Implementing a Loosely-Coupled Integrated Assessment Model in the Pegasus Workflow Management System

Patrick J Clemins
patrick.clemins@uvm.edu
University of Vermont
Burlington, Vermont, USA

Morgan Rodgers
morgan.rogers@preventiongenetics.com
Prevention Genetics
Marshfield, Wisconsin, USA

Scott Turnbull
scott.turnbull@uvm.edu
University of Vermont
Burlington, Vermont, USA

Asim Zia
asim.zia@uvm.edu
University of Vermont
Burlington, Vermont, USA

ABSTRACT

Integrated assessment models (IAMs) are commonly used to explore the interactions between different modeled components of socio-environmental systems (SES). Most IAMs are built in a tightly-coupled framework so that the complex interactions between the models can be efficiently implemented within the framework in a straightforward manner. However, tightly-coupled frameworks make it more difficult to change individual models within the IAM because of the high level of integration between the models. Prioritizing flexibility over computational efficiency, the IAM presented here is built using a loosely-coupled framework and implemented in the Pegasus Workflow Management System. The modular nature of loosely-coupled systems allows each component model within the IAM to be easily exchanged for another component model from the same domain assuming each provides the same input / output interface. This flexibility allows researchers to experiment with different models for each SES component and facilitates smoother upgrades between each version of the independently developed component models.

CCS CONCEPTS

• **Applied computing** → **Environmental sciences**; • **Software and its engineering** → *Software system models*; **Data flow architectures**; **Abstraction, modeling and modularity**; **Software design tradeoffs**.

KEYWORDS

integrated assessment modeling, workflow management systems, loosely-coupled software systems, modular software systems

ACM Reference Format:

Patrick J Clemins, Scott Turnbull, Morgan Rodgers, and Asim Zia. 2020. Implementing a Loosely-Coupled Integrated Assessment Model in the Pegasus Workflow Management System. In *Practice and Experience in Advanced*

Research Computing (PEARC '20), July 26–30, 2020, Portland, OR, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3311790.3396638>

1 INTRODUCTION

Integrated assessment models (IAMs) are a modeling framework widely used to explore the interactions between different components of socio-environmental systems (SES) [10, 23, 29, 32]. They are particularly popular in climate change impacts studies in which climate models are linked to terrestrial process models such as hydrological or lake models to determine impacts of changes in climate on water resources and water quality. Social systems models such as land use land change and economic models are included in more sophisticated IAMs [35]. IAMs are often used to inform policy makers and create policy recommendations because of their ability to show the interactions between components of an SES over a suite of future scenarios with varying climate change mitigation or land use policy.

Most IAMs are implemented in tightly-coupled frameworks because of the sophisticated nature of the feedbacks between the various models within the IAM. A tightly-coupled framework is characterized by each component model executing simultaneously, often accessing the current SES state in system memory as shown in Figure 1. While each component model may not operate on the same time scale, they remain as an active process while the simulation time step increments at the interval required by the model with the smallest time step requirement. An in-memory control process is usually used to coordinate the execution of and communication between the component models. This approach allows the component models to access the current state of the

Figure 1: Tightly-Coupled Design

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PEARC '20, July 26–30, 2020, Portland, OR, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6689-2/20/07.

<https://doi.org/10.1145/3311790.3396638>

Figure 2: Loosely-Coupled Design

system at any arbitrary time step and allows software developers to jointly optimize the efficiency of the component models. While these tightly-coupled frameworks are efficient, their high level of integration between the models means that a large amount of effort is required to switch between different community models for the same domain, that is, to switch between different lake or hydrology models.

The use of a loosely-coupled architecture addresses this shortcoming. In a loosely-coupled architecture, each model runs independently, without access to the current state of the SES in system memory or each other as shown in Figure 2. Instead, the initial conditions of the SES required by each component model are provided by input files. The model then runs to completion for a given time frame. After execution, the model updates the SES system state by providing updated output files. Within each time frame, the models execute in a cascade, with the models that are most sensitive to the current SES state executing at the end of the cascade. The order of the cascade is an important design consideration and is discussed further in section 4.1. A workflow management system is often used to manage the sequence of execution of the models and start each subsequent cascade after the previous cascade finishes.

