



US 20050188268A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2005/0188268 A1**
Verbowski et al. (43) **Pub. Date: Aug. 25, 2005**

(54) **METHOD AND SYSTEM FOR TROUBLESHOOTING A MISCONFIGURATION OF A COMPUTER SYSTEM BASED ON CONFIGURATIONS OF OTHER COMPUTER SYSTEMS**

Related U.S. Application Data

(60) Provisional application No. 60/545,799, filed on Feb. 19, 2004. Provisional application No. 60/547,608, filed on Feb. 24, 2004.

Publication Classification

(51) **Int. Cl.⁷** **G05B 15/00**; G06F 11/00; G06F 15/00; G06F 11/30

(52) **U.S. Cl.** **714/37**; 702/186

(57) **ABSTRACT**

A method and system for identifying a configuration parameter of a "sick" computer system that is at fault for causing an undesired behavior based on analysis of configuration parameters from other computer systems is provided. In one embodiment, a troubleshooting system collects "suspect" values for "suspect" configuration parameters used by a "sick" application when the undesired behavior was exhibited by the sick computer system. The troubleshooting system then compares the suspect values to sample values of the suspect configuration parameters retrieved from sample computer systems. The troubleshooting system uses that comparison to identify one or more suspect configuration parameters that are likely at fault for causing the application to exhibit the undesired behavior.


(75) Inventors: **Chad E. Verbowski**, Redmond, WA (US); **Jiahe Helen Wang**, Issaquah, WA (US); **John C. Platt**, Redmond, WA (US); **Ruyun Zhang**, Beijing (CN); **Yu Chen**, Beijing (CN)

Correspondence Address:
PERKINS COLE LLP/MSFT
P. O. BOX 1247
SEATTLE, WA 98111-1247 (US)

(73) Assignee: **Microsoft Corporation**, Redmond, WA

(21) Appl. No.: **10/918,786**

(22) Filed: **Aug. 13, 2004**






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

Likelihood of being wrong	entry name	current value	most popular value	second most popular value
.8	e2	.htc/content	text/x-comp	not exist
.1	e1	.jpg/content	image/jpg	--
.1	e3	url-visited	--	--

select a value for one or more entries

update entries reset to current

FIG. 8

METHOD AND SYSTEM FOR TROUBLESHOOTING A MISCONFIGURATION OF A COMPUTER SYSTEM BASED ON CONFIGURATIONS OF OTHER COMPUTER SYSTEMS

CROSS-REFERENCE TO RELATED APPLICATION(S)

[0001] This application claims the benefit of U.S. Provisional Application No. 60/545,799, filed on Feb. 19, 2004, entitled "FRIENDS TROUBLESHOOTING NETWORK: TOWARDS PRIVACY-PRESERVING, AUTOMATIC TROUBLE-SHOOTING" and U.S. Provisional Application No. 60/547,608, filed on Feb. 24, 2004, entitled "METHOD AND SYSTEM FOR TROUBLESHOOTING A MISCONFIGURATION OF A COMPUTER SYSTEM BASED ON CONFIGURATIONS OF OTHER COMPUTER SYSTEMS," which are hereby incorporated by reference.

TECHNICAL FIELD

[0002] The described technology relates generally to identifying a configuration parameter whose value is causing an undesired behavior.

BACKGROUND

[0003] An ever-increasing number of applications (i.e., computer software) with various features are available to users of personal computers. Users can tailor the operation of these applications to suit their needs by specifying various configuration parameters. For example, a browser application may have a configuration parameter that provides a URL of a web page that is displayed initially whenever the browser application starts (i.e., "a home page"). The browser application may also have configuration parameters that identify programs to be invoked to process certain types of content (e.g., a "jpeg" file) and that specify passwords to be used when the application connects to various servers. The values of the configuration parameters can be stored in application-specific configuration files such as UNIX resource files or in a central registry such as the Windows® registry file. The application-specific configuration file for an application may have an internal format that is specific to that application. With a central registry, many different applications can share the same configuration parameters. The applications access these files to retrieve the values of their configuration parameters.

[0004] If certain configuration parameters have incorrect values, then the applications may exhibit an undesired behavior. For example, if the value of a home page configuration parameter is not set correctly, then when the browser application starts, it will exhibit an undesired behavior by not displaying a home page or displaying the wrong home page. If a configuration parameter incorrectly indicates a certain text editor should be invoked to process a graphics file, then the undesired behavior will be the incorrect display of the graphics content.

