

# The LuMPB Key: A Multiple View Interface to Explore High Dimensional Mountain Pine Beetle Simulation Data

Lothar Schlesier  
Faculty of Computer Science  
Otto-von-Guericke University  
Universitätsplatz 2  
39106 Magdeburg, Germany  
lothar@divnull.net

Josie Hughes  
Department of Zoology  
University of Toronto  
25 Harbord Street  
Toronto, Canada M5S 3G5  
josie.shoes@gmail.com

Andrew Fall  
Gowlland Technologies Ltd.  
220 Old Mossy Road  
Victoria, Canada V9E 2A3  
fall@cs.sfu.ca

M. Sheelagh T. Carpendale  
Department of Computer Science  
University of Calgary  
2500 University Dr. NW  
Calgary, Canada T2N 1N4  
sheelagh@cpsc.ucalgary.ca

## Abstract

*In this paper we present the Landscape Unit Mountain Pine Beetle Key (LuMPB Key) - a tool for assessing different mountain pine beetle (MPB) management scenarios on a landscape scale under various conditions. Based on a conceptual model of the simulation data we develop a multiple view interface which provides access to and exploration of the data. Different management scenarios can be compared and a given management scenario can be considered under varying simulation conditions, such as forest type, climate, and degree of initial beetle infestation. On the one hand we use a multiple view interface to provide understandable chunks of information. On the other hand we use it to communicate that different views of the data exist. Views are optionally linked together to support this awareness of possible different ways of looking at this data.*

## 1 Introduction

Driven by the current devastating outbreak of mountain pine beetles (MPB) in the interior of British Columbia (BC), Canada, many efforts are being taken to get more insight to the beetle behavior and to provide forest managers with information on the beetles impact on the pine forests. Com-

puter simulations are an important part of these efforts.

The behavior of MPB in forest stands is relatively well understood, and forest risk rating systems [18] as well as MPB population models for forest stands [15] exist. More recent projects [6, 9, 10] have adopted simulations to the landscape scale using the MPB SELES model [11]. While these models have proven to be useful for the landscapes they were applied to, many more landscape types exist. Modeling and simulating all the different types of landscapes in BC would be an extremely time consuming process, and since the beetle problem is rapidly spreading it is important to make decisions about management scenarios now.

For this reason simulation models which simulate a great variety of typical landscape *patterns*, together with different types of MPB outbreak *patterns*, are being used rather than “real” landscapes in the “Projection of the Efficacy of MPB Management at the Landscape Scale” [13]. For the simulations the Spatially Explicit Landscape Event Simulator (SELES)[7, 8] is being utilized. The models build and simulated with SELES are raster-based, stochastic landscape models in which changes are described as landscape events. The overall behavior of a model is formed by events from individual landscape cells which are one hectare square. A landscape is built from a set of layers, including vegetation cover, climate zone, elevation and more. MPB behavior is projected from a forest stand scale onto the landscape scale

by taking each landscape cell's state and looking up its response in a table of forest stand simulations.

The data generated in the simulations is complex in its volume, dimensionality, and relations. Baldonado et al. propose a set of guidelines when and how to use multiple view interfaces [2] – among them: “Use multiple views when there is a diversity of attributes, ...” and “Partition complex data into multiple views to create manageable chunks and to provide insight into the interaction among different dimensions.”

The aim of the *LuMPB Key* is to provide access to the simulation data so that it can serve as a decision key for landscape managers regarding MPB. We develop a multiple view interface to address the problem of the data's complexity. Our tool allows landscape managers to compare different management scenarios that were simulated on the same set of simulation conditions, as well as to compare results between different simulation sets.

## 2 Related Work

SELES output data can be imagined as a block of spatial-temporal changes stacked on top of each other. The block itself is made of many small component blocks, each one representing the state of a landscape cell at a certain point of time.

Carpendale et al. [4] use lens techniques to probe this block. Moving a probe device into the data block displaces the component blocks with which the spatial-temporal block has been assembled. This gives access to data dynamics represented in the block. Browsing through the data block can, for example, reveal start and spread of a landscape event, such as fire.

The Tardis [3] is a tool used to show spatial-temporal changes in the simulation data produced by SELES. The Tardis provides different interaction techniques to acquire information from the data block. Sampling planes can be moved through the block revealing the time of the emerging of a landscape event, its duration, and its spatial expansion. Another way to access the data block is to open it in a book like manner allowing a person to see two adjacent temporal data pages.

Duta et al. presented VICO [5], a tool that supports visual comparisons of different MPB management approaches in terms of infested area. Two simulations that either are different in the management approach or that are different in the observed time frame can be compared. Maps show the cells that will probably be attacked by the beetles. Color histograms are used to indicate the number of cells to be attacked and the likelihood of attacks.

## 3 Organizing the MPB Simulation Data

### 3.1 Simulation Space

A SELES simulation requires landscape layers as input. A landscape layer holds the start data for a particular data dimension spatially plotted on the landscape map. Some of these layers, such as elevation, are static in the simulation and some such as beetle distribution are dynamic. The landscape patterns used by the MPB simulations have been developed during this project to best reflect typical landscapes in BC [12]. Several data dimensions are of particular importance when considering MPB infestations. These dimensions are listed below. The number in parentheses indicates the length of the dimensions.