One important feature that emerges from a loosely-couple design is modularity. Modularity has been embraced in the software design [3, 24] and computational intelligence [1, 25] fields because it allows the system to break a complex problem into smaller, solvable subsystems and allows those subsystems to specialize in their task. While it is certainly possible to maintain modularity within a tightly-coupled design, it is not as critical to the design philosophy and thus, in practice, it is often given less priority than other design elements more intrinsic to the tightly-coupled approach such as computational efficiency and optimized memory usage. A significant practical benefit of modularity is that it allows researchers to switch between and experiment with different community models within the same domain assuming both models provide similar output and rely on similar input parameters. In addition, because the underlying code for each component model is completely independent, models can be written in a variety of programming languages and frameworks and be updated to new versions with relative ease. The two main drawbacks to the loosely-coupled approach are the lack of access to the current SES state by all component models at any arbitrary time step and the inability to jointly optimize computational efficiency between the models, although, each model can be optimized independently.

2 THE BREE IAM

The research team for the current Vermont National Science Foundation (NSF) Established Program to Stimulate Competitive Research (EPSCoR) Research Infrastructure Improvement (RII) Track-1 grant, titled “Basin Resilience to Extreme Events” (BREE), has built an IAM to evaluate the effects of future climate and policy on water quality in the Lake Champlain Basin. This BREE IAM is currently comprised of statistically downscaled climate projections [34], a hydrology model [22], a land use land change model [30], and a lake model [18, 19]. An economic model and governance model are slated to be added soon. Figure 3 shows the proposed architecture of the completed BREE IAM. Each box represents a different domain-specific component model and the arrows show current or planned feedback between the component models. Note that the majority of the interactions between the models are bidirectional.

The BREE IAM is an expansion and refinement from the IAM built for the previous NSF EPSCoR RII Track-1 grant [7]. The hydrology model has already been replaced with a newer version of the same community model, the Regional Hydro-Ecologic Simulation System (RHESSys), and expanded to include nutrient transport for phosphorus and nitrogen [26, 28]. In addition, completely new climate [16, 27], land use land change [2, 11, 12], and lake models [17, 21] will be exchanged for the previous models, all while preserving most of the core software infrastructure that was developed for the previous grant. The economic [13, 33] and governance models [4] will be new domain models in the IAM. The ability of the BREE IAM to quickly adapt to new and updated models is due to its loosely-coupled, modular architecture.

3 DATA INTERFACE BETWEEN MODELS

The loosely-coupled architecture of the BREE IAM precludes in-memory data exchange that is common in more tightly-coupled IAM architectures. Because the component models of the BREE IAM execute independently of each other, all data transfer between the models is file-based. The state of the Lake Champlain Basin SES at any given point in time is represented by the collection of output files from each of the component models. When a component model begins execution, it collects the state information it needs from the appropriate files, converts the data into its required input format, and then begins execution. Consequently, each model has “prep” stage which performs the data collection and conversion tasks required to prepare the input files for the model.

Figure 3: The BREE IAM

Figure 4: The Data Interface

To manage data flows within the IAM, an input / output (I/O) interface is defined for each type of model (i.e. climate, hydrology, etc.). Figure 4 shows this I/O interface for the hydrology model. By defining these interfaces for each modeling domain, much of the IAM that controls data flow can be designed and implemented without knowledge of the specific domain model that will be used. Small adjustments to the I/O interface for each specific domain model are often necessary, but the concept of inheritance to expand or adjust the base I/O interface has proven to be an efficient strategy for addressing these small adjustments. For instance, if the RHESys hydrology model was replaced with another hydrology model such as SWAT which does not operate on a 2D grid, but instead on a per basin basis, the methods that provide temperature, precipitation, and land use for SWAT could inherit from the 2D grid versions and then perform the aggregation necessary to provide those input parameters on basin scale instead of through a 2D grid.

4 IMPLEMENTATION IN PEGASUS

The BREE IAM is implemented in the Pegasus Workflow Management System, developed by the Information Sciences Institute (ISI) at the University of Southern California [9]. A Pegasus workflow is built by defining the input requirements and outputs of each task in the workflow. Pegasus then plans the workflow by constructing a directed acyclic graph of the workflows tasks and executes each workflow task as its input requirements become available. The execution environment for the workflow is flexible with support for HTCondor [5], commercial cloud providers, and a selection of national HPC resources.

