stories. (The squares in the diagram are the gardens, not the buildings.)



Step 5. The grid-like array of square gardens is characterized by two parameters, courtyard diameter, and building thickness. However, its regularity is not essential or even good: rather, the layout needs to be adapted to the site boundaries in such a way as to generate coherently shaped courtyards and pedestrian space, while leaving building volumes simple. This transformation requires the use of ROUGHNESS, LOCAL SYMMETRIES, DEEP INTERLOCK, NOT SEPARATENESS, GOOD SHAPE and INNER CALM. Application of these transformations achieves the necessary computation at this stage: it suggests buildings on either side of the promenade, forming the space of the promenade, with different shaped gardens on either side, filling the space as comfortably as possible.



Step 6. The schematic arrangement of the previous cycle now gives way to a series of shapes that pay more detailed attention to each individual garden as a shape in itself, using GOOD SHAPE, LOCAL SYMMETRIES, POSITIVE SPACE, THE VOID, and ALTERNATING REPETITION, so that the whole is coherent, and feels like one thing.

The boundary where this neighborhood abuts other neighborhoods is modified by THICK BOUNDARIES, DEEP INTERLOCK, and ALTERNATING REPETITION, so that it becomes a thickened semi-permeable membrane, allowing people to talk a stroll, to pass in and out comfortably, yet also maintaining a certain privacy for the interior of the neighborhood.

In figure 37 we see the framework for a further harmony-seeking process. The plan shows 200 households on a 5.5 acre site, each house unique, and where the patterns of gardens, walkways and road access have been laid down so as to protect the harmony of the adjacent neighborhood, and project the immediate environment for 500 people.

10.2. The uniqueness of each region in the generated structure

If we examine the structure generated for Somerville by the harmony-seeking computations, we see that each part, though similar in structural character to others, is *unique* (figures 38 and 39). That comes about because the application of the computation to even slightly differing contexts inevitably produces morphologically different results. After a few cycles of wholeness-extending transformations, the difference is magnified, and leads to substantially different configurations. This is the origin of uniqueness in the world. A world with genuine, and profound, uniqueness at every spot arises whenever harmony-seeking computations govern the unfolding.



Figure 37. Final plan, after the six steps I have described, and approximately another twenty, comparable steps dealing with more detailed fine structure. (For lack of space, these could not be illustrated here)

For example, at the largest level of scale, each of the dozen or so courtyard gardens (figure 38) has a different shape and configuration. Differences are not arbitrarily inserted to be cute: the application of the principal transformations (POSITIVE SPACE, LOCAL SYMMETRIES, ECHOES and GRADIENTS, ALTERNATING REPETITION, CONTRAST) generates a different configuration for each one according to its starting point, according to its context.



Figure 38. Detailed view of the houses and apartments, showing the detailed adaptation and variety that arises from the use of harmony-seeking computations




Figure 39 Different gardens in the project, each with its unique character and atmosphere

The same quality of uniqueness and subtle differentiation continues to smaller and smaller scales. The detailed configuration of terraces, entrances, paths, lawns, stairs and archways, produces unique results in each part of the larger whole, and in each part of the individual gardens (figure 39). This is not from a shallow desire to make each thing different for its own sake (sometimes the driving force behind the more commercial postmodern housing developments); it occurs because the effect of harmony-seeking computations, on only slightly different starting conditions, is to generate entirely new and different configurations, but all members of the same family.

11. WHOLENESS-EXTENDING TRANSFORMATIONS

There is a single theme present in all these studies: atomistic, bottom-up computations cannot adequately describe what is really happening in the world, and, do not describe the especially important processes that *enhance the structure* the world, that bring order into configurations in the land, in nature, in buildings deeply adapted to the land and deeply adapted in their internal structure.

It is not only large-scale configurations that benefit from this insight. Every configuration, at every level, is working to help the coherence of some larger configuration in which it is embedded. This runs up and down the ladder of scales, and must do so in any living world. A successful computation emulates this upwards- and downwards-reaching process. Purely bottom-up forms of calculation, not invoking this principle, will always remain too sterile to be real or profound.

