PROBLEM SOLVING

Kevin Dunbar

Department of Psychology, McGill University

Published in Dunbar, K. (1998). Problem solving. In W. Bechtel, & G. Graham (Eds.). A companion to Cognitive Science. London, England: Blackwell, pp 289-298.

Address all correspondence to
Kevin Dunbar
Department of Psychology,
McGill University,
1205 Docteur Penfield Avenue
Montreal, Quebec, Canada H3A 1B1

Email: dunbar@ego.psych.mcgill.ca
Phone: (514) 398-6112
FAX: (514) 398-4896

PROBLEM SOLVING

Kevin Dunbar

Department of Psychology, McGill University, Montreal, Quebec, Canada H3A 1B1

In the movie “The Gold Rush” Charlie Chaplin and his “friend” are stranded in a log cabin in the middle of winter while a blizzard rages. The cabin is isolated and they have a *very* big problem -there is nothing to eat. They pace around wondering what to do. Charlie’s friend starts to see Charlie as a chicken and he tries to kill him. He chases Charlie around the cabin many times. Eventually they hit upon the idea of boiling an old boot and eating it for dinner. With great delicacy they sit at the table and eat the boot as if it were a gourmet meal. They solved the problem of having nothing to eat. While their solution to the problem did not result in a culinary feast, this example reveals two crucial features of Problem Solving. First, a problem exists when a goal must be achieved and the solution is not immediately obvious. Second, problem solving often involves attempting different ways of solving the problem. Put more formally, a problem has four components. First, there is an initial state. This is the person’s state of knowledge at the start of a problem. Second, there is the goal state, this is the goal that the person wishes to achieve. Third are the actions or operations that the problem solver can use to get to the goal state. Fourth, is the task environment that the solver is working in. The task environment consists of the features of the physical environment that can either directly or indirectly constrain or suggest different ways of solving a problem. In this article I will sketch out the main currents of thinking in research in this area. I will begin this review by giving a history of research on problem solving, then I will focus on a number of important issues in problem solving research. Finally, I will give an overview of some recent developments in problem solving research.

A Miniature History of Problem solving

Research on problem solving has a long and varied background. Many of the early psychologists at Würzburg such as Oswald Külpe, Karl Bühler, and Otto Selz investigated the mental processes that are engaged in during complex reasoning and problem solving. While they made a number of interesting discoveries on the nature of thinking and problem solving, that had a large effect on the later Gestalt school of Psychology, their research has been almost forgotten by contemporary researchers. The Gestalt psychologists overlapped with and continued research on complex thinking and problem solving. In the 1940's and 1950s gestalt psychologists investigated how people solve difficult problems. This research resulted in a number of classic problems that have been used extensively in problem solving research. For example, Karl Dunker's radiation problem in which subjects were asked to find a way to destroy a stomach tumor without destroying the surrounding tissue. The types of problems that the Gestalt psychologists used often were problems in which subjects must discover a crucial element and once this element is discovered all the other elements fall into place and the problem is solved. These are insight problems. Gestalt psychologists argued that rather than problems being solved by trial and error (as the behaviorists argued), problem solvers gain "insight" into the problem. The Gestalt psychologists often referred to the four different stages that a person might go through in solving a problem (preparation, incubation, insight, and verification). What insight is, and the mechanisms underlying insight had to wait until more detailed accounts of problem solving and languages for capturing problem solving were invented. One further approach to problem solving is the one taken by Sir Fredrick Bartlett in his 1958 book Thinking. In this book, Bartlett characterizes problem solving as a form of exploration.

It was not until the 1960s when Herbert Simon and his colleagues began investigating human subjects solving difficult problems that problem solving research took its current form. There were several distinctive and influential features that characterized Simon's approach: First, he used complex problems in which there was no one key element that led to the solution of a problem, and thus the focus was not on insight, but on characterizing the processes underlying all problem solving. Second, Simon used concurrent verbalizations (rather than introspections) obtained from subjects to identify the mental operations, representations, and strategies that people use when they solve problems. Third, Simon and his colleagues built a series of computer programs that simulate human problem solving. Using both protocols and computational modeling, Newell and Simon (1972) were able to propose a comprehensive theory of problem solving that continues to be at the heart of contemporary theorizing about problem solving.

