Causal Tracing: One Technical Solution to the Modeling Dilemma

Robert L. Eberlein
David W. Peterson
William T. Wood

Ventana Systems, Inc.
60 Jacob Gates Road
Harvard, Mass 01451 USA

Abstract

A central dilemma for system dynamics is the fact that the same human limitations that motivate the use of models also make models hard to create, debug, and even hard to use. Two commonly proposed escapes from this dilemma are education and generic models. We propose a third, technical, approach and give an example. The example approach, "causal tracing," is a computer tool that makes it much easier to find the feedback loop or input responsible for a given variable's behavior. Correctly implemented, this tool reduces the time required for causal tracing by a factor of 10 to 100. The payoff is faster and more accurate creation of models and use of models.

The Dilemma

A good model is: a laboratory for exploration, hypothesis testing, forecasting, and experimentation; an embodiment of structural knowledge; and a consistent way of deriving behavior from feedback structure. To be these things, a model must incorporate richness, detail and novelty. The model must also be as free of errors as possible. Unfortunately, the creation of a good model is subject to four serious hazards: time, expense, skill requirements, and bugs. That is, making a model "from scratch" is likely to take too long, cost too much, absorb gifted people, and end up with embarrassing errors anyway.

Ironically, the difficulties of dealing with models stem from the same reasons that models are necessary in the first place: the inability of the human mind to accurately come to grip with the consequences of dynamic interaction. For example, to understand the behavior of a model, one must do quite a bit of mental simulation of fragments of the model. This is obviously easier than mentally simulating the whole model, but the unreliability of mental simulation has been well documented, even for trivial models.

Thus, the human frailties which motivate modelling are the same human frailties which impede the development of good models. It is natural, therefore, that the modelling process should be full of inefficiency, errors and unwanted surprises. This is a central dilemma affecting the entire field of system dynamics.

Responses to the Dilemma

The recent popularity of generic models is one natural reaction to the dilemma. Generic models have, in this context been seen as a way of distributing some high-quality benefits of system dynamics with less risk and expense. But generic models don't help with new problems. Further, understanding the behavior of any model is still subject to the human limitations mentioned above. And, of course, the creation of generic models in the first place is subject to all the hazards of model building, debugging, and understanding.

Education is another approach toward resolving the dilemma. However, the teaching of system dynamics attempts to develop skills which are fundamentally lacking in most people -- the ability to mentally simulate small dynamic structures.

This paper begins a third, technical approach to the problem. Rather than restricting models, or trying to develop skills for which the human animal is ill-suited, we propose the use of additional modeling tools to compensate for human weaknesses. In other words, we propose tools to enable models of increasing complexity, novelty and correctness. We believe these tools allow the development of quality models in a short time, mitigating the requirements for simple generic structures and imposing none of the restrictions.

There are a large number of techniques for automatically assisting the modeler. We have implemented a number of different tools and techniques with varying degrees of success. Rather than try to review all of these techniques, we will show one especially helpful and pleasing tool. We refer to this as "causal tracing", a partially automated technique for determining important causes of a variable's behavior.

Background

The existence of interesting, aberrant or incorrect behavior is a normal part of the modeling process. Determining the source of such behavior is typically a slow and error prone process. Identification of the variables involved is a time consuming and mindless task of tracking backwards, variable by variable, using equation listings, documenter listings, diagrams, and voluminous output or repeated simulations. The various forms of documentation may be located elsewhere and not immediately accessible. Even worse, they may be out of date and, therefore, misleading. The simulations may be slow in coming, especially if a large quantity of output is required.

The traditional process of looking for a source of behavior, then goes something like this:

1. Observe that a variable is doing something interesting.

2. Print and paste a graph in a notebook.

3. Look at the model source for the equation for the variable

4. Find the behavior of each right-hand-side variable in turn (this alone may require new simulations).

5. When a variable also has the interesting behavior print it and put it in the notebook.

6. Repeat steps 3 to 5 until the chain has terminated or gone full circle.

7. Draw a diagram and start to think about what is going on.

Even at a glance it is clear that steps 3 through 5 are time-intensive, but not thought-intensive. The process is good at uncovering important feedback and important errors, but the mechanical difficulties of tracing chains of cause and effect may absorb most of the (highly skilled) human time. Also, the slowness of the process tempts the modeler to take short cuts and ignore areas where he or she "knows" nothing interesting is happening. The solution is simple in concept, and is illustrated below.

