Automation and the Map Label Placement Problem: A Comparison of Two GIS Implementations of Label Placement

Jill Phelps Kern

Department of Geography
Pennsylvania State
University
jill.kern@comcast.net

Cynthia A. Brewer

Department of Geography
Pennsylvania State
Universty
cbrewer@psu.edu

The placement of feature name labels on maps has challenged mapmakers throughout history. Before the development of mapping software, placing labels in manual map production could consume up to half or more of overall map production time. This paper explores the extent to which current GIS software can place labels legibly, without overlap, and with good visual association between features and labels. This evaluation takes place in the context of a densely featured municipal sewer utility map book. The primary research objective is to evaluate the ability of current GIS software to automate label placement; the research also identifies factors that make manual refinement of automated label placement necessary in order to complete the labeling process. The research compares map-labeling tools from ESRI™'s ArcMap™ 9.2: the Standard Labeling Engine and the Maplex™ labeling extension. Label placement success is assessed by both quantity and quality metrics, using a methodology developed and tailored specifically for evaluation of sewer map label placement. The research found that Maplex placed almost seven percent more labels overall than the Standard Labeling Engine. For the labels they did place, both products provided equally good quality label placement: About 93 percent of labels were placed with no overlap, and virtually 100 percent of labels were placed in their preferred position. After conversion to annotation, manual label position refinement eliminated all overlaps but at the cost of a nine percent decline in the preferred position metric.

Key words: Map label placement, automated label placement, utility map labeling, map design, GIS mapping.

INTRODUCTION

The placement of feature name labels on maps has posed a significant challenge throughout cartographic history, consuming up to 50 percent or more of overall map production time (Yoeli 1972). Different types of features–point, line and area–involve different labeling challenges. Point features have received specific attention because of the difficulty of placing labels legibly and without overlap in densely featured maps while still maintaining unambiguous visual association of labels with their features (Hirsch 1982; Wu and Buttenfield 1991; Christensen, Marks, and Shieber 1995).

The literature on automation of map label placement presents three primary themes: rules for label placement, the development of automated algorithms applying those rules to maps, and the measurement of label placement quality when employing automated algorithms. Each of these

themes is examined in turn in the sections that follow. The paper then describes research evaluating two current GIS tools for high-quality labeling in the context of production of a densely featured municipal sewer utility map book. The primary research objective is to evaluate automated label placement in commercial, off-the-shelf GIS software. The research also identifies factors that make manual refinement of automated label placement necessary in order to complete the labeling process.

Label Placement Rules

Imhof (1962, 1975) compiled an initial set of map label placement rules that have influenced the development of label placement automation from its earliest days. Imhof's guiding principles included legibility, clear association between labels and the features to which they apply, and avoidance of overlapping or obscuring other labels or other map features, while also satisfying aesthetic criteria. Imhof provided separate, specific rules and examples for point, line, and area features, as well as overall design considerations.

Freeman and Ahn (1987) revisited Imhof's rules for placing labels on maps with an eye to automating map annotation from a rule-based, expert system perspective. An expert system is "an artificial intelligence application that uses a knowledge base of human expertise . . . and a set of algorithms or rules that infer new facts from knowledge and from incoming data . . . to aid in solving problems. The degree of problem solving is based on the quality of the data and rules obtained from the human expert" (Howe 1996).

Freeman and Ahn expanded upon and modified Imhof's rules, creating a system of name placement rules applicable to automated labeling of point, line, and area features, while recognizing that an automated system must also allow flexibility in modifying or defining additional rules specific to particular applications. Area feature name placement being most difficult, Freemen and Ahn placed those labels first, followed by point and then line feature labels. They also permitted backtracking to resolve any overlaps that may have arisen.

Wood (2000) also extended Imhof's rules, supplying detailed rationale and sample illustrations regarding name placement for a variety of specific feature types not addressed by Imhof. For example, where Imhof provided general guidance on placement of areal feature names, Wood extended the discussion with specific suggestions for labeling lakes, islands and island groups, capes and points, channels, gulfs, bays, and lagoons.

Yoeli (1972) proposed a scheme to prioritize eight potential positions of labels around a point feature, with top priority for placement above and to the right of a feature (Figure 1). This is now considered standard practice in cartography texts (e.g., Robinson et al. 1995; Dent 1998; Slocum et al. 2005). Yoeli referenced Imhof and others as providing useful directions for placing point feature labels but does not offer specific support for his label placement position priorities. Imhof himself simply referenced Krummil and Eckert: "Where space allows, it is best to have the name beginning to the right of the symbol or sign" (Imhof 1975, 131). Freeman and Ahn (1987) provide the clearest rationale for preferring placement above and to the right:

> It is preferable for the name to read away from the feature (e.g., for the first character to be the one closest to the feature), as this achieves the closest possible association between the feature and its name. Since in the English language there are more ascenders than descenders, it is preferred for a name to be placed above rather than below a feature.

"The primary research objective is to evaluate automated label placement in commercial, off-the-shelf GIS software."

"Imhof (1962, 1975) compiled an initial set of map label placement rules that have influenced the development of label placement automation from its earliest days."

"Freeman and Ahn expanded upon and modified Imhof's rules, creating a system of name placement rules applicable to automated labeling . . ."

"Yoeli (1972) proposed a scheme to prioritize eight potential positions of labels around a point feature, with top priority for placement above and to the right of a feature. This is now considered standard practice in cartography texts."

This suggests that the most desirable position of a point-feature name is to the right and slightly above the feature. (132)
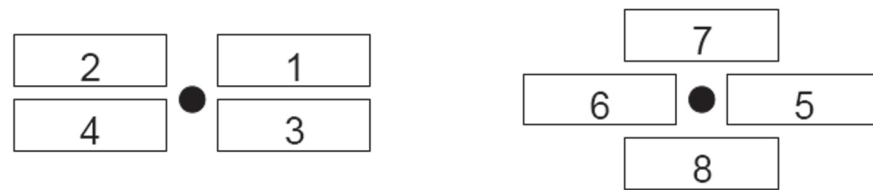

Figure 1. Discrete point-feature label position priorities (Yoeli 1972).

Wu and Buttenfield (1991) revisited Yoeli's point feature label placement rules to determine whether the label positions and priorities on which they were based were valid, through examining the label placement practices of three road map publishers. They found that only four of the eight Yoeli positions were used by a majority of labels and concluded that the complexity of the many other map features affecting name placement demands much greater flexibility in label positioning than allowed in Yoeli's system.

Guidance is also provided in the literature regarding sizes of lettering used for map labels. Shortridge (1979) showed that map readers can reliably discriminate between font sizes which vary by 34 percent or more, or differing by 2 to 2.5 points at typical map label sizes. For example, 10-point type is about 34 percent larger than 7.5-point type. Shortridge also found that linework interrupted for lettering does not reduce the ability to discriminate type sizes and that providing a window or mask around letters superimposed on graphic patterns enables map readers to maintain their ability to distinguish font size differences.

"Shortridge (1979) showed that map readers can reliably discriminate between font sizes which vary by 34 percent or more . . ."

Sadahiro (1995) applied visibility and legibility ratios to measure the loss of information resulting from varying label font size, providing "a basis for determining the size of labels to be used in a GIS" (Sadahiro 1995, 39). The study identified ratio values of 95 percent for visibility and 90 percent for legibility as being desirable. A 95 percent visibility ratio indicates that 95 percent of the letters of the labels are visible on-screen (not lying off-screen on the GIS display); a 90 percent legibility ratio means that 90 percent of the labels are not overlapping other labels. Sadahiro's legibility ratio is a useful guide for determining an acceptable threshold of overlap. The visibility ratio seems a more limited metric, however, relating as it does to on-screen displays; for printed maps, labels are generally constrained to be placed entirely within the map's margins.

Algorithms to Automate the Label Placement Process

"Following articulation of rules for label placement, rules-based algorithms were introduced . . ."

Algorithms for automated label placement have developed significantly since exploration of possible approaches began in the early 1970s. Following articulation of rules for label placement, rules-based algorithms were introduced, and they began to evolve toward expert systems. Simulated annealing and genetic algorithms emerged as viable research directions in the 1990s, while exploration of slider-based label placement began around the turn of the millennium. More recent developments have emphasized force-directed methods for ensuring labels are not placed too closely to one another, as well as ways of speeding label placement for dynamic on-screen map displays. The following sections explore each of these developments in turn.

Early Research

Early research into automation of label placement on maps looked into the feasibility of automation and explored the potential for development of appropriate algorithms and their implementation in software. Yoeli (1972) proposed basic logic to automate the placement of labels for point, line, and area features. Hirsch (1982) addressed point feature label placement in the context of Yoeli's three-step label placement process: selection, layout, and final placement. Hirsch simulated the label layout phase with an algorithm that sought to place names according to Imhof's rules. The algorithm allowed feature names to be placed in any position around the circumference of a point feature, in contrast to Yoeli's approach of limiting label placement to eight specific positions around each feature. The Fortran algorithm Hirsch developed for the simulation implemented an iterative process to resolve label conflicts and demonstrated the feasibility of automating point feature label placement. This work is now seen as a precursor to the force-directed approach to the slider model, which will be discussed in a later section.

