


Experiences Using Windows Azure to Calibrate Watershed Models

Marty Humphrey, Norm Beekwilder
University of Virginia

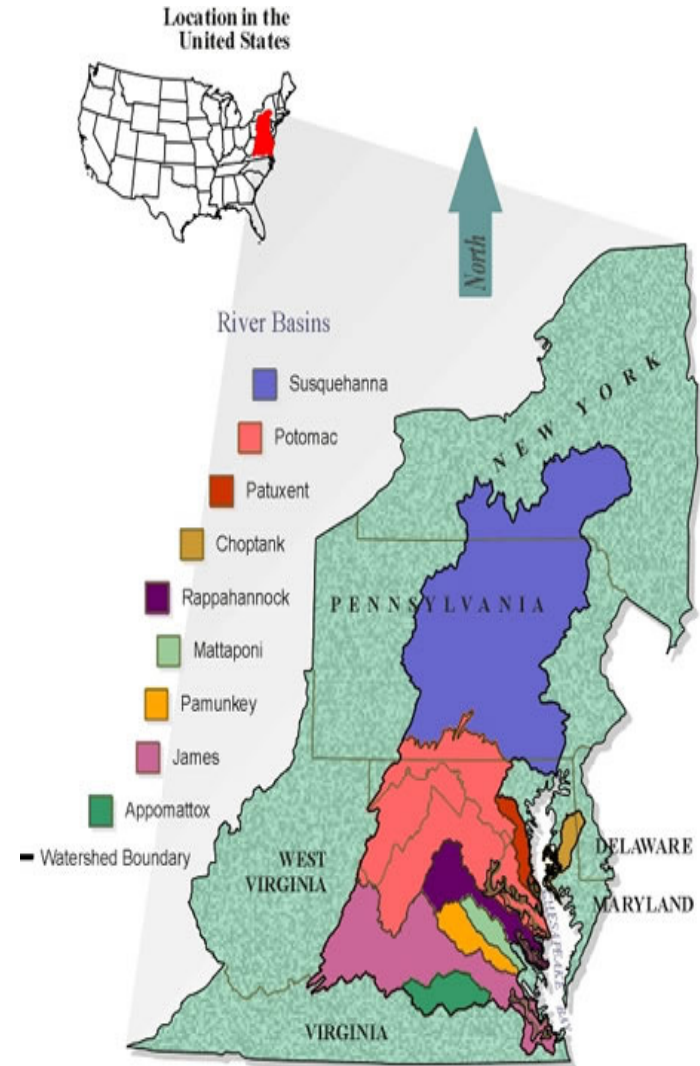
Jon Goodall, Mehmet Ercan
University of South Carolina

