An Update on Parsl Sustainability

Daniel S. Katz (d.katz@ieee.org, @danielskatz@fosstodon.org)
Chief Scientist, NCSA
Associate Research Professor, CS, iSchool
University of Illinois at Urbana Champaign

With contributions from Parsl team & community

ILLINOIS
NCSA | National Center for Supercomputing Applications

26 September 2024
ParslFest 2024

Parsl's external stakeholders

- Direct users: use Parsl for science/etc.
  - E.g., LSST DESC
- Platforms: platform developers use Parsl as a component of a platform/application used by end users
  - E.g., QC Archive, Globus Compute (was funcX)
- CI providers: support Parsl on their HPC/etc. system
  - E.g., Argonne, NCSA, NERSC
- Linked contributors: Link naturally complementary components with Parsl
  - E.g., Parsl provides interesting ways to describe related tasks, and Work Queue provides interesting ways to schedule those tasks: WorkQueue -> Parsl WorkQueue executor
- Funders
  - E.g., NSF, CZI, DOE, collaborating projects

Parsl history

- Initially supported by an NSF SI2 award from 2016-2022 (5 years + 1-year NCE)
- Released version 1.0 in 2020, now releasing weekly
  - Focus since v1.0 mostly maintenance rather than adding new features
- Initial funded development team
  - 2-4 people/FTEs per year
- Current funding (new NSF award, CZI award, contributions from projects that require Parsl) supports
  - ~1 FTE/year maintenance and development
  - ~0.5 FTE/year community management

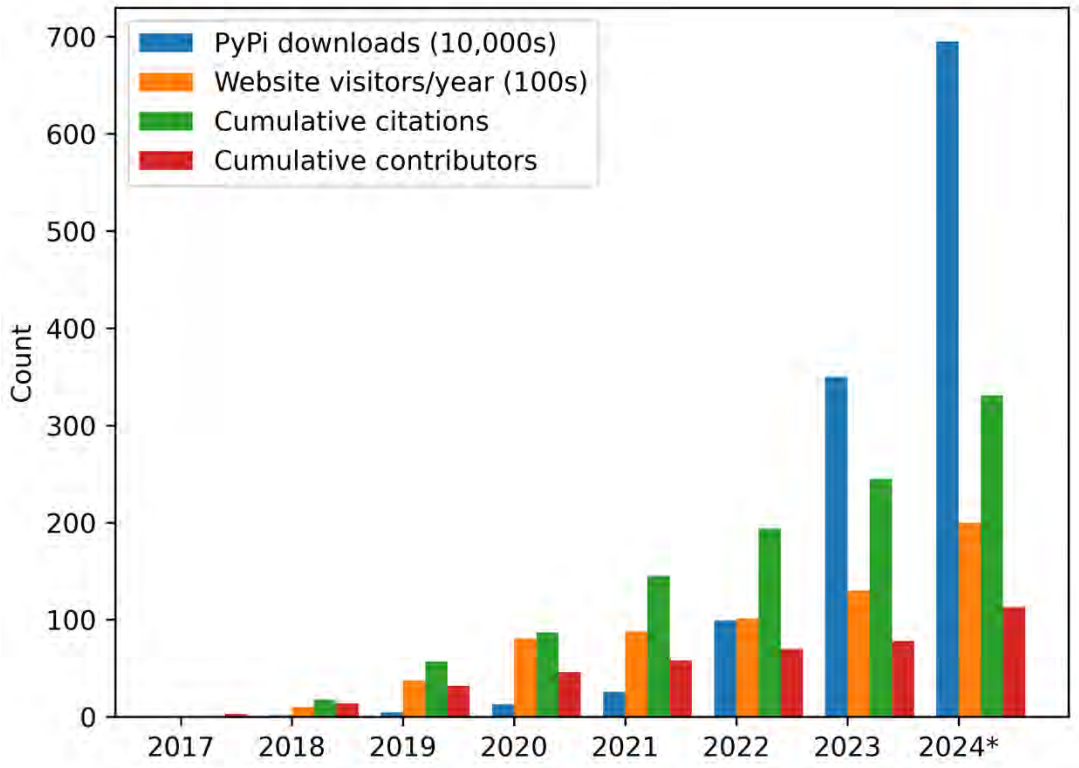

* Through ~1 Aug 2024

Want to balance resources & work

- Parsl resources
  - Grants
  - External funding from projects that depend on Parsl
  - Volunteer (in-kind) effort from groups that develop tools that use Parsl
  - Companies that use Parsl in their services
- Funded Parsl team does the core work, such as
  - Managing the community
  - Reviewing code contributions
  - Fixing bugs
  - Supporting users
  - Developing new features
  - Releasing new versions of the software
- Provided/volunteers resources can add features to Parsl and support some limited number of use cases
  - But aren’t currently sufficiently coordinated or aligned to fully support Parsl's core needs over multiple years

