

## **SIMULATING WATER, INDIVIDUALS, AND MANAGEMENT USING A COUPLED AND DISTRIBUTED APPROACH**

Jonathan Ozik  
Nicholson Collier  
John T. Murphy

University of Chicago  
Computation Institute  
5735 S Ellis Ave  
Chicago, IL 60637, USA

Richard B. Lammers  
Alexander A. Prusevich

University of New Hampshire  
Institute for the Study of Earth,  
Oceans, and Space,  
8 College Road,  
Durham, NH 03824, USA

Mark Altaweel

University College London  
31-34 Gordon Square  
London, WC1H 0PY, UK

Andrew Kliskey  
Lilian Alessa

University of Idaho  
709 S. Deakin Street,  
Moscow, ID 83844, USA

### **ABSTRACT**

Water is a key issue in sustainable urban development. SWIM (Simulating Water, Individuals and Management) is an agent-based model of water supply, management structure, and residential water consumer perception and behavior. Initial work applied data mining on newspaper articles to map networks of water management institutions and structures. SWIM extends this by linking an agent-based model of residential water consumption connected via networks of water managers to a global-scale hydrological model. In our case study, we focus on Tucson, Arizona, where management and social behaviors are well documented. Census data are used to create synthetic populations of consumers endowed with price sensitivity and behaviors impacting water use. Social networks, including those based on geographic proximity, allow water use behaviors to spread to others. We examine possible factors leading to recent attested declines in per-capita water use, leveraging ensemble runs on high-performance computing resources using the Swift parallel scripting language to strategically explore complex parameter spaces.

### **1 INTRODUCTION**

Modern societies face increasing challenges in managing water, as the frequency and intensity of problems such as floods, water contamination, and water shortages continue to rise (UNESCO 2009). In many countries, including the United States, water management agencies are the forefront of managing and attempting to solve a variety of water issues. The prevailing systems of water management in the U.S. West and Southwest are fragmented and include a myriad of local, state, and federal agencies. Often authority in enforcing decisions differs, and even where there are strong authorities involved, jurisdiction of any agency is normally based on geographic boundaries that do not correspond to watersheds; alternative models include one large agency governing a given community (Kenney 1997). On the one

hand, aridity and growing cities would seem to create large-scale water shortages for communities of the West and particularly the Southwest. However, over the last two decades per capita water usage has, in fact, declined in some areas (Breyer et al. 2012).

Understanding why water use is declining in some communities is difficult, as a myriad of reasons for decreases in per capita consumption are possible (Olmstead & Stavins 2009). Nevertheless, this is an important issue, as managers need to know what factors, such as pricing mechanisms or conservation campaigns, are effective if water conservation success is to be replicated in cities across the West and beyond. Furthermore, water systems are inherently complex, coupled social-natural systems, requiring careful examination to understand how they are affected by the physical constraints of climate, hydrology, and by social factors such as water use behaviors, perceptions, and management decisions. In this paper, we propose a method that addresses such complex problems using Tucson, Arizona as our test case. The test case introduces the Simulating Water, Individuals, and Management (SWIM) agent-based model and demonstrates distributed ensemble Repast Symphony (North et al. 2013) agent-based modeling on large high performance clusters utilizing the Swift parallel scripting language (Wilde et al. 2011). This approach also applies the coupling of the agent-based model to a physical water balance model. We address possible reasons for the decline in per capita water use, and social and environmental reasons that lead to this decline.

## 2 METHODS

### 2.1 SWIM Model

The SWIM model simulates residential water consumption at the household level, linked with a regional model of watershed hydrology. Households pass consumption norms to geospatially bounded collections of neighbors and via non-geographic social networks. Primary factors driving water usage by consumers include water prices and dynamics of social influence and interactions in a manner comparable to the technique employed in Athanasiadis et al. (2005). Water pricing is determined using historical data from the case study of Tucson. The physical system is provided water by the Water Balance Model (WBM) (Wisser et al. 2009), which is driven by local and regional weather processes and geography.

#### 2.1.1 Water Consumer Agents

The primary entities in the SWIM model are water consumer agents. The consumption behavior of these agents is driven by water prices and social influences; social influence is transmitted via social networks, which encompass both geographically limited networks of neighbors and non-geographic associations. Exogenous factors, such as abrupt policy changes by water companies, can also lead to adjustments in water use rates. Neighbors influence social perceptions of water use, using a bounded confidence approach found in Weisbuch et al. (2002) that mimics the way that watching and perceiving what others in your neighborhood do could lead to an adjustment in one's own water use. A household social network affects water use by converging social attitudes between the overall network and the household's. In the current model the KE network is used (Klemm and Equiluz 2002). Price affects water use via an adjustable weight similar to that used in Athanasiadis et al. (2005). Additionally, there are seasonal and climatic factors that affect usage. The water use calculation is given by:

$$\begin{aligned}
 n &= p_{inf} * p_w + s_{f_{inf}} * s_{f_w} + e_{inf} * e_w \\
 c &= 1 + \left( d * (T_{annual} - T_{baseline}) - \left( r * \left( \frac{P_{annual} - P_{baseline}}{10} \right) \right) \right) \\
 u &= b * s * c * (1 - n) * (1 + o)
 \end{aligned} \tag{1}$$

where the overall water conservation “nudge” factor ( $n$ ) is based on the price influence ( $p_{inf}$ ) and associated price weight ( $p_w$ ), social factors influence, via social networks and local neighborhood influences ( $sf_{inf}$ ), and their associated weight ( $sf_w$ ), and exogenous effects ( $e_{inf}$ ) and weight ( $e_w$ ). The climatic factor ( $c$ ) is affected by seasonal and regional variations using degree percent change ( $d$ ), average annual temperature ( $T_{annual}$ ) of the current simulation year, and the baseline temperature ( $T_{baseline}$ ) over all the simulation years, percent decrease of water based on precipitation ( $r$ ), annual precipitation ( $P_{annual}$ ) of the simulation year, and baseline precipitation ( $P_{baseline}$ ) over all the simulation years. Water usage is then calculated based on baseline water use ( $b$ ), a seasonally adjusted water use factor ( $s$ ), climatic factor, nudge factor and water use noise ( $o$ ). The seasonally adjusted water use factor is derived from data on the average monthly deviation from a yearly average over the course of the simulation time span. The resultant factor is roughly sinusoidal function yielding a fractional increase or decrease for each simulated month. The water noise factor is based on typical monthly household water usage noise profiles, as described by Tucson Water. The noise factor used in the current model runs yields an up to 30% variation for individual households in month-to-month water usage. Figure 1 depicts the influences between the factors affecting household water use.

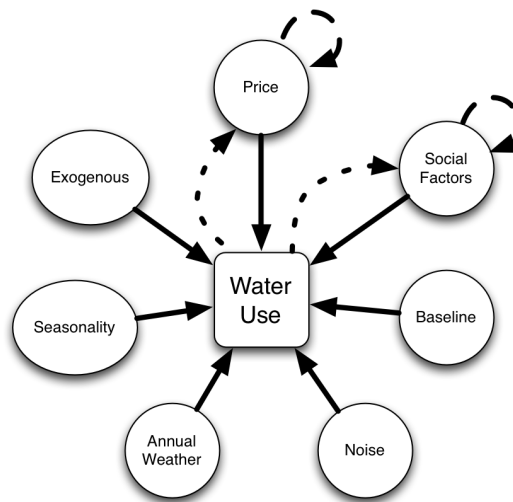


Figure 1: The water consumer agent water use influence diagram. The solid arrows are the direct influences in calculating an agent’s water usage. The dashed lines heading from water use to price and social factors indicate the influence that water use exerts, via water bills for price and via habituation for social factors. The self loops indicate the persistence, or inertia, of price and social factors within each agent over time.

### 2.1.1.1 Social Factors, Social Networks, Neighborhoods and Bounded Confidence

Household agents have a “social factors” property ( $sf_{inf}$  (1)) that reflects the social perceptions on water use held by the household and can be thought of as the propensity of the household to conserve water. While this propensity to conserve reflects a ‘conservation ethic,’ it is important to note that we consider this property of household agents quite broadly: it subsumes, for example, the physical properties of a home, such as the presence and size of a yard or the efficiency of faucets, toilets, and irrigation systems. For this study, we consider all of these in the aggregate and counterpose them collectively against the price factor (see below). Social factors are captured in a single value that ranges from -1 to 1; as a household’s social factors property increases, the household’s water use will tend to decrease and vice versa. For the simulation runs presented in this paper the households are initialized with random  $sf_{inf}$  drawn from a normal distribution with mean 0 and standard deviation 0.5, clipped to the -1 to 1 range.

There are three ways in which social factors evolve during a simulation; household-to-household (HH) interactions, household-to-neighborhood (HN) interactions, and habituation. In the SWIM model, households are connected through social networks. As part of the model initialization process, households are passed to a network constructor, which connects the households based on a specified network generation algorithm. While the current model utilizes the KE network (Klemm and Equiluz, 2002), it is a simple matter of specifying an alternative network generator to create other types of networks. The network connections define the HH's allowed interactions. At a specified interval households connected to each other will interact and, in doing so, will potentially modify their social factors properties. As part of the model initialization, households are also assigned neighbors based on the existence of other households within a specified radius. The presented model uses a radius of 400 meters. These neighbors are the basis of HN interactions. Unlike the dyadic HH interactions, HN interactions depend on the relationship of a household's social factors property to the average value of a neighborhood's social factors property.