11.1. What is the underlying process involved in these harmony-seeking computations?

I am proposing not only that we consider these many real world systems as computational processes. We need to find out how they work, and how, in particular, they work to allow the unfolding of structure, under the impact of wholeness. This requires a form of representation that is new. It has been sketched, in broad brushstrokes, in [1][2]. We now need to find ways of describing this unfolding in more well-defined mathematical and computational terms. That will be a long job. My colleagues and I have, in the last few years, gained insight into the nature of the harmony-seeking computations, but we are far from understanding them in detail. That task must now be undertaken, hopefully by many dedicated scientists together.

The phenomena I am talking about cannot simply be grouped under what is loosely called emergence. The emergence of wholes in the world does not come about by autonomous processes that happen to aggregate themselves to form wholes. Rather, the phenomena are the wholeness-extending processes I describe.

11.2. Wholeness-Extending Transformations

In all the real-world examples, the common phenomenon is that the computational steps are wholeness-extending transformations. Each wholeness-extending transformation operates on one wholeness to produce another in such a way as to preserve, or embellish, or enhance the global structure of the first.

$$W_1 \xrightarrow{WE_1} W_2 \xrightarrow{WE_2} W_3 \xrightarrow{WE_3} W_4 \xrightarrow{WE_4}$$

The mathematical description of a wholeness-extending transformation is the subject of future work. However, there is abundant evidence that the concept of being wholeness-extending is well-defined and objective, in the sense that different observers largely agree among that different possible transformations of a given whole, which ones are more wholeness-extending and which are less so.

11.3. Models of the wholeness in a given configuration

To establish the character of wholeness-extending transformations, we begin with five postulates about the structure of wholeness:

Postulate A1. In any configuration we see certain salient wholes, or centers. Each of these wholes is an identified, spatially contiguous subset of the configuration that corresponds to something we see, or experience, as an entity.

Postulate A2. The sub-configurations may be spatially nested, or overlapping, or disjoint.

Postulate A3. Each sub-configuration has a measure associated with it, the degree of coherence or saliency within the larger whole.

Postulate A4. Latent centers may be low in saliency, almost invisible, but nevertheless coherent configurations in their own right, created by others in the configuration. It is these latent centers that play the biggest role in the evolution of the configuration, and in any wholeness-extending process.

Postulate A5. The wholeness is defined as the system of configurations, each one specifying its coherence, and each connected with other configurations which are part of it, or of which it is a part.

We also state four postulates about the definition of wholeness-extending transformations:

Postulate B1. A transformation is considered to be wholeness-extending if it elaborates existing centers or latent centers, and does not introduce new centers that violate or ‘cut across’ existing centers.

Postulate B2. The best latent center to work on is the latent center whose improvement or repair, would (probably) do the most to increase the coherence of the whole configuration.

Postulate B3. A successful transformations transformation must always have positive morphological impact on the structure and coherence of some still larger whole.

Postulate B4. Once a transformations transformation has been applied, new latent centers, not present or not visible before, make their appearance on the backs of the newly introduced structures. It is this process that continually generates newness in the evolution of the configuration.

11.4. A few examples of harmony-seeking computations

Each of the examples in this section (figures 40-45) is relatively simple: the same paradigmatic cycle is followed each time.

- Frame 1: a configuration
- Frame 2: identifies the locus and extent of some latent center in that configuration – one that, if consolidated, would help the whole to become more coherent.
- Frame 3: establishes smaller centers to embody and solidify that latent center.
- Frame 4 (when present): establishes yet further enhancement of new latent centers that have been generated.

The steps taken in going between each frame in each example are some conglomerate of the 15 wholeness-extending transformations.




Figure 40. STRONG CENTER → LEVELS OF SCALE → DEEP INTERLOCK
→LOCAL SYMMETRIES




Figure 41. → THICK BOUNDARIES, THE VOID → LEVELS OF SCALE, ALTERNATING REPETITION, LOCAL SYMMETRIES




Figure 42. → STRONG CENTERS, LEVELS OF SCALE → LOCAL SYMMETRIES




Figure 43. → THICK BOUNDARIES, THE VOID → GRADIENTS, LEVELS OF SCALE, ALTERNATING REPETITION, ECHOES




Figure 44. → POSITIVE SPACE → GRADIENTS, LEVELS OF SCALE → POSITIVE SPACE → MORE LEVELS OF SCALE




Figure 45. →GRADIENTS → SIMPLICITY → POSITIVE SPACE → LOCAL SYMMETRIES → CONTRAST

We can also see how symmetry elaboration is happening at each step. The space is differentiated, holistically, by an injection of LOCAL SYMMETRIES. The symmetries that appear, in most cases, appear because other symmetries are struck out, leaving the coherent form we see. Because of this symmetry-breaking aspect, the injecting of local symmetries is wholeness-extending, not wholeness-destroying. However, it obscures the core of what is going on to think of it in this way *exclusively*. It is more illuminating to focus on the fact that local symmetries of various kinds are being created. The kinds of local symmetry being created, when created well, enhance the underlying structure.