[0005] Because of the complexity of applications and their large number of configuration parameters, it can be very time-consuming to troubleshoot which configuration parameters are at fault for causing an application to exhibit the undesired behavior. Most users of personal computers have difficulty performing this troubleshooting. As a result, users

typically rely on technical support personnel to assist in the troubleshooting. This troubleshooting not only is expensive but also users may experience a significant productivity loss as a result of their inability to effectively use an application that is exhibiting an undesired behavior.

[0006] Typically, technical support personnel use an ad hoc approach to troubleshooting configuration problems. Because some central registries store over 200,000 configuration parameters and some computer systems have over 100,000 files, the personnel using knowledge gained from experiencing similar problems will try to narrow in on the at-fault configuration parameter. This ad hoc approach can take a considerable amount of time and even longer if it is a combination of configuration parameters whose values are incorrect. In some cases, the technical support personnel may compare the values of the configuration parameters to "ideal" values for that application. It can be very difficult to identify the configuration parameters used by an application, and this identification requires application-specific knowledge. Moreover, because of the large number of configuration parameters available and the large number of possible values for each configuration parameter, many of the configuration parameters will have no "ideal" value. Thus, technical support personnel still need to review those values of the application that are different from the ideal values.

[0007] It would be desirable to automatically identify a configuration parameter that is at fault for causing an application to exhibit an undesired behavior. It would also be desirable, after such a configuration parameter is identified, to identify an appropriate value for that configuration parameter.

SUMMARY

[0008] A method and system for identifying a likely cause of a component (e.g., application or hardware device) to exhibit a certain behavior is provided. A system collects values for configuration information (e.g., configuration parameters of an application) that may be causing certain behavior and retrieves values for the configuration information from other occurrences of that component (e.g., other computer systems that host the same application). The collected values may be for all the configuration information or a subset of the configuration information that is actually accessed by the component when the certain behavior was exhibited. The system then performs a statistical analysis over the collected values and the retrieved values to determine which configuration information is likely causing the certain behavior of the component.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 is a block diagram that illustrates components of the troubleshooting system in one embodiment.

[0010] FIG. 2 is a flow diagram that illustrates processing of the tracer function in one embodiment.

[0011] FIG. 3 is a flow diagram that illustrates processing of the tracer function in another embodiment.

[0012] FIG. 4 is a flow diagram that illustrates processing of the troubleshooting engine in one embodiment.

[0013] FIG. 5 is a flow diagram that illustrates processing of the canonicalize component in one embodiment.

[0014] FIG. 6 is a flow diagram that illustrates processing of the identify sample set component in one embodiment.

[0015] FIG. 7 is a flow diagram that illustrates processing of the rank suspects component in one embodiment.

[0016] FIG. 8 is a display page that illustrates the results of the troubleshooting engine.

DETAILED DESCRIPTION

[0017] A method and system for identifying a configuration parameter of a “sick” computer system that is at fault for causing an undesired behavior based on analysis of configuration parameters from other computer systems is provided. In one embodiment, a troubleshooting system collects “suspect” values for “suspect” configuration parameters used by a “sick” application when the undesired behavior was exhibited by the sick computer system. The troubleshooting system then compares the suspect values to sample values of the suspect configuration parameters retrieved from sample computer systems that may be peer computer systems. The troubleshooting system uses that comparison to identify one or more suspect configuration parameters that are likely at fault for causing the application to exhibit the undesired behavior. For example, if a suspect configuration parameter has a value that indicates a different text editor should be used to process a graphics file, but the corresponding sample values for all the sample computer systems indicate a certain picture editor should be used, then that suspect configuration parameter might be at fault for causing a graphics file to be displayed incorrectly. This might be especially true if the values for all other suspect configuration parameters are the same for the sick computer system and the sample computer systems. The troubleshooting system may apply various statistical analysis techniques to the sample values when identifying the suspect configuration parameter that is likely at fault. For example, Bayes’ rule can be applied to derive a probability for each suspect configuration parameter indicating its likelihood of being at fault. Once a likely at-fault suspect configuration parameter has been identified, the troubleshooting system can identify likely correct values based on analysis of the sample values. Thus, the troubleshooting system may be considered to use the “typical” sample value as the “ideal” value for an at-fault configuration parameter. In this way, the troubleshooting system can quickly and automatically identify the configuration parameter that is likely at fault and propose a sample value that may be correct.

