

Heuristic Search and Information Visualization Methods for School Redistricting

Marie desJardins, Blazej Bulka, Ryan Carr, Eric Jordan, and Penny Rheingans

Department of Computer Science and Electrical Engineering

University of Maryland Baltimore County

1000 Hilltop Circle, Baltimore MD 21250

{mariedj,bulka1,carr2,ejordan1,rheingan}@umbc.edu

Abstract

We describe an application of AI search and information visualization techniques to the problem of school redistricting, in which students are assigned to home schools within a county or school district. This is a multicriteria optimization problem in which competing objectives must be considered, such as school capacity, busing costs, and socioeconomic distribution. Because of the complexity of the decision-making problem, tools are needed to help end users generate, evaluate, and compare alternative school assignment plans. A key goal of our research is to aid users in finding multiple qualitatively different redistricting plans that represent different tradeoffs in the decision space.

We present heuristic search methods that can be used to find a set of qualitatively different plans, and give empirical results of these search methods on population data from the school district of Howard County, Maryland. We show the resulting plans using novel visualization methods that we have developed for summarizing and comparing alternative plans.

Motivation and Overview

This research focuses on developing decision support tools for the problem of school redistricting. In this domain, the goal is to assign the students from each geographic region (neighborhood or *planning polygon*) in a county or school district to a home school at each level (elementary, middle, and high school). We are working with the Howard County, Maryland, school system to develop tools that will aid in generating, evaluating, and comparing alternative school assignment plans. Related applications include emergency response planning, urban planning and zoning, robot exploration planning, and political redistricting.

The school assignment plan should ideally satisfy a number of different goals, such as meeting school capacities, balancing socioeconomic and test score distributions at the schools, minimizing busing costs, and allowing students in the “walk area” of a school to attend that home school. Since these objectives are often at odds with each other, finding the best plan is a complex multicriteria optimization problem. It is also often desirable to create several alternative plans for consideration; these plans should be *qualitatively different*—that is, they should represent different tradeoffs among the

evaluation criteria. Finally, because of the complexity of the problem, it is difficult for users to fully understand these tradeoffs. Therefore, developing effective visualizations is an important challenge.

The contributions of our work are (1) a computational formulation of the school redistricting problem as a multicriteria optimization problem; (2) novel heuristic local search techniques for generating high-quality, diverse (i.e., qualitatively different) plans; (3) visualization methods¹ for comparing alternative plans; and (4) empirical results demonstrating the effectiveness of our search methods on actual Howard County school data.

The rest of this paper is organized as follows. We first describe the current redistricting process in Howard County, and present some example plan visualizations that we have developed. Next, we describe the search methods and present empirical results comparing manually and automatically generated plans in terms of plan quality and diversity. Finally, we summarize related work, then present our future work and conclusions.

Redistricting Process

The Howard County Public School System (HCPSS) serves a rapidly growing county in suburban Maryland. The pace of development and population growth has necessitated the opening of 25 new schools in the last 14 years, turning the adjustment of school attendance areas into an almost annual event. Under the current process, candidate plans and feasibility studies are generated manually² by school system staff. These plans are evaluated and refined by a committee of citizens, then presented at regional meetings for public comment. A small set of candidate plans is forwarded to the Superintendent, who presents two or three recommended alternatives to the Board of Education. The Board has final decision-making authority, and will typically select one of the recommended plans, sometimes making minor modifications in response to concerns raised by parent groups or

¹These visualization methods are summarized here; they are described in detail in an earlier publication (Shanbhag, Rheingans, & desJardins 2005).

²Map-based tools are used to show the proposed school districts, and a set of spreadsheets is used to generate evaluation data. No other decision support tools are used in the current process.

staff. Note that this process is specific to Howard County; other school districts may have different processes and models.

Candidate plans are evaluated according to eleven measured criteria: (1) the educational benefits for students, (2) the frequency with which students are redistricted, (3) the number and distance of students bused, (4) the total busing cost, (5) the demographic makeup and academic performance of schools, (6) the number of students redistricted, (7) the maintenance of feeder patterns (i.e., the flow of students from elementary to middle to high school), (8) changes in school capacity, (9) the impact on specialized programs, (10) the functional and operational capacity of school infrastructure, and (11) building utilization. Some of these criteria can be clearly quantified (e.g., building utilization and busing costs), while others are harder to quantify (e.g., educational benefits and impact on specialized programs).

In practice, the process is primarily driven by building utilization, but serious consideration is given to feeder patterns, the number of students redistricted, demographic makeup, busing costs, and the frequency with which students are redistricted. Ideally, building utilization should be between 90% and 110% of program capacity and should stay in that range as projected population and capacity changes occur. Desired feeder patterns ensure that there is a critical mass of students who move together from one school level (elementary, middle, and high school) to the next. For instance, the students from a particular middle school should constitute at least 15% of the population of any high school that they feed into. Consideration of the demographic makeup of schools helps to ensure that economically and academically disadvantaged children are not unnecessarily segregated into a few schools.

The specific redistricting problem that we focus on in this paper is one that the county faced during the 2004-2005 school year, that of developing a school assignment plan for a twelfth high school (Marriotts Ridge) that opened in Fall 2005. Figure 1 shows the partitioning of planning regions into school attendance areas before the new school was opened. A colored circle shows the location of each school, and each planning polygon is colored according to the high school attended. For instance, all students in the northwest region of the county are assigned to Glenelg High School (labeled G; bright green region) in the original plan. The unused capacity at a school is shown by a black area (wedge) inside the school circle. Schools that are over capacity are outlined with a ring. Glenelg is slightly under capacity; Marriotts Ridge (labeled MR) has zero utilization (i.e., is completely black), since no students have yet been assigned to this school in this plan; and Mt. Hebron (labeled MH; bright pink region) is significantly over capacity.

Figure 2 shows a comparison picture of two alternative plans that include Marriotts Ridge. The school assignments for the closest-school plan—generated by assigning every polygon to the school closest to its geographic center—are indicated by the outer colored rings in each planning polygon. This plan provides a useful baseline, because it optimizes both walk usage and busing costs, but may be undesirable in terms of capacity and demographics. The in-

ner ring shows the school assignments for the recommended (“green”) plan that was recommended by the superintendent’s office. The “tree ring” effect allows the user to easily see the planning polygons where the two plans make different recommendations. Polygons with a single color are assigned to the same school in both plans.