Since the 1970's researchers in problem solving have tended to use the approaches of Simon, or have used a more descriptive approach of problem solving in the Gestalt tradition. One key aspect of research on problem solving has been the use of verbal protocols (see Ericsson chapter). Using this approach, subjects are asked to state out loud what they were thinking while they are solving a problem. These "think aloud protocols" become the data for formulating models of problem solving. The researchers use the protocols, and whatever actions the subjects took, to build a model of the problem solving strategies that the subjects used. The early work on problem solving was concerned with problems that were puzzles or games such as the Tower of Hanoi task (see below for a description of the task). Later research has tended to focus on more complex 'real world' tasks taken from the domains such as science and writing.

Understanding Problem solving: Searching Problem Spaces

Newell and Simon (1972) proposed that problem solving consists of a search in a problem space. A problem space has an initial state, a goal state, and a set of operators that can be applied that will move the solver from one state to another. Thus, the problem space adds the notion of an operator to the definition of problem solving presented earlier. The complete set of states that can occur when the operators are applied is known as the problem space. A classic task that has been used to investigate problem solving is the Tower of Hanoi task. The initial states and goal states for this task are shown in Figure 1. In this task, a subject is given a board with three rods on it. Three disks of decreasing size are placed on the leftmost rod. The goal of the subject is to place all the disks on the rightmost rod. There are two rules for moving the disks -only one disk can be moved at a time and a larger disk can never be placed on a smaller disk. The problem solver will have to perform many actions (either physical actions, or mental operations) to get from the initial state to the goal state. In the Tower of Hanoi example, the only operation that a solver can perform is to move disks.

Insert Figures 1& 2 about here

In the Tower of Hanoi problem, the problem space consists of the initial state when all the disks are on the Peg A, the goal state when all the disks are on Peg C, and all other possible states that can be achieved when the operation of moving a disk is applied. The complete problem space is given in Figure 2. The complete problem space consists of 27 states. Note that each state is linked to each other state by the movement of one disk. In the standard version of this task, the initial state is state 1 and the goal state is state 8. This problem can be most efficiently solved by going from states 1 to 8 in increasing order. The problem space shown in Figure 2 also reveals that there are many different ways of getting from state 1 to state 8. The problem space thus provides a way of understanding the different ways that a problem can be solved.

One important comment about problem spaces is that the problem solvers are not presumed to have the entire problem space represented in their mind when they are solving a problem. Often, problem solvers will only have a small set of states of the problem space represented at any one point in time. Furthermore some problem spaces, such as that for chess, are so large that it is impossible to keep the entire space in mind. In many problems, the problem solver will not be able to consider all possible problem states and will have to search the problem space to find the solution. Thus, one of the most important aspects of problem solving becomes one of searching for a path through the problem space that will lead to the goal state. Problem solvers will use strategies or heuristics that allow them to move through a problem space.

People often use heuristics for searching problem spaces. In problem solving research, a heuristic is a rule of thumb that will generally get one at the correct solution, but does not guarantee the correct solution. An example of a heuristic might be that "If I start playing tic-tac-toe by putting an X in the middle square, I will win. This heuristic does not always work; sometimes I lose even with this strategy! Heuristics can be contrasted with algorithms, where application of the algorithm always guarantees the correct answer (e.g., the rules of addition). I will discuss some of the different heuristics for searching a problem space in order of increasing complexity. The most simple search technique is to randomly pick a next step. Often people will use this strategy when they have no idea what will lead them to the goal state. A slightly more complex strategy is to move to the state that looks most like the goal state. In this situation, the solver just looks one move ahead and chooses the state that most closely approximates the goal state. This is known as a hill climbing technique. This strategy can be useful if it is impossible to look more than one move ahead, but can lead subjects astray. That is, a move that locally may look like it is bringing the solver closer to the goal state, may in fact be taking the solver further away from the goal state. For example, in the Tower of Hanoi task, moving from state 5 to state 23 may look like the best move as state 23 more closely resembles the goal state than state 6. However, state 6 would lead the problem solver faster to the goal state than state 23. For problems such as this, a more effective strategy is means-ends-analysis. Using this strategy, a solver looks at what the goal state is and sees what the difference is between the current state and the goal state. If the solver cannot apply an operator that will get to the goal state, because the operation is blocked or cannot be executed, then the solver sets a subgoal of removing the block. Using means-ends analysis, the problem solver then decomposes the difference between the current state and the goal state into another subproblem and sets a goal of solving that problem. In this situation, removing the blocked state becomes the new goal. If that subproblem cannot be solved using the current operators, the problem is further decomposed until an operator can be applied. This strategy can be applied recursively until the problem is solved. When the solver can solve one of the subproblems the solver can then solve the higher level problem, and ultimately reach the goal state. Means-ends-