Such a procedure is creating structure in a new way. It may be done by an artist, with an intuitive grasp of the underlying latent structure, at each step in the unfolding of the whole. Or it may be done by an engineer, consciously. Socially speaking, this is a new kind of process, a new function for artists, and a new challenge to engineers, architects and planners. This new kind of work demands enormous concentration and attention.

The intuitive act is nevertheless a computation, and we may be able to pin down what *kind* of computation it is. Then, if we can succeed in making a harmony-seeking computation, perhaps helped by a computer working to achieve similar holistic results, it will be because the thing we recognize intuitively as coherent or whole, is, mathematically, a particular recursively generated structure of symmetries and centers that have the 15 properties in them. It is this underlying structure that allows the human mind, and natural processes, both, to follow this path and to seek wholeness in the way they do. Most important, we may become conscious about this process, and consciously use this kind of computation to improve the coherence and harmony of our world.

12. CONCLUSIONS

12.1. Fact and Value

The architectural examples, and the example of the cosmological structure of voids and filaments, bring the potentially extraordinary nature of harmony-seeking computations into sharp focus. In the parlance of 20th century thought, designing a building is a matter of personal *taste*; it has essentially nothing to do with *fact*. On the other hand, the dynamics that lie behind the distribution of matter in the universe is undoubtedly a matter of fact, although the facts may be still poorly understood and much debated.

No matter how subtle our perception of a building design and its harmony might be, the idea that we treat its structure as a matter of fact may seem absurd to present day scientists, if they follow 20th century thought. It is, for the moment, still a contemporary article of faith that the goodness of a building is a matter of *taste*, not of fact. The harmony of the building cannot be (according to 20th century thought) a matter of fact in any sense. The mental protocols of 20th century thinking have forbidden it.

Here we come to a profound change lying in wait for 21st century science. It is a matter of historical record that scientists of earlier eras – indeed, many of the great scientists of earlier eras – had no difficulty in thinking of the great harmony that exists in the world. Pythagoras’s phrase “the harmony of the spheres” was not an idle one. Newton considered the progress of the universe, in the large and in the small, to be entangled, inevitably, with a movement towards harmony, and with the greater harmony of the world as a necessary underpinning for the discoveries of science [47]. Leibniz, Kepler, and others thought the same.

When we consider the quality of a building merely as a matter of taste, this strongly limits our ability to understand its harmony. When we consider the evolution of the universe as a mechanical process that can best be modelled through value-neutral means, this also strongly limits our ability to understand the ensuing harmony, and creates a real possibility that we shall ultimately fail to understand the physics.

The possibility of harmony-seeking computations ventures into this forbidden domain. It is undoubtedly challenging to define and expand this concept so that it can become an effective part of the way we think about the world. But it may play a considerable part in reopening a door that has been closed for too long.

12.2. A new science of harmony-seeking computation: when and where?

A new science of harmony-seeking computation and wholeness-extending transformations can make amenable to computation phenomena that are beyond the reach of conventional computational methods. It can help to open doors to the global qualities of harmony, figural goodness, ecological health and structural coherence – even value itself – as possibly computable features of configurations.

One first practical item on the agenda is to provide well-defined versions of the 15 transformations in *NOO*. This is a difficult task, for three reasons. Firstly, the 15 properties, though defined with some precision, remain somewhat elusive. Defining computational operations that can induce these properties in arbitrary configurations is challenging. Secondly, it is difficult to define them as transformations, since this presupposes a language of configurations that is amenable to transformation. Thirdly, some of the transformations are harder than others to define operationally in sufficiently concrete terms. For example, LOCAL SYMMETRIES, THICK BOUNDARIES and LEVELS OF SCALE are relatively easy to define. POSITIVE SPACE and ECHOES are harder. SIMPLICITY AND INNER CALM and NOT SEPARATENESS are among the most difficult.