In both HH and HN interactions, social factors are adjusted based on the bounded confidence mechanism (Weisbuch et al. 2002). The idea behind bounded confidence is that if the opinions held by two entities about a concept are within a certain threshold of each other, they will tend to converge in opinion over time. If, on the other hand, the opinions are further apart than the threshold, they are considered sufficiently divergent and, as a result, they remain unchanged. In the case of HH interactions, two households connected via their social network will interact and depending on their social factors at that specific time, will either move slightly closer to each other in opinion or remain unchanged. When HN interactions occur, a household will move slightly closer to the mean social factors property of its neighbors, or maintain its previous social factors.

As a household uses water, social factors are adjusted due to habituation as well. At specified intervals, yearly in the current model, the social factors are adjusted to reflect the notion that a household's concept of normal water consumption will be affected by what a household is habitually consuming.

#### **2.1.1.2 Household Data**

Each household agent has the following demographic properties: income, household size, age of the head of household, race and the location (latitude and longitude) of the household. These properties, as well as an additional id property, are initialized from a synthetic population produced by RTI International for the Models of Infectious Disease Study (MIDAS; Wheaton 2012). RTI's synthetic population database (Wheaton et al. 2009 and Wheaton 2012) includes geospatially explicit persons and households (as well as the workplaces and schools of these persons, although that information is not utilized in the current model). The persons and households are synthesized from US Census Bureau TIGER Data, Summary File 3 Data, and the Public Use Microdata Sample. Each synthetic household is geo-located. The location data is synthesized from Census block locations, using the LandScan population density dataset to assign each household a location within its Census block. The SWIM model uses the synthetic household population for Pima County Arizona, filtered to include only those households within the area served by Tucson Water.

Of the properties initialized from the synthetic database, only id and location are currently used by the model. The id is used to uniquely identify households and the location is used when constructing the geospatial neighbors of each household (see Figure 2). The inclusion of the additional properties opens up avenues for exploring water use models that can use household size, income and other factors in future work.

The households' initial values for baseline water use ( $b$  in (1)) are randomly initialized using a right skewed beta distribution. The particular parameters used for the beta distribution ( $\alpha = 1.045$ ,  $\beta = 3.919$ ) were chosen to most closely match the distribution of households in the 5 pricing tiers (see Section

2.1.1.4) and the mean household water consumption in Tucson for the year 1997, the first year for which we had access to the distribution of households in pricing tiers. One of the model parameters adjusts the maximum value of the beta distribution,  $b_{max}$ , and was varied in our numerical experiments to calibrate the model (see Section 3).




Figure 2: The spatial distribution of households in the RTI International synthetic population database restricted to Pima County, Arizona and filtered to include only those households within the area served by Tucson Water.

### 2.1.1.3 Anchor and Adjust Price Model

Agents have a memory of bills received; this affects the price factor calculated. Upon receiving a bill, the household agent will update the current value for its price effect, and will add the bill to its memory. The relationship of the bills that a household receives with the price effect can depend on a number of factors. The agent compares the value from the current bill to the average of the preceding  $N$  bills. If the current bill's value is higher than this average, the price factor  $p_{inf}$  is increased according to:

$$p_{inf_{t+1}} = \text{MAX}\left(p_{inf_t} \left(\frac{C}{C_N} - 1\right) M, 1\right) \quad (2)$$

where  $C$  is the current bill value,  $C_N$  is the average bill value over the preceding  $N$  months, and  $M$  is a scaling factor. If the bill's value is lower than the calculated average, price factor is allowed to decay according to:

$$p_{inf_{t+1}} = p_{inf_t} \left(\text{MIN}\left(1 - \left(1 - \frac{C}{C_N}\right)^\gamma, \Gamma\right)\right) \quad (3)$$

where  $\Gamma$  is a default decay multiplier and  $\gamma$  is a parameter of the model. For the simulations presented in the paper, the values for  $\Gamma$  and  $\gamma$  are .95 and 1.0 respectively. (Other values for these variables are possible and ranges of them may be explored in future work, but for the model runs presented here they are held constant.) The effects of equations (2) and (3) are that if the price increases above the recalled average, the price effect will be increased by a fraction that is equal to the fractional price increase multiplied by the scaling factor  $M$ ; if the price is below the recalled average, the price effect will decay,

minimally by  $\Gamma$  but perhaps by more if the drop is large. In either case, the new price factor is limited to values between zero and one.

#### **2.1.1.4 Water Prices**

Water pricing is one of the key levers that water managers can use to try to manage or reduce consumption (Olmstead 2013). In practice, however, the impact of pricing on water consumption is complex. While price increases can reduce consumption in some areas (for an example see Athanasiadis et al. 2005), in other places water prices are so low that even large percentage changes in price are irrelevant to customers (Martin and Kulakowski 1991). Moreover, managers' ability to price water is limited by the fact that water is essential and may be considered a public good; pricing water out of reach of poorer citizens would be unacceptable, and pricing it at its 'real' cost may be inappropriate or politically impossible (Olmstead and Stavins 2009).

