101

102

103

104

105

106

107

108

109

110

Your next e-graph framework looks like Datalog

Yihong Zhang University of Washington USA yz489@cs.washington.edu

Recent years have seen a rejuvenation of e-graphs in many areas, ranging from floating-point arithmetic [14] to machine learning compiler [16] and from computational fabrication [11] to rule synthesis [12]. Several e-graph frameworks emerge as a result [3, 6, 15]. Among these frameworks, egg, the e-graph framework that first introduced the idea of rebuilding and e-class analyses, greatly expands the capability of the e-graph data structure by scaling up congruence maintenance with novel algorithms and by providing an extensible interface based on e-class analyses.

18 However, the wide variety of e-graph applications is plac-19 ing new requirements on the capability of e-graph frame-20 works. Some of these requirements are difficult to fulfill in 21 existing frameworks like egg. For example, applications like 22 e-graph-based tensor graph optimizations [16] use a stan-23 dard extension to e-matching called multi-patterns, which 24 egg does not readily support (as of mid-Mar 2022). A pull 25 request was made recently in an attempt to support multi-26 patterns. However, it requires modifications to the current 27 rewrite interface, and the proposed implementation does not 28 benefit from recent progress in e-matching [17]. As another 29 example, while equational reasoning is efficiently supported 30 in egg, non-symmetric reasoning like the logical implication 31 relation is fairly non-trivial and potentially inefficient in egg. 32 Yet practical reasoning may interleave symmetric relations 33 with non-symmetric ones. To remedy these deficiencies, we 34 need a radical change to the current interface. We call the 35 new interface design egg 2.0. 36

What could egg 2.0 look like? Advances in e-graph and database researches give us some clues. In previous work [17], we improved e-matching by reducing it to queries over relational databases. This hints at the underlying connection between e-graphs and relational databases. Moreover, Datalog, a fixpoint-based relational language, is able to express various non-symmetric relations (e.g., reachability) and can compute them using efficient algorithms (e.g., the semi-naïve evaluation algorithm [7]). Modern Datalog engines are also being extended to support efficient equational reasoning. For example, Soufflé has first-class support for efficient equivalence relations [13]. Finally, the semiring [1, 8] and lattice [4, 10] semantics of relational databases precisely capture the monotonic nature of e-class analyses.

We argue that egg 2.0 should be a Datalog language. This will at least have the following advantages.

55

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

- 1. A relational representation of e-graphs in Datalog could improve the performance of e-matching asymptotically in many cases, as shown in our previous work on relational e-matching.
- 2. The efficient evaluation algorithms designed for Datalog, such as semi-naïve evaluations, could benefit rule rewriting in e-graphs.
- 3. Rules in Datalog are naturally multi-patterns. This will allow first-class support for multi-patterns, whose performance will also benefit from relational e-matching.
- 4. The well-studied lattice semantics of relations supports and in fact generalizes e-class analyses in egg.

As our first step, we built egg^{Q1} , a relational e-graph framework on top of SQLite. In previous work, we build a prototype implementation for e-graphs with SQLite as well [17]. However, that prototype only supports basic e-graph operations like insertion and merging. egg^{Q} significantly expands its usability with support for match-apply iterations and multipatterns, which allows users to write real-world applications like equality saturation. To our knowledge, this is the first full-fledged e-graph framework fully on top of a relational database.

egg^Q can be viewed both as a multi-pattern equational reasoning language for e-graphs and as a Datalog language with an internalized notion of congruence. egg^Q has a Datalog-like surface language, and it translates e-graph operations into SQL statements, which are executed in SQLite. In egg^Q, we use relation Add(x_1, x_2, c) to represent a term Add(x_1, x_2) with e-class id *c*. The associativity rule of Add can be specified as follows:

 $(\operatorname{\mathsf{Add}} \alpha_1 (\operatorname{\mathsf{Add}} \alpha_2 \alpha_3)) @\alpha_4 \Rightarrow (\operatorname{\mathsf{Add}} (\operatorname{\mathsf{Add}} \alpha_1 \alpha_2) \alpha_3) @\alpha_4.$