Past, current, and sustainable resources

Current NSF & CZI awards aim at making Parsl *sustainable*:

Resources (from various sources) will be able to perform all project activities well into the future

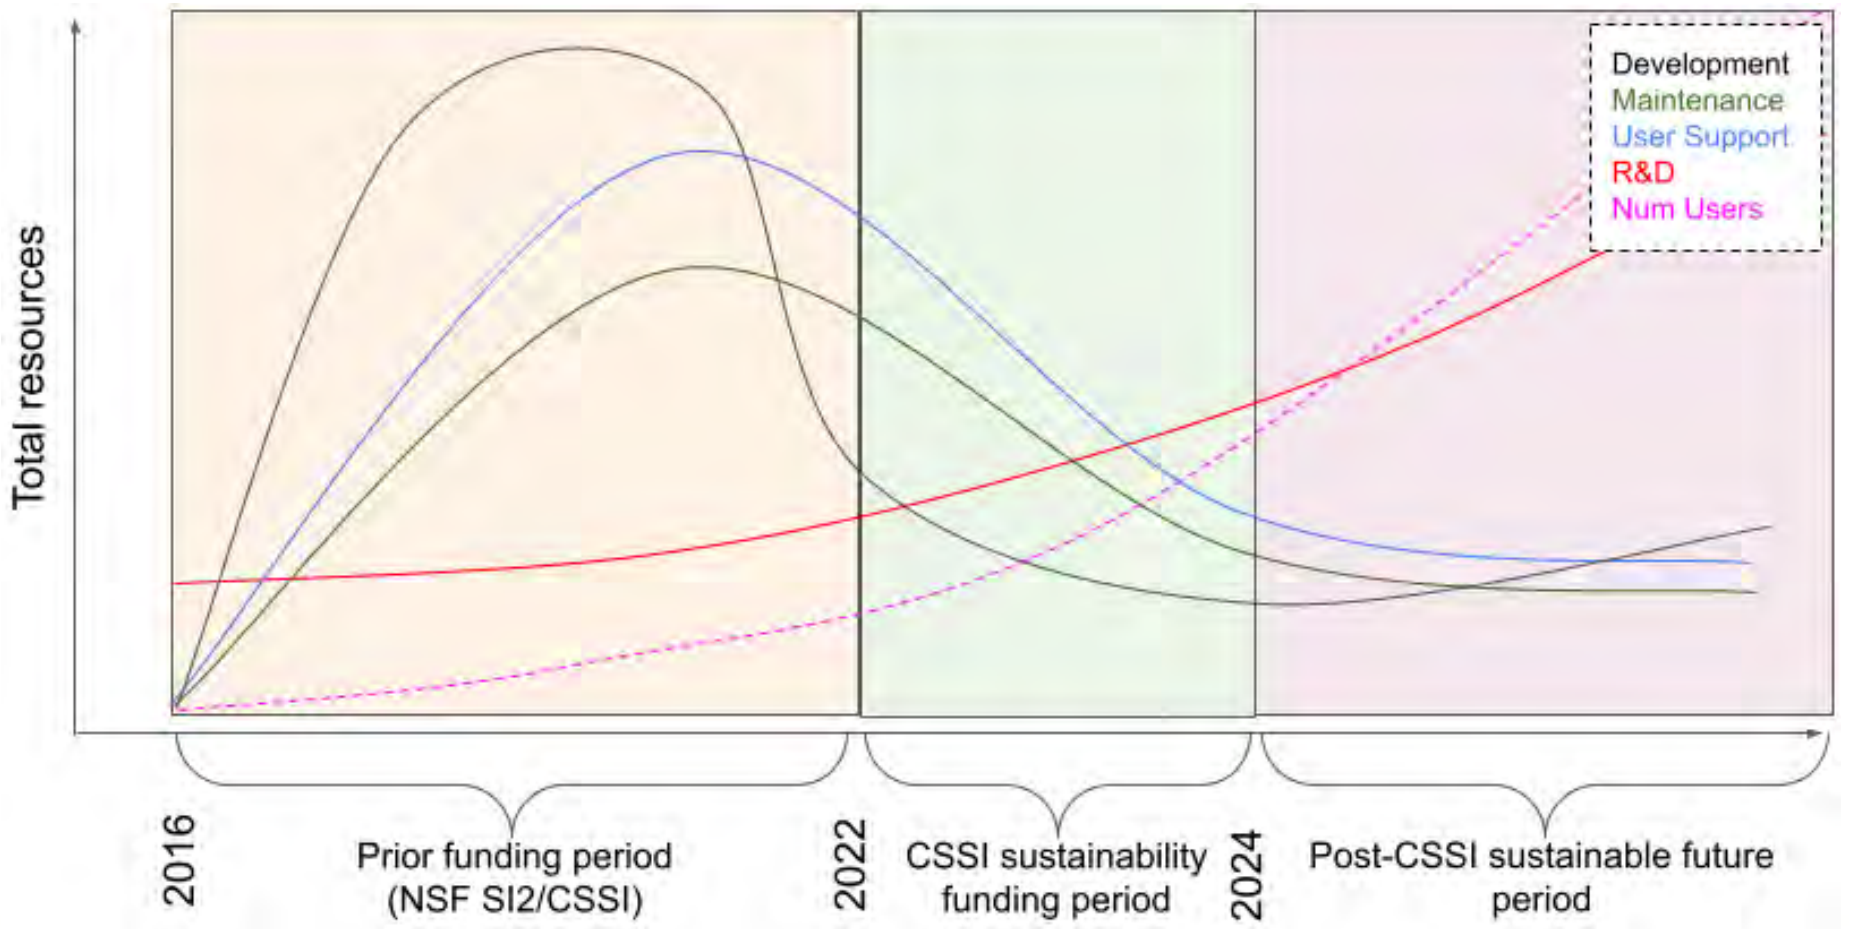

Path to sustainability

- At a high level, we want to follow the successful sustainability model of AstroPy, yt, etc.
- Work on community, governance, funding streams, innovation, training, outreach/engagement
- Work with other related community (software sustainability, workflows)
- Capture, document, and share sustainability lessons

- Also reduce technical debt
  - The easier it is to do things, the less resources are needed to do them

Changing developer work

- Current needs (e.g., maintenance, outreach, and support) differ from earlier in the project, leading to a need for new types of contributions
- Development activities include
  - Maintain the Parsl codebase
    - Including adding additional tests to improve code quality
  - Respond to issues, including supporting deployment on different computing resources
    - As community grows and diversifies, range of use cases and range of challenges also grow
  - Review contributed code
    - Leads us to develop minimal requirements on contributed code, starting with a pre-coding discussion and including plans for future maintenance and support of contributed code
- See Ben's previous talk

Changing community work