4.1 Bidirectional Data Exchange in Directed Acyclic Graphs

Pegasus, along with most other workflow management systems, use directed acyclic graphs to plan and execute their workflows. However, as Figure 3 shows, the BREE IAM is designed with bidirectional data exchange between models. The linked, sequential cascade is designed to overcome this disparity between the one-directional nature of directed acyclic graphs and the two-directional design of the BREE IAM. The native direction of the cascade implements one direction of the BREE IAM data exchanges between component models, and for the less time sensitive data, the results of the previous cascade are used to implement the other direction of the data exchange. A set of initial conditions are used for the

SES system state for the first cascade. Therefore, the choice to use a workflow management system and their one-directional directed acyclic graphs to implement the IAM was a significant impetus to creating the cascading model for data exchange between IAM component models.

4.2 Execution Environment

The BREE IAM makes use of an heterogeneous computing environment consisting of local compute resources as well as NCAR's Cheyenne HPC cluster [8] to execute the Pegasus workflow. There are two main local compute resources, a traditional compute workstation with 32 CPU cores and 256GB of system memory and a NVIDIA DXG-1 with 40 CPU cores, 8 V100 GPUs providing 40,960 CUDA cores, and 512GB of system memory. The local compute resources are managed by HTCondor. A lightweight PBS submission server was created so that HTCondor could submit PBS batch jobs to Cheyenne using SSH for communication and job management and Globus [31] for file transfer. Generally, the hydrology and lake models are run on Cheyenne while many of the other models, due to dependencies and/or resource constraints, are run on local resources.

The component models and data prep tasks are written in several different programming languages at various versions including Fortran, R, Python, C, C++, and Java. For compiled languages, the workflow tasks are compiled for Linux, specifically for each platform when possible. However, the current lake model consists of two pre-compiled, proprietary Windows binaries and are run under Wine, a Windows capability layer for POSIX operating systems [20]. We are beginning a transition to the use of containers for the modeling tasks to standardize the execution for each model on each platform. The wide range of hardware and software platforms used in the BREE IAM highlights the flexibility of the loosely-coupled architecture.

5 DESIGN CONSIDERATIONS

5.1 Order of the Cascade

The order of execution of each component model within a time frame, or the cascade, is an important design consideration in a loosely-coupled IAM architecture. Unlike tightly-coupled IAMs where each component model has instant access to the system state, component models in a loosely-coupled IAM do not have this instant access because they all execute independently and often times, in sequence to each other instead of in parallel. The resulting impact of placing a model near the start of the cascade (upstream) is that the model only has access to outputs from the previous cascading time frame (see Figure 2) while models toward the end of the cascade (downstream) have access to data from models that were executed earlier in that cascade.

A good guideline for determining the order of the cascade is to place models that simulate boundary conditions and external drivers near the start of the cascade and to place models that largely respond to those boundary conditions and external drivers near the end of the cascade. Another consideration is the time scale of the process that the component model is simulating. Models that operate on longer time scales can often be placed near the start with minimal impact on their results (i.e. climate, land use land change),

Figure 5: The BREE Cascade

while models that operate on shorter time scales (i.e. hydrology, lake) should be placed near the end of the cascade so that they have access to more recent data from upstream models. Figure 5 shows the current cascade used for the BREE IAM, designed following these guidelines. Note that there are component models shown in Figure 3, the BREE IAM, that are not yet implemented in the IAM.

5.2 Determining the Cascade Time Frame

The time frame of each cascade determines the maximum age of any input data within the IAM. Equivocally, it determines the longest delay in feedbacks between the component models. There are trade offs in choosing a longer or shorter cascading time frame. The longer the time frame, the longer each component model can run within the cascade, reducing the number of times each model must be restarted and spun up or recover from a save point. Longer time frames also mean fewer cascades are required to span the overall simulation time period. The shorter the time frame, the shorter the delay in feedback from the downstream to the upstream models in the cascade.