Zoraster (1986) took an optimization approach to point label placement. He developed an automated algorithm which used integer programming to resolve label overlaps and applied it to petroleum industry basemaps containing both point and line features. Five iterations on a map with 273 point labels were required to resolve 170 pairwise overplots; the algorithm also successfully placed over 2000 labels, resolving more than 700 conflicts.

Rule-Based Label Placement Automation

Freeman and Ahn (1984, 1987) compiled rules for label placement and developed a rule-based "expert" automatic name placement system in Fortran, called AUTONAP. Testing on small-scale maps produced acceptable results while falling short of the quality that a skilled cartographer could achieve. The program placed labels for area features first, then point, and finally line features without backtracking, which the authors identified as a limitation on its ability to effectively label high-density maps with high quality. In a production setting, a small amount of interactive post-editing was expected to be required.

Jones (1989) and Cook and Jones (1990) extended Freemen and Ahn's approach by applying a logic programming language, Prolog, to the label placement problem, with the goals of "maintaining clear graphic association and avoidance of overlaps" (Jones 1989, 46). This type of language provided the ability to implement a set of rules for name placement, identifying trial positions for labels and resolving conflicts among these. The language included an inference mechanism that sought a solution to satisfy all rules, in contrast to other approaches using, for example, iterative techniques as in Hirsch (1982) or Zoraster (1986). Jones was able to reduce overall processing time using Prolog, creating priority-order lists of potential positions for each name in advance of final placement

As noted above, in the 1980s rule-based systems such as these began being referred to as expert systems. Zoraster (1991) disputed this appellation due to the lack of inclusion of knowledge engineering and the fact that mathematical optimization techniques can substitute for rule-based approaches to the map label placement problem. Among those applying the term *expert system* were Ebinger and Goulette (1990), who reported on the automated name placement system used for the 1990 US Census. For the Census to produce an estimated 1.3 million map sheets under strict deadlines with limited resources, they required a non-interactive approach. Because development began in 1985 before rule-based approaches had

"Yoeli proposed basic logic to automate the placement of labels for point, line, and area features."

"Freeman and Ahn compiled rules for label placement and developed a rule-based "expert" automatic name placement system . . ."

developed sufficiently, the Census Bureau instead employed a completely automated process coded in Fortran 77, a non-recursive (no backtracking or iterations) language. The software used point, line, and area feature placement algorithms to place names according to a hierarchy of label-type importance, testing alternative placements sequentially to find the first non-overlapping position.

Doerschler and Freeman (1992) continued applying rule-based systems to maps of increasing feature density. The three-part system they described included:

- a map database containing all map data to be processed and labeled;
- a rule database containing the order of and rules for name placement, as well as placement quality measurements; and
- a rule processor which applies the rules, tests the aesthetic quality of the resulting placement, and continues applying additional rules and placements as needed to achieve acceptable quality.

The system was implemented in Fortran 77 and was able to place over 2000 characters, labeling half of the 400 features on a 1:19,500-scale street map of Troy, New York. On a denser 1:875,000-scale regional map of Central New York State, the program placed approximately 18,000 characters.

Mower (1993) applied the emerging technology of parallel computing to the map label placement problem, developing a point-feature label placement algorithm for the CM-2 massively parallel computer from Thinking Machines, Inc. He sought to overcome the lengthy execution times of labeling algorithms for large data sets by assigning each name or feature its own processor. He found that increasing feature density became the driving factor in increasing execution time, rather than increases in the overall number of labels to be placed, as with serial-processing systems.

Simulated Annealing

Christensen, Marks, and Shieber (1994, 1995), Edmondson et al. (1996), and Zoraster (1997) applied simulated annealing to the label placement problem. Earlier methods took an iterative or recursive approach to finding locally optimal label placement but did not allow for temporarily worse label placement in order to find a globally better solution. Simulated annealing is defined by the National Institute of Standards and Technology (NIST) as "a technique to find a good solution to an optimization problem by trying random variations of the current solution. A worse variation is accepted as the new solution with a probability that decreases as the computation proceeds" (Black 2004).

Christensen, Marks, and Shieber (1994) found that a test map with 120 point features resulted in forty-two conflicting labels using iterative local improvement, versus two conflicting labels using simulated annealing. They concluded that this method finds better results at all label densities and provides competitive execution times as well. Christensen et al. (1995) proposed methods based on discrete gradient descent and simulated annealing, and compared these and other existing labeling algorithms. Gradient descent is defined by the authors as choosing "from among the set of available operations the one that yields the most immediate improvement" (Christiansen, Marks, and Shieber 1995, 209), whereas the term "operation" refers to the placement of a single label. Gradient descent repeatedly applies the most-immediate-improvement operation to significantly improve the original trial label placement solution. The authors concluded that simulated annealing was preferred when the quality of the labeling solution was important, and that it was also a relatively easy algorithm to implement.

"Christensen, Marks, and Shieber, Edmondson et al., and Zoraster applied simulated annealing to the label placement problem."

"The authors concluded that simulated annealing was preferred when the quality of the labeling solution was important, and that it was also a relatively easy algorithm to implement."

Edmondson et al. (1996) sought to overcome the fact that the more powerful, recently developed name placement algorithms were too inefficient for production use by proposing a general algorithm that combines "expert" cartographic rules with effective label placement optimization. The rules were summarized into a scoring function to evaluate the quality of alternative individual label placements before overall placement is optimized using simulated annealing. Zoraster (1997) followed this proposal by applying simulated annealing to the oil field base maps on which he previously (1986) had tested an integer programming approach. He found that simulated annealing both computed results more quickly and resulted in fewer deletions to resolve overlapping labels.

Wagner et al. (2001) provided an approach to label placement that was independent of feature type and of label size and shape. Their algorithm first applied a series of rules in order to label as many features as possible while reducing candidate-label sets for those remaining; it then reduced the number of candidates to a maximum of one per feature. A comparison of this approach with five other methods, using datasets up to 3000 points, showed that their rules-based method was equivalent to simulated annealing in quantity of labels placed, but much faster, showing potential for application to fast Internet labeling.

Slider-Based Label Placement Algorithms
Van Kreveld, Trijk, and Wolff (1999) took the approach of relaxing the Yoeli-style requirement of limiting point feature label positions to a few fixed locations, allowing continuously sliding labels. The algorithm's objective was to optimize the number of points receiving non-overlapping labels. In a comparison with Christensen, Marks, and Shieber (1995), Van Kreveld's algorithm "yields almost equal results as simulated annealing for less than 750 points, and is always better beyond 750 points" (Van Kreveld et al., 43), while also running considerably faster.

Kameda and Imai (2003) presented a refined slider algorithm, designed to avoid packing labels so tightly as to be difficult to read by separating labels as much as possible within a continuous labeling space for each point or line feature. With the objective of maximizing the number of labels placed, they found that more labels could be placed using continuous labeling spaces. The authors applied an additional algorithm for labels with leader lines in densely featured areas where there is no labeling space for a particular point.

Force-Directed Label Placement Algorithms
Ebner et al. (2003) developed a force-based simulated annealing algorithm for maximizing the number of labels placed. Their approach "uses repulsive forces between labels … labels are not placed close to each other if possible and the method achieves a good distribution of the labels in the available space" (Ebner et al. 5). This force-directed method is combined with simulated annealing, which allows worse intermediate label placements to avoid being trapped in local minima. The results showed label placement numbers close to optimal, with better label distribution than algorithms that maximize total number of labels placed. The algorithm also solved large problems quickly, with results that "often look similar to those of a human cartographer" (Ebner et al.,11).

Stadler et al. (2006) applied a different two-step approach, using morphological image processing for initial point-feature label placement and an iterative force-directed method for final placement. The first stage involves pixelizing the map and excluding regions around point and line features to avoid overlaps before placing labels to produce an initial, fea-

"A comparison of this approach with five other methods, using datasets up to 3000 points, showed that their rules-based method was equivalent to simulated annealing in quantity of labels placed, but much faster, showing potential for application to fast Internet labeling."

"With the objective of maximizing the number of labels placed, they found that more labels could be placed using continuous labeling spaces."

sible configuration. The second step resulted in placement of most labels closer to their associated point features and farther from other point features and labels. The authors indicated that the use of simulated annealing would further improve the outcome, particularly for label- and feature-dense maps where "the force-directed method might not exhibit enough flexibility to rearrange the labels." (214)

Genetic Label Placement Algorithms
Genetic algorithms are also being applied to the challenge of map labeling. Verner et al. (1997) summarized the approach as follows: "The genetic algorithm operates as an iterative procedure on a fixed size population or pool of candidate solutions. The candidate solutions represent an encoding of the problem into a form that is analogous to the chromosomes of biological systems." (4). Their algorithm outperformed other labeling algorithms, including placing up to 7 percent more labels without overlapping than simulated annealing. Yamamoto and Lorena (2005) apply a variant, constructive genetic algorithm and reported additional improvements of 2 percent over the results of Verner et al.

