Working Paper

No. 2004.17

Kristian Kreiner

**Knowledge Management
as Attention**

INSTITUT FOR ORGANISATION OG ARBEJDSSOCIOLOGI
Handelshøjskolen i København
Solbjerg Plads 3
2000 Frederiksberg
Tlf: 38 15 28 15   Fax: 38 15 28 28

KNOWLEDGE MANAGEMENT AS ATTENTION

Principles and Practice of Product Development

by

Kristian Kreiner

Copenhagen Business School
Department of Organization

Main distinctions and concepts: Intention versus attention; dispositional versus episodic words; translation versus transference; the labyrinth.

Introduction

This article explores the case of product development for insights into the potential role of knowledge management. Current literature on knowledge management entertains the notion that knowledge management is a specific set of practices – separate enough to allow specialization of responsibility. By common standard, the proclaimed responsibility of knowledge management is shared knowledge, saved learning costs and coordinated action in an organization. The significance of the practices of knowledge management is *the intention* of shared knowledge, saved learning costs and coordinated action.

In a fundamental sense we may confront the notion of knowledge management as a specific set of practices. If we assume that management is a dispositional word, in the way Gilbert Ryle used this notion, then management is a dispositional word that signifies "abilities, tendencies, or pronenesses to do, not things of one unique kind, but things of lots of different kinds. ... The temptation to construe dispositional words as episodic words and this other temptation to postulate that any verb that has a dispositional use

[managing] must also have a corresponding episodic use [specific acts of management] are two sources of one and the same myth" (Ryle, p.114). Thus, knowledge management is not a separate set of specific acts of knowledge management; management is a quality of the ways in which all sorts of acts are performed; to know is a "capacity verb … of that special sort that is used for signifying that the person described can bring things off, or get things right." (Ryle, p. 128-29)

Let us assume that knowledge management is a disposition to "add together" knowledgeable people. By adding together I am referring to Drucker's notion of adding knowledge to knowledge as the wealth-creating process of the knowledge society.[1] Building on the knowledge of others in pursuit of one's own aims is of course an important ambition – and a necessary one in most daily endeavors as well as in high-technological product development. It is my contention that the interface between knowledgeable people, and in general between knowledge domains, is the point of attention that we may identify as knowledge management. We do not specify the actor; and we do not specify the range of actions that would qualify as knowledge management. All kinds of actors can adopt such a focus, and all kinds of acts can be performed with a focus on the interface between knowledgeable people, i.e. people with an ability to bring things off within a certain area.

I model the focal interface of knowledge management in terms of two alternative processes: a process of translation and a process of transferring. The translation takes place when the abilities of one party is black-boxed and put at the disposal of some other party as a tool, a problem etc. Transferring takes place when the abilities are transferred so that both parties to a certain degree have the same abilities and pronenesses. Interfaces may thus be

---

[1] Kreiner & Tryggestad (2002).

characterized in terms of different degrees of redundancies in terms of abilities – of overlapping knowledge domains. We may depict the two models by a Venn diagram.

--- insert Figures here ---

This article illustrates – as a contrasting image of knowledge management to the one we receive in the knowledge management literature – that the translation strategy is a viable one under certain circumstances – that the boundary objects linking significantly different and scientifically separate knowledge domains may serve perfectly well to translate the abilities of one domain into platforms of excelling performance of the other. The necessary redundancy in terms of ability, understanding and knowledge, is thus an empirical question, not a matter of definition and logical implication.