- The *percentage of suitable forest* includes any kind of forest with pine trees regardless of age, climate zone, or anything else that influences how well will MPB do. (5)
- The *percentage of highly suitable forest* indicates how much forest is a really good habitat for beetles. (8)
- The *proportion of the landscape forested* indicates how much land is currently forested. This proportion also influences the amount of forest allowed to be logged each year on that landscape unit. (3)
- The *initial infestation* names how much forest is infested by MPB at the beginning of the simulation. (4)
- The *beetle distribution* describes whether beetles are clumped (concentrated in relatively small regions) or dispersed through the forest. (4)
- *Weather* conditions are based on data analyzed over the past 30 years from a variety of weather stations across BC. Throughout a simulation run the weather can be either the average of the observed conditions, or above (slightly warmer than) these conditions, in which case beetles suffer less mortality during the winter. Weather conditions in the real world are subject to high uncertainty – especially over a long period of time. This uncertainty is removed from the simulations, by keeping the weather constant throughout a simulation run. (2)
- The *beetle pressure* describes the ratio of immigrating MPB from neighboring landscape units to emigrating MPB. While a simulation is run for a particular region, this region is surrounded by other regions which may contain many beetles (high beetle pressure) or few beetles. This dimension provides one method of introducing context. (3)

- *Management scenarios* that target the MPB focus management capacities either on areas with new infestations (to possibly stop MPB spreading), or on areas with the highest detected populations (to reduce the MPB population). Different versions of these scenarios exist which vary in the level of effort that is applied. If less aggressive management is able to control beetles, then more intensive strategies are not assessed. Another management scenario targets on salvaging killed wood (to reduce economic losses). Two reference scenarios exist, one does no management at all, the other does usual forest management – cutting old trees in easily accessible areas – without caring about the MPB. (12)

In addition there are meta dimensions of time and replication. Each simulation is run for 20 years in one year time steps. The *replication* dimension results from the fact that each set of conditions is simulated five times in a Monte-Carlo manner.


This description allows us to distinguish between the individual simulation runs. Each combination of parameters from those abstract dimensions that feasibly could represent a landscape condition was simulated. The above list actually represent a hierarchy. That is as choices are made in the order they are listed, each choice affects the available options for the subsequent data dimension.

The first seven dimensions describe the forest state and environmental conditions. A set of parameters on these dimensions will be denoted as a *case*. Within a case different management scenarios are simulated, each five times and for 20 years. The data set consists of 1754 cases in total, while between four to twelve different management scenarios have been simulated in each case. As a management scenario in a case is simulated five times more than 68000 individual simulations were run in total. with each between four to twelve management scenarios have been simulated.

### 3.2 Simulation Output

While these simulations based on landscape patterns can provide information about the amount of pine trees killed, it does not provide information specifying the precise location of the damage on a “real” landscape. For that reason the simulations report the percentage of forest volume for each year then, through use of the replicate runs, can be combined to provide a mean and standard deviation.

Initially pine trees and non-pine trees are standing in the forest. MPBs attack and kill pine trees. The wood of killed trees still can be used in lumber production. Over time the value of the killed trees decreases, so that it becomes less valuable and eventually is no longer usable. Forest management scenarios can cut both living and killed trees.



**Figure 1. Organization of a “case”.**

The simulation output for a case can be arranged as a three-dimensional data block (see Figure 1) with one axis representing the forest volume, one the management scenarios, and one time. Moving a plane perpendicular to the time axis shows the forest attributes over management scenarios for the given year. Moving a plane in the same way along the management scenario axis shows a time series of the forest attributes for a selected management scenario.

## 4 Accessing the Simulation Space

### 4.1 Presentation

We looked at different ways of presenting the long list of simulated cases. As the dimensions that form a case hierarchically depend on each other, tree visualizations such as node-edge trees or tree maps [16] come to mind. Node-edge trees would either require more space than is available on the screen, or considerable user interaction to expand or collapse branches in the tree in order to find the cases of interest.

A tree map could be constructed for all possible cases by the decision hierarchy as indicated in the data dimension list. However, in this tree map, similar cases are not adjacent, and since there are 1754 leaf nodes they tend to be small and as a result hard to read. In addition for both types of tree techniques cases cannot easily be sorted in a way that those with the same value for a certain dimension are displayed side by side.

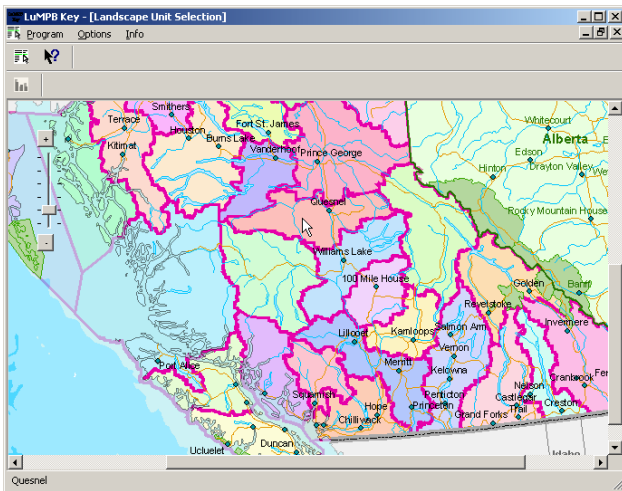
Using a table to display the list of cases allows us to sort them by any dimension (or column of the table). In a table lens [14] like interface continuous values are encoded by length of bars in an attribute’s cell and by position in the cell for discrete values. Filtering for certain attributes can remove rows from the table without problems, as values are

not encoded by the position in the table. Selecting a row increases its height, providing space to display the values in the row's cell as known from the original table lens [14].