A common practice that is relevant for our Tucson case is tiered pricing. Residents pay a low rate for an initial volume of consumption per month; once a threshold is reached, a higher rate is applied to the amount consumed above the threshold. Tucson's pricing structure has had multiple tiers for over two decades. The tiers used have allowed the average customer to pay the lowest rate; for a long time customers whose usage went above 15 Hundred Cubic Feet (CCF) per month begin to pay at a higher rate, and there are now three blocks above this with thresholds at 30, and 45 CCF. More recently in 2012, the lowest threshold has dropped further, to 10 CCF. The increase in rates at a threshold boundary can exceed 300%. The intent, clearly, is to discourage high usage while keeping normal and necessary use of water inexpensive for the majority of residents.

Tucson has implemented a number of additional billing components that make the typical resident's water bill even more complex. The bill includes a service fee, which is a flat fee that is based on a residence's tap size. It also includes a sewage charge. Because sewage is not directly metered, the volume of sewage is estimated from the water usage; however, because usage in the summer rises in ways that are not directly connected to sewage, the sewage estimate is calculated for the months of December – February, and this value is used for 12 months beginning the following July. Also, Tucson charges an additional fee if water usage in summer months (beginning in May) exceeds a fraction (e.g., 145%) of the metered winter use during the previous November-April.

The exact fractions, thresholds, and rates have changed several times over the past 20 years. The SWIM model implements the Tucson billing framework and uses values from historical data, adjusted for inflation to real dollars. Customer usage, including the usage needed for calculation of summer penalties and for sewage rates, is tracked based on each household's water use.

## **2.2 Model Validation via Ensemble Modeling**

As presented in the previous sections, the SWIM model has a number of parameters. A large portion of the current model development work involves sensitivity analyses of the model parameters for calibration and validation. Given the large number of simulation runs that are needed for proper parameter space coverage and adequate statistics, two batch simulation methods have been employed for large scale concurrent ensemble modeling.

### **2.2.1 Distributed Batch**

The Repast Symphony toolkit (North et al. 2013) comes with built-in distributed batch capabilities (Collier and Ozik 2014). Repast Symphony's batch run functionality iterates through parameter space combinations that the user provides and performs a run using each combination. These runs can be performed in parallel on a local machine (e.g., a laptop or desktop), on remote machines (secure shell [ssh] accessible resources), in the cloud (e.g., Amazon EC2, Microsoft Azure) or on a combination of the

three. Utilizing available desktop computing resources, we ran up to 10 computing instances in parallel, with each responsible for a number of simulation runs, the parameter space having been divided between the computing instances. While this is computationally adequate for exploratory work, the size of the parameter space forced us into considering more powerful computational resources for our more extensive parameter space explorations.

### **2.2.2 HPC Cluster with Swift**

Swift (Wilde et al. 2011) is a free and open source parallel scripting language for running applications in parallel on multi-core, cluster, cloud and supercomputing resources (<http://swift-lang.org>). Swift reduces the effort needed to create parallel workflows. Short scripts are written to define analysis applications along with inputs and outputs of those applications. Swift can then run a large number of applications concurrently as soon as the necessary inputs are made available. In this way, in addition to large scale parameter sweeps, dynamic and adaptive workflows can be defined. The same Swift script can be run on multiple types of machines, with minimal configuration file differences.

We developed a Swift script (Swift version 0.94.1) for launching Repast Symphony InstanceRunner jobs, along with bash scripts to wrap the invocation of application calls (see `repast.swift` and `repast_instance.sh` here <https://gist.github.com/jozik/5233968494b1fd16b639>). Our intention was to develop complementary capabilities with Swift that most directly utilized existing batch mechanism in Repast Symphony. Repast Symphony batch runs utilizes a SessionsDriver object that distributes parameters to local and remote machines and monitor the progress of launched LocalDrivers on each of the resources. The LocalDrivers, in turn, launch a specified number of concurrent InstanceRunner jobs, which process a set of parameter combinations. That is, each InstanceRunner receives what amounts to a set of tasks in the form of parameter combinations and completes each one of these tasks before proceeding to the next, until all the assigned tasks are completed. The Swift script that we developed essentially replaces the functionality of the SessionsDriver for high-performance clusters, while also providing scheduling, load balancing and fail over recovery capabilities. We are also developing the ability to run adaptive parameter space exploration for use in optimization (e.g., simulated annealing) and general computational experimentation with Swift and Repast Symphony.