Van Dijk et al. (2004) examined the proposal that design rules can be applied to the development of competent selecto-recombinative genetic algorithms, in the context of a map-labeling case study. Such algorithms are based on finding (selecting) and combining building blocks and are considered competent if they reach good quality (e.g., 97 percent of optimum), with reasonable (e.g., linear) scale-up of solution time with size of problem. The authors laid out a series of design rules on which they based development of their genetic algorithm, then compared the performance of their algorithm to simple genetic algorithms and other types of labeling algorithms such as simulated annealing, finding that the scale-up behavior of their algorithm matched that predicted by theoretical models.

Label Placement Quality Measures
Van Dijk et al. (1999) reviewed existing rules for map labeling (e.g., Imhof 1975; Yoeli 1972), identified quality criteria relevant to automated label placement, and developed a quality function to measure how well a particular algorithm placed labels on a map. Their four resulting quality parameters were (1) aesthetics; (2) label visibility; (3) feature visibility; and (4) label-feature association. They also provided specific evaluation criteria for each of the four parameters.

- <u>Aesthetics</u> refers to the quality of a line or area label's shape as it follows the shape of the associated feature. More than one inflection point and excessive curvature represent poor quality. Point features, having no curvature, are not evaluated for aesthetics.
- <u>Label visibility</u> refers to how visible a label is given other features and labels in its vicinity. Quality is defined as the percent of the label's text block that is not overlapped by other labels, or by features.
- <u>Feature visibility</u> for line and area features is defined as the percent of its line or area not overlapped by labels, excluding the feature's own label; for point features, any label intersection equals poor quality.
- <u>Association quality</u> defines how clear the association is between a feature and its label. Van Dijk et al. provide separate, increasingly complex criteria for the quality of point, line, and area feature association.

Based on the approach of Van Dijk et al. (1999), Table 1 shows the type of quality criteria used to evaluate the effectiveness of various name place-

""The genetic algorithm operates as an iterative procedure on a fixed size population or pool of candidate solutions. The candidate solutions represent an encoding of the problem into a form that is analogous to the chromosomes of biological systems.""

"Van Dijk et al. reviewed existing rules for map labeling, identified quality criteria relevant to automated label placement, and developed a quality function to measure how well a particular algorithm placed labels on a map. Their four resulting quality parameters were (1) aesthetics; (2) label visibility; (3) feature visibility; and (4) label-feature association."

Research Article	Type of Label Quality Criteria Applied in Article			
	Aesthetics	Label Visibility	Feature Visibility	Association
Hirsch (1982)		Y	Y	Y
Freeman and Ahn (1984, 1987)	Y	Y	Y	Y
Zoraster (1986, 1991)		Y		
Jones (1989)		Y	Y	Y
Ebinger and Goulette (1990)		Y		
Doerschler and Freeman (1992)	Y	Y	Y	Y
Mower (1993)		Y		
Christensen et al. (1994)		Y	Y	
Christensen et al. (1995)		Y		
Freeman (1995)		Y	Y	Y
Sadahiro (1995)		Y		
Edmondson et al. (1996)	Y	Y	Y	Y
Pinto and Freeman (1996)	Y			Y
Zoraster (1997)		Y	Y	Y
Wagner and Wolff (1997)		Y		
Van Kreveld et al. (1999)		Y		
Chirie (2000)		Y	Y	Y
Barrault (2001)	Y	Y		Y
Wagner et al. (2001)		Y		
Huffman and Cromley (2002)		Y	Y	Y
Ebner et al. (2003)	Y	Y		Y
Kakoulis and Tollis (2003)		Y		
Kameda and Imai (2003)		Y		Y
Freeman (2004)		Y	Y	Y
Van Dijk et al. (2004)		Y		
Been (2006)		Y		
Kakoulis and Tollis (2006)		Y		Y
Ribeiro and Lorena (2006)		Y		
Stadler et al. (2006)		Y	Y	Y
Mote (2007)		Y	Y	

Table 1. Types of quality criteria applied in map labeling research articles.

ment algorithms presented in the literature. The most frequently evaluated quality criterion has been label visibility, which in practice became label-to-label overlap. Specific overlap measures employed included number of features labeled without overlap (e.g., Van Kreveld et al. 1999; Ebner et al. 2003; Kameda and Imai 2003), percent of features labeled without overlap (e.g., Sadahiro 1995), and number of conflicting labels (e.g., Christensen et al. 1994). Least frequently evaluated has been aesthetics, an admittedly challenging criterion to measure objectively. One pair of authors also opined, "nowadays there is an increasing need for large, especially technical maps for which legibility is more important than beauty" (Wagner and Wolff 1997, 388). Van Dijk et al. concur, stating that for technical maps "the visibility of labels and a good label-feature association is more important than aesthetics or the visibility of objects that constitute the map background" (62).

Nevertheless, development of labeling quality measures continues. Barrault (2001) proposed a new quality measure for evaluating how well

". . . for technical maps "the visibility of labels and a good label-feature association is more important than aesthetics or the visibility of objects that constitute the map background"".

the shaping as well as the placement of an area map label fits, spreads across, and thus effectively represents that feature. The author, referencing Freeman (1995) and Pinto and Freeman (1996), among others, described criteria for evaluating area label coverage of the feature, a process for computing alternative label support lines (circular arcs) for labels and for assessing the area coverage each support line provides against those criteria. The algorithm was tested in labeling a variety of simple, complex, and extremely long area features, as well as in situations where the algorithm had to work around other obstructing labels, producing legible and aesthetically pleasing automated label placements.

Others continue to pursue different objectives in label placement. Huffman and Cromley (2002) developed a model for applying labels to point features with the goal of placing the maximum number of labels possible. The model incorporated several labeling criteria and allowed for the relative weighting of the criteria. Ribeiro and Lorena (2006) addressed the issue of minimizing label overlap when all features must be labeled. They presented a binary integer linear programming model and examined three constraint relaxation approaches that provided near-optimal solutions to problems up to 1000 points. Kakoulis and Tollis (2006) extended their 2003 framework for automatically labeling any set of graphical features, including maps and other types of diagrams, with specific requirements related to the placement of multiple labels per feature. Applying two different algorithms to circular, symmetric, and orthogonal drawings produced similar results, with their flow method performing faster with same-sized labels; a more flexible, iterative approach performed better for labels of variable size or with strict constraints on order.

Dynamic Map Labeling and the Need for Speed

Whereas much prior research has focused on automating and optimizing label placement quality and/or quantity in static maps, modern dynamic displays demand *speedy* labeling. Freeman (2005) reviewed the evolution of cartographic labeling rules as applied to automated labeling software, closing with a look forward to the challenge of dynamically labeling electronic map displays with pan and zoom.

Been et al. (2006) introduced a series of desired characteristics for providing consistency in dynamic map labeling, as well as a labeling framework to address the additional dimension of scale in the dynamic labeling situation. Their algorithm included a dynamic placement and selection phase in pre-processing and a filtering and display phase during interaction.

Mote (2007) provided a fast, efficient, scalable method for real-time point feature labeling on dynamic maps without pre-processing. Mote subdivided the map space into a trellis structure of rows and columns; each trellis cell was associated with the features within its boundaries, significantly limiting the search for label conflicts. Labels were prioritized, with priorities revised upward as alternative candidate positions for a feature's label were eliminated by other labels. Re-prioritization was also weighted for proximity when zooming to increased view magnification. Resulting labeling speed increased by orders of magnitude over most other approaches in the literature and was up to 10 times faster than the fastest results reported to date while testing up to 130,000 features to be labeled, where previously reported tests examined 20,000 or fewer.

Literature Review Summary

The literature demonstrates ongoing interest in advancing the speed of map label placement along with the quality of the results. Freeman (2007)

"Barrault proposed a new quality measure for evaluating how well the shaping as well as the placement of an area map label fits, spreads across, and thus effectively represents that feature."

"Whereas much prior research has focused on automating and optimizing label placement quality and/or quantity in static maps, modern dynamic displays demand speedy labeling."

surveyed the history of problems faced when attempting to automate map label placement, including leading, key numbering, labeling short or divided roads, as well as special-purpose labeling such as for elevation contours and soil maps. Freeman highlighted a variety of specific automation challenges, concluding that dedicated efforts to address such apparently simple yet computationally complex issues have produced impressive advances over the past quarter century.

Since Yoeli first postulated that map label placement could be automated via algorithms based on cartographic best practices, researchers continue to apply new techniques while building on existing knowledge. In 2006, for example, Stadler applied image-processing techniques along with force-directed methods from the literature of 2003, with the potential to derive further benefit in future work from simulated annealing which emerged in the mid-1990s.

Development has advanced to the point that sophisticated label placement tools are available in commercial mapping software. ESRI's Standard Labeling Engine and Maplex's extension to its Arc™ suite of products represent one family of such offerings. The basic capabilities of Maplex were laid out in a white paper (ESRI 1998) and have continued to be developed throughout ESRI's releases of new versions of ArcGIS.