Analysis is a particularly useful strategy for solving the Tower of Hanoi problem. An initial goal might be to get the largest disk onto the rightmost rod, but this goal is blocked by the medium sized disk. A subgoal is set to get the medium sized disk out of the way, but this is also blocked, so a new subgoal is set of removing the smallest disk. This goal can be achieved, and the solver can then achieve the goal of removing the medium sized disk.

Strategies such as hill climbing and means-ends analysis involve an incremental search of a problem space. For example, the problem solver moves from state 1 to state 8 via the other 6 states in the Tower of Hanoi problem. Not all search heuristics are like this. For certain problems it is possible to jump from one part of a problem space to another part bypassing many of the intermediate states. One heuristic for jumping from one part of a problem space to another is to reason analogically. If the problem solver has solved a similar problem in the past, she or he can go directly to the solution by mapping the solution to the old problem onto the current problem. A good example of this is the way that subjects solve Duncker's radiation problem (mentioned above) when they have been pre-exposed to a problem with a similar solution. Keith Holyoak and Mary Gick (1983) told subjects about a problem where an army is attacking a fortress. All the roads leading to the fortress have landmines that will explode when the weight of an entire army walks on it. The general breaks up the army into smaller groups to avoid setting off landmines. Each small group takes a different road to the fortress and they converge at the fortress. Subjects were then given Duncker's radiation problem to solve. Subjects that have been exposed to the fortress problem map the army onto the radiation and will solve the radiation problem by proposing to break up the radiation into a converging set of rays. We can interpret the Holyoak and Gick results from a problem solving perspective. In their task, subjects are shown the correct path through the problem space by being given the solution to the fortress problem. The subjects can then use a previous problem (the fortress problem) to solve a new problem (Duncker's radiation problem) and bypass conducting an incremental search of the problem space. Thus, analogical reasoning can be seen as a powerful strategy for making a more efficient search of large problem spaces.

Using analogy to search large problem spaces is very efficient and is frequently used in science (Dunbar, 1996).

Types of Problems

Researchers have distinguished between two main types of problems; well defined and ill defined. Well defined problems have a definite initial state and the goals and operators are known. Examples of a well defined problem might be solving an equation or addition of numbers. Ill defined problems are ones in which the solver does not know the operators, the goal, or even the current state. Examples of an ill defined problem might be finding a cure for cancer, or writing the first great twenty-first century novel. For ill defined problems the solver must discover operators, define more specific goal states, or perhaps even the initial state. Most research has been conducted using well defined problems. Researchers such as Simon have argued that many of the heuristics that are used in well-defined problems will also be used in ill-defined problems, and that the distinctions between the two types of problems may not be as important as they appear at first glance.

Is problem solving restricted to solving puzzles? While much of the early work on problem solving was concerned with solving puzzles, much of human thinking and reasoning can be regarded as a form of problem solving. Problem solving strategies can be used in many different domains, regardless of what the domains are. The goal of much research in problem solving has been to identify these domain general heuristics, as they are clearly of great importance and can be used to understand problem solving in general. Domain general strategies are contrasted with domain specific strategies which are applicable only a very specific domain such as chess, or designing an experiment using a particular hormone in endocrinology. Herbert Simon and his colleagues have characterized scientific thinking and concept acquisition as forms of problem solving and have identified a number of domain general strategies such as “Finding A Constant” that can be applied in a large number of different domains (Kaplan & Simon, 1990). The problem solving framework has been applied to many domains ranging from Architecture to Medical reasoning to Scientific Reasoning. The reason that it has been possible to regard these different activities as a form of problem solving is due to the concept of searching a problem space. Thus, Simon and Lea (1974) and Klahr & Dunbar (1988) have argued that much of concept acquisition and scientific thinking can be thought of as a search in two problem spaces: A space of hypotheses and a space of Experiments. What researchers have done is to identify the structure of these spaces and the heuristics that are used to search these spaces.