## 4.2 Filtering

Our tool provides three ways to apply a filter on the simulation space. The first way is to load a landscape classification file created with a SELES model. This model transforms landscape layers into the abstract notion of suitable forest, highly suitable forest, proportion forested, and optionally (if that data is available) initial infestation as used in the simulations [12]. This allows the use of up-to-date and detailed data as a classification for a landscape unit and as a filter on the simulation space.

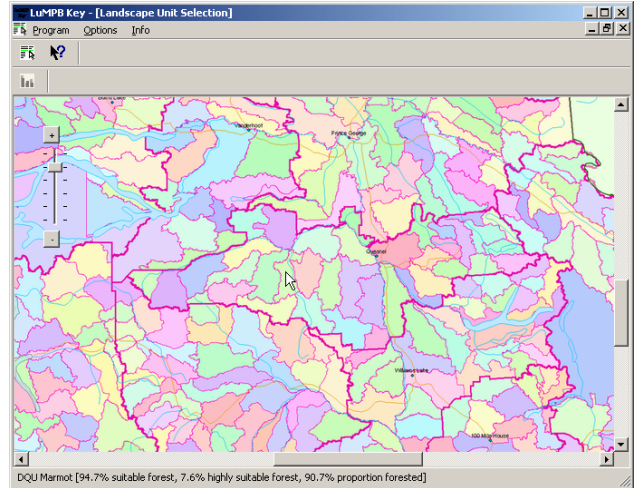
The second way is to select a landscape unit in British Columbia, Canada, from a map interface (see Figures 2 and 3). Each of the landscape units is associated with a classification based on rough observation data with a resolution of 16 hectares, collected in 2005. If new landscape classification data becomes available, it can be fed into the *LuMPB Key*.



**Figure 2. Selecting a district zooms into the map.**

Finally filters can be modified and applied directly on the table lens interface (see Figures 4 5). For the first four columns pull down menus allow the selection of simulated values, any other values can be typed into the menus' text fields. For the other three columns options can be enabled or disabled with check boxes.

After filtering the simulation space, only those cases that are closest to the applied parameters are shown in the table. In the first four columns vertical black lines show the



**Figure 3. Selecting a landscape unit from the map. Moving the mouse over a landscape unit displays its classification in the status bar. Left clicking on a landscape unit applies its classification as a filter on the case list.**

selected values, thus indicating which of the cases may be most relevant. For example in Figure 4 a classification for a landscape unit selected from the map has 94.7% suitable forest. The closest simulated options are 95% and 85%, with the first one being more relevant. Options selected for the last three columns are aligned under the appropriate check boxes.

## 5 Visually Comparing Simulation Results

Our tool provides two general ways to visualize and compare simulation results. The user can compare management scenarios for a single case as well as for different cases. Either way, the window layout, visualizations, and interactions follow the same principles. Bar charts are used to show a single forest attribute while stacked bar charts are used to display the attributes making up the total forest volume. We consider the two types of views on the data blocks introduced in Section 3.2: attribute vs. management scenario and attribute vs. time.

### 5.1 Visualization Components

A *case window* in the *LuMPB Key* is separated into a main area and three context areas (see Figure 6). The main area contains one of the four charts. The context area in the lower left part of the window contains the attribute vs. scenario charts, whereas attribute vs. time charts are located




Figure 4. List of cases after selecting a landscape unit on the map. The filtered list contains 208 cases. Not all can be displayed on the screen. Black lines in the first three columns represent the forest classification.




Figure 5. After applying additional filtering the list is reduced to 20 cases. The user also sorted the list by the initial infestation. Selecting a row increases its height, allowing to print the values into it. To open a case the user presses the button in the last cell of the row.

in the lower-right area. A text browser is located to the right beside the main area and is minimized by default. If the user requests context information about a management scenario or forest attribute, this area is enlarged showing a textual description. The size of all areas can be changed by moving the separators between them.

In addition to all the components of a case window, a *case-comparing window* has a bar full of tabs, each representing a case that is currently opened in the tool. Any number of the currently open cases can be selected, via the check box on its tab, to be displayed simultaneously in the case-comparing window. The currently selected tab determines which case is in focus. Mean value and standard distribution are printed on top of the bars of the focus case. Multiple case comparing windows can be created to allow the comparison of different sets of cases.

Charts show a different level of detail depending on their size. At their maximum size, charts display all components such as user interface components, labels, numeric values on top of the bars, error bars, and a grid. If less space is available, less important components are removed. For example, no values and error bars are shown in the attribute vs. time chart in Figure 6. If even less space is available, the charts display only the bars. Even though very little can be read from such a chart, it creates a sense of awareness that different views on the data exist – especially since changes on charts are linked to each other.