[0018] In one embodiment, the troubleshooting system identifies the suspect configuration parameters by monitoring the execution of the application when it exhibits the undesired behavior. A tracer program may be used to trace the application’s access to the configuration parameters that may be stored in a configuration file for the application or in a central registry. The tracer program logs the configuration parameters accessed by the application and their corresponding values. The logged configuration parameters are considered the “suspect” configuration parameters because they are the ones that the application used when it exhibited the undesired behavior. The tracer program may execute on the sick computer system and then provide its log to a troubleshooting computer system for identification of the suspect configuration parameter that is at fault.

[0019] The troubleshooting system can retrieve the sample values for configuration parameters from the sample com-

puter systems before or after the suspect values for the sick computer system are collected. The sample computer systems may be “peer” computer systems in the sense that they are related in some way. For example, the peer computer systems can be the set of user computer systems on the local area network of a company or a subset that shares a common characteristic (e.g., used by sales representatives). If the sample values are retrieved before the suspect values are collected, the troubleshooting system may store the sample values in a configuration database. The configuration database may have, for each peer computer system, an entry that contains the value of each configuration parameter retrieved from that computer system. In addition, the troubleshooting system may maintain an index of which peer computer systems have which applications installed. The troubleshooting system may use the index to rapidly select as the sample computer systems those peer computer systems that have the sick application installed. The troubleshooting system uses the sample values of the suspect configuration parameters in its analysis. When the sample values are retrieved before the suspect values are collected, the configuration parameter values for as many different applications that can be troubleshot need to be retrieved and stored. In addition, these sample values may need to be retrieved periodically to reflect the then-current configuration of the sample computer systems.

[0020] If the sample values are retrieved after the suspect values are collected, then the troubleshooting system can limit the retrieval of values to those of the suspect configuration parameters as identified by the tracer program. Because the retrieval is performed at the time of the analysis, the sample values represent the then-current values of the suspect configuration parameters. To speed up the retrieval of the sample values, the troubleshooting system may maintain an index of which peer computer systems have the sick application installed. The troubleshooting system can use the index to identify those peer computer systems that have the sick application installed and retrieve the values for the suspect configuration parameters from only those computer systems.

[0021] In one embodiment, the troubleshooting system may preprocess various configuration parameters to place them in a canonical or normal form. Some configuration parameters may have names or values that are system-specific (e.g., different for every peer computer system). Such system-specific configuration parameters may include password parameters, user name/ID parameters, machine name parameters, device ID parameters, and so on. The troubleshooting system may set the value of each of these parameters to a canonical form, such as “username” for a user name parameter. Alternatively, the troubleshooting system may simply disregard the system-specific configuration parameters when performing its analysis. Other configuration parameters may have values that need to be normalized. For example, a configuration parameter with a URL as its value may be case-insensitive, in which situation the troubleshooting system may normalize values to a standard case to assist in identifying which values are the same (i.e., reference the same resource). As another example, the value of a number in single quotes (e.g., ‘1’) and in double quotes (“1”) may represent the same behavior to the application. In such a situation, the troubleshooting system may replace all double quotes with single quotes to normalize the values.

The troubleshooting system may place both the suspect values and the sample values in a canonical or normal form.

[0022] One skilled in the art will appreciate that many different mathematical analysis techniques may be used to identify the at-fault configuration parameter. Those techniques may include a nearest neighbor analysis, a Bayes net, a neural network, a decision tree, a support vector machine, and so on. In one embodiment, the troubleshooting system calculates a probability that each suspect configuration parameter is at fault using Bayes' rule as described below in more detail.

[0023] FIG. 1 is a block diagram that illustrates components of the troubleshooting system in one embodiment. The troubleshooting system 110 is connected to computer systems 120 via a communications link 130. The troubleshooting system includes a troubleshooting engine 111, a canonicalize component 112, an identify sample set component 113, and a rank suspects component 114. The troubleshooting system also includes a configuration database 115. Each peer computer system includes one or more applications 121, a configuration data store 122, and a tracer program 123. In operation, when a sick application is identified, the tracer program is run on that computer system to log the suspect configuration parameters that the application accesses from the configuration data store. The tracer program may then provide that log to the troubleshooting system via the communications link 130. When the troubleshooting system receives the suspect configuration parameters and their values, it invokes the troubleshooting engine. The troubleshooting engine invokes the canonicalize component to place the values of system-specific configuration parameters in a canonical form and the values of other configuration parameters in a normal form. The troubleshooting engine then invokes the identify sample set component. The identify sample set component identifies the peer computer systems whose values for the suspect configuration parameters are to be used in identifying the at-fault suspect configuration parameter. The identify sample set component may use an index to identify which peer computer systems have the sick application installed

devices), output devices (e.g., display devices), and storage devices (e.g., disk drives). The memory and storage devices are computer-readable media that may contain instructions that implement the troubleshooting system. In addition, data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communications link. Various communications links may be used, such as the Internet, a local area network, a wide area network, or a point-to-point dial-up connection.