The BREE IAM is currently configured to use a 10-year time step for each cascade. This time step was chosen by analyzing the feedbacks from downstream models to upstream models and determining which feedback was the most time-critical. In the case of the BREE IAM, this was the feedback from the downstream hydrology model to the upstream land use land change model. Considering the trade offs between computational simplicity and the desired maximum delay in hydrology inputs to the land use land cover change model, a 10-year time step was chosen. The ability to quickly modify the cascading time step to determine its effects on the IAM output is a current work in progress.

6 DISCUSSION

The BREE IAM and its previous iterations has been used to generate both a wide array of research results [14, 15, 35] and provide insight to Vermont EPSCoR's Policy and Technical Advisory Committee (PTAC) consisting of policy makers and other stakeholders across the state through a series of mediated modeling workshops [6]. In addition to its broader impacts, the BREE IAM's combination of technologies, implementation techniques, and design choices are quite novel in the SES and IAM communities and can be applied singularly or in concert with each other to a broad range of systems modeling projects. The loosely-coupled, modular design of the BREE IAM is a primary reason the IAM can be quickly reconfigured for a wide array of simulation experiments that explore diverse research and policy questions.

This same design choice also allows researchers to iterate on the component models quickly. Features and complexity can be added to individual component models without concern that they will interfere with other component models. Furthermore, the updated models can be independently calibrated and tested outside of the IAM using techniques standard in each domain. These improvements to the component models can then be contributed back to the community, resulting in the BREE IAM not only contributing to research on extreme events in the Lake Champlain Basin, but also to the research domains that make up the BREE IAM.

Finally, as design and implementation choices for the BREE IAM are explored and evaluated, it serves as experimental platform for complex software systems design, cyberinfrastructure implementation innovations, and advances in computational intelligence through novel approaches to integrating artificial intelligence and machine learning within the BREE IAM. Thus, the BREE IAM has emerged as the focal point of the current Vermont EPSCoR RII Track-1 research project because of its flexibility to explore a wide range of research questions and the interactions between the sub-systems under study within the Lake Champlain Basin SES.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under grants EPS-1101317 and OIA-1556770. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] Mohammed Amer and Tomás Maul. 2019. A review of modularization techniques in artificial neural networks. *Artificial Intelligence Review* 52, 1 (jun 2019), 527–561. <https://doi.org/10.1007/s10462-019-09706-7>
- [2] Kevin Andrew. 2019. Modeling the Cooperative and Adversarial Behaviors of Farmer and Regulator Agents in Vermont's Missisquoi Bay Area. In *Second Northeast Regional Conference on Complex Systems (NERCCS)*. Binghamton, NY.
- [3] Carliss Y. Baldwin and Kim B. Clark. 2000. *Design Rules, Volume 1: The Power of Modularity*. MIT Press, Cambridge, MA. 484 pages.
- [4] Patrick Bitterman. 2018. Modeling water quality governance networks on the Missisquoi River Watershed. In *Lake Champlain Research Conference*. Burlington, VT.
- [5] Center for High Throughput Computing. 1988. HTCondor: High Throughput Computing. <http://htcondor.org>
- [6] Christopher Koliba, Asim Zia, and Brian H Y Lee. 2011. Governance Informatics: Utilizing Computer Simulation Models to Manage Complex Governance Networks. *The Innovation Journal: Innovations for the Public Sector* 16, 1 (2011), Article 3.
- [7] Christopher Koliba, Asim Zia, Andrew W Schroth, Arne Bomblied, Judith Van Houten, and Donna M Rizzo. 2016. The Lake Champlain Basin as a Complex Adaptive System: Insights from the Research on Adaptation to Climate Change (RACC) Project. *Vermont Journal of Environmental Law* 17, 4 (2016), 533–563.
- [8] Computational and Information Systems Laboratory. 2020. *Cheyenne: HPE/SGI ICE XA System (University Community Computing)*. National Center for Atmospheric Research, Boulder, CO. <https://doi.org/10.5065/D6RX99HX>
- [9] Ewa Deelman, Karan Vahi, Gideon Juve, Mats Rynge, Scott Callaghan, Philip J Maechling, Rajiv Mayani, Weiwei Chen, Rafael Ferreira da Silva, Miron Livny, and Kent Wenger. 2015. Pegasus: a Workflow Management System for Science Automation. *Future Generation Computer Systems* 46 (2015), 17–35.
- [10] Thomas Dietz. 2017. Drivers of Human Stress on the Environment in the Twenty-First Century. *Annual Review of Environment and Resources* 42, 1 (oct 2017), 189–213. <https://doi.org/10.1146/annurev-environ-110615-085440>
- [11] Elizabeth MB Doran. 2018. Unpacking intention: Using agent based models to predict adoption of best management practices in the Missisquoi River Watershed. In *Lake Champlain Research Conference*. Burlington, VT.
- [12] Elizabeth MB Doran, Asim Zia, Stephanie Hurley, Yu-Shiou Tsai, Christopher Koliba, E Carol Adair, Rachel E Schattman, V Ernesto Mendez, and Donna M Rizzo. 2019. Social-Psychological Determinants of Farmer Intention to Adopt Nutrient Best Management Practices: Implications for Resilient Adaptation to