The annotation syntax $p@\alpha$ means that α denotes the id of (sub)pattern p. The semantics of rules is as follows: for each pattern on the left-hand side, substitute and populate patterns on the right-hand side and unify the patterns annotated with the same id.

egg^{\heartsuit} is performant. It uses a novel algorithm for performing batched rewrites. Such an algorithm is reminiscent of the chase [2], a procedure used in relational databases for repairing databases using functional dependency. It also uses an efficient algorithm adapted from egg for maintaining congruence. Both algorithms lay a solid foundation for efficient computations in egg 2.0. egg^{\heartsuit} also demonstrates how far we

EGRAPHS '22, June 14–15, 2022, San Diego, CA, USA
2022.

¹Pronounced as egg lite.

118

119

120

121

122

155

156

157

11	egg	egg [♯]
12	Equational rewrites	Datalog rules
13	Congruence rules	Functional dependencies (FD)
14	E-classes	User-defined sorts
5	E-class merges	FD repair through unification
6	E-class analyses	User-defined lattices
7	E-class analysis maintenance	FD repair through lattice joins

Figure 1. The correspondence between constructs in egg and in egg[♯].

can push the limits of an in-memory database system and use it as an e-graph engine. The preliminary benchmark shows that, even with the overheads for interpretation, parsing, and communication, egg^{\heartsuit} is within one order of magnitude slower than egg, which is highly customized for the e-graph workload.

129 egg^{∇} also exhibits some interesting designs. For example, 130 different from a traditional e-graph implementation, egg 131 does not have a global union-find data structure. Instead, a 132 local union-find is created transiently during each rebuild-133 ing. This challenges the traditional view that an e-graph is 134 a DAG of terms and an equivalence relation over the terms 135 [5]. In fact, the design of egg^{\heartsuit} demonstrates that the equiva-136 lence relation in an e-graph can be avoided with aggressive 137 normalization. Moreover, although our current egg^V imple-138 mentation only supports congruence rules, we realized that 139 congruence is just a special kind of functional dependency 140 over the database. This again challenges the long-held belief 141 that congruence is an essential ingredient to e-graphs, inspir-142 ing future designs for egg 2.0. We have discovered several 143 interesting non-congruent functional dependencies so far. 144 One of them is the destructor. For example, snoc, the dual 145 to cons, takes a list l and returns x and l', the head and tail 146 of the list. In egg^Q, it is represented as snoc(l, x, l') and this 147 multi-output operator naturally has a functional dependency 148 $l \longrightarrow x, l'$, meaning that each l uniquely determines both x 149 and l'. Unification (with injective rules) is another kind of 150 functional dependency that we find interesting. For example, 151 the unification closure of a type theory with function type 152 \rightarrow and type variables will have the following rule (besides 153 transitivity, reflexivity, and commutativity of \equiv): 154

$$\frac{e_1 \rightarrow e_2 \equiv e_1' \rightarrow e_2'}{e_1 \equiv e_1' \land e_2 \equiv e_2'}$$

In egg $^{\heartsuit}$, \rightarrow is represented as a relation with three columns e_1, e_2 , and c, and the above rule is equivalent a functional dependency from c to e_1 and e_2 . Unification is known in the literature to be the dual to congruence [9]. The functional dependency captures this duality: reversing the arrows in functional dependencies in a congruence theory produces its unification dual. 166

167

168

169

170

171 172

173

174

175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

Based on the experience with egg^{\heartsuit} , we started a potential design for egg 2.0, which we call $egg^{\sharp 2}$. egg^{\sharp} is an extension to Datalog that provides a unified language for describing congruence reasoning and e-class analyses based on functional dependencies. Relations in egg^{\sharp} are annotated with functional dependencies. Values in egg^{\sharp} are divided into sorts and lattices. Sorts are uninterpreted and values of the same sort can be unified. Relations whose dependent sets (columns determined by other columns) have only sort values generalize e-classes in an e-graph, and relations whose dependent sets have only lattice values generalize e-class analyses. In other words, in egg^{\sharp} , e-classes and e-class analyses are just relations with different dependent sets.