Problem solving and representation

One of the key elements to solving a problem is finding a good way of representing the problem. When a problem is represented, the solver will forefront certain features of the problem and use these features to choose what to do when searching a problem space. This is often the case in politics. Frequently, rival political parties will form different representations of what is the cause of an economic problem, and propose very different operators to apply when solving the problem. Thus, tax cuts will be a solution offered by one party and tax increases will be the solution offered by the rival party. What is at the root of these different solutions is an underlying difference in the way that the problem is represented. What this means is that the same problem will have a number of different representations, and that some representations may be more beneficial to solving a problem than other representations. This often happens in science. For example, molecular biologists recently discovered the breast cancer gene (BRCA1). When different laboratories were in search of the breast cancer gene, each laboratory represented the problem in different ways, and it turned out that one way of representing the problem was the most efficient in discovering the gene. Thus rival laboratories having different representations of the problem space that they were searching in resulted in the use of different strategies for finding the breast cancer gene.

How do problem solvers find a representation? In experiments using fairly simple problems such as the Tower of Hanoi, the subjects construct their representation based upon the problem statement and features of the task environment. The solver isolates what she or he thinks are the relevant features of the problem and then constructs a representation of the problem using those features. By varying the types of instructions given and monitoring their effect on problem solving behavior, it is possible to discover the effects of different types of representation on solving a problem. For example, in the Tower of Hanoi problem, researchers have used different isomorphs of the problem. One version might be the disk version discussed above. Another version that has been used is the tea ceremony problem. In this problem, three people must perform an oriental tea ritual. There are a set of rules that specify what order the steps can be carried out in, and the problem solver must discover the sequence of steps that can be used to complete the ritual. The underlying structure of the tea ceremony problem and the Tower of Hanoi problem is identical, however, the cover story and task environments are different. What researchers have found is that the cover story can have a very large effect on the way that a person will represent a problem, even minor differences in wording of a problem can lead problem solvers to very different types of representations of a problem.

General Models of Problem solving

Researchers have proposed many different models of how particular problems are solved, however, there have been few general models of the problem solving process. The general models of problem solving that have been proposed have been based upon the problem space hypothesis and have been instantiated as production systems (see Schunn & Klahr this volume). Knowledge is represented in these models as symbols, and the production systems operate on the symbols to produce new knowledge and solve problems. The main place that problem solving occurs is in "working memory." Working memory is the part of memory in which computations on the currently active symbols take place. What happens is that a production (i.e. a rule) will replace one symbol with another, and will incrementally search through the problem space. When solving a problem the production systems construct many temporary representations in working memory as a problem space is searched.

One of the most influential early models was the GPS model of Newell, Shaw, and Simon. In this model, they proposed a comprehensive model of problem solving that incorporates many different search strategies and can solve problems from many different domains. One central component of GPS is means-ends analysis. When the GPS program is given a Tower of Hanoi problem to solve, it will use means-ends analysis to solve the problem. In 1990 Newell proposed a Unified Theory of Cognition using the SOAR architecture, in which all human behavior can be thought of as a search in a problem space. Newell and his colleagues have applied the SOAR architecture to a wide variety of domains, and have shown how search in a problem space can be used to understand many different aspects of cognition. John Anderson (1983, 1993) has proposed the ACT family of models that he has used to account for cognition in general and problem solving in particular.

Other approaches to Problem solving

Two other computational approaches to modeling problem solving have been the connectionist approach and a hybrid approach that is a combination of a symbolic system and a connectionist system. There are few connectionist models of problem solving processes, more than likely due to the fact that problem solving frequently involves the use of many temporary representations. These type of processes have proved somewhat difficult to model in connectionist systems, though it is an important topic of contemporary research. When connectionism has been used to model problem solving, it has been used in hybrid models in which there is a symbolic level that is connected to a connectionist layer. The Barnden and Holyoak (1991) book has an interesting collection of hybrid models. This will also be an area of problem solving research that will rapidly change over the next few years.

Recently, there has been much debate in the problem solving literature regarding the role of the task environment in problem solving. Some researchers, such as Jim Greeno and Lucy Suchman, have argued that one of the most important factors in problem solving is the task environment and that the task environment is the major determinant of how a person will solve a problem. This “situated” viewpoint stresses the role of the objects and physical features of the environment and how the environment constrains what a problem solver can and will do. A summary of different views on this topic can be found in a special edition of the journal Cognitive Science edited by Norman (Norman, 1993).