Causal Tracing: A Conceptual Overview

Causal tracing is a collection of computer tools for automatically finding and viewing the causes of interesting behavior.

The computer implementation that we use is based on a windowing environment using a pointing device (such as a mouse). Figure 1 shows a simplified view of the environment. At any time, one variable from the model is listed as the "Focus" variable (in this case, Population). Above the Focus display are two tools for causal tracing: one that generates tree diagrams of the variables which cause the Focus variable, and one that displays compact time graphs of the Focus variable and its causes. CAUSAL GRAPHS are collections of graphs sharing a common time axis stacked one on top of another; the uppermost graph is the Focus variable, and the following graphs show the causes of the focus variable. CAUSAL TREE diagrams are "word and arrow" charts that branch out, showing successive levels of causes for a variable. Causal tree diagrams are purely structural. They simply show paths of influence with no reference to behavior. Causal graphs bring up the behavior, simultaneously showing the paths of influence. Because tree diagrams contain less information, they allow the compact display of multiple layers of causality.

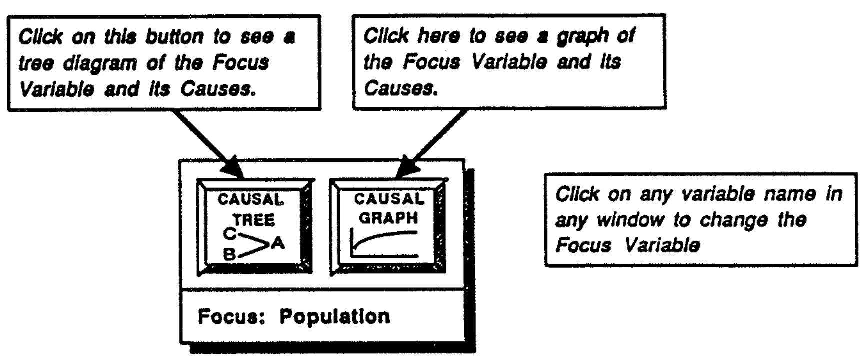

Figure 1: Command Window, showing the Focus variable and two buttons that can be clicked to see a tree diagram or graph of the Focus variable and its causes.

Figure 2 shows a simplified view of the use of the causal tracing tools. First, the user selects Population as the Focus variable, applies the causal graph tool to see stacked time plots of the Focus variable and its causes (the "causes" being the variables on the right-hand side of the Focus variable's equation). Clicking on one of the causal variable's names in the graph (in the example, deaths) causes that variable to become the Focus variable. Thus, using only the mouse (no typing, no looking at equations), the user can identify the causes of the causes of the causes, and so on until the chain reveals one or more feedback loops, an input, or an error that causes the observed behavior.

