

COMET: An Application of Model-Based Reasoning to Accounting Systems

*Robert Nado, Melanie Chams,
Jeff Delisio, and Walter Hamscher*

■ An important problem faced by auditors is gauging how much reliance can be placed on the accounting systems that process millions of transactions to produce the numbers summarized in a company's financial statements. Accounting systems contain *internal controls*, procedures designed to detect and correct errors and irregularities that can occur in the processing of transactions. In a complex accounting system, it can be an extremely difficult task for the auditor to anticipate the possible errors that can occur and evaluate the effectiveness of the controls at detecting them. An accurate analysis must take into account the unique features of each company's business processes. To cope with this complexity and variability, the COMET system applies a model-based reasoning approach to the analysis of accounting systems and their controls. An auditor uses COMET to create a hierarchical flowchart model that describes the intended processing of business transactions by an accounting system and the operation of its controls. COMET uses the constructed model to automatically analyze the effectiveness of the controls in detecting potential errors. Price Waterhouse auditors have used COMET on a variety of real audits in several countries around the world.

Auditors have the task of determining whether the financial statements of a company are a fair presentation of the company's financial position. An important problem faced by auditors is gauging how much reliance can be placed on the accounting systems that produce the numbers summarized in the financial statements. Accounting systems contain *internal controls*, procedures designed to detect and correct errors and irregularities that can occur in the processing of transactions. In a complex accounting system, it can be an extremely

difficult task for the auditor to anticipate the possible errors that can occur, determine their downstream effects in the accounting system, and evaluate the effectiveness of the controls at detecting them. An accurate analysis must take into account the unique features of each company's business processes. To cope with this complexity and variability, the COMET system applies a model-based reasoning approach (cf. Hamscher, Console, and de Kleer 1992) to the analysis of accounting systems and their controls.

COMET supports the creation of hierarchical flowcharts that ultimately describe the processing of business transactions in terms of a set of primitive activities for operating on records and a set of controls for detecting and correcting errors that might occur in the processing. Using knowledge of the basic ways in which the primitive activities can fail, COMET finds potential failures that can occur in the accounting system and uses the structure of the flowchart to analyze the impact of these failures on the validity of the accounts. COMET then matches each potential failure to the set of controls capable of detecting it and evaluates the effectiveness of the controls in reducing the risk that the potential failure will go undetected. Finally, COMET ranks the controls with respect to their relative contribution to reducing the risk of undetected failures and selects a subset of key controls whose proper operation should be tested.

Task Description

In the United States, the Securities and Exchange Commission requires a yearly independent audit of the financial statements of public companies. Other countries have simi-

lar requirements. An accounting firm that is engaged to perform an audit of a public company has the task of issuing an opinion on whether the financial statements are a fair characterization of the financial position of the company and follow generally accepted accounting principles. The numbers that appear in the financial statements are typically the accumulated results of thousands, even millions, of detailed financial transactions in which the company has participated over the previous year.

Two main approaches can be taken to assessing the accuracy of financial statements. The *substantive approach* attempts to obtain evidence of the validity of financial statements by examining records of detailed transactions and applying analytic methods to gauge the reasonableness of the reported numbers. By contrast, the *systems-reliant approach* focuses not on verifying the numbers themselves but on assessing the adequacy of the accounting systems that produced the numbers. In taking a systems-reliant approach, an auditor looks at the internal controls that are in place in the accounting systems and evaluates their effectiveness in detecting and correcting errors that can occur in processing transactions.

For example, a company's *purchases and payables system* handles transactions involving the purchase of goods from suppliers. Such a system is designed to receive and record purchase orders; transmit them to suppliers; and ensure that goods are received, payables recorded, and the supplier eventually paid for goods received. In auditing such a system, it is important to focus not so much on the computer system itself but on the business processes that it supports. A business process usually contains both manual and computerized steps and is partially performed by parties outside the company.

Many things can go wrong in a purchases and payables system. For example, an invoice can be received from a supplier for goods that were never ordered or received; the quantity or price of goods listed on the invoice can be incorrect, either because of an error at the supplier or because of an operator error in entering the invoice into the computer system; or a fictional invoice can be entered into the system as part of an attempt to defraud.

To detect and correct such problems, a purchases and payables system should contain a number of internal controls. For example, invoices that have been entered on the computer system should be matched to corresponding purchase orders and records of goods

received, with quantities and prices agreed (although the matching process can be computerized, any discrepancies will generally need to be investigated manually and resolved), access to the computer system for data entry should be restricted to authorized personnel by means of an appropriate security system, and data entry of an invoice should not be performed by the same person who later authorizes or reviews the invoice.