Recent developments in Problem solving research

One important issue in problem solving problem solving is how previous experience with a problem, or related problems, influences current performance on a problem. Gestalt psychologists proposed the concept of functional fixedness to account for the negative effects of previous experience with a problem. However, it is only recently that the effects of experience have been integrated into a problem space view of problem solving. Marsha Lovett and John Anderson (1996) have looked at the way in which previous experience with a particular class of problems might influence performance on the current problem. They have shown that problem solvers will use both the current state of the problem and their previous history of success at using specific operators when deciding what to do next while solving a problem. For example, if applying an operator, such as moving a piece in chess, theoretically leads one closer to the goal state, but past experience with that particular operator leads to failure, a problem solver must incorporate these two sources of knowledge when attempting to solve the problem. They have developed a model in ACT-R that combines both experience with the particular operators and the current state of the problem in an additive manner to predict performance on a problem solving task. This combination of previous experience and state of the problem is a consequence of the cognitive architecture and not necessarily accessible to consciousness. Indeed, Schunn and Dunbar (1996) have investigated how solutions to one problem can be primed by solving a similar problem, and have shown that subjects are often unaware that earlier experience on a problem is having a predictable effect on their current problem solving efforts.

Much cognitive research on problem solving has focused on the ways that an individual problem solver sets about solving a problem. Furthermore, much work on problem solving has involved the use of puzzles in very artificial domains. Recently, a number of researchers have started to investigate more complex real-world problems and problem solving in groups (e.g., software design, engineering, and science). These researchers have found that much real-world problem solving takes place in groups -rather than individuals. Work on problem solving in groups indicates that groups encourage the generation of alternate representations of a problem. When solving a problem, a group can potentially examine a number of possible representations of a problem and decide which representation appears to be the best. Furthermore, in group problem solving, steps such inductions, deductions, and causal reasoning can be distributed among individuals. Thus, an important component of group problem solving is distributed reasoning (Dama & Dunbar, 1996). Group problem solving is not always successful. If all the members of a group are from the same background, they tend to represent the problem in the same way. If their representation is incorrect, they fail to solve the problem. If, however, the members of the group are from different backgrounds, but also share similar goals and have overlapping knowledge bases, many representations of a problem are generated (Dunbar, 1996).

Research on problem solving has tended to take place independent of research on other higher-level cognitive activities, such as concepts, decision making, induction, deduction and causal reasoning. However, each of these areas could be regarded as a form of problem solving (Simon & Lea, 1974). While a number of the models such as John Anderson's and Allan Newell's have incorporated problem solving into a general account of cognition, problem solving and search in problem spaces has been regarded as only applicable to well defined problems such as puzzles. Consequently, a challenge for researchers in problem solving is integrating problem solving with other cognitive activities such as memory, reasoning, and decision making.

Overall, research on problem solving has centered around the notion of the representation of knowledge and the concept of working in a problem space. Early research focused on puzzles, and more recent research has focused on more complex domains. The shift to more complex domains has necessitated the postulation that problem solvers search in multiple problem spaces, rather than one problem space (Klahr & Dunbar, 1988), and has forced researchers to give much more explicit accounts of the role of the task environment in problem solving. Thus, one of the goals of current research is to determine how people generate new representations and problem spaces as they work on a problem (Kaplan & Simon, 1990). The next decade of research should see models of problem solving that will incorporate theoretical constructs from other aspects of cognition. Finally, researchers in problem solving are now beginning to tackle the question of what is the role of the brain in problem solving and what different parts of the brain mediate what aspects of problem solving. Many new discoveries await researchers within this field.

Acknowledgment

I would like to thank Christian Schunn and David Klahr for their comments on an earlier version of this manuscript.

Recommended Readings

Anderson, J.A. (1993). *Rules of the Mind*. Hillsdale, NJ: Erlbaum

Klahr, D., & Dunbar, K. (1988). The psychology of scientific discovery: Search in two problem spaces. *Cognitive Science*, **12**, 1-48.

Newell, A., & Simon, H.A.(1972) *Human Problem solving*. Englewood Cliffs, NJ:Prentice-Hall, Inc.