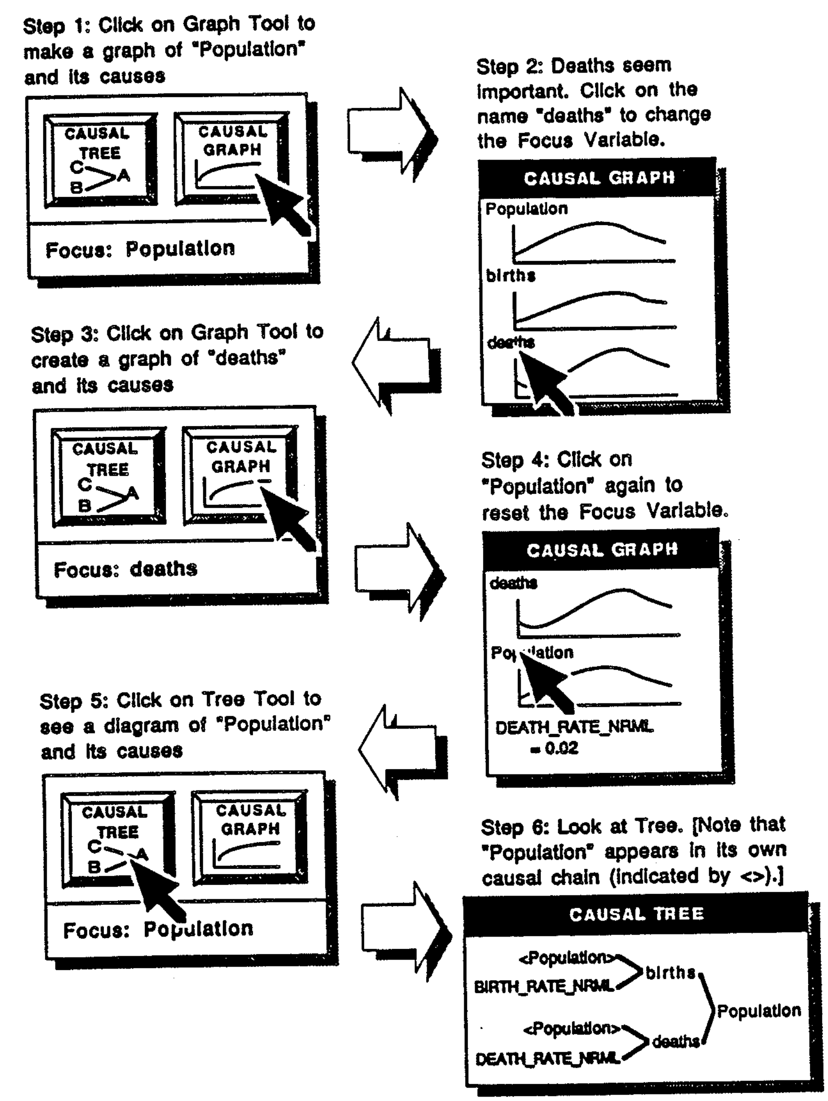

Figure 2: Tracing causality using just the mouse.

Causal Tracing: A Specific Example

The best way to proceed in the description of causal tracing is to give specific examples of how it can be used. To do this, we will use the World Dynamics model

(Forrester 1973), and look at its base overshoot and decline behavior. The example serves as a basis for discussing the tools and also shows their usefulness, even for small models. Our own experiments with these tools have been done with models even smaller than this, as well as models containing many thousands of variables.

The following description is necessary for clarity and completeness, but its length is misleading, because much more time is required to read the description than to actually perform the tasks.

Figure 3 shows the base simulation of the world model as described in Forrester (p. 70 - where it is labeled Figure 4-1). Suppose we begin with population, and ask the obvious question: Why does population level off and then fall continually? By selecting Population as the Focus variable and clicking on the CAUSAL GRAPH tool, we get the graph shown in Figure 4. Population falls because deaths exceed births. (This seems obvious, but it is important to verify it. As we all know, models may do reasonable-looking things for the wrong reasons.).Now lets select deaths as the Focus variable to figure out why these are so high. The answer to this is shown in Figure 5. The key contributor to the high number of yearly deaths is Population itself, with minor, and sometimes mutually cancelling effects from crowding, food, pollution and the material standard of living.

