to day, and even during the day based on weather and rainfall. Schedules created at the start of the day can have little relevance later on if the weather changes. The Aquadapt software generates a new optimized solution every half an hour as it adapts to changing conditions. The Aquadapt software also automates the entire water distribution process, providing consistent and reliable operation 24 hours a day.

WATERONE IMPLEMENTATION

The Aquadapt software runs on a pair of standard rack mounted server PCs. All user access is via thin client Windows applications running on the WaterOne Wide-Area Network (WAN). Figure 5 is a screen shot of one such application, the standard operator interface at WaterOne called the 'Operator Panel'.

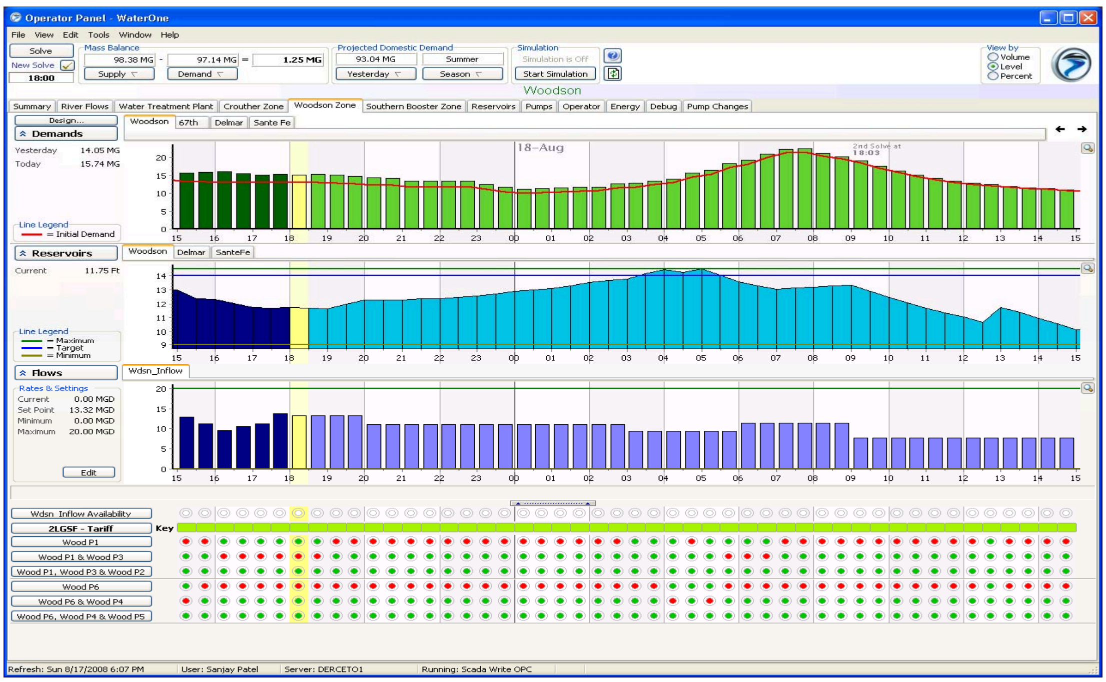

Figure 5. Derceto Aquadapt Operator Panel

This is a screenshot for a single pressure zone, and there are many similar pages for each of the other zones as well pages for the treatment plants, raw water sources, energy reporting data and miscellaneous supporting data. The top of the screen in Figure 5 reports summary information such as total system demand and supply, the time of the last update as well as access to operator inputs such as changing seasons or manually overriding setpoints. The Aquadapt software seeks to balance the supply and demand starting with the projected demand and then applying seasonal demand patterns for each zone with variations for weekday and weekends. The upper green graph in Figure 5 shows the actual and projected demand for water for this pressure zone for a 24 hour period. The yellow bar is the current time, with everything to the left of the bar being history and everything to the right being projected. The middle portion shows reservoir levels and flows in half-hour increments for the next 24 hours. The bottom portion shows pump schedules for each pump in half-hour increments, with in WaterOnes standard red meaning ON and green means OFF.