Norman, D.A. (1993). (ED.) Special Issue on Situated Action. *Cognitive Science*, 17, 1-147.

References

Barnden, J. A., & Holyoak, K.J. (Eds). (1994). Advances in connectionist and neural computation theory, Vol. 3. Ablex Publishing Corp; Norwood, NJ.

Dama, M., & Dunbar, K. (1996). Distributed reasoning. When social and cognitive worlds fuse. *In Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society*. 166-170

Dunbar, K. (1996). How scientists think: Online creativity and conceptual change in science. In T.B. Ward, S.M. Smith, & S.Vaid (Eds.) *Conceptual structures and processes: Emergence, discovery and Change*. APA Press. Washington DC

Gick, M. & Holyoak, K. (1983). Schema induction and analogical transfer. *Cognitive Psychology*. 15, 1-38.

Kaplan, C.A., & Simon, H.A. (1990). In search of insight. *Cognitive Psychology*, 22, 374-419

Lovett, M, & Anderson, J. A. (In Press) History of Success. *Cognitive Psychology*.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press

Schunn, K. & Dunbar, K. (1996). Priming, Analogy, & Awareness in complex reasoning. *Memory and Cognition*., 24, 271-284.

Simon, H.A., & Lea, G. (1974) Problem solving and rule induction: A unified view. In L.W. Gregg (Ed.), *Knowledge and cognition*. Hillsdale, NJ: Erlbaum

Figure Captions.

Figure 1

The Tower of Hanoi Task

Figure 2

The Problem space for the three Disk version of The Tower of Hanoi Task

Initial State

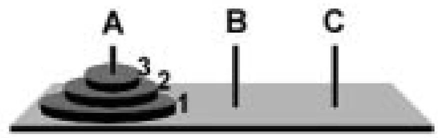

Goal State

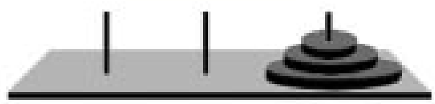

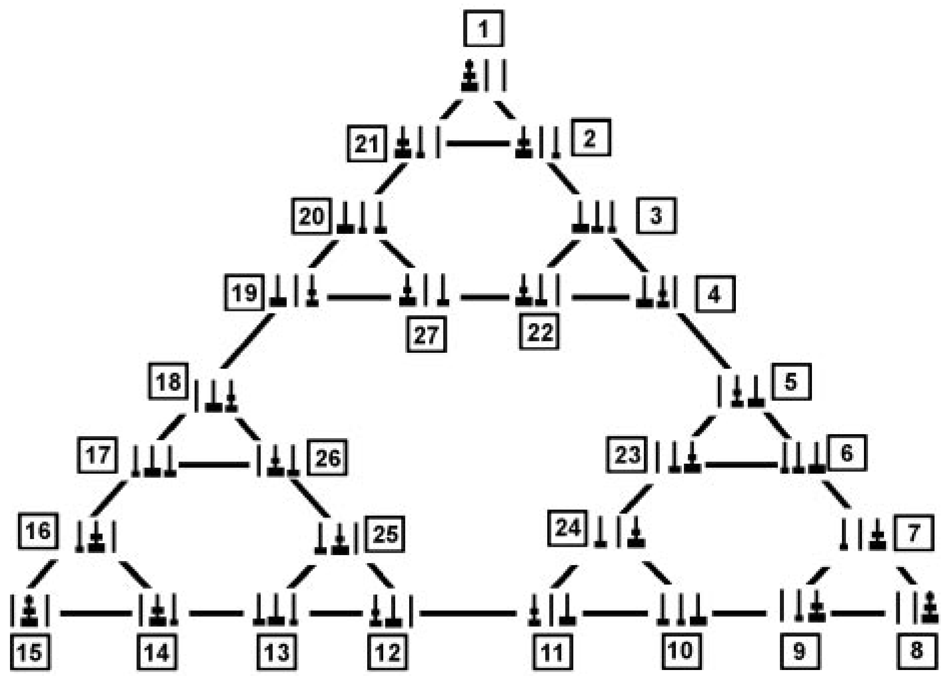

1
21
2
20
3
19
4
27
22
18
5
17
26
23
6
16
25
24
7
15
14
13
12
11
10
9
8