Example Large Watersheds

Mississippi River Watershed

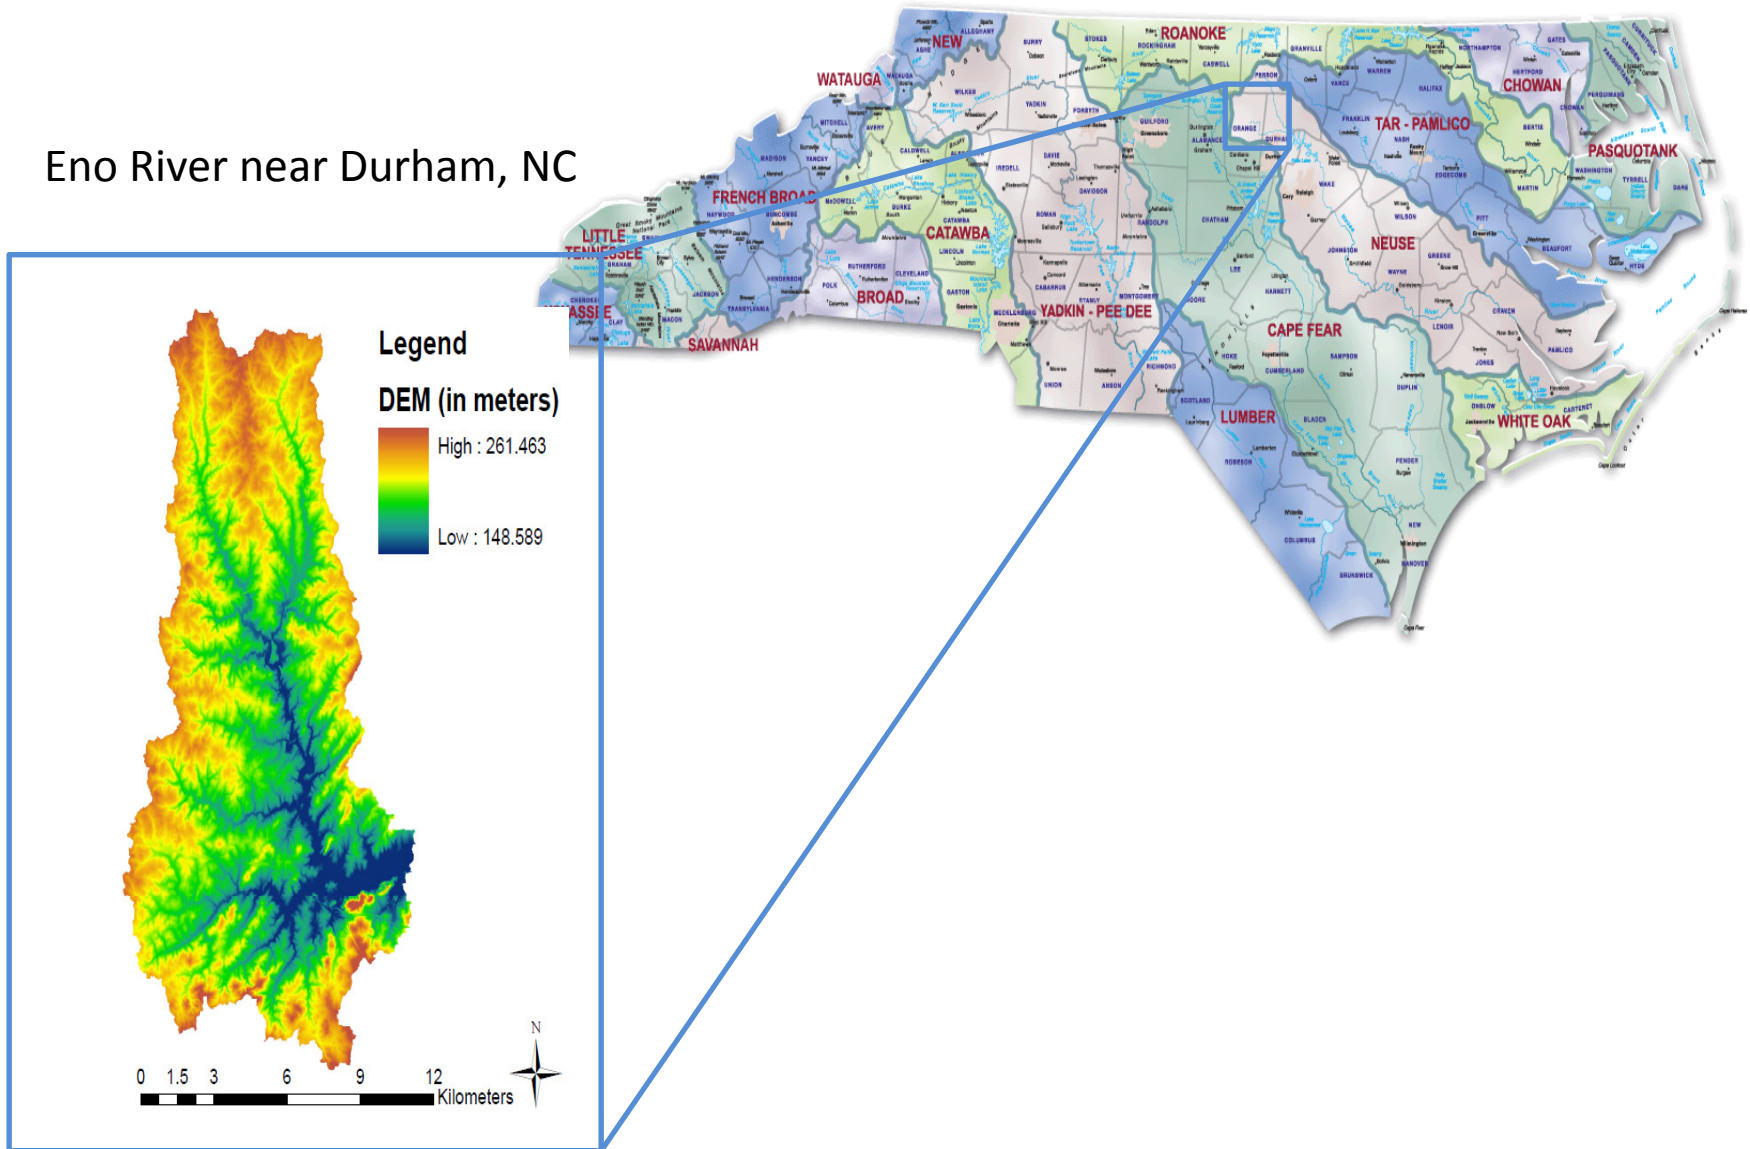


Chesapeake Bay Watershed



Typical Scale of Watershed Models

Eno River near Durham, NC




Scale

- Mississippi: 1,245,000 sq mi (3,220,000 km²)
- Chesapeake: 64,000 sq mi (166,000 km²)
- Eno (our case study watershed): 66 sq mi (171 km²)

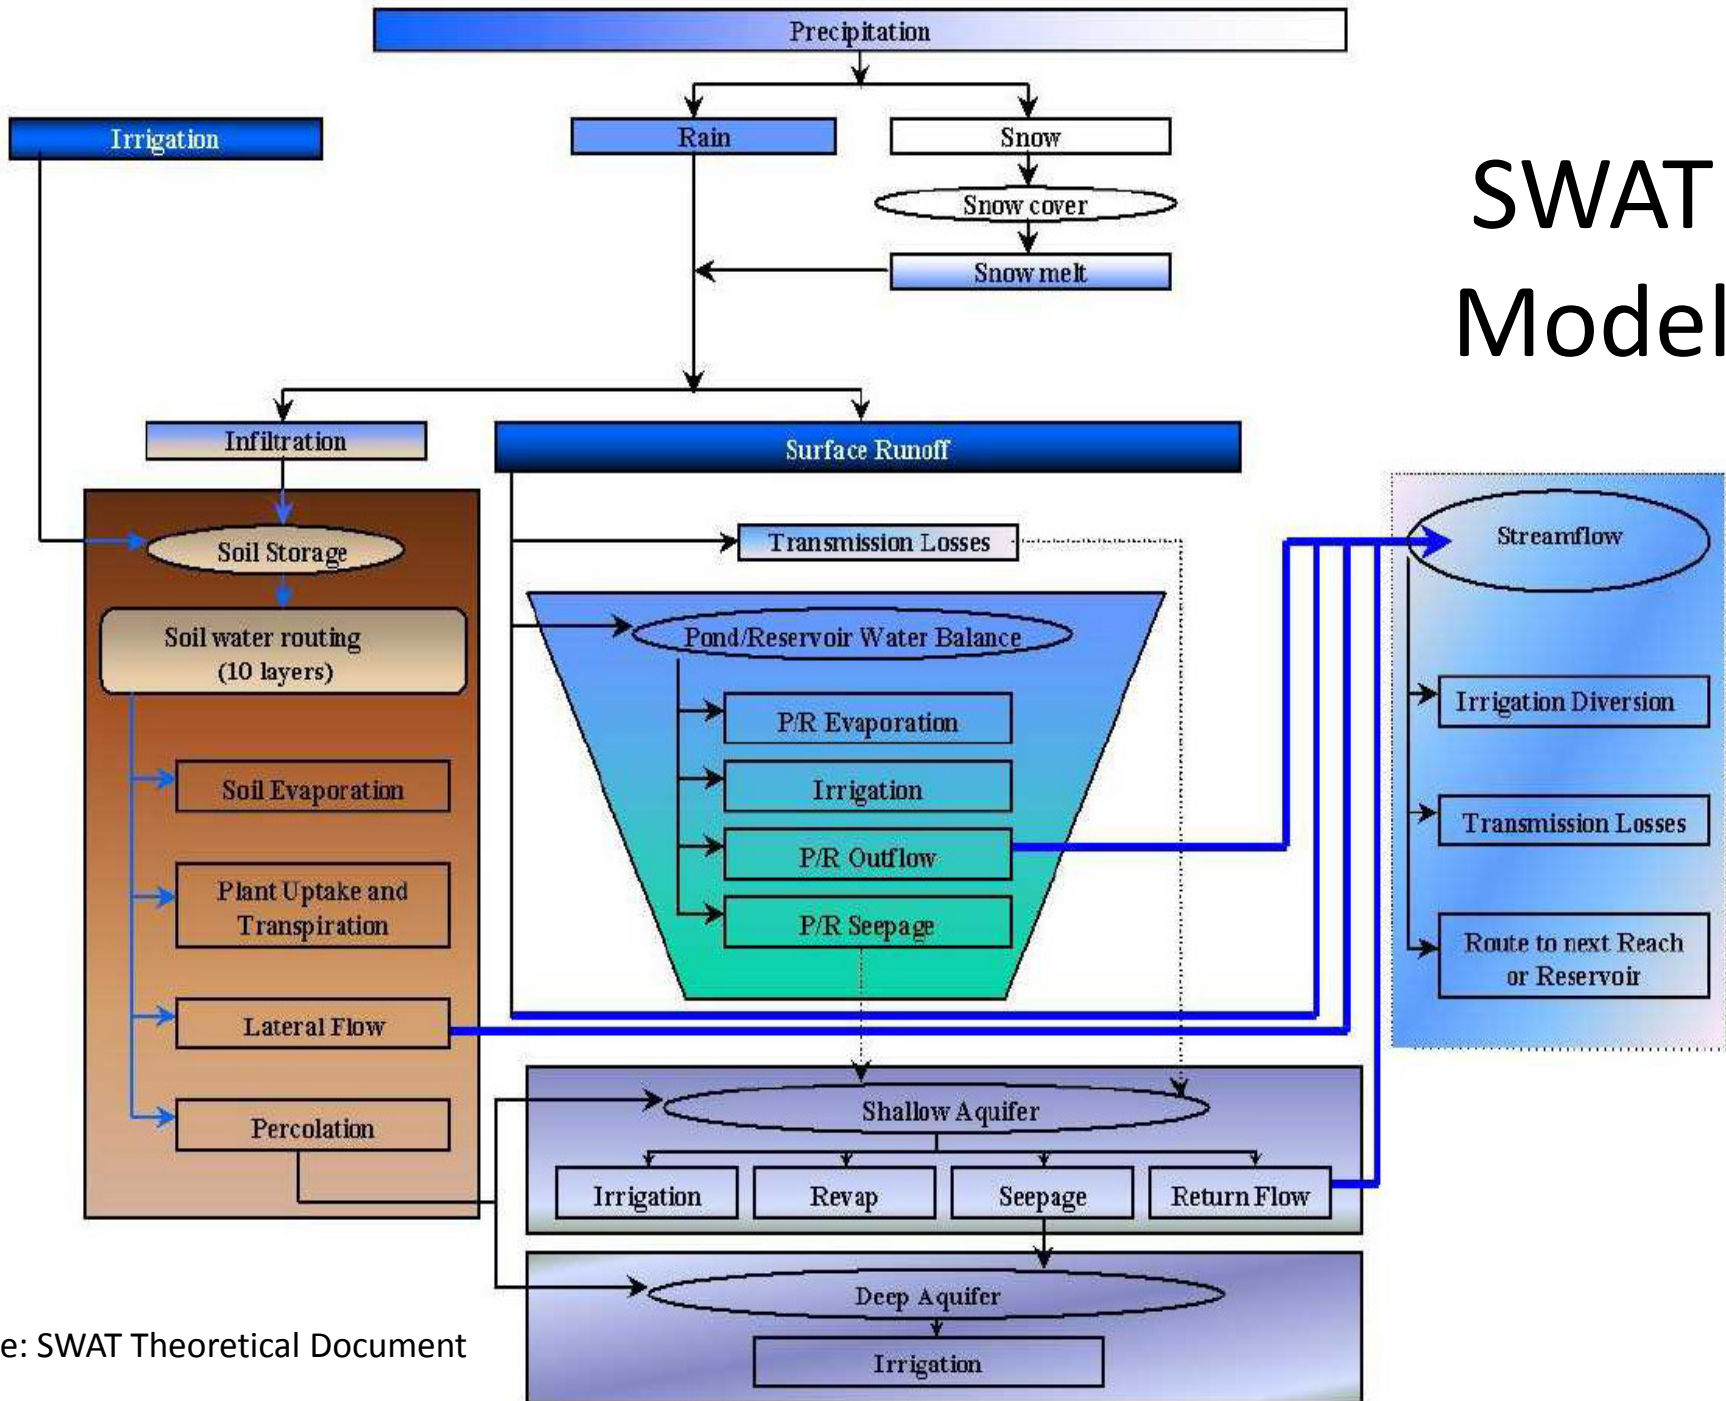
Eno to Chesapeake (~ 1,000 times)

