A Planning Support System to Construct and Evaluate Alternatives

Insung Lee and Lewis D. Hopkins

(January 25, 1999)

1. Iteration and Procedural Expertise in Planning Support Systems

The underlying premise of a Planning Support System is that it should complement human planners by recognizing comparative advantages and biases of computers and humans. Harris and Batty (1993) argue that PSSs should provide for sketching alternatives and tracing out their consequences. Klosterman (1997) argues that approaches to PSSs have and will continue to evolve in concert with conceptions of the problematic tasks of planning. Here we describe a particular planning support system and its application to a bus routing problem to demonstrate the potential to support the iterative process of constructing and evaluating alternatives. This PSS, called PEGASUS (Procedural Expertise for Generation and Analysis of Spatial Urban Systems), supports sketching alternatives, tracing out their consequences, and evaluating these consequences. It also supports this process in the kind of collaborative context identified by Klosterman as the most recent conception of the planning task.

The development PEGASUS focused on two particular concerns. First, we sought to support the iterative process of constructing and evaluating alternatives in a tightly integrated, coherent environment. Such iteration is pervasive in both descriptions of and prescriptions for problem solving (e.g. Alexander 1992; Black 1990; Bryson 1995). In particular, such iteration recognizes that values are discovered, formed, and changed in the process of creating and analyzing alternatives and that alternatives are recognized in the process of assessing values. Second, we sought to support the specific techniques or procedures—to provide procedural expertise—so as to decrease the effort required of planners when using these technical procedures. In this PSS, procedural expertise supports one multiattribute evaluation technique, but it is conceptually possible to design procedural expertise for any partially structured procedure used in planning.

A bus routing problem based on realistic data demonstrates the implementation of the planning support concepts and tools in PEGASUS. The bus routing problem is particularly appropriate for our purposes because of its spatial characteristics that allow display and manipulation of alternatives, its complexity that precludes direct analytical solutions, and its evident multiple objectives. In section 2, after briefly describing our particular version of the bus routing problem, we describe the components of PEGASUS. In Section 3 we summarize our PSS development process, which itself iterates between construction and evaluation, and initial observations about such systems in use.

2. Components of PEGASUS

PEGASUS has two major parts: alternative construction and alternative evaluation. The first part provides tools to edit an existing bus route, to add a new route, and to combine routes into networks. Models included in PEGASUS then trace out the performance of each network and generate a multiattribute table of performance measures. The second part provides tools to use Multiattribute Utility Theory (MUT), the Analytical Hierarchy Process (AHP), or Computation of Equivalent Alternatives (CEA) alone or as complementary techniques to evaluate alternative networks. These evaluations provide diagnostic feedback that suggests improved routes and networks[1]

2.1 The Bus Routing Example

The bus routing example consists of a portion of Champaign-Urbana, Illinois using a "virtual" network of 133 nodes and 412 links. This virtual network comprises all links (street segments) that can potentially be included in a set of bus routes that make up a network of routes. Seven network attributes are analyzed for a given route network: ridership (total riders), capital cost (cost of required number of busses), revenue (farebox revenue), operating cost (variable costs of operating busses on network), average walking time (average walking time per trip to and from bus stops), average travel time (average total travel time per trip), and transfer rate (percentage of trips requiring transfers).

The performance of alternative networks is based on standard analyses with a few simplifications. See Lee (1993) for a more detailed description of the bus routing, mode choice, and trip distribution models used. Trip generation data were derived from a survey conducted by students at the University of Illinois at Urbana-Champaign, and trip distributions were computed using TRANPLAN (Urban Analysis Group 1993). A Trip-Interchange Modal-Split Model (see e.g., Dickey 1983) was coded to determine the transit share of trips given a particular bus network. This model considers travel time, travel cost, and income level in computing the travel behavior and thus the performance characteristics of a particular network. The following must be provided as data for the models: transit fare, transfer cost, capital cost per bus, operating cost per bus per mile, personnel cost per bus per mile, maximum walking distance, average waiting time for bus, average transfer time, tradeoff between walking time and transferring, tradeoff between walking time and transit time, and private car occupancy rate.

Figure 1 Computation of the Attribute Levels

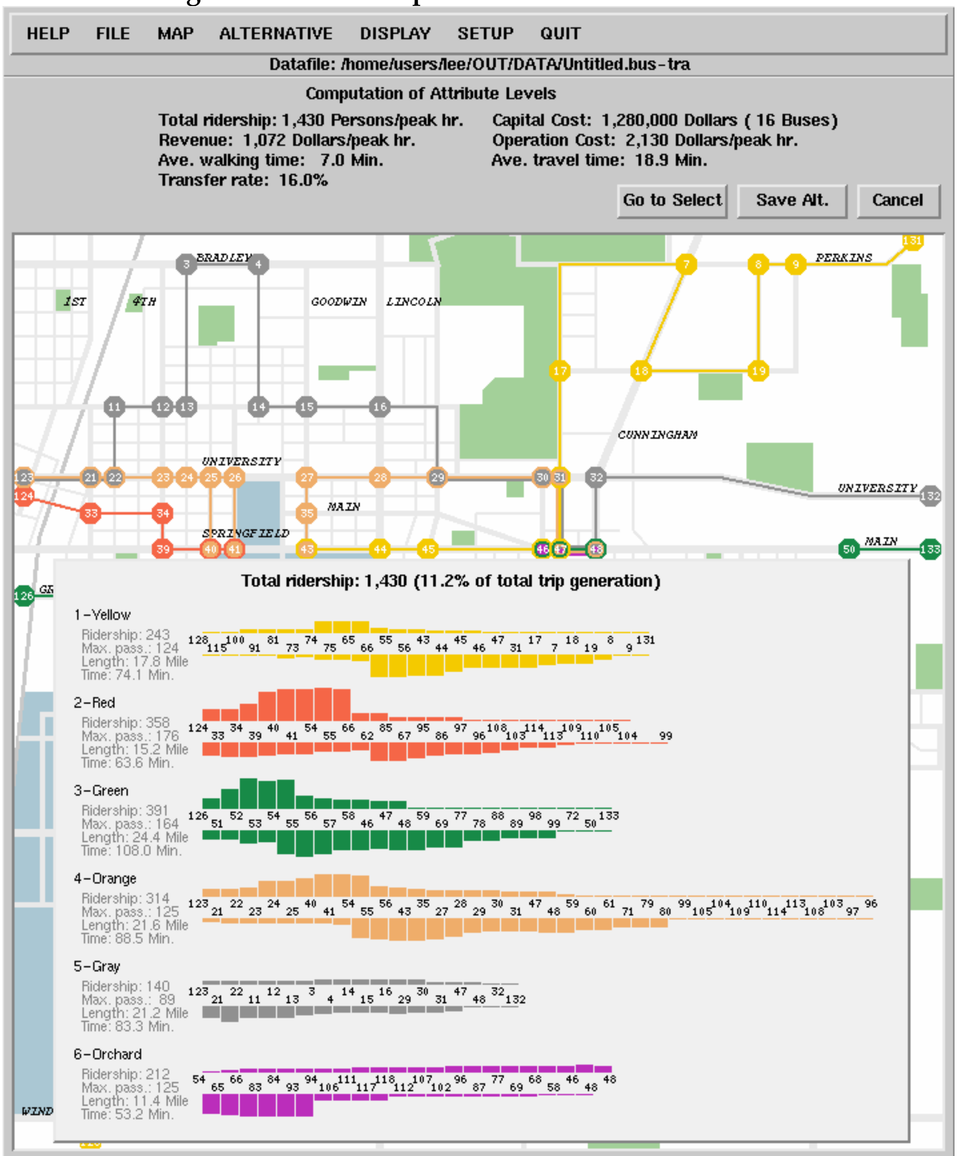

Figure 1 shows the report of performance for a network. PEGASUS reports total ridership, capital cost of required busses, fare revenues, operating costs, average walking time, average travel time, and transfer rate. PEGASUS also reports a map of the network and a diagram of passenger loading on each link for each route. The actual computer screens are in color and convey more information than is possible in the black and white images presented here of some screen images. For example, routes are color coded with consistent colors throughout all windows. Note that these results are computed for a network of routes, not for individual routes, because the entire network of available routes affects rider choice with respect to individual routes. These data provide the effectiveness matrix, which is the input for evaluation and a major source of feedback for users to consider in constructing alternatives. The map display supports consideration of recognizable characteristics of a network that may not be quantified in the modeling, such as equity of access among neighborhoods. A user can also request that PEGASUS display desire line diagrams of trip demand and maps of population characteristics, which provide additional information for constructing improved alternatives.