PROJECT MANAGEMENT AND PERFORMANCE AUDITING

WaterOne started a three phase implementation of the Derceto Aquadapt software in June 2004. The timeline for the phases are as follows:

- Phase 1 - Feasibility Study - June 2004 to August 2004
- Phase 2 - Detailed Design - September 2004 to April 2005
- Phase 3 - Configuration and Testing - July 2005 to Mar 2006
  - Delivery, Implementation and Site Testing - April/May 2006

The feasibility Study determined that WaterOne could achieve $500,000 annual savings from peak demand reduction and lowest cost source selection. In particular, by reducing electrical demand at the facilities supplied by BPU during the hours of 10:00 a.m. and 8:00 p.m. on weekdays, WaterOne can lower the demand charge.

In order to accomplish the electrical demand reduction, WaterOne would need to vary its treatment plant production by as much as 100 MGD during each 24 hour period, increasing the flow as much as possible during night time hours and cutting back as much as possible between 10:00 a.m. and 8:00 p.m. on weekdays. This approach is contrary to conventional wisdom and operator training which desires to keep the treatment plant flow as steady as possible to prevent upsets in the treatment process.

WaterOne decided to proceed with the Detailed Design to investigate further whether or not changes to the operation of the treatment process could be made to accommodate the large flow swings, without adversely impacting treatment processes and water quality, and fine-tune savings estimates. The results of the Detailed Design showed that an annual savings of $810,000 could be achieved by shifting demand, the payback for implementing the changes would be less than 2 years, and that minimal hardware and equipment changes were necessary for integration with the SCADA system.

Implementation timing was critical to getting the 2 year payback. The software needed to be installed by May 2006 which is the beginning of the summer season for the BPU electrical demand. Otherwise the majority of the savings would be lost for 2006/07. The Derceto software was installed by and running effectively in May 2006.

WaterOne created and maintained an audit tool to measure savings on a monthly basis. The tool used a step function to determine, based historical data, what the demand in kW would have been without Derceto software and compared it to the actual demand incurred when using the Derceto software. The difference was considered to be the monthly savings achieved with the Derceto software. Four years of historical data was used to develop a relationship between demand in kW and treatment plant flow in MGD. This approach to monitoring the savings allowed for changes in tariff in future years to be normalized in the evaluation.

RESULTS

Derceto helped WaterOne achieve a 20% electrical demand reduction resulting in a $90,000 savings that very first month. Figure 6 shows estimated daily demand reductions for the first month.