COMPUTERIZED HEARING INSTRUMENTS: A CASE STORY[2]

> 1995 was the year when hearing instruments turned into a computer. It took four years of concentrated development work and more than ten years of audiological research to achieve this. (Oticon Annual Report 1995 – my translation)

This was the jubilant announcement of major technological breakthrough by the Danish manufacturer, Oticon, who had

---

[2] My account is based primarily on publicly available material, e.g. annual reports and (Pedersen. The Genesis of a Digital Hearing Instrument. Hearing Instruments [March], 38-39. 1996.). In connection with another study ((Kreiner & Tryggestad op.cit.)) I interviewed a number of people inside and outside Oticon who were involved in the DigiFocus project. The logical structure of the problem presentation is documented in this material, but the data collection does not allow me to claim detailed insights into the particular processes, the intentions of the actors, and the historical contingencies that impinged upon the development process.

hitherto been famous mainly for its exotic organizational design – the so-called 'spaghetti organization'[3]. The new hearing instrument with a 'digital audio processor' was named DigiFocus.

DigiFocus was meant to *improve the quality of life* for the hearing impaired population. The needs of this population were construed in ways that translated into three specific requirements that the new hearing instrument aimed to fulfill:

- **Miniaturization** – to satisfy aesthetic demands

- **High fidelity** sound reproduction – to satisfy the functional demands in often chaotically changing sound and noise environments. Since hearing impairment is highly idiosyncratic, this requirement included a need for adapting the instrument to the individual user.

- **Usability** – to satisfy the need for forgetting the instrument (and the hearing impairment) in daily life, including avoiding a too often recharging of batteries and an automatic adjustment of volume etc.
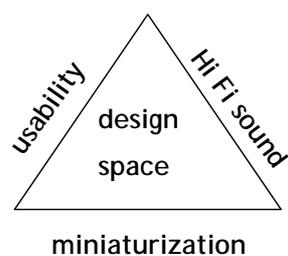

Figure 1: The Design Space for Hearing Instruments

The exploration of this design space, which ultimately materialized in the form of the DigiFocus, was in many ways unpredictable and indirect. The project team continued to hit upon technical

---

[3] (Peters, Tom. *Liberation Management. Necessary Disorganization for the Nanosecond Nineties*. 1992. New York, Alfred A. Knopf.)

problems that needed to be solved before further progress could be made. On one occasion, one that will be the focal theme of this case account, the project got sidetracked to a quite foreign place.

Some researchers have likened the process of technology development to moving around in a labyrinth.[4] At each turn, new obstacles appear that block the direct way to the goal. Solving the problems requires you to take a detour that initially distances you further from the goal, but which eventually allows you to proceed.

I will describe the labyrinth in some detail. I will also describe the competencies and strategies that enabled Oticon, with ingenuity and luck, to find its way in the labyrinth towards the goal.

The technical obstacles of hearing instrument design

To put a small computer into a hearing instrument was the overarching concept of DigiFocus. The heart of a computer is a chip. The high performance of the chip was critical for achieving Hi Fi sound reproduction. Computer chips operate on electricity. In the case of hearing instruments, batteries were (and so far are) the only available source of electricity.

Increasing the quality of the sound reproduction could be translated into an increasing number of operations that the chip needed to perform. E.g. traditionally, compression is done in three frequency bands, but DigiFocus was conceived to compress in seven bands. The more operations required of the chip, the higher its power consumption. Everything else being equal, the increase in sound reproduction quality could be translated into a need for higher battery capacity.

However, the capacity of the batteries is positively correlated to their size. Simply adding capacity by increasing battery size was not

---

[4] Latour, Bruno. "Morality and Technology. The End of the Means". *Theory, Culture & Society* (2002), Vol. 19:5-6; pp. 247-260.

an option in view of the miniaturization constraint. Simply reducing capacity by increasing the frequency of recharging batteries was not an option in view of the usability constraint. An adequately small battery, with a correspondingly small capacity, was not an option because of the Hi Fi sound reproduction constraint. The only logical way out of this design impasse was *to invent* a chip with lower energy consumption.

