NATURAL HIERARCHIES pdf

1: Natural Hierarchies in Power Pivot and Tabular - SQLBI

Creating a hierarchy in a Power Pivot or Tabular data model is very easy, but you might experience performance issues if the hierarchy is not "natural". This article explains how to create hierarchies avoiding these problems. UPDATE issue described in this article is fixed in Power.

For more with Jim, check out this , and this , and this , and this. Everyone talks about building a meritocracy, but few companies actually operate that way. When you think about how we coordinate human behavior, we tend to do it hierarchically. During the growth of large corporations many organizations adopted the military philosophy: But without a formal hierarchy, how does a reasonable size company operate effectively? Red Hat grew out of the open source movement. In open source, code talks. The best source code ultimately emerges. You can apply that same principle to how you make decisions. The best ideas emerge--if you let them. One person can toss out the same idea as a person who has been an outstanding contributor for 10 years, and the longer tenured person tends to get listened to more. Influence is often based on a history of outstanding contribution. Leaders build their own followership. Why do they follow you? In an opt-in, they follow you because they want to. People who are brilliant, inspiring, contribute a lot--have a solid track record. They naturally are seen as leaders. Think about it this way. In every company, everyone knows who the jerks are. Everyone knows who the kiss-ups are. And everyone knows who the real leaders are. The analogy we used at Boston Consulting Group is thermometers and thermostats: You need the thermostats on your side because they set the temperature. Informal thought leaders are thermostats. Unfortunately most companies jam the leaders in instead of letting them emerge. A definite shift in the traditional career arc is taking place. In addition to careers of advancement there are now increasingly careers of achievement. We try to reflect that in our compensation. As long as you continue to contribute, voice your opinions, make things happen. You can continue to build a reputation and a career without having to change jobs. You always hear from older, traditional managers, "These kids all want to get promoted in six months, they all want to be CEOs in six months. With many ambitious people the outlet for their ambition is having more influence on the company and being a great contributor and building a followership. I always say, "If you want to be a leader, prove it by finding people who will follow you. We want a degree view, not just the view from above. Everyone knows who are the people are who do great work, who create great outcomes, who are the leaders. Use that and not the traditional, hierarchical system of promotion. We still have an org chart, but that org chart does its best to defer to the culture and to the leaders who make great decisions. The people who are leaders are the people you would call because you trust them, because they get stuff done, because they make things happen--on their own and through other people. In Econ , you learn about supply demand curves and equilibrium, and in the last class the professor says, "Those are all simplified explanations. The real world is a lot messier. Yet many businesses still operate as if those assumptions are true. Traditional hierarchies ignore all the real-world, messy stuff. A meritocracy recognizes the subjective--and it recognizes people who are true contributors and leaders. Check out other articles in this series:

NATURAL HIERARCHIES pdf

2: Twenty signs that you have hired the wrong SSAS consultant

Natural Hierarchy is a unique Rotfeather Talisman Rotfeather Talisman Talisman Tier: 3 ()% increased Damage Corrupted Death met with the First Ones and demanded they hunt with mercy. For while pain might delight the mind.