- In addition to developers, members of the community sometimes do (kind of in order of where community contributions actually happen)
  - Answer user questions
  - Share experiences (configs) for specific computing platforms
  - Share expertise in applying Parsl to different scholarly disciplines
  - Support outreach activities (e.g., presenting tutorials, hosting summer students, developing training materials for various domains)
  - Coordinate the yearly user meeting
  - Apply for funding
  - Advertise success stories in blogs
  - Manage social media
- Parsl community manager does/coordinates this

Where we are now

- Good news
  - Community growth – at least in part due to Sophie as community manager
    - Lots of contributors (see Ben's slides from earlier & many of the other talks)
    - Lots of users (see many of the other talks, plus lots of people who aren't here)
  - Parsl code has gotten better – at least in part due to Ben as core maintainer
    - Moving to more plug ins to reduce what the core code has to do
    - Removing old code that isn't used, doesn't work, etc.
    - More and better tests
- Less good news
  - Unclear how to sustain Sophie's and Ben's work
    - This core community and maintenance work is hard/impossible to rely on volunteers to do, at least for a project of Parsl's size

Acknowledgements

- Parsl has been supported by NSF (1550588, 2209919, 2209920) and the Chan Zuckerberg Initiative, and DOE, collaborating projects, and a community of developers, maintainers, and users