Eno to Mississippi (~ 20,000 times)

Watershed Hydrology



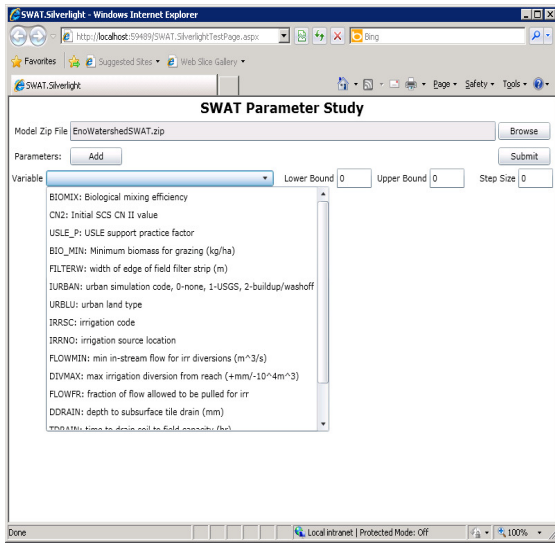
SWAT Model





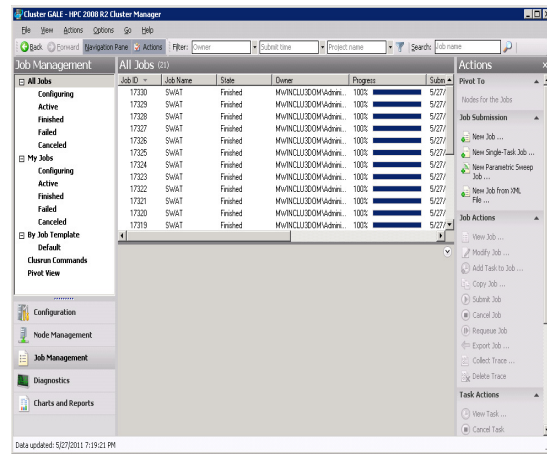

Source: SWAT Theoretical Document

Challenges in Watershed Modeling

- Data Preparation
 - Data exists, but files are large and require preprocessing
- Model Calibration
 - Requires running the model multiple times with varying parameters
- Scale up Model to Large System (Chesapeake, Mississippi)
 - Impractical using current approaches



HPC Cluster



Early Results

(Cloud Futures, June 2011)

	Stage-in	Compute	Total
Scientist laptop	0	55 sec	55 sec
Win2	5 sec	60 sec	65 sec
Azure (ex-large)	53 sec	32 sec	85 sec

Issues

(Cloud Futures, June 2011)

- Windows Azure only or cloudbursting?
- Data storage – where, how?
- Data sharing/reuse policy?
- Task granularity / coding?
- Task synchronization (e.g., MPI)?