Lowering the energy consumption of a chip can be achieved by lowering its voltage. At the time, the standard voltage in all modern electronic equipment was 5 volt, but DigiFocus became envisioned to be equipped with a 1-volt chip, which would reduce the power consumption to 1/25! Oticon had never done it before, and existing design tools and libraries were of little use for the chip designers in arriving at a functional design. Work had to be done at the transistor level all along, which made the design job very complex.[5] At the time of taking this turn in the labyrinth it was not at all certain whether such a detour would lead to success, i.e. that such a functional design could be made within the time parameters of the project. This became the more uncertain when further obstacles were encountered.

While reducing the power consumption, reducing the voltage has also less fortunate implications. First of all, it reduces the speed of the chip, which translates directly into a loss in performance. The strategy of lowering the chip's voltage might prove self-defeating unless the chip designers were able to increase the efficiency of the chip itself. Many new design features were invented and built into the chip, e.g. in the form of new ways of parallel processing. However, immediately the choice of a 1-volt chip simply redefined the problem from one of providing sufficient battery power to one

---

[5] (Kreiner & Tryggestad op.cit.)

of increasing the efficiency of the chip, i.e. to perform more functions within a given capacity.

The choice of a lower voltage solution had other implications. With a 1-volt chip the whole spectrum of sound (in terms of frequencies) has to be represented on a vastly reduced scale. This would require a level of precision that was unattainable with the envisioned chip. Thus, the development of DigiFocus arrived at a new impasse. Saving energy by lowering the voltage of the chip not only slowed it down; in the context of hearing instruments it also added new and unattainable processing requirements. It seemed that this strategy led from one impasse to another even worse impasse.

Logically, the need for precision would be relaxed if not all frequencies needed to be represented on the 1-volt scale, i.e. if some frequencies could be skipped. The implied logic is heretic, however, because according to commonsense that would also reduce the quality in sound reproduction. Nonetheless, the question was framed in this way to escape the impasse: can some frequencies be skipped without the human ear noticing a loss of sound reproduction quality – and if so, *which* frequencies could be skipped?

This turn in the labyrinth was dramatic in a different way than the previous ones had been. It shipped the problem out of the hands of Oticon's own experts in chip design and electro acoustics, and into the hands of experts in psychoacoustics. Perceived sound quality is highly subjective, and psychoacoustics conducts experiments on human beings to collect data on their perceptions. Thus, the problem was transported not only to a foreign university where the experiments were conducted over several years; it was also transported to a foreign knowledge domain – foreign in terms of both expertise and methodology.

Eventually, the psychoacoustics experts provided the IC designers with the algorithms and specifications they needed to reduce the replication of frequencies. With this break-through of the hitherto technical impasse, their chip design proved feasible and the DigiFocus became a functional hearing instrument.

***

This is far from the whole story, and importantly, it is *my* story: my reconstruction of the logic behind the impasses encountered along the way – and the logic behind the detours of developing a hearing instrument. It cannot be claimed that these logical steps correspond to the episodic steps in the process. Certainly, it cannot be claimed to correspond to the experience and current memory of the involved experts. In is quite possible, even likely, that my imputation of an underlying logic in the detours was experienced as annoying obstacles and frustrating delays – even as human obstacles and unnecessary delays.[6]

<u>Case analysis</u>

How were knowledge domains mobilized, and how were they coordinated to produce coherent product architecture? What was the content of the intersection of the domains, and by what dynamics did the redundancy of knowledge ebb and flow? We might be tempted to apply hindsight logic and claim that such a brilliant design *could only* have been created consciously and through rationally managed processes. For sure, many aspects of the project were planned. However, consider the following divergent perspectives on what participation meant. Psychoacoustics experts were mobilized on the idea that frequencies could be eliminated without loss of quality. The chip designers were mobilized on the idea of increasing the number of operations with a slower chip.

---

[6] The methodological implications of a logical case derived from an empirical process will need to be discussed.