For example, in the center of the county, several polygons are assigned to Marriotts Ridge by the recommended plan, but to the nearby Mount Hebron (labeled MH in Figure 1), River Hill (RH), and Centennial (C) High Schools by the closest-school plan. The inner ring (center) of these polygons is displayed as sky blue (Marriotts Ridge); the outer ring is either bright pink (Mount Hebron), dark blue (River Hill), or orange (Centennial). Along the southeast border of the county, the closest-school plan assigns a number of polygons to Hammond High School (H: medium blue outer ring) that are assigned to Reservoir (R: green inner ring) by the recommended plan. This difference occurs because assigning them to Hammond would cause that school to be over capacity; also, in this case, those polygons help to balance the socioeconomic distribution at Reservoir.

Search Methods

The search space for the redistricting problem is very large. For p polygons and s schools, there are

$$s^{(p-s)}$$

possible assignments of schools to polygons (since polygons containing a school are constrained to be assigned to that school). Requiring that school attendance areas be geographically contiguous reduces the number of possible plans, but the number of plans still grows exponentially with the number of schools and polygons. Because of this complexity, we have chosen to use heuristic local search methods, which do not guarantee optimality, but which can be used to find good solutions reasonably quickly.

Our basic approach is a two-stage process: first, we generate an initial “seed” plan using one of several methods described below; second, we use local search to “hill-climb” to a local optimum. Because of the multicriteria nature of the redistricting problem, we have designed several different variations of hill-climbing search that can be used to find qualitatively different alternative plans in the solution space, as discussed later.

Before describing the methods for finding seed plans and for performing local search, we first introduce the evaluation criteria that we use to measure the quality of a school assignment plan.

Evaluation Criteria

A school plan assignment can be evaluated along multiple dimensions. The measured criteria used by the Howard County schools were summarized earlier. We have defined five quantitative criteria, f_1, \dots, f_5 , based on these measured criteria. Each of these criteria is scaled and normalized so that the value for a given plan will always fall in the range $[0, 1]$, with a lower value being preferred.

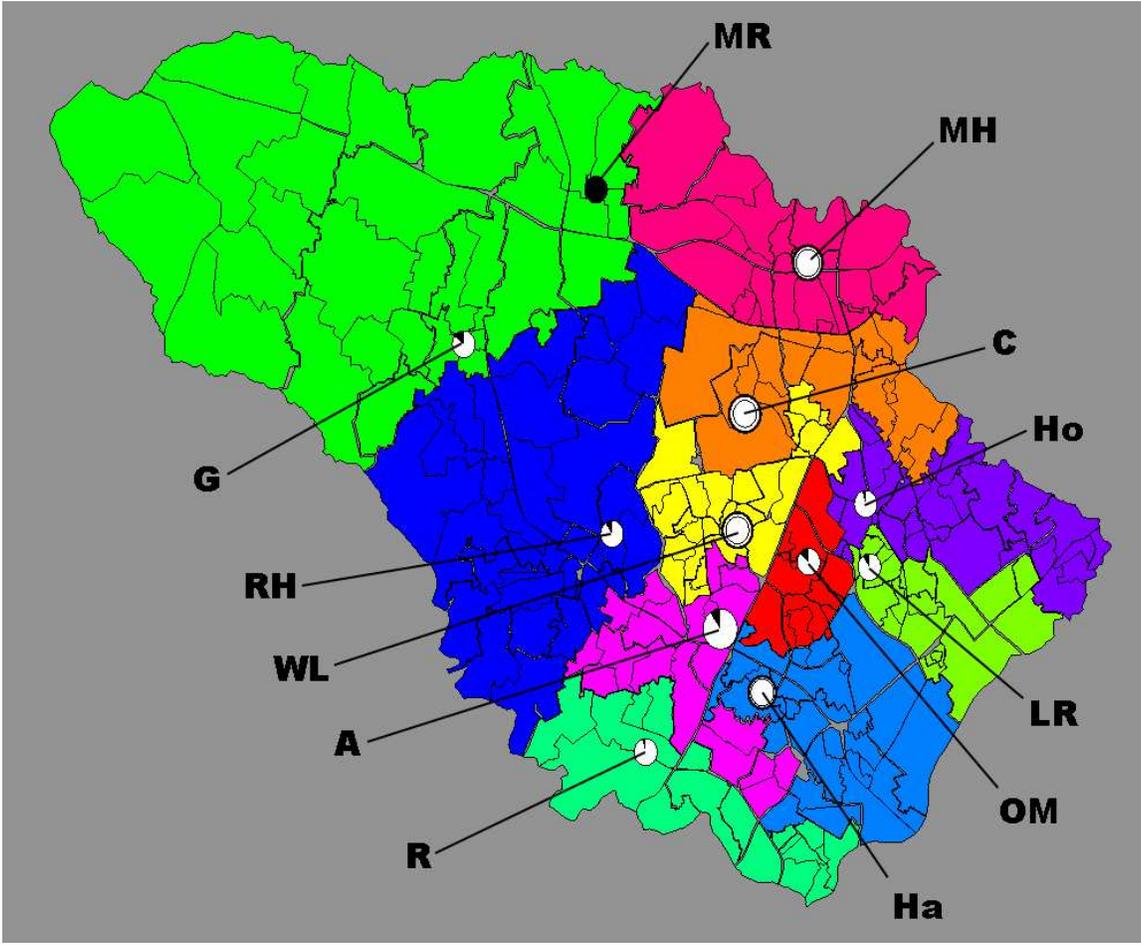


Figure 1: Original 2004-2005 school assignment plan. Each planning polygon is colored according to the school to which it is assigned. School glyphs show underutilization as a “pie wedge” (percentage of black fill). Overutilization is shown by a ring outlining the glyph; the diameter of the ring is proportional to the degree of overutilization.

Using concepts from the multiattribute decision theory literature (Keeney & Raiffa 1993), a plan P_1 is said to *dominate* a second plan P_2 if P_1 is better along some dimension, and no worse along any dimension, than P_2 :

$$Dom(P_1, P_2) \Leftrightarrow \exists i f_i(P_1) < f_i(P_2) \wedge \forall i f_i(P_1) \leq f_i(P_2).$$

Two plans are *incomparable* if neither plan dominates the other.

It is often desirable to define a combined score that incorporates all of the evaluation criteria. For this purpose, we use a simple linear combination, F :

$$F(P) = \sum_i w_i f_i(P),$$

where the weights $w_i \in [0, 1]$ represent the relative importance of each of the criteria. Since each f_i ranges from 0 to 1, F will range from 0 to 5.

We introduce the following notation:

- $Pop(p)$ is the number of students in planning polygon p .
- $Pop(s, P)$ is the number of students assigned to school s by plan P .