Each chart has an associated context menu that can be accessed by right-clicking on the chart. The context menu allows the user to enable and disable certain visualization options, such as grid lines, error bars, and mean and standard deviation labels.

Furthermore, images of the selected chart or its window can be saved. Along with each image thus saved a web page is created that embeds the image and lists the simulation conditions for that particular visualization. Finally, the context menu has an option to move the selected chart into the main area of the window. Charts can also be moved into the main area using accelerator keys, making it very easy to switch between the different views.

### 5.1.1 The Bar Charts

The bar charts allow the user to view a single forest attribute which has been selected from a drop-down menu. For logged volume it is possible to view either the percentage of annual logged volume or cumulatively logged volume. Either the total forest volume or the total pine volume for all pine-related attributes can be used as a reference (similarly, non-pine volume can be used). This distinction is useful, for example, to see that in a forest with only some pine trees, if they have all been killed, this has little effect on the total forest volume but a huge impact on the forest's

pine population. The attribute and volume reference can be selected from drop-down menus that have been integrated into the description of the chart.

Color is used to encode the forest attribute types (a blue tone for non-pine values, green for living pine trees, etc.), and pattern encodes the attributes state (no pattern for trees standing in the forest, a vertical line pattern for annual logged trees, a crossed line pattern for cumulatively logged trees). The height of the bars encodes the mean value of the displayed attribute over the five replicate simulation runs. Error bars are displayed by default, and mean values and standard deviations are printed directly above the bars.

In the management chart time can be changed using a slider or a spin box. The initial amount of living pine standing in the forest is indicated by a dashed rectangle.

In the time-series chart for a single case, one or more management scenarios can be selected from a tab-bar-like interface. Checked scenarios are visualized in the chart. One scenario can be put into focus by selecting the associated tab. For this scenario, the mean values and standard deviations are printed centered above the corresponding bars. In the case-comparing window, only one management scenario can be visualized at once. The scenario can be selected from a drop-down menu.

We integrated bifocal display techniques [19] in the charts. If the bounds are within a reasonable range, handles for the upper and/or lower bounds of the currently displayed attribute are placed on the x-axis. Moving these along the axis allows the user to expand the size of the bars. Thus, differences between bars can become more apparent.