[0025] FIG. 1 illustrates an example of a suitable operating environment in which the troubleshooting system may be implemented. The operating environment is only one example of a suitable operating environment and is not intended to suggest any limitation as to the scope of use or functionality of the troubleshooting system. Other well-known computing systems, environments, and configurations that may be suitable for use include personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0026] The troubleshooting system may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, and so on that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments. The term "application" refers to any type of executing software such as drivers, kernel-mode code, operating systems, system utilities, web servers, database servers, and so on. The functions of the troubleshooting system can be performed by each peer computer system or by a central troubleshooting computer system, or distributed between the peer computer systems and a central troubleshooting system.

TABLE 1

Name	Suspect	Sample1	Sample2	Sample3	Sample4	Sample5
.jpg/contentType	image/jpeg	image/jpeg	image/jpeg	image/jpeg	image/jpeg	image/jpeg
.htc/contentType	null	text/x-comp	text/x-comp	text/x-comp	text/x-comp	text/x-comp
url-visited	yahoo	hotmail	Nytimes	SFGate	google	friendster

and are thus eligible to be a sample computer system. As described above, the identify sample set component may retrieve the sample values from the configuration database or may dynamically retrieve the sample values from some sample computer systems. After the sample values are retrieved, the troubleshooting engine invokes the rank suspects component. The rank suspects component performs a Bayes' rule analysis to calculate a probability that each suspect configuration parameter is at fault and then ranks the suspect list based on this probability. The top-ranking configuration parameter is most likely at fault.

[0024] The computing devices on which the troubleshooting system may be implemented include a central processing unit, memory, input devices (e.g., keyboard and pointing

[0027] Table 1 illustrates an example of a troubleshooting case. The table shows suspect values and sample values for various suspect configuration parameters. The tracer program logged the suspect configuration parameters and their suspect values when the application exhibited an undesired behavior. In this example, the tracer program identified three configuration parameters as being suspect (e.g., ".jpg/contentType"). The troubleshooting system retrieved sample values from five sample computer systems for each suspect configuration parameter. Analysis of the sample values indicates that the most likely at-fault configuration parameter is the second configuration parameter (i.e., ".htc/contentType") in the table. All the sample values for that configuration parameter are the same (i.e., "textx-comp"),

but different from the suspect value (“null”). Therefore, it might be assumed that the suspect value of the second configuration parameter is incorrect. Since the suspect value and all the sample values of the first configuration parameter are the same, the suspect value is probably correct. Moreover, because the suspect value and all the sample values for the third configuration parameter are completely different, it would be difficult to assess the correctness of the suspect value from the sample values. In addition to identifying that the second configuration parameter may be at fault, the correct value can be inferred as being the common sample value (i.e., “textx-comp”).

[0028] Table 1 illustrates two different types of configuration parameters: configuration state and operational state. The first and second parameters are configuration states, and the third parameter is an operational state. Operational state parameters have values that are typically set by the application itself and include timestamps, usage counts, caches, seeds for random number generators, window positions, most recently used related information, and so on. The configuration state parameters can generally be set by users to control the operation of an application and are thus more likely than operational state parameters to be at fault for causing an undesired behavior of an application.

TABLE 2

N	Number of sample computer systems
t	Number of suspect configuration parameters
i	Index of the suspect configuration parameters (from 1 to t)
V	Value of a suspect configuration parameter
c	Number of possible values for a suspect configuration parameter
m	Number of sample values that match the suspect value
P(S)	Prior probability that a suspect configuration parameter is at fault (i.e., “S” for sick)
P(H)	Prior probability that a suspect configuration parameter is not at fault (i.e., “H” for healthy)
P(S V)	Probability that a suspect configuration parameter that is at fault has the value V
P(V S)	Probability that a suspect configuration parameter with a value V is at fault
P(V H)	Probability that a suspect configuration parameter with the value V is not at fault