Hierarchies have their own special vocabulary. These terms are easiest to understand when a hierarchy is diagrammed see below. In an organizational context, the following terms are often used related to hierarchies: Multiple hierarchies are possible per dimension taxonomy or Classification-system , in which selected levels of the dimension are omitted to flatten the structure Terms about Placement Hierarch , the apex of the hierarchy, consisting of one single orphan object or member in the top level of a dimension. The root of an inverted-tree structure Member , a member or node in any level of a hierarchy in a dimension to which superior and subordinate members are attached Orphan , a member in any level of a dimension without a parent member. Often the apex of a disconnected branch. Orphans can be grafted back into the hierarchy by creating a relationship interaction with a parent in the immediately superior level Leaf , a member in any level of a dimension without subordinates in the hierarchy Neighbour: Auburn In a mathematical context in graph theory , the general terminology used is different. Most hierarchies use a more specific vocabulary pertaining to their subject, but the idea behind them is the same. For example, with data structures , objects are known as nodes , superiors are called parents and subordinates are called children. Degree of branching[edit] Degree of branching refers to the number of direct subordinates or children an object has in graph theory, equivalent to the number of other vertices connected to via outgoing arcs, in a directed graph a node has. Hierarchies can be categorized based on the "maximum degree", the highest degree present in the system as a whole. Categorization in this way yields two broad classes: In a linear hierarchy, the maximum degree is 1. Note that this is referring to the objects and not the levels; every hierarchy has this property with respect to levels, but normally each level can have an infinite number of objects. An example of a linear hierarchy is the hierarchy of life. In a branching hierarchy, one or more objects has a degree of 2 or more and therefore the minimum degree is 2 or higher. The broad category of branching hierarchies can be further subdivided based on the degree. A flat hierarchy is a branching hierarchy in which the maximum degree approaches infinity, i. Therefore, a flat hierarchy is often not viewed as a hierarchy at all. For example, diamonds and graphite are flat hierarchies of numerous carbon atoms which can be further decomposed into subatomic particles. An overlapping hierarchy is a branching hierarchy in which at least one object has two parent objects. History of the term[edit] Possibly the first use of the English word "hierarchy" cited by the Oxford English Dictionary was in , when it was used in reference to the three orders of three angels as depicted by Pseudo-Dionysius the Areopagite 5thâ€"6th centuries. Since hierarchical churches, such as the Roman Catholic see Catholic Church hierarchy and Eastern Orthodox churches, had tables of organization that were "hierarchical" in the modern sense of the word traditionally with God as the pinnacle or head of the hierarchy , the term came to refer to similar organizational methods in secular settings. This is an example of a hierarchy visualized with a triangle diagram. For example, the few Directors of a company could be at the apex , and the base could be thousands of people who have no subordinates. An example of a triangle diagram appears to the right. An organizational chart is the diagram of a hierarchy within an organization , and is depicted in tree form below. Examples include fractal maps, TreeMaps and Radial Trees. Visual hierarchy[edit] In the design field, mainly graphic design, successful layouts and formatting of the content on documents are heavily dependent on the rules of visual hierarchy. Visual hierarchy is also important for proper organization of files on computers. An example of visually representing hierarchy is through the Nested clusters. The Nested clusters represents hierarchical relationships by using layers of information. The child element is within the parent element, such as in a Venn diagram. This structure of representing hierarchy is most effective in representing simple relationships. For example, when directing someone to open a file on a computer desktop, one may first direct them towards the main folder, then the subfolders within the main folder. They will keep opening files within the folders until the designated file is located. For more complicated hierarchies, the stair structure represents hierarchical relationships through the use of visual stacking. Visually imagine the top of a downward staircase beginning at the left and descending on the right. The child elements are towards the bottom of the stairs and the parent elements are at the top. This structure is effective when representing more complicated hierarchies where steps are not placed in obvious sequences. Further steps are concealed unless all of the steps are revealed in sequence. In the computer desktop example, a file that is being sought after can only be found once another file is opened. The link for the desired file is within another document. All the steps must be completed until the final destination is reached. In plain English, a hierarchy can be thought of as a set in which: The first requirement is also interpreted to mean that a hierarchy can have no circular relationships ; the association between two objects is always transitive. The second requirement asserts that a hierarchy must have a leader or root that is common to all of the objects. Hierarchy mathematics Mathematically, in its most general form, a hierarchy is a partially ordered set or poset. Within this system, each element shares a particular unambiguous property. Objects with the same property value are grouped together, and each of those resulting levels is referred to as a class. Operations such as addition, subtraction, multiplication and division are often performed in a certain sequence or order. Usually, addition and subtraction are performed after multiplication and division has already been applied to a problem. The use of parenthesis is also a representation of hierarchy, for they show which operation is to be done prior to the following ones. In this problem, typically one would multiply 5 by 7 first, based on the rules of mathematical hierarchy. But when the parentheses are placed, one will know to do the operations within the parentheses first before continuing on with the problem. These rules are largely dominant in algebraic problems, ones that include several steps in order to solve. The use of hierarchy in mathematics is beneficial in order to quickly and efficiently solve a problem without having to go through the process of slowly dissecting the problem. Most of these rules are now known as the proper way into solving certain equations. Nested hierarchy[edit] Matryoshka dolls , also known as nesting dolls or Russian dolls. Each doll is encompassed inside another until the smallest one is reached. This is the concept of nesting. When the concept is applied to sets , the resulting ordering is a nested hierarchy. A nested hierarchy or inclusion hierarchy is a hierarchical ordering of nested sets. Each doll is encompassed by another doll, all the way to the outer doll. The outer doll holds all of the inner dolls, the next outer doll holds all the remaining inner dolls, and so on. Matryoshkas represent a nested hierarchy where each level contains only one object, i. The general concept is both demonstrated and mathematically formulated in the following example:

NATURAL HIERARCHIES pdf

3: Attribute Relationships | Microsoft Docs

Last week I wrote an article about Natural Hierarchies in Power Pivot and Tabular where I show why it is important creating natural hierarchies for performance reasons. I'd like to spend a few words about the reasons why this is so important, so you might read the article in advance before moving.