- $Cap(s)$ is the capacity of school s .
- $N_{schools}$ is the number of schools in the county (12 in this particular planning problem).
- $N_{students} = \sum_p Pop(p)$ is the number of students in the county.
- $s(p, P)$ is the school to which polygon p is assigned by plan P .
- $p \in s$ refers to the set of polygons assigned to school s by a given plan. This notation is shorthand for the more space-consuming set notation, $\{p : s(p, P) = s\}$.

The five evaluation criteria, which are defined in the following paragraphs, are: school capacity (f_1), socioeconomic distribution (f_2), test score distribution (f_3), busing costs (f_4), and walk area usage (f_5).

School Capacity (f_1). A plan that utilizes any school at less than 90% or greater than 110% of its capacity is considered to be highly undesirable. Therefore, we compute the

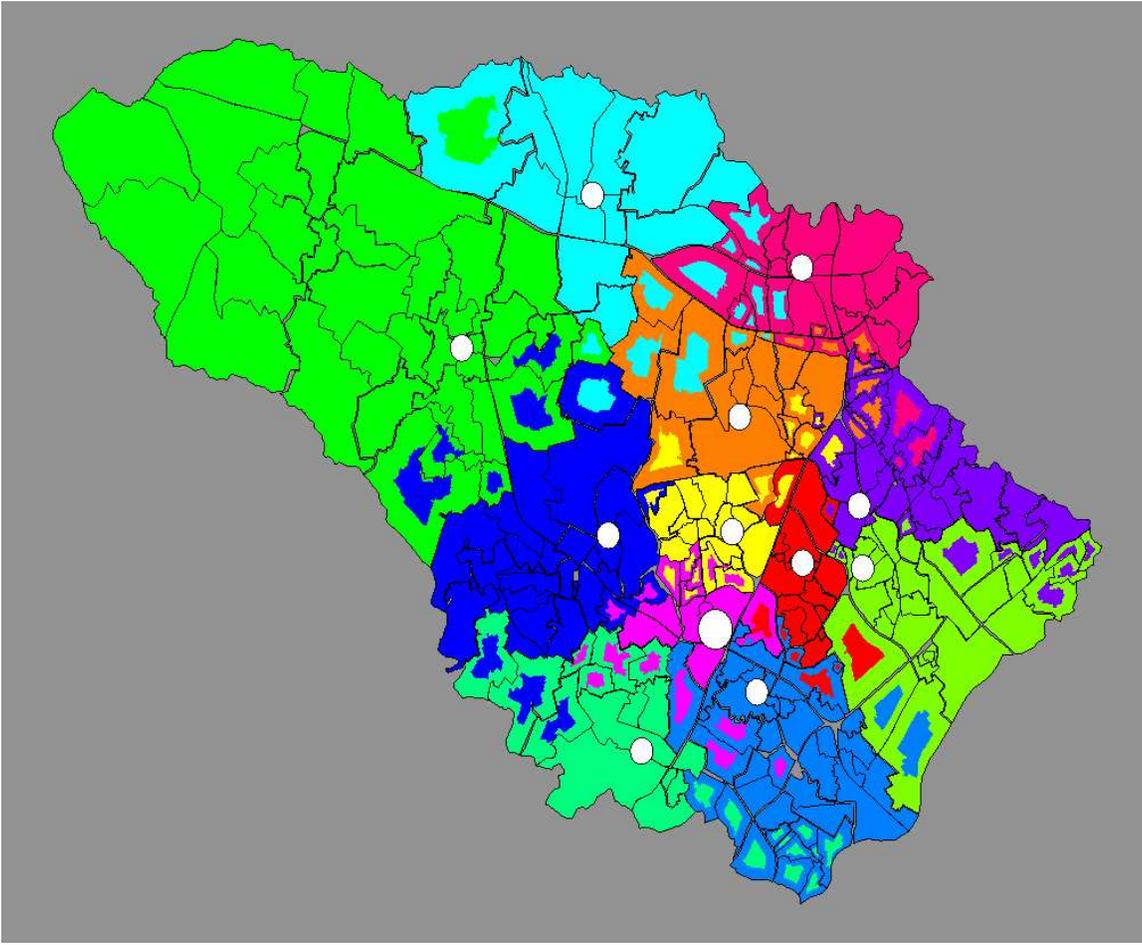


Figure 2: A comparison of the recommended (“green”) plan to the closest-school plan. The color of the inner and outer rings in each planning polygon indicate the school assignments for the recommended and closest-school plan, respectively.

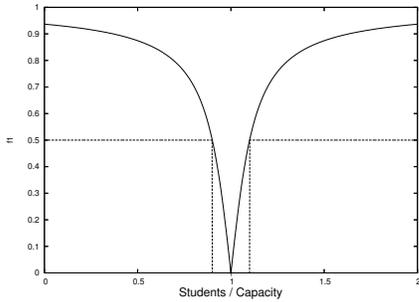


Figure 3: Penalty function for school utilization.

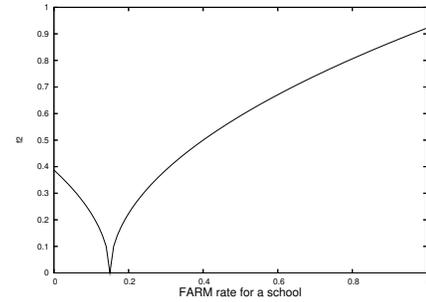


Figure 4: Penalty function for socioeconomic distribution, for a county FARM rate of 0.15.

school utilization (i.e., the ratio of proposed school enrollment to school capacity), and map it to a penalty function with its minimum at 100%. For this purpose, we used a scaled arctan function (Figure 3):

$$f_1(P) = \frac{\sum_s \left| \arctan \left(\sigma \left(1 - \frac{Pop(s,P)}{Cap(s)} \right) \right) \right|}{\frac{\pi}{2} N_{schools}}$$

where σ is a scaling factor that causes the scaled arctan to be equal to 0.5 when the school utilization is either 0.9 or 1.1 ($\sigma = 10.0$). As seen in Figure 3, the penalty function increases rapidly away from the ideal capacity of 100%, and assigns high values for values that are significantly outside the target range, $[0.9, 1.1]$.