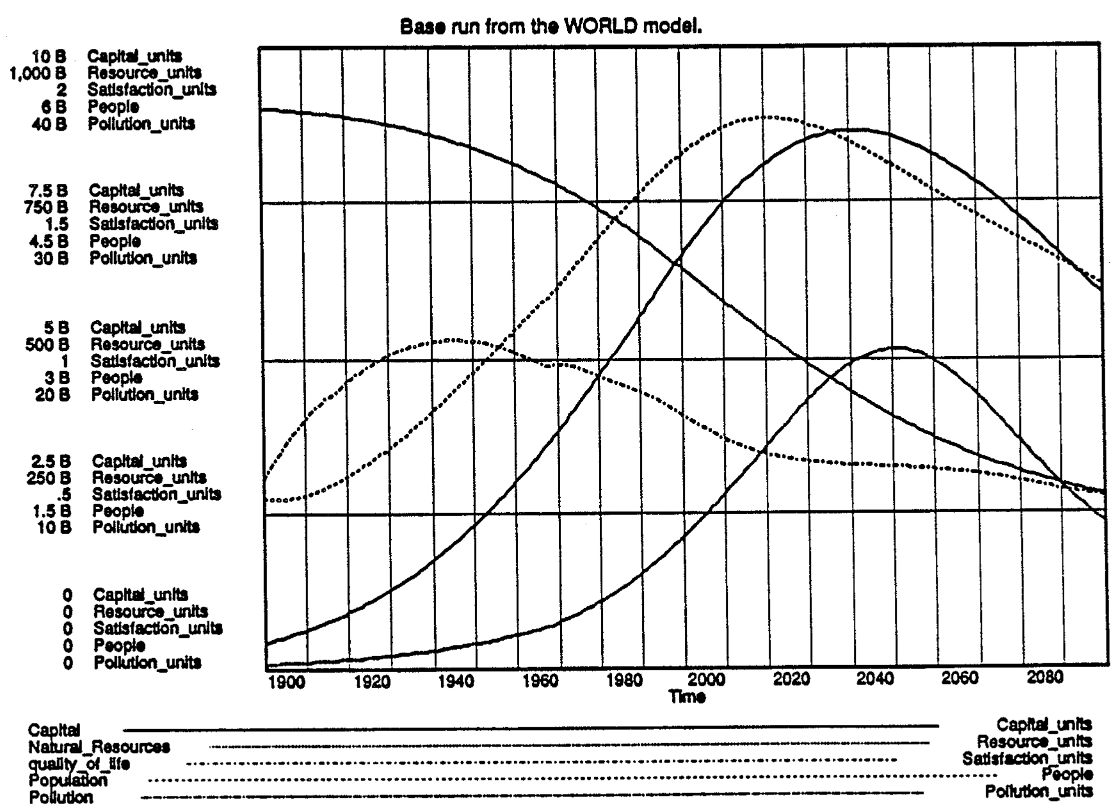

Figure 3: Standard graph from base simulation from the WORLD model.

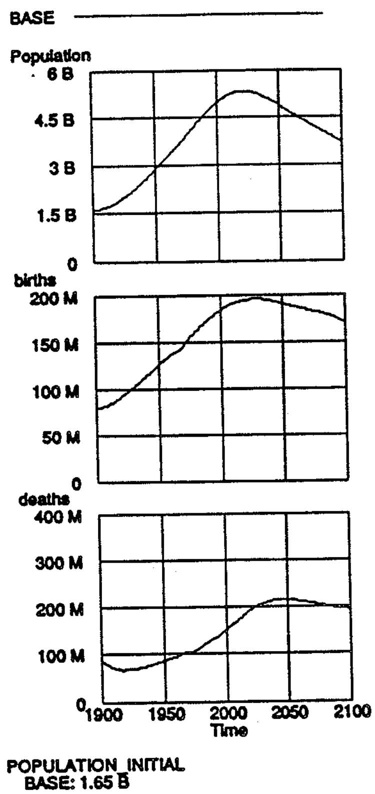

Figure 4: Causal Graph for Population.

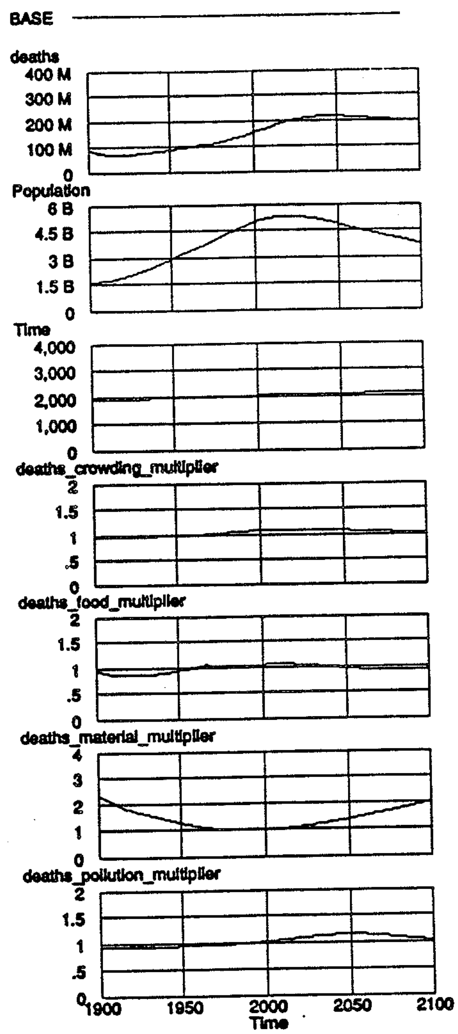

Figure 5: Causal graph for deaths.

So far we have uncovered the standard feedback loop from population to the number of deaths back to population. That feedback loop is fundamental to population dynamics, but can only explain constant growth or decline rates in population. Since we are interested in the change from growth to decline we look elsewhere. From the timing it is clear that crowding, food and pollution are key in turning population change around, while the material standard of living is responsible for the continued decline.