In practice, any given audit will combine elements of both the substantive and systems-reliant approaches, with the relative emphasis dependent on the particular characteristics of the business and its components. With large companies that have complex, computerized accounting systems processing vast numbers of transactions, the systems-reliant approach is becoming increasingly important, both to obtain adequate audit evidence and reduce the cost of the audit. A specialized category of auditor, called a *computerized information systems (CIS) auditor*,¹ brings to bear skills in both accounting and systems analysis to carry out a systems-reliant audit approach.

To take a systems-reliant approach, a CIS auditor must obtain and document an understanding of how an accounting system processes business transactions and what internal controls are in place. In preparing this "model," the auditor can make use of available systems documentation from the client. However, systems documentation generally is not prepared from an audit point of view. It can explain how the system works in great detail but, generally, does not contain adequate information on controls, does not have a business-process focus, and omits the manual components of the business process. The auditor must supplement information obtained from documentation with observation of the system in operation and interviews with key personnel.

In determining the effectiveness of controls, it is important to distinguish the role of a control in the design of an accounting system from its performance in practice. By analyzing the processes and data flows of an accounting system, an auditor attempts to determine those controls that play key roles in the prevention and detection of errors that might affect the validity of the financial statements. To obtain sufficient comfort that the system is actually operating as designed, the key controls need to be tested to ensure that they are being performed properly.

For complex accounting systems, a thorough and accurate controls evaluation is almost impossible to perform efficiently with-

out some form of computer-based support. There are many different possible sources of error, some of which can be overlooked. It is extremely difficult to manually trace the effects of possible errors through the transaction processing to determine whether they are significant to the audit. There can be redundancy in the coverage of errors by controls, but detailed analysis is required to determine this redundancy with confidence. Because systems evolve rapidly, it is costly to determine the impact of system changes on controls effectiveness. Most importantly, human fallibility in the face of complex systems can lead to costly consequences.

Prior to COMET, CIS auditors used a combination of flowcharting software and controls checklist software in their evaluation of controls. Commercial flowcharting software can be used to document major activities carried out in an accounting system, but the result is not in a form that allows automated analysis. Checklist software is populated with libraries of controls that could be expected to be found in a client's system to address the major areas of risk. Although different libraries of controls can be developed for the major components of generic accounting systems, as well as for different accounting software packages, it is difficult to tailor checklist software to reflect the varying characteristics of different industry sectors and the idiosyncratic aspects of a particular client's implementation. Furthermore, controls checklist software takes no advantage of the information captured in flowchart documentation.

The development of COMET was motivated by the intuition that an accounting system can hierarchically be decomposed into a structure that bottoms out in instances of a small set of primitive types of actions for processing records and implementing internal control. Provided that the behavior of the primitive activity and control types can be characterized suitably, a model-based approach can be taken to the analysis of failures and their detection by internal controls. As a consequence, the auditor can concentrate on developing an accurate model of the accounting system under review, with COMET automating the more burdensome aspects of controls evaluation.

Application Description

Although model-based reasoning has previously been applied to financial domains, the models have generally consisted of equations and constraints representing the relationships between financial and microeconomic quan-

ties (Hamscher 1994; Bridgeland 1990; Hart, Barzilay, and Duda 1986; Bouwman 1983). COMET is novel in its application of a model-based approach to analyzing systems for processing financial records.

Basic Modeling Concepts

Accounting systems process records of business transactions through activities that create, use, alter, and store those records. COMET represents the processing performed by an accounting system as a hierarchically structured flowchart graph. The two most important kinds of node in a COMET flowchart are collection nodes and activity nodes. *Collection nodes* represent repositories of records, which can be in either paper or electronic form. *Activity nodes* are represented hierarchically, starting with nodes representing activities at a high level of abstraction and progressively decomposing them until nodes representing primitive activities are obtained.

Figure 1 shows the top-level flowchart of PURCHASE, a model of a simple purchases and payables accounting system. The top-level flowchart is intended to give a high-level overview of the system, indicating the major activities performed by the system, the relevant general ledger accounts, and important collections of records that are accessed and updated by the processing of a transaction. Activity nodes are distinguished by having a rectangular icon in their lower-left corner. Collection nodes have a trapezoidal icon for paper records and a cylindrical icon for electronic records. Nodes representing general ledger accounts contain a "boxed T" icon. Two kinds of arc connect nodes in the flowchart. The solid arcs represent data flow relationships between activities and collections. The dashed arcs represent precedence relationships between activities; the activity at the tail of a dashed arc must be completed before the activity at the head of the arc can proceed.