All dimensions have parent child hierarchies. There is absolutely no excuse for this. Use parent-child hierarchies only when you have no other option. The data source view has named queries with a lot of transformations, name changes of columns, derived tables, sub queries and unions in the same mess. The ETL system has moved in to the cubes and the source data warehouse If there is one is not properly structured. It is a clear indication of that the back end data warehouse is a mess. The dimension sources consist of 5 to 15 views joined together and have several layers of views on top of each other, the last with different naming standards. This makes it very hard to follow the ETL flow back to the source tables. The dimension keys are text based, not integers. All dimensions have no attribute relations in the user or natural hierarchies. The natural hierarchies in the source dimension tables or views have undetected many-to many relations that will make the distribution of data across dimension members random. This can happen when you have hidden relations in the source table between attributes that are not primary keys. The cube dimensions have a lot of report hierarchies, which is a hierarchy where the attributes do not have a one to many relations from the top level to the subordinate attributes. Color and the size of a product is an example. The cube has no aggregations. The cube has measures in a single measure group. Sometimes dimensions moves in to measures like actual, budget and forecast measures. The cube has rows of MDX code when you copy and paste it in to word and use 8 as the font size. To be fair you can be forced to do this by business requirements but you should warn about it from a maintenance perspective. It is very hard to debug. Measures have null for many dimensions. This is a clear sign that the data model behind the cube is not properly designed. The cube has no check for division by zero errors in calculated measures so users will see ugly error messages in the table cells. The consultants do not know what UAT is, has never seen a test protocol and do not know how to do tests while developing the cube. The consultant thinks that Kimball and Inmon are two soccerteams. The consultants do not know how to write documentation for a cube except for dropping all the MDX scripts in tables in word. Dimensions key errors, while processing have been set to ignore errors and discard records in a production system. This is a sign of no handling of orphan dimension foreign keys in the fact table. SSAS has a great feature to take care of this with missing dimension keys. The cube has customer dimensions, not one customer dimension with natural hierarchies. This is a migration from AS without a review of the cube model despite all the changes that came in SSAS and later. There are no partitions in the cube when the fact table has million records. Presently he is working in the telecom industry in Sweden, with a data warehouse based on SQL Server His blog can be found here:

NATURAL HIERARCHIES pdf

4: Natural and Unnatural Hierarchies in the SSAS Tabular Model

What is the natural hierarchy? When referring this notion, BOL always raise an example of country--state--city, and when talking abut the non-natural hierarchy, gendor-age is the right instance.

Attributes are exposed to end users through attribute hierarchies. A hierarchy defines the relative position of attribute members in a dimension. By default, in UDM, every attribute can form its own attribute hierarchy and dimensions are the containers for these hierarchies. Dimension may also contain user hierarchies which act as navigational paths to help the user in exploring the data from various angles. Now we are going to see here the different types of user-defined hierarchies that can be created based on the relationship of attribute members with each other. Natural Hierarchy â€" When we create a user-defined hierarchy, we define relationship between hierarchy levels. When levels in a hierarchy are linked in a natural relationship like one-to-one or many-to-one, such hierarchies are known as Natural Hierarchies. Like in a Calendar hierarchy, Day level related to the month level, month level to the Quarter and so on. Natural hierarchies are also known as Balanced hierarchy. In a Balanced Hierarchy, each attribute is a member property of the attribute one level down. Such hierarchies are known as Non-Natural or Unbalanced hierarchies. Ragged Hierarchies â€" Are the type of unbalanced hierarchies wherein, the logical parent member of at least one member is not in the level immediately above the member. In this case the hierarchy descends to different levels for different drilldown paths. Take an example of geographic hierarchy. But some countries like Greece do not have provinces thus creating a ragged hierarchy. In dimension table supporting a ragged hierarchy, the logically missing members can be represented in using nulls or empty strings, or they can contain the same value as their parent to serve as a placeholder. Parent-Child Hierarchies â€" When an attribute in a dimension table has the parent attribute which is related using a self-referencing relationship, it is known as Parent-Child Hierarchy. Such hierarchies are constructed from a single parent attribute. Please refer to the DimEmployee dimension schema in the AdventureWorks database. The ParentEmployeeKey column in the table is related with EmployeeKey primary key column with foreign key relationship. This means each record in a table is related with another record in the same table through a parent-child relationship. In this kind of structure the data can be derived using self-join. When we create a dimension out of such table having self-referencing relation set, Analysis Services recognizes the parent-child hierarchy and accordingly build the dimension. Look at the parent-child icon in the attributes pane of the dimension designer. The attribute Employee in nothing but the Employee Key whose Name Column is set as Employee Full Name which is a named calculation defined in the data source view. After compiling the dimension we can use the browser tab to check how the data will be displayed as shown in the following figure.