Despite these difficulties, I am reasonably sure that the task of dealing with all 15 transformations can be accomplished, and open the door to a fully fledged, though first-draft, science of harmony-seeking computations. First steps have already been taken (for example, [42] [48] [49] [50]). Harmony-seeking computation may then sit alongside other methods as a new tool in an armory of well-founded computational techniques. It is likely to be appropriate whenever a computational task is defined by adaptation and wholeness, with reference to the position some system has in some still larger whole, or by a desire for beauty, or life, or elegance. Complexity theory, architecture, biology, ecology, physics, cosmology – and computation – may then all be more firmly rooted.

13. ACKNOWLEDGMENTS

I should like to record my very great thanks and my debt to Susan Stepney, for her patience and insight in editing and discussing this paper with me, while it was prepared for publication. I would also like to thank the anonymous referees who gave many constructive and helpful comments.

14. REFERENCES

- [1] C. Alexander. (2001) *The Phenomenon of Life*, Book 1 of *The Nature of Order*. Center for Environmental Structure Publishing, Berkeley.
- [2] C. Alexander. (2002) *The Process of Creating Life*, Book 2 of *The Nature of Order*. Center for Environmental Structure Publishing, Berkeley.
- [3] C. Alexander. (2005) *A Vision of a Living World*, Book 3 of *The Nature of Order*. Center for Environmental Structure Publishing, Berkeley.
- [4] C. Alexander. (2004) *The Luminous Ground*, Book 4 of *The Nature of Order*. Center for Environmental Structure Publishing, Berkeley.
- [5] M. Golubitsky, I. Stewart. (2001) *The Symmetry Perspective*. Birkhauser.
- [6] C. Alexander, A.W.F. Huggins. (1964) On changing the way people see. *Perceptual And Motor Skills*. **19**, 235-253
- [7] C. Alexander. (1965) A City is Not a Tree. *Architectural Forum*, **122**(1):58–62 (Part I), **122**(2):58–62 (Part II)
- [8] L. Wolpert *et al.* (2002) *Principles of Development*, 2nd edition. Oxford University Press.
- [9] GardenWorld. <http://www.gardenworld.co.uk/project-willow.asp>
- [10] M. Wertheimer. (1924) *Gestalt Theory, Über Gestalttheorie* [an address before the Kant Society, Berlin, 7th December, 1924], Erlangen, 1925. In the translation by Willis D. Ellis, *Source Book of Gestalt Psychology*, Harcourt, Brace and Co, 1938.
- [11] W. Köhler. (1930) *Gestalt Psychology*, G. Bell & Sons Ltd..
- [12] K. Koffka (1935) *Principles of Gestalt Psychology* Lund Humphries, London.
- [13] J. Hochberg, E. McAlister (1953) A Quantitative Approach to Figural "Goodness". *Journal of Experimental Psychology* **46**, 361-364.
- [14] C. Alexander, S. Carey. (1968) Subsymmetries. *Perception And Psychophysics*. **4** (2):73–77.
- [15] Aproveitamento de energia eólica.
<http://www.fem.unicamp.br/~em313/paginas/eolica/eolica.htm>
- [16] Miguel Pontes. *Mare Nostrum*.
<http://marenostrum.org/vidamarina/algalia/verdes/acetabularia/>
- [17] B. Goodwin. (1994) *How the Leopard Changed Its Spots: the evolution of complexity*. Simon & Shuster.
- [18] G. Webster, B. Goodwin. (1998) *Form and Transformation*, Cambridge University Press.
- [19] Edmund Bacon. (1974) *Design of Cities*. Viking.
- [20] S.A. Shechtman, S.D. Landy, A. Oemler, D.L. Tucker, H. Lin, R.P. Kirshner, P.L. Schechter. (1996) The Las Campanas Redshift Survey. *Astrophysical Journal*, **470**, 172.
- [21] Greg L. Bryan and Michael L. Norman.
<http://csep10.phys.utk.edu/astr162/lect/gclusters/soap.html>
- [22] G. Bothun. (1998) *Modern Cosmological Observations and Problems*. Taylor & Francis.
- [23] R. P. Feynman, R.B. Leighton, M. Sands. (1965) *Feynman Lectures on Physics vol 3: Quantum Mechanics*. Addison Wesley