The order-processing activity prepares a purchase order, which is sent to a supplier to be filled and also recorded in the orders database. When the supplier fills the order, it sends a goods received note (GRN) and an invoice along with the goods. The goods-receiving activity records the GRN in the GRN database and tries to match it with a corresponding record in the orders database. The purchase-invoice activity records the invoice from the supplier in the invoices database and compares it with the corresponding record on the orders database. If a matching order can be found, the purchase-invoice activity posts a credit to the creditors account

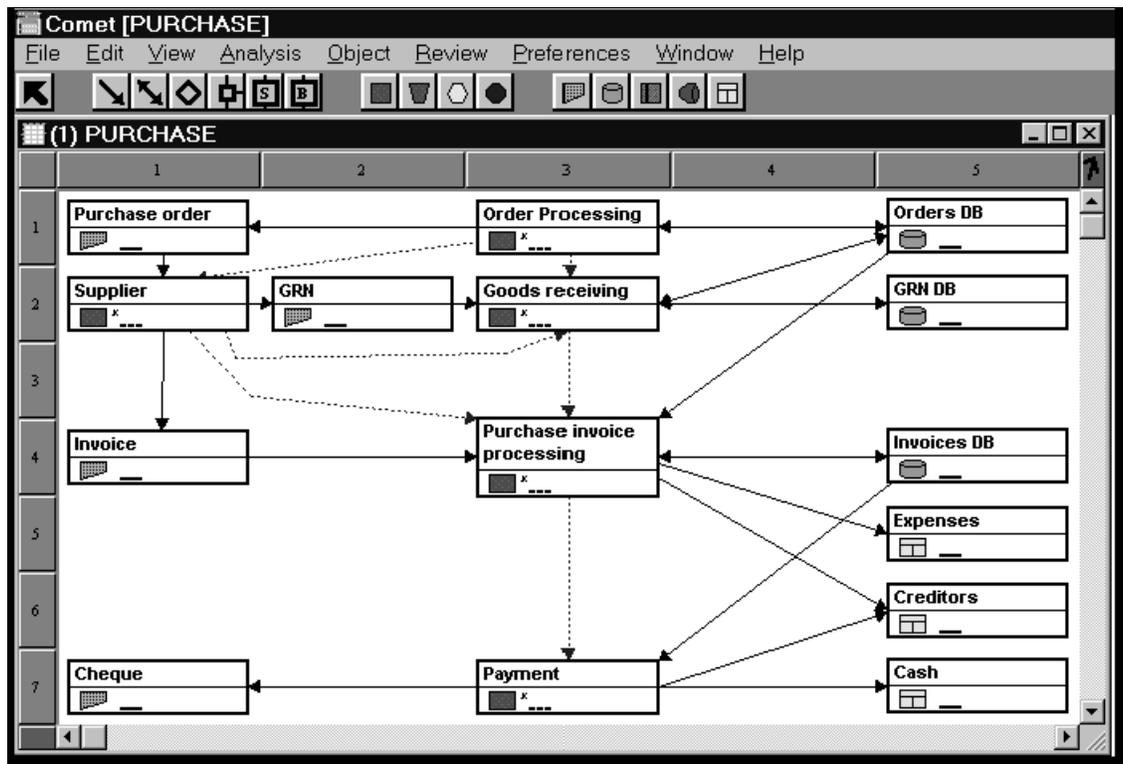


Figure 1. A Top-Level Flowchart.

and a debit to the expenses account. The payment activity periodically extracts invoices that are due for payment, prepares checks for payment to suppliers, debits the creditors account, and credits the cash account.

Because the top-level flowchart of PURCHASE gives a high-level overview of the system, it contains no primitive activities or controls. Each of the top-level activity nodes has a decomposition into a subflowchart that gives more detail about how the activity is performed. Figure 2 shows the flowchart for the decomposition of the payment activity; it can be reached by double clicking on the payment node in the top-level flowchart. The nodes in figure 2 that have dashed light-gray borders are called *reference nodes*; they refer to collections whose primary depiction is elsewhere in the flowchart. When an activity node is decomposed, each collection node to which it is directly connected has a reference node automatically created in the subflowchart. The reference nodes allow the input and output collections of the top-level activity to be referenced by the activities in the subflowchart.

COMET contains a predefined vocabulary of activity and control types, called *verbs*, that are used as a focal point for organizing the knowledge that COMET contains about accounting systems and their controls. Some

verbs, such as *transfer*, *copy*, *create*, *merge*, *find*, *compute*, and *copy-field*, represent typical operations on records that are treated as primitive by COMET. Other verbs, such as *maintaining-data* and *data-entry*, represent processing patterns that are common enough that COMET provides automatic decompositions for nodes using these verbs. For describing internal controls, COMET provides a set of control verbs, including *authorize*, *compare-agree*, *grant-access*, *reconcile*, and *review*. The verb associated with an activity or control node is indicated in the display of that node using a one- or two-letter code inside the icon in the lower-left corner.