5: Hierarchy - Wikipedia

A natural (user-defined) hierarchy is where each attribute is a property is a member property of the attribute above it. For example, Date (Year -> Quarter -> Month -> Day, or Country -> State -> City).

This article explains how to create hierarchies avoiding these problems. We have a table with the following columns: Year Quarter 4 values Month 12 values In order to facilitate navigation, you can create a hierarchy with these levels maybe adding the Date as a fifth level. However, the same month appears in different branches of the hierarchy one per year. For example, January is the same string for both and in the following screenshot. This is an example of an unnatural hierarchy. When a value of the level of the hierarchy e. January is not unique for the entire level and is present in more than one branch of the hierarchy, then the identification of that member requires the knowledge of one or more parents of the same hierarchy. In other words, you have to describe the entire path January or January in order to identify a single month. You can define a natural hierarchy by using attribute that have unique values for the entire level of the hierarchy, without requiring the entire path to identify a member of the hierarchy in a correct way. For example, instead of using January, you can use January , January and so on. This requires a higher number of distinct values in the Month column, because you will have 12 values for each year. You can create these attributes using calculated columns: Navigating the natural hierarchy, you can see that each month, each quarter and each semester has a unique name across the hierarchy. Using a natural hierarchy has many advantages: When you only have 12 months, it is hard to combine year and month in order to obtain a single continuous line for more than one year. Any PivotTable in Excel generates queries in MDX, so this optimization is very important also if you are using a relatively small data model in Power Pivot. For example, you can import a few tables from Adventure Works DW: You create a first PivotTable based on this model, putting the Sum of SalesAmount as a measure and the following attributes on the rows of the PivotTable: The performance are very good and you should not wait more than one second for every refresh. This is what you should see just before adding the DimGeography-City attribute to the rows of the PivotTable. The reason for this behavior is that the MDX query generated by Excel requires a heavy computation that consumes a lot of memory. Maybe Microsoft will improve this in the future, but it is a best practice presenting natural hierarchies to MDX, so this is important for Excel PivotTables. You can repeat the same test using the Natural hierarchy: Any unnatural hierarchy might generate bad performance in MDX queries, even with relatively small amount of data. The data model and the Excel PivotTables shown in this article are included in the sample workbook you can download. Learning DAX from scratch?

NATURAL HIERARCHIES pdf

6: Hans-Hermann Hoppe in 10 Great Quotes

The static hierarchy of the pyramid gives way to fluid natural hierarchies, where influence flows to people who have the most expertise, passion or interest. Freed from the rigidity and sluggishness of a command and control structure, Teal organizations can be more responsive and more energized.