Focussing on the continued decline let us select the effect of the material_standard_of_living on deaths. Using a different view of causal tracing we can ask for the tree shown in Figure 6. From this figure it is clear that the effect of the material_standard_of_living on deaths is determined completely by the effective_capital_ratio. The other inputs being a constant and a table as indicated by full capitalization of names (all such formatting is, of course, automatically achieved by the computer environment). The effective_capital_ratio, however, has a number of inputs and we need to see which is important. Selecting this, we see in Figure 7 that the real problem is the natural_resource_extraction_multiplier. A causal tree (not shown) identifies the feedback to Natural_Resources, and a strip graph of its causes is shown in Figure 8.

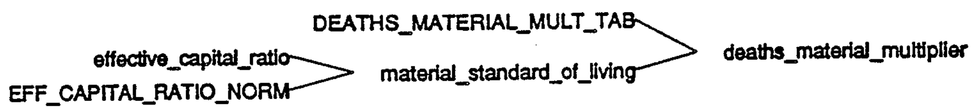

Figure 6: Causal tree for deaths_material_multiplier *(depth 2).*

We have just been lead back to the caption on Figure 4-1 in *World Dynamics* (our Figure 1) for we are looking at the "basic behavior of the world model, showing the mode in which industrialization and population are suppressed by falling natural resources." The path by which we reached the conclusion was moderately direct. It would have been more direct if we had assumed a strong working knowledge of the model. It would have been less direct if we had gone off on a few dead ends, and checked all behavior that is the least suspicious. The latter is more common, because with the causal tracing tools the cost of doing the checking is so low.

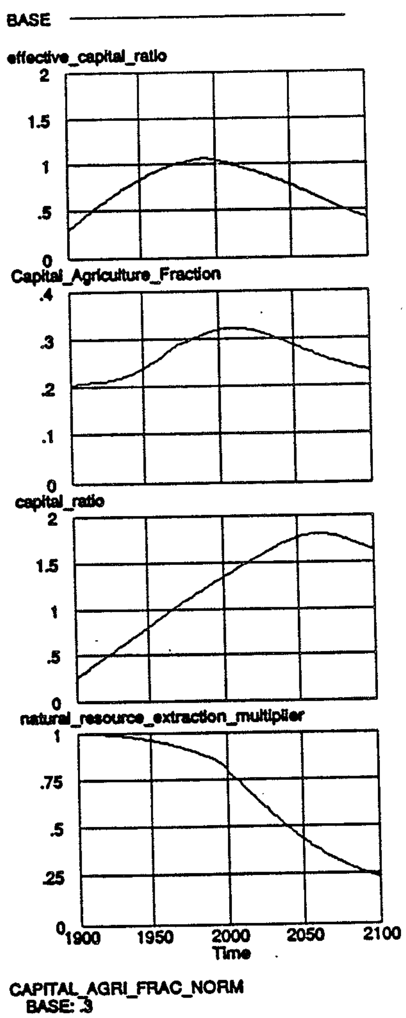

Figure 7: Causal graph for effective_capital_ratio

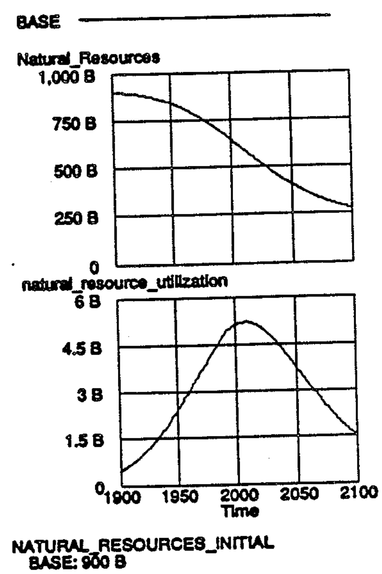

Figure 8: Causal graph for Natural_Resources.