2.2 Construction of Alternatives

PEGASUS supports five aspects of constructing alternatives: editing a route, constructing a route, constructing alternative routes that are as different as possible from a given route, combining routes into a network of routes, and estimating whether a network is likely to be better than alternative networks already identified.   These tools help in creating entirely new routes and networks or in modifying already proposed routes or networks as part of an iterative process.

Figure 2 shows the route editing window. At the top, the nodes comprising a route are shown in sequence for the busses traveling in a route in each of its two directions. This distinction between directions of travel allows for loops and one way streets. The user can select nodes to be edited either by pointing to the node number in the sequence list or by pointing to the node on the map. Identifying a “from node” and a “to node” deletes or adds a link to the sequence. Once a route has been edited, clicking on the “Calc. Ridership” button computes the effectiveness matrix of performance measures for the modified network of routes.

Figure 2 Interface Window to Edit Existing Route

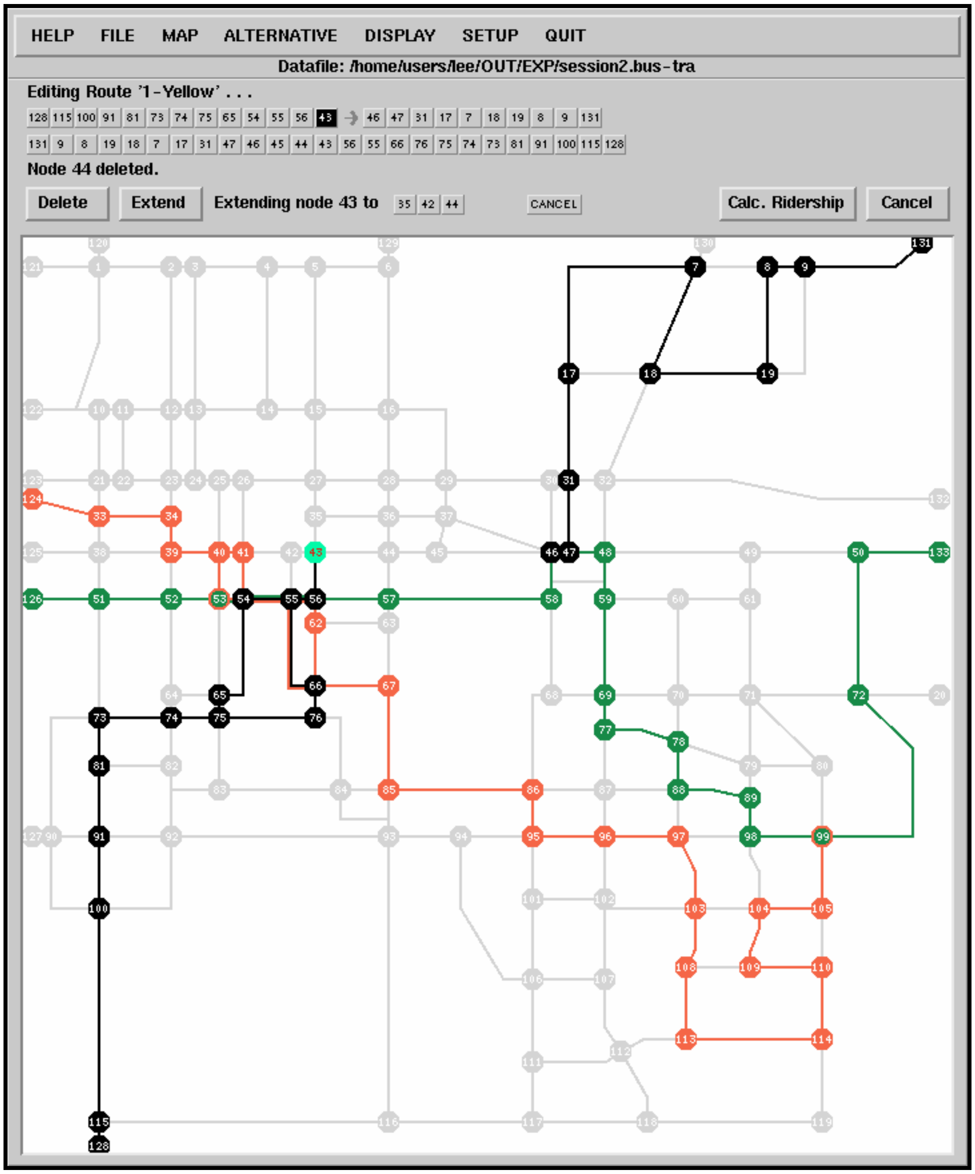

Figure 3 shows the route creation window. A user can create a new route by specifying the entire sequence of nodes in both directions, by specifying a beginning and ending node and any number of intermediate nodes, or by specifying a degree of difference from an existing route. The shortest path algorithm used follows Winston (1984) and considers such subtleties as right turns versus left turns in minimizing travel time. Travel time is, however, only one objective among many, so the system also supports the Modeling to Generate Alternatives (MGA) approach (Brill et al. 1990; Hopkins, Brill, and Wong 1982). MGA generates good routes that are as different as possible from routes already generated. To use MGA, the user 1) identifies a route already proposed by specifying its beginning, ending, and (optionally) some of its intermediate nodes, 2) a degree of difference from that route, and 3) the number of alternative

Figure 3 **Interface Window to Add New Route** Error! Bookmark not defined.