Jim Gray, Microsoft Research,
San Francisco, California, July 2003

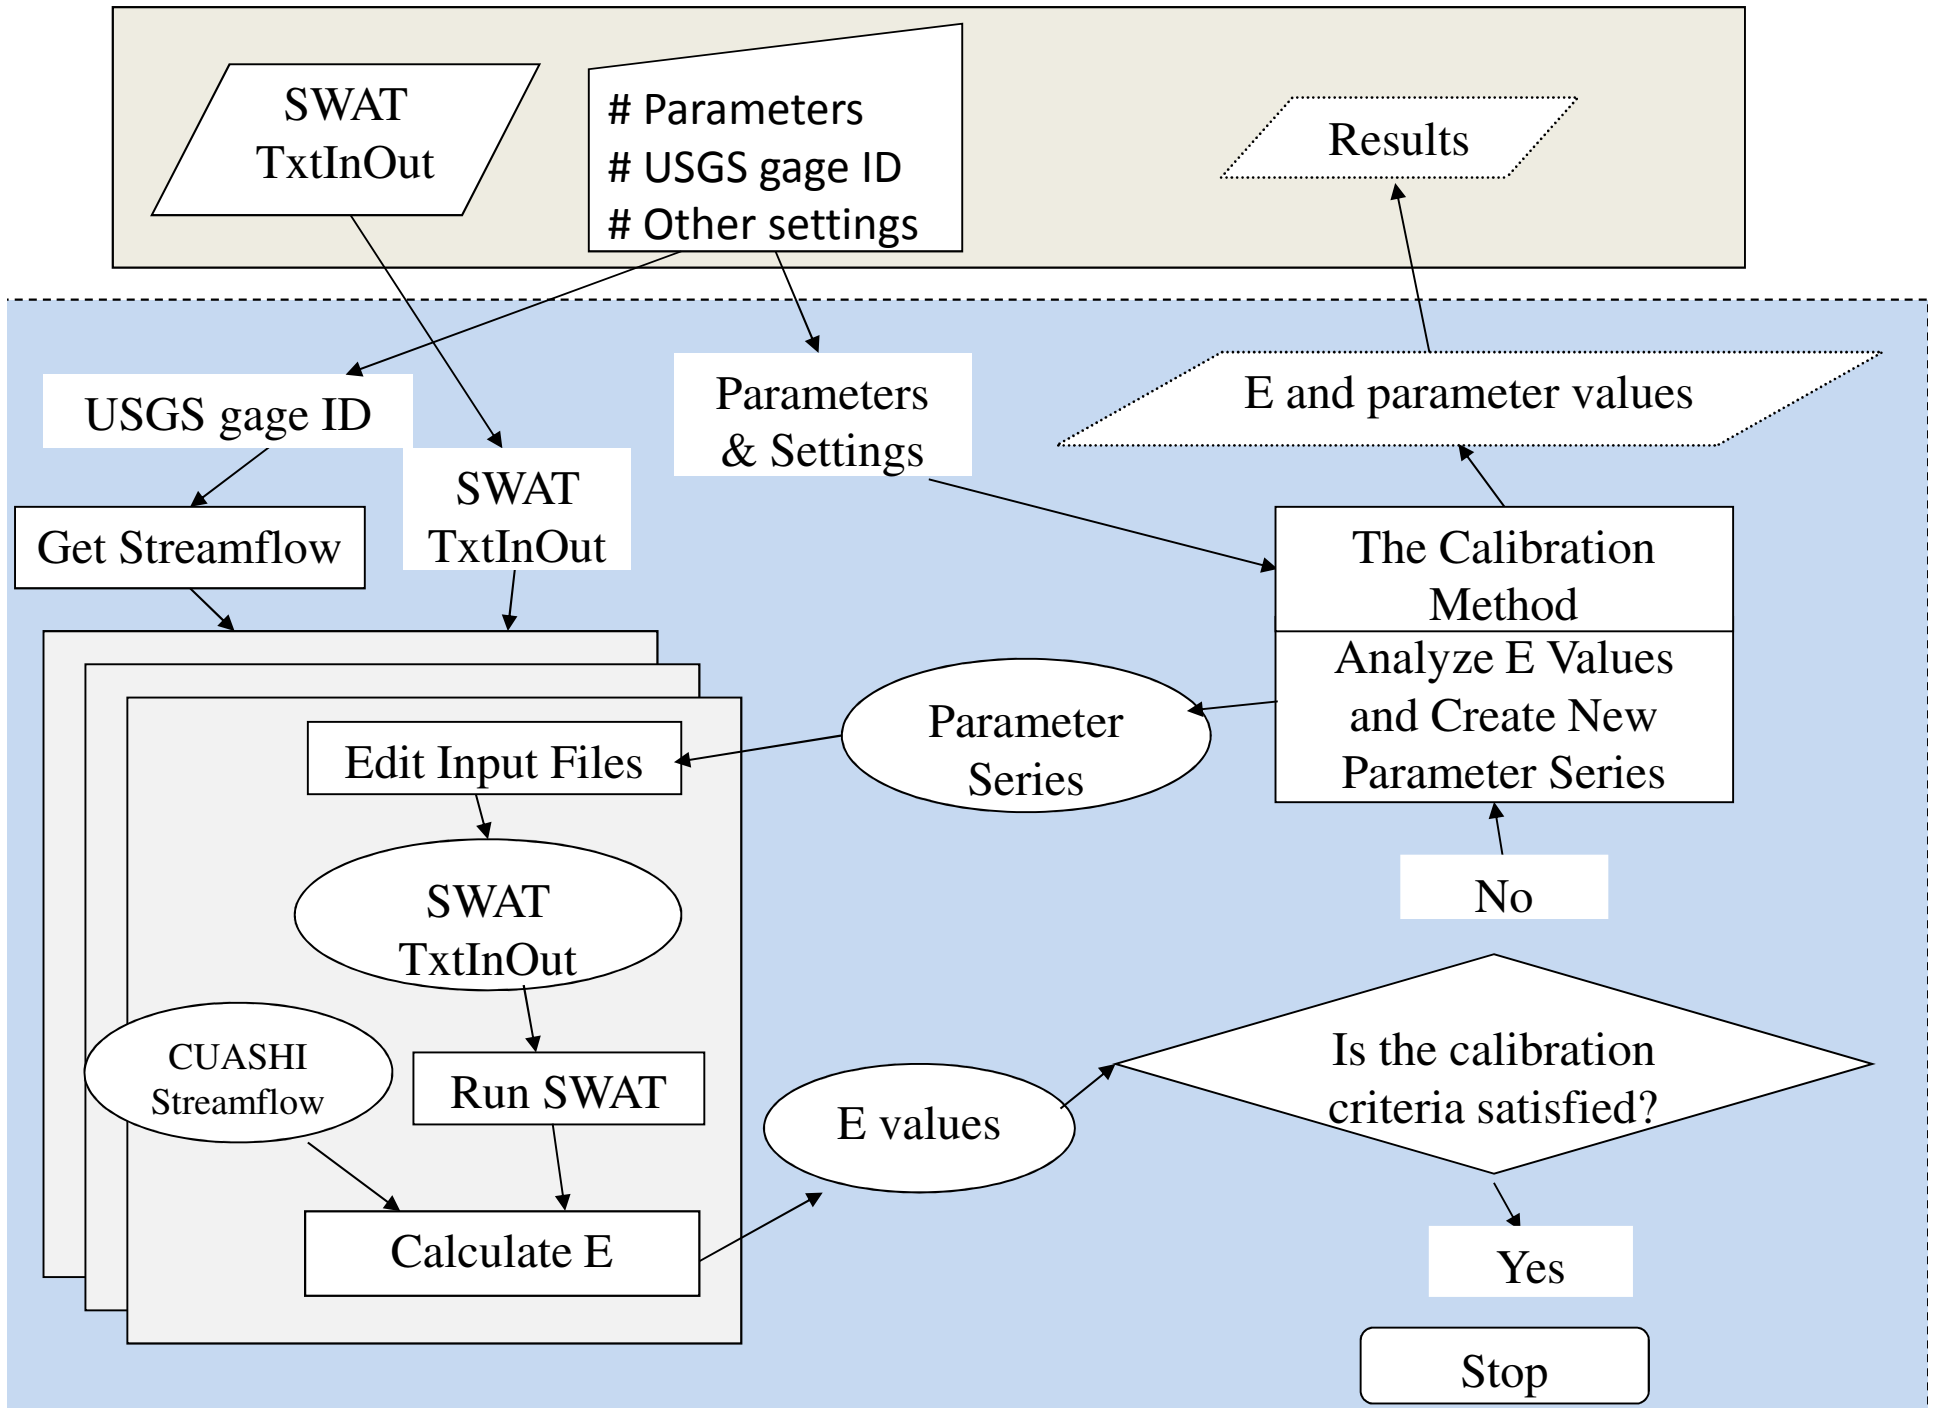
Distributed Computing Economics

JIM GRAY

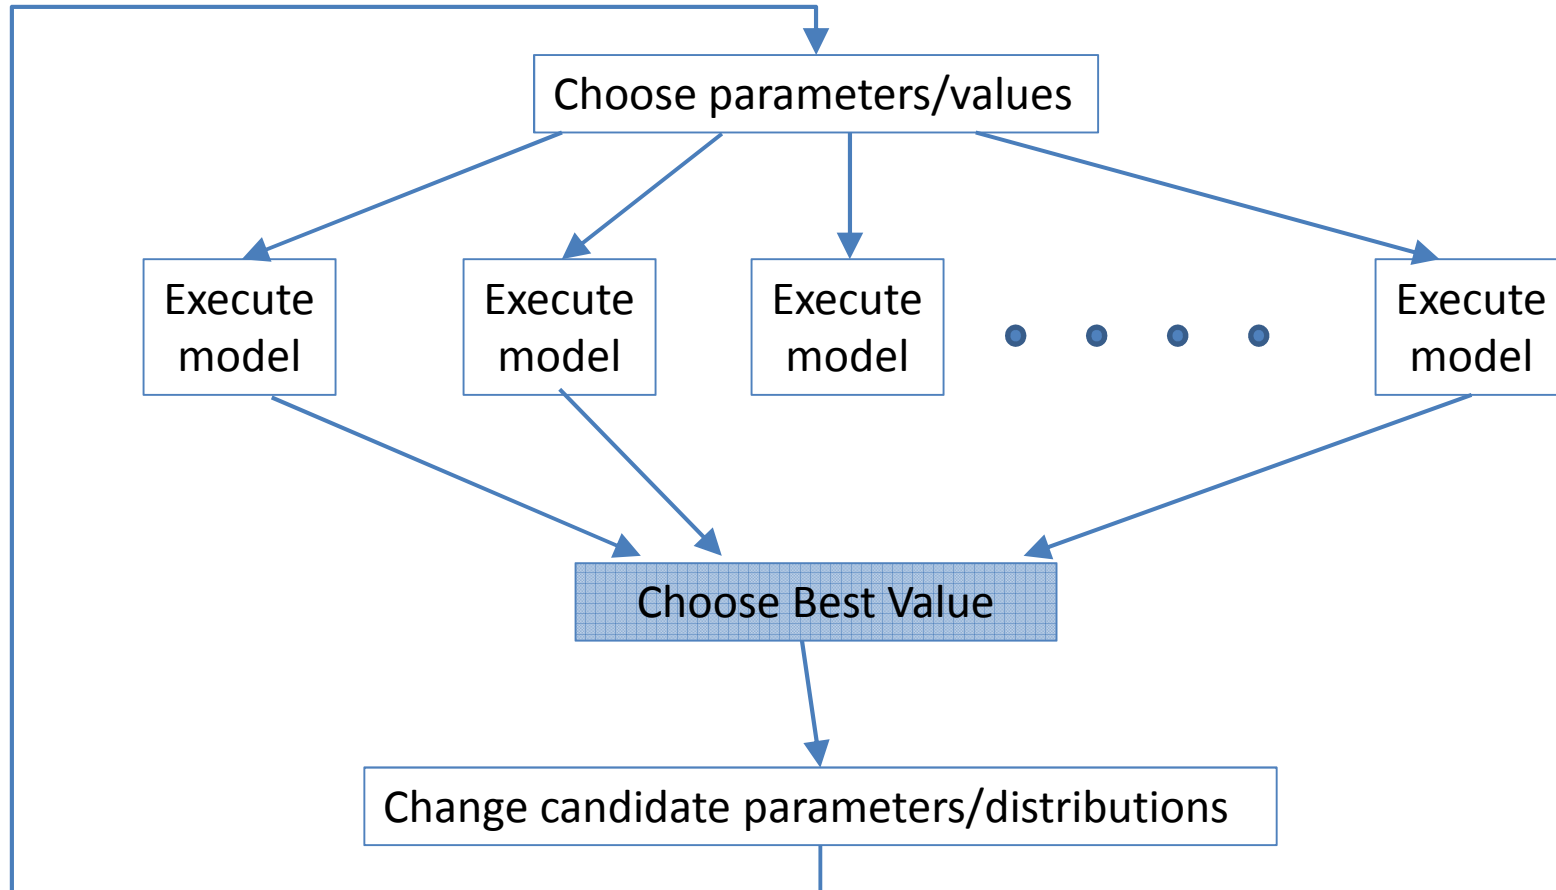
Computing economics are changing. Today there is rough price parity between: (1) one database access; (2) 10 bytes of network traffic; (3) 100,000 instructions; (4) 10 bytes of disk storage; and (5) a megabyte of disk bandwidth. This has implications for how one structures Internet-scale distributed computing: one puts computing as close to the data as possible in order to avoid expensive network traffic.

THE COST OF COMPUTING

Computing is free. The world's most powerful computer is free (SETI@Home is a 54-tera-flop machine).¹ Google freely provides a trillion searches per year to the world's largest online database (two petabytes). Hotmail freely carries a trillion e-mail messages per year. Amazon.com offers a free book-search tool. Many sites offer free news and other free content. Movies, sports events, concerts, and entertainment are freely available via television.



P-DDS



SWAT Calibration

Submit Jobs Job Listing Job Status Job Status Text Search Plot

Model Zip File

NW-DDS GA

Fixed Iterations Best E Threshold

Total number of iterations Concurrency there are currently 16 idle cores

Verify Simulation flow values against USGS Data

Parameters:

Variable Lower Bound Upper Bound Modification Mode

Variable Lower Bound Upper Bound Modification Mode

- ALPHA_BF
- BIOMIX
- blai
- CANMX
- Ch_Cov
- Ch_Erod
- CH_K2
- CH_N2
- CN2
- EPCO
- ESCO
- GW_DELAY
- GW_REVAP
- GWQMN
- Nperco
- Phoskd
- Pperco
- Rchrg_Dp
- REVADMN

SWAT Calibration

Submit Jobs Job Listing Job Status Job Status Text Search Plot

Load/Refresh

Job ID	Status	Input File	Parameters Calibrated	Best Daily E
69	Done	SWATnoCal_2012_2_21_17_2_49.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0
70	Done	AutoCal_test_2012_2_22_11_17_46.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.950795400