- Climate Change in the Lake Champlain Basin. In *International Symposium of the North American Lake Management Society (NALMS)*. Burlington, VT.
- [13] William Gibson. 2018. Modeling the impact of extreme events on the water quality in Lake Champlain. In *Lake Champlain Research Conference*. Burlington, VT.
- [14] Jory S. Hecht, Asim Zia, Donna M. Rizzo, Andrew W. Schroth, Patrick J. Clemins, Matthew C. H. Vaughan, and Kristen E. Underwood. 2018. The systematic underestimation of nutrient load variability in coupled streamflow-water quality models: Effects on lake cyanobacteria blooms. In *American Geophysical Union Fall Meeting (AGU)*. Washington, DC.
- [15] Jory S. Hecht, Asim Zia, Donna M. Rizzo, Andrew W. Schroth, Patrick J. Clemins, Matthew C. H. Vaughan, and Kristen E. Underwood. 2018. The systematic underestimation of nutrient load variability in coupled streamflow-water quality models: Effects on lake cyanobacteria blooms. In *2018 American Geophysical Union (AGU) Fall Meeting2*. Washington, DC.
- [16] Huanping Huang, Jonathan M. Winter, Erich C. Osterberg, Janel L. Hanrahan, Cindy L. Bruyère, Patrick J. Clemins, and Brian Beckage. 2018. Simulating Extreme Precipitation in the Lake Champlain Basin using a Regional Climate Model: Limitations and Uncertainties. In *2018 American Geophysical Union (AGU) Fall Meeting*. Washington, DC.
- [17] HydroNumerics. [n.d.]. Hydrodynamic-Aquatic Ecosystem Model (AEM3D). <http://www.hydronumerics.com.au/software/aquatic-ecosystem-model-3d>
- [18] Peter D F Isles, Courtney D Giles, Trevor A Gearhart, Yaoyang Xu, Greg K Druschel, and Andrew W Schroth. 2015. Dynamic internal drivers of a historically severe cyanobacteria bloom in Lake Champlain revealed through comprehensive monitoring. *Journal of Great Lakes Research* 41, 3 (2015), 818–829. <https://doi.org/10.1016/j.jglr.2015.06.006>
- [19] Peter D. F. Isles, Yaoyang Xu, Jason D. Stockwell, and Andrew W. Schroth. 2017. Climate-driven changes in energy and mass inputs systematically alter nutrient concentration and stoichiometry in deep and shallow regions of Lake Champlain. *Biogeochemistry* 133, 2 (apr 2017), 201–217. <https://doi.org/10.1007/s10533-017-0327-8>
- [20] Alexandre Julliard. 1993. Wine. <http://winehq.org>
- [21] Clelia L. Marti, Andrew W. Schroth, and Asim Zia. 2019. Physical and Biogeochemical Processes across Seasons in Missisquoi Bay, Lake Champlain: Insights from a Three-dimensional Model. In *American Geophysical Union Fall Meeting (AGU)2*. San Francisco, CA.
- [22] Ibrahim Nourain Mohammed, Arne Bombliès, and Beverley C. Wemple. 2015. The use of CMIP5 data to simulate climate change impacts on flow regime within the Lake Champlain Basin. *Journal of Hydrology: Regional Studies* 3 (mar 2015), 160–186. <https://doi.org/10.1016/j.ejrh.2015.01.002>
- [23] Finn Müller-Hansen, Maja Schlüter, Michael Mäs, Jonathan F. Donges, Jakob J. Kolb, Kirsten Thonicke, and Jobst Heitzig. 2017. Towards representing human behavior and decision making in Earth system models – an overview of techniques and approaches. *Earth System Dynamics* 8, 4 (nov 2017), 977–1007. <https://doi.org/10.5194/esd-8-977-2017>