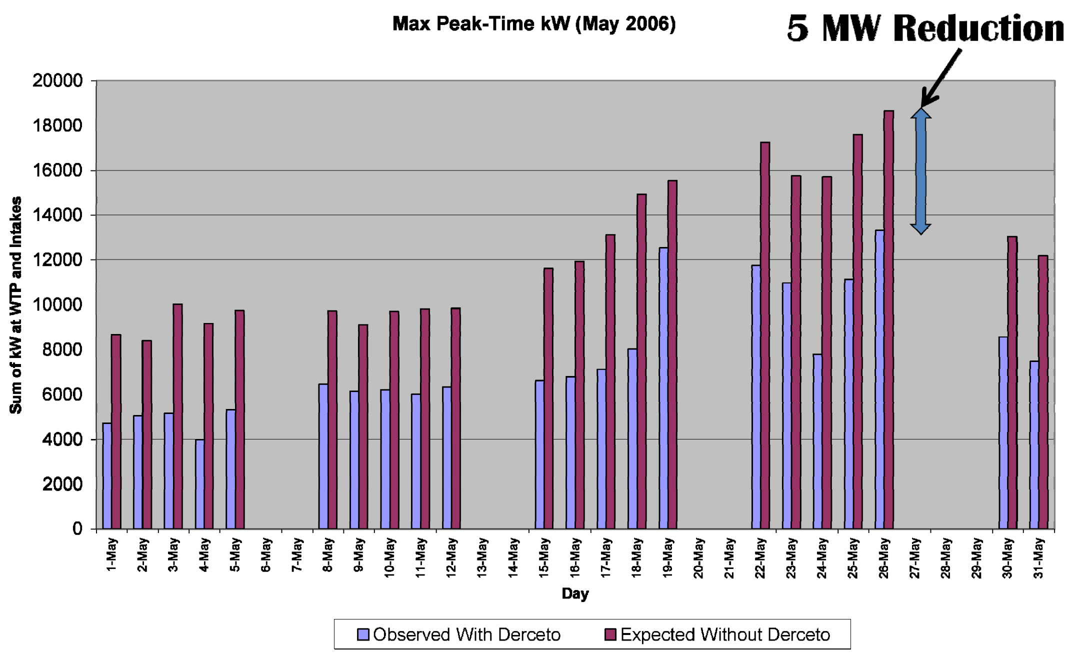

Figure 6. Electrical Peak Demand Reduction May 2006

Using the software, WaterOne was able to save over $1 million in electrical costs in the first 18 months of operation, reduce greenhouse gas emissions by 4,800 tons/year, identify areas where further savings can be made, and streamline operation of its water treatment plant. The reduction in energy use was so dramatic that BPU double checked their electric meters to make sure they were working properly and called WaterOne to verify that we had actually reduced our usage that much.

Operational changes were necessary to accommodate the flow swings that lead to the savings. The operators were involved from the very beginning to help ease the cultural change and shift in operating strategies. Before installing the software, Operators tested alternative ways to make the large flow swings and work out any bugs and gain confidence that the new way of operating would work. Ongoing training and product enhancements have lead to even greater efficiencies and smoother operations. Although not required, some electric valve operators were installed to help with the flow swings. This minimal cost enhancement increased operator acceptance and more than paid for itself in manpower savings and efficiencies. Operators have gained so much confidence in their abilities to handle the flow changes that they have pushed the flow swings to a greater amount than originally anticipated and are achieving even greater energy savings.

WaterOne is currently using the Derceto Aquadapt software to evaluate energy savings by increasing distribution system storage and other system modifications.

CONCLUSIONS

The considerable investment in a Scada system is often hard to justify financially, but it is an essential tool. Utilization of the data and interfaces presented by a Scada system in areas such as planning and operations optimization can provide additional financial benefits. Even in an energy market that looked unfriendly to energy optimization, significant benefits were achieved.

Buying advanced software “off-the-shelf” reduces risk. Use robust systems to measure the benefits to avoid doubt. Operators need to be involved from start to finish; they know the treatment and distribution systems best. This is just one example of taking SCADA one step forward. Keep an eye on costs and benefits of integrating SCADA in other areas to best serve your customers.

REFERENCES

Electric Power Research Institute, Energy Audit Manual for Water/Wastewater Facilities, (Palo Alto: 1999), Executive Summary.

ARC Population and Employment Allocation Disaggregator - A Simple GIS-based Tool for Parcel-Scale Population Projection

Stephen Bourne, PE [1], Mike Alexander, AICP [2], and Wei Wang [2]

[1]PBS&J, 2198 Goodwood Blvd, Smyrna, GA 30080; PH (404) 895-0753; FAX (770) 426-5316; email: sfbourne@pbsj.com
[2] Atlanta Regional Commission, 20 Courtland St., NE, Atlanta, GA, 30303; PH (404) 463-3100; FAX (404) 463-3205; email: MAlexander@atlantaregional.com, WWang@atlantaregional.com

ABSTRACT

Continued development of urban areas requires careful planning to ensure adequate resources are provided for growing industry and population while sustaining a healthy environment. This paper will discuss the on-going development of a GIS-based software tool, developed by the Atlanta Regional Commission (ARC), for projecting future development of the Atlanta Metropolitan Region. With this tool, the ARC can model growth of projected populations at fine scales, and thereby estimate future demands for transportation infrastructure, water, and other utilities.