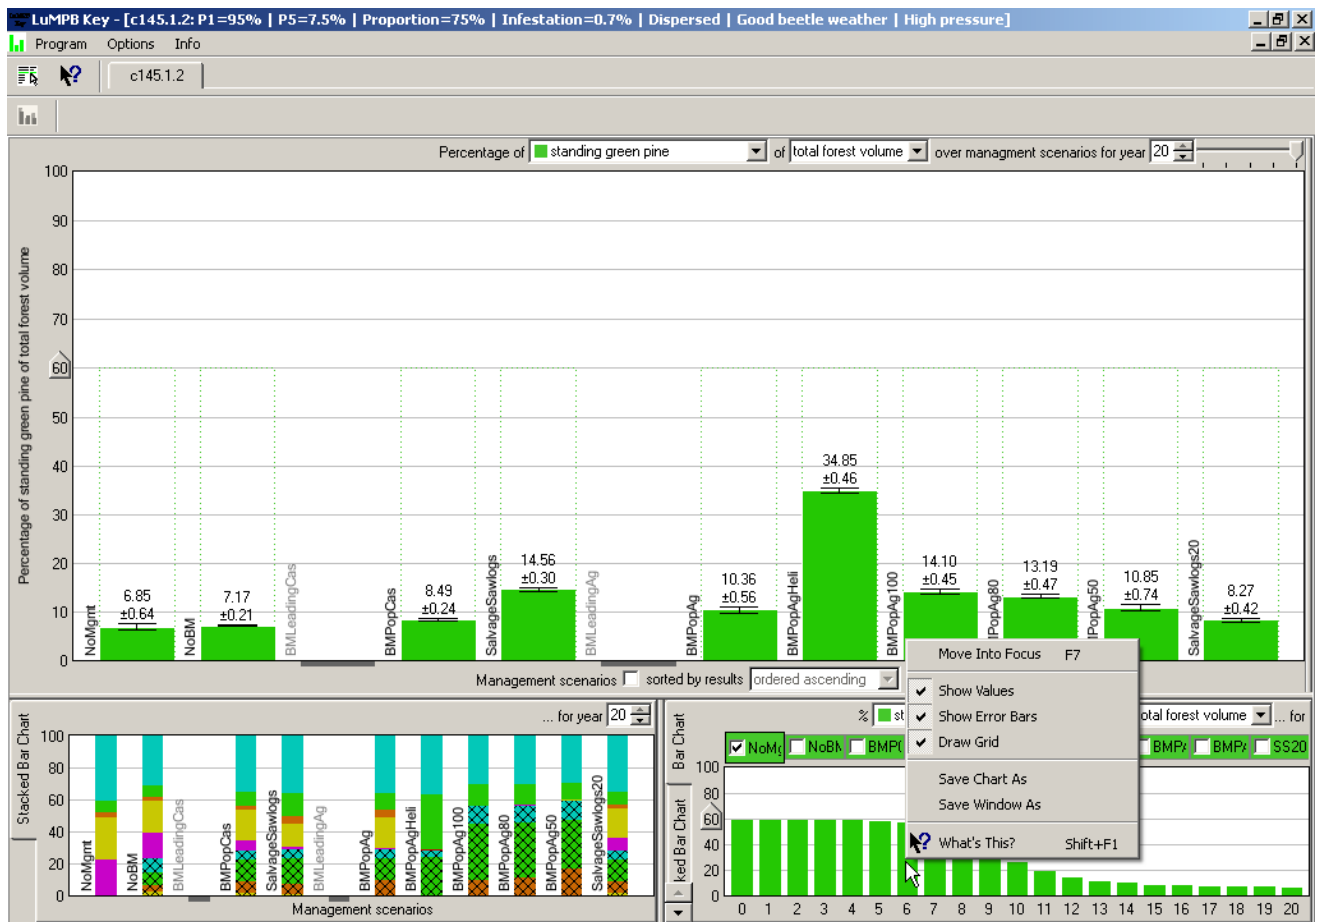
### 5.1.2 The Stacked Bar Charts

The stacked bar charts show the amount of standing and cumulatively logged pine and non-pine volume. The interface is similar to that of the bar charts, except for the attribute selection. The order of stacking of the attributes can be changed using drag-and-drop interactions on the chart's legend. Selecting one (left mouse click) or more (left click + Ctrl) legend items brings them into focus. The charts are divided into three parts: the focus area in the middle, and a context area above and below it (see Figure 7). The focus area can be enlarged using the handles on the x-axis.

## 5.2 Further Options

Further detail-on-demand techniques are included in our tool. Left clicking on a block of bars shows a tool tip printing the values of the selected set of bars (see Figure 7). Descriptions of the management scenarios (see upper right corner of Figure 7) and the forest attributes can be displayed.

Certain properties of charts in one window can be linked together:



**Figure 6. Initial view of the “case window”. Right clicking on a chart shows a context menu that – among other options – allows to move the chart into the main area.**

- Changing the displayed attribute in the management-scenario bar chart, also changes it in the time-series bar chart and vice versa.
- Changing the order or selection of attributes in one stacked bar chart is reflected in the other stacked bar chart.
- Changing the time in one of the management-scenario charts also changes it in the other one.
- Changing the selection of scenarios to be visualized in one time chart also changes it in the other one.
- Changing the focus area in one (stacked) bar chart also changes it in the other (stacked) bar chart.

These connections help to communicate the relationships between the different views on the same data set. By

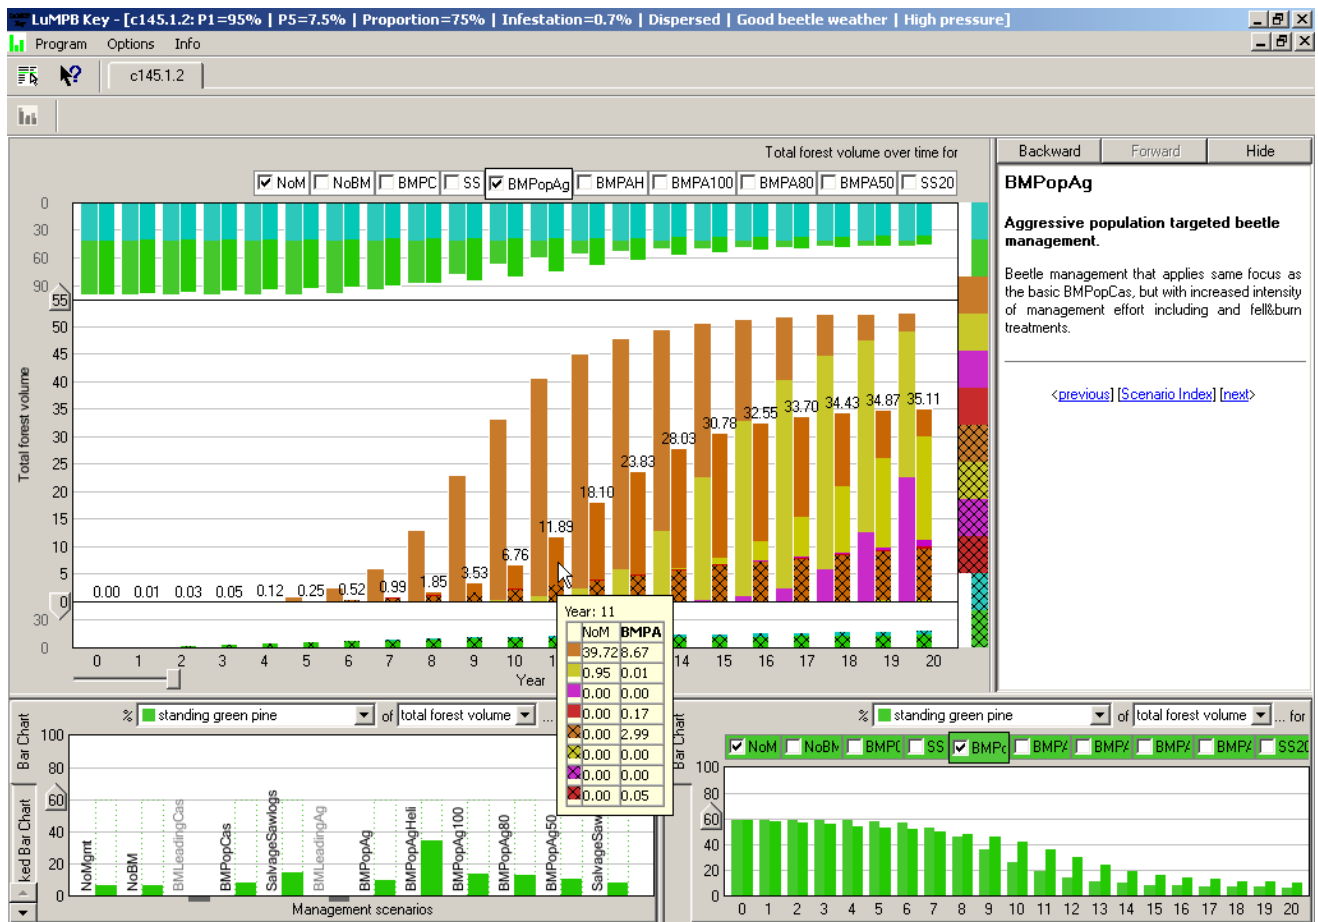
default they are enabled, but the user can disable these connections.