71	Done	AutoCal_test_2012_2_22_11_44_40.zip	CANMX, CN2, Ave. AW Incl. Rock Frag, ESCO	0.812586847
72	Failed	SWATnoCal_2012_2_22_13_13_29.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.751560458
73	Done	SWATnoCal_2012_2_23_10_55_57.zip	ALPHA_BF	0.592561975
74	Done	SWATnoCal_2012_2_23_11_4_28.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.753263733
75	Done	SWATnoCal_2012_2_23_12_13_16.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.785117606
76	Done	SWATnoCal_2012_2_24_10_7_34.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.752704232
77	Done	SWATnoCal_2012_2_24_11_53_5.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.761718524
78	Done	SWATnoCal_2012_2_24_14_17_52.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.749799309
79	Done	SWATnoCal_2012_2_24_14_54_10.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.785899864
80	Done	SWATnoCal_2012_2_24_16_55_55.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.785873733
81	Done	SWATnoCal_2012_2_27_11_7_31.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.718144797
82	Failed	SWATnoCal_2012_2_28_14_3_56.zip	ALPHA_BF, REVAPMN, GWQMN, CANMX, CN2, Sol_Z, Ave. AW Incl. Rock Frag, ESCO, SURLAG, CH_K2	0.467343344
84	Failed	SWATnoCal_2012_2_27_11_7_31_2012_3_2_19_24_38.zip	BIOMIX	0
85	Failed	SWATnoCal_2012_2_27_11_7_31_2012_3_3_6_4_3.zip	BIOMIX	0
86	Failed	SWATnoCal_2012_2_27_11_7_31_2012_3_3_6_7_29.zip	BIOMIX	0
87	Done	SWATnoCal_2012_2_27_11_7_31_2012_3_3_6_27_57.zip	BIOMIX	0.578020512
88	Done	SWATnoCal_2012_2_27_11_7_31_2012_3_3_7_9_21.zip	BIOMIX	0.577775339
90	Done	SWATnoCal_2012_2_27_11_7_31_2012_3_3_7_18_8.zip	BIOMIX	0.578035445
91	Done	SWATnoCal_2012_2_27_11_7_31_2012_3_3_8_16_13.zip	BIOMIX	0.108922727
92	Done	SWATnoCal_2012_2_27_11_7_31_2012_3_3_8_31_15.zip	CANMX, Ave. AW Incl. Rock Frag, BIOMIX	0.644082067
93	Done	SWATnoCal_2012_2_27_11_7_31_2012_3_3_9_25_38.zip	CANMX, Ave. AW Incl. Rock Frag, BIOMIX	0.643990437