Causal Tracing: A Second Example

To continue the story, and show two small refinements on what we have seen, we add a second simulation, with decreased the use of natural resources. Again we start with Population, asking for its causes for both the reduced-resource simulation and the base simulation, as shown in Figure 9. The dramatic difference is immediately noticeable: a huge spike in the number of deaths results if Natural_Resources are used more efficiently. We focus on this spike by narrowing the time range of interest (done quickly by clicking and dragging the mouse across the desired fraction of any of the graphs), and this turns Figure 9 into Figure 9a. Doing a strip graph for the causes of deaths results in Figure 10. It is clear that Pollution is the most important influence, with food in second place. From Figure 11 (via the CAUSAL TREE tool) either pollution_absorption or pollution_generation is to blame. Guessing that pollution_generation is at fault we do a strip graph for this as shown in Figure 12. We are hot on the trail; the cause is the pollution_capital_multiplier. Figures 13,14 and 15 show trees along which the most likely path is to capital_investment, the material_standard_of_living and the effective_capital_ratio. A causal graph of the effective_capital_ratio (Figure 16) shows that the natural_resource_extraction_multiplier is much bigger, because (Figure 17) the fraction of natural resources remaining is so much bigger.

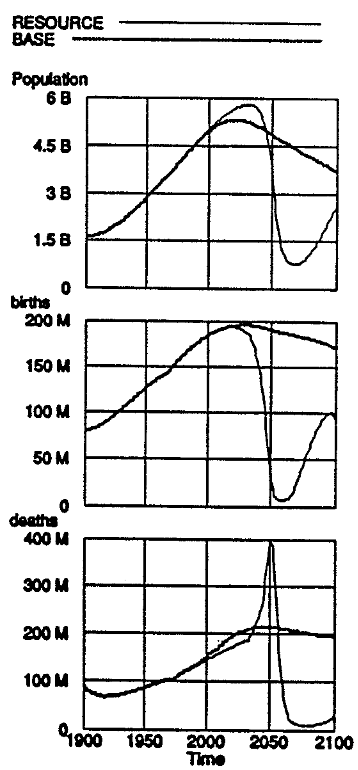

Figure 9: Causal graph for Population *(two simulations).*

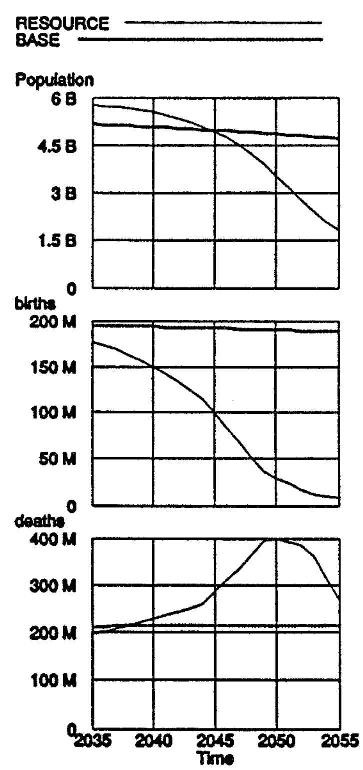

Figure 9a: Same as Figure 9, with time axis zoomed.

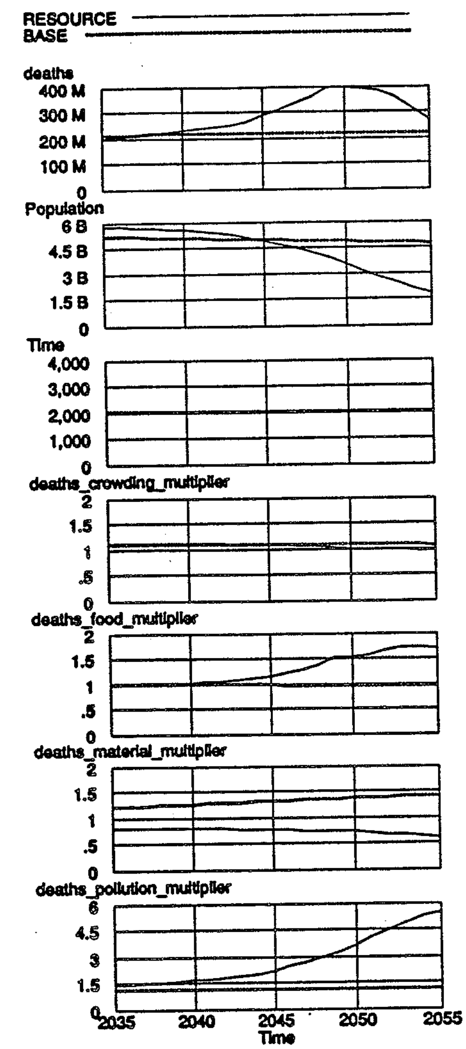

Figure 10: Causal graph for deaths.

