April 17, 2012

Denver Community Members,

Until late 2011, Denver Public Schools' (DPS) enrollment process included over 60 different paths for students to enroll in school. These multiple timelines, forms, and expectations caused inequities among Denver residents trying to enroll their children [(Enrollment Inequities)](). In response to these inequities, and with the help of Get Smart Schools, DPS began collaborating with its multiple constituents – including school leaders – to develop a coordinated and unified enrollment system called SchoolChoice. The SchoolChoice enrollment system was designed to allow all families the opportunity to select their highest priority schools without having to go through multiple forms and timelines.

Though streamlined and simplified, the new SchoolChoice process allows schools – including charter schools – to maintain their own unique school priorities and autonomies, while adhering to core principles throughout the process. Over the past two years of analysis development and rollout, third-party stakeholders were engaged with DPS in policy development, funding and community engagement to support this large effort. One of the outcomes that came out of this process is the SchoolChoice Transparency Committee.

In January 2012, the SchoolChoice Transparency Committee was created to receive and interpret two separate third party (The University of Colorado Denver) reports analyzing the SchoolChoice process and data. The committee is made up of school leaders (with representation from charter, magnet, and traditional schools), district leaders, and third-party community stakeholders. Members were selected by A+ Denver and DPS and have SchoolChoice background knowledge.

This letter addresses the University of Colorado Denver's first analysis, which examined the methods of assigning students to schools (see Appendix). The purpose of the initial algorithmic analysis was to ensure the SchoolChoice system upheld all entrance criteria and priorities for all schools and that the procedures set serve the needs of both students and schools in a fair and transparent manner.

Dr. Gary Kochenberger, a professor and co-director of UCD's Decision Science program, spearheaded the first analysis. He was selected based on his experience in designing and testing algorithms for large-scale optimization problems. Dr. Kochenberger spent a number of days inside DPS's offices working with the team that designed the algorithm, watching the algorithm as it was tested for functionality, and asking related questions. His report (see Appendix) provides a detailed description of his process, what he was looking for, what he observed and when he observed it. Dr. Kochenberger also included analysis of the coding of the algorithm for completeness and explained the procedures of the algorithm in a step-by-step process.

On Wednesday, March 21ˢᵗ, 2012 the SchoolChoice Transparency Committee met to hear and receive Dr. Kochenberger's memo and analysis.

Dr. Kochenberger identified three inputs (student preferences, school capacities and school priorities) and laid out the steps which the tool uses to analyze the information. Additionally, Dr. Kochenberger determined the system was run correctly and without special treatment, that schools' priorities and capacity were properly included in the algorithm, and that the data collected from students was accurately reflected through the input process with minor data entry error.

The Transparency Committee discussed problems encountered in the SchoolChoice process, including human error in data entry and confusion around schools' allocation of mid-year entry slots. Overall feedback from the schools represented on the committee was positive with a request for more raw data to include the entire district and not just their school bubble. One issue of concern for Dr. Kochenberger and the Transparency Committee was that the average form submitted had only about 2.5 choices listed. Because of this, we know families did not fill out all 5 slots. The group discussed this information and concluded families were only interested in choicing into 1-3 schools. Dr. Kochenberger also communicated the process was run a total of nine times, which included eight practice runs to resolve system errors.

Overall, the SchoolChoice Transparency Committee accepts Dr. Kochenberger's analysis of the process and his detailed description of what he observed, when he observed it and what he was looking for. His analysis provides an in-depth look at the steps followed within the algorithm and a determination that the system accurately reflects the inputs. Dr. Kochenberger's work and discussion with the Transparency Committee demonstrates to us that the tool performed strongly, as DPS communicated it would. His identification of issues related to human data entry error and partially completed forms are of concern to the Transparency Committee and will need to be addressed by DPS through further automation and development of precise process and procedures.

We look forward to receiving the second portion of the University's data analysis in the coming month. If you have any questions or comments, please do not hesitate to contact Van Schoales at van@aplusdenver.org or 303.725.1151.

Sincerely,

SchoolChoice Transparency Committee

University of Colorado
Denver | Anschutz Medical Campus

April 9th, 2012

To: Transparency Committee
From: Dr. Gary Kochenberger
Re: Assessment of Assignment Tool

Overview:

The new assignment tool (algorithm) employed by DPS this year is a version of the Gale-Shapely algorithm which is designed to solve the two-sided assignment problem. This procedure has been known in the literature since the 1960's where it became popular as the solution for the hypothetical "stable marriage" problem. In recent years the algorithm has found numerous practical applications including its use in student assignment settings such as the DPS student choice program. The general intent of the procedure is to produce a set of assignments that serve the needs of both the students and the schools in a fair and transparent manner.