- [24] D. L. Parnas. 1972. On the criteria to be used in decomposing systems into modules. *Commun. ACM* 15, 12 (dec 1972), 1053–1058. <https://doi.org/10.1145/361598.361623>
- [25] David Poole and Alan Mackworth. 2017. *Artificial Intelligence: Foundations of Computational Agents* (2nd ed.). Cambridge University Press.
- [26] Linyuan Shang. 2018. *Climate Change and Land Use/Cover Change Impacts on Watershed Hydrology, Carbon, Nutrient Dynamics - A Case Study in Missisquoi River Watershed*. Doctoral Dissertation. University of Vermont.
- [27] W. C. Skamarock, J. B. Klemp, J. Dudhia, D. O. Gill, Z. Liu, J. Berner, W. Wang, J. G. Powers, M. G. Duda, D. M. Barker, and X.-Y. Huang. 2019. *A Description of the Advanced Research WRF Version 4*. Technical Report. National Center for Atmospheric Research (NCAR). 145 pages. <https://doi.org/10.5065/1dfh-6p97>
- [28] C. L. Tague and L. E. Band. 2004. RHESys: Regional Hydro-Ecologic Simulation System—An Object-Oriented Approach to Spatially Distributed Modeling of Carbon, Water, and Nutrient Cycling. *Earth Interact.* 8, 19 (2004), 1–42.
- [29] The Intergovernmental Panel on Climate Change. 2014. *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Technical Report. The Intergovernmental Panel on Climate Change, Geneva, Switzerland. 151 pages.
- [30] Yushiou Tsai, Asim Zia, Christopher Koliba, Gabriela Bucini, Justin Guilbert, and Brian Beckage. 2015. An interactive land use transition agent-based model (LUTABM): Endogenizing human-environment interactions in the Western Missisquoi Watershed. *Land Use Policy* 49 (dec 2015), 161–176. <https://doi.org/10.1016/j.landusepol.2015.07.008>
- [31] University of Chicago. 1997. Globus. <https://globus.org>
- [32] Detlef P. van Vuuren, Jason Lowe, Elke Stehfest, Laila Gohar, Andries F. Hof, Chris Hope, Rachel Warren, Malte Meinshausen, and Gian-Kasper Plattner. 2011. How well do integrated assessment models simulate climate change? *Climate Change* 104, 2 (jan 2011), 255–285. <https://doi.org/10.1007/s10584-009-9764-2>
- [33] Léon Walras. 1877. *Elements of Pure Economics*.
- [34] Jonathan M. Winter, Brian Beckage, Gabriela Bucini, Radley M. Horton, Patrick J. Clemins, Jonathan M. Winter, Brian Beckage, Gabriela Bucini, Radley M. Horton, and Patrick J. Clemins. 2016. Development and Evaluation of High-Resolution Climate Simulations over the Mountainous Northeastern United States. *Journal of Hydrometeorology* 17, 3 (mar 2016), 881–896. <https://doi.org/10.1175/JHM-D-15-0052.1>
- [35] Asim Zia, Arne Bombliès, Andrew W. Schroth, Christopher Koliba, Peter D F Isles, Yushiou Tsai, Ibrahim N Mohammed, Gabriela Bucini, Patrick J Clemins, Scott Turnbull, Morgan Rodgers, Ahmed Hamed, Brian Beckage, Jonathan Winter, Carol Adair, Gillian L Galford, Donna Rizzo, and Judith Van Houten. 2016. Coupled impacts of climate and land use change across a river-lake continuum: insights from an integrated assessment model of Lake Champlain's Missisquoi Basin, 2000–2040. *Environmental Research Letters* 11, 11 (nov 2016), 114026. <https://doi.org/10.1088/1748-9326/11/11/114026>