Furthermore, the application’s state – including all opened windows and arrangements made on the charts of each window – can be saved into a session file. This allows a user to save certain views on a data set for later reviewing or for sharing it with other forest managers.

## 6 Walk-through Example

To illustrate the use of the *LuMPB Key*, we will step through an example. Fred is a forest manager who is interested in how MPB would affect the landscape unit for which he is responsible. He decides to select his landscape unit from the tool’s map interface rather than loading a classification file. The map shows the forestry districts of BC (see Figure 2). Clicking on the district in which his landscape unit is located causes the map to zoom in, revealing





**Figure 7. Stacked bar chart comparing the time series for two management scenarios moved into focus. In this chart the attributes combined to cumulative kill are put into focus. Left clicking on a set of bars opens a tool tip showing the mean values of the different bar sections. Right beside the main area is a description of one management scenario.**

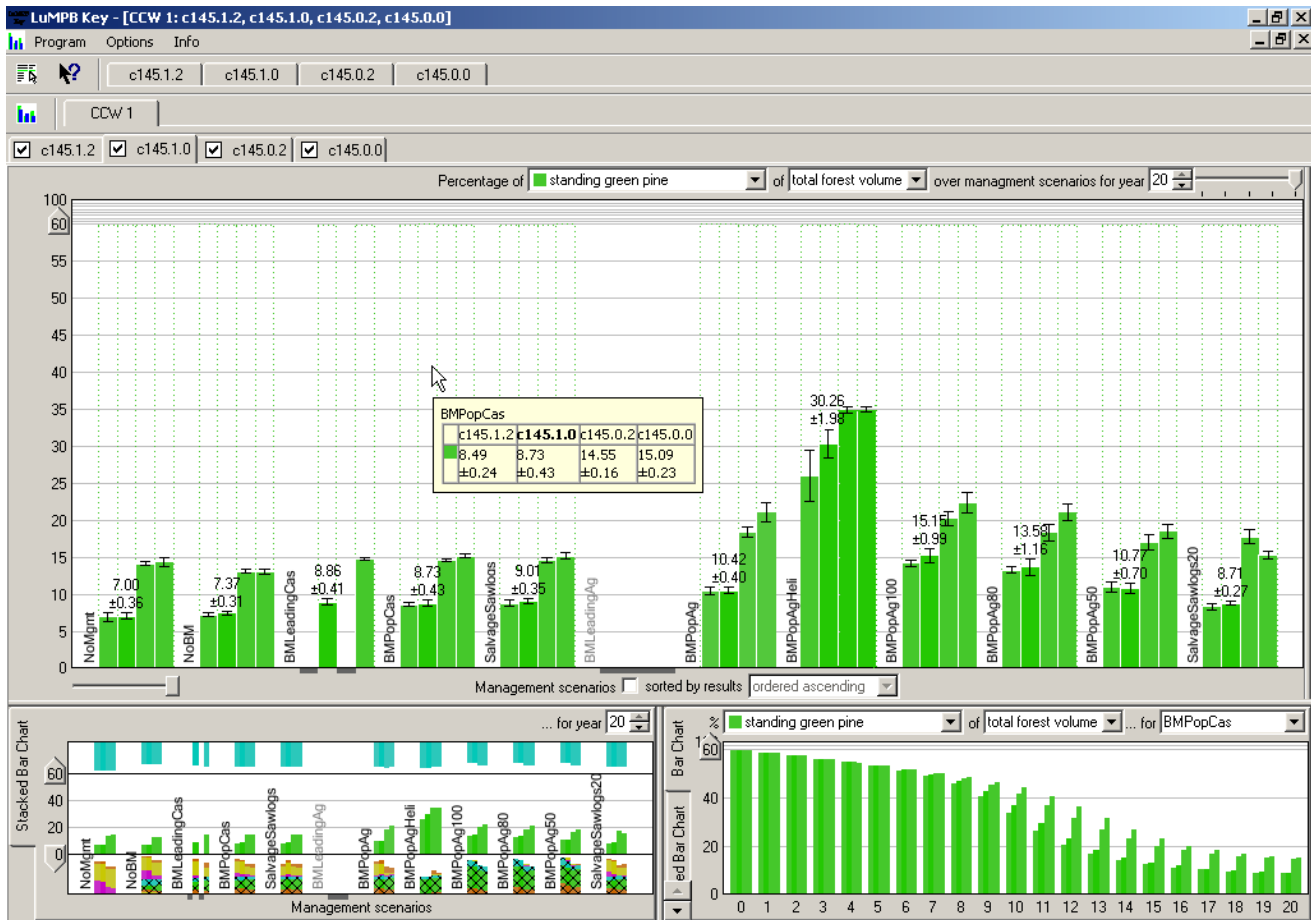
the individual landscape units and allowing him to select the landscape unit of interest. Moving the mouse over the map shows the classification for the unit under the mouse (see Figure 3). Clicking on it switches the view to the case selection table (see Figure 4). The landscape unit's classification is used as a filter on the list of cases, reducing the list from 1754 cases to 208 potentially interesting cases. As his classification has values very close to actual simulated ones for suitable and highly suitable forest, Fred decides to concentrate on those. As the last winters were warmer than in the previous decades he wants to focus on good beetle weather only. Furthermore he disables the simulation option that does not consider beetle immigration and emigration. Applying these criteria reduces the list to 20 cases (see Figure 5).