Figure 2 contains four primitive activity nodes that have the verbs (1) *extract* (EX), (2) *copy* (CO), (3) *debit* (DR), and (4) *credit* (CR). Figure 2 also contains four control nodes (the nodes with the circle icon) using two different control verbs: (1) *authorize* (AU) and (2) *review* (R). In addition to nodes representing collections, activities, and controls, figure 2 contains three smaller, rectangular nodes, called *selectors*. Selector nodes are used to indicate the fields of records that are accessed or modified by activities. For example, the selector node between the debit creditors activity and the creditors account indicates that the debits involve a field called *Value*.

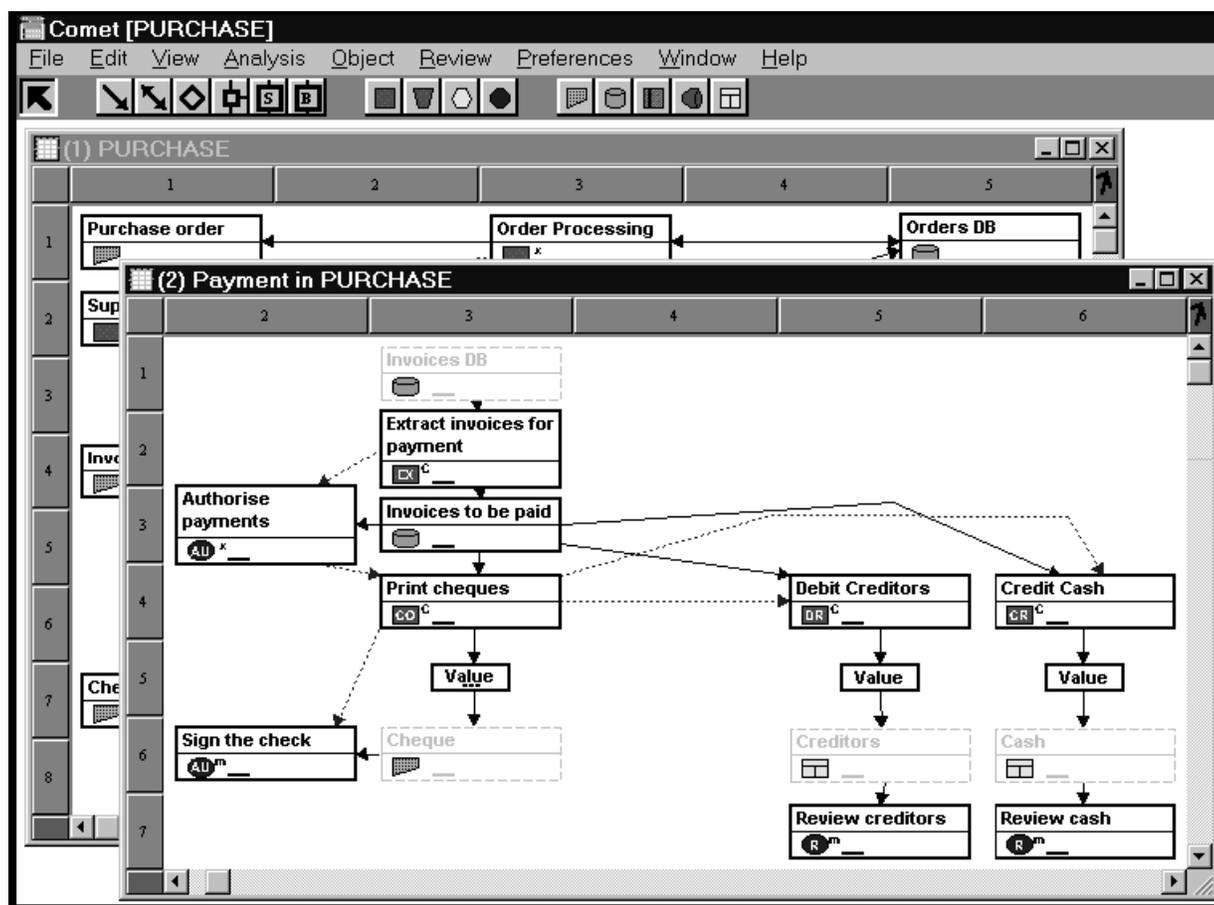


Figure 2. Expansion of the Payment Activity.

Model-Building Support

The analysis performed by COMET depends on the accuracy of the models that it operates on for its validity. Auditors attempt to verify the accuracy of a model by walking through the transaction-processing steps specified in the model, checking for matching steps performed in the modeled system. Ideally, the walk-through is performed by a person not involved in the model preparation. Although COMET cannot ensure that the models constructed by users are, in fact, accurate representations of the modeled accounting systems, COMET incorporates a number of tools to aid in the construction of models that are at least internally consistent and that contain enough detail to support COMET's analysis.