The algorithm produces a set of assignments based on three major inputs—student preferences, school capacities, and school priorities. Given these inputs, the tool iterates in the following manner until no further assignments can be made:

Step 1:
- Each student requests their 1st choice (their # 1 preference)
- Each school uses its priority order to *tentatively* assign seats to students

Step 2:
- Students not getting their 1st choice now request their 2nd choice
- Each school considers a pool of students consisting of
  - Students already tentatively accepted
  - New students
- Each school uses its priority order to *tentatively* assign seats to students

Step 3:
- The algorithm repeats this process, moving down the student preference order, until
  - All students have been assigned; or
  - Capacity has been reached at the requested schools

Note:
- The output of the algorithm is a set of assignments where students are assigned to a school they prefer unless that school is full to capacity with higher priority students.
- Students assigned to a school lower than their top preference school will be waitlisted at each school they ranked higher.

Some Basic Results:

A complete presentation and analysis of the results will be undertaken by others. Here I present some key results that highlight the performance of the algorithm:

Placement of Participants Who Made at Least One Choice						
	Assigned to 1st Preference	Assigned to 1st or 2nd Preference	Assigned to 1st, 2nd, or 3rd Preference	Assigned to 1st, 2nd, 3rd, 4th Preference	Assigned to 1st, 2nd, 3rd, 4th or 5th Preference	Not Assigned to One of Their Preferences
ECE	61.00%	70.28%	74.85%	76.82%	78.32%	21.68%
Kinder	79.93%	87.68%	92.07%	92.07%	93.01%	6.99%
1st-5th Grade	54.34%	63.20%	66.99%	69.53%	70.43%	29.57%
Middle School (6-8)	70.14%	82.25%	86.55%	87.98%	88.50%	11.50%
High School (9-12)	73.26%	84.93%	88.91%	90.11%	90.57%	9.43%
DPS Total (EC-12)	69.64%	79.59%	83.44%	85.10%	85.95%	14.05%

* Non participants were removed from the above totals

EC-12 Round 1 Participants That Listed 5 Preferences and Were Assigned to None of Them	
Grade	Number of students
ECE	224
Kinder	**108**
1st	21
2nd	14
3rd	7
4th	5
5th	0
6th	**8**
7th	2
8th	4
9th	**0**
10th	0
11th	0
12th	0
DPS Total	393

Performance of the Assignment Tool:

In preparation for assessing the performance of the algorithm, I reviewed the relevant literature on the two-sided assignment problems (aka, the Stable Marriage Problem) and read accounts of this approach to school choice. In all, the literature here is substantial in size with notable pieces such as "Practical Market Design," by Atila Abdulkadirogula, Parag Pathak, and Alvin Roth, "The Design of School Choice Systems in NYC and Boston: Game Theoretic Issues," by A. Roth,   Atila Abdulkadirogula, Parag Pathak, and T. Sonmes, and "A Work Plan for Enrollment Reform in DPS," by the Institute for Innovation in Public School Choice, providing especially useful background for the task at hand. Several other papers were read as well.

My direct involvement with DPS started in December, 2011, and continues to the present. Key formal interactions are re-counted below:

> December 22, 2011: Conference call involving key players from DPS, the Transparency Committee, and UCD. This meeting provided an overview of the entire project, introduced the team, and generally set the tone for all that followed.

> February 3, 2012: Met with the DPS staff (Shannon Fitzgerald, Cathy Kidder, and Linnea Lewis) to discuss the algorithm and the preparation/testing underway for the actual running of the assignment tool later in the month. During this meeting we reviewed the overall project and talked about a broad set of topics ranging from data collection, to data entry and output testing.

> February 16, 2012: Met with DPS staff and Atila Abdulkadirogula (representing the team Implementing the assignment tool) to discuss the status of the data collection and entry, current testing of the algorithm, and general preparation for the final runs to be made in the weeks ahead.

> February 21, 2012: Meet with Tyler Lyons and Van Schoales to discuss a variety of issues of concern to the transparency committee.

> February 22, 2012: Met with the DPS team along with Neil Dorosin (from the Institute for Innovation in Public School Choice) and (by phone) Parag Pathak (representing the team Implementing the assignment tool) to discuss last minute implementation issues and general readiness for live running of the assignment tool.

> March 21, 2012: Met with the transparency committee to give an oral report on my assessment of the performance of the assignment tool.