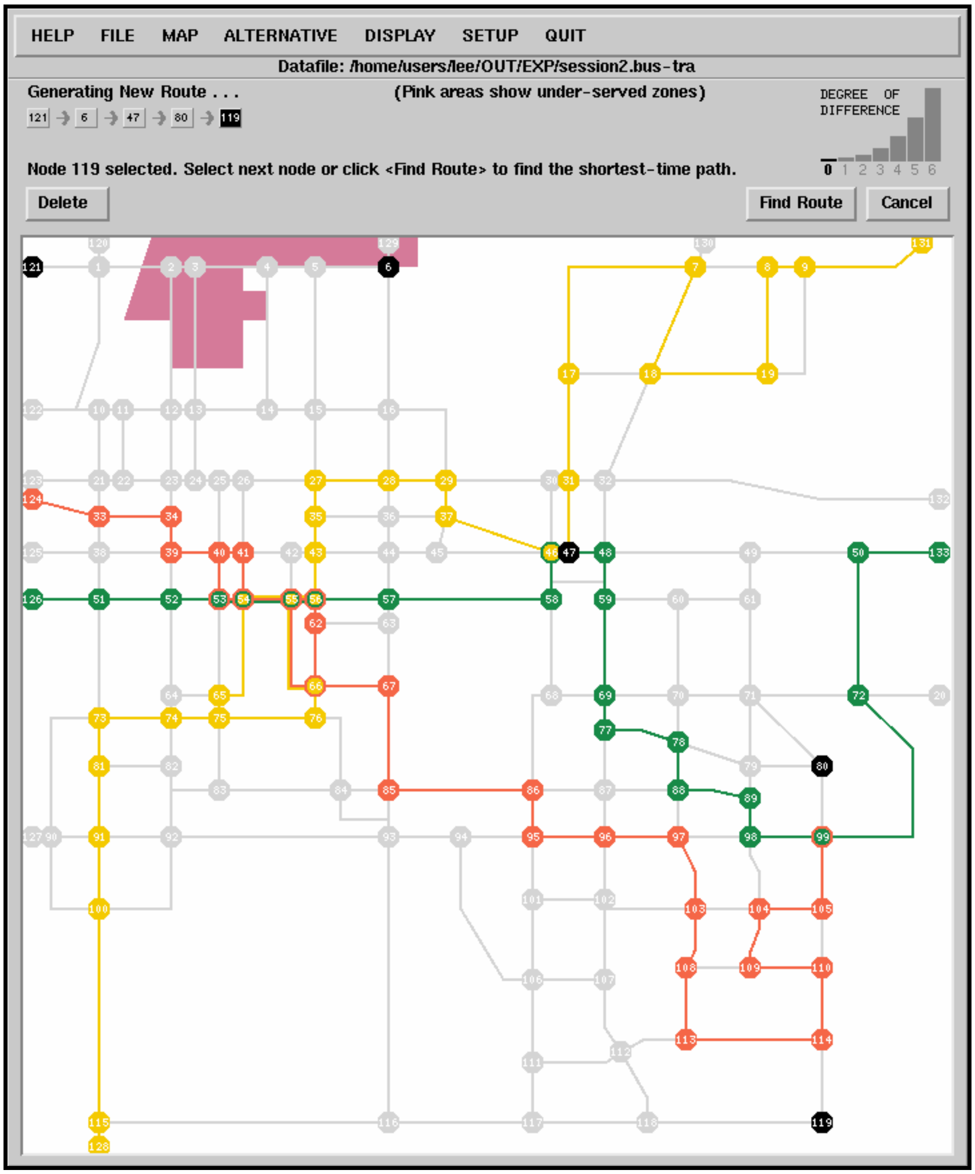

routes to be generated. PEGASUS then displays alternative routes as shown in Figure 4. The user can then select from these alternative routes a route to be incorporated into a network for evaluation.

Figure 4 An Example Output of Modeling to Generate Alternatives (MGA)

Degree of Difference: 0 (the sortest-time route) | **Degree of Difference: 2** | **Degree of Difference: 4**

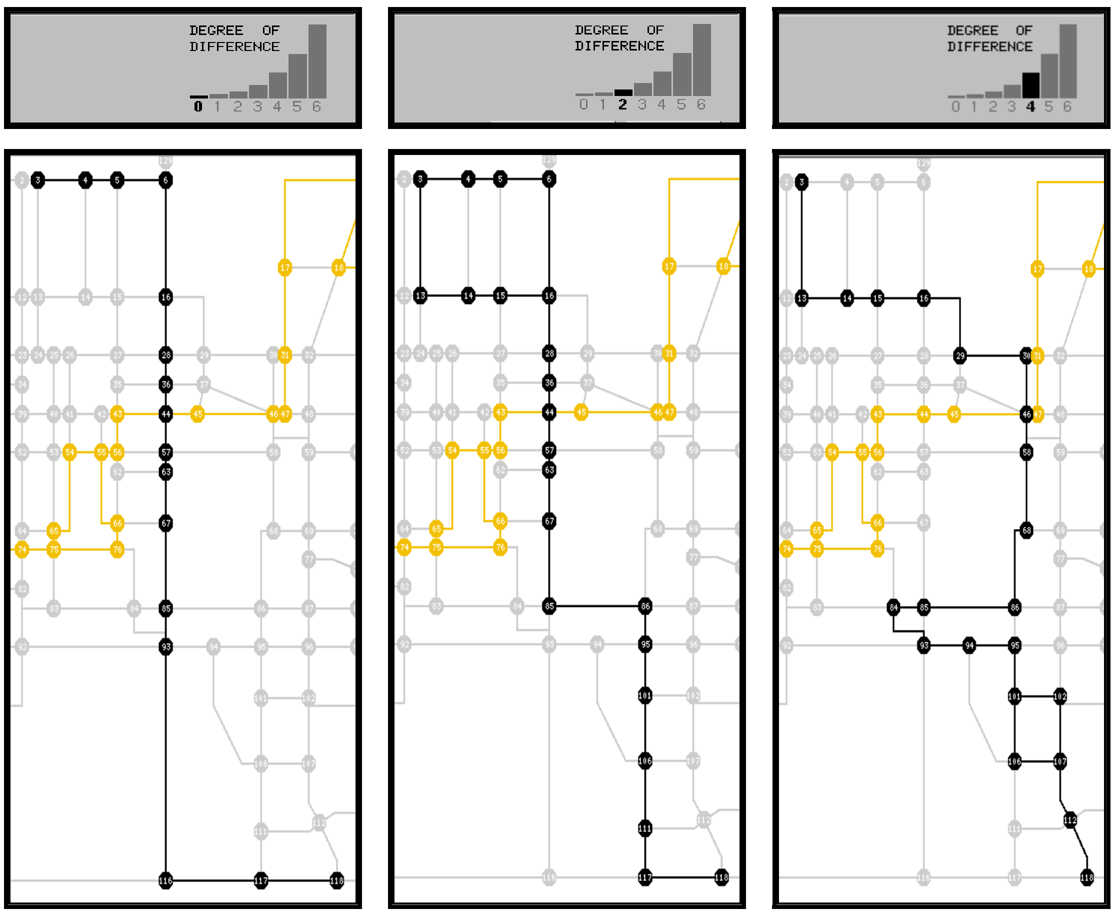

Figure 5 shows the window used to select routes to be combined into networks and to select networks to be evaluated. The left side of the screen displays the effectiveness matrix for all the selected networks and attributes. This matrix can be edited to add performance measures for attributes in addition to those that are generated from the models included in PEGASUS. Note the icons on the left, which represent each network by a simplified graphic image of its routes. These icons are memory devices to assist users in keeping track of which alternatives are which.

Figure 5 **Alternative Table Window** Error! Bookmark not defined.