Each type of node has an associated set of declarative constraints on the ways that a node of this type can correctly be connected by arcs to neighboring nodes. For example, a credit activity node must have exactly one input collection and at least one output collec-

tion. Every output collection must be an account. Finally, there must be selector nodes intervening between the credit node and each of its output accounts giving the fields that are posted to the accounts. As the user edits a model, COMET monitors the constraints on each node and draws a red flag on those nodes whose constraints are not satisfied. For any node with a red flag, the user can obtain an explanation of the unsatisfied constraints.

COMET contains a number of review commands for examining the completeness and consistency of a model: finding all nodes with violated syntactic constraints, finding all unexpanded generic activity nodes, finding control nodes that have been described incompletely, finding inconsistencies between the fields read from a collection node and the fields written to it, finding activity nodes that access records from a collection node without having a preceding activity node that creates records on the collection, and finding inconsistencies between the input and output specified for an activity node and its subflowchart.

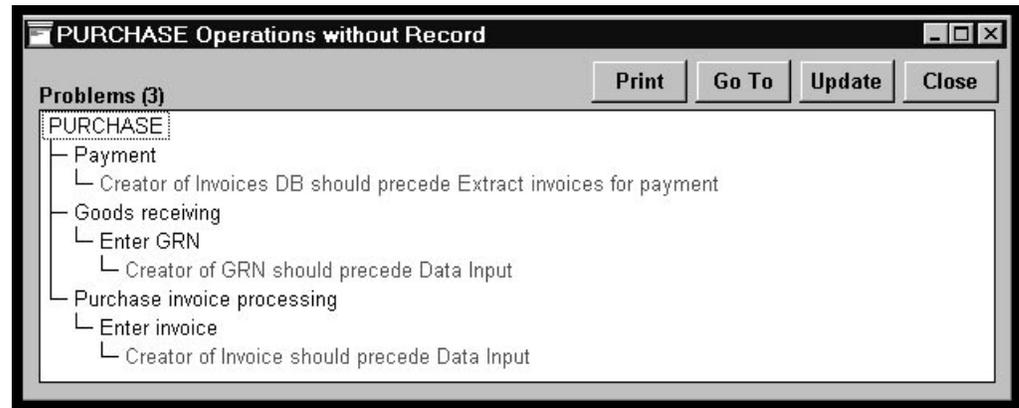


Figure 3. Example Review Dialog Box.

The results of the review commands are presented in the form of dialog boxes that allow convenient navigation to the points where problems occur in the model (cf. figure 3).

Failure Generation and Propagation

COMET categorizes the errors and irregularities that can occur in an accounting system into three broad categories of failure, corresponding to the focus on the processing of records: (1) missing, (2) spurious, and (3) incorrect. A *missing* failure occurs when an activity that should have produced a record as output fails to do so. A *spurious failure* occurs when an activity produces an unauthorized or duplicate record as output. Finally, an *incorrect failure* occurs when an activity produces an incorrect value for a field in a record. An incorrect failure is associated with the name of the affected field. Each primitive activity type has associated with it the categories of failure to which it can give rise.

The first stage of COMET's analysis of a model generates the set of potential failures corresponding to each of the primitive activities in the model. COMET then determines which of the potential failures have audit significance. A failure has audit significance if its downstream effects in the flowchart model could cause any of several types of disagreement between the transactions that actually occurred and the way that they are recorded in the accounts. COMET works backward in the flowchart from the account collections using a few fixed simple rules for the different primitive activity types to determine how failure effects on an output collection can be produced from failure effects on input collections. The result of this stage is to determine for each potential failure the impact that it might have on the validity of each account collection.

When constructed at a level of detail appropriate to the control-evaluation task, a COMET model typically contains on the order of hundreds of primitive activities. Because each of these activities can fail only in a small number of ways, it is a tractable task to enumerate the set of potential failures and determine their effects on the validity of accounts.

Control Evaluation

To evaluate the controls documented for an accounting system, COMET assesses, for each potential failure with audit significance, the likelihood that if it occurs, it will not be detected by any control in the system. We call this likelihood, for a given failure, its *failure-coverage risk*. To determine whether the potential failures are adequately covered by detecting controls, a CIS auditor using the system is required to associate with each account an *allowable risk level*. The allowable risk is the highest level of risk the auditor is willing to accept that any failure that occurs and that is relevant to the account is not detected by any control.