> *Note: the above meetings were supplemented by phone calls, email interactions and additional face-to-face meetings with the DPS Staff as needed.*

In my discussions/investigations regarding the assignment algorithm, my focus was on the results obtained from various stages of testing and whether or not the results matched

my expectations. In the weeks leading up to "going live," several tests were made designed to check the output and provide a reality check regarding algorithm performance. A comprehensive set of tests (known at DPS as the "play book") were applied addressing issues ranging from the correct coding of student preferences, proper handling of school priorities, and proper handling of sibling relationships. Additional test on the output side of the tool checked to see, for instance, if every student assigned to a school had higher or equal priority than every other student who was not assigned to that school.

Discussions with DPS staff regarding this testing and the results obtained led me to the conclusion that the algorithm was both well-conceived and well implemented. Unexpected outcomes encountered in various tests of the algorithm raised red flags that were investigated and resolved, prompting another run of the assignment tool. Some examples of issues that necessitated the re-running of the algorithm are:

- Most boundary schools want to accept all of their "currently attending" choice applicants, even though the Board of Ed Policy does not guarantee acceptance for these kids. After an early run, the DPS team identified all currently-attending kids who were waitlisted, called the principals at the schools, and asked if they'd like to accept the students (in other words, increase their number of seat offers to accommodate them). In the vast majority of cases, the principals said yes, so we increased the seat offer number in those grades and ran the algorithm again. In the second run, the desired kids were then approved.
- After an early test run, the DPS team checked each school's list of approved and waitlisted kids to make sure that the appropriate students got in. In the early runs of the algorithm, we identified some errors in our prioritization schemas (for instance, siblings at the bottom of the waitlist) at a few schools. We then corrected the priorities and re-ran the algorithm to make sure the appropriate kids were, indeed, being approved.
- At schools with HGT (highly gifted and talented) programs, the seat capacity for the programs was driven by the number of HGT kids who wanted to get in, which we did not know with certainty until several runs of the algorithm had been made. For example, Teller Elementary told us in January that they could take 7 new 3rd graders, with HGT kids receiving "first dibs" on the seats. Based on earlier runs, we could tell that there were 5 HGT kids who needed to be accommodated, which left 2 new seats for the general education kids. Thus, we changed the seat offer numbers in this grade at this school to 5 HGT and 2 Traditional and re-ran the assignment tool.

In all cases where the tool was re-run, the motivation was to "get it right" to the extent possible for both the students and the schools.

In the end, the performance of the algorithm depends on the student preference data, school capacity data, and school priority data. There is no way to guarantee that all the data handling was error proof. Given the magnitude of it all, I'd be surprised if all the input coding was in fact completely error proof. I'm convinced, however, that any errors that weren't detected were small in number.

Separate from my focus on the functioning of the algorithm, the UCD team (Tracey Obrien and others) is working on an extensive analysis of the algorithm output and its full range of educational implications. Independently of the UCD effort, the DPS team (Shannon Fitzgerald and her staff) are undertaking their own assessment of the results. While I think things went very well in this first year of implementation, it's possible that the output analysis teams will identify a few assignment characteristics that may motivate small adjustments for the algorithm for next year. Any such lessons learned can only lead to even better performance in the coming years.

Overall Assessment:

In my opinion, the results strongly support the conclusion that the assignment algorithm performed as advertised. Namely, considering school priorities and the restrictions imposed by school capacity limitations, the algorithm assigned students to one of their preferred schools in a fair manner. Areas for improvement lie not so much with the algorithm itself but with issues pertaining to algorithm inputs. For instance, only slightly more than 6,000 students listed a preference for five schools and the average number of schools listed was roughly 2.5. Having more students state preferences for a complete set of five schools would enable the algorithm to consider a broader set of possible assignments, leading to even more desirable outcomes.

Once the output performance teams have completed their analysis, certain additional areas for improvement may be revealed. I'm confident these will be small in number and magnitude, but important nonetheless. I'm also confident that the algorithm is flexible enough to readily accommodate any such changes that might be needed.

Finally, I would recommend that the process of collecting student preferences and inputting the data be automated. While the number of "data errors" this year was remarkably small, moving to a higher level of automation would not only speed up the process, but would also lessen the likelihood of errors moving forward.

About the Author:

Dr. Gary Kochenberger is a professor of Decision Science at the University of Colorado at Denver where he is a co-director of the Decision Science program. His research focuses on designing and testing algorithms for large scale optimization problems. He has co-authored more than 70 refereed papers and three books.