SQL Server Analysis Services Azure Analysis Services Dimensions are collections of attributes, which are bound to one or more columns in a table or view in the data source view. Key Attribute Each dimension contains a key attribute. Each attribute bound to one or more columns in a dimension table. The key attribute is the attribute in a dimension that identifies the columns in the dimension main table that are used in foreign key relationships to the fact table. Typically, the key attribute represents the primary key column or columns in the dimension table. You can define a logical primary key on a table in a data source view which has no physical primary key in the underlying data source. When defining key attributes, the Cube Wizard and Dimension Wizard try to use the primary key columns of the dimension table in the data source view. If the dimension table does not have a logical primary key or physical primary key defined, the wizards may not be able to correctly define the key attributes for the dimension. An attribute is always bound to one or more key columns, which determines the members that are contained by the attribute. By default, this is the only column to which an attribute is bound. An attribute can also be bound to one or more additional columns for specific purposes. For more information, see Dimension Attribute Properties Reference. Attribute Hierarchies By default, attribute members are organized into two level hierarchies, consisting of a leaf level and an All level. However, if the IsAggregatable property is set to False, the All level is not created. Attributes can be, and typically are, arranged into user-defined hierarchies that provide the drill-down paths by which users can browse the data in the measure groups to which the attribute is related. In client applications, attributes can be used to provide grouping and constraint information. When attributes are arranged into user-defined hierarchies, you define relationships between hierarchy levels when levels are related in a many-to-one or a one-to-one relationship called a natural relationship. For example, in a Calendar Time hierarchy, a Day level should be related to the Month level, the Month level related to the Quarter level, and so on. Defining relationships between levels in a user-defined hierarchy enables Analysis Services to define more useful aggregations to increase query performance and can also save memory during processing performance, which can be important with large or complex cubes. Attribute Relationships, Star Schemas, and Snowflake Schemas By default, in a star schema, all attributes are directly related to the key attribute, which enables users to browse the facts in the cube based on any attribute hierarchy in the dimension. In a snowflake schema, an attribute is either directly linked to the key attribute if their underlying table is directly linked to the fact table or is indirectly linked by means of the attribute that is bound to the key in the underlying table that links the snowflake table to the directly linked table.

7: not so natural hierarchies in tabular model

So most hierarchies are natural anyway but the most commonly requested hierarchy is 3 or more levels with a middle level as a slowly changing attribute. The scenario is tracking jobs. The job has many attributes which are all static but the Debtor attribute (i.e. who's paying the invoice) can change over the course of the job.

These are not meant to be the top three best practices to follow, but rather three among many very important best practices you should follow. I believe following these three best practices will make a difference in your solution. Create Hierarchies with Attribute Relationships In my opinion, creating natural hierarchies are the single most beneficial thing an SSAS developer can do to improve the performance and usability of a cube. There are several reasons correctly defined user hierarchies are beneficial, but here are a couple of the top reasons. Increased Query Performance Creating attribute relationships between attributes that are included in a user defined hierarchy improve the performance of queries using these attributes for a couple different reasons. The Calendar hierarchy has five levels with the Calendar Year attribute at the top level and the Date attribute at the bottom level of the hierarchy. We can also see the following attribute relationships created to give SSAS an understanding of how the members of these attributes relate to one another. Once created, these relationships give SSAS the ability to understand that a given date member is related to a single month member, a given month member relates to a single quarter, and so on and so forth. This also means that during processing, special indexes are created that map the relationships between each member of each level in the hierarchy. This means that before a query is written, SSAS also knows to which month, quarter, semester, and year each date aggregates. These indexes are only created if you correctly define the attribute relationships between the attributes in the hierarchy. A user defined hierarchy does two things for a user: And 2 the hierarchy organizes the data based on your business requirements. The user can focus on reading their reports, understanding the data, and making business decisions. Partition Measure Groups to Separate Volatile and Static Data Measure groups that are larger than about one million records should be separated into multiple partitions. There are several advantages to separating larger measure groups into multiple partitions. Increased Processing Performance Partitions in a measure group are processed in parallel. This means that a measure group containing three years worth of data separated into one partition for each month will process faster than a measure group with all three years worth of data in a single partition. Your processing strategy for each measure group may vary depending on the amount of data. For instance, a large telecommunication company may collect hundreds of millions of records per day, dictating a more complex and granular partitioning strategy. Also, we should consider which partitions contain data that is changing vs. If our businesses will continually log transaction in the current month, we can partition our data by month effectively separating our volatile data from the static data. This means we only need to process one months worth of data in our cube to pick up the latest changes instead of having to processing the whole measure group. Increased Query Performance By partitioning our data at the month level, certain queries will also perform better. Because we have partitioned our measure group with each month being loaded into a separate partition, this means that the entire measure group does not have to be scanned. Only the partition containing the data for the month in question needs to be queried. Aggregations are typically used for measure groups that are large and take a considerable amount to query. Improved Query Performance Aggregations contain the data of a measure group at a summary level typically higher than the lowest level of data included in the measure group. Aggregations are populated during the processing phase of the partition. You can think of Aggregations as exactly the same as the lowest level of the measure group just at a summary level. This means that when SSAS receives a query that can be answered using an Aggregation, SSAS does not have to spend the additional time required to retrieve the measure group data from the lowest level and roll up the data to the requested level because the Aggregation design already contains the data at the desired level. Beware Too Many Aggregations Because Aggregations are built during the processing phase, this means that for every aggregation you define more time is required to build the aggregations. This also means that it is especially important to only build useful aggregations that are necessary to improve performance. By creating aggregations that are not often used, you can degrade query performance and increase processing time with little to no benefit. Like indexes on a SQL Server table, too many aggregations or the wrong aggregations can actually hurt performance, so make sure you test, test, and then test to ensure your aggregations are helping your query performance. I hope you have found this information useful. Please leave me a comment and let me know what you think. Do you have any ideas on something I left out or should have included? Please let me know!