Figure 4 shows a table generated by COMET of those potential failures generated for the PURCHASE model that have audit significance and the failure-coverage risks that have been determined for them. Certain controls in a COMET model can be designated as proposed; *proposed controls* are used to explore the effects of recommending to the client that additional controls be added to the accounting system to address control weaknesses. The failures table in figure 4 contains two columns listing failure-coverage risks in percentage terms. The first column (Prop) gives the failure-coverage risk, taking into account both proposed controls and controls that are actually present in the modeled system; the second column (Act) takes into account only controls that are actu-

#	Failure	Risk (Prop)	Risk (Act)	GE	PA	RE	OP
6	Incorrect Value in Debit Creditors in Payment in PURCHASE	2%	3%	.	X	.	.
7	Incorrect Value in Debit Expenses in Purchase invoice processing in PURCHASE	[55%]	[100%]	.	X	.	.
8	Incorrect Value in Generate purchase requisition in Order Processing in PURCHASE	0%	0%	.	X	.	.
9	Incorrect Value in Print purchase order in Order Processing in PURCHASE	0%	0%	.	X	.	.
10	Incorrect Value in Supplier - Invoice in Supplier in PURCHASE	0%	1%	.	X	.	.
11	Missing in copy invoice to DB in Purchase invoice processing in PURCHASE	2%	3%	.	.	X	.
12	Missing in Credit Cash in Payment in PURCHASE	3%	3%	.	.	X	.
13	Missing in Credit Creditors in Purchase invoice processing in PURCHASE						

Figure 4. Failure-Coverage Risks.

ally present. If a failure-coverage risk is above the allowable risk level for one of the accounts that the failure affects, the failure-coverage risk is highlighted by enclosing it in brackets. A failure with a bracketed failure-coverage risk indicates a potential control weakness in the accounting system that the CIS auditor should carefully examine.

In determining the failure-coverage risk for a failure, COMET first determines the set of controls in the flowchart model that are relevant to the detection of the failure and then assesses, for each relevant control, the likelihood that the control will fail to detect the failure, called the *control-detection risk*. The failure-coverage risk for a failure is determined by multiplying together the control-detection risk for each control that could detect the failure. The risks are multiplied together because we assume that the controls operate independently, and for a failure not to be detected, all the potentially detecting controls would have to miss it.

In assessing the control-detection risk for a given control and potential failure, COMET takes into account three different factors: (1) control strength, (2) control defeat, and (3) control attenuation:

Control strength is an assessment of the in-

trinsic effectiveness of the control, based on its type and its performance. In COMET, the control strength is initially determined from the answers supplied by the modeler to a generic series of questions about how the control is performed. The control strength can be adjusted later as a result of testing the control.

Control defeat is an assessment of the degree to which a control is rendered ineffective by problems with the maintenance of reference data on which it depends. For example, a control cannot be relied upon if the maintenance process for a database of information that it uses has potential failures that are not sufficiently mitigated by controls.

Control attenuation is a measure of the degree to which the effectiveness of a control is reduced by the distance in the flowchart between the control and the primitive activity whose failure it might detect. Control attenuation varies with the type of control and the types of activity along the path from the control to the failing activity.

Key-Controls Selection

A set of *key controls* is a subset of controls in the model that is sufficient to adequately mitigate the risk of all those potential failures that have audit significance and are adequately

mitigated by the full set of controls. Because placing reliance on a set of controls requires that the controls be tested for proper operation, testing costs can be reduced by choosing a minimal set of key controls. Unfortunately, the problem of finding a minimal set of key controls is a computationally intractable minimal set-covering problem. COMET uses a greedy algorithm that works well in practice but does not guarantee a minimal set.

In selecting a set of key controls for testing, COMET uses a relative measure of the importance of a control in reducing the failure-coverage risk of potential failures; this measure is called *control contribution*. The control contribution for a control is relative to a set of failures, F , to be covered and a set of controls, C , to be compared. At each point in the selection process, the set F consists of those potential failures whose risk is sufficiently mitigated (with respect to allowable risks) by the complete set of controls in the model but not yet by those controls already selected for testing. The set C consists of those controls not yet selected for testing. If there are any failures in F that have unique detectors in C with a control-detection risk that is less than 1, all these unique detectors are added to the set of key controls. Otherwise, the next control selected for addition to the key controls is that control with the highest control contribution relative to F and C . The algorithm terminates when the set F is empty, or there are no controls in C with nonzero control contributions.

Performance

COMET has been used successfully by Price Waterhouse CIS auditors to construct and analyze models of complex client accounting systems. A representative example is a stock trading room system whose COMET model has a total of 934 nodes, including 217 primitive activities, 104 composite activities, 118 collections, and 139 controls. COMET's analysis produced 709 potential failures, of which 338 were found to have an impact on the validity of accounts, and 68 were potential defeaters of controls. Of these relevant failures, all but 17 were found to be covered adequately by the controls in the system. COMET found 60 controls to be key and, therefore, candidates for inclusion in a plan for testing controls. The total time required for the analysis was under 30 seconds on a 66-megahertz Pentium PC.