When two values of the same sort are unified, they are no longer distinguishable in egg^{\sharp} . Such unification could potentially break the integrity of functional dependencies, i.e., multiple distinct tuples with the same determinant set (columns that determine other columns) can exist after unification. egg^{\sharp} remedies these violations. For each sort column in the dependent set, egg^{\sharp} unifies the sort values in that column, which makes values in the column indistinguishable and therefore unique again. For each lattice column in the dependent set, the new, unique value for each column is computed as the join of the lattice values in that column. As a result, functional dependency repair in egg^{\sharp} unifies the semantics of e-classes and e-class analyses.

egg[#] greatly expands the expressivity of egg: It has natural support of multi-patterns thanks to the relational representation. As an extension to Datalog, it also supports various kinds of reasoning expressible in Datalog, including nonsymmetric ones. Finally, lattice values in egg[#] generalize e-class analyses. As an example, egg[#] allows interdependent analyses, which are naturally expressed as egg[#] rules over several analysis relations. In contrast, e-class analyses in egg are not composable. With the new expressive power, egg[#] is able to express the classical type inference algorithm for Hindley-Milner type systems.

egg^{\sharp} also benefits from the efficient evaluation algorithms in Datalog. For example, a straightforward semi-naïve evaluation algorithm for egg^{\sharp} exists, which may be unintuitive or inefficient in the traditional representation of e-graphs.

In the talk, I will first describe issues with the current egg interface. Next, I will introduce egg^Q and describe the algorithms that make it practical. In particular, I will go through the chase-like algorithms for match-apply iterations in equality saturation and the rebuilding algorithm in egg^Q. I will also go through some alternative designs in egg^Q. Finally, I will present egg^{\sharp}'s early design, some example egg^{\sharp} programs, and its evaluation algorithms.

²Pronounced as egg sharp.