[0029] In one embodiment, the troubleshooting system uses Bayes’ rule in calculating a probability that each suspect configuration parameter is at fault. Bayes’ rule can be stated as follows:

$$P(S|V) = \frac{P(V|S)P(S)}{P(V|S)P(S) + P(V|H)P(H)} \quad (1)$$

[0030] Table 2 contains a description of the terms used in the equation. To simplify the equations, the index i for the suspect configuration parameters has been omitted from these equations (e.g., $P_i(S|V_i)$ is represented as $P(S|V)$). If there is only one at-fault suspect configuration parameter and all the suspect configuration parameters have an equal prior probability of being at fault, then

$$P(S) = \frac{1}{t} \quad (2)$$

-continued

$$P(H) = 1 - \frac{1}{t} \quad (3)$$

[0031] If all the possible values of a suspect configuration parameter have an equal prior probability of being at fault, then

$$P(V|S) = \frac{1}{c} \quad (4)$$

[0032] As an example, if a configuration parameter can have the values of red, green, or blue, then it has three unique sample values. The troubleshooting system can calculate the cardinality by counting the number of unique sample values for a suspect configuration parameter (including a null value if that occurs) and adding one to account for all possible values that do not occur in the sample set. In this example, the cardinality is 4 (i.e., 3+1). Thus, the prior probability of Equation 4 is $\frac{1}{4}$.

[0033] The probability that a suspect value of a suspect configuration parameter is not at fault can be derived from the sample values. In particular, the probability can be assumed to be the number of times the suspect value occurs in the sample values divided by the number of values. The maximum likelihood of this probability is given by the following equation:

$$P(V|H) = \frac{m}{N} \quad (5)$$

[0034] Substituting equations (2)-(5) into equation (1) results in the following equation:

$$P(S|V) = \frac{N}{N + cm(t-1)} \quad (6)$$

[0035] The maximum likelihood, however, may have undesirable properties when the number of sample computer systems is limited. For example, when there are no matching values to V in the sample values, then $m=0$ and $P(S|V)=1$, which expresses complete certainty that is unjustified. For example, in Table 1, the maximum likelihood would indicate that the second and third configuration parameters are both at fault with complete and equal confidence. Bayesian estimation of probabilities may be more appropriate when the number of samples is small. Bayesian estimation uses a prior estimate of $P(V|H)$ before the sample values are examined and produces a posterior estimate of $P(V|H)$ after the sample values are examined. Therefore, $P(V|H)$ is never 0 or 1.

[0036] If it is assumed that $P(V|H)$ is multinomial over all possible values V, the multinomial has parameters p_j . Each p_j is the probability that the value V_j occurs and

$$\sum p_j = 1 \quad (7)$$

[0037] The probabilities have prior and posterior values p_j that can be represented in a Dirichlet distribution. Dirichlet distributions are conjugate to multinomials. That is, combining observations from a multinomial with a prior Dirichlet yields a posterior Dirichlet.

[0038] Dirichlet distributions are characterized by a count vector n_j , which corresponds to the number of possible counts for each value V_j . To perform Bayesian estimation of $P(V|H)$, a prior set of counts n'_j that reflects a prior belief about the likelihood of various values V_j is used. The count m_j for each unique value from the sample values can be collected and the mean of the posterior Dirichlet yields the posterior estimate of the following equation:

$$P(V_j | H) = \frac{m_j + n'_j}{N + \sum_j n'_j} \quad (8)$$

[0039] Since this probability is only needed for the suspect value, m_j can be replaced with m , the number of samples that matches the suspect entry. Furthermore, if it is assumed that all values V_j have the same a priori probability, then n'_j can be replaced with some value n and the sum $\sum_j n'_j$ can be replaced with cn . These assumptions result in the following equation:

$$P(S | V) = \frac{N + cm}{N + cnt + cm(t - 1)} \quad (9)$$

[0040] The parameter n is proportional to the number of sample values that is required to overwhelm the prior and to move the estimated $P(V|H)$ probabilities away from the estimated probability that p_j is $1/c$. The parameter n indicates the strength of the prior. In other words, the higher n is the less confidence is placed in the sample values. A higher n leads to a stronger prior, which requires more evidence (i.e., samples) to change the posterior. If $n=1$ is selected for a prior, which is equivalent to a flat prior, then all multinomial values p_j are equally likely a priori. This is known as an “uninformative” prior.

[0041] FIGS. 2 and 3 are flow diagrams that illustrate processing of a configuration parameter retrieving function and a tracer function in one embodiment. FIG. 2 is a flow diagram that illustrates processing of a function that is invoked by the sick application to retrieve a value for a configuration parameter from the configuration data store. The function is passed an identifier of the configuration parameter to be retrieved. In block 201, the function locates in the configuration data store the entry for the identified configuration parameter. In decision block 202, if the entry is located, then the function continues at block 203, else the function continues at block 204. In block 203, the function retrieves the value for the identified configuration parameter. In block 204, the function invokes a tracer function passing the identification of the configuration parameter and the retrieved value. If the configuration parameter was not located, then the value may be null. The function then returns the retrieved value to the sick application. This function may be implemented as part of an application programming interface of the configuration data store.