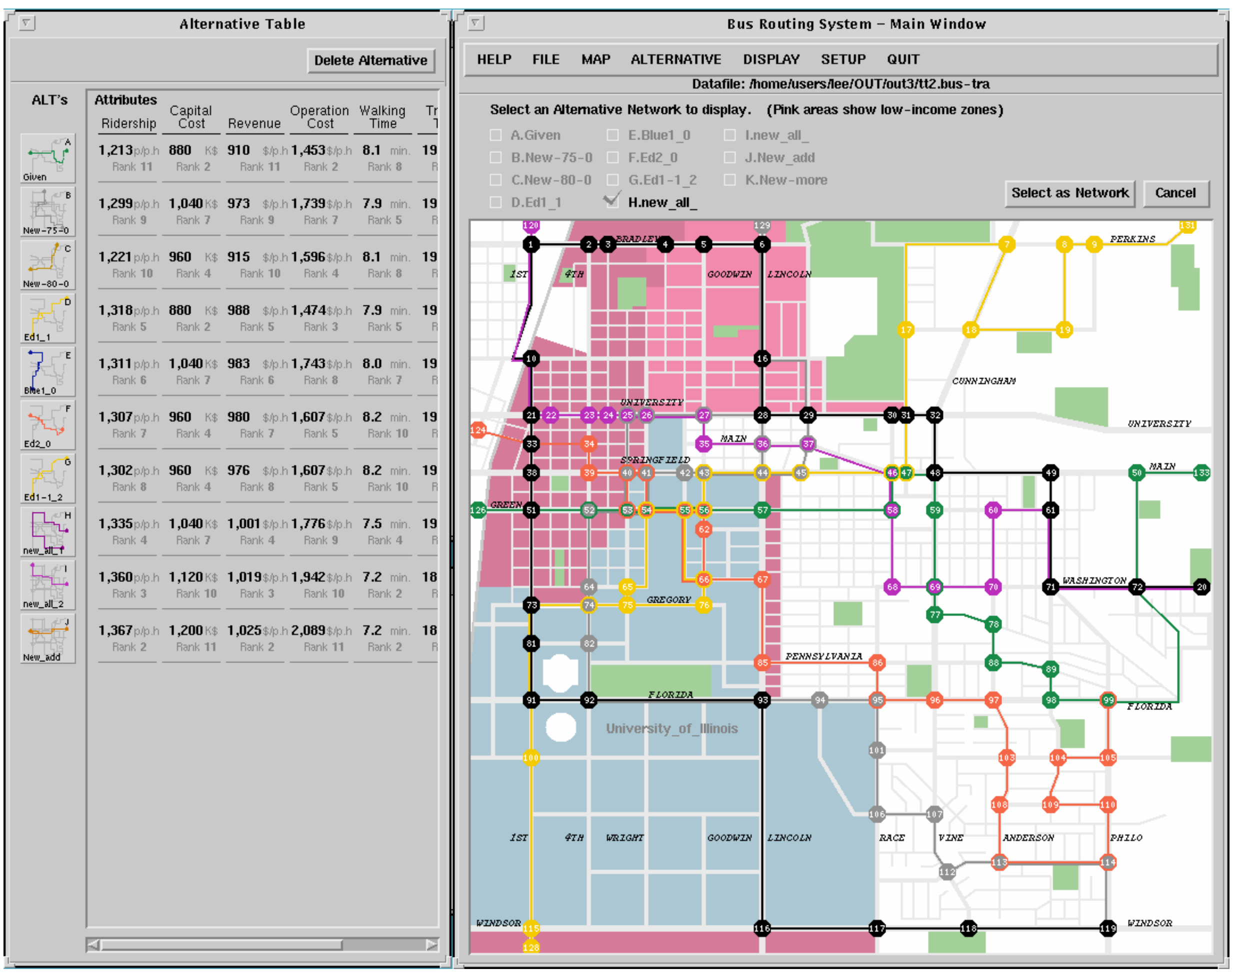

Finally, a tool for scanning alternatives can check whether a newly proposed network is dominated by an already proposed network, meaning that the already proposed network is as good as or better than the newly proposed network on all performance attributes. This tool can also estimate the likelihood of dominance based on preference information elicited during a previously executed evaluation. This tool enhances the iterative process of alternative construction and evaluation by avoiding the time required to fully evaluate a network that is very unlikely to be preferred to a network that has already been evaluated.

2.3 Evaluation of Alternatives

PEGASUS is designed to support complementary use of four different multiattribute evaluation techniques because each has advantages and disadvantages, adherents and detractors. A user can choose one technique or use all four techniques as multiple estimates of preferences, which may be appropriate given the difficulties in applying any one of them accurately (Lai and Hopkins 1995). PEGASUS supports Multiattribute Utility Theory (MUT), the Analytic Hierarchy Process (AHP), Computation of Equivalent Alternatives (CEA), and direct determination of linear value functions, all of which are described in detail in Lai and Hopkins (1989).

Figure 6 shows the window for MUT. The upper part is a gallery from which attributes can be selected for valuation. The icons indicate which attributes have been assessed previously and by which technique. The middle part displays the value function based on the mid-value splitting technique (Keeney and Raiffa 1976). For MUT, the system also supports exponential value functions and provides for assessment of weights given value functions.

Figure 6 Midvalue Splitting Error! Bookmark not defined.

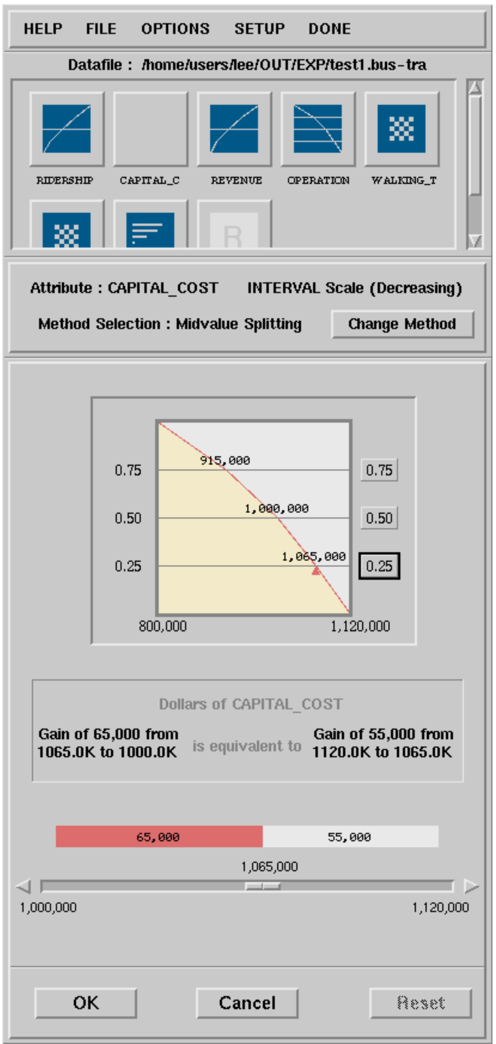

Figure 7 Interface Window of AHP

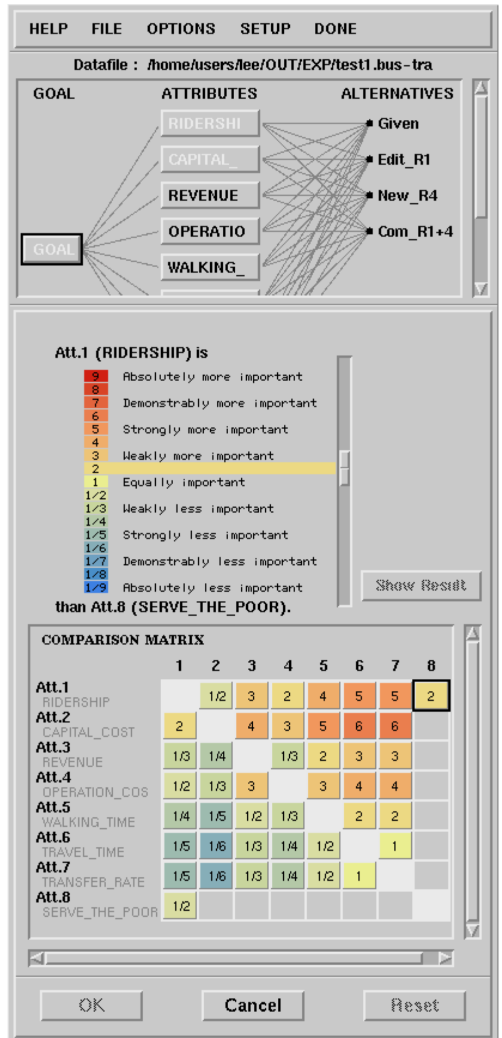

Figure 7 shows the judgment window for pairwise tradeoffs in AHP. Another window supports definition of the AHP hierarchy. PEGASUS also supports one-way and two-way sensitivity analyses of value functions and weights, which are described in detail in Lee (1993).

PEGASUS also supports CEA and, most importantly, provides procedural expertise for implementing the CEA technique. CEA involves a sequence of tradeoffs chosen to discover the preference order among a given set of alternatives based on as few tradeoff judgments as possible (Patton and Sawicki 1993; Stokey and Zeckhauser 1978). The selection of judgments to make and the order in which to make them is, however, left to the artful skill of the decision maker. The objective of each tradeoff judgment is to create a hypothetical alternative equivalent in value to a real alternative, but with attribute performance levels that dominate or are dominated by another alternative. The dominated alternative can then be eliminated from consideration. The procedure is repeated until only one alternative remains. Using Monte Carlo simulation of large numbers of preference patterns, Lee and Hopkins (1995) devised rules of expertise for selecting a sequence in which to consider value tradeoffs.[2]

PEGASUS implements these rules as procedural expertise to support CEA as shown in Figures 8 and 9. Given a profile of performance for a set of alternatives on a set of attributes, as shown in the upper part of each of these windows, the user must choose a pair of alternatives and a pair of attributes on which to make a tradeoff judgment: How much am I willing to give up of attribute 1 for how much gain in attribute 2, given the current levels of performance on all attributes, so as to consider the new (hypothetical) alternative equal in value to the original? Based on the results of our simulations, PEGASUS suggests pairs of attributes and alternatives and reports the rule used to select them. The user can accept these suggestions, request use of a different rule to generate a suggestion, or reject all suggestions and make a direct selection of pairs on which to make tradeoff judgments. The user is then asked for the tradeoff judgment as shown in the lower part of the window shown in Figure 9. The system uses the new tradeoff judgment to check for dominance among alternatives and deletes any dominated alternatives. The process continues until one alternative remains.