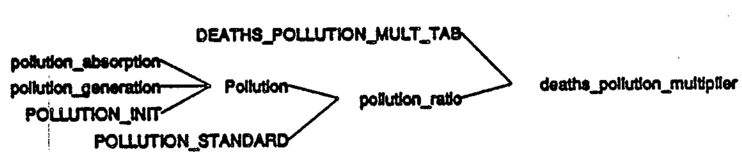

Figure 11: Causal tree for deaths_pollution_multiplier.

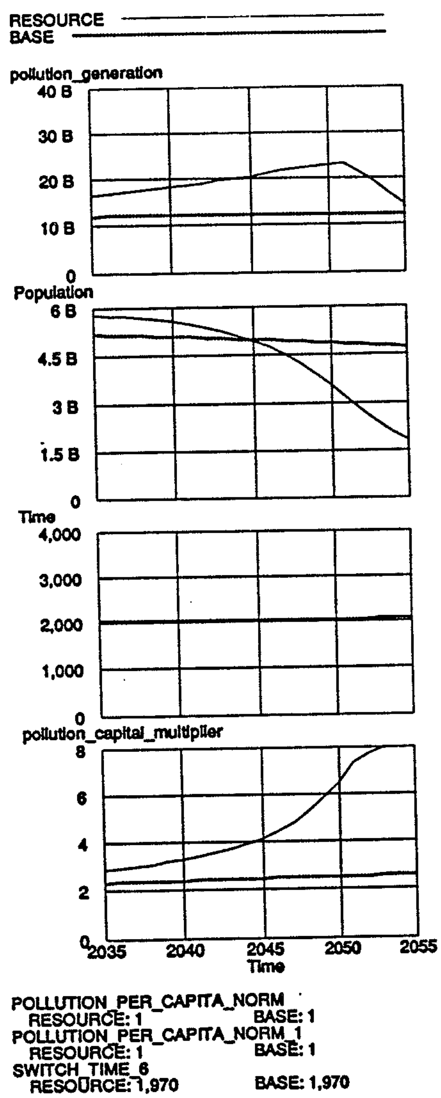

Figure 12: Causal graph for pollution_generation.

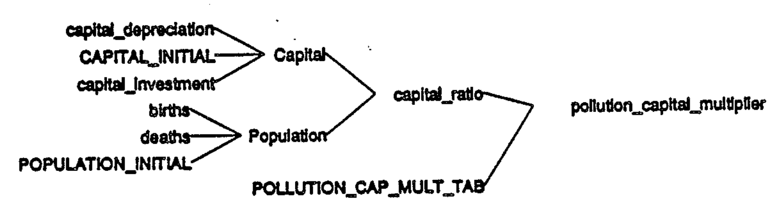

Figure 13: Causal tree for pollution_capital_multiplier.

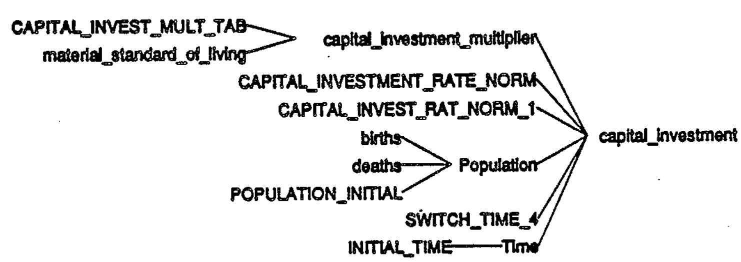

Figure 14: Causal tree for capital_investment.

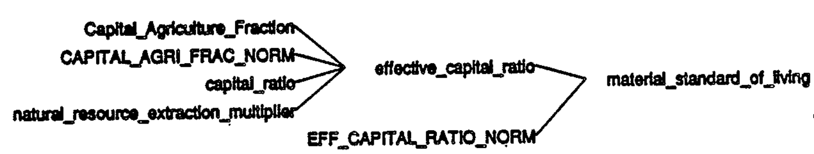

Figure 15: Causal tree for material_standard_of_living.