The Atlanta metropolitan area is a 22-county region in the northwest corner of Georgia, with a burgeoning population of approximately 5 million. By land area, Atlanta is the fastest growing city in the world with some of the highest commute times in the United States. Continued development strains Atlanta's resources more and more. The ARC is dedicated to unifying the region's collective resources to prepare for a prosperous future. It does so through professional planning initiatives, the provision of objective information and the involvement of the community in collaborative partnerships.

To help adequately plan for future growth, the ARC has developed the ARC Population and Employment Allocation Disaggregation tool, an ESRI ArcMAP extension. The tool employs a two-step disaggregation procedure. Starting with a single region-wide population and employment projection (2000-2050) provided by a Regional Economic Models, INC (REMI) model, the tool first disaggregates this projection to large planning regions called super-districts. The ARC then adjusts the planning level projections manually to match their expectation for growth – and those of Atlanta's stakeholder community; city governments, developers, etc. The tool then disaggregates the superdistrict-scale projections to parcel scale by using map-based factors such as major roads and expressway ramps to estimate likelihood of development and allocating new growth to the most likely areas first. When the projection is at parcel scale, statistics for demands on the travel and utility networks can be calculated and used in forecasting future travel demands on the transportation system. The tool also provides a calibration procedure that compares modeled growth to actual growth and finds the set of model parameters that minimizes differences between the two.

INTRODUCTION

Continued development of urban areas requires careful planning to ensure adequate resources are available for growing industry and population while sustaining a healthy environment. This paper will discuss the on-going development of a GIS-based software tool, developed by the Atlanta Regional Commission (ARC), for projecting future development of the Atlanta Metropolitan Region. With this tool, the ARC can model growth of projected populations and employment at fine scales, and thereby estimate future demands for transportation infrastructure, water, and other utilities.

The Atlanta metropolitan area is a 20-county region in the northwest corner of Georgia, with a burgeoning population of approximately 5 million. By land area, Atlanta is the fastest growing city in the world with some of the highest commute times in the United States. Continued development strains Atlanta's resources more and more. The ARC is the municipal planning organization (MPO) dedicated to unifying the region's collective resources to prepare for a prosperous future. It does so through professional planning initiatives, the provision of objective information and the involvement of the community in collaborative partnerships.

This paper will discuss on-going development of the ARC Population and Employment Allocation Disaggregator tool in two sections, 1) algorithm description, and 2) technology description.

ALGORITHM DESCRIPTION

Step One: Disaggregation from Region-wide to Superdistrict scale.

The first disaggregation step starts with a regional projection provided by the ARCs Regional Economic Models, Inc (REMI) model, which includes 1) projections for 18 employment sectors, based on the North American Industry Classification System (NAICS) categories, and 2) projections for population by age, where there are 17 age categories. The NAICS employment projections are disaggregated to the 78 superdistricts by forcing the distribution of jobs across the superdistricts to follow the distribution observed in the base year – the first year of the time horizon.

The REMI model provides population projections as 17 population-by-age categories, which are first aggregated to a total population projection. To translate population into land use, the total population is translated to household size/income groups, where there are 6 household sizes (size is number of people in the household) and 4 income categories – giving a total of 24 household size/income categories. The households are also distributed across the superdistricts according to observed distributions in the base year.

Interactive Adjustment of Superdistrict-scale Projections

Once the tool has disaggregated the REMI projections to the superdistrict scale, users can make adjustments to the projections using the tool. The need for manual adjustments comes from 1) the tool assumes the base year distribution of jobs and housing will remain constant at the superdistrict scale over the projection time horizon, which may not be true, and 2) often, the ARC planners and the local stakeholders have accurate information about how development will occur in the short-range horizon (0-10 years) and medium-range horizon (10 -20 years). By allowing for manual adjustments, the tool can ingest this information. In the process, the tool can be used to build consensus on superdistrict-scale growth over the time horizon and can therefore raise confidence in ARC projections.