PEGASUS provides additional tools for complementary use of the evaluation techniques, such as the comparison report shown in Figure 10. Such comparisons are useful because even careful and considered judgments elicited through the different multiattribute evaluation

Figure 8 Alternative Selection Window Error! Bookmark not defined.

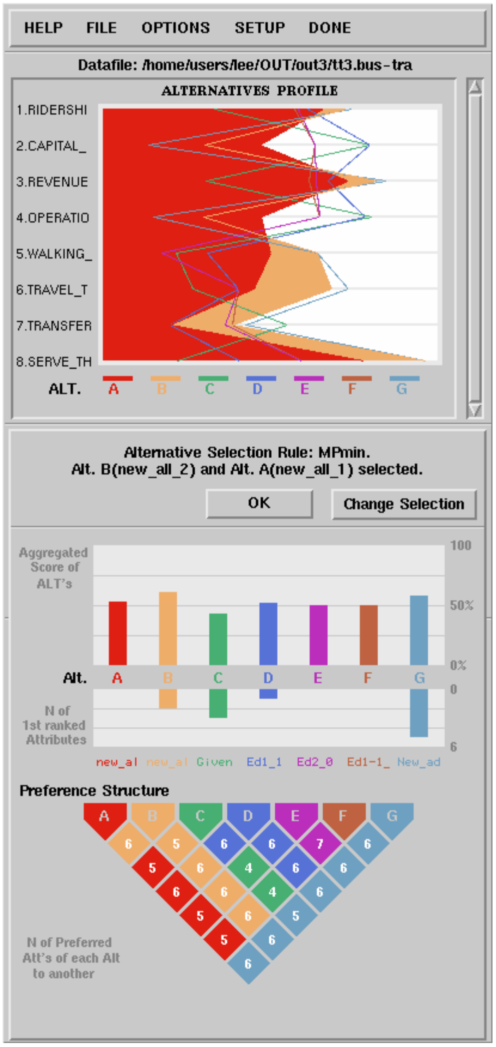

Figure 9 Attribute Selection Window Error! Bookmark not defined.

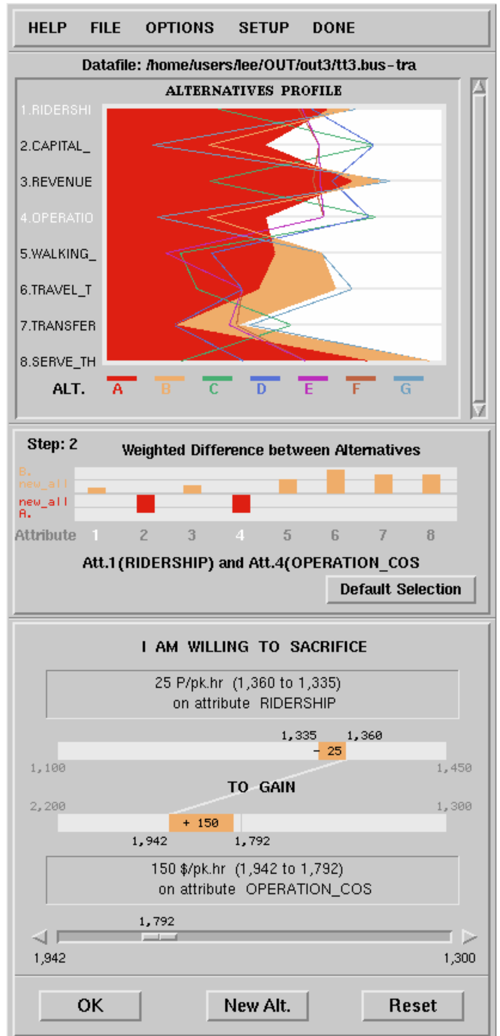

techniques are likely to yield different results because of the differences in underlying theories and operational characteristics among the techniques (Lai and Hopkins 1989; Lai and Hopkins 1995). These reports are designed to encourage a user to consider how the assessments are different and what further considerations might reconcile such differences or lend confidence to a particular choice.

Figure 10 Cross Comparison and Report

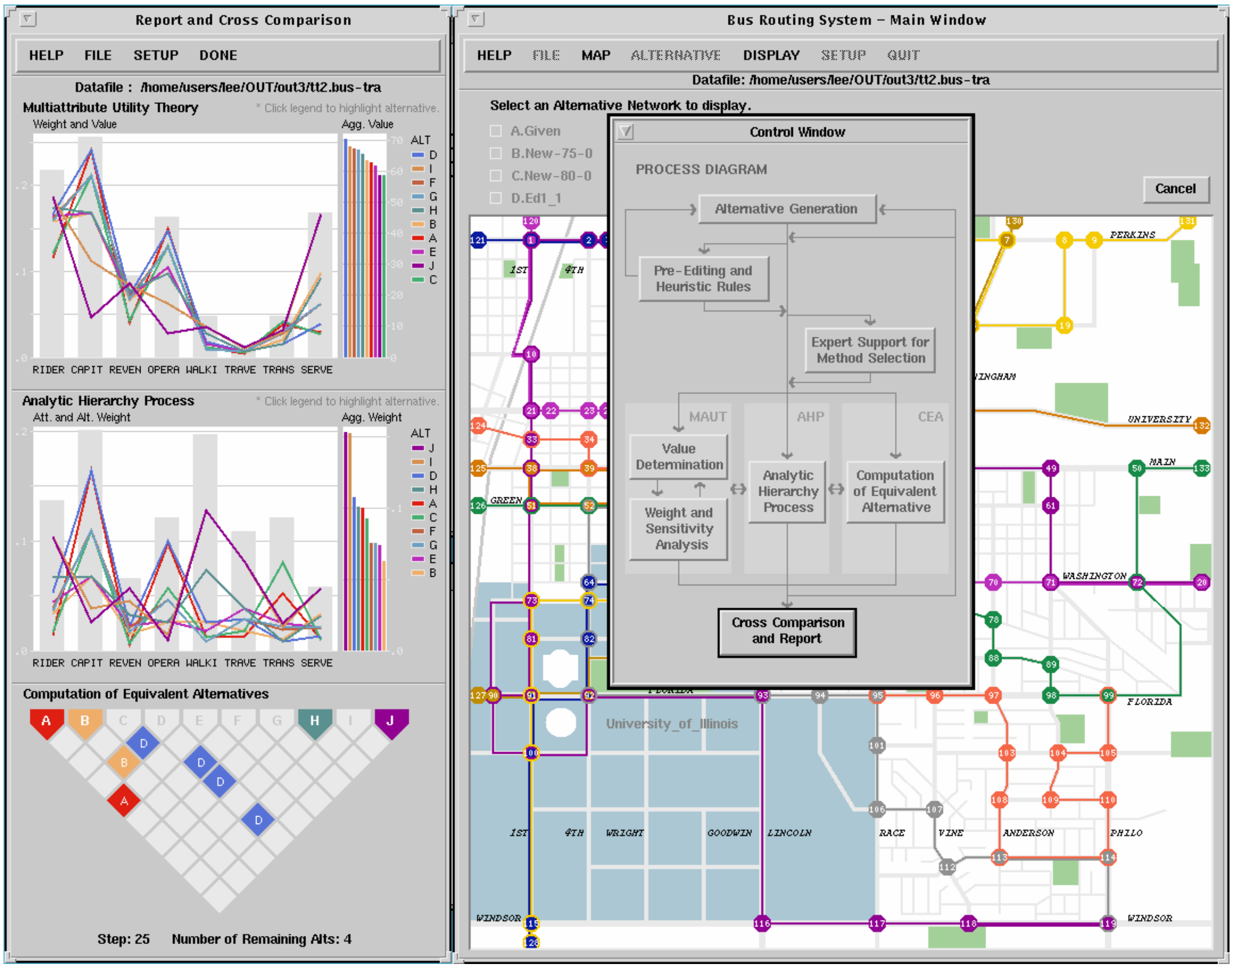

3. Assessment of the Planning Support System

We implemented PEGASUS in the form reported here primarily to consider two performance characteristics of planning support systems: iteration between construction of alternatives and evaluation of alternatives and the effectiveness of procedural support for CEA. Informal assessment of the system was continual throughout its development, but once it was reasonably usable and reliable we turned to a more systematic, though still formative assessment of its intended performance.