Socioeconomic Distribution (f_2). The school system uses the percentage of students that qualify for free and reduced meals (FARM) as a measure of socioeconomic distribution. The goal in creating a school assignment plan is to equalize this distribution across the county: ideally, each school will have the same FARM rate as the county as a whole. FARM rates are given in the data on a per-polygon basis, denoted as $FARM(p)$. We compute $FARM_C$, the average FARM rate for the county as a whole:

$$FARM_C = \frac{\sum_p Pop(p) FARM(p)}{N_{students}} \quad (1)$$

and the FARM rate for each school:

$$FARM(s, P) = \frac{\sum_{p \in s} Pop(p) FARM(p)}{Pop(s, P)}. \quad (2)$$

In order to penalize greater deviations from the average FARM rate more heavily, we take the square root of the difference between the school and county FARM rates (which both range from 0 to 1), and then average this over the schools in the county to compute the overall socioeconomic distribution criterion:

$$f_2(P) = \frac{\sum_s \sqrt{|FARM(s, P) - FARM_C|}}{N_{schools}}$$

Figure 4 shows what one school’s contribution to the f_2 penalty function would look like for a county rate of $FARM_C = 0.15$, as a function of the school’s FARM rate, $FARM(s, P)$.

Test Score Distribution (f_3). The test score rate is measured by the percentage of students in a given polygon or school who achieve a score at the Proficient or Advanced level on the Maryland State Assessment (MSA) standardized test. As with the FARM data, these rates are given on a per-polygon basis in the input data. This criterion is defined analogously to f_2 :

$$f_3(P) = \frac{\sum_s \sqrt{|MSA(s, P) - MSA_C|}}{N_{schools}},$$

where MSA_C and $MSA(s, P)$ are computed as in Equations 1 and 2, respectively.

Busing Costs (f_4). Busing costs would be minimized by sending all students to the geographically closest school. Our evaluation criterion for busing costs is based on the “average excess busing distance”—that is, the average distance traveled beyond the minimum required. To simplify the computation, we treat each polygon as a group of students all traveling from the geometric centroid of the polygon. The distance from this point to the assigned school is calculated; we then subtract the distance to the closest school (denoted by $CS(p)$) to determine the excess busing distance. This distance is normalized by the excess busing distance that would be needed to bus those students to the

fourth closest school (denoted by $4CS(p)$).³ These computations are weighted by the population of each polygon:

$$f_4(P) = \frac{\sum_s \sum_{p \in s} Pop(p) \frac{Dist(p, s) - Dist(p, CS(p))}{Dist(p, 4CS(p)) - Dist(p, CS(p))}}{N_{students}}$$

Walk Area Usage (f_5). It is preferable to send children who are within walking distance of some school to their neighborhood (“walk”) school. This criterion is based on the percentage of such students who are in fact assigned to a school within walking distance.⁴

$$f_5(P) = 1 - \frac{w'(P)}{w}$$

where w is the number of students who are within walking distance of the closest school:

$$w = \sum_{p: Dist(p, CS(p)) \leq k_w} Pop(p)$$

and w' is the number of students who are within walking distance of their assigned school:

$$w'(P) = \sum_{p: Dist(p, s(p, P)) \leq k_w} Pop(p).$$

Generating Seed Plans

In this paper, we use two seed plans: the closest-school plan and the current (“original”) redistricting plan.

The closest-school plan simply assigns each polygon to the school that is geographically closest to the centroid of the polygon. This plan minimizes busing costs and allows all students in the walk area of a school to attend that school. However, as seen in Table 1, it results in poor school utilization, since some schools are overcrowded and others are underutilized. The closest-school plan also performs somewhat poorly on test-score and socioeconomic distribution, since it reinforces geographic “clustering” of income levels.

The original plan, for this data set, is the high school assignment plan before the new Marriotts Ridge High School was built. It assigns no polygons to the new high school, so its utilization is not particularly good (Table 1). However, in terms of the other criteria, it is approximately as good as the final accepted plans.

We have also experimented with several other seed methods, including random seeds, “breadth-first” assignment of polygons to schools, and a minimum-spanning-tree assignment based on distance. The latter two are similar to the

³The reasoning behind this normalization is that it there is very rarely a need to send students further than the third closest school. As a result, normalizing by the most distant school in the county would result in extremely small values for any “reasonable” school assignment, making it difficult to differentiate among alternatives.

⁴The walking distance, k_w , depends on the age of the children; for high schools, we use 1.5 miles. Note that this computation is an approximation, since the actual assignment of walk areas is more complicated, using actual distance traveled, and taking into account “walkability” (e.g., sidewalks are required, and busy streets must be avoided).

closest-school assignment, but make less sense to an end user. Random seeds are useful for a baseline, but are actually nonsensical from an application perspective, since there is no geographic contiguity at all. In our preliminary experiments, we also found that these seeds do not yield good search performance; in particular, the random seeds cause the local search to take a very long time to converge (since they are so far from a reasonable plan), and the resulting plans are not particularly good. Therefore, we omit these alternatives from the results that we present here.

Local Search Methods

We have developed and applied three basic search methods: basic hill climbing, biased hill climbing with blind bias, and biased hill climbing with diversity bias.

Basic hill climbing. This method simply performs a variation of hill climbing on the combined score, $F(P) = \sum_i w_i f_i(P)$. At each step, the search algorithm considers moving a single polygon to a different school. The branching factor for this search is high (there are $(s-1)^{(p-s)}$ possible actions), so rather than evaluating every possible move, we consider each of these moves in a random order, and take the first one (if any) that improves the combined score. In other words, at each step, a random polygon is selected and assigned to a randomly selected neighboring school. If this improves (reduces) the combined score, then the change is made to the plan. This process continues until a local minimum is reached; that is, until there is no individual polygon that can be moved in order to improve the score.

Biased hill climbing with blind bias. Biased hill climbing is a novel technique that we introduce here. It uses the notions of dominated and incomparable solutions in the multicriteria optimization space to find multiple alternative plans. As with basic hill climbing, this search method tries to move a random polygon to a randomly selected neighboring school. The difference is that this move will only be accepted if it results in a *dominating* plan – that is, if some f_i is improved by the change, and no f_i is made worse. If the change is strictly worse, it is ignored. However, if the change results in an incomparable solution (that is, one that is better with respect to some f_i , but worse with respect to some other f_j), then the resulting plan is placed on a list of “incomparable” solutions, which we refer to as \mathcal{I} .

When a local optimum is reached in “dominated plan space,” under the *blind bias* option, this plan is added to the solution list, \mathcal{S} , then a plan is selected randomly from the incomparable solution list \mathcal{I} and used as the seed for a new search. This process is repeated until a prespecified number, k , of alternative plans is found. The set of k plans, \mathcal{S} , is returned.