For our current work, the main impetus for using Swift was to carry out large-scale model calibration parameter sweeps on the computing resources available at the University of Chicago, specifically the Midway cluster at the University of Chicago Research Computing Center ([http://rcc.uchicago.edu/resources/midway\\_specs.html](http://rcc.uchicago.edu/resources/midway_specs.html)). The cluster scheduling is managed by the SLURM resource manager (<https://computing.llnl.gov/linux/slurm/>). However, since Swift is responsible for most of the direct interactions with resource managers, the specific type of resource manager only minimally affects the configuration “site.xml” file that is provided to Swift when launching a job. Exploiting the capabilities of Swift and the Midway cluster, we were able to launch a parameter sweep job of 4320 parameter combinations, with each simulation run taking around 30 minutes, and run it efficiently on 512 computing resources in under 5 hours.

## **2.3 Water Balance Model Integration**

The SWIM model is integrated with a global-scale hydrology model called the Water Balance Model (WBM), developed at the University of New Hampshire. The integration of SWIM with the WBM allows bridging local-scale water consumption with its impact on regional hydrology; in the future this may be extended to understanding human impacts on the global water cycle.

The WBM models the global movement of water through the hydrologic cycle. It simulates both the movement of water from land to the atmosphere and the flow of water horizontally throughout the land surface. It employs grid system that can be set to scales as fine as 6” and operates at a daily time step. Its core accounting system tracks water entering and leaving each grid cell, and permits the specification of

reservoirs, dams, snow and glacial melt, irrigation, and inter-basin transfers. A separate river network can be specified to more accurately model surface flow.

The WBM provides to the SWIM model point sources of water. These correspond in the WBM to the water available on a single grid cell; in the ABM they are translated into water sources that can be made available to Water Managers, thus allowing the water tracked by the WBM to enter the simulated flow of the SWIM system.

SWIM and the WBM exchange information in a handshake that occurs on the daily time step at which both models run. SWIM initiates the exchange by sending a message to the WBM; this message contains the list of points (specified by latitude and longitude) at which SWIM will need to know the value for water available. The WBM responds with a list of these points and the water that is available at them, corresponding to the amount of water (in cubic meters) that is available at that point for that simulated day. SWIM calculates the net amount of the water available that will be extracted at that point and responds to the WBM with a list of points at which water was taken, and the amount taken from each. All of the SWIM messages, including the initial one, use the same format, so that this response acts as the prompt for the next message from the WBM. If SWIM adds a new point to the list, the WBM will respond with the water available at that point during the next simulation day. The WBM considers the amount taken to be urban use, and calculates its movement through the rest of the hydrological cycle (e.g. transpiration to the atmosphere) using its built-in algorithms. Future SWIM/WBM integration work will explore the possibility for SWIM to supply figures to WBM based on more specific, not solely restricted to urban, behaviors of agents.

### **3 RESULTS AND DISCUSSION**

The initial goals for the SWIM model have been to both calibrate it to the Tucson case study data and to determine the salient mechanisms required to adequately characterize water usage patterns at an aggregate level, specifically for the purposes of understanding aggregate water usage trends. In order to do this, we sought to first determine model parameter combinations that yielded water usage consistent with water usage data from Tucson.

The model was run for a time span of two decades, from January 1, 1993 to December 31, 2012, with a temporal resolution of one simulation ‘tick’ per day (7305 total ticks). Due to the large number of simulation runs that were anticipated and the exploratory nature of the initial investigations, rather than use the full set of Tucson households, we uniformly sampled 20% of the population, yielding in the vicinity of 52,000 households per simulation. At this sampling level, each simulation run took about 30 minutes to complete. The parameters that were varied were the price weight,  $p_w$  in (1) (0.0 to 0.2 in steps of 0.04), the social factors weight,  $sf_w$  in (1) (0.0 to 0.2 in steps of 0.04), the maximum value of the initial beta distribution of household baseline usages,  $b_{max}$  in Section 2.1.1.2 (1450 to 1600 in steps of 50 Gallons / day), the price scaling factor,  $M$  in Section 2.1.1.3 (0.3 to 1.5 in steps of 0.3). For each parameter combination we also ran 6 different random seed runs, for a total of  $6 \times 6 \times 4 \times 5 \times 6 = 4320$  simulation runs. Had we restricted ourselves to desktop computers for these runs, assuming on the order of 10 parallel processes, the job would have taken at least  $4320 / 10 \times 30$  minutes = 216 hours. Instead, distributing the model runs over 512 parallel processes on the Midway cluster at the University of Chicago’s Research Computing Center, enabled these runs to be completed in under 5 hours.

The model runs yielded a large spread of aggregate daily water usage outputs. At this stage of our investigations we have been most concerned with being able to match the general year over year trend seen in residential consumer water demand in Tucson. As such, the model outputs were aggregated into average monthly usage and further into 12-month rolling averages. The 4320 simulation runs were collapsed into 720 unique model parameter combinations by averaging the 6 stochastic variations for each model parameter combination.