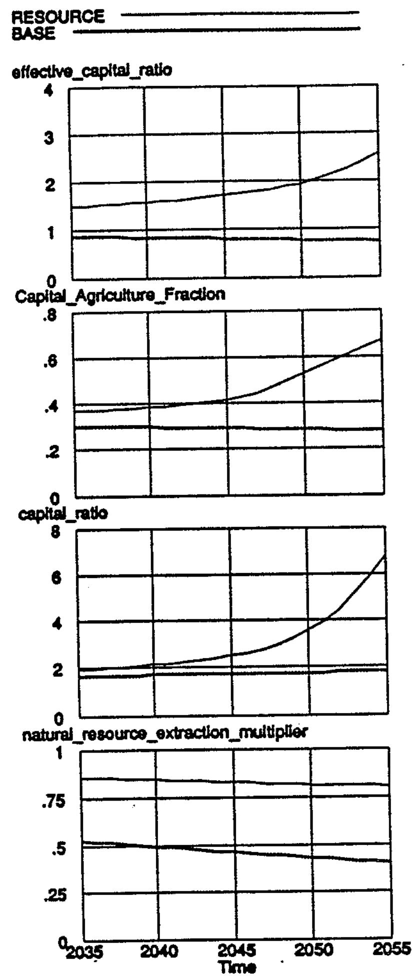

Figure 16: Causal graph for effective_capital_ratio.

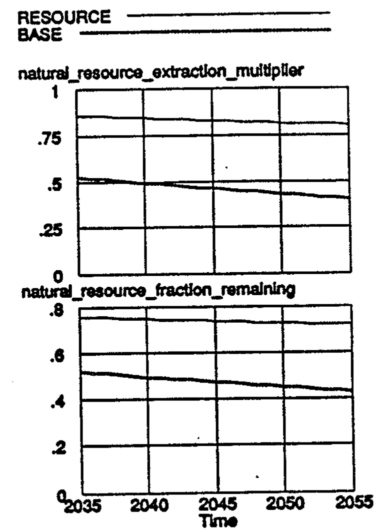

Figure 17: Causal graph for natural_resource_extraction_multiplier.

Unwinding, the lower consumption of resources allows increased capital buildup, which increases material well being and, in turn, pollution. With the pollution eventually overwhelming. With lots of pictures, and more than the usual amount of words we quickly arrive at the interesting results. At this point you may want to say, "Hold on; the choice of the proper trail is easy when the book has been written, but what about exploring things blind? Aren't the dead ends still there?"

The answer is an emphatic yes, the dead ends are still there. The causal tracing tools we have described do not eliminate futile searching, they simply make the search happen quickly. The dead ends are quickly explored and quickly dismissed. In addition, less obvious opportunities are also more likely to be pursued. When getting sidetracked means taking just a few minutes to check something out, it is an easy and worthwhile thing to do. It is no longer necessary to dismiss something as not worth looking into. If there is any odd or interesting behavior observed, it is typically just a few button presses to determine if the causes are interesting. If something important or unexpected is going on, it is simple to try to figure out what it is, not requiring hours of staring at equations and diagrams, hand-tabulating numbers, and repeating simulations to obtain more output.

Conclusions

Since 1960 the speed of commonly available computers has increased a thousandfold, but modelers have benefitted only in limited ways from the additional power. The increase in raw computing power has simply not translated into an equivalent increase in modeling efficiency. There are a number of reasons for this, and many evolve around the high degree of human interaction in the modeling process. But much of the continued slowness of modeling is simply anachronistic. We are in the era of fast computers and powerful languages, and our modeling work should benefit directly from the increased computing power available to us.

The tools we have discussed (and others not discussed) offer about two orders of magnitude improvement in the speed of analyzing simulations. At first glance, a factor of 100 may not seem revolutionary, but it is. For a tangible comparison, consider the difficulty of walking from Rome to Beijing, compared with flying by jet. Jet travel is about 200 times as fast as traveling by foot, and this increase in speed has opened up the far reaches of the globe to many people.

Tools for causal tracing, combined with other techniques - some of which are available and some of which are dreams, may give us eventually a 1000-fold boost in speed for several of the most time-consumptive modeling tasks. That boost should not only improve the speed and quality of existing modelers' work, it should also bring the power of dynamic modeling into the hands of many more people.

References

Forrester, J.W. 1971. *World Dynamics*. Cambridge, Mass.: Wright-Allen Press.