Application Use and Payoff

A beta release of COMET has been used on a pilot basis by Price Waterhouse CIS auditors on a variety of real audits in several countries around the world, including Australia, Argentina, Brazil, India, Malaysia, Mexico, the United States, the United Kingdom, and much of Western Europe. The pilot audits have involved clients from a representative cross section of different industries, including banking, insurance, oil and gas, manufacturing, and entertainment. The official 1.0 version of COMET was released in May 1996.

The CIS audit partners and managers who have supervised the pilot audits believe, based on their experience, that the use of COMET will lead to a significant improvement in auditor productivity. It is difficult at this point to reasonably estimate the size of the gain because a number of factors must be taken into account:

First, the nature of the work performed changes with the use of COMET. Business processes and their controls are documented to a greater level of detail and more rigorously than they would have been previously. Thus, the documentation cost was increased, but the analysis performed by COMET allows the auditor to spend much less time anticipating possible errors and thinking about the controls available to detect and correct them. The increased detail and rigor of the models, in conjunction with the analysis performed by COMET, allows a greater reliance to be placed on controls with a comparable level of auditor effort.

Second, a nontrivial learning curve applies to the efficient use of COMET to model and analyze systems. Experience on the pilot audits suggests that it takes a typical user three to four jobs before he/she becomes truly proficient in the use of COMET. Part of what a user needs to learn through experience is the choice of an appropriate level of detail at which to model a system. Enough detail needs to be added to allow a useful COMET analysis to be performed; too much detail adds to the modeling cost without an additional payoff from the analysis.

Third, the cost of using COMET to model a system and its controls can be amortized more effectively over several years than previous methods of documenting the system. COMET is most appropriately used in a "year of change," either when a new or substantially updated system has been installed by the client or with a new client. In subsequent years, when minor system updates occur, the COMET model can be updated quickly and the impact of the

changes on controls' effectiveness analyzed. Thus, somewhat greater initial modeling effort in the year of change can be justified because the work that needs to be performed in subsequent years is reduced.

Fourth, the use of COMET can reduce the cost of testing. Because of the difficulty of manually performing a thorough and precise evaluation of controls, there is a temptation to perform more detailed testing of transaction records than would be required if the controls work could be done more efficiently. COMET's ability to automatically generate lists of key controls also leads to more focused controls testing because each control to be tested has been determined to make an important contribution to mitigating the risk of possible failures in the system.

Finally, COMET's rigorous analysis can uncover both control weaknesses and control redundancies, leading to recommendations to the client that are a key value-added function of the audit.

Application Development and Deployment

In 1991, the SAVILE Project was begun at the Price Waterhouse Technology Centre to examine the potential of applying a model-based approach to evaluating accounting systems and their internal controls. An initial prototype, also called SAVILE, was developed in Lucid Common Lisp running on a UNIX workstation to establish proof of concept. The record processing performed by an accounting system was described using an imperative programming language called SPLAT. Expressions in SPLAT were transformed into a causal network to support the evaluation of controls (Hamscher 1992).

The CIS audit community within Price Waterhouse responded enthusiastically to the SAVILE prototype, and resources were authorized to implement the SAVILE approach on the standard platform found in Price Waterhouse practice offices—IBM PC clones running Microsoft WINDOWS. In late 1992, work began on developing a more graphic form of representation for SAVILE models that would support both a highly interactive flowcharting system and the analysis of failures and evaluation of controls. Franz Inc.'s Allegro Common Lisp for WINDOWS was chosen as the implementation language to support rapid application development in the WINDOWS environment.

Since early 1993, an average of three full-time programmers have worked on the devel-

opment of COMET. In addition, the involvement of CIS auditors was critical to developing a system that matched the requirements of the CIS audit task. A senior CIS manager was assigned to the Price Waterhouse Technology Centre for two months in 1994, two months in 1995, and one month in 1996 to work intensively with the COMET developers to refine the system design.

The CIS audit staff has developed a training course in the effective use of COMET in response to increasing worldwide demand. To date, approximately 20 percent of the total number of CIS auditors in Price Waterhouse firms worldwide have taken the course. In the European firm, all CIS auditors with more than one year of experience are being trained in the use of COMET, and it is the recommended tool for use with relatively complex client systems.