The resulting 720 rolling average trajectories were compared to the Tucson average household 12-month rolling average for the same time period as the model runs. The SWIM model works by initializing



the synthetic population of agents at one point in time and evolving the population while observing how climatic, economic, social and other factors affect the water usage. As such, our interest was in finding simulation parameters that yielded good initial agreement first, and out of the best candidates observe if the general trends of the data are being followed. The motivation is that if, under the currently applied mechanisms, the simulation is not able to follow observed trends, we are forced to consider additional mechanisms since the model is effectively incomplete. We would thus build complexity into the model out of necessity, an important guidepost when considering the potential for a combinatorial explosion of parameter and mechanism combinations.

Figure 3 shows the top 10 best fitting runs in magenta and the top 40 best fitting runs in yellow where the fit was evaluated on the first 6 years (1994-1999) of the rolling average Tucson data, which is in blue. While we do find that there exist trajectories within the parameter space that we covered that closely mirror the first 6 year span of Tucson's average household water use, it is qualitatively evident that the general downward trend in the data after 1999 is not observed. It is possible that further parameter sweeps involving alternative values for the parameters that were already varied as well as for the parameters that were kept constant in these model runs would generate the overall downward trend. However these experiments suggest, and in the authors' opinion this is the more likely scenario, that price, climatic effects and passive social factors diffusion are simply not enough to capture the dynamics observed in the empirical data.

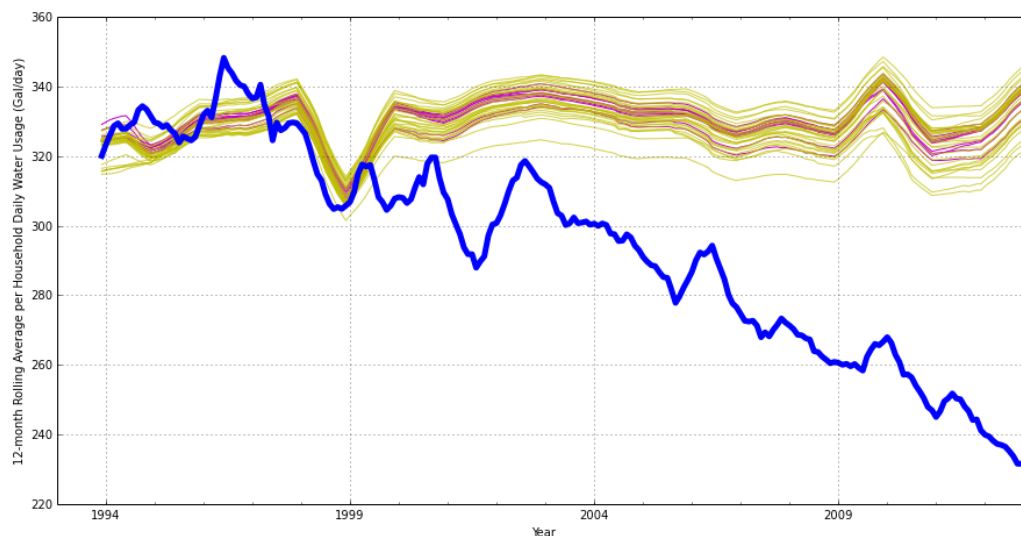


Figure 3: Top 10 (magenta) and top 40 (yellow) simulation runs in terms of first 6 year fits to the Tucson data (blue).

Given these results, there are a number of logical next steps to pursue for further development and validation of the SWIM model. First, Tucson Water and other water managers often employ informational campaigns to promote water conservation efforts. These campaigns can target the general population (e.g., public service announcements, billboards, campaign mascots [see [http://www.tucsonaz.gov/water/pete\\_the\\_beak](http://www.tucsonaz.gov/water/pete_the_beak)]) or they can be more targeted (e.g., messaging through specific neighborhood associations, focused on specific water use activities). They can promote conservation practices (e.g., xeriscaping, gray water recycling) as well as the adoption of conservation technologies (e.g., low-flow shower heads; also see Millock and Nauges 2010). We will exploit the disaggregated nature of the SWIM agent-based model to run experiments that implement active attempts at influencing household social factors. We would potentially see that the households affected by the informational campaigns would act as social factors diffusion sources, acting against or in concert with

the other factors affecting household water usage in the model. When the SWIM model is adequately validated for a locale, it may open up the possibility to be used as an *in silico* laboratory for evaluating future potential water conservation campaigns.

Second, water related legislation can also sharply curtail water use, especially in peak summer periods. For example, many municipalities forbid activities such as watering of lawns and filling pools when certain emergency conditions are met. While the emphasis of the current paper has been on 12-month rolling averages and annual temperature and precipitation effects, it would be possible to increase the resolution of the experimentation to account for seasonally sensitive restrictions to see if sharp declines in water consumption can be tied to these events.