- [24] C.J. Polinsky. (1999) *Flight Simulation of Flocking Geese Using Particle Set Animation*, Swarthmore College. <http://www.cs.swarthmore.edu/~polinsky/cs97/>
- [25] D. Gordon. (1999) *Ants at Work: how an insect society is organized*. Free Press.
- [26] Cornelis Weijer. (1999) Morphogenetic cell movement in *Dictyostelium*. *Seminars in Cell & Developmental Biology*, **10**:609-619.
- [27] S. Johnson. (2001) *Emergence: the connected lives of ants, brains, cities, and software*. Scribner .
- [28] D. Niehoff. (2005) *The Language of Life: How Cells Communicate in Health and Disease*. Joseph Henry Press.
- [29] C. Reynolds. (1987) Flocks, Herds, and Schools: A Distributed Behavioral Model. *Computer Graphics* **21**(4):25-34. ACM SIGGRAPH
- [30] Henri Weimerskirch, Julien Martin, Yannick Clerquin, Peggy Alexandre, Sarka Jiraskova. (2001) Energy saving in flight formation, *Nature* **413**, 697-698 (18 October).
- [31] F. Reed Hainsworth, *Precision and Dynamics of Positioning by Canada Geese Flying in Formation*, Department of Biology, Syracuse University, USA
- [32] Larry Gedney. (1982) Why Birds Fly In Vees, *Alaska Science Forum* August 16, 1982 Article #559.
- [33] NASA. (2001) NASA'S Autonomous Formation Flight: Follow The Leader And Save Fuel, October 29, 2001, Release: 01-61
- [34] M. Resnick *et al.* (2004) *StarLogo 2.1*, Media Lab, Massachusetts Institute of Technology
- [35] aerospaceweb.org. V-Formation Flight of Birds.
<http://www.aerospaceweb.org/question/nature/q0237.shtml>
- [36] NASAexplores. <http://media.nasaexplores.com/01-035/PlaneTrail.jpg>
- [37] L.K. Hjertager, B.H. Hjertager, N.G. Deen, T. Solberg. (2003) Measurement of turbulent mixing in a confined wake flow using combined PIV and PLIF. *Can. J. of Chem. Eng.* **81**, 1149-1158
- [38] R McNeill Alexander.(2004) Hitching a lift hydrodynamically - in swimming, flying and cycling. *J Biol.* **3**(2): 7.
- [39] M. Cross. (2000) Pattern Formation in non equilibrium systems. Course notes from Cal Tech Physics 161b.
<http://www.cmp.caltech.edu/~mcc/Patterns/index.html>
- [40] R. Mech, P. Prusinkiewicz (1996) Visual Models of Plants Interacting with Their Environment. *Proc SIGGRAPH 96*. pp. 397-410 ACM SIGGRAPH
- [41] I. Stewart, M. Golubitsky. (1992) *Fearful Symmetry: is God a Geometer?*, Blackwell
- [42] N.A. Salingeros. (1997) Life and Complexity in Architecture from a Thermodynamic Analogy. *Physics Essays* **10**, 165-173
- [43] W.A. Bentley, W.J. Humphreys. (1931) *Snow Crystals*, McGraw-Hill. (reprinted, Dover 1962).
- [44] K.G. Libbrecht, P. Rasmussen. (2004) *The Snowflake: Winter's Secret Beauty*. Colin Baxter Photography Ltd
- [45] K.G. Libbrecht. *SnowCrystals.com*
<http://www.its.caltech.edu/~atomic/snowcrystals/>

- [46] U. Nakaya. (1954) *Snow Crystals: Natural and Artificial*. Harvard University Press
- [47] M. White. (1997) *Isaac Newton: The Last Sorcerer*. Fourth Estate
- [48] N. Salingaros. (1998) A Scientific Basis for Creating Architectural Forms. *Journal of Architectural and Planning Research* **15**, 283-293
- [49] C. der Groot, C. Alexander. (2004) First Steps in a Site Layout Program written in Squeak. (unpublished)
- [50] C. Alexander, S. Cowan. (2005) The Field of Wholeness: Notes Towards A Mathematical Model. (unpublished)