Benchmark Data

As development of map label placement algorithms has progressed, the use of benchmark data has evolved as well. Early researchers utilized what might be termed internal benchmark data by presenting papers over time that tested increasingly sophisticated algorithms against the same or similar datasets. Zoroaster (1991, 1997) applied different approaches to petroleum industry basemaps, while Kakoulis and Tollis (2003, 2006) utilized test graphics devised by earlier researchers. Standard labeling tasks have also been undertaken, such as labeling 1000 or so US cities with their names (van Kreveld et al., 1999 and Stadler et al., 2006).

More recently, the Internet has enabled fast and easy sharing of standard datasets, such as those used by Ebner et al. (2003) and Ribeiro et al. (2006). Researchers also make use of publicly available Internet-based data (Barrault 2001), which can then be incorporated into other studies. While many researchers generate their own test data, and others (including the authors of this paper) work in the context of locally available real-world data, benchmark data sets are emerging as a way to provide apples-to-apples comparisons between different approaches to labeling.

Research Objectives

The primary goal of this research was to evaluate the extent to which current GIS software can automate the placement of feature labels on densely labeled maps. The research also attempted to identify factors that make manual label placement necessary in order to complete the labeling process and satisfy the map's intended use.

The research focused on the development of a GIS-based sewer utility map book from the Town of Concord, Massachusetts, and compared the labeling capabilities of two map labeling engines found in ESRI's ArcMap 9.2. Real-world data was utilized rather than benchmark data in order to conduct research wherein the findings would be applicable in a practical setting. The balance of this article will review the research context, process, and findings, including recommendations for future research.

"Freeman highlighted a variety of specific automation challenges, concluding that dedicated efforts to address such apparently simple yet computationally complex issues have produced impressive advances over the past quarter century."

"More recently, the Internet has enabled fast and easy sharing of standard datasets . . ."

". . . benchmark data sets are emerging as a way to provide apples-to-apples comparisons between different approaches to labeling."

"The primary goal of this research was to evaluate the extent to which current GIS software can automate the placement of feature labels on densely labeled maps."

Research Context
At the start of this project, 11-by-17-inch, printed copies of the Town of Concord's sewer map book were used daily by the dozen or so members of the Water & Sewer Division crew to identify the correct sewer mains or manholes on which to work. The map books are printed in black and white from 1970's-era drawings on Mylar. Though the drawings are updated manually when new or replacement sewer mains or ties are installed, the crew's paper versions are reprinted only every year or two, when they get too tattered to use. The Water & Sewer Division would like to start printing the map book from the GIS to allow for more timely and accurate updates. Ultimately, the Water & Sewer Division intends for their crews to take a rugged laptop into the field with the latest GIS data available.

Sewer Infrastructure Features and Labels
The sewer infrastructure features on these maps include:
- sewer manholes, which provide access to the underground sewer mains;
- the large sewer main pipes themselves; and
- smaller pipes called sewer service ties, which link buildings into the sewer system.

At least one and as many as four labels must be placed on sewer manholes. Each manhole has:
- A unique facility identification number;
- A station number, which shows the distance in feet down a particular sewer branch in which the manhole is located;
- A rim elevation, which is the elevation of the top of the manhole; and
- An invert elevation, which is the elevation of the bottom of the manhole.

Key attributes to be labeled for the main sewer pipes are:
- Size, which is the diameter of the pipe; and
- Material, or the composition of the pipe, such as vitreous clay or PVC.

For sewer mains, an additional label is required for slope. An arrow indicates the direction of flow, and a numerical slope label must also be placed that reads in the direction of flow (Figure 2). The placement of numerical slope labels reading in the direction of the sewer's flow may be said to defy cartographic labeling conventions, which require labels to be placed more-or-less right side up for legibility. This is, however, standard, accepted, and indeed expected practice within the utility community. The numerical label in this case can be considered a symbol, the placement of which provides a critical visual cue to utility workers: A single glance at the arrow and slope immediately conveys the direction of flow.

The final element of sewer infrastructure is the service tie to each building. The only attribute label for this pipe is the service number, which is the crew's index to all related customer information.

Sewer Labeling Quality Metrics and Evaluation Methodology
The quality metrics for this research project are based on label placement rules from the literature, modified by the needs and expectations of the users of the map–the sewer field crew. Three general types of metrics are employed, relating to:
- Quantity of labels placed;
- Preferred positioning of labels; and
- Overlap.

"The Water & Sewer Division would like to start printing the map book from the GIS to allow for more timely and accurate updates."
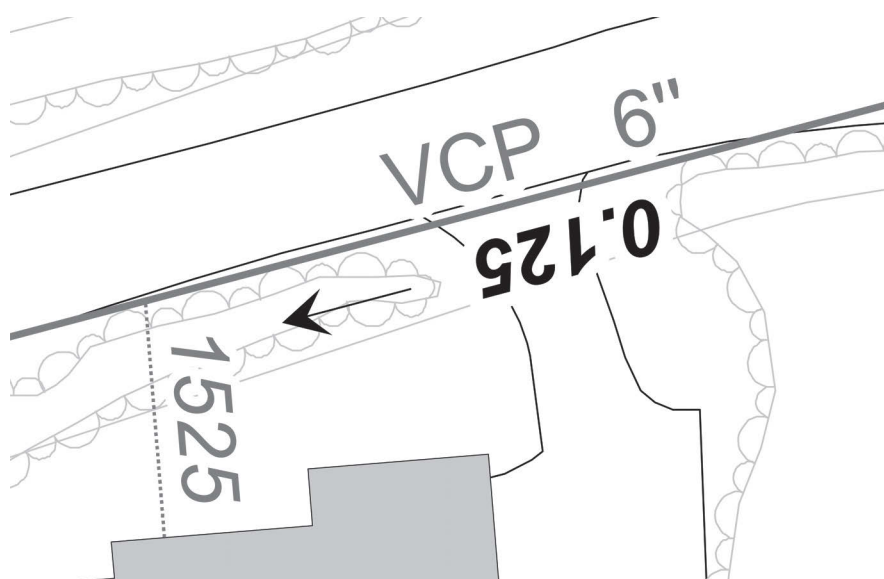
Figure 2. Sewer main with inverted slope label and arrow. (see page 87 for color version)

Regarding the first category, Quantity, the following data are gathered for each iteration of the labeling process:
- Total number of labels placed; and
- The percent of labels placed versus the ideal number (that is, all possible labels needing placement).

In the category of Labels in Preferred Position, each type of sewer infrastructure feature–manhole, main and service tie–has a specific preferred placement criterion:
- <u>Manhole</u>: Up to four labels (facility ID, invert elevation, rim elevation, station number) to be placed in a north-south-east-west configuration around the point feature (Figure 3).
- <u>Main</u>: Up to three labels (size, material, and slope) to be placed parallel to and either above or below the line feature.
- <u>Service Tie</u>: One label (service number) to be placed parallel to and either above or below the line feature.

Finally, two Overlap criteria are also tracked throughout the research process:
- Label-label overlap; and
- Labels overlapping service ties, which must be fully visible to the crew.

As noted earlier, Van Dijk et al. suggest four categories of label placement quality metrics: label visibility, feature visibility, association, and aesthetics. The metrics employed in this project, apart from the quantity measurements, correspond to these categories as follows:
- <u>Label visibility</u> is measured by *label-label overlap;*
- <u>Feature visibility</u> is measured by *labels overlapping service ties;*
- <u>Association</u> is measured by the ability to place labels in their *preferred position* without the use of leader lines.

Aesthetics are not measured explicitly, although the preferred position criteria also reflect the desired aesthetics from the perspective of the field crew using these maps.

During the course of the research and measurement process, the density of the features and labels on the maps involved were such that the same label was often involved in overlapping both a key feature (service tie) and one or more other labels. In consequence, the label visibility and
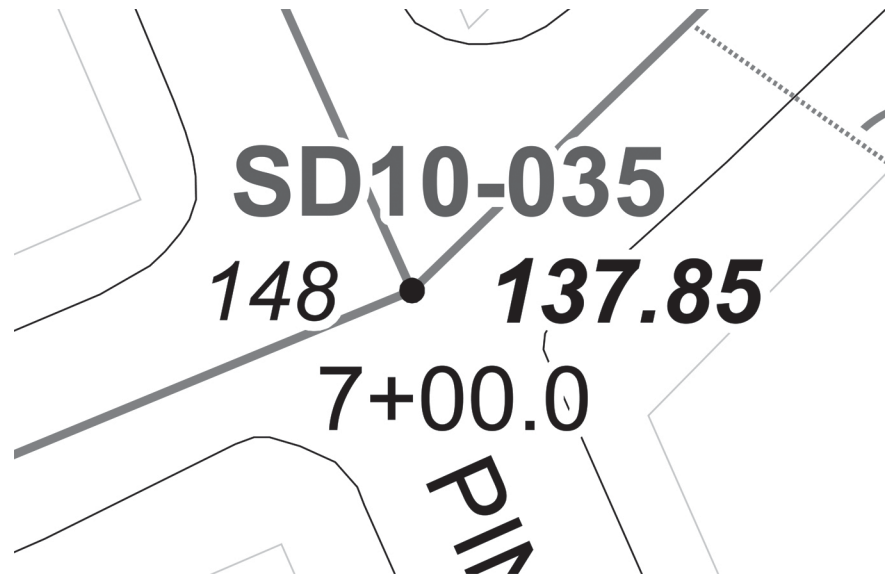
Figure 3. Ideal sewer manhole label positioning. (see page 88 for color version)

feature visibility measures were combined into a single overlap metric, to avoid overstating the degree of overlap through double counting.