References 221

- 222 [1] Mahmoud Abo Khamis, Hung Q. Ngo, Reinhard Pichler, Dan Suciu, and Yisu Remy Wang. 2022. Convergence of Datalog over (Pre-) Semirings. 223 In Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium 224 on Principles of Database Systems (Philadelphia, PA, USA) (PODS '22). 225 Association for Computing Machinery, New York, NY, USA, 105-117. 226 https://doi.org/10.1145/3517804.3524140
- 227 Michael Benedikt, George Konstantinidis, Giansalvatore Mecca, Boris 228 Motik, Paolo Papotti, Donatello Santoro, and Efthymia Tsamoura. 2017. Benchmarking the Chase. In Proceedings of the 36th ACM SIGMOD-229 SIGACT-SIGAI Symposium on Principles of Database Systems (Chicago, 230 Illinois, USA) (PODS '17). Association for Computing Machinery, New 231 York, NY, USA, 37-52. https://doi.org/10.1145/3034786.3034796
- 232 [3] Alessandro Cheli. 2021. Metatheory.jl: Fast and Elegant Algebraic Computation in Julia with Extensible Equality Saturation. Journal of 233 Open Source Software 6, 59 (2021), 3078. https://doi.org/10.21105/joss. 234 03078 235
- [4] Neil Conway, William R. Marczak, Peter Alvaro, Joseph M. Hellerstein, 236 and David Maier. 2012. Logic and lattices for distributed programming. 237 In ACM Symposium on Cloud Computing, SOCC '12, San Jose, CA, USA, October 14-17, 2012. 1. https://doi.org/10.1145/2391229.2391230 238
- [5] Peter J. Downey, Ravi Sethi, and Robert Endre Tarjan. 1980. Variations 239 on the Common Subexpression Problem. J. ACM 27, 4 (1 Oct. 1980), 240 758-771. https://doi.org/10.1145/322217.322228 241
 - [6] Kiran Gopinathan. [n.d.]. ego: e-graphs in OCaml. https://gitlab.com/ gopiandcode/ego
- [7] Todd J. Green, Shan Shan Huang, Boon Thau Loo, and Wenchao 243 Zhou. 2013. Datalog and Recursive Query Processing. Found. Trends Databases 5, 2 (nov 2013), 105-195. https://doi.org/10.1561/1900000017
- [8] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. 2007. Prove-246 nance Semirings. In Proceedings of the Twenty-Sixth ACM SIGMOD-247 SIGACT-SIGART Symposium on Principles of Database Systems (Beijing, China) (PODS '07). Association for Computing Machinery, New York, 248 NY, USA, 31-40. https://doi.org/10.1145/1265530.1265535 249
- Paris C. Kanellakis and Peter Z. Revesz. 1989. On the relationship of [9] 250 congruence closureand unification. Journal of Symbolic Computation 251 7, 3 (1989), 427-444. https://doi.org/10.1016/S0747-7171(89)80018-5 252 Unification: Part 1.
- [10] Magnus Madsen, Ming-Ho Yee, and Ondřej Lhoták. 2016. From Datalog 253 to Flix: A Declarative Language for Fixed Points on Lattices. SIGPLAN 254 Not. 51, 6 (jun 2016), 194-208. https://doi.org/10.1145/2980983.2908096
- 255 [11] Chandrakana Nandi, Max Willsey, Adam Anderson, James R. Wilcox, 256 Eva Darulova, Dan Grossman, and Zachary Tatlock. 2020. Synthe-257 sizing Structured CAD Models with Equality Saturation and Inverse Transformations. In Proceedings of the 41st ACM SIGPLAN Conference 258 on Programming Language Design and Implementation (London, UK) 259 (PLDI 2020). Association for Computing Machinery, New York, NY, 260 USA, 31-44. https://doi.org/10.1145/3385412.3386012
- 261 [12] Chandrakana Nandi, Max Willsey, Amy Zhu, Yisu Remy Wang, Brett 262 Saiki, Adam Anderson, Adriana Schulz, Dan Grossman, and Zachary 263 Tatlock. 2021. Rewrite Rule Inference Using Equality Saturation. Proc. ACM Program. Lang. 5, OOPSLA, Article 119 (oct 2021), 28 pages. 264 https://doi.org/10.1145/3485496 265
- [13] Patrick Nappa, David Zhao, Pavle Subotić, and Bernhard Scholz. 2019. 266 Fast Parallel Equivalence Relations in a Datalog Compiler. In 2019 267 28th International Conference on Parallel Architectures and Compilation 268 Techniques (PACT). 82-96. https://doi.org/10.1109/PACT.2019.00015
- [14] Pavel Panchekha, Alex Sanchez-Stern, James R. Wilcox, and Zachary 269 Tatlock. 2015. Automatically Improving Accuracy for Floating Point 270 Expressions. SIGPLAN Not. 50, 6 (June 2015), 1-11. https://doi.org/10. 271 1145/2813885.2737959
- 272 [15] Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha. 2021. Egg: Fast and Extensible 273 Equality Saturation. Proc. ACM Program. Lang. 5, POPL, Article 23 (jan 274
- 275

242

244

245

2021), 29 pages. https://doi.org/10.1145/3434304

- [16] Yichen Yang, Phitchaya Mangpo Phothilimtha, Yisu Remy Wang, Max Willsey, Sudip Roy, and Jacques Pienaar. 2021. Equality Saturation for Tensor Graph Superoptimization. In Proceedings of Machine Learning and Systems. arXiv:2101.01332
- [17] Yihong Zhang, Yisu Remy Wang, Max Willsey, and Zachary Tatlock. 2022. Relational E-Matching. Proc. ACM Program. Lang. 6, POPL, Article 35 (jan 2022), 22 pages. https://doi.org/10.1145/3498696

276

277

278

279

280

281

282

283

328 329 330

318

319

320

321

322

323

324

325

326

327