

I, Robot, Architect

John Demme Simha Sethumadhavan
Department of Computer Science, Columbia University
Email: jdd2127@columbia.edu, simha@cs.columbia.edu

2035: Ode to the Architect

*Architects, did we thoroughly enjoy the times;
a revolution unmatched did you bring.*

*This profession we lay to rest today
leaves behind the legacy of children:
computational science,
communications,
and indeed all technology
can trace their lineage to you;*

*However, parricide it did come to,
as the old cannot keep up with the new.*

*Lost, like the factory workers of yore,
to the tides of automation were you.*

2025: First Automatically Designed Processor Released

Bengaluru, India— Computer chip manufacturers rejoiced today at the release of the first automatically designed chip with several specialized functional units. A leading manager remarked, “This marks the beginning of a new design methodology. We no longer require the skills of an architect; only access to code pattern statistics is necessary to advance the state of the art.” Ever since the controversial decision to record and mine information about the code being run on their processors, manufacturers have been better able to produce custom chips that closely match the usage patterns. “The addition of an automated system to design the chips based on runtime information has enabled us to further reduce our time to market, and increase the speed and power efficiency of our customers applications” said the technology’s proud father. In a related story, local bars saw an increase in traffic and a good deal of complaints about “damned soulless software.”

2020: Hardware Phones Home Code Statistics

Cambridge, MA— An unnamed popular hardware vendor was accused today of building backdoors into their hardware intended to collect and transmit information on common code sequences executed on their processors. Unnamed sources inside the company claim that this action was “a move motivated by business managers who wanted to undercut the profits made by a certain semantic code search organization.” The blogosphere is reporting this act as one of pure desperation, however. One venerable architect commented, “The seeds for this discovery were sown in 2000 AD when tools capable of continuous program monitoring, such as ATOM and PIN were first developed.”

2017: End of Moore’s Law

San Francisco, CA — Technology insiders are marking today as the end of Moore’s Law. Initially a date feared by the computer technology community as the end of industry growth, one chip company executive claims “Government power caps killed silicon technology scaling years ago, forcing us to specialize our chips.” These specialized chips feature several heterogeneous units aimed at increasing computing efficiency by including hardware support for very specific functions. By mining the COMPASS database, manufacturers have begun producing processors with these accelerators, increasing power efficiency by an order of magnitude.

2012: Beginning of Semantic Code Search

New York, NY — Researchers today announced a milestone in code-based search technology. Computer scientists believe that

using their semantic code search technology, programmers will be able to identify code semantically equivalent to their needs, leading to drastic increases in code re-use and programmer productivity. Initially, the search technology will especially be a great boon to programmers attempting to parallelize their applications, as it is being used in the COMPASS system.

2009: The COMPASS Project

Washington, DC — Computer scientists from Columbia University announced that they are developing a Community-driven Parallel Advisor for Sequential Software (COMPASS). Designed to assist programmers in adapting their code to parallel systems, COMPASS monitors expert parallel programmers, recording their code changes (before and after), summarizing this information and storing it in a centralized Internet-accessible database. When a relatively inexperienced new user wants to parallelize his/her code, the system first identifies the regions of code most warranting performance improvement (determined by profiling typical executions), and targets those regions which are most amenable to parallelization (by consulting its database of previously parallelized code). COMPASS then presents a stylized template, or “sketch”, that can be used as a starting point for parallelization by the new user. To effectively provide these capabilities, COMPASS uses an alias-free program segment representation as the basis for a semantic code search engine used to locate parallelization solutions contributed by users; It also uses a code recommendation engine that uses approximate graph matching to find similar code segments.

To the best of our knowledge, we are the first to propose leveraging collective human wisdom — crowdsourcing — to parallelize code. This is a new way of thinking about (multicore) performance engineering enabled by the connectedness brought about the Internet and the wide availability of parallel machines. COMPASS is not just a reactive solution to fundamental changes in computer architecture; it provides an infrastructure that can proactively influence and aid in the design of new computer systems. For example, in the steady-state COMPASS could be data-mined to determine the most frequently requested optimizations and that information utilized by computer architects to create new hardware extension units or by compiler writers to generate targeted compiler analyses that speedup these frequently executed regions.

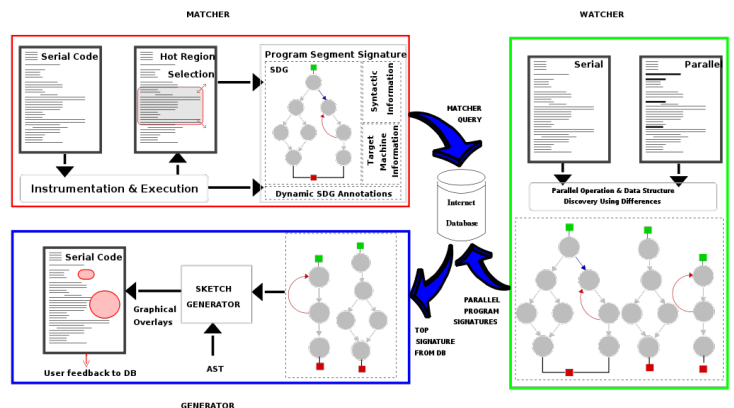


Figure 1: The COMPASS System

COMPASS Project Members: Simha Sethumadhavan, Gail Kaiser, Nipun Arora, Ravindra Babu Ganapathi, John Demme