Fred opens a case with moderate initial beetle infesta-

tion. A new window is created with a bar chart showing the amount of pine trees left living in the forest after 20 years for ten different management scenarios. He immediately can see which management scenario is best in protecting the pine trees under the selected conditions for his landscape (see Figure 6). Fred now wants to see how the forest is affected by beetles over time. He moves the stacked bar chart version of the time series into the main area and brings all the attributes adding up to cumulatively killed volume into focus. He selects an aggressive management scenario and compares it against no management as a reference (see Figure 7). Fred also enlarges the focus area by moving the handles on the x-axis, thus making the differences between bars more apparent.

Now Fred wants to know what would happen under more optimistic conditions. He chooses cases with bad weather





**Figure 8. Comparing four cases across weather conditions and external beetle pressure reveals that for these simulations weather plays a more important role.**

conditions for beetles and equal external MPB pressure. Fred opens a new case-comparing window and enables comparing over all currently open cases.

Figure 8 shows this view, comparing four cases. For this set of cases Fred can see that weather has more influence on the general outcome than the beetle pressure.

In general, case comparing allows a forest manager to see the range of results of management for extreme conditions, and which conditions have more importance.

## 7 Evaluation

Parts of the *LuMPB Key* have been evaluated with information visualization heuristics [21]. For the evaluation three different sets of heuristics have been used: Zuk and Carpendale’s selection of perceptual and cognitive heuristics [20], Shneiderman’s “Visual Information-Seeking Mantra” [17], and Amar and Stasko’s knowledge

and task-based framework [1].

At the time of the evaluation the *LuMPB Key* used a different saturation for each management scenario to distinguish between them. The heuristic “consider gestalt laws” indicated that this can mistakenly indicate a grouping or an order among the scenarios. After the evaluation the tool has been changed to use the same saturation for all management scenarios. The heuristics “integrate text wherever relevant” and “provide details on demand” both revealed that there should be ways to query the exact value of a stacked bar’s component. For components that are not aligned to the chart’s base line, it is hard to read the value from their size. In the version of the *LuMPB Key* presented in this paper this problem has been addressed by implementing tool tips for the charts.

Even though the evaluation focused just on visualization techniques some usability suggestions, such as providing drag-and-drop rearrangement for management scenarios in

the charts were also mentioned.

Beside the mentioned issues the evaluation resulted in favorable comments about the usage of visualization techniques and components. Of particular note in this regard are the tool's ability to save session files, and the variety of methods used to provide different levels of detail.

## 8 Conclusion

Mountain pine beetles are a severe problem in British Columbia – about ten million hectares of forest have been affected by the end of 2005. Using simulations to get a better understanding of the beetle's behavior and to assess different management scenarios has proven to be useful. The huge amount of data created in the simulations is hard to access and makes it difficult to gain information from the data.

In this paper we presented a conceptual model of the simulation data. This model allows us to access the data and to compare data sets based on different aspects. For instance, a user can compare different management scenarios within a case, or simulation results can be compared across several cases.

Based on the conceptual model we implemented the *LuMPB Key*, which has been introduced in this paper. The tool will officially be provided to forest managers in BC shortly. In a first informal presentation at the Pacific Forestry Centre, Victoria, the tool has been favorably accepted.

## Acknowledgements

We are grateful for the funding provided by the Pacific Forestry Centre, Victoria, to support this research. Furthermore we thank Terry Shore, Bill Riel and Marvin Eng for fruitful discussions and support. Finally we thank Petra Neumann, Marc S. Hancock, and Jim Young from the Innovations in Visualization and Interaction Laboratory at the University of Calgary for helping in editing this paper.