Both efforts were highly explorative and it would be virtually meaningless to try to define specific demands on the eventual solution. The engaging quality of these problems was a result of their character of "insight problems" that directs our energies towards new ways of *presenting* problems, rather than straightforward searches for solutions. [7]

The new way of presenting the problems seemed to have organizing power. Little communication and mutual understanding was required for the two domains to establish order.[8] A simple idea justified in a language that even we as lay readers understand, and awaiting an answer in less simple, yet directly operational terms, formed the point of tangency. The time it took for the psychoacoustics experts to devise their solution required neither coordination nor social communication. It was a completely decentralized activity, independent of all the other concurrent struggles fought on other technological frontiers. The algorithm proved successful because it had the character of a boundary object. A boundary object has an identity, but it means quite different things to different communities. Visually, it is exactly the point of tangency between communities (or knowledge domains in the present article). The same thing (the algorithm) meant less instructions and higher speed to IC designers, less power consumption to the electrical engineers, and a whole new paradigm for studying the human ear and perception of sound to the psycho-acoustics experts. To the sales people it meant a highly convincing argument (even scientifically accountable) for introducing a revolutionary first-mover product in a very profitable segment of

---

[7] (Simon . Learning to Research about Learning. 1999.), P.26.

[8] There is no doubt that considerably more social interaction and communication took place across the domains of knowledge. Some of this communication had no doubt the function of creating mutual trust and respect. I do not propose to cut out such forms of communication. In the present context, however, my aim is to understand how little communication is required for transmitting the knowledge from one domain to the other.

the market; to the production people it meant new alliances with highly specialized IC fabs. The organizing potency of a boundary object is the number of contexts in which it is enacted and rendered meaningful. The idea of less accuracy in reproducing hi-fi quality was fuzzy, yet meaningful. It gave the psychoacoustics experts the license to work on a new and enlarging set of problems. When they returned with a 'solution' in the form of algorithms, this solution was immediately useful to the chip designers. The chip designers did not need to understand the first thing about psychoacoustics.[9] They were not dependent on the intentions and premises that these foreign experts worked on. They might as well have read the algorithms in a book, had this book existed, which it did not, of course. They imagined that such a book could be written and instigated the search for somebody to write it.

Likewise, the psychoacoustics experts did not need to know anything about chip design, power-consumption and the aesthetic demands of customers. They were handed a problem that made immediate sense within their specialized and isolated knowledge domain. They responded to the problem not because of an appreciation of its significance to others, but because of an appreciation of its significance to their own domain. They produced a solution in a form that did not reflect the needs of the chip designers, but that reflected their own way of working. It was "readable" for the chip designers, not because of planning and coordination, but because a solution came in a form of packaging that reflected the problem addressed.

In conclusion, my account of the new product development process does not resemble the ideal of project teams with its insistence on

---

[9] In fact, they probably did know quite a lot. Also, they probably communicated more with the psychoacoustics experts in Sweden that we have described here. But the chip designers' ability to add knowledge to the knowledge of the psychoacoustics experts would seem not to hinge on such communication and insights.

shared structural capital, high degree of coordination and an architectural design that is frozen from the outset. It is easier to recognize the networking traits of the process, with an opportunity-driven interaction, and a spontaneous order growing out of the outcomes conceived and produced locally by the participants. There is obviously a high productivity of knowledge when e.g. the IC-designers put the insights of psychoacoustics experts to work in the 1-volt design. Coordination is not planned, but seemingly unproblematic to attain by the IC-designers when implementing the new knowledge. The mutual social relationship is one of resourcefulness rather than constraints. The technical relationships between miniaturization, Hi Fi sound reproduction and usability are ones of constraint. But since the social relationships do not replicate the technical ones, and because the social knowledge domains do not act as mutual constraints, the need for coordinating them is lessened – almost not present and definitely not pressing.

CONCLUSION

In this conclusion I will return to the overall issue of knowledge management. What lessons may the case study of new product development offer for the ways in which we manage knowledge in organizations?