[0042] FIG. 3 is a flow diagram that illustrates processing of the tracer function in one embodiment. The function is passed an identifier of the configuration parameter and its value. The function logs the passed identifier and value. In block 301, the function increments the number of configuration parameters that have been retrieved. In one embodiment, the function may log only one value for each configuration parameter. If the sick application accesses the same configuration parameter multiple times, this function may log the last configuration parameter and value pair. In such a case, this function would need to check for duplicate accesses to a configuration parameter and increment the number of configuration parameters only when a duplicate is not found. In block 302, the function stores the suspect configuration parameter identifier. In block 303, the function stores the suspect configuration parameter value and then returns.

[0043] FIGS. 4-7 are flow diagrams that illustrate processing of the troubleshooting system in one embodiment. FIG. 4 is a flow diagram that illustrates processing of the troubleshooting engine in one embodiment. The troubleshooting engine receives the values for the suspect configuration parameters from the sick computer system hosting the sick application. In block 401, the engine invokes the canonicalize component passing the suspect configuration parameter identifiers and their values. In block 402, the engine invokes the identify sample set component passing an identifier of the sick application and the suspect configuration parameter identifiers and receiving the sample values in return. In block 403, the engine invokes the rank suspects component passing the suspect values and sample values and receiving a ranking of the suspect values based on the probability that each suspect value is at fault. The rank suspects component calculates the probability based on Equation 9. The engine then completes.

[0044] FIG. 5 is a flow diagram that illustrates processing of the canonicalize component in one embodiment. The component is passed the suspect configuration parameter identifiers and their values. The component canonicalizes the suspect values of the system-specific configuration parameters. The component can also be adapted to normalize the suspect values. In block 501, the component initializes an index into an array of suspect configuration parameters. In blocks 502-507, the component loops processing each suspect configuration parameter. In block 502, the component selects the next suspect configuration parameter. In block 503, if all the suspect configuration parameters have already been selected, then the component returns, else the component continues at block 504. In block 504, the component selects the next system-specific configuration parameter. In decision block 505, if all the system-specific configuration parameters have already been selected, then the component loops to block 502 to select the next suspect configuration parameter, else the component continues at block 506. In block 506, if the selected system-specific configuration parameter is the same as the selected suspect configuration parameter, then the component continues at block 507, else the component loops to block 504 to select the next system-specific configuration parameter. In block 507, the component sets the value of the selected suspect configuration parameter to the canonical value for the selected system-specific configuration parameter and then loops to block 502 to select the next suspect configuration parameter.

[0045] FIG. 6 is a flow diagram that illustrates processing of the identify sample set component in one embodiment. The component is passed an indication of the sick application and the suspect configuration parameters. In this embodiment, the component accesses the sample configuration database containing sample values that have been previously retrieved from peer computer systems. As part of the retrieval process, the troubleshooting system may have placed the sample values in a canonical or normal form, for example, using the canonicalize component of FIG. 5. In block 601, the component initializes an index of the sample computer systems. In blocks 602-607, the component loops identifying sample computer systems until a designated number have been identified. In block 602, the component selects the next peer computer system from the sample configuration database. In decision block 603, if all the peer computer systems have already been selected, then the component returns an indication that not enough sample computer systems have been found, else the component continues at block 604. In decision block 604, if the sick application is installed on the selected peer computer system, then it is a sample computer system and the component continues at block 605, else the component loops to block 602 to select the next peer computer system. In block 605, the component increments the index of the sample computer systems that have been identified. In block 606, the component retrieves the sample value for each suspect configuration parameter for the selected peer computer system. In decision block 607, if enough sample computer systems have been identified, then the component returns the sample values, else the component loops to block 602 to select the next peer computer system.