Third, while initial investigations have focused almost exclusively on highly aggregate model outputs, there are a number of areas that a deeper investigation into water usage patterns may help calibration and validation efforts. One such area is the examination of households' movement between price tiers over time. Most of the price increases in Tucson have mainly targeted the highest tiers of water users, while also affecting, to a lesser degree, the lower tiers. Validating the model by comparing price effects on the distributions of households in these pricing tiers with empirical data on the distributions over time may help create a more accurate picture of price effects on consumption.

#### **4 CONCLUSIONS**

High-performance computing and the use of the Swift language have opened a new window into the investigation of the determinants and dynamics of residential, urban water use. Here we have presented the SWIM model of household water consumption and its deployment on a high-performance computing platform. Our test case has focused on the household water consumption agents in SWIM, and has been grounded in our initial case study of Tucson, Arizona. We have shown that the use of the high-performance platform has allowed us to productively explore a very large parameter space very rapidly, allowing us to experimentally calibrate our model and work toward an insight into the factors that drove the real-world case on which the model has been based.

Future work will include examining factors such as informational campaigns by water authorities, the diffusion of water conservation technologies, and water related legislation. This future work will also include extending the model to other urban regions in the West to investigate why some regions are experiencing water use declines while others may be increasing. For instance, urban heat islands have been postulated as one set of influences on residential water use (Guhathakurta & Gober 2007). Building on the methods presented here, such spatial and behavioral interactions could be investigated further.

In all of these future directions, powerful computational resources and techniques will need to be employed. Swift will form a key component of this, providing the ability to move our Repast simulations onto platforms that permit greater numbers of runs to be performed quickly and to efficiently evaluate the results produced. We will continue to utilize high-performance computing resources and develop more sophisticated adaptive parameter space exploration. The results from these cases will help us to better understand the relationship of technology and water use, such as the role of technology induced environmental distancing (Alessa et al. 2010), whereby one can begin to address how and which technologies may promote or detract from water conservation.

#### **ACKNOWLEDGEMENTS**

We are grateful to Tom Arnold at Tucson Water for his invaluable insights and data on Tucson area water consumption and billing. We would also like to thank Michael Wilde and the Swift team for the input they provided to the development of the Swift based parameter sweeping. This work was completed in part with resources provided by the University of Chicago Research Computing Center. This material is based upon work supported by the National Science Foundation under Grant No. BCS-1114851. This work is supported by the U.S. Department of Energy under contract number DE- AC02-06CH11357.

## REFERENCES

- Alessa, L. N., A. Kliskey, and P. Williams. 2010. "Forgetting Freshwater: Technology, Values and Distancing in Remote Arctic Communities." *Society and Natural Resources* 23:254–268.
- Athanasiadis, I. N., A. K. Mentis, P. A. Mitkas, and Y. A. Mylopoulos. 2005. "A Hybrid Agent-Based Model for Estimating Residential Water Demand." *Simulation* 81:175.
- Breyer, B., H. Chang, and G. H. Parandvash. 2012. "Land-use, Temperature, and Single-Family Residential Water Use Patterns in Portland, Oregon and Phoenix, Arizona." *Applied Geography* 35:142-151.
- Collier, N. T. and J. Ozik. 2014. "Repast Symphony Batch Runs Getting Started." *Repast Development Team*. <http://repast.sourceforge.net/docs.php>.
- Guhathakurta, S., and P. Gober. 2007. "The Impact of the Phoenix Urban Heat Island on Residential Water Use." *Journal of the American Planning Association* 73(3):317-329.
- Hintjens, P. 2013. *ZeroMQ Messaging for Many Applications*. Sebastopol, CA: O'Reilly Media.
- Kenney, D. S. 1997. *Resource Management at the Watershed Level*. Boulder: University of Colorado School of Law.
- Klemm, K. and V. M. Eguíluz. 2002. "Growing Scale-Free Networks with Small-World Behavior." *Physical Review E* 65:057102.
- Martin, W. E. and S. Kulakowski. 1991. "Water Price as a Policy Variable in Managing Urban Water Use: Tucson, Arizona." *Water Resources Research* 27(2):157-166.
- Millock, K., and C. Nauges. "Household Adoption of Water-Efficient Equipment: The Role of Socio-Economic Factors, Environmental Attitudes and Policy." *Environmental and Resource Economics* 46(4):539-565.
- North, M. J., N. T. Collier, J. Ozik, E. R. Tatar, C. M. Macal, M. Bragen, and P. Sydelko. (2013). Complex adaptive systems modeling with Repast Symphony. *Complex Adaptive Systems Modeling*, 1(1), 3.
- Olmstead, S., and R. Stavins. 2009. "Comparing Price and Non-Price Approaches to Urban Water Conservation." *Water Resource Research* 45(4).
- Olmstead, S. M. 2013. "Climate Change Adaptation and Water Resource Management: A Review of the Literature." *Energy Economics*. DOI: [10.1016/j.eneco.2013.09.005](https://doi.org/10.1016/j.eneco.2013.09.005).
- Schwarzenbach, R. P., T. Egli, T. B. Hofstetter, U. v. Gunten, B. Wehrli. 2010. "Global Water Pollution and Human Health." *Annual Review of Environment and Resources* 35:109-136.
- UNESCO 2009. *Water in a Changing World: The United Nations World Water Development Report No. 3*. UNESCO Publishing.
- Weisbuch, G., G. Deffuant, F. Amblard, and J-P. Nadal. 2002. "Meet, Discuss, and Segregate!" *Complexity* 7(3):55-63.
- Wheaton, W. D., J. C. Cajka, B. M. Chasteen, D. K. Wagener, P. C. Cooley, L. Ganapathi, D. J. Roberts, and J. L. Allpress. (2009). Synthesized Population Databases: A US Geospatial Database for Agent-Based Models. RTI Press publication No. MR-0010-0905. Research Triangle Park, NC: RTI International. Retrieved November 2012 from <http://www.rti.org/rtipress>.
- Wheaton, W. D. (August, 2012). 2005-2009 U.S. Synthetic Population Ver. 2. RTI International. Retrieved from <https://www.epimodels.org/midas/Rpubsyntdata1.do>.
- Wilde, M., M. Hategan, J. M. Wozniak, B. Clifford, D. S. Katz, and I. Foster. 2011. "Swift: A Language for Distributed Parallel Scripting." *Parallel Computing* 37:633–652 .
- Wisser, D., S. Frolking, M. E. Douglas, B. M. Fekete, A. H. Schumann, and C. J. Vörösmarty. 2010. "The Significance of Local Water Resources Captured in Small Reservoirs for Crop Production: A Global Scale Analysis." *Journal of Hydrology* 384(3-4):264-275.