It is tempting to use the "dispensing with frequencies" as a symbol for the new knowledge management agenda. It reminds us that the circulation of as much information and knowledge as possible – and always defining and planning in advance what needs to be done – is not the only feasible strategy. This strategy will tax the capacity of the organization – possibly overtax it, and possibly tax it in vain. The case study suggested that the need to share knowledge, to circulate structural capital, to plan interfaces and interactions, etc. is considerably less than much literature and commonsense want us to believe. The strategy of transferring abilities – of

redundancy of information and skills – is never the only option, and probably often not the most obvious one. An alternative strategy for knowledge management might start from the question of "how little" rather than "how much" redundancy. It would entail experimentation with reducing the pre-designed intersection of knowledge domains to a level of requisite redundancy.

To be sure, the requisite redundancy is not definable in absolute and quantitative terms. Any search for requisite redundancy would also have to search for *new types* of redundancy – redundancy with high potency to engage and organize. The conclusions of the case study should be carefully assessed. The 'teamwork' between the IC-designers and the psychoacoustics experts involved little communication and planned coordination. But it was the efficient boundary object – the "insight problem" – that enabled the connection to be made and the coordination of the distributed efforts. By all indications, the formulation of this insight problem was not a trivial task – rather the result of long experience and dedicated experiments with advanced chip design. Nor did it start as an authoritative definition of the problem – rather a tentative formulation that *was given* authority by the psychoacoustics experts' reception of it. The true nature of the labyrinth is the surprising impasses that are encountered. No one has the layout of the labyrinth in mind; one's destiny of taking some detour is uncertainty; and the aim of the travel, i.e. the exit from the labyrinth, may easily be found retrospectively.

It is surprising that knowledge management, and theories about new product development, has not given the requisite redundancy more thoughts. Instead of mindlessly advocating more knowledge sharing, more front-end planning, and more detailed coordination of imagined constraints, at least the specialization of knowledge and the mechanisms for applying knowledge to knowledge should

enter the discussion as potential complications. It is furthermore surprising that knowledge extraction and structural capital, control and manageability have apparently been given higher priority than the struggle to increase the productivity of knowledge. At the level of product development it seems paradoxical that a predefined and preconceived architecture is not confronted with the implication that such architecture prevents the participants from taking advantage of the knowledge creation that defines the project. The case study illustrated that the architecture (at least in some non-trivial aspects) was easily adaptable to the actually solutions being produced locally. Another organizational structure might conceivably put more emphasis on maximizing the knowledge productivity – the provision of boundary objects that enlarge the solution spaces rather than constrain them. As we learned from the case study, carefully formulated questions may better allow people to coordinate their actions than pre-specified solution parameters. A knowledge management more focused on the distillation of problems and issues than on the circulation of readymade solutions, more focused on stimulating search than on providing information, more focused on enabling local knowledge creation than on making it redundant - such are some of the ingredients to a new agenda for a knowledge management that I envision would be less reactionary and more realistic than the presently prevailing one.

Knowledge management may come to represent a type of attention more than a type of intention.[10] It may become focused on the ways in which translations take place at the boundaries of knowledge domains. Management is not about bringing different domains on speaking terms – to subject them to some common understanding and coordinate their efforts by design. The ideal of management of a smooth and direct route to the lofty aims of

---

[10] This is similar to Weick's suggestion that design is a matter of attention more than intention.

rationality is naïve and counterproductive. The ideal route of technological development is one of traveling far and wide in the knowledge labyrinth. Impasses are more important than open gates, since they help us explore the new world of possibilities, rather than traversing the already known territories. The knowledge management may be productive in engineering the points of tangency between knowledge domains – the ideas, problems, questions and objects that each domain, within their own idiosyncratic understanding of reality, may relate to meaningfully. Designing the "algorithm" as a boundary object that allowed the chip designers and the psychoacoustics experts to collaborate without understanding each other – without intending to collaborate and without the perception that they do – may replace the current futile images of knowledge management as providing common understanding.

Figure 1: The transferring strategy

Figure 2: The translation strategy