Maintenance

As a model-based application, COMET does not contain a large knowledge base encoding expert experience in the domain of CIS audit. Thus, the often-difficult issues surrounding knowledge base update and maintenance are eliminated. Rather, the behavior of COMET's analysis engine is a product of the properties of a small set of primitive activity and control types and the structure of the particular accounting system model being analyzed. The set of primitive activity and control types has been remarkably stable over the course of COMET's development and has been found adequate to model a large variety of different client systems encountered during the pilot audits.

After the official release of COMET, responsibility for evolutionary development will transfer from the Research and Development Group in the Price Waterhouse Technology Centre to a Price Waterhouse organization responsible for supporting audit-related software.

Conclusion

Most applications of model-based reasoning have been to engineering domains. COMET applies model-based-reasoning techniques to a new task domain, the analysis of the effectiveness of controls in accounting systems. Because of the complexity and variability to be found in realistic accounting systems, CIS auditors have difficulty evaluating controls to the level of detail required to place a high degree of reliance on systems when performing an audit of a company's financial statements. COMET allows a CIS auditor to focus on build-

ing a model that accurately describes the accounting system, then makes use of this model to automate the analysis of the adequacy of the controls for detecting potential errors in the system. Demand from the Price Waterhouse CIS audit community for deployment of COMET has been high because it is an effective tool in support of delivering high-quality audits to clients.

Acknowledgments

We would like to acknowledge the contributions of two Price Waterhouse CIS auditors who have been instrumental in the development and deployment of COMET. Pat Russell gave us early guidance on the issues of importance to CIS audit and has been a tireless champion of COMET within the firm. Robert Halliday worked intensively with us on key design issues and has been our day-to-day liaison with the CIS audit practice, giving us numerous suggestions on how to make COMET more useful to the CIS auditor.

Note

1. They are also called electronic data processing (EDP) auditors or information systems risk management (ISRM) auditors.

References

- Bouwman, M. J. 1983. Human Diagnostic Reasoning by Computer: An Illustration from Financial Analysis. *Management Science* 29(6): 653-672.
- Bridgeland, D. M. 1990. Three Qualitative Simulation Extensions for Supporting Economics Models. In Proceedings of the Sixth Conference on Artificial Intelligence Applications, 266-273. Washington, D.C.: IEEE Computer Society.
- Hamscher, W. C. 1994. Explaining Financial Results. *International Journal of Accounting, Finance, and Management* 3(1): 1-19.
- Hamscher, W. C. 1992. Modeling Accounting Systems to Support Multiple Tasks: A Progress Report. In Proceedings of the Tenth National Conference on Artificial Intelligence, 519-524. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Hamscher, W. C.; Console, L.; and de Kleer, J., eds. 1992. *Readings in Model-Based Diagnosis*. San Francisco, Calif.: Morgan Kaufmann.
- Hart, P. E.; Barzilay, A.; and Duda, R. O. 1986. Qualitative Reasoning for Financial Assessments: A Prospectus. *AI Magazine* 7(1): 62-68.



Robert Nado has been a member of the Research and Development Group at the Price Waterhouse Technology Centre since 1988. He has been the COMET project manager since early 1994, and he has also been responsible for refining

COMET's model representations and analysis engine. He previously worked on developing software for international tax planning and supporting knowledge-based-system development. Prior to joining Price Waterhouse, he was a developer of the KEE and OPUS knowledge-based-system development tools at IntelliCorp. Inc. He received an M.S. in computer and communication sciences from the University of Michigan.



Melanie Chams has been a member of the Research and Development Group at the Price Waterhouse Technology Centre since 1991. She was one of the principal developers of COMET and was responsible for the development of its user interface. Chams was formerly a senior member of the technical staff at Symbolics Inc. She received her certificate in AI from the University of California at Los Angeles and her Masters in applied mathematics from the Swiss Institute of Technology. Her interests include innovative user interfaces and human-computer interaction.



Jeff Delisio is the technical web master for the Price Waterhouse World Firm. He has been a member of the Research and Development Group at the Technology Centre since 1988. He was a principal developer on the COMET Project and contributed to the development of algorithms for the analysis engine. He previously was project manager for PLANET, an expert system for assisting audit partners and managers in strategic planning, now deployed in Price Waterhouse firms worldwide. He is currently involved in exploring emerging Internet technologies. Delisio received an M.S. in computer science from the University of Massachusetts at Amherst.



Walter Hamscher is a Boston-based Price Waterhouse Technology strategy consultant to Fortune 1000 companies in the computing, telecommunications, manufacturing, and financial-services sectors. He was previously the director of research and development at the Price Waterhouse Technology Centre, where he initiated projects in model-based reasoning, natural language understanding, and Internet computing. He was the principal investigator on the COMET Project for its first two years. Hamscher holds a Ph.D. in electrical engineering and computer science from the Massachusetts Institute of Technology.