## References

- [1] R. Amar and J. Stasko. A Knowledge Task-Based Framework for Design and Evaluation of Information Visualizations. In *Proc. of IEEE InfoVis*, pages 143–149, Los Alamitos, USA, 2004. IEEE Press.
- [2] M. Q. W. Baldonado, A. Woodruff, and A. Kuchinsky. Guidelines for using multiple views in information visualization. In *Advanced Visual Interfaces*, pages 110–119, 2000.
- [3] M. S. T. Carpendale, D. J. Cowperthwaite, M. Tigges, A. Fall, and F. D. Fracchia. The Tardis: A visual exploration environment for landscape dynamics. In *Proceedings of SPIE Conference on Visual Data Exploration and Analysis VI*, volume 3643, pages 110–119, 1999.
- [4] M. S. T. Carpendale, A. Fall, D. J. Cowperthwaite, J. Fall, and F. D. Fracchia. Case study: Visual access for landscape event based temporal data. In *VIS'96: Proceedings of the Seventh IEEE Conference on Visualization*, pages 425–428. IEEE Computer Society and ACM, 1996.
- [5] A. C. Duta, M. S. T. Carpendale, and K. Barker. Vico: A tool for supporting visual comparisons of different pine-beetle management approaches. In *Proceedings of IEEE Conference on IEEE Information Visualization IV'02*, pages 234–242, 2002.
- [6] A. Fall, M. Eng, T. Shore, L. Safranyik, B. Riel, and D. Sachs. Mountain pine beetle audit project: Kamloops forest district landscape model. final documentation. BC Ministry of Forests internal report, 2001.
- [7] A. Fall and J. Fall. Seles: A spatially explicit landscape event simulator. In *Proceedings of the NCGIA Third International Conference on GIS and Environmental Modeling*, pages 104–112, Santa Barbara, CA, USA, 1996. National Center for Geographic Information and Analysis.
- [8] A. Fall and J. Fall. Beauty and the beast: Separating syntax from semantics for models of landscape dynamics. In *Modelling of Complex Systems Conference, New Orleans, LA*, 1998.
- [9] A. Fall, D. Sachs, T. Shore, L. Safranyik, and B. Riel. Application of the mpb/seles landscape-scale mountain pine beetle model in the lakes timber supply area. final report. BC Ministry of Forests internal report, 2002.
- [10] A. Fall, D. Sachs, T. Shore, L. Safranyik, and B. Riel. Application of the mpb/seles landscape-scale mountain pine beetle model in the morice timber supply area. final report. BC Ministry of Forests internal report, 2003.
- [11] A. Fall, T. Shore, L. Safranyik, B. Riel, and D. Sachs. Integrating landscape-scale mountain pine beetle projection and spatial harvesting models to assess management strategies. In T. L. Shore, J. E. Brooks, and J. E. Stone, editors, *Mountain Pine Beetle Symposium: Challenges and Solutions*, pages 114–132. Canadian Forest Service - Pacific Forestry Centre, 2003.
- [12] J. Hughes, A. Fall, T. Shore, and B. Riel. Generating realistic forested landscapes to compare alternative management strategies. *to appear in Landscape Ecology*, 2005.
- [13] Project website, 2006. URL: [http://mpb.cfs.nrcan.gc.ca/research/projects/2-01\\_e.html](http://mpb.cfs.nrcan.gc.ca/research/projects/2-01_e.html).
- [14] R. Rao and S. K. Card. The table lens: merging graphical and symbolic representations in an interactive focus + context visualization for tabular information. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 318–322. ACM Press, 1994.
- [15] L. Safranyik, H. Barclay, and W. G. Riel. A population dynamics model for the mountain pine beetle, *dendroctonus ponderosae*, hopk. (coleoptera: Scolytidae). Pacific Forestry Centre Information Report BC-X-386, Victoria, BC, Canada, 1999.
- [16] B. Shneiderman. Tree visualization with tree-maps: A 2-D space-filling approach. *ACM Transactions on Graphics*, 11(1):92–99, 1992.

- [17] B. Shneiderman. The eyes have it: A Task by Data Type Taxonomy for Information Visualizations. In *Proc. of the IEEE Symposium on Visual Languages*, pages 336–343, Washington, Sept. 3–6 1996. IEEE Press.
- [18] T. Shore and L. Safranyik. Susceptibility and risk rating system for the mountain pine beetle in lodgepole pine stands. Pacific and Yukon Information Report BC-X-336, 1992.
- [19] R. Spence and M. Apperley. Database navigation: An office environment for the professional. In *Proc. ACM Conf. Human Factors in Computing Systems, CHI*, pages 43–54, 1982.
- [20] T. Zuk and M. S. T. Carpendale. Theoretical Analysis of Uncertainty Visualizations. In *Proc. SPIE & IS&T Conf. Electronic Imaging, Vol. 6060: Visualization and Data Analysis 2006*, 2006.
- [21] T. Zuk, L. Schlesier, P. Neumann, M. S. Hancock, and M. S. T. Carpendale. Heuristics for information visualization evaluation. In *Proc. ACM of the Workshop on Advanced Visual Interfaces.*, 2006. To appear.