Of course, REMI-based regional totals of jobs and households – also know as control totals - must be matched. Manually adjusting projections will result in either a net increase or decrease in the total regional number of households or jobs. To conserve the control totals, the tool implements a re-balancing algorithm, which evaluates the total number of jobs or households moved by the user to a superdistrict and then compensates by removing a weighted fraction of jobs or households from the remaining superdistricts such that the total removed equals the total re-allocated.

Step Two: Disaggregation from Superdistrict to Parcel Scale

Once superdistrict scale disaggregation has been completed and approved, the second step of the algorithm is to further disaggregate projected growth to the parcel scale. At this point, the analysis becomes raster-based. In a year-by-year loop, the algorithm allocates projected growth first to employment land uses and then residential land-uses, emulating the actual development process, in which available land is occupied progressively over time. Figure 1 shows a schematic of how the algorithm proceeds through step two of the calculation.

The priority for allocation of new growth is driven by likelihood of development, which is evaluated for each raster cell according to a linear combination of map-based factors. For example, for a given scenario, the user might specify that commercial land use is most likely to develop in close proximity to interstate ramps, major roads, and in neighborhoods where commercial land use already exists. By bringing in map layers of interstates and major roads, the tool can evaluate a likelihood of development raster – called an L-raster. This L raster is calculated as the weighted sum of three rasters: 1) proximity to the closest major road, 2) proximity to closest freeway ramp, and 3) density of commercial land use in a neighborhood of raster cells around each grid cell. Figure 2 shows an example of an L-raster calculation for commercial landuse. The L-raster is the figure on the far right. Areas in blue represent the highest likelihood areas for commercial development to occur. New growth will be allocated to these areas first.

Each map factor is associated with a weight, which gives the importance of one factor relative to the other factors. The weights must add up to 1.0, and can be set by the user. The choice of map factors and their associated weights is decided by the user, making the tool flexible for doing what-if scenarios and identifying the set of factors that are most influential to new growth – see the calibration section of this paper for more detail.

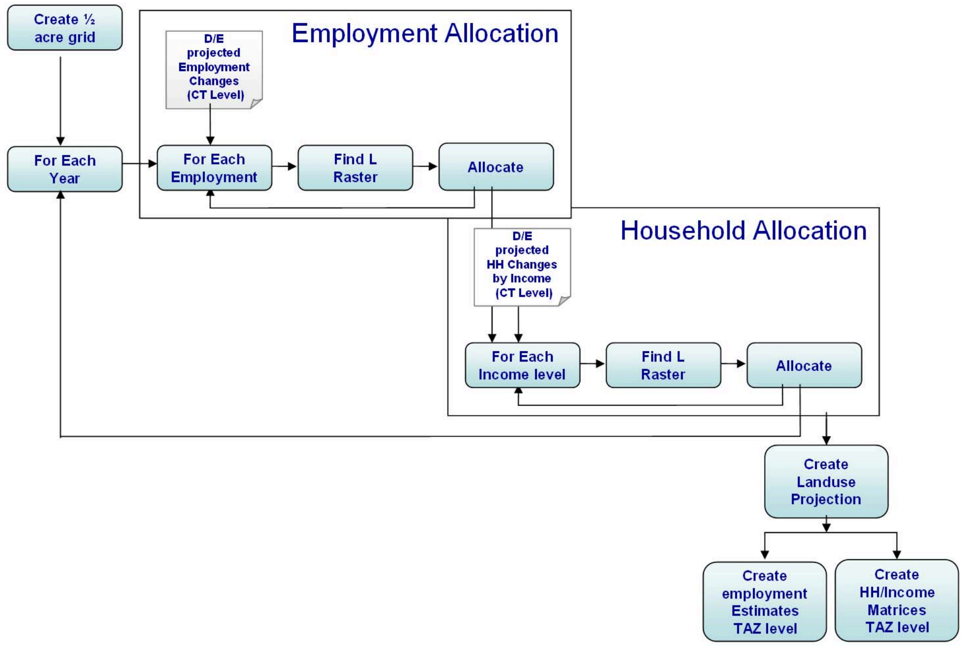

Figure 1: Schematic of Algorithm for Growth Modeling.