Biased hill climbing with diversity bias. This method uses the same basic “dominated hill climbing” approach as the previous method to find the initial solution. However, to find subsequent solutions, a *diversity bias* is applied. Specifically, after a local optimum is found, the new seed will be

Plan	f_1 Cap.	f_2 FARM	f_3 MSA	f_4 Bus	f_5 Walk	F $\sum f_i$
Closest	0.64	0.26	0.24	0.0	0.0	1.15
Original	0.47	0.25	0.22	0.13	0.17	1.24
Green	0.40	0.24	0.21	0.16	0.16	1.17
Red	0.40	0.24	0.21	0.15	0.17	1.19
Closest/Basic	0.44	0.26	0.24	0.04	0.04	1.02
Closest/Blind	0.64	0.26	0.24	0.00	0.00	1.15
Closest/Div	0.64	0.26	0.24	0.00	0.00	1.15
Original/Basic	0.26	0.24	0.21	0.14	0.17	1.02
Original/Blind	0.37	0.22	0.20	0.13	0.17	1.09
Original/Div	0.44	0.23	0.20	0.11	0.12	1.10

Table 1: Average evaluation criteria values (f_i and combined score F (with all weights w_i set to 1) for seed plans and search results. Search results represent the average of 10 runs.

the plan from the incomparable list \mathcal{I} whose average Euclidean distance (in the evaluation space) from \mathcal{S} , the local optima found so far, is greatest. That is,

$$\text{seed}_{i+1} = \operatorname{argmax}_{P' \in \mathcal{I}} \frac{\sum_{P \in \mathcal{S}} \text{EvalDist}(P', P)}{|\mathcal{S}|},$$

where

$$\text{EvalDist}(P', P) = \sqrt{\sum_i (f_i(P') - f_i(P))^2}. \quad (3)$$

Since each of the f_i ranges from 0 to 1, the maximum possible pairwise distance between any two solutions for the five-dimensional evaluation space is $\sqrt{5} = 2.24$. Of course, it is unlikely that we would find locally optimal plans with such extreme values, so typically the pairwise distance will be much smaller.

Empirical Results

In this section, we present results using the Howard County school data for the 2004–2005 redistricting process. The first experiment was designed to compare the quality of the plans that are produced by different heuristic search methods, using different seed plans, to those produced by the redistricting committee. The second experiment was designed to assess the diversity (with respect to the evaluation criteria) of the sets of plans produced by different search methods.

Plan Quality. In Table 1, we compare the average f_i and combined F values for the original, closest-school, recommended (green), and alternate (red) plans to the plan generated by each search method. (The alternate plan was proposed by the Superintendent’s office as an alternative to the green plan.) To generate this data, we ran each search method ten times, since some of the search steps are stochastic. In general, the search process is heavily influenced by the choice of seed. The results of basic hill climbing search when starting from the closest-school plan does better than the closest-school plan with respect to capacity (f_1), but still worse on this measure than the other manually generated plans (original, recommended, and alternate). Similarly, the search results for basic hill climbing starting with

Method	Diversity	f_1	f_2	f_3	f_4	f_5	F
Basic	0.044	0.22	0.24	0.21	0.14	0.15	0.96
Weighted	0.048	0.23	0.24	0.21	0.14	0.14	0.96
Blind	0.032	0.44	0.20	0.18	0.11	0.11	1.04
Diversity	0.389	0.61	0.19	0.15	0.18	0.23	1.36

Table 2: Average diversity (average pairwise distance in evaluation space) and evaluation measures for different methods.

the closest-school plan perform much better with respect to busing (f_4) and walk usage (f_5) than the other plans.

The biased hill climbing methods show a pathological behavior when starting from the closest-school seed. Because the closest-school plan is already optimal with respect to busing costs and walk area usage, the local area contains many incomparable plans, so the search is unable to make any progress.

The overall combined plan quality (F) is better for all of the search methods than for any of the manually constructed plans. This is a good sign, since it means that we are able to find high-quality plans using our search methods. However, we have not yet performed a user study to determine whether plans that appear better with respect to these criteria are, in fact, seen to be better by end users. The result is certainly promising, though, since the general framework can easily be used with different evaluation criteria that are “tuned” to the end users’ actual preferences.

Plan Diversity. In Table 2, we compare sets of three plans generated by each of our search methods to a group of hand-generated plans (recommended, alternate, and closest-school). We give the diversity (average pairwise distance in evaluation space) and the average f_i and F scores for each of the sets. All numbers are the average of 10 runs of the specified search method, using the original plan as the seed.

Examining three manually generated plans (closest, green (recommended), and red (alternate)) gives us a baseline for the diversity measure. The pairwise evaluation distances of the manually generated plans are 0.337 (closest vs. green), 0.333 (closest vs. red), and 0.019 (red vs. green), for an average pairwise distance of 0.223.

Not surprisingly, basic hill climbing with all weights set to 1.0 produces sets of very similar plans, with almost no diversity (0.044 on average—“Basic” in Table 2). Therefore, we experimented with performing hill climbing three times, each time with different weights. (The three runs assign weight 1.0 to f_1 (capacity), f_2 (socioeconomic distribution), and f_4 (busing costs), respectively, and weight 0.5 to all of the other evaluation criteria.) This process (referred to as “weighted” in Table 2) yields only slightly higher diversity (0.048 on average, still below the baseline).

The plans using both hill climbing methods have very low combined F values. A typical plan (with $w_1 = 1.0$ (capacity) and the other $w_i = 0.5$) is compared to the recommended plan in Figure 5. The outer color in each planning polygon shows the school assigned by the recommended plan; the inner color shows the school assigned by

the weighted hill climbing plan. The overall F measure for this plan is 0.91, lower than any of the manually generated plans or any of the average search results in Table 1. Not surprisingly, this plan performs extremely well on the capacity measure ($f_1 = 0.20$), while maintaining fairly good performance along the other dimensions ($f_2 = 0.21$, $f_3 = 0.20$, $f_4 = 0.16$, $f_5 = 0.15$). Comparing these measures to those given in Table 1, it is clear that this plan is better than the manually generated plans with respect to *all* of the evaluation criteria. The figure shows that the weighted hill climbing plan assigns a number of polygons to a more distant (but still nearby) school than the recommended plan. However, although it is not stated explicitly in the evaluation criteria, the geographic “pocketing” that this plan displays is likely to be undesirable.

Figure 6 compares three plans produced by a single representative run of weighted hill climbing. The variation in the plans produced by the weighted hill climbing method is primarily in the capacity, busing, and walk usage measures. Although one of the weight assignments emphasizes socioeconomic distribution, there is not much difference in the f_2 values for these plans. This may be because significantly decreasing socioeconomic distribution would require busing students a very long distance, entailing a severe penalty in busing costs and walk usage. Interestingly, two of the plans are quite similar, so only one planning polygon shows three different school assignments for the three plans. (This polygon is in the southern part of the county, and is assigned to Reservoir (green) by the first plan (outer ring), River Hill (dark blue) by the second, and Atholton (pink) by the third.)