- Where one label overlapped one service tie feature, one overlap occurrence was recorded.
- Where multiple labels overlapped each other and/or overlapped service tie features, one overlap occurrence was recorded for each label involved.

From the 121-page sewer map book, three representative pages (Figures 4-6) were selected. Criteria for selection included a range of feature and label density, from low to very high on each page and a variety of orientations of sewer mains to fully exercise the labeling software's capabilities. The streets (and hence the sewer mains–symbolized as red lines–buried beneath them) on page D10 (Figure 4) lie mostly at right angles to one
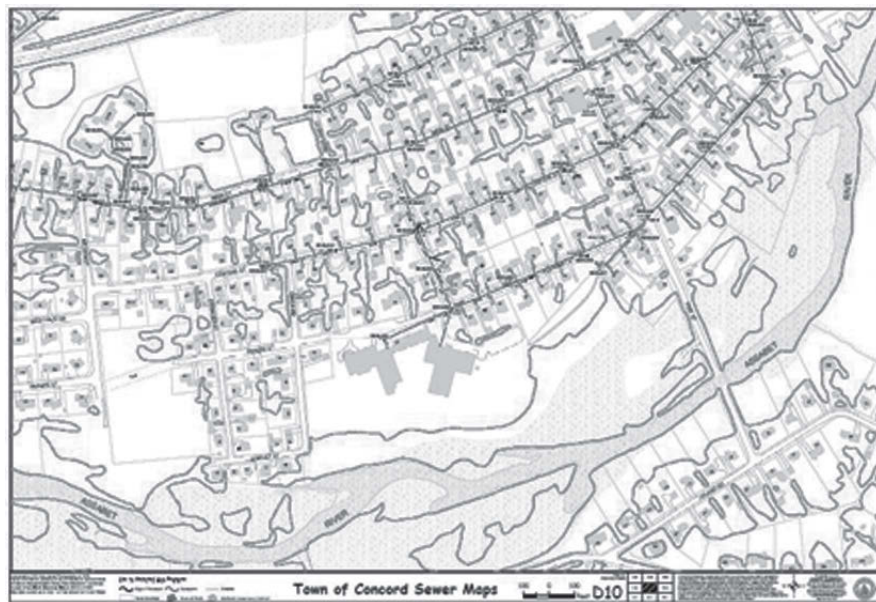

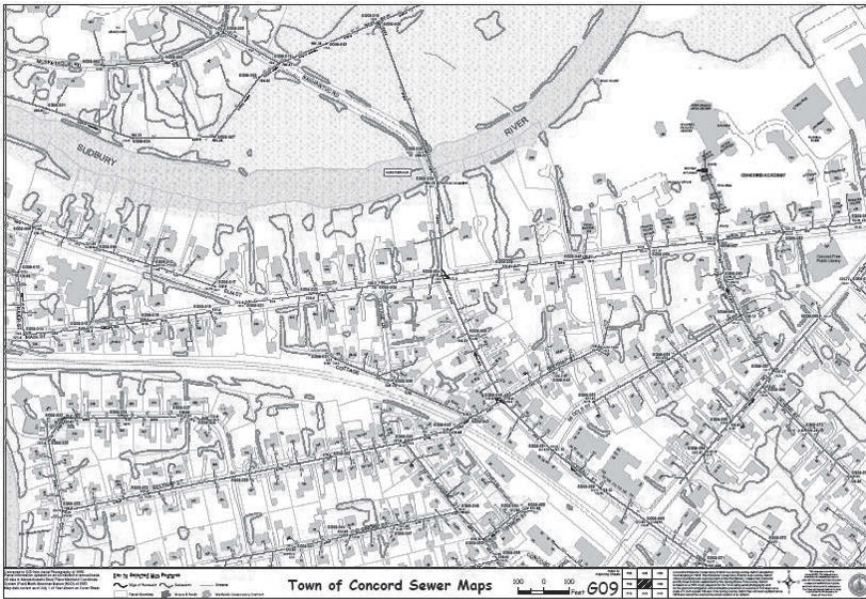*Figure 4. Sewer map book page D10. (see page 88 for color version)*

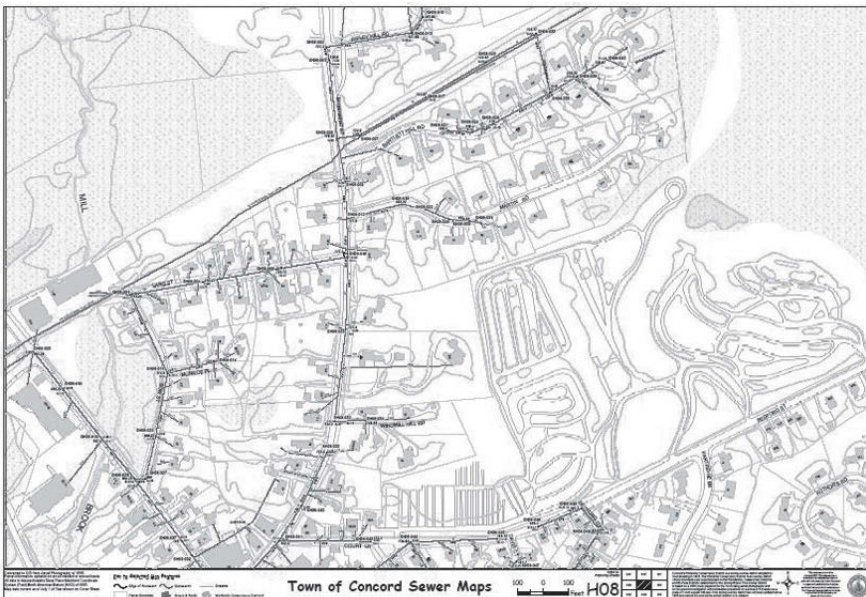*Figure 5. Sewer map book page G09. (see page 89 for color version)*

Figure 6. Sewer map book page H08. (see page 89 for color version)

another, while pages G09 (Figure 5) and H08 (Figure 6) each encompass a wider range of sewer main orientations.

Label density varies considerably across each of these maps. For example, page H08 (Figure 6) contains an area of very low label density near the top of the page: a long straight stretch of sewer main (red line) with widely spaced manholes (black points) and no service ties. In contrast, the label density for the neighborhood just to the south is nearly double, due to more closely spaced manholes and numerous service ties (dashed red lines).

As is evident from the three figures just discussed, these real-life maps contain large areas with no sewer infrastructure, where no labels are to be placed. In order to provide useful label density measurements, den-

sity calculations have been limited to buffered zones around the sewer mains. These "buffer segments" represent the areas where virtually all sewer infrastructure labels will be placed and provide a more consistent baseline for density comparisons than using the entire map area. A buffer width of 80 feet (map units) along both sides of sewer mains encompasses 97.8 percent of sewer infrastructure labels to be placed. Only labels falling within these buffer segments were considered in the density and quality measurements.

For each buffer segment, the total number of feature labels available for placement was calculated. This figure represents the Ideal Label Count for each buffered segment and served as the baseline for the quantity metrics. A total of 1474 labels were available to be placed across these three map book pages, and ideal label densities ranged from 196 to 424 labels per million square feet on the ground, or from 2110 to 4564 labels per square kilometer.

Research Process
The research was conducted in three stages: (1) automated labeling; (2) conversion to annotation; and (3) manual refinement of label placement. The first, automated portion of the process applied two suites of automated ESRI labeling tools to the three selected case study maps in parallel, ArcGIS 9.2's Standard Labeling Engine and its Maplex labeling extension. ESRI's products were utilized for this research primarily due to their ready availability. These products are widely used in GIS applications in a variety of business sectors, including government and utility, and specifically by the Town of Concord. No other automated labeling products were available to the researcher at the time of this study.

After developing initial label engine settings, the research proceeded through four automated labeling iterations. After each iteration, the quality of the labeling results produced by each of the two engines was measured, adjustments were made to the label engines' settings to attempt to improve label placement quantity and quality, and the labeling process was repeated. Once automated labeling results could be improved no further, the higher-quality result was converted to annotation and taken forward into the final, manual refinement stage of the process. Details of these iterations and their impact on metrics are presented next.