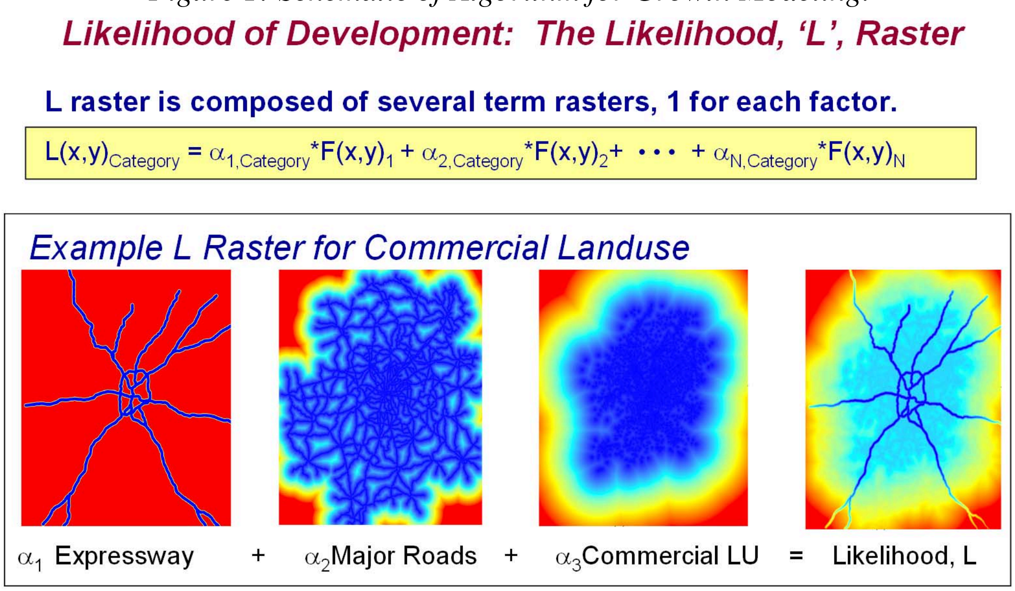

Figure 2: The L-Raster defines the likelihood of commercial or residential development at any given location by creating a weighted sum of user specified factors.

Translation from Employment and Households to Landuse

The superdistrict scale projections come in the form of numbers of new jobs and numbers of new households. As the algorithm allocates new growth by changing land use from vacant to developed, a suitable method for translating employment and households to land use is required. Using an ultimate build-out density map layer, the tool has an estimate at each location of the maximum density of jobs and households. The tool multiplies the number of new jobs or new households for any given year by the ultimate build-out density to arrive at a total number of acres required to contain the jobs or households.

Choosing the Raster Cell Size

In step two, the algorithm is converted from a vector-based to a raster-based calculation, where the vector-based shapes are converted to a grid. Using a raster-based approach is necessary because allocation during the second step must emulate the actual development process, where typically a developer will purchase and divide a large parcel into smaller ones and then develop multiple small parcels at a time in each given year. Converting to a raster-based analysis allows the algorithm to subdivide large parcels into smaller pseudo-parcels, in the form of raster grid cells.