Biased hill climbing with a blind (random) bias also gives very little diversity (0.032 on average). However, biased hill climbing with diversity bias gives the highest diversity of any method, and much higher diversity than the baseline set of plans (0.389 on average). Note that this is accomplished at some loss of quality: the average combined measure (F) for diversity-biased search is 1.36, compared to averages ranging from 0.96–1.04 for the other search methods. Three plans produced by a typical run of diversity-biased hill climbing are shown in Figure 7. Although the average diversity is high, the three plans are not equidistant in evaluation space. Rather, one of the plans (plan 1) falls “between” the other two plans, with pairwise distances of 0.36, 0.37, and 0.67.

Note that pairwise distance is a somewhat naive notion of diversity. Recently (desJardins & Wagstaff 2005), we have studied measures of set diversity in the context of preference modeling. We plan to explore whether alternative measures yield better performance in the biased search process.

It remains to be seen whether the “diverse” plans that we are generating are useful for the end user. However, on initial inspection, they appear to be reasonable plans that effectively show some of the key tradeoffs in the evaluation space.

Related Work

The problem of school redistricting is related to that of political redistricting. Several software packages (such as Mapitude (Caliper Corporation 2006)) are available for build-

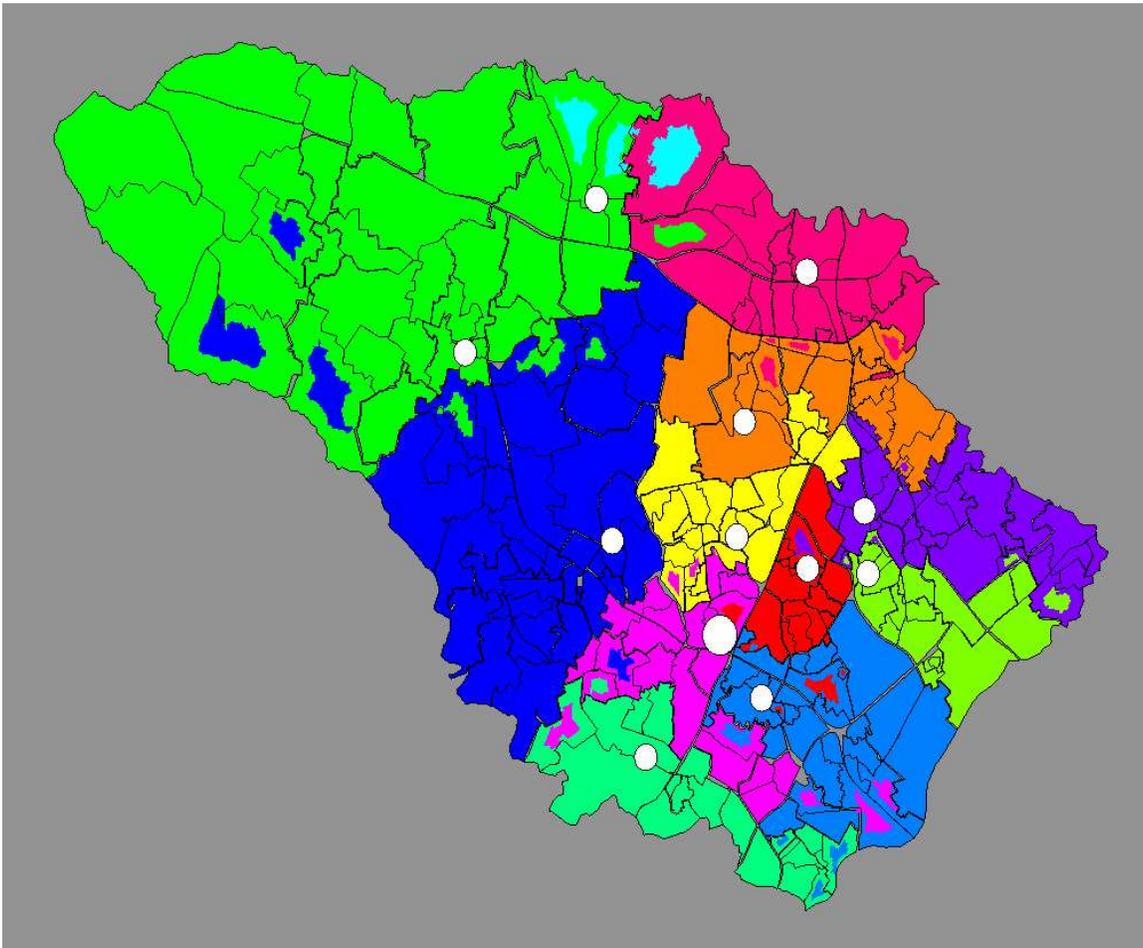


Figure 5: A comparison of one of the best plans (according to the F measure) found by weighted hill climbing (inner ring) to the recommended (green) plan (outer ring).