Automated Label Placement Iterations
Initial label engine settings for the Standard and Maplex labeling engines are shown in Tables 2 (Standard) and 3 (Maplex), as are settings for all subsequent automated iterations; results for this and all subsequent iterations appear in Table 4. The objective of the initial settings was to place all sewer feature labels automatically, while satisfying the No Overlap and Preferred Position metrics described earlier. While Maplex placed 91% of the ideal number of labels compared to the Standard Engine's 79.4%, both engines produced relatively high-quality results in both overlap and preferred position metrics.

In the second automated iteration, labeling expressions were added for the sewer main labels for both engines to suppress placement of labels on extremely short (>80 feet) lengths of sewer, observed to be a key source of overlap in the first iteration. Orientation of sewer size and material labels in Maplex were changed to align with the map rather than the direction of the pipe; sewer slope labels remained oriented to pipe flow direction. Also in this iteration, street name labels were placed automatically by both engines rather than as static annotation (first iteration), to provide additional flexibility for label placement. While these changes resulted in

"For each buffer segment, the total number of feature labels available for placement was calculated. This figure represents the Ideal Label Count for each buffered segment and served as the baseline for the quantity metrics."

"The research was conducted in three stages: (1) automated labeling; (2) conversion to annotation; and (3) manual refinement of label placement."

"The objective of the initial settings was to place all sewer feature labels automatically, while satisfying the No Overlap and Preferred Position metrics described earlier. While Maplex placed 91% of the ideal number of labels compared to the Standard Engine's 79.4%, both engines produced relatively high-quality results in both overlap and preferred position metrics."

Label Type	Label Manager Category	Label Placement Factor	Standard Labeling Engine Iteration 1	Iteration 2	Iteration 3
Sewer Manhole	Placement Properties	Point Setting	Offset label horizontally around point		
			Top Right Only (Facility ID)		
			Top Left Only (Invert Elevation)		
			Bottom Right Only (Station Number)		
			Bottom Left Only (Rim Elevation)		
	Conflict Detection	Label weight	High		
		Feature weight	High		
		Buffer	0		
		Place overlapping labels	No		
Sewer Slope	Placement Properties	Orientation	Parallel		
		Orientation system	Page		
		Position	Below		
		Offset (map units)	2		
		Remove duplicate labels	Yes		
		Location along line	At Best		
	Conflict Detection	Label weight	High		
		Feature weight	Low		
		Buffer	0		
		Place overlapping labels	No		
	Label Expression			Do not label if length < 80 feet	Label < 80 feet with leader line
Sewer Size and Material	Placement Properties	Orientation	Parallel		
		Orientation system	Page		
		Position	Above		
		Offset (map units)	2		
		Remove duplicate labels	Yes		
		Location along line	At Best		
	Conflict Detection	Label weight	High		
		Feature weight	Low		
		Buffer	0		
		Place overlapping labels	No		
	Label Expression			Do not label if length < 80 feet	Label < 80 feet with leader line
Sewer Service Number	Placement Properties	Orientation	Curved		
		Orientation system	Line		
		Position	Left or Right		
		Offset (map units)	2		
		Remove duplicate labels	Yes		
		Location along line	N/A		
	Conflict Detection	Label weight	High		
		Feature weight	High		
		Buffer	0		
		Place overlapping labels	No		
Street name	Placement Properties	Orientation	Static annotation	Curved	
		Orientation system		N/A	
		Position		On the line	
		Offset (map units)		N/A	
		Remove duplicate labels		Yes	
		Vary font size		Yes (via expression)	
		Location along line		N/A	
	Conflict Detection	Label weight		High	
		Feature weight		None	
		Buffer		0	
		Place overlapping labels		No	

Table 2. Standard Labeling Engine Settings for Automated Labeling, by Iteration.

fewer labels being placed by both engines, the quality of label placement increased for both.

In the third iteration, leader line labeling was attempted for both engines, with Maplex producing more aesthetically pleasing results for the few such labels that were able to be placed in this fashion. As noted in Freeman (2007) with regard to leadering, the task is more complex than it might appear, as both space for the label and a path for the leader must be found, and the leader must touch the feature without crossing more than one or two other features in the process. While the Standard engine seemed constrained to connect leaders only to the center point of the line feature, Maplex offered more flexibility in both leader and label place-ment.

Label Type	Label Manager Category	Label Placement Factor	Iteration 1	Iteration 2	Iteration 3	Iteration 4
			Maplex			
Overall setting	Maplex Labeling Mode		Fast			Best
Sewer Manhole	Label Position	Position options	North (Facility ID) East (Invert Elevation) South (Station Number) West (Rim Elevation)			
	Label Fitting Strategy	Stack label	No			
		Reduce font size	No			
		Abbreviate label	No			
	Conflict Resolution	Feature weight	100			
		Background label	No			
		Remove duplicate labels	No			
		Never remove label	Yes		No	
		Label buffer (% font height)	15			
Sewer Slope	Label Position	Position options	Street placement - offset straight			
		Label offset (map units)	2			
		Label orientation	Align label to direction of line			
		Repeat label	No			
		Spread characters	No			
		Spread words	No			
	Label Fitting Strategy	Stack label	No			
		Overrun feature	No			
		Reduce font size	No			
		Abbreviate label	No			
		Minimum feature size for labeling (map units)	0			
	Conflict Resolution	Feature weight	100			
		Background label	No			
		Remove duplicate labels	No			
		Never remove label	Yes		No	
		Label buffer (% font height)	15			
	Label Expression			Do not label if length < 80 feet	Label < 80 feet with leader line	Do not label if length < 80 feet
Sewer Size and Material	Label Position	Position options	Street placement - offset straight			
		Label offset (map units)	2			
		Label orientation	Align label to direction of line	None		
		Repeat label	No			
		Spread characters	No			
		Spread words	No			
	Label Fitting Strategy	Stack label	No			
		Overrun feature	No			
		Reduce font size	No			
		Abbreviate label	No			
		Minimum feature size for labeling (map units)	0			
	Conflict Resolution	Feature weight	0			
		Background label	No			
		Remove duplicate labels	No			
		Never remove label	Yes		No	
		Label buffer (% font height)	15			
	Label Expression			Do not label if length < 80 feet	Label < 80 feet with leader line	
Sewer Service Number	Label Position	Position options	Street placement - offset straight			
		Label offset (map units)	2			
		Label orientation	None			
		Repeat label	No			
		Spread characters	No			
		Spread words	No			
	Label Fitting Strategy	Stack label	No			
		Overrun feature	Yes			
		Reduce font size	No			
		Abbreviate label	No			
		Minimum feature size for labeling (map units)	0			

Table 3. Maplex Settings for Automated Labeling, by Iteration.

Feature Class	Category	Setting		Value	Value
	Conflict Resolution	Feature weight		100	
		Background label		No	
		Remove duplicate labels		No	
		Never remove label		No	
		Label buffer (% font height)		15	
Street Name	Label Position	Position options	Static annotation	Street placement - center curved	
		Offset along line		At best position along line	
		Label orientation		None	
		Repeat label		No	
		Spread characters		No	
		Spread words		No	
	Label Fitting Strategy	Stack label		No	
		Overrun feature		No	
		Reduce font size		No	Yes
		Abbreviate label		No	
		Minimum feature size for labeling (map units)		0	
	Conflict Resolution	Feature weight		0	
		Background label		No	
		Remove duplicate labels		No	
		Never remove label		No	
		Label buffer (% font height)		15	
	Labeling Options - Maplex	Enable connection of line segments into continuous features		No	Yes - Unambiguous
All other features (no labels)	Conflict Resolution	Feature weight		0	

Table 3 (continued). Maplex Settings for Automated Labeling, by Iteration.

Standard Engine	**Automated Iteration**					
	1	**2**	**3**			
% of ideal number placed	79.4	78.7	78.7			
% placed without overlap	90.8	92.4	93.7			
% placed in preferred position	98.4	99.5	99.7			

Maplex	**Automated Iteration**				**Conversion to Annotation**	**Manual Refinement**
	1	**2**	**3**	**4**		
% of ideal number placed	91.0	87.8	87.3	85.3	85.7	99.7
% placed without overlap	91.1	93.1	90.9	93.5	93.5	100.0
% placed in preferred position	89.8	97.9	96.8	100.0	100.0	91.4

Table 4. Labeling Metric Results by Iteration.

Also in this iteration, street name font size was varied by street width for both engines in an attempt to reduce overlap with edge-of-pavement lines, and placement of duplicate street names was eliminated. These changes improved quality metrics for the Standard engine, while reducing both quantity placed and quality metrics for Maplex (leader line labeling reduces the preferred position metric).

The dual iteration process was stopped at this point, as no further opportunities for significant improvement in the Standard labeling engine were identified. One further iteration was undertaken with Maplex on parameters with no equivalent settings in the Standard engine: The overall Fast versus Best toggle was changed from Fast to Best, and the setting for Never Remove Labels was set to No for all label classes.