[0046] FIG. 7 is a flow diagram that illustrates processing of the rank suspects component in one embodiment. The component is passed the suspect values and the sample values. The component returns a probability that each suspect configuration parameter is at fault and may return an indication of a possible correct value for the configuration parameters with the highest probabilities. In blocks 701-708, the component loops selecting each suspect configuration parameter and calculating the probability that it is at fault. In block 701, the component selects the next suspect configuration parameter i . In block 702, if all the suspect configuration parameters have already been selected, then the component returns, else the component continues at block 703. In block 703, the component counts the number of different values c_i in the samples for the selected suspect configuration parameter. The troubleshooting system may calculate and store this count when the configuration parameters are retrieved from the peer computer system. In block 704, the component sets the number of suspect configuration parameters t , which may be received from the sick computer system. In block 705, the component sets the number of samples N , which may be derived from the identify sample set component. In block 706, the component sets the number of sample values of the selected suspect configuration parameter that matches the suspect value m_{ij} . The troubleshooting system may also calculate and store the number of occurrences of each sample value for each configuration parameter retrieved from the peer computer systems. In block 707, the component sets an indicator of the confidence n that the system has in the samples. In block 708, the component calculates the probability for the suspect con-

figuration parameter $P(S|V)$ using Equation 9 and loops to block 701 to select the next suspect configuration parameter.

[0047] FIG. 8 is a display page that illustrates the results of the troubleshooting engine. This display page may be displayed by the sick computer system to assist a user in correcting the at-fault configuration parameter. The display page includes a table 801 that has a row for each suspect configuration parameter. Each row includes the identification of the suspect configuration parameter, the probability that that suspect configuration parameter is at fault, the suspect value of that suspect configuration parameter, and sample values that may be correct. Those sample values may be the most popular of the sample values. The display page also includes button 802 for updating the value of a suspect configuration parameter. When a user selects a sample value for one or more suspect configuration parameters and then selects the update button, the system changes the values of the suspect configuration parameters in the configuration data store on the sick computer system to the selected values. The user can then execute the application to determine whether it still exhibits the undesired behavior. If so, the user can select the reset button 803 to reset the suspect configuration parameters to their suspect values. The user can repeat this process for various combinations of suspect configuration parameters and their values until the application no longer exhibits the undesired behavior.

[0048] One skilled in the art will appreciate that although specific embodiments of the troubleshooting system have been described for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. The troubleshooting system can also be used to identify hardware configuration problems. For example, if the peer computer systems include special-purpose signal processing hardware with configuration parameters, then suspect values of those configuration parameters for a computer system that is exhibiting an undesired behavior can be compared to sample values as described above. More generally, the troubleshooting system can be used in an environment with multiple configuration parameters, such as settings for television set-top boxes, cell phones, automobiles, and so on. The techniques of the troubleshooting system can also be used to identify information generally that may be causing a certain behavior, whether desired or undesired. For example, the execution of an application may be adversely affected by the overall configuration of the computer system on which it is executing. As an example, the undesired behavior may be caused by a missing operating system component, an outdated driver, insufficient main memory, interactions with a user, URL parameters, API parameters, and so on. The techniques of the troubleshooting system can be used to analyze such information collected from sample systems to identify the cause of the behavior. The techniques can also be used to identify the cause of a desired behavior. For example, a complex system may have hundreds of "parameters" with many different possible values that may affect its behavior. In such a case, it may be impractical to predict the behavior of each possible combination of parameter values. However, once a desired behavior is identified, the technique can be used to identify the parameters and their values that are likely causing the desired behavior by comparing it to a sample set that is not exhibiting that behavior (or even one that is exhibiting that behavior). The described technology is related to U.S. Provisional Application No. 60/545,799

(Attorney Docket No. 41826.8015 US), entitled "Friends Troubleshooting Network: Towards Privacy-Preserving, Automatic Troubleshooting" and filed on Feb. 19, 2004, and U.S. Provisional Application No. 60/547,607 (Attorney Docket No. 41826.8015 US1), entitled "Method and System for Collecting Information from Computer Systems based on a Trusted Relationship" and filed on Feb. 24, 2004, which are hereby incorporated by reference. These applications describe technology that can be used to retrieve the sample configuration parameter values. Accordingly, the invention is not limited except by the appended claims.

We claim:

1. A method in a computer system for identifying a configuration parameter that is likely causing a component to exhibit an undesired behavior, the method comprising:

collecting suspect values for suspect configuration parameters of the component;

providing sample values for the suspect configuration parameters for sample components; and

comparing the suspect values with the sample values to identify a suspect configuration parameter whose suspect value is likely to be causing the component to exhibit the undesired behavior.

2. The method of claim 1 wherein the comparing includes calculating a probability for each suspect configuration parameter.

3. The method of claim 1 wherein the suspect values are collected during execution of the component in which the undesired behavior was exhibited.