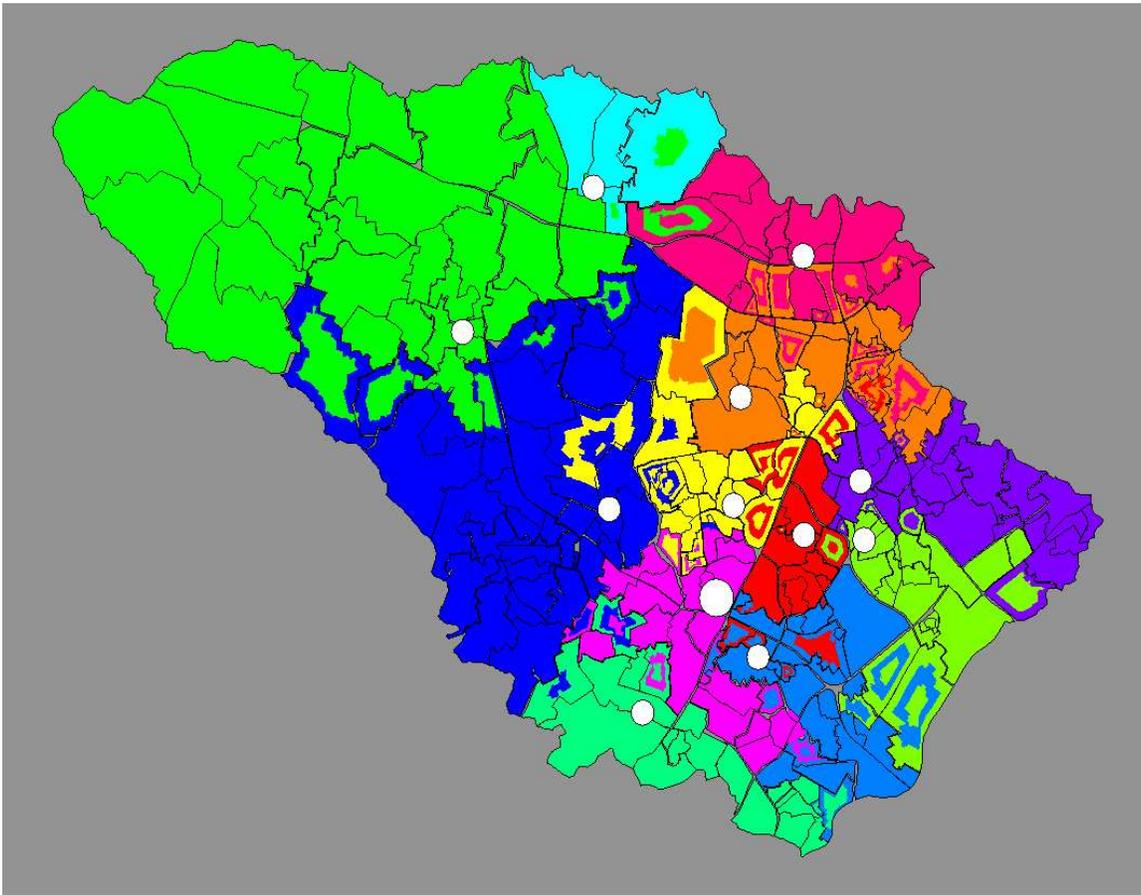


Figure 6: Three-way comparison of plans from a representative run of weighted hill climbing.

ing and analyzing political and school redistricting plans. These packages do not generally provide automated or interactive search methods, do not provide visual comparison techniques such as our “tree-ring” comparison, and do not facilitate the discovery of qualitatively different plans.

Heldig, Orr, & Roediger (1972) were among the earliest researchers to discuss computational approaches to political redistricting. The focus of their approach is on geographic criteria (compactness, contiguity, and “preservation of natural and/or political boundaries”) and population balancing, although they also mention the possibility of considering other criteria, such as demographic distributions. Their approach is based on linear programming, minimizing an objective function that is specifically designed to maximize geographic compactness of the districts, subject to a population-balancing constraint. Variations of this basic approach form the core of most of the more recent computational approaches to redistricting.

Altman’s (1998) dissertation discusses the objective principles that should ideally be used in political redistricting, including population equality, compactness, and contiguity. He analyzes the computational complexity of political redistricting, and shows that different measures of geographic compactness can produce very different plans, supporting

our claim that it is important to generate multiple plans from different perspectives.

School redistricting differs from political redistricting in several important ways. First, although compactness is an important factor (both for community building and to minimize busing costs), it is not as important as in political redistricting. Second, the walk usage and feeder issues complicate the scenario for school redistricting. Third, redistricting occurs more frequently (at least in Howard County) than in most political districts, and students are greatly affected by the process. As a result, minimizing the number of students who are redistricted is also an important criterion. Finally, the nature of the decision-making process, in which alternative plans are explicitly compared and contrasted to each other, raises the desirability of generating multiple plans that represent different tradeoffs.

Although we do not yet address all of these issues in our work, we believe that the general optimization framework we have developed, based on local search methods, is more applicable than those that are commonly used for political redistricting, which typically use specialized optimization algorithms that focus primarily on geographic constraints.

The problem of multicriteria optimization (also referred to as “multiobjective” and “multiattribute” optimization) has

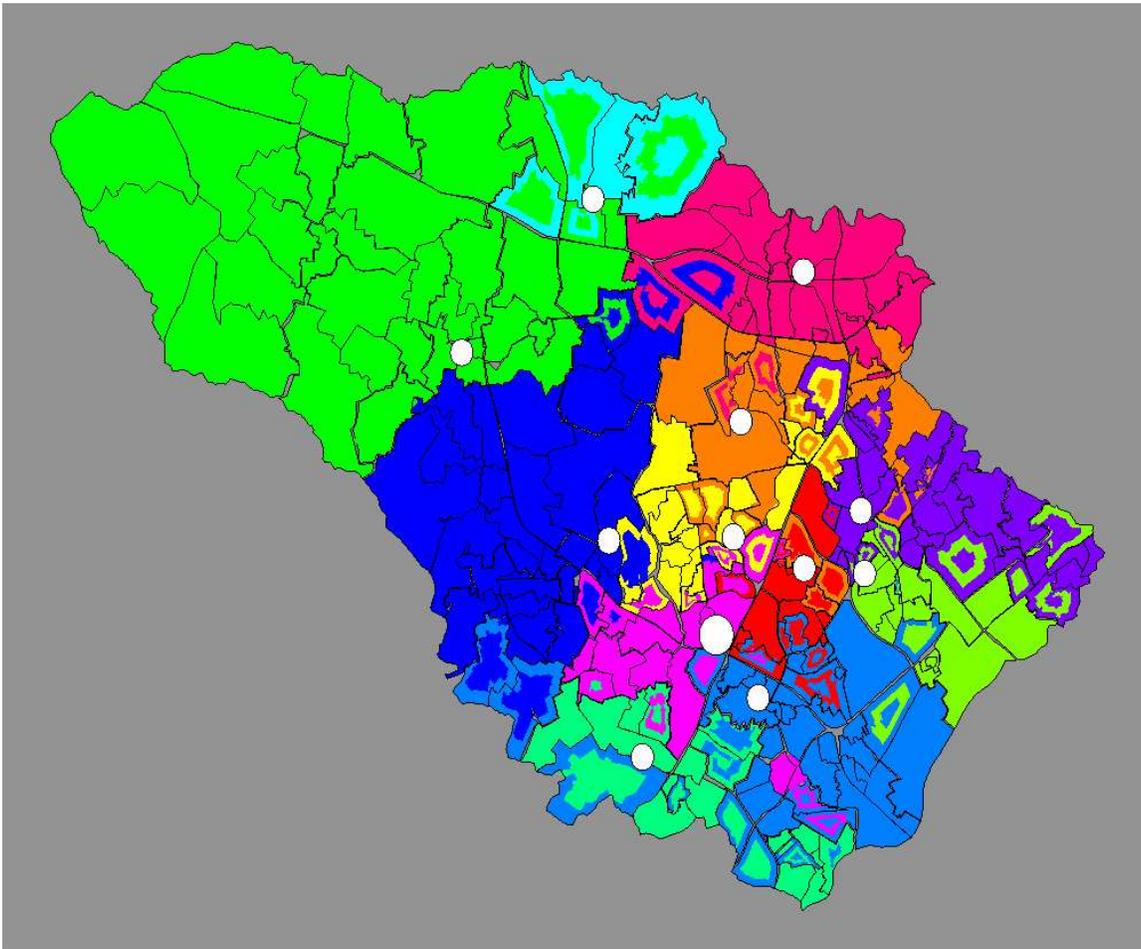


Figure 7: Three-way comparison of plans from a representative run of diversity-biased hill climbing

been explored by researchers in artificial intelligence, economics, and operations research. Keeney and Raiffa (1993) discuss a variety of approaches for combining multiple objectives into a single multiattribute utility function. They focus largely on methods for eliciting a single combined utility function. However, in practice, as Keeney and Raiffa point out, it is not always easy to create a single utility function. This supports our claim that a key component of a decision-making system for school redistricting is to provide tools that help the user to understand the nature of the evaluation criteria and the tradeoffs among them.