After completion of these iterations, Maplex had placed almost 7 percent more labels overall than the Standard Labeling Engine–85.3 percent

"After completion of these iterations, Maplex had placed almost 7 percent more labels overall than the Standard Labeling Engine–85.3 percent and 78.7 percent of ideal, respectively. For the labels they did place, both products provided equally good quality label placement: About 93 percent of labels were placed with No Overlap, and virtually 100 percent of labels were placed in their Preferred Position."

and 78.7 percent of ideal, respectively (Figure 7). For the labels they did place, both products provided equally good quality label placement: About 93 percent of labels were placed with No Overlap, and virtually 100 percent of labels were placed in their Preferred Position. Both labeling engines placed nearly 100 percent of point feature labels (Figure 8), while Maplex placed about 15 percent more line and area feature labels than the Standard Labeling Engine.

Although publicly available information about these two labeling engines does not reveal the algorithms that drive them, a comparison of Tables 2 and 3 makes it clear that Maplex offers considerably more opportunities for adjusting label placement parameters. The user must decide when the time devoted to (or potentially wasted in) tweaking and fine-tuning the many options has yielded sufficient improvement in the automated labeling stage before moving on to manual placement and editing.

Conversion to Annotation

Because the Maplex engine placed more labels than the Standard Labeling Engine, and with equally high quality, the Maplex labeling was taken forward into the manual refinement stage of the research, through conversion
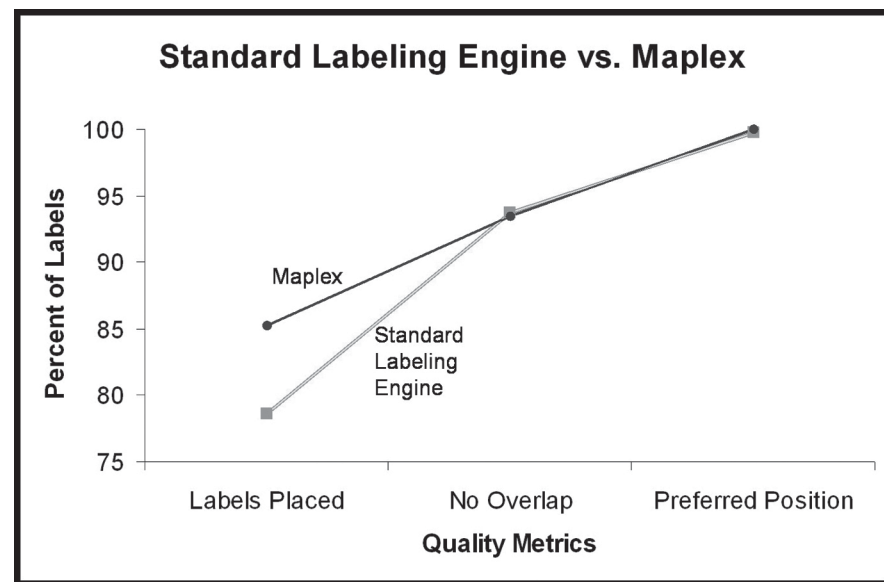

Figure 7. Comparison of quality metrics after automated labeling: ESRI Standard Labeling Engine versus Maplex Labeling Engine.
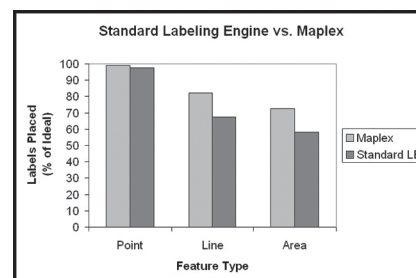

Figure 8. Comparison of percent labels placed by feature type after automated labeling: ESRI Standard Labeling Engine versus Maplex Labeling Engine.

of labeling to annotation. Two aspects of the conversion process proved critical in preserving label position for these dense maps:

- Converting all the labels on each page in one step rather than layer by layer, since each new layer of annotation creates a barrier to the labeling of subsequent layers; and
- Performing the conversion in Layout View (rather than Data View), so that the labels of all features on the page (the current extent) would be converted at once.

In fact, the conversion process automatically placed a few additional labels.

After conversion to annotation, all missing labels were added manually to the three case study maps. Individual label positions were then adjusted in two rounds of refinement, to (1) eliminate overlaps; and (2) maximize the preferred position quality metric.

Manual label position refinement resulted in placement of 99.7 percent of the ideal number of labels, with six redundant street name labels not placed (Figure 9). After 477 manual position adjustments, the No Overlap metric was improved to 100 percent. The tradeoff was a 9 percent decline in the Preferred Position metric, to 91.4 percent, primarily due to use of leader lines for labels unable to be placed next to their feature due to crowding. With regard to feature types, manual refinement resulted in placement of 100 percent of both point and line feature labels (Figure 10). As noted above, a choice was made not to place six repetitive street name labels, but area feature label placement still improved to 90.9 percent.

"Manual label position refinement resulted in placement of 99.7 percent of the ideal number of labels . . ."
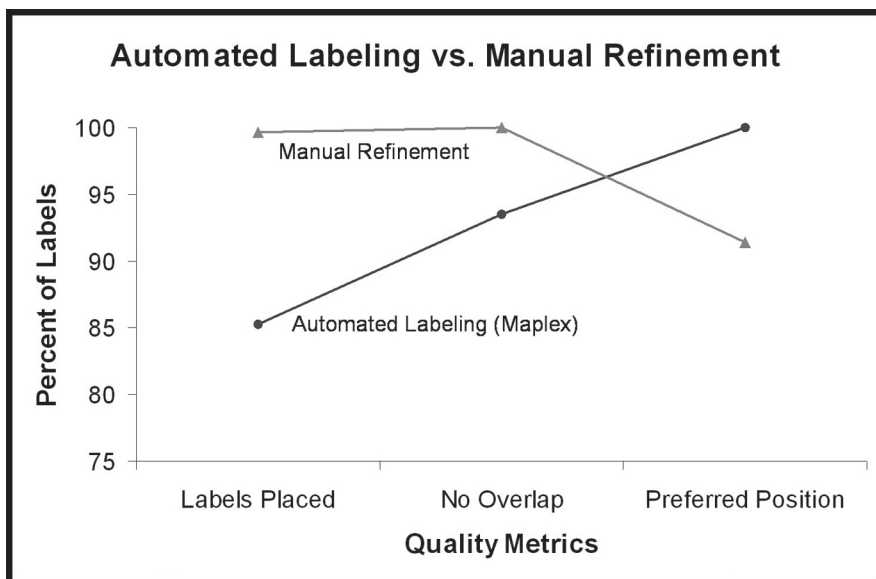

Figure 9. Comparison of quality metrics after manual refinement of label placement: Automated label placement using ESRI Maplex Labeling Engine versus manual refinement.
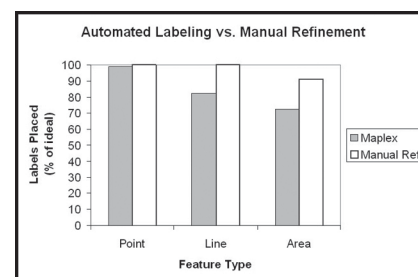

Figure 10. Comparison of percent labels placed by feature type after manual refinement of label placement: Automated label placement using ESRI Maplex Labeling Engine versus manual refinement.

Research Limitations and Potential for Future Research

This research was conducted on labeling engines within ESRI's ArcGIS 9.2 only. Future research could compare different software products (for example, those offered by MapText, Inc., MapInfo Corp., Spatial Projects and Avenza Systems, Inc.) as well as testing follow-on versions of ArcGIS, such as the recently released 9.3. Also, only one type of map was tested, leaving open the question of how well these labeling engines perform across a variety of types of densely labeled maps.

An additional avenue worthy of exploration revolves around the question of time. As noted earlier, Yoeli (1972) estimated that manual labeling of paper maps could consume up to 50 percent or more of total map production time. How might one construct a test of this proportion when constructing and labeling a map via software? The primary time sinks in this research, though not quantified, lay in (a) developing initial label classes and label engine settings; (b) manually eliminating overlaps remaining after automated label placement; and (c) manually placing leader-lined labels. Additional potential time sinks for many labeling projects will be the initial learning curve associated with the many label engine settings available, particularly with Maplex, and the development of a clean database of label names to be applied to map features.

"This research also highlighted possibilities for further development of labeling software capability to further reduce the necessity of manual label placement."

This research also highlighted possibilities for further development of labeling software capability to further reduce the necessity of manual label placement. These include automated leader-line labeling in situations where:

- Feature spacing is too tight to permit placement of a single legibly sized label in the preferred position;
- Feature spacing is too tight to permit placement of a legibly sized label <u>cluster</u> in the preferred position (e.g., around manholes, as in Figure 3);
- Line feature length is too short for legible label placement;
- Area feature width is too narrow for legible label placement; or
- Labels must not obscure "non-feature" elements (e.g., sewer main-service tie junctions).