An important aspect of converting to raster-based calculation is to choose the grid cell size of the rasters so that emulated development matches reality. Figure 3 shows the result of an analysis that found the size of actual changes in land use from 1999 to 2001. The rationale for the analysis is that changes in the land use polygons from year to year represent new development. We can set the grid cell size of the raster to match the most likely size of these changes. The implication is that 1) the growth modeling will be more accurate, as the chunks of land allocated to new development will match typical development, and 2) the calculation will proceed more quickly, as the basic grid cell size will be much larger than a single parcel.

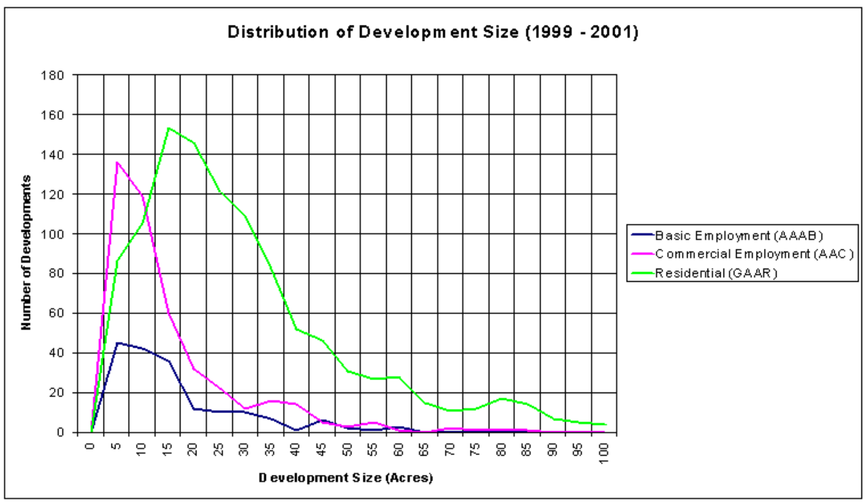

Figure 3: Distribution of Actual Landuse changes from 1999-2001.

For basic and commercial employment and residential landuse polygons, the difference in polygon sizes was calculated by essentially subtracting the 2001 polygons from the 1999 polygons. The distribution of the size of the changes was charted. From the Figure 3 on the preceeding page, we can see that for residential land use, the most likely size of new development is around 15-20 acres. Smaller average sizes for basic and commercial employment land use are found, 5-10 acres respectively. The grid cells size can be specified by the user in the ARC Population Disaggregation tool. The typical size used in our calculations has ranged from 500 ft (5.7 acres) to 1000 ft (23 acres).

Calculating Travel Demand Statistics

Once the algorithm has completed the two-step disaggregation process, a post-processing step is required to evaluate travel demand statistics for use in the ARC's travel demand model. The travel demand model requires employment and population estimates at the traffic analysis zone (TAZ) scale. The U.S. Census defines a TAZ as a special-purpose geographic entity delineated by state and local transportation officials for tabulating traffic related data from the decennial census, especially journey-to-work and place-of-work statistics. In the 13-county Atlanta metropolitan area, there are 1767 TAZs, with an average size of 1688 acres.

The tool calculates TAZ-scale totals for employment and households by income by translating the projected land use in each TAZ back to the NAICS employment and 24 category household size/income categories, ensuring that the distribution of jobs and housing matches that of the superdistrict.

Similar methods for aggregation of results to find demands for water, electricity, gas, and other resources can be added to the tool in the future.

Scenario-based Growth Modeling

In addition to the basic input data (map layers, REMI input, etc.), the user must specify the following to complete a run of the algorithm.

1. The time horizon of the analysis
2. The adjustments to superdistrict-scale projections
3. The map factors that influence development at the parcel scale for basic employment, commercial employment, and residential land uses.
4. The weights for the influential map factors

Collectively, this set of user specifications can be called a growth scenario. The ARC uses the tool to create multiple scenarios and compare the results. The intent can be to evaluate sensitivities to a certain map factor (eg. the effect of major roads), the effect of changing superdistrict scale projections to implement higher growth trends in one area versus another (eg. west Gwinnett County growing faster than east Gwinnett), and so on.