Multiattribute optimization techniques include weight-based optimization (where each criterion is assigned a weight, and a combined objective function is optimized), priority-based optimization (where the most important criteria are optimized first), and goal programming (in which one objective is minimized while constraining the others to be within a given range). None of these methods are ideal for our application, where the tradeoffs are difficult to prioritize or quantify. They also do not yield multiple qualitatively different solutions; to our knowledge, this problem has not been explicitly addressed.

School redistricting is somewhat analogous to the NP-complete problem of multi-objective graph partitioning (Selvakumaran & Karypis 2006), which attempts to optimize multiple objectives, each of which can be expressed as a sum of edge weights in a graph. Research on this problem has primarily used priority-based and weight-based optimization. The analogy to graph partitioning breaks down in the case of some of our criteria (walk usage, demographic distribution), and as mentioned above, priority-based and weight-based optimization methods do not help us with our goal of finding multiple qualitatively different solutions.

Future Work

Our future work focuses on four primary areas: usability, the criteria for plan evaluation, visualization of feeder systems and gradients, and new methods for multiattribute optimization.

Usability We have demonstrated the current prototype to the HCPSS Superintendent's Office, and received a very positive response. We plan to use our system to produce visualizations that highlight the tradeoffs made by the alter-

native proposed plans in the 2006–2007 redistricting process. Our goal is to make a web-based version of our tool available to the 2007–2008 redistricting committee for viewing, modifying, and evaluating proposed redistricting plans. We are also planning a more formal user study.

Plan Evaluation Criteria Additional evaluation criteria that we have explored include feeder quality, compactness, and robustness to future development.

In a pure feeder system, each elementary school would move as a group to a single middle school, and each middle school would move to a single high school. Because of geographic and capacity constraints, a pure feeder systems is impractical; nonetheless, maintaining “feeds” of reasonable size is an important criterion in the process. We have developed a feeder criterion to add to the set of evaluation criteria.

In most political redistricting approaches, compactness is one of the key factors—and, in fact, is sometimes the only factor that is considered. In school redistricting, compactness is desirable, but not explicitly mentioned in the HCPSS measured criteria. Furthermore, compactness is already captured indirectly in the busing cost and walk area criteria. However, we plan to explore whether significantly different plans are produced if compactness is added as an explicit criterion.

Finally, one ongoing problem with the redistricting process is that new developments (and new schools) are continually being built. This results in shifting demographics, and can mean that today’s best plan may lead to overcrowding at certain schools in a few years. The county does produce population projections, but these are inherently uncertain. Future research includes using these projections to measure the “robustness” of a plan to future planned development.

Visualization We are currently working on new visualization techniques that will show feeds more clearly. The problem is a challenging one from a visualization perspective, because showing multiple school regions simultaneously produces a significant amount of “clutter” and corresponding cognitive overload.

Another focus of our visualization research is on gradient displays. The *gradient* along a boundary between two school regions characterizes how the plan would change if a polygon on one side of the boundary were moved to the other school. If the gradient is positive, then the plan tends to “pull” the polygons along the boundary towards the first school. If it is negative, then the plan tends to “push” the polygons away from the first school. Each evaluation criterion can have its own gradient, so some of the gradients may pull, while others push. Developing effective ways to visualize these “tensions” on the plan will give the user insight into how local changes to the plan would change its evaluation. Supporting these displays will also require novel computational techniques for computing and summarizing gradients.

Multiattribute Optimization We are also exploring different methods for multicriteria optimization. We have im-

plemented a simulated annealing approach, which is a common technique for such problems. The preliminary results for this approach are not particularly encouraging. Part of the problem is that there are many regions of the search space that are quite poor (e.g., plans that create islands of students traveling a large distance). Adding an annealing step tends to lead to poor plans in this search space unless the cooling schedule is extremely slow.

We also plan to implement an evolutionary search approach, and are developing a novel multi-agent approach. This approach builds on recent research in combinatorial auctions in a scenario where each polygon can “bid” on schools to attend.

Conclusions

School redistricting is an interesting and challenging problem both computationally and from an application perspective. We have developed a prototype system that uses novel heuristic search and visualization techniques to aid an end user in generating, evaluating, and comparing alternative plans. These tools should provide end users with significant insights into the tradeoffs among alternatives.

The school redistricting problem is closely related to the resource positioning problem of deciding where to build schools, locate fire or police substations, or position emergency response equipment. Our optimization framework, search methods, and visualization tools can all potentially be applied to these other application domains.

Acknowledgements

This work was supported by the National Science Foundation (grant #IIS-0414976). We would like to thank David Drown, Joel Gallihue, and Larry Weaver of the Howard County Public School System’s Superintendent’s Office for their support, data, and valuable feedback. Thanks also to Poonam Shanbhag, who wrote the original visualization software, and to Andrew Hunt and Priyang Rathod, who also worked on the implementation.

References

- Altman, M. 1998. *Districting principles and democratic representation*. Ph.D. Dissertation, California Institute of Technology.
- Caliper Corporation. 2006. Maptitude for redistricting. <http://www.caliper.com/mtredist.htm>.
- desJardins, M., and Wagstaff, K. 2005. DD-PREF: A language for expressing preferences over sets. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*.
- Heldig, R. E.; Orr, P. K.; and Roediger, R. R. 1972. Political redistricting by computer. *Communications of the ACM* 15(8):735–741.
- Keeney, R. L., and Raiffa, H. 1993. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press.
- Selvakkumaran, N., and Karypis, G. 2006. Multi-objective hypergraph partitioning algorithms for cut and maximum

subdomain degree minimization. *IEEE Transactions on Computer Aided Design* 25(3). Forthcoming.

Shanbhag, P.; Rheingans, P.; and desJardins, M. 2005. Temporal visualization of planning polygons for efficient partitioning of geo-spatial data. In *Proceedings of InfoVis 2005*.