Done

Internet | Protected Mode: On

110%

SWAT Calibration

Submit Jobs Job Listing Job Status Job Status Text Search Plot

Job ID Refresh

Current best E (daily) Current best parameters

Best run output file: 7 of 1008 iterations complete. Estimated completion in 01:52:11.0099857.

Job ID	Task ID	Daily Verification Value	Monthly Verification Value	Verification Result
215	1	-1.79769313486232E+308	-1.79769313486232E+308	
215	2	0.435030079124287	0.430177564846467	[Subbasin = 9, Monthly E = 0.430177564846467, Monthly R^2 = 0.492689608758142, Daily E = 0.435030079124287
215	3	-1.79769313486232E+308	-1.79769313486232E+308	
215	4	0.448744819957035	0.440016242695082	[Subbasin = 9, Monthly E = 0.440016242695082, Monthly R^2 = 0.496065138442328, Daily E = 0.448744819957035
215	5	0.51956086430255	0.43099698157476	[Subbasin = 9, Monthly E = 0.43099698157476, Monthly R^2 = 0.503963742286991, Daily E = 0.51956086430255, D
215	6	-1.79769313486232E+308	-1.79769313486232E+308	
215	7	-1.79769313486232E+308	-1.79769313486232E+308	
215	8	0.515377393104625	0.415546753302534	[Subbasin = 9, Monthly E = 0.415546753302534, Monthly R^2 = 0.500532711453686, Daily E = 0.515377393104625
215	9	0.518694240600354	0.426704982246484	[Subbasin = 9, Monthly E = 0.426704982246484, Monthly R^2 = 0.500508918265719, Daily E = 0.518694240600354
215	10	-1.79769313486232E+308	-1.79769313486232E+308	
215	11	-1.79769313486232E+308	-1.79769313486232E+308	
215	12	-1.79769313486232E+308	-1.79769313486232E+308	
215	13	-1.79769313486232E+308	-1.79769313486232E+308	
215	14	-1.79769313486232E+308	-1.79769313486232E+308	
215	15	0.508337718856126	0.431533626688847	[Subbasin = 9, Monthly E = 0.431533626688847, Monthly R^2 = 0.505526188610535, Daily E = 0.508337718856126
215	16	-1.79769313486232E+308	-1.79769313486232E+308	
215	17	-1.79769313486232E+308	-1.79769313486232E+308	
215	18	0.52094989783067	-1.79769313486232E+308	[Subbasin = 9, Monthly E = 0.420503562403376, Monthly R^2 = 0.507040772193582, Daily E = 0.52094989783067,
215	19	-1.79769313486232E+308	-1.79769313486232E+308	
215	20	-1.79769313486232E+308	-1.79769313486232E+308	
215	21	-1.79769313486232E+308	-1.79769313486232E+308	

SWAT Calibration


Submit Jobs | Job Listing | Job Status | Job Status Text | Search Plot

Job ID Refresh

Current best E (daily) Current best parameters

Best run output file: 174 of 1008 iterations complete. Estimated completion in 00:39:34.7430258.

Job ID	Task ID	Daily Verification Value	Monthly Verification Value	Verification Result
215	1	0.516288540295505	0.40803986589935	[Subbasin = 9, Monthly E = 0.40803986589935, Monthly R ² = 0.504028669720751, Daily E = 0.516288540295505,
215	2	0.435030079124287	0.430177564846467	[Subbasin = 9, Monthly E = 0.430177564846467, Monthly R ² = 0.492689608758142, Daily E = 0.435030079124287
215	3	0.501021560831151	0.445400065711808	[Subbasin = 9, Monthly E = 0.445400065711808, Monthly R ² = 0.500649426165851, Daily E = 0.501021560831151
215	4	0.448744819957035	0.440016242695082	[Subbasin = 9, Monthly E = 0.440016242695082, Monthly R ² = 0.496065138442328, Daily E = 0.448744819957035
215	5	0.51956086430255	0.43099698157476	[Subbasin = 9, Monthly E = 0.43099698157476, Monthly R ² = 0.503963742286991, Daily E = 0.51956086430255, D
215	6	0.514446933166626	0.420894260411255	[Subbasin = 9, Monthly E = 0.420894260411255, Monthly R ² = 0.502975755573573, Daily E = 0.514446933166626
215	7	0.452667041875028	0.45123090796741	[Subbasin = 9, Monthly E = 0.45123090796741, Monthly R ² = 0.503218946219286, Daily E = 0.452667041875028,
215	8	0.515377393104625	0.415546753302534	[Subbasin = 9, Monthly E = 0.415546753302534, Monthly R ² = 0.500532711453686, Daily E = 0.515377393104625
215	9	0.518694240600354	0.426704982246484	[Subbasin = 9, Monthly E = 0.426704982246484, Monthly R ² = 0.500508918265719, Daily E = 0.518694240600354
215	10	0.493568916566423	0.456231549622343	[Subbasin = 9, Monthly E = 0.456231549622343, Monthly R ² = 0.504628920854873, Daily E = 0.493568916566423
215	11	0.416955966312226	0.424203622164639	[Subbasin = 9, Monthly E = 0.424203622164639, Monthly R ² = 0.492996863333701, Daily E = 0.416955966312226
215	12	0.445690509060064	0.449105269715507	[Subbasin = 9, Monthly E = 0.449105269715507, Monthly R ² = 0.500259193622934, Daily E = 0.445690509060064
215	13	0.49355397331126	0.447015525582635	[Subbasin = 9, Monthly E = 0.447015525582635, Monthly R ² = 0.500406035407888, Daily E = 0.49355397331126,
215	14	0.50128257724122	0.422206967197406	[Subbasin = 9, Monthly E = 0.422206967197406, Monthly R ² = 0.4969032734016, Daily E = 0.50128257724122, Da
215	15	0.508337718856126	0.431533626688847	[Subbasin = 9, Monthly E = 0.431533626688847, Monthly R ² = 0.505526188610535, Daily E = 0.508337718856126
215	16	0.471129147945864	0.442785011019014	[Subbasin = 9, Monthly E = 0.442785011019014, Monthly R ² = 0.497083230851034, Daily E = 0.471129147945864
215	17	0.496680554462471	0.443531708959966	[Subbasin = 9, Monthly E = 0.443531708959966, Monthly R ² = 0.499525215170893, Daily E = 0.496680554462471
215	18	0.52094989783067	0.420503562403376	[Subbasin = 9, Monthly E = 0.420503562403376, Monthly R ² = 0.507040772193582, Daily E = 0.52094989783067,
215	19	0.512845721168434	0.441613896012705	[Subbasin = 9, Monthly E = 0.441613896012705, Monthly R ² = 0.503041894175349, Daily E = 0.512845721168434
215	20	0.513870788314959	0.405570792007069	[Subbasin = 9, Monthly E = 0.405570792007069, Monthly R ² = 0.496075647095762, Daily E = 0.513870788314959
215	21	0.46108884434218	0.434403512317391	[Subbasin = 9, Monthly E = 0.434403512317391, Monthly R ² = 0.493896842020705, Daily E = 0.46108884434218,