## **AUTHOR BIOGRAPHIES**

**JONATHAN OZIK**, Ph.D., is a computational scientist at the Center for Complex Adaptive Agent Systems Simulation within the Decision and Information Sciences Division of Argonne National Laboratory and a Senior Fellow at the Computation Institute at the University of Chicago. His email address is [jozik@uchicago.edu](mailto:jozik@uchicago.edu).

**NICHOLSON COLLIER**, Ph.D., is a software engineer at the Center for Complex Adaptive Agent Systems Simulation within the Decision and Information Sciences Division of Argonne National Laboratory and a staff member at the Computation Institute at the University of Chicago. His email address is [ncollier@anl.gov](mailto:ncollier@anl.gov).

**JOHN T MURPHY**, Ph.D., is an anthropologist and a computational engineer at the Center for Complex Adaptive Agent Systems Simulation within the Decision and Information Sciences Division of Argonne National Laboratory and a staff member at the Computation Institute at the University of Chicago. His email address is [johntmurphy@uchicago.edu](mailto:johntmurphy@uchicago.edu).

**MARK ALTAWHEEL** is a Lecturer at University College London. He received a Ph.D. from the University of Chicago in Archaeology focusing on ancient land use and water systems. His research deals with how human societies adapt and change in order to respond to environmental stress in both the distant past and modern societies. His email address is [m.altaweel@ucl.ac.uk](mailto:m.altaweel@ucl.ac.uk).

**RICHARD B. LAMMERS**, Ph.D., is a hydrologist, earth system scientist and co-Director of the Water Systems Analysis Groups in the Earth Systems Research Center, Institute for the Study of Earth, Oceans, and Space at the University of New Hampshire. His email address is [Richard.Lammers@unh.edu](mailto:Richard.Lammers@unh.edu).

**ALEXANDER A. PRUSEVICH**, Ph.D., is a geophysicist and numerical modeler at the Water Systems Analysis Groups in the Earth Systems Research Center of the University of New Hampshire. His email address is [alex.proussevitch@unh.edu](mailto:alex.proussevitch@unh.edu).

**ANDREW KLISKEY**, Ph.D., is a Professor in the Department of Forest, Rangeland, and Fire Sciences at the University of Idaho. His research interests include human response and adaptation to environmental change and spatial methodologies for understanding coupled natural human systems and place-based science. His email address is [akliskey@uidaho.edu](mailto:akliskey@uidaho.edu).

**LILIAN ALESSA**, Ph.D., is the Director of the Center for Resilient Rural Communities at the University of Idaho and Professor of Resilient Landscapes. Her expertise is in social-ecological systems science and the development of highly integrated approaches to better understanding complex environmental processes. Lil is a member of the National Science Foundation's Advisory Committee for Environmental Research and Education. Her email address is [alessa@uidaho.edu](mailto:alessa@uidaho.edu).