Assessment of a PSS should occur not only after a system is developed, but also throughout the process of system development (Cohen and Howe 1989; Keen and Morton 1978; Landauer 1995). During the design and development of PEGASUS, a small number of colleagues tested each module of the system. These tests focused on interface issues and proper operation of the system. Some of the features of PEGASUS are also included in and evolved together with other systems we were developing simultaneously and thus benefited from tests of those related systems (Johnston and Hopkins 1994).

Although it is impractical to seek direct, conclusive, empirical evidences of the effectiveness of a support system for ill-structured problems (Hopkins 1984), it is important to assess how a support system works in use. We did not implement controlled experiments to test specific hypotheses about the effect of the system on problem solving performance because such questions are extremely difficult to address experimentally as demonstrated in Hopkins (1973), Trybus and Hopkins (1980), and Lai and Hopkins (1995). Instead, we explored possible opportunities for further development of the system. For this purpose, we asked a relatively small number (18) of testers to use the system in such a way that we could observe patterns of use. We analyzed the observations qualitatively rather than quantitatively focusing on two questions that were most central to this version of PEGASUS:

> Did the system enable users to iterate between constructive and iterative tasks?
>
> Did users accept the procedural expertise provided for CEA?

3.1 Observation Methods

The 18 users worked on the bus network design problem described above. All the basic data of the bus networking problem—traffic zones, population distribution, income distribution, trip generation and distribution—were given, and users were not allowed to change the data. Users were also not allowed to change certain parameters that might invalidate the computation of attribute levels: transit fares, costs, and rider behavior. PEGASUS computed the levels of the seven attributes for alternative networks: Ridership, Capital cost, Revenue, Operation cost, Average walking time, Average travel time, and Transfer rate. Users could not delete any of these attributes and could not change the reported performance levels for the attributes modeled in PEGASUS. Users were required to consider one unmodeled attribute—service to lower income residents—and could choose to add additional unmodeled attributes if they wished. For these unmodeled attributes, users entered their assessments of the performance levels for each alternative because the system had no models to generate such measures.

We gave users a network that included the three existing bus routes, and did not allow users to delete this alternative from the effectiveness matrix. Users could use any of the methods of alternative generation—Edit existing route, Add new route, MGA, and Reorganize routes—to generate alternative bus networks. Users were required to generate at least five alternatives before using the multiattribute evaluation techniques and to use three multiattribute techniques: MUT, AHP and CEA. Users chose the sequence in which to use the three multiattribute techniques.

Our observations were organized into three major categories:

1) Process Observations

- Sequence of using the sub-procedures of the system
- Time spent in each sub-procedure
- Time spent to reach a decision

2) Performance Observations

- Number of preference judgments to reach a decision
- Improvement in decision quality through iterative, constructive process
- Acceptance of procedural supports by the user

- Efficiency or effectiveness of the procedural supports

3) Perceptual Observations

- Understanding and learning of the judgment techniques
- Ease of answering the tradeoff elicitation questions
- Acceptance of the results of the judgment techniques and heuristic decisions
- Consistency among the judgment techniques and heuristic decisions
- Explanations of inconsistent results among multiattribute techniques

We used automatic recording by the computer—"system logging"—to record each keystroke and each button or menu selection with the mouse. This technique provides rich and accurate data for the process observations and some of the performance observations without disturbing users. We also used an attitude survey in which users answered short questions at particular stages in their process. Most of the perception observations were based on the attitude survey. Finally, we analyzed the output datafiles at each stage of the process to assess problem solving performance.

The assessment design involved a relatively small number of users, each for a relatively long time. Users learn about the bus routing problem and the problem solving methods by using the system, and thus must use it long enough to gain some expertise. Each user participated in three sessions, and users were supposed to spend two hours for each session. A tutorial manual and a detailed instruction sheet were given to the users in each session as detailed in Lee (1993).

In Session I (Walk Through) users answered a short questionnaire about their computer literacy, knowledge of the multiattribute evaluation techniques, and knowledge of transportation planning models and methods. Then each user worked through the tutorial, which introduces all basic functions of PEGASUS and briefly explains the multiattribute techniques.

In Session II (First Iteration) some advanced functions of the system were explained to the user. Each user then generated five alternative bus networks. Before using any of the evaluation techniques users were asked to select the one or two of these alternatives that they thought “intuitively” were best. Then the users evaluated their alternatives using three multiattribute techniques. Users reported their perceptions of the results of their evaluations and of the techniques in a questionnaire.

In Session III (Second and Third Iterations) users were asked to refine their attribute tradeoff determinations from the first iteration and to select a set of tradeoff weights to be used to check whether a new alternative was likely to be dominated by an existing alternative. Users deleted some inferior alternatives, and added new alternatives. Users then reassessed their preferences using the new set of alternatives. Then, the users were asked to repeat the whole process once more—delete inferior alternatives, add new alternatives, and make value judgments using the multiattribute techniques. Users then reported their perceptions as in the second session.

3.2 Results and Discussion

The participants in this assessment were faculty or students in planning or civil engineering at the University of Illinois at Urbana-Champaign. Students were paid five dollars per hour for up to six hours for the three sessions. Among the 18 participants, 15 finished all three sessions. The observation data were carefully reviewed before analysis. Low motivation and insincerity of subjects often cause noise in evaluation results (Bunn and Wright 1991; O'Connor 1989). The system log files were particularly useful for this purpose. One user was excluded from our analysis because the log showed that this user spent less than three seconds per tradeoff judgment, chose extremely skewed value functions without apparent reason, and made pairwise comparisons without trying to enhance their consistency. We thus analyzed data from 14 users.

3.2.1 Time Spent in the Decision

The users spent an average of five hours (300.7 minutes) for the three iterations included in sessions II and III. On the average, more than one third of the time was spent in constructing alternatives. Designing a good bus network was a challenging task to most of these users, both difficult and engaging. The length of time for generation in the first iteration may be due in part to the complexity of the alternative generation procedure. Even with the graphic user interfaces, it took considerable time for the users to become proficient in creating alternatives.

In the second iteration, apparently, the users were more aware of the characteristics of good solutions, and some users spent an even longer time constructing alternatives in the second iteration than in the first iteration. For the multiattribute techniques, however, users spent much less time in the second iteration than in the first iteration.

Users spent less total time in the third iteration. The third iteration was typically used for confirmation of previous decisions rather than for a new level of exploration. Users spent more time in the third iteration on the report procedure and iterative use of judgment procedures than for alternative generation.

3.2.2 Impact on the Problem Solving Process

The observation data show that the problem exploration styles were quite different from user to user. Users differed in the sequence of using the procedures and the time spent in each procedure. In general, exploration processes seemed to fall into three styles. Process charts of the three styles are shown in Figure 11. The upper part (first three hours) of each chart shows the process of session II (first iteration), and the lower part (second three hours) shows session

III (second and third iterations).