NATURAL HIERARCHIES pdf

8: natural hierarchy | Data and Analytics with Dustin Ryan

Natural hierarchies are also known as Balanced hierarchy. In a Balanced Hierarchy, each attribute is a member property of the attribute one level down. In other words all branches in the hierarchy descend to the same level and each member's logical parent is the level immediately above the member.

The corresponding Analysis Services dimension has seven attributes: Customer based on CustomerKey, with CustomerName supplying member names Age, Gender, Email, City, Region, Country Relationships representing natural hierarchies are enforced by creating an attribute relationship between the attribute for a level and the attribute for the level below it. For Analysis Services, this specifies a natural relationship and potential aggregation. In the Customer dimension, a natural hierarchy exists for the Country, Region, City, and Customer attributes. The Country attribute as an attribute relationship to the Region attribute. The Region attribute as an attribute relationship to the City attribute. The City attribute as an attribute relationship to the Customer attribute. For navigating data in the cube, you can also create a user-defined hierarchy that does not represent a natural hierarchy in the data which is called an ad hoc or reporting hierarchy. Users do not see any difference in how the two hierarchies behave, although the natural hierarchy benefits from aggregating and indexing structures â€" hidden from the user â€" that account for the natural relationships in the source data. The SourceAttribute property of a level determines which attribute is used to describe the level. The KeyColumns property on the attribute specifies the column in the data source view that supplies the members. The NameColumn property on the attribute can specify a different name column for the members. To define a level in a user-defined hierarchy using SQL Server Data Tools SSDT , the Dimension Designer allows you to select a dimension attribute, a column in a dimension table, or a column from a related table included in the data source view for the cube. For more information about creating user-defined hierarchies, see Create User-Defined Hierarchies. In Analysis Services, an assumption is usually made about the content of members. Leaf members have no descendents and contain data derived from underlying data sources. Nonleaf members have descendents and contain data derived from aggregations performed on child members. In aggregated levels, members are based on aggregations of subordinate levels. Therefore, when the IsAggregatable property is set to False on a source attribute for a level, no aggregatable attributes should be added as levels above it. Defining an Attribute Relationship The main constraint when you create an attribute relationship is to make sure that the attribute referred to by the attribute relationship has no more than one value for any member in the attribute to which the attribute relationship belongs. For example, if you define a relationship between a City attribute and a State attribute, each city can only relate to a single state.

9: Natural hierarchies | Learning Change

A flat hierarchy is a branching hierarchy in which the maximum degree approaches infinity, i.e., that has a wide span. Most often, systems intuitively regarded as hierarchical have at most a moderate span.

NATURAL HIERARCHIES pdf

Oracle 12c data warehousing guide Employment and people with mental illness Geoff Waghorn and Chris Lloyd Spotsylvania-See Fredericksburg Cat people Margaret and Michael Korda. Cub scout meeting guide Code of banking practice australia Curious George feeds the animals Nuclear medicine imaging in diagnosis and treatment It All Fits Together Winter Quarter Administrators Guidebook: Gods Story ICT and classroom management Mary Webb Vet behind the ears Statistical analysis using excel 2013 Getting into America The Common Reader Why they are called chopsticks 13./tRelief against Forfeiture/t/t67 5 Finding your niche : making a living with medicinal plants Life of fred cats Crisis in the family The Hindu ruins in the plain of Parambanan Powers, Possessions and Freedom Chinese rugs designed for needlepoint Residential building structural design Science, philosophy, and the history of knowledge : Husserls conception of a life-world and Sellarss mani Numbers and Sense of Time Journal of information technology The Complete Wordstudy Dictionary Has completed the entire core curriculum of the graduate training program Tanni Kents (1920) 3rd International Exhibition of Miniature Textiles Will it fly thomas k mcknight The Conquerors and their Consciences Our American nation. Not without my neighbour The universal mechanism: motion and its transformations. Imperial Guptas and their times Development of American literary criticism Thomas and the Fat Controllers engines King of bollywood Amorous Fiammetta.