From: mehmet ercan <mehmetbercan@gmail.com>
To: GOODALL, JONATHAN L
Cc: Norm Beekwilder; Marty Humphrey
Subject: Re: update on the cloud version

Sent: Mon 10/17/2011 12:22 PM

I have tested new CN2 parameter on the cloud. 1008 times SWAT run (60 days time step) took 1 hour 54 minutes on the cloud (job: 136) and about 700 SWAT run (60 days time step) took about 10 minutes for automatic calibration implemented in SWAT. Best E value was 0.86 for the cloud and 0.89 for the SWAT automatic calibration. I can't see best CN2 result in percentage on cloud. Best CN2 value was + 9.99 % for SWAT automatic calibration.

Mehmet.

On Fri, Oct 14, 2011 at 4:43 PM, GOODALL, JONATHAN L
<GOODALL@cec.sc.edu> wrote:

 mehmet ercan



1 hour 54 minutes on the cloud

I have been comparing the results of the automatic calibration on the cloud with the results of the automatic calibration implemented in SW AT. The SW AT was run for 20 days (the days were about 20 minutes for automatic calibration implemented in SW AT). The R² value was 0.80 for the cloud and 0.80 for the SW AT automatic calibration. I can't see how the CTI2 result in percentage on cloud. The CTI2 value was + 0.00% for the SW AT automatic calibration.

Subject

On Fri, Oct 14, 2011 at 4:43 PM, GUYONNEAU, RICHARD L.
(<mailto:RGUYONNEAU@comcast.net>) wrote:

1 hour 54 minutes on the cloud

10 minutes for automatic
calibration implemented in SWAT.

Re: update on the cloud version - Message (HTML)

File Message

Ignore X Delete Reply Reply All Forward Meeting IM More Quick Steps Move Rules OneNote Actions Mark Unread Categorize Follow Up Translate Find Related Select Zoom

From: mehmet ercan <mehmetbercan@gmail.com> Sent: Fri 10/21/2011 11:13 AM
To: Norm Beekwilder
Cc: Marty Humphrey; GOODALL, JONATHAN L
Subject: Re: update on the cloud version

Dr. Beekwilder,

I understand. Then the total run time on the cloud for the model is 2 hours 44 minutes versus about 10 hours on our computer (ce-congaree).

I know that the model now changes parameters on Hydrologic Respond Unit (**HRU**) level. We can run our model with this way or the way it changes with same range for all HRUs. The problem is that I am not able to find out what are the ways to change the remaining of flow parameters with the way going into all individual files. However, I have found out a way that is implemented in SWAT which let us

mehmet ercan

2 hours 44 minutes

Dr. Bookwala,

I understand. Thus the total run time on the cloud for the model is 2 hours 44 minutes versus about 10 hours on our computer (two computers).


I know that the model uses changes parameters on Hydrologic Response Unit (HRU) level. We can run our model with this way or the way it changes with some range for all HRU's. The problem is that I am not able to find out what are the ways to change the remaining of these parameters with the way going into all individual files. However, I have found out a way that is implemented as SW AT which let us


A screenshot of a software interface, likely a simulation or modeling tool. The interface has a top toolbar with various icons and a main workspace. A white text box is overlaid on the workspace, containing the text "2 hours 44 minutes".

2 hours 44 minutes

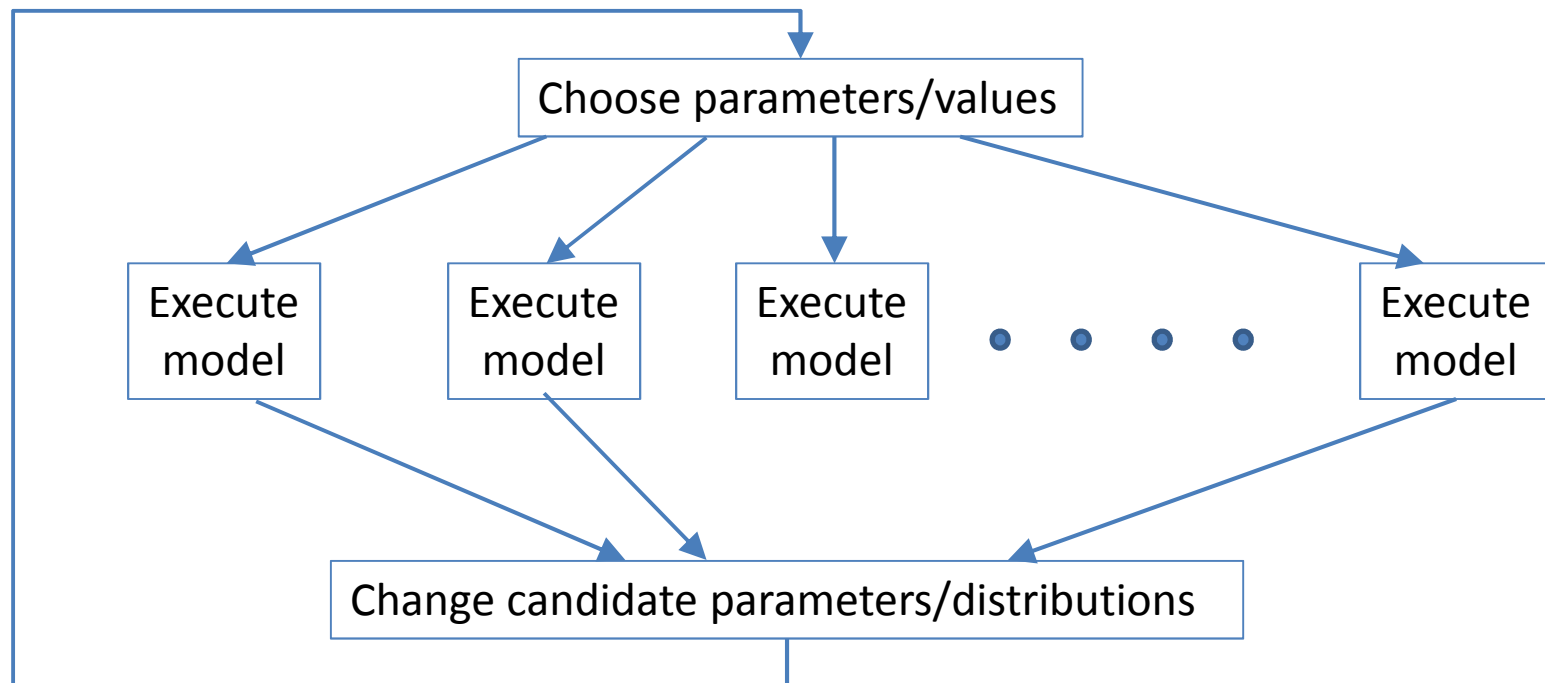
10 hours on our computer

P-DDS








NW-DDS




NW-DDS









Efficiency

	Duration	Speedup
Desktop, SWAT internal calibration (single threaded)	11.4 hours (41040 sec)	--
Windows Azure, Ex-large VM, 16 cores, NW-DDS	43.32 min (2599 sec)	15.78x
Windows Azure, Ex-large VM, 64 cores, NW-DDS	11.76 min (706 sec)	58.13x
Windows Azure, Ex-large VM, 256 cores, NW-DDS	5.03 min (302 sec)	135.89x