Such nuances of automated label placement have, for the Town of Concord, become somewhat less critical. In the time that has elapsed since this research commenced, the sewer crew has begun using a rugged laptop in the field to view the sewer map book data. Because ArcMap provides unlimited zoom-in capability, the crew can always zoom in close enough for live automated labeling to appear–a tribute to the development of speedy labeling in dynamic mapping applications, as referenced in the literature review. A few crew members are not yet completely comfortable with the laptop, however, so the paper map book remains a useful tool for the time being.

CONCLUSION

The ultimate goal for label placement within GIS software should be live automated labeling, without going to annotation and manual refinement, that equals both the quantity and quality of manual label placement in significantly less time overall. The software tested here performed quite well on dense utility maps, placing 85.3 percent of labels overall, with high quality: 100 percent in preferred positions; 93.5 percent with no overlap. However, opportunities remain to further perfect the automated labeling process, particularly through automatically leader-lining labels in spots where features are too close together for legibly sized labels to be placed in ideal positions.

ACKNOWLEDGMENTS

REFERENCES

Barrault, M. 2001. A methodology for placement and evaluation of area map labels. *Computers, Environment and Urban Systems* 25: 33-52.

Been, Ken, Daiches, Eli, and Yap, Chee 2006. Dynamic map labeling. *IEEE Transactions on Visualization and Computer Graphics* 12(5): 773-780.

Black, Paul E. (2004). Simulated annealing. National Institute of Standards and Technology. http://www.nist.gov/dads/HTML/simulatedAnnealing.html.

Chirie, F. 2000. Automated name placement with high cartographic quality: City street maps. *Cartography and Geographic Information Science* 27(2): 101-110.

Christensen, Jon, Marks, Joe, and Shieber, Stuart M. 1994. Placing text labels on maps and diagrams. In Heckbert, P. S. (Ed.), *Graphics Gems IV*, 497-504. Cambridge, MA: Academic Press.

———. 1995. An empirical study of algorithms for point-feature label placement. *ACM Transactions on Graphics* 14(3): 203-232.

Cook, A. C. and Jones, C. B. 1990. A Prolog interface to a cartographic database for name placement. *Proceedings of the International Symposium on Spatial Data Handling*: 701-710.

Dent, B. D. 1998. *Cartography: Thematic Map Design*, 5th edition, Boston, MA: William C. Brown.

Doerschler, Jeffrey S., and Freeman, Herbert. 1992. A rule-based system for dense-map name placement. *Communications of the ACM*, 35(1): 68-79.

Ebinger, Lee R., and Goulette, Ann M. 1990. Non-interactive names placement for the 1990 decennial census. *Cartography and Geographic Information Science* 17(1): 69-78.

Ebner, Dietmar, Klau, Gunnar W., and Weiskirscher, Rene. 2003. *Force-based label number maximization.* Vienna, Austria: Vienna University of Technology.

Edmondson, Shawn, Christensen, Jon, Marks, Joe, and Shieber, Stuart M. 1996. A general cartographic labeling algorithm. *Cartographica* 33(4): 13-23.

ESRI, *Maplex: Automated cartographic name placement software,* August 1998, <http://www.esri.com/library/whitepapers/pdfs/maplexwp.pdf>.

Freeman, Herbert, and Ahn, John. 1984. AUTONAP – an expert system for automatic name placement. *Proceedings of the International Symposium on Spatial Data Handling*: 544-569.

———. 1987. On the problem of placing names in a geographic map. *International Journal of Pattern Recognition and Artificial Intelligence*, 1(1): 121-140.

Freeman, Herbert. 1995. On the automated labeling of maps. In Dori, D. and Bruckstein, A., (Eds.) *Shape, Structure and Pattern Recognition*, 432-444. Singapore: World Scientific.

———. 2005. Automated cartographic text placement. *Pattern Recognition Letters*, 26(3): 287-297.

———. 2007. On the problem of placing names on a map. *Journal of Visual Languages and Computing*, 18(5): 458-469.

Hirsch, Stephen A. 1982. An algorithm for automatic name placement around point data. *The American Cartographer*, 9(1): 5-17.

Howe, Denis, Expert system, May 29, 1996, http://dictionary.reference. com/browse/expert%20system.

Huffman, F. T., and Cromley, R. C. 2002. An automated multi-criteria cartographic aid for point annotation. *Cartographic Journal*, 39:51-64.

Imhof, Eduard. 1962. Die anordnung der namen in der karte. *Internationales Jahrbuch fur Kartographie,* 2:93-129.

Imhof, Eduard. 1975. Positioning names on maps**.** *The American Cartographer*, 2(2): 128-144.

Jones, Christopher B. 1989. Cartographic name placement with Prolog. *IEEE Computer Graphics and Applications*, 9(5): 36-47.

Kakoulis, Konstantinos G., and Tollis, Ioannis G. 2006. Algorithms for the multiple label placement problem. *Computational Geometry-Theory and Applications*, 35(3): 143-61.

Kakoulis Konstantinos G., and Tollis, Ioannis G. 2003. A unified approach to automatic label placement. *International Journal of Computational Geometry and Applications*, 13(1): 23-59.

Kameda, Ttakayuki, and Imai, Keiko. 2003. Map label placement for points and curves. *IEICE Transaction Fundamentals* E86-A(4): 835-840.

Mote, Kevin. 2007. Fast point-feature label placement for dynamic visualizations. *Information Visualization*, 6(4): 249-260.

Mower, James E. 1993. Automated feature and name placement on parallel computers. *Cartography and Geographic Information Science*, 20.2: 69-82.

Pinto, I. and Freeman, H. 1996. The feedback approach to cartographic areal text placement. In Perner, P., Wang, P. and Rosenfeld, A. (Eds.) *Advances in Structural and Syntactical Pattern Recognition*, 341-50. Berlin: Springer-Verlag.

Ribeiro, Glaydston Mattos, and Loreno, Luiz Antonio Nogueira. 2006. Heuristics for cartographic label placement problems. *Computers and Geosciences*, 32(6): 739-748.

Robinson, Arthur Howard, Morrison, Joel L., Muehrcke, Phillip C., Kimerling, A. Jon, and Guptill, Stephen C. 1995. *Elements of Cartography*, 6th edition, New York: John Wiley & Sons.

Sadahiro, Yukio. 1995. Size of map labels used in GIS and loss of literal information. *Cartographica*, 32(4): 29-41.

Shortridge, Barbara G. 1979. Map reader discrimination of lettering size. *The American Cartographer*, 6(1): 13-20.

Slocum, Terry A., McMaster, Robert B., Kessler, Fritz C., and Howard, Hugh H. 2005. *Thematic Cartography and Geographic Visualization*, 2nd edition, Upper Saddle River: Pearson Prentice Hall.

Stadler, Georg, Steiner, Tibor, and Beiglbock, Jurgen. 2006. A practical map labeling algorithm utilizing morphological image processing and force-directed methods. *Cartography and Geographic Information Science* 33(3): 207-215.

Van Dijk, Steven, Marc Van Krefeld, Tycho Strijk, and Alexander Wolff. 1999. Towards an evaluation of quality for label placement methods. In Keller, C. P. (Ed.) *Proceedings of the 19th International Cartographic Conference and 11th General Assembly*: 57-64.

Van Dijk, Steven, Thierens, Dirk, and de Berg, Mark. 2004. On the design and analysis of competent selecto-recombinative GAs. *Evolutionary Computation*, 12(2): 243-267.

Van Kreveld, Marc, Tycho Trijk, and Alexander Wolff, A. 1999. Point labeling with sliding labels. *Computational Geometry* 13:21-47.

Verner, Oleg V., Roger L. Wainwright, and Dale Schoenefeld. 1997. Placing text labels on maps and diagrams using genetic algorithms with masking. *INFORMS Journal on Computing*, 9(7): 266-275.

Wagner, Frank, and Alexander Wolff. 1997. A practical map labeling algorithm. *Computational Geometry* 7: 387-404.

Wagner, Frank, Alexander Wolff, Vikas Kapoor, and Tycho Strijk. 2001. Three rules suffice for good label placement. *Algorithmica* 30(2): 334-349.

Wood, Clifford H. 2000. Descriptive and illustrated guide for type placement in small scale maps. *The Cartographic Journal*, 37(1): 5-18.

Wu, Chyan Victor, and Barbara P. Buttenfield. 1991. Reconsidering rules for point-feature name placement. *Cartographica*, 28(1): 10-27.

Yamamoto, Missae, and Antonio Nogueira Lorena. 2005. A constructive genetic approach to point-feature cartographic label placement. *Meta-heuristics: Progress as Real Problem Solvers*, New York, NY: Springer.

Yoeli, Pinhas. 1972. The logic of automated map lettering. *The Cartographic Journal*, 9(2): 99-108.

Zoraster, Steven. 1986. Integer programming applied to the map label placement problem. *Cartographica*, 23(3): 16-27.

———. 1991. Expert systems and the map label placement problem. *Cartographica*, 28(1): 1-9.

———. 1997. Practical results using simulated annealing for point feature label placement. *Cartography and Geographic Information Science*, 24(4): 228-238.