4. The method of claim 1 wherein the providing of values includes retrieving the sample values for the suspect configuration parameters from each sample component.

5. The method of claim 4 wherein the retrieving of the sample values is performed before the collecting of the suspect values wherein the sample values are retrieved for configuration parameters that include the suspect configuration parameters.

6. The method of claim 4 wherein the retrieving of the sample values is performed after collecting the suspect values.

7. The method of claim 6 wherein the sample values are retrieved from sample computer systems that have the component installed.

8. The method of claim 7 wherein an index is maintained that indicates computer systems that have the component installed.

9. The method of claim 1 wherein the collecting includes running a program that intercepts and records accesses by the component to a configuration data store.

10. The method of claim 9 wherein the configuration data store is a registry file.

11. The method of claim 9 wherein the configuration data store is a resource file.

12. The method of claim 1 wherein the values for configuration parameters that are system-specific are placed in a canonical form.

13. The method of claim 1 wherein the values for configuration parameters are placed in a normal form.

14. The method of claim 1 wherein the component is a hardware device.

15. The method of claim 1 including aggregating the sample values prior to collecting the suspect values to facilitate comparing to the suspect values.

16. The method of claim 15 wherein the aggregating includes counting the number of distinct sample values for each suspect configuration parameter.

17. The method of claim 16 wherein the aggregating includes counting the number of sample components that have each distinct value.

18. The method of claim 1 including identifying a likely correct value for the sample configuration parameter that is identified as likely causing the component to exhibit the undesired behavior.

19. The method of claim 18 wherein the likely correct sample value is the value that is most common.

20. The method of claim 18 including changing the suspect value of the suspect configuration parameter that is identified as likely causing the component to exhibit the undesired behavior to the identified likely correct value.

21. The method of claim 18 including presenting to a user an indication of likely correct values for the suspect configuration parameter that is identified as likely causing the component to exhibit the undesired behavior.

22. The method of claim 18 including presenting to a user a ranking of suspect configuration parameters that may be causing the component to exhibit the undesired behavior.

23. A method for storing values for configuration parameters of computer systems, the method comprising:

retrieving from the computer systems sample values for configuration parameters; and

generating aggregate values based on the sample values for the configuration parameters.

24. The method of claim 23 wherein the retrieved sample values and aggregate values are stored in a data store.

25. The method of claim 23 wherein the retrieved sample values are placed in a normal form.

26. The method of claim 23 wherein an aggregate value is generated by counting the number of distinct sample values for each configuration parameter.

27. The method of claim 26 wherein an aggregate value is generated by counting the number of computer systems that have each distinct sample value.

28. The method of claim 23 wherein each configuration parameter is associated with an application and wherein each computer system has one or more applications and including generating an index that maps each application to the computer systems associated with that application.

29. A computer-readable medium containing instructions for controlling a computer system to identify a cause of a component that is exhibiting a certain behavior, by a method comprising:

collecting suspect values for suspect information of an instance of the component;

providing sample values for the suspect information, the sample values being collected from other instances of the component; and

identifying a suspect information whose suspect value is likely causing the certain behavior based on analysis of the sample values.

30. The computer-readable medium of claim 29 wherein the component is an application and the information is a configuration parameter of the application.

31. The computer-readable medium of claim 30 wherein collecting includes executing the instance of the application in which the certain behavior is exhibited and recording accesses to the configuration parameters of the application.

32. The computer-readable medium of claim 31 wherein the behavior is an undesired behavior.

33. The computer-readable medium of claim 30 wherein the analysis includes calculating a probability that each suspect configuration parameter is causing the certain behavior.

34. The computer-readable medium of claim 33 wherein the probability is calculated based on Bayes' rule.

35. The computer-readable medium of claim 30 including changing the suspect value of the identified configuration parameter and executing the instance of the application.

36. The computer-readable medium of claim 29 wherein the analysis includes calculating a probability that each suspect value is causing the certain behavior.

37. A computer system for identifying a configuration parameter of an application that is causing an undesired behavior, comprising:

a data store of suspect values for suspect configuration parameters, the suspect values being collected from an execution of the application when the undesired behavior was exhibited;

a data store of sample values for the suspect parameters, the sample values being collected from other instances of the application; and

means for identifying a suspect parameter whose suspect value is likely causing the undesired behavior based on analysis of the sample values.

38. The computer system of claim 37 wherein the analysis includes calculating a probability that each suspect parameter is causing the undesired behavior.

39. The computer system of claim 38 wherein the probability is calculated based on Bayes' rule.

* * * * *