As number of cores increase, dominated by [a] stragglers, and [b] re-running best param set to get output files

Re-visiting the Issues (Cloud Futures, June 2011)

- Windows Azure only or cloudbursting?
 - *Only Windows Azure (with head node inside enterprise)*
- Data storage – where, how?
 - *Only Windows Azure, so [a] “where” is largely non-issue, and [b] “how” is as little as possible*
- Data sharing/reuse policy?
 - *Not as much as an issue as we would like*
- Task granularity / coding?
 - *Hmm... not exactly an issue (see next slide)*
- Task synchronization (e.g., MPI)?
 - *Removed in an uninteresting way*

Next steps: Calibration

- Additional watershed models (e.g., HSPF)
- Additional calibration algorithms
- Better control of search (e.g., stopping criteria)
 - Visual results presentation and visual steering
 - Trading off fast-vs.-more-exhaustive
- Better sharing / non-sharing
 - Insta-virtual-cluster a la HadoopOnAzure
- *How does insta-calibration change the watershed model design process? (continual calibration?)*

Have something unsolvable? It's not a problem!

Request a new cluster

DNS name

DNS name

http://cloudapp.net

Cluster size

- Small**
4 nodes
2 TB disk space
Available
- Medium**
8 nodes
4 TB disk space
Available
- Large**
16 nodes
8 TB disk space
Available
- Extra large**
32 nodes
16 TB disk space
Available

Cluster login

Username

Password

Confirm password

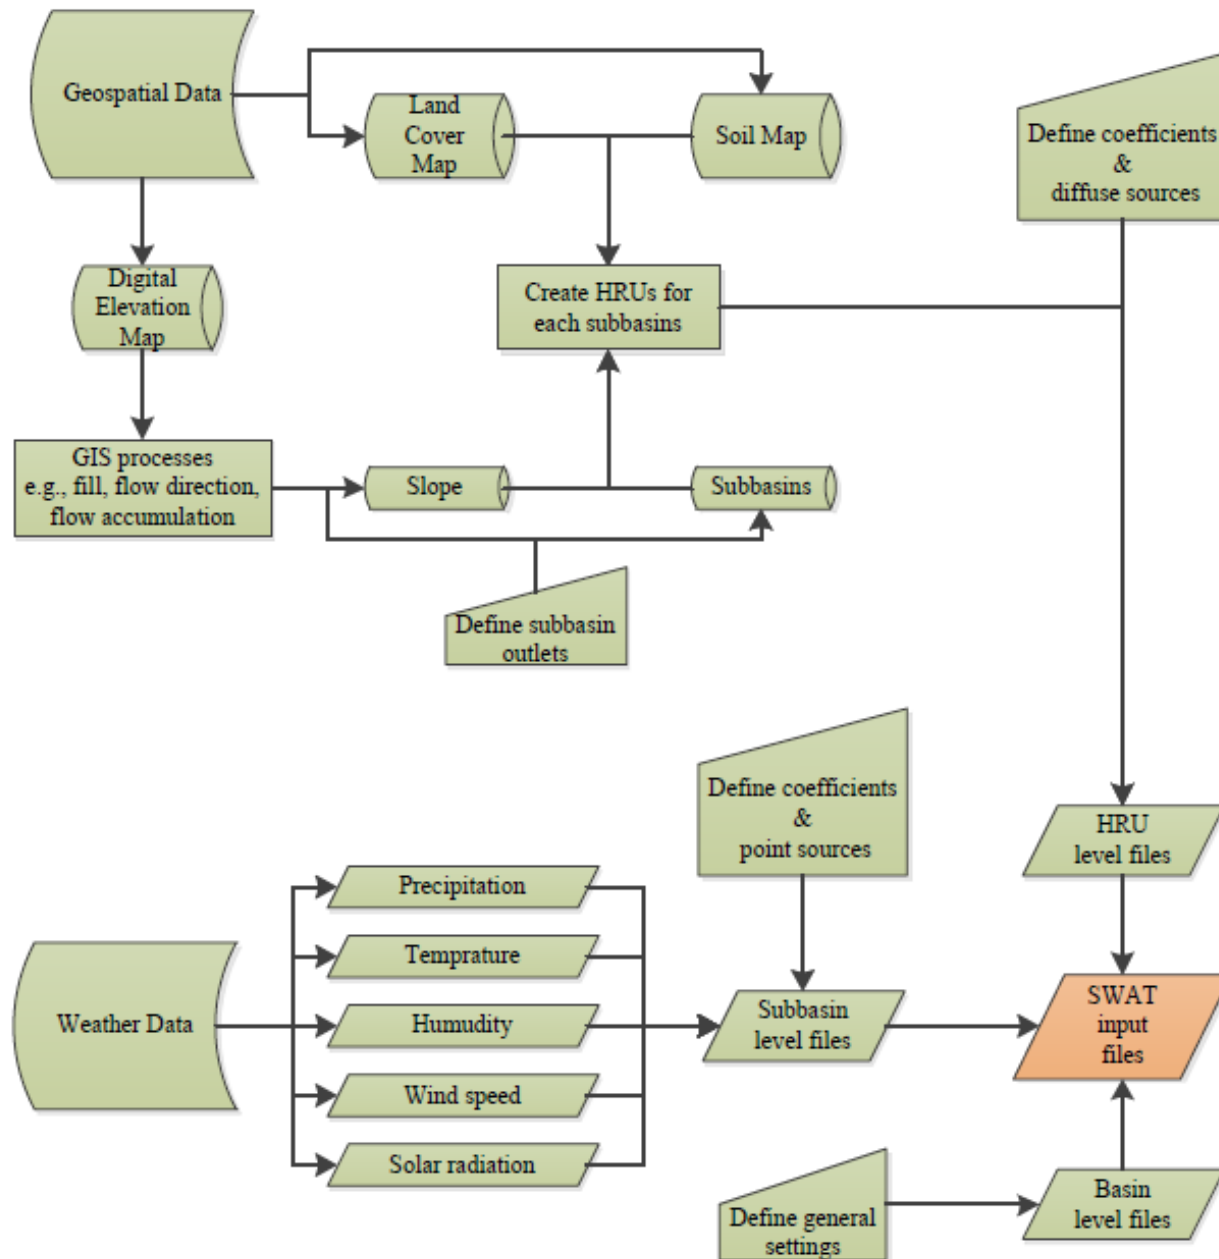
Valid characters: A-Z, 0-9, underscore (_), hyphen (-)

Password must be between 7 and 14 characters, contain both upper- and lowercase letters, at least one number, and no symbols.

Please ensure your passwords match.

SQL Azure is optional

Next steps: Data preparation



Observations on Windows Azure

- Does Windows Azure (more broadly: cloud computing) substantially change the challenges of collaboration between domains? No.
- What do we like?
 - Cloudbursting mechanism (head node, VPN).
 - Predictability is sufficient (and continues to improve).
- What would we like to see?
 - Faster boot time for VMs.
 - Non-round-up cost-model.
 - Faster speed of light.

Summary

- Goals: Watershed calibration, data prep, and large-scale modeling
- Good progress on watershed calibration, data prep next
- Basic issues of collaboration across domains will continue to dominate