Columbus, Ohio
📱 614 407 8980
✉ kaan@bgenc.net
🌐 bgenc.net
⚫ SeriousBug
📅 January 4, 1995

Kaan Genç

Education

M.S. **The Ohio State University**, *Computer Science & Engineering*, August 2017 – Dec 2021

B.Sc. **İzmir University of Economics**, *Software Engineering*, September 2013 – June 2017, Graduated with high honors

Work

Platform Engineer, *Tailwind*, June 2022 – current

I joined the Platform Team at Tailwind.

Software Engineer, *Dendron*, May 2021 – June 2022

Dendron is a note taking and personal knowledge management solution that is built for developers. At Dendron I designed and lead projects to make the user experience richer and more straightforward. I worked on both editor and web frontend features such as tags, text highlighting, search capabilities, and more. I also wrote documentation, built testing harnesses, measured performance and usage metrics, and helped with customer support. My work involved TypeScript, React, and Next.js for the software; and Amplitude and Sentry for telemetry and error monitoring.

Graduate Assistant, *The Ohio State University*, August 2017 – May 2021

I worked during grad school, both teaching and doing research.
I was given full responsibilities for the "Advanced C Programming" class for 2 years where I taught classes, and prepared lecture materials and exams. Student evaluations ranked my teaching well above average compared to the rest of the department.
My research was focused on dynamic program analysis for data race detection and persistent memory. Work I led or assisted resulted in several publications at top journals, which are listed at the end of my CV.

Personal Projects

Click on the names of these projects for more details.

Bulgur Cloud Simple cloud storage that's easy to self host. Bulgur cloud is a solution similar to Google Drive or Nextcloud, but can be deployed without any databases or complex configuration. It is built with a Rust backend and a React Native frontend using TypeScript and Redux.

gandi-live-dns A dynamic DNS system that uses Gandi's APIs to automatically update your DNS entries. Makes it possible to self host web services at home without a static IP address. Released as a self contained binary, with docker containers and systemd services available.

Image Survey A web application to run surveys where users compare two images side by side. Built with a Python backend using sqlite, and a React frontend.

Awards

2021 **Graduate Research Award**, *Ohio State University Department of Computer Science and Engineering*

2020 **Distinguished Artifact Reviewer**, *OOPSLA*
ACM Conference on Object-Oriented Programming, Systems, Languages & Applications

Research

My research focuses on efficient, scalable, persistent (durable) transactions combining commodity non-volatile memory with hardware transactional memory, and predictive data race detection methods. This research allows scalable, efficient programs taking advantage of new hardware to be developed easily. Below is all of my publications, including links to our open source implementations.

PLDI 2020 **Crafty: Efficient, HTM-Compatible Persistent Transactions**, *Kaan Genç, Michael D. Bond, and Guoqing Harry Xu*, ACM SIGPLAN Conference on Programming Language Design and Implementation, Online, June 2020
Extended Paper   Paper   Talk   Implementation

Non-volatile memory combines byte-addressability of DRAM with durability of persistent storage, but it presents many challenges to consistency. Prior works providing consistency incur significant performance costs or require hardware modifications. We propose a new method that provides fully ACID transactions efficiently on existing hardware using our novel logging method, and our efficient adaptive method for providing consistency.

OOPSLA **Dependence-Aware, Unbounded Sound Predictive Race Detection**, *Kaan*
2019 *Genç, Jake Roemer, Yufan Xu, and Michael D. Bond*, ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, Athens, Greece, October 2019

Extended Paper   Paper   Talk   Implementation

Data races can cause bugs that are hard to diagnose. Predictive data race analyses can find races from a multitude of program executions by analysing a single execution, but prior works miss many races. We introduce two new analyses incorporating data and control dependence, improving data race detection capabilities compared to prior works and advancing what can be predicted from a single execution.

PLDI 2020 **SmartTrack: Efficient Predictive Race Detection**, *Jake Roemer, Kaan Genç, and Michael D. Bond*, ACM SIGPLAN Conference on Programming Language Design and Implementation, Online, June 2020

Extended Paper   Paper   Implementation

Predictive data race analyses detect hard-to-find races in programs, but cause large performance impacts compared to widely used happens-before analysis. Our work introduces a variety of optimizations, bridging the performance gap for two analyses from prior work and a new analysis we propose.

PLDI 2018 **High-Coverage, Unbounded Sound Predictive Race Detection**, *Jake Roemer, Kaan Genç, and Michael D. Bond*, ACM SIGPLAN Conference on Programming Language Design and Implementation, Philadelphia, PA, USA, June 2018

Extended Paper   Paper   Implementation

Prior predictive data race analyses either miss some races or can not analyze full program executions. Our work presents a new analysis that finds all predictable data races from a single execution, but also finds some false races. We then present our novel algorithm which filters out false races, making the approach sound overall.