Figure 11 Three Types of Decision Making Process

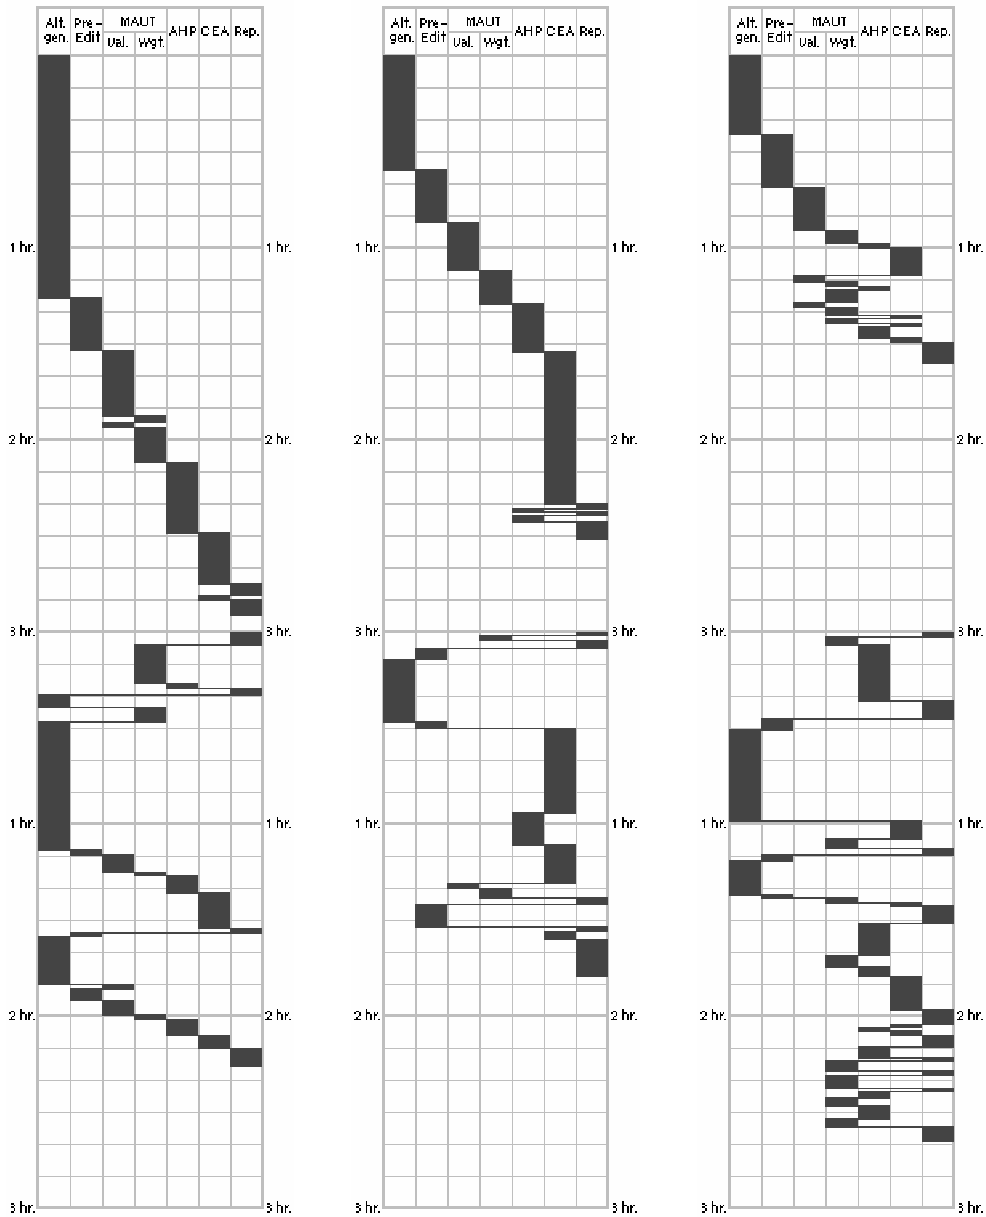

Users of the first style showed no specific preference among multiattribute techniques and used all techniques one after another, spending a similar amount of time for each method, as shown in the first diagram of Figure 11. The second style concentrated attention on a particular multiattribute technique, using other techniques only for confirmation. The third style used a highly iterative process, using one technique for a relatively short time, then another, then back to the previous technique. Although it is hard to say that one of these styles is better than the others, there was a tendency for the process to change gradually from sequential (style 1) to iterative (style 3) as a user's familiarity with PEGASUS, the supported procedures, and the bus routing problem grew. The iterative alternative construction and evaluation style involves more cognitive feedback loops in the exploration, which should increase the likelihood of discovering better alternatives. The iterative style also spends more time in the comparison and fine-tuning of the results from various multiattribute techniques, which should generate more consistent results across different multiattribute techniques. Our informal assessment was insufficient, however, to draw specific conclusions about the effects of these styles.

These observations demonstrate that PEGASUS accommodates problem exploration processes that iterate between alternative construction and evaluation and iterate among evaluation techniques. Users could move easily and naturally between alternative generation and evaluation as well as between different evaluation techniques because PEGASUS facilitated such behavior within a common interactive environment.

3.2.3 Improvement in Solution Quality

We provided an alternative that consisted of three routes and asked the users not to eliminate or change this alternative throughout the evaluation. We used the ratio of the aggregate value of the alternative selected by the user to that of the given alternative (as assessed by the particular user's value functions) as an index to measure the improvement in solution quality each user was able to obtain. Users were able to find better networks. Figure 12 shows that the average ratio increased with each iteration. The results from CEA are not shown in the figure because the CEA does not generate a complete preference function and thus does not generate such an index.

Figure 12 The Ratio of the Aggregate Value of Selected Alternative to Given Alternative

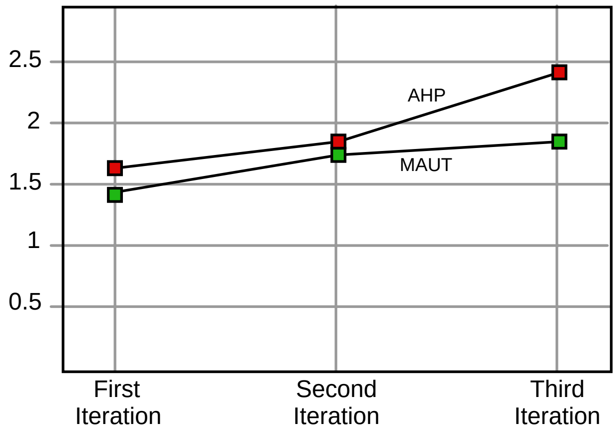

Several factors might contribute to this improvement. First, the feedback of value information to the design phase could help the users structure new alternatives based on a clearer view of their value preferences. The enhanced integration of alternative generation and multiattribute evaluation seemed to encourage exploration of the problem, which should improve the quality of the resulting solution. In addition, the system helped users judge the likely quality of new alternatives in the later iterations. The dominance estimation technique was used to scan an average of 4.8 alternatives prior to applying the multiattribute techniques in the second and third iterations, and the users generally accepted the results. Users discarded 87% (41 out of 47) of the alternatives estimated to be dominated by alternatives they had already created. Thus the solution search process was conducted with better estimates of quality, which should have contributed to the improved solutions.

Finally, a user's ability to interpret information improved with iteration. Users became more aware of the causal relationships between the characteristics of alternative networks and the performance of these networks on particular attributes. For example, one user observed the ridership report window carefully and found that a heavy ridership on any link for a long route caused high capital cost and high operation cost because it required additional busses. This user tried to distribute this ridership by adding an additional route that would draw some of these riders to a shorter route that could be served by fewer busses (because of the shorter time to complete the route). This user was able to construct a better alternative network with lower capital cost, lower operation cost, and higher ridership. This demonstrates the importance of having several different displays to explain the performance measures in relation to the network design, not just to report overall performance as an index.

Identifying feasible network improvements is difficult, which suggests that procedural expertise could be extremely valuable in this task. Such expertise should facilitate learning the relationships between the characteristics of alternatives and performance, and thus increase the joint cognitive ability between human and machine to generate good solutions. Sengupta and Abdel-Hamid (1993) refer to such useful information as “cognitive feedback.”

3.2.4 Use of Procedural Expertise

The users accepted the procedural expertise provided by PEGASUS. The impact of procedural support was particularly noticeable in the CEA procedure. For the CEA, the system provided procedural support for selecting alternatives and attributes for tradeoff judgments. The users accepted 96% (115 out of 120) of the default alternative selections suggested by the system. The default rule for attribute selection was also accepted (94%, 396 out of 421). With the procedural supports, the users could and did use the CEA procedure even thought they were not skilled in the procedural judgments.

For the AHP, the system helped the user enhance the consistency of the comparison matrix by suggesting which pairwise comparisons to reconsider in order to increase consistency.

These suggestions were also usually accepted (88%, 176 out of 200) and were used successfully to increase consistency of the pairwise comparisons.

The system also provided a tool for comparing the results from different multiattribute techniques. This comparison diagram was frequently used in the evaluation process as recorded in the log of keystrokes and mouse clicks. The users utilized the tool frequently (an average of 9.7 times per user per session) to see the effect of changes within methods as well as to compare the results from different methods.

The sensitivity analysis tools, which show the sensitivity of the current decision to changes of attribute weights, were used less often than we expected. The main reason, we think, is that the diagram was not easy to understand for most of the users. Unlike the comparison diagram, it is not simple enough for the users to comprehend its meaning without further explanation, and most users did not know what to do with the analysis. A more effective communication tool and procedural support for responding to information implied in the sensitivity analyses will apparently be needed.

4. Conclusions

The assessment of PEGASUS showed that in general it accomplishes what it was designed to do. It effectively supported iteration between alternative construction and evaluation. The enhanced integration of basic data display, network design, and multiattribute evaluation in a coherent computer interface gave users greater facility with the whole problem solving process by allowing them to find iterative processes that fit their cognitive styles and the given problem situation. Users were able to improve on a given alternative and they took advantage of procedural expertise provided by the system. PEGASUS thus shows by example that procedural expertise can be identified and that such expertise can be brought to bear in an operational system.

The implementation and assessment of PEGASUS raised a number of implications for future research and development. First, our experience in applying PEGASUS to a specific planning situation suggests that the approach and elements of this system can be incorporated into other planning support systems focused on other problem domains. Second, the evaluation suggested that procedural expertise can and should be extended to other aspects of problem exploration. Alternative generation and sensitivity analysis seem to be the most promising areas in which the development and implementation of procedural expertise might yield the greatest gains.

Acknowledgments: This work was in part coincident with the development of TRAINER (a planning support system for management of training installations) under contract to the US Army Construction Engineering Research Laboratory. Douglas Johnston provided important feedback throughout the development and assessment of this system. T. John Kim was particularly helpful in the implementation of the transportation data and modal split modeling.

References:

Alexander, E. R. 1992. *Approaches to Planning: Introducing Current Planning Theories,* *Concepts and Issues*. 2nd ed. Philadelphia: Gordon and Breach Science Publishers.

Black, A. 1990. The Chicago Area Transportation Study: A Case Study of Rational Planning. *Journal of Planning Education and Research* 10 (1):27-37.

Brill, E. D., Jr., J. M. Flach, L. D. Hopkins, and S. Ranjithan. 1990. MGA: A Decision Support System for Complex, Incompletely Defined Problems. *IEEE Transactions on Systems, Man, and Cybernetics* 20 (4):745-757.

Bryson, J. M. 1995. *Strategic Planning for Public and Nonprofit Organizations: A Guide to Strengthening and Sustaining Organizational Achievement*. revised ed. San Francisco: Jossey-Bass.

Bunn, D., and G. Wright. 1991. Interaction of Judgmental and Statistical Forecasting Methods: Issues and Analysis. *Mangement Science* 37:501-518.

Cohen, P. R., and A. E. Howe. 1989. Toward AI Research Methodology: Three Case Studies in Evaluation. *IEEE Transactions on Systems, Man, and Cybernetics* 19:634-646.

Dickey, J. W. 1983. *Metropolitan Transportation Planning*. 2nd ed. New York, NY: McGraw-Hill.

Hammond, J. S., R. L. Keeney, and H. Raiffa. 1999. *Smart Choices: A Practical Guide to Making Better Decisions*. Boston, MA: Harvard Business School Press.

Harris, B., and M. Batty. 1993. Locational Models, Geographic Information and Planning Support Systems. *Journal of Planning Education and Research* 12 (3):184-198.

Hopkins, L. D. 1973. Design Method Evaluation: An Experiment with Corridor Selection. *Socioeconomic Planning Sciences* 7:423-430.

Hopkins, L. D. 1984. Evaluation of Methods for Exploring Ill-Defined Problems. *Environment and Planning B: Planning and Design* 11:339-348.

Hopkins, L. D., E. D. Brill, Jr., and B. Wong. 1982. Generating Alternative Solutions for

Dynamic Programming Models of Water Resources Problems. *Water Resources Research* 18:782-790.

Johnston, D. M., and L. D. Hopkins. 1994. TRAINER: A System for Training Requirements Assessment and Integration with Environmental Resources, contract report to U.S. Army Corps of Engineers Construction Engineering Research lab. Urbana: University of Illinois at Urbana-Champaign.

Keen, P. G. W., and M. S. Morton. 1978. *Decision Support Systems: An Organizational Perspective*. Reading, MA: Addison-Wesley.

Keeney, R., and H. Raiffa. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley & Sons.

Klosterman, R. E. 1997. Planning Support Systems: A New Perspective on Computer-Aided Planning. *Journal of Planning Education and Research* 17 (1):45-54.

Lai, S.-K., and L. D. Hopkins. 1989. The Meanings of Tradeoffs in Multi-attribute Evaluation Methods: A Comparison. *Environment and Planning B: Planning and Design* 16:155-170.

Lai, S.-K., and L. D. Hopkins. 1995. Can Decision Makers Express Multiattribute Preferences Using AHP and MUT? An Experiment. *Environment and Planning B: Planning and Design* 22 (1):21-34.

Landauer, T. K. 1995. *The Trouble with Computers: Usefulness, Usability, and Productivity*. Cambridge, MA: The MIT Press.

Lee, I. 1993. Development of Procedural Expertise to Support Multiattribute Spatial Decision Making. PhD, Regional Planning, University of Illinois at Urbana-Champaign, Urbana, Illinois.

Lee, I., and L. D. Hopkins. 1995. Procedural Expertise for Efficient Multiattribute Evaluation: A

Procedural Support Strategy for CEA. *Journal of Planning Education and Research* 14 (4):255-268.

O'Connor, M. 1989. Models of Human Behavior and Confidence in Judgment. *International Journal of Forecasting*:159-170.

Patton, C. V., and D. S. Sawicki. 1993. *Basic Methods of Policy Analysis and Planning*. Second ed. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.

Sengupta, K., and T. K. Abdel-Hamid. 1993. Alternative Conceptions of Feedback in Dynamic Decision Environments: An Experimental Investigation. *Management Science* 39 (4):411-428.

Stokey, E., and R. Zeckhauser. 1978. *A Primer for Policy Analysis*. New York: W. W. Norton & Company.

Trybus, T. W., and L. D. Hopkins. 1980. Humans versus Computer Algorithms for the Plant Layout Problem. *Management Science* 26:570-574.

Urban Analysis Group. 1993. TRANPLAN User Manual. Danville, CA: Urban Analysis Group.

Winston, P. H. 1984. *Artificial Intelligence*. 2nd ed. New York, NY: Addison Wesley.

---

[1] The system was implemented on a Sun4 SPARC™ station using the C language and X windows programming libraries, including X library, X toolkit, OSF/Motif™, and BuilderXessory™.

[2] Hammond, Keeney, and Raiffa (1999) reprise the CEA technique under the label "Even Swap Method" as the only method of multiattribute evaluation they recommend. They emphasize a tradeoff strategy of trying to create equivalent hypothetical alternatives so as to make all alternatives perform equally on a particular attribute. This attribute can then be dropped from consideration. After eliminating an attribute, they check for dominance among the alternatives, as in the traditional CEA technique. We have neither incorporated this strategy in our simulations nor investigated whether it is likely to be more efficient. They acknowledge the difficulty of choosing which tradeoffs to make, but claim that it becomes easy with practice. The CEA technique has been used very little,

however, since it was introduced 20 years ago in Stokey and Zeckhauser (1978), which suggests that procedural expertise of the kind we provide is needed to enable use. Hammond et al.'s recent recommendation of CEA confirms